azalient

azalient.com

software for nanosimulation modelling

Commuter

User Guide - Manual

Table of Contents

Contents

Reference

Plugins

The Basics

User Guide

Why use Simulation Modelling?

A computer simulation is an attempt to model a real-life process or system on a computer so that it can be studied to see how the system works. By changing parameters and variables, predictions may be made about the behaviour of the system. Computer simulation is often used to model systems for which simple closed form analytical solutions are not possible. A computer simulation generates a sample of representative scenarios for a model in which a complete enumeration of all possible states would be prohibitive or impossible. Graphical computer simulations provide the added benefit of visualisation of the system as it changes over time.

In the analysis of people travelling, we use simulation:

- To analyse the effect of congestion on throughput, delay and smoothness (number of stops).
- To measure queue lengths, especially when queues extend back to the previous road or pathway junction.
- To measure environmental factors, such as emissions which can be estimated more precisely if we simulate each individual engine in the system.
- To measure trip statistics for a complete person-trip, over several modes of travel.
- As a test-bed for external control systems, such as traffic lights and variable speed limits.
- As an educational tool, as the graphical display makes it easier to visualize the behaviour of the system, even counter-intuitive effects such as backwards-travelling shock waves.

In the analysis of people travelling, simulation is particularly useful in congested conditions, where the free-flow relationships between flow, speed and density no longer apply.

Introduction to Commuter

Commuter is software for modelling people, vehicles or freight moving on a transport network. The people may be driving, or walking, or cycling, or be passengers on a train, or a bus. Commuter models these people all the way from start to finish.

Why is this useful? Today, a typical transport analysis study needs to look at *all* modes of transport, not just private cars. Following each person from person-origin to person-destination gives us a much clearer insight into the total cost of each trip. And if we know the total cost, we are better able to assess the potential benefit.

Build – Run (x N) – Analyse

To use Commuter, you first build a model, then you run the model and then you analyse the results. Some of the input parameters to the model are selected at random from distributions. You can use different values within these distributions each time you run the model, so normal practice is to run the model a number of times and analyse the distribution of results.

Visualization: Seeing is Believing

Commuter allows you to visualize your model, including all the roads, paths, and the people travelling on them, in three dimensions. You can follow a person through any part of their trip, and you can see how they are affected by other people, traffic and any control or information devices in your model. Visualization of this type makes it easier to understand the results, and very useful for illustrating the effects of proposed designs to your colleagues, clients or the public.

Agents

Commuter is a multi-agent simulation or agent-based model.

In simple terms, an agent in Commuter is a **person** or a **vehicle**.

A more detailed classification divides agents into several groups:

- people walking = pedestrians
- people cycling = cyclists
- people in public transport = passengers
- people in private transport = driver or passenger(s)
- private vehicles
- public transport vehicles
- units of freight

Note that a driver or passenger travelling within a vehicle retains all its behavioural characteristics. It is passive while travelling as a passenger. The behaviour characteristics of a vehicle are inherited from its driver.

A cyclist is a special type of private vehicle that has a driver, no passengers (and a zero-emission engine!) It is not particularly different from any other type of single-occupant vehicle, but is reported separately because within transport modelling it is traditionally considered as a separate class.

Freight units are a bit like people in that they aim to get from an origin area to a destination area by the lowest cost route. The main difference is that a freight unit cannot move by itself, and relies on moving channels (walkways) or being carried in dynamic-route (private) vehicles or fixed route (public transport) vehicles.

A Commuter Model

A Commuter model is saved as a single file with extension "**.aza**". This file contains all the input components, and can also contain the results. A single file can also contain variations to the model, for example for the base case and design proposals. So if you use it like this, you can think of the single file being like a "workspace" containing all resources. This makes it easier to transfer the work to a colleague or client.

Input – Components

A Commuter model consists of several input components, including:

- **Parameters:** defining how people behave in the model, for example on mode choice, route choice and speed choice.
- **Network**: a representation of the paths and roads used by the people in the model and the geographical areas defining their origins and destinations.
- **Control:** timing and sequencing information for devices such as traffic signals and pedestrian crossings.
- **Demand:** one or more tables of *aggregated* travel demand, defining how many people want to travel between each origin and destination, the time profile of this demand and a distribution by person-type.
- **Trips:** a *disaggregated* list of trips, defining the Agent character, origin, destination and departure time of every trip for a particular scenario.
- **Validation:** tables of observations that are used to validate the performance of the model. These may take many forms, but typical examples include turn counts and travel times.

These are the principal components in a model, but this is not an exhaustive list.

Output – Results

A Commuter model generates results in tables, similar to database tables. These tables can be stored within the **.aza** file, or can be exported in several formats, including **.xls** (Microsoft Excel compatible) and **.csv** (Comma-separated value text files).

There are many built-in results tables, or you can create plug-ins to extract and format results to suit a particular project. Built-in result tables include:

- Summaries for People/ Cyclists/ Public Transport/ Private Vehicles/ Freight:
    - Distance, time, stops for each mode
    - Modes includes walking, passenger, driving, waiting
    - Lane changes
    - Loop activations
    - Emissions (CO_2, NO, PM10)
- Detailed Public Transport Information
- Economic Evaluation
- Level of Service Reports

You can choose to analyse your results externally, or you can use the tools within Commuter, for example, you can compare two tables side-by-side where any differences (perhaps from different runs) will be highlighted.

The Commuter Main Window

When you start commuter, a window is opened. We call this window the Commuter Main Window. It contains:

- the **menu bar** along the top of the window
- the **left pane** containing a file tree, and other tabs
- the **right pane** containing a graphics panel, initially empty

The left pane is used:

- to browse files, or the components within a **File**
- to edit the model, by executing an **Action**
- to modify the graphical view, by changing the **Layer**, or moving the **View**
- to access or enable a **Tool** or **Plugin**
- to **Undo** an edit

The right pane shows a **Graphics** panel, containing a three-dimensional view of the model, showing the roads, paths, control devices, and the people and traffic as they travel in the simulation. The Graphics panel also contains the simulation **Toolbar**, which can be used to control the simulation - play, pause, rewind, etc.

[New for Commuter 5] The table view of the data is shown on the File tab, underneath the graphics panel. Slide the divider to the right, or use the small arrows at the top of the divider to view the table panel.

Navigation in 3-D

Navigation is possible using the mouse buttons, the keyboard, a joystick [NOTE 3] or a specialist 3-D mouse [NOTE 4].

Pan, Zoom, Tilt and **Bearing** are used to describe movement in a 3-D space. Definitions of these terms are given later in this section.

Mouse-1 means Mouse Button 1: for a standard set-up this is the left mouse button. You may have inverted the mouse buttons if, for example, you are left-handed.

Pan	Mouse-3-Drag Cursor Keys [NOTE 1]
Zoom	Mouse-Scroll-Wheel [NOTE 1] Mouse-1 on Control Bar on Right Edge of Window [NOTE 1] Page-Up Key / Page-Down Key [NOTE 1, NOTE 2]
Bearing	Mouse-1 on Control Bar on Top Edge of Window [NOTE 1] Mouse-1 + Mouse-3 (on left third or right third of window) Home-Key then Page-Up Key / Page-Down Key [NOTE 2]
Tilt	Mouse-1 on Control Bar on Left Edge of Window [NOTE 1] [CTRL]-Mouse-3 Mouse-1 + Mouse-3 (on central third of window) End-Key then Page-Up Key / Page-Down Key [NOTE 2]
Select	**Single Item:** Mouse-1-Double-Click or Mouse-2-Click **Additional Item:** [CTRL]-Mouse-1-Double-Click or Mouse-2-Click **Multiple Items:** Mouse-2-Drag or Mouse-1-Drag **Additional Items:** [CTRL]-Mouse-2-Drag **Clear Selections:** [ESC]
Drag All Items Selected	**Straight-Line-Horizontal:** Mouse-1-Drag + Tilt < 45 **Rotate-Horizontal:** Mouse-1 Control-Bar Bottom-Edge **Straight-Line-Vertical:** Mouse-1-Drag + Tilt > 60

Pan: To move the cursor and viewpoint horizontally, parallel to the ground plane. While panning the mouse cursor changes to a 4-way arrow.

Zoom: To move the viewpoint vertically. The cursor position does not change. The Zoom value can change from 2 to Infinity, representing the height of the viewpoint, in metres, above the ground plane.

Tilt: To tilt the view relative to the vertical axis. A Tilt value of 0 means you are looking straight down at the network in plan view, whereas a Tilt of 90 degrees means you are looking horizontally towards the horizon. Double-click on the Tilt bar to snap the view back to vertical.

Bearing: The bearing of your view, in degrees, clockwise from North. Zero degrees is due North, 90 degrees is due East, etc. Moving the mouse to the right turns the map clockwise, which means that the “straight ahead” direction from your point of view turns anti-clockwise. This is why the bearing decreases as you move the mouse to the right. Double-click on the rotation bar to snap the view back to due North / 0 degrees.

NOTE 1: Can be used with SHIFT key for fine resolution movement.

NOTE 2: **Page-Up** and **Page-Down** apply to the “Current Page Function” which can be cycled Zoom → Tilt → Bearing using the **Home** and **End** keys.

NOTE 3: Suitable joysticks include Logitech Extreme 3D Pro and Saitek Cyborg Evo. These can be used in combination with a mouse for two-handed operation. The Logitech is a right-hand-only moulding, but the Saitek can be set up for left-handed operation allowing simultaneous right-handed mouse operation.

NOTE 4: Suitable 3-D controllers include 3DConnexion Space Navigator , Space Pilot, or Space Explorer.

Selecting Objects

When objects are selected, they are highlighted in a different colour, which is orange by default. You can change the colour used for highlighting in the Layer Tab.

You can select any number of objects at any time; all objects selected form the Selection Set. The list of objects contained in the Selection Set is decribed at the top left corner of the graphics panel.

Normally selections operate only on visible objects, so you may need to change to another Layer, or click on the Layer tab and make the object that you want to select visible.

Action	Mouse Pointer	Selection Effect
Double-Click Mouse-Button-1	Over object	Select that object only.
CTRL-Double-Click Mouse-Button-1 or Single-Click Mouse-Button-2	Over object	Select that object if not selected, or de-select it, if already selected.
Mouse-Button-2	Dragging box	Clear selections, then select all objects within the box.
CTRL-Mouse-Button-2	Dragging box	Select all objects within box, or de-select any already selected.
Mouse-Button-1 (No current selections)	Dragging box	Select all objects within the box.

Modifier	Selection Effect
SHIFT	Modifies all actions above to select or de-select objects even if they are not currently visible.

Actions for Editing

An **Action** is something you do to edit the currently active model.
For example:

- Road **>** New Road at Cursor
- Zone **>** New Zone at Cursor
- Walkway **>** New Crossing on Lane
- Lane **>** New Loop
- Road **>** New Intersection
- Adjust

Some actions can be applied at any time. Other actions become available only when you have selected one or more objects.
Some example objects are a Road, Walkway, Node, Crossing, Loop or Intersection.

Commuter does not have separate editing modes. **All** actions are available at **all** times, as long as you have selected the objects to which the actions apply. That is, you can edit an Intersection with one click, then with the next mouse click, you can select a Loop and edit that. You do not need to switch between Loop-editing mode and Intersection-editing mode.

Use Action Tab or Right-Mouse Menu

Actions can be accessed either from the Action Tab or by a right-mouse menu on the graphics window. The right mouse menu shows only the currently available actions; the Action Tab always shows all actions, but some may be greyed out.

Demand and Trips

In Commuter, Trips must be generated from the Demand before the simulation begins, if there are to be any people or traffic in the model. Each generation of Trips is called a Trips Page. All the randomness required to effectively simulate people and traffic is contained within the Trips page. One key benefit of this approach is that a page of trips can be generated once, then used for several different simulation runs, ensuring repeatability, even if the underlying network changes.

The modelling process in Commuter is:

1. Build Network – Roads, Walkways, Crossings
2. Add Areas for people origin / destination
3. Add Zones for traffic origin /destination / parking
4. Define Demand matrices (directed O-D) or volumes (undirected)
5. **Generate Trips from Demand**
6. Run Simulation, collecting Results
7. Analyse Results

The Essentials

User Guide

Components

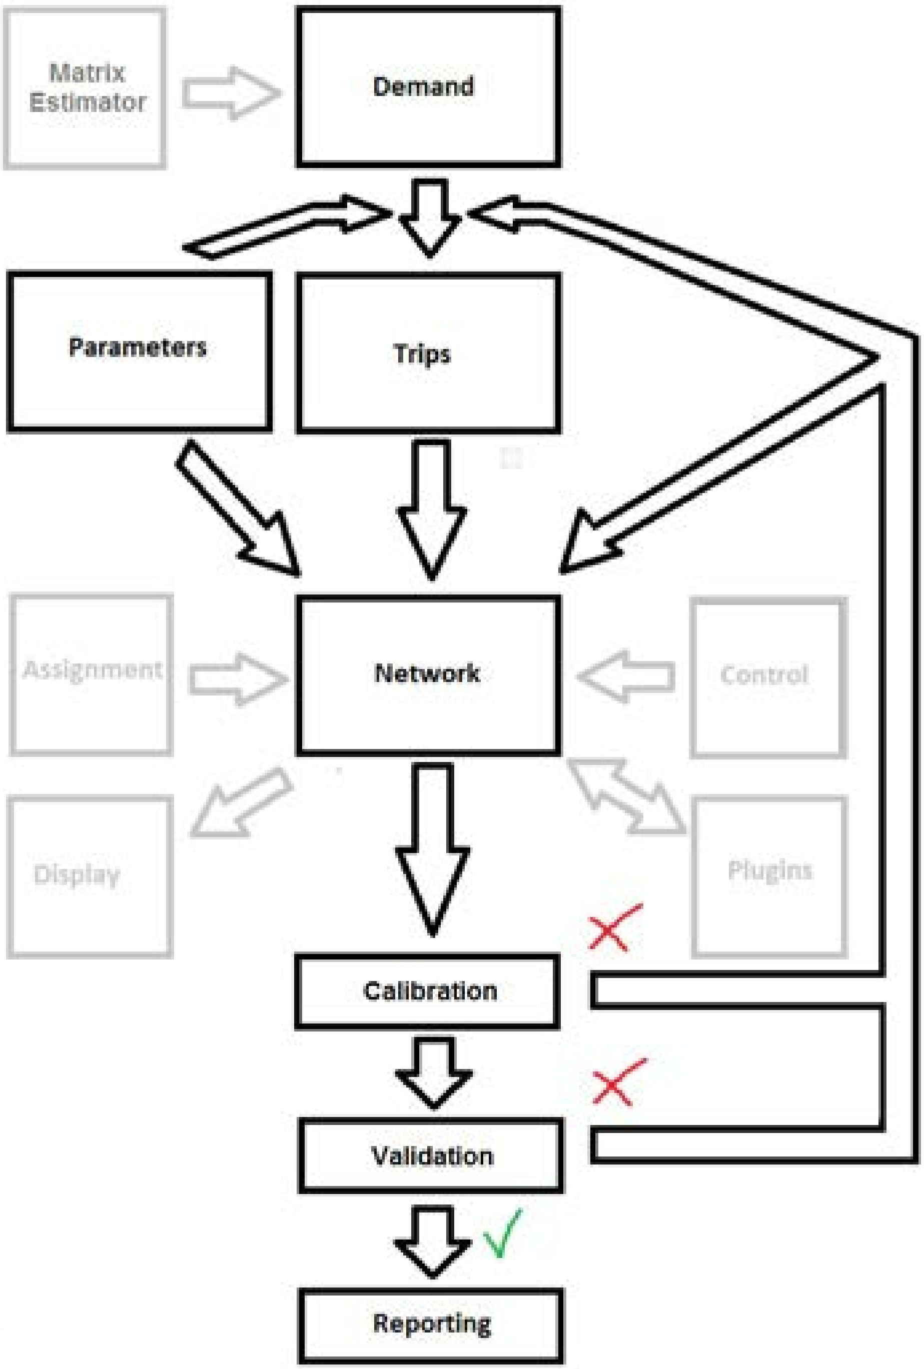
Matrix
Estimator
Demand
Parameters
Trips
Assignment
Network
Control
Display
Plugins
Calibration
Validation
Reporting

Components

The diagram shows the principal building blocks or Components of a Commuter model. The components shown in grey are not as central to the model and are discussed in the Advanced section.

The menus in Commuter are laid out according to these Components and Tables – the interface has been structured to provide a guide to the best order of defining the data for the model.

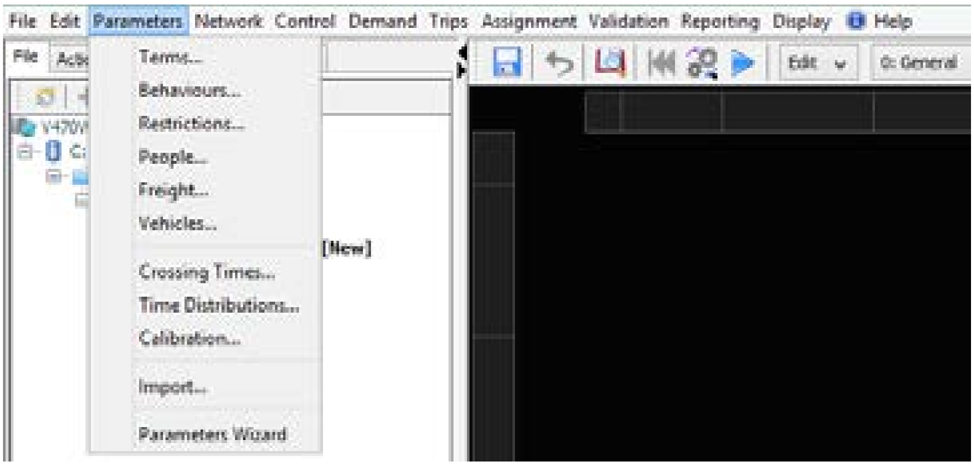

Parameters

Within Commuter, the Parameters component contains variables that control both trip generation and the model behaviour.

- **Terms -** A Term is a multi-purpose window in time. You can have as many terms as you like and they can overlap. Terms can be used to split the day into periods, but they can be used for many other things too. A Term is defined by { Start Time, End Time, Day of Operation }.

- **Behaviours -** A Behaviour contains parameters that control decisions made by each person in the model. These decisions include **mode and route choices,** lane selection and parking selection. A Behaviour defines weights that describe the relative **value of time, distance and price.** A Behaviour also contains parameters controlling temperament and driving style.

- **Restrictions –** A Restriction is used to limit the use of a walkway or a lane. Restrictions are based on mobs (groups of behaviours). A Restriction can be set up to operate all the time, or only during certain times (during a Term). As a restriction is defined by mob, the same restriction can be applied to both a walkway and a lane on the road. If you had two types of users - Staff and Visitors - You could use a "Staff Only" restriction to apply both to lanes on the road, for a parking area, and to walkways that accessed certain doors of a building.

- **People** (Person Types) - A Person Type defines the physical aspects of a person, such as size and speed. A single Behaviour is assigned to each person type, but there can be multiple person types with the same behaviour. For example, you might define Child, Adult and Elder, where the Child was smaller than the adult, and the Elder was slower than the Adult.

- **Vehicles** (Vehicle Types) - A Vehicle Type defines the physical aspects of a type of vehicle, such as "Small Car" "Medium Car", "Truck" etc. Each Vehicle Type must also be assigned a Behaviour, describing how the **driver** behaves. If specific behaviour is required for a type of vehicle, for example a lane restriction, then it will be necessary to define a behaviour for that type of vehicle. For example, to create a truck restriction, first define a **Truck Driver** behaviour, then assign that behaviour to the Truck vehicle type. A restriction can then be defined allowing only the Truck Driver behaviour. It is common to define multiple vehicle types for each behaviour. For example, if behaviours are “Business” and “Leisure”, and there are two physical sizes of car “Small” and “Large” then it is likely that there would be 4 types of vehicle defined: “Small – Business”, “Large – Business”, “Small – Leisure”, “Large – Leisure”.

- **Crossing Times** - these are used to define the **Walk**, **Flash** and **Don't Walk** times for pedestrian crossings.

- **Time Distributions** - these are used for boarding and alighting times on public transport, and also for delay times at entry and exit boom gates, or at toll plazas.

Demand

Demand consists of

- Demand Divisions
- Profiles
- Matrices (Directed) or Origin Volumes & Splits (Undirected)

Demand Division

A Demand Division defines a group of Types and assigns a proportion to each Type, with the proportions summing to 100%. The demand division is then associated with a demand matrix or origin volume.

Profiles

A Profile can be used to vary the rate at which demand is released over time. A Profile has a Term, defining start time and end time, a number of intervals, and a weight for each interval, specified as a percentage. The weights for all intervals must sum to 100%.

For example, in a motorway model with a term of 08:00 to 09:00, automated traffic count data is used to determine the profile of the northbound traffic released from the motorway South zone. This data shows:

- 800 vehicles 08:00 to 08:15,
- 2000 vehicles 08:15 to 08:30,
- 800 vehicles 08:30 to 08:45 and
- 400 vehicles 08:45 to 09:00

The total number of northbound vehicles in the modelled period is 4000. The Profile applied to the South zone would have a Term of 0800 to 0900, interval of 4, and demand release percentages of:

20% (0800 to 0815)
50% (0815 to 0830)
20% (0830 to 0845)
10% (0845 to 0900).

Directed Demand (OD Matrices)

Directed demand is a set of trips, where each trip has an origin and a destination. This demand is entered into the model as a two-dimensional table known as an Origin-Destination matrix (OD matrix).

In Commuter a matrix can be for people, vehicles or freight. A person matrix allows each person to use all modes of transport available to them, a vehicle matrix is used to generate vehicle trips. Vehicle OD matrices can be used to generate background traffic where no mode switch is required.

Freight is modelled as an agent in Commuter with no walk speed, The freight has route choice and moves along moving walkways and on available modes of transport.

Commuter has a two-level demand that allows you to enter as many person matrices, freight matrices and vehicle matrices as required.

Undirected Demand (Origin Volumes and Splits)

Undirected demand is made up of two elements: a one-dimensional origin volume and a set of splits. The origin volume specifies the number of people or vehicles to be released from each area or zone, respectively. The splits are used for route assignment at each route choice location – for example at intersections or walkway junctions. Undirected demand can be used for smaller networks to speed up the demand building process. When using undirected demand, it is important to recognise that each person or vehicle does not have a destination while travelling through the network; it discovers its destination only when it arrives.

Each route choice location automatically creates a “Split” object, that by default assigns an equal proportion of all incoming traffic to each of the available exit options. At a 4-way intersection, an incoming agent has a choice of left, right or straight ahead, and the default split object will stochastically assign the agent to one of those exits, each with equal probability of 1/3. This means that if there is a loop in the network, an agent may be directed around the loop, arriving back at the same point. With typical splits, the probability of this happening is low, and in some networks, following a loop around may reflect reality, for example where a pedestrian is window shopping. Some worked examples on the following pages illustrate the probabilities of following a loop.

The Split function is very simple; if more complex routing rules are required, then the route choice tool allows you to create rules that assign the exit (left, ahead, right) based on origin area, agent type, sign settings, etc. The outcome of a route choice rule can be a single exit or a stochastic distribution across multiple exits, similar to a Split.

Undirected Demand - Agents May Get Stuck in Infinite Loop

Assume a square grid network, often referred to as a Manhattan grid. Assume a vehicle has the choice of going left, ahead or right.

Example 1:

proportion of traffic turning left = P_L = 15%

proportion of traffic turning right = P_R = 10%

proportion of traffic going straight ahead = P_A = 75%

probability of left loop (four sequential left turns)

= 0.15 * 0.15 * 0.15 * 0.15

= 0.00050625

= 1 vehicle in 1,975

probability of right loop

= 0.1 * 0.1 * 0.1 * 0.1

= 0.0001

= 1 vehicle in 10,000

Example 2:

P_L = 25%
P_R = 20%

P_A = 55%

P(left loop) = 0.25^4 = 0.00390625 = 1 vehicle in 256

P(right loop) = 0.2^4 = 0.0016 = 1 vehicle in 625

Trips

Within Commuter, the Trips component contains a specification of all the trips that will use, or try to use, the capacity provided by the network. There are several tables of trips within the Trips component:

- Undirected person-trips (Origin only)
- Undirected vehicle-trips (Origin only)
- Undirected freight-trips (Origin only)
- Directed person-trips (Origin-Destination)
- Directed vehicle-trips (Origin-Destination)
- Directed freight-trips (Origin-Destination)
- Transport vehicle-trips (Service Route)

The Trip components are used to generate Trips to simulate on the model network.

Every trip contains the following fields:

- A unique name
- A departure time
- A person, freight or vehicle type
- Origin and/or Destination, or Service
- A DNA string defining the decision-making behaviour of the agent making this trip

Network

The Network component defines the Walkways, roads, crossings, Areas, Zones and their related devices. The network is the supply side of the model, specifying the capacity. (The demand side specifies how many people and vehicles want to use that capacity.)

- **Areas** – an Area is a 3-dimensional shape defining an origin and/or destination for person-trips or freight-trips. It has a boundary defined by a set of points and a base and a ceiling. It can be designated as a person area (the default) or a freight area, but not both.
- **Bounds** – the boundary of the network.
- **Centres** – used to define circular arcs for Lanes or Walkways.
- **Controllers** – see Intersections.
- **Courses** – a connection between two or more Walkways.
- **Crossings** – a "special" Walkway that crosses a road.
- **Groups** – a signal group is defined as a set of turns.
- **Intersections** – an Intersection is a collection of one or more Nodes, and the road surface on those Nodes on which streams are defined. If an Intersection is signalised it has a Controller. Intersections are required only if traffic streams conflict.
- **Lanes** – a Lane is a surface on a road that can be used by vehicles. It can be any width, and vehicles can use any part of a Lane, not just its centreline. A Lane is on one Link only, and may connect directly to another Lane on the same Link, or directly to a Lane on another Link, or to another Lane via a stream or an Intersection.
- **Links** – a Link is the basic routing element between Nodes. A Link can hold any number of Lanes.
- **Loops** – a device on a Lane for detecting vehicles.
- **Nodes** – there is a Node at each end of every Link.
- **Sectors** – a Sector is a set of Zones or Areas.
- **Services** – a Service is used to define a public transport route, and is made up of a set of Trails.

- **Signal Displays** – a signal display holds information about the signalling equipment at an Intersection, for example whether a signalised approach has arrows or roundels.
- **Signs** – a Sign is a device beside or above a Lane that can be used to convey information to vehicles.
- **Stands** – a Stand is a device on a Lane that can represent a bus stop or a train station, or any other location where a public transport vehicle stops to pick up and set down passengers.
- **Streams** – a Stream is a connection between two Lanes over an Intersection. One Lane is the entry, or approach, and the other is the exit.
- **Trails** – a Trail is a sequence of Links or Walkways used to control or trace the routing of vehicles or people.
- **Turns** – a turn is a set of Streams which always have the same signal applied.
- **Walkways** - a Walkway is a path for pedestrians, and it can be either inside or outside a building. It is sometimes referred to as a channel.
- **Zones** – a Zone is is a 3-dimensional shape defining an origin and/or destination for vehicle-trips. It has a boundary defined by a set of points and a base and a ceiling.

Once the model network has been created,it is important to check the model before starting the calibration process. Surface colouring can be used to check lane index, lane and walkway speed and gradient etc. The route choice for all modelled origins and destinations can be checked using Area Routes and Zone Routes.

Validator: Calibration & Validation

In simulation modelling, calibration is the process of adjusting the parameter and demand inputs to a model, so that a set of measurements from the model - the “modelled values” - match a set of measurements taken on-site, in the real world – the “observed values” or “observations”. The calibration process is an iterative one: make model demand and network adjustments; check the model results against observed data; repeat until satisfied.

Validation is a subsequent process, checking that a second, independent set of modelled values matches the corresponding observed values. The use of a second set of measurements avoids the situation where the model is “over-fitted” to one data-set.

There is no “golden rule” for the correct data to use. This will be different for each model, depending on the data that already exists or can be collected for the site being modelled. Models may be calibrated and validated using pedestrian or vehicle counts, public transport boarding and alighting counts, travel times, queue lengths, levels of service, parking occupancy, or any other observable data source that can also be measured in the model.

The Validator tool can be used to compare several types of modelled values with observations. These types include:

- **Cordon Flows:** flows in and out of areas or zones
- **Link Flows:** flows counted for all lanes on a link
- **Turn Flows:** flows counted for an approach-exit pair of links at an intersection
- **Walk Flows:** directional flows at the exit of a walkway
- **Screen-Line Flows:** link flows summed over a screen line
- **Trail Flows:** flows counted for vehicles that complete all links in a sequence (a "Trail")
- **Trail Times:** travel times measured for all people or vehicles that complete all links in a sequence, in order
- **Traverse Times:** travel times measured for all people or vehicles that pass two points on the network, by any route.

The Validator tool shows arithmetic error difference (modelled-observed), error ratio (modelled / observed) and also GEH.

The GEH statistic is used to compare observed and modelled count data. It is a better judge of goodness of fit than error difference or error ratio alone, as the GEH takes into account the absolute size of a count.

The following classification is in general use for GEH:

- GEH less than 5: good match
- GEH between 5 and 10: partial match, could be better
- GEH more than 10: data not matching, not acceptable.

The Validator Tool is not a one-stop solution for verification of the model. It should be used in combination with other methods, some quantitative, some more qualitative:

- has all the demand been released into the model
- do the routes taken in the model look sensible
- do the parking patterns look realistic

Reporting

The Reporting component specifies when results should be saved, and what type of results should be saved. Result types include:

- **Person Trip Summary and/or Detail**
- **Freight Trip Summary and/or Detail**
- **Cycle Trip Summary and/or Detail**
- **Private Vehicle Trip Summary and/or Detail**
- **Public Transport Trip Summary and/or Detail**

By default, results will be saved at the end of the simulation term, but additional times for saving results can also be specified (Perhaps every hour during the model run).

Some tools may reset their running totals when saved: please see the help pages for each individual tool.

Results can be saved within the AZA model file, or externally as XLS or CSV. The advantage of saving the results internally is that the whole model, including results, can be transferred to a colleague or customer as a single file. The advantage of saving the results externally, especially for larger models, is that the model file is kept to a manageable size.

When saving the results externally, either CSV or XLS format can be selected. The advantage of using CSV (comma-separated value) format is that different runs can be compared using a text-comparison difference tool.

Advanced

User Guide

Assignment

Within Commuter, the Assignment component contains route and lane choice parameters, to allow the user to vary how people and vehicles are "assigned" to links or lanes in the network.

Varying these parameters allows for easier calibration, by moving people and vehicles in the model away from walkways or links that are above their observed counts and towards those which are under their observed counts.

When modelling parking, similar techniques can be used to direct parking vehicles towards particular parking zones, or set a preferred path through a car park.

As well as point-location decision rules, it is also possible to define rules to make people or vehicles follow a sequence of links. This technique is known as trail following. Trail following is useful for getting vehicles into the correct lane for approaching consecutive intersections. The rules can reflect the lane marking and road signs.

The assignment parameters can also include a set of route classes, which partition the network into a hierarchy - for example major routes, local routes, minor routes. This type of hierarchy can be used to improve route choice, for example increasing the cost of minor routes to reduce “rat-runs” or restricting certain users from certain road types (such as excluding trucks from minor routes).

Control

Within Commuter, the Control component defines objects that control traffic:

- **Phases** – a phase is a set of signal groups that run together. A signal group may apply to one or more vehicle movements, or to one or mode pedestrian crossings.
- **Plan** – a signal plan defines a set of phases, the order in which they are applied, the green time for each phase (the “splits”) and the offset of the first phase in the ordered list. The offset is measured against a global synchronisation time, normally the start time of the simulation term.
- **Signal Rules** – a signal rule defines an action and a logical test. The action is applied to a phase if the logical test evaluates to “true”. Example rule actions are "Skip", “Terminate” and “Extend”.
- **Playback** - a table of timing information allowing the simulation to "play back" a recorded sequence of signal group or phase changes. This can be a useful method of modelling the effect of an adaptive traffic control system, using a snapshot of its output. One of the advantages of using playback is that the each run in the model will be repeatable.

Display

Within Commuter, the Display component contains information about how the model should be displayed. In general, data stored in this component is for visual effects only, and will not affect the results of the simulation. (However, one exception is the structure, which can be configured as an obstacle. An obstacle structure will affect the trajectory of pedestrians moving in its vicinity.)

The display contains objects such as:

- **Aspects:** a definition of which layers are switched on, and the colours for each layer. The Aspects are numbered 0-20
- **Fixtures:** Shapes for static objects, such as buildings, furniture, trees, etc, imported from 3DS format.
- **Labels:** Static annotation.
- **Symbols:** simple 3-D shapes, extruded polygons of a single colour.
- **Shapes:** for moving objects, to display animated characters.
- **Tags:** coloured annotation for people and vehicles.

Plugins

Within Commuter, the Plugins component contains information about any Plugins that have been enabled. A **Plugin** is a standalone utility that can be switched on for any model but is switched off by default. This has the advantage of keeping the software package lighter and quicker for simple models, as it loads code for more complex tasks only when it is required.

A Plugin can be added by any Commuter user by writing a software module that conforms to the Commuter Application Programming Interface.

The Plugins component provides a framework for data storage. Each Plugin can store data in this framework, and retrieve it when the model is next opened.

Matrix Estimation

The Matrix Estimation component of Commuter is used to create origin-destination matrix demand.

The minimum requirements for this process is that there must be some observed counts for a reasonably comprehensive coverage of the model, and there should be also be departure and arrival volumes for each area or zone.

The Matrix Estimation process starts with a prior matrix. This may be taken from an older or higher level model, or if that is not available then one can be created using the departure and arrival demands.

There are several stages to the Matrix Estimation process:

1. set the departure and arrival volumes
2. create the initial input demand matrix;
3. manually edit the demand matrix, if required.
4. run the input demand on model;
5. compare the modelled counts with the observed counts
6. redistribute trips between origins and destinations while retaining the departures and arrival volumes

This is an iterative process. Steps 3-6 are repeated until the observed and modelled counts are relatively close and sense checks have been performed on all origin-destination trips.

Steps 4-6 automate the redistribution process, and can be run several times in succession with the aim of converging to an acceptable solution.

Chapter 1

Tutorial

Main Window

When you start Commuter, the **Main Window** opens. It contains:

- the **Menu Bar** along the top of the window
- the **Left Pane** containing several **Tabs:** File, Action, Layer,...
- the **Right Pane** containing two **Panels:** Tables and Graphics

The Left Pane Tabs:

- **File:** to browse files, or the components within a file
- **Action:** to edit the model, by executing one of the actions
- **Layer:** modify the colour or visibility of features in graphical view
- **View:** move to a pre-defined viewpoint, or save a new one
- **Plugin:** access or enable one of the plugins or special tools, including ones you have created yourself.
- **Undo:** to reverse an editing change

The Right Pane Panels:

- **Tables:** a panel showing the tables of data that make up the model. This is for advanced use only.
- **Graphics:** a panel showing a three-dimensional view of the model, visualizing the roads, paths, control devices, and the people and traffic as they travel in the simulation.

  The Graphics panel contains the simulation **Toolbar**, which can be used to control the simulation - play, pause, rewind, etc.

For basic operation, let's first concentrate on the Tabs highlighted above, the **Action** and **Layer** Tabs on the left side, and the **Graphics** Panel.

Action Tab

An **Action** is something you do to edit the model. An **Object** is something in your model to which an action is applied, for example, a Lane, a Crossing or an Intersection.

Some Actions are **active** all the time, some become active only if you have an Object of the correct type selected. If an Action is active, it will be shown in black text, if it is inactive it will be greyed out.

There are some options on the Action Tab available from a small drop-down menu above the tab:

- Single Click / Double Click: single click is faster, but it is more difficult to set up key bindings
- View as Table / Tree: The table view shows only the currently available actions, the tree view shows all actions, with unavailable ones greyed out. IN tree view there are [+] and [-] buttons to expand and collapse parts of the tree
- Simple / Standard / Advanced: controls which actions are visible

Exercise: Using an Action to create a Road

- In the Action Tab, expand the **Road** section by clicking on it once
- Click on **New Road at Cursor** (either single-click or double-click)
- A window will pop-up asking you to confirm the directions of travel, number of lanes, lane width and speed limit. Change these value if you want, or use the defaults. Press **OK**
- A road will be added starting at the X and ending at the T. The **Out** direction is from the X to the T - away from you.

The road you have just created is automatically selected. You can tell it is selected because it is highlighted in a different colour. Only the **Out** direction is selected - if you chose to create a 2-directional road, the **In** direction will not be selected. Because a Road is selected, more actions in the **Road** section are now active. These actions include **Bisect, New Intersection** and **New Roundabout.**

Key Bindings for Actions ("Hot Keys")

Some actions have a code in square brackets after the description, for example **New Intersection [F7].** This **hot-key** code [F7] indicates that you can implement the action by pressing that key on the keyboard.

Exercise: Create an Intersection using a hot-key

- Select the road just created in the previous exercise
- Click on the Graphics panel (this is called giving it the focus)
- Press the function key F7
- A window titled **New Intersection** will pop up
- There is a pull-down selector that allows you to choose the intersection layout. Click on the selector and choose one of the alternatives. Press **OK**

(Now use your mouse scroll wheel to move to the next page of Help).
You can define your own hot-keys, to speed up the editing of a model.

Exercise: Defining your own hot-key for an Action

- In the **Action** Tab, click on the Road section to highlight the word Road. If the section is now closed, click it again so that you can see all the Actions, including **Connect**
- Using the arrow keys on the keyboard, move the focus down to the word Connect
- Press the C key on the keyboard
- The action should now say **Connect [C]**

Hot keys are effective only if an action is active, and the keyboard **focus** is in the graphics window. However, you can define or re-define the hot-key for an action when it is not active.

Right Mouse Menu: Actions can be accessed either from the Action Tab or via a menu raised on the graphics window if you press and release the right mouse button. In either case, the effect of the action is the same. The only difference is the presentation of the actions: the right mouse menu shows only the currently available actions; the Actions Tab shows all actions.

Exercise: Using the Right-Mouse Menu to call an Action

- Click on the Graphics panel to give it the focus
- Press **Escape [Esc]** to ensure you have no selections
- Click the Right Mouse to raise the Action menu
- Scroll down to **Walkway** and select **New Walkway at Cursor**
- Rename the Walkway, and change the width, or accept the defaults. Press **OK.**

Layer Tab

The Layer tab is a two-dimensional matrix of coloured toggle buttons

- named **Layers** are listed down the left hand side
- numbered **Aspects** (0-20) are displayed along the top.

A Layer is a group of objects of the same type, like Lanes or Signs or Crossings. Each Layer can be switched on or off in any Aspect by pressing the toggle button. You can also change the colour of any Layer in any Aspect by right-clicking the button.

Exercise: Define your own settings for Aspect 10

- Click and drag the divider to the right of the Layer Tab to make the Layer Tab bigger
- Change the selector at the top of the Tab to say **All**
- Click on the orange number in the table heading, and slide it right so that Aspect 10 is the active Aspect
- Click on the cell in column 10 opposite the Layer **Road Edges**
- Right-click on the same cell to raise the colour editor. Move the colour slider down to blue, and drag the small circle in the colour panel to from the top corner to the centre. Press **OK**
- Repeat for other Layers

You can define any combination of Layers and colours in each Aspect. Some solid objects can also be made transparent in the colour editor. By default

- Aspect 0 is a general purpose view
- Aspect 1 is for editing walkways and crossings
- Aspect 2 is for editing roads
- Aspect 3 is for intersections
- Aspect 4 is for routing and assignment

You can redefine these standard Aspects if required and define your own layers too. Your Aspect definitions will be saved with your model.

Aspect Import/Export: If you want to exchange Aspect definitions with other models, you can export the Aspect definitions to a file, by pressing the **Export Aspect Definitions...** button at the top left corner of the Layer Tab. In another model, you can import these layers using the adjacent **Import Aspect Definitions...** button.

Graphics - Navigation

The Graphics panel by default shows a black background with grey **Control Bars** around the edges and a blue **Cursor** in the centre. If the Cursor is not visible, switch it on in the Layer tab.

Cursor: the cursor has an cross **[X]** in the centre of the graphics panel and a vertical line rising to the a horizontal bar **[T]** in the centre of the top half of the window. We use the [X] and [T] points to position objects. There is also a small vertical line to the left of the [X] which is used when creating zones and areas.

Dimensions: both the vertical and horizontal bars can be used as rulers, showing dimensions. You can also switch on a measuring grid by selecting the option in the toolbar below the graphics window.

Control Bars: the standard control bars are **Tilt** on the left, **Bearing** on the top and **Zoom** on the right. On the bottom, the blue bar controls the rotation of the cursor, and the grey bar is used to rotate the current selection.

Navigation Method

Unfortunately, there is no single, universal navigation method for moving around a 3-D model. Over the years, several methods have been created by different applications, and the method you prefer may depend on the application you have used most, or most recently. In the top right hand corner, the navigation method selector allows you to switch between the Commuter method and several others. The alternatives available are:

- **Commuter**
- **CAD:** inverted scroll wheel zoom, middle-mouse pan, basic View-Cube functionality
- **Google:** inverted scroll wheel, left-mouse pan, middle-mouse rotate, right-mouse zoom

If you have a suggestion for another navigation method, please let us know.

Exercise: Become Familiar with the Navigation Controls (using the Commuter Method)

- Using the left mouse button, drag the mouse downwards in the **Zoom** Control Bar to set the Zoom height to about 100. The exact value is not important, but if you want to make fine changes, you can do this by holding down the [SHIFT] key
- Right-click and select **Area > New Area at Cursor.** If you are prompted to make the Areas feature visible, press Yes
- You will see a purple square in the centre of the screen
- On the Graphics Toolbar, change the Draw selector to **Solid**
- Use the **Zoom** Control bar again to zoom in and out. The value on the Zoom bar is the height of the viewpoint above the cursor
- Use the **Bearing** Control bar to rotate the view. You can double-click near the ends of the bar to rotate by 90 degrees
- Use the **Tilt** Control bar to tilt. Double-click on the bar to return the tilt angle to zero.
- Using the **Right Mouse** button, drag to **Pan** the view. You can also Pan using the arrow keys on the keyboard.

Switching Navigation Mode

A drop down menu within the main drawing pane allows switching the navigation mode to one that may be more familiar.

There are three alternative navigation modes provided with Commuter. These are:

- **CAD:** mouse button actions modified to provide CAD-style zooming and panning. A basic "view cube" allows manipulation of the view.
- **Google:** mouse button actions modified to provide GoogleMap-style zooming, panning and rotation.
- **Paramics**: mouse button actions modified to provide Paramics-style zooming and panning

Graphics – Selection

Commuter incorporates 3D line-of-sight selection of any object. You can select a single object by clicking the mouse, select multiple objects at a time by dragging out a box, or add or remove objects to the current selection using various selection gestures.

Clear current selection: at any time you can clear all current selections by pressing the **ESC key**. You can also clear selections by double clicking on the background of the model, where there are no objects

Selecting a single object, discarding other selections: you can select a single object by pointing at it and **double-click left-mouse** button. This will de-select any objects that were previously selected

Selecting a single object, keeping other selections: you can also select an object using **single-click middle-mouse** button. This will add the current object to the selection set. This gesture also acts as a toggle: if the object is currently selected this will de-select the object. An equivalent gesture which can also be used, and is perhaps more standard on Windows, but requires two hands is **[CTRL] + double-click left-mouse**

Selecting multiple objects: you can select any number of objects in one go using **drag middle-mouse** to drag out a selection box. All visible objects inside the box will be selected. You can also use drag left-mouse, but exercise caution with this gesture because if you have anything selected, it will move those objects, which may not be what you want.

Move selected objects horizontally: any number of objects can be dragged horizontally using **drag left-mouse** while the **Tilt** control is at a setting of less than 45 degrees

Move selected objects vertically: any number of objects can be dragged vertically using **drag left-mouse** while the Tilt control is at a setting of more than than 60 degrees

Rotate selected objects around [X]: any number of objects can be rotated in a horizontal plane around [X] using the Control bar at the bottom of the screen.

Exercise: Become Familiar with the Selection Controls

- Create a number of objects using
    - **Area > New Area at Cursor**
    - **Road > New Road at Cursor**
    - **Walkway > New Walkway at Cursor**
- Try selecting a single object and then moving it to a different position
- Try selecting multiple objects by dragging out a selection box, and then rotating them around the centre point
- Tilt the view to more than 60 degrees, then select an object and move it vertically
- Try using the action **Edit > Undo** to reverse any moves applied

Graphics – Handles

Commuter uses a mode-less editor. This means that any object can be selected at any time, you do not need to change to Road mode to select Roads, or Zone mode to select Zones. There are no modes.

To make it easier to select or move the objects you want, many objects have one or more selection handles. A handle is displayed using a small arrow, box, or other icon. There are several types of handle:

- **End:** a handle at the end of the centreline of a linear object such as a Walkway or Lane. The centreline feature for an object must be visible for these handles to be visible. For example, to see the End Handles for Lanes, both the Lane Centreline feature and the Handle – End feature must be visible.
- **Curve:** a handle controlling the curvature of a curved linear object, such as a Lane or Walkway. There are two curve handles, one at ⅓ of the length of the curve, the other at ⅔ of the length.
- **Width:** a handle controlling the width of one end of a linear object, such as a Lane or Walkway.
- **Angle:** a handle controlling the angle of the line drawn at the end of a linear object, such as a Lane or Walkway. This angle is for visualisation only, it does not affect the simulation
- **Other:** all other handles, not falling in to the categories above. These handles are used for the corners of three-dimensional shapes, such as zones and areas.

Exercise: Become Familiar with Handles

- Zoom to about 200
- Create a new multi-lane road using **Road > New Road at Cursor**
- Switch on the Feature **Lane Centrelines**
- Switch on the Features **Handle – End** and **Handle – Curve**
- Zoom in and pan until you can see the handles, then select a single handle
- Move the handle and notice how it moves that lane
- Use **Undo** ([CTRL]-Z) to move the end of the lane back
- Drag out a box to select all end points for the road, then move them all at the same time
- Use action **Geometry > Make Curvable**
- Move the end points again to change the curvature of the lanes
- Select one of the Curve Handles and change the curvature
- Continue experimenting with other objects and handles

Chapter 2

Tutorial

Model Specification

The basic steps in creating a model are:

1. **Collate the data you will use to run and validate the model**
2. **Create a zoning plan**
3. Create a network of roads and walkways
4. (optional) Add public transport services and timetables
5. (optional) Add parking lanes and zones
6. Add demand data, then generate trips
7. Validate, by running model and adjusting parameters
8. Collect results and report

This tutorial describes how to do the steps highlighted above: collating the data and creating a zoning plan

The data you will need is as follows:

- A scale map of the area you want to model
- A set of trip-end locations, to use as origins and destinations.
- Demand data, specified either as an origin-destination matrix, or for simpler models, as the volume from each origin.
- (Optional) Behavioural data, such as value of time information from stated preference surveys
- Population data: physical sizes of vehicles, etc
- Details of the network layout, both for roads and walkways.
- Public transport routes and timetables
- Validation data, such as counts, travel times and queue lengths

To create a zoning plan

- Import an ***overlay*** – a technical plan drawing or an aerial photograph. You will use the overlay like electronic tracing paper, as a guide to add areas and zones, and later to add roads, walkways and other parts of the network
- Draw the trip-end locations for person-trips as **Areas**
- Draw the trip-end locations for vehicle-trips as **Zones**

Collating The Data

This section covers the data that you will need before you begin to create a model in Commuter.

1. **Scale Map.** This might seem obvious, but without an accurately scaled map of the area you want to model, all calculations of distance, speed and travel time will be inaccurate. Scale map data is available in many electronic forms and the process of importing an overlay will be covered later in this tutorial.

2. **Trip Ends.** A set of trip-end locations will be used as the trip origins and destinations. These may be point locations on the boundary of the model area, or geographical areas within the model, such as workplaces, shopping centres, etc. Commuter has been designed to model person-trips, so in the ideal case, the trip ends should specify origins and destinations for **people,** not vehicles. On the boundary of the model, it is possible to create single-occupant vehicles instantly, by linking a person-origin to a vehicle-origin zone, so it is best to create one Area and one Zone for each entry and exit point to the model.

3. **Demand Data.** This can be specified either as an origin-destination matrix (directed demand), or for simpler models, as the volume from each origin combined with turning counts at each junction of roads or walkways (undirected demand). In the ideal case, this data will all be specified in terms of person-trips. However, depending on the coverage of the person-trip data, you may also want to add additional “fill-in” vehicle demand.

4. **(Optional) Behavioural Data.** Commuter can be used as a dynamic mode choice model. Mode choice for each user group is controlled by a generalised cost equation that applies weights for time, distance and price. For example, if you wanted to model two groups labelled business and leisure, it is likely that the business group would have a higher value of time. This data is optional: it is possible to run the model using a single standard distribution of weights that is applied to all people within the model.

5. **Population Data.** For accurate modelling of road space usage, it is important to define a fleet of vehicles that reflects the distribution of size (particularly length) in your local region.

6. **Road and walkway network.** This includes any gradients and speed limits, intersection priorities and signal control timing, turn lane information, usage restrictions, pedestrian crossing types, and any other relevant information that may affect network performance. Much of the spatial data can often be gathered from aerial photographs. Other data, such as signal timing might require a site visit or collaboration with relevant authorities.

7. **Public Transport.** This includes all routes, the location of all stops, and timing information, particularly the departure time of each vehicle. Commuter can model rail and road transport; a vehicle can be made up of a leader and any number of trailers, so it is possible to model multi-carriage heavy rail, light rail and trams as well as buses. A Stand is the name given to the location where a public transport vehicle stops: this can represent a bus-stop, or a train platform, or a ferry wharf.

8. **Validation data.** Typical validation data includes mid-link counts, turning counts, travel times and queue lengths. This data allows you to check that the model is an accurate representation of the real world.

Importing an Overlay Map or Aerial Photograph

It is recommended that you import a scale drawing map or aerial image from a suitable file as an electronic “tracing-paper” overlay when creating your model. This makes it easier to verify that your roads and walkways are the correct size and position.

Suitable overlay sources, in order of preference:

1. CAD drawing, either 2D or 3D, in DWG format
2. Auto-scaled aerial photograph from internet map server
3. Manually scaled aerial photograph

Importing CAD drawing in DWG format

DWG is the recommended **first choice** for overlay. This format often contains multiple layers: road edges, walkways, markings, signs, etc.

Advantages of DWG

- precise, complete coverage
- automatically scaled to the correct size
- multiple layers can be coloured or switched on or off

Exercise: Import a DWG image

- Select **Display - DWG Drawing**
- In the **DWG Drawing** window, select [+] and browse to open the required DWG file.
- The DWG drawing will now be displayed in the main graphics panel. In the Layer tab click on the selector at the top, and at the bottom of the pop-up list you should see the DWG file name
- Select the DWG file name to open the section showing all the drawing layers included in the DWG file.
- Click on the buttons in the Layer tab to turn individual DWG layers on or off. You can also change colour of each layer
- Make the “Handle- Other” Feature visible. If you zoom in, you should see the handle for the bottom left-hand corner of the DWG drawing. You can use this handle to move the DWG object
- You can also move or rotate the DWG object by changing the values in the DWG Drawing window

The DWG file may be large, and it is not incorporated within the .AZA Commuter file. If you send the model to a colleague or client, you should also send the DWG file, if you want it to be visible.

Importing Auto-scaled Aerial Photo

The **second choice** for an overlay is an automatically-scaled aerial photograph image.

Advantages of JPG

- automatically scaled to the correct size
- photographic detail

Disadvantages of JPG

- all information in a single layer: “all-or-nothing” overlay
- tall buildings or bridges can obscure road details
- no access to walkways inside buildings

Exercise: Import an aerial photo image from the internet

- Action **Display - New Map at Cursor**
- Enter main street and cross street
- Enter, suburb, state, country
- Modify map size as required. The smaller the map size, the higher the detail in the photograph
- Press OK, and wait a few moments. The process can take up to a minute, depending on your internet access speed.

Importing Manually-scaled Aerial Photograph

The **third choice** for an overlay is a manually-scaled aerial photograph image. The image format is not important - most common image formats are supported by Commuter, including JPG, GIF, TIF, PNG and BMP.

Advantages

- quick and relatively easy

Disadvantages

- Single feature layer – “all on” or “all off”
- For a large area, an image that provides enough detail when zoomed in will be a large file, requiring a lot of memory
- Licensing: “Free” images from some internet sources are not usually free.
- Scaling inaccuracy – the “Scale Legend” included on images from internet mapping sites are usually for guidance only
- Parallax – an aerial view is pieced together from multiple images. There is only one point in each image that is directly underneath the camera, all other points have parallax errors.
- Obstructions – Tall Buildings can obscure roads and walkways

If you have no CAD drawings in DWG format, and you want to use a manually scaled image as an overlay, you should ensure that:

- The image contains something of known size that you can use as a scale
- Use a scale map to measure the distance between two known points (for example intersection centres) and mark this distance on your image using an image editor
- Use images that are no larger than 5MB in size. Larger images will be accepted, but may severely degrade the performance of Commuter unless you are using a high specification graphics workstation
- It is better to take several small but high-resolution images of each intersection or other area of conflict and a larger-area low-resolution image of the model area rather then a single high-resolution image of the entire model area, as this will soak up resources and lead to an unresponsive system
- For image editing outside of Commuter we recommend http://irfanview.com

Exercise: Import a JPEG image

- From the **Display** Menu, select **Images / 3D**
- Select the **Images** Tab
- Press the **Import** button, second from the top on the left hand side
- Browse to the image file you want to import.
- Press OK to accept the default options of Maps / Horizontal / 1.0m
- Press OK in the 2D Images and 3D Structures window
- On the Graphics panel, use **Display > New Image at Cursor**
- Select the newly added image from the Image Data list, press OK. The Image should now be visible in the centre of the graphics panel. If not, check you have the Images Feature switched on in the Layer Tab
- Select the newly added image so that its border is highlighted
- Switch on the Grid using the toggle button below the graphics panel. Adjust the grid size to suit the measured distance on your image
- Right-Click the mouse in the graphics panel, select the action **Display > Scale Image**, then adjust the image scale to match the measured distance against the grid. You may need to rotate the image so that the measured distance is parallel to the grid lines
- Once the image is the correct scale, select the image by double clicking on it, select **Adjust** from the right-mouse menu, which will open the Adjust Parameters window. In this window, look for the column labelled **Fixed** and set this option to on (ticked). This will fix the image to prevent you from accidentally moving it in future.
- To select an image for unfixing. Hold down [SHIFT] and drag out a box around the bottom left hand corner of the image. This will select the image, even though the handles are invisible. Once selected you can use **Adjust** and then clear the **Fixed** option
- Another way of selecting a fixed image is to switch on the layer **Handles - Other** and then select the handle at the bottom left corner.

Creating A Zoning Plan

A zoning plan is a drawing on a map of all the origins and destinations for the trips that will be modelled. These trip-ends may be point locations on the boundary of the model area, or geographical areas within the model, such as workplaces, shopping centres, etc. Commuter has been designed to model person-trips, so in the ideal case, the trip-ends should specify origins and destinations for people, not vehicles.

There are several types of trip-end locations:

- **Areas** represent origins and destination for person trips
- **Parking zones** represent the locations where cars can park. These contain parking lanes, each parking lane is marked into bays which can be parallel to traffic flow or angled
- **Cordon zones** are used on the edge of a network to connect an area directly to a zone. This allows person trips to be transferred instantly to or from moving vehicles.
- **Vehicle zones** are used to add fill-in vehicle demand. Vehicle-only zones will generate vehicles that do not contain people. These can be used to add extra vehicles on roads with your model so that you have realistic traffic flow, and those vehicles which do contain people experience realistic delays, queues and travel times.

All of these locations are modelled in Commuter as a multi-sided box with vertical sides, that has a floor and a ceiling. This allows you to stack them vertically, for example, to create one Zone for each floor of a building

Exercise: Creating a Shaped Area

- Set Bearing to 0, Zoom to about 100, Tilt to 0, Draw mode to **Solid**
- In Layer Tab, make **Person Areas** and **Handle–Other** visible
- Action **Area > New Area at Cursor**
- Tilt to 45, and pan backwards to see the whole of the Area box
- Select the Area just added
- Action **Adjust.** Change Base to -2.0 and Height to 20.0. Press OK
- Make **Handle–Other** visible in Layer tab
- Select one of the handles on the floor of the box
- Action **Area > New Boundary Point.** Select new point and move
- Continue creating more points to create a suitable shape

If you want to model a person trip that includes driving, parking and then walking to the ultimate (person-trip) destination, but that person trip originates outside your model boundary, you will need generate person trips on the boundary of your network that are immediately transferred into a moving vehicle. To achieve this, you need to create an Area and a Zone, and link them together.

Exercise: Creating a Boundary Area linked to a Cordon Zone

- Set Bearing to 0, Zoom to about 100, Tilt to 0
- In the Layer Tab, make **Person Areas**, **Parking Zones**, **Transition Zones** and **Vehicle Zones** visible
- Action **Area > New Area at Cursor**
- Move the cursor to one side of the new Area
- Action **Zone > New Zone at Cursor**
- Select both the new Area and the new Zone
- Action **Area > Connect to Zone**. The Zone should change colour and a line will be drawn between the Area and the Zone

To create a parking zone, you must first create a road with a suitable lane, and then switch parking on for that lane. A parking zone can contain any number of roads with parking lanes. The roads can be one-way or two-way and the parking bays can be parallel or angled to the traffic flow.

Exercise: Creating a Parking Zone

- Set Bearing to 0, Zoom to about 100, Tilt to 0
- Action **Road > New Road at Cursor.** Select one-way, 2 lanes, OK
- Select the kerb-side lane
- Action **Lane > New Zone**
- Save the model to force recalculation of Lane in Zone tests
- Select the Zone and move it slightly to the kerb side, but still partly covering the kerb-side Lane
- Select the kerb-side Lane
- Action **Lane > Parking On.** Select On-street, OK.
- Accept default settings for bay size and angle. Press OK.
- In Layer Tab, make **Lane Marking** and **Lane Parking** visible
- Notice that a walkway has been added beside the parking lane. This will be used by pedestrians walking to and from parked vehicles.

Chapter 3

Tutorial

Network Creation

The basic steps in creating a model are:

1. Gather the data you will use to run and validate the model
2. Create a zoning plan
3. **Create a network of roads and walkways**
4. (optional) Add public transport services and timetables
5. (optional) Add parking lanes and zones
6. Add demand data, then generate trips
7. Validate, by running model and adjusting parameters
8. Collect results and report

This tutorial describes how to do the step highlighted above: creating the network of roads and walkways that is used for travel by the people and vehicles in your model.

To create a network:

- Add pathways, then connect them together
- Add roads and lanes, and connect them together
- Add pedestrian crossings on lanes where walkways cross roads
- Set up turning lanes using streams
- Set road priorities or traffic signals at intersections

Walkways & Crossings

The first stage in creating a network for pedestrians is to create some walkways. Walkways are bi-directional: all walkways allow two-way flow by default, although this can be modified – see **Delay**. All walkways have a centreline and a width, which may used either as a limit or a guideline – see **Walled**. The width may be uniform or vary linearly from start to end.

A **Standard Walkway (or "Channel")** is a walkway that does not cross a road. It can represent a marked linear area, such as a footpath or pavement, or it can represent the preferred route across an open area. A standard Walkway can be curved. A **Crossing** is a specialised walkway that crosses a road. It is always straight.

Walkway-related actions

Walkway > New Walkway at Cursor: Creates a new Walkway at the current cursor position, starting at the X point and finishing at the T point. By default, a window will be raised each time to allow you to set parameters for the Walkway.

> **Display this window before adding a new Walkway:** Normally this is on, but if you switch it off, then any new Walkway added will use the same parameters as for the previous one. You might do this if you are adding several Walkways of the same type. If you switch it off, you can get it back using **Walkway > New Walkway Parameters**

Walkway > New Connection: This action will connect together the two closest ends of any pair of selected walkways that are not already connected. If you have 3 or more walkways meeting at a point, you can quickly connect all of them together by dragging out a box to select all the walkways, and then call this action.

Lane > New Crossing: Creates a new Crossing. There are several types of crossing, described later in this section. Crossings are added by selecting one of the lanes that they cross, so before creating any crossings, you must first create a Road – see later sections

Adjust: You can use the Adjust action on a selected walkway to set the parameters of the walkway.

Walkway Parameters

Width: The width of the walkway. This can be used to reset the width to be uniform along the entire length of the walkway. However, if the width handles at either end of the walkway are moved, this value is not used.

Step Length: If set, the Walkway will be displayed as steps if it has a gradient. If there is no gradient, this has no effect.

Walled: If this is set, people will stay within the boundaries of the walkway. If this is not set (the default) the width acts more as a guideline than a rule: if traffic density is low, pedestrians will stay close to the centreline of the Walkway; if traffic density is high, they may walk outside the width of the Walkway to get past other pedestrians moving in either direction.

Delay CD/ DC: A Walkway is bi-directional, its ends are labelled C and D The field Delay CD is used to add a delay to traffic travelling in the direction CD. This delay will act like a tollbooth, delaying each person for the exact time specified, and will also be used in routing calculations. Enter **inf** to enter infinite delay, making the Walkway one-way.

Exercise: Create three new Walkways and connect them together

- Zoom to about 100m. Bearing 0. Press [ESC] to clear selections
- Change to a Layer where **Walk Edges** are visible
- Action **Walkway > New Walkway at Cursor.** In the pop-up window change Width to 4, and clear the tick box **Show this window every time**
- Rotate the view Bearing to about 120, and move the cursor to be about 10m away from the end of the Walkway just added
- Action **Walkway > New Walkway at Cursor**
- Rotate the view Bearing to about 240, and move the cursor to be about 10m away from the end of the Walkway just added
- Action **Walkway > New Walkway at Cursor**
- Drag out a box to select the ends of all 3 Walkways
- Action **Walkway > New Connection**
- In Layers, make **Walk Centrelines** and **Walk Connections** visible. Notice that 3 new connections have been added.
- Select a Walkway, then Action **Geometry > Make Curvable**
- In Layers, make **Handle-End** and **Handle-Curve** visible
- Move the handle to adjust the shape and position of the Walkway

Crossing Types

There are several types of crossing:

- **Signalised, Linked to Phase:** Use this type at a signalised intersection where the crossing is assigned to a signal group. You can use this type of crossing to simulate both "protected" crossings and "unprotected" crossings where turning vehicles must wait for pedestrians to cross before proceeding over crossing.
- **Signalised, Pedestrian Activated (Pelican):** Use this type of crossing at a mid-link position on a road, away from intersections. This type of crossing is commonly used in areas where there is high volume pedestrian traffic and medium to high vehicle traffic, such as outside schools, hospitals or other public buildings.
- **Unsignalised, Pedestrian Right-of-way (Zebra):** Use this type of crossing at any position on a road to model the situation where an arriving pedestrian will halt traffic immediately. This type of crossing is more common within areas of medium to high pedestrian traffic and low-medium vehicle traffic.
- **Unsignalised, Vehicle Right-of-way (Unmarked, suburban):** Use this type of crossing at any position on a road to model the situation where an arriving pedestrian will not halt traffic and must wait for a gap in traffic to cross. Any vehicle traffic arriving once the pedestrian has started crossing will stop until the pedestrian has crossed. This type of crossing is more common within areas of low to medium pedestrian traffic.

Roads: Links & Lanes

Conventionally, traffic networks are represented by points and lines, known as **nodes** and **links**. Commuter uses a node-link network for routing, but the surface that vehicles use is represented by a set of lanes, where **a lane can start and end anywhere, not just at a node**.

