ISSN 0280–5316
ISRN LUTFD2/TFRT--5609--SE

Stop & Go Controller for Adaptive Cruise Control

Mikael Persson

Department of Automatic Control
Lund Institute of Technology
December 1998

Department of Automatic Control
Lund Institute of Technology
Box 118
S-221 00 Lund Sweden

Document name
Master Thesis

Date of issue
December 1998

Document Number
ISRN LUTFD2/TFRT--5609--SE

Author(s)
Mikael Persson

Supervisor
Fredrik Botling and Erik Hesslo Volvo TD, Rolf Johansson LTH

Sponsoring organisation

Title and subtitle
Stop & Go Controller for Adaptive Cruise Control. (Regulator för start och stopp till adaptiva farthållare.)

Abstract

In the field of vehicle, conventional cruise control systems have been available on the market for many years. During the last years, modern cars include more and more electronical systems. Since computers can quite easily be programmed with powerful software, it opens a new dimension to what services a car can provide to a driver. One of those new services is Adaptive Cruise Control (ACC) (or Autonomous Intelligent Cruise Control (AICC)), which extends the conventional cruise control system to include automated car following when the preceding car is driving at a lower speed than the desired set-speed. The focus of ACC has mainly been on highway traffic (high-speed), but to improve the comfort to the driver also low-speed situations must be considered. This thesis presents an ACC system that is capable of car following in low-speed situations, e.g. in suburban areas, as well as in high-speed situations. The system is implemented in a test car and the result is evaluated.

Key words
Adaptive Cruise Control, Vehicle Control, Car Model, Driver Model

Classification system and/or index terms (if any)

Supplementary bibliographical information

ISSN and key title
0280–5316

ISBN

Language
English

Number of pages
46

Recipient's notes

Security classification

Department of Automatic Control, Lund Institute of Technology
Volvo Technological Development Corporation, Gothenburg

VOLVO

Master Thesis

Stop & Go Controller for Adaptive Cruise Control

981207

Author:
Mikael Persson

Keywords: Adaptive Cruise Control, Vehicle Control, Car Model, Driver Model

Abstract

In the field of vehicle control, conventional cruise control systems have been available on the market for many years. During the last years, modern cars include more and more electronical systems. These systems are often governed by a computer or a network of computers. Since computers can quite easily be programmed with powerful software, it opens a new dimension to what services a car can provide to a driver. One of those new services is Adaptive Cruise Control (ACC) (or Autonomous Intelligent Cruise Control (AICC)), which extends the conventional cruise control system to include automated car following when the preceding car is driving at a lower speed than the desired set-speed. The focus of ACC has mainly been on highway traffic (high-speed), but to improve the comfort to the driver also low-speed situations must be considered. This thesis presents an ACC system that is capable of car following in low-speed situations, e.g. in suburban areas, as well as in high-speed situations. The system is implemented in a test car and the result is evaluated.

Sammanfattning

Dagens bilar innehåller allt fler elektroniska komponenter. En dator, eller oftare ett nätverk av datorer, styr och övervakar de olika systemen i bilen. På så sätt kan till exempel en effektivare bränsleutnyttjning erhållas jämfört med ett helt mekaniskt system (gamla bilar). Styrkan ligger i att datorer kan enkelt programmeras med avancerad mjukvara. Detta har lett till att också helt nya system har utvecklats som förr var mer eller mindre omöjliga att genomföra.

Ett av dessa nya system är adaptiva farthållare (eng. Adaptive Cruise Control, ACC), som är en utvidgning av den konventionella farthållaren. En konventionell farthållare kan automatiskt hålla en förinställd hastighet vald av föraren. Denna hastighet hålls oberoende av omgivningen, t.ex. andra fordon på vägen. Om framförvarande fordon kör långsammare än den förinställda hastigheten måste föraren för eller senare ta över kommandot över bilen och antingen bromsa eller köra om fordonet. En adaptiv farthållare skulle däremot automatiskt sänkt hastigheten på bilen och sedan hållt ett visst lämpligt avstånd till fordonet framför. Systemet skulle alltså byta läge från att hålla en viss hastighet till att hålla ett visst avstånd. Om det framförvarande fordonet senare skulle accelerera, skulle den adaptiva farthållaren också ha ökat hastigheten, förutsatt att den inte överstigit den förinställda hastigheten. Om det inte finns något framförvarande fordon fungerar den adaptiva farthållaren precis som en konventionell farthållare.

För att den adaptiva farthållaren ska fungera måste framförvarande fordons hastighet och avståndet mellan fordonen vara kända. Denna information erhålls vanligtvis genom en frontmonterad radar. Annan utrustning som krävs är elektronisk gas och broms (möjliggör elektronisk styrning) samt en dator som styr och övervakar det hela.

Den adaptiva farthållaren kan delas upp i två separata problem. Det första problemet är att styra gas och broms så att önskad acceleration eller retardation uppstår. Det andra problemet är att på något sätt modellera en förares beteende givet den situation bilen befinner sig i. Denna modell ska utmynna i förarens önskade acceleration på bilen. Resultatet kombineras sedan ihop med lösningen på det förra problemet.

De första systemen på adaptiva farthållare finns idag redan i kommersiellt bruk. Dessa system är gjorda för motorvägstrafik, dvs vid höga hastigheter, och är ganska enkla. Motorvägstrafik kännetecknas av att de nödvändiga reaktionerna (gasa eller bromsa) är ofta relativt små. Vid lägre hastigheter (ej motorvägstrafik) blir det mer komplicerat eftersom de nödvändiga reaktionerna blir kraftigare. Speciellt måste förarmodellen vara mer detaljerad.

Detta examensarbete presenterar ett system som klarar av hastigheter allt från stillastående till motorvägstrafik. Systemet testas också i praktiken och resultatet utvärderas.

Table of Contents

1. INTRODUCTION **6**

1.1 DIFFERENCES BETWEEN HIGH-SPEED TRAFFIC (HIGHWAY) AND LOW-SPEED TRAFFIC (SUBURBAN AREAS) 6

2. BASIC CONCEPTS **7**

3. SYSTEM OVERVIEW **8**

3.1 THE COORDINATION 8
3.2 THE ACC SYSTEM 8

4. CAR MODELLING **10**

4.1 DATA COLLECTION 10
4.2 STATIC ANALYSIS 10
4.2.1 Correlation between acceleration and brake pressure *10*
4.2.2 Correlation between acceleration and throttle position *11*
4.3 DYNAMIC ANALYSIS 12
4.4 DATA EXAMINATION 12
4.5 MODEL STRUCTURE 13
4.6 MODEL ESTIMATION 13
4.7 VALIDATION 14
4.8 CONCLUSIONS 16

5. ACTUATOR CONTROL **17**

5.1 DESIGN OVERVIEW 17
5.2 THE NOISE 18
5.3 CONTROLLER DESIGN 18
5.3.1 Gain scheduling *19*
5.3.2 Modifications *23*
5.4 IMPLEMENTATION 25
5.5 SUMMARY 26
5.6 CONCLUSIONS 27

6. DRIVER MODELLING **28**

6.1 DATA COLLECTION 28
6.2 LINEAR REGRESSION 28
6.3 DANGEROUS SITUATIONS 29
6.3.1 Small headways *29*
6.3.2 Time-To-Collision (TTC) *30*
6.3.3 Minimum deceleration to avoid collision (a_{min}) *30*
6.3.4 Modifications *31*
6.4 APPROACHING SITUATIONS (NEGATIVE ΔV) 31
6.4.1 Non-dangerous situations *33*
6.5 SEPARATING SITUATIONS (POSITIVE ΔV) 35
6.6 UPPER AND LOWER BOUNDS 36
6.7 CONCLUSIONS 37

7. IMPLEMENTATION AND VALIDATION **38**

7.1 TYPICAL SITUATIONS 38
7.2 SPECIAL SITUATIONS 39
7.2.1 Approach situations *39*
7.2.2 Dangerous situations *41*
7.3 CONCLUSIONS 44

8. ACKNOWLEDGEMENTS **45**

9. REFERENCES **45**

APPENDIX A: THE TEST CAR **46**

1. Introduction

In a conventional cruise control system, the driver can set a desired speed and the car will maintain this speed as soon as it has been established. This is done independently of the environment, e.g. other vehicles on the road. When the vehicle ahead is travelling slower than the desired speed, the driver must at some point intervene with the brake pedal to avoid a collision. Alternatively, he must overtake the vehicle. The ACC concept extends the conventional cruise control system to include car following. In the scenario above, the ACC would have automatically lowered the speed of the car to match the speed of the vehicle ahead and to maintain an appropriate distance. If the preceding vehicle would have later on increased its speed, the ACC system would have automatically increased the speed (thereby following the car), unless it becomes greater than the desired speed set by the driver.

The ACC concept is currently being introduced by several car manufactures in their latest car models. These systems consist of a sensor, mounted in the front of the car, that measures the preceding vehicle's velocity and distance. The sensor could either be of optical or radar type, but the radar sensor is often preferred since it is much less influenced by the weather conditions than the optical sensor. The sensor information is transmitted to an ACC controller (a computer) that controls the engine and brake systems. The first generation of ACC will only allow gentle acceleration and deceleration. A major reason for this is that the driver should never be surprised by the actions of the ACC system and the driver should always be able to intervene if the system does not comply with the driver's intentions. These systems will only work in highway traffic, where the needed speed changes are moderate. The second generation of ACC will allow a greater acceleration and deceleration, which is necessary in suburban areas where the speed and distance is relative low, but the relative speed on the other hand can be rather high.

It is important to remember that the ACC is only a service to help the driver, <u>not</u> a replacement of the driver. The driver is still in charge of the car at any moment, regardless if the ACC system is active or not.

1.1 Differences between high-speed traffic (highway) and low-speed traffic (suburban areas)

In highway traffic, the vehicles are travelling at a high speed, but they are all travelling at approximately the same speed, i.e. the relative speed is small. The vehicles are only accelerating and decelerating in gently fashions, i.e. all speed changes are moderate and are performed rather slowly. Hence, the traffic flow is quite smooth on highways.

In suburban areas, on the contrary, the vehicles are travelling at a relative small speed, but the relative speed between two vehicles can be quite high. For example, one car can be travelling at 20 m/s, whereas the car ahead is almost not moving at all (perhaps because of a queue). Another major difference is that the vehicles are often stopping, e.g. due to a red light, and then accelerating to a modest speed, but after a short time stopping again. The traffic flow is therefore much less smooth than in the highway situation. The vehicles have to shift gears quite frequently, in contrast to the highway case, where shifting gears is almost not necessarily at all.

A difficulty that arises in suburban areas is that it might be hard to find a model of the car dynamics that is good enough. Since the controller must be able to accelerate and decelerate quite hard, it must be carefully designed so it does not behave in an inconvenient manner, i.e. uncomfortable to the driver. There can be only small overshoots and even if the acceleration is hard it must be smooth. Hence, the design of the controller requires a good model of the car dynamics. But even if a good model is known, it still might be hard to design a suitable controller.

A more practical difficulty is that the radar must have good resolution, even at small distances. Typically, the radar must be able to measure targets in ranges of 1-150 m with an error smaller than 0.5 m. The relative speed must also be measured with high resolution, typically with an error less than 1.0 m/s.

2. Basic concepts

This section describes some basic concepts that are used throughout this text.

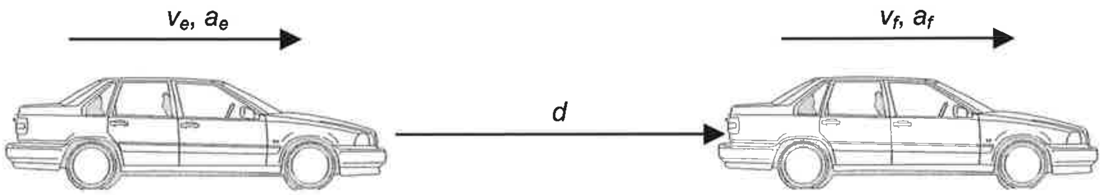

Figure 2.1: A car is following another car at a certain distance.

Figure 2.1 shows a situation where one car is following another car. The speed of the preceding and following vehicle are denoted v_f respectively v_e. The corresponding accelerations are denoted a_f and a_e and the distance between the vehicles is denoted d. The relative speed is defined as:

$$\Delta v = v_f - v_e$$

The following relation then holds:

$$\Delta v = \frac{d}{dt} d$$

i.e. the relative speed is the time derivative of the distance. The following vehicle wants to follow the vehicle ahead at a preferred distance. This distance is denoted d_{set}. If only v or a is written, i.e. without suffix, it is equivalent to v_e respectively a_e.

3. System overview

3.1 The coordination

The ACC system must be coordinated with the ordinary cruise control system in some way. Both systems generate a desired acceleration of the vehicle, but only one value can be passed on to the accelerator controller. This could be solved in many different ways. One solution is to switch between the two systems when some preferred situation has occurred, e.g. when d and Δv equals some pre-defined values. The solution here uses another approach. Both systems are always running in parallel. If a target vehicle is existing, the overall desired acceleration (a_{set}) equal the minimum of the desired accelerations generated by the two systems. If there is no target ahead, the overall desired acceleration equal the desired acceleration generated by the conventional cruise control (a_{CCC}). Figure 3.1 shows this coordination.

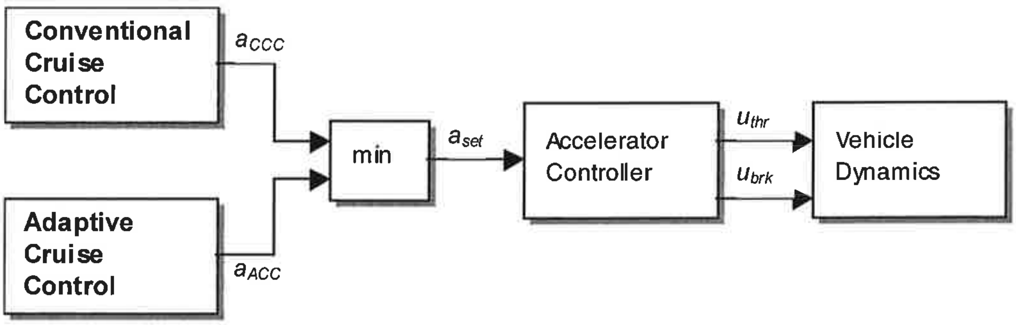

Figure 3.1: The coordination of the conventional and the adaptive cruise control system. The minimum acceleration is chosen.

There will always be smooth transitions when a_{set} changes from a_{CCC} to a_{ACC} or vice versa. This solution has proved to work very well. Further on in this text, only the car following regime (ACC) will be considered. Therefore, it will always be assumed that $a_{set}=a_{ACC}$.

