TITLE A STUDY IN THE USE OF FUZZY LOGIC IN THE MANAGEMENT OF AN AUTOMOTIVE HEAT ENGINE/ELECTRIC HYBRID VEHICLE POWERTRAIN

AUTHOR Simon FARRALL

DEGREE Ph.D

AWARDING BODY Warwick University

DATE 1993

A study in the use of fuzzy logic in the management of an automotive heat engine/electric hybrid vehicle powertrain.

Simon Farrall.
Department of Engineering
University of Warwick

Thesis submitted to the University of Warwick for the degree of Doctor of Philosophy

Contents

Acknowledgements

I would like to acknowledge the help and support of many friends and colleagues in the work described in this thesis, and in the preparation of the thesis itself. Firstly, I would like to thank my academic supervisors, Dr. Peter Jones and Mr. Aubrey Corbett, of the Department of Engineering, at the University of Warwick, for their guidance and for giving me the benefit of their considerable experience.

I would also like to thank the Rover Group, with whose financial support this work was carried out. In particular, I would like to thank Mr. Ken Lillie and Mr. Carl Bourne at the Rover Advanced Technology Centre, at the University of Warwick, for their support and assistance.

Many other friends and colleagues in the Engineering Department, too numerous to mention individually, have been very helpful, of which, two people stand out. Mr. Alan Hulme has been responsible for the smooth operation of the computing facilities on which this work was performed, and has occasionally resurrected seemingly irretrievable information. Mr. Julian Matthews has been conducting research for a PhD over the same period as I have. Our contrasting backgrounds have lead to some "full and frank exchanges of opinion", but, on the whole, we have had some extremely informative and stimulating discussions, and I have greatly enjoyed working with him.

The help of other friends and family should also be acknowledged. I am grateful for their patience and support during the production of this thesis.

Declaration.

This thesis is original work by the author, parts of which have been published elsewhere, as stated in the text.

Summary.

This thesis addresses the problem of the instant-by-instant control of the powertrain of a hybrid heat engine/electric vehicle. In the absence of a prototype vehicle on which the work could be carried out, the work has taken the form of computer simulation experiments.

In order to develop the powertrain control strategies, a computer model of a conceptual hybrid vehicle is then developed, containing components from real, production and prototype vehicles. The use of this component based modelling approach allows the models to be validated by comparing their predictions with the performance of the real vehicles in which the components are used.

The previous work conducted in the field of hybrid vehicle powertrain control is then reviewed. It is found that fuzzy logic could potentially provide a means of controlling the hybrid powertrain in a realistic manner, in which some of the disadvantages of previous hybrid powertrain control strategies could be overcome. The results of initial simulation experiments are then reported, finding that whilst the basic method appears to have the potential to successfully control the powertrain, there is a need for an adaptive fuzzy powertrain controller.

A review is then presented of previous work conducted in the field of adaptive fuzzy control, finding that none of the reported adaptive fuzzy control methods are capable of being easily applied in the case of the hybrid powertrain. An adaptive fuzzy controller is then developed, whose rule modification strategy is specifically designed to work in the hybrid powertrain control problem.

This initial adaptive powertrain controller is then modified to improve its ability to control the overall performance of a hybrid vehicle, whilst maintaining vehicle driveability. It is found that this controller is able to adapt to the different driving styles of individual vehicle users within the space of a few simulated urban journeys.

Experiments are then performed in which improvements in the overall efficiency of the vehicle powertrain are investigated. It is found that significant improvements in the operation of the powertrain are impossible, due to some of the features of the vehicle model and constraints placed upon the control strategy.

Conclusions are then drawn, for the work done in the field of hybrid vehicle powertrain control and, also, for the work done in adaptive methods of fuzzy control. The most significant contribution in the field of hybrid powertrain control is the development of a controller that can adapt to the habits of different users. The most significant contribution in the field of fuzzy control is the form of the basic hybrid powertrain controller and the use of small fuzzy controllers in the powertrain controller adaptation strategy.

When I was eighteen I thought my father was a fool, by the time I was twenty-one, I was amazed how much the old man had learned in three years.

Mark Twain

To Julie.

Chapter 1

Introduction.

1.1 Vehicle emissions.

It has become generally accepted that the environmental effects of the way in which life is currently lead by the developed world are too severe to be allowed to continue. This is causing changes in behaviour at personal and national levels, as individuals alter their habits voluntarily and governments implement legislation aimed at reducing environmental destruction. In the case of the private passenger car, the environmental harm caused by pollution is not immediately obvious to the owner and modifications to make vehicles less harmful are expensive. The need to reduce the harmful effects of the use of cars is being tackled by the introduction of environmental legislation.

The car has been developed over a period of some one hundred years and in that time has gone from being the unreliable expensive possession of the very wealthy to the reliable and indispensable requirement of the majority of the population of the developed world. However, cars have obviously had a considerable environmental effect on the world, in their manufacture, use and destruction. It has been estimated [1] that approximately 20% of the fuel that can be associated with the life of a particular car is used in its manufacture. The procurement and processing of raw materials and waste products of manufacturing processes also

create pollution. When the life of a car is over, many of the metal components are recycled, but, much of the interior goes to landfill and more seriously, in the U.K., sometimes the lubricants have been simply emptied into the earth, although recent regulations make this far less likely.

The largest environmental effect of the use of cars is the emission of harmful exhaust gases. For a car with a petrol engine these gases are carbon monoxide, carbon dioxide, oxides of nitrogen and hydrocarbons. In the case of the diesel engine which is to be found in many cars and most commercial vehicles the variety of pollutants is far greater since the combustion process is very inhomogeneous. A detailed guide to the this subject can be found in [2].

The following very brief discussion of some aspects of vehicle emission regulations is included to give background information and show the external forces which are causing research in this area. It should be realised that emissions regulations are not the only force acting to cause radical changes in the use of vehicles. There are additionally regulations regarding safety, congestion, recycling, manufacture and so on that all have a bearing on the subject.

Vehicle manufacturers are forced to make their vehicles comply with emission regulations which become increasingly stringent as time goes on. The strictest emissions have always been mandated by the state of California because a combination of geography, climate and very high levels of vehicle usage have lead to the formation of a photochemical smog over the basin in which Los Angeles sits. Legislation which is enforced in California tends to be adopted by the rest of the United States and influences regulations over the rest of the world.

The regulations that have given impetus to the research reported in this thesis and a vast amount of other work are embodied in [3]. This regulation sets out various emission related types of vehicle and also dictates the proportion of the total sales volume which shall be made up of each type, see Table1.1.

	1997	1998	1999	2000	2001	2002	2003
Transitional low emission vehicle	73	48	23	0	0	0	0
Low emission vehicle	25	48	73	96	90	85	75
Ultra low emission vehicle	2	2	2	2	5	10	15
Zero emission emission vehicle	0	2	2	2	5	5	10

Table 1.1: Proportion of total sales (in percent) of different emission categories of vehicle for the six years after 1997 in California.

The details of the performance of the various types of vehicle can be found in [3], the important points are the volume figures for the ultra low emission vehicle and the zero emission vehicle. It is currently felt that it may not be possible to meet the requirements of the ultra low emission vehicle using a conventional internal combustion engine alone for propulsion and that it may be necessary to supplement the engine of the vehicle with some other means of propulsion that is less emissive. In the case of the zero emissions vehicle it will obviously not be possible to use a combustion engine burning fossil fuel at all. Whilst the regulations do not state that a zero emissions vehicle will be an electric vehicle, there is currently no other practical means of zero emission propulsion.

The ultra low emission vehicle is thought likely to be a hybrid vehicle whose power requirements will be met by the combined use of a heat engine and an electric motor powered by batteries.

Californian regulations have been structured to improve air quality and this has been done by limiting the emissions of the noxious exhaust gases, carbon monoxide, oxides of nitrogen and hydrocarbons. In Europe it is proposed that the emissions of carbon dioxide will be taxed, see [4]. Complete combustion of carbon based fuels produces carbon dioxide and water, the other exhaust gases being pro-

duced by imperfect combustion or undesirable side reactions. A regulation taxing the emission of carbon dioxide, therefore, effectively taxes fuel consumption.

Additionally, some European towns and cities are starting to consider imposing restrictions on the vehicles that enter their centres. It is felt that such restrictions could lead to zero emissions regulations in the future.

The reasons for restricting the emissions of carbon dioxide are different from restricting the noxious gases. Carbon dioxide contributes strongly to the greenhouse effect, and this motivates a limit to its production. Conversely, regulations that have been implemented to limit noxious gas production, to improve air quality, have had the effect of increasing carbon dioxide production. The regulatory bodies have tended to write emission regulations that moderate the pollution problem which they see as being most important. A consequence of the Californian regulations is that the emissions which will no longer occur in vehicles will be replaced by emissions in electric power generation. Whilst Californian electricity is very clean, this is not always the case for the rest of the world. There is considerable debate amongst environmental groups, regulatory bodies, vehicle manufacturers and electricity suppliers as to the extent of the environmental improvement that will be achieved by the various legislative proposals. A comparison of references [5] and [6] illustrates arguments that currently exist particularly where there are commercial vested interests.

Emission regulations have never been universally accepted as being beneficial in all aspects, however, when a proposal becomes law, the regulation has to be met regardless of the argument that will inevitably continue. This means that, if the regulations that are currently being discussed are adopted, the world wide motor industry will be forced to develop electric vehicles and possibly hybrid electric/heat engine vehicles.