In Commuter, **Road = Link + Lanes**

A Road supplies traffic capacity for vehicles and consists of:

- **Link:** the routing component, joining two Nodes and holding travel time, distance and price information used for calculating routes
- **Lanes:** the drivable surface of a Road. Each Lane has its own set of parameters, including speed limit, width and restrictions

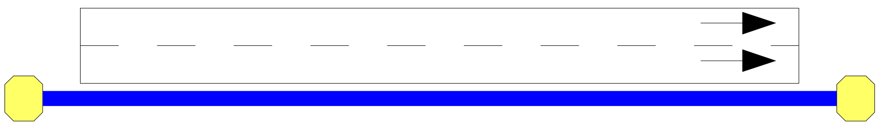

Road A-B: A simple uniform 2-lane road

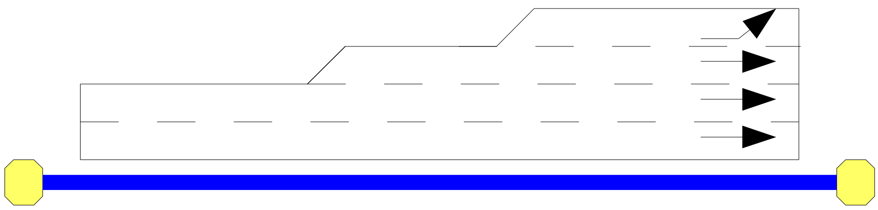

Road C-D: A road that flares from 2 lanes to 4

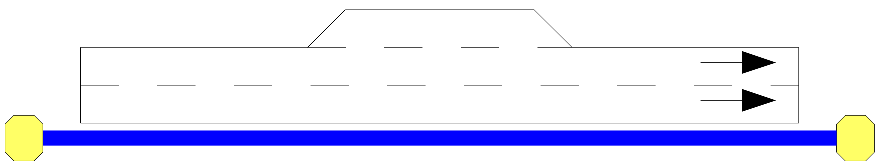

Road E-F: A road that widens to form a bus bay

All three of the roads in the diagram above are represented in Commuter by a **single link** between **two nodes**

A road defines a connection in one direction only.

A one-way road requires only one link

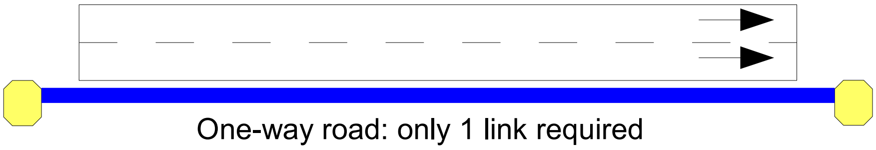

One-way road: only 1 link required

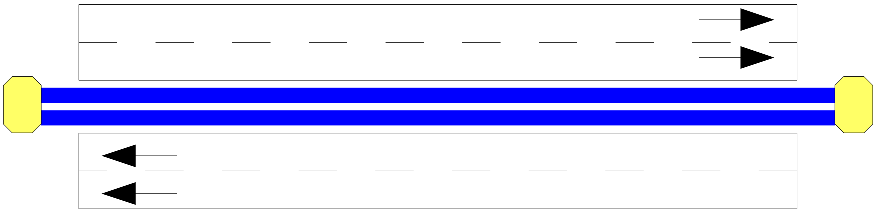

Two-way road: 2 links required

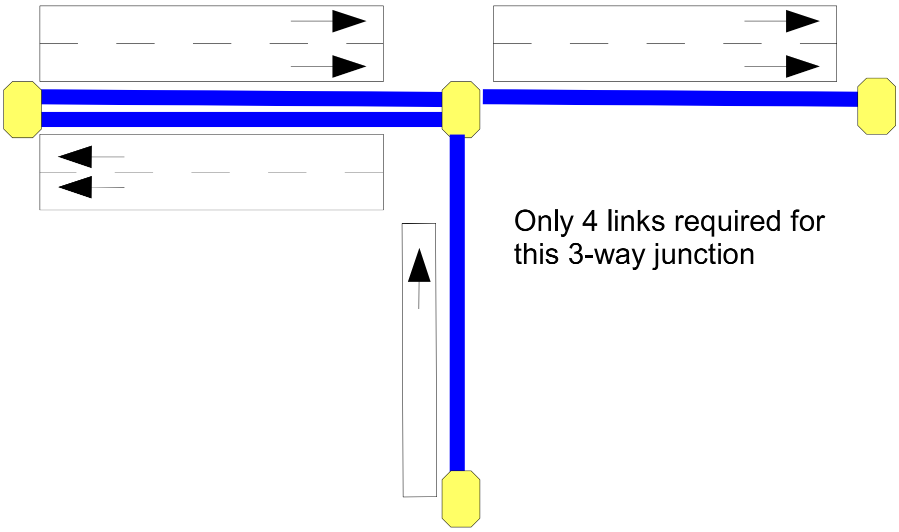

Only 4 links required for this 3-way junction

***** For fastest simulation, create as few links as possible. *****

Exercise: Create a city block with a 1-way circulatory system and 2-way roads entering and leaving

- Zoom to 500m
- Action **Road > New Road at Cursor**
  - Select 2 out-lanes and 2 in-lanes
- Action **Road > New Intersection** On pop-up window:
  - Select 4-way intersection
  - Set road to left to have 2 out-lanes and 2 in-lanes
  - Set road ahead to have 3 out-lanes and no in-lanes
  - Set road to right to have no out lanes and 3 in-lanes and change length to 100m. Other roads should be 200m
  - Press OK
- Select one of the lanes of the road created ahead
- Action **Road > New Intersection** On pop-up window:
- Set road to left to have 2 out-lanes and 2 in-lanes
  - Set road ahead to have 2 out-lanes and 2 in-lanes
  - Set road to right to have 3 out lanes and no in-lanes
  - All roads should be 200m in length
  - Press OK

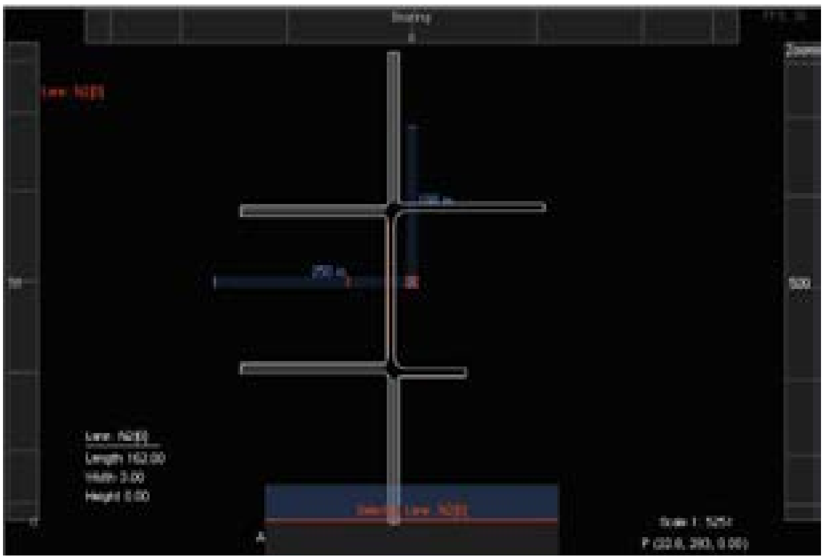

- Select one of the lanes of the road created to the right
- Action **Road > New Intersection**. On pop-up window:
    - Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 2 out-lanes and 2 in-lanes
    - Set road to right to have 3 out lanes and no in-lanes
    - Length of all roads should be 200m
    - Press OK
- Select one of the lanes of the road created to the right
- Action **Road > New Intersection**. On pop-up window:
    - Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 2 out-lanes and 2 in-lanes
    - Set road to right to have 3 out lanes and no in-lanes and change length to 100m. Other roads should be 200m
    - Press OK
- Select the lanes of the 100m road to the right and the road beyond it on the other side of the gap

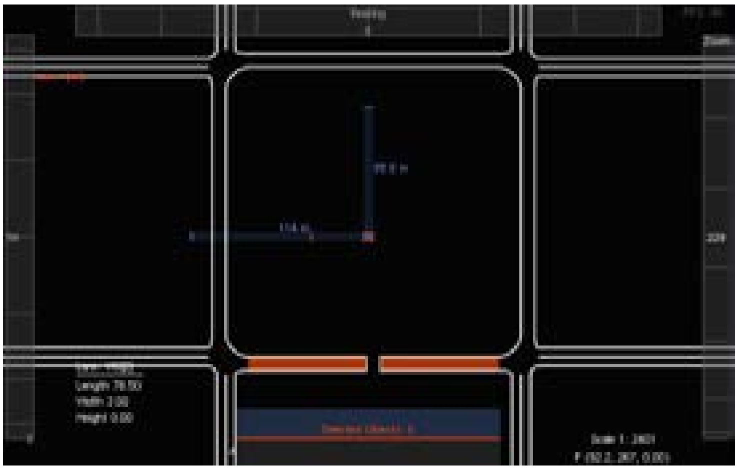

- Action **Road > Connect**
- Change Layer so that you can see Nodes and Links. If you have the default Layer configuration, then Layer 4 will show these objects. You should see that there is a node in the middle of the link, which is not required
- Action **Road > Combine.** There will now be a single link running along the bottom of the square block.
- **Save the model**

Roads: Creating New

The first stage in creating a network for vehicles is to create some roads, and for this purpose there are several available actions.

Road > New Road at Cursor: This is the simplest action for creating a new Road, and is available (active) if you have nothing selected. If this action is greyed out, press [ESC] to clear selections and it will become active. This action creates a new Road at the current cursor position, starting at the [X] point and finishing at the [T] point. When you call this action, a window will pop up to allow you to set parameters for the new Road:

- **Class:** the route class is used for routing calculations. We will get to this later, when discussing assignment. For the moment use the Standard class.
- **Direction (one-way / two-way):** The In and Out switches allow you to specify a one-way link away from the cursor (Out), or a one-way link towards the cursor (In) or a two-way link
- **Number of Lanes:** This field allows you to set the initial number of lanes that run the full length of the road, in each direction. You can add and remove lanes later, or move their end positions, but this field provides a quick way of getting you started.
- **Lane Width:** Each lane on a road can have its own width, and the width at the start can be different from the width at the end, but this field is a quick way to assign an initial width to all the lanes as they are created
- **Speed Limit:** This is the speed limit applied to the lanes created initially. The value in this field is taken from the Route Class, but you can modify it here.
- **Display this window before adding a new road:** Normally this is on, but if you switch it off, then any new road added will use the same parameters as for the previous road. If you are adding several roads of the same type, this can be a useful time-saver. If you switch it off, you can get it back using **Road > New Road Parameters**

The following actions require a road to be selected. You can select a road by selecting any of the Lanes on the Road, or by selecting the Link object.

Road > New Intersection: This creates a new Intersection at the end of the selected Road. The intersection can be 3-way or 4-way. A window is raised to set parameters for the new Intersection, which can be signalised, priority all-way Stop or a mini-roundabout. This action is intended for priority or signalized intersections. There are separate actions for creating free-flow intersections such as Merge or Diverge, which we will get to later.

Road > New Roundabout: Create a new roundabout at the end of the Road. The Roundabout can be 3-way or 4-way. A pop up window is raised to set parameters for the new roundabout and its arms, including inner radius, circulating lanes, spiral lanes, etc.

Road > New Road to T: This action is active if you have a Road selected and creates a new Road from the end of the selected Road to the [T]. This allows you to build outwards from a starting point, creating roads as you go. This action can be used when the selected Road is open-ended (no outward Roads) or if already has one or more outward Roads.

Road > New Road from X: This action is similar to the previous action but in reverse. It is active if you have a Road selected but in this case creates a new Road to the start of the selected Road from the [X]

Road > New Lane: Creates a new Lane to one side of the selected Road. A window is raised to allow you to select which side of the Road, and parameters for the new lane, including its length and position.

Exercise: (next page)

Exercise: Create a city block with a 1-way circulatory system and 2-way roads entering and leaving

- Zoom to 500m
- Action **Road > New Road at Cursor**
    - Select 2 out-lanes and 2 in-lanes
- Action **Road > New Intersection** On pop-up window:
    - Select 4-way intersection
    - Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 3 out-lanes and no in-lanes
    - Set road to right to have no out lanes and 3 in-lanes and change length to 100m. Other roads should be 200m
    - press OK
- Select one of the lanes of the road created ahead
- Action **Road > New Intersection** On pop-up window:
- Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 2 out-lanes and 2 in-lanes
    - Set road to right to have 3 out lanes and no in-lanes
    - All roads should be 200m in length
    - press OK

- Select one of the lanes of the road created to the right

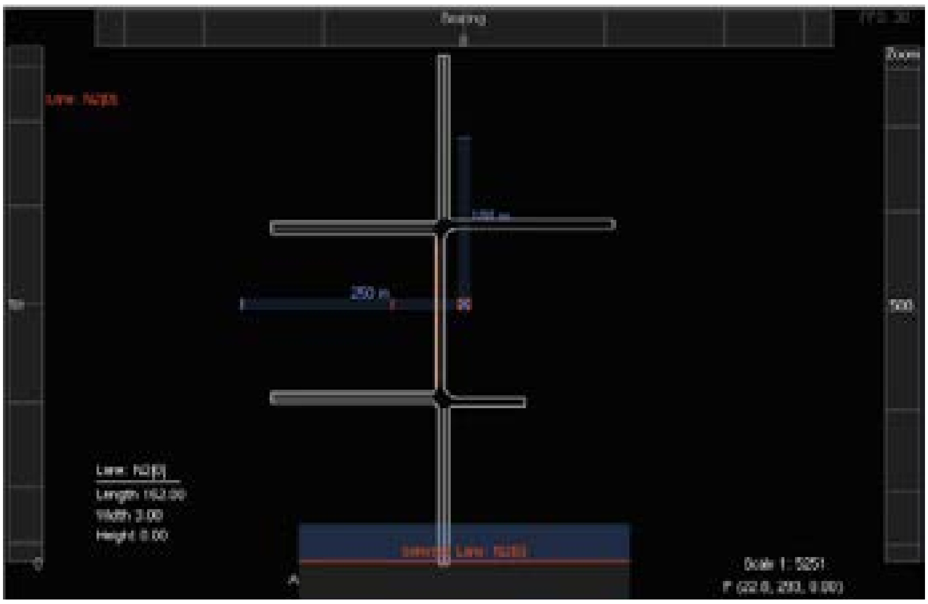

- Action **Road > New Intersection**. On pop-up window:
    - Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 2 out-lanes and 2 in-lanes
    - Set road to right to have 3 out lanes and no in-lanes
    - press OK
- Select one of the lanes of the road created to the right
- Action **Road > New Intersection**. On pop-up window:
    - Set road to left to have 2 out-lanes and 2 in-lanes
    - Set road ahead to have 2 out-lanes and 2 in-lanes
    - Set road to right to have 3 out lanes and no in-lanes and change length to 100m. Other roads should be 200m
    - press OK
- Select the lanes of the 100m road to the right and the road beyond it on the other side of the gap

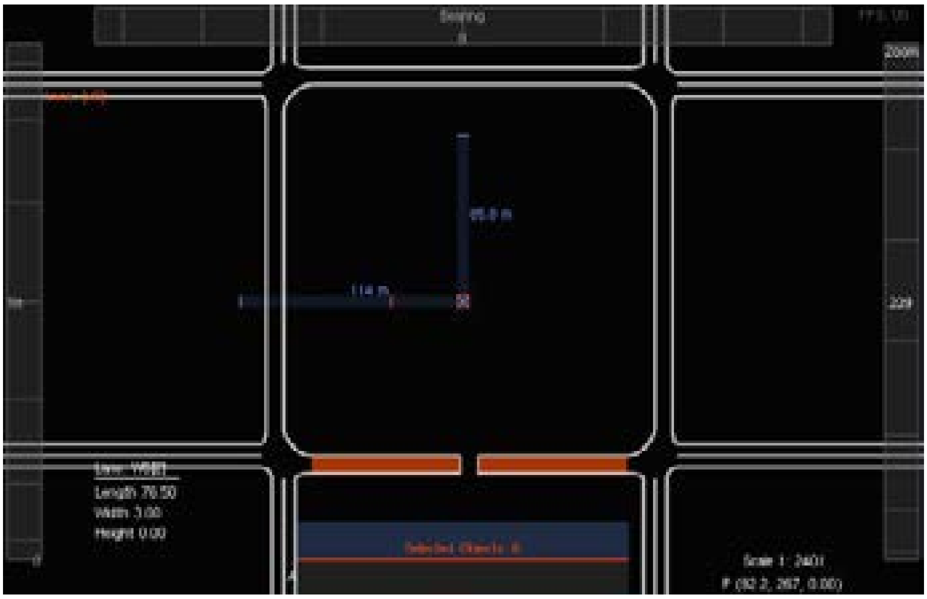

- Action **Road > Connect**
- Change Layer so that you can see Nodes and Links. If you have the default Layer configuration, then Layer 4 will show these objects. You should see that there is a node in the middle of the link, which is not required
- Action **Road > Combine.** There will now be a single link running along the bottom of the square block.
- **Save the model**

Roads – Connect, Divide & Combine

In this section we look at operation that allow you connect roads, and to add and remove nodes in existing roads.

Connecting Roads to Existing Nodes or Intersections

The standard technique we recommend for building networks is to start at a central point and build outwards, creating roads as you go. However, unless you are building a simple corridor model, there will come a time when you will need to connect a road you have created to a Node or Intersection that already exists. We will discuss later how a node differs from an intersection, but at this stage, for simple intersections, you can think of the two objects as being equivalent.

Road > Connect

(2 Roads selected) Connect the end of one selected Road to the start of the second selected Road, deleting the node that was at the end of the first Road.

(Road + Node selected) Connect the end of the selected Road to the selected Node, deleting the Node that was previously at the end of the selected Road.

Exercise: Connect a new road to existing Roads at an Intersection

- Action **Road > New Road at Cursor.** Choose a two-way road
- Action **Road > New Intersection** Select a T-junction shape with no road ahead
- Press [ESC] to clear selections
- Move the Cursor [X] to just above the head of the T-junction
- Action **Road > New Road at Cursor.** Choose a two-way road
- Select a Lane in the opposite direction on the road just created, that is, in the direction towards the intersection.
- Using middle-mouse button, select a lane in the road leaving the intersection down the leg of the T junction. You should now have two roads selected, and the Road > Connect action will be active
- Action **Road > Connect** will connect the road to the intersection, to create a 4-way intersection.
- Try more examples like this, but also try selecting a node and a road and connecting the road to the node.

Dividing (Bisecting) and Combining Roads

In Commuter, we use the term **Bisect** to describe the operation where you select Road A-B and **divide it into two equal parts** A-C and C-B. Bisect and Combine are complementary operations – one is the opposite of the other.

Road / Bisect: If you have selected Road A-B this adds a new node C to create A-C-B. There are several variations:

- **Bisect:** divide the road into two equal parts as described above
- **Bisect Diverge:** divide the selected Road into two equal parts, and add a one-way exit Road flowing out of the new central node at 45º
- **Bisect Merge:** divide the selected Road into two equal parts, and add a one-way entry Road flowing into the new central node at 45º
- **Bisect Side Road Left/Right:** divide the selected Road into 2 parts, and add a two-way Road connected to the new central node at 90º

Road / Combine: The opposite of Road / Bisect. If you have Roads A-C and C-B selected, Action Combine removes node C to leave you with A-B

We recommend using Road / Combine to remove extra dummy nodes whenever you can. This is good model building practice, to keep the network as simple and clean as possible, but it also increases the speed of simulation.
***** For fastest simulation, create as few nodes as possible. *****

Exercise: Bisect a Link and then Combine the two halves to remove the central node

- Action **Road > New Road at Cursor**
- Action **Road > Bisect.** When asked **Both Directions?** Press Yes. If you bisect a link in one direction only you will have a link A-C and C-B in one direction, but retain the original link B-A
- Press [ESC] to clear all selections
- Select one lane on the road leading to the new node just inserted
- Select one lane on the road leading away from the new node
- Action **Road > Combine**. This will remove the central node taking you back to the original configuration

Lanes

To recap, a Lane is the surface on which vehicles drive. A Road between Nodes A and B can have any number of Lanes, and the Lanes can start and stop anywhere. However, when a new Road is created, it has one or more Lanes along its entire length. Lets do some exercises to create roads that have lanes that are shorter than the Road.

Exercise: Create a Road with a Left and Right Turning Lane

- Zoom to about 200
- Action **Road > New Road at Cursor.** Select a 2-way road, with 2 lanes in each direction. Once created both lanes in the out direction will be selected
- Drag both lanes a small distance to the side away from the median by dragging with the mouse
- Select the kerb side lane only (use middle mouse button to deselect the median lane)
- Action **Lane > New Lane Beside.** In the pop up window, select the options **On Kerb Side** and **At End** and set Length to **25**
- Select the median lane only
- Action **Lane > New Lane Beside.** In the pop up window, select the options **On Median Side** and **At End** and set Length to **25**
- Make the lane centrelines and all handles are visible**.** In particular, display **Handles - Width** and **Handles - Angle**
- Increase the width of the kerb side lane at the end
- Move the start of the median side lane so it is a few metres longer
- Use the Angle handles to change the drawn angle of the start of both the median and kerb side turning lanes
- To improve the way the lanes are drawn, switch on the Feature **Lane Marking**, then select all Lanes on the Road and call action **Geometry > Make Curvable**

Exercise: Create a Road with a Bus Bay

- Zoom to about 200
- Action **Road > New Road at Cursor.** One-way road, 3 lanes
- Select the kerb side lane only
- Action **Lane > New Lane Beside.** In the pop up window, select the options **On Kerb Side** and **At Start** and set Length to **20**
- Select the newly created lane
- Drag it towards the end of the road
- To improve the way the lanes are drawn, switch on Lane Marking, and Action Geometry > Make Curvable for all Lanes on the Road

Connecting Lanes

Each end of a lane has one of three possible configurations
- Connected to another Lane
- Connected to an Intersection
- Not Connected

If a lane is connected directly to another lane, that means there is no conflict of traffic streams and traffic can flow freely. For best results, we recommend the use of direct Lane to Lane connections wherever possible. They are especially effective at merge and diverge points.

Lane / Connect: This connects the end of one lane directly to the start of another. This action is possible if two lanes are selected, and they are both on the same road, or they are on different roads but their roads share a common node. If the Roads are not connected use Road > Connect first.

Exercise: Connect two Roads then connect two Lanes
- Action **Road > New Road at Cursor.** Select Out only, 2 out-lanes
- Move the cursor to a position beyond the end of the new Road
- Action **Road > New Road at Cursor.** Select Out only, 2 out-lanes
- Select one lane on each of the new roads
- Change to a Layer where you can see the Road Edges and the Nodes and Links
- Action **Road > Connect.** Notice that the links are connected and you now have 3 nodes instead of 4.
- Lanes on two connected roads may not necessarily be connected. After the steps above, the lanes will have been connected on the roads just connected.
- Select two lanes, one on each road, Action **Lane > Disconnect**.
- A line will be drawn across the ends of the lanes, showing that they are not connected. This means that traffic cannot flow from one lane to the other. This might be the case if one is a parking lane for example.
- Reconnect the lanes using Action **Lane > Connect**. This connects a Lane on one Road directly to a Lane on the other Road

Dividing and Combining Lanes

Lane / Bisect: Replaces the selected lane with two lanes, each half as long as the original, with a direct lane to lane connection in the middle. Both lanes will be on the original link; no new Node is created. This is probably not an action you will require very often, but it is useful for some situations:

- if you want to create a short section where lane changing is not allowed. To achieve this, bisect a lane, then select one part of it, and the lane beside it, then action **Lane > Lane Changing Barred**
- if you want to create a curve that is more complex than is possible with four control points. Bisecting a lane into two parts gives you an object with 5 moveable control points rather than just 2

Lane / Combine: The opposite of Lane / Bisect. Replaces each pair of Lanes in the selection with a single Lane. Both Lanes must be on the same Link, and one Lane must be connected directly to the other.

Exercise: Bisect a Lane and then re-Combine the two halves

- Action **Road > New Road at Cursor.** Select 2 out-lanes
- Press [ESC] to clear all selections
- Select one lane on the road just added
- Action **Lane > Bisect.** Notice that you now have 3 lanes on the road, but there is no node at the central point
- Select both of the new lanes just created
- Action **Lane > Combine**. You should now be back to the original configuration

Adding Pedestrian Crossings

Where a walkway crosses a road, you should add a pedestrian crossing so that vehicles and pedestrians can “see” each other, and the crossing can be controlled. Crossings are described earlier, under “Walkways”.

Exercise: Add Crossing at the end of a 2-Lane Road

- Set Zoom to about 200. On toolbar set Draw mode to **Flat**
- Set Layer so that **Road Edges and Walk Edges** are visible
- Action **Road > New Road at Cursor**. Select 2-way, 2- lanes each
- Action **Lane > New Crossing**. In pop up, select **Position = Middle** and **Type = Unsignalised, Pedestrian Right of Way (Zebra)**
  In subsequent pop-up, change width to 3.0 and clear the tick box **Display this window...**
- Select the newly added Zebra Crossing and drag to the desired position.
- Press [ESC] to clear selections. Move the cursor so that it is about 5m from the end of the crossing
- Action **Walkway > New Walkway at Cursor.** In the pop-up window change Width to 4, and clear the tick box **Display this window...**
- Drag out a box to select the end the crossing and the end of the walkway
- Action **Walkway > New Connection**
- Continue, adding more walkways and crossings to create a walkway network.

Some types of pedestrian crossings have timing parameters that can be modified. For example, a signalised, pedestrian activated (pelican) crossing has parameters specifying the green time (WALK), the flashing green time (flashing DON'T WALK) and the red time (DON'T WALK).

Exercise: Crossing Times

- Open the exercise model. It is a four way intersection with a long arm towards the left.
- Background traffic has been prepared in this model. Run simulation and you should see traffic appearing.
- There are two areas and two walkways, one set on each side of the road.
- In order for people to walk between these two areas, we need to build a crossing first.

Nodes

Nodes are Route Decision Points – create as few as possible

In Commuter, a Node is a route decision point. For best operation, create as few nodes as possible. You don't need to create dummy nodes when a lane is gained or dropped. If you have extra dummy nodes created during building, you can remove them using **Road > Combine**

Exercise: Use Road > Combine to remove a dummy node

- Action **Road > New Road at Cursor**
- Action **Road > Bisect.** When asked **Both Directions?** Press Yes.
- Press [ESC] to clear all selections
- This gives you a section of road with a dummy node in the middle. There is no routing decision at the central node
- Select one lane on the road leading to the new node just inserted
- Select one lane on the road leading away from the new node, in the same direction of travel as the first road selected
- Action **Road > Combine**. This will remove the dummy central node

Node-First Network Building

When building a network, you can use a technique where you first create nodes at route decision points then join them together. Some people find this easier than the “standard” technique of building out from the centre.

Node > New Node: This adds a new node at the cursor. A circle is drawn around the node to show that it is not connected to any others

Node > New Road Between: This adds a new road between two selected nodes. You must select two nodes for this action to be active

Exercise: Create two Nodes then join them with a Road

- Action **Node > New Node.** By default you can use F5 as a hot-key
- Move the cursor to a new position
- Action **Node > New Node.** You should now see two nodes, each with a circle around them.
- Select both nodes. A quick way to do this is to drag out a box containing both nodes
- Action Node > New Road Between**.** The New Road window will pop up. Select the number of lanes, width, etc., that you require, and press OK

Other Node Actions

Node > Combine: This combines two selected Nodes to be a single Node, retaining all inward and outward roads. This is another way of joining a "hanging" link onto an existing node

Node > New Road From X: Create a new Road to the selected Node from a location under the cursor [X]. For best effect, ensure that the selected Node is not too near to the [X] otherwise the new Road created will be very short.

Node > New Road To T: Create a new Road from the selected Node to a location under the far end of the cursor [T]. For best effect, ensure that the selected Node is not under or near to [T], otherwise the new Road created will be very short.

Node > Reset Location: Move each selected node to a position at the end-median corner of one of its entry roads. This is useful when lanes have been moved by dragging and Nodes are no longer adjacent to the ends of the roads to which they are connected.

Intersection – Streams

An Intersection is an area on the road network where traffic movements conflict. If there are no conflicting traffic movements, then you should create direct Lane to Lane connections, and there is no need for an intersection. When we talk about conflict we mean both merging traffic and crossing traffic.

You do not create an Intersection directly: it is created automatically if you create a Stream. A Stream is a connection between two Lanes across an Intersection. Lets create an intersection, and then manipulate its Streams.

Exercise: Create a priority Intersection and modify its Streams

- Zoom to about 400. Set Bearing to zero.
- Action **Road > New Road at Cursor**. Two way, 2 lanes each way
- Action **Road > New Intersection.** Accept defaults, press OK
- In Layer Tab, display Road Edges, Lane Streams and Turn Arrows
- Zoom in to about 50 so that you can see the Streams on the intersection. Notice that by default, each approach has a shared left and through lane, and a shared right and through lane.
- Select the stream from the median lane on the southern approach that goes straight ahead to the northern exit.
- Action **Edit > Delete** (or press [DELETE] key). The median lane is now a right-only turn lane
- Select the kerb side lane on the southern approach and the kerb side lane on the eastern exit
- Action **Lane > New Stream**. The kerb side lane is now enabled for right turns as well as left and ahead.
- Select the Intersection by clicking on the intersection surface anywhere except on one of the stream lines.
- Action **Intersection > Turn Lanes.** This will raise the turn lanes window. Move it to one side so that you can see both it and the graphics panel. The Turn Lanes window shows a table with one row for each approach.
- Click in the first column on the name of each approach in turn. The view in the graphics panel will rotate automatically so that the approach you have selected is at the bottom of the view.
- In columns 3 and 4, try changing the turn-lane configuration for some of the approaches. Notice that streams are created or deleted automatically to match the selected turn lane configuration.
- If you create a conflicting configuration, a warning sign will appear, and you can press **Repair** to automatically repair these conflicts

There are several parameters you can set for each Stream to control detailed behaviour such as gap acceptance. For more information on these parameters, use the reference section **Actions /Adjust/ Streams.** One parameter we will cover here is the **Advance** distance. This is the distance beyond the normal stop line that a vehicle turning across an oncoming traffic stream will travel while waiting for a gap. This is sometimes referred to as the "stacking" distance.

Exercise: Setting the Advance (stacking) distance on a Stream

- Use the intersection from the previous exercise.
- Set Bearing to Zero
- Select the right-turning Stream from the southern approach
- Action **Adjust.** Set Advance field to 10, press OK. Notice that there is a broad band drawn at right angles across the stream, at 10m from the start.
- Make Feature **Handles – Curve** visible
- Drag the curve handles towards the centre of the intersection to make the stream travel forward in a straighter curve, until then turning sharply after the advance line.

Intersection – Nodes

An Intersection is an area on the road network that contains one or more Nodes, where traffic Streams can conflict. A Stream is a connection between two Lanes on an Intersection.

We looked earlier at the action **Road > New Intersection** which creates a simple Intersection containing a single Node. In many cases, a single-Node Intersection will be sufficient, but in some cases we need to create a construction with several nodes which acts as a single intersection. A roundabout is an example of a multi-Node intersection. A 4-way roundabout is a single intersection that contains 4 nodes.

Exercise: Create a 5-way roundabout Intersection

- Zoom to about 400. Set Bearing to zero.
- Action **Road > New Road at Cursor**. Two way, 2 lanes each way
- Action **Road > New Roundabout.** Select 4-way roundabout, inner radius 30m, 2 circulating lanes, not spiral
- Change to a Layer displaying Road Edges, Nodes and Links only
- Double-click on the roundabout surface just north of where the southern approach joins the roundabout. This should select the intersection. The name of the selected intersection will be displayed in the top left corner of the graphics panel. Notice that four nodes are highlighted each surrounded by a small circle. Also notice that all 4 conflict areas are highlighted on the roundabout, showing that these are all part of the same intersection.
- Select one of the lanes on the North-East quarter of the roundabout
- Action **Road > Bisect Side Road Left.** Select 2-way, 2 lanes each way. A new approach road will be added to the roundabout.
- Select the Intersection again, in the same way as before. Notice that there are still only 4 nodes highlighted. The new node just added is not highlighted as part of the intersection.
- Using the middle-mouse button, select the newly added node so that you now have that node and the intersection selected.
- Action **Intersection > Include Node.** The node will now be highlighted as part of the intersection, but it will be surrounded by a square, where the others have circles.
- Action **Adjust.** In the Adjust Parameters window, change the Form field of the Node from Priority to **Roundabout.** Press OK. The node will now be highlighted with a circle

Intersection – Signals

An Intersection is an area on the road network that contains one or more Nodes, where traffic Streams conflict or merge. Each stream on an intersection forms part of a turning movement, known in Commuter as a Turn. A Turn may contain more than one Stream. Each Turn may be signalised or unsignalised. If it is unsignalised it is assigned a fixed priority, such as Free Flow, Give Way, Stop, etc. In this section we will lok at signalised turns on an intersection. Note that an intersection marked as signalized may also contain unsignalised turns, for example, a slip or filter turn may have a fixed Give Way priority, rather than being controlled by a variable signal.

A signalised intersection has a Controller, a device that holds information about the signal control at an intersection. A Controller retains a record of all Turns, Groups, Phases, Plans and Rules for an intersection, and all of these can be edited using the Intersection / Controller Editor window. For more information, see the section Reference > Menus > Control > Controller.

Exercise: Basic intersection signal configuration:

- Zoom in to 500.
- Action Road> New Road at Cursor. Accept the default setting to create a two way road with two lanes each approach.
- With the road still selected, create a four way intersection using Action Road> New Intersection. Use the default settings.
- In the layer tab, turn on Lane Centrelines and Lane Streams so you can see all the possible movements.
- We will first build some background traffic to see how to intersection works without being signalized.
- Create a zone at each end of the road, you can create one and then copy and paste it to other arms.
- Then create a demand matrix with 1200 trips in total from these zones. Then generate these trips.
- Start simulation running and notice cars running as free flow, give way or yield.
- Now select the intersection and adjust it. Click on the Signalize button.
- You will see a number of new tabs are created. By default, group tab is the active one.

- There is already a signal group created for you. Let's add three more groups.
- You can have any number of groups; it depends on how you want to set up the phases.
- To keep this exercise simple, we will use four groups)
- After adding the groups, click on "Apply" and then go to the "Turns" tab.
- This tab shows all the possible movements. We will put them into groups . Later, we will associate groups with phases.
- Start from the through movements as these are usually the major ones. We will assign each of them with a group.
- It is likely that when the through movement is running, their associated left turn movement can run simultaneously.
- Therefore, we will assign the left turns with the same group as for the through movements.
- You can press the text of "Direct Group" to sort them by number.
- Now there should have four right turn movements remaining.
- Because our roads have only two lanes each approach and speed limit is 50km/h.
- And the relatively low demand we put in won't cause heavy traffic, therefore, cars turning right can filter through the opposing traffic.
- Rather than assigning a movement to a direct group, the turning movement is assigned to a filter group. Select the appropriate group for the movement to filter.
- Click on "Apply" and then switch to the Phases tab. For this simple intersection, we will create two phases.
- One phase is for North/South, the other for West/East movements. Left turns will run simultaneously and right turns will filter.
- In Phase A, set G1 and G2 (Group 1 and 2) to green signal so they can be activated in that phase.
- Assign green signal for G3 and G4 in Phase B. Then press "Apply".
- If you click on the Phase letter, you will see the movements in that phase highlighted in the graphic window.
- Solid lines indicate direct group, whereas dashed lines indicate filter movements.
- Now click on the Timing tab to set phase time. We will allocate 30 seconds green time for each phase.
- When finished, click OK to close the window.
- Start simulation running again, turns drawn in green indicate traffic

is controlled by signals.

- In the layer tab, turn on “Controllers”, this will display the controller board showing all the phases and timing.
- Make sure in the layer tab, “Handle-other” is turned on. Adjust its size by dragging the corner handle.
- Many intersections have detectors count traffic volume.
- These are often induction loops cut into the road surface, but video or other technologies may also be used.
- In Commuter, all detectors are referred to as “Loops” regardless of the technology used on site.
- Loops can be added for individual lanes, or for all approach lanes to an intersection by selecting the intersection then Intersection> New Loops at Stoplines.
- Notice the loop rectangles are drawn at all the approach lanes.
- Selecting the loop causes its name to be displayed in the upper left corner of the graphic window.
- Their names can also be displayed by turning on layer “Loop Names”.
- These loops will light up when occupied. Occupancy of loops can be used for vehicle-actuated signal control.
- For more information on vehicle-actuated signal control, please refer to Chapter 10, signal rules.

Chapter 4

Tutorial

Public Transport Services

The basic steps in creating a model are:

1. Collate the data you will use to run and validate the model
2. Create a zoning plan
3. Create a network of roads and walkways
4. **(optional) Add public transport services and timetables**
5. (optional) Add parking lanes and zones
6. Add demand data, then generate trips
7. Validate, by running model and adjusting parameters
8. Collect results and report

This tutorial describes how to do the step highlighted above: creating the public transport services. [While steps 4 and 5 are both marked as optional, we'd expect at least one of them to be present. If you have neither public transport services or parking then you don't have a multi-modal model, and Commuter might not be the best choice of tool.]

To create public transport services

- Add stands (stops) for public transport, to create a location that connects a walking path to the road network.
- Add sequences of roads called “trails”
- Add the route for each service built from a sequence of one or more trails
- Specify which stands along the route are used by the service
- Add a timetable for each public transport service to specify when each public transport trip will depart

Public Transport – Stands

A stand represents a location where a public transport vehicle stops to pick up or set down passengers. It is the location at which people change mode from walking to public transport, and vice versa. As a mode change location, it is a decision point in the multi-modal route choice network.

Some of the more notable parameters for a Stand include:

Size: The length of the platform or standing area for passengers. This is used in combination with the number of doors to calculate the position of the points at which groups of passengers gather while waiting for the next transport.

Doors: The number of doors expected on the next transport to arrive. This is used in conjunction with the size of the stand to calculate the position of the points at which groups of passengers gather while waiting for the next transport. Note that the type of the transport can vary by service, and even within the service.

Exercise: Create a public transport Stand and set its parameters

- Open the Model Four-City-Blocks
- Save the Model as Four-City-Blocks-Services
- Set Zoom to about 50. Draw mode to Flat
- Select the kerb side lane of a Road
- Action **Lane > New Stand**. A new Stand will be created, represented by a diamond shape with a line connecting to the lane
- Select the newly created stand. It will be highlighted
- Action **Adjust**
- Change Size to 5.0 and Doors to 2. Press OK
- Press [ESC] to clear selections
- Move the cursor to be 5m from the stand, away from the road
- Action **Walkway > New Walkway at Cursor**
- Select the new Walkway and the Stand
- Action **Walkway > New Connection to Stand**
- Continue creating stands and connecting them to walkways

Public Transport – Trails

A Trail is a sequence of route elements (links or walkways) used to control routing. With public transport, trails are used to define the route of each Service. This makes creating a Service route a two-level exercise:

- first, create a trail made up of two or more roads (links)
- second, create a service route made up of one or more Trails

For very simple networks, this two-level system adds some unnecessary complexity, but for most networks, the two-level system allows us to create multiple Services more quickly. It is often the case that in a model with multiple services, several will use a common sequence of links through the centre of the model. So we can define that common sequence of links once as a Trail, and then re-use that Trail in all of the Services.

Exercise: Add a new Trail

- Open the Model Four-City-Blocks-Services
- Select the Road where you want the Trail to start
- From the Network menu, open the Trails window
- Press **[+ New]** to start a new trail, and OK the default name
- The trail will be highlighted, and the view will move to the first decision point on the trail, and prompt you with a choice of colour
- The graphics panel will not be accessible, but you can zoom in and out using the Zoom control on the New Trail window
- Press one of the coloured buttons to select which of the exits you want to follow: red for the red exit, green for the green exit etc.
- The view in the main window will then move to the next decision point on the trail, and prompt you again with a choice of colour
- You can select “Advance Step = Single” to move along a trail by a single link at a time. The default setting “Advance Step = Multiple” always moves the end of the trail to the next decision point.
- Press the “Delete” (Red X) button to reverse back along the trail if you have gone too far.
- Continue pressing coloured buttons until you reach the end of the trail
- When you reach a dead end, you have the choice of accepting the trail as ending on the current dead-end link, or you can press delete to reverse back
- Repeat the exercise to create another two trails

Public Transport – Services

A Service defines the route taken by a public transport vehicle, which might be a bus, a tram or a train, and also defines the Stands, or locations where the vehicle will stop to pick up or set down passengers.

The Services window is accessed from the Network menu, and contains three panels:

- The Service list on the left, lists all services. Click on a row in this list to select a Service
- The Stand list in the centre, list all stands on the route of the service selected in the service list
- The Timetable on the right, listing all the departures for the service selected in the service list

Exercise: Add a new Service

- Open the Model Four-City-Blocks-Services
- From the Network menu, open the Trails window, and select the Trail you want to form the first part of the route for the service. The trail will be highlighted in the Graphics Panel – verify that it follows the links you want.
- From the Network menu, open the Services window and press **[+]** button ("New...") to start a new Service
- Select the Trail you just highlighted in the Trails window
- Enter a Name for the Service when prompted
- If any follow-on trails are defined which start immediately after the last link of the trail, you will be offered to choose between these follow-on Trails as the next part of the service.
- If there are no follow on trails, press OK to accept the end of the service at the end of the current Trail
- The new service will be added at the end of the list. Change the colour of the new service by clicking on the cell in the Colour column next to the service name
- The Stand list should show all stands that are passed by this service. If there are no stands in the central panel, you need to create some stands along the route and then return to this window
- under **Dwell,** put a tick against every stand at which you want this service to stop to pick up or set down passengers. Press Apply.
- Change the minimum dwell time at some of the stands.
- In the Timetable panel, press the [+] button and OK the defaults. This will create a set of departures at the stated interval. Press OK

Associated with each service are one or more transport trips, defining the departure time and initial occupancy of each vehicle on that service. A network model seldom covers the full extent of all services, so in many cases, the “departure time” is actually the time at which the transport vehicle first enters the network. This is why the transport trip defines an initial occupancy; obviously any public transport vehicle leaves the terminus empty, but from the point of view of the model, the departure point is the point at which the vehicle is first modelled, and it may already contain passengers at this point.

Chapter 5

Tutorial

Parking

The basic steps in creating a model are:

1. Collate the data you will use to run and validate the model
2. Create a zoning plan
3. Create a network of roads and walkways
4. (optional) Add public transport services and timetables
5. **(optional) Add parking lanes and zones**
6. Add demand data, then generate trips
7. Validate, by running model and adjusting parameters
8. Collect results and report

This tutorial describes how to do the step highlighted above: creating parking lanes and zones.

Parking zones represent the locations where people change mode from walking to driving and vice-versa. There are several types available:

- **On-street parking** allows vehicles to park on either side of the road, in numbered bays. There must be at least one non-parking lane for through traffic. This type of parking is used on links where there are through trips to other zones
- **Off-street parking** can be used to model dedicated parking facilities, with either one-way or two-way aisles. (Any link within an off-street parking zone will have zero cost for routing.)
- **Drop-off/ pick-up parking** allows vehicles to stop for a short time to allow passengers to get in or out of the vehicle.
- **Taxi rank parking** is for taxi type vehicles only, and allows only pick-up. It differs from other zones in that vehicles will park in a first-in, first-out order, and all will move forward when one departs
- **Instant**: Models a parking zone using a quick, crude method, where capacity is modelled, but bays and circulation are not.

Transition zones are used on the edge of a network to connect an area directly to a zone. This allows person trips to be transferred instantly to or from moving vehicles.

To create parking

- Create parking zones and transition zones
- Connect external areas to transition zones
- Add parking lanes inside parking zones
- Connect each parking lane to a walking pathway
- Connect walking pathways to areas

Creating a Parking Zone

To create a parking zone, you must first define the zone boundary, then mark at least one lane inside the zone as being a parking lane. The sequence of operations is as follows:

- Add a zone, defining the perimeter of the zone by moving and adding boundary points, as you would for any zone defining the location of a vehicle-trip origin or destination
- Make sure that the zone contains at least one lane. If you want the zone to be an origin and a destination, it must contain at least one inbound lane and one outbound lane.
- Select the lane (or lanes) you want configured for parking in the zone, and then select the action **Lane > Parking On**.
- A pop-up window will ask you what type of parking zone you want to have, from the types listed above.

Once the parking zone has been created, the following options are available from **Zone > Adjust**, on the **Zones** tab:

- **Parking Type:** On-street, Off-street, etc.
- **Price Base; Price/ Hour:** the price for parking in this zone as a sum of a fixed base price and an hourly rate, specified in the minor currency unit (pence, cents, etc.)
- **Expected Walk:** for large parking zones, an estimate of the walk time from the expected parking bay to the exit of the parking zone. That is, for large parking zones, where a driver expects the zone to be quite full, and for all the spaces near the exit to be taken, this parameter is an estimate of the time that will be required to walk from the parking bay that the driver expects to use. This parameter is set to zero by default, but with large parking zones, this may bias trips towards large parking zones which have one exit close to the destination, even though there is little chance of parking there.
- **Circles:** The number of times the driver should circle this parking zone, looking for a free bay. This might be used in a city centre, to generate circulating traffic. After the defined number of circles, the driver will divert to any other parking zone that has access to the ultimate destination, or if there is no other parking zone, the driver will return to the origin.
- **Wait** (true/ false): Set to true if a person must wait for a suitable vehicle to be available. If this is set to false, and no suitable vehicle is parked when a person arrives at the zone, a new vehicle will be generated, at a random location within the parking zone. [See also: Generation of Initial Vehicles}

(See also **Adjust / Zone**)

Once a lane inside a parking zone has been marked as a parking lane, the following options are available from **Lane > Adjust**, on a separate **Parking** tab:

- **Park Angle:** The angle of the parking bays (0, 45, 60 or 90) and whether face-in, face-out or both.
- **Bay Size:** The length (for angle = 0) or width (otherwise) of the parking bays. A vehicle will use a bay only if it will fit.
- **Park Time:** The time required to park or "unpark" from any bay on this lane. This value affects the time that a parking vehicle will obstruct other traffic in the through lane.
- **Transition (Time):** The time taken by the driver to get into or out of the car. This time will form part of the walking trip time. This value does **not** affect the time that a parking vehicle will obstruct other traffic in the through lane.
- **Dwell (Time):** For drop-off zones **only**, this is the time that a vehicle will park in order after the passenger gets out.
- **Popularity:** The popularity of this lane within the parking zone, expressed as a whole number. Higher numbers mean more popular. It is recommended that if you want N levels of popularity that you use numbers 1 to N. There is no upper limit on popularity, but the displayed key will be more readable if you use 1 to N. More popular lanes would normally be those that required a shorter walk to the destination. By default, lanes have popularity 0.
- **Best Bay:** The index number of the best bay on a lane. Similar to the popularity index, the best bay defines which bay on this lane would be chosen first, if all bays were empty. By default, the best bay value is zero, which means that all bays are equally popular. You can set this to a high value, such as 999, to ensure that the highest number bay is always selected even if the lane is stretched and more bays are added.
- **Cluster**: The name of the cluster to which this parking lane belongs. The parking cluster name is used by the Parking Directions plugin, to provide a way of displaying the number of remaining spaces in a cluster on a sign. The cluster is used only after a Sign has been added to the model, and the Sign associated with a Cluster in the Parking Direction plugin window.

Generation of Initial Vehicles in Parking Zone

If you want to generate an initial population of vehicles within a parking zone, you can do this using the undirected demand tab on the Demand Editor. Select **Add Volume** then **Mode = Vehicles – Parking**. A table will be displayed with one row for each parking zone. Enter in each row the number of vehicles you want to be generated in the first time step of simulation. The vehicles generated will be distributed randomly around the parking zone.

Behaviour Parameters Used Exclusively for Parking

- **Can Drive** (true/ false): If true, then a person has a car available and can drive to a parking zone.
- **Can Drop** (true/ false): If true, then a person has a car and a driver available and can be driven to a drop-off parking zone.
- **Must Pickup** (true/ false): If true, then a person has arranged a pick-up, and must use that mode of transport for the trip.
- **Parking Vehicles:** (Division): This defines a set of vehicles that a person will look for when selecting a car in a parking zone. For example, if you have two groups of people, staff and customers, you might want to direct staff towards a set of vehicles defined as staff-cars, and customers towards customer cars.
- **Park Duration** (Hours): The planned duration of the parking stay, used with the zone parking price parameters to calculate the cost of the parking, and consequently to steer the mode choice algorithm.

Behaviour Parameters Related to Parking

- **Walk Cost/ second; Walk Cost/ (distance-unit):** The cost of time and distance for walking. A high cost value will steer a person's choice away from a parking zone where a long walk is expected

- **Drive Cost/ second; Drive Cost/ (distance-unit):** The cost of time and distance for driving. A high cost value will steer a person's choice away from a parking zone where a long drive is expected

- **Off-Network Distance, Time and Price:** Each area can hold values for the off-network distance, time and price associated with the travel between the area and the ultimate origin or destination. These parameters are used in conjunction with behaviour-specific prices to calculate the cost of a journey by other modes, when deciding if parking will be used. For example, if the total cost of using a taxi is greater than the cost of parking, then parking will be used in favour of taking a taxi, and vice versa. The total cost of taking a taxi includes the taxi fare and any cost associated with walking between the origin and a taxi rank and walking between the the drop-off zone and the ultimate destination.

Parking – Areas and Zones

To add parking to your model, you will need up to 3 types of area or zone:

- **Areas** represent the origins and destination for person trips. Some may be shaped to represent internal trip-ends such as workplaces or shopping within your model boundary. Others may be simple rectangles representing external trip-ends
- **Parking zones** represent the locations where cars can park, and contain parking lanes, each parking lane is marked into bays
- **Cordon zones** are used on the edge of a network to connect an area directly to a zone. This allows person trips to be transferred instantly to or from moving vehicles.

Exercise: Create 2 Areas, Cordon Zone and Parking Zone

- Open the Model Four-City-Blocks
- Save as Four-City-Blocks-Parking
- Set Bearing to 0, Draw mode to **Solid**
- Select any eastbound lane of the central road on the west side
- Action **Lane > New Zone.** Then move the Zone so that it still covers the Road but does not protrude into the intersection
- Move the cursor to the west of the zone, zoom to about 40
- Action **Area > New Area at Cursor**
- Select the Area and the Zone.
- Action **Area > Connect to Zone.** The zone will change colour
- Clear selections, and move the cursor to a road in the centre of the network, and select a kerbside lane
- Action **Lane > New Lane Beside**. Select **Full Length, Kerb Side**
- Select the Lane just added
- Action **Lane > New Zone.** You can move the Zone a little to make it easier to select the lane, but the Zone must cover at least one kerb-side corner of the Lane just added
- Save the model to force recalculation of Lane in Zone tests
- Change Draw mode to Flat, and Select the Lane just added
- Action **Lane > Parking On.**
- Select On-street OK
- Select **Walkway Beside**, leave other settings as default
- Clear selections, and move the cursor to the end of the lane

- (continued)

- Zoom to about 30
- Action **Area > New Area at Cursor**
- Action **Walkway > New Walkway at Cursor**
- Select the Walkway just added and the Walkway beside the parking lane
- Action **Walkway > New Connection**
- From the Assignment Menu, open Person Routes
- Check that a route exists between the two areas

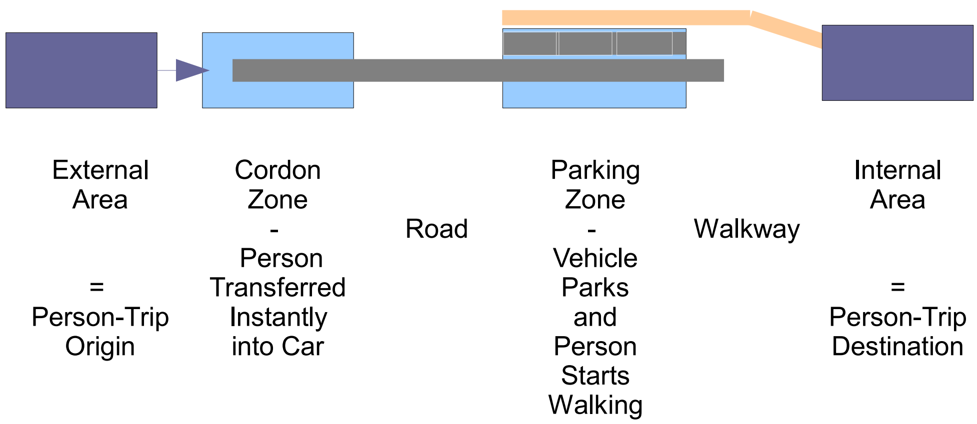

Parking – Lanes

Parking Lanes are special lanes that are marked with parking bays. You can mark a lane as a parking lane using action **Lane > Parking On**

A Parking Lane has the following attributes:

- it is contained within a Zone. Remember that a Zone has a floor and a ceiling, so if the Action **Lane > Parking On** is not active, it may be that the lane is below or above the zone
- it must allow lane changing to an adjacent moving traffic lane. Check that the lane boundary line is drawn as a dashed line, not a solid line, and that there is no gap between the parking lane and the moving traffic lane
- it must have no connections to other lanes at the start or the end
- it must have no stream at its entry or exit

The parameters for a parking lane can be modified by selecting the lane then action **Adjust:**

Angle: The parking bay angle, either 0 or 90 degrees.
Bay Size: The size of the bay in the direction of travel. That is, the length of each bay, for parallel parking, or the width of the bay for 90-degree parking.
Transition (time): The number of seconds required to occupy or unoccupy a parking bay; moving traffic is delayed during this time.
Dwell Time: (Applies to drop-off zones only). The time that a vehicle stays in a bay having dropped off a passenger. You can enter either a number of seconds, and all vehicles will dwell for that time, or you can enter the name of a Time Distribution.
Popularity: An integer value defining the popularity of all bays in this lane, where a higher number indicates a more popular lane. There is no upper limit, and the lower limit is zero. The choice of bay can be further refined using the Best Bay parameter.
Best Bay: An integer parameter indicating the bay that is the most popular within the lane. A value of zero (the default) indicates that all bays are equally popular. Popularity is assumed to decline at a linear rate away from the best bay. So if bay 7 is defined as the most popular, bays 6 and 8 will be the equal second most popular bay. Entering a value higher than the highest number bay indicates that the highest numbered bay is the most popular. It can be useful to enter 99 or 999 in the best bay field to make sure the highest numbered bay is the most popular, even after a lane has been stretched.

Chapter 6

Tutorial

Defining Demand

The basic steps in creating a model are:

1. Collate the data you will use to run and validate the model
2. Create a zoning plan
3. Create a network of roads and walkways
4. (optional) Add public transport services and timetables
5. (optional) Add parking lanes and zones
6. **Add demand data, then generate trips**
7. Validate, by running model and adjusting parameters
8. Collect results and report

This tutorial describes how to do the step highlighted above.

In simple terms, the demand is the number of people that want to travel.
The network of roads and pathways is the supply for this demand.
If demand is larger than supply, then there will be congestion.

Commuter is primarily a tool for modelling Person-Trips, but can also be used to analyse Vehicle-Trips. There are two "levels" of demand:

- Upper Level: Person-Trips are defined between Areas. Each trip defines a door to door journey for a person, which uses one or more modes of transport
- Lower Level: Vehicle-Trips are defined between Zones. Vehicle trips may exist to carry people on trips defined by the upper level of demand, or they may be generated as background traffic, or because the model is being used to study vehicle traffic only.

Within this two-level hierarchy, demand can be defined in three ways:

- Directed demand defines a number of trips, with each trip having an explicit origin and destination. Directed Demand can be used for Person-Trips or Vehicle-Trips.
- Undirected demand defines a number of trips from each origin, but no explicit destination. Undirected demand is often accompanied by turn counts (splits) at each junction. The turn counts define the probability of any trip taking each of the possible exits at a junction. Undirected Demand can be used for Person-Trips or Vehicle-Trips.
- Public Transport (Fixed route) demand defines public transport vehicle trips on the network. This type of demand is for Vehicle-Trips only. These trips use capacity on the lower-level road network, and supply capacity for person-trips to the upper level.

Before defining demand you must create Areas or Zones or both. This is described in an earlier chapter.
For an Area to be effective, one of the following must be true:

- The Area must contain at least one walkway OR
- The Area must be connected directly to a Zone OR
- The Area must be connected directly to a Stand

For a Zone to be effective, one of the following must be true:

- The Zone must contain at least one running Lane OR
- The Zone must contain at least one parking Lane OR
- The Zone must be directly connected to an Area

Exercise: Create an Origin-Demand Matrix of Person-Trips

- Create a Walkway
- Create two Area, one at each end of the Walkway
- Select the Walkway, Action **Walkway > Bisect** (A Walkway can only be associated with one area, this is why you need to bisect the walkway, so you have two walkways, one for each Area)
- File / Save
- Menu **Demand / Demands...**
- Press Add Matrix...
- Select **Mode=People, Profile=Flat and Division=D1**
- Change the name of the matrix if desired, and press OK
- The next window, Select Origins and Destination, allows you to use only some of the Areas as origins or destinations. Here we have only two areas, so we need both. Press OK.
- In the Demand Editor window, you will now have a grid, with column totals and row totals. Type 200 as the number of trips between Area 1 and Area 2, and 100 from Area 2 to Area 1. When you select a cell in the table, the origin and destination will be highlighted with an arrow on the Graphics panel.
- You can type values directly into the row total (all trips from this origin), the column total (all trips to this destination), or into the matrix total (all trips). The value you type will be distributed to each of the possible cells, in proportion to the existing values
- It is not possible to type values into the diagonal cells, as this represents intra-Area trips (from Area X to Area X).
- If you have Excel or another spreadsheet application on your computer, open that, and try copy-pasting parts of the table to the spreadsheet and vice-versa.

Generating Trips

Trip “generation” here refers to the process of creating individual disaggregated trips from aggregated demand specification. Trip generators employ pseudo-random number generators to select values from within parameter distributions. For example, a trip may be generated with a smaller than average physical size but a larger than average propensity to change route.

Commuter separates trip generation from simulation. That is, trips are generated once, before simulation begins, and the trip table is saved as input data. This means that the simulation can be repeated with exactly the same sequence of trips any number of times, even if the model changes. If you have a base model, and a proposed design model, it is often useful to know that changes in measurements resulting from a run of the design model are due to the changes imposed in the design, not due to variations in the sequence of random numbers generated.

The trips generated can be saved as part of the model, or can be saved in an external file. Saving externally is useful in larger models, where the number of trips is large (as a guide, 100,000 +).

Exercise: Generate Trips

- Menu Trips > Browse / Generate...
- Select Trip Generation Tab
- Change the “seed” for the random number generator, if required, or use the default
- Press Generate Trips
- Check that all trips you expected have been generated (Success Rate = 100%). If some trips are missing, it is likely that no route exists between one or more origin-destination pairs. Use Menu Assignment > [Person|Vehicle] Routes... then Route Analysis / Unavailable Routes to locate the problems.
- If all trips generated, press Yes to save
- You can now run the model by pressing Play on the Simulation toolbar

Demand Profiles

In Commuter, a Profile is a set of weights that sums to 100%, where each weight applies to a time interval. The weights are applied in sequence to generate a time varying quantity. The Term defines the start and end time. A profile is use to vary the number of trips generated over the course of a term. It is often lower at the start and end, and higher in the middle to represent a peak hour.

A profile can have any number of intervals; all of the time intervals are of the same length. If for example, you have an hourly demand, but want to weight the demand by some 5-minute counts, you would define a profile with 12 x 5-minute intervals, and set the corresponding 12 weights according to the count values.

Exercise: Demand Profiles

- Open the exercise model of a roundabout. It has four arms, each with two way traffic and two lanes each way. There is a zone located at each end of the arm to generate traffic flow.
- The first step of this tutorial is to add a demand matrix to this model.
- In the Profile section, there is only one option available, it is called "Flat". We will use this profile for now.
- Leave the default settings for the Origins and Destinations.
- Set a total demand, for example, 1200. Then save and generate trips.
- Start simulation and you should see traffic flowing.
- Notice the volume of traffic is rather consistent and evenly distributed.
- Take a look at the demand we created and we know this is because the "Flat" profile was chosen.
- Now let's create a new profile to simulate traffic flow increasing and decreasing over a period of time.
- In the "New Profile Parameters" window, we will choose "Simulation" as the Term, because we only want this profile effective during the simulation
- Set the number of intervals to 5. This will divide our simulation term into 5 parts, we will be able to control the intensity of demand in each part individually.

- The new window shows the control bars. Since our term is one hour and we have 5 intervals, each represents 12 mins.
- You can manually input exact value as well. The sum of these values should be 100, indicating 100%.
- Change the values to 10,20,40,20,10 so that we have a gradual increase to a peak flow followed by a decrease.
- After typing the value, tick the box on the right hand side, which will automatically adjust the totals so that they sum to 100%, and then click OK.
- In the demand window, change the profile from Flat to the one we just created, and re-generate trips.
- Start simulation and notice that traffic flow is light to begin with, but as time passes, the flow increases.
- High flow is particularly noticeable when simulation is around half way, which is the peak in the profile.
- In the demand window, you may have noticed that you can also choose a profile for demand from a certain origin.
- So that only traffic from certain origins will generate changing flow, while other origins generate consistent flow or vice versa.
- Profiles can be applied to all kinds of demands. Let's experiment applying the same profile to pedestrian movements.
- Firstly, create walkways next to the roads, you can do this by selecting the road and then use Lane> walkway beside.
- Turn on Walk Centrelines and Walk Connections plus Handle-End.
- Move the end of the walkway towards the roundabout, then create a perpendicular walkway and connect them.
- Create an area at the end of the perpendicular walkway; we will generate some pedestrian trips from here.
- Do the same at another adjacent arm.
- Now we will create a demand of 200 pedestrian trips between the two areas.
- In the profile section, choose the new one we created. When finished, save and generate trips.
- Run simulation, you should see people appearing and walking. Turn off Walk Centrelines and Walk Connections.
- If they are too hard to be seen, change them to a different colour. For example, fluorescent green.
- We can also add tags to people so they can be observed easier.
- Create a tag, choose a bright colour, and set origin to Area 1, this will tag people coming from Area 1.

- Create another tag, use the same colour, and set origin to Area 2 to tag people from the other way.
- In the Layer tab, turn on “Tags”. Run simulation and you should be able to observe pedestrian movements a lot easier.
- At this moment pedestrian movement is not very heavy.
- Turn off the Vehicles layer so we can focus on observing only pedestrians.
- With simulation running towards half way, notice the flow of pedestrians becomes heavier.
- And toward the end of the simulation, the flow declines as expected.

Route Classes

A route class is defined to encode parameters for weighting route costs. It may be used, for example, to bias routing decisions towards major routes. A route class contains only routing information, and is associated only with branches on the routing network, and not with any physical surfaces used for carrying agents.

Exercise: Route Classes

- Zoom to 300.
- Create a new road. In the parameters window, there is now only one class to be chosen. Continue to create the road with the default settings.
- Open the Route Class window from the Assignment menu, notice there are three items listed. But “Standard Road” is the only one associated with road,
- Create a new class by clicking the green “+” sign; you can change the description as required, e.g. change it to “highway”.
- Set surface to “Road” so that this becomes a route class for roads and then set the route speed, e.g. 110.
- You can also change the “Price Per Unit Distance” value, we will just leave it at as default.
- In the “Following” tab, there are a number of car-following algorithms available. This gives you flexibility, for example, to use one car following algorithm in urban streets, and another on highways. We will leave it to blank as default.
- Click OK to exit from the window and then create another road. This time, in the class cell, there are two options to be chosen from, Standard Road and highway.
- Notice the max speed automatically changes to 110 once you choose highway, click OK, and the new road is created.
- Select a lane from the road we created first and check its parameters, notice the speed and route speed are [50]. The “[]” around the speed value indicates it is the default value. You can manually change it.
- Select a lane from the new road, and check its parameters, notice the speed and route speed are set to [110], which is taken from the default highway route class.
- Select the route speed value and change this to another number, e.g. 90. Once you press enter or click Apply notice the brackets are gone.

Indicating it is no longer the default value.

- Change the route speed back to 110, and then switch to the Lanes tab, notice the speed may still be set as 90. Click on the Apply button to refresh it, it should now become 110, same as the route speed.
- Lane speed can be set higher or lower than the route speed; it is the actual speed limit of the lane.
- The route speed is used for assignment, the lane speed is used for vehicle motion.
- On a congested road, the effective route speed used for assignment may be less than the speed limit.
- On an uncongested road, the effective route speed may be higher than the posted speed limit.
- Experiment with changing the value, e.g. 120 or 90. You can also change it back to the default value of 110. When you are done, click OK to close the window.
- Now pan the view a bit so we have a blank space to create some new objects.
- This time, create a walkway, and notice in the window there is also only one Route class to choose. Create the walkway with the default settings.
- In the route class window, create a new class and give it a name, e.g. “new walkway”. Then choose “Walkway” as surface. Change the route speed to a suitable value. E.g. 6.
- Create a new walkway with the new route class. Leave all other parameters to default.
- Select the walkway we created first, and check its parameters, notice the route speed is 5.
- Check parameters for the new walkway, and notice it is 6, as defined for the new walkway route class.
- Route classes can also be created for rails. Experiment with creating a rail with the default setting.
- In the route class window, add a new class and assign Rail as surface, select a suitable speed, such as 70.
- Create another set of rail, in the parameters window, notice once you choose the new class, the speed changes.
- Select the old rail and check its parameters, the speed and route speed should be 60.
- Select the new rail and check its parameters, the speed and route speed should be 70.
- And the route class should be showing the name of the newly added rail route class.