3.2 The ACC system

The ACC system should control the car in a safe and comfortable way and ensure that the desired distance to the vehicle ahead is maintained. The overall control problem can be divided into two separate parts, see figure 3.2.

The outer control loop generates a desired acceleration (a_{set}), given the speed of the car (v_e) and the distance (d) and relative speed (Δv) to the vehicle ahead. It is assumed that the desired acceleration is a static function of these inputs, i.e. it includes no dynamics. Several investigations supports this assumption (see e.g. [1]). The inner control loop should control the brake pressure (by using u_{brk}) and the throttle position (by using u_{thr}) in such way that the desired acceleration is obtained quickly and with little overshoot.

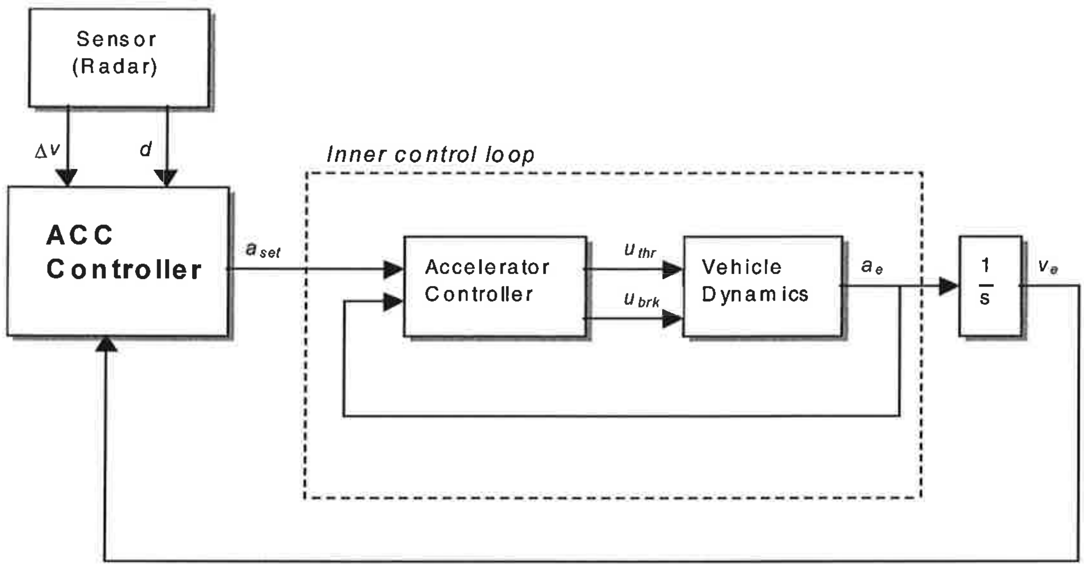

Figure 3.2: The overall control problem is divided into two separate parts.

The major advantage with this separation into two loops is that the loops are independent of each other in the following sense: The outer loop represents the driver behaviour and is independent of the vehicle to be controlled. The inner loop, on the other hand, is highly dependent of the vehicle dynamics, but is independent of the driver behaviour (i.e. the outer loop). This separation makes it possible to change the algorithm in the outer loop without having to change the algorithm in the inner loop. If a car with different dynamics is to be controlled, only the inner loop has to be changed. It is though important that the inner loop is well designed, i.e. that the desired acceleration is obtained quickly and with little overshoot. Otherwise, the outer loop cannot be evaluated in a fair way, since the overall performance will be different compared to using a well-designed inner loop. If the performance is not satisfactory, one might draw the conclusion that the outer loop is poorly designed even if it is well designed. If the overall performance is satisfactory even though the inner loop is poorly designed, the outer loop will be vehicle dependent and has to be changed if a different vehicle is to be controlled, which is an undesirable situation.

4. Car Modelling

To be able to design a good accelerator controller, a good model of the car behaviour must be determined. The true car dynamics is unfortunately complex and includes many non-linearities. Experiments show for instance that the static gain is dependent of the current gear. Unfortunately was direct gear data not available information. The relevant car behaviour is how the throttle angle and the brake pressure influence the acceleration of the car. The throttle angle and the brake pressure can be regarded as inputs and the acceleration of the car as the output of a system.

4.1 Data collection

Data was initially collected by simply driving around in an urban area (in this case, in Gothenburg). Stop and go situations as well as cruising situations were considered. During all time, the speed was fairly low (not above 70 km/h), i.e. in the speed range of the future operation. The reason why many situations were considered was that the goal was to find an average model that describes the behaviour fairly well in the whole range of operation. Data was collected with a sampling period of h=0.10 s.

The speed of the car will further be noted v, the acceleration with a, the throttle position with u_{thr} and the brake pressure with u_{brk}.
Note: u_{thr} is more correctly speaking the control signal to the servo that manages the throttle, not the throttle position itself. But since the true throttle position is practically proportional to u_{thr}, it is often viewed as the throttle position in a different scale. u_{brk} is also in reality different from the real brake pressure, since u_{brk} is actually the set-value to the braking system. But several experiments have shown that in practice the difference is minor.

4.2 Static analysis

4.2.1 Correlation between acceleration and brake pressure

The first analysis of the data was a simple correlation analysis.

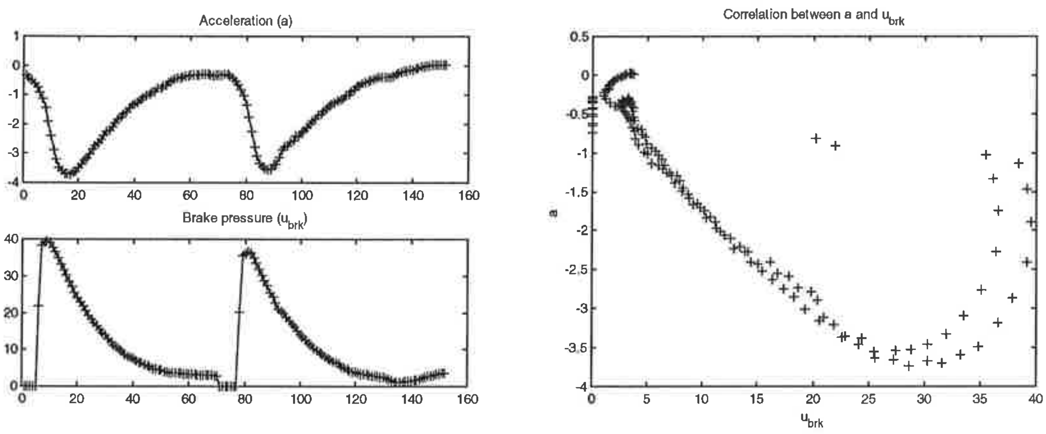

Figure 4.1: The static correlation between u_{brk} and a.

Assume that the following relation holds:

$$a = k_{brk} u_{brk}$$

where k_{brk} is a constant. The leftmost diagram in figure 4.1 shows some data from a brake experiment. The acceleration was filtered with a zero-phase discrete filter to attenuate the noise. As seen, the car was braked twice. **Note**: The data shown in the diagram is actually a concatenation of two subparts of the original data. In reality, the car had a speed of 50 km/h initially in both cases and a speed of 10 km/h when the braking ended.

The rightmost diagram in figure 4.1 depicts the correspondence between the acceleration and the brake pressure. After an initial transient (corresponding mainly to the data points where the pressure is above 25), the acceleration is as good as proportional to the brake pressure. The proportional constant is approximately:

$$k_{brk} = -0.14$$

4.2.2 Correlation between acceleration and throttle position

The same investigation can be done with the throttle position, i.e. assume that the following relation holds:

$$a = k_{thr} u_{thr}$$

The acceleration and throttle position data is shown in figure 4.2.

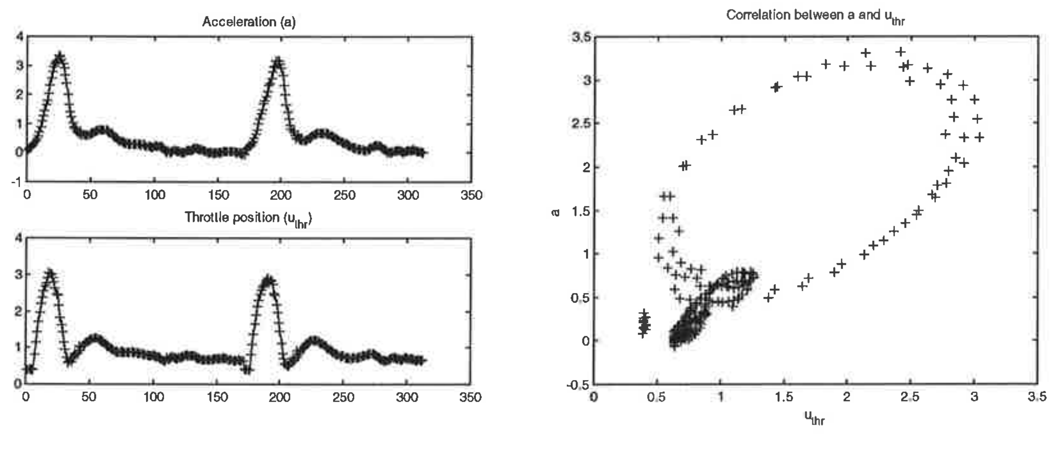

Figure 4.2: The static correlation between u_{thr} and a.

The acceleration is also here filtered with a discrete zero-phase filter. Similarly, the car was accelerating twice in this data, which in turn also is a subpart of a larger original data. In both acceleration cases, the car had initially a speed of 10 km/h and a speed of 50 km/h when the acceleration ended.

The rightmost diagram in figure 4.2 depicts the fact that there is not a simple linear static relationship between the throttle position and the acceleration.

The major difference between the throttle position signal and the brake pressure signal, is that the brake pressure affects the wheels almost directly, whereas the throttle position only affects an air stream. This air stream affects in turn an engine combustion, which in turn affects a transmission system, which finally affects the wheels (all this simplified speaking). Since these systems include dynamics, there is not a simple static linear relationship between the throttle position and the acceleration.

4.3 Dynamic analysis

As indicated in the previous section, there is dynamics involved in the relation between the inputs (u_{brk} and u_{thr}) and the output (a). Therefore, assume that the relation between the inputs and the output is described with a dynamic linear system, i.e. that the following relation holds:

$$a = H'_{thr} u_{thr} + H'_{brk} u_{brk}$$

Since v is the actual measured signal and not a and to reduce the influence of the noise, the following relation is analysed instead:

$$v = H_{thr} u_{thr} + H_{brk} u_{brk}$$

The analysis is further done in discrete time and in the forward shift operator q:

$$v(k) = H_{thr}(q) u_{thr}(k) + H_{brk}(q) u_{brk}(k)$$

4.4 Data examination

Figure 4.3 shows the signals involved in the analysis. The rightmost diagram shows the signals used to validate the model and they are therefore never used in the model derivation. The time axis is sample number in all diagrams.

There exist one problem in the original data. When the car is standing still and you want to keep it still, you have to press the brake (this was the case with the test car used, other car models could have different behaviour). In this situation, there is no deceleration even though the brake pressure is non-zero (the car cannot go backwards). To remove this “non-linearity”, u_{brk} was set to zero if the velocity equated zero.

As seen, the driver used the brake pedal more rarely than the throttle pedal. That depends partly on the fact that one can decelerate a car by either decrease the throttle position or step on the brake pedal. To accelerate you can only use the throttle pedal. It also, of course, depends on the given traffic situation and on the characteristics of the driver.

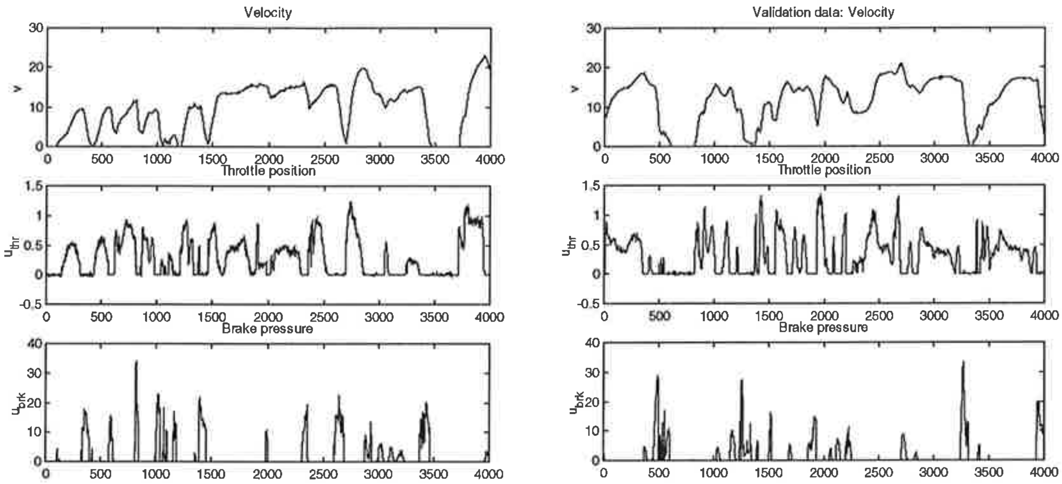

Figure 4.3: The data used in the model derivation. The rightmost diagram shows the validation data.

4.5 Model structure

The model structure chosen is very general, since little knowledge about the system was known *a priori*. The model structure is:

$$A(q)v(k) = \frac{B_{thr}(q)}{F_{thr}(q)} u_{thr}(k) + \frac{B_{brk}(q)}{F_{brk}(q)} u_{brk}(k) + \frac{C(q)}{D(q)} e(k)$$

where *e(k)* is a white noise sequence with zero mean value. The *A*-polynomial represents the dynamics that is in common to all inputs, including the noise.

4.6 Model estimation

The next step is to select appropriate orders of the polynomials in the model and then estimate the coefficients of the polynomials. The coefficients were estimated using a predictive error method. After several iterations the following model was accepted to be the most proper one:

$$(q-0.9984)v(k) = \frac{0.0279q^2}{q-0.7812} u_{thr}(k) - 0.0103q\; u_{brk}(k) + \frac{q^2+0.9548q}{q+0.9766} e(k)$$

The common polynomial *A* is essentially functioning as an integrator. The small deviation compared with a pure integrator (*q*-1), could be explained with the air resistance and other losses. The losses will reduce the speed even if all inputs equals zero. If all inputs equals zero, the following equation holds:

$$(q-0.9984)v(k) = 0 \Leftrightarrow v(k) = 0.9984v(k-1)$$

i.e. the velocity will decline (slowly) exponentially. Let

$$a(k) = \frac{v(k) - v(k-1)}{h}$$

where h is the sample period (0.10 s), be an approximation of the acceleration. The above equation could then be rewritten as:

$$a(k) = -0.017\, v(k)$$

i.e. the losses gives a deceleration that is proportional to the velocity. The entire model could also be reformulated into:

$$a(k) = \frac{0.278q}{q - 0.781} u_{thr}(k) - 0.103\, u_{brk}(k) - 0.017\, v(k)$$

Here the noise model is neglected (the noise model will not be further analysed or commented). As seen, the u_{thr} signal is low-pass filtered before it affects a, whereas the u_{brk} signal affects the acceleration directly. The proportional coefficient -0.103 can be compared with the value of k_{brk} found in the static analysis (k_{brk}=-0.14).

