Paramics Technical Report

Car-Following, Lane-Changing and Junction Modelling

Gordon Duncan

Introduction

The Paramics car-following and lane-changing model has been developed over a period of 5 years from 1992 to 1997. It is loosely based on a number of other models, principally [1], but in most respects it was created from scratch, with the primary objectives being to demonstrate validity from two points of view:

- using iterative simulation it should show a close correlation to an array of observed numerical data for urban and inter-urban roads in the UK (objective validation)
- using computer graphics it should show a close correlation to visual observations, both on video and "in the mind's eye" (subjective validation)

This report will describe key parts of the Paramics model, but for reasons of commercial confidentiality may not reveal every detail.

In addition to the car-following and lane-changing models, the method of modelling turning vehicles on road intersections is described. Under congested conditions, effective modelling of all types of intersections – including priority junctions, signalised junctions and roundabouts, as well as grade-separated intersections – is vital to the accuracy of a simulation model, as congestion almost always starts at an intersection and then blocks back onto its inward links.

Contents

1 Car-Following: Target Headway

Each Driver-Vehicle Unit (DVU) in the Paramics simulation has a target headway h. The mean value for target headway is typically around 1 second, and this is the default, but other values can be set by the user.

The target headway for each DVU varies around the mean depending upon the value of certain parameters assigned to the DVU, as shown in Table 1.

Condition / Parameter	Variation Factor
Vehicle type: e.g. Heavy goods vehicles	1.8
Presence of single-lane highway (no lane changing possible, as in road-works),	1.5
Weather, Lighting (fog, rain, darkness) (unvalidated, just for example)	1.1 - 3.0
Close to motorway merge (accept smaller headway for limited time)	0.5
Close to traffic signals: straight ahead	0.5
Close to traffic signals: turning left	1.1
Aggressiveness (everywhere)	
$Ag \leq 4$	$\frac{6-Ag}{2}$
$Ag \geq 4$	$\frac{12-Ag}{8}$
Awareness(Near Lane-Drop)	
$Aw \leq 4$	$\frac{Aw+4}{8}$
$Aw \geq 4$	$Aw - 3$

Table 1: Variation of Target Headway Parameter

In terms of driver behaviour, a high aggression value will cause drivers to accept a smaller headway. Similarly, a high awareness value will effect the use of a longer headway when approaching a lane drop in order to allow DVUs in other lanes to merge more easily.

If not constrained by an approaching junction, a DVU varies its speed so as to attain its target headway. The reaction time of the driver is modelled by basing the calculation of the necessary acceleration on the speed at which the DVU in front was travelling at some time in the past. In Paramics, a mean reaction time of one second is used, and this is modelled by giving each DVU a short memory, so that it carries with it not only its current speed and position, but also a record of its speed and position at a number of points in its past. More accuracy is achieved by retaining a greater number of historical values for speed and position, but as is often the case there is a trade-off between accuracy and performance. Three historical values are used in the calibrated model.

The introduction of a reaction time results in the effective simulation of backward travelling shock waves. In tests, these shock waves were found to travel upstream at approximately 11 km/h.

Intelligent Cruise Control (also referred to as Automated Highway System) can be simulated by reducing the reaction time to a very small value, almost zero. When this is done, shock waves are no longer visible, and total throughput on congested roads is greatly increased.

2 Car-Following: Modes of Acceleration

A DVU in Paramics changes its speed according to its perception of the speed of the DVU ahead. These changes are normally smooth, following linear functions, but may be abrupt following the detection of one of two binary signals. These signals are visible brake lights and perceptible acceleration (of the DVU immediately ahead). There are therefore three modes of following within the Paramics model, referred to as *braking, cruising* and *acceleration* modes.

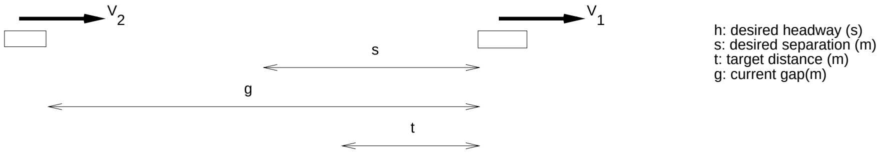

Figure 1: Diagrammatic representation of target point for car-following algorithm

For all modes of following, the concept of a target point is used. This point is based on a position at a distance s behind the leading DVU, calculated on the desired headway and the current perceived speed of that DVU, as shown in Figure 1. Note that the dimension of h is seconds, and the dimensions of s, t and g are metres. Note also that the current perceived speed is the *actual* speed at some time in the past, due to the influence of the reaction-time modelling.