Route Class Cost Factors

A route class is defined to encode parameters for weighting route costs. Each route class may have one or more cost factors associated with it. The cost factors for each route class can apply to an individual behaviour type, or to all behaviour types. Cost factors are multiplicative, so a cost factor of 1.0 has no effect.

Exercise: Route Class Cost Factor

- Open the model "Route_Cost_Factor_before.aza". It contains a road network linking two zones.
- The expressway has four lanes, there is also a single lane shortcut connecting the expressways. Traffic using the shortcut will diverge from the expressway and merge back to it at the end of the shortcut.
- The model has two behaviours (Parameters> Behaviours...), "Local" and "Tourist". Local people are familiar with the network and know the shortcut. While the tourists are not sure where the shortcut leads to and thus keep to the expressway.
- To make it easier to observe the effect, two vehicle types are also created (Parameters> Vehicles…), associated with the behaviours. The Local cars will be drawn in green, and the tourist cars will be drawn in red.
- The demand division (Demand> Demand Divisions…) setting shows the simulation includes 30% local cars and 70% tourist cars. And the OD matrix (Demand> Demands…) tells us there will be a total of 3000 trips going from zone 1 to zone 2.
- From the route class window, (Assignment> Route classes…) two route classes are defined for road; the "express" class has route speed of 90, and 50 for "standard" road class.
- Open the surface colouring window from the display menu, choose "RouteSpeed" for the Lane Colouring. Set the minimum and maximum value and lanes should be coloured differently based on their route speed. In this case, we will set them to 0 and 90.
- You should be able to tell all the expressways have the same route speed, and the shortcut is using a different route speed.

- Select a lane on the expressway and check its parameter (right click > adjust), route class indicates it is an express way with route speed of 90. Select the shortcut lane and check its parameter, it's under the standard road class with route speed of 50.
- In the surface colouring window, set lane Colouring to Normal, this will remove the colouring on all lanes.
- Start simulation running, and notice all vehicles are using the shortcut; no one is using the expressway.
- Open the "Vehicle Routes" window from the Assignment menu. Then select Zone 1 as origin and Zone 2 as destination.
- The route should be highlighted in the graphic window; it is using the shortcut rather than the expressway. And the cost to destination is 31.41.
- Let's temporarily close the shortcut by selecting the lane and then right click and choose adjust, in the parameters window, click on the "Lanes" tab, choose "Closed" for the "Closed" option and apply. Then open the "Vehicle Routes" window again.
- With Zone 1 still selected for origin and Zone 2 for destination, click on "Rebuild" button.
- Notice the shortcut is no longer highlighted. The expressway is now the preferred route and the cost is 39.42.
- This shows that using the expressway costs more than using the shortcut, that's why everyone is using the shortcut.
- Select the shortcut lane and adjust the parameter so it's no longer closed. (Leave the "Closed" option to blank)
- In real life, due to unfamiliarity of the local road network, tourists are likely to stay on the expressway. But local people are likely to choose the shortcut so they can get to destination quicker.
- In this case, we want to make tourist stay on the expressway. But allow local people to use the shortcut.
- We can do this by using the "Route Cost Factors" parameter. Open the window from the Assignment menu.

- Create a new factor by clicking on the green "+" sign.
- Choose "Standard Road" as the Route Class. We know from previous the shortcut is under "Standard Road".
- Choose "Tourist" for the behaviour since we want to keep them off the shortcut. When finished, click "OK".
- Type a value that's greater than the default, e.g. 10 for the Route Distance Factor, and then press "Apply".
- This causes greater cost for tourists to use the shortcut so that they will choose the expressway.
- Save the model and start simulation again.
- All the red cars (tourists) should start using the expressway and all the green cars (locals) should take the shortcut.
- Open the "Vehicle Route" window again and set zone 1 as origin and zone 2 as destination.
- When the Behaviour is set to "Local" notice the highlighted route uses the shortcut.
- Once switch the Behaviour to "Tourist", the expressway becomes the highlighted route.

Sectors

Sectors are groups of zones or areas. They are used for aggregating demand, for defining assignment and choice rules, for tagging, and for filtering in reports.

Exercise: Sectors

- Open sector_before.aza. It contains two networks; a road network on the left and a walkway network on the right.
- In the road network, there are four vehicle zones to generate traffic. Towards the bottom part there is also a parking zone.
- The walkway network has four areas to generate pedestrian trips.
- Check the demand and there are two OD matrices, one for “Traffic” and the other one for “Walking”.
- The traffic matrix has 25 trips between each pair of zones, 300 trips in total, whereas the walking matrix has 75 trips between each pair of areas, 900 trips in total.
- Open the sectors window from the Network menu, there are two sectors created by default, one for all areas and the other for all zones.
- To create a sector, first select the zones or areas. Let’s select the bottom zone, right zone and the left zone.
- Bring up the sector window, and click on the green “+”, you will see the selected zones listed on the left, click OK and a sector is now created.
- It also indicates all the zones that are included.
- If “Display by Colour” is ticked, the zones in sectors will be rendered using the sectors colour in the graphic window once you activate the graphic window.
- The default colour is yellow; you can click on the colour cell to choose a different one.
- To edit the sector, for example, exclude the left zone from it. First, click on the sector in the list, it selects all included zones. Then use the middle mouse button to click on the left zone, (Ctrl+left mouse button if using a two button mouse). Now you have two objects selected rather than three. Bring up the sector window and click on the edit button. Now two zones are listed on the left, click OK and the sector now includes only these two zones.
- A sector can also contain different kind of zones, for example, a

parking zone and a vehicle zone.

- Create a new sector includes the parking zone and the zone on the right. (Zone 5 + Zone 2)
- After created and if “Display by Colour” is ticked, you will see the zone on the right have a rectangle around it. This indicates this zone is included in two sectors. Bring up the sector window and we can confirm this.
- Now create a sector for areas, select the left and the bottom areas in the walkway network.
- And then create a sector, area 3 and 4 are included in the sector, change colour of the sector to green.
- Click “Apply”, in the graphic window, you should see those two areas drawn as green squares.
- Select the area on the right and add another sector. And then assign a bright blue colour to it.
- Now you are familiar with how to crate sectors, we will proceed to learn when and how to use them.
- For the next step, to keep it simple, we will keep sector 1 (zone 1 + 2), plus sector 3(area 3 + 4). Delete other sectors.
- The next thing we want to do is to use sectors in lane choice.
- We will target on vehicles coming from zone 1 and zone 2 going to zone 3.
- Normally they will drive on the kerb side lane since they will be turning left at the next intersection.
- For demonstration purpose we will force them to use the median lane using lane choice rules.
- First, let’s have a look how vehicles behave without the rules. To make observation easier, we will tag these vehicles.
- Create a tag (Menu> Display >Tags…) and set destination to zone 3, change the colour of the tag to something easier to see.
- Play simulation and notices them staying on the kerb side lane, which is correct.
- Now we will force them to use the median lane instead, by using lane choice rules. First, select a lane.
- Then right click and choose Road> Rules for Lane Choice. Set so that vehicles from zone 1 going to zone 3 take lane 2.Because we also want to force vehicles from zone 2 to use the median lane, therefore, we will create another lane choice rule.
- Now we have set two rules, save, rewind and start simulation again, notice vehicles are forced to take the median lane.
- If necessary, change the speed by using the real time speed

multiplier.

- We have created two lane choice rules because we want to control vehicles from two zones. This is ok but can be time consuming to create and manage if there are multiple zones.
- Since we created a sector containing these two zones, we can simply use that sector as origin, so we need only one lane choice rule.
- Remove these rules we added previously and then create a new one, this time, choose Sector 1 as origin and zone 3 as destination. Set lane 2 for the outcome lane range as previously.
- Rewind and start simulation again, notice our tagged cars still try to stay on the median lane, which is the same as with the two separate lane choice rules.
- Not only for lane choice rules, sectors can also be used along with other rules as well.
- Sectors can also be used with tags. Create a sector for zone 3 and name it “From Zone 3”and tag vehicles departing from this sector.
- In the tag window, create a new tag and set origin to the sector just created, in this case, “From Zone 3”, the cell is case sensitive.
- Run simulation and notice vehicles leaving this sector are tagged.
- This is useful when you want to tag vehicles from multiple origins. Similarly to the lane choice rule, you can tag a sector that includes a group of zones rather than create multiple tags.
- For the next part we no longer need the sector “From Zone 3” and the tag we used just now, so remove both of them.
- The next thing we will learn is to use sectors for report filtering.
- Open the private vehicle trips window from the Report menu, set origin to Sector 1. Speed to fastest and start simulation.
- After the simulation is finished, notice it showing completed distance of 55.34, trips departed 150 and arrived of 149.
- Rewind simulation and now change the origin sector to blank, run simulation from start to end again.
- Notice this time the departed volume is 300 with 296 arrivals. The completed distance is also a lot higher.
- Leaving the filter blank shows the result for the total trips, whereas setting the filter will only show filtered results.
- This works for person trips as well, open the person trips window and start simulation from start to end. When it stops notice there are 900 departed trips and 860 arrived trips. The total distance travelled is 223.41.
- Rewind and set origin to Sector 3 and run simulation again, now

both the departed and arrived volume is lower. The completed distance is lower as well.

- You can also report trips from certain origin to certain destination, too.
- From the sector list we know we already have one sector containing the lower and left area in the walkway network.
- Let's create another sector including the area at the top.
- In the person trips window, request the display of trips from sector 3 to sector 2 only by setting the filter for origin and destination and then play simulation from start to end, notice the volume for departed trips is 150.
- In the demand matrix we can see it is the sum of trips from Sector3 to Sector 2 = area 3 ->2 +area 4->2 = 75+75 =150.

Chapter 7

Tutorial

Validation

The basic steps in creating a model are:

1. Collate the data you will use to run and validate the model
2. Create a zoning plan
3. Create a network of roads and walkways
4. (optional) Add public transport services and timetables
5. (optional) Add parking lanes and zones
6. Add demand data, then generate trips
7. **Validate, by running model and adjusting parameters**
8. Collect results and report

The Validator tool displays comparisons of modelled and observed traffic counts and travel times, to help validate the operation of a model. To use this tool you will need observed pedestrian counts, traffic counts or travel times from your study area.

The Validator tool expects pedestrian or traffic counts to be associated with a Term.

Exercise: Add a set of traffic turn counts at an Intersection

- For this exercise, we will assume that the counts are for the period 08:00 to 09:00 and are for all vehicles
- In the Graphics Panel, create a four-way intersection, with a zone at the outward end of each approach arm. Create demand from these zones, and generate trips. For more information on these steps, refer to previous chapters.
- Menu **Parameters > Terms...** Check that a Term exists corresponding to the time period 08:00 to 09:00. If none exists, create one now
- In the Graphics Panel, drag out a box around the intersection, so that all entry and exit roads are selected
- Menu **Demand > Observations...** select **Turn Counts** tab
- Press the **[+] Add Turn Flow** button
- Select the Term created above and the **All Vehicles** Division
- Type a value into the **Observed** column. If the count values are held in a spreadsheet, and the turns are in the correct order, you can use standard copy and paste (Control-C and Control-V) to transfer values from the spreadsheet
- Save the model
- (continued)

- Open the Validator window. You may need to press the **Refresh** button to see newly added observations.
- Run the model, and notice how the value in the Count column increases by 1, every time a vehicle makes that turn
- The value in the **Normalised** column will be scaled up to match the Term
- The Diff field shows absolute difference
- The % field shows percentage difference
- The GEH field shows the GEH statistic. For a description of GEH see Wikipedia
- **Note that GEH is valid only for hourly flows.**

All the other types of observation can be used in a similar way. For more information see the Reference section under Validation.

Travel Times can also be used for validation, and these can be applied over Trails or Traverses. A trail is a sequence of roads or walkways, as used with roads to create public transport route

A Traverse is a pair of branches, the start branch and the end branch. Travel time is measured for all agents that use the start branch and then use the end branch some time later. The route that an agent takes between the start and the end is not important. An agent may be on any number of traverses at any one time. The travel time includes the time for both start and end branch. The start branch can be the same as the end branch, in which case a traverse time is the time to travel along a single branch. A Traverse can be thought of as an "in-gate" and an "out-gate", with no concept of where the vehicle goes in between, like a pair of CCTV cameras or license plate readers.

Exercise: Add a Traverse Time Validation for a Road Traverse

- In the Graphics Panel, select a lane on the road at the start of the traverse, and a lane on the road at the end.
- Open the Observations window
- Select Traverse Times tab
- Press the **[+] Add Traverse Time** button
- Select a Term and a Division
- Enter a value (in seconds) in the Observed column
- Save the model

Making Adjustments for Validation

It is unlikely that the validation tool will show perfect scores – all green lights on the GEH column – when you first run the model. If it does – well done – you are finished! If not however, read on.

There is no magic formula for validating a model, but in this section we will give you some pointers to things you can try. For all of the following adjustments, it is good practice to document your changes, giving justifications where necessary for modifying any parameter from its default, to make it easier for you or others to audit the model.

1. **Adjust the demand matrices**
   - In most cases, the demand data is the single biggest source of error in the model.
   - Demand data is often estimated from a combination of pattern matrices derived from old surveys, updated with traffic or pedestrian counts.
   - Demand estimation is an under-specified multi-variable optimisation problem. That means that in almost all networks (except for the very simplest) there will be many possible solution matrices that all generate the correct counts. There are even more solutions that generate the wrong counts!

2. **Check the routing in the network, if there are route choices.** This applies to both Roads and Walkways

   - Menu Assignment > Person Routes...
   - Select Route Viewer tab
   - Select an Origin and Destination
   - Examine the route. Using your knowledge of the network in the real world, does the route computed by the simulation look valid?
   - Change the **Route Spreading** slider display value. If

you get a more realistic spread of routes with a higher or lower value, you may need to adjust the route parameter specified for each Behaviour Type Group. (Menu Parameters > Behaviours > Routing).

- Change the **cost parameters** (value of time) for each Behaviour, to adjust the weight given to each mode selection (Menu Parameters > Behaviours > Walking Costs / Transport Costs / Driving Costs).
- Modify the **Route Length, Speed, Price** or **Cost Factor** for a Road or Walkway to increase or decrease its attractiveness.
- Add **Route Choice Rules** to force certain groups of people or vehicles to make a particular route choice, in locations where you know something that has not been captured by the model.

3. **Check the lane usage** (Roads only)
    - Compare per-lane Loop Measurements against observed values
    - Modify headway or reaction time factors on each lane, to compensate for areas where drivers will accept smaller or larger gaps
    - Modify the Lane Choice distances for each Road to control where lane choice decisions are made
    - Create Lane Choice Rules to force certain groups of vehicles to make a particular lane choice, in locations where you know something that has not been captured by the model.

4. Adjust the Physical Space for one or more Person Type
    - Smaller physical space may be accepted by people under certain conditions
5. Adjust the Driving Parameters for one or more Behaviour
    - Gap, headway and reaction time parameters can be changed for all roads on the network
6. Adjust the Calibration Parameters for Vehicles or Pedestrians
    - This is a last resort, as it will affect all Vehicles or all Pedestrians.

Making Adjustments for Validation

It is unlikely that the validation tool will show perfect scores – all green lights on the GEH column – when you first run the model. If it does – well done – you are finished! If not however, read on.

There is no magic formula for validating a model, but in this section we will give you some pointers to things you can try. For all of the following adjustments, it is good practice to document your changes, giving justifications where necessary for modifying any parameter from its default, to make it easier for you or others to audit the model.

1. **Adjust the demand matrices**
   - In most cases, the demand data is the single biggest source of error in the model.
   - Demand data is often estimated from a combination of pattern matrices derived from old surveys, updated with traffic or pedestrian counts.
   - Demand estimation is an under-specified multi-variable optimisation problem. That means that in almost all networks (except for the very simplest) there will be many possible solution matrices that all generate the correct counts. There are even more solutions that generate the wrong counts!

2. **Check the routing in the network, if there are route choices.** This applies to both Roads and Walkways

   - Menu Assignment > Person Routes...
   - Select Route Viewer tab
   - Select an Origin and Destination
   - Examine the route. Using your knowledge of the network in the real world, does the route computed by the simulation look valid?
   - Change the **Route Spreading** slider display value. If

Chapter 8

Tutorial

Collecting Results

The basic steps in creating a model are:

1. Collate the data you will use to run and validate the model
2. Create a zoning plan
3. Create a network of roads and walkways
4. (optional) Add public transport services and timetables
5. (optional) Add parking lanes and zones
6. Add demand data, then generate trips
7. Validate, by running model and adjusting parameters
8. **Collect results and report**

Usually, when a model is built, visualization is not sufficient: a set of numerical results is required to summarise the findings of the study.

Results collection is off by default for a new model. This is because results are not usually required until the model has been built and is running correctly.

Exercise: Collect Summary Results for Person Trips in Excel format

- Menu Reporting > Results...
- Select **Save as XLS files in Resources**
- In Table Selection, select **Person Summary**
- Run Model through to end of Simulation Term.
- Open Windows Explorer in the directory:
  (Model Path)/(Model Name)-Resources/Results/(Scenario)

If you choose not to save the results in XLS format, then they will be saved in the internal AZA format, within the model file. While this can be convenient for smaller models, in that you have only a single file to manage, it can become less convenient for models that generate large quantities of results, as the model file will be slow to open and save.

If you choose to save the results internally, you can export them in XLS format at any time by selecting a results page in the File tab on the left pane, and then pressing the Excel button on the Toolbar.

Generating Results by Scenario

Frequently, there are two or more options in a model to test or compare. Within Commuter, a Scenario can be created for each of these options. A Scenario is a short-cut to selecting all the relevant optional pages in a model. It is also used to name the results. Once a scenario is created, it is necessary only to select that scenario in order to select all the relevant input pages.

Exercise: Create two scenarios and generate results

- In the left panel, switch to the File Tab
- This Tab uses a "Tree" display. If a branch of the tree is closed, it has a small + button to the left of the name. You can click on this button, or double-click the name to open it. To close, double-click again or press the - button.
- Create a simple network, for example, a walkway between two Areas.
- Set demand and generate trips.
- Under the Components/Network section of the tree, select the Base network item, then press the green [+] button on the tool bar to create a new Derivative, call it Design1.
- Select the Design1 item in the tree, then right click and press the “Light Bulb” button to activate it.
- Change something in this design network so that it is different from the Base network.
- Create a derivative demand named Design1 and then generate trips, you will notice you have a new trips table listed.
- Activate the Design1 network, Design1 demand, and the newly generated trips table.
- Click on Scenarios and press the green [+] to create a new Scenario, name it Design1.
- Select the Base Network, Base demand and the old trips table and then create another scenario called Base1.
- Now switch between these scenarios and activate them, you will notice they are associated with different sets of components (network, demand and trips).
- Menu > Reporting > Results, select “Save in AZA file” and tick “Person Summary”.
- Keep Design1 Scenario active, run the model to completion, notice that the results page includes the name Design1.

- Rewind the simulation and you will notice the results page is no longer highlighted,
- Now right click on that choose “Export as XLS”, a window will open with the result in it.
- Activate the Base1 scenario and run simulation to completion, rewind and then export result.
- Compare the results in Excel.

Batch Runs

Once a base model has been built, and you are satisfied that it is an accurate representation of the real world, you may want to run it several times, with each run using a table of Trips generated from a unique seed. The objective of this type of testing is to examine the variability of results. If there is a low variability of results, this generally implies that the model is quite stable, and will be a good predictor of time, distance and other measurements for your designs.

To create and run a batch file:

- decide how many batch runs you require. There is no hard and fast rule for this, but we recommend at least 5. For complex models you may need many times more than this
- generate a trip table for each run, using a different seed for each. You type the seed value in to the trip generation window. The actual value of the trip seed does not really matter, but seed values over 10000 apparently produce better ("more random") results.
- it is a good idea to name the trip table based on the seed. So for example, if it is an AM model, and you are using seed 27293, then name the trip table AM-27293. This is just a convention, you can choose any name
- after creating each trip table, create a similarly named scenario, including the trips file.
- when all scenarios are created, right click on the scenarios "node" in the file tree, and select **Save Batch File.** This will save a windows batch file (.BAT) adjacent to the model file, containing commands that will run Commuter for each scenario defined. If you want to edit the batch file (for example delete some scenarios) you can edit it in a text editor.
- before starting the batch runs, make sure that you have selected which type of results you want to save, in the Reporting/Results window
- save the model file and close the Commuter window
- In Windows Explorer, locate the batch file, double-click to run it.
- a small text window will be raised showing the message log output, and the simulation clock time. When the simulation term is completed, this window will disappear. If you are running more than one scenario. The window will be raised again for each scenario in turn

Level of Service Reporting

Level of service (LoS) is a measurement of performance, commonly dividing all possible measures into bands, and assigning a single letter to each band. In the Los Tool, it is applied to delay on intersection approaches and also to delay, space and speed on walkways. There are many different definition of the bands; a common definition uses 6 bands, each indicated by a letter, and ranging from A (best) to F (worst).

Exercise: Level of Service

- Open the LoS_before.aza model, it has two intersections, with a pelican crossing in the middle of the long road that connects the intersections.
- The crossing allows people to walk between the two areas. Select the crossing and right click then choose adjust; The Crossing Time parameter indicates it's using the "Long" Timing.
- Open the Crossing Time window from the Parameters menu and we can see the "Long" timing has 50 seconds Walk, 5 seconds Flash and 100 seconds Don't Walk.
- Check the demand of this model, there are 150 pedestrian trips going each way using the crossing.
- And there are 2700 vehicle trips between all the 6 zones.
- Start simulation and zoom in to the crossing, people arriving wait until the signal turns to green to cross.
- Open the "Level of Service" window from the Reporting menu.
- This window is divided into two parts, the upper focuses on Intersections, and the lower focuses on Walkways.
- You can click the Commuter icon on the LoS window to make it stay on top.
- To observe level of service of an intersection, select the intersection in the graphics window, then press the green "+" in the LoS window.
- It will add all the approaches to that intersection for observation. Clicking on an approach in the list will high light it in the graphic window.
- Select the other intersection and add it to the level of service window.
- Once intersections are added, select one from the list in the LoS window and the view will pan to focus on that intersection.
- Now let's add the pelican crossing to the walkways list as well.
- Same as for intersections, select the crossing and click the green

"+" sign in the lower section of the LoS window to add it.

- You may have noticed there are a number of cells in the table highlighted in green, these indicate the LoS. Green indicates A, which is the highest.
- Start the simulation running and notice these cells change colours and LoS to represent the current performance.
- In this particular model, we can see intersection 3 has averagely low LoS after the model has been running for a while. The LoS is calculated based on the "Control Delay".
- And "Control Delay" is the sum of "Approach Delay" and Acceleration Delay".
- Now let's switch focus to intersection 2 by selecting it from the drop down list in the LoS window, and depends on progression of simulation, you may notice it has higher performance with shorter delay.
- This is because this intersection has only two phases, but a similar cycle time.
- This means there is less "dead time", also known as inter-green or all-stop time, so the intersection is more efficient.
- In real life, each intersection requires specific considerations to achieve the highest efficiency and safety.
- Another way to check intersection configuration is to turn on the controller. Controllers can be resized by dragging the handle. (Ensure handle-other is visible) They can also be rotated.
- Open the LoS window, there are a number of methods to choose for LoS marking, they can be chosen from the drop down list after clicking in the Bands cell.
- Click on the tab "Service Level Bands" next to "the Intersections" you will see how the band is calculated.
- Choose a different band from the drop down list to see their settings.
- In the LoS window, you can also tick "Extended" to show more statistics.
- And tick "Intersection Average" to show the average delay for all the approaches for the selected intersection.
- It is also possible to visually represent the LoS in the graphic window, to do so, tick Road Colours and Walk Colours.
- This will highlight the road and walkway added to the observation with corresponding colour to their LoS.

Economic Evaluation Tool

The Economic Evaluation Tool creates a report listing detailed costs for each trip, and a summary of all trips within the simulation term. This can be used to assess the Benefit-Cost Ration (BCR) of a proposed design, or to compare the BCR of two or more designs.

Economic Evaluation Exercise

- Open the FourCityBlock-Economic-Assessment.aza model in your Tutorial Chapter
- Before doing anything let us inspect some of the configuration of the model
- If we look at Parameters>Behaviours… there are a couple of different behaviours
- The main difference lies under the Driving Costs tab
- Where the Leisure behaviour has lower driving costs and the Business behaviour with a higher one
- In Parameters>Vehicle… we have various vehicle types with different behaviours applied to them
- This will help us get a more realistic data in our reports
- One more thing to look at, is the Demand Divisions window
- Go to Demand>Demand Divisions…
- In the Private Vehicles tab to see the list of Divisions that we have setup
- This will help with the report, as Commuter creates a sheet for each division that you have
- We are going to generate a report using the Economic Evaluation Assessment Tool
- The first thing to do is to go to Reporting>Results
- Move the checkbox to “Save as XLS files” or “Save as CSV files”
- If it is on Save No Results, Commuter will not generate any reports
- Close the Results window when done
- Now go to Reporting>Economic Evaluation…
- Here we have some editable fields and some other options
- More information about each option can be found in the help files
- For this exercise we will leave the defaults values in the Parameters tab
- Under the Services tab you should see that there are 3 services in

the model

- This is where you pick the services that you would like to assess for the evaluation
- For this exercise we are going to check Service1
- Press “Apply” then “Close”
- Set the speed to “Fastest” then press “Play”
- Allow the simulation to run the 1hr term until it stops, this should take less than a minute
- Once finished Commuter will generate a report with the selected options
- Go to your examples folder which should be located in C:\Aalient\Commuter\models\Examples…
- Find and open the Resources folder\Results\NoScenario@time-date\Economic-Measures-time-date.xls
- Here you have the Economic Evaluation Report
- Depending on what you want to measure, some fields might be more relevant that others
- There will always be an Economic Parameters tab…
- … An Assumptions & Definitions tab…
- … A tab for all public services…
- … And one tab for each division that you have in the model
- Each tab contains detailed information to help you evaluate the economic benefits of a proposed design

PT routes that have been selected (toggled) will be reported in output statistics in a grouping denoted by the letter “S” (marked as being part of the "Study" group).
All other routes will be reported in a grouping denoted by the letter “X” (not part of the "Study" group).

Chapter 9

Tutorial

Stairs and Escalators

If you change the elevation of one of the end points of a walkway, so that it is higher or lower than the other end, then you will create a sloping walkway. (As a reminder, you can change the elevation of an end point, or any handle, by tilting the display to more than 60 degrees, and dragging the point upwards or downwards. You can also use the Action **Geometry > Change Height**.)

If you want to create steps on a sloping walkway, select the walkway and then choose **Action > Adjust** and edit **Step Height**. Set the height of a step as appropriate, for example, you might choose 0.2 (metres, if units are metric). The walkway will now be drawn as steps.

If you use the Action **Geometry > Make Curvable**, you can create spiral stairwyas, or stairs with changing gradient. The path of the stair way can be curved in the vertical plane or the horizontal plane.

This Step Length parameter affects the visualisation only – it has no effect on the movement of people. You may want to change other parameters on the walkway, to reflect the nature of a stair way. For example, you might want to change the travel speed, or restrict its use to certain user types.

Once a walkway has been converted to a stair, you can make it an escalator by setting a forward speed in the parameter **Speed CD**. First, select the walkway, to identify the ends. One end will be labelled 'C', the other end will be labelled 'D'. If you want the escalator to move in direction CD, then set Speed CD to a positive value, if you want it to move in direction DC, set Speed CD to a negative value.

The **Speed CD** parameter is added to the walking speed for a pedestrian, so this will affect the time it takes for a person to travel on the escalator. Approximately half of all people will stand still on the escalator, while the other half of the population will try to keep walking. If a walking person reaches a standing person, and there is no room to pass, then the walking person will also stop.

Exercise: Stairs and Escalators:

- Stairs are walkways with steps and a slope.
- So start by creating a short walkway. To make it easier to observe, let’s zoom in to 100.
- Create a 5m wide walkway.
- Select the walkway, right click and choose Geometry> Change height. We will enter two different values to give it a slope.
- Enter 0, 10 so one end of the walkway sits at 0 and the other is 10 metres above.
- Notice if you type 10, 0 for the height rather than 0, 10 the slope will run in the opposite direction.
- Tilt and rotate the view and you will see now we have a sloping walkway that looks more like a ramp.
- We will now add steps to make it stairs, select the walkway (if not already), and right click then choose adjust.
- Look for a parameter named Step Height, type 0.15 for 15 cm, and then click Apply. The ramp instantly becomes a stair.
- Curvature can be applied to make non-straight staircases, such as semi-circular stairway.
- To do so, select the walkway, and then choose Geometry> Make Curvable.
- Ensure that the Handle-Curve and Walk Centrelines layer is visible; select the handle point to shape the stairs.
- When a staircase is applied with speed, it becomes an escalator.
- Select the stairway, choose Adjust, set Speed CD to 0.5 so that the stairs move at half a metre per second from C to D.
- To make the escalator run the other way, type a negative value for Speed CD, e.g. -0.5.
- The steps are not shown when the walkway is not selected; to show the steps all the time, make Walk Surfaces layer visible.

Lifts

To understand how a lift works in Commuter, you should first be familiar with public transport services. A lift is similar to a vertically travelling bus service – it has a number of Stands – one at each floor, and a single Transport vehicle – the Lift itself. However, there is no timetable for the lift: it moves up and down according to calls made on each of its floors.

To add a new lift, select Action **Walkway > New Lift at Cursor**. You will be prompted to enter some parameters:

- a description for the lift. This will be used in reporting
- size: width, depth and height. The lift is always square in horizontal cross-section
- maximum speed
- capacity
- inter-floor heights

Of these, only the inter-floor heights parameter needs much explanation, and it is explained on the pop-up window. The key thing to remember is that there should be one less height value than there are floors. If there are 3 floors serviced by the lift, then 2 inter-floor height values are required.

Once you have added the lift, there will be two stands at each floor. One stand is for lifts going down, and the other is for lifts going up. There are two stands on each floor, even the top and bottom, because for example, people going down and getting out at the bottom use one stand and people going in at the bottom and going up use the other. There will be short walkway attached to the stands on each floor. You can modify this walkway, but it is recommended that you do not try to attach more walkways directly to the stands.

Initially, you may need to create another method of access for travel between floors/areas, in order to create the trips. That is, create a stair or ramp connecting the floors, generate the trips, then remove the stairs if not required. This is because the lift is like a service, but initially has no "departure times" scheduled. These will be scheduled as soon as there are people that want to use the lift, but until there is a departure, the lift shaft looks to the trip generator like a service with no buses, so it will appear that the other floors are unreachable.

Exercises: Lifts:

- Open the model “Lift_Before.aza”, it has a group of stairs linking 3 levels, floor height of this model is 4 metres.
- There are two person types in this model; green ones go to Level 2 and red ones heading to Level 3.
- We will now create a lift so that people have an alternative route.
- In the view tab, click on “Start” to activate that view and press “o” key in the viewport to switch to orthographic view.
- Pan the view so that the vertical cursor points to the areas.
- Action> Walkway> New Lift at Cursor.
- We will leave everything by default except the Inter-Floor Heights.
- There are two ways of defining the lift, either by entering number of floors multiply by floor height if floor heights are universal,
- Or, specify inter-floor heights individually. Remember, if you are doing this way, the number of inter-floor heights is one less than number of floors, for example, 2 inter-floor heights for a 3 level building.
- We will use the first method as all of our floor heights are 4 metres, type in 3x4 since we have 3 levels.
- Now the lift is created, connect the walkway to the ones come from the areas.
- When all done, re-generate trips and run simulation, you will see people start to use the lift rather than stairs.

Checkpoint Barriers

A checkpoint barrier is a place where the flow of pedestrians is constrained, by applying a delay to each person at that place. A checkpoint is used to model a place where documents are checked, or where some other process is carried out that delays each person by a certain time. Examples include airport check-in, security screening, or even a coffee-shop queue.

Exercises: Checkpoint Barriers:

- Open “Checkpoint_Barriers_before.aza”. It includes two pedestrian networks, both with pedestrians travelling from South to North.
- Play simulation and notice people walking through smoothly.
- To create a checkpoint on the left model, select the middle walkway, notice point D is the one not connected to the first walkway.
- Adjust it and give it a delay value, for example, 5 seconds. This will make people delay at point D.
- If end C is the delay end, then enter the value in Delay DC.
- Set “Private” on. This ensures that only one person (or person group) at a time will use the walkway.
- This models an instruction like “stand behind the red line until called forward”.
- If there are two or more processing areas, then you can add further walkways parallel to the second walkway.
- And connect to the first walkway. Don’t forget to set the “Private” flag on the other processing areas. As you see for network on the right.
- Notice when the queue builds up, people are pushed out of the walkway. Set them “Walled” can keep them in the walk edges.
- It is assumed that all processing walkways have the same permissions, so if you want to separate processing area, “Local residents” and “Overseas residents” for example, then this separation should happen before the first queue walkway.

Queueing at Stands

There are a number of parameters you can change to modify how people queue for buses and trains, how they board and how they alight.

- **Transport Type / Doors**: The number of doors on a transport vehicle controls the rate of **alighting** (off-loading). If there are 20 passengers wanting to alight, and the alighting time is set to an average of 2 seconds, and there is a single door, then the total time required is 40 seconds. If there are 2 doors then the time will be reduced to 20 seconds.
- **Stand / Doors:** The number of doors on a stand reflects a passenger's expectation of how many **boarding** doors there will be on the next transport vehicle to arrive. This is more relevant to trains – a suburban train might have 8 or 12 doors for example. It does not matter if the next vehicle to arrive does not have a matching number of doors – all passengers queuing at the stand will be allowed to board, up to the capacity of the vehicle. This is also the case if the stand is longer than the vehicle. Setting Stand/Doors to a value larger than one will help to spread out the queue of passengers waiting on busy stands.
- **Stand / Slots:** At bus stands, sometimes there is space for more than one bus. If you want two or more buses to be able to stop at a stand simultaneously, you should set **Stand / Slots** accordingly, but the length of the bus stop bay should also be set using **Stand /** Size, allowing a bit of extra space for a gap between the buses. For example, if buses are 12m long, and you want 2 to stop at once, set Stand / Size to 26m and Stand /Slots to 2.

At busy stands, use a separate walled walkway as a way of forming an orderly passenger queue. Often, it is better to create something like a “two-lane” passenger area, using two separate walkways, with only the queuing walkway connected to the stand so that people coming from the other direction have to go to the end of the queue and double back

Exercises: Queuing at stands:

- Open the model "Queuing_at_Stands_Before.aza", It is a road network with three stops, the stop in the middle with a long platform is connected to an area by a walkway.
- The bus service runs north bound every 10 minutes, it stops at the middle stop to load and unload passengers.
- Within the simulation term, there are 160 passengers coming South to the middle stop, and 120 from the middle stop to North.
- Run simulation and notice when bus pulls over, this one connection is the only route for passengers to board and alight.
- This is not quite efficient, let's model it so the bus has two doors for alighting, so the time required will halve.
- To do so, go to Network> Service and click on the "Transport Types…" button, set Doors to 2, and click OK.
- Now run simulation again, you will notice passengers alight from the front and the back. This will significantly reduce alighting time.
- In some cases, the service allows passengers to board using multiple doors.
- Similar to the action we just did, but the number of doors is set on the stand this time. Set Doors to 2 for the stand.
- Remove connection from the middle stand to the area. Select the stand and the walkway and choose Walkway> New Connection to Stand.
- Notice this time two walkways connections are created.
- Run model and notice passengers now board and alight the bus from both doors.
- The Doors parameter on a stand is useful when modelling trains or subways, for example, a train with multiple carriages can be modelled as having several doors, so that when people come to the platform, they will distribute themselves across all the doors rather than all gathering at the front door.
- At busy stops, it's often seen that multiple buses pull over simultaneously, this can be done by setting the number of slots for a stand.
- But make sure the stand is long enough to accommodate multiple buses.

Pedestrian Flow

There are several parameters in Commuter you can change to modify the way that pedestrians flow and interact with others within your model:

Parameters / People / Space: The “personal space” parameter for a person-type defines the minimum distance that person will aim to create between itself and another person in front or to the side. A semi-circle in front is drawn in front of the person to illustrate this distance (in Flat view mode). The radius of the semi-circle is the person's body size plus the personal space distance. If a person is forced to accept a smaller distance, by another person moving into its space, then the semi circle changes colour to yellow.

Walkway / Adjust / Space Factor: This multiplicative factor (default 1.0) changes the space parameter for that walkway. So, if there is an area where you want people to bunch more closely, set the value less than 1, perhaps 0.8 or 0.5. If you want to space people out a bit more, set the value greater than 1, perhaps 1.5 or 2.0

Walkway / Adjust / Speed Factor: This multiplicative factor (default 1.0) changes the speed at which people walk on a walkway. If you want to make people move a bit more quickly, for example to walk briskly across a road crossing, you should set this to a value higher than 1.0

Walkway / Adjust / Capacity: This controls the number of people allowed on a walkway, regardless of direction of travel. Additional people will be held at the end of the previous walkway or crossing.

Connection / Adjust / Capacity: This controls the number of people allowed on a connection, regardless of direction of travel. Additional people will be held at the start of the previous connection. Use this to create a central crossing island with limited space, where a crossing is split over two directions of traffic.

Walkway / Adjust / DelayCD, DelayDC: Set either of these to “inf” to create a one-way walkway

Parameters/ Calibration / Pedestrian Motion: calibration parameters

Display / Diagnostics / Obstacles: Switch this on to understand the relationships between pedestrians.

Exercises: Pedestrian Flow

- Open the model “Pedestrian_ flow_Before.aza”. Here we have two networks.
- Let’s start with the left one. Demand is set so that people walk from the South area to the North.
- They will split and use both diverges. Play simulation and notice pedestrian flows smoothly.
- We want to model a ticket gate, so people are delayed for 10 seconds just before they merge into the walkway to destination.
- Select the two short upper walkways of the diverge, notice the end closer to destination is D, adjust the walkways and set Delay CD to 10.
- This will make everyone wait for 10 seconds at the D point before proceeding.
- Now play simulation again and notice the delay take places.
- Quite often we want pedestrians to wait until they are called to the gate, so they form a queue before the processing area.
- Let’s do this by selecting the two walkways we previously set delay for, and type 1 for their capacity settings, so that once one person is allowed in that walkway, once that person leaves, the next person is permitted to enter.
- Compared to use private setting, this gives you more flexibility as you can define any number for the capacity. However, it is an approximate value, a nominal capacity, as it can only hold pedestrians back at the end of the previous walkway (not on a crossing), so if there is one remaining space, but two people are already on connections, but not yet on the walkway, then the capacity will be exceeded by 1 when both enter, because by this time they have both got past the point at which they would be held.

- (continued over...)

- Now let's move to the network on the right. It has two crossings; one covering each direction of traffic, the two crossings are linked by a walkway connection, serving as a pedestrian island.
- Pedestrians are generated from the area on the left; they will walk towards the area on the right.
- Run simulation and see them walk at the same speed and stay the same distance from one another on crossings as on normal walkways.
- Let's modify this so that they move a bit faster while on crossings, and also stay a bit closer than they normally would.
- To do this, select the crossings, set space factors to 0.8, and speed factor to 1.5.
- This will result them requiring 80% of their usual personal space and walk 50% faster than they normally would when on crossings.
- The next step is to reduce population on the island and increase safety. By setting capacity on the walkway connection.
- As for walkways, connection capacity is a nominal figure, an approximate value, so you may need to set it a bit lower than your absolute maximum for safety numbers.
- Set it to a low number, for example, 5, run simulation again and now only five people are allowed to wait at the island, while other people will be held back before the crossing.

Chapter 10

Tutorial

Signal Rules

Signal rules are used at signalised intersections to switch phases depending on vehicle demand. Demand is measured using loops.

In this chapter we shall focus on four of the most common uses of signal rules, allowing:

- Phases to be skipped
- Phases to “gap out”
- Phases to be called on demand
- Phases to continue running into subsequent phases

For Signal rules to be implemented we require a signalised intersection. We will use a a simple 4-way junction as shown below. To create this intersection:

- In the graphic panel, right click mouse, select Road/New Road at cursor. In the pop-up window, select two-way road, and set number of lanes for “out” to 3 and “in” to 2.
- With the road just created still selected (it should be by default), right click the mouse again and select **Road > New Intersection**. In the pop up window, make sure the icon of the intersection is a four way intersection (like a big '+'). You can make this intersection a signalised one by selecting **Signalised** under the intersection icon. You will see that because we are creating a four way intersection and we have already created a road before we come here, we now have three groups of tables on each corner, each of them represent a road. Make sure each of them have both directions on, and number of out lanes being 2 and number of in lanes being 3. Click **OK** to create the intersection.
- Now in order to make the intersection work, you will need to create zones for each end of the intersection. To create a zone, pan the view so that the cursor is near the outward end of an approach road, right click mouse and select **Zone > New Zone At Cursor**, this will create a zone. Repeat this step for the other approaches
- To add loops, you can zoom in a bit in the graphic panel if you can’t see the centre of the intersection quite clearly. Double click left mouse button at the centre of the intersection, this will select it,

you will notice that lanes around the intersection changes to red indicating it is selected. Right click mouse and select Intersection/ **New Loops At Stoplines**. You will see that there are green rectangles added to the intersection. If not, you might need to switch on the **Loops** layer in the Layer Tab. If they are not showing the loop numbers within the rectangle, then turn on the **Loop Names** layer.

- Now create some trips. From the main menu choose **Demand / Demands**. This brings up a window for you to define your demand and generate trips later. Click on **Add Matrix**, you can name your matrix with a different name or simply accept the default, Click OK. The next window asks you to Select Origins and Destinations, for this tutorial, because we want trips to be generated from all the directions so we don't have to modify anything in this window, just click OK. Now you will notice that there is a table created in the Demand Editor window. This is where you can define how many trips you want Commuter to generate. The number 1,2,3,4 represents the names of the 4 zones we created earlier, you can specify how many trips there will be between each of these zones by typing a value you want. But for this tutorial, we will simply type a total number (e.g. 1200) at the lower right, above the "Close" button. Press "Enter" key on keyboard after you type the value, you will see the 1200 is automatically divided evenly for all the separate trips. Now you can click on "Save". To generate trips, click on "Trips…" it will bring up another window, click on the button "Generate Trips…" . Another window named "Accept Generated Trips?" will appear, you should see you have a Success Rate=100.0%. Click Yes to accept these trips. This will close the windows and bring you back to the main Commuter graphic window.
- Now, you can start the simulation running by pressing the play button on the tool bar above the graphic window.

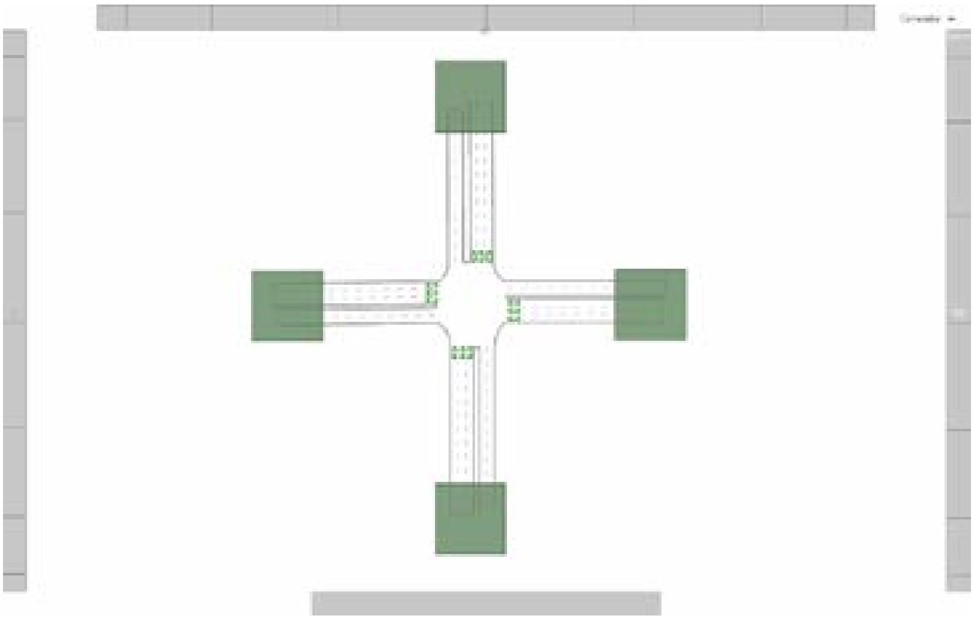

If loops are added it is important they are correctly named to reference their relationship to the intersection in question, for instance if we were working with intersection 2 and added a loop number 10, we would need to rename this loop to “2_10” for the intersection controller to recognise the loop. If you select a loop, then Action **Adjust** you can view or modify the controller to which a loop is attached.

The image below shows the loop numbering for this intersection:

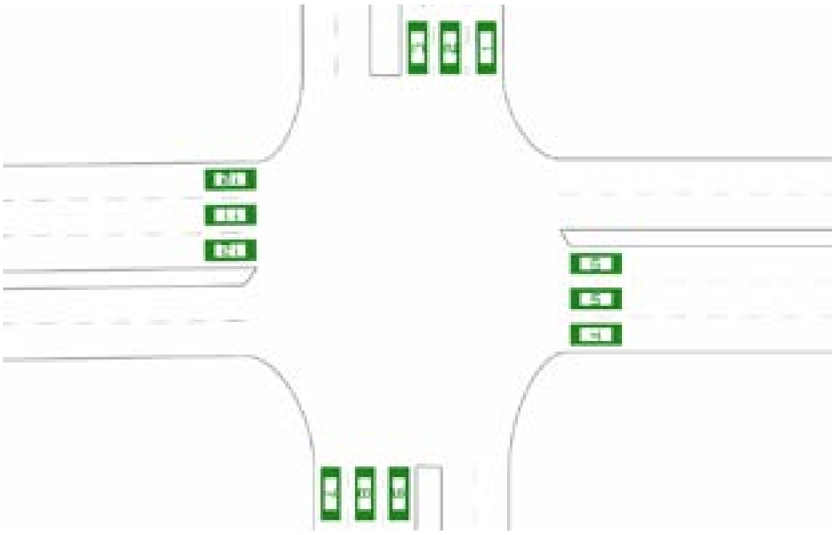

Note that the names do not include the “2_” in the graphic – we can easily check associated loops in the “loops” tab of the signal controller dialogue).

The possible phases for this traffic control site are shown in the following image:

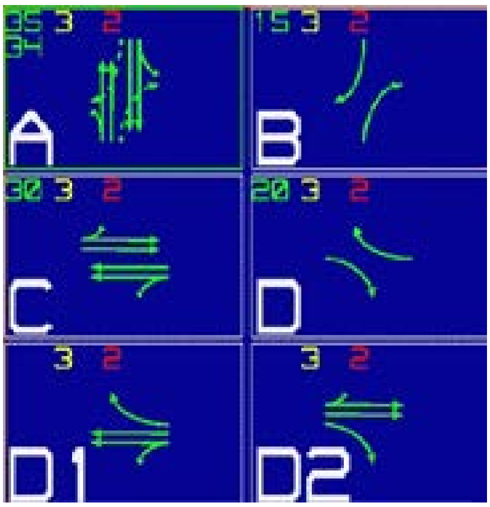

In terms of the rules, the three core actions available are:

1. Terminate
2. Extend
3. Skip

The two key tests we will use relate to:

1. Gap
2. Occupancy (occ)

The available operators are:

>, <, =, >=, <= (for numerical comparison)
AND, OR (for logical)

The relevance of this will become apparent shortly

Skipping Phases

The requirement to skip phases is common when there is no demand for the relevant movements. We can see in the above image that A Phase is running the turning movement from the north and south approaches as filtered movements – Phase B then runs which operates these same movements in a protected state for up to 15 seconds of green time (+ 5 seconds intergreen).

Skipping usually results in the surplus time being allocated to a phase that is yet to run in the plan.

As an example, assume that B Phase is not required (as all of the traffic has managed to discharge during A Phase when those movements are filtered) and we want to give this surplus time to Phase C. Referring to the detector diagram, we can see that loop 3 is the relevant loop for the southbound right turn and loop 9 is related to the northbound approach right turn. This can be phrased as:

> *Skip B = When Phase B is about to start, if loops 3 and 9 are not occupied then do not run Phase B and give its time to Phase C.*

To code this as a rule

- From the main menu, open Control / Intersections / Controllers
- Select the tab called "Rules".
- Press [+ Rule] - the "New Phase Rule" window will be raised
- Type the new name for the rule: "Skip B"
- Set the Action to **Skip** - the action phase is the phase to which the rule will be applied
- Set Balance phase to C - the balancing phase for this type of rule is the phase that will receive the surplus time taken from B
- Set the Test to ***occ3=0 and occ9=0***

Gapping Out

Gapping-out is a term used when a phase is terminated if there is a gap in the traffic that is serviced by that phase. It is an important part of locally optimising the way a traffic control site operates and is a common requirement where some movements experience volatile or low demands.

Working again with B Phase in the above, lets say that B Phase is running. However, after a few seconds it is no longer required as demand for the movement falls and we want to give this extra time to Phase C

Again the action phase represents the phase during which the conditions are checked for, in this case Phase B. And again, Phase C is our balancing phase.

The *Action* this time is “Terminate”.

In this instance the rule we would look to test is:

> *If as Phase B is running and the gap between subsequent vehicles on loops 3 and 9 exceeds an acceptable threshold (say 5 seconds) then* terminate Phase B (with yellow and red time) and give its remaining time to Phase C.

The logic for the test is:

“*gap3>5 and gap9>5*”

Calling a Phase

Although we are referring to this as “calling a phase” this is not strictly the case which will become apparent.

The most common time we may want to do this in the above example is where we have Phase D which can also run as D1 or D2. In this case we may want to check to see which of these phases is actually required and run it appropriately.

Note that according to the signal plan that D Phase is allocated 20 seconds of green time and 5 seconds of intergreen and while we say this is calling a phase what we will actually do is skip D and give this time to D1 or D2 as required.

We are calling this rule “Call D1”.

The action phase represents the phase during which the conditions are checked for, in this case the split second Phase D runs.

We will discuss the necessary “test” shortly.

The *Action* in this instance will be “Skip”

The balancing phase this time will be D1.

Looking now at the actual Rule to test we refer to our detector diagram and see that loop 6 is the relevant loop for the westbound right turn and loop 12 is related to the eastbound approach right turn movement.

D1 refers to the scenario where the westbound right turn is needed but the eastbound is not.

In this instance the rule we would look to test is:

If at the split second Phase D commences loop 6 is occupied and loop 12 is not currently occupied then do not run Phase D and instead run Phase D1. The logic looks like:

“*occ6>0 and occ12=0*”

To cover the alternative scenario we would then create a second rule which would consider the need for D2.

This is essentially the same except as you can see the balancing phase is now D2 and the test would be:

“*occ6=0 and occ12>0*”

We could also employ the techniques learnt above and have a third scenario that considered whether any form of D Phase is needed and if not we could allocate this time to a more deserving movement – essentially Skipping D Phase all together.

This would again be a Skip rule and the logic quite simply would be:

“*occ6=0 and occ12=0*”

The phase we allocate as the Balancing Phase would then receive this time.

Phase Extensions

Considering the phase extension scenario – we can see in figure 3 that A Phase is running the turning movement from the north and south approaches as filtered movements – Phase B then runs which operates these same turns in a protected state for up to 15 seconds of green time (+ 5 seconds intergreen).

At the end of phase A we may want to check to see if B Phase is actually required and if it isn't, continue to run the A Phase instead as clearly it is discharging the necessary traffic.

In this case we are calling our Rule "Extend A"

The action phase represents the phase during which the conditions are checked for, in this case during A Phase.

We will discuss the necessary "test" shortly.

The *Action* in this instance will be "Extend"

The balancing phase for this type of rule represents the phase the extension will occur into – in this case we are stealing time from Phase B for Phase A if Phase B isn't required.

Looking now at the actual Rule to test we refer to our detector diagram and see that loop 3 is the relevant loop for the southbound right turn and loop 9 is related to the northbound approach right turn movement.

In this instance the rule we would look to test is:

If the gap on loop 3 and gap on loop 9 is greater than a certain value and those same loops are not currently occupied then allow A Phase to continue. The logic looks like:

"*gap3>5 and gap9>5 and occ3=0 and occ9=0"*

The value of 5 in the above represents 5 seconds of time, thus if there has been more than 5 seconds between the subsequent passing of vehicles and this isn't a result of a vehicle being stopped due to an inability to filter (loop is not occupied), then the test is true and A Phase can continue to run until it has used up all of the allocated time for Phase B as well as its own time.

The Extend rule is checked every second at the end of a phase as at any point while A Phase is running the allocated time for Phase B, there may still be a need to run B Phase.

Another useful thing to consider at this point is the fact that the "occupied" (occ) parameter refers to *time* occupied and is not a true/false value – this means we may want to include a value for occupancy that is greater than zero. For instance in the above we may want to say "if the loop has been occupied for more than x seconds" as a further consideration.

Other Considerations

Understanding "Time"

The timing tab is where we nominate phase times along with their intergreen (yellow and red) and the Minimum and Maximum times. Below we will discuss the effect of minimum and maximum times on signal rules.

Minimum Time (Min): No running phase can gap out before its minimum time has expired, that is if B phase were running and its minimum time is set as above at 5 seconds, it is not possible to gap this phase out until that 5 seconds has elapsed (provided the phase actually runs).

Maximum Time (Max): If extra time is attempted to be allocated to a phase (through gapping or extensions etc) that would cause that phase to have a green time that exceeds is maximum green time (as nominated in the timing tab of the controller dialogue – above all are set at 300 second), this rule will have no effect and in essence will be "false".

Note that if a phase has 0 green time as shown for phases D1 and D2 that even if minimum or red and yellow times are allocated these will not be run if the phase is not run – these only become active if the phase actually runs (has green time >0).

The Role of the Stretch Phase

Also note the ability to nominate a phase as the "Stretch Phase". This becomes important in instances where no balancing phase is nominated in the rule – in this case the extra time would automatically be allocated to Phase A in the case above.

Rule Order is Important

The order in which rules are specified can be important. Rules can be easily re-ordered as shown below:

Summary

The key behind the rule system in Commuter is allowing for many actions on various phases during any given cycle without altering the overall cycle time of the site. This means progression and “green waves” can be maintained despite what might or might not be happening at any one signal site in the system.

It should be clear by now that a lot of flexibility is provided with only a few simple concepts. Central to the approach is understanding that each traffic control site can have a multiple number of rules as required to replicate true localised optimisation.

As an example when we discussed phase calling we could also imbed rules that allowed those called phases to gap-out – we may want for instance D to gap into D1 or D2 and if D1 or D2 is called, for that phase to likewise be able to gap out into another higher demand phase.

Each rule can have complex tests – those discussed above are simple, but we may want to for instance consider which phase to actually give our unused time to so we may not just be interested in loop occupancy for our particular movement – we may also want to check the occupancy for the loops of the phase we are considering using as the balancing phase.

One of the other powerful elements is an ability to take action with respect to phases that are due to run in the next cycle. That is in the above if we were running D phase and gapped it out we may want to allocate this unused time to B phase in the next cycle – this is also possible where the phase nominated as the balancing phase does not occur in the current cycle (i.e. it has already passed).

Boom Gates

The boom gates tool allows you to add a movable boom gate, like those found on parking entrances and exits, or on restricted access lanes, such as bus-only lanes. The boom gate device will delay each vehicle for a variable time, taken from a specified time distribution.

Exercise: Create a boom gate

- Open Boom Gates window from the menu: **Control / Boom Gates**
- Select a lane
- Press the green [+] button on the Boom Gates window
- Press OK
- In the Layer Tab, switch on the **Signs** Layer

After creating a boom gate, there are some parameters you can modify:

Stop Time: A new stop time distribution for a boom gate can be defined in Parameters / Time Distributions, or you can use an existing distribution. For example, if you select “Uniform Standard A” every vehicle will be stopped for 2 seconds. However, you may want to create a distribution where 40% of vehicles stop for 2 seconds, 30% for 3 seconds, 20% for 4 seconds, and 10% for 15 seconds. Note that the total stop time is the sum of the stop time taken from the distribution and the (fixed) time for the gate to open.

Time to Open: This time is the same for all vehicles, and is added to the variable stop time. It can be set to zero, if not required.

Limited Height: A visualisation option, splits the boom gate bar into two parts so that it does not collide with the roof in a multi-level parking area.

Ramp Metering

The Ramp Metering Tool allows you to model traffic signal control used to limit the flow of traffic from an on-ramp onto a multi-lane road. The tool uses a logic module that takes its input from one or more loops and a set of variable parameters. The output from the logic module is the cycle time of the control signal. That is, each ramp meter generally has a fixed green time, and a variable red time, and the length of the red time is determined by flow or occupancy detected on one or more loops.

Exercise: Ramp Metering

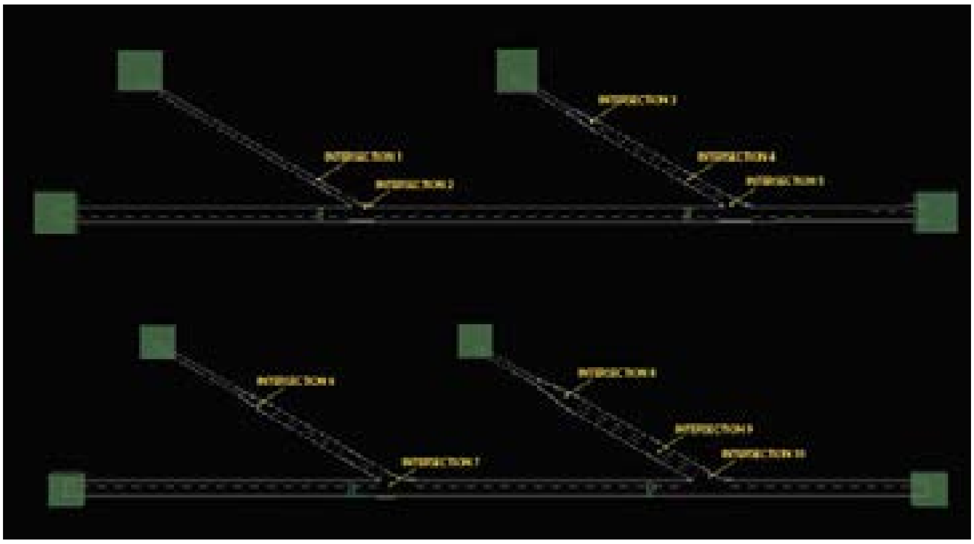

- Open the model "Ramp_Metering_Before". In this model, there are two groups of networks. Each with a main road and two arms merge into them.
- There are five intersections on each group of network. Two of them are at the merging intersections. One is on the single lane arm and two on the arm that starts with one lane but then splits into multiple lanes.
- Zoom in to the left merging intersection on the upper network. Make "Lane Centrelines", "Lane Steams" and "Handle – End" visible. You will see there is a short lane on the arm near the

merging intersection, and a long lane on the arm away from the merging intersection. These two lanes are connected by intersection 1.

- Loops are created on the main road on the left of the merging intersection, select them and check their parameters. They are all assigned to Controller 1.
- Now pan the view toward the right to see the other merging intersection on this network.
- There are also three loops added, and they are associated with intersection 4.
- This arm is similar to the arm we saw previously. The difference is that it starts with one lane and then splits into two. An intersection set at the split point. Whereas the other arm just has a single lane all the way to the intersection.
- Pan the view to see the other network at the bottom. It also has two merging arms. The configuration is very similar to the one we just checked.
- Intersection 18 connected the lanes on the merging arm near the merging point. And the loops on the main road near the merging point are associated with intersection 18.
- Similarly intersection 16 connects the lanes on the merging arm near the merging point on the right. And the loops near the merging intersection are associated with intersection 16.
- There are two OD matrices already set in the model, one for each network.
- Each matrix contains through flow of 1000 and merging traffic of 200 each arm. The other network has the same flow volume.
- Start simulation running and you should be able to see traffic flowing.
- Select intersection 2 and check its configurations. The through flow is set to free flow and the merging to give way. And flow to the main road from this arm is continuous.
- We want to control the ramp flow so that the volume released to the main road is limited. We will do this by using the “Ramp Metering” tool. Open the window from the Control menu.
- Once the window is open, you will see there are three tabs, each indicating a metering type.
- Let’s start from the left intersection on the upper network and experiment with the “Fixed Rate” type.
- Select the short lane on the arm and a loop on the main road. Then in the Ramp Metering window click “New”. Click “OK” to the

default setting. A new meter is added, indication of meter and its type also appears in the graphic window.

- Rewind simulation and start it running again, notice the meter switches between red and green. When it is red, vehicles stops before the intersection and while it's green, one vehicle is released.
- Check the parameter of this meter from the ramp metering widow. Under V/Cycle it is set as 1. This means only 1 vehicle is allowed to be released in every cycle.
- Let's change this value to 3 so that 3 vehicles can be released. Rewind and start simulation again.
- Notice the setting is not making a difference, there is still only 1 vehicle released each time.
- This is because we have not changed the green time for the meter. 1 second is not enough to release 3 vehicles. Change the green time to 6 second for example and then start simulation again.
- Notice more vehicles are allowed to pass. When there are three of more vehicles waiting, three of them are released. This is true regardless of the number of lanes. If there were three lanes, the one vehicle per lane would be released.
- Now pan the view to the next arm on the same network. It has a single lane from the zoom and then splits into two lanes. This can be seen more clearly by turning on "Lane Centreline" and "Lane Streams".
- Select a lane on the arm near the merging intersection and a loop on the main road. Under the Fixed Rate tab, create a new meter, accept the default settings. Notice now the V/Cycle is set to 2 automatically because there are two lanes.
- Rewind and start simulation and notice two vehicles are released each time.
- Now pan and zoom to the left merging intersection of the bottom network. It also has one single lane from the zone that splits into two lanes.
- This time, we will create another type of meter, called ALINEA (*Asservissement Lineaire d'entrée Autoroutiere).*
- Select a lane on the merging arm and a loop on the left of the merging intersection.
- Ensure you are in "ALINEA" tab and create a new meter.
- Rewind and start simulation to see it work in action.
- Experiment changing the V/Cycle to 10 and Green to 5 to allow more vehicles to be released every cycle.

- Click the “ALINEA-Q” tab in the ramp metering window. We will experiment this with the right merging intersection.
- The difference between ALINEA-Q to ALINEA is that is has an additional loop to detect the queue length. Once the queue on the arm exceeds limit. The meter will be disabled and flow will be uncontrolled.
- The first thing to do is add the required additional loop; we will add this to the lane from the zone.
- Set its controller to 16, which is the intersection on the arm prior merging into the main road.
- This associates the loop to that intersection.
- Select a short lane near the merging intersection on the arm and a loop on the main road. Then create a new meter.
- Notice the ALINEA-Q meter comes with an “Off Occ” setting.
- This is the threshold value, when occupancy reaches this value, the meter will be disabled.
- The Occupancy is shown in “Occ[2]”.
- Start simulation to see the meter running. “Occ[2]” is currently 0.11, set the “Off Occ” to a value smaller than this.
- Re-start simulation; notice meters are on at the beginning, but goes to OFF when the “Occ [2]” is higher than “Off Occ”.
- The intersection signals are drawn as three green arrows, there is no longer controlling of vehicle flow.
- Increase the “Off Occ” to 0.30 and notice the meter control comes back when the “Occ[2]” is lower than “Off Occ”.
- Change “Off Occ” back to 0.05 and soon the signals go to all green if “Occ[2]” is more

Trail Following

The Trail Following Tool allows you to control the route assignment of all or some of the agents in the model, by forcing them to use one or more predefined trails. A Trail is a sequence of roads or walkways and can contain the same branch more than once (i.e. circular routes are allowed), and it can start and finish anywhere.