4.7 Validation

The model can and should be validated in several ways. One way is to look at the locations of the poles and zeros. Figure 4.4 shows the pole-zero plots for the transfer functions H_{thr} and H_{brk}, which are defined as:

$$H_{thr}(q) = \frac{B_{thr}(q)}{A(q)F_{thr}(q)}$$

$$H_{brk}(q) = \frac{B_{brk}(q)}{A(q)F_{brk}(q)}$$

As noted, there are no zeros in either pulse-transfer function. One can also see the low-pass pole in H_{thr}, which is non-existing in H_{brk}.

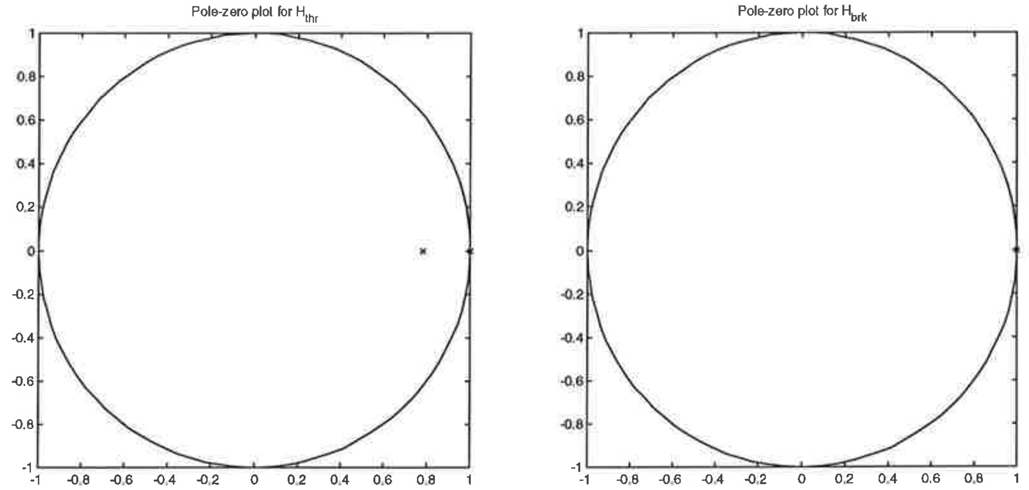

Figure 4.4: Pole-zero plots for the transfer functions from u_{thr} and u_{brk} to a. There is one low-pass pole in H_{thr}, which is non-existing in H_{brk}.

Another important way to validate the model is to use simulation. Figure 4.5 shows the simulation results when the noise model is excluded (inclusion of the noise model proved to give very little difference). The upper ones is the result when using the data used in the model derivation, and the lower ones is the result when using the validation data.

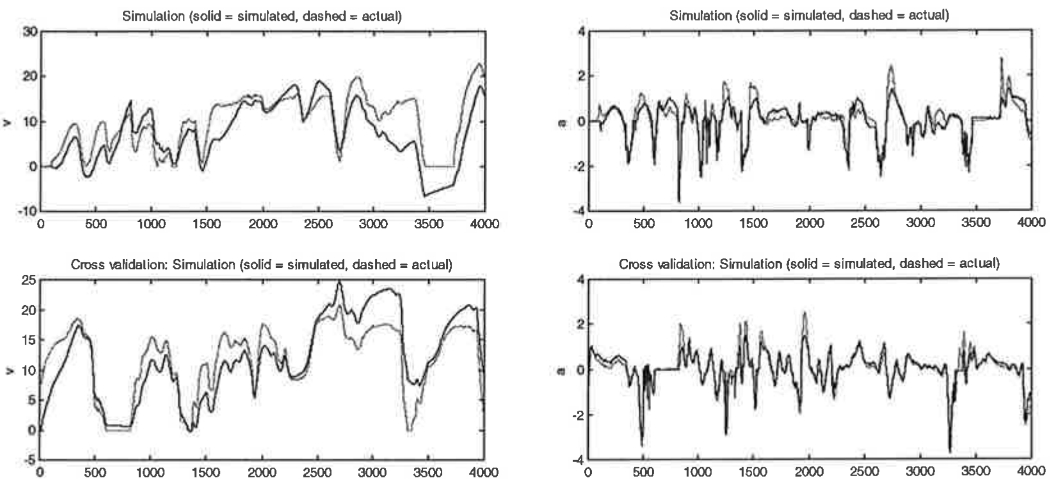

Figure 4.5: Simulations of the model. The simulated output (solid line) is compared to the actual value (dashed). The fit is better in the acceleration signals than in the velocity signals.

The leftmost diagrams depict the correspondence between the actual and the simulated velocity (v), whereas the rightmost diagrams depict the correspondence between the actual and simulated acceleration (a). Since the car model is to be used in the design of an accelerator controller, it is more important that the fit is better in the acceleration case. The correspondence is rather good, except at high accelerations.

A further analysis of the residuals from the cross validation simulation gives the results shown in figure 4.6. All diagrams give an indication that the residuals are white noise. Note that in this analysis the noise model is included.

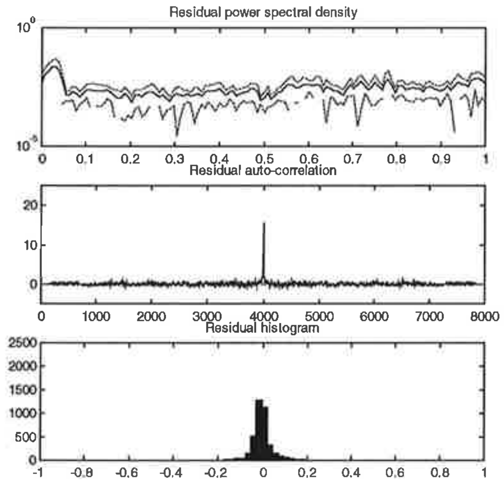

Figure 4.6: Different analysis of the residuals. The top one shows the spectral density with a 95% confidence interval. The mid one shows the autocorrelation and the bottom one shows the residual histogram.

4.8 Conclusions

The derived car model mirrors the actual system in the big picture. The model does not describe the system well when different gears are used, due to the fact that each gear has an individual static gain. This shortcoming does not have to be severe, because an integrator in the control loop could take care of this error well. The model does not either describe the non-linearities that inherently exist, e.g. that the engine torque is a non-linear function of the throttle position and the engine speed. The non-linear functioning of the converter is neither described.

5. Actuator control

The innermost control loop is the actuator control loop. The actuator controller should control the acceleration (a) of the car to match the desired acceleration (a_{set}). That should be accomplished by using two control signals: u_{brk}, which controls the brake pressure, and u_{thr}, which controls the throttle position.

5.1 Design overview

It is clear that u_{thr} and u_{brk} never should be both non-zero at the same time (you do not want to step on the gas and brake pedal at the same time). The general idea is therefore to calculate a single control signal u and then let u_{thr} and u_{brk} be a function of u. If u is positive, u_{thr} should be non-zero and u_{brk} should equal zero, and if u is negative, u_{thr} should equal zero and u_{brk} should be non-zero:

$$u_{thr} = \begin{Bmatrix} 0, & u < 0 \\ K_{thr}u, & u \geq 0 \end{Bmatrix}$$

$$u_{brk} = \begin{Bmatrix} -K_{brk}u, & u < 0 \\ 0, & u \geq 0 \end{Bmatrix}$$

K_{thr} and K_{brk} are positive constants introduced to allow scaling of u. The calculation of u could be done using ordinary control theory. Figure 5.1 shows an overview of the controller.

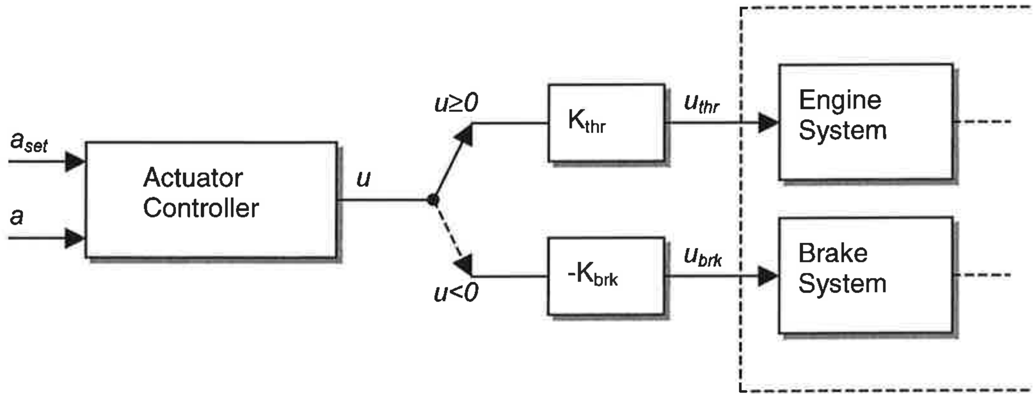

Figure 5.1: The actuator controller either controls the throttle position or the brake pressure, but not both at the same time.

5.2 The noise

Although the noise in the velocity signal (v) does not seem to be significant, the picture becomes quite different when looking at the acceleration signal. If the acceleration is approximated with:

$$a(k) = \frac{v(k) - v(k-1)}{h}$$

where h is the sampling period (0.10 s), the following diagrams results:

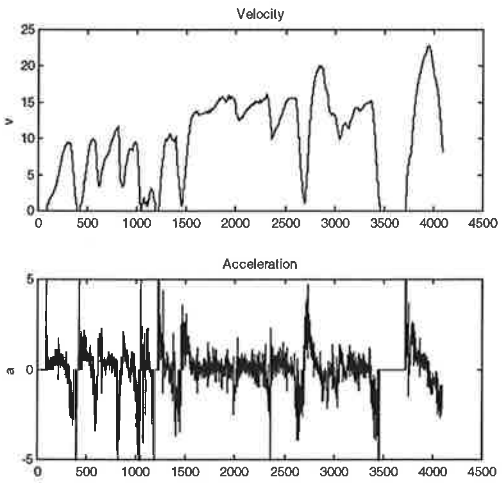

Figure 5.2: The noise in the velocity signal is magnified many times in the acceleration signal.

As seen, the noise in the velocity signal is magnified many times in the acceleration signal. This is a consequence of that the noise is differentiated in the same way as the actual velocity. Low-pass filtering the signal will attenuate the noise, but the inevitably introduced phase-lag will lower the phase margin and the system will be more oscillatory. This is not acceptable since the system must be well damped for comfort reasons. Experiments and simulations have shown that an explicit pre-filtering of the acceleration before it is fed back to the controller causes more problems than it solves. The noise has to be handled internally by the controller. But even so, there are conflicts. On one hand, the noise should be attenuated well by the controller which implies that the bandwidth of the controller cannot be too high. On the other hand, the response to set-point changes should be fast which implies that the bandwidth of the controller should be rather high.

5.3 Controller design

Since the true car dynamics is not known in detail, the controller must be reliable and robust. Therefore, a simple PI-controller was the starting point in designing a suitable controller:

$$u(t) = K\left(e(t) + \frac{1}{T_i}\int_0^t e(\tau)d\tau\right) = Ke(t) + K_i\int_0^t e(\tau)d\tau$$
$$e(t) = a_{set}(t) - a(t)$$

or in discrete time (*h* is the sample period):

$$u(k) = Ke(k) + \frac{K_i h}{q-1}e(k)$$
$$e(k) = a_{set}(k) - a(k)$$

Experiments and simulations have shown that the noise in the acceleration signal is of that great magnitude that a proportional part in the PI-controller is impossible. More correct, to attenuate the noise enough the proportional constant K has to be so small that the improvement in system speed is neglectible. Even if a low-pass filtered $e(k)$ is used in the proportional part (but not in the integral part), the improvement in system speed is neglectible since the bandwidth of the filter must be quite low to attenuate the noise enough.

A pure I-controller removes stationary errors, but it cannot provide the required system speed. The closed system must be well damped so K_i cannot be too large. To acquire the desired speed, the I-controller is complemented with a feed-forward term:

$$u(k) = K_{ff}a_{set}(k) + \frac{K_i h}{q-1}e(k)$$

The feed-forward term has the effect that u will respond directly on set-point changes.

Assume that a_{set} is positive and that u should approximately equal u_0 to accomplish this acceleration, i.e. to make a equal a_{set}. The idea is then to choose K_{ff} so that $K_{ff}a_{set}$ is close to u_0. The integrator should then tune up or down u so that a=a_{set} is finally achieved.

5.3.1 Gain scheduling

Since the dynamics from u_{thr} to a is quite different from the dynamics from u_{brk} to a (see section 4.6), it is not reasonable to use the same control law in the entire scope of operation. The scope of operation is therefore divided into two parts: one when u is positive (u_{thr} is non-zero) and one when u is negative (u_{brk} is non-zero). Different values on K_{ff} and K_i are used in the two parts.

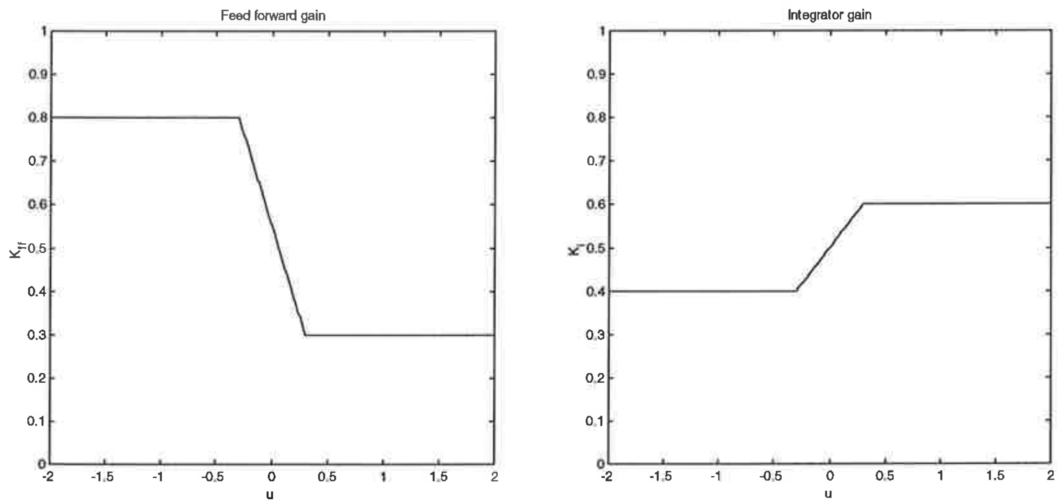

Figure 5.3: The value on the feed-forward gain and the integrator gain varies with the control signal. Linear interpolation gives a smooth transition.