1.2 Electric Vehicles.

In the early 1900s road vehicles could be bought powered by internal combustion engines, steam engines and electric motors. It is worth noting, in passing, that the earliest reference to a hybrid drive system is also around this time [7]. Initially, electric vehicles enjoyed greater popularity than internal combustion engine vehicles due to their quiet, reliable operation and ease of use [8]. By 1920 a reversal had taken place and internal combustion engine vehicles were being sold in their millions, whilst electric vehicle sales amounted to only a few thousands. The reasons for the reversal in fortunes were the increase in reliability of the internal combustion engine and the low performance and range of the electric vehicle. Ironically, the electric motor contributed to the downfall of the electric vehicle by its use in the electric starting systems of the internal combustion engine, the improvement in starting being one of the reasons for the increased popularity of the IC engine.

Since this time, the use of electric vehicles has been mainly limited to enthusiasts, see [9], leisure vehicles and industrial vehicles such as fork lift trucks. The main exception to this has been the large scale use of the electric milk delivery van in the U.K. which has given the U.K. something of a lead in the use of road going electric vehicles.

The reasons why electric vehicles have not been popular are that they are, in general, slow and have a very limited range when compared with their internal combustion powered counterparts. Additionally, battery lifetimes are short and, whilst electric vehicles are cheap to run over their lifetimes, batteries are expensive items to replace in one payment. In many other ways the electric vehicle has advantages over the IC engine vehicle, being quieter, simpler, more reliable and less directly emissive.

The extent of the problems that have to be overcome before electric vehicles are

comparable in performance with IC engine vehicles is illustrated by considering the rate of energy transfer when refuelling an IC engine vehicle. During the minute that it usually takes to refuel a conventional vehicle, energy flows into it by means of adding fuel, at a rate whose electrical equivalent is 10MW, the output of a reasonable sized power station [10].

The fundamental limitation of the electric vehicle is that no sufficiently dense means of on-board energy storage has yet been found. Sufficient power can be generated in an electric vehicle using lead acid batteries to give the vehicle acceptable performance [11]. Such a vehicle will then have a very limited range, especially if the maximum performance is used continuously. If other types of batteries are used which have a higher energy density, a much greater range can be achieved (though still inadequate) but the performance of the vehicle is lowered unacceptably.

In response to proposed emissions regulations, the electric vehicle, and in particular its battery technology, has been the subject of intense development. Despite the combined efforts of the motor industry, battery industry and electric power generation industry, no feasible technological solution has yet been found to overcome the limitations of charging and energy storage, nor is there any sign that such a solution will be found.

A further problem that will seriously affect the use of electric vehicles is that of heating and air conditioning which requires a power source of at least 5kW. Normally, this power is available in the engine coolant for heating or is taken off the crankshaft for air conditioning. Electric motors are highly efficient when compared with a heat engine, so very little waste heat is available for heating. Also, the power requirement cannot easily be met by the batteries as they tend to have an installed power of only around 20kW all of which is needed to propel the vehicle. Since there is a legal requirement to be able to heat a vehicle, and to

be able to demist the windows, some early electric vehicles have been fitted with diesel fuel burning heaters, which unfortunately prevent the vehicle from being a true zero emission vehicle.

Where a regulation requires the use of an electric vehicle the drawbacks will have to be accepted. Where a regulation requires the use of a zero emission vehicle only in certain regions or requires a vehicle whose emissions are lower than those of a conventional vehicle, the concept of a hybrid vehicle has some advantages.

1.3 Hybrid Vehicles.

A hybrid vehicle, as the hybrid electric/heat engine vehicle will be called in the rest of this thesis, is powered by a heat engine and an electric motor. The term heat engine is used to indicate that the possible sources of power from fossil fuel in a hybrid vehicle should not necessarily be limited to reciprocating internal combustion engines. The essential feature of this type of vehicle is that it can be used without creating any emissions by drawing power from the battery alone. When used in this manner, the vehicle will have the low range and performance of the electric vehicle. The heat engine can be used, where regulations allow, to give the vehicle acceptable range and performance.

This alternative obviously has the advantages of both the electric vehicle and the conventional vehicle. It should be pointed out that there are also disadvantages to this approach, for example, a hybrid vehicle will be more expensive and complex than either a conventional vehicle or an electric vehicle. The extra, physically large and heavy components make a hybrid vehicle heavier and more difficult to design.

Attempts have been made to justify the use of hybrid vehicles on grounds other than purely as a means of meeting emission requirements [12] but have generally met with a lack of interest by the traditionally conservative motor industry. Given the above disadvantages, there is no way that, without the use of legislation or a

very large change in the attitude of the public, vehicle manufacturers would ever be persuaded to make hybrid vehicles.

Hybrid vehicles tend to be found in two types, the serial hybrid vehicle and the parallel hybrid vehicle, [13] gives a very readable general introduction to the two types.

1.3.1 The serial hybrid vehicle.

In the serial hybrid vehicle the heat engine is used to drive a generator which in turn feeds power into a circuit consisting of batteries, control electronics and a motor which drives the vehicle. There is no direct mechanical connection between the heat engine and the wheels of the vehicle. A diagram illustrating this concept is shown in Figure 1.1.

The serial hybrid concept has a number of advantages. The heat engine can be run at a constant speed since it is not coupled to the wheels. This allows its design to be more optimal and its operation to be more efficient in terms of the power generated per unit of emission and per unit of fuel. The vehicle designer has some freedom over the types of heat engine and generator since the speeds at which they rotate are not fixed. Also, the dynamic response of the heat engine is not important, again allowing other types to be considered. The absence of a linkage between the engine and the wheels of the vehicle allows some freedom in the layout of the vehicle and finally, the heat engine does not have to be able to meet the instantaneous vehicle power requirements as the battery acts as a "load leveller" in transient operation.

The drawbacks of the serial hybrid vehicle are that it will be expensive since it has two electric machines and associated power electronics. The control task in a hybrid vehicle is complex as two interacting electrical machines, an engine and a set of batteries have to be managed. Possibly the largest disadvantage of the

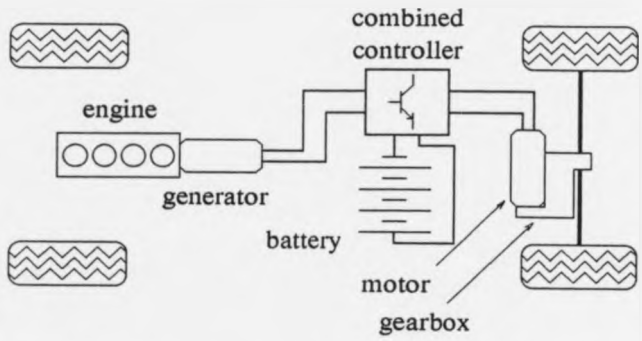

Figure 1.1: Schematic diagram of a serial hybrid vehicle.

serial hybrid is the energy path potentially passes through the heat engine, and then through two electric machines before reaching the road. The efficiencies of the devices are then multiplied together resulting in a low overall efficiency for the powertrain.

The serial hybrid concept is most sensibly embodied using a small heat engine that can be worked hard in order to operate efficiently in terms of fuel used and emissions created. The vehicle is always propelled by an electric motor and, consequently, the serial hybrid vehicle can be thought of as a range extended electric vehicle.

1.3.2 The parallel hybrid vehicle.

In the parallel hybrid vehicle both the engine and the motor of the vehicle can be used to drive the wheels. This may be done via some shared driveline components such as gearboxes, differentials and driveshafts or independently by driving different sets of wheels. A diagram showing one implementation of this concept (there are many) is included as Figure 1.2.

The advantages of the parallel hybrid vehicle include relatively low cost and complexity since there is only one electric machine with associated, reasonably

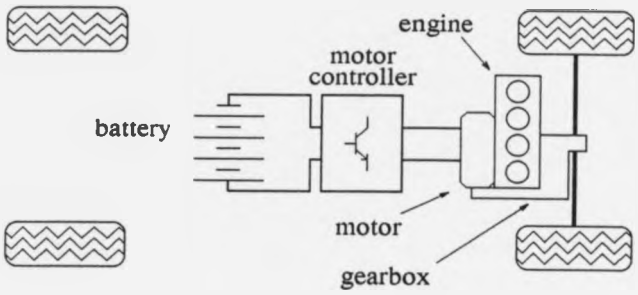

Figure 1.2: Schematic diagram of one type of parallel hybrid vehicle.

simple, electronics. The engine is able to drive the wheels directly, allowing the fuel to be used efficiently. A parallel hybrid will be relatively straightforward to design since it has a great deal in common with both the conventional vehicle and the electric vehicle, and it shares the load levelling ability that batteries give the serial hybrid vehicle.

There are only two real disadvantages of the parallel hybrid vehicle. The engine turns at a speed dictated by the speed of the vehicle and, as in a conventional vehicle, its design must be compromised to work over a large speed range. There is little freedom about the location of components since they are all mechanically connected to the wheels.

Practical parallel hybrid vehicles have tended to use conventional vehicles with additional electric motors. This reflects the fact that parallel hybrid vehicles tend to be seen as conventional vehicles which have the ability to operate in a zero emission zone.

1.4 Hybrid vehicle operation.

Under many circumstances, the correct method of operating a hybrid vehicle will be obvious. For example, when the vehicle is being used in a zero emission zone power should be drawn from the batteries alone. Conversely, when the vehicle is being used for long journeys on motorways which require long range and large amounts of power, the power for propulsion will be drawn from the fuel. In addition to this, there may well be a need to allow the driver of the vehicle to force it to operate using only the batteries and thereby obtain the benefits of electric propulsion including quiet and cheap operation and possible or perceived environmental responsibility.