Exercise: Trail Follower

- Open the FourCityBlock-Trail-Follower.aza Model in your Tutorial Chapter
- This is a pedestrian network with trips already generated, so click on the “Play” button
- By default, people will always pick the lowest route cost
- Press “Pause” then “Rewind”
- We are going to use the Trail Follower Tool to make them take a different route
- First we need to create a trail
- Select the bottom walkway that comes out of “Area 1”
- Open the trail window under Network>Trails…
- Click on the “New… “ button to create a new trail
- You can specify a name here for the Trail here, when done, press OK
- A New Trail window appears
- Use the zoom value to zoom out and see the different colour coded options that you have
- We will make them walk in an “S” route shape, so click on blue 3 times, red, green 2 times, and blue
- Press OK to create the new trail
- Press Apply then OK to close the Trails window
- What we need to do now is to apply this trail to the network
- To do this, go to Assignment>Trail Follower to open the Trail Following window
- Click on the add button (Green Arrow) to add a new rule
- Here you can rename the Rule Name and choose the trail that you want to use
- Select the trail you just created and click on OK when done
- This will apply the trail to the walkways and affect everyone on it

- Press Apply and OK to close the Trail Following window
- Go to File>Save As… and name it Trail-Following-1
- Press “Save”
- Press “Play” and watch the people follow the trail that you have just created!
- Press “Pause” then “Rewind” when you are done

Menus

Reference

File Menu

New		Start building a new model. This is available only after you have saved changes to new or existing model
Open...		Browse the file system for a model, then open it
Open [Most Recent]		Open the most recent model. This is active only after a model has been saved
Open Previous ▶		A sub-menu containing the most recently opened models.
Save		Save the current model
Save As...		Save the current model as a different file
Restore		
Wizards▶	**Parameters**	Step-by-step setting of parameters
	Demand	Step-by-step creation of demand
	Service	Step-by-step creation of a public transport service
Import ▶	**Paramics**	Import a model from Q-Paramics format
	VISSIM	Import a model from VISSIM INP format
	SCATS	Import data from SCATS format
Save Snapshot ...		Save a JPEG image snapshot of the current 3D view
Exit		Exit Commuter

File → Import from SCATS

Before You Begin – Set Up SCATS Data Files for Import

This type of import assumes you can open the Central Manager database file for SCATS, normally called SCMS.mdb. The MDB format of this file is Microsoft Access database format. If you do not have this MDB file, then it is not possible to import SCATS data.

In addition to the MDB file, other files can provide supplementary information. These additional files are:

- LX files, one per region
- SFT files, one per intersection

If you have LX and SFT data, we suggest you structure the files as follows:

1. Create a base directory, and copy the MDB file to that base directory. You may have data for different years. The example shows data for 2005
2. In the base directory, create sub-directories LX and SFT
3. In the LX directory, create a sub-directory for each region. In our example, the regions are called DY and NSYD. You should use exactly the same region codes that you use within SCATS for each region
4. In each LX region directory, create a sub-directory called sys, and copy the files sys.lx, sys.ram and sys.tc from your SCATS installation to those folders
5. Repeat steps 3 and 4 for each region
6. In the SFT directory, create one sub-directory for each region, and copy all the SFT files for that region into that sub-directory
7. Repeat step 6 for each region

The suggested file structure is shown in images on the next page.

Suggested File Structure for SCATS Import Data

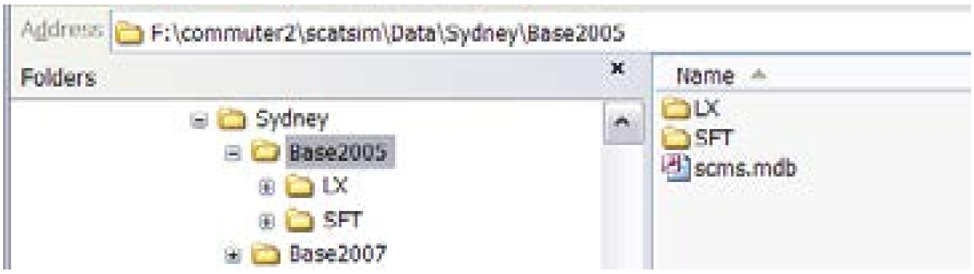

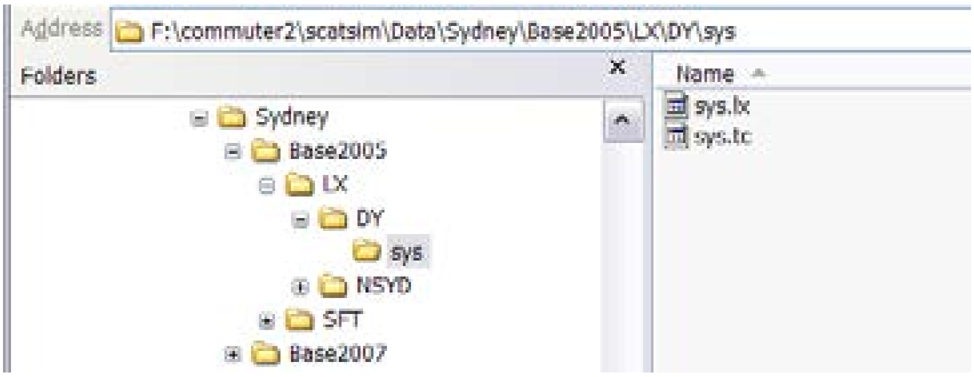

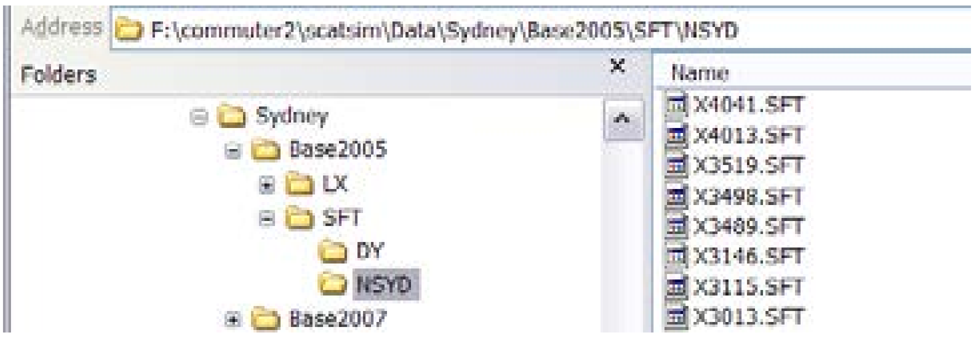

Import from SCATS Data

The first window you see asks you to enter a location for the SCMS.mdb file. Browse to the MDB file, and if you have set up the LX and SFT data in the standard way, as suggested, the Commuter SCATS importer tool will automatically locate all your LX and SFT data. If you have set up the data in a different, non-standard, way, you can browse to the LX and SFT root directories manually. The root directory is the directory in which all the region subdirectories are located.

Import Step 1

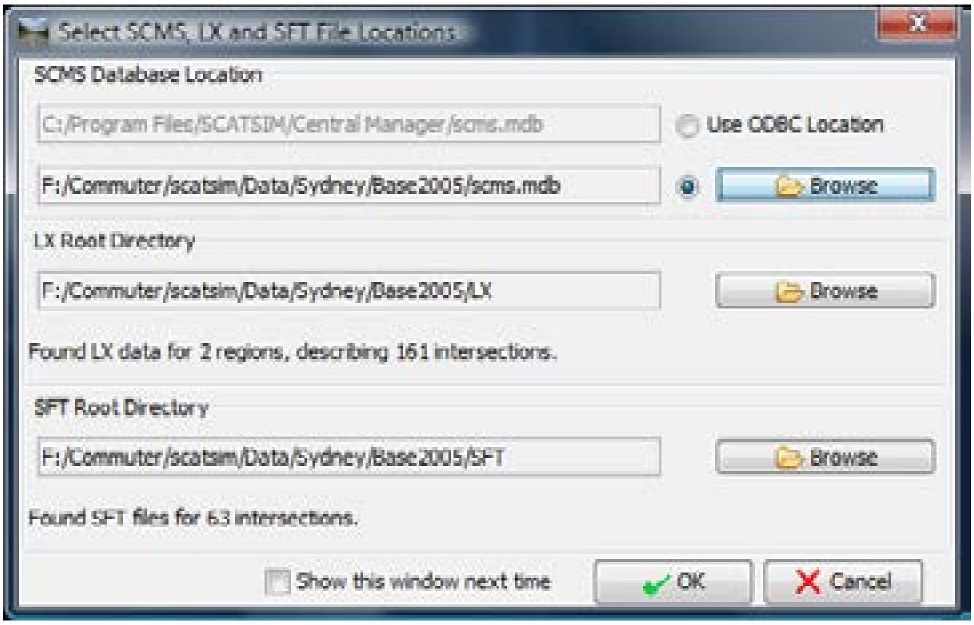

Press OK when the data locations are correct.

Another window will be raised and there will be a pause while it reads the SCATS data.

Import Step 2

Once the data has been read, a list of all available intersections will be shown. There are a number of filter boxes along the top of the window – use these to select a subset of intersections. For example, the image below shows all intersections in region DY:

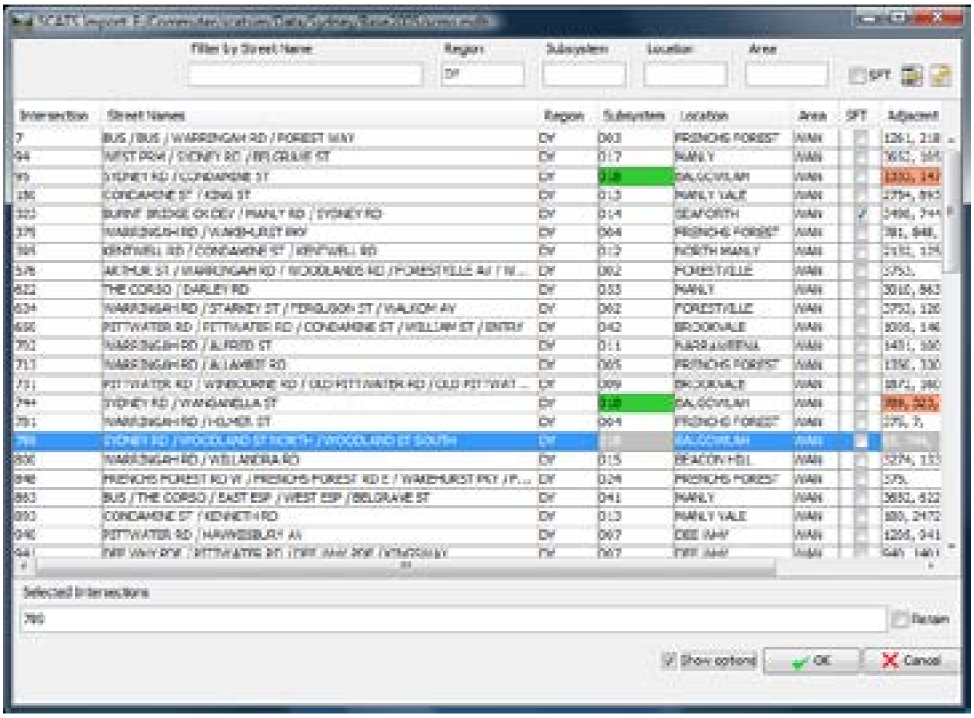

If you are extracting data for later connection to SCATS using the SCATSIM plugin, you should choose only those intersections for which there is both LX and SFT data. If you are using the Data import solely to construct a network, and do not intend to connect the simulation model to SCATS, then you can pick any intersection.

When you select an intersection, other intersections in the same sub-system will be highlighted in green, while adjacent intersections will be highlighted in pink. Once you have selected the intersections you require, click to select 'Show options' and press OK.

Import Step 3

This window allows you to set various options for the import process. In most cases it will be sufficient to accept the default settings and press OK.

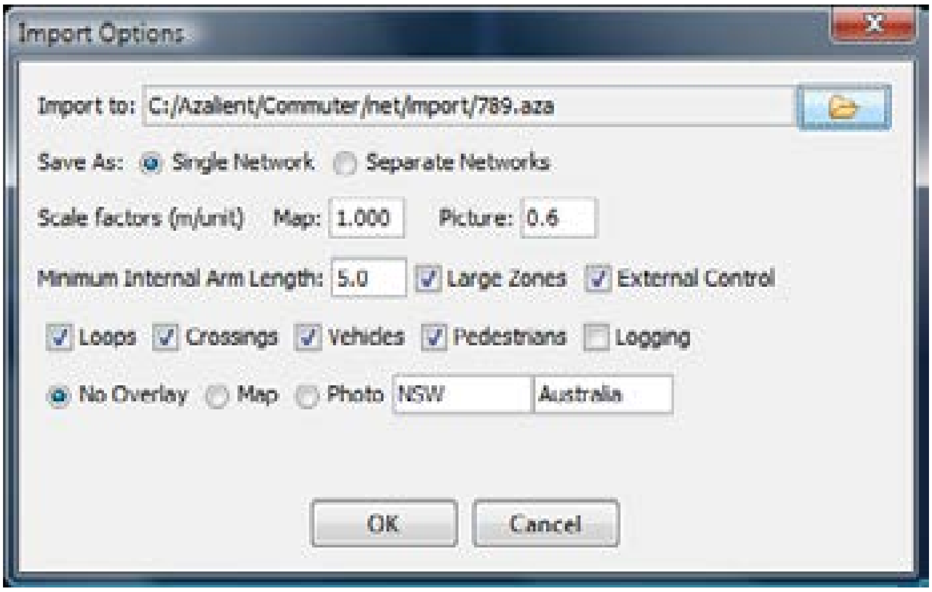

Single / Separate Networks: selecting **Single** will create a single commuter model containing all selected intersections. Selecting separate networks will create as many Commuter models as there are selected intersections, with one intersection per model. Once all models have been created, the Commuter window will remain open on the last model.

Scale Factors: the SCATS map data and the SCATS picture data contain implicit scale factors. The values shown above are the defaults. If you change the values from the default, this will be stored and used on subsequent imports.

Minimum Internal Arm Length: where an intersection has been drawn in SCATS with offset arms, this parameter controls the length

of the smallest internal arm that will be converted into a road in Commuter. Internal arms shorter than this value will not cause a road to be created within the intersection.
Large Zones: if this is selected, the zone created at the end of each arm of the intersection will extend to beyond the half-way point of the arm. In this way, if the network is subsequently exported as a Paramics model, the zone will be large enough to included the mid-link point as required by Paramics to included a link within a zone. The **Vehicles** option must also be selected for Zones to be created.

External Control: select this if you wan the external control flag to be set for each intersection. If you intend to connect the network to SCATS via the SCATSIM plugin in Commuter, or through an exported model in Paramics, you should select this option.

Loops/ Crossings/: select these options if you want the new model to contain loops and pedestrian crossings respectively.

Vehicles/ Pedestrians: select these options if you want the model to be generated with Zones and Areas respectively. Even though no demand matrices or volumes will be created, the model will generate undirected trips when started using the default volume for each zone or area.

Logging: select this option to request information on the progress of the import process to be spooled to the Log window.

No Overlay/ Map/ Photo: The Commuter SCATS import can automatically import Google maps or photos if these options are selected. The map or photo will be read from a list of street names. The Image Fetch window allows you to select the appropriate pair of street names per intersection.

Press OK when you have set the options as required.

The data will now be imported and the model will be open in the main Commuter window. You can now run the model, save it, or export it to Paramics format.

Edit Menu

Undo	Undo the last Action. The Undo Tab contains a list of all undoable actions
Copy	Copy all the selected objects to the application clipboard, storing their position relative to the cursor
Paste	Paste all objects on the clipboard relative to the current cursor position
Copy To File	Browse to a file, then copy all the selected objects to that file, storing the position of each object relative to the cursor. You can use this type of file as a template: if you have a pattern that is repeated in your model, like a common intersection, you can copy the intersection to a file, then paste it later into this model, or other models.
Paste From File ►	Paste all objects from a file relative to the current cursor position. The sub-menu contains a list of files used recently for pasting. Where the objects are pasted depends on how the file has been saved. If the file was saved using **Copy To File** then this will have saved the position of each object relative to the cursor, so you can paste the same file in multiple times, moving the cursor in between. However, if the file is a normal AZA mode file, the objects will be pasted at a fixed position, so if you paste this type of file more twice, then the second copy will be in the same position as the first, regardless of the cursor position.

Preferences ...	Open a window that allows you to set preferences for • File locations: Open/ Save, Import/ Export, Copy/Paste • International Settings: Driving Side, Walking Side, Measurement Units, Currency, Language • Options: Multiple Threads for Pedestrians, Vehicles, Public Transport

Edit / Preferences

Files and Directories

Home Directory for Open / Save	Set the default location for opening and saving model. You can set this to your own user directory, or one specific to your project. You can use the "Home" button in the top right hand corner of the file browser to go to this location
… for Import	As for Open / Save, but for imported models
… for Export	As for Open / Save, but for exported models
… for Copy	As for Open / Save, but used for saving data when using **Edit / Copy To File**
… for Paste	As for Open / Save, but used for saving data when using **Edit / Paste From File**
Backups Kept	When a model is saved, the previous save of that model is retained as backup file 1, the previous backup file 1 is moved up to backup file 2, and so on. By default there are two previous backups, but you can increase or decrease this number. The backup files are normally hidden files. The **File > Restore** function reverts to the most recent backup, and moves all the backup files down by a number.
Size of "Open Previous" List	The number of models listed in the menu for **File > Open Previous** Setting this to -1 shows all previous models

International Settings

Driving Side of Road	The side of the road used for driving in your country (Left or Right)
Walking Side	The preferred walking side for walkways. This can be changed on each walkway using the **Keep** field (keep left/ keep right)
Measurement Units	Metric, UK or US units
Currency Symbol	The currency symbol used in reports
Language	The language used for the user interface. The menu will show the language packs that have been installed. Please request a language pack if yours is not shown.

Options

Multiple Threads	If you have a computer with multiple cores (e.g. Intel Core-i7), using multiple threads can make the simulation run much faster. The **Automatic** setting is recommended.

Skins

Skin (Theme)	The “Skin” setting controls the “look and feel” of the windows in the user interface. The recommended setting is **System**. If you use another Skin, and you experience display issues, you can reset it to the default by switching on Caps Lock before you start Commuter.

Terms

A Term is a named window in time. It has

- a start time, normally specified as HH:MM
- an end time, also normally specified as HH:MM
- a day of the week, by default “Any Day”, but can be something more specific

A Term is considered to be “on” or applicable for any time between these end points on a matching day. The day can be a particular day or it can be a day type (weekday or weekend). The association of a day can be useful, for example, if you want to define a demand for a weekday and a different demand for the weekend. When you run the simulation, or generate trips, the Terms associated with profiles or other time-varying functions will be checked to ensure you have all the correct simulation components enabled for any particular run.

You can have as many Terms defined for your network as you require. They can overlap, or be sequential. A lane restriction can be defined to have one term (for example “buses only” defined for 07:30 to 09:30) while a demand profile is defined to have another (for example “AM Peak” defined for 07:00 to 10:00).

There are two special terms, which are always present:

- **Always** – Where a term is required, selecting Always ensures that the parameter to which the term is attached applies all the time, even if the simulation term is changed. The Always term cannot be modified.
- **Simulation** – This is editable, although the name cannot be changed. This term defines the start and the end of the simulation period.

Behaviours and Mobs

A Behaviour defines a number of parameters, controlling decision made by each person. Behaviours are assigned to person-types; each person generated has one of these types.

A group of Behaviours is called a **Mob** (Mix Of Behaviours), and this is described at the end of this section. The Mobs window can be accessed by a button at the bottom of the Behaviours window

Behaviours are also associated with Vehicle Types, but this is relevant for the following types of Vehicles only:

- Taxis
- Drop-Off Vehicles
- Pick-Up Vehicles
- Zone to Zone (background traffic) vehicles

For other Vehicles, the Behaviour is adopted from the Person that gets into the Vehicle and becomes the driver.

Separating Behaviour from the Type of the Person or Vehicle allows you to separate the physical constraints from behavioural constraints, and to define a range of types that incorporate combinations of ranges of each. For example, you could define an aggressive behaviour and a passive behaviour, and then go on to associate each of those behaviours with a range of physical vehicle types, such as car, van and truck.

Another advantage of separating behaviour from physical constraints is that the Behaviour can be retained when modelling a multi-stage journey. For example, if a journey is made up of a car trip followed by a pedestrian trip (walking from the car park to the station) followed by a public transport trip, all three trips can be modelled using an agent with a consistent behaviour. See the help on Types for more information on assigning Behaviour to Types.

The Behaviours window is divided into several tabs, as described below.

Mode Choices Tab

This tab contains the “master switches” that control which modes are available to each behaviour. If you want to define a behaviour that will always drive, rather than taking public transport, then this is the place to do it. Similarly, you can define behaviours for people who prefer to cycle, or take a taxi, or are lucky enough to have someone available to drop them off at their destination. The switches available are:

- **Can Walk:** True if a person of this behaviour can use walking as a mode of travel in the model.
- **Can Drive:** True if a person of this behaviour can drive to a parking zone or a transition zone in the model.
- **Can Ride:** True if a person of this behaviour can take public transport
- **Can Taxi:** True if a person of this behaviour can take a taxi, if taxis are available. Even if this is true, the cost of taking a taxi will still be taken into account in the calculation of the lowest-cost route to the destination. Setting this to false will rule out the option of taking a taxi.
- **Can Cycle:** True if a person of this behaviour will cycle, and a cycle vehicle type exists
- **Can Be Dropped Off:** True if a person of this behaviour can be dropped off at a designated drop-off zone. This implies that this person has a car and driver available to drive them to their destination, and thus will not need to spend money on parking. This may often be the cheapest option for travelling to a destination with pay parking, but in many cases only a small fraction of the population will have a car and driver available to them.
- **Will Be Picked Up:** True if a person of this behaviour will be picked up at a designated pick-up zone. If this flag is True then no other mode will be considered (apart from walking from the origin to the point of pick-up).

Walking Costs

- **Can Walk:** True if a person of this behaviour can use walking as a mode of travel in your model.
- **Walk Cost/ second:** The value of time for walking, specified in small currency units (pence, cents, etc.) per second.
- **Walk Cost/ distance:** The value of distance for walking, specified in small currency units (pence, cents, etc.) per long distance unit (km, mile).
- **Walk Cost Base:** a base walk cost, used to bias short trips against long ones.

Driving Costs

- **Can Drive:** True if a person of this behaviour can drive to a parking zone or a transition zone in the model.
- **Drive Cost/ second:** The value of time for travelling in a private vehicle (car, van, truck, etc.) The cost value is specified in small currency units (pence, cents, etc.) per second.
- **Drive Cost/ distance:** The value of distance for travelling in a private vehicle, specified in small currency units (pence, cents, etc.) per long distance unit (km, mile).
- **Drive Cost Base:** a private vehicle cost, used, for example, to model the base cost of ownership

Transport Costs

- **Can Ride:** True if a person of this behaviour can use public transport as a mode of travel.
- **Ride Cost/ second:** The value of time for travelling on public transport, specified in small currency units (pence, cents, etc.) per second.
- **Ride Cost/ distance:** The value of distance for travelling on public transport, specified in small currency units (pence, cents, etc.) per long distance unit (km, mile).
- **Ride Cost Base:** a base public transport cost, used, for example, to model the base fare payable.
- **Wait Cost/ sec:** The value of time for waiting for on public transport, specified in small currency units (pence, cents, etc.) per second

Parking

- **Park Duration:** A time, in hours, used to calculate the cost of parking. The premise for this parameter is that a person with this behaviour is planning an outward-return journey, and intends to park for the specified duration. For example, a person working in a city centre may park for 8.5 hours during a working day. A business traveller going to an airport may be staying at their (air) destination overnight, and be planning a stay of 36 hours at the airport, from 07:00 Monday to 19:00 Tuesday.
- **Parking Vehicles:** A (demand) division of vehicle types to use when selecting from parked vehicles. A demand division has % values to generate vehicles in proportion when
    - there are no parked vehicles to choose from in a "normal" parking zone
    - the zone is a transition zone or an instant parking zone
- **Optimist**: if this is on, the driver will always assume that there is still a free bay in the most popular area. If it is off, the driver will take the first available space

Driving

- (Speed) **Compliance**: A multiplier, by default 1.0. This is used for vehicles to specify the “perceived” speed limit on any road. Multiply the signed speed limit by this parameter to obtain the perceived speed limit that will be used as the maximum speed a vehicle will attain if unconstrained by other vehicles or lane geometry.
- **Minimum Gap:** A distance, in metres, used to specify the minimum spacing between stationary vehicles in congestion.
- (Target) **Headway:** a time, in seconds. This is used in the vehicle-following algorithms
- **Reaction Time:** a time, in seconds, used in the vehicle-following algorithms
- **Safety Margin:** used to calculate stopping distance. The absolute minimum stopping distance is derived from current speed and (constant) maximum deceleration is multiplied by this safety margin.
- **Lane** (Change) **Gap:** a time, in seconds, used for the lane changing algorithm. A larger value means that a larger gap will be required, making lane changing more conservative. A smaller value for lane changing gap will make lane changing more aggressive.
- **Variability:** a number between 0.0 and 1.0. The variability of: (speed) compliance, reaction time, target headway and minimum gap distance defined in terms of standard deviations relative to the mean. So a value of 0.1 would mean that a distribution would have -1SD to +1 SD from 0.9x mean to 1.1x mean. Set this to a non-zero value to make the behaviour of your agents less uniform.

Routing

- **Drive Spreading:** for road routes, the relative increase in cost of an alternative route over the least cost route that an agent may consider acceptable in choosing a route [Sometimes referred to as the “perturbation” value.]
- **Walk Spreading:** as for Drive Spreading, but for walkway routes
- **Cost/ price:** the perceived value of money for any agent with this behaviour. This is applied to any tolls or other charges along the length of the route before these are included in the generalised cost equation.
- **Dynamic Routing:** If this is on, then any agent with this behaviour will receive cost feedback updates, if cost feedback is enabled in the appropriate routing window (Assignment → Vehicle Routes or Assignment → Person Routes). You might use this to model two sets of vehicles; one whose drivers receive updates during their trips (perhaps by radio) about the state of the traffic, and another whose drivers have a static view of the cost information on the network.

Temperament

- **Aggression, Awareness, Patience, Familiarity**: Values in the range 0 to 100%, by default 50%. These are not used in the core algorithms, but are available to the API for use in plug-ins. Some software packages use abstract parameters like these as indirect proxies for a set of other hidden parameters - headway, speed limit compliance, etc. However, in Commuter, those parameters are controlled directly and independently, without the need for a proxy.

Mobs

The Mobs window can be accessed by a button at the bottom of the Behaviours window, opened using **Parameters / Behaviours** then **Mobs...**

A Mob (Mix Of Behaviours) is a group of behaviours. A Mob can contain any number of behaviours, and Mobs can overlap. That is, a behaviour can be in any number of Mobs. When a new behaviour is created, a new Mob will be created automatically, with the same name, containing only that behaviour. You don't need to keep each Mob, but it is often useful to have one for each behaviour.

Restrictions are defined in terms of Mobs. If you want to create a restriction that applies to more than one behaviour, you should create a Mob that contains those behaviours.

For example, imagine you are modelling a hospital that has four behaviours: Doctor, Nurse, Patient and Visitor. If you wanted to create a “Staff Only” restriction, you would first create a Mob called “Staff” and include the behaviours Doctor and Nurse in that Mob.

Then you would define the restriction, selecting Staff as the Mob, and Allowed as the specification on the restriction.

Restrictions and Speed Controls

A **Restriction** is applied to a surface (a lane or a walkway), in order to control the types of people or vehicles that move on that surface. The people or vehicles are selected for a restriction by behaviour.

A related concept is that of a Speed Control which applies only to roads, and selects vehicles by behaviour

A Restriction has the following fields:

- **Name**: A numeric name for the restriction which is created automatically and is unique.
- **Description**: A descriptive name for the restriction, of any length, which need not be unique, but for clarity, a unique description is better.
- **Colour:** A Colour for the restriction, used when the appropriate Feature is enabled in the Layer Pane.
- **Term:** A time period defining when the restriction applies.
- **Mob:** A single behaviour or group of behaviours to which the Restriction applies.
- **Permission:** Either **Barred, Allowed** or **Mandatory.** This setting means that you can create a restriction to ***bar*** all behaviours in the mob or to ***allow*** all behaviours in the mob. Mandatory applies to lanes only, and forces all behaviours in a mob to use that walkway or lane, even if alternatives exist. Thus you can create a bus lane and make it mandatory for bus drivers to make sure that all buses drive in that lane.
- **Pocket:** (Roads Only) A distance from the end of any lane at which any vehicle may use the lane. For example, if the left hand lane is a bus lane, but pocket is defined as 50 metres, left-turning cars may use that lane for the last 50 metres before the stop line.

To create a restriction for a group of behaviours, you must first create a Mob to give a name to the selected group of behaviours. If you want to apply a restriction to a single behaviour, it is likely that a Mob containing only that behaviour already exists, as one is automatically created, with the same name, when a behaviour is created.

For example, to create a "Staff-Only" restriction

- In the Behaviours window, create all the behaviours you wish to model. For this example, assume you have six behaviours, A,B,C,D,E and F, but only types C & D are classified as "Staff"
- create a Mob named "Staff" containing behaviours C & D
- create a Restriction and change the Description to "Staff-Only"
- set the colour for this restriction
- if the restriction applies only at certain times, select a term for the restriction, otherwise leave this at "(Always On)".
- under "Mob", select "Staff"
- under "Permission", select "Allowed"

You could create an equivalent restriction by creating a Mob called "non-Staff", containing behaviours A,B,E & F, and then creating a restriction where Mob=non-Staff and Permission=Barred

Speed Controls

A Speed Control can be applied to a lane to change the speed limit on that lane, by user class, or by time of day.

Each has several fields:

- **Name**: A unique number, created automatically, not editable
- **Description**: A descriptive name, of any length, which need not be unique, but for clarity, a unique description is better.
- **Colour:** A colour, used when the appropriate Layer is enabled in the Layer Pane.
- **Term:** Start, end and day of operation.
- **Mob:** The behaviours to which the speed control applies.
- **Speed:** The speed limit that is applied

It is normally the case that you would set either a Mob or a Term for a Speed Control, or both, so that it applied only to some of the vehicles in the simulation, or only for a part of the simulation period. If neither of these are set, then the speed control will be applied at all times, to all vehicles, which is the same effect as changing the speed limit directly for that lane.

There are some standard speed controls defined, which have neither field set, but these are intended for use by tools or plugins, where some external logic would apply a speed control at a specific time. For example, the Lane Control plugin uses this functionality to implement managed motorways.

Person Types

Each person in the model has a Type. The Type specifies parameters which control the size, the movement and the display of each person in the model. Each person may be a driver or a passenger in a vehicle or be walking, as a pedestrian. For very simple models, only a single person type is required, using the Standard Behaviour, but most models will have several Types. In any model there should be at least one Type for each Behaviour, where a Type is a Person Type, a Vehicle Type, or a Freight Type. For each Behaviour you may want to define several physical types: for example, one type for lone people, another for parent+child, etc.

The Person Type parameters are edited in several tabs:

Description Tab

Name: A unique number, created automatically, not editable

Description: A textual description for the type, of any length, which need not be unique, but for clarity, a unique description is better. An example might be “Generalised Car” or “Small Truck”.

Behaviour: The Behaviour associated with this type of person. This defaults to the “Standard” Behaviour. Where there is more than one user class in the model, there should be one behaviour for each user class. The Behaviour defines the parameters used for the mode choice and route choice decisions within the simulation.

Colour: A colour for identifying vehicles of this type. If the colour is black, then the vehicle is coloured according to the layer.

Avatar Group: (Visualisation Only) Selects the 3-D shape used in the Detailed display mode. In some cases, you might want to define an avatar group for each trip purpose. For example when modelling a hospital, one person type, “Doctor”, might use an avatar group called “Doctors”, containing various 3D models of male and female doctors. Similarly for other person types, “Patient”, or “Nurse”.

Size Tab

Length, Width, Height: The mean dimensions of the agent (localized short distance units). To generate a distribution of sizes, see Size Variation. In general, regardless of any visualisation rendered onto the agent, a vehicle is modelled as a cuboid, and a pedestrian is modelled as an upright cylinder. Therefore, the width parameter for a pedestrian is always assumed to be the same as the length.

Mass: The mass (or "weight") of the person (kg or pound).

Space: The "comfort zone" around a person when walking. This is the preferred minimum distance to another person. People will temporarily accept smaller distances, but will try to move so that the this minimum distance is regained as soon as possible. This space is used between people, but is not used between a person and a wall.

Size Variation: A value specified as a fraction of each of dimension, used to specify the standard deviation of a normal distribution for each dimension. For example, if the mean length of a vehicle type is 4.0m and the size variability is 0.1 (10%), the lengths of vehicles of this type will follow an approximate normal distribution of mean 4.0 metre and standard deviation (SD) 0.4 metre. (In practice, the distribution is clamped to 3xSD around the mean. A true normal distribution would have approximately 0.3% of values outside 3xSD, but using such a clamping mechanism avoids unexpected values such as negative or zero lengths.)

Motion Tab

Maximum Speed (m/s): The mean maximum speed an agent of this type will attain when unconstrained by other agents or any speed limits on the current surface. You can apply a distribution of speeds by setting the Speed Variation field to a non-zero value, similar to Size Variation. This value is used by the routing algorithm. The route speed for a behaviour B is taken to be the lowest of:

- the route speed specified for a walkway
- the lowest maximum speed for ***any*** type using B

Speed Variation: A value specified as a fraction of the maximum speed, defining the standard deviation of a normal distribution, as with Size Variation above. (This parameter is not available for vehicles, as in almost all cases, vehicles are constrained by road speed limits before they are constrained by a physical maximum speed. If you want to define a range of maximum speeds for vehicles, you can use the “Speed Compliance” parameter in the Behaviour.)

Maximum Acceleration: The absolute maximum acceleration of this type (in m/s^2). This maximum value is modulated by the a function of the current speed and the gradient. The standard modulation function returns the absolute maximum value at half the maximum speed on a zero gradient, and reduced values at higher or lower speeds. The function uses a parabolic curve that reduces the value to 0.5x maximum acceleration at zero speed and 0.25x maximum acceleration at maximum speed.

Maximum Deceleration: The absolute maximum deceleration of this type, always specified as a positive value (in m/s^2). This maximum value is modulated by the a function of the current speed and the gradient. The standard modulation function returns the absolute maximum value at speed zero, reducing in a parabolic curve to 0.25x absolute maximum at maximum speed.

Attachment Tab

Controls: The leading agent has this field set, describing the type that it is pulling or controlling.

Attachment Angle: The angle of attachment, clockwise in degrees, of the attached type, relative to the direction of motion. For a child to the right hand side it would be 90, and -90 or 270 to the left. For an attachment that is pushed, such as a shopping trolley or a child's buggy, the attachment angle is zero.

Attachment Distance: the distance from the periphery of the leading type at which the attached type will pivot. For people this is normally arms length, about 0.5 to 1.0 metres

Attachment Count: This is set to 1 by default if there is an attachment, but it can be set higher if you want to repeat the attachment. For example, if you want to create a parent with two identical children, you can set this to 2.

Count: Set this to true for an Attachment that should be counted as a person. For example this would be **True** if the attachment is a child travelling with a parent, but **False** for a suitcase.

Person Types

Each person in the model has a Type. The Type specifies parameters which control the size, the movement and the display of each person in the model. Each person may be a driver or a passenger in a vehicle or be walking, as a pedestrian. For very simple models, only a single person type is required, using the Standard Behaviour, but most models will have several Types. In any model there should be at least one Type for each Behaviour, where a Type is a Person Type, a Vehicle Type, or a Freight Type. For each Behaviour you may want to define several physical types: for example, one type for lone people, another for parent+child, etc.

The Person Type parameters are edited in several tabs:

Description Tab

Name: A unique number, created automatically, not editable

Description: A textual description for the type, of any length, which need not be unique, but for clarity, a unique description is better. An example might be “Generalised Car” or “Small Truck”.

Behaviour: The Behaviour associated with this type of person. This defaults to the “Standard” Behaviour. Where there is more than one user class in the model, there should be one behaviour for each user class. The Behaviour defines the parameters used for the mode choice and route choice decisions within the simulation.

Colour: A colour for identifying vehicles of this type. If the colour is black, then the vehicle is coloured according to the layer.

Avatar Group: (Visualisation Only) Selects the 3-D shape used in the Detailed display mode. In some cases, you might want to define an avatar group for each trip purpose. For example when modelling a hospital, one person type, “Doctor”, might use an avatar group called “Doctors”, containing various 3D models of male and female doctors. Similarly for other person types, “Patient”, or “Nurse”.

Freight Types

Each unit of freight in the model has a Freight Type. The Type specifies parameters which control the size, the movement and the display of each freight unit.

The Freight Type parameters are edited in tabs, outlined below.

Description Tab

Name: A unique number, created automatically, not editable

Description: A textual description for the type, of any length, which need not be unique, but for clarity, a unique description is better. An example might be “Generalised Car” or “Small Truck”.

Behaviour: The Behaviour associated with this type of person. This defaults to the “Standard” Behaviour. Where there is more than one user class in the model, there should be one behaviour for each user class. The Behaviour defines the parameters used for the mode choice and route choice decisions within the simulation.

Colour: A colour for identifying vehicles of this type. If the colour is black, then the vehicle is coloured according to the layer.

Shape Group: (Visualisation Only – Detailed Display Mode) Selects a 3-D shape at random from the shapes making up this shape group. For example, a group named “Containers” might include several colours or texture map renderings of a shipping container.

Size Tab

Length, Width, Height: The mean dimensions of the freight unit (metres or feet). Each freight unit is modelled as a cuboid, all units of the same type are exactly the same size.

Mass: The mass (“weight”) of the freight unit (kg or pound).

Vehicle Types

Each vehicle in the model has a Vehicle Type. The Type specifies parameters which control the size, the movement and the display of the vehicle in the model. In any model there should be at least one Type for each Behaviour, where a Type is a Person Type or a Vehicle Type, or a Freight Type.

A group of Vehicle Types is called a **Fleet**. Fleets are described in more detail at the end of this section. The Fleets window can be raised using a button at the bottom of the vehicle types window.

The simplest of Commuter models can use a single private vehicle type, the standard car, measuring 4.0m x 2.0m. However, any model that includes goods vehicles will need at least one other type for a generalised truck.

More commonly, a model will have several vehicle types defined, a typical set would include:

- Small Car
- Medium Car
- Large Car
- Van
- Small Truck
- Large Truck
- Cycle
- Motorbike

The parameters for a vehicle type are edited in several tabs, as described below.

Description Tab

Name: A unique number, created automatically, not editable.

Description: A textual description for the type, of any length, which need not be unique, but for clarity, a unique description is better. An example might be “Generalised Car” or “Small Truck”.

Behaviour: The Behaviour associated with this vehicle type, which is used only if this vehicle is not being driven by a Person object. A vehicle will adopt a new behaviour when a Person gets into that vehicle as its driver. However, the behaviour associated with a vehicle is important for the following types of vehicles:

- Taxis
- Drop-Off Vehicles
- Pick-Up Vehicles
- Zone to zone (background traffic) vehicles

Colour: A colour for identifying vehicles of this type. If the colour is black, then the vehicle is coloured according to the layer. It is also possible to colour a vehicle by its driver type, using an option on the Tags window, accessed from the Display menu.

Solid Shape: This is used to set the 3D shape used in Solid display mode. If it is not set, a solid shape will be assigned based on the size of the vehicle and the name of any shape group

Shape Group: This is used to associate a type with a set of 3D shapes, in order that all instances of a type do not look the same in the Detailed display mode. When a vehicle of this type is created, it will select one of the shapes from the group at random.

Taxi: A taxi type can queue in a taxi rank, pick up a passenger at the head of the rank, and drop off a passenger at a drop-off zone.

Cycle: Cycle vehicles are counted separately to other vehicles.

Emergency: Emergency vehicles have different rules for lane use, and can exceed the posted speed limit.

Size Tab

Length, Width, Height: The mean dimensions of the agent (localized short distance units). To generate a distribution of sizes, see Size Variation. In general, regardless of any visualisation rendered onto the agent, a vehicle is modelled as a cuboid

Mass: The mass (“weight”) of the vehicle (kg or lb).

Size Variation: A value specified as a fraction of each of dimension, used to specify the standard deviation of a normal distribution for each dimension. For example, if the mean length of a vehicle type is 4.0m and the size variability is 0.1 or 10%, the lengths of vehicles of this type will follow an approximate normal distribution of mean 4.0 metre and standard deviation (SD) 0.4 metre. (The distribution is clamped to 3xSD around the mean to avoid unexpected values such as negative or zero lengths.)

Side Gap: the minimum gap to the side of a vehicle when passing another vehicle in the same lane

Extras: the number of extra people to count for this vehicle. It is not necessary to include a value here for the driver of a taxi or a drop-off or pick-up vehicle, as these are counted automatically. This field should be used for additional people, that you want to count in the Extras section of the trip summary statistics.

Load Base: The vertical distance up from the road surface defining the base of the load are, for freight or people carried by this vehicle.

Load Front, Back: The horizontal distances from the leading edge of the vehicle to the load, and from the trailing edge of the vehicle to the load.

Load Side: The horizontal distance in from the sides of the vehicle to the load area. *Set this to any non-zero value to display people*.

Load Capacity: The maximum mass of the load (kg or lb).

Dynamics Tab

Engine: The Engine defines the maximum speed, acceleration and braking rates, and the emissions (CO_2, NO*, PM10) for vehicles of this type. There are a number of standard engines defined, using data from the UK Transport Research Laboratory (TRL) but if you have data you wish to use, you can define your own engines.

Tyre Friction: This value controls the maximum speed at which a vehicle can turn a corner. The maximum speed is calculated from:

$$v_{max} = \sqrt{g\, u\, R}$$

where g = 9.81 m/s/s
u = the tyre friction co-efficient
R = the turn radius
The default value for this parameter (u) is 0.8

Gap Factor: This multiplicative factor is used to increase or decrease the gap to other vehicles *behind* vehicles of this type. You might use this to increase the distance of queuing vehicles behind large trucks. In effect, this parameter makes vehicles of this type appear longer than they are to other drivers.

Conflict Factor: This multiplicative factor is used to increase or decrease the cross and merge times where streams intersect, on a per-type basis. So you can use it to adjust the gap acceptance at intersections on a type-by-type basis. See also the Cross and Merge time parameters for a Stream.

Keep: the preferred side of the lane. Normally, this is set to **None**, which indicates that a vehicle will aim to drive along the centre of the lane in the absence of other traffic. However, for some vehicle types, such as cycles, you may want to set this to **Left** or **Right**

Attachment Tab

Controls: The leading agent has this field set, describing the type that it is pulling or controlling.

Attachment Angle: The angle of attachment, clockwise in degrees, of the attached type, relative to the direction of motion. For a trailing vehicle, this would be 180.

Attachment Distance: the distance from the periphery of the leading type at which the attached type will pivot. An articulated truck might behave best if this is set to a negative number, which makes the trailer overlap the cab. The total length of this cab-trailer combination is the sum of the lengths of the two types, plus the attachment distance. So for example, if the cab is 5 metres long and the trailer is 12 metres long, and the attachment distance is -1m, the total length will be 5 + 12 – 1 = 16 metres.

Attachment Count: This is set to 1 by default if there is an attachment, but it can be set higher if you want to repeat the attachment. For example, if you want to create a road train with one truck and 2 trailers, you should:

- define a Trailer type
- define an truck type with fields:
  - Controls = Trailer
  - Attachment Angle = 180
  - Attachment Distance = 0
  - Attachment Count = 2

Engines

An Engine defines a maximum speed, acceleration and braking rates, and emissions levels for a range of exhaust gases (CO2, NO*, PM10). Commuter contains a number of standard engine definitions, using data from the UK Transport Research Laboratory (TRL) but if you have data you wish to use, you can define engines according to your own specifications.

Each vehicle type has an associated engine definition. A simple model might define a single car type, with one of the standard petrol car engines, and a single truck type (“heavy goods vehicle” or HGV) with the standard diesel truck engine. A more detailed model might define six vehicle types: small, medium and large car types, a van or light goods vehicle type, and a two heavy goods vehicle types. It is impossible to say what the correct level of detail is for any particular model: it is up to you to decide how much detail you want, after deciding what is important for your particular area of study.

An Engine definition contains several tables of information, in “matrix” form:

- **acceleration rates:** values in m/s^2 for a range of speeds and gradients. Typically, a vehicle has its highest acceleration in the lower middle range of its speed. Also, the effect of gravity means that the maximum possible acceleration is higher on a down slope (negative gradient) and lower on an up slope (positive gradient). The standard engine definitions modify acceleration to take account of both these effects, but it is impossible to provide a definition that is correct for all circumstances, and you should verify that the engine specifications you use are of an acceptable accuracy for your model. In many cases, it is a comparison between two or more options that is important, and the absolute value of any particular output may be of lesser importance. It is also worth noting that the car following model employed will not necessarily use the maximum possible physical acceleration. Just because your car can do 0-100 km/h in 7.8 seconds does not mean that you floor the accelerator at every green light.

- **braking rates:** maximum braking or deceleration rates are almost always a function of the braking system on a vehicle, typically hydraulic pads onto wheel discs, and not a function of engine braking, which is used only for lower values, and when "coasting". However, the engine definition is a suitable place to store maximum deceleration values. These are specified in the same way as acceleration, in a matrix of speed by gradient.
- **exhaust emissions:** CO_2, NO_x and PM10 emissions are tabulated by speed and acceleration. The constant speed emission values for the standard engine definitions in Commuter are taken from a study commissioned by UK Transport Research Laboratory (TRL). The figures are *"Based on review and assessment of new factors for Euro I and II vehicles given in TRL Database of Emission Factors, September 2001 (Barlow, Hickman and Boulter). These factors refer to 'ultimate CO2', from all the carbon in the fuel emitted at the tailpipe as CO2, CO, unburned hydrocarbons and particulate matter which ultimately have the potential in forming CO2."* These figures have been extrapolated to positive and negative accelerations, assuming 50% more emissions at maximum acceleration and 50% less when decelerating. It should be stressed that these estimates for the change in emissions under the influence of acceleration are crude estimates, and were not included in the cited study.

Taxis

A taxi is a private vehicle that has some additional capabilities:

- It can be used by any person travelling to a parking zone marked as "drop-off"
- It can park in a taxi rank, a special type of parking lane.
- After dropping off passenger(s) at a drop-off zone it will go to the nearest taxi rank with spaces, or if none exists, then it will return to the originating zone.

If taxis are available, a taxi may be used by a person to travel between two parking zones, if the expected cost of the taxi journey is less than other alternative travel modes, when weighted according to the parameters defined for the behaviour associated with the taxi.

If you are going to use taxis, we would recommend you define a behaviour specifically for taxis, and set the drive cost parameters for distance, time and base according to your local taxi fare parameters.

Example 1, using rates for Sydney, 2009:

- base cost = $3.10: enter as 310 (cents)
- distance cost = $1.85 / km: enter as 185 (c / km)
- time cost = 80c / minute: enter as 1.333 (c / second)

Example 2, using rates for London, 2009:

- base cost = £2.20 enter as 220 (pence)
- distance cost = 20p / 143.7m: enter as 139.2 (pence / km)
- time cost = 20p / 30.9s: enter as 0.647 (pence / second)

Unfortunately, there is no universal method of taxi fare calculation – some methods charge simultaneously for distance and time, others charge for time only if the speed is less than a certain threshold. The algorithm in commuter will always calculate the sum of costs for distance and time, so if this is not the method used in your locality, you may want to reduce the values to compensate for this, perhaps even setting the time cost to zero.

When calculating the cost of any trip, including taxi trips, it is possible to take the cost of the trip outside the area of the network in our model. For example, if the cost of taking a taxi to the airport is compared to the cost of parking at the airport, it is necessary to have distance and time values for the whole journey, to compare against the parking cost. For this reason, each cordon area can hold values for the off-network distance, time and price associated with the travel between the area and the ultimate origin or destination. There is also a facility for entering a standard deviation value, so that the distance and time values follow a normal distribution.

The expected cost of a trip in a taxi is calculated using distance and time for the trip, both off-network and on-network.

The taxi fare is calculated as

```
fare = {base_cost} + {off_network_cost} + {on_network_cost}
```

where

- base_cost is a parameter for the taxi behaviour,
- off_network_cost is calculated from the stated off-network {time, distance, price} parameters for origin and/or destination areas, multiplied by the taxi behaviour driving cost parameters.
- on_network_cost is calculated from the network-derived {time, distance, price} values for the expected route, multiplied by the same taxi behaviour driving cost parameters.

Fleets

The Fleets window can be accessed by a button at the bottom of the Vehicle Types window, opened using **Parameters / Vehicles** then **Fleets...**

A Fleet is a group of vehicle types. A Fleet can contain any number of vehicle types, and Fleets can overlap. That is, a vehicle type can be in any number of Fleets. When a new vehicle type is created, a new Fleet will be created automatically, with the same name, containing only that vehicle type. You don't need to keep each Fleet, but it is often useful to have one for each vehicle type.

Lane choice, stream choice and route choice rules use fleets,as do some types of reporting functions, such as those for Loops. If you want to create a lane choice rule that applies to more than one vehicle type, you should create a Fleet that contains those vehicle types.

For example, imagine you have a model that has five vehicle types: Small Car, Large Car, Taxi, Bus, Truck. If you want to create a rule that applies only to cars, then you would create a fleet containing those two vehicle types, perhaps calling the fleet “Cars”.

Then you would select that fleet in the Lane Choice rules window.

Crossing Times

A Crossing Time definition specifies “Walk”, “Flashing” and “Don't Walk” times for pedestrian signals at crossings, and can also have a Term, or time interval, over which the definition applies.

Walk (W): during this part of the cycle, pedestrians have uncontested priority on the crossing. This state is often signalled by a solid green signal, either a walking figure or text. At an intersection, you can choose to allow turning vehicles to conflict with a pedestrian movement (depending on the road rules for your country) in the way that you define the groups (and perhaps also phases) in the definition of your controller. In this situation, a turning vehicle that approaches a crossing on Walk or Flashing will wait at the edge of the crossing until pedestrians have finished crossing its lane.

Flashing (F): during this part of the cycle, pedestrians have priority on the crossing if they have started to cross, but will not start to cross. otherwise it is the same as the Walk state. In some countries, this state is signalled as a flashing red light, while in others it is signalled as a flashing green light. The time for this part of the cycle is often set according to the length of the crossing, as it is the “clearance” time required for all pedestrians to finish crossing.

Don't Walk (DW): during this part of the cycle, pedestrians do not have priority on the crossing, and any vehicles crossing it will not look for pedestrians. This state is often signalled by a solid red signal, either a standing figure or text. The Don't Walk time is used to define the minimum period between activations, and is used only for timed signals that are not linked to an intersection. For example, if the crossing is push-button activated, and you define a Don't Walk time of 80 seconds, the time between the end of the flashing signal and the next walk signal will be at least 80 seconds, even if a pedestrian arrives within that time. After the end of the 80 second period, if a pedestrian has arrived, the crossing will go to walk immediately. otherwise, the crossing will remain in the Don't Walk state until the arrival of the next pedestrian.

Term: the timing for a crossing may vary over the day – for example, a timed crossing may come on more frequently during rush hour, before or after school hours, etc. For this reason, you may want to define more than one crossing time object for any one crossing. For example, you might define a peak crossing time object of [30, 10, 80] (total cycle time = 120 sec) for peak hour, and another of [30, 10, 140] (total cycle time 180 sec) for inter-peak. On crossings which have an associated group at an intersection, Crossing Times values are used to put an upper limit on the walk time for a crossing, but the controller has ultimate control over the state of the crossing. That is, even if you set the walk time to 60, the controller can terminate the group at any time before that. A Controller termination causes the crossing to go to the “flash” state.

To add multiple crossing times to a crossing: First, define all the crossing times you require in the Crossing Times window. Then, select the crossing, then choose Adjust. In the Timing column, first select '*' crossing time in the drop down box. The Multiple Crossing Times window will appear to enable a combination of crossing time distributions to be selected.

Application to Route Choice

Crossing time values are used by the routing algorithm on some types of crossing as an estimate of the expected wait time.

The expected wait times at crossings are:

- Phased Crossing: the intersection cycle time
- Pelican crossing: F + DW
- Zebra crossing: zero
- Unmarked crossing: F + DW

Time Distributions

A time distribution is a general-purpose method of defining a discrete distribution of whole-second times. For example, a distribution of stop times at a toll plaza might be as follows:

Probability	8%	15%	13%	9%	8%	7%	6%	5%	5%	4%	8%	7%	5%
Time (sec)	1	2	3	4	5	6	7	8	9	10	12	15	30

While this is not as smooth as a continuous distribution, remember that a time-interval simulation such as Commuter will often round a time period to a multiple of the simulation interval, so a discrete time distribution is often a practical solution.

The Time Distributions window allows you to define as many **Discrete** times as you require for your distribution, each in a column, and then you enter the probability for that discrete time in the relevant cell in the table. If you visualize the distribution as a bar chart, the number of columns is the number of bars, and the value in each cell is the height of the bar.

Press **New Time Column** to add a new discrete time interval. The time must be a whole number of seconds. The value can be entered in the form S, M:S, or H:M:S.

Press **Add...** to add a new distribution. The sum of all values cannot be greater than 100, so if values currently sum to 100, you will need to type a lower value over an existing value before continuing. it may be easier to set values in all columns to zero first, then enter the values you want. You can press Apply to validate that your values add to 100. If you values sum to less than 100, the value in the first column corresponding to zero seconds will be adjusted to make the total up to 100.

Time distributions can also be entered in **Gaussian** form with a **Mean** and **Standard Deviation** values

Calibration

The Calibration window allows you to modify parameters used for the agent dynamics models, both pedestrian and vehicle.

The **Pedestrian** Calibration Parameters are

- **Preferred Side:** The side of the walkway that pedestrians will tend towards in the face of opposing pedestrian flow
- **Grid Size**: This is used for the Pedestrian collision algorithm. A larger grid size means that each Person will “see” further, at the expense of more computation. If the Grid Size is too small, there may be too many grid cells for a very large network. The best setting for this varies by network size and complexity.
- **Pedestrian Vision Range:** the distance each pedestrian can “see” ahead, when looking for other pedestrians or obstacles to avoid. Increasing this value will cause the simulation to run more slowly, as more area must be scanned for possible obstacles
- **Pedestrian Vision Angle:** The field of vision angle ahead of a pedestrian, used when looking for other pedestrians or obstacles.
- **Leader Search Max Vision Angle:** The field of vision angle used when looking for another pedestrian to follow
- **Leader Search Max Delta Bearing:** When looking for another pedestrian to follow, a possible leader will be rejected if its current heading differs by more than this angle (a pedestrian will follow another only if it is going in the same direction)
- **Leader Search Max Distance:** When looking for another pedestrian to follow, a possible leader will be rejected if it is further away than this distance.

- **Following Time To Equilibrium:** When following another pedestrian, this value is used to set the acceleration used to match the speed of the leader

- **Following Target Headway:** When following another pedestrian, the follower will aim for a point at this distance behind the leader.

- **Crossing Occupation Footprint:** The distance ahead and behind a pedestrian that is “closed” to traffic as the pedestrian crosses

- **Release Slot Width:** The width of the “slots” into which pedestrians are released at the end of walkways

- **Avoidance Radius Factor:** Another pedestrian will be avoided only if it is within a circle of radius that is larger than the space radius by this factor. For example, if this factor is 4.0 and the space radius for this person is 0.8, then another pedestrian will be avoided if its centre is closer than 3.2

- **Avoidance Sideways Speed:** If another pedestrian is to be avoided, this is the maximum sideways speed used in the avoidance manoeuvre

- **Avoidance Delta Bearing:** If another pedestrian is to be avoided, this is the steering angle used in the avoidance manoeuvre

- **Route Choice Model:** Where there are two possible routes to a destination, with costs c_1 and c_2, this model determines which route will be chosen.
  - The **Perturbation** model adjusts c_1 and c_2 by adding a random noise value from a uniform distribution (-P, P), then selects the lower of the two adjusted costs.

  - The **Logit** model chooses a random value R from a uniform distribution [0, 1) and selects the higher cost route if R > P(c_2), where

$$P(c_2) = \frac{1}{1 + e^{\frac{c_2 - c_1}{c_1}}}$$

The **Vehicle** Calibration Parameters are:

- **Drive Side:** The side of the road used for driving, according to national rules for country of model

- **Car Following**: The default car following model. This can be modified by route class and also by lane, if there is a requirement to use multiple car-following models.
  A sub -menu allows modification of the parameters for the selected car-following model.

- **Lane Changing:** The default Lane Changing Model.
  A sub -menu allows modification of the parameters for the selected lane-changing model.
  See also Commuter5 Lane Changing Parameters.

- **Avoidance**: Switch the collision-avoidance functions on everywhere in the model. This requires much more computation of collision vectors, and will cause a significant slow down in the speed of the model for larger networks. It is possible to switch collision avoidance on for individual roads, which may be a better solution if there are one or two areas where traffic conflicts with pedestrians.

- **Unreleased Store**: Normally **On,** if this is **Off,** then vehicles will be generated even if there is no room on the lane for them and they will form a queue back off the end of the lane.

- **Stream Friction on Merge:** When On, this prevents a vehicle from advancing on an intersection if another vehicle, with a higher priority, is approaching or crossing that intersection using an adjacent stream that does not conflict. Although there is no actual conflict, switching this parameter off tends to over-estimate the capacity of an intersection.

- **Route Choice Model:** Perturbation or Logit, see Pedestrian Calibration section for explanation

- **Curve Resolution:** Not specifically for vehicles, this defines the number of points in each Bezier curve. It may be reduced in very large road networks, where detailed curves on every lane can cause computing performance problems on cheap computers with limited memory.

- **Stop Definition Lower Speed & Higher Speed:** A stop is defined as the sequence of events when a vehicle is observed travelling at more than the higher speed and then observed travelling at less than the lower speed. If these two values are equal then a vehicle travelling at about that speed may record many stops as it is observed travelling slightly faster then slightly slower.

- **Leader Search Max Vision Angle:** The field of vision angle used when looking for another vehicle to follow

- **Leader Search Max Delta Bearing:** When looking for another vehicle to follow, a possible leader will be rejected if its current heading differs by more than this angle (a vehicle will follow another only if it is going in the same direction)

- **Leader Search Max Distance:** When looking for another vehicle to follow, a possible leader will be rejected if it is further away than this distance.

- **Standard Signposting Distance**: Signposting is used to look ahead from intersection for the correct lane choice. It can also be set on a per-road basis for intersection approaches.

- **Avoidance Delta Bearing:** If another vehicle is to be avoided on an intersection, this is the steering angle used in the avoidance manoeuvre

Commuter5 Lane Changing Parameters

Look Time & Look Distance: When looking for a gap in an adjacent lane, these values (call them T and D) specify the maximum distance for the search, using the larger of vT and D, where v is the current speed of the vehicle. Any other vehicle that is further than this distance, either ahead or behind, will be ignored.

Minimum Speed: A lane change will not be attempted unless the vehicle is travelling at the minimum speed or more.

Cut-off Speed & Speed Ratio: (Call these C & R) This vehicle will attempt to change lane to pass its leader if it is travelling at more than C and the leading vehicle is travelling at less R.C

Maximum Speed Ratio: (Call this M) If the R.C condition above is not true, but this vehicle's maximum speed is more than M times the maximum speed of the leader, then this vehicle will attempt to change lane to pass its leader

Standing Distance: If a vehicle is less than this distance from the stop line, then lane changing will not be attempted, unless the vehicle is outside it preferred lane range.

Force Change Distance: If a vehicle is less than this distance from the end of its lane, and there is no valid exit from that lane, then its lane changing will be given higher priority

Last Change Distance: If a vehicle has travelled less than this distance forward since its last lane change then it will not attempt to change lane again, unless it is outside its preferred lane range

Adjacent Advantage Distance: If a vehicle can see more than this free advantage distance in an adjacent lane that is also within its preferred lane range, then it will try to change to that lane. If this distance is too short, or zero, then vehicles may “hunt” between lanes.

Term Variations for Parameter Values

A Term Variation is used to set a new value for a parameter at a particular time during the simulation run. For example, it could be used to set a new time distribution for delay on a walkway during peak hour, or a change to the space factor at times when it is known to be crowded.

Term Variations can be used for any network object, for any field that can be modified using the “Adjust Parameters” window.

To specify a Term Variation:

- select one or more Network objects (e.g. Lane, Walkway)
- raise the Adjust window
- select a row, or select Apply To All
- If it is possible, also select the field to which the Term Variation will be applied. For some types of fields, such as on/off fields, it is not possible to select the column, but it does not matter, the field can be selected at the next stage.
- Now press the Term Variation button.

In the window that is raised

- select a Term
- verify or select the parameter that you want to change
- type the new value for that parameter. Make sure you spell the value correctly, if it is a name
- If you want the parameter to revert to the original value at the end of the term, click the Revert box.

Term Variations can be modified or deleted using the window raised from **Parameters > Term Variations.** However, it is not possible to add a new Term Variation from this window, you must do this using the Adjust window.

Importing Parameters

All parameters can be imported from another model.

This includes:

Terms
Behaviours
Mobs
Types (people, vehicles, freight, transport) & Engines
Divisions
Fleets
Calibration parameters (Pedestrian and Vehicle)
Restrictions & Speed Controls
Crossing Times
Patch Definitions
Assignment – Route Classes
Assignment – Route Class Behaviours

One of the best ways to ensure that all models are built using the same parameters is to save a standard model in the central location, and import it when a new model is created.

The suggested central location is

(Installation Folder)/Standard_***.aza

For example

C:/Azalient/Commuter/Standard_NSW.aza

Trails

A Trail is a sequence of route elements (links or walkways) used to control routing. A Public Transport Service is a route definition specified as a sequence of Trails; each Public Transport Vehicle follows the Service to which it has been assigned.

Introducing the Trail object into the definition of Services (bus, tram or train routes) often allows us to create multiple Services more quickly. It is often the case that in a model with multiple services, several will use a common sequence of links through the centre of the model. With the system used in Commuter, you would define the common sequence of links once, as a Trail, and then re-use that Trail in all the Services.

To add a new Trail:

- select the link where you want the Trail to start
- open the Trails window
- press “New...” to start a new trail
- the trail will be highlighted in yellow, and the view will move to the first decision point on the trail, and prompt you with a choice of colour
- press one of the coloured buttons to select which of the exits you want to follow: press red for the red exit, green for the green exit and so on
- the view in the main window will then move to the next decision point on the trail, and prompt you again with a choice of colour
- you can select “Advance Step = Single” to move along a trail by a single link at a time. The default setting “Advance Step = Multiple” always moves the end of the trail to the next decision point.
- press the “Delete” (Red X) button to reverse back along the trail if you have gone too far.
- continue pressing coloured buttons until you reach the end of the trail
- when you reach a dead end, you have the choice of accepting the trail as ending on the current dead-end link, or you can press delete to reverse back

To Extend an Existing Trail:

- Select a Trail by its name in the Trails window
- Press “Extend >” to add to the end of the Trail or “Extend<” to add to the beginning
- Add links to the trail following the instructions for adding a new Trail

To Shorten (Prune) an Existing Trail:

- Select a Trail by its name in the Trails window
- Press “Prune”
- Choose to remove the first link on the Trail or the last link

You can Join two existing trails to make a new third trail. The second trail must start on the link immediately following the last link in the first trail.

To Join Two Adjacent Trails:

- select the first Trail by its name in the Trails window
- if the “Join...” button is enabled, then there is at least one valid trail which follows on directly from the selected one
- press the Join button
- select the trail that you want to form the second part of the new trail

To Delete a Trail:

- select the trail by name in the table
- press Delete

Offset Calculator for Signalized Intersections

This function calculates and applies simple timing offsets for any number of signalised intersections along a corridor.

- Create a trail along the corridor starting at the first intersection and ending at the last
- Select the trail in the Trails window
- Press the **Offset** button. (The first link of the trail must start at a signalised intersection.)
- Offsets will be calculated and displayed. Press **OK** to apply the timing offsets to the intersections, or **Cancel** to discard.

Services

A Service defines the route taken by a public transport vehicle, which might be a bus, a tram or a train, and also defines the Stands, or locations where the vehicle will stop to pick up or set down passengers.

Associated with each service are one or more transport trips, defining the departure time and type of each vehicle on that service. A network model seldom covers the full extent of all services, so in many cases, the “departure time” is actually the time at which the transport vehicle first enters the network. If the vehicle would contain passengers at this point, then they should be added using a “Transition” Stand on the first link of the service route.

Before Creating a new Service: First you must create at least one Trail, using the Trails window, then you will use one or more Trails to define the route for each new service. This two-stage process is slightly more time-consuming for the trivial case of a single Service, but in a typical network, it saves a lot of time in the long run, because you can re-use a trail in many different service routes.

Important: If you are creating trails to make a new service you must confirm the new trails by pressing “OK” in the Trails window ***before*** starting to create a new service.

Important: A service can be made up of any number of trails, but the trails must not overlap, and must have no gaps. Thus if a service is made up of two trails, the first trail must end on the same node that is the start of the second trail.