Figure 5.3 shows how K_{ff} and K_i could vary with u. They have different values when u is positive and when u is negative. Below some breakpoint, K_{ff} and K_i are interpolated with a line to make a smooth transition between the values. **Note**: Since $u(k)$ is a function of K_{ff}, K_{ff} cannot be a function of $u(k)$. This problem is easily solved by letting K_{ff} be a function of $u(k-1)$ instead. If the sampling period is short, the difference is minor.

Since the deceleration is almost proportional to the brake pressure (see 4.2.1), it is reasonable to use a fairly high feed-forward. Consider the case when u is negative and assume that a pure feed-forward is used as control law (K_i=0). Also assume that the deceleration is proportional to the brake pressure:

$$u = K_{ff} a_{set}$$
$$a = k_{brk} u_{brk}$$

The control signal is given by:

$$u_{brk} = -K_{brk} u = -K_{brk} K_{ff} a_{set}$$

The acceleration then evaluates to:

$$a = k_{brk} u_{brk} = -k_{brk} K_{brk} K_{ff} a_{set}$$

Since a should equal a_{set}, the following relation should hold:

$$K_{ff} = -\frac{1}{k_{brk} K_{brk}}$$

A pure feed-forward cannot compensate for disturbances and model errors since there is no feedback. Therefore a small integral part should still be used, i.e. K_i should equal some small positive value. Experiments show that K_{brk}=10, K_{ff}=0.7 and K_i=0.2 gives good results.

When u is positive, i.e. when u_{thr} is non-zero, the situation becomes quite different since there is more dynamics involved. Figure 5.4 shows the step response from u_{thr} to a using the model found in section 4.6.

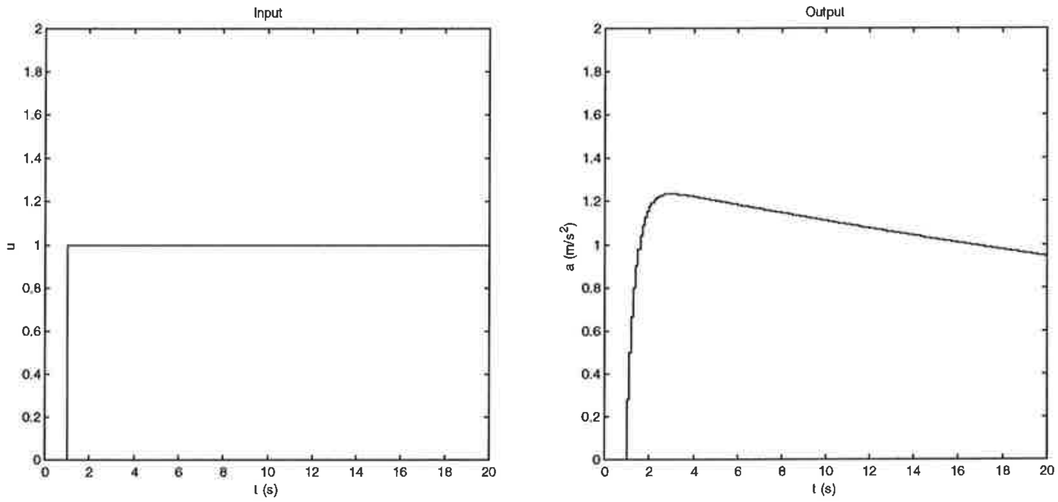

Figure 5.4: The step response when u_{thr} is stepped at t=1.0 s.

Let k_{max} denote the maximum of a in this step response (k_{max} ≈1.23). One can view this value as the maximum gain of the input signal. The output is declining from this value due to frictions and other losses. Now, assume that there is no losses, i.e. the maximum value is maintained once it is achieved, and that a pure feed-forward is used as control law. After the initial transient, the following relations hold:

$$a = k_{\max} u_{thr}$$
$$u = K_{ff} a_{set}$$

The control signal is given by:

$$u_{thr} = K_{thr} u$$

The acceleration then result in:

$$a = k_{\max} u_{thr} = k_{\max} K_{thr} K_{ff} a_{set}$$

Hence, the feed-forward gain should be chosen as:

$$K_{ff} = \frac{1}{k_{\max} K_{thr}}$$

Due to the losses and model errors a pure feed-forward in not enough. Adding an integral part to the control law would reduce those effects. The price is that the system will be more oscillatory. K_i should be chosen so that the errors are taken care of well, but without causing an unacceptably oscillatory system. The noise must also be well attenuated. The choice of K_{ff} is also affected by a non-zero K_i. The greater K_i, the lesser than the suggested value above must K_{ff} be. Otherwise, the control signal would get too large.

Figure 5.5 shows the simulation result using K_{thr} =1.0, K_{ff}=0.6 and K_i=0.5.

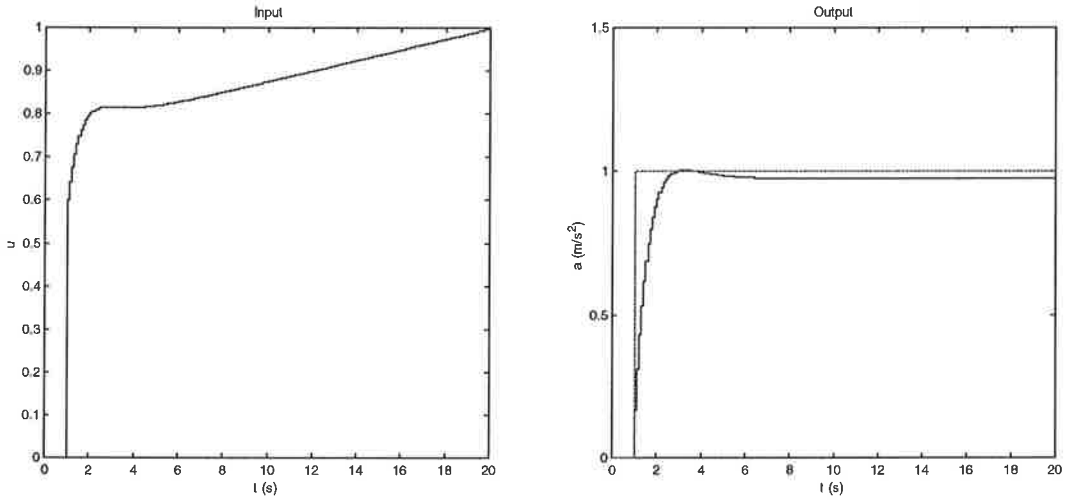

Figure 5.5: The simulation result when the set-value (a_{set}) is stepped at t=1.0 s. The performance is fast and the output is well damped. The control signal (u_{thr}) also shows good behaviour.

The overshoot is quite small and the system is well damped. As seen, a stationary error exists. If it is required that this should not exist, the control law has to include a double integrator, but that would have caused a more oscillatory system. Since the error is small, it is concluded that it is within the bounds of acceptance. Another important aspect is how the control signal (input) looks like. If the control signal oscillates, so will the engine speed. That is not acceptable, since it feels uncomfortable and unnatural to a real driver.

The closed loop behaviour can be derived in the following way. Approximate the acceleration with:

$$a(k) = \frac{v(k) - v(k-1)}{h} = \frac{q-1}{qh} v(k)$$

where h is the sample period (0.10 s). Using the results found in section 4.6, the following relations can be derived:

$$a(k) = \frac{0.0279q^2}{(q-0.9984)(q-0.7812)} \frac{q-1}{qh} u_{thr}(k)$$

$$u_{thr}(k) = K_{thr}u(k) = K_{thr}K_{ff}a_{set}(k) + \frac{K_{thr}K_i h}{q-1}(a_{set}(k) - a(k))$$

Inserting the second equation into the first one and solving for a the result becomes:

$$a(k) = \frac{0.279K_{thr}K_{ff}q^2 + 0.279K_{thr}(0.1K_i - K_{ff})q}{q^2 + (0.0279K_{thr}K_i - 1.778)q + 0.780} a_{set}(k)$$

Inserting K_{thr}=1.0, K_{ff}=0.6 and K_i=0.5 results in:

$$a(k) = \frac{0.167q^2 - 0.154q}{q^2 - 1.766q + 0.780} a_{set}(k)$$

The poles are located at:

$$q = 0.883 \pm 0.024i$$

As seen, they are well damped as required. The zero is determined by K_{ff} and is located at:

$$q = 0.917$$

Note that the zero almost cancels one of the poles. That will make the system response faster. Hence, the feed forward term plays a major role in the achievement of the fast response. Experiments have shown that K_{ff} and K_i should be a little lesser than the previous suggested values. It depends partly on the different static gains of the gears. More suitable values are: K_{ff}=0.5 and K_i=0.4.

5.3.2 Modifications

The controller designed in the previous section showed good results in practical experiments, but its performance could be even further improved by some small modifications. The integral part was originally updated using the following formula:

$$I(k+1) = I(k) + K_i h\, e(k), \quad e(k) = a_{set}(k) - a(k)$$

To reduce the effect of a sudden temporal huge noise signal it was modified to:

$$I(k+1) = I(k) + K_i h\, sat(e(k))$$

where *sat* stands for saturation and is defined as:

$$sat(e) = \begin{cases} e_{low}, & e \le e_{low} \\ e, & e_{low} < e < e_{high} \\ e_{high}, & e \ge e_{high} \end{cases}$$

With this modification, $e(k)$ is upper and lower bounded before it affects $I(k+1)$. This modification sets a limit on how fast $I(k)$ can change from sample to sample. Suitable values of the bounds are found to be e_{low}=-1.0 and e_{high}=1.0.

A general conclusion drawn from several experiments was that the acceleration often was too high when the control signal u_{thr} was rather huge. That can be explained by studying the engine map. The engine map shows how the engine torque varies with the engine speed and throttle position. The engine map of the test car used showed a non-linear mapping between the torque and the throttle position and the engine speed. Large throttle positions and high engine speeds gives a relative higher torque than small throttle positions and low engine speeds. The idea was therefore to use a relative smaller control signal for large computed values of u. The following formula was used instead:

$$u_{thr} = K_{thr}\, 2\left(\sqrt{u+1}-1\right) \quad u \ge 0$$

Figure 5.6 compares this formula with the original one (u_{thr}=$K_{thr}u$).

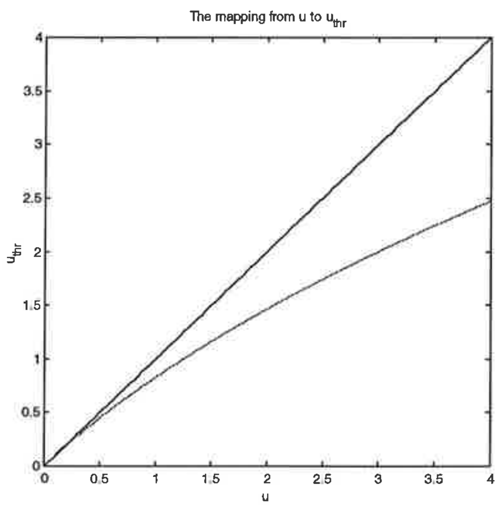

Figure 5.6: The actual control signal to the throttle (dashed line) is relative lower at high control values due to non-linearities.

For small u, both formulas give approximately the same result. The larger u gets, the greater is the difference.

Experiments have shown that the modifications result in more comfortable accelerations and it will be better matched with the desired one.

5.4 Implementation

The derived controller was implemented and tested in practise. The acceleration generated by the conventional cruise control system was used as set-value (a_{set}). Figure 5.7 shows the result when the car is accelerating from v=2 m/s to v=12 m/s.

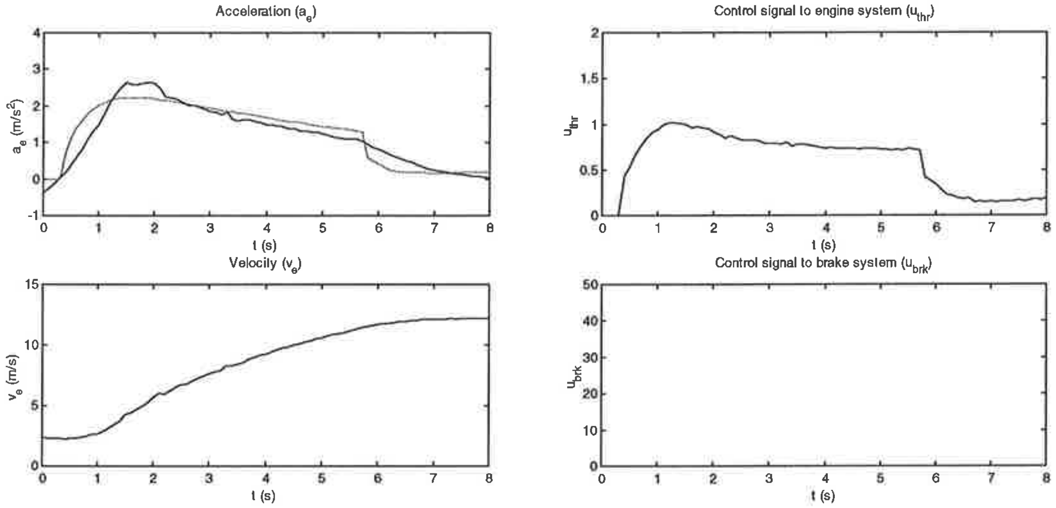

Figure 5.7: The controller was tested in practise using the conventional cruise control system. The actual acceleration (solid line) is close to the desired acceleration (dashed line).

The overshoot in the acceleration signal is small and there are no oscillations. The match with the desired acceleration is also rather good. The noise in the control signal is well attenuated as required. Figure 5.8 shows the result when the car is decelerating from v=16 m/s to v=7 m/s.