The previous paragraph discusses the situations that have forced the hybrid vehicle to be seriously considered. In general use, there will be situations in which the vehicle should make use of both sources of energy simultaneously. This would occur, for example, when the maximum acceleration of the vehicle was required for overtaking. A further need for simultaneous operation would be if the vehicle were approaching a zero emissions zone with a low battery state of charge which would need to be increased to allow the vehicle to cross the zone. Increasing the battery state of charge could be done by working the motor against the engine, whilst the engine powers the vehicle.

Many users of cars tend to have patterns of usage that are essentially the same from one day to the next. For such users there may well be optimal ways of using the energy stored in the vehicle. In general, electric propulsion provides a cheap means of transport and it might therefore be desirable to use the batteries as much as possible in daily use.

When the two sources of energy on the vehicle are being used simultaneously the driver will not be able to individually control the motor and the engine during normal driving. Some means of automatically controlling the powertrain will

have to be provided. Engine and motor control would be achieved by means of controlling the throttle on the engine and the currents flowing in the motor.

This controller should cause the stored energy in the vehicle to be used in some energetic or cost optimal way. The control actions should also increase or decrease the battery state of charge, as required by the driver. As these requirements change, the controller should function such that the vehicle drives in a consistent manner and has good driveability at all times. The requirements of various users may be such that different control actions should take place, so some adaption of the controller may be desirable. Lastly, the controller should allow the vehicle to meet the emission regulations of the market into which it is being sold.

1.4.1 Hybrid vehicle powertrain control.

The aim of the work presented in this thesis is to develop methods of controlling or managing hybrid vehicle powertrains, in such a way that meets the requirements of the items listed in the previous section.

A detailed discussion of the previous work done in the field of hybrid powertrain control is given in Section 3.2.4. At this point it is sufficient to note that there is no established procedure or technique that has been used in this field. The reasons for this are twofold. Firstly, very few hybrid vehicle prototype cars have been built and consequently, as a research area, the control of hybrid powertrains is in its infancy. Secondly, the commercial success or otherwise of a particular hybrid vehicle will be heavily influenced by the performance of the powertrain controller, making its form extremely commercially sensitive. Where the research has been carried out by industry the results are, therefore, not widely known.

Section 3.2.4 contains a detailed consideration of the form that a hybrid powertrain controller should take. It is found that, because of the unconventional problem formulation and its highly intuitive nature, the use of fuzzy logic is helpful.

As will be seen in Section 1.5 fuzzy logic is increasingly being used in automotive control and this also makes the approach attractive.

1.5 Automotive Control.

The purpose of this section is to place in context, the work that is subsequently reported and give background information to assist in its understanding.

1.5.1 Application Areas.

Reference [14] provides an excellent general source of automotive engineering information and would be of use to anyone interested in automotive control.

The most commonplace controller found on vehicles is undoubtedly the engine controller or controllers. The modern engine controller or engine management system, as it is more commonly known, has evolved over the last ten years or so. Petrol engine control operations can be broken down into fuel control and ignition control. The first improvement in engine control resulting from the application of electronics was the introduction of programmed ignition. Unlike breakerless ignition, which switched the ignition using solid state devices, programmed ignition also controlled the ignition timing..

Another early improvement was the introduction of devices aimed at improving the cold starting and cold running of engines. These devices were normally stepper motors which acted on the mixture control and throttle stop of the carburettor. Since the control involved was so minimal very little electronic complexity was required and this really was a method of increasing the convenience of the car rather than improving it.

A very large increase in complexity occurred when the carburettor was replaced by fuel injection. The first systems were analogue constant flow systems, whose chief advantage was that the air-fuel ratio could be set at any demanded value

electronically rather than mechanically by carburettors. Constant flow systems were soon replaced by systems in which the injectors fired at intervals, injecting a measured quantity of fuel into the air flowing into the cylinder. Initially, fuel injection systems were kept separate from ignition systems, and it is interesting to note that typically fuel injection controllers had around 20 connections and ignition controllers had around 10 connections, giving some idea of the relative complexity of the control tasks.

Fuel injection systems and ignition systems were combined to create the engine management system. Subsequently, the most significant increase in complexity has been the use of closed loop fuelling and the three way catalyst (so called because it greatly reduces the amounts of the three major exhaust gas pollutants). In practice, an exactly stoichiometric air fuel mixture is required to reduce the emission of noxious gases to extremely low levels. Very accurate air fuel ratio control is a demanding task because the controller has to maintain an exactly stoichiometric air fuel ratio, in the presence of abrupt changes in air flow rate caused by the driver suddenly opening or closing the engine throttle. The air fuel ratio is indicated by sensing the composition of the exhaust gas using a lambda sensor to provide a feedback signal. The control problem is further complicated by the evaporation of fuel sitting in the inlet manifold which changes the air fuel ratio in the cylinder.

In addition to the control of air fuel ratio, engine management systems control the engine idle speed. Again, this takes place in the presence of disturbing loads such as power assisted steering pump loads, air conditioning pump loads and sudden electrical loads. Recent additions to engine management have included turbocharger control, variable cam timing and variable inlet manifold geometry. Future engine management systems will have to comply with [15] which requires that misfires are detected as the engine is used and the operating conditions

prevailing at the time are stored. Misfire detection is relatively straightforward with engines which have a small number of cylinders at high loads and low speeds, but is very difficult to detect in engines with many cylinders at high speeds and low loads without introducing additional sensors. A recent advance is the introduction of engine management systems to control the fuel injection of diesel engines.

Of the other vehicle control applications the next most complex and common are anti-lock braking systems which work by means of calculating the angular acceleration of each road wheel. If the wheel acceleration corresponds to an effective vehicle deceleration of greater than around 1g, this is interpreted as the wheel locking, and the brake hydraulic pressure to that wheel is then momentarily reduced. The function of the anti-lock braking system is sometimes complicated by its use in combination with the engine management system to form a traction control system that limits wheelspin on slippery surfaces. The torque output of a vehicle differential is equal, or very nearly equal, on both sides. If the brakes are applied to a spinning wheel then engine torque equal to the braking torque may be passed to the other wheel allowing it to drive the vehicle.

Other common areas of control in vehicles are in automatic transmission control, climate control, cruise control and suspension control. Cruise control, climate control and transmission control are, when compared to modern engine control, relatively straightforward. Suspension control is very dependant upon the types of suspension system being considered, and these vary from fully controllable systems using active elements to conventional suspension systems containing some variable elements. Examples of the latter include variable rate damper systems which are being produced by a few manufacturers, including Rolls-Royce. A discussion of various suspension systems is given in [16].

1.5.2 Automotive control techniques.

Despite large amounts of research having been conducted, see [17] and [18] for a comprehensive list of references, automotive control has seen very few production applications of conventional linear or advanced linear control methods. The reason for this lies in the nature of automotive control problems. In the application areas discussed above, closed loop fuelling, idle speed control, suspension control, cruise control and climate control can be posed in the form of a conventional control problem. Conventional control being taken to mean controlling the output of a system to track a reference value and, in the case of feedback control, using the error between the demanded output and the actual output as the control signal.

Cruise control and climate control are generally handled by relatively straightforward PID type controllers. Conventional control work on suspensions, whilst being an area of active research, has not, to the authors knowledge, been used in a production vehicle, though it has been used in racing cars. Conventional control could make a practical contribution in the field of engine control, however, in this and other areas, conventional linear control methods, especially the advanced linear control methods, have a huge disadvantage. Automotive engineers tend to be specialists in restricted areas of automotive technology and use heuristic methods which reflect their experience and understanding in preference to highly theoretical techniques, see [19].

Automotive control applications have therefore tended to use rule based methods or maps, which have been successful because their methods of operation are obvious, making them quick to develop and relatively easy to understand. The same considerations motivated the development of fuzzy logic and related intelligent control techniques and it is therefore unsurprising to see the growth of automotive fuzzy logic controllers at both a practical and research level, see [20] and [21].

1.5.3 Automotive control hardware.

Emissions regulations have caused huge growth in the number and complexity of engine management systems and this, in turn, has added impetus to the development of the microcontroller. The commercial effect of this should not be underestimated since the number of cars with engine management systems must make this the largest single control application in the world.

The microcontroller differs from the microprocessor in that as well as containing computational and storage areas it also has, on the same device, functions such as analogue to digital conversion, digital to analogue conversion, timers, serial communication ports and useful amounts of memory. Whilst these devices are not controllers on a chip, they do go a long way towards this. The purpose of these peripherals is to allow the processor to restrict its activity to the calculation of output variable values and interpretation of inputs, rather be involved in the exact timing of the output actions which is so critical in engine management and other applications. As requirements of engine management systems increase, the ability of the peripheral devices to run the engine with minimal involvement of the processor core, allowing the core to concentrate on more computationally intensive tasks, will become ever more important.

The sheer volumes of the automotive industry have meant that microcontrollers have become tailored to meet certain control requirements, engine management, of course being the most important. Details of one family of microcontrollers can be found in [22]. Some idea of the increase in complexity that has occurred is given by considering that, in around fifteen years, engine management systems have gone from using analogue circuits, to 32 bit microcontrollers with very large numbers of peripherals on the same device.