To Create a new Service:

- Press “New...” in the bottom left corner to start a new Service
- In the resulting Select Trail window use the drop-down menu to select a trail and press OK
- Enter a Name for the Service when prompted
- On the last link of the trail, the software will look for any other follow-on trails, and if any exist, will offer you a choice of next trail.
- if there are no follow on trails, press OK to accept the end of the service at the current position
- if there are follow-on trails, press one of the coloured buttons to select which of the exits you want to follow: press red for the red exit, green for the green exit and so on
- if there is only one possible trail, then there will be only one coloured button. (When building services, the extension mode is always in “single advance” mode.)
- the view in the main window will then move to the next decision point, and prompt you again with a choice of colour
- press the “Delete” (Red X) button to reverse back along the service if you have gone too far.
- continue pressing coloured buttons until you reach the end of the service
- once the sequence of trails matches the route you want for the new service, press OK to accept it.
- The new service will be added at the end of the list. If necessary, scroll down to reveal it. If necessary change the colour of the new service by clicking on the cell in the Colour column next to the service name

To Extend an Existing Service:

- Select a Service by its name in the Services window
- If there are any follow-on trails suitable for use as an extension of this service, then the Extend button will be enabled
- Press “Extend...”
- Select a Trail using the coloured buttons. If only one trail is available then only one button will be active.

To Shorten (Prune) an Existing Service:

- Select a Service by its name in the Trails window
- If a Service is comprised of two or more Trails then the Prune button will be enabled. (You cannot prune individual links from a service, only trails. To prune individual links from the last trail in a service, use the Prune function in the Trails window.)
- Press “Prune...”
- Choose to remove the first Trail or the last Trail
- Repeat as required, then press OK

To cause service vehicles to halt at a Stand:

- In the Service window, select the service you want to modify
- The table in the centre of the window will show all the stands on the route of this service.
- Click to put a tick next to each stand at which you require the transport vehicle to stop.
- When you are finished press Apply under the table, then press OK.
- You can sort the table alphabetically by stand name, by clicking on the column header.

To Add a Layover at a Stand (A Layover defines a variable dwell time at a stand, to bring that vehicle back to its timetable.)

- In the Service window, select the service you want to modify
- The table in the centre of the window will show all the stands on the route of this service.
- Double-click in the Layover column
- Press the [+] button
- In the New Layover window, first select the Term, during which this Layover will apply
- Then set the time interval (hours, minutes, seconds) after the departure time from the first stand at which the bus or train should leave this (layover) stand. For example, if you want the bus or train to leave the layover stand 15 minutes after its departure time from the first stand on the route, then enter 00:15:00 in the first box.
- Press OK

To add a departure time for a Service

- In the Service window, select the service you want to modify
- The table on the right hand side shows all the transport trips for this service, defining departure time and vehicle type for each trip.
- After you have added all the transport trip data, press Apply to save this information in the Service.
- You can sort the table by any of the columns, by clicking on the table header.

Stands By Service Window

The Stands By Service window shows a table with all Services as rows, and all stands as columns. Each cell of this table can hold one of three icons:

- [Pause] this stand is on the route of this service; all vehicles on this service will halt at this stand
- [Right Arrow] this stand is on the route of this service; all vehicle on this service will **not** halt at this stand
- [Blank] this stand is **not** on the route of this service

Clicking once on a cell will move the cursor in the Graphics Panel to be above the stand corresponding to that cell.

Double-clicking on a cell that contains an icon will toggle the state at that stand. That is, if the icon shows that the vehicles on a service do not currently stop at a stand, you can make them stop by double clicking on the icon.

Hovering over a cell will raise a tool-tip window that shows some information about that service/stand pair.

Stands are listed in the order in which they were originally added. You can re-order them in the table by dragging the table header for each column, but this re-ordering is temporary.

Sectors

A Sector is a group of Zones or a group of Areas. A Sector of Zones can be a mixture of Parking, Transition and Vehicle Zones.

A Sector can be used to match groups of zones or areas with rules for choosing lane, route or destination, or for defining tags on vehicles.

You can assign a colour to a sector, as a visual check to make sure you have grouped the zones you intended.

To add a new Sector: select all the areas or zones required for your sector in the main window. Then press the “[+] button, at the top left corner of the Sectors window. This will raise a window showing a list of areas or zones that were selected. You can include or exclude from the sector using the horizontal arrows, then press OK. After a sector has been added you can edit the name and colour as required

To change the zones or areas in a Sector: Select the Sector in the Table and press the **Edit** button to modify the list of Areas and/or Zones.

The **Parking** option for a sector can be used to mark a set of parking zones as *equivalent.* If two parking zones are equivalent, then no vehicle trips will be considered between those two zones.

The **Display By Colour** option highlights all the

(Public) Transport Types

Each public transport vehicle (or simply “Transport”) in the model has a Type. The Type specifies parameters which control the size, the movement and the display of the vehicle in the model.

These types are used to model public transport vehicles that move on pre-defined routes, or “Services”. One type is required for each size of public transport vehicle. A simple model including buses may find a single bus size adequate, perhaps 10m x 2.5m x 2.5m, but for more detailed models, you may want to represent buses, trains and trams of various dimensions.

The Transport Type information can be accessed from the Services Window

The parameters are edited in a number of tabs:

- **Description:** description, colour and behaviour
- **Size:** length, width, height, weight, variation, space and gaps
- **Dynamics:** engine, factors for gaps, turning etc.
- **Dwell**: dwell times at stands
- **Attachment:** for multi-part types, such as parent & child, or truck and trailer

Description Tab

Name: A numeric name for the type which is created automatically and is unique.

Description: A textual description for the type, of any length, which need not be unique, but for clarity, a unique description is better. Examples are “Standard 45-Seater” and “Bendy Bus”.

Behaviour: The Behaviour associated with this type. By default, there a single “Standard” Behaviour defined, but you can define multiple behaviours if required. The Behaviour encapsulates the behavioural aspects of the driver of this vehicle, so if it is a bus type, it is recommended to create a “Bus Driver” behaviour, and use that.

Solid Shape: This is used to set the 3D shape used in Solid display mode. If it is not set, a solid shape will be assigned based on the size of the vehicle and the name of any shape group

Shape Group: This is used to associate a type with a set of 3D shapes, in order that all instances of a type do not look the same in the Detailed display mode. When a vehicle of this type is created, it will select one of the shapes from the group at random.

Size Tab

Length, Width, Height: The mean dimensions of the agent (localized short distance units). To generate a distribution of sizes, see Size Variation. In general, regardless of any visualisation rendered onto the agent, a vehicle is modelled as a cuboid, and a pedestrian is modelled as an upright cylinder. Therefore, the width parameter for a pedestrian is always assumed to be the same as the length.

Mass: The mass (“weight”) of the vehicle (kg or lb).

Size Variation: A value specified as a fraction of each of dimension, used to specify the standard deviation of a normal distribution for each dimension. For example, if the mean length of a vehicle type is 4.0m and the size variability is 0.1 or 10%, the lengths of vehicles of this type will follow an approximate normal distribution of mean 4.0 metre and standard deviation (SD) 0.4 metre. (The distribution is clamped to 3xSD around the mean to avoid unexpected values such as negative or zero lengths.)

Side Gap: the minimum gap to the side of a vehicle when passing another vehicle in the same lane

Capacity: The total passenger capacity of a transport vehicle, including standing capacity.

Standing: The capacity of a transport vehicle for standing passengers. The seating capacity of the vehicle is derived from the total capacity minus the standing capacity.

Doors: the number of **exit doors** on the vehicle. This may be different to the number of entry doors, which is set on the stand, as this affects how the passengers queue at the stand. There are two numbers to provide finer control over boarding and alighting times, and also because some transport vehicles have fewer doors for entry, for example if all boarding passengers have to board at the front door to pay the driver.

Extras: the number of extra people to count for this vehicle. It is not necessary to include a value here for the driver of a taxi or a drop-off or pick-up vehicle, as these are counted automatically. This field should be used for additional people, that you want to count in the Extras section of the trip summary statistics.

Load Base: The vertical distance up from the road surface defining the base of the load are, for freight or people carried by this vehicle.

Load Front: The horizontal distance from the leading edge of the vehicle to the load, defining the front edge of the load area.

Load Back: The horizontal distance from the trailing edge of the vehicle to the load, defining the back edge of the load area.

Load Side: The horizontal distance in from the sides of the vehicle to the load area. **Set this to a non-zero value to display passengers.**

Load Capacity: The maximum mass of the load (kg or lb).

Dynamics Tab

Engine: The Engine defines the maximum speed, acceleration and braking rates, and the emissions (CO_2, NO*, PM10) for vehicles of this type. There are a number of standard engines defined, using data from the UK Transport Research Laboratory (TRL) but if you have data you wish to use, you can define your own engines.

Tyre Friction: This value controls the maximum speed at which a vehicle can turn a corner. The maximum speed is calculated from:

$$v_{max} = \sqrt{g\,u\,R}$$

where g = 9.81 m/s/s
u = the tyre friction co-efficient
R = the turn radius
The default value for this parameter (u) is 0.8

Gap Factor: This multiplicative factor is used to increase or decrease the gap to other vehicles *behind* vehicles of this type. You might use this to increase the distance of queuing vehicles behind large trucks. In effect, this parameter makes vehicles of this type appear longer than they are to other drivers.

Conflict Factor: This multiplicative factor is used to increase or decrease the cross and merge times where streams intersect, on a per-type basis. So you can use it to adjust the gap acceptance at intersections on a type-by-type basis. See also the Cross and Merge time parameters for a Stream.

Keep: the preferred side of the lane. Normally, this is set to zero, which indicates that a vehicle will aim to drive along the centre of the lane in the absence of other traffic. However, for some vehicle types, such as buses which drive in shared bus & cycle lanes, you may want to set this to -1, to indicate driving to the left side of the lane, or +1 to prefer the right side.

Dwell Section

Boarding Times: A named time distribution defining the spread of boarding or loading times for a transport vehicle of this type. See Parameters / Time Distributions. The number of **entry doors** is also significant in the calculation of total boarding time, which is defined on the stand. The total boarding time is calculated from the sum of the individual boarding time for all boarding passengers, divided by the number of entry doors.

Alighting Times: A named time distribution defining the spread of alighting or unloading times for a transport vehicle of this type. See Parameters / Time Distributions. The number of **exit doors** is also significant in the calculation of total boarding time, which is defined in the Size tab. The total alighting time is calculated from the sum of the individual alighting time for all alighting passengers, divided by the number of exit doors.

Load Mob: A mob defining the behaviours that can board this transport vehicle. This can be used to model different ticket types, which may not be valid on all services, or different freight types, which cannot be loaded onto all types of vehicles.

Attachment Section

Controls: The leading agent has this field set, describing the type that it is pulling or controlling.

Attachment Angle: The angle of attachment, clockwise in degrees, of the attached type, relative to the direction of motion. For a trailing vehicle, this would be 180.

Attachment Distance: the distance from the periphery of the leading type at which the attached type will pivot. An articulated truck might behave best if this is set to a negative number, which makes the trailer overlap the cab. The total length of this cab-trailer combination is the sum of the lengths of the two types, plus the attachment distance. So for example, if the cab is 5 metres long and the trailer is 12 metres long, and the attachment distance is -1m, the total length will be 5 + 12 – 1 = 16 metres.

Attachment Count: This is set to 1 by default if there is an attachment, but it can be set higher if you want to repeat the attachment. For example, if you want to create a train with one engine and 8 wagons, you should:

- define a Wagon type
- define an Engine type with fields:
    - Controls = Wagon
    - Attachment Angle = 180
    - Attachment Distance = 0
    - Attachment Count = 8

Engines

For a description of Engines, see Vehicle Types.

Intersection / Controller

An intersection is a set of one or more nodes where there is conflict between traffic streams. An intersection may be signalised or unsignalised. If it is signalised, then it will have a Controller.

Unsignalised Intersection

An unsignalised intersection can be edited by selecting any part of the intersection surface and then selecting the **Adjust** action. It is easier to select the intersection if Streams are not displayed.

The Intersection Editor for an unsignalised intersection has a single tab containing all the turns on that intersection.

Unsignalised Turn: a Turn is a collection of Streams, all having common entry and exit *roads*, where a Stream is the connection between an entry and an exit *lane* across an intersection. The parameters for each Turn on a Unsignalised Intersection are:

Fixed Signal: The sign or signal visible to traffic on the approach to this turn. Select from the following options, **listed in decreasing order of priority:**

- **Free Flow:** Approaching traffic is on the main road, has priority, and does not need to check for conflicting traffic.
- **Yield:** Approaching traffic is on the main road, but must cross an opposing stream, and must slow down on approach.
- **Give way:** Approaching traffic has a Give Way sign, and slows down accordingly
- **Stop Sign:** Approaching traffic has a STOP sign, and slows down to a stop before proceeding
- **Barred:** The turn is closed to all traffic

Restriction: a restriction for the turn, for example to allow buses and taxis to make a turn, but not other traffic.

Preference: (For Turns inside a Parking Zone) Set this to 1 for the preferred turn for a vehicle that is looking for a parking bay.

Signalised intersection & Controller

A Controller is a device that holds information about the signal control at an intersection. There is one controller per intersection, and one intersection per controller. However, this is not really a limitation, because an intersection in the model is defined to be a group of one or more nodes, and does not necessarily have to be limited to a single *physical* intersection.

The key limitation is that any signalised movement may be controlled directly by one controller only. This does not exclude the possibility of controllers communicating in order to synchronise, and in that sense a controller may influence the setting of a signal belonging to a neighbouring controller.

A Controller retains a record of all Turns, Groups, Phases, Plans and Rules for an intersection, and all of these can be edited using the Intersection / Controller Editor window. The Controller is also aware of a list of Loops that are connected to the intersection. However, a Loop “belongs” to a Lane, not to a controller, so if you want to add a Loop, you should do this by first selecting a Lane. You can edit the properties of a Loop by calling the Adjust action on a selected Loop.

For a signalised intersection, it is often easier to select and adjust the parameters by selecting the controller. This can be made visible using the Controllers Layer.

Signalised Turn: a Turn is a collection of Streams, all having common entry and exit *roads*, where a Stream is the connection between an entry and an exit *lane* across an intersection. The parameters for each Turn on a Signalised Intersection are:

- **Fixed Signal:** set if this turn does not have a variable signal, but has a signal fixed for all time. An example is where a turn to the kerb side has a Give Way or Yield sign.
- **Direct Group:** unopposed movement - vehicles given a green signal will proceed without checking other traffic streams. This is sometimes known as the “primary”.
- **Filter Group:** opposed movement – vehicles given a green signal will proceed if there is a suitable gap in opposing traffic. This is sometimes known as the “secondary”.
- **Filter Signal:** if the turn has a Filter Group, then you can set the signal that is applied to the turn when the Filter Group is active. Normally, the Filter Signal is Green Yield, which causes vehicles to slow before turning across an opposing traffic stream, if a suitable gap exists. You can also set it to Green Stop, which causes vehicles to come to a complete stop before turning when the Filter Group is on red. This allows you to code Kerb-Turn-On-Red signals.
- **Restriction:** a restriction for the turn, for example to allow buses and taxis to make a turn, but not other traffic.

In almost all cases, there is only one Turn for each entry-exit pair of roads. However, occasionally, an entry-exit pair will have two turns (sometimes known as “auxiliary” group). An example is where there is a special bus-signal for a movement, and this would have its own Turn to differentiate it from other traffic streams making that movement. To create this type of extra Turn, select the Stream that has its own signal group, and then call the Action **Stream > New Turn**. This turn can then be given its own signal group, so that the bus lane can proceed on green while other traffic on the same approach is held on red.

(Signal) Group

A Group is either
- a set of (vehicle) Turns, or
- a pedestrian walk

If a Group is defined as a Walk, it can have an associated Walk ID number, or it can apply to all walks. The latter case is known as a “scramble” phase, because it is used to stop all vehicle traffic and allow pedestrians to walk in any direction across an intersection.

A Group has a signal state defined for each phase, and may also have an initial (“T=0”) state. If the initial state is not specified, it defaults to red.

You can add or delete Groups on the Groups tab of the editor. Groups are numbered sequentially from 1, but you can change the number if required. it may be your convention to number vehicle groups from 1 to 8, and pedestrian groups from 9 to 16. Once you have created the correct number of groups you need, you can then add phases, and set the state of each group in each phase. To select a Phase for editing or deletion, use the Phases tab.

Once you have created the Groups, you should return to the Turns tab and assign a direct group or filter group, or both to each turn.

If an intersection is controlled by an external controller (such as SCATSIM), then the groups do not require any phase specification, as groups will be controlled directly by the external controller.

Phase

A Phase is a set of signal states for all groups on a controller. Commonly, a phase is thought of as a collection of groups that are at the “green” state, but more formally, a phase encodes a set of groups that are green, a set that are off, and the remainder as red. While there is no implementation reason that a phase could not encode other signal states (Green Yield etc.), these are not offered as alternatives in the user interface in order to simplify the range of choices.

(Phase) Timing

The timing for each phase is shown in a separate tab in the user interface. The timing parameters are:

Green: the default value, in seconds, for the green time, which can be varied by each plan. You cannot edit the green time here, you should edit it in the Plan tab.

Yellow: Yellow time for this phase for all plans.

Red: Red time for this phase for all plans.

Min: The minimum green time for this phase, which is zero by default, but can be set higher to stop the vehicle actuation rules from reducing the green time below that minimum.

Max: The maximum green time for this phase. This can be set to any value that is larger than the minimum,

Stretch: Set to true if the green time of this phase can be extended by a vehicle-actuation rule, where the balancing phase is not explicitly defined.

[Note for SCATS users: 'Stretch' applies here to the phases described as the stretch phase and time gain phases in SCATS.]

(Phase) Plans

A Plan specifies the cycle time, offset, phase order and green splits of a sequence of phases.

The ***cycle time*** (in seconds) specifies the total time for all phases in the plan, including time for red and yellow. The total green time per cycle or “cycle-green” equals the cycle time less the sum of all yellow and red times. The green split percentage values are expressed as a fraction of the cycle-green time.

The ***offset time*** (in seconds) is useful for synchronising a set of intersections all having the same cycle time. For example, if intersection 1001 has an offset of 0 and intersection 1002 has an offset of 30, then intersection 1002 will start its cycle 30 seconds after intersection 1001. If intersections do not have the same cycle time, then the offset time is not particularly meaningful.

The ***phase order*** must be at least one phase, but has no maximum length, and may contain repeats. For example, if the phases defined for an intersection were A, B, C and D, a phase plan could define the sequence [A, B, A, C, D]. You can have as many plans for an intersection as you like, and switch between them using the “Current Plan” selector.

The ***green splits*** for each phase in a plan are specified as a sequence of numbers. For example, if a plan has three phases A,C,D then the green time specification should have three numbers, in order, for example 40, 35, 25. The values are expressed as percentages, so should sum to 100%. However, for convenience, you can enter absolute values, and they will be normalised when you press Apply. The percentages refer to the proportion of the cycle-green time given to each phase, not a proportion of the total cycle time.

Loops

There can be any number of loops associated with an intersection and its controller. The loop does not have to be on a lane leading up to the intersection, but commonly this is the case. These loops can be used for signal control logic, particularly where the intersection is controlled by an external controller. You can associate loops with an intersection by naming them according to the convention (intersection)_(index). For example 1001_2 associates a loop with intersection 1001, giving the loop an index of 2 on that intersection.

If you select a loop, then use Action **Adjust,** you can set the controller for a loop, or check its current assignment. If you change the controller to which a loop is assigned it will automatically be given a new loop ID, and rename the loop according to the convention described.

The fields in the Loop tab are:

Name	The name of the loop
Lane	The lane on which the loop is located
To Stop	The distance from the downstream edge of the loop to the stop line on its lane
Occupied	Has a value of 1 if occupied, 0 otherwise
Occupancy	The occupancy ratio (time occupied / total time) since the last loop reset
Override	Select this to mark the loop as permanently occupied, which can be used to test signal control

Rules – for Vehicle-Actuated Signals

There can be any number of signal rules defined for a controller. These rules can be used to implement demand-sensitive changes to the signal control, such as vehicle-actuation or pedestrian-actuation. Each rule has a **Plan**, an **Action Phase**, an **Action,** a **Test** condition, and may have a **Balancing Phase**

The ***Action*** defines what will happen if the test is true

- **Terminate**: terminate the action phase (with yellow and red). The remaining green time may be given to another phase, determined in order as follows
  (a) a specified Balancing Phase (even if not stretch)
  (b) the first following phase marked as stretch
  If no phase is identified, the remaining time is discarded
  When Tested: every second during the Action Phase green

  Example: if the phase order is A,C,E,E1,E2, with both E1 and E2 having zero green time by default, and a rule is defined for E with balance phase E2, then the rule will terminate E and go directly to E2, missing out E1

- **Skip**: skip the action phase. This action is similar to Terminate. It uses the same procedure for identifying the Balancing Phase.
  When Tested: just before the Action Phase is due to run.

- **Extend**: extend the Action Phase by one second (plus any Time Parameter value). If a balancing phase can be determined then subtract the extension time from that balancing phase. As above, a balancing phase is determined as:
  (a) a specified Balancing Phase (even if not stretch)
  (b) the first following phase marked as stretch
  When Tested: just before the Action Phase green finishes, when there is one second of remaining green time.

The ***Test*** for each rule is a conditional expression using parameters.

Loop	gap[N]	the time, in seconds, since loop N was last occupied. If the loop is occupied, this parameter returns zero.
	occ[N]	the time, in seconds, that loop N has been occupied. If this loop is unoccupied, this parameter returns zero.
	flow[N]	the flow rate, in vehicle per hour, on loop N
	speed[N]	the mean speed of vehicles crossing loop N
	density[N]	the mean traffic density on loop N
	count[N]	the number of vehicles that have passed loop N
	people[N]	the number of people that have passed loop N
Crossing	push[P]	the number of button pushes registered for crossing with walk ID = P since the last time the crossing was open. Each push corresponds to one person arriving. For example, push2 > 0
Phase	elapsed (lap)	the elapsed green time on the phase to which the rule applies. This will be zero if the phase is not running
	remaining (rem)	the remaining green time on the phase. If the phase is not running, this will hold the green time for the next running of this phase, including any adjustments made by giving or taking time to or from another phase
Controller	pi[C]	the current phase index of controller C, returning 1 for phase A, 2 for phase B and so on. This test can be used to co-ordinate intersections.

Examples

push2 > 0	tests if there are any people waiting on crossing 2
gap3 > 5	tests if loop 3 has been unoccupied for 5 seconds or more
occ7 < 10	tests if loop 7 has been occupied for less than 10 seconds
pi2004 = 3	tests if controller for intersection 2004 is currently on phase C

Available operators for number comparison are: > < = >= <=
Operators available for logical combination are: AND OR

An example test is

(gap7 > 5 AND gap8 > 5) OR occ7 < 10

The ***Plan*** specification can limit the rule to a single plan, or if set to zero, causes the rule to be tested during all plans.

The ***Action Phase*** determines when the rule is tested, and which phase will receive the Action if the Test condition is true.

The **Balancing Phase** can be used to retain the cycle time, as described under Action.

The order of the rules is significant, and the rules can be re-ordered using the up & down arrows beside the rules table.

The **Log Tests** and **Log Actions** options can be selected while testing the rules. These will generate diagnostic messages in the Log window, which can be accessed from the lower right corner of the main window. These options are not saved with the model. They generate a lot of diagnostic text, which can slow down the model significantly, so it is necessary to select them whenever the diagnostic output is required.

Rules Example:

Consider the following intersection, with phase B active. To get this phase to gap out when there is no demand for it, then this requires a test for a gap on all of loops 6, 7, and 8. If an acceptable gap threshold is 5 seconds, then the rule would be defined as follows:

Description	**Gap-Out-B**
Plan	**1**
Phase	**B**
Action	**Terminate**
Test	**gap6 > 5 AND gap7 > 5 AND gap8 > 5**

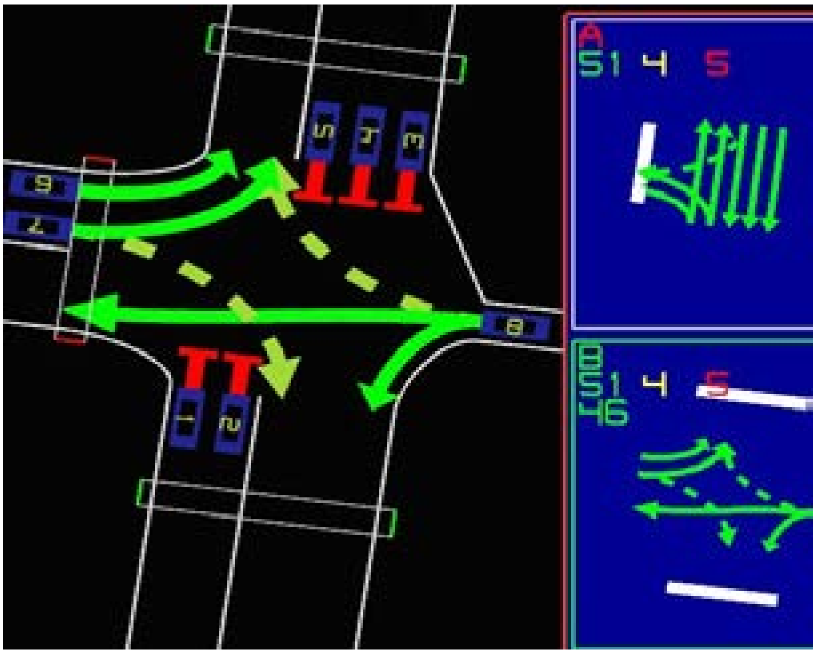

If there was a separate phase for the right turn movement on loop 7, and a rule was required to make that phase terminate on no demand, then the test would need to look at gap7 only.

Boom Gates

The boom gates tool allows you to add a movable boom gate, like those found on parking entrances and exits, or on restricted access lanes, such as bus-only lanes. The boom gate device will delay each vehicle for a variable time, taken from a specified time distribution.

Location: select a lane, to add a boom gate at the end of that lane. If you want a boom gate at a distance back from an intersection, bisect the approach lane into two parts, andput the boom gate at the end of the first lane.

Stop Time: A new stop time distribution for a boom gate can be defined in Parameters / Time Distributions, or you can use an existing distribution. For example, if you select “Uniform Standard A” every vehicle will be stopped for 2 seconds. However, you may want to create a distribution where 40% of vehicles stop for 2 seconds, 30% for 3 seconds, 20% for 4 seconds, and 10% for 15 seconds (the “where did I put my ticket” group). Note that the total stop time is the sum of the stop time taken from the distribution and the (fixed) time for the gate to open.

Time to Open: This time is the same for all vehicles, and is added to the variable stop time. It can be set to zero, if not required.

Limited Height: A visualisation option, splits the boom gate bar into two parts so that it does not collide with the roof in a multi-level parking area.

Ramp Metering

The Ramp Metering Tool allows you to model the control of any number of single-entry, single-exit signalised intersections, to limit the flow of traffic from an on-ramp onto a multi-lane road. The tool uses a logic module that takes its input from one or more loops and a set of variable parameters. The output from the logic module is the cycle time of the control signal.

The logic module is an implementation of any suitable algorithm. A typical algorithm might vary the cycle time in proportion to the occupancy on the main road - the higher the occupancy, the higher the cycle time, and the lower the amount of additional traffic being allowed to join the main road from this ramp, thus averting congestion and the sudden drop in capacity associated with flow breakdown.

There can be more than one logic module operating with the Tool on a road network, allowing different algorithms to be used for different types of ramps.

Three logic modules are supplied as standard, but more can be supplied on demand.

Preliminary Explanation of Meter Timing and Flow

The timing of the each meter is controlled by its cycle time, T, measured in seconds. If N vehicles are allowed to pass on each cycle, the total hourly flow allowed to pass by a meter is equal to N x 3600 / T. There will always be at least one green per cycle, but there may be more, for example if the meter is set up as a two-lane alternating meter, then there will be two greens per cycle. The lengths of the red and green time components are not important for the purpose of simulation: the signal will be coloured green while a vehicle is allowed to pass, but the number of vehicles allowed to pass on each green depends on the value of N, not on the length of the green time.

Standard Input Parameters

Meter ID: Each meter is identified by a unique name

Term: A meter may operate all the time, or it may be limited to operation during a chosen term.

Sign: A meter may be linked with a Sign, and the sign used to switch the ramp meter on and off. This can be used to let plugins and API users control a ramp meter. If a ramp meter is linked to a sign, and that sign displays “OFF”, then the ramp meter will switch off. Any other message will cause the ramp to switch on.

V / Cycle: The number of vehicles allowed to pass per cycle. This is a whole number, which is normally 1, but can be higher. Where the meter is located at the end of a link with two or more lanes, the value will probably be equal to the number of lanes.

Green: The green time for each cycle, in seconds.

1st Cycle: The initial length of the cycle, in seconds. The cycle length is the sum of the red and the green times.

Min Cycle: The minimum length of the cycle, in seconds. This value controls the maximum flow that can pass the meter. If the cycle is set to the default minimum of 5 seconds, and 1 vehicle per cycle can pass, then the maximum flow through the meter is 720 vph.

Max Cycle: The maximum length of the cycle, in seconds. This value controls the minimum flow that can pass the meter.

Interval: The time between calls to the logic module, in seconds. A short interval will make a meter more responsive to changes in the main flow, but will also make the output value less stable, and may make it more prone to oscillation between large and small values. The call interval is also used as the time period over which the logic calculates the percentage occupancy.

Additional Input Parameters

In addition to the parameters listed above, there may be further input parameters specific to the logic module applied to the meter.

Ramp Meter Output Values

Cycle Time: The principal output from the meter, this controls the flow of traffic past the meter

V / Hour : The flow expressed in terms of vehicles per hour. This value is derived from N x 3600 / T, where T is the cycle time and N is the number of lanes

[Additional Values]: There may be additional values to display, specific to the logic module applied to the meter.

Adding a Ramp Meter

1. If not already created, create a signalised intersection at the location of the ramp meter. This intersection must have a single entry link and a single exit link, and have 2 phases.
2. Add loops to your model that will be used by the meter for input. *Each loop must be associated with the controller* for the ramp meter "intersection". If you want to add a multi-lane loop, first select the controller, then select one of the lanes, and use Action **Lane > New Multi-Loop**. This type of loop registers as occupied if any one of the lanes is occupied.
3. Select the loops that will be used by the meter. Each ramp metering logic has its own specification for the number of loops required.
4. On the Ramp Meter Status window, select the tab for the logic module that is required for the new meter.
5. Press "New..." and edit the name of the meter, as required.
6. Verify that the correct link and loop names are shown in the text fields, then press "OK".
7. Modify the parameters as required in the new row that appears in the Ramp Meter Status Window. Changes are applied immediately.
8. Save the model.

Logic Module

Each logic module takes one or more loops and perhaps some other input parameters and produces a value for the cycle time. The module may also display a number of additional values in the status window.

Module 1: ALINEA

This logic module is based on the ALINEA method proposed by Markos Papageorgiou et al. It requires a single loop.

	O_G is the target occupancy
$R_T = R_{T-1} + K_R (O_G - O_{T-1})$	O_{T-1} is the measured occupancy
	K_R constant of proportionality

This expression states that the metered flow rate R for the next interval, at time T, should be derived by modifying the flow rate for the previous interval, at time T-1, by an amount that is proportional to the difference between the target occupancy and the measured occupancy. The flow rate values are measured in vehicles per hour and the occupancy values are unit-less values between 0 and 1 and are calculated over an interval. The value 1 means the loop has been occupied for the entire interval, and 0 means it has been unoccupied for the entire interval.

The interval, the target occupancy and the constant of proportionality can be changed in the Ramp Meter window. The length of the loop affects the measured occupancy and the distance of the loop from the end of the ramp also has an impact on the effect of the ramp meter control. Loop length and position can be changed in the main Commuter window. T

Suggested parameter values are:

$O_G = 0.2$ $K_R = 70$ interval = 30

Module 2: ALINEA-Q

This is a modification to ALINEA to limit the effect of a queue from the ramp meter spilling back onto the rest of the network. This module requires an additional loop, to detect the back of the queue, and a threshold value to trigger a "switch-off" signal sent to the ramp meter. When a meter is switched off, the flow through the meter is uncontrolled, and is as high as the geometry allows. This is probably not a very effective procedure to use in practice, as it results in a system that oscillated between controlled and uncontrolled states, but it illustrates the use of additional input parameters.

The second loop is placed upstream from the meter, just after the point at which a tail back from the meter might start to create "grid-lock". The parameter "Off Occ" is the threshold value. If the occupancy on the second loop is greater than the "Off Occ" value, the meter is switched off. The current value of the occupancy can be seen in the status window under the heading "Occ[2]".

Module 3: Fixed Rate

This is a simpler module that requires no input loops. It creates a ramp meter with a fixed cycle length.

To add a fixed rate ramp meter, create the intersection as described, select the controller, and press [+].

Then edit the flow rate, which will automatically calculate the cycle time required.

Timing Playback

The Timing Playback tool allows you to read signal control timing information from an external spreadsheet in XLS format, and use that as the basis for controlling the signals in a model during a simulation run. This allows you to "replay" the signal control strategy that was used on a particular day, or to fine tune signal control by specifying the start time of each phase.

You must first define phases for the intersection, as all timing information is specified according to phase.

There are two possible formats.

Importing Playback Data: To import playback data:

1. Select menu item Control> Timing Playback…

2. On the Advanced Controller Parameters window, switch to the Playback tab.

3. Press **Import Timing**, and browse to the Excel file

4. Select the format, Cycle-By-Cyle or Phase-By-Phase (see below for a description of each format)

This copies the data from the XLS file and saves it in the AZA model file. After this the XLS file is no longer referenced so if you change the spreadsheet you will need to re-import the data.

WARNING: Although there is an intersection selector at the top of the Controller parameters window, the **Import Timing** imports data for all intersections. Any import operation automatically clears any previous data held. The intersection selector is used for the Gap-Outs tab only (see below)

Format 1: Cycle - By - Cycle

	A	B	C	D	E	F	G	H	I	J
1	Time	Cycle	A	B	C	D	E	F	G	Order
2	15:00:00	110	63		27			20		ACF
3	15:01:50	90	63		27					AC
4	15:03:20	90	63		27					AC
5	15:04:50	120	63		27	10		20		ACDF
6	15:06:50	90	63		27					AC
7	15:08:20	90	63		27					AC
8	15:09:50	113	63		20	10	20			ACDE
9	15:11:43	90	63		27					AC
10	15:13:13	110	63		27		10	10		ACEF
11	15:15:03	90	63		27					AC
12	15:16:33	90	63		27					AC
13	15:18:03	90	63		27					AC
14	15:19:33	90	63		27					AC
15	15:21:03	78	63		15					AC
16	15:22:21	90	63		27					AC
17	15:23:51	90	63		27					AC
18	15:25:21	90	63		27					AC
19	15:26:51	73	63		10					AC
20	15:28:04	90	63		27					AC
21	15:29:34	90	63		27					AC
22	15:31:04	110	63		27		20			ACE
23	15:32:54	73	63		10					AC
24	15:34:07	90	63		27					AC
25	15:35:37	90	63		27					AC
26	15:37:07	90	63		27					AC
27	15:38:37	90	63		27					AC
28	15:40:07	90	63		27					AC
29	15:41:37	90	63		27					AC
30	15:43:07	90	63		27					AC

2013 2077 [illegible]

Ready

Intersection number

If a phase is not used, the time can be set to zero, or left blank.

In Excel, you can make it easier for yourself by creating equations:

Cycle: in cell B2 **=SUM(D2:P2)**
Time: in cell A3 **=A2+(B2/(24*60*60))**

The Cycle value is the sum of all phase times. The Time value is the sum of “Previous Time” + “Previous Cycle”. [In Excel, time is stored as a decimal number of days, so you need to divide the time in seconds by the number of seconds in a day]. Then copy an dpaste the formula to the whole column.

Format 2: Phase - By - Phase

	A	B	C	D	E
1	Start	End	Duration	Phase	
2	00:00:00	00:01:44	104	A	
3	00:01:44	00:02:04	20	B	
4	00:02:04	00:02:23	19	C	
5	00:02:23	00:06:05	222	A	
6	00:06:05	00:06:24	19	C	
7	00:06:24	00:06:39	15	A	
8	00:06:39	00:06:57	18	B	
9	00:06:57	00:09:15	138	A	
10	00:09:15	00:09:34	19	B	
11	00:09:34	00:11:45	131	A	
12	00:11:45	00:12:05	20	C	
13	00:12:05	00:17:34	329	A	
14	00:17:34	00:17:52	18	B	
15	00:17:52	00:19:50	118	A	
16	00:19:50	00:20:07	17	B	
17	00:20:07	00:22:02	115	A	
18	00:22:02	00:22:21	19	B	
19	00:22:21	00:22:37	16	A	
20	00:22:37	00:22:54	17	B	
21	00:22:54	00:23:09	15	A	
22	00:23:09	00:23:29	20	B	

45 | 85 | 146 | 180 | 1005

Ready — Intersection number

As with format 1, you can create equations to work out all the values in columns A and B after the first row. Then you just need to enter values for columns C and D. If these phases repeat, you can copy the timing from the cycle above in columns C and D.

With the correct tool, timings can be extracted from SCATS IDM data into this format.

Gap-Outs

A gap-out is similar to a signal rule, but for use with playback timing.

A gap-out definition is used to terminate a phase early if there is no demand for that phase, as detected by one or more loops.

In a typical case, the phase would be of the form D1 or 1D, with the master phase being D. The form “1D” is often used for a “bus queue-jump” operation. This alternative phase does not appear in the table of playback data directly, rather the phase is a variation of a master phase and is conditional on the loop test.

The gap-out data comprises:

- **Min:** a whole number of seconds. If this phase has a master, then this specifies the minimum time that the master phase (e.g. D) must run before swapping to this phase (D2)

- **Gap:** a whole number of seconds. This is the time that all loops must be unoccupied to trigger the start of this phase.

- **Loops:** a set of up to 8 loops. All loops in the set must be unoccupied to trigger this phase.

A gap-out specification works **only** with Playback timing, not with a fixed controller timing.

(Demand) Divisions

A Demand Division is a weighted group of Types. You can define any number of divisions, each of which can contain any number of types. Divisions are used by the trip generation process to apply a demand matrix or a volume in given proportions to a group of vehicles. Divisions are also used to select vehicles for parking. This is also a type of demand allocation, because it decides the proportions of vehicles to generate in a parking zone if none are available.

Divisions may overlap, that is, any one type may be contained in any number of divisions. A division may contain a single type.

A division may be assigned a mode to restrict it to vehicle types of that mode. For example, if you create a new division with a pedestrian mode, only pedestrian types may be included in that division. However, a division can be mixed-mode, so that a single division can contain a combination of vehicle and pedestrian types.

The Division Window shows a table with all types as columns, and all divisions as rows.

To create a new division: press the “New Division” button in the top left corner of the Division Window. Choose a name for the division and select a mode, or one of the mixed-mode combinations. Then select the types you want in the division by ticking the boxes in the column for each type.

To delete a division: select the division in the table by clicking on its name, then press the “Delete Division” button.

To rename a division: click on the name of the division and type the new name.

Profiles

A Profile is a set of weights that sums to 100%, where each weight applies to a time interval. The weights are applied in sequence to generate a time varying quantity. The Term defines the start and end time. A profile is use to vary the number of trips generated over the course of a term. It is often lower at the start and end, and higher in the middle to represent a peak hour.

A profile can have any number of intervals; all of the time intervals are of the same length. If for example, you have an hourly demand, but want to weight the demand by some 5-minute counts, you would define a profile with 12 x 5-minute intervals, and set the corresponding 12 weights according to the count values.

The Term applied to a Profile can be independent of all other Terms. For example, it does not need to coincide with the Terms defined for restrictions or signal timings, or even for other profiles.

The default profile, named **Flat** applies a constant weight, of value 100%, all the time.

Apart from their main use with Demand, Profiles can also be used by Plugins, for example to set the toll cost on a time-variable road pricing model.

1 x 100% is not the same as 4 x 25%

If you define an hour in the simulation (for example AM Peak) to have 4x 15-minute intervals, with 25% assigned to each, you will not get the same results as 1x 60 minute interval @ 100%. The difference will be most noticeable where there are small numbers of trips. If you have 1 trip defined in your matrix between A and B, then you are asking for a quarter of a trip every 15 minutes, which is not possible. At the end of the matrix term, extra trips are generated to try to make the row totals correct, but it is not possible to make all 4 15-minute interval totals correct as fractional trips cannot be created. By creating 4 intervals in your profile, you are effectively creating 4 matrices, one for each interval, each scaled by the factor applied to that interval.

Using the Profiles Window

To define a new profile:

- press the “New Profile” button in the top left corner
- Enter a name, for example “AM Peak”
- Choose a Term for the Profile, or define a new one using the drop-down selector
- Define the number of intervals within the Profile

The list on the left allows you to select a profile for editing.

The number on the right hand side below the buttons shows you how many intervals are in the current profile. You may need to scroll the central pane left or right to see all the weights.

To accept a new or edited profile, you must tick the “Sum to 100%” ($\sum = 1$) box. This will normalise the weights so that their sum is 100%.

After creating a Profile, you can

- add an interval using the “New Interval” button top right
- delete an interval
- reset all the weights to be equal using the “Reset” button. For example if you have 4 intervals in the profile, reset will set all 4 weights to 25%
- adjust the weight for any single interval by moving a slider, or by typing in the box below the slider.
- If you move the slider and the $\sum = 1$ box is ticked, then all the other weights will be adjusted to keep the sum at 100%
- If you type a value into the box below a slider, the $\sum = 1$ box will be cleared for you.
- you can paste a row of values from a spreadsheet into the weight value boxes. Select the first (leftmost) box and paste using Control-V.

Demands – Directed & Undirected

Undirected Demand (Zone Counts and Turning Counts)

Undirected demand is simpler than directed (or O-D) demand, as you need only a single value for each origin zone, rather than a set of values, one for each destination. As there is only a single demand value for each origin zone, this type of demand requires a one-dimensional column of values (a **volume**) rather than a two-dimensional **matrix**. However, if there are any route choices in your network, you need to specify the turning counts at each route choice location for each possible destination. If your model is fairly simple, and has only one route between each origin and destination pair, for example a corridor model, then this method of specifying demand may be suitable.

Each route choice location automatically creates a “Split” object, that by default assigns an equal proportion of all incoming traffic to each of the available exit options. At a 4-way intersection, an incoming agent has a choice of left, right or straight ahead, and the default split object will stochastically assign the agent to one of those exits, each with equal probability of 1/3. This means that if there is a loop in the network, an agent may be directed around the loop, arriving back at the same point. With typical splits, the probability of this happening is low, and in some networks, following a loop around may reflect reality, for example where a vehicle is looking for a parking space, or a pedestrian is window shopping. Some worked examples on the following pages illustrate the probabilities of following a loop.

The Split function is very simple; if more complex routing rules are required, then the route choice tool allows you to create rules that assign the exit (left, ahead, right) based on origin zone, agent type, current lane, sign settings, etc. The outcome of a route choice rule can be a single exit or a stochastic distribution across multiple exits, similar to a Split.

Directed Demand (OD Matrices)

Directed demand requires a two-dimensional **origin-demand matrix (OD matrix)**. This fully specifies all trips between each origin and each destination. This data is often the most difficult to collect for the modelling process, because an accurate matrix requires either

- a roadside survey, asking drivers to disclose the start and end point of their journey, and if possible also the time of depature. This option is expensive, and normally only surveys a small sample of the total network traffic
- an electronic monitoring system, for example one recording license plates or automatic toll payment cards. These systems are often unpopular due to privacy concerns, and, like roadside surveys, seldom provide more than a patchwork of cover.

In most cases, because collection of accurate data for a full matrix is not possible, often due to cost, a matrix is estimated. This process commonly starts with a pattern matrix, which is an initial estimate of relative demand between pair of zones. The pattern matrix is often derived from historical data. The pattern matrix is then modified using growth factors for each zone and verified against measured counts on the network. Commuter includes a matrix estimation tool, for more information see the relevant section.

Matrix estimation can never produce a result that is known to be correct, because for any realistic network, there are too many unknowns. The demand matrix is frequently the largest source of error in all of the data used for a simulation model.

Therefore, before deciding to use a matrix, make yourself aware of the error margin that you will introduce, and consider the alternatives, such as undirected demand in combination with turning counts.

Observations

The Observations window is used to enter and display “real-world” measurements of flow counts or travel times. Each measurement is for a single location, and for a single Term. Each measurement can be either for vehicles or for people. These observed values can then be used either for **Matrix Estimation** or for model **Validation**.

Entering Data

- Select the links, walkways or other locations
- Press the [+] button.
- Select a term. For example, if you had a 4-hour model, with counts for each hour, you should define 4 terms then enter each set of counts in turn.
- Select a division. By default this is set to “all types”, but you can associate any other division. For example, if you had separate traffic counts for cars and buses for each hour, you could define a division for each then import each set of counts in turn, selecting the appropriate division when prompted while entering the data. It is also possible to use a **Fleet** (like a division, but with no proportions).

Importing Observed Data

Observed data for the most common types of counts can be imported from external CSV files, rather than entering each value individually. Data should be contained in plain-text comma-separated-value (CSV) files, with these columns:

- Mid-Block Counts (link/walkway name), (count)
- Turn Counts (inward link name), (outward link name), (count)
- Cordon Counts (zone/area name), (in count), (out count)

Scaling: You can copy and paste between the observations window and a spreadsheet, so if you need to scale the values, you can do this by copying the values to a spreadsheet, apply the scale factor to an adjacent column, and then copying the new values back.

Mid-Block Counts: These are counts of all the traffic on a link, or pedestrians on a walkway, regardless of exit. Each link is one-way, so for a two-way street you will need two link counts. A walkway is two-way; the count is the sum of people travelling in both directions

Turn Counts: These are traffic counts at an intersection, and each is defined as the number of agents making a turn out of a link to a numbered exit. For example, if a link has 3 exits, then there would normally by 3 turn counts for that link, corresponding to left, ahead and right. The turn counts are specified in order, from kerb side to median side. Where traffic drives on the left, the first count is for the left turn, the second is for the ahead turn, and the third is for the right turn. Where traffic drives on the right, the first count is for the right turn.

O/D Counts (Origin/Destination): These are counts derived from zone or area cordons. If you have a count of the number of trips starting from a zone, you can enter that as an “outward” cordon count, and if you have the value for trips ending in that zone, you can enter that as an “inward” cordon count.

Screen Line Counts: A Screen Line Count is similar to a Mid-Block Count but is an aggregated value of counts on a set of two or more of links. To create a screen line, select all the links that are part of the screen line, and press the Add button on the Screen Line tab.

Trail Counts: A Trail Count counts only those pedestrians or vehicles that follow the entire trail from start to finish. Trails can be added using the Trails window accessed from the Network menu.

Traverse Counts: A Traverse is defined by two links or walkways: one at the start, the other at the end. A Traverse Count measures the number of vehicles or people that go from the start link to the end link, by any route. Start and end can be the same link.

Trail Times / Traverse Times: as for counts, but measuring the travel time, including both the start and end links.

Visualisation of Inconsistent Turn Counts

This feature allows you to view and edit inconsistencies in turn counts.

- Open the Demand / Matrix Estimator window
- Select a matrix that matches the Term and Division of the turn counts to be verified. If the Term or Division is null for the turn counts, they will match any Term or Division in the Matrix Estimation window.
- In the Layers window, select the Estimation section, and switch on Inconsistent Counts

There will be a coloured rectangle connecting each link that has at least one turn count in and one turn count out. This is a turn count comparison object. If there is an inconsistency, the rectangle will be coloured. The hotter/colder the colour, the greater the inconsistency, relative to other turn count comparisons in the model.

Select the coloured rectangle, and use the Adjust action, or triple-click, as normal. A small window will be raised with text fields allowing you to type in modified values for the turn counts.

Why would you want to change turn counts? The truth is that observations are rarely consistent. Sometimes this is because counts have been taken on different days, or it is because the model does not take all movements into account. For example there may be a small side road on the link between two intersections that is not being modelled and this accounts for the inconsistency between turn counts in and out. Sometimes, the best solution is to modify the turn count values.

Demand Editor

The Demand Editor has three tabbed panes, one for undirected demand, one for directed demand and the other for (Public) Transport demand. Before proceeding, if you do not understand the difference, please refer to the section that explains the difference.

Directed and Undirected Demand Tabs

These two tabs have the same general layout. Directed demand uses OD matrices, where demand between all OD pairs is known.

Undirected demand uses traffic volumes from origin zones in combination with turning counts. Undirected demand with turning counts can be used to estimate origin-destination demands by running the model and saving the arrivals matrix.

List Panel: on the left side, contains a list of the demand sources within the model. You can add a new source of demand by pressing the Add button above the list, or delete a source by selecting it and pressing the Delete button below the list.

Contents Panel: in the centre of the panel, shows a table containing details of any source currently selected in the List Panel. This table contains a breakdown of the demand by origin (and destination, for directed demand only).

Mode Indicator: shows the mode of the demand source, either Pedestrians or Vehicles

Profile Selector: allows you to vary the release of the demand over time.

Division Selector: allows you to restrict the demand to a subset of all available types for the Mode.

(Public) Transport Demand Tab

Each public transport service contains a timetable, but the vehicle trips for a service must be generated at the same time as all other trips.

The Transport Demand Tab repeats much of the information that is available in the Services window. However, it can be useful to see all the demand information in a single window.

The Transport Demand Tab lists all public transport services on the left side. The departure times and associated vehicle types are are displayed for each service on the right hand side. You can add or delete entries to the timetable using the buttons below the time table.
To edit the route of a transport Service, you should use the Services window, accessible from the Network menu on the Main Window.

Add Demand Page: this button allows you to add a completely new page of demands to the model. For example, you may create one page of demands for morning peak, and another for the evening peak. You may also create base year and future year variations for each of the AM and PM peaks, giving you a total of four demand pages. You can create as many demand pages as you require, but only one can be active when you generate trips.

Multiple Sources are additive: All the sources listed on a demand page are added together when trips are generated, and this includes undirected demand, directed demand and transport demand sources. If you have 3 volumes and 2 matrices defined, then trips will be generated based on the sum of all these sources. Typical reasons for having multiple sources are that you might have one matrix for private cars, another for goods vehicles. Also, it is common to break down a single peak period into 3 hours, and have a separate demand source for each hour.

Matrices do not require full coverage: When you create a matrix, you are asked to enter a list of origins and a list of destinations. By default, both these lists contain all possible origins and destinations, but you can trim them to suit. For example, if you have one matrix covering highway zones, and another covering city blocks, you can add these separately. [Volumes always assume full coverage, as there is only one cell for each origin. To “trim” an origin in a volume, just enter zero in its cell.]

Trim out one-way zones: Models commonly contain some areas or zones that operate only as origins or as destinations, not both. Trimming out areas or zones from either the origins or the destinations can reduce the size of the matrix and make the data easier to manage. For example, if you have a motorway off-ramp that is a destination, but does not act as an origin for any trips, then remove it from the list of origins.

Demand source divisions must be single-mode: while it is possible to create a mixed-mode division, any division you select for a demand source must be a single mode, and must be either People, or Vehicles. It is not possible to create demand for Public Transport vehicles. Public Transport trips are created in the service window, by defining departure times for each Service; typically this is done from the timetable for that Service.

Demands – Matrix Estimation

The Matrix Estimation Tool start with a pattern matrix, and changes each cell value to get the best possible fit with observed traffic counts.

Creating a Prior Matrix

Given departure and arrival volumes for all zones in the network, the Demand editor allow the creation of a basic pattern matrix.

1. Open the Volumes tab of the Demand Editor
2. Add a Volume, choose *Direction = Departures*.
3. Choose Profile (Term) and Division
4. Choose the sector for the origin zones for which departure counts exist, or the default “All Zones” sector
5. Enter the departure counts in the demand column
6. Repeat steps 2-5 but this time, choose *Direction = Arrivals*, and enter arrival counts
7. Open the Matrix Estimator Tool, press [+],
8. Switch to the *By Departure/Arrival* tab
9. Select the Departures and Arrivals volumes just defined, and press OK.

This creates a matrix using Iterative Proportional Fitting (IPF, also known as Furnessing). The totals should be close to or equal to the original departures and arrivals entered. The IPF method does not use the distance or cost of travelling between zones, as used in Gravity or Entropy models, it is based purely on departure and arrival counts. So while this pattern matrix is an acceptable place to start, it will need further tuning, that is, redistributing trips within the matrix, while retaining row and column totals. This will use observed counts, such as turn counts or mid-link counts.

Trip Simulator / Matrix Estimator

This uses dynamic traffic assignment, without creating individual agents. That is, there is a time axis, and each trip departs at its departure time, and travels along the links in the network, making route choices as it goes. As it passes each counting location in the model it is recorded, and these modelled counts are later compared against observed counts. The comparisons are then used to produce a better estimate of the matrix by redistributing trips. The redistribution aims to retain the row and column totals, while reducing the difference between modelled and observed counts.

One important thing to understand is that trips are created by the tool “on the fly”: you do not need to generate trips.

The steps in the process of matrix estimation by redistribution are listed below. These are described for vehicles, zones and roads, but can equally be applied to people, areas, and any combination of walkways, parking, services, etc.

1. Open Demands / Matrix Estimator.
2. Start by creating a pattern matrix, either using Departure and arrival volumes, as described, or by creating an empty matrix of the right size, and pasting in the cell values from a spreadsheet application. The pattern matrix can also be created in the Demand Editor window
3. Enter all the observed counts in the Observations window, with a term and division that match the matrix. Observed counts can be turn counts or mid-block counts.
4. On the Matrix Estimator window, press the Play button
5. Press the redistribute button (small gear wheels), the clock should advance through the simulation interval
6. On completion of the interval, press the rewind button
7. Repeat steps 4-7 until a good solution has been achieved.

After the trips have been passed through the system, cells will be coloured according to the modelled/observed comparison, either blue (too low) or red (too high). If you click the “delta” radio button you will see the difference values. When Edit mode is off (see below), double clicking a cell will raise another window which shows you more information about that cell and why it is shown as too high or too low.

There is a lock button (black key) which will lock any selected cells so they will not be adjusted. [The green key is for “Automatic lock” which will lock any cells that have only one trip passing through them, and they can be automatically locked to that value. However, this does not do anything yet]

The “Edit” button switches on edit mode, which allows you to manually edit any cell you like. At the moment, the changes you make through manual editing will not be redistributed through the rows and columns, so your changes will cause the row and column total to be inconsistent with the departure and arrival volume.

This uses assignment based on lowest cost to destination, but takes no account of reduced levels of flow caused by congestion – it does not use dynamic assignment or “cost feedback”. If the modelled count at a location is too high you may want to consider modifying the model to bias traffic away from that link, either by adding a speed or route cost factor, or by adding a route choice rule.

If you are unsure about your turn counts, go to the manual pages for the Validator tool, and look for a section entitled *Visualisation of Inconsistent Turn Counts*

Trips – Introduction

In Commuter, a page of Trips must be generated before the simulation begins, if there are to be any agents (pedestrians or vehicles) in the model. All the randomness required to effectively simulate people and traffic is contained within the Trips page. One key benefit of this approach is that a page of trips can be generated once, then used for several different simulation runs, ***ensuring repeatability, even if the underlying network changes.***

For example, imagine you have a Base network and two options, Option1 and Option2. You can generate the Trips page once, then run exactly the same trips on all three networks, knowing that any variation you see is due entirely to the variations in the networks, and not due to an artefact of the random number generator in the simulator.

A page of trips contains all the information for a simulation, with separate sets of tables for pedestrians and vehicles, and within those sets, separate tables for directed and undirected trips. (A directed trip has a definite destination, an undirected trip does not.)

Each table contains a number of rows, with one row for each trip. Each row defines all required fields for a trip, including:

- a unique name for the trip
- the departure time (to the millisecond)
- agent type (for vehicles - car, van, truck, etc.; for pedestrians – adult, child, etc.)
- trip origin (area/zone identifier)
- trip destination (area/zone identifier, for directed trips only)
- the trip “DNA”, a block of randomly distributed variables, for use by the core simulation engine and any plugins

Trips - Browser

The Trips Browser window shows details about all available Trip pages. Details include

- When the Trips page was generated, and its random seed
- The Demand page used as the source of the trips table
- The number of Directed and Undirected Trips (D + U) for each category of trip – People, Freight and Vehicles
- The number of scheduled fixed route transport trips
- The time of the first and the last trip in the table.

You can create **multiple** trip pages in any network, each containing a unique set of trips. One reason for creating several trip pages is to test your model for stability. If all other parameters are the same, then for each page to be unique, you must specify a different “seed” value before trip generation begins. The seed value is a way of “jumbling up” the random numbers used to set variable fields in each trip – for example, departure time, vehicle type etc.

Another reason for creating multiple trip pages is when you have multiple demand pages, and you want to create one or more trip pages for each demand page.

The **active demand page** is used as the source of information for generating trips. You may only have a single demand page, in which case you do not need to worry about this. However, it is more often the case that you will have multiple demand pages: for example, one demand page for an AM peak model, and another demand page for a PM peak model. You may also have base and predicted future demands. So, before generating trips, ensure that the demand page you want to use as the source is set as the active page in the File Tab on the Left Pane. Right-click on a Demand page and select Activate Page.

Trip Generator

The Trip Generator allows you to generate new Trip Pages

You can create **multiple** trip pages in any network, each containing a unique set of trips. One reason for creating several trip pages is to test your model for stability. If all other parameters are the same, then for each page to be unique, you must specify a different “seed” value before trip generation begins. The seed value is a way of “jumbling up” the random numbers used to set variable fields in each trip – for example, departure time, vehicle type etc.

Another reason for creating multiple trip pages is when you have multiple demand pages, and you want to create one or more trip pages for each demand page.

The **active demand page** is used as the source of information for generating trips. You may only have a single demand page, in which case you do not need to worry about this. However, it is more often the case that you will have multiple demand pages: for example, one demand page for an AM peak model, and another demand page for a PM peak model. You may also have base and predicted future demands. So, before generating trips, ensure that the demand page you want to use as the source is set as the active page in the File Tab on the Left Pane. Right-click on a Demand page and select Activate Page.

Pre-generation parameters

Before trips are generated, you must specify a number of parameters.

Page Name: use a name that is meaningful to you. By default the name of active demand page is used as a prefix, so, for example, if your demands are called BaseAM, the default names will be BaseAM-T1, BaseAM-T2 etc, but you can change the suggested name as required.

Seed The random seed is a number used to initialise the pseudo-random number generator in Commuter. This is used to "randomise" many of the parameters within a trip, including the type and behaviour of the agent, and for directed trips, the destination. The seed can be any positive integer value.

Demand weight: A percentage of the total demand to be generated. The default is 100% meaning that all trips will be generated. You may want to use this to quickly increase or decrease the total demand, for example, if the future test option has 5% extra trips then using a Demand Weight of 105% will increase the demand for all sources. For finer control, there is also a demand weight for each source in the Trip Sources window (see below).

Destinations / Departure Times

- **stochastic_stochastic**: each trip from an origin A, has a departure time chosen at random and has a destination chosen probabilistically, based on the share of trips from A to each destination. If you have specified 600 trips departing zone A in one hour, then exactly 600 trips will be generated from A but the exact number of trips to each destination may vary from the specification. That is, the row total for each row in the matrix will be an exact match, but the individual cell counts in the matrix may not match. For departure times, all times are equally likely between the start and the end of the simulation Term. Each trip's departure time is independent of all other trips. Departure times of trips will be irregular. For the example above, the average time between departures will be 6 seconds, but the randomness introduced will ensure they depart in "clumps".
- **exact_stochastic:** each trip has a departure time chosen at random as above. The difference here is that each A-B cell count of trips will match the demand specification. If you have specified 600 trips departing zone A and arriving at zone B in one hour, then exactly 600 trips will be generated between A & B, but, as above, their departure times will be irregular. The type of each trip will be selected individually at random, according to the % split specified in the division.

- **exact_stochastic_exact_types**: as for exact_stochastic above, but the exact proportion of each type specified in the division will be generated. For example, if you have specified 600 trips departing zone A and 25% of the division is type T, then exactly 150 of type T will depart from A.
- **exact_regular**: the trips will be evenly spaced in the simulation Term. So, as above, if you specify 600 trips leaving a zone in one hour, then exactly 600 trips will depart from that zone and these will be exactly 6 seconds apart.

Trip Sources: This button open a window which allows you to select any combination of the matrices, volumes and services defined within the current demand page.

- **On**: Select this to include a source in the trip generation
- **%:** Set this value to factor the number of trips from this source. The value can be larger than 100%. For example, it may be set to 105 if demand is expected to increase by 5%
- **Invert**: Use this button to invert the current selection. This is useful if you have two sets of sources, for example, one for AM and one for PM, and you want to switch between them.

[Note: some simulators use a probability method of trip generation, but this is not available in Commuter. While this method is arguably better for some types of simulations, in practice it is not popular in traffic simulation, because it introduces more variability into a part of the model – the demand - that often contains significant errors already. For example, if the time step is one second, and you have specified 600 trips leaving a zone in one hour, then using this method, the probability of a trip being released in each second is 600/3600 or 1/6. The most probable number of trips departing from that zone is 600, but any single run of this method may produce any number of trips between 0 and 3600. In practice, trip counts of less than 540 or more than 660 might be experienced, introducing an error of more than 10%, which many users find unacceptable.]

Trips - Matrix Estimation

See: Demands - Matrix Estimation

[The simple Trip Estimation tool provided with Commuter 3 and 4 has been superseded by a more comprehensive Matrix Estimation Tool in Commuter 5]

Route Classes

A route class is defined to encode parameters for weighting route costs. It may be used, for example, to bias routing decisions towards major routes. A route class contains only routing information, and is associated only with branches on the routing network, and not with any physical surfaces used for carrying agents. Speed parameters in each route class should not be confused with parameters representing posted speed limits on the physical surfaces. For example, if route class “A” is applied to a set of roads, and is given a routing speed of 60 km/h, this value is independent from the speed limit posted on any lanes on links which have route class A. It is up to you to verify that the route class speed is appropriate for each link. The Auditor Tool allows you to check that the speeds you have applied on your surface network are reasonably consistent with the speeds you have applied on the underlying routing network.

Route classes apply to both vehicle routes on links, and pedestrian routes on walkways.

Standard Classes: There are always three route classes defined, one for pedestrians, one for vehicles and the other for rail.

Route Class Fields

Name: A unique name, which can be a number.

Description: A description, which does not need to be unique.

Surface Walk, Road or Rail

Colour: A colour used to highlight route branches assigned to the class, useful for auditing networks. That is, if you define one route class for major routes, and another for minor routes, displaying route branches by route class colour may illustrate where any branches have been assigned the wrong class.

Route Speed: The mean travel speed for routes of this class. This speed is used as the default initial speed for all branches that are assigned to this route class, but each branch can then modify the speed value from that initial default. The units used are the “local” speed units. This speed is used for route calculations only, and not for the control of the speed of individual vehicles or people in the simulation.

Price Per Unit Distance: The price in your preferred currency unit, to travel one long distance unit on routes of this class. For example, in Australia, this would be dollars per kilometre. This setting is useful only if you are modelling a scheme that uses road pricing based on distance travelled. In most other cases, it will be more appropriate to set prices on individual links. Using a price per unit distance is not the same as using a higher route distance factor, because each agent behaviour type can have its own time, distance and price weighting factors, and agents with that behaviour will then react in different ways to a cost made up of distance and price.

Following: This can be used to set a car-following algorithm by route class, that is different from the default algorithm chosen for the model in the Parameters/ Calibration window. For example, the Wiedemann algorithm may be more appropriate for freeway car-following behaviour while Gipps may be better at modelling the car-following of urban roads.

Route Class Cost Factors

A route class is defined to encode parameters for weighting route costs. Each route class may have one or more cost factors associated with it. The cost factors for each route class can apply to an individual behaviour type, or to all behaviour types. Cost factors are multiplicative, so a cost factor of 1.0 has no effect.

One example of how you may want to use this feature is to define two route classes (“major roads”, “minor roads”) and two behaviour types (“familiar”, “unfamiliar”). The cost factor applied to the minor roads for the unfamiliar drivers would be greater than 1.0, causing their route choices to be biassed towards the major routes.

Fields

Route Class: The route class to which the cost factor applies

Behaviour: The behaviour type to which the cost factor applies, or the wild card (All) in which case it will apply to all behaviour types.

Route Distance Factor: When calculating the route cost, the length of any branch of this route class is multiplied by this factor. It may be more appropriate to set the length of an individual branch explicitly, rather than setting a multiplication factor for all branches of this class.

Route Time Factor: When calculating the route cost, the travel time calculated for any branch of this route class is multiplied by this factor. Travel time is calculated from branch length divided by branch speed. It may be more appropriate to set the length or the speed of an individual branch explicitly, rather than setting a cost factor for all branches of this class.