The distance s is calculated as follows:

$$s = h\Delta V$$

where

$$\Delta V = V_1 - V_2$$

However, in order to pull DVUs together at a faster rate than would be the case with linear acceleration alone, the target point position is calculated by:

$$t = \frac{s^2}{g}$$

where g is the current distance between the DVUs.

In addition to the use of a target point, a *bunching* acceleration, c, is also used to bring DVUs together rapidly.

$$c = k_1 \frac{g - 2.0}{g}$$

where $k_1 = 1.0s^{-2}$. Clearly this term decreases to zero as DVUs close in toward the minimum separation of 2 metres.

2.1 Cruising Modes

There are five discrete areas A, B, C, D and E in the headway / velocity-difference phase space, as shown in Figure 2. Each of these regions has a separate expression for acceleration, expressed as a_A to a_E, and derived below. Of these five, three correspond to conditions where the DVU ahead is cruising:

- In region A, the following DVU has overshot the target point (the headway is less than the desired value), and an attempt is made to achieve the desired speed as quickly as possible, i. e. as fast as the physical constraints of the DVU allow.
- In region B, the leading DVU is pulling away from the following DVU
- In region C, the DVUs are at a constant separation or coming together.

The base acceleration values (with dimension ms^{-2}) for each of these regions are as follows:

$$a_A = k_2 \Delta V$$

$$a_B = k_2 \Delta V + k_1 \frac{g - t}{t}$$

$$a_C = c - \frac{(\Delta V)^2}{g - t}$$

where $k_1 = 1.0 s^{-2}$ as before, and $k_2 = 1.0 s^{-1}$.

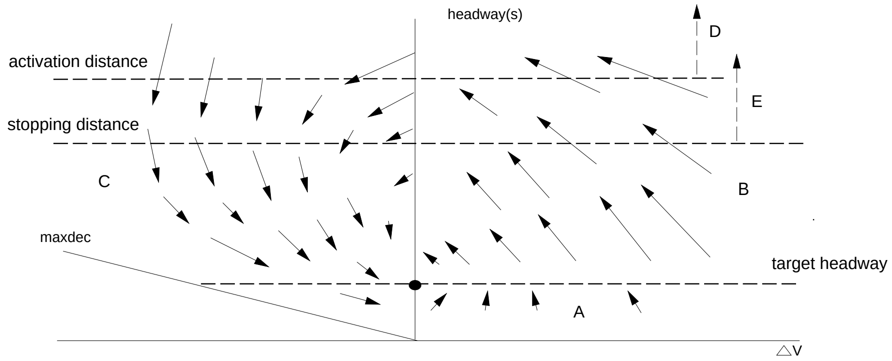

Figure 2: Pictorial representation of headway / velocity-difference phase space, and the resultant acceleration values produced by the Paramics car-following algorithm. The magnitude of an arrow is a representation of the acceleration derived by the car-following algorithm for a given $(\Delta V, h)$ pair. Note that the arrows are not to scale.

2.2 Braking Mode

When the DVU ahead is perceived to be braking (its deceleration is greater than a certain threshold), its perceived speed is decreased by an amount dependent on its maximum deceleration rate. This models a driver's expectation that if the DVU ahead is braking, its speed in the next time step will be considerably less than at the current time-step. The method of application of speed difference and current separation to acceleration ensures that a DVU will over-compensate if the DVU ahead is braking, and that this over-compensation will increase as the distance between the DVUs decreases. This, and the time-lag introduced by modelling reaction time results in shock-wave characteristics as seen typically in highway traffic flow.