The other important advance in automotive control, is the communication controller. Details of one of many communication protocols may be found in [23]

with decsriptions of devices that meet this requirement being found in [22] and [24]. The purpose of these communication controllers is to allow the interaction of co-operating systems such as engine and transmissions controllers. The maximum speed of the communication is, in automotive terms, very fast at around 1Mbit/s and this allows genuinely interactive control rather than simple information sharing.

1.5.4 Costs.

It is well known that the automotive industry is very cost driven. In control terms this is reflected in the prices of actuators, sensors and controllers. Most automotive sensors have a component cost in the region of pence to a few pounds, examples including reluctance sensors, pressure sensors and temperature sensors. Actuators, which tend to be devices such as solenoid valves, stepper motors, fuel injectors and spark plugs are rather more expensive, however, the most expensive of these is still only around fifteen to twenty pounds. The exceptions to this tend to be in the area of suspension control and it is the high cost of devices such as accelerometers and actuators that has contributed to the lack of production applications. The control electronics used in vehicles tend to have piece prices from a few tens of pounds to a few hundreds of pounds depending on the application.

1.6 Objectives of this thesis.

Research effort in hybrid vehicles has tended to be concentrated in areas such as battery technology, motor technology, alternative engine technologies and vehicle design and construction. Considerable work has been done on component sizing and matching, but this has tended to be at a rather high level. The instant-by-instant operation of a hybrid vehicle has not been the subject of a great deal of research, the work that has been done is discussed in Section 3.2. In practice, it is

impossible to operate a hybrid vehicle without some form of powertrain controller and, as hybrid vehicle prototypes are produced in increasing numbers, the need for understanding of their operation will increase.

This thesis describes work done in developing control strategies for a particular hybrid vehicle powertrain. The practical aspects of automotive control are considered throughout the work reported. This means that the control methods considered must be capable of being implemented on a low cost microcontroller, they must not involve expensive sensors or actuators and they must be capable of being included in an automotive development cycle. The controller decided upon must be capable of being developed quickly and robustly. The operation of the controller must be easily understood without complex theory both because automotive engineers tend to be component specialists and not systems specialists and also because straightforward software operation is advantageous from a product liability point of view.

The absence of a suitable vehicle on which to develop control strategies meant that the work was carried out using computer simulations. During the course of the work it has been realised that the correct control approach will vary depending upon the type of vehicle being studied. The use of computer simulations allows vehicle models to be quickly and easily changed so that control strategies for different types of vehicle may be developed. Additionally, all the usual advantages of modelling controllers before implementing them in hardware have been gained. It is intended that the work described will subsequently be used in hardware on the powertrain of an actual hybrid vehicle under development.

It will be seen in Section 3.3 that the control problem has aspects of intelligent decision making, and cannot be formulated in a conventional manner. This makes fuzzy logic a natural choice for the control strategy. Fuzzy logic also has the advantages of being an excellent tool for automotive component and system

specialists, allowing them to embody their experience and understanding in the control strategy. It can be shown, again see Section 3.4.3, that a fuzzy controller can be reduced to a form that is identical to the maps used in modern engine management systems, making the method readily implemented by automotive software engineers. As a result of the use of fuzzy logic in the controllers, some conclusions have been drawn about the role of fuzzy logic in modern control engineering and, in particular, in automotive engineering.

A first requirement before controllers could be developed was a suitable simulation environment. In order to simulate a hybrid vehicle component models had to be developed. These models, their combination to form vehicle models and the validation of these vehicle models formed are the subject of the next chapter.

Chapter 2

Hybrid vehicle powertrain modelling.

2.1 Introduction.

This chapter describes the development of a set of hybrid vehicle component models whose purpose is to provide an environment in which the control of a hybrid vehicle powertrain can be investigated. The aims and objectives of the model development are first stated and the modelling approach is explained. A potential, practical hybrid vehicle concept is then described which is the basis of the modelling work.

Component models are then developed for the major components in the vehicle powertrain and the vehicle itself. The completed model is then discussed and details of validation exercises are presented. Some of the work discussed in this chapter has been published elsewhere, see [25, 26].

2.2 Aims and objectives.

Figure 2.1 is a schematic diagram of a parallel hybrid vehicle in which it is seen that the role of the powertrain controller is to interpret the movements of the driver's accelerator pedal and generate demand signals for the engine throttle angle and

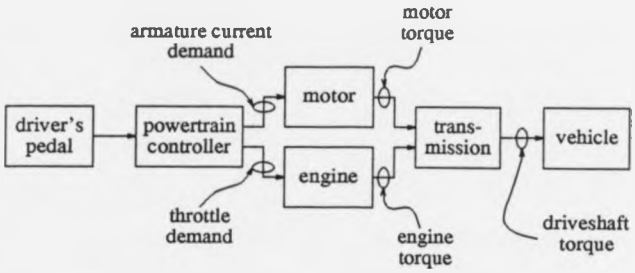

Figure 2.1: Schematic block diagram showing the role of the powertrain controller in a parallel hybrid vehicle.

the motor armature current. These signals should blend the operation of the motor and the engine so that the vehicle performance meets various requirements without the driver being involved in the detailed control of both the engine and the motor.

Since emission regulations have forced hybrid vehicles to be seriously considered, a primary requirement of the powertrain controller will be to ensure that the vehicle which it controls will meet emissions standards which are defined over drive cycles that have durations of tens of minutes. As will be seen in Section 2.4, the modelling of heat engine emissions is extremely difficult and involves processes that have timescales in the order of tens of milliseconds. In an ECE+EUDC European emissions test, about 70% of the carbon monoxide emissions occur during the first two minutes of a twenty minute test. This is because fuel enrichment must be used until the engine is warm and the emissions are not closely controlled until the catalyst "lights off", or begins converting, which occurs when it reaches a temperature of about 290°C. Currently, very little work has been done on cold running emissions or thermal models of engines and catalysts. Also, as described below, feedback fuel injection systems have been modelled, but the models that result are complex and also give little indication of the transient emissions performance of the engine. The inherent difficulty in modelling the processes that lead

to the formation of pollutants, and the mismatch between the process timescales and the duration of the test cycles, makes the modelling of vehicle emissions over drive cycles excessively complex. A modelling study of very much higher complexity than that reported here would be required to predict vehicle emissions or, more sensibly, work should be carried out on a prototype vehicle.

This is not as great a drawback as might at first be thought. The type of vehicle considered is likely to find itself in either zero emission zones or areas in which the heat engine is intended to be the primary source of motive power and would therefore be able to meet the relevant emission requirement. Also, current petrol engine emissions are dominated by cold starts and cold running, which contributes strongly to the difficulty in modelling engine emissions, since the actions of the engine controller are very complex under these circumstances, again see Section 2.4.

Having removed the impractical aspects of the modelling exercise the more tractable tasks of the powertrain controller can now be defined. The vehicle should perform so that its energy resources are used in the correct proportions relative to one another. For example, for economic reasons, the driver of the vehicle may wish the vehicle to obtain most of its energy for propulsion from the batteries rather than from fossil fuel. Alternatively, the driver may wish to increase the battery state of charge, at the expense of using extra fuel, to allow the vehicle to cross a zero emissions zone.

A second requirement of the powertrain controller is to meet the demands of the driver in a manner which is energetically efficient. This means that operation of the motor and the engine are combined in such a way that the vehicle uses a minimum of energy (or cost) to meet the requirements of the journey being undertaken.

The last requirement of the powertrain controller is that the vehicle should

drive in a consistent and smooth manner regardless of the bias in the use of energy. The subjective aspects of the way in which a vehicle drives are generally termed driveability, a full consideration of which requires an understanding of factors such as transmission lash, driveline compliance and other vehicle characteristics which are undefined until the vehicle is actually designed. The timescales over which the effects that influence driveability occur are, again, very short. The consideration given to driveability in this work is that the driveshaft torque in Figure 2.1 should be the same for any controller when the same accelerator pedal value is used at the same engine speed.

The requirements set out above, and the timescales of journeys, define the level of detail of the powertrain component models that should be used. To determine whether a vehicle will meet the requirements of a journey, the primary modelling requirement is to be able to relate the use of the energy resource to the generation of power. As will be seen in Sections 2.4 and 2.5, this can be done without considering the dynamic operation of the components. This would not be the case if engine emissions were to be considered explicitly.

2.3 Modelling philosophy.

In powertrain modelling, there have generally been two approaches adopted which can be described as forward and reverse dynamic models. Figure 2.2 shows, in block diagram form, the information flows in the hybrid vehicle model developed in this chapter.

2.3.1 Forward dynamic models.

In the forward dynamic modelling approach, the input-output information flow is essentially causal, the inputs to a component model tending to be the equivalent input to the component in the real world. An example of this might be an engine

model, in which the inputs might include throttle angle, engine speed and coolant temperature and the outputs might include engine torque and engine fuel flow rate. The engine torque would then cause a torque to be applied to the input of the gearbox which would cause a torque to appear in the driveshafts and this, in turn, would cause the vehicle to accelerate.

In this work, the forward dynamic modelling approach has been taken. This is because, at the time the work was begun, the form of the powertrain controller was completely unknown, and it was important that the powertrain model did not in anyway constrain the type of controller that was subsequently to be developed (see Section 2.3.2). Additionally, if the model functions in the forward dynamic sense, then so will the controller. This means that the controller used in the model will operate in a very similar way to the controller on the vehicle, making the vehicle controller design more straightforward. The forward dynamic aspects of the model operation are shown, in Figure 2.2, by the arrows moving towards the vehicle dynamics block.