Route Price Factor: When calculating the route cost, the price of travel (from toll charges) for any branch of this route class is multiplied by this factor. It may be more appropriate to set the price of an individual branch explicitly, rather than setting a cost factor for all branches of this class.

Route Splits for Undirected Trips

Undirected demand is simpler than directed (or OD) demand, as you need only a single volume value for each origin area. However, if there are any route choices in your network, you need to specify the turning counts or “route splits” at each route choice location for each possible destination. If your model is fairly simple, and has only one route between each origin and destination pair, for example a corridor model, then this method of specifying demand may be suitable.

Branch: This is a collective term used for both links and walkways. Undirected trips and route splitting can be applied to both vehicle and pedestrian traffic.

Route Choice Locations on Links: There is a route choice location at the end of any link that has two or more exits.

Route Choice Locations on Walkways: Walkways are bi-directional, so there can be a route choice location at both ends, depending on the number of exits for foot traffic travelling in that direction. The nominal “forward” direction of a walkway is shown by arrows on its centre line, but this designation is for convenience only, there is no difference experienced by a pedestrian travelling in one direction, compared to the other.

Default Split: Each route choice location automatically creates a Split object, assigning an equal proportion of all incoming traffic to each of the available exit options. At a 4-way intersection, an incoming agent has a choice of left, right or straight ahead, and the default split object will stochastically assign the agent to one of those exits, each with equal probability of 1/3. This means that if there is a loop in the network, an agent may be directed around the loop, arriving back at the same point. With typical splits, the probability of this happening is low, and in some networks, following a loop around may reflect reality, for example when a driver is looking for on-street parking.

To modify the Route Splitting Parameters: Select the branch in the graphics panel, then for links, select the action:

- select the action *Road* > *Route Splitting*

for or walkways select the action

- *Walkway* > *Route Splitting (Forward)* OR
- *Walkway* > *Route Splitting (Reverse)*

In the Adjust Splits window that is raised, there will be one column for each exit, and one value in each column. After you change any of the values in one of these exit columns, you should press apply to see the normalized value.

Term: Each Split has a term, a time period over which it is active. You may want to define several splits for each route choice location, for example one for the AM peak, and another for the PM peak.

Mob: Each Split can be restricted to apply to only those people or vehicles whose behaviours are included in a mob. For example, you might want to define one split for cars (car driver behaviours), and another for trucks (truck driver behaviours). If no mob is specified for a split then it applies to all people.

Override: Use this to override the normal lowest-cost-to-destination route choice algorithm for people or vehicles with a destination. That is, route split values are normally for undirected trips only, but selecting this option uses the values for directed trips as well.

Normalization
% : the sum of the values in the exit columns is 100%. So for example, if there are three possibilities, and you want 20% to turn left, 70% to go straight ahead, and 10% to turn right, you should enter the values 20, 70 and 10.
Absolute : this setting allows you to enter absolute turning count values, which will be converted to percentages on pressing Apply. You can also paste in values from a spreadsheet.

Display: show the route split values by colouring the streams

Area & Zone Routes

Route Parameters Tab

Feedback Interval: This parameter specifies how often routes are recalculated. The cost of travelling each link is modified according to the travel times experienced by agents in the model and as a result, the recalculation may identify a different route as being the lowest cost between any origin-destination pair. Updated costs are visible only if **Behaviour / Routing / Dynamic Routing** is on

Smoothing If route cost feedback is on, this parameter, **α,** is used as the smoothing factor for simple exponential smoothing.

$$\begin{aligned}\text{Smoothed Cost} &= \boldsymbol{\alpha}\,(\text{New Cost}) + (1-\boldsymbol{\alpha})\,(\text{Old Cost}) \\ &= (\text{Old Cost}) + \boldsymbol{\alpha}\,((\text{New Cost}) - (\text{Old Cost}))\end{aligned}$$

α = 0.9 results in a more responsive system
α = 0.1 results in a more stable system

Decay: If route cost feedback is on, the decay parameter, **β**, is used to cause the recorded cost to move back towards the base free flow cost over time

$$\begin{aligned}\text{Decayed Cost} &= \boldsymbol{\beta}\,(\text{New Cost}) + (1 - \boldsymbol{\beta})\,(\text{Base Cost}) \\ &= (\text{Base Cost}) + \boldsymbol{\beta}\,((\text{New Cost}) - (\text{Base Cost}))\end{aligned}$$

Look Ahead: This parameter allows control over the number of branches which will be interrogated ahead of the current position to find if there have been any changes in cost to the destination since the last route table rebuild. That is, this can “peek” at the changing costs, without requiring a full table rebuild. However, use of a high value will cause a serious degradation in run-speed.

Gateways Sub-Tab

You can mark any branch as being a gateway. Where a Gateway is defined, any trip that passes through the gateway will be split into two hops, and the route spreading algorithm will be applied to each hop individually. This creates a more realistic assignment pattern.

There can be as many gateway branches as required, but best performance will be achieved with a low number. To create a gateway, first select a branch in the Graphics Panel, then press Add Gateway.

Delays Sub-Tab (Zones Only)

Defining turn time delays can give more realistic routes and a better assignment. You can define turn times:

(a) for individual turns, for all behaviours
(b) for individual turns, for individual behaviours
(c) for all turns of the same turn movement priority (Signalised, Yield, Give Way, Stop Sign). For example, setting a time of 00:02:00 for Stop Signs states that an average delay of 2 minutes will be experienced at each Stop Sign. A typical use of this is to keep drivers on main roads as any back street which ends in a Stop Sign will be penalised.

For (a) and (b), to define a turn time for an individual turn, first select a lane, then press 'Add...', and enter turn times, in seconds, for all exits. The turn times added will apply to all behaviours. You can then select a behaviour for the turn times just added if you want to make them behaviour-specific.

If the default time for Signalised Turns is set to zero, then the system will estimate a delay for each signalised turn equal to half the cycle time at that intersection. This is intended to be an estimate of the delay that a driver would expect if a route through the signalised intersection was taken. An expression to calculate the mean delay at an intersection would be more complex.

Snapshot: The Snapshot button (camera icon) records the current recorded turn times for each routing branch, and stores these values in the model. You can use this to store what you believe to be a realistic pattern of delay which will then be used as the base line (the expected delay) in future simulation runs. After pressing the snapshot button, press Apply then Save the model.

Export Delays: Press this button to export the current delays to a spreadsheet. If exported files follow a naming scheme and are stored in the correct location, they can be imported automatically. For example, you might save typical delays for 07:00, 08:00 and 09:00 from you base model, and have these read in automatically to your design model.

Import Delays: Press this button to manually import delays from a spreadsheet in the same XLS format that is exported above.

To import turn delays automatically, store XLS files in:

(Model Resources Folder)/TurnTimes/VTT_HH_MM.xls

Route Viewer Tab

Origin & Destination: Select an origin in the list on the left and a destination in the list on the right. If both an origin and destination are selected, and a route is available between them, then the route will be highlighted in the graphics panel. If there is no route highlighted, check the Message Log for more information.

Cost of Selected: This displays the name and cost (in SCU) of the origin branch. If a zone or area with more than one departure branch was selected, only the first of these is displayed. To see the cost of others, select it in the graphical window.

Cost from Next: This displays the name of the lowest-cost exit for the chosen destination from the origin branch. The cost to the destination, including this branch is displayed below the name

Cost to Destination is the sum of the two values to the left.

Behaviour: if there is more than one behaviour, you can change the selector to show the routes taken by each behaviour.

Route Spreading: this slider illustrates how routes would be distributed for a given spread value. When a behaviour is selected, the slider is set to the current spread value for that behaviour. Moving the slider only visualizes the effect of changes to the spread. To change the value, use the Behaviour window.

Use Agents: Selecting this option will display routes by creating a set of dummy agents and calling the route choice logic for those agents. Using agents allows display of routes modified by tools and plugins, such as route choice rules. If there are generated trips, then those trips will be used as the basis for the dummy agents, so the % splits shown by the colouring will be representative of the actual trips generated in the model. If no trips have been generated, then dummy trips will be created on the fly, but these will not be representative of the demand, they will just choose the first type that matches the behaviour.

The route spreading slider is disabled if you select this option. The spread will defined by the value defined for the behaviour assigned to each of the dummy agents.

Lane Choice

The Lane Choice Tool allows you to define rules that control the lane usage on the current link for groups of vehicles in the model, by defining a range of suitable lanes.

Add New Rule, or Edit Existing Rules
Select a lane, then select action **Road > Rules for Lane Choice**.

A table of rules defined for this link will be displayed. You can delete, re-order or edit existing rules for this link. If there are no rules defined for this link, the table will be empty. Press **Add** to add a new rule, or **Edit** to edit an existing rule. You can double click an existing rule to edit it. To copy an existing rule, select it, then **Add**.

New Lane Choice Rule
This window has several panels, as described below. In all panels other than the outcome, you can leave the selectors blank, which has the effect of matching all vehicles.

Term: Activate a rule for a specific term (time period).

Distance (Start / End): If these values are set, the lane choice rules will be applied at a point somewhere between start and end, that will be different for each vehicle. Use this if you want to spread out the effect of a lane choice rule so that all vehicles do not try to change lane at the same point. If these values are blank, the lane choice rule will be active at the start of the link. See also the Link/ Adjust parameters.

Person Trip: Select vehicles by the behaviour (Mob) of the person driving, or by origin or destination Area.

Vehicle Trip: Select vehicles by type (Fleet), origin or destination. A Fleet is a set of Vehicle Types. If you require to filter on a single Type, create a Fleet that contains only that Type. Define a Sector (a group of zones) if you want to collectively address all vehicles travelling to or from a region larger than a single zone.

For public transport, you can match by service, or by next stand.

Network: Activate rule using a message on a sign. The sign can be anywhere on the network, it does not need to be attached to the link for which the rules are applied. Plugins can control the messages displayed on signs. For example, the Parking Directions plugin can display on a sign the number of spaces available in a cluster. Then, a rule can control lane choice according to that information.

Location: Select vehicles according to their current location and route selection:

- **Intersection Exit:** This will filter vehicles based on the exit they will take from the next N intersections, where N is any number greater than or equal to 1. For the simple case of next intersection only (where N=1): enter the intersection exit in the box. For advanced use (where N >= 2): enter intersection exits separated by semi-colons. For example 2;1;3 would select vehicles taking exit 2 at the first intersection, then exit 1 at the second intersection, and finally exit 3 at the third intersection. For drive-on-left networks, exits are numbered clockwise starting at 1 from the leftmost exit available. [For drive-on-right networks, exits are numbered anti-clockwise.]
  If you have already selected vehicles by destination, in the Trip Criteria, and you know that this means that all selected vehicles will be making a particular turn, because there is no alternative route for those vehicles, then you do not need to specify Intersection Exit.
- **Link Exit:** This will filter vehicles based on the exit that the vehicle will take from the current link. Link exits are numbered as for intersection exits.
- **Current Lane:** In most cases, this will be blank. If it is set, it will restrict the match to vehicles that are already in this lane.
- **Open Lanes:** This is used when the number of open lanes may change during the simulation, as is the case if there is a tidal flow control applied to the selected link. Unless you are using the tidal flow Plugin, it is best to leave this field blank.

Outcome

Lane Range: The target lane range for any vehicle that matches the criteria. The lane range may be a single lane, for example “2”, if a vehicle must be in lane 2, or it may be more than one lane, for example “3 – 4”, if either lane 3 or lane 4 is suitable. (This section is mandatory- the lane range must be set, or the rule has no effect.)

Defining Origin and Destination Zone Groups by Sector: The Sectors feature can be used to achieve this. See the Help on Sectors in the Network Section. It is not a requirement to have any sectors defined in the network. If no sectors are defined, vehicles can still be grouped by individual origin or destination zone.

Performance Tip

Define as few criteria as are necessary to select the group of vehicle you want to match. If you have selected a group of vehicles by trip, by selecting origin and destination, and you have a simple network with only one route available for that trip, then you do not need to set the “Exit Index” and “Next Exit Index” criteria.

View Rules: Displaying rules that are already defined

On the Layers pane, select **Lane Choice** to display all rules. If there are many types of rules on the network it may be easier to select **Assignment > Lane Choice** then **Filter View of Rules**. This will reveal an array of selectors that will allow you to display a subset of all current rules. If there are multiple rules for a link, the text will be prefixed by a count of the form **(1 / 4)** to indicate how many rules are defined, and which one is currently being displayed. If there are multiple rules on a link, you can cycle through the rules by pressing the Refresh button to the right of the Filter View check box.

Rule Order & Cascading Logic

The order of definition of rules for each link is significant. A vehicle will be controlled by the last rule which matches its parameters. This can make the rule definitions seem complex, but it is not a requirement that this cascading logic is used. If rules are defined such that each vehicle matches only one rule (or no rules) then the order of rule definition is irrelevant.

For example, if you have a 6-lane link, and you want vehicles to use either lane 1 or lane 6, and you want this to happen such that vehicles that enter the link in lanes 1 to 3 use lane 1, and those entering the link in lanes 4 to 6 use lane 6 then this can be defined as:

```
current lane = 1 →  outcome = 1-1
current lane = 2 →  outcome = 1-1
current lane = 3 →  outcome = 1-1
current lane = 4 →  outcome = 6-6
current lane = 5 →  outcome = 6-6
current lane = 6 →  outcome = 6-6
```

The order of definition of these rules does not matter, because they do not overlap, and each vehicle can match only one rule at a time. However, if you are comfortable with the idea of cascading logic, you can define the rules as follows:

```
(any vehicle)    →  outcome = 1-1
current lane = 4 →  outcome = 6-6    [overrides rule 1]
current lane = 5 →  outcome = 6-6    [overrides rule 1]
current lane = 6 →  outcome = 6-6    [overrides rule 1]
```

In this case, fewer rules are defined, but at the expense of the rules being more complex to understand.

Stream Choice

The Stream Choice Tool allows you to define rules that control the choice of stream leading to the next link (choosing the “next lane”)

Add New Rule, or Edit Existing Rules
Select a lane, then select action **Road > Rules for Stream Choice**. A table of rules defined for this link will be displayed. You can delete, re-order or edit existing rules for this link. If there are no rules defined for this link, the table will be empty. Press **Add** to add a new rule, or **Edit** to edit an existing rule. You can double click an existing rule to edit it. To copy an existing rule, select it, then **Add**.

New Stream Choice Rule
This window has several panels, as described below. In all panels other than the outcome, you can leave the selectors blank, which has the effect of matching all vehicles.

Term: Activate a rule for a specific term (time period).

Person Trip: Select vehicles by the behaviour (Mob) of the person driving, or by origin or destination Area.

Vehicle Trip: Select vehicles by type (Fleet), origin or destination Zone. For public transport, you can match by service or by next stand.

Network: Activate rule using a message on a sign. The sign can be anywhere on the network, it does not need to be attached to the link for which the rules are applied. Plugins can control the messages displayed on signs. For example, the Parking Directions plugin can display on a sign the number of spaces available in a cluster. Then, a rule can control stream choice according to that information.

Location: Select vehicles according to their current location and route selection:

- **Intersection Exit:** This will filter vehicles based on the exit they will take from the next N intersections, where N is any

number greater than or equal to 1. For the simple case of next intersection only (where N=1): enter the intersection exit in the box. For advanced use (where N >= 2): enter intersection exits separated by semi-colons. For example 2;1;3 would select vehicles taking exit 2 at the first intersection, then exit 1 at the second intersection, and finally exit 3 at the third intersection. For drive-on-left networks, exits are numbered clockwise starting at 1 from the leftmost exit available. [For drive-on-right networks, exits are numbered anti-clockwise.]
If you have already selected vehicles by destination, in the Trip Criteria, and you know that this means that all selected vehicles will be making a particular turn, because there is no alternative route for those vehicles, then you do not need to specify Intersection Exit.

- **Link Exit:** This will filter vehicles based on the exit that the vehicle will take from the current link. Link exits are numbered as for intersection exits.
- **Current Lane:** This will be set to the index of the approach lane selected for the rule. You cannot change this value, as the stream choice rule applies to a single approach lane only.

Outcome

This section is mandatory- you must select a value, or the rule has no effect.

- **Next Lane:** This is the lane the vehicle will use on the next link. Choose between a single value, in which case all vehicles will use that value, or multiple values, where you can specify the split between all available options. For example, if there are two available lanes on the next link, and you want a 40%/60% split you would put 40 in the first line of the table and 60 in the second line.

Defining Vehicle Fleets: a Stream Choice Rule can select on type using a Fleet, which is a set of Vehicle Types. If you require to filter on a single Type, you should use a Fleet that contains only that Type.

Defining Origin and Destination Zone Groups by Sector: It may be useful to be able to create groups of zones, to collectively address all vehicles travelling to or from a region larger than a single zone. The Sectors feature can be used to achieve this. See the Help on Sectors in the Network Section. It is not a requirement to have any sectors defined in the network. If no sectors are defined, vehicles can still be grouped by individual origin or destination zone.

Performance Tip

Define as few criteria as are necessary to select the group of vehicle you want to match. If you have selected a group of vehicles by trip, by selecting origin and destination, and you have a simple network with only one route available for that trip, then you do not need to set the “Exit Index” and “Next Exit Index” criteria.

View Rules: Displaying rules that are already defined

On the Layers pane, select **Stream Choice** to display all rules. If there are many types of rules on the network it may be easier to select **Assignment > Stream Choice** then **Filter View of Rules**. This will reveal an array of selectors that will allow you to display a subset of all current rules. If there are multiple rules for a link, the text will be prefixed by a count of the form **(1 / 4)** to indicate how many rules are defined, and which one is currently being displayed. If there are multiple rules on a link, you can cycle through the rules by pressing the Refresh button to the right of the Filter View check box.

Route Choice

The Route Choice Tool allows you to define rules that control the route choice decisions of groups of vehicles or pedestrians in the model. You can control the exit taken by all vehicles (or people) in any group you define, at any route choice location.

The Route Choice Tool is switched off by default. To switch it on for a model, select Assignment / Route Choice. This will switch the tool on, and also raise the rules window. Then follow the process below to create rules for a particular road or walkway.

Add New Rule, or Edit Existing Rules

For Roads
- select a lane, then Action **Road > Rules for Route Choice**

For Walkways
- select a walkway, Action **Walkway > Rules for Route Choice**

A table of rules defined for the selected road/ walkway will be displayed. You can delete, re-order or edit existing rules.. If there are no rules defined, the table will be empty. Press **Add** to add a new rule, or **Edit** to edit an existing rule.

New Route Choice Rule:

This window has three panels for walkways, and five for roads (vehicle-only panels shown in green). In all panels other than the outcome, you can leave the selectors blank, which has the effect of matching all vehicles or people.

Person Trip: Select by behaviour (mob), origin or destination of the person. For a route choice rule on a road, this selects a vehicle according to its driver (or passenger, for drop-off). [This filter does not apply to public transport vehicles]

Vehicle Trip: Select by type (Fleet), origin or destination of the vehicle.

Network: Activate a rule using a message on a sign. The sign can be anywhere on the network, it does not need to be attached to the link for which the rules are applied. Plugins can control the messages displayed on signs. For example, the Parking Directions plugin can display on a sign the number of spaces available in a cluster. Then, a rule can control lane choice according to that information.

Location / Current Lane: If this is set, it will restrict the match to vehicles that are already in this lane.

Term: Activate a rule for a specific term (time period).

Outcome / Exit: This section is mandatory. You must select a value, or the rule has no effect. This is the exit any matching vehicle or person will take. Choose between a single value, for exclusive use of that exit, or multiple values, where you can specify the split between all available options. So for example, if there were two available exits, and you wanted 40% to take exit 1 and 60% to take exit 2, you would put 40 in the first line of the table and 60 in the second line. For vehicles, the exits are numbered from kerb side to median side, but for walkways no ordering can be assumed.

Trail Following

Trails can be created for both vehicles and pedestrians ("agents").

The Trail Following Tool allows you to control the route assignment of all or some of the agents in the model, by forcing them to use one or more predefined trails.

A Trail is a sequence of roads or walkways, which can be used for Trail Following or for Validation. For Vehicles, Trails can also be used to define public transport Services. A Trail can contain the same branch more than once (i.e. circular routes are allowed), and it can start and finish anywhere.

The Trail Follower window shows a table, containing one row for each rule. The columns in the table are:

Name: A descriptive name for the rule.

Trail: The trail to follow if an agent matches this rule.

Follow %: This is a number between 0 and 100. This controls the proportion of agents which match the type and trip parameters which will follow a trail. If you do not want any vehicles to follow a trail, set Follow % to zero.

Mob: This allows you to select one Mob for each trail. As an example, you might define a Mob called “Business" containing all the business people, and the Trail Follower tool can then be used to force all matching people to use that route. Selecting "[Any]" causes an agent of any behaviour to follow this trail, as long as it also matches the origin and destination parameters.

Origin/Destination: These values allow you to select the agents that use the trail by origin-destination trip information. You can select either a single zone or area, or a sector, where a sector is a set of zones or areas. You can set either the origin or the destination or both.

Parking Choice

The Parking Choice Tool allows you to define rules that control the parking zone destinations of vehicles. It is useful if there is more than one possible parking zone destination which allows the occupant to reach the ultimate person-trip destination. It can be used to divert vehicles when parking areas become full, or if congestion build up during the model changes the relative attractiveness of one parking zone over another.

To create Parking Choice Rules, you must first enable the Parking Choice Tool, in the Tool Tab on the Left Pane of the Main Window.

Add New Rule, or Edit Existing Rules

Select a lane, then select action **Road > Rules for Parking Choice**.

A table of rules defined for this link will be displayed. You can delete, re-order or edit existing rules for this link. If there are no rules defined for this link, the table will be empty. Press **Add** to add a new rule, or **Edit** to edit an existing rule.

New Parking Choice Rule

This window has five panels. In all panels other than the outcome, you can leave the selectors blank, which has the effect of matching all vehicles.

Term: Activate a rule for a specific term (time period).

Person Trip: Select by behaviour (mob), origin or destination of the person. For a route choice rule on a road, this selects a vehicle according to its driver (or passenger, for drop-off). [This filter does not apply to public transport vehicles]

Vehicle Trip: Select by type (fleet), origin or destination of the vehicle.

Network / Sign: Activate a rule using a message on a sign. The sign can be anywhere on the network, it does not need to be attached to the link for which the rules are applied. Plugins can control the messages displayed on signs. For example, the Parking Directions plugin can display on a sign the number of spaces available in a cluster. Then, a rule can control lane choice according to that information. If you enter “ABC” in the text field beside the sign selector, the rule will be applied if the message on the sign equals ABC. If the sign displays a numeric message, you can begin the comparison message with either '>' or '<'. For example, if the sign says 10 and you enter “> 4” in the text field, then this will match.

Network / Parking Spaces: Activate a rule based on the number of spaces remaining in any parking zone. Select a zone, enter a number of spaces, and specify if the rule should trigger when the number of spaces is less than or greater than the threshold.

Location/ Intersection Exit: This will filter vehicles based on the exit they will take from the next N intersections. See Lane Choice for a full explanation.

Location/ Link Exit: This will filter vehicles based on the exit that the vehicle will take from the current link. Link exits are numbered as for intersection exits.

Location / Current Lane: If this is set, it will restrict the match to vehicles that are already in this lane.

Outcome / Destination: This is the new parking destination that will be applied to any vehicle that matches this rule. You can specify either a single destination, in which case all vehicles that match will be assigned that parking destination, or you can specify a split amongst up to 8 parking zones.

If there are more than 8 parking zones in your network, you can select which ones can be used here by selecting the “Choice” option in Zone / Adjust.

Gateway Choice

The Gateway Choice Tool allows you to define rules that add an intermediate destination, or “gateway”, for a person or a vehicle. This causes the person or vehicle to go **via** that gateway.

Create Gateways

Before defining a rule, you must have at least one gateway defined.

To define a person-trip gateway:
- select a walkway in the main window
- select Assignment / Person Routes
- select Route Parameters Tab
- press [+] to add a gateway on the selected walkway

To define a vehicle-trip gateway:
- select a lane in the main window
- select Assignment / Vehicle Routes
- select Route Parameters Tab
- press [+] to add a gateway on the selected lane

Add New Rule, or Edit Existing Rules

For Roads
- select a lane, then Action **Road > Rules for Route Gateway**

For Walkways
- select a walkway, Action **Walkway > Rules for Gateway Choice**

A table of rules defined for this link will be displayed. You can delete, re-order or edit existing rules for this link. If there are no rules defined for this link, the table will be empty. Press **Add** to add a new rule, or **Edit** to edit an existing rule.

New Gateway Choice Rule

This window has three panels for walkways, and five for roads (vehicle-only panels shown in green). In all panels other than the outcome, you can leave the selectors blank, which has the effect of matching all vehicles or people.

Term: Activate a rule only for a specific time period.

Person Trip: Select by behaviour (mob), origin or destination of the person. For a gateway choice rule on a road, this selects a vehicle according to its driver (or passenger, for drop-off). [This filter does not apply to public transport vehicles]

Vehicle Trip: Select by type (fleet), origin or destination of the vehicle.

Network / Sign: Activate a rule using a message on a sign. If a sign is used, it can be anywhere on the network, it does not need to be attached to the link for which the rules are applied. Plugins can control the messages displayed on signs. Then, a rule can control lane choice according to that information. If you enter “ABC” in the text field beside the sign selector, the rule will be applied if the message on the sign equals ABC.

Network / Parking Spaces: Activate a rule based on the number of spaces remaining in any parking zone. Select a zone, enter a number of spaces, and specify if the rule should trigger when the number of spaces is less than or greater than the threshold.

Location / Current Lane: If this is set, it will restrict the match to vehicles that are already in this lane.

Gateway: This is the new intermediate gateway destination that will be applied.

Auditor

The Auditor tool has several panels which allow you to check if some of the standard parameters have been changed from the default values.

Parameters

This panel show simulation time step and behaviour parameters, showing default value, value used in model, and a status icon. A red X indicates the value has been changed from the default

Links

This panel shows where route length, speed, price or cost factor have been modified in order to manipulate the routing within the model.

Connectivity

This panel allows you to control the display of unconnected roads, walkways and stands and to highlight zones or areas that contains no roads or walkways, respectively. This is useful to verify that the model network is connected as you believe. It is also good practice to remove any extra network objects that are not connected to the main model, to improve performance.

Validator

The Validator tool displays comparisons of modelled and observed values, where the values may be counts or travel times. The data can be either for vehicles or for people.

The observed values are entered in **Demand / Observations**. For more information on each type of count or travel time, including Visualisation of Inconsistent Turn Counts, refer to that window.

As the simulation runs, the modelled value for at each observation point will be updated, as will the comparison metrics, such as % difference and GEH. The absolute count is shown as well as a normalized value, displaying the equivalent count at the end of the measurement term. For example, if the term is an hour in length, then the normalized counts are vehicles or people per hour.

Filter: The filter panel allow the display of only those measurements that match the selection

Table Refresh: The frequency of table refresh can be selected as every minute or every second. If every second is selected, this may cause the simulation to run significantly slower, if there are large numbers of observations.

Display: This panel enables colouring of the roads or walkways according to the GEH or difference values.

Save: The observed, modelled and comparison values can be saved either in a single large table at the completion of the Simulation Term or in smaller individual tables at the end of each observation Term.

GEH requires Hourly Flows: While any term can be used, it is recommended that you use hourly flows, if you use the GEH statistic, as GEH applies to hourly flows only.

Obstruction Monitor

The Obstruction Monitor removes:

I. any vehicle that is obstructing traffic and causing the model to lock up.
II. any person obstructing the movement of people

It is designed to be used as part of model checking and calibration so that excessive delays are not allowed to build. The overall aim is that this tool is used to identify problem areas which the modeller then rectifies. A final (calibrated) model should not use this tool as it is an indication that the assignment is wrong.

In the early stages of building a simulation model, it is often the case that one or two wayward agents can cause total gridlock by becoming blocked and obstructing key links or walkways in the network. This tool monitors the progress of all agents and “lifts out” any blocked agent to a “quarantine area”. A log and visual display highlight all quarantined agents, showing origin, destination, type, etc., aiding the modeller to fix the root of the problem while keeping the model running.

Threshold Delay: This is the key parameter for this Tool, as it determines how long an agent may be stationary before being quarantined. Each agent has a counter that is reset when it moves onto a new link or walkway. If the agent is still on that link after a threshold delay, then the agent is removed from the model. There is an entry in the table for each agent removed.

Choose a threshold delay that is long enough to allow normal travel time along the longest or slowest link. A value of zero switches the function off. To change the threshold, open the Obstructions window and use the sliders to set the threshold times for vehicles and pedestrians.

As the model runs, vehicles that have been on any link for more than the threshold delay minutes are cleared from the model. A log of all quarantined vehicles, is output to the Lost & Obstructing Agents window; a row for each cleared vehicle.

- Trip – Trip definition from origin to destination
- Start – The time the agent first loaded onto the network (HH:MM:SS)
- End - The time the agent was cleared from the network (HH:MM:SS)
- Type – Agent type
- Location – The location where the agent was cleared from the network
- Distance – Distance from the end of the link or walkway showing the exact position where the agent was cleared (distance 0 means that the agent was cleared as it waited at the stop line)
- Exit – The exit index and the next exit index (for example 2/3 means the cleared vehicle was going to exit 2 from the current link and onto exit 3 from the next link)

The information in this window can also be saved to a log file.

Results

The Results window allows you to control when data is saved, and what type of data is saved as the simulation runs.

Master Control

- **Save No results:** If selected, this option switches off the saving of results. This is the initial setting for each new model, because in general, you will not want to save results until you have built the model and watched it run it a few times, changing parameters as it runs, in order to fine-tune the performance of the model.
- **Save in AZA file:** If selected, this causes results to be saved within the model file itself (with file extension .AZA). This is a suitable choice for small to medium size models that will be run in interactive mode only, or in single batch runs. (Concurrent batch runs cannot save data to the model file, as that might result in two processes conflicting when trying to save results to a single file.)
- **Save as XLS [CSV] File in Resources:** If either of these options is selected, the results will be saved as separate files, in a folder named Resources/[Scenario]@[Date/Time]/. The XLS option produces Microsoft Excel files, and condenses several tables into each file, using multiple worksheets. The CSV option saves each table as a separate comma-separated text file. Which format you use depends on how you will process the results. You may prefer to compare text files with Text-file comparison and processing tools, other people prefer to analyse the data with Excel.

Timing of Results

By default, results are saved once, on completion of the model run at the end of the Simulation term. In addition, you can specify times for intermediate saves during the run. These intermediate times can also be used to clear accumulated values.

For example, if you are studying a peak hour of 08:00 to 09:00 with a 15 minute warm up, so that the simulation term is 07:45 to 09:00, use:

Time	Save	Clear
08:00		x
09:00	x	

Table Selection

You can choose to select any combination of the following types of results, according to your objective, and the nature of your model

Person Summary: Aggregated time, distance and other information for all person-trips in the model. This is reported as total PHT and PKT, and also reported for each mode – for example time and distance walking, in private car, on public transport. Waiting time is also reported. The tables saved to the results file match the layout and content of the **Reporting / Person Trips** window

Freight Summary: Aggregated time, distance and other information for all freight-trips in the model. This is reported as total FHT and FKT, and also reported for each mode. Waiting time is also reported. The tables saved to the results file match the layout and content of the **Reporting / Freight Trips** window

Cyclist Summary: Aggregated time, distance and other information for all cycling trips, as displayed in the **Reporting / Cycle Trips** window

Vehicle Summary: Aggregated time, distance and other information for all private vehicle trips (cars, trucks, taxis, etc), as displayed in the **Reporting / Private Vehicle Trips** window.

Transport Summary: Aggregated time, distance and other information for all transport vehicle trips (buses, trains, trams, etc) as displayed in the **Reporting / Service Vehicle Trips** window.

Person/ Cyclist/ Vehicle Detail: Individual time and distance information, reported per trip. Two tables are created, one for completed trips, the other for incomplete trips.

Transport Detail: Information on individual transport service trips, as well as comprehensive information by stand and by passenger. This saves all the information that is displayed in the **Reporting / Public Transport Detail** window.

Person Trips Summary Window

The Person Trips window show aggregated information for all person-trips, and origin-destination tables of counts of successful departures, unsuccessful (unreleased) departures and arrivals.

Summary Tab: the person trip summary reports time, distance and other measures for all people in the model, as a total for all modes used, and broken down by mode.

The first five rows show

- PKT Person-Kilometres Travelled
- PHT Person-Hours Travelled

These are shown for C, I and U trips:

- C: Completed trips
- I: Incomplete trips (that have started)
- U: Unreleased trips (not started, due to congestion at origin)

For example, PKTC means PKT for Complete trips.

Time and distance are displayed for:

- Walking
- Driving
- Passenger (on public transport)
- Waiting (time only, no distance)
- Cycling
- Taxi
- Drop-off trips (passenger in private car)
- Pick-up trips (passenger in private car)

The time measured for unreleased trips is not included in either the completed trip time or the incomplete time.

Display

If you tick any of the boxes in the Display column, this measure will be displayed on the main graphics window. You can vary the Text Size using the control at the bottom of the window, and also choose between the Total value and the Mean value

Filters

By default, the window shows information for all trips, but you can set filters to display only a subset of all trips. The filters must be set before simulation commences.

You can filter by origin or destination, but you must first define a Sector (a set of Areas or Zones). If you wanted to look only at trips going to a single Zone, create a Sector containing that Zone only, and then select that sector on the destination selector in the filter bar.

You can also filter by Mob, allowing you to display trips for a group of behaviours only.

Save

By default, the information will be saved at the end of the simulation, or at any other time specified in the Display / Reporting window. Pressing the Save button causes the information to be saved in tis current state.

Freight Trips Summary Window

The Freight Trips window show aggregated information for all freight-trips, and origin-destination tables of counts of successful departures, unsuccessful (unreleased) departures and arrivals.

Summary Tab: the freight trip summary reports time, distance and other measures for all freight units (or Packages) in the model, as a total for all modes used, and broken down by mode.

The first five rows show
- PKT Package-Kilometres Travelled
- PHT Package-Hours Travelled

These are shown for C, I and U trips:
- C: Completed trips
- I: Incomplete trips (that have started)
- U: Unreleased trips (not started, due to congestion at origin)

For example, PKTC means PKT for Complete trips.

Time and distance are displayed for:
- Processing (moving along channels at fixed speed)
- Driving (in dynamic-route private vehicle)
- Transport (on fixed-route “public” transport)
- Waiting (time only, no distance)
- Drop-off trips (in dynamic-route private vehicle)
- Pick-up trips (in dynamic-route private vehicle)

The time measured for unreleased trips is not included in either the completed trip time or the incomplete time.

Display

If you tick any of the boxes in the Display column, this measure will be displayed on the main graphics window. You can vary the Text Size using the control at the bottom of the window, and also choose between the Total value and the Mean value

Filters

By default, the window shows information for all trips, but you can set filters to display only a subset of all trips. The filters must be set before simulation commences.

You can filter by origin or destination, but you must first define a Sector (a set of Areas). If you wanted to look only at trips going to a single Area, create a Sector containing that Area only, and then select that sector on the destination selector in the filter bar.

You can also filter by Mob, allowing you to display trips for a group of behaviours only.

Save

By default, the information will be saved at the end of the simulation, or at any other time specified in the Display / Reporting window. Pressing the Save button causes the information to be saved in its current state.

Cycle Trips Summary Window

The Cycle Trips window shows aggregated information for all cycle-trips, and origin-destination tables of counts of successful departures, unsuccessful (unreleased) departures and arrivals.

Summary Tab: the cycle trip summary reports time, distance and other measures for all cyclists in the model:

- **Completed Trip Cycle Distance:** the sum of all incremental movements of distance by all vehicles. This may not be exactly the same as the length of all lanes and streams travelled, as a vehicle may not always follow the centre line of a lane or stream
- **Completed Trip Cycle Time**: the sum of all time increments from the time at which the vehicle was created in the model to the time of arrival. The time of creation may be after the trip start time, due to congestion at the origin. The additional time will be recorded in the unreleased time.
- **Unreleased Time:** the sum of all time increments for vehicles which could not be released at the trip start time because of congestion at origin
- **Incomplete Cycle Distance/Time:** reported for all trips that have not yet arrived at the destination

Filters / Display / Save – see Person Trips window

Private Vehicle Trips Summary Window

The Vehicle Trips window show aggregated information for all vehicle-trips, and origin-destination tables of counts of successful departures, unsuccessful (unreleased) departures and arrivals.

Summary Tab: the vehicle trip summary reports time, distance and other measures for all private vehicles in the model:

- **Completed Trip Drive Distance:** the sum of all incremental movements of distance by all vehicles. This may not be exactly the same as the length of all lanes and streams travelled, as a vehicle may not always follow the centre line of a lane or stream
- **Completed Trip Drive Time:** the sum of all time increments from the time at which the vehicle was created in the model to the time of arrival. The time of creation may be after the trip start time, due to congestion at the origin. The additional time will be recorded in the unreleased time.
- **Unreleased Time:** the sum of all time increments for vehicles which could not be released at the trip start time because of congestion at origin
- **Extra Distance / Time:** if extra people have been defined for a vehicle type, this is the sum of the distance and time for those extra people. This can be used to get a value for PHT and PKT for the driving part of a trip without defining person-trips.
- **Incomplete Drive Distance/Time:** reported for all trips that have not yet arrived at the destination

Filters / Display / Save – see Person Trips window

Service Vehicle Trips Summary Window

The Service Vehicle Trips window shows aggregated information for all public transport service vehicle trips. These are fixed-route vehicles, so no origin/destination departure/arrival information is recorded here. For service-specific information, refer to the Public Transport Detail window.

Summary Tab: this tab reports time, distance and other measures for all public transport vehicles in the model:

- **Completed Trip Transport Distance:** the sum of all incremental movements of distance by all vehicles, for all trips where the vehicle has reached the last stand on the service. This may not be exactly the same as the length of all lanes and streams travelled, as a transport vehicle may not always follow the centre line of a lane or stream
- **Completed Trip Transport Time:** the sum of all time increments from the time at which the vehicle was created in the model to the time of arrival at the last stand. The time of creation may be after the trip start time, due to congestion at the origin. The additional time will be recorded in the unreleased time.
- **Unreleased Time:** the sum of all time increments for vehicles which could not be released at the trip start time because of congestion at origin
- **Completed Passenger-Trip Distance / Time:** distance and time summed for all passengers on each transport vehicle, for each completed transport vehicle trip. (Passenger-trip distance and time for completed person-trips is shown on the person-trip summary window)
- **Extra Distance / Time:** if extra people have been defined for a transport type, this is the sum of the distance and time for those extra people. This includes time and distance for the driver of the vehicle, but does not include any other staff. These values can be used to get PHT and PKT for the driving part of a trip without defining person-trips.
- **Incomplete Transport/Passenger/Extra Distance/Time:** reported for all transport vehicle trips that have not yet arrived at the last stand on the service.

Filters / Display / Save – see Person Trips window

Parking Window

The Parking Summary Window shows information on parking bay usage. For advice on how to create Parking, see the Parking chapter in the Tutorial..

Zone: The name of the zone

Type: The type of the parking zone: on-street, off-street, etc

Bays: The total number of parking bays in the zone

Full: The number of bays currently occupied, including where there is a vehicle on approach, or a vehicle just leaving.

Spaces: The number of bays currently unoccupied. (= Bays – Full)

Occupancy: The occupancy ratio

Arrivals: A count of the number of vehicles to have occupied a parking bay in this zone since the start of the simulation term.

Departures: A count of the number of vehicles to have vacated a parking bay in this zone since the start of the simulation term.

Stay Duration: The average duration of a stay in a parking bay in this zone

Popularity by Colour: Colour the lanes according to popularity. Best bay numbers can be displayed using the Parking Names Layer in the Layer Tab.

Route: If this is selected, and the **Priority Turns** Layer is also selected, then the preferred route around a parking area will be shown. The preferred route can be set using the **Preference** field for a Turn.

Transport Service Reports

The Service Reports window has eight tabs, described below:

Services

This table has a row for each service route, and each column shows a statistic summarised for all passengers and all vehicles on that service. Selecting a service highlights all the links on the service route.

Name: the short service name, often just a number

Description: the long descriptive name for the route

Transports: the number of vehicles started on this service

Unexpected Wait: the mean and standard deviation of times spent waiting after the expected arrival time of the first bus. This is non-zero if a bus is late, or if the first bus is full.

Travel Time: the mean and standard deviation of passenger times spent travelling on one of the vehicles on this service.

Crawl Time: the total time all vehicles on this service have spent at less than the crawling speed (8 km/h). This includes time spent waiting at traffic signals and in congestion.

Dwell Time: the total time spent at stops on the route, for all vehicles on this service.

Cruising Time: for all vehicles on this service, the proportion of time spent travelling at or above the cruising speed, as defined in the vehicle types window.

Average Speed: the average speed of all vehicles on this service, calculated by dividing the total distance travelled by the total travel time, which includes time spent at stops, but not time on layovers.

Boardings: the total number of boardings, for all vehicles on this service

Alightings: the total number of alighting, for all vehicles on this service

Distance-Averaged Load: the average of the distance-averaged load for all vehicles on this services, calculated as a fraction of the capacity of each vehicle, and expressed as a percentage

Time-Averaged Load: the time-averaged load (Passenger-Seconds / Total Trip Time)

Vehicles

This table has a row for each vehicle (bus / train/ tram, etc).

Name: A unique identifier for the vehicle

Service: The service this vehicle follows

Trip: the vehicle identifier, expressed as (route number)_(identifier). The vehicle identifier is a counter that starts at one for the first vehicle on the timetable for each service.

Travel Distance: the distance travelled

Travel Time: the total travel time, including dwell time, but not including any layovers

Crawl Time: the time spent travelling at less than the crawl speed(8 km/h). This includes time spent waiting at traffic signals and in congestion.

Dwell Time: the total time spent at stops on the route, for all vehicles on this service.

Layover Time: the total time spent at layovers

Cruising Time: the proportion of time spent travelling at or above the cruising speed, as defined in the vehicle types window.

Average Speed: the average speed, calculated by dividing the total distance travelled by the total travel time, which includes time spent at stops, but not time on layovers.

Boardings: the total number of passengers getting on at all stops

Alightings: the total number of passengers getting off at all stops

Passengers: the number of passengers currently on the vehicle

Minimum Load: the minimum number of passengers on the vehicle since it started

Maximum Load: the maximum number of passengers on the vehicle since it started

Distance-Averaged Load: the distance-averaged load, as a fraction of the capacity of each vehicle, expressed as a percentage. (Passenger-Distance / Total Distance)

Time-Averaged Load: the time-averaged load (Passenger-Seconds / Total Trip Time)

Stands

This table has a row for each stand. A stand is a physical location, and there may be several Stops associated with any one Stand.

Name: the name of the stand

Q Now: the total number of passengers waiting for all services at this stand

Q Avg: the time-averaged number of passengers at this stand

Q Max: the maximum number of passengers at this stand, since the start of the simulation

Vehicles: the total number of vehicles to have stopped at this stand

Boarding: the total number of passengers to have boarded any vehicle from this stand

Alighting: the total number of passengers to have alighted from any vehicle at this stand

Exp. Wait: the mean and standard deviation of expected waiting times.

The expected waiting time is the time from when a passenger is generated at the origin stop to the expected arrival time of the first vehicle. This is essentially an input value, as passengers are generated at a time relative to the first arrival.

Late Wait: the mean and standard deviation of times spent waiting after the expected arrival time of the first bus. This is non-zero if a vehicle is late, or if any vehicles were full.

Dwells

This table has a row for each dwell (a service/stand combination).

Name: the name of the stop, which is of the form (service)~(stand)

Q Now: the number of passengers waiting at this stop

Q Avg: the time-averaged number of passengers waiting at this stop

Q Max: the maximum number of passengers at this stop, since the start of the simulation

Vehicles: the total number of vehicles to have stopped at this stop

Next V: the time until the next vehicle is due, according to the timetable. The time counts down until the scheduled arrival time, prefixed by **Due**, and then if the bus is late, the time counts up, prefixed by **LATE**

Passengers

This table has a row for each passenger trip

Name: the name of the trip [Trip]@[Service]~[Stand Name]

Agent: the unique identifier for the person / freight unit

Service: the name of the service, appended with an underscore and the departure number once the passenger has boarded. For example, 23A_2 means the second timetabled departure on service 23A.

Board Stand: The name of the stand at which the passenger boarded

Board At: the time at which the passenger boarded a vehicle at the boarding (origin) stand

Alight Stand: The name of the stand at which the passenger alighted

Alight At: the destination stop

Expected Wait: the time between the passenger arriving at the stand and the ***expected*** (timetabled) arrival time of the next service vehicle.

Unexpected Wait: the time spent waiting after the expected arrival time of the first vehicle. The Unexpected Wait is non-zero if the vehicle is late, or it is full, and the passenger has to wait for a later vehicle.

Travel Time: the time spent travelling on the vehicle

Trip Time: the total trip time, including all waiting time and travel time

Trip Distance: the distance travelled while Onboard

Average Speed: the average speed of the passenger, calculated by dividing the travelled distance by the **Trip Time**

Walking Reports

This window shows all individual pedestrian trip components. Each pedestrian component may form the whole trip, but in many cases forms only part of the trip. For example, where a person arrives into the model in a train, then walks from the train station to a bus stop, then takes the bus, then walks from the bus stop to the destination area, a single person-trip has generated two pedestrian trip components, one from train to bus, the other from bus stop to area. There are three type of ends of possible trip components: an area (A), a stand(S), and a parking bay (P). A parking bay may be a drop-off point or pick-up point, or it may be a taxi rank. So there are 9 possible combinations: AA, AS, AP, SA, SS, SP, PA, PS, PP

Individual Walks

Name: the name of the trip

Type: the type of the person walking

Start/Finish: the locations of the ends of the walking component

Departure: the time the person left the start location

Arrival: the time the person arrived at the finish location

Time: the time taken to walk between start and finish, including any waiting time. If the finish location is a stand, this will include the time spent waiting for a bus or train to arrive.

Distance: the distance walked by the person between start and finish

Speed: the speed calculated from distance over time

Summary Statistics

Start/Finish: the locations of the ends of the walking component

Time: the mean walking time for this component

Distance: the mean distance walked for this component

Speed: the mean speed for this component

Colour / Width: These selectors can be used to highlight the display, showing a coloured “heat map” or “band-width” between each of the start-finish pairs.

Traffic Reports

This window shows traffic measurements for marked roads, loops and trails. To measure any of these objects, set the appropriate flag:

- Select Road (Link) / Adjust / Roads Tab / Measure [x]
- Select Loop / Adjust / Measure [x]
- Network / Trails / Measure [x]

Road Measurements

Vehicles: count of vehicles that have completed the link

People: count of all people in vehicles counted above

Mean Time: mean travel time for the link

Mean Flow: mean flow calculated from count of vehicles in measurement period

Mean Speed: mean speed calculated from travel time and link length

Mean Density: mean density, calculated from mean entry interval time over mean travel time

Density: instantaneous density, calculated from number of vehicles on all lanes of road divided by length of road

Loop Measurements

Occupancy Time: Instantaneous occupancy time. Zero if unoccupied

Non-occupancy Time: Instantaneous gap time. Zero if occupied

Occupancy Ratio: Ratio of total occupancy time to total time

Non-occupancy Ratio: Ratio of total non-occupancy time to total time

Flow: Flow calculated from count of vehicles between first and last event

Vehicles: Count of vehicles passing loop

People: Count of people in vehicles passing loop

Headway: The headway in seconds) between the last two vehicles, measured head to head

Gap: The head-to-tail gap (in seconds) between the last two vehicles (headway minus vehicle length/speed)

Vehicle Length: Mean vehicle length

Density: Density calculated form mean vehicle length

Gap Speed: Speed calculated from occupancy and mean vehicle length

Gun Speed: Mean of instantaneous speeds of vehicles at loop

Trail Measurements

All measurements apply to vehicles that have completed the trail – traversing all links, in sequence, from first to last.

Time: The mean travel time to complete the trail

Delay: The mean delay, calculated from the mean travel time, minus the free-flow time calculated from the Route Speed and link length

Vehicles: Count of all vehicles

People: Count of people in all vehicles

Mean Speed / Std Dev: Mean speed and standard deviation

Min / Lower Qt / Median / Upper Qt / Max (Speed): From the distribution of speeds, the minimum, lower quartile, median, upper quartile and maximum speed

Signal Reports

This window allows you to save a report of signal activity on selected intersections. It is possible to save signal changes by Phase, by Group or by Turn.

If selected, each method will create a report with four columns:

ID: a unique integer identifier, counting upwards from 1

Time: the time of the signal change, in format HH:MM:SS

Phase/Group/Turn: the phase, group or turn identifier, according to the selection made in the Signal Reports window

Signal: The signal setting for the phase, group or turn, recorded as G, Y, R, for green, yellow red, or combinations

Level of Service Reporting Tool

Level of service (LoS) is a measurement of performance, commonly dividing all possible measures into bands, and assigning a single letter to each band. In this tool, it is applied to delay on intersection approaches and also to delay, space and speed on walkways. There are many different definition of the bands; some are listed at the end of this section. A common definition uses 6 levels, each indicated by a letter, and ranging from A (best) to F (worst). There are several standard definitions of levels, for example, HCM-S = Highway Capacity Manual Vehicles at Signalised Intersection.

Level of Service for Intersections

The level of service for an "Approach" to an intersection is measured using the control delay on that approach. The control delay is defined as the additional time which a vehicle needs to traverse an intersection due to both the geometry and the controlling systems in place at the intersection (signals, stop signs or yield signs). So, in order to measure the level of service, one or more Approaches must be defined.

Approach: This may be just the link leading immediately to the intersection, or it may be longer, defined as a series of links. In a congested urban network, it may be difficult to determine automatically where the approach to an intersection starts, and for this reason, the user must specify a start link for each approach. From the start link, a sequence of links is built automatically, and this sequence of links is used to calculate the approach delay on the corresponding approach.

Adding Approaches: There are several methods of adding approaches:

- To add a single-link approach to an intersection, first select a lane on a link connected directly to an intersection, then press **+**
- For a longer approach, select a lane on the most upstream approach link,then select the intersection, then press **+**
- To add all single-link approaches for an intersection, select only the intersection, then press **+**

Once an approach is created, you cannot change the start link for an approach – you must delete it and add a new approach.

Exit: Modify this field to report the Level of Service for a single turning movement. If this is blank, all (remaining) turning movements are combined.

Measurements in Intersections Table

LoS: The Level of Service, usually a letter from A to F, set according to the average delay.

Control Delay: This is a time in seconds representing the total delay for vehicles in all lanes on an approach. The level of service is calculated from this delay. This value is the sum of the approach delay and the acceleration delay - see below for the definition of each.

Approach Delay: This is the difference between the time that a vehicle takes to traverse the entire approach to the intersection and the “free-flow time”. The free-flow time is the time that it would have taken to traverse that approach under free-flow conditions and no controlling systems. “Free-flow conditions” assume the vehicle is constrained only by the speed limit and its own maximum speed, and it will travel at its “free-flow speed” which is the lower of these two speeds. The speed limit may be different on each link on the approach, so the free-flow time for the approach will be summing the individual times for each link, calculated by dividing the link length by the free-flow speed for each link. Where the free-flow speed changes on the approach, the free-flow time will account for time required for the vehicle to change speed. The change in speed will always happen on the link which has the higher free-flow speed, and it is assumed that the vehicle will change speed at a constant rate - its maximum acceleration or deceleration rate. Approach Delay (and therefore Control Delay) is a rolling average, updated every time a vehicle moves from an approach onto the intersection.

Acceleration delay: This is the time needed by the vehicle to accelerate back up to the exit free-flow speed after it has crossed the stop line. This is approximated by assuming that the vehicle accelerates at a constant rate (its maximum acceleration rate) from its current speed up to the exit free-flow speed. The exit free-flow speed is the free-flow speed on the link before the intersection, because any other change in speed that the vehicle might need to make is not caused by the intersection control.

Queue Delay: The average delay experienced by queuing vehicles, where a queuing vehicle is one travelling at less than the queue threshold speed. The queue threshold speed is defined as 0.25 x the route speed on the first link of the appraoch. The queue delay is a rolling average value, updated every second.

Queue Size: The average queue size, in number of vehicles, counting those vehicles on the approach that are travelling at, or less than the queue threshold speed. The queue threshold speed is set as 25% of the free-flow speed on the first link of the approach. The average queue size is a rolling average, updated every second, and counts vehicles in all lanes. So if there are 3 lanes, and "Queue Size" is 21, then this equate to 7 vehicles per lane.

Queue Max: The maximum queue size, in number of vehicles, counting vehicles in all lanes on the approach.

Stops: A count of the number of stops experienced on this approach by all vehicles. Each vehicle may stop more than once on each approach. A stop is defined by the upper and lower threshold speeds defined in the Vehicle Motion section of the Parameters > Calibration window.

Turned: The total number of vehicles that have turned from this approach on to the intersection.

Approaching: The number of vehicles currently on the approach. This includes all vehicles, regardless of their speed.

Total: The sum of **Turned** + **Approaching**

Level of Service for Walkways and Crossings

The level of service for a walkway is measured either by the number of people per unit area on the walkway (the density), or by the average speed of walking. For a crossing, by the delay caused by the crossing controls.

To measure the level of service on a walkway or a crossing, first select it, and then press **[+].** You can select more than one at a time.

Measurements in Walkways Table

People: The number of people on the walkway

Length/Width/Area: The dimensions of the walkway

Density: The number of people per unit area. This is either sq.m. or sq.ft. depending on the unit system in use.

Space: The space in area units per person, inverse of density.

Speed: The average walking speed, in m/s or ft/s

Flow: The average number of people per hour passing a point

LosS: The level of service based on Space

LosV: The level of service based on speed (Velocity).

LosT: The level of service based on Time delay.

Standards used for Level of Service - Intersections

HCM Standard, Vehicles, Signals and Roundabouts

Level Of Service	Control Delay (seconds)	
	Lower Limit	Upper Limit
A	0	10
B	10	20
C	20	35
D	35	55
E	55	80
F	80	∞

HCM Standard, Vehicles, Stop Signs and Give Way (Yield) Signs

Level Of Service	Control Delay (seconds)	
	Lower Limit	Upper Limit
A	0	10
B	10	15
C	15	25
D	25	35
E	35	50
F	50	∞

NSW, Australia Standard, Vehicles, All Intersection Types

Level Of Service	Control Delay (seconds)	
	Lower Limit	Upper Limit
A	0	14.5
B	14.5	28.5
C	28.5	42.5
D	42.5	56.5
E	56.5	70.5
F	70.5	∞

Standards used for Level of Service - Crossings

HCM Standard, Pedestrians, Signalised Crossings

Level Of Service	Control Delay (seconds)	
	Lower Limit	Upper Limit
A	0	10
B	10	20
C	20	30
D	30	40
E	40	60
F	60	∞

HCM Standard, Pedestrians, Unsignalised Crossings

Level Of Service	Control Delay (seconds)	
	Lower Limit	Upper Limit
A	0	5
B	5	10
C	10	20
D	20	30
E	30	45
F	45	∞

Economic Evaluation Tool

The Economic Evaluation Tool creates a report listing detailed costs for each trip, and a summary of all trips within the simulation term. This can be used to assess the Benefit-Cost Ration (BCR) of a proposed design, or to compare the BCR of two or more designs.

Parameters Tab

Report Title, Report Subtitle (optional): These are displayed on the front page of the report spreadsheet that is produced. They can be modified directly in the spreadsheet, but saving the values here saves you typing them every time.

Total Trips for Normalisation (optional): If you run the simulation with a different trip set, that is, one produced with a different random seed, the simulation will produce different answers. The normalisation value is applied to the total values of distance, time, stops and the derived costs and this allows easier comparison when several different trip sets are used. For simple demand configurations, this value should be set to the sum of all values in the OD matrix. For complex demand configurations, with multiple demand matrices, demand profiles etc., it may be difficult to determine the exact value. In this case choose a value that is close to the average value for all the network options.

Save in CSV format: Save the data in CSV format, as several CSV files compressed into a single zip archive file, rather than a single XLS file. You may want to use this if you do not have access to a program that can read XLS format. For reading XLS, you may use Microsoft Excel, or an alternative package such as OpenOffice.

Save per-Type Pages: If this option is selected, a separate page (XLS: worksheet, CSV: file) will be produced for each agent type in the simulation, in addition to a summary page that aggregates the information for all types. If this option is not selected, only the summary page is produced.

Record Cost Values: If this option is selected, a copy of the cost input parameters page is saved with the output data, as a record of the prices that were used.

Record Parameters: If this option is selected, a copy of the general input parameters page is saved with the output data.

Record Assumptions: If this option is selected, a page is produced showing the assumptions and definitions used in the calculations.

Services Tab

This tab contains optional parameters only, which allow you to nominate certain public transport service routes as being of special interest, worthy of study. Whereas the Zones tab partitions the OD trips into four, the Study Routes selection creates 2 partitions in the report for public transport service vehicles:

- study routes (S)
- other routes (X)

2D Images & 3D Structures

This window displays the images and structures that have been loaded into the current model, and also shows a list of all available images and structures in the central shape database.

Shape: A collective term used for both two-dimensional images and three-dimensional structures.

Image: A 2-dimensional image, imported from a range of image formats, including JPG, PNG, GIF, BMP, TIF.

Structure: A 3-dimensional structure, used as scenery or a fixed obstacle or to give a body to an agent. Import formats include 3DS.

Central Shape Database: this binary-format database contains compressed image and structure data that can be used in any network. When uncompressed, it consumes memory and processor resources, so by default, none of the images or structures are included in a new network. Include only the structures you need for each model, otherwise you may find the application becomes sluggish, depending on your hardware.

Central vs Local Shapes: A central shape is one that is loaded into the database, and available for all networks on the installation on one computer. A local shape is one that is contained in the model definition file (the .AZA file) and can be easily transferred to another computer or another user in that file. An example of a local shape might be a JPEG-format aerial photograph image used as a background underneath a network. This is compressed automatically before being included in the .AZA file. The standard installation includes several hundred standard images and structures (including road signs, pedestrians, vehicles and buildings) but it is also possible to import additional images or structures to the central database on your installation. However, be aware that if you do this, and then use one of those shapes in a model, those shapes will not be visible if you transfer the model file to another computer, unless you also import the additional shapes into the central database on that installation.

Pedestrians Tab

To use a new character frame from the central database in your model, press the **Add** button in the top left corner. In the window that is raised, to include all frames for a character, select all rows in the table with the same character name (for example Character1) and press OK. To include all characters, you can use Control-A to select all frames of all characters. The standard database contains a number of characters, each having 11 frames. Of these 11 frames, the first, numbered 00, is a standing figure, while the others, numbered 01-10, are walking figures in 10 phases of a “walk cycle”. When moving, the walk cycle will be adapted to the forward moving speed of the pedestrian.

The individual cells of the table on the Pedestrians tab are not editable, they are there for information on the size of each frame structure. The first column, Status, is the most important. After loading it should show a green tick. On first opening a network that has already included some shapes, the icon may be a red cross, but this may be only because the shape has not yet been accessed. Try clicking on the row and this may cause the status to change. All shapes are not automatically loaded when a network is opened, as this would delay network loading unnecessarily. Such shapes will be loaded on demand.

Select a row in the table to temporarily display that character frame in the graphics panel. The figure will be drawn “life size” so you may need to adjust the zoom level. Recommended settings are Tilt = 45 and Zoom = 3. Use the arrow keys on your keyboard to move the selected row up or down the table. When you do this you should see each character “walk”.

Use Control-Select to de-select the character frame so that it is no longer drawn on the graphical display.

Vehicles Tab

To use a new vehicle shape from the central database in your model, press the **Add** button in the top left corner. Select the class of shape, then select the vehicles you want to include. As with any selection, you can use the shift and control keys to select a range or toggle selections on and off. You can also use Control-A to select all.

Most of the cells in the table on the Vehicles tab are not editable, they are there for information on the size of each shape. The first column, Status, is the most important. After loading it should show a green tick. [If not, see Pedestrians Tab section for explanation.]

Select a row in the table to temporarily display that vehicle shape in the graphics panel. The shape will be drawn "life size" so you may need to adjust the zoom level. Recommended settings are Tilt = 75 and Zoom = 10. Use the arrow keys on your keyboard to move the selected row up or down the table.

Group: this column is editable, and sets the shape group for the shape. Vehicle shapes are assigned to vehicle types in the model according to the group name. If the shapes you chose to include were cars, they will automatically be in the "Car" group. If you open the Types window from the Parameters menu on the main window, and switch to the "private vehicles" tab, then scroll the window to make the last few columns visible, you will see that one of them is "Shape Group", and in the row for the Standard type definition, contains the word "Car". This means that any vehicle of the Standard type released will be assigned a shape in the "Car" group. This matching is done by text matching alone – a group can be named anything you want it to be, and is created automatically when you change the group of a shape to the new name. So for example, to create three car groups, named "Small Car", "Medium Car" and "Large Car", you would assign one of those names to each of the car shapes you have included, then create three types in the type window, and perhaps also change the length parameter on the type definition to suit.

Structures Tab

To use a new fixed structure from the central database in your model, press the **Add** button in the top left corner. Select the class of shape, then select the structures you want to include.

Structures can be used as scenery, taking no active part in the simulation (for example buildings, or trees, beside a road), or they can be used as obstacles, to be avoided by pedestrians.

The individual cells of the table on the Structures tab are not editable, they are for information only. The first column, Status, is the most important. After loading it should show a green tick. [If not, see Pedestrians Tab section for explanation.]

Select a row in the table to temporarily display that structure in the graphics panel. The shape will be drawn “life size” so you may need to adjust the tilt and zoom settings. Use Control-Select to de-select the character frame so that it is no longer drawn on the graphical display.

Import: This button, below the Add button allows you to import a shape (for example, in 3DS format) into the local model file only, and not in the central database. If you have a 3-D model that you would like included in a particular model, perhaps a local building or landmark then you should use this option. You need to do this only once; after that the data for the model is included in your .AZA model file.

Important: Adding a definition of a structure in this window does not add any instances of the structure to your network. Once you have added a definition, you should then use the action **Display > New Structure at Cursor** to add an instance of the structure.

Images Tab

Images can be used for map-style overlays, road markings (bus lane etc.) or as flat vertical road signs for viewing in 3D. Images take no active part in the simulation, but are used for annotation and illustration.

Important: This tab allows you to import a definition of an image but does not add any instances of the image to your network. Once you have added a definition, you should then use the action **Display** > **New Image at Cursor** to add an instance of the image. In some cases, you may want to add multiple copies of an image (for example a common road sign) to your network, and this two-stage process cuts down on the amount of storage that would be required.

Important: While you can use this feature to import an aerial photograph image (for example JPEG), to use as a "tracing" overlay for building your network, we strongly recommend that you use scaled raster images (for example ECW format) or scaled vector graphics (for example DXF format) instead. JPEG images often do not have a very accurate scale, and you might build your network then discover all your measurements are wrong by 5% or more.

To use a (vertical) Road Sign or (horizontal) Map or Road Marking in your model, press the **Add** button in the top left corner. Select the class of image, then select the images you want to include. Road Markings are in the Map class, as this relates to horizontal images.

The individual cells of the table on the Images tab are not editable, they are for information only. The first column, Status, is the most important. After loading it should show a green tick. [If not, see Pedestrians Tab section for explanation.]

Select a row in the table to temporarily display that image in the graphics panel. The shape will be drawn "life size" so you may need to adjust the tilt, zoom and bearing settings.

For road signs, suggested settings are
Tilt = 90, Bearing = 0 and Zoom = 5.

For road markings suggested settings are
Tilt = 0, Bearing = 0 and Zoom = 5.

Use Control-Select to de-select the character frame so that it is no longer drawn on the graphical display.

In many cases, the same image appears twice in the list, with the suffix -Sign to indicate it will be drawn vertically, and with the suffix -Mark to indicate it will be drawn horizontally. However, you may move the images (using the handles on the surrounding box to be at any position in the network, or rotate them around a vertical axis, to suit your requirements.

Once you have added an image definition, go back to the main window and select the action **Display** > **New Image at Cursor** to add an instance of one of the images just defined. This instance can then be dragged, rotated, scaled, etc. just like any other object.

If the image is the wrong size, select the image and use the action **Display > Scale Image** to change the size. A popup window in the top-right hand corner of the main window will allow you to type in a new scale, or you can press and hold the arrow buttons to change the scale.

Advanced Tab

The Advanced Tab allows you to modify the central shapes database or to change settings for the 3-D rendering engine.

Shape Class: Each image or structure has a shape class, to make it easier to find in the database. This attribute is primarily a search key for the database, but is also used to identify particular types of structures for specific purposes.