However, because the speed of the DVU ahead is predicted, and may have a resultant value of zero, a threshold is used to test whether the following DVU is close enough to be in danger of collision. if not, the acceleration is set to a positive value. This corresponds to region D in Figure 2. So

$$a_D = k_3$$

where $k_3 = 1.0 ms^{-2}$

2.3 Acceleration Mode

If the DVU ahead is perceived to be accelerating at a high rate, and is more than the following DVU's safe stopping distance away, acceleration is set to the maximum value. This corresponds to region E.

$$a_E = a_{MAX}$$

3 Lane Changing

Lane changing in the Paramics model is done using two devices:

- a gap-acceptance policy
- a historical record of suitable gap availability

The gap acceptance policy is linked with the car-following model, in that the accepted gap is based on the target headway. In the following description, the term *target lane* will be used to describe the lane that a DVU is aiming to move into from its current lane.

If traffic in both the current lane and the target lane For a DVU is moving at a constant speed, a gap must exist both in front and behind the position it would occupy that is at least as large as the target headway for that DVU. If there is a speed difference between the two lanes, this expression is extended to take into account the time it would take for the DVU to attain the speed of the DVU it would follow in the target lane.

So, if DVU_0 is the DVU under consideration, and DVU_1 and DVU_1 are the vehicles in the target lane that are ahead and behind of the position DVU_0 would occupy, then both of the following must be true for the lane changing manoeuvre to take place:

$$g_1 > d_{\Delta V_1} + hv_1 \qquad\qquad g_2 > d_{\Delta V_2} + hv_2$$

where:

$$\begin{aligned} d_{\Delta V_1} &= t_{r_0} + \frac{\Delta V_1}{D_0} \\ d_{\Delta V_2} &= t_{r_0} + \frac{\Delta V_2}{D_2} \end{aligned}$$

$$\begin{aligned} \Delta V_1 &= v_1 - v_0 \\ \Delta V_2 &= v_0 - v_2 \end{aligned}$$

v_N is the current speed of DVU_N
D_N is the maximum deceleration (braking rate) of DVU_N

g_1 is the gap between the back of DVU_1 and the front of DVU_0
g_2 is the gap between the back of DVU_0 and the front of DVU_2

If both these conditions are true continuously for a period of T_{LC} seconds then the lane changing manoeuvre will take place. The value of T_{LC} varies depending upon behaviour and location parameters, but is typically in the range 3-6 seconds.

4 Intersection modelling

Simulation of DVUs on straight or curved network links in Paramics is carried out essentially in one dimension only, i.e. by their distance (and speed, and acceleration) along the link. At road junctions or intersections there is a need for a much greater detail of modelling, and this section describes a method of modelling vehicles in 2 dimensions in intersection areas.

Modelling vehicles in two dimensions is obviously much more complex and compute-intensive, with calls to trigonometric functions such as $cos(x)$ $sin(x)$ and $tan(x)$. However, in a congested network the total number of vehicles on a junction in any time-step will only ever be a small fraction of the total number of vehicles being modelled.

Paramics uses located unit vectors (directed points) to describe a junction. That is, a triple (x,y, bearing), to describe not only the position of a point to which a vehicle must head for any particular exit from a junction, but also the required angle of orientation once it gets there. Paramics employs an algorithm that defines a general purpose method to steer a vehicle over a realistic path between its current position to any target position, taking angles of orientation and steering limits into account. Complex junctions can be described using such vectors, in a concise but powerful way. The method is much easier than, for example, specifying centre-points for turning arcs, as only one vector is needed for each exit lane from a junction, irrespective of the entry point to the junction. Also, the nature of the curve produced by the iterative algorithm is comparatively much more flexible than a simple arc.

The method is based on calculating iteratively the position of an imaginary point at which the vehicle should aim at the next time-step in order to arrive at the ultimate target point at the correct angle of approach. The rate of change of bearing (i.e. the radius of the turning circle) is regulated by both the physical attributes of the vehicle and its current speed.

4.1 Constraints on Turning Vehicles

A turning vehicle is constrained in the maximum angle it can turn though in a fixed time period by both the limit of its steering mechanism (full lock) and the friction of its tyres on the road surface given that it is travelling at a certain forward speed. Both of these constraints are considered here.