2.3.2 Reverse dynamic models.

A reverse dynamic model reverses the causality of the original system. When using a reverse dynamic model, the output of the original system is specified or calculated by some means, the system model then calculates the physical input that would be required to cause the physical output to occur. Taking the example of the engine once again, the output torque of the engine would need to be calculated and then a reverse dynamic model would determine the throttle angle required to cause the engine to generate the output torque, at the engine speed concerned.

Clearly, coping with dynamic systems is more difficult when using reverse dynamic models and, since the aim of much simulation and modelling work is

to address dynamic problems, reverse dynamic modelling is not commonly used. However, as will be discussed in Section 2.5, reverse dynamic modelling has been used extensively in the simulation of electric and hybrid vehicles by other authors. It was not used in this work because it was felt that it might compromise the controller design.

Aspects of reverse dynamic modelling appear in Figure 2.2, since the rotational speeds of the various vehicle components are propagated back through the powertrain.

2.3.3 A conceptual hybrid vehicle.

Prototype or production vehicles and components have been used as a basis for a conceptual hybrid vehicle, which does not actually exist but which could, relatively straightforwardly, be constructed. By using existing vehicles as a modelling base, component models can be validated against vehicles by comparing performance predictions from the models with actual vehicle performance. The imaginary vehicle that forms the basis for the work reported in this thesis is a small parallel hybrid vehicle, which would be designed to meet a zero emission regulation in towns and cities but which has only to meet a more conventional emission regulation elsewhere.

The electric powertrain components are the same as those used in the Metrolec electric vehicle which is described below. The heat engine is the engine used in the Rover Metro GTi, a 1400cc, 4-cylinder, 16-valve fuel injected petrol engine whose emissions are controlled using a 3-way closed loop catalyst. It was chosen because of its recent design which makes it as similar as possible to an engine that might be found in a hybrid vehicle of the future.

The Metrolec electric vehicle is described in detail in [27]. It is a Rover Metro car that has been converted to run on electric energy. The vehicle is propelled

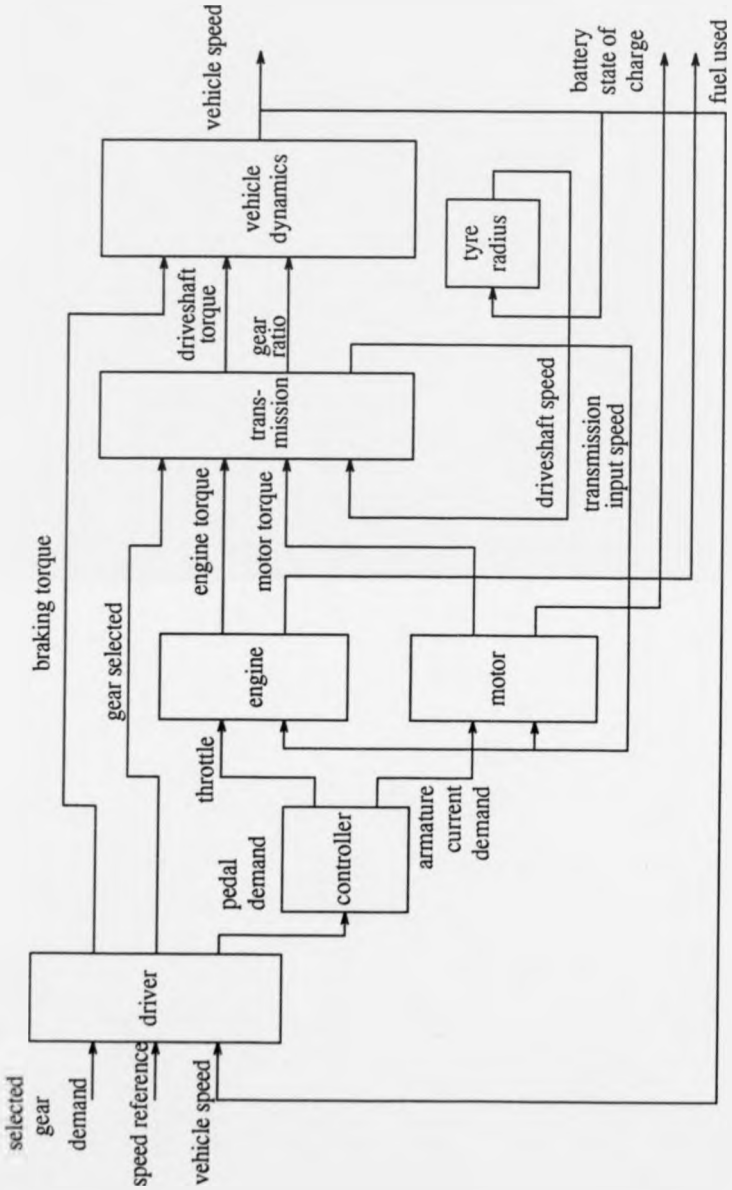

Figure 2.2: Block diagram of complete hybrid vehicle model.

using a separately excited brushed DC motor of a type described in [28]. The operation of the motor will be discussed in detail in Section 2.5.

A novel feature of the Metrolec is its use of sodium sulphur batteries which were first investigated in 1967 see [29]. It is a measure of the development effort brought about by Californian regulations that this technology has only recently started to be used in prototype applications. The major disadvantage of the batteries used in the Metrolec [30] and other high temperature battery types, [31, 32] is that they must be maintained at a temperature of around 300°C. At the time at which the Metrolec was built, and the work described here was begun, high temperature batteries looked quite promising, however, recently the safety and reliability of these battery systems has been questioned.

2.4 Hybrid vehicle engine model.

Engine models vary enormously in complexity and method of operation, due to the differing applications for which they are developed. At the most numerically intensive end of the spectrum, there are models which simulate combustion by modelling the gas flows in the cylinder head and combustion chamber. This can then be used as a tool in the design of the cylinder head, camshaft and so on. Some of these models use computational fluid dynamics techniques, and models of combustion and pollutant formation, to calculate the properties of each cell in a mesh that covers the combustion space. Such models are extremely numerically expensive, involving varying meshes of upto a million cells and having an integration step size of one degree of crankshaft rotation. Simulating three hundred degrees of crankshaft revolution takes in the order of days on modern supercomputers. See [33] for a description of this approach. There are other less numerically expensive methods of flow prediction that were used before computers advanced to the stage which made computational fluid dynamics feasible, see

[34, 35].

Another reason for the development of engine models has been to assist in the design of engine controllers. A comprehensive review of the models used in this application is given in [36], whilst [37] also includes a discussion of the subject. The models described in these references have tended to be used for the optimal calibration of engines over driving cycles, a typical requirement being to minimise the use of fuel in the presence of pollutant inequality constraints whilst maintaining acceptable driveability, [38, 39, 40]. Models used for controller design tend to be less numerically intensive than flow and combustion models and are often continuous time or discrete time models with between two and ten degrees of freedom.

Since the work that was reviewed in [36, 37], emissions regulations have been made much stricter so that very different emission control methods have been required. Emission control methods using three way catalysts require a degree of accuracy in air fuel ratio control that was previously completely unachievable and unnecessary. Much of the current engine modelling work has been motivated by this need to improve air-fuel ratio control. Inlet manifold wall wetting and subsequent fuel evaporation form a complex fuelling dynamic, and there are also dynamics in the air flow into the manifold. These effects have been modelled by work such as [41, 42]. An alternative low emission engine management strategy is lean-burn technology and an example of modelling and control work in this area is given in [43]. Another major current engine control problem is idle speed control, an example of engine modelling used to design idle speed controllers being found in [44].

Recently, the use of turbochargers to give performance improvements has become popular. However, turbocharged petrol engines have an inherent lag problem and [45] describes modelling and simulation work that attempts to reduce lag to

improve the engine transient response.

In [46], which describes engine modelling work used in hybrid vehicle modelling, the accurate closed loop control of engine torque was felt to be important. Since measurement of engine torque on a practical vehicle is very difficult, a dynamic model of the engine was developed for identification purposes and predicted transient responses that compared well with observations. The transients introduced by sudden throttle openings had a duration of around 200 ms.

This degree of detail was not considered necessary in the engine model used in this work for the reasons given in Section 2.2 and, because the engine used in the work reported here uses fuel injection rather than carburetion, which makes its transient response quicker.

The engine modelling approach adopted in this work uses static engine maps and has previously been used in [47, 48]. Of the vehicle fuel consumption predictions given in [47], the least accurate agrees with the official figure for the vehicle concerned to within 10%, the next worst being 3%. Similarly accurate results have been achieved by vehicle manufacturers using the same approach.

The case for the use of this method for the hybrid vehicle engine model is improved because of changes in engine technology since the work reported in [47] was carried out. The engine used in [47] had a carburettor, whereas the engine which forms the basis for this work had a closed loop fuel injection system. The use of feedback fuel injection means that variations in air-fuel ratio between different engines are almost eliminated. A possible source of errors in using the map based engine model are that the map is obtained at steady state operation but the engine undergoes considerable transient operation on fuel economy cycles. In the case of the carburettor engine, these transients result in air-fuel ratio variations which would not be predicted using the steady state map. The purpose of the closed loop fuel injection system is to maintain close control over air-fuel ratio under all

conditions and, consequently, better results should be expected by the model in this work, than that in [47].

The engine torque, T_e and fuel flow rate, $\dot{m}_f$ are functions of engine speed, ω_e and throttle angle, θ.

$$T_e = f(\omega_e, \theta) \tag{2.1}$$

$$\dot{m}_f = g(\omega_e, \theta) \tag{2.2}$$

The data used to form the static Metro GTi engine maps which represent these functions are shown in Figures 2.3 and 2.4.