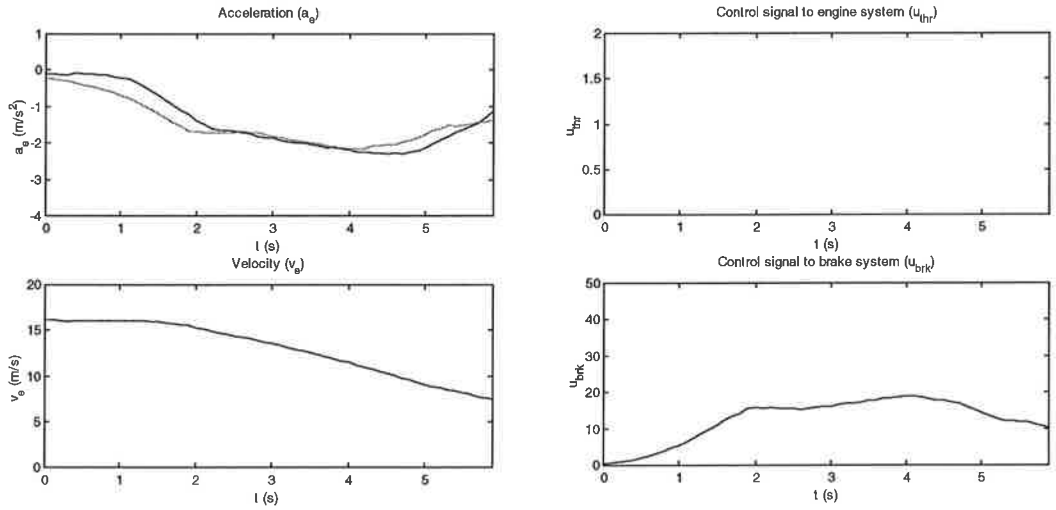

Figure 5.8: Testing the deceleration performance. The actual deceleration (solid line) is close to the desired one (dashed line).

The match with the desired deceleration is here also rather good. There is always an initial time delay before the deceleration starts increasing. The reason for this is that it takes a little time to achieve the desired brake pressure. Remember that u_{brk} is not the actual brake pressure, but rather the desired brake pressure (see section 4.1).

5.5 Summary

The final controller is summarised in this section. The feed forward and integrator gain is determined by figure 5.9. They are a function of the pervious control signal, *u(k-1)*.

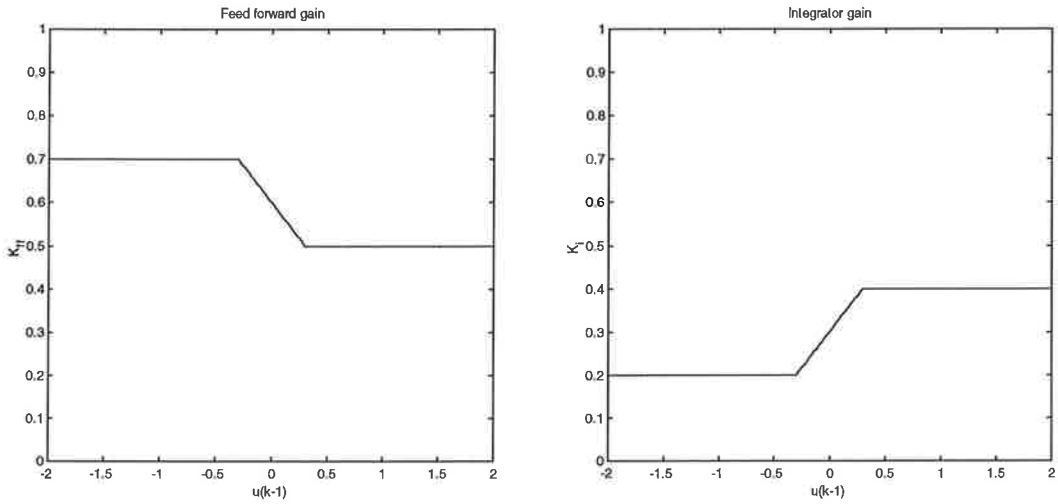

Figure 5.9: The controller uses gain scheduling. Different parameter values are used depending on controlling the throttle position (u>0) or the brake pressure (u<0). Linear interpolation gives a smooth transition.

The breakpoint where the linear interpolation begins is located at $u(k\text{-}1)=\pm 0.3$.

The control law is given by:

$$u(k) = K_{ff} a_{set}(k) + I(k)$$

The integral part is updated in the following way:

$$\begin{aligned} e(k) &= a_{set}(k) - a(k) \\ I(k+1) &= I(k) + K_i h \; sat(e(k)) \\ sat(e)&: \quad e_{low} = -1.0, e_{high} = 1.0 \end{aligned}$$

The actuator control signals are calculated from u and are given by:

$$u_{thr}(k) = \begin{cases} 0, & u(k) < 0 \\ 2K_{thr}(\sqrt{u(k)+1} - 1), & u(k) \geq 0 \end{cases}$$

$$u_{brk}(k) = \begin{cases} -K_{brk} u(k), & u(k) < 0 \\ 0, & u(k) \geq 0 \end{cases}$$

The scaling coefficients are:

$$K_{thr} = 1$$
$$K_{brk} = 10$$

5.6 Conclusions

Several experiments have shown that the controller behaves in a comfortable and fast way, but it is in no way perfect. A better controller could have been designed if the engine speed and the current gear were available information to the controller. Unfortunately, that was not possible in the test car used for the experiments. Unknown model errors, e.g. non-linearities, also debase the performance of the controller. Particular when the car should accelerate from a non-moving position (v=0), the performance is poorer compared to other situations.

6. Driver Modelling

The driver model should model the real driver behaviour in a satisfactory way. The model must conform to the safety needs a real driver experience, i.e. keeping an appropriate distance to the vehicle ahead. It is also very important that the model conforms to the need of comfort, so that the overall system performs smoothly and well. Small deviations from the desired distance should not lead to big reactions (e.g. huge brake pressures), but rather react in such way that the desired distance is obtained slowly and smoothly. The acceleration and deceleration of the car must conform to a real driver's behaviour in every situation. The fundamental approach is that a driver's desired acceleration is a static function of v, d and Δv. Several investigations supports this assumption (see e.g. [1]).

6.1 Data collection

Data was collected in the same way as described in section 4.1. Since the inner control loop should ensure that desired acceleration is obtained quickly, the desired acceleration generated by the model should match the actual acceleration of the car when a real driver is driving.

6.2 Linear regression

The first approach was to find a simple linear relationship between the desired acceleration (a_{set}), and the speed of the following car ($v=v_e$), distance (d) and relative speed ($\Delta v=v_f-v_e$):

$$a_{set} = k_0 + k_1 v + k_2 d + k_3 \Delta v$$

The coefficients were estimated using a least square method, given data of a real driver's behaviour. The desired distance could then be written as:

$$d_{set} = T_g v_e + d_0 \qquad or \qquad d_{set} = T_g v_f + d_0$$

where T_g is the desired time gap and d_0 is the desired zero-speed distance. Analysis of several data shows that T_g is typically in the range of 0.7-1.2 s, and that d_0 is in the range of 2.5-4.0 m. Which formula to choose depends on viewpoint. Figure 6.1 shows the correspondence between the simulated and the actual value. The lower diagram used validation data, i.e. data that had not been used earlier.

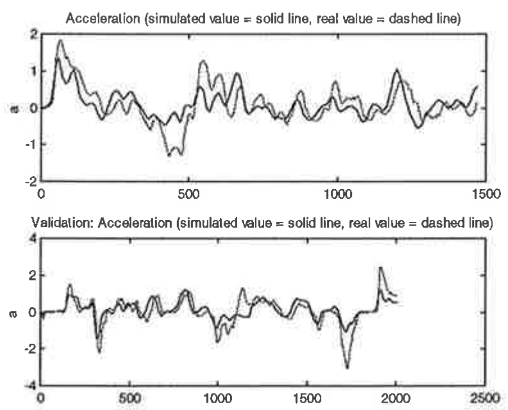

Figure 6.1: The first approach is a simple linear relation. The simulated value (solid line) is compared with the actual acceleration of a real driver (dashed line). The validation (lowest diagram) uses data that not has been used earlier.

The correspondence is rather good when the magnitude of the acceleration is small, but is less accurate when the magnitude is large.

6.3 *Dangerous situations*

One major drawback with the regression model is that it does not ensure that collisions cannot occur. The reason for this is that it reacts to gently in dangerous situations. Dangerous situations arise when the relative speed is large negative and the distance to the vehicle ahead is short. Hence, it seems valuable to try to get some kind of measurement of how dangerous a given situation is. Therefore, a new variable D is introduced that should reflect the danger of a situation:

$$D = D(v, d, \Delta v), \quad D \geq 0$$

Note that D only assumes non-negative values. For non-dangerous situations, D should be small and then increase with higher danger. There are many possible definitions of D. Some suitable are defined below.

6.3.1 Small headways

Intuitively, small headways are dangerous because the following vehicle has little time to react if the preceding vehicle changes its speed. It is especially desired to maintain the desired distance at small distances, because a distance error of a couple of meters can feel very uncomfortable and possibly be hazardous. Therefore, let D have the following form:

$$D(v, d, \Delta v) = \frac{1}{d}$$

For small d, D becomes huge (high danger) and for huge d, D becomes small (low danger). The major disadvantage with this formula is that it is only a function of the distance. A more appropriate measurement should also be a function of Δv, at least when Δv is negative. The greater magnitude of Δv, the more dangerous is the situation. Thus, D should increase with the magnitude of Δv when it is negative.

6.3.2 Time-To-Collision (TTC)

A convenient and intuitively clear measurement on how dangerous a situation is, is the Time-To-Collision (TTC) concept. Assume that both vehicles are continuing with their current speed, then TTC is the time it takes until a collision occurs. If Δv is positive, a collision will never occur and TTC is defined as positive infinity:

$$TTC = \begin{Bmatrix} -\dfrac{d}{\Delta v}, & \Delta v < 0 \\ \infty, & \Delta v \geq 0 \end{Bmatrix}$$

TTC is small for dangerous situations and huge for non-dangerous situation, i.e. the direct opposite behaviour of D. Therefore, define D as the reciprocal of TTC:

$$D(v, d, \Delta v) = \frac{1}{TTC} = \begin{Bmatrix} -\dfrac{\Delta v}{d}, & \Delta v < 0 \\ 0, & \Delta v \geq 0 \end{Bmatrix}$$

The advantage with this formula is that D depends on Δv. If the magnitude of Δv is small, D will remain fairly small even if d is rather small.

6.3.3 Minimum deceleration to avoid collision (a_{min})

Another suitable concept to evaluate the danger of a situation is the minimum required constant deceleration to avoid a collision. Assume that the preceding vehicle is continuing with its current speed and that the following vehicle is approaching. The minimum required constant deceleration to avoid a collision can de derived in the following way (v_f and a are constants):

$$v_e = v_e(0) + at$$
$$d = d(0) + \int_0^t (v_f - v_e(\tau))d\tau = d(0) + v_f t - v_e(0)t - \frac{at^2}{2}$$

Solving the equations:

$$v_e = v_f$$
$$d = 0$$

results in:

$$a = -\frac{(v_f - v_e(0))^2}{2d(0)}$$

This quantity can be calculated at any moment:

$$a = -\frac{(v_f(t) - v_e(t))^2}{2d(t)} = -\frac{\Delta v^2}{2d} = -a_{\min}$$

The formula above states that given the momentaneous values of v_e and v_f and that v_f is continuing constant in the future, a_{min} is the minimum constant deceleration required to avoid a collision. This value can be used as a measurement of the danger:

$$D(v, d, \Delta v) = \begin{cases} \frac{\Delta v^2}{2d}, & \Delta v < 0 \\ 0, & \Delta v \geq 0 \end{cases}$$

When Δv is positive, D is defined as zero just as in the previous section (TTC). The major difference with this definition and the previous one is that this one considers huge Δv more dangerous than the previous one (since it is using the square of Δv).

6.3.4 Modifications

Since the desired distance never drops below d_0, one can regard the "collision point" when d equals d_0 instead of zero. To do so, replace every location of d with $(d-d_0)$ in the formulas for D defined in the previous sections, e.g.:

$$D(v, d, \Delta v) = \begin{cases} \frac{\Delta v^2}{2(d - d_0)}, & \Delta v < 0 \\ 0, & \Delta v \geq 0 \end{cases}$$

When Δv is negative and d is less than or equal to d_0, D is defined as positive infinity. The effect of this modification is that the safety margin will be increased, since the collision point is moved backwards.

6.4 Approaching situations (negative Δv)

Since a real driver often brakes with a rather constant deceleration during the brake procedure (see e.g. [2]), the model should generate a rather constant a_{set} when Δv is large negative (fast approaching). For large negative Δv the desired acceleration should approximately equal:

$$a_{set} \approx -\frac{\Delta v^2}{2(d - d_{set})}, \qquad \Delta v < 0, d > d_{set}$$

This is not achieved with the linear regression model. Consider the following scenario. The following car is approaching the vehicle ahead. The vehicle ahead is keeping a constant speed of 5 m/s. The following car has initially a speed of 20 m/s and a distance of 200 m. It is assumed that $a_e = a_{set}$, i.e. that the desired acceleration is achieved instantly by the inner control loop. Figure 6.2 shows the simulation results when the linear regression model (model 1) is used.

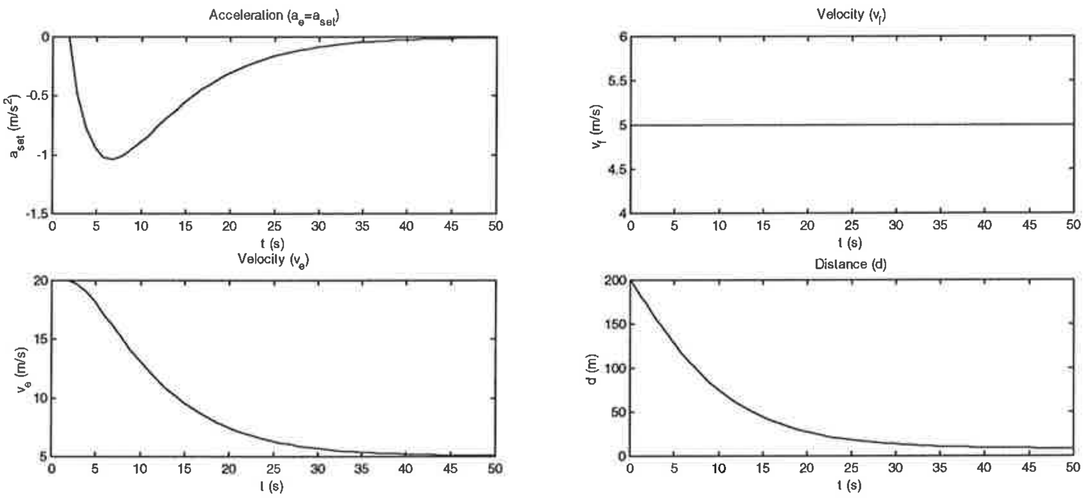

Figure 6.2: Simulation of an approach scenario using the linear regression model. The deceleration is not constant.