Import 2D: press this to add a 2-dimensional image to the central shapes database. Browse to the file you want to use (in format JPG, PNG, GIF, BMP or TIF). Set the shape class, orientation and initial scale factor, and press OK. After successful importation, the image will be available for selection if you switch to the Images tab and press Add.

Import 3D: press this to add a 3-dimensional structure to the central shapes database. Browse to the file you want to use. Set the shape class according to the intended use for the structure. For example, if you want to use the structure as a shape for cars in the simulation, you select type “Car”. For certain shape classes, you may also be asked to set the spin rotation around a vertical axis (so that the car shapes know which end is the front). After successful importation, the image will be available for selection on either the Pedestrians, Vehicles, or Structures tab.

Refresh: Press to scan the database for up-to-date contents
Merge: this merges two databases into one.
Modify: this allows you to delete or rename shapes in the database.

Drawing: Set the level of detail in drawing images and shapes. To view images, such as aerial photographs, you must set this to “Detailed”. You can also change this setting on the tool bar on the main window.

Shading: This calculates normals to points or triangles, allowing the OpenGL rendering engine to apply shading to 3-dimensional objects. This is a computationally intensive operation.

Modify Shapes

This window allows you to rename or delete the shapes contained within the central database. You cannot change any of the other attributes of a saved shape.

Rename: To rename a shape, type the new name in the rename column. Renaming a shape requires the entire database to be re-indexed, so it is faster to enter all new names at one time, and then press OK, so that the indexer has to run only once. After you have entered all new names, press OK.

Delete: Select all the shapes you want to delete, then press OK.

Polygon Reduction Tool

Any 3-dimensional shape rendered onto a computer screen is composed of flat faces. Even a surface which appears curved, will be formed from flat faces if you look closely enough. The faces are commonly known as polygons, although in almost all cases the polygons are in fact triangles, having only three vertices. The speed at which a graphics engine draws a shape is directly proportional to the number of faces on the shape. As the shape gets further away from the viewpoint, the number of faces drawn can decrease, as it is not possible to distinguish the detail. It is therefore useful to have a tool that can reduce the number of faces on any shape, while still retaining the general form of that shape.

It is common practice to create several variations of a shape at decreasing levels of detail. For example, if the most detailed shape has 4800 faces, the next level of detail might have 1200, followed by 300 then 75. This is using a reduction target ratio of 25%. Which of these 4 **Level-of-Detail** or **LOD** shapes is displayed depends on the distance of the shape from the viewpoint.

Remove Duplicate Vertices: A shape often uses the same vertex for several faces – for example, the vertex at the corner of a cube is used for three square faces. When a shape is first imported, it may not use the same vertex object for each of these faces, and removing such duplication will reduce the space used to store the shape on disk and in memory. This operation can be quite time-consuming, so it must be selected as an option before reductions starts.

Auto-Promote after Reduction: “Promotion” means moving the second LOD shape up to position 1, to become the most detailed shape. If the Auto-Promote option is selected, the first round of reduction will have a single target, move that up to position 1, then queue another round of reduction to several levels of detail. Selecting this option allows you to queue a batch of newly imported files for reduction, by selecting each in turn using the up/down keys in the Shapes window.

Allow 2-Edge Vertices: Normally, each vertex in a shape will have at least 3 edges. If a vertex has only 2 vertices, it means that it is on the corner of a flat surface, which may be a mistake. These vertices are removed by default, but selecting this option causes them to be retained, if the edge selection algorithm allows.

Retain Material Edges: If a shape is composed of more than one material (different colour or texture) then this option reduces the likelihood of the removal of edges between two different materials.

Visible Level of Detail: The level of detail shown in the display when you select a shape in the Shapes Window.

Target: The target ratio for the number of faces of the next level of detail. For example, if the most detailed level has 4000 faces, and this is set to 30% the second level will have 1200 faces and the third level will have 360 faces. A setting of 100% is valid: it signifies a target of N-1, where N is the number of faces in the higher LOD.

Alpha: Increase this to force removal of all edges where curvature is less than the chosen value of Alpha. All edges with curvature below the alpha threshold will be given a very small, but non-zero curvature, ensuring that the edges are still sorted by length.

Beta: Decrease this to force comparison of edges by length alone, for all edges where curvature is more than the chosen value of Beta. All edges with curvature above the Beta threshold will be given the maximum curvature of 1, then sorted by length. Set Beta to 0 to select edges by length only, removing small details first.

Reduce: (Re-)Run the reduction algorithm on the shape selected in the Shapes window, with the (modified) parameters. The reduction algorithm will be run automatically when a shape is first selected in the window.

Promote: Promote level of detail 2 to level of detail 1, and re-run reduction algorithm with current parameters.

Save: / (over)

Save: Save modified shapes to the shapes.azb file and LODs to their own files. Shapes are marked as modified if they have been “promoted” - the original will be deleted from the database, and the new one added to the end. The LODs are not stored in the database, but in their own file. An LOD will be re-written if it has been re-reduced or promoted, or the LOD file is not already present.

How to add new pedestrian characters:

1. In Poser 7:
   (a) Figure / Delete Figure to delete standard character
   (b) load a character (low res)
   (c) hair / choose hair piece / double-click to apply
   (d) Window / Libraries / Universal / Waiting / Waiting 01 (or some other from Waiting or Standing)
   (e) Double-click to apply
   (f) save as man_00.3ds (or whatever), must end in _00.3ds
   (g) open walk designer (no walk path - use walk on spot)
   (h) use default walk style
   (i) one walk cycle is 30 frames
   (j) apply walk designer to generate 30 frames
   (k) camera must be in the same position for all frames
   (l) save all frames using multi-save
2. To save frame(s) as 3DS
   (a) File/Export/3DS (ALT-F, ALT-E, ALT-3)
   (b) Multi Frame - OK
   (c) De-select Ground object
   (d) OK
   (e) OK (Blank = "Export object groups for each body part")
   (f) choose file name - as man1.3ds - auto suffixed to man1_0, man1_1 up to man1_29
   (g) (Commuter import assumes _00, then _0, _1 etc.

Tags

Tags are used to highlight some people or vehicles of interest. A tag is drawn as a small triangle above the person or vehicle that is tagged.

There are two steps to using tags:

1. Define the tag by origin, destination, fleet or mob or a combination

2. Turn on the Tag Layer in the Layer pane

Defining Tags

- select **Display / Tags** to raise the Tags window
- Press [+] to add a new tag
- Set the colour
- Set an annotation label (optional)
- Set one or more of Origin, Destination, Fleet or Mob
- The Origin and Destination fields can accept Zone, Area or Sector names or descriptions. Type in the name, then press Apply to check that it has been recognised. If the cell is blanked out then it has not been recognised.

Tag Controls

- Single Tag / Multiple Tags: If single tag is selected, each vehicle or person will have at most one tag. If multiple tags is selected, a vehicle or person may have more than one tag, according to the match criteria for each Tag definition.
- Annotation: select this option to display the Annotation label for each tag

Other Controls

- Colour by Driver: If a vehicle has a driver, and the driver has a colour defined in the Person Type window, then the vehicle will be coloured according to the driver colour. This option is not directly dependent on tags, but it is another way of identifying or “tagging” vehicles.

Text Size and Font

This window allows you to change the font and the size of the text displayed in the graphical window.

Font

The Basic font is a vector font, designed for Commuter to be be fast to draw. It is not particularly pleasing to the eye, but it is effective and fast.

The list of other fonts depends on your computer, but should contain at least one Serif font (commonly Times New Roman) and one Sans-Serif font (commonly Arial). Any font other than the Basic font requires a lot of graphical computation, for a process known as Tessellation. It is recommended that fonts other than Basic are used only when simulation speed is not particularly important, for example when taking screen snapshots for reports. It is also possible to skip the Tessellation process by viewing in Outline mode.

Text Size

The Size controls allow you to change the size of the text for a range of different objects in the network. The text for these objects is normally off, but it can be switched on in the Layer panel if you select **All** or **Text** from the selector at the top of the panel, then switch th relevant Text Layer on in the current Aspect.

Surface Colouring

This window allows you to change the way that surfaces are coloured, to help understand the behaviour of the model. Surface colouring will not be effective in Outline mode, use Flat, Solid or Detailed mode.

For each colouring option, the object selected will be coloured in a temperature-styled colour ranging from blue (cold) through green (neutral) to red (hot). The numeric values to the right hand side of the selector control the extents of the range. Any surface whose value is equal or lower than the minimum value will be blue; any value equal or higher than the maximum will be red. Any surface whose value is equal to the midpoint value between minimum and maximum will **not be coloured**

Lane Colouring Options

- Height: use lane centre point height
- Gradient : use lane gradient
- Speed: use lane speed (speed limit)
- Route Speed: use route speed (for free-flow time calculation)
- Speed Diff: difference between lane speed and route speed
- Price: price (normally zero)
- Cost Factor: route cost factor (normally 1.0)
- Route Class: uses the index number of route class
- Lane Index: lane index number, used for lane choice and stream choice rules. This option is recommended to check lane index continuity.
- Headway Factor: headway factor (normally 1.0)
- Reaction Factor: reaction time factor (normally 1.0)
- Following Model: uses the index number of the following model applied to the link

Walk Colouring Options

- Height: use lane centre point height
- Gradient : use lane gradient
- Route Speed: use route speed (for free-flow time calculation)
- Price: price (normally zero)
- Cost Factor: route cost factor (normally 1.0)
- Route Class: uses the index number of route class

Draw Walls: Select this option to display walls on walled walkways

Depth: This option, and the variable value allows you to display surfaces with vertical sides, with the given depth (height). This is useful when capturing screen shots, but causes many more polygons to be drawn, so will tend to slow down the progress of the simulation, unless you have a very fast graphics chip.

Display When Closed: Continue to Display the Surface Colouring, even after the Surface Colouring window has been closed.

Reset: Reset Lane Colouring and Walk Colouring selections to Normal

Agent Log

The Agent Log, or Dashboard, window shows a table containing speed, gap to leader, acceleration and other values related to the following algorithm for any currently selected vehicle or person. This window is enabled only if you have the API license.

Diagnostics

This window contains a number of check boxes that you can use to better understand the interactions between people and vehicles in the model, by drawing lines between pairs of objects, or from objects to points of interest. Most should be self-explanatory.

People

- Selected: draw lines only from the currently selected person
- Trip: lines to the trip origin and destination
- Blocker: line to any current blocker
- Avoid: the target vector from the avoidance algorithm
- By State: following = blue, blocked = red, overlapping = green
- Path: line to end of current pathway
- Leader: line to any current leader
- Obstacles: lines to all obstacles

Vehicles

- Selected:
- Trip
- Blocker
- Avoid: [all as above]
- By State:
- Path
- Leader
- Obstacles
- Zip: indicates part of a lane merging zip operation
- Lane Change: indicates any current desire for lane change
- Rank, Lane: the rank of this in the current lane, and lane index
- Overlapper: line to any overlapping vehicle or person

Parking

- Selected [as above]
- Bay to Car: line from bay to car that has reserved the bay
- Bay to Person: line from bay to person that has reserved the car
- Car to Occupant: line from car just parked to occupant not yet arrived

Transport

- Next Stop: draw a line from a transport vehicle to its next stand

DWG Drawing Tool

DWG is a CAD data file format developed by Autodesk.

The DWG Drawing Tool allows you to open one or more DWG files, and display them as layers in the graphics panel.

Unlike the 2D Images accessed through the shapes database window, a DWG file is always external and is not imported into the Commuter .AZA file. Therefore, if you use any DWG files, you should make sure they are available to anyone to whom you pass your model. If the model is opened on another machine, and the path to a DWG file is not the same, it will not be visible.

Open: To open a DWG file, press the “Add” button on the top left of the window.

West/ South/ Base: The location of the western and southern edges of the image, and its base elevation, in the units used for your model. You can also drag the DWG object around the model using the “handle” at the south-west corner.

Scale: A scale factor applied equally to East-West and North-South. You can also stretch the DWG image using the handle at the North-East corner.

Show: If this is selected, the image will be drawn on the graphics panel. You can switch it off when you do not require the image, to improve the speed of response of the application.

Layers: After successfully adding a DWG image, there will be an extra set of layers in the drop-down box on the Layers pane on the left of the main window. There will be as many layers as were defined in the DWG file. You can make each of these layers visible in any of the Commuter layers. the default colour for each layer is white. If you leave it at the default, the colour for a layer will be taken from the definition in the DWG file. Otherwise you can change the colour of each layer to suit.

Export Frames to OpenMicroSim Format

This window allows you to record and export a sequence of frames, where each frame contains the location and direction of the moving agents in the simulation, for visualisation in an external package, such as Forum8. VR-Studio or UC-win/Road

Area By Symbol: By default the entire model area will be exported. However, if you define a symbol around the area of interest, then select that symbol, only those agents which are inside the skin of the symbol will be included. This applies to height as well as lateral position.

Duration & Interval: These values define the number of frames to be exported. If the Duration is 00:04:00 and the interval is 5, and the current time step is 0.2 seconds, then the number of frames exported will be (4 x 60) / (5 x 0.2) = 240 frames.

Pedestrians: Select this to export pedestrian locations as well as vehicles and traffic signal states

Save As & Record: Press the Record button to open a file browser, then choose the name of the destination file. This file name will be saved in the Save As field.

Once the Record button is pressed and the file name is set, recording is ready to start, and will commence when the simulation Play button is next pressed. Recording will continue until the specified number of frames have been saved.

For more information on the export format, see openmicrosim.org

Tabs

Reference

File Tab

Visible Folders and Files: The file tab within the left pane shows the structure of the folders on your computer, and the Commuter files within that folder structure. Files of other types are not shown: only Commuter model files are shown in this view, and, for reasons of speed, the initial view will show only the folders from the root of the disk to the location of the currently open file. If you want to explore for sub-folders or other files, you should click on one of the folder icons, and all the sub-folders in that folder will then be shown. When you click on a folder icon, any Commuter files in your chosen folder that are not already showing will also be shown. The Files Tab also shows the structure of the contents of the currently open Commuter file. The contents are divided into Books, Chapters, Pages and Tables, as described below.

Opening a Model: Right clicking on a Commuter file will show a menu that allows you to open that model. A Commuter file can be recognized by its icon. The currently open Commuter File is shown in bold font. You can also open a model using File/Open, or using drag and drop or double-clicking in Windows Explorer.

Model Hierarchy: A Commuter model has a hierarchical structure:

- each Model contains 3 "Books": Scenarios, Components & Results
- each Book contains a number of Chapters
- each Chapter contains a number of Pages
- each Page contains a number of Tables
- each Table contains a number of Records

Scenarios / Components / Results

- Scenarios provides a method to quickly change from one set of inputs to another. For example, scenarios might be "Base", "Design1", "Design2"
- Components defines all the inputs for the model and contains a fixed number of chapters: Parameters, Network, Demand, etc.
- Results stores the output and contains a variable number of Chapters, where each Chapter contains the results from one simulation run. A Results Chapter name includes the date and time at which the simulation run was started.

Components Chapter (e.g. Parameters) Each chapter can contain any number of pages, where each page can be thought of as an alternative;

only one of the pages is active at any time. For example, the Parameters chapter may contain 3 pages named “Parameters-1”, “Parameters-2” and “Parameters-3”. Any run of the simulation will use only one of these pages. To use Parameters-3, you would set that page to be the active page and then run the simulation. When the Model Book is not expanded in the tree view, the chapter label shows the currently active page.

Results Chapter: each is generated by a run of the simulation, and contains results for that run only. A Results Chapter contains a number of Pages, where each Page contains results relating to a particular area of interest. It is often the case that a Results Page is generated by a tool or a plugin. For example, if validation reporting is selected, one of the Results pages will contain all the validation data.

Page/ Tables: A Page can contain any number of tables. For example, the Network page contains tables describing lanes, crossings, zones, etc., while the Trips page contains tables describing the trips that will be generated on a network during simulation.

Table/ Records: A Table can contain any number of records. A table is similar to a database table, in that it contains records, and each record contains a value for a number of fields. Unlike a standard (“third normal form”) database, the records in a Commuter table can contain “multi-fields”. A multi-field is one in which there are a variable number of values, all having the same field name and type. For example, a Zone is defined as having a variable number of boundary points. In the conventional view, each record is shown as a row in the table, and each field is shown as a column. However, the view may be transposed, which can be useful for tables that have a small number of records and a large number of fields.

Tool Bar Actions

- Refresh: (F5) refresh the view to include recently created items
- New: add a new scenario, or a new derivative page
- Activate: activate an existing scenario or page
- Rename: rename a scenario or page
- Delete: delete a scenario or page
- Export to Excel: Export results page as Excel files

Tables Panel

The Tables Panel is to the right of the Files Tab, underneath the graphics panel. You can move the divider to the left of the graphics panel to see more of the tables pane, or click on the double arrow at the top to move the divider completely to the right side.

The Tables Panel contains a list of table names on the left hand side, a tool bar along the top, while the main part of the panel contains a table showing the data in the currently selected table.

Tool Bar – Save: Allows you to save any changes you have made to data in the Tables Panel, and synchronize this view with the graphical view. **!!** If you have also made changes in the graphical view, and not yet saved them, those changes will be lost if you save from here. **!!**

Tool Bar – Transpose: Allows you to transpose the table in the view, which can be useful if there are a small number of records and a large number of fields.

Tool Bar – Compare: Allows you to compare a table in the selected page with the equivalent table in another page. The other page can be in the same file, or in another file.

Tool Bar – Add Record: Allows you to add a record to the currently selected table.

Editing Fields: It is possible to edit any data value in the Tables Panel, but take care not to enter an out-of -range value, as there is limited data validation in this view. Standard copy and paste mechanisms can be used, for single or multiple cells. So for example, if you want to modify all loops to be 4.5m in length, rather than 4.0m, type 4.5 into the first cell then copy this value using Control-C and paste it into all the other cells in the column using Control-V. Editing is possible in both normal and transposed view modes.

Action Tab

An **Action** is something you do to edit the currently active model. Most actions are reversible, using the Undo command. Example actions are

- Area/New Area at Cursor
- Walkway/New Crossing on Lane
- Lane/New Loop
- Road/New Intersection
- Adjust

Actions Tab & Right Mouse Menu: Actions can be accessed either from the Actions Tab or via a menu raised on the graphics window if you press and release the right mouse button. In either case, the effect of the action is the same. The only difference is the presentation of the actions: the right mouse menu shows only the currently available actions; the Actions Tab shows all actions.

An **Object** is something in your model to which an action is applied. A Crossing is an object, as is a Loop or an Intersection.

Actions, Objects & Adjust: Most actions apply to one type of object only. For example, Lane/New Stand creates a Stand object (for public transport) at a position as near as possible to the current cursor position on the currently selected Lane. It is not possible to add a Stand to any other type of object. The **Adjust** action is special in that it can be applied to many different type of objects.

Actions & Selections: Most actions apply to selections and are not available if no object of the correct type is selected. For example, Lane/New Stand is available as an action only if a Lane is selected. However, a few actions are available all the time, even if no object is selected. Examples of always-available actions are

- Area/New Area at Cursor
- Road/New Road at Cursor
- Walkway/New Walkway at Cursor

Always-available actions are the starting point to building a network – add Walkways to create a pedestrian network, add Roads to create a vehicle network, add Areas and Zones to define the origin and destination locations of the trips that use those networks.

Action Groups: All Actions apart from Adjust are in a group. The first group is Edit, whose actions apply to any selected objects. The Edit group contains generic actions such as Delete, Copy, Paste, Rename. Other groups are named for object types, and contain actions relevant to those objects. A group is always enabled, but its icon will be greyed out if none of the actions in the group is currently available.

Expanding/ Collapsing Groups: Each action group can be expanded by clicking on the icon next to the group name. The group can be expanded even when it appears disabled, which is useful when looking for a particular action, as it does not require the selection of an object of the appropriate type to enable the action group.

Parameter Actions: These actions can be recognised by their icon, “(x)”. A parameter action raises a window that allow you to modify the parameters used when adding new objects, but does not add any objects itself, nor does it make any change to the model that is saved. These actions modify only the parameters in memory relating to the next object to be added. For example, Road/New Road Parameters changes the parameters used when adding a new road. If you are adding several new roads of the same type (for example 2 lanes, 60 km/h speed limit) it is useful to be able to set those parameters in advance, then add all the new roads in quick succession, without having to OK the road parameters window for each road. The road parameters window contains a check box entitled “Display this window before adding a new road”. If this box is cleared, then roads of the same type can be added in quick succession.

Layer Tab

A **Layer** is a set of objects of the same type in the context of making those objects visible in the main graphics view. For example, the Areas layer, when enabled, causes all Areas to be visible in the graphics window, as long as the viewpoint is in a suitable position.

An **Aspect** is a collection of Layers, defining a colour for each Layer. So for example, Aspect 0 might show Areas in blue and Walkway surfaces in yellow, while Aspect 1 shows Walkway surfaces in orange and Intersections in green. In this way, it is possible to define up to 20 Aspects, each showing a unique combination of Layers and colours.

Different Colour for Each Aspect: Each Layer can have a different colour for each **Aspect**.

Selection Layers Always Visible: Some Layers, for example those relating to selections, are always on. This means that any currently selected object is always visible, even if the Layer for that object is off.

Left-Click Toggles Layer On-Off: Clicking with the left mouse button toggles a Layer on and off in any particular Aspect.

Right-Click Colour Editor: The colour for a Layer can be changed by pressing the right mouse button over the relevant coloured square.

Drag & Drop Colours: The colour for a Layer can be brushed sideways to other Aspects or to other Layers by pressing and dragging with the left mouse button. The “off” state can also be dragged in this way.

Currently Selected Aspect: The currently selected Aspect is shown by two vertical orange lines, and the number of the Aspect is highlighted in the header row of the table.

Layer Set Selector: The Layer set selection drop-down menu allows you to filter the set of Layers shown in the Layers Tab according to a number of pre-set criteria: All, Standard, Surfaces, Network, etc.

Reset: The Reset button allows you to restore the Aspect definitions to the “factory defaults”.

Import/ Export: The Import and Export buttons allow you to load the Aspect definitions from a file, or to save the current definitions to a file, respectively. This allows you to define a “standard” set of Aspect definitions for your organisation, or even for your own preferences, which is particularly useful if you have any form of colour blindness.

View Tab

A **View** is a predefined viewpoint position and view direction. It is analogous to defining a fixed camera position at a given Northing, Easting and Elevation, with the camera aligned to a given bearing angle and attitude angle.

You should create a view for each position of interest in your network. For example, you might want to create a view for each intersection, or each crossing, or each Stand on a Transport service.

Initial View: If the View List contains any Views, the first one in the list is used to initialise the viewpoint when a model is first loaded.

View List: All defined views are shown in a vertical list below the buttons on the View Tab. Clicking on the name of any View in the list moves the viewpoint to that view.

Add View: The Add View button defines a new view at the current viewpoint position. You are prompted to enter a name for the view, which can be edited later.

Up/Down: Use these buttons to change the position of a View in the list. The first View in the list is used when a model is first loaded.

Time: If a time (HH:MM) is set for a view, the viewpoint will move to the view automatically when the simulation clock reaches that time.

Smooth: If this is set the view will move more quickly between views, changing to flat drawing mode under some conditions.

Flights: This button opens the Flights or “Fly-Through” window.

A **Flight** is a sequence of **Flight Views**. A Flight View is a a named View, with a transition time and a hold time, which are real-time values, not simulation time. Once started, using the Play button, the Flight will progress through the Flight views, pausing for the Hold time. The FPS value controls how many intermediate points there are on the straight line path between one flight view location and the next.

Plugin Tab

A **Plugin** is a standalone utility that can be switched on for any network but is switched off by default. This has the advantage of keeping the software light and quick for simple networks, only loading code for more complex tasks when it is required.

Each Plugin has a configuration window can be raised from the Plugins Tab. A Plugin can be added by any Commuter user - the list of Plugins can be extended by the user, by writing a module that conforms to the Commuter Application Programming Interface.

Connect: To switch a Plugin on, click on the Plugin name, and select Connect. This will switch the tool on, and automatically raise the configuration window for that tool.

Configure: Once a Plugin has been switched on, its configuration window can be raised using the Configure option on its pop-up menu.

Disconnect: To switch a Plugin off, select Disconnect on its pop-up menu.

Connection Status Indicator: To the left of the Plugin name is a plug & socket icon. If the Plugin is enabled, the plug will be in the socket.

Operational Status Indicator: Between the plug & socket icon and the Plugin name is the operational status indicator. If the Plugin is switched off, this will be blank. If the Plugin is switched on, this will show a green tick if the Plugin is working properly, or a red cross if not. If a red cross is visible, check the Message Log window for any error messages. To raise the message log, click on the notepad icon in the bottom right corner of the main window.

Undo Tab

The Undo Tab holds a record of all actions that can be undone. The list is displayed with the most recent at the bottom. Pressing the Undo button undoes the last action and removes it from the list.

The Undo button on this pane has the same effect as Edit / Undo on the main window menu, and likewise, the same as Undo on the Action Tab menu.

The Undo list is similar to an audit trail, holding all the actions executed by the user since the network was opened.

Graphics Tab

The Graphics tab contains a Toolbar and a 3-D OpenGL visualization panel (the Viewport) which resizes automatically when you resize the main window.

There is a virtual slider bar along each edge of the Viewport. Open any model, then try dragging with Mouse-Button-1 on each of these bars to see the effect. You can change the Zoom, Tilt and Bearing settings like this. For more information on these settings, see the User Guide under Basic / Navigation In 3-D.

The virtual slider bar on the bottom edge of the window is a rotation bar, but it changes function if you have selected any objects. Without selection, the bottom edge bar rotates the view, but with selections, it is the selected objects that are rotated around the cursor centre point.

If the cursor is visible in the current layer, there will be a second bar above the bottom-edge rotation bar, the cursor bar. This cursor bar will be the same colour as the cursor: normally blue. Rotating the cursor bar rotates the far end of the cursor (the T-Point). This means that you can add new objects at any angle.

In the lower right corner, a display shows the (east, north, height) co-ordinates of the cursor. If the display is in Orthographic mode (by default, the action is bound to letter O on the keyboard) and the view is not tilted, then a map scale is also shown, which is corrected for your monitor's size and dot-pitch.

Graphics Panel Tool Bar

Save: Allows you to save any changes you have made to data in the graphical view and synchronize this with the data view on the Tables Panel. Important Note: If you have also made changes in the table view, and not yet saved them, those changes will be lost if you save from here.

Undo: Undo the last editing action. The Undo tab on the left pane displays a list of all undoable actions.

Find: Allows you to find and move the view point to a named object. Several types of objects can be used as the target [Crossing / Intersection / Stand / Node / Link / Loop / etc.]. it is also possible to enter raw location data [Easting and Northing] and move to that location.

Rewind: Allows you to reset the simulation clock to the start time for the model, and clear all the people and vehicles from the simulation.

Timing Mode [Continuous / Single Step / Timing]: Switch between continuous and single step simulation, or raise Simulation Timing window.

Simulation Timing Window

> **Simulation Play Interval:** an interval of any length for the simulation. If the interval entered is 1 minute, the simulation will run for 1 minute of simulation time, and then stop. Single step mode is the same as setting the play interval to the simulation time step.
>
> **[N] x Real Time:** constrain the simulation to run no faster than N times faster than real time. If the model is very large, or the computer is not powerful enough, or the graphics are in Detailed mode, the simulation may not be able to keep up with the specification.
>
> **Simulation Time Step:** the interval used to advance the simulation time, variable from 0.01 to 0.5 sec in increments of 0.01
> **Warning:** Changing this parameter will change the results
>
> **Simulation Date:** a date for the simulation. This may be used, for example, when applying economic assessment data, or for applying traffic control that varies by day or season.

Play/Pause: Allows you to start the simulation when it is paused, and to pause it while it is running. In single-step mode, pressing this button causes the simulation to advance by one step or interval. Pressing the space bar key has the same effect as pressing this button.

Layer[0 – 20]: The currently visible layer, as defined in the Layer tab.

Draw [Flat / Outline / Solid / Detailed]: The detail used to render 3-D shapes onto the graphics view. Memory and processing overhead increases as this setting is increased. Depending on your hardware, higher values may cause noticeable falls in simulation speed.

- **Flat:** Draw 3-D shapes as flat 2-D shadows (where possible)
- **Outline:** Draw line edges of 3-D shapes only
- **Solid:** Draw solid 3-D shapes, using uniform colours
- **Detailed:** Draw solid 3-D shapes, applying images to surfaces. This is also known as texture-mapping.

This parameter does not affect the results of the simulation, but may speed up the simulation process, as the computer spends more time calculating, and less time drawing.

Speed [Standard / Fast / Faster / Fastest]: A “master” control that applies a pre-set value to each of the Draw, Refresh and Scope controls. Use this to quickly speed up or slow down the simulation. Changing this control will have no effect on the results produced by the simulation.

◄ A small arrow at the right hand end of the tool bar displays some “Advanced Controls”.

Advanced Controls

Refresh [Default 1]: Changes the interval at which the display is refreshed. For example, if this is set to 15, the display will be refreshed after every 15 seconds. This parameter does not affect the results of the simulation, but may speed up the simulation process, as the computer spends more time calculating, and less time drawing.

Scope [Default 10] This control has a range of 0 to 10. Increasing this value draws objects at further distance from the cursor, using a logarithmic scale. A value of 0 means only draw objects which are very close to the cursor. A value of 10 means draw everything. This parameter does not affect the results of the simulation, but may speed up the simulation process, as the computer spends more time calculating, and less time drawing.

Exp: Exploded view magnification. This parameter expands all objects in the vertical direction. This makes it easier to “see inside” some constructions, such as multi-level parking structures. If your network is flat, with all heights set to zero, this control will have very little visible effect.

Depth: A horizontal slice of this depth (in metres) through the network, centred on the Base height. In combination with the **Base** control, this can be used to look only at one floor of a multi-floor building. Set this value to zero to see all depths.

Base: The height of the cursor above “ground-level” (the zero-height horizontal plane).

Actions

Reference

Adjust

The Adjust action can be applied to many types of objects. First, select the objects you want to adjust, then press the button to implement this action, which will raise the **Adjust Parameters** window.

You can select as many objects as you want at the same time.

Tip: **Triple-Click = Adjust.** While double-clicking on an object will select it, a quick way to raise the Adjust Parameters window is to triple-click on an object.

The Adjust Parameters window contains multiple tabs, one tab for each type of object. A tab will be visible only if one or more objects of that type are currently selected. In each tab, there will be one row in the table for each object of that type that is currently selected.

If you select multiple objects, open this window, then select the "Apply to All" option, then any new value you type for any object will be copied to all other objects.

Any new values you type will not be applied to the objects until you press **Apply** or **OK**. The **Apply** button is equivalent to **OK**, except that **Apply** does not close the window. **Apply** validates your input, so you can use this button to check that the values you have typed are valid and will be accepted before closing the window.

Neither the **Apply** button nor the **OK** button causes the network to be saved. To save the network, you need to press the **Save** button on the toolbar or on the File menu.

Adjust – Areas

Parameter	Effect
Name	A unique numeric name for the Area, assigned by the system. This is not editable.
Description	A descriptive name for the Area, which does not need to be unique.
Freight	Select this option to mark this Area as an origin or destination for freight. This will cause the Area to be drawn in a different layer, which may be a different colour. If this option is not selected, then this Area is an origin or destination for people.
Base	The height of the base of the Area box above the nominal “sea level”. An Area is a vertical extrusion of a polygon, with a finite height. This means that one Area can be directly above another, which is useful for example, for defining one Area for each level of a multi-storey building.
Height	The height of the Area above its base. By default, this is 10m. The height must be large enough so that the Area contains all Walkways which are to be used as origins or destinations.
Volume	The hourly volume of people or freight produced by this Area in the absence of any demands. This is used as a fall-back in simple models, to get some people moving on the network. In most cases, it is recommended that you define demands explicitly for each Area, rather than relying on this value to generate trips.
Cordon	Set to true if this is marked as a cordon Area, for reporting purposes. A cordon Area is normally one on the edge of the network, which does not represent a geographic Area. Commonly, this flag will be set for any Area that is linked to a Transition Zone, but this is not a requirement.

Parameter	Effect
Off Distance/ Off Time/ Off Price/ Off Variation	The off-network distance, time, price and variation associated with this Area. It is possible to model people arriving from, or going to, a set of off-network Areas all connected to the same transition zone or stand. For example, if you were modelling an airport that had four distinct catchments, you might want to create 4 separate cordon Areas, and define a (mean) distance, time and price for each of them. Then each of these 4 Areas would be connected to a transition Zone at the edge of the model Area, where vehicle trips would start. These fixed off-network values are then added to the variable on-network costs when calculating the cheapest route to the destination for each behaviour type. The variation value specifies a standard deviation, and can be used to create a distribution around the mean for each measurement.

Adjust – Connection

Parameter	Effect
Name	A name for the Connection, which can be modified.
Restriction	A behaviour-based restriction for the Connection (for example: staff only). The restriction does not need to apply to either the entry or exit Walkway.
Capacity	A limit to the number of people using the Connection at any time, in either direction. Additional people will be held at the end of the previous Walkway, until space is available.
Straight Edge	Make the edge of the Connection a straight line. This is not just a visualisation effect, it defines the boundary of the Connection, so if it is walled then it will affect the space available for movement. A Connection can be walled by walling the Walkway at either end.
Texture	(Visualisation only, no effect on model results) Select a Texture Map image for the Connection between Walkways. This will be used in Detailed display mode only. For example, a Connection could be displayed as bricks or gravel using pre-defined textures. Other texture maps, such as paving or cobbles can be added by saving the texture map in the azb/tex/ directory in the installation folder.

Adjust – Crossings

Parameter	Effect
Name	A name for the Crossing, which must be unique.
Intersection	(Signalised Intersection Crossings Only) The Intersection with which this Crossing is associated.
Group	(Signalised Intersection Crossings Only) The signal group that activates this Crossing.
Timing*	One or more Crossing time definitions for this Crossing. A Crossing time object defines walk times and waiting (don't-walk) times.
Sensor ID	(Signalised Intersection Crossings Only) A "special" field, to encode information required for an external system. For example, in some traffic control systems, some Crossing push buttons are treated as detectors, and the sensor ID is used when the button is pressed.
Route Class	The Route Class is used to assign cost parameters to a group of similar Crossings.
Route Length	The length of a Crossing for routing purposes.
Route Speed	The “free flow” speed of a Crossing for routing purposes.
Route Price	A monetary cost that applies to this Crossing, used in the generalised cost equation.
Cost Factor	A multiplication factor used to increase or decrease the cost for this Crossing as calculated using the generalised cost equation.

Parameter	Effect
Restriction	A restriction on this Crossing, controlling the people (by behaviour) that can use it. Restrictions are specified by mob, so you may need to define a mob containing the behaviours you want to allow or bar from this Crossing.
Keep	The side of the Crossing to which people will keep when walking. Can be set to Left, Right or None. If set to None, opposing flows will be lower, especially for high volumes.
Walled L	Creates a “wall” on the left side of the Crossing, looking in direction CD. If this is set, people on the Crossing will not go outside the left boundary.
Walled R	As above, but on the right side, in direction CD.
Measure	Adds this Crossing to the list that is recorded in the **Traffic Reports / Walkways** window.
Following	Select the following behaviour for this Crossing. A blank value means use the default “automatic following” where people follow a leader in the face of an oncoming crowd. The “On” setting is good for walkways which form queues. Use Display / Diagnostics / People / Leader to display follow-leader lines.
Width	The width of the Crossing. This can be used to reset the width to be uniform along the entire length of the Crossing. However, if the width handles at either end of the Crossing are moved, this value is not used.

Parameter	Effect
Space Factor	A multiplicative factor for the size of the space around each person. Setting this to less than 1.0 causes people to pack more densely than normal.
Speed Factor	A multiplicative factor for the speed of walking in both directions. Use this to make people walk faster or slower than normal.
Texture	(Visualisation only, no effect on model results) Select a Texture Map image for the Crossing. This will be used in Detailed display mode only. For example, a Connection could be displayed as bricks or gravel using pre-defined textures. Other texture maps, such as paving or cobbles can be added by saving the texture map in the azb/tex/ directory in the installation folder.

Adjust – Dwells

Parameter	Effect
Name	A unique name, derived from the stand and service. A dwell object contains information for a particular service at a given stand.
Description	A text description, which does not need to be unique.
Dwell Minimum	The minimum dwell time for the relevant service at this stand. Each vehicle arriving at the stand will dwell for at least this time, regardless of the number of passengers boarding or alighting.
Dwell Maximum	The maximum dwell time for the relevant service at this stand. Each vehicle arriving at the stand will not dwell for more than this time, regardless of the number of passengers boarding or alighting. That is, if a passenger has not yet boarded, it will miss the vehicle, and must wait for the next one. The board time can be extended by passengers arriving just as the vehicle is ready to depart, and each passenger will add another few seconds of delay. The maximum board time is useful on busy stations to ensure that the train leaves promptly.

Adjust – Fixtures

Parameter	Effect
Name	A unique name for the Fixture.
E/W **N/S** **Elevation**	The location of one corner of the bounding box that defines the position and size of the Fixture.
Length **Width** **Height**	The dimensions of the Fixture.
Gradient	A tilt angle, relative to horizontal for the Fixture.
Rotation	A rotation, relative to North, in the horizontal plane.
Flip	Allows rotation around either X axis or Y axis.
Normals	(Visualisation – Detailed Mode Only) Computes normals for each face, or each vertex, depending on the option selected. This results in better shading of the object, but at the expense of more computation. If you switch this on, your model will run more slowly. Normal computation is off by default.
Hidden	Set to **true** if the Fixture should not be shown in the model. This is useful for creating an invisible Obstacle.
Fixed	Set to **true** if the Fixture cannot be selected or moved. This is useful for background Images: set this flag once you have the Image in the correct position as a backdrop, or overlay. To reselect the Fixture once you have set it as fixed, either display **Handles – Other** and zoom in to the corner, then select the handle, or hold down shift and drag out a box around one of the corners. Holding down shift during box selection selects all Objects even if they are not currently visible.

Parameter	Effect
Reduced	(Visualisation Only) **True** if the object should be drawn faster by reducing the polygon count when the Scope control is less than 10. This can speed up the display for very complex shapes.
Obstacle	Set to **true** if pedestrian agents should avoid this Fixture. Setting this to true increases the amount of computation required for the avoidance algorithm.

Adjust – Labels

Parameter	Effect
Name	A name for the Label, which can be modified.
Text	The visible text on the Label. The text will be displayed on a single line and cannot contain formatting characters.
GoTo	(Optional) A hyperlink target. This may be a file on the local system, a URL for an internet resource, or an Object within the model. The field should have two parts, separated by the '~' character. The first part determines the type of the target, the second part defines the name of the target. Examples: • File~C:/Azalient/Commuter/net/Results2.xls • Web~traffic.gov/standards1.pdf • Intersection~345 • Controller~345 • Node~1001 If the target is an Object in the model, then the viewable layer of the Label will be changed according to the target. Thus if the Label's target is a Node, and Nodes are not currently visible, then the Label will not be visible.
Colour	The colour of the text.
Size	The height of the text, in metres.
Rotation	The rotation in degrees, relative to North.
Magnify	Select this to magnify the Label as the cursor is moved near to its location.

Parameter	Effect
Texture	(Visualisation only, no effect on model results) Select a Texture Map Image for the Label. This will be used in Detailed display mode only. The texture map will be displayed on a rectangle defined by the area required for the text on the Label. If no text is required, space characters can be used to create a rectangle of the correct size. This feature can be used to add texture map Labels, rather than ones which use a vector font.

Adjust – Lanes

Parameter	Effect
Name	The name is determined by the link to which this lane belongs
Index	An external index value, normally 1 signifies the kerb lane, increasing towards the median. This value is used for lane choice matching rules. The external index value does not need to be unique on a link: if a lane is bisected, both new lanes will have the same external index value. The external index value is not the same as the internal index value, which is unique for all lanes on a link. The internal index value is also used in the name of the lane.
Speed	The maximum speed for the lane, in the preferred units (km/h or mph). Normally, you should set this to be equal to the posted speed limit. There are several display formats: **[50]** The speed is the same as the route class speed for the link. Like this, if you change the route class then the lane speed will change automatically **45** The speed is different from the route class. Changing the route class will not affect the maximum speed on the lane. **45 (38)** The maximum speed has been reduced because of the curvature of the link. This reduction is for all vehicle types, using a standard tyre friction co-efficient, unlike the turning speed on streams which is calculated per-type.

Parameter	Effect
Width	The width of the lane along its entire length. If you move the width indicator at either end, so that a lane has a different width at one end to the other, this value is not used. However, setting this value (then press OK) resets the width indicators at both ends so that the lane has a uniform width along its entire length.
Restriction	A restriction for the lane, or blank, if there is no restriction. Restrictions are based on mobs (groups of behaviours) so you may first need to create a mob containing behaviours you want to allow, or bar. Selecting '*' from the drop down menu allows you to set multiple time-based restrictions. There can only be a single restriction at any one time on a lane, but there may be different restrictions at different times of day. For example, a lane may be used for a bus lane in peak hours, and for parking at other times.
Speed Control	A speed control definition applied to the lane. These are defined in the second tab of the Restrictions window. These can be applied to all vehicles, or by selecting a Mob, only to a subset. For example, a speed control on heavy vehicles. Speed controls can also be defined to apply at certain times of day only. Like restrictions, there may be multiple speed controls on a lane, but they should not overlap in time.

Parameter	Effect
Closed	The ultimate restriction: closed to all traffic.
Hold	A distance from the end of the lane at which no more lane changing will take place. That is, if the value is 10m, and the front of a vehicle is less than 10m from the stop line, then that vehicle will not attempt to change lane. By default, the lane hold distance is 10% of the length of a link but not more than 20m.
Box Speed	If this is non-zero, a vehicle approaching an intersection will not enter the intersection if the vehicle ahead of it is travelling at the box speed or less. This attribute takes its name from the term “box junction” which is used in some places to describe an intersection where approaching traffic must not enter the intersection until it is clear there is room for them to exit the intersection without stopping.
Avoidance	If this is on, then all vehicles on this lane will scan the area ahead for other vehicles or pedestrians. While this leads to more detailed avoidance behaviour, it requires more computing power, so you should balance the use of this parameter with your requirement for speed of simulation. The global parameter “Avoidance” in the Parameters/ Calibration window can also be used to set the default state for all links.
Zip	Allow zip-style merging on this lane. This is where each vehicle increases its normal following gap to allow one other vehicle to merge into its lane.

Parameter	Effect
Discrete	(Visualisation only, no effect on model results) Set for lanes which are not connected to an intersection, and should not be included when drawing the intersection outline. This might be used, for example, for an offset cycle lane.
Margin	An invisible margin to each side of the lane used to control lane changing. Lane changing is possible if the **Lane Change Barred** action is not used, and the distance between the lanes is less than this margin. If lane changing is possible, the edge of the lane will be drawn as a broken line.
Headway Factor	A multiplicative factor applied to the headway time used for car-following and lane changing on this lane. The default value is 1.0. Example uses are a value of 0.8 to apply a shorter headway near a merge, where vehicles tend to have a higher density, or a value of 1.2 in a tunnel to apply a larger headway and reduce the maximum density.
Reaction Factor	A multiplicative factor applied to the reaction time used for car-following & lane changing. This has different effects for each car following model . A value of less than 1 will generally lead to higher saturation flows at signalized intersections.
Lane Friction	A value between 0.0 and 1.0 used to lower the effective speed limit in the presence of slow-moving traffic in an adjacent lane. For example, if the speed limit on this lane is 70, the average speed on the adjacent lane is 30, and the lane friction value is 0.75, then the applied speed limit on this lane will be: 70 – 0.75 x (70 – 30) = 40

Parameter	Effect
Following	Select a car-following algorithm to be used on this lane. For example, one algorithm might be applied to highways, and another to urban streets
Texture	(Visualisation only, no effect on model results) Select a Texture Map image for the lane.

Adjust – Loops

Parameter	Effect
Name	A unique name for the Loop. By convention, Loops used for traffic control at an Intersection, use a name of the format 1234_5, where 1234 is the Intersection ID and 5 is the Loop ID.
Controller	The Intersection controller to which this Loop is attached. The controller has the same name as the Intersection.
Fleet	If this field is blank, then the Loop will count all vehicles. If a fleet is defined (use Parameters > Vehicles > Fleets), then the Loop will count only vehicles that are in that fleet.
Distance	The distance of the start of the Loop from the start of the Lane.
Size	The length of the Loop.
Width	The width of the Loop.
Offset	The sideways offset of the Loop, with a negative value meaning to the left of the centreline, looking in the direction of travel.
Reset	A reset interval for the Loop, in format HH:MM:SS. This resets all counts, average flows etc., to zero.
Measure	Select this option to add this Loop to the Traffic Reports window.
Log	Select this option to save data for this Loop in the results file for the Traffic Reports.

Adjust – Nodes

Parameter	Effect
Name	A unique name for a node. Can be modified by the user, but it might be more useful to rename an intersection rather than the nodes it contains
Form	• **Priority**: a junction where each movement is controlled by a fixed sign, for example Give Way or Stop • **Signalised**: a junction controlled by traffic signals • **Roundabout**: a special type of priority junction node, where variation to the vehicle behaviour apply. If the intersection is selected, a roundabout node will be outlined with a circle, nodes of other Form values will be highlighted with a square • **Ramp**: a node where traffic merges onto a free-way midway along the main line link. This type of ramp can be created only by importing a model from Paramics. • **Merge**: this form triggers a change to the keep-side behaviour for approaching vehicles, to improve the merging behaviour • **Mini-Roundabout:** a node in a single-node intersection where there may be only one vehicle on the intersection at any time, and priority is given to the approach to the opposite of the driving side. • **All-Equal:** A special type of priority junction in which all approaches have the same priority. Only one vehicle will be allowed on the intersection at a time. There should be only one node in the intersection.

Parameter	Effect
	• **Limited:** Only one vehicle from a fleet can use the node at any one time. This can be used where there is a width restriction, for example, only one truck at a time.
Fleet	A fleet of vehicles for use with Form = Limited

Adjust – Parking

(Tab is enabled if Lane is enabled for parking)

Parameter	Effect
Name	The name of the Lane, which is not editable as it is derived from the Link.
Park Angle	The parking bay angle, expressed relative to the direction of travel, and any direction specification: • Face-In: park facing the kerb; • Face-Out: park facing the street; • Face-Either: driver can choose.
Bay Size	The size of the bay in the direction of travel. That is, the length of each bay, for parallel parking, or the width of the bay for 45, 60 or 90-degree parking. For angled parking, the Lane width parameter can be set to control the length of the parking bays.
Park Time	The time, in seconds, required to manoeuvre into or out of a parking bay. This controls the amount of time that other traffic is delayed while a vehicle is parking.
Transition (Time)	The time, in seconds, required for a driver to get into or out of a vehicle. This time does not affect traffic in the moving Lane, but does add to the total trip time for the person parking.
Dwell (Time)	(Applies to drop-off zones only). The time that a vehicle stays in a bay having dropped off a passenger. You can enter either a number of seconds, and all vehicles will dwell for that time, or you can enter the name of a Time Distribution.
Popularity	An integer value defining the popularity of all bays in this Lane, where a higher number indicates a more popular Lane. There is no upper limit, and the lower limit is zero. The choice of bay can be further refined using the Best Bay parameter.

Parameter	Effect
Best Bay	An integer parameter indicating the bay that is the most popular within the Lane. A value of zero (the default) indicates that all bays are equally popular. Popularity is assumed to decline at a linear rate away from the best bay. So if bay 7 is defined as the most popular, bays 6 and 8 will be the equal second most popular bay. Entering a value higher than the highest number bay indicates that the highest numbered bay is the most popular. It can be useful to enter 99 or 999 in the best bay field to make sure the highest numbered bay is the most popular, even after a Lane has been stretched.
Cluster	A name given to the cluster of parking Lanes to which this parking Lane belongs. The cluster name is used by the Parking Guidance Tool.
Block	If this is set, a vehicle approaching this Lane with an intention to park will wait until a space becomes free. If this option is not set, and alternatives are available, then the vehicle will continue driving and try to find a space elsewhere.

Adjust – Rail Links

This section describes parameters that apply to the routing Link for selected Rail. Both Roads and Railways have routing Links, referred to as Road-Links and Rail-Links. Layer "Links" displays the routing Links for both Roads and Railways.

Parameter	Effect
Name	A unique name for the Rail-Link. This will be changed if the Description is changed.
Description	A description for the Rail-Link that does not need to be unique. A new description will change the name to match. For example, to rename all the Rail-Links making up Northern Line, select them all, **Adjust**, select **Apply To All,** type **Northern** into the description field, then **Apply**. Links will be renamed as “Northern_1”, “Northern_2”, etc.
Route Class	A Route Class is used to assign cost parameters to a group of similar Links. For example, you might define route classes as “Suburban Line”, “Freight Line”, “Inter-City”.
Route Length	The Length of a Rail-Link for routing purposes. If the value in this field is enclosed in square brackets, then it is the default value, calculated from the length of the Tracks. The default will change if you move the end points of the Tracks. The Route Length is **not** calculated from the Node positions. If you set this value to be different from the default it will be shown without any brackets, and will not change if the lane geometry is changed. You might change this value from the default if you want to bias the route away or towards this Rail-Link. To reset the value to the default, simply blank the field and press Apply.

Parameter	Effect
Route Speed	The expected or perceived "free flow" speed of a Rail-Link for routing purposes. If this Rail-Link is using the default value taken from the Route Class, then the value will be enclosed in square brackets like this: [50]. If the value has been changed from the Route Class value, then there will be no square brackets. Whether or not it uses the Route Class default, the Route Speed may be lower or higher than the Lane Speed. The Lane Speed is the physical speed limit on the road, the Route Speed is the expected, or perceived speed.
Route Price	A monetary cost or charge, that applies to this Rail-Link, used in the generalised cost equation. If the value displayed is enclosed by square brackets, then it is the default value, derived from: (Route Price per Unit Distance) x (Route Length) The Route Price per Unit Distance is defined in the Route Class. If you define a different value for Route Price and apply it, then the value shown will not be enclosed in square brackets.
Cost Factor	A multiplication factor used to increase or decrease the cost for this Rail-Link as calculated using the generalised cost equation.
Release Rate	This field is editable only for Rail-Links that are inside zones. If it is not set (blank), then trips will be released within a Zone using the Rail-Link that has the lowest cost to destination. If this field is set to a non-zero value, then it controls the proportion of traffic generated on this Rail-Link.

Adjust – Roads

This section describes parameters that apply to the routing Link for selected Roads. Both Roads and Railways have routing Links, referred to as Road-Links and Rail-Links. The Layer named "Links" displays the routing Links for both Roads and Railways.

Parameter	Effect
Name	A unique name for the Road-Link. This will be changed if the Description is changed.
Description	A description for the Road-Link that does not need to be unique. A new description will also change the Name to match. For example, to rename all the Roads making up Smith St, select them all, **Adjust**, select **Apply To All,** type **Smith St** into the description field, then **Apply**. Roads will be renamed as “Smith St_1”, “Smith St_2”, etc.
Route Class	A Route Class is used to assign cost parameters to a group of similar Road-Links. For example, you might define route classes as “Highway”, “Major Road”, “Minor Road”, “Local”.
Route Length	The Length of a Road-Link for routing purposes. If the value in this field is enclosed in square brackets, then it is the default value, calculated from the length of the Lanes. The default will change if you move the end points of the Lanes. The Route Length is **not** calculated from the Node positions. If you set this value to be different from the default it will be shown without brackets, and will not change if the Lane geometry is changed. You might change this value from the default if you want to bias the route away or towards this Road-Link. To reset the value to the default, simply blank the field and press Apply.

Parameter	Effect
Route Speed	The expected or perceived "free flow" speed of a Road-Link for routing purposes. If this Road is using the default value taken from the Route Class, then the value will be enclosed in square brackets like this: [50]. If the value has been changed from the Route Class value, then there will be no square brackets. Whether or not it uses the Route Class default, the Route Speed may be lower or higher than the Lane Speed. The Lane Speed is the physical speed limit on the Road, the Route Speed is the expected, or perceived speed.
Route Price	A monetary cost, commonly a toll, that applies to this Road-Link. This will be used in the generalised cost equation. If the value is displayed enclosed by square brackets, then it is the default value, derived from: (Route Price per Unit Distance) x (Route Length) The Route Price per Unit Distance is defined in the Route Class. If you define a different value for Route Price and apply it, then the value shown will not be enclosed in square brackets.
Cost Factor	A multiplication factor used to increase or decrease the cost for this Road-Link as calculated using the generalised cost equation.
Lane Choice	Two distances, measured from the start of the Road, in the format Start;End, specifying where decisions about Lane choice should be taken. For example, if the value is 200;400 then a Lane choice decision will first be made at a point between 200 and 400 (metres or feet) from the start of the Road-Link. The distance for an individual vehicle will be defined by its DNA, using a flat distribution between the Start and End values.

Parameter	Effect
Look Distance	A distance, specified in short distance units (metres or feet) beyond the end of this Road-Link over which a driver will look to decide which exit is required at the next intersection.
Look Links	The maximum number of Road-Links ahead that will be used to determine which range of Lanes are suitable. For example, if a vehicle is turning ahead, and there are several intersections in close proximity, it may be better to get in Lane several Road-Links upstream.
U-Turn	Select this option if it possible for a vehicle to make a U-Turn at the end of this Road-Link, or on this Road-Link, if it is used for Parking.
Release Rate	This field is editable only for Road-Links that are inside zones. If it is not set (blank), then trips will be released within a zone using the Road-Link that has the lowest cost to destination. If this field is set to a non-zero value, then it controls the proportion of traffic generated on this Road-Link.
Measure	Select this option to mark this Road for measurement in the Traffic Reports window.

Adjust – Signals

Parameter	Effect
Left2 / Left / Ahead / Right / Right2	The signalling equipment in place for the given turn movement on the selected approach signals. For example: • green + yellow + red roundels • green arrow • green + yellow + red arrows Left2 and Right2 are used only if the approach has 4 or 5 exits.

Adjust – Signs

Parameter	Effect
Name	A unique name for the Sign. This is a number by default, but can contain letters.
Message	The message displayed on the Sign.
Offset	The sideways offset of the Sign from the "kerb-side" edge of the Lane.
Size	The width of the Sign. A value of zero or less sets the Sign to the default width.
Height	The height of the top of the Sign above the Base.
Base	The height of the bottom of the message area above the ground; the length of the supporting posts.
Text Size	The height of the letters on the Sign.
Colour	The colour of background of the message area.

Adjust – Stands

Parameter	Effect
Name	A unique name for the Stand. This is a number by default, but can contain letters.
Description	A descriptive name that does not need to be unique.
Size	The length of the Platform or Standing area for passengers. This is used in combination with the number of doors to fix the position of the points at which groups of passengers gather while waiting for the next public transport vehicle.
Offset	The sideways offset of the Stand from the "kerb -side" edge of the Lane.
(Boarding) Doors	The number of doors expected on the next transport vehicle to arrive. This controls the entry rate for passengers boarding a transport vehicle. The higher the number of doors, the faster the speed of boarding. The number of doors on a Stand is used in conjunction with the size of the Stand to fix the position of the points at which groups of passengers gather while waiting for the next transport vehicle. The type of the vehicle can change by service, and even within the service, and each vehicle may have a different number of doors. For example, an arriving train may have 4 cars or 6 cars, each with 2 doors each. There is no requirement that the number of doors defined for a vehicle matches the number of doors on the Stand.
Slots	(Buses Only) This defines the number of stopping positions at a Stand. For example, if the length (Size) of the Stand is 48m, and there are 4 slots, this will create 4 x 12m stopping positions within the Stand area, so that up to 4 buses will be able to stop at the same time, as long as each of the buses is less than 12m in length (including minimum separation gap).

Parameter	Effect
Move In	(Buses Only) This is the distance before the Stand at which a bus will move in to the kerb side Lane. If there are several Stands close to each other, set this to a small value (for example, 10 or 15 metres) to keep buses in the running Lane until the last minute. Without this value, a bus going to a downstream Stand may get stuck behind a bus at an upstream Stand.

Adjust – Streams

Parameter	Effect
Name	A name for a Stream derived from the entry and exit Lanes, which cannot be edited.
Turn Speed	The maximum speed for this Stream, as determined by its curvature. It can be further reduced by the user, but it cannot be increased, other than by "straightening out" the Stream path.
Advance	A (positive) distance used to control where filtering traffic waits. This controls the throughput of a filtered turn under congested conditions. Setting this to a non-zero value will show a wide bar on the display across the direction of travel. This distance controls the number of vehicles that will make it through a filter turn in a congested Intersection. The number of vehicles per green that will get through will be directly proportional to this distance divided by the average vehicle length.
Clearance	A (positive) distance used to control where a vehicle is judged to have cleared the Intersection and no longer conflicts with turning traffic. Setting this value will show a triangle on the display to indicate the clearance point for this Stream. A higher clearance value will, in general, increase the throughput on Streams that conflict with this Stream.
Merge	A time, in seconds, used by the gap acceptance algorithm when a vehicle is merging with another Stream. The value is zero by default and can be negative or positive. A positive value increases the perceived time required to merge safely in the presence of another vehicle of higher priority. Conversely, a negative value decreases the perceived time required. Thus a positive value will reduce throughput under congested conditions, and a negative value will increase throughput. Typically, values will be in the range -4 to +4.

Parameter	Effect
Cross	Similar to the Merge time value, but used where a vehicle is crossing a conflicting Stream, rather than merging. As with the merge time value, a positive value will reduce throughput under congested conditions, and a negative value will increase throughput. Typically, values will be in the range -4 to +4.
Keep	Use this to control the lateral position of vehicles – that is, which side of a Stream a vehicle uses. Vehicles will use the corresponding side of the Lane leading into and out of the Stream. • **(None) -** No preference for lateral position • **Left -** Stay to left side of Lane • **Centre -** Keep to centre of Lane • **Right -** Keep to right side of Lane
Texture	(Visualisation only, no effect on model results) Select a Texture Map Image for the Stream. This will be used in Detailed display mode only.

Adjust – Symbols

A Symbol is a shape that can be flat or can be a box with vertical sides (in mathematical language, a vertically extruded polygon). Its base height can be moved, like other objects, and if it is a box, its height can be changed. It is added to a model for illustration or annotation purposes. It has no effect on the results of the model. For example, it may be used to illustrate the position of a building or a water feature. For more complex buildings, consider using a Structure.

Parameter	Effect
Name	A name for the Symbol, which can be modified.
Colour	The colour of the Symbol.
Line Width	The line width of the Symbol edges. Set to zero for no edge highlighting. Edge highlighting may not be required if shading is set to on.
Filled	Turn this off to have the Symbol displayed as a wire frame for all display modes.
Base	The elevation of the base face.
Height	The height of the top face above the base.
Texture	(Visualisation only, no effect on model results) Select a Texture Map Image for the Symbol. This will be used in Detailed display mode only. The texture map will be displayed on the vertical sides of the Symbol. This can be used to display walls, fences or other vertical partitions.

Adjust – Tracks

A Track is one "Lane" of a Rail-Link.

Parameter	Effect
Name	The name is determined by the Rail-Link to which this Track belongs.
Index	An index value, normally 1 signifies the Track nearest to the Platform.
Speed	The maximum speed for the Track, in the preferred units (km/h or mph). Normally, you should set this to be equal to the posted speed limit. There are several display formats: **[50]** The speed is the same as the route class speed for the Link. Like this, if you change the route class then the Track speed will change automatically. **45** The speed is different from the route class. Changing the route class will not affect the maximum speed on the Track. **45 (30)** The maximum speed has been reduced because of the curvature. This reduction is for all vehicle types, using a standard friction co-efficient, unlike the turning speed on streams which is calculated per-type. **[50] (30)** The default would be used, but the speed has been reduced because of the curvature.
Width	The width of the Track in the units appropriate to your settings (metres or feet).

Parameter	Effect
Restriction	A restriction for the Track, or blank, if there is no restriction. Restrictions are based on mobs (groups of behaviours) so you may first need to create a mob containing behaviours (types) you want to allow, or bar. Selecting '*' from the drop down menu allows you to set multiple time-based restrictions. There can only be a single restriction at any one time on a Lane, but there may be different restrictions at different times of day.
Speed Control	A speed control definition applied to the Lane. These are defined in the second tab of the restrictions window. These can be applied to all vehicles, or by selecting a Mob, only to a subset. Speed controls can also be defined to apply at certain times of day only. Like restrictions, there may be multiple speed controls on a Lane, but they should not overlap in time.
Avoidance	If this is on, then all trains on this Track will scan the area ahead for other vehicles or pedestrians. While this leads to more detailed avoidance behaviour, it requires more computing power, so you should balance the use of this parameter with your requirement for speed of simulation. The global parameter “Avoidance” in the Parameters/ Calibration window can also be used to set the default state for all Links.
Texture	(Visualisation only, no effect on model results). Select a Texture Map Image for the Lane. This will be used in Detailed display mode only. There are some pre-defined textures, more can be added by saving the texture map in directory azb/tex/ in the installation.

Adjust – Turns

Tip: while it is possible to edit parameters for Turn in this window, it is easier to use Intersection Editor to edit these parameters.

Tip: To select a Turn, you should first switch to the Layer tab and display the **Signalised Turns** or **Priority Turns** Layers . Then select a Turn by clicking on one of the bold coloured arrows. Selection is easier if the Lane Streams Layer is not displayed.

Parameter	Effect
Name	A name for the Turn, constructed by concatenating the names of the entry and exit links for this Turn. This name will be unique, but visual inspection will be required to distinguish between Turns to opposite directions of the same street. For example, at the Intersection of George St and Park St, each George St approach will have two Turns with names that contain George and Park. Often these will have the form George1>Park1, George1>Park2, etc.
Fixed Signal	If the Turn is not signalised, this field can be set to indicate the fixed signal that applies, for example, free flow, stop, yield, etc. The direct and filter group fields must be blank for this field to be used.
Direct Group	At a signalised Intersection, this is the signal group that controls this Turn directly (this is sometimes also called the “primary group”). The drop down selector in this field contains all the vehicle groups defined for an Intersection. If there are no groups to choose, you should first define the signal groups in the groups panel of the Intersection Editor (select Intersection, then Adjust).

Parameter	Effect
Filter Group	At a signalised Intersection, this is the signal group that controls this Turn indirectly (this is sometimes also called the “secondary group”). That is, if the ahead or through movement is associated with group G, and this Turn filters through oncoming traffic when G is green, then G is the filter group for this Turn. This Turn may also have a direct group (indicating a green arrow for this Turn when oncoming traffic has been stopped) but note that it is not necessary for a Turn to have a direct group if it has a filter group, and you do not need to define a “dummy” group for the direct “primary”.
Filter Signal	If this Turn has a filter group, then this is normally set to “green yield”, but it is also possible to set it to another signal state in special circumstances. For example, if traffic on a Turn can proceed without conflict in either of two groups, then the filter signal would be “green”. If this is a Turn to the kerb side (LTOR or RTOR) then “green stop” might be more appropriate.
Restriction	A mob-based restriction applied to the Turn (for example: Buses can Turn left). The restriction does not need to apply to either the entry or exit lane.
Preference	(Inside Off-Street Parking Zone Only). Set this field to 1 to indicate the preferred Turn for a vehicle that is looking for a parking bay. To display the preferred path around a parking zone, switch on the **Priority Turns** Feature in the Layers tab, and select the **Route** option in the Reporting/ Parking window.

[Note on “**Turn**” vs. “**Movement**”: In almost all cases, a Turn is the same as a movement, the exception being when the Intersection is signalised and you define an extra Turn for a bus phase or similar, so that there are two Turns for a single movement. Some might say that Turn is not a good term, particularly for the straight-ahead movement, as the vehicles do not “Turn”; another view is that straight-ahead is a valid Turn, just as zero is a valid number.]

Adjust – Walkways

Parameter	Effect
Name	A name for the Walkway, which must be unique.
Route Class	The Route Class is used to assign cost parameters to a group of similar Walkways.
Route Length	The length of a Walkway for routing purposes.
Route Speed	The “free flow” speed of a Walkway for routing purposes.
Route Price	A monetary cost that applies to this Walkway. This will be used in the generalised cost equation.
Cost Factor	A multiplication factor used to increase or decrease the cost for this Walkway, as calculated using the generalised cost equation.
Restriction	A restriction on this Walkway, controlling the people (by behaviour) that can use it. Restrictions are specified by mob, so you may first need to define a mob containing the behaviours you want to include or exclude.
Width	The width of the Walkway. This can be used to reset the width to be uniform along the entire length of the Walkway. However, if the width handles at either end of the Walkway are moved, this value is not used.
Step Height	If set to a non-zero value, the Walkway will be displayed as steps if it has a gradient. The number of steps is calculated from the total height difference between the ends of the Walkway divided by the step height. In combination with Speed CD, this field can be used to create an escalator.
Wall Height	The height of any walls. If the Walkway has two walls they must be the same height.