The data in these maps was obtained by measurements from engines. When mapping engines there are often areas in which measurements are not taken because of instability of engine operation or because data is not needed in that particular area. In the case of the maps used in this work, the low throttle angle-high engine speed area of the map had no data. In order to populate the matrices used in the simulation work the data in this area was fitted by eye. This simple method of adding data can be justified because, in practice, engines hardly ever enter the low load high speed region. When an engine does enter this region, the car is usually being slowed down by the brakes and so the small amounts of negative torque (known as overrun torque) produced in this area are not important. Furthermore, in this area of operation, the fuel flow rate is generally very low and is zero above a certain engine speed, when the throttle is closed because the engine management system chooses not to fire the fuel injectors under these conditions.

The engine load axis is expressed in terms of throttle angle. A more conventional way of expressing engine load is inlet manifold pressure, which is generally preferred because it is independant of the throttle body fitted on the engine and, also, is used as input information in some engine management systems. However, inlet manifold pressure cannot be easily manipulated as a controlling input be-

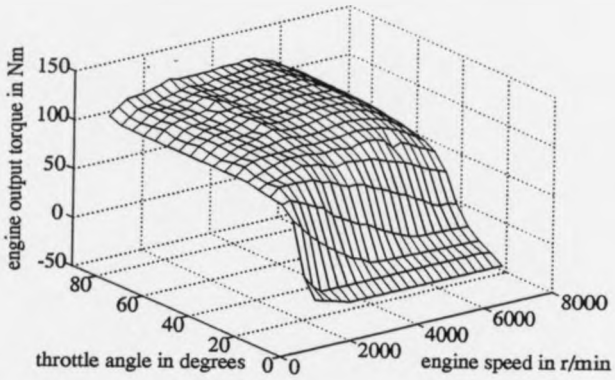

Figure 2.3: Engine torque map for K-series 1400cc 16 valve engine.

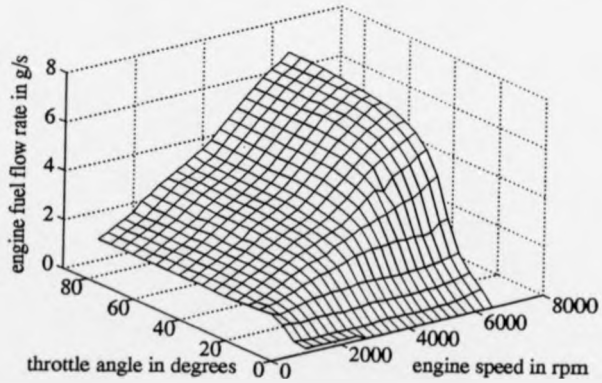

Figure 2.4: Engine fuel map for K-series 1400cc 16 valve engine.

cause it is a function of engine speed as well as throttle angle. Throttle angle can be easily and directly controlled by a simple position controller, and this is frequently done in vehicle cruise control systems.

It can be seen that the "origin" of the maps is a throttle angle of zero degrees and an engine speed of 1000 r/min. This condition is intended to represent the engine at idle and, as such, is a considerable simplification. At idle, the throttle is a few degrees open, and this opening will vary, depending on the idle speed demand (idle speeds are higher at low coolant temperatures), additional loads at idle (such as power steering, electrical loads, air-conditioning and engine internal friction) and will also vary from engine to engine, being dependant upon fine mechanical dimensions. Using a notional throttle angle of zero at idle allows an evenly spaced vector of throttle angle values to be used in the definition of the space over which the engine map is defined and, in practice, the actual idle throttle angle could be used as an offset into the throttle angle input value. Also, the idle speed is generally lower than the 1000 r/min implied in Figures 2.3 and 2.4 but the same arguments apply for the engine speed aspect of the idle condition as for the throttle angle. The important aspects of idle in this work are the engine torque of zero and the idle fuel consumption, both of which are adequately represented by this method.

Interpolation is used to obtain output variables from the maps, taking engine throttle angle and engine speed as input variables. These two input variables form a model input space which is segmented using vectors of throttle angle, and engine speed, and form a grid over which the output variables are defined. Once the region of the grid in which the operating point is defined has been determined, bilinear interpolation, see [49] is used to obtain an output estimate. A smoother method of interpolation would be given by bicubic interpolation, however, the data in the maps was obtained from engine tests and, therefore, contains measurement

noise. Using bicubic interpolation would be inadvisable when there may well be very poorly conditioned data for polynomial interpolation, and the improvement in accuracy would be negligible when it is borne in mind that considerable variations between the full load torque of individual engines, are typically around 5%. The data used in the engine model is included in Appendix A.

2.5 Hybrid vehicle traction motor modelling.

The majority of work in the area of electric and hybrid vehicle powertrain modelling has used a map of the motor as the motor model. Use of efficiency maps is made in [47, 48, 50, 51, 52, 53], and use of torque maps is made in [54, 55], while [56] uses a combination of the two types of map. However, [57, 58, 59, 60] use a method that considers the physical processes that take place in the motor, though only [60] describes the model in any detail. In [60], the work was motivated by the desire to be able to predict the motor torque in order to be able to accurately control it. Attention was paid to representing the dynamic aspects of the model and, when compared with measured values, the results are very encouraging. The motor takes about 300 ms to settle after responding to a step change in set point. For the reasons set out in Section 2.2, it was decided not to attempt to model the dynamic aspects of the motor operation in this work.

Much of the material below is similar to [60] but it differs in that the actions of the motor controller are different and the traction batteries appear in the model. Also, the regenerative actions of the motor are considered in more detail in this work. An excellent general text on DC machines is [61].

The purpose of the traction motor model described below is to predict the steady state torque generated by the motor and also the total current that the motor draws from the battery. The inputs to the model are a demanded armature current and the motor speed.

The fundamental equations that govern the operation of a DC motor are:

$$T_g = k_m \Phi(I_f)\, I_a \tag{2.3}$$

$$E = k_m \Phi(I_f)\, \omega_m \tag{2.4}$$

where T_g is the motor gross torque (the torque developed in the armature before any mechanical losses), $\Phi(I_f)$ is the field flux, which is a non linear function of field current, I_a is the armature current, E is the armature voltage, ω_m is the motor speed and k_m is the motor constant. If all the quantities are measured in SI units, then k_m is the same in both equations.

The function $\Phi(I_f)$ is non-linear because of magnetic saturation of the material that makes up the armature and the field cores of the machine. In practice, obtaining values for $\Phi(I_f)$ is very difficult, and unnecessary, since the term $\Phi(I_f)$ does not appear on its own in Equations 2.3 and 2.4. Conversely, the measurement of $k_m \Phi(I_f)$ is relatively straightforward, and is done by spinning the motor at various speeds whilst applying a field current to the field winding and measuring the armature voltage, V_a, which is the same as the back EMF, E for zero armature current. Using Equation 2.4, the function $k_m \Phi(I_f)$ can be obtained. This procedure, and various other measurements, were carried out on the motor fitted to the Metrolec vehicle and are described in [62]. The function $k_m \Phi(I_f)$ is plotted against I_f in Figure 2.5.

The motor model in this work is, really, a model of the entire electric powertrain and has been developed by considering the circuit formed by the motor, the motor controller and the batteries. This circuit is shown in Figure 2.6.

The symbols in Figure 2.6 are explained as follows. The battery voltage and internal resistance are V_s and R_s, and C is a smoothing capacitor. The field and armature winding resistances and inductances are represented by R_f, L_f, R_a and L_a. The values of these components in the Metrolec are given in Table 2.1. The

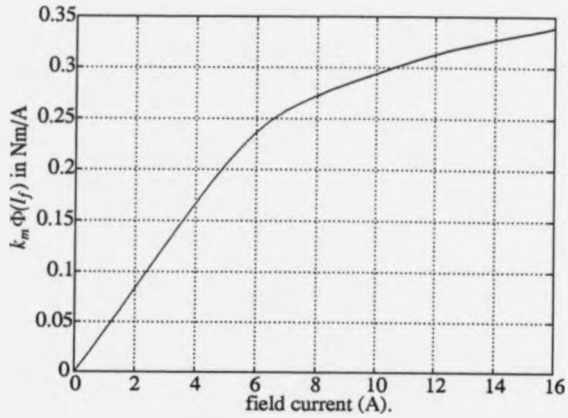

Figure 2.5: $k_m\Phi(I_f)$ against field current for Metrolec motor.

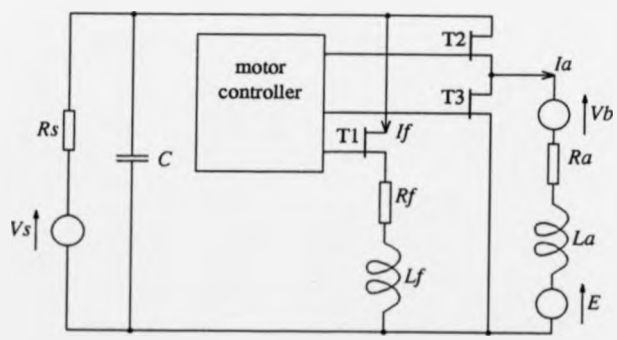

Figure 2.6: Diagram of the circuit formed by the motor, the motor controller and the batteries.