The deceleration becomes high quickly, but decreases slowly after have reached a maximum. The entire brake procedure is quite extended in time. The maximum deceleration is 1.0 m/s^2. The braking procedure does not conform well to how a real driver approaches a slower-moving car. At higher approach speeds, the behaviour is more unacceptable than in this case. The model was therefore expanded with another variable F that had the property that if the linear regression model was multiplied with F, the above relation was approximately achieved. Other modifications were also made. Figure 6.3 shows a simulation of the approach situation with these modifications included.

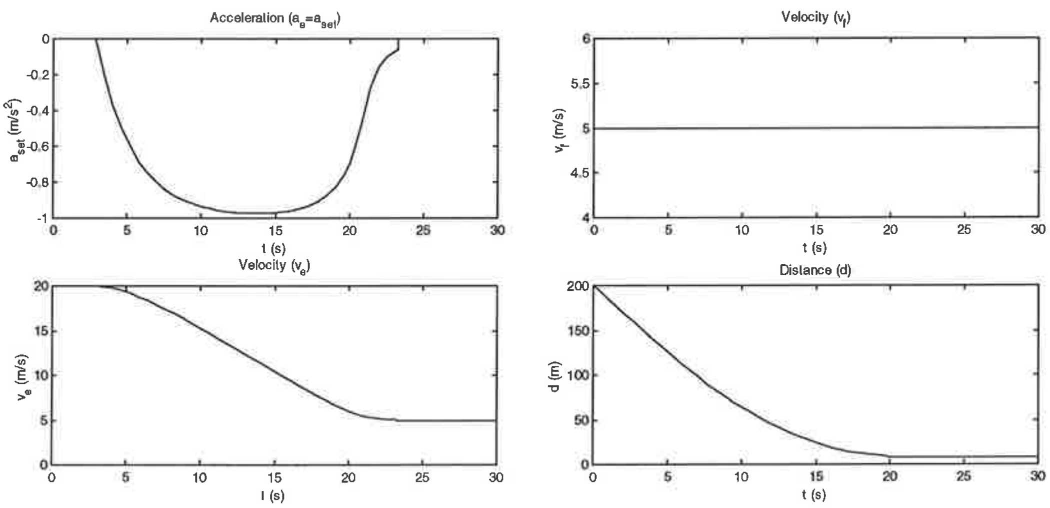

Figure 6.3: Simulation of a fast approach situation. The deceleration is fairly constant during the brake procedure.

The reason why the linear regression model does not have an appropriate behaviour is that small distance errors are penalised too little with the consequence that the deceleration becomes too small. Therefore, it takes plenty of time before the desired distance is achieved. Large distance errors are on the other hand penalised too hard with the consequence that the deceleration becomes too high in fast approaching situations. The modified model shows an appropriate penalty on distance errors and the overall behaviour is quite similar to a real driver's behaviour.

6.4.1 Non-dangerous situations

Even though the model behaves satisfactory in the approach situation, it behaves unsatisfactory in other situations. Assume for instance the following scenario. The vehicle ahead keeps a constant speed of 10 m/s. The following vehicle has an initial speed of 11 m/s and a distance of 50 m. Figure 6.4 shows the simulation results:

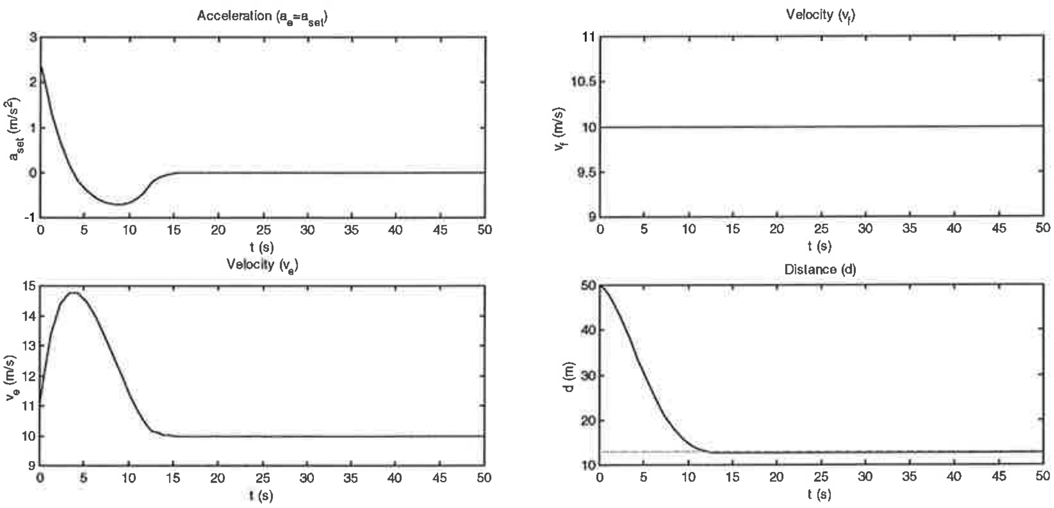

Figure 6.4: Simulation of a non-dangerous situation. The acceleration is initially too high to be comfortable.

The maximum acceleration is 2.3 m/s^2 and the maximum velocity is 14.8 m/s. The initial acceleration is too high to feel comfortable. Since the deviation from the desired distance is not particular high, a real driver would have only accelerated gently so the later needed deceleration also could be performed gently.

The scenario above depicts a general flaw of the model. In non-dangerous situations, the magnitude of the acceleration is generally too high. Therefore, a new variable Q was introduced that is a function of how dangerous a situation is. The model was then multiplied with this factor. Q should typically be less than 1 for non-dangerous situations and greater than 1 for dangerous situations. If D is a measure of how dangerous a situations is, then Q should increase from a lower bound (Q_{min}) when D is small (non-dangerous situations) towards an upper bound (Q_{max}) when D gets high (dangerous situations). Figure 6.5 shows an example of a suitable relation between Q and D.

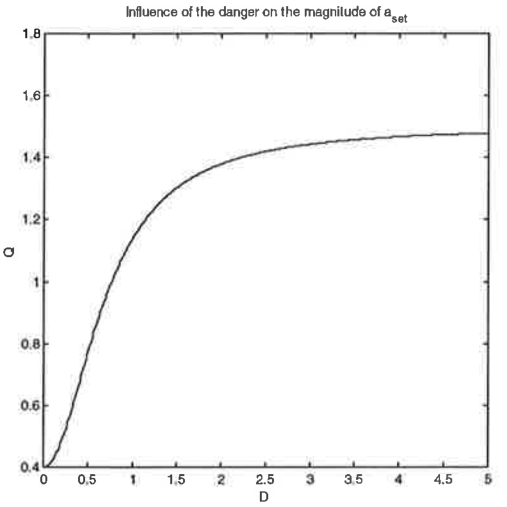

Figure 6.5: The magnitude of a_{set} is attenuated in non-dangerous situations. In dangerous situation, it is on the other hand magnified.

D could be defined as one of the alternatives in section 6.3. Simulation with the new model (with the Q-factor included) results in:

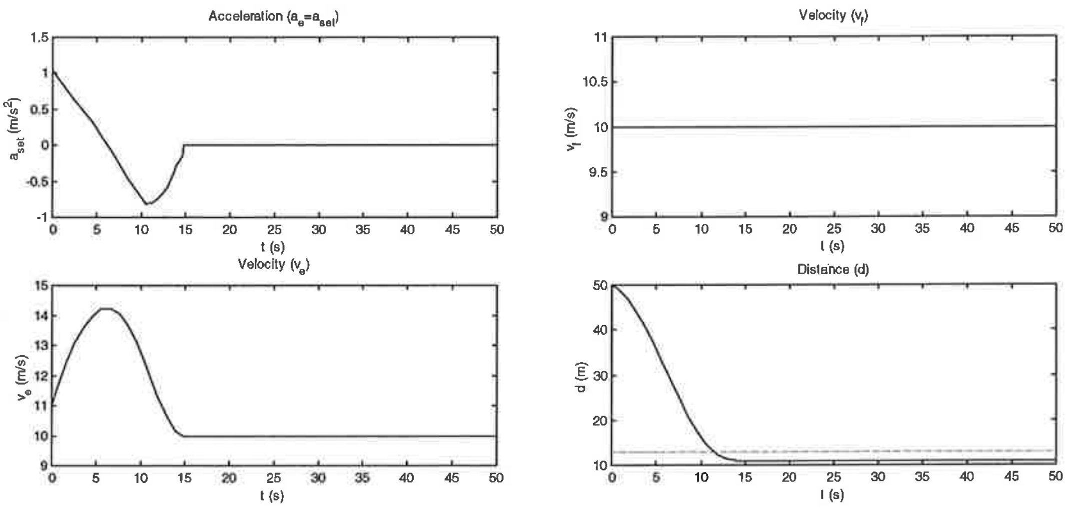

Figure 6.6: The introduction of the Q-factor results in a gentler behaviour. The initial acceleration is now on an acceptable level.

The maximum acceleration is now 1.0 m/s^2 and the highest speed is 14.2 m/s. The highest acceleration is more than halved compared to the case when the Q-factor is excluded. The price is that the undershoot has increased from 0.4 m to 2.1 m, but that is not a particular severe drawback since the situation never becomes real dangerous (there is still almost 11 m to go before a collision).

6.5 Separating situations (positive Δv)

Situations with positive Δv are quite different from situations with negative Δv. The major difference is that there is no imminent danger for collision, since the vehicle ahead is moving away from the vehicle behind. Whereas the problem with negative Δv is to avoid a collision, the problem with positive Δv is to avoid that the distance becomes much higher than the desired one, i.e. to avoid that the following vehicle lag behind. A natural requirement is that a_{set} should be continuous when Δv changes from negative to positive values.

Remember that F was introduced to make the deceleration fairly constant during an approach. There is no such analogous situation when Δv is positive. Hence, F is defined as:

$$F = 1, \qquad \Delta v \geq 0$$

Since there is no imminent risk for collision, Q is also modified. Define Q_0 as:

$$\Delta v = 0 \Rightarrow Q = Q_0$$

and let Q equal Q_0 for all positive Δv:

$$Q = Q_0, \qquad \Delta v \geq 0$$

In this way, a_{set} will be continuous when Δv changes from negative to positive values. Now, assume the following scenario. Both vehicles are initially standing still with the initial distance of $d = d_0 = 3$ m. The preceding vehicle is then accelerating with 1.0 m/s^2 until the velocity has reached 15 m/s. Figure 6.7 shows the simulation result using this model.

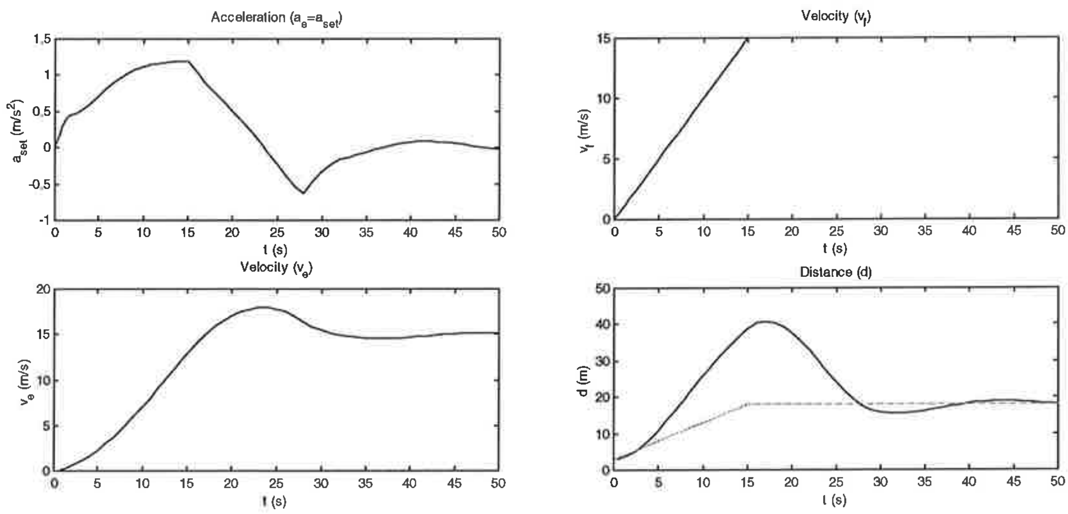

Figure 6.7: Simulating a scenario where the preceding vehicle is accelerating with 1.0 m/s^2 until a speed of 15 m/s is achieved. The performance is poor. The distance gets much higher than the desired one (dashed line).

The result is very poor. The acceleration is too low in the beginning, with the consequence that the vehicle will lag behind. The maximum distance is 40.8 m and the maximum speed is 17.9 m/s.

The reason why the acceleration is too low is that the magnitude of a_{set} is not amplified when Δv is large. To improve the performance, Q was altered to increase with increasing Δv. Figure 6.8 shows the result when simulating the same scenario, using this modified formula instead.

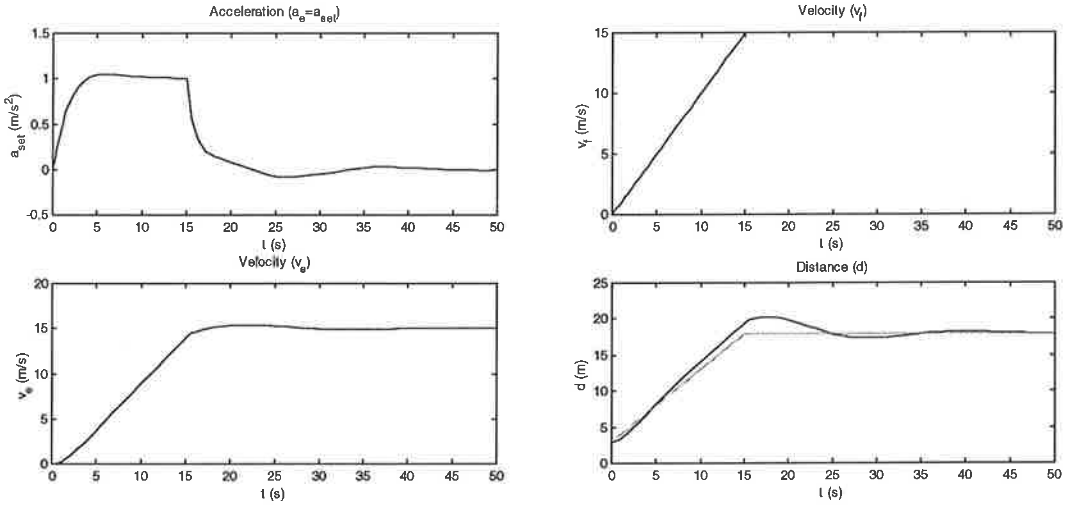

Figure 6.8: The performance is much better with the modification included. The distance is now close to the desired one (dashed line).