Parameter	Effect
Keep	The side of the Walkway to which people will keep when walking. Can be set to Left, Right or None. If set to None, opposing flows will be lower, especially for high volumes.
Walled L	Creates a “wall” on the left of the Walkway, looking in direction CD. If this is set, people on the Walkway will not go outside the left boundary.
Walled R	As above, but on the right side, in direction CD.
Shared	If set, signifies this Walkway is shared with vehicles, normally cycles, but this can also be used for spaces shared between people and cars. The vehicle pathway should be created as a single-lane road, occupying the same space as the Walkway. Avoidance should be enabled on the Lane, so that vehicles can 'see' pedestrians.
Private	If set, only one person at a time can use the Walkway. This can be used to set up a checkpoint or barrier. If this Walkway is occupied by a person, other people will be held at the end of the previous Walkway, and will not be transferred until this Walkway is empty. This can be used with the Delay parameter to create a timed barrier or gate. See also Capacity.
Measure	Adds this Walkway to the list that is recorded in the **Traffic Reports / Walkways** window.
Wait Here	This can be used only on a walkway that is connected to a stand. If it is set, people will wait on the walkway until the next service arrives. This feature can be used to model a simple waiting room. Each person waiting will be allocated to a “seat”. The seats are aligned in rows on either side of the walkway. More rows will be created as the waiting room fills up. The waiting room can be given a capacity, like any walkway, but subsequent arrivals will queue outside the waiting room.

Parameter	Effect
Following	Select the following behaviour for this Walkway. A blank value means use the default "automatic following" where people follow a leader in the face of an oncoming crowd. The "On" setting is good for Walkways which form queues. Use Display / Diagnostics / People / Leader to display follow-leader lines.
Release Rate	Where there are multiple Walkways within an Area that can all be used to "release" new people at the start of a trip, this percentage value controls how many people start on each Walkway. For example, if there were two Walkways leaving an Area, you might set one to 80%, and the other to 20%. If there is only one Walkway in an area, this value is ignored.
Delay CD	A Walkway's ends are referred to as C and D. This field represents a delay, in seconds, at the "D" end of the Walkway. This delay will act like a check point, delaying each person for the time specified. The delay will also be used in routing calculations. Entering "inf" creates an infinite delay, resulting in a one-way Walkway in the other direction, from D to C. It is also possible to enter the name of a time distribution. Each person will select a time at random from that distribution. Delay CD can be used in conjunction with **Private** to create a control barrier, where only one person at a time approaches the barrier (for example passport control or security check). It is possible to use the action **Walkway > New Delay** to create a value in Delay CD.
Delay DC	As above, but for the opposite direction.

Parameter	Effect
Speed CD	A speed on the Walkway, used to model a moving Walkway or an escalator. Set this to a negative value for motion in direction DC.
Space Factor	A multiplicative factor for the size of the space around each person. Setting this to less than 1.0 causes people to pack more densely than normal.
Speed Factor	A multiplicative factor for the speed of walking in both directions. Use this to make people walk faster or slower than normal.
Capacity	A limit to the number of people using the Walkway at any time, in either direction. Additional people will be held at the end of the previous Walkway, until space is available.
Standing	(Escalators Only) The proportion of people, expressed as a percentage who will stand on an escalator; the remainder will walk. For example, a value of 40% means 4 out of 10 people will stand on the escalator, 6 out of 10 will walk.
Texture	(Visualisation only, no effect on model results). Select a Texture Map Image for the Walkway. This will be used in Detailed display mode only. For example, a Connection could be displayed as bricks or gravel using pre-defined textures. Other texture maps, such as paving or cobbles can be added by saving the texture map in the azb/tex/ directory in the installation folder.

Adjust – Zones

There is a Zone at each end of a private vehicle trip. A Zone can be one of three types: Vehicle-Only, Transition or Parking.

Non-Parking Zones

Parameter	Effect
Name	A unique name for the Zone, assigned by the system. This is not editable.
Description	A descriptive name for the Zone, which does not need to be unique.
Base	The height of the base of the Zone box above the nominal “sea level”. A Zone is a vertical extrusion of a polygon, with a finite height. This means that one Zone can be directly above another, which is useful for example, for defining one Zone for each level of a multi-storey parking area.
Height	The height of the “roof” of the Zone above its base, default = 10m.
Volume	The hourly volume of traffic produced by this Zone in the absence of any demands. This is used as a fall-back in simple models, to generate some traffic on the network. It is recommended that you define either directed (OD matrix) or undirected demands for each Zone, rather than relying on this value to generate trips.
Cordon	True if this is a cordon Zone, on the edge of the network, not representing a geographic Zone area. This is used in some reporting mechanisms, to report external-internal trips, etc.
Choice	Set if this Zone can be used for parking choice.

Parking Zones

Parameter	Effect
Parking Type	The type of parking for this Zone, selected from: • **On-street parking** allows vehicles to park on either side of the road, in numbered bays. There must be at least one non-parking lane for through traffic. This type of parking is used on links where there are through trips to other Zones. • **Off-street parking** can be used to model dedicated parking facilities, with either one-way or two-way aisles. (Any link within an off-street parking Zone will have zero cost for routing.) • **Instant parking** is used as a quick solution when there is no requirement to model the parking manoeuvres within the parking Zone. For more accurate trip times, it is recommended that a fixed time delay is added to the walkway entrances and exits, to model the amount of time that a person would require to walk to or from the parking bay. • **Drop-Off/ Pick-up parking** allows vehicles to stop for a short time to allow passengers to get in or out of the vehicle. Drop-off Zones can be used by private vehicles and taxis, but Pick-Up Zones can be be used only by private vehicles. • **Taxi rank parking** is for taxis only, and allows only pick-up. It differs from other Zones in that vehicles will park in a first-in, first-out order, and all will move forward when one departs. A vehicle is marked as a taxi in the “Description” section of the Types/Vehicles parameters.
Price Base	The base price of parking in this Zone, in the small currency unit (e.g. cents).

Parameter	Effect
Price/ Hour	The time-based price per hour for parking in this Zone, which is added to the base price. The parking duration for each vehicle is set in the parking section of behaviour parameters.
Exp. Walk	The expected walk time from the chosen parking bay to the exit of the parking Zone. This is the time required to exit the parking area on foot as perceived by a driver when selecting a car park. This walking time is added to any other walking times when calculating the total cost of a parking option, taking into account the relative value of time spent walking, as defined by the weights in the behaviour/walking section.
Circles	If this parking Zone has been selected, but no free bays have been found, then the circles parameter controls how many times the driver/vehicle will circle before giving up and returning to the origin Zone. This is often used for taxi ranks to control how many times the taxis will circulate near to a rank if there is no space within the rank.
Wait	If this is true, a person arriving at this Zone on foot will wait for a suitable vehicle to be available before claiming it, walking to it, and then departing. If this flag is false, a new vehicle will be generated in a bay selected at random.
Capacity	For Instant Parking, this capacity constraint will prevent vehicles from choosing the car park as a destination if the car park is full. However, this choice of destination normally occurs only *at the Zone of origin*, so for a vehicle trip where the destination parking Zone becomes full after the trip is started, the vehicle will still continue to the original destination, unless the Parking Choice Tool is used to add a rule that diverts such vehicles to an alternative Zone.
Initial	For Instant Parking, this specifies the number of vehicles in the parking Zone at the start of the Simulation Term. The number of available spaces is the capacity minus this number. The available spaces in an instant parking Zone will increase if people arrive by foot and leave by car from this Zone.

Edit Actions

Some of the Edit actions (Copy, Cut, Paste) use a temporary storage place known as the application clipboard. This is for holding one set of objects for later pasting in the model. This is not the same as the system clipboard, so objects will not be available for pasting in other applications, or other instances of Commuter.

Action	Selection	Effect
Copy	1+ objects	Copy selected objects to the application clipboard, for pasting in another location. The position of the selected objects relative to the cursor is stored, and used when pasting.
Cut	1+ objects	Cut selected objects from the model, and place on the application clipboard. This can be used to move objects a long distance, rather than dragging them.
Delete	1+ objects	Delete selected objects from the model. Objects will not be placed on the clipboard.
Deselect	1+ objects	Clear all selections, so that nothing is selected. The escape key 'Esc' can also be used to deselect.
Paste	(none)	Paste all objects previously copied or cut to the application clipboard into the model. The position of the pasted objects is relative to the current cursor position; if an object was directly underneath the cursor when copied to the clipboard, then it will be pasted at the cursor.

Action	Selection	Effect
Rename	1+ objects	Rename the selected objects. If you want to rename several objects, it may be easier to use the Adjust action, which shows a single window for all objects.
Switch	1+ objects	A multi-purpose action. It can be used to switch the state of a Loop, or a Crossing, or for a user-defined action, through the application programming interface (API).
Undo	None required	Undo the last action. An ordered list of all actions is visible in the Undo tab on the left pane of the main window.

Agent Actions

An Agent is either a person or a vehicle in the model. These actions apply to both types of Agent, unless specified otherwise. An Agent can be selected and moved by dragging, just like any other object.

Action	Selection	Effect
Apply Speed Profile	1 Agent	Apply a controlled speed profile, using the API. This can, for example, be used to make a leading vehicle speed up then slow down, in order to study the effect on following traffic.
Cancel Speed Profile	1 Agent	Cancel the controlled speed profile.
Halt	1 Agent	Halt an agent, freezing it on the spot.
Resume	1 Halted Agent	Resume an agent after a halt.
Tracking On	1 Agent	Follow an agent through the model, by moving the view automatically as the agent moves.
Tracking Off	1 Tracked Agent	Turn off tracking.

Area Actions

An Area is a bounded space used as an origin and/or destination for person-trips. It is like a Zone but for people.

Action	Selection	Effect
Combine	2+ Areas	Combine 2 or more Areas into a single Area, modifying all demand matrices and volumes accordingly. See also Divide.
Connect Stand	1 Area 1 Stand	Add a connection between a Stand and an Area. If the Stand is on a link at the edge of the network then this action will produce an 'arrival' Stand for passengers leaving the network or 'departure' Stand for passengers entering the network. At the edge of a network, the Stand does not represent a real bus-stop or station but is the equivalent of a Zone connector for traffic.
Connect To Zone	1 Area 1 Zone	Add a connection between an Area and a Zone, normally on the edge of a network. This type of connection allows a person to transfer to or from a private vehicle directly, without the need to drive or walk between parking bays.
Create Symbol	1+ Areas	Create a new Symbol with the same size and shape as each of the selected Areas.
Disconnect From Arrival Stand	1 Area	Remove any connections to an arrival Stand.
Disconnect From Departure Stand	1 Area	Remove any connections to a departure Stand.
Disconnect from Zone	1 Area	Remove connection to Zone.
Divide	1 Area	Divide this Area into 2 or more Areas, modifying all demand matrices and volumes. See also Combine.
Label Auto-Centre	1+ Area	Centre the label for the selected Areas.

Mode Freight	1 Area	Change this Area to make it an origin and destination for Freight. See also People.
Mode People	1 Area	Change this Area to make it an origin and destination for People. See also Freight.
New Area at Cursor	(None)	If Tilt is zero, add a new square Area at the cursor. If the view is tilted, the Area will be rectangular along the line of the cursor. Two small check marks on the cursor indicate the size of the new Area.
New Boundary Point	1 Area Boundary Point	Add a new point on the boundary midway between the selected boundary point and the next point, travelling clockwise around the boundary.
Remove Boundary Point	1 Area Boundary Point	Remove the selected point from the boundary.

Display Actions

Action	Selection	Effect
New Hyperlink at Cursor	1 Target object	Create a hyperlink to a target network object. The target object can be a Controller, Intersection, Loop, Sign, Node or Stand. To use this action, first select the object you want to be the target destination of the hyper link. Move the cursor to the position at which you want to place a clickable label, that when clicked, will move the cursor to the target object.
New Image at Cursor	(None)	Add a new Image centred at the cursor, choosing from Images included in the Image library for this network. To add to the Image library select Images/3D from the Display menu.
New Image on Sign	1 Sign	Add a new Image from the Image library to the selected Sign.
New Label at Cursor	(None)	Add a new Label at the current cursor position. A Label can be just annotation, or it can be a hyperlink to an external file or a web address.
New Map at Cursor	(None)	Add a new map or aerial photo at the cursor position, by downloading from the internet.
New Structure at Cursor	(None)	Add a new Structure centred at the cursor, choosing from Structures included in the Structure library for this network. To add to the Structure library select Images/3D from the Display menu.
New Symbol at Cursor	(None)	Add a new Symbol at the cursor position. A Symbol is a shape, similar to a Zone or an Area, used for annotation only. It can be a flat polygon or extended vertically to form a prism.
New Symbol Boundary Point	1 Symbol Boundary Point	Add a new point on the boundary midway between the selected boundary point and the next point, travelling clockwise around the boundary.

Action	Selection	Effect
Remove Boundary Point	1 Symbol Boundary Point	Remove the selected point from the boundary of the Symbol.
Remove Image from Sign	1 Sign	Remove an Image from a Sign.
Scale Image	1 Image	Scale a horizontal Image by equal values in the width and depth dimensions. This is useful for scaling an aerial photo while retaining the aspect.
Toggle Blending Mode	(None)	The blending mode affects how detailed 3D shapes are drawn. The effect of changing this mode depends on your graphics card hardware.
Toggle Brightness	(None)	Toggle between normal mode and bright mode, where lines are drawn thicker, for presentations.
Toggle Cost	(None)	Toggle display of cost for selected Agent.
Toggle Dashboard	(None)	Switch the "Head-Up Display" on or off.
Toggle Link Following	(None)	If Link Following is on, the cursor will move to the end of a new road or rail object when it is first added.
Toggle North/ South	(None)	Cycle through three display modes, displaying All → Northbound Only → Southbound only.
Toggle Orthographic / Perspective	(None)	The projection method changes your view of the model as it is tilted. Toggle between Orthographic and Perspective projection.
Toggle Proximity Selection	(None)	Proximity selection automatically selects the object nearest the cursor, without the need for a separate gesture, such as a mouse-click.
Go To Object By Name	(None)	This action is available on the toolbar, and allows you to move the cursor to a named network object, such as an Intersection, Lane or Walkway.
Simulation	(None)	This action is available on the toolbar, and

Action	Selection	Effect
Continuous		enables the normal continuous mode of simulation.
Simulation Single Step	(None)	This action is available on the toolbar, and enables the single-step mode of simulation. Once in single step mode, press the Play button or the space bar to move to the next time step.

Geometry Actions

Geometry actions are used to change the height or curvature of roads or walkways. Two classifications of network elements for this purpose are:

- **Way:** Lane or Walkway
- **Bezier**: Lane, Walkway, Stream or Connection

Each Bezier element has 4 control points that control the curvature. The curve is known as a Cubic Bezier curve, also sometimes referred to as a spline. Each Bezier element may have one of three "sticky" attributes: straight, arced or smoothed.

Action	Selection	Effect
Change Height	1+ Way	Change the height of all 4 points by the same amount.
Curve Best Fit	1+ Bezier	The control points are moved to be 1/3 of the distance along the curve, while retaining the tangents and gradients at each end of the curve. This operation is not "sticky" - it does not apply any attributes.
Curve Free	1+ Bezier	Remove the "smoothed" attribute from this curve, allowing any of its control points to be moved independently of each other and adjacent elements.
Curve Smooth	1+ Bezier	Set the "smoothed" attribute on this curve, such that when the Bezier element is connected to another, the points on this Curve are moved so that the end of the curve is tangential to the adjacent element. This is a "sticky" operation - control points will subsequently be adjusted if the adjacent curve is moved.
Free / Fixed Centre	1 Centre	For Bezier curves that are arcs, associated with this centre, this changes the way in which the arc is changed as the mouse is dragged. If the centre is fixed, then only the radius will change, but if the centre is free then the other end of the arc is fixed
Import Elevations	Base Tile	Import elevation data from a CSV file, and apply it to the Boundary Base Tile.

Action	Selection	Effect
Make Arc Left	1+ Way	Set the “arced” attribute on this curve, create a centre point to the left, and move the control points so that they lie on the arc of a circle centred at that point.
Make Arc Right	1+ Way	Set the “arced” attribute on this curve, create a centre point to the right, and move the control points so that they lie on the arc of a circle centred at that point.
Make Curvable	1+ Bezier	Remove “straight” or “arced” attributes, allowing all 4 control points to be moved independently.
Make Straight	1+ Bezier	Add “straight” attribute, making the 2 intermediate control points invisible.

Intersection Actions

Action	Selection	Effect
Adjust	1+ Intersections	Raises the Intersection Editor even if other objects are selected. The main Adjust action raises the Intersection Editor only if there are no other selections. Binding a key to this action makes it easy to raise the Intersection Editor.
Combine	2+ Intersections	Combine all the selected Intersections to form a single Intersection. All Nodes, Streams and Turns in all the selected Intersections will become part of a single Intersection. Groups and phases will be retained from the lead Intersection only, so it is best to combine Intersections before signalization.
Exclude Node	1+ Nodes	Exclude the selected Node(s) from the Intersection to which it is currently attached. All Streams and Turns over this node will no longer be part of the Intersection. A new Intersection will be created for every Node excluded.
Include Node	1 Node and 1 Intersection	Include the selected Node in the selected Intersection. The Streams and Turns associated with this Node will then be part of the Intersection, and can be edited in the Intersection Editor.
External Control	1+ Intersections (signalized)	Set the selected Intersection(s) as having external signal control. Commuter's own internal Intersection controller will be deactivated. All group or phase changes are expected from an external controller.
Internal Control	1+ Intersections (signalized)	Set the selected Intersection(s) as having internal signal control. Commuter's own Intersection controller will be activated, cycling through phases in the current phase plan, activating the signal groups in each phase.
New Loops at Stop Lines	1 Intersection	Create a new Loop at the end of each entry Lane to the Intersection. This is the default configuration for some signal controllers.

Action	Selection	Effect
New Road Between	2 Intersections	Create a new Road between the selected Intersections. No Streams are created by this action at each end of the new Road; Streams must be added afterwards.
New Road from X	1 Intersection	Create a new Road to the selected Intersection from a location under the cursor - the X point. For best effect, ensure that the selected Intersection is not under or near to the X point, otherwise the new Road created will be very short.
New Road to T	1 Intersection	Create a new Road from the selected Intersection to a location under the far end of the cursor – the T point. For best effect, ensure that the selected Intersection is not under or near to the T point, otherwise the new Road created will be very short.
Rebuild	1+ Intersections	Rebuild the curves that outline the Intersection. This is a drawing operation only, the Intersection out line does not affect the simulation behaviour.
Turn Lanes	1 Intersection	Raise the Turn Lane Editor window, allowing you to create and remove Streams quickly by selecting from a range of pre-set Turning lane options.
Turn Lanes (Auto)	1 Intersection	Automatically “repair” Turn Lanes at an Intersection, by creating and deleting Streams so that no Stream from any approach conflicts with any other Stream from the same approach.

Lane Actions

Action	Selection	Effect
Align	1+ Lanes	Align the ends of all selected Lanes to the median Lane of the Link. This is useful if Lane end points have been moved and there are gaps between the Lanes near an Intersection.
Bisect	1+ Lanes	Divide the selected Lane into two equal parts. Both Lanes will be on the original Link; no new Node is created. This is useful if you want to create a curve that is more complex than is possible with a single 4-point cubic Bezier curve.
Combine	2+ Lanes (in pairs)	The opposite of Bisect. Replace each pair of Lanes in the selection with a single Lane. Both Lanes must be on the same Link, and one Lane must be connected directly to the other.
Connect	2 Lanes or Lane + Intersection	Connect a pair of Lanes, where both Lanes are on the same Link, or where each is on a Link that is already connected to the other at a Node. If the Links are not connected use Road > Connect. For a Lane that has become detached from an Intersection, this action connects it to the Intersection so that it can then have Streams added. The Lane must belong to a Link that is already connected to a Node on the Intersection.
Disconnect	2 Lanes	Disconnect a pair of Lanes from each other. The Lanes must be on the same Link or on Links that share a Node.
Lane Changing Allowed	2 Lanes (adjacent)	Enable Lane changing on a pair of Lanes. Lane Changing Allowed is the default state– new roads created will have Lanes in this state. Lanes in this state are drawn separated by a dashed line.
Lane Changing Barred	2 Lanes (adjacent)	Disable Lane changing on a pair of Lanes. Lanes in this state are drawn separated by a solid line.

Action	Selection	Effect
Make Filter	1 Lane	Convert the selected Lane to a Filter Lane. A Filter Lane goes directly to the kerb side exit from the Intersection at the end of the Lane. Traffic exits a Filter Lane by moving sideways into the kerbside Lane on the exit Link.
Make Non-Filter	1 Lane (Filter)	Convert the Filter Lane back to a normal Lane.
New Crossing	1 Lane	Add a pedestrian Crossing to the road to which the selected Lane belongs. The New Crossing Parameters window will be raised to allow you to set the parameters for the new Crossing. If you have elected to skip this step, you can raise the parameters window using Walkway > New Crossing Parameters.
New Filter	2 Lanes (entry Lane + exit Lane)	Create a new Filter Lane at the end and to the side of the selected entry Lane. The exit Lane must be on the first exit Link from the entry Lane's Link.
New Lane Beside	1 Lane	Create a new Lane beside the selected Lane. A window is raised to allow you to select parameters for the new Lane, including its length and position.
New Lane Parameters	(none)	Set the parameters (width etc.) that are used when a new Lane is created. This action is active only if you clear the check box on the window that is displayed when you add a new Lane.
New Loop	1+ Lanes	Add a Loop on each selected Lane where the downstream edge of the new Loop is at the nearest point on the Lane to the cursor.
New Multi-Loop	1 Lane + 1 Controller	Add a multi-loop spanning all Lanes on the road, and associate it with the selected Controller. A multi-loop is used for Signal Control.
New Sign	1+ Lanes	Add a Sign on each selected Lane where the Sign is located at the nearest point to the cursor.
New Stand	1+ Lanes	Add a Stand on each selected Lane where the Stand is located at the nearest point to the cursor.

Action	Selection	Effect
New Stream	2 Lanes	Create a Stream between two Lanes, where one is the entry Lane to an Intersection, and the other is an exit Lane.
New Walkway Beside	1+ Lanes	Create a walkway to the kerb side of a Lane, if it has a Lane index of 1. If you select several roads this will create a footpath at the side of each road.
New Zone	1 Lane	Create a new Zone covering the selected Lane. Only one Zone will be created at a time, covering the first selected Lane. This can be used to create a Zone that will be used for parking on a single Lane.
Parking Off	1+ Lanes (Parking)	Return these Parking Lanes to normal Lanes. Any connections to Walkways will be removed.
Parking On	1+ Lanes (in Zone)	Configure the selected Lanes as parking Lanes. A parameter window will be raised to allow you to set parameters for bay size and orientation. If this is the first parking Lane in a Zone, then a window will be raised to allow you to specify the type of the parking Zone.
Shorten End	1+ Lanes	This action reduces the length of each of the selected Lanes at the end by a distance equal to the width of that Lane.
Shorten Start	1+ Lanes	This action reduces the length of each of the selected Lanes at the start by a distance equal to the width of that Lane.
Single Stream	1+ Lanes	This action allows you to reduce the number of Streams on an entry by limiting each entry Lane to have a single Stream. This is useful only if there are at least as many Lanes as exits, otherwise at least one exit will have no Stream. You can also use the Intersection > Turn Lanes action to adjust.

Loop Actions

Action	Selection	Effect
Connect	2 Loops	Connect 2 Loops together using the “diode” connection method used by some traffic control systems. This connection method causes an occupancy trigger on the upstream loop to signal occupancy on the downstream loop. The Loops should be in the same Lane.
Disconnect	1 Connected Loop	Disconnect the selected Loop from its connected Loop.

Node Actions

Action	Selection	Effect
Combine	2 Nodes	Combine the two selected Nodes to be a single Node, retaining all inward and outward Roads. The total number of Roads into and out of the new Node cannot exceed 8 in each direction.
New Node	(none)	Create a new Node at the cursor X position. The new Node will be outlined with a large circle to show it has no Roads in or out.
New Node Between	2 Nodes	Create a new Road between the selected Nodes. No Streams are created by this action at each end of the new Road; Streams must be added afterwards.
New Road From X	1 Node	Create a new Road to the selected Node from a location under the cursor - the X point. For best effect, ensure that the selected Node is not under or near to the X point, otherwise the new Road created will be very short.
New Road To T	1 Node	Create a new Road from the selected Node to a location under the far end of the cursor – the T point. For best effect, ensure that the selected Node is not under or near to the T point, otherwise the new Road created will be very short.
Reset Location	1+ Nodes	Move each selected Node to a position at the end-median corner of one of its entry Roads. This is useful when Lanes have been moved by dragging and Nodes are no longer adjacent to the ends of the Roads to which they are connected.

Rail Actions

Action	Selection	Effect
Bisect	1 Rail	Divide the selected Rail into two equal parts, with a new Node at the central point. The new Rails will have the same attributes as the original.
Bisect Diverge	1 Rail	Divide the selected Rail into two equal parts, with a new Node at the central point, and add a new one-way exit Rail leaving at an angle of 45º from the central point.
Bisect Merge	1 Rail	Divide the selected Rail into two equal parts, with a new Node at the central point, and add a new one-way entry Rail joining at an angle of 45º to the central point.
Connect	2 Rails or Rail + Node/ Intersection	Connect the end of one Rail to the start of the second Rail, deleting the Node that was at the end of the first Rail. Connect the end of the Rail to the Node, deleting the Node that was at the end of the Rail. If an Intersection is selected, and it is more than one Node, then you will be prompted to choose a Node.
New Level Crossing	1 Rail	Create a Level Crossing at the end of the selected Rail. This is equivalent to a 4-way cross roads, where 2 Links are Rails and 2 are Roads.
New Rail Parameters	(none)	Set the parameters (width etc.) that are used when a new Rail is created.
New Rail at Cursor	(none)	Create a new Rail at the current cursor position, starting at the X point and finishing at the T point. By default, a window will be raised each time to allow you to set parameters for the Rail, but you can choose to skip this step on subsequent calls to this action. You can raise the window at any time using the New Rail Parameters action.

Action	Selection	Effect
New Rail From X	1 Rail	Create a new Rail to the start of the selected Rail from a location under the cursor - the X point. For best effect, ensure that the selected Intersection is not under or near to the X point, otherwise the new Rail created will be very short.
New Rail To T	1 Rail	Create a new Rail from the end of the selected Rail to a location under the far end of the cursor – the T point. For best effect, ensure that the selected Intersection is not under or near to the T point, otherwise the new Rail created will be very short.
Route Split	1 Rail	Set the Route Split Values for the selected Rail. For undirected Routing, Route Splits are used to control how many trains take each exit. A Split can be defined for a mob, or for all types (if no mob is selected). You can specify % values, which must sum to 100, or absolute values, which will be converted to % values on Apply. If you set the override flag for a Route Split, it acts like a simple Route choice rule for Directed Routing, overriding the Routes selected by the cost-to-destination algorithm.

Road Actions

Action	Selection	Effect
Bisect	1 Road	Divide the selected Road into two equal parts, with a new Node at the central point. The new Roads will have the same attributes as the original
Bisect Diverge	1 Road	Divide the selected Road into two equal parts, with a new Node at the central point, and add a new one-way exit Road leaving at an angle of 45º from straight ahead from the new Node
Bisect Merge	1 Road	Divide the selected Road into two equal parts, with a new Node at the central point, and add a new one-way entry Road joining at an angle of 45º to the selected Road into the new Node
Bisect Side Road Left	1 Road	Divide the selected Road into two equal parts, with a new Node at the central point, and add a new two-way Road on the left side at an angle of 90º from the selected Road connected to the new Node
Bisect Side Road Right	1 Road	Divide the selected Road into two equal parts, with a new Node at the central point, and add a new two-way Road on the right side at an angle of 90º from the selected Road connected to the new Node
Combine	2 Roads (common Node)	This action is the opposite of Bisect, removing a central Node and creating one Road out of two. The Roads selected must have a common Node, and that Node can have no other Roads attached

Action	Selection	Effect
Connect	2 Roads or Road + Node/ Intersection	Connect the end of one Road to the start of the second Road, deleting the node that was at the end of the first Road. Connect the end of the Road to the Node, deleting the Node that was at the end of the Road. If an Intersection is selected, and it is more than one Node, then you will be prompted to choose Node
New Intersection	1 Road	Create a new Intersection at the end of the Road. The intersection can be 3-way or 4-way. A window is raised to set parameters for the new Intersection, which can be signalised, priority all-way Stop or a mini-roundabout. For a full-size roundabout use Road > New Roundabout.
New Lane	1 Road	Create a new Lane to one side of the selected Road. A window is raised to allow you to select which side of the Road, and parameters for the new lane, including its length and position.
New Road Parameters	(none)	Set the parameters (width etc.) that are used when a new Road is created. This action is active only if you clear the check box on the window that is displayed when you add a new Road.
New Road at Cursor	(none)	Create a new Road at the current cursor position, starting at the X point and finishing at the T point. By default, a window will be raised each time to allow you to set parameters for the Road, but you can choose to skip this step on subsequent calls to this action. You can raise the window at any time using the New Road Parameters action.
New Road from X	1 Road	Create a new Road to the start of the selected Road from a location under the cursor - the X point. For best effect, ensure that the selected Intersection is not under or near to the X point, otherwise the new Road created will be very short.

Action	Selection	Effect
New Road to T	1 Road	Create a new Road from the end of the selected Road to a location under the far end of the cursor – the T point. For best effect, ensure that the selected Intersection is not under or near to the T point, otherwise the new Road created will be very short
New Roundabout	1 Road	Create a new Roundabout at the end of the Road. The Roundabout can be 3-way or 4-way. A window is raised to set parameters for the new Roundabout and its arms, including inner radius, circulating lanes, spiral lanes, etc.
New Zone	1+ Roads	Create a new Zone that encloses the ends of the selected Road. This will enclose the Road only if it is straight. This action is intended for zone connector link Roads. The Road must have an intersection at one end only.
Route Split	1 Road	Set the route split values for the selected Road. For undirected routing, route splits are used to control how many vehicles take each exit. A Split can be defined for a mob, or for all types (if no mob is selected). You can specify % values, which must sum to 100, or absolute values, which will be converted to % values on Apply. If you set the override flag for a route split, it acts like a simple route choice rule for Directed Routing, overriding the routes selected by the cost-to-destination algorithm.
Rules for Gateway Choice	1 Lane / 1 Road	Raises a window that allows you to create or edit gateway choice rules for the selected Road. A gateway choice rule creates an intermediate destination for a vehicle, without changing its final destination.
Rules For Lane Choice	1 Lane / 1 Road	Raises a window that allows you to create or edit lane choice rules for the selected Road. Lane choice rules allow you to control the target lane range for any group of vehicles, selected by mob, origin, destination, intersection exit, etc.

Action	Selection	Effect
Rules for Parking Choice	1 Lane / 1 Road	Similar to the Lane Choice Rules function, this allows you to specify rules that control the choice of **destination** zone for vehicles. The selection criteria for these rules include the number of remaining spaces in the current destination parking zone. That is, when a parking zone nears capacity, you can define rules that redirect traffic to alternative parking
Rules For Route Choice	1 Lane / 1 Road (2+ exits)	Similar to the Lane Choice Rules function, this allows you to specify rules that control the choice of **exit** from the selected Road. These rules override the standard routing logic
Rules for Stream Choice	1 Lane	Similar to the Lane Choice Rules function, this allows you to specify rules that control the choice of stream from a lane to the chosen exit. These rules do not affect the choice of exit road, so for a stream choice to be valid there must be two or more streams from the chosen lane to the chosen exit road. It will not be valid if there are multiple streams form the lane but each goes to a different exit road.
go to End	1 Road	Move the cursor to the end of the current Road. Binding this action to a key allows you to move smoothly to the end of a Road, where you may add a new Intersection or add a new Road to an Intersection that is already there.

Sector Actions

A Sector can be used for aggregating measurements by grouping Zones or Areas by location or function.

Sectors can be edited using the Sectors window from the Network menu.

Action	Selection	Effect
New Sector	1+ Area or Zone	Create a Sector containing all the selected Zones or Areas.

Signal Actions

Action	Selection	Effect
New Signal	1 Signal	Create a copy of the selected Signal equipment. This copy is made purely for visualization purposes, for example to place a Signal equipment box on the far side of the Intersection as well as the near side. Multiple Signal equipment boxes for each approach have no effect on the simulation.
Reset Position	1 Signal	If a Signal equipment box has been moved from its default position, this action resets it to its default position.

Stand Actions

Action	Selection	Effect
Arrival Split	1 Stand (with 2+ exit walkways)	Set the route split values for the selected Stand for passengers alighting ~~arriving~~ at the Stand. For undirected routing, route splits are used to control how many people take each exit. A Split can be defined for a mob, or for all types (if no mob is selected). You can specify % values, which must sum to 100, or absolute values, which will be converted to % values on Apply. If you set the override flag for a route split, it acts like a simple route choice rule for Directed Routing, overriding the routes selected by the cost-to-destination algorithm.
Boarding Split	1 Stand (with 2+ Services)	As for Arrival Split, but splitting the number of boarding passengers between Services that stop at the selected Stand.
Destination Split	1 Stand (with 2+ subsequent Stands on Service)	As for Arrival Split, but splitting the number of passengers that alight at each of the next Stands. If there were 4 Stands on the Service, and the selected Stand was the first, then a 3-way Split would be possible, dividing passengers between each of the possible “destination” Stands.

Stream Actions

Action	Selection	Effect
Delete	2 Lanes	Delete any Stream that connects the two selected Lanes. This can be used instead of the more general Edit > Delete action on a selected Stream. On a complex Intersection where there are a large number of Streams, it can be difficult to click on the required Stream to select it, and it may be easier to select the corresponding pair of Lanes.
New Turn	1 Stream	Normally, there is only a single Turn for each entry-exit pair of Links. In some cases, it may be necessary to create a second Turn for an entry-exit pair, in order to provide Vehicle-specific or Lane-specific signalling for that turn. An example of this on a signalised intersection is a bus-priority signal group that shows a green signal to buses for a movement while other traffic is shown a red signal. This action creates a second turn for the entry-exit pair defined by the selected Stream. This duplicate Turn will then appear in the Intersection editor and can be assigned to Groups and Phases as required.
Switch Turn	1 Stream (2+ Turns available)	Where more than one Turn is available for a Stream, as described for New Turn, this action allows you to switch a Stream from one Turn to another. For the bus priority example, the Stream leaving from the bus Lane should be associated with the Turn that is controlled by the bus group.

Walkway Actions

A Walkway is a bi-directional Path for pedestrians, who will, in general stay within the boundaries, but might go outside the boundaries in congestion or if there is a sharp turn. A Crossing is a special type of Walkway associated with a set of vehicle Lanes, and represents a place where there is some type of traffic control for pedestrian Crossing. A (plain) Walkway can also cross a road, but pedestrians will be seen by vehicles on the road only if the avoidance flag is enabled for each Lane. If the avoidance flag is not set, pedestrians will not be seen by vehicles, and vehicles will not slow down to avoid them.

Action	Selection	Effect
Bisect	1 Walkway	Create 2 Walkways from 1 by splitting the selected Walkway in the middle, and adding a Connection in the middle.
Change Crossing Type	1 Crossing	Change the type of a Crossing while retaining all the Connections. Crossings of different types are different objects, so this will delete the old Crossing and create a new one, while attempting to retain all the parameters.
New Area	1 Walkway	Create a new Area covering the unconnected end of a Walkway. The Walkway must be connected to another Walkway at the other end.
New Connection	2+ Walkways	Create a Connection between the nearest ends of each pair in the set of selected Walkways. In the simplest case, if you have 2 Walkways selected, this will create 1 Connection between them.
New Connection to Parking	1 Walkway + 1 Parking Lane	Create a Connection between the selected Walkway and the selected Parking Lane. This also aligns the Walkway to run beside the Lane. This type of Connection allows people to transfer between walking and driving.
New Connection to Stand	1 Walkway + 1 Stand	Create a Connection between the selected Walkway and the selected Transport Service Stand. This type of Connection allows people to transfer between walking and riding on public transport.

Action	Selection	Effect
New Crossing Parameters	(none)	Set the parameters (width etc.) that are used when a new Crossing is created. A Crossing is created using Lane > New Crossing.
New Delay	1 Walkway	Set a delay at either end of the selected Walkway. This will cause each person to pause for that time before proceeding. Use “inf” (infinity) to block the Walkway in that direction.
New Lift at Cursor	(none)	Create a new Lift object at the cursor position. A pop-up window will prompt for the size and capacity of the lift, the number of floors and the heights between the floors. Once created, there will be a pair of Stand objects on each floor. Walkways can be connected to these Stands as would be done for bus stops or train platforms.
New Path Junction	1 Walkway	Create an “Intersection” of Paths at the unconnected end of the selected Path. If both ends are unconnected then the new Paths and Connections will be created at the D end.
New Walkway Parameters	(none)	Set the parameters (width etc.) that are used when a new Walkway is created.
New Walkway at Cursor	(none)	Create a new Walkway at the current cursor position, starting at the X point and finishing at the T point. By default, a window will be raised each time to allow you to set parameters for the Walkway, but you can choose to skip this step on subsequent calls to this action. You can raise the window at any time using the New Walkway Parameters action.
Parking Alignment	1 Walkway (connected to parking)	For a Walkway that is connected to a Parking Lane, align the Walkway to the side of the Lane, changing its length and shape as required.
Parking Disconnect	1 Walkway (connected to parking)	Disconnect the Walkway from a Parking Lane. If there are no Walkways connected to a Parking Lane, the parking bays will be highlighted in red.

Action	Selection	Effect
Route Split CD / Route Split DC	1 Walkway (2+ exits)	Set the route split values for the selected Walkway in the direction CD or DC. For undirected routing, route splits are used to control how many people take each exit. A Split can be defined for a mob, or for all behaviours (if no mob is selected). You can specify % values, which must sum to 100, or absolute values, which will be converted to % values on Apply. If you set the override flag for a route split, it acts like a simple route choice rule for Directed Routing, overriding the routes selected by the cost-to-destination algorithm.
Rules for Destination Choice	1 Walkway	Raises a window that allows you to create or edit destination choice rules for the Walkway. These rules control the choice of **final destination** Area. These allow you to switch destinations according to congestion or other factors, but this will mean that the arrival OD will be different from the departure OD.
Rules for Gateway Choice	1 Walkway	Raises a window that allows you to create or edit gateway choice rules for the selected Walkway. A gateway choice rule creates an **intermediate destination** for a person, without changing its final destination. These rules create a **via** instruction, and are particularly useful for directing people towards particular car parks if there are several to choose from.
Rules for Route Choice	1 Walkway (2+ exits)	Raises a window that allows you to create or edit route choice rules for the selected Walkway. These rules allow control over the choice of exit from the selected Walkway, overriding the standard lowest-cost route selection. Selection criteria for the rules include mob, origin, destination and time of day.

Zone Actions

A Zone is a bounded space used as an origin and/or destination for vehicle trips.

Action	Selection	Effect
Combine	2+ Zones	Combine 2 or more Zones into a single Zone, modifying all demand matrices and volumes accordingly. See also Divide.
Divide	1 Area	Divide this Zone into 2 or more Zones, modifying all demand matrices and volumes accordingly. See also Combine.
Label Auto-Centre	1+ Zone	Centre the label for the selected Zones.
New Zone at Cursor	(None)	If Tilt is zero add a new square Zone at the cursor. If the view is tilted, the Zone will be rectangular, along the line of the cursor. Two small check marks on the cursor indicate the size of the new Zone.
New Boundary Point	1 Zone Boundary Point	Add a new point on the boundary midway between the selected boundary point and the next point, travelling clockwise around the boundary.
Remove Boundary Point	1 Zone Boundary Point	Remove the selected point from the boundary.

API

Reference

Commuter API – Simple Plugin Template

To create and install a plugin:

- Ensure you have a license for creating user-defined plugins. The Help / About window should display "User Plugins"
- Write your plugin java source code: the plugin class must extend com.azalient.api.BasePlugin, and implement any event interfaces for which you want to register
- Your classes should be in an appropriate package to avoid name clashes. By convention, reverse your internet domain: for example, if your domain is thu.edu.cn, the package should be cn.edu.thu
- In the constructor of the plugin class, add a call to register the plugin for each event
- The event interfaces are in packages com.azalient.api.event.model.* and com.azalient.api.event.agent.*
- For example if you implement ModelEventTimeStep, and you register the plugin using addModelEventTimeStepListener(this); your plugin will be called on every simulation time step. Similarly if you implement AgentEventLoop, and add the relevant registration call, your plugin will be called whenever a vehicle activates a loop.
- Once you have written the Java code for your plugin, compile it into a class file. Follow standard Java conventions: the class file should be in a folder that agrees with the package name. For example, if your plugin is MyPlugin..java in package cn.edu.thu, the class file should be in a file cn/edu/thu/MyPlugin.class
- Package up the class file(s) into a JAR file, and put that jar file into the bin directory of your Commuter installation, beside the other jar files. This is normally C:/Azalient/Commuter/bin. It will be automatically added to the class path from here by the Commuter.exe launcher.
- Add a 3-letter acronym, the name of the plugin and the "main" class (this class) to commuter.plugins.csv

  For example: "XYZ,Test Plugin,cn/edu/thu/MyPlugin"

  (The 3-letter acronym is historical, it is required but the value is not important)
- When you next start Commuter, your plugin will be visible in the Plugins Tab

```
package com.azalient.test;


import java.awt.Color;
import java.awt.Window;
import java.io.File;
import java.io.IOException;
import java.io.PrintWriter;

import javax.swing.JFrame;
import javax.swing.JLabel;

import com.azalient.api.BasePlugin;
import com.azalient.api.b.control.IController;
import com.azalient.api.b.control.IGroup;
import com.azalient.api.b.control.IPhase;
import com.azalient.api.draw.IDrawing;
import com.azalient.api.event.model.ModelEventViewLegend;
import com.azalient.api.quick.Api;
import com.azalient.apo.SC;
import com.azalient.apo.enums.Signal;


/** This is a template for a Commuter plugin.
 *
 * <p>
 * The steps you need to take to install a plugin are as follows:
 * <ul>
 * <li> The new plugin must extend BasePlugin, and implement any event interfaces
 * for which you want to register
 * <li> Add a 3-letter code, the name of the plugin and the "main" class (this class)
 *  to commuter.plugins.csv <br>
 *  For example:  "TTT,Test Plugin,com/mydomain/test/TestPlugin".
 * The 3-letter code is historical, it is not important.
 * <li> Package up the class files into a JAR file, and put that jar file into the bin directory
 * of your Commuter installation, beside the other jar files. This is normally
 *  C:/Azalient/Commuter/bin. It will be automatically added to the class path from here
 *  by the Commuter.exe launcher.
 *
 * </ul> */


public class TestPlugin  extends BasePlugin
        implements ModelEventViewLegend
        /* implement the event handlers corresponding to the events you want
        to be sent to this plugin */
{
    /**
         * In the constructor, register for the events you want to be sent to this
         * plugin. See the packages com.azalient.api.event.agent and .model for all
         * the events. There are both model events, corresponding to significant
         * transitions in the model loading and handling cycle, and agent avents,
         * which correspond to significant events in the life cycle of an agent
         * (vehicle, pedestrian etc) in any simulation
         **/
```

```
    public TestPlugin()
    {
        addModelEventViewLegendListener(this);

        frame.getContentPane().add(new JLabel("Add your own controls here, to allow user
interaction with your plugin."));
        frame.pack();

    }

    /** This is the implementation of the event handler registered above */
    public void viewLegend(IDrawing drw)
    {
        drw.colour(Color.WHITE.getRGB());
        drw.string("Test Plugin", 0.5, 0.7, 0, 0.5);
    }

    /** This is called for all plugins, after the network "document" has been loaded */
    public void pluginOpen()
    {
            File file = new File("c:/temp/phases.csv");
            try
            {
                    PrintWriter pw = new PrintWriter(file);


pw.println("Phase,Group*,Group*,Group*,Group*,Group*,Group*,Group*,Group*,Group*,Grou
p*,Group*,Group*,Min,Gap,Loop*,");

            for (IPhase phase: Api.model().control().phases())
            {
                    Signal[] signals = phase.signalArray();

                    IController c = phase.controller();
                    IGroup[] ga = c.groups();

                    int ns = signals.length;
                    int ng = ga.length;

                    pw.print(phase.name()+SC.CM);
                    for (int i = 0; i < ng && i < ns; i++)
                    {
                            if (signals[i] == Signal.Green)
                            {
                                    pw.print(SC.N+ga[i].index()+SC.CM);
                            }
                    }
                    pw.println();
            }

            pw.close();
            }
            catch (IOException iox) {}

    }
```

```
    /** This is called for all plugins, when a network "document" is closed */
    public void pluginClose()
    {
        frame.setVisible(false);
        frame.dispose();
    }

    /** This is called for all plugins,  when "File/Save" is selected, to save any data in the
plugin */
    public void pluginSave() {}

    private JFrame frame = new JFrame("Test Plugin User Interface");

    /** called when the "Configure" button is pressed on the Plugins tab */
    public Window pluginWindow()
    {
        return frame;
    }



}
```

Plugins

Parking Directions Plugin

The main function of this plugin is to allow route choice rules to be defined based on the number of empty bays remaining within a section (or "cluster") of a parking zone. Using this information, vehicles will be re-routed within a zone to another part of that zone.

To be effective there needs to be at least two clusters in a zone.

To set up a model to use this plugin:

1. Divide the parking zone into clusters. A cluster is identified by its name, which can be letters or numbers or a combination of both. To set the cluster for a parking lane, Select the Lane, then Action > Adjust, go to Parking Tab, change value for Cluster. As with any change in the Adjust window, you can select multiple lanes and change them all simultaneously, using "Apply To All". A parking lane can be in no cluster, in which case the cluster field will be blank.
2. Add a sign for each cluster. Sign location is not that important, as they vehicles can "see" a sign for the route choice rules wherever it is, but it makes sense to create it on the link where the route choice rule will be defined. Each sign will display the number of spaces in the associated cluster, once the association has been made.
3. Save and re-open the model.
4. All distinct clusters will be listed on the Clusters tab of the plugin window, showing the number of bays and a colour for the cluster. The colour is editable.
5. Switch to the Signs tab and associate one sign with each cluster, by selecting the cluster name from the list. The sign will display the number of spaces in the associated cluster.
6. Create Route Choice rules, where each rule uses the name of a sign as a variable to represent the number of spaces. If your sign is called SignXYZ, and you included a comparison phrase "< 12" (without the quotes) then the rule will be applied if the number of spaces shown on SignXYZ is less than 12.

Lane Control

The Lane Control plugin allows you to model time-varying changes to lane controls, such as variable speed controls and lane closures. You can define one or more lane control areas within the same model. Within each of the lane control areas, a set of states is defined, and a set of transitions defining how the system will move from one state to another. The transition may be instantaneous, or may be phased, for example in tidal flow systems a lane will be gradually closed for both directions of flow, before being opened in the opposite direction. The transitions can be initiated at pre-set times, using a “program”, or when a set of conditions are met, using a rule (similar to signal control rules). The transitions can also be tested by the user using an interactive user interface.

For example, in a tidal flow scenario which has two states **4_4** (4 lanes south / 4 lanes north) and **5_3** (5 lanes south / 3 lanes north) a transition could be made between these states

- at the simulation time of 07:30 (program)
- if the southbound flow is more than 20% higher than the northbound flow (rule)
- when the user presses a button to activate the transition, based on the user’s judgement about flow levels

The specification for a lane control area contains the following objects:

- gantries
- messages
- barriers
- states
- sequences
- transitions
- programs
- rules

Description of Lane Control Objects

The lane control specification consists of its physical objects, such as gantries, and its logical objects, such as states and transitions.

A **gantry** is a device that controls lane behaviour on a single link. Its location is specified by link and a position from the start of the link, but the position is for visualisation only: it always controls all lanes along the entire length of the link. The link must have a uniform number of lanes along its length.

A **message** is displayed on a gantry, above one of the lanes spanned by the gantry. A message:

- has a descriptive name
- has a display type: either constant or flashing, and a colour
- has a string of up to 3 display characters which is aligned centrally. This is commonly a number for a speed control, or special values '^' for open and 'X' for closed. It can be the special value “off” which causes nothing to be displayed
- has either an associated lane restriction or an associated speed control rule. The restriction and/or speed control rule is selected from those defined in the model.

A **barrier** is a movable longitudinal device, also referred to as a “movable median” or “flipper”. These are used at the entrances and exits to tidal flow systems to guide vehicles to the appropriate lanes. Each barrier has a descriptive name and a list of positions, each of which has an associated pattern of open lanes and next lanes .

A **state** is a specification of the messages on all gantries, the positions of all barriers, and messages on all signs. A state

- has a descriptive name
- has a list of {gantry, lane} pairs with an associated message.
- has a list of {barrier, position} pairs
- has a list of {sign, message} pairs

A **sequence** is a series of messages applied to a gantry lane. A sequence:

- has a descriptive name
- has a series of messages
- can have a duration specified for each message, specified in intervals, and if specified must be a whole number of 1 or more. If no duration is specified, a time of 1.0s is assumed
- will cause the gantry lane sign to rest in the final message

A **transition** specifies how a lane control system moves between two defined states. A transition:

- has a descriptive name
- has a source state and a destination state
- can be called **only** if the subsystem is currently resting in the source state
- is a time series consisting of {gantry lane, sequence} pairs where a sequence is applied to the matching gantry lane at the appropriate time
- does not need to list all gantry lanes
- can have waiting times listed at any point in the series, and if present, these are used to cause the transition to happen as a travelling wave. If no wait times are specified, all sequences will commence at the same time.
- can contain sequences of unequal length

For any control area, only one transition can be active at any time.

A **program** causes a transition to happen at a predetermined time. A program:

- has a descriptive name
- has a named transition
- has a start time in the format HH:MM
- requests that the named transition is called at the time specified, but **only** if the control area is resting in the transition's source state

A **rule** causes a transition to happen when a test is true. A rule

- has a descriptive name
- has a named transition
- has a test expression (for example flow1 > 1000)
- requests that the named transition is called whenever the test is true **and** the control area is resting in the transition's source state

Lane Control Window

Control Area: This line allows you to select one of the lane control areas, to add a new area, rename or delete an existing area.

There are four Edit buttons:

Gantries / Edit...: Allows you to add and delete Gantries and messages, and to view the current state of all the gantries.

Barriers / Edit...: Allows you to add and delete Barriers (lane guidance devices) and define the various positions for these barriers.

States / Edit...: Allows you to add and delete States in the system, and to set the gantry, barrier and sign messages within those states.

Transitions / Edit...: Allows you to add and delete Sequences, Transitions, Rules and Programs

To the right of the Edit buttons are controls that allow you to set each gantry sign or barrier individually, to set a state manually, or to select a transition and initiate that transition.

Toll Plazas

A toll plaza is a location where vehicles can be delayed for a period of time before advancing. On the road chosen as the location for the toll barrier, a toll booth is created at the end of each lane. Each vehicle is delayed at the the toll booth for a period of time taken from a stop-time distribution. The same stop-time distribution can be defined for all booths, or several distributions can be defined, where each represents a different method of payment.

Building a Suitable Location

For best operation, before connecting the toll plaza plugin, a road should be created with the following attributes:

- Make the road short, around 20 metres in length is good
- If there are two or more lanes, use the Action **Lane > Lane Changing Barred** to prevent lane changing
- There should be a single link entry into the toll barrier link. This entry link will be used as the “gateway”, where vehicles become aware of the toll plaza ahead.

Creating Restrictions for Multiple Payment Types

If you want a simple toll plaza, with a single payment type, then you can skip this step. If you want to model more than one payment method, then you should create a restriction for each payment method. For example, if you have electronic payment as one method, but only some vehicles are enabled for electronic payment, then create a vehicle type for those vehicles, perhaps named “E-Tag” and then restriction that allows only E-Tag vehicles to pass.

You may also need to create some lane choice rules upstream of the toll plaza, to move vehicles into the correct payment lane. These lane choice rules model the effects of advance signing, informing drivers of the payment types available in each lane.

Stop Time Distributions

A stop time distribution assigns a percentage share to each of a number of possible stop times. Depending on the data you have collected, you can define as many different stop times as you like. However, each stop time must be a whole number of seconds.

If you want all vehicles to stop for the same time, then simply define a share of 100% for that time.

Modelling Electronic Payment Lanes with no Stop Time

Zero seconds is a valid stop time value; if you define a stop time of zero, create a stop time profile, then assign the zero stop time a 100% share, then no vehicle will stop in any lane which has that profile. This is often used to model electronic payment. If you want to reduce the through speed of vehicles on that lane, without actually causing them to stop, then you can define a speed limit for the toll booth lane.

Creating a New Toll Plaza

Before creating a new plaza, you should build a suitable location, create and assign lane restrictions, if you have more than one payment type, and create stop time distributions, as described above. Once you have completed these preliminary steps, select one lane on the road where you want to locate the toll barrier, select **Edit > New** Plaza and then enter a name and description for the toll plaza.

In the Toll Barriers window, there are three tabs:

- **Name & Gateway:** verify the name, description and the location of the gateway link, where vehicles will first become aware of the toll plaza
- **Approaches & Barriers:** modify the booth names as required
- **Stop Time Assignment:** Set the stop time profile for each vehicle type for each restriction/ payment type.

Managed Motorways

This plugin allows the simulation of variable speed control, activated by measurements of flow, speed and/or occupancy on detectors.

1. Divide links into lanes corresponding to the sections of road that will be controlled by each gantry. That is, if there is a 1000m road, controlled by two gantries at 500m separation, then the lanes on the link should be bisected into 500m length. There is no need to bisect the links, just the lanes.
2. Create gantries, using Lane > Add Gantry
3. Create Multi-Loops, using Lane > Add Multi-Loop. A Multi-Loop creates a Loop on each lane, collects measurements for each of these individually, and reports the average values of flow, speed and occupancy to the management rules.
4. Open the Managed Motorways Plugin, and press [+] on the top row to create a new Area. It is possible to have multiple areas, for example northbound and southbound.
5. Select gantries and multi-loops and add them to the device table. You can use the up and down arrows on the left to sort the devices into order in the table. The flow of traffic is up the table; the first device in the table is furthest downstream. The order is visual, and is also used for the speed gradient.
6. Open the Linkages window and create linkages. A linkage is a connection between a loop and a gantry, indicating that the loop may modify the speed control on the gantry. A loop may be connected to more than one gantry, and a gantry may be connected to more than one loop, but each connection must be set up as a separate linkage. Assign a type letter or word or number to identify similar linkages. For example a linkage between a loop and the immediately upstream gantry might be type A, a linkage to a loop further upstream might be type B, and so on.
7. Create linkage rules that will apply a (speed) sign to all lanes of a gantry if the rule matches the type of the linkage and the flow speed and occupancy of a loop.
8. Leave a field blank to match any value. That is, if the type field is blank, it will match any linkages. If the flow field is blank it will match any level of flow.

SCATSIM Plugin for Commuter

The SCATSIM Plugin for Commuter automates the process of setting up a SCATSIM model, initiates the connection, and allows you to view the information passed as messages between the simulator and SCATS.

The quickest way to build and run a Commuter model with SCATSIM is to import the data directly from the SCATS data files, using the **File → Import from SCATS** action. it is assumed that before you begin the process of connecting to SCATSIM, you have a model that contains one or more SCATS intersections, either imported from SCATS data or built from scratch.

Configuring your model for SCATSIM operation requires the creation of an **Automation** Script. The steps below will walk you through this process.

Before you Begin Creating an Automation Script

Before the automation will work for the first time, it is necessary to set up port numbers and SCATS license information. Select **Action → SCATS Licensing/ Ports.** Check that the customer name and license code for the central manager and the region are set correctly. if you have already run the SCATS configuration tool to enter the license data on this computer, then these will al be set. Check also that the port numbers agree with the settings for the other modules. The window is pre-filled with default values, so unless you have a non-standard configuration, you will not need to make any changes.

You will also need to make sure that you have created an ODBC data source for the temporary working copy of the SCMS.mdb database used for SCATSIM. There is list of steps explaining how to create an ODBC data source at the end of this section.

Creating an Automation Script

The Plugin uses a “wizard”-style user interface to create an automation script. Many of the values you input are stored for the next time, so many of the following steps are required only the first time you run the wizard. The steps where you need to provide input every time are highlighted in red. To begin, press the button with the small green plus on the left side of the SCATSIM Plugin window.

1. On the first page of the wizard, you will be asked to locate a source for the SCMS.mdb file. if you have already imported SCATS data from this file using File → Import from SCATS, then this will have been completed for you. If not, you should browse to the location of the SCMS.mdb file.
2. Check that the destination for the copy has been correctly located as the ODBC data source. If you received an error about querying the registry, and the **Destination for Copy** box is blank, then press Cancel and set up the ODBS data source as explained at the end of this section. Then start again at Step 1.
3. Press Next
4. Browse to each of the 5 executable files (*.exe) required for the SCATS side of SCATSIM. The location of these depends on how you have installed SCATSIM.
5. Select the controller version number (VC) used for the LX/SFT data
6. Select the Simulation: Commuter/ Q-Paramics or Other
7. If you are not using Commuter as the simulation, select the location of the simulation model file or directory.
8. Press Next
9. Check that you have two green ticks, informing you that LX and SFT data has been located for all intersection. If you do not have green ticks then check that the location for the LX and SFT data root directories is correct. If you cannot locate LX and SFT data, then you will have to Cancel and exit.
10. Press Next
11. If you want to export this automation script for use outside of Commuter, choose a location and name for the script file
12. Press Finish.

Saving Automation Scripts

To save any new automation scripts, press the Save button on the left side of the SCATSIM Plugin window, or any Save button on the application window.

Running an Automation Script Manually

1. Click on a Script in the table in the centre of SCATSIM plugin window. The Start Now button will become active
2. Press Start Now
3. Once all processes are started and connected (usually around 5 to 10 seconds), press Play on the Commuter window to start the simulation. Or if you are using another simulator, press the button that starts that simulation.

Running an Automation Script Automatically

If you want to run a script automatically when you open a Commuter model, select the script you want to run in the drop-down box below the table and select the option labelled **Start Automatically**. The next time you open the model, the script will run.

Controlling and Stopping an Automation Script

When an Automation Script has been started, there will be a small **Automation Script Window** in the lower half of the Plugin panel, containing a button for each process. The window title will have the name of the script that you are running.

- **Process Buttons:** The Automation control window is like a mini-taskbar: pressing any of the buttons will make the process window visible, although depending on your operating system and settings, it may not raise the window to the top. Look in the Windows taskbar for the icon representing the process, and click on that to raise the process window.
- **Stop Button:** Pressing the Stop button will stop all of the SCATSIM processes, after a confirmation check.

SCATSIM Plugin for Commuter Menu Options

Action → SCATSIM Intersections

This raises a window showing a list of all intersections in the model, and indicates which intersections will be controlled by SCATS. if you do not want an intersection to be controlled by SCATS, remove the tick from the appropriate box. If the SFT Signal ID or SFT file is not correct for an intersection, you cannot modify these here – go to Action → SFT Files.

Action → LX Regions

This raises a window showing the short code name and number for each region. There is also a column for a description. This description is optional, and defaults to the SCATS short code name. However, you can enter a longer description for clarification.

Action → SFT Files

This raises a window showing a list of all the SFT files that have been discovered. In some cities, such as Sydney, the Signal ID is always the same as the intersection ID, and the SFT file is named based on a convention, so that it is easy to map an SFT file to an intersection and a controller. However, in other cities, no such convention exists, and it may be necessary for you to enter the Signal ID and Intersection ID that correspond to each of the discovered SFT files.

Action → SCATS Licensing/ Ports

This raises a window showing the Customer Name and License Code pairs for the central manager and the region. It is assumed that the same license code is used for al regions, but that this may be (and normally is) different from the central manager code. if you are using non-standard port numbers, you should enter the port numbers in this window, and the automation script will write those port numbers into all the configuration files required for SCATSIM.

Action → SCATS Access

This raises the SCATS Access window, which can be opened “off-line” to examine an intersection layout in SCATS graphics window.

Action → Disconnect

If the SCATSIM Plugin is connected to SimHub, this action will be enabled. it is also available as a button on the SCATSIM Plugin window. Disconnecting from SimHub and then re-connecting may fix a problem if the communications have become unsynchronised following loss of data or another type of communication problem. If the simulation is “hanging” when you press the **Play** button, it may be necessary to Disconnect and restart the communications.

Action → Offer Connection

Select this action after the plugin has been disconnected to re-initiate the communication protocol sequence. This action causes the SCATSIM Plugin to open a port on the agreed port number and wait for the Win-Traff application to connect. Note that it is assumed here and elsewhere in the SCATSIM configuration that all SCATSIM processes are running on the same computer. Although the socket technology used for the connection can operate across a network, the user interface has no facility to enter computer names or internet addresses for the various processes, only port numbers.

View → Messages

This raises a window that allows you to inspect the data messages being sent to and received from the SCATSIM processes. Select Log → Log messages to have a row of information added in the table for each message. Select Options → Reverse order to force the most recent messages to be added on the first line of the table.

View → Highlight Intersections

If this is selected an 80-meter diameter yellow circle is drawn around each intersection that is controlled by SCATSIM.

Option → Send Detector Data
If this option is off, no data messages containing detector data will be sent to SCATSIM

Option → Send Message 254 on Completion
Message 254 (see SCATS-SP38 documentation) causes the SCATSIM system to write the modified data to file, overwriting previous values. You may want to de-select this option to prevent this finalisation action from happening. Completion is defined as the simulation reaching the end of the Simulation term defined for the model.

Option → Disconnect on Completion
This option causes the SCATSIM Plugin to disconnect from the SCATSIM system if the simulator successfully reaches the end of the Simulation Term.

Option → Disconnect Delay
This option allows you to add a delay to wait for any finalisation messages before disconnecting the communication socket.

Option → Auto Connect
This option causes the SCATSIM plugin to initiate the connection mechanism as soon as the model is opened. If you are going to use an automation script, it is better to switch this off.

Option → Use Consoles
This option causes the SCATS and Central manager processes to be run from consoles. The default setting is on. If SCATS or Central Manager is not starting, you can try setting this option to off.

Option → Close All Consoles on Shutdown
This option causes **ALL** **Command Prompt** console windows (cmd.exe processes) to be terminated when an Automation Script is stopped. This is not limited to just those windows that ere started for SCATSIM, this will close all other consoles. If you do not use Command Prompt windows for other tasks you can leave this option set to on.

Option → Run in scatsim_temp

This option causes all SCATSIM processes to write their output files to a central temporary directory which is cleared before any subsequent runs. This prevents the build up of large amounts of data if you are running the simulation many times. If this option is not set, a new directory will be created each time for the SCATSIM processes, and all data will be permanently saved, regardless of the success of the simulation run.

Option → Move Data on Shutdown

This option is available only if the **Run in scatsim_temp** option is selected. If the data is saved to a central temporary file, this option causes it to be move to a safe place for storage on successful completion of a simulation run.

Option → System Time → Set Time on Connect

This option causes the Plugin to set the System time on the computer to equal the time in the simulation. This may assist in running SCATS in a repeatable way, as some actions within SCATS may use the time of day on the system clock rather than the simulation time transmitted from the simulator via SimHub and WinTraff.

Option → System Time → Set Date and Time on Connect

This option is as above, but also sets the system date. be careful using this option, as it may cause your software licenses to expire prematurely if dongles or other devices check for the system clock being set back to a previous date.

Creating an ODBC Data source for SCMS.mdb

1. Create a directory where the working copy of your SCMS database will be kept for SCATSIM operation. The default location is C:/Azalient/SCATSIM/central/SCMS.mdb
2. Copy the SCMS.mdb file to this location. This copy will be modified by SCATSIM as it runs. A fresh copy will be made on each run and moved to this location.
3. From the Windows Start Menu, select
   Administrative Tools → Data Sources (ODBC)
4. Select the **System DSN** tab
5. Select **Add..**.
6. Select **Microsoft Access Driver (*.mdb)** / **Finish**
7. Type **SCMS** for Data Source Name
8. Type **SCATS Central Manager Service** for Description
9. In the Database panel, press **Select...**
10. Browse to the working copy of SCMS.mdb created in step 2
11. Press OK on the ODBC Microsoft Access Setup window
12. press OK on the ODBC Data Source Administrator window