Component	Value
V_s	148.9V
R_s	280 mΩ
C	12.0 mF
R_f	5.93 Ω
L_f	399 mH
R_a	59 mΩ
L_a	0.76 mH

Table 2.1:

back EMF is represented by E. There is a resistance associated with the brushes and the commutator which is non-linear, and tends to fall as the current through the commutator rises, giving a voltage drop that is roughly constant at around 2 to 2.5 V, and is represented by V_b. I_a and I_f are the currents in the armature and field windings of the motor. The functions of the power field effect transistors T1, T2 and T3 will be explained below.

The brushed, separately excited, DC machine can be used either as a motor or a generator. Motoring action will be considered first.

2.5.1 Operation of the DC machine when acting as a motor.

On Metrolec, the input signal for the motor controller comes from a potentiometer on the accelerator pedal and is interpreted as armature current demand. At low motor speeds (less than around 470 rad/s, about 4500 r/min) the field current is set to be the nominal field current of 13.5 A. This field current is controlled by the motor controller by applying a pulse width modulated (PWM) signal to the gate of T1. T3 is turned off, and the armature current is controlled by PWM applied to the gate of T2 to meet the armature current demand. This is known as armature control. In practice, T2 and T3 are banks of transistors since they have to switch currents in the order of a few hundred amperes.

As the motor speed rises, with the field current set at its nominal value, the back EMF increases in proportion to the speed. At a certain speed, the back EMF and the voltage drops in the rest of the circuit reach the battery voltage. When this occurs, a technique called field weakening is used to allow the motor to rotate faster. In field weakening, the field current is reduced in order to reduce the back EMF, so that the demanded armature current can continue to flow. The motor is now being controlled by the field since T2 is on continuously. This technique is used until the motor speed reaches its maximum value (6700 r/min) which is set to limit sparking in the commutator.

The operation of the circuit for low motor speeds will be considered first. The problem is most easily approached by first considering the armature circuit during field weakening, when the following equation holds:

$$V_s = (I_a + I_f)R_s + V_b + I_a R_a + E \tag{2.5}$$

This can be rearranged equating voltages on either side of the fully on transistor T2:

$$V_s - (I_a + I_f)R_s = V_b + I_a R_a + E \tag{2.6}$$

During armature control, T2 is used to control the armature current by PWM. The armature PWM ratio, α_a is given by:

$$\alpha_a = \frac{V_b + I_a R_a + E}{V_s - (I_a + I_f)R_s} \tag{2.7}$$

Calculating the motor torque is straightforward using Equation 2.3, in which the field current is set to the nominal value, the value of $k_m \Phi(I_f)$ being obtained from I_f by table look up and linear interpolation. The current drawn by the motor is simply the sum of the armature and field currents.

In order to model field weakening operation, the first task is to determine whether field weakening is required to meet the armature current demand at the

current motor speed. The most straightforward way of doing this is to evaluate Equation 2.7 with the field current at the nominal value. If the value of α_a is greater than one, then the motor should be in field weakening. In field weakening operation, T2 is on continuously, therefore, Equation 2.5 holds and has to be solved for I_f. Combining Equations 2.5 and 2.4 yields:

$$I_f R_s + k_m \Phi(I_f)\omega_m = V_s - I_a(R_a + R_s) - V_b \tag{2.8}$$

which has to be solved for I_f. Since $k_m\Phi(I_f)$ is not analytic, a numerical method is used to find the value of field current, involving calculating the value of the right hand side and creating a vector of values of the left hand side corresponding to the values of I_f at which the function $k_m\Phi(I_f)$ has been evaluated. Linear interpolation is then used within the vector to obtain a solution to 2.8. Again, knowing I_f allows the motor gross torque and the battery current to be calculated.

2.5.2 Operation of the DC machine when acting as a generator.

Pressing the brake pedal on the Metrolec vehicle is interpreted as a demand for negative armature current. Since Equation 2.3 holds for both positive and negative values of I_a, negative motor torques are generated when negative armature currents flow. Negative armature currents will obviously flow when the armature voltage exceeds the battery voltage, however, as seen in the previous section, this only occurs at high motor speeds. At lower motor speeds a circuit called a step up chopper is used to cause negative currents to flow into the battery. This circuit is created by T3, which is controlled by pulse width modulation. For regeneration, the T2 transistors are turned off, the current path back to the battery being via free wheel diodes incorporated in the devices.

When E is high enough for the demanded current to flow back into the battery without the use of the step up chopper the operation of the circuit is analogous to

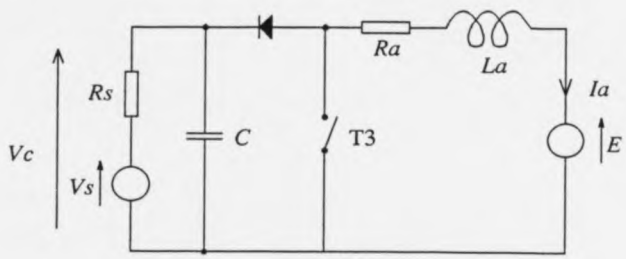

Figure 2.7: Simplified circuit diagram used in the development of the motor regenerative model.

field weakening, the correct field excitation being applied to cause the demanded current to flow. A similar numerical method is used to solve the circuit in this case as in the motoring field weakening case.

When the step up chopper is used, a detailed consideration of the actions of the T3 transistors is needed. Figure 2.7 is a simplified circuit diagram including T3 which is shown as a switch.

The objective of the work described below is to derive a steady state description of the motor operation when the step up chopper is used. In order to be useful in simulations lasting hundreds of seconds in real time, a detailed model of the chopper, whose PWM frequency is 15.6 kHz, would not be appropriate. However, the action of the chopper is considered in detail in order to derive a simple steady state description.

The state Equations for the armature circuit are as follows. With T3 off and the diode conducting:

$$\dot{I}_a = \frac{V_c - E - I_a R_a}{L_a} \tag{2.9}$$

$$\dot{V}_c = \left(\frac{V_s - V_c}{R_s} - I_a\right)/C \tag{2.10}$$

With T3 off and the diode no longer conducting:

$$\dot{I}_a = 0 \tag{2.11}$$

$$\dot{V}_c = \frac{V_s - V_c}{R_s C} \tag{2.12}$$

With T3 on Equation 2.12 still holds and:

$$\dot{I}_a = \frac{-E - I_a R_a}{L_a} \tag{2.13}$$

In order to simplify the circuit the approximation $V_c = V_S$ is made. This can be justified as follows. From Equation 2.10, if $V_c = V_s$:

$$\dot{V}_c = \frac{-I_a}{C} \tag{2.14}$$

With the component value for C given in Table 2.1 the most negative value of I_a (-200A) would result in a change in V_c of just over 1V if it were to flow for one PWM period. In practice, the current in C would not flow in the same direction for a whole period and, since 1V is small when compared with the battery voltage, the assumption $V_c = V_S$ is valid.

In the steady state the current profiles that the armature current follows are shown in Figures 2.8 and 2.9. The total steady state current that flows will be the integral of the current flowing over one PWM cycle divided by the PWM period. T3 is on for a period $\delta\tau$ and off for a period $\tau - \delta\tau$. In the case shown in Figure 2.8, the maximum current flows when E is $-V_s$ (this is found by differentiating the expression for the current with respect to E). The current which flows is then 1.58A, however, because of the action of the chopper only half of this current flows into the battery. This means that the case where the current falls to zero can be neglected.

In the case where the current does not fall to zero, the contribution to the total current flowing from the area beneath the dashed horizontal line in Figure 2.9 is even less (the same procedure as above indicates that, as I_a increases, the contribution from this part of the area becomes smaller). The current flowing in the armature when the step up chopper is used is dominated by the current

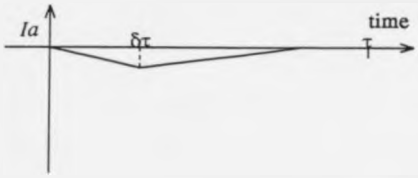

Figure 2.8: Current profile when I_a falls to zero.

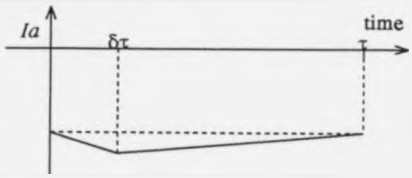

Figure 2.9: Current profile when I_a does not fall to zero.

flowing above the dashed horizontal line in Figure 2.9. Effectively, the dashed line is at the value I_a. The current flowing for time less than $\delta\tau$ is flowing in the loop formed by the step up chopper, the rest flowing back into the battery. The ratio of the gradients is used to determine the amount of the current that flows back into the battery and, combining Equations 2.9 and 2.13, the battery current, I_{bat}, is found:

$$I_{bat} = \frac{E + I_a R_a}{V_s} I_a \tag{2.15}$$

The motor gross torque is found using Equation 2.3.

2.5.3 Modelling the motor mechanical losses.

The motor gross torque, or the airgap torque, is greater than the actual output torque, or the nett torque, T_n, because of mechanical losses. These losses tend to be divided into friction losses, such as bearing drag, commutator brush drag and aerodynamic drag and "iron losses", which arise from magnetic hysteresis and eddy currents. The report [62] contains an investigation into these losses for the Metrolec motor. The upper plot in Figure 2.10 is taken from [62] and shows the friction power losses in the motor which are determined by measuring the power required to spin the motor at various speeds with no armature or field currents flowing. The lower plot converts the data to radians per second and the lost power to lost torque, since torques are used in the rest of the modelling.