4.2 Tyre Friction Constraint

The speed of turning round a corner is constrained by

$$v_{max} < \sqrt{f_s g r} \tag{1}$$

where g is 9.81 m / s / s, f is the tyre friction co-efficient, typically between 0.2 and 0.4 and r is the radius of the turn.

Suppose that with a constant steering angle, the time taken to move from A to B is Δt. (See Figure 3.)

In that time, travelling at the maximum speed v_{max} at which the tyres will hold on a curve of radius r, the distance travelled is clearly $v_{max}\Delta t$. Thus the angle of rotation of the vehicle, ϕ, in this limiting situation is

$$\phi = arcsin(\frac{v_{max}\Delta t}{r}) \tag{2}$$

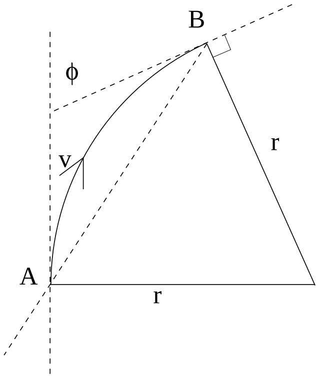

Figure 3: Transition from A to B, with fixed speed and steering angle

Substituting for r from (1) gives:

$$\phi = arcsin(\frac{f_s g \Delta t}{v_{max}})$$

Thus in the general case:

$$\phi < arcsin(\frac{f_s g \Delta t}{v}) \quad (3)$$

4.3 Minimum Turning Circle Constraint

If the steering wheels of a vehicle are offset by an angle of θ from the straight-ahead position, the radius of the turning circle, r can be found approximately from

$$r = \frac{L}{\sin\theta} \quad (4)$$

where L is the length of the vehicle's wheelbase.

If the maximum steer angle is θ_{LIMIT}, then from (4), the radius of the minimum turning circle is given by:

$$r_{min} = \frac{L}{\sin\theta_{LIMIT}} \quad (5)$$

Suppose a vehicle is turning with its minimum turning circle, at speed v. In a time Δt, it is clear that

$$\phi_{max} = arcsin(\frac{v\Delta t}{r_{min}})$$

This defines the second constraint on ϕ, the general angle turned in time Δt, when travelling at speed v

$$\phi < arcsin(\frac{\Delta t \sin\theta_{limit}}{L} v) \quad (6)$$

The combination of the two constraints is shown on Figure 4. The allowable values for θ are shown in the shaded part.

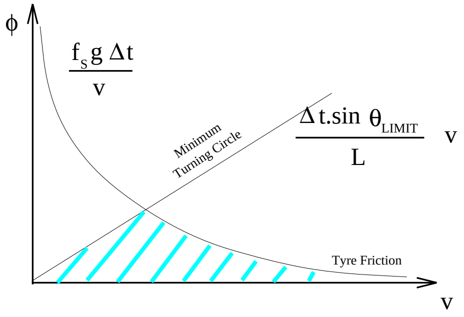

Figure 4: Constraints on the steering angle of a vehicle

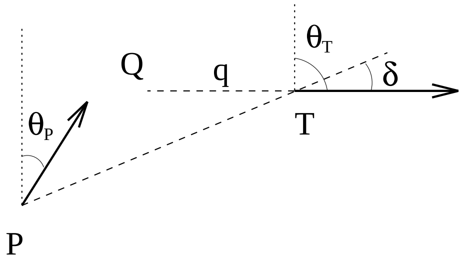

Figure 5: Steering Heuristic for Turns less than Quarter-Circle

4.4 Steering Heuristic

Suppose a vehicle is at a point P with bearing θ_P, and requires to navigate itself to a second point T, and arrive there on bearing θ_T. (see Figure 5).

With the steering angle constrained according to physical limits and the current speed of the vehicle, as described previously, the heuristic attempts to aim the vehicle at an imaginary point, Q, projected along a line through T at a bearing of $-\theta_T$. The distance of this Q-point along the back-projected line, q, is proportional to $\sin\delta$, where δ is the difference in bearing between a direct approach from the current position, and the desired heading θ_T. The distance q is also proportional to the distance between P and T.