The performance is now much improved. The acceleration is quickly approaching the acceleration of the preceding vehicle (1.0 m/s^2) and the distance is close to the desired distance. The maximum velocity is 15.4 m/s and the maximum distance is 20.2 m.

6.6 *Upper and lower bounds*

If the distance is huge and the relative speed is large positive, the driver model will generate a very high desired acceleration. But a real driver never wants to accelerate too hard, because above some level it feels uncomfortable. Experiments have shown that this level depends on the current velocity. The higher velocity, the lower level. The reason for this is that an acceleration of 1.0 m/s^2 at a speed of 25 m/s is percepted to be much higher than the same acceleration at a speed of 10 m/s. For the reasons above, a_{set} is upper bounded by a_{set_max}. This upper bound is speed dependent and decreases with higher speed. Figure 6.9 shows an example how this relation could be like.

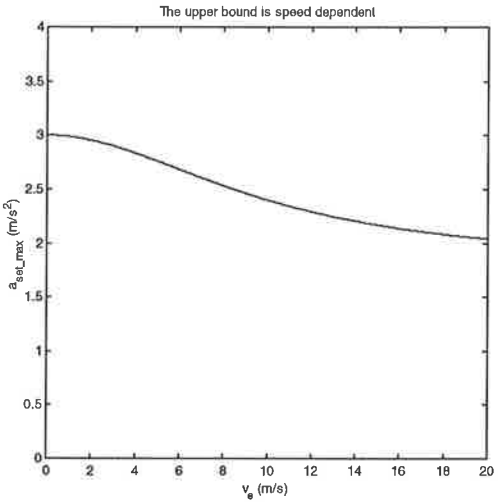

Figure 6.9: The upper bound on the acceleration should be speed dependent. The higher speed, the lower bound.

It also makes sense to use a lower bound (a_{set_min}). The value of this bound is mainly determined by how much the ACC system should be allowed to brake. It should not be too high (not negative enough), since collisions can occur if the system is not allowed to brake hard enough. If a_{set_min}=-3.5 m/s^2 most situations will be handled well.

6.7 Conclusions

The driver model is derived step-by-step. Starting with a simple linear relation, the model is modified and expanded piece by piece. The changes are driven by unacceptable behaviour in certain situations of the current model. The proposed changes are often a result of intuitive reasoning. Most effort has been made on situations where the relative speed is negative.

The final driver model includes several parameters. The advantage with many parameters is that the model could be tuned to fit a desired behaviour very well. The disadvantage is that it might be difficult to understand the model and the influence of the different parts.

A more numerical approach in deriving a driver model could also be chosen (see e.g. [1]). The idea here is to fit the parameters of a general model using data from a real driver. The problem with this approach is that the data must be very extensive to get an appropriate behaviour in all situations. It is also hard to find a suitable general model and it would probably include even more parameters than in the approach used here.

It is concluded that a combination of both approaches would be optimal.

7. Implementation and validation

The driver model in the previous section was derived using many intuitive reasoning and simulation results. It is therefore very important that the results are verified in practise. The driver model must be compared to the behaviour of real drivers. One must though remember that there will seldomly be a perfect match between the behaviour of the driver model and the real driver. A real driver can react in several acceptable ways in the same situation. The actual behaviour might change from time to time and from driver to driver. Hence, the driver's perception of the driver model must also be noted.

7.1 Typical situations

First, the driver model was implemented and tested during normal driving. The behaviour was tested in normal traffic with no special preparations. The car was driving in both urban and sub-urban traffic. The behaviour was tested in many typical traffic situations that required a fairly low acceleration or deceleration. Figure 7.1 shows an example of how the behaviour could be like.

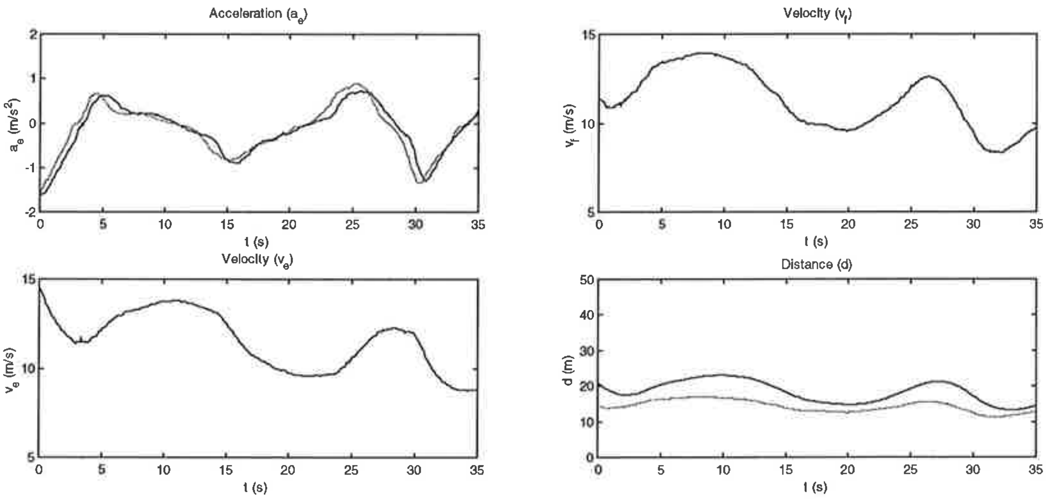

Figure 7.1: An example of how the behaviour could be like in simple car following.

The general conclusion was that the driver model conforms very well to the behaviour of a real driver. This conclusion was drawn by observing how many situations that felt uncomfortable or unnatural. Most situations felt very comfortable and natural, whereas some few felt only fairly comfortable and natural. The behaviour was in no case unacceptable. The evaluation of the behaviour is here done very qualitative, but since the result was so good no further investigations were made (mainly due to limited time resources).

7.2 *Special situations*

The behaviour in typical situations is very important, but other situations must be examined as well. Typical situations are often simple situations because the required reaction is often limited. To examine the behaviour in non-typical situations, e.g. dangerous situations, special arrangements were made. First, the tests were not performed in real traffic, but rather at a non-public piece of road. Second, the driver of the vehicle ahead was instructed to drive in a certain way depending on test case. Each test case was performed twice. One with the ACC system activated and one with the ACC system deactivated, i.e. manual driving. In this way, the behaviour of the ACC system could be compared with the behaviour of a real driver.

7.2.1 Approach situations

The first test case was the following scenario. The vehicle ahead should drive with a constant speed of 40 km/h (=11.1 m/s) and the following vehicle should use a set-speed of 60 km/h (=16.7 m/s). Figure 7.2 shows the deceleration behaviour when the following car approaches the car ahead.

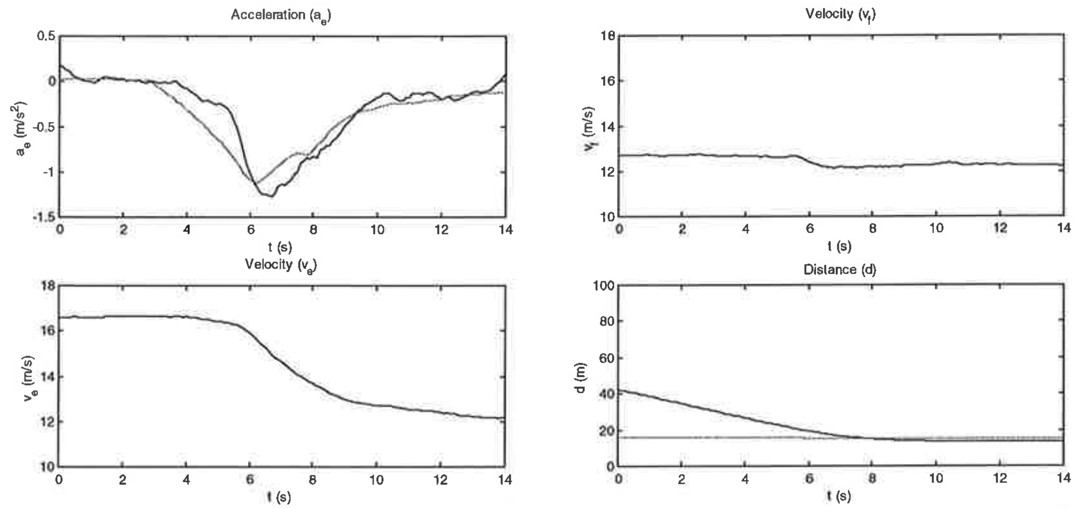

Figure 7.2: Testing a scenario where the following vehicle is approaching slowly. The dashed lines are the desired values.

As noted, the desired and actual acceleration does not match very well. This makes the evaluation of the driver model more uncertain. The same test was then repeated, but this time with the ACC system deactivated. The driver of the following car should instead drive manually and approach the car in a normal way. Figure 7.3 shows the result.

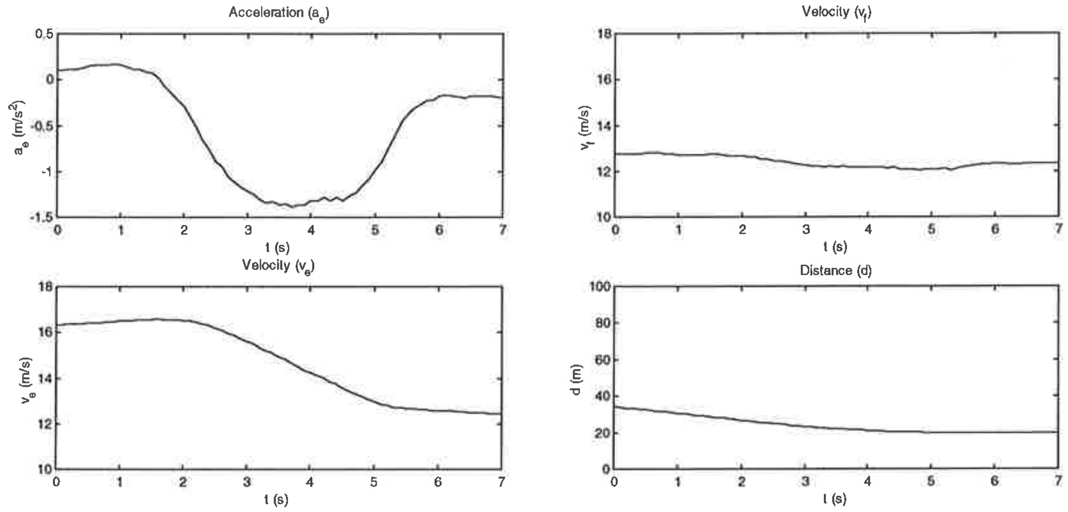

Figure 7.3: The behaviour of a real driver. The deceleration is fairly constant.

The deceleration is fairly constant during the brake procedure. The behaviour of the driver model does not seem to fit well with the behaviour of a real driver in this situation. But as noted before, it is hard to evaluate the driver model since the actual acceleration does not match the desired acceleration very well. A simulation of this scenario shows that the driver model generates a more constant deceleration than in figure 7.2. Even though there is a difference between the desired and actual acceleration one can draw one conclusion. The distance never drops much below the desired distance, which is an indication that the overall system is robust.

The same scenario was tested a second time, but this time the initial speed of the following vehicle was increased from 60 km/h to 80 km/h (=22.2 m/s). Figure 7.4 shows the result when the ACC system was activated.

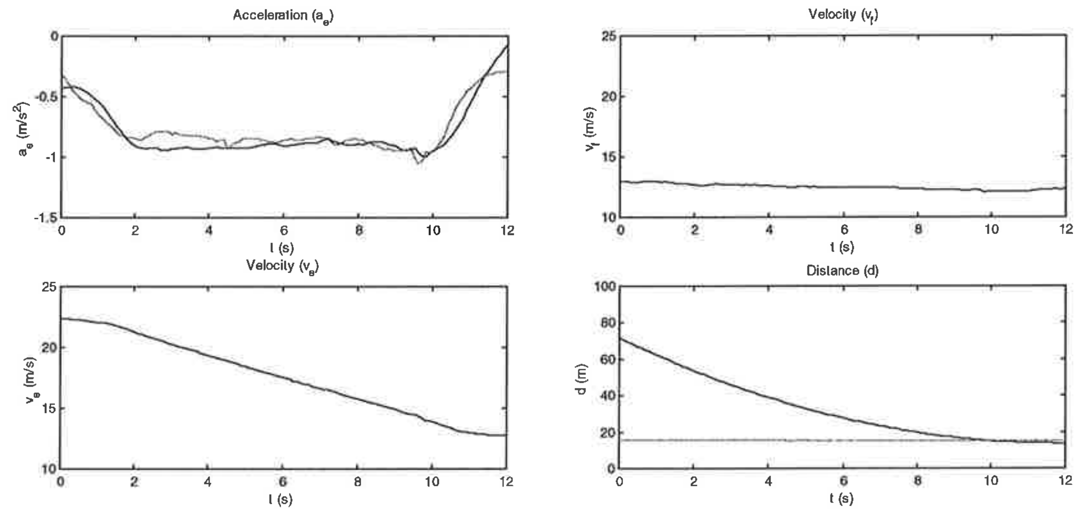

Figure 7.4: The scenario is now tested with a higher approach speed. The deceleration is as expected rather constant during the brake procedure.

The deceleration is as expected constant during the brake procedure and has a magnitude close to the desired one (1.0 m/s^2). The behaviour of a real driver is shown in figure 7.5.