The friction torque can be represented by a cubic polynomial in the motor

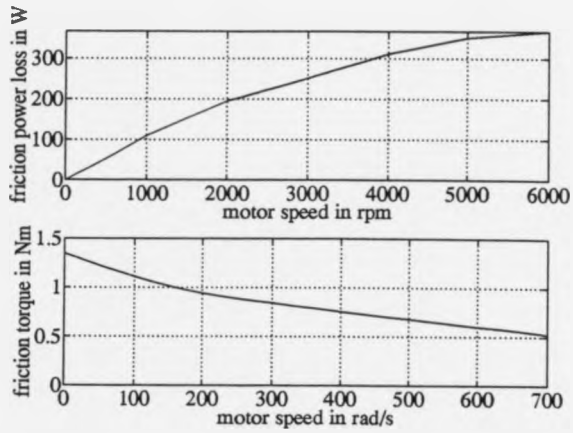

Figure 2.10: Power and torque lost to friction in the Metrolec motor.

speed. It is interesting to note that, as the speed increases, the friction torque decreases. A possible explanation for this unexpected result is that at high motor speeds aerodynamic forces cause the commutator brushes to press less firmly on the commutator, see [63, 64].

The iron losses in the motor are shown in Figure 2.11 and arise from the magnetic field caused by the field current. They are therefore measured by spinning the motor whilst various different values of field current flow .

In [61] it is stated that the power loss due to hysteresis, P_h, and the power loss due to eddy currents, P_e, may be expressed:

$$P_h = k_h \omega k_m \Phi^x(I_f) \tag{2.16}$$

$$P_e = k_e \omega^2 k_m \Phi^2(I_f) \tag{2.17}$$

where $0.8 \leq x \leq 2.3$. The constants of proportionality, k_h and k_e, can be found by finding least squares best fits of the following function for all three field current values:

$$P_{fe} = \alpha_2 \omega^2 + \alpha_1 \omega \tag{2.18}$$

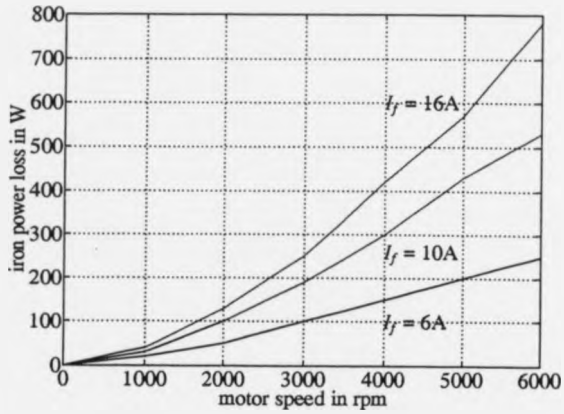

Figure 2.11: Iron losses in the Metrolec motor.

in which P_{fe} is the total power of the iron losses and α_1 and α_2 are again constants. Note that the curves in Figure 2.11 have a shape that will fit Equation 2.18 well. Combining Equations 2.16, 2.17 and 2.18, gives:

$$\alpha_1 = k_h k_m \Phi^x(I_f) \tag{2.19}$$

$$\alpha_2 = k_e k_m \Phi^2(I_f) \tag{2.20}$$

By using least squared error fits again, on the values of I_f, and the corresponding values of α_1 and α_2, estimates for k_h, k_e and x are found. The intermediate results in this process are not particularly good. This is, possibly, due to the small number of field current data points that are available. Also, in [62] it is stated that the instrument that was used to measure the loss powers was not well matched to the powers it was used to measure, leading to the introduction of experimental errors which are unspecified. However, when estimates for the values of k_h, k_e and x are found they can be used to predict the iron losses. Figure 2.12 compares these predictions with the measured values.

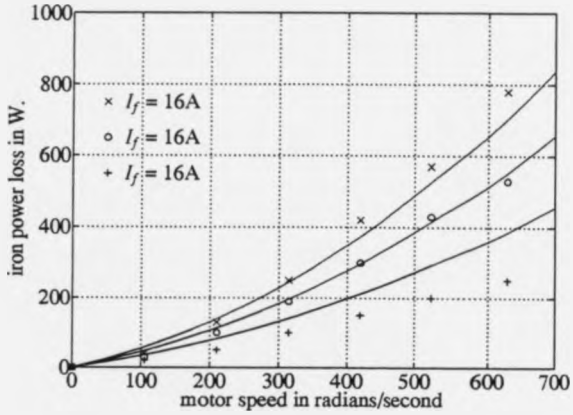

Figure 2.12: Comparison of measured iron loss values (shown with crosses) and predictions from iron loss modelling (shown with solid lines).

The predictions are very good for I_f = 10A, and acceptable for I_f = 16A. For I_f = 6A the results are rather poor, however, field currents of 6A are only generally used at motor speeds above 650 rad/s (about 6300 r/min) and this condition occurs very infrequently.

2.6 Hybrid vehicle transmission models.

In general, the subject of transmission modelling has received very little attention, especially when compared to engine modelling. Some modelling work has been done for the purposes of controlling automatic shifting, [65] and [66] being examples which address the issue of how to change gear using a novel gear changing mechanisms. A more pressing problem has, generally, been the scheduling of gear changes which has tended to be achieved either using hydraulic control or a mixture of both electronic and hydraulic control. This problem has seen little use of modelling work, perhaps due to the very subjective nature of gear change

timing. Reference [67] considers the problem of controlling a continuously variable Perbury transmission to optimise vehicle driveability and emissions in which both the "how to shift" and "when to shift" issues are considered.

In this work, it is important that the torque losses in the transmission are accurately modelled. Again, this area of transmission modelling has received very little attention in the literature. In the case of manual transmissions, the power lost in the transmission is very low compared to the power of the engine, so the effect is not felt to be particularly important. Also, manual transmissions are relatively straightforward and knowing the losses in the transmission has not been felt to be useful when there is no obvious method of reducing them by design improvements to the unit. Note, that exactly the reverse is true for engines.

Automatic transmissions are known to be relatively inefficient due to the use of a torque converter, instead of a clutch, to couple the engine and transmission. Improvements in this system have been made by locking up the torque converter once the vehicle has reached a certain roadspeed. For the most part, though, poor fuel economy and acceleration has been accepted as a necessary consequence of the convenience of using an automatic transmission.

The introduction of zero emissions legislation and carbon dioxide incentives has made the issue of transmission efficiency much more important. In an electric vehicle, the motor power is frequently less than a third of the power of the petrol engine that would normally be installed in the same vehicle, whilst the batteries tend to store less than a seventh of the energy that is available from burning the fuel stored in a conventional car. This has meant that some measure of transmission efficiency is important in the modelling of electric vehicles.

In some cases, the transmission efficiency used in electric vehicle modelling work is ignored and, presumably, set at a fixed value of 100%. More commonly, the transmission efficiency is set at a fixed value over all operating conditions, ex-

amples of work taking this approach include [48, 56, 51]. A more refined approach is given in [53], in which tables of efficiency value are used and in [47], where an efficiency function is used.

The results of two studies that limit their attention to the efficiency of transmissions are to be found in [68] and [69]. Reference [68] is of particular interest, since it includes the results of experiments carried out on a transmission very similar to the R65 type transmission used in the electric vehicle and the conventional vehicle, whose components make up the hypothetical hybrid vehicle on which this study is based.

The modelling approach taken in this work could have used raw efficiency data, interpolated by some means as a model of the transmission. However, it was felt that analysing the results presented in [68] might give some physical insight into the efficiency data and would also allow the losses in the transmission to be expressed as torques rather than efficiencies. Using torques in this manner would be more closely aligned with the philosophy of forward dynamic modelling than the use of raw efficiency data.

The author is very grateful to Leo van Dongen for allowing him to use the material first published in [68]. The purpose of the following analysis is to take the results presented in [68] and express them in terms of transmission input torque, transmission input speed and gear ratio selected. The resulting model can then be used for any transmission of the type considered in [68] if the ratios of the specific transmission being considered are known.

The ratios that are used in [68], and in the work described here, are defined as:

$$gear\ ratio = \frac{engine\ speed}{wheel\ speed}$$

This means that high values of this ratio indicate low gears (in the sense of first gear, second gear, etc).

In the conclusions to [68], it is stated that the losses in a manual transmission consist of losses due to oil churning and bearing and seal drag, and losses that are due to friction between gear teeth. A frictional torque will be present between the gear teeth in the transmission even at no load since the input shaft has to transmit torque to the other shafts (the layshaft and the output shaft) in order to overcome the frictional torques that are present at these shafts. However, this torque will be very small, compared to any real input shaft torques, and is neglected in the next part of the analysis.

Modelling the transmission no load torque loss.

The curves in Figure 2.13, reproduced from [68], show the no load power losses in a manual passenger car transmission, and can be approximated by a quadratic function of the form:

$$P_L = k_2\,\omega^2 + k_1\,\omega \tag{2.21}$$

where P_L is the power loss, ω is the transmission input speed, in revolutions per minute and k_1 and k_2 are the coefficients of the quadratic function. Note that there is no k_0 term, by definition, since this would imply a power loss at zero speed. Equation 2.21 can be replaced by:

$$T_l = c_2\,\omega + c_1 \tag{2.22}$$

where T_l is now the torque loss. Also, c_1 and c_2 are not the same as k_1 and k_2 because of the non-SI units in which gearbox input speed is expressed, for which the author apologises, (rad/s being almost unheard of in the automotive industry).

Estimates of the input power loss at various transmission input speeds and gear ratios were obtained by measurements taken