Thus:

$$q = PT \sin\delta$$

This technique results in a smooth trajectory for most combinations of initial and target positions and bearings. However, if the points are placed such that the change in bearing exceeds a certain threshold ($\delta > \pi/4$), the heuristic is slightly modified. To allow for a “swing out” the Q-point is offset a perpendicular distance r_{min} from the back-projected line. As soon as the change in bearing becomes less than the threshold, the Q-point is relocated to the back projected line. This swinging-out effect is not enabled on some types of junction, because of lack of space, and instead the DVU will head directly for the aim point and turn to the correct bearing in a circle with radius r_{min}.

Typical output from the steering heuristic is shown in Figure 7 which illustrates the paths taken for a vehicle from a fixed initial point to a fixed destination point, but with a variety of required bearings at the destination point.

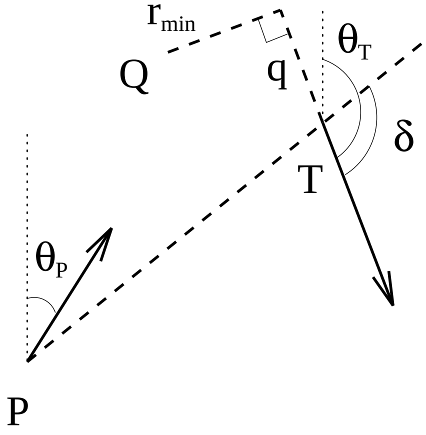

Figure 6: Steering Heuristic for Turns more than Quarter-Circle

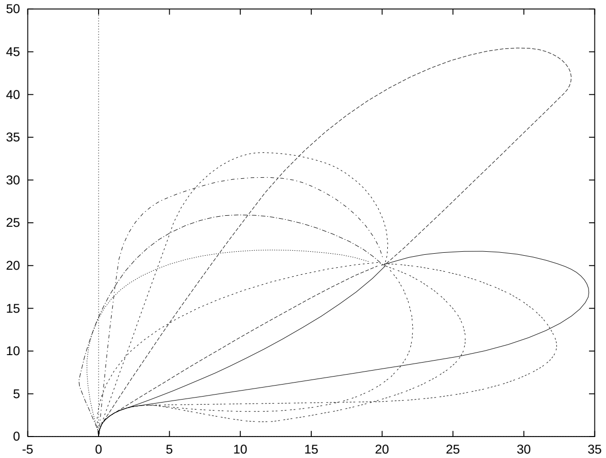

Figure 7: Trajectories for a vehicle for varying destination bearings

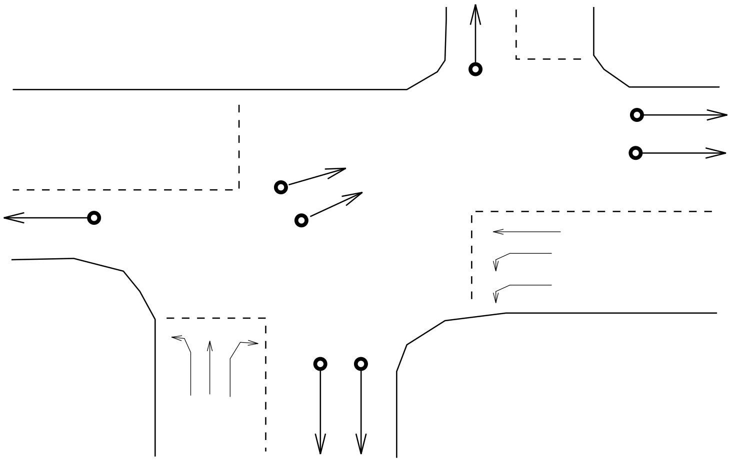

Figure 8: Example Junction Description showing Control Points

4.5 Junction Description

Each junction is described by a set of locus points. A vehicle entering a junction, must steer itself smoothly through at least one, but possibly a sequence, of these control points. Each exit lane from a junction is described by a single point in the centre of the lane at the point at which it becomes straight. Additional points may be used to guide traffic around islands, or other physical constraints. An example junction is shown in Figure 8.

References

[1] Hans-Thomas Fritzsche, "A model for traffic simulation," *Traffic Engineering and Control*, pp. 317–321, May 1994.