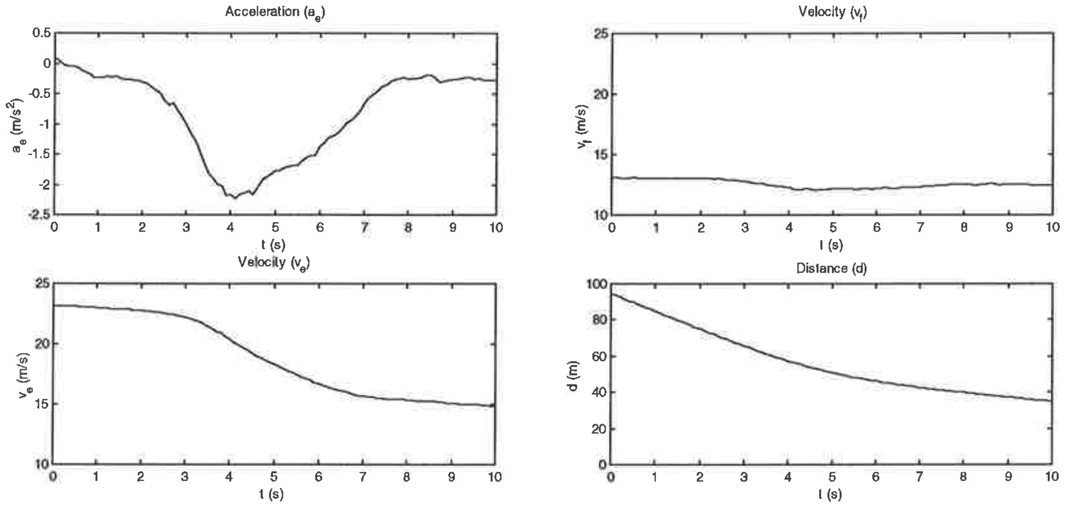

Figure 7.5: The behaviour of a real driver does not necessarily show a constant deceleration. In this case, it is slowly declining.

The deceleration is not very constant and the magnitude is also higher compared to the behaviour of the driver model. Even though the behaviour of the driver model is different than the behaviour of the real driver, it does not make the driver model unacceptable. Tests have shown that most manual brake procedures of an approach situation either shows a constant deceleration or a slowly declining deceleration. Both are acceptable to a driver. The driver model includes parameters that determine the level of deceleration and they could easily be chosen so the deceleration is not constant, but rather conforms to the deceleration in the figure 7.5.

7.2.2 Dangerous situations

Dangerous situations arise when the approach speed is high and the distance is short. In these situations, the behaviour of the ACC system must ensure that collisions are avoided. Still, the reaction should not be unnecessary hard, i.e. the behaviour must be both safe and comfortable. Figure 7.6 shows the driver model behaviour in the following situation. The distance is initially 12.6 m, the speed of the following car is 58 km/h (=16.1 m/s) and the speed of the preceding car is 45 km/h (=12.5 m/s). At this situation, the ACC system is activated (corresponding to t=1.4 s in figure 7.6).

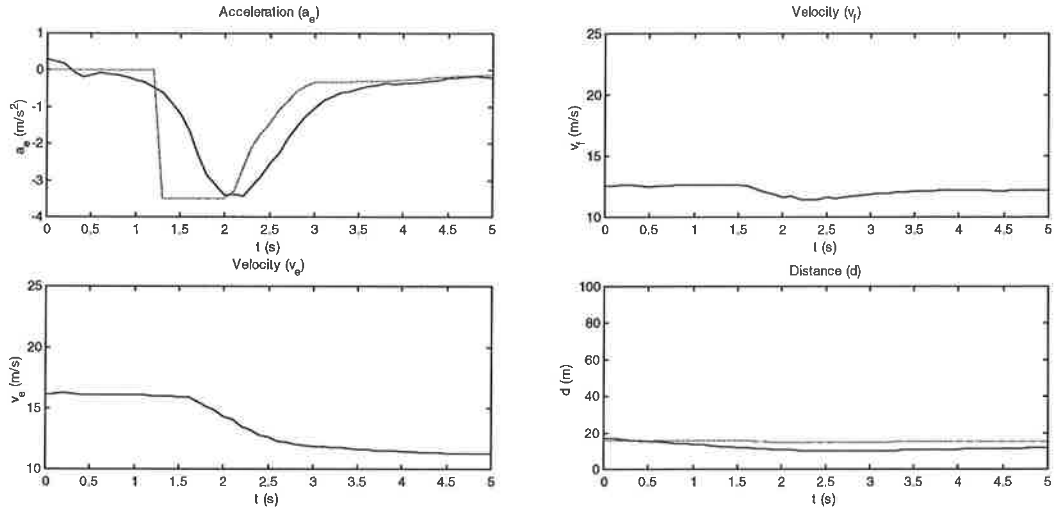

Figure 7.6: A fairly dangerous situation is tested. The desired deceleration equals the maximum allowed. The match between the desired and actual deceleration is not very good.

The driver model instantly wants to decelerate with maximum allowed deceleration (3.5 m/s^2). It takes a little time before the required brake pressure is achieved, so the actual deceleration will be a little bit delayed. The minimum distance is 10.2 m. Figure 7.7 shows the behaviour of a real driver in approximately the same situation.

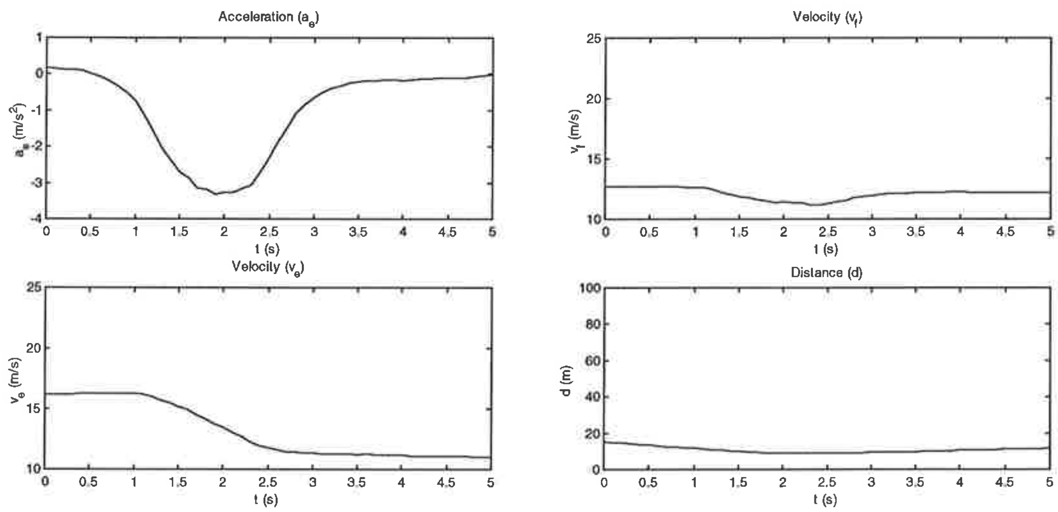

Figure 7.7: The behaviour of a real driver in approximately the same situation. The deceleration is hard, but fairly constant.

The deceleration matches well the actual deceleration of the ACC system. The minimum distance is 8.9 m. It is hard to evaluate if the behaviour of the driver model is appropriate or not, since the actual deceleration does not match well with the desired acceleration. There is though indications that the driver model wants too decelerate too hard in this situation. That could be remedied by tuning the parameters.

Consider a similar scenario but where the speed of the following car is increased to 77 km/h (=21.4 m/s). The initial distance is 46.8 m and the speed of the preceding car is

45 km/h (=12.5 m/s). Figure 7.8 shows the behaviour of the driver model in this situation.

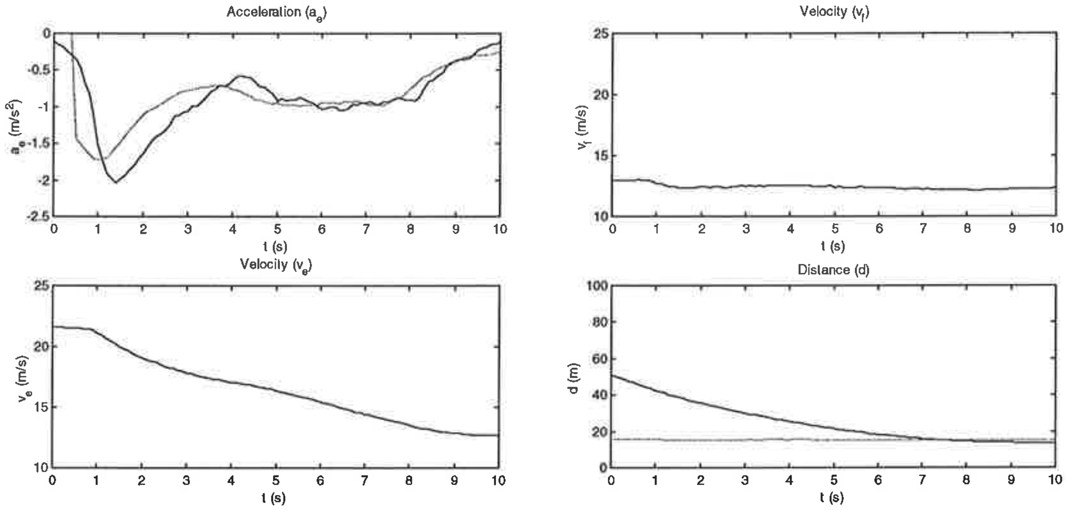

Figure 7.8: Testing a scenario where the approach speed is rather high and the distance is relative low. The deceleration is initially harder, but approaches the desired one (1.0 m/s^2) after a while.

The deceleration increases first to 2 m/s^2 and then it decreases to 1 m/s^2. This is a general property of the driver model. The driver model tries to decelerate with the desired deceleration, but if this is not enough to achieve $\Delta v=0$ and $d=d_{set}$, it first decelerates harder until the desired deceleration is enough. The behaviour of a real driver in a similar situation is shown in figure 7.9.

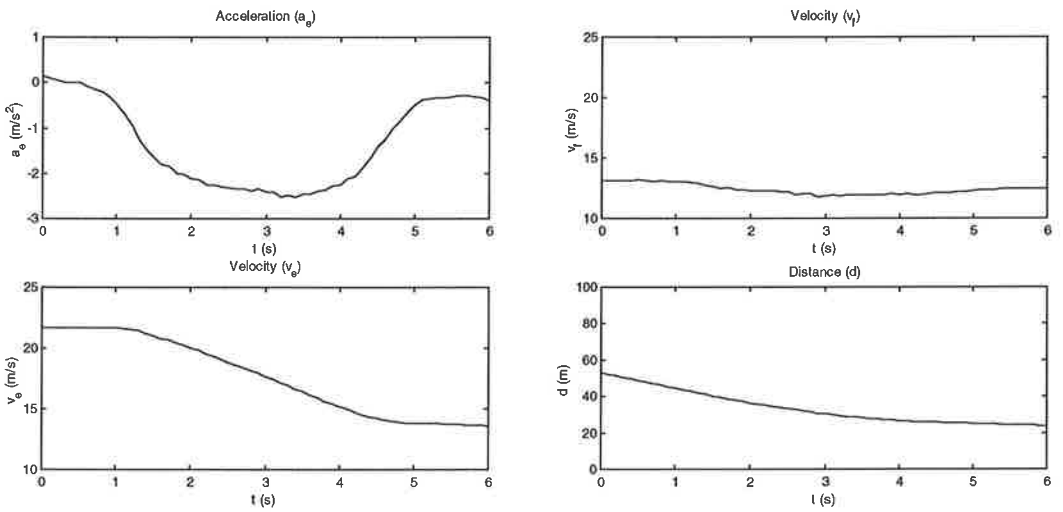

Figure 7.9: The behaviour of a real driver in the same situation. The deceleration is rather constant during the brake procedure.

The deceleration is fairly constant with a value close to 2 m/s^2.

7.3 Conclusions

The general conclusion of the driver model is that it behaves in a comfortable and safe way in most situations. The driver model is complex with many parameters, but most parameters have a clear intuitive meaning, e.g. desired deceleration. Due to many parameters, the driver model could be tuned to fit the behaviour of a real driver well. Implementation has also shown that the driver model is quite robust, i.e. even if the desired acceleration is not achieved quickly by the actuator controller, the desired distance is achieved rather comfortable. Due to limited time resources, there has been only very little tuning of the parameters. More time has to be spend on this to get an even better behaviour of the driver model.

8. Acknowledgements

This thesis was a joint project with the Department of Automatic Control, Lund Institute of Technology and Volvo Technological Development Corporation. Some people contributed to my work a great deal. I would like to thank Prof. Rolf Johansson at the Department of Automatic Control for being my advisor and for giving me guidelines. I would also like to thank the Department of Driving Support Systems at Volvo Technological Development Corporation, where my work was conducted. Especially Fredrik Botling and Erik Hesslow gave me much assistance.

9. References

1. Bleile, Thomas [Robert Bosch GmbH]: *A New Microscopic Model for Car-following Behaviour in Urban Traffic*
2. van der Horst, A.R.A. (1990): *A Time-Based Analysis of Road User Behaviour in Normal and Critical Encounters*
3. Johansson, Rolf (1998): *System Modeling and Identification*, ISBN 0-13-482308-7
4. Åström, K.J & Wittenmark, Björn (1990): *Computer Controlled Systems: Theory and Design*, ISBN 0-13-172784-2

Appendix A: The Test Car

The test car used in the experiments was a Volvo 850 Turbo. It was equipped with an automated transmission, which is a requirement for an ACC system designed for stop and go situations due to the huge speed range of operation. A radar was mounted in the front, which measured the distance and the relative speed to the vehicle ahead. A yaw rate sensor determined in which direction the radar should scan. In this way, the radar could track the vehicle ahead during curve driving. The ordinary throttle was replaced with an electronical throttle and a brake booster was installed. This is necessary to be able to control the throttle position and brake pressure electronically. The actuator controller and the ACC controller were implemented in a PC computer. The signals from the sensors and the control signals from the computer were all communicated via a CAN bus.

The driver has the ability to easily activate and deactivate the ACC system. The desired speed and time gap (T_g) can also be changed at any time. For safety reasons, the driver could always override the ACC system by pressing the brake pedal.

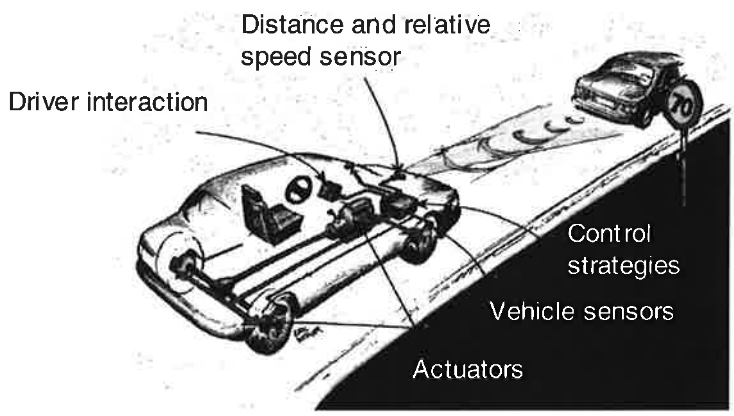

Figure A.1: The different parts of the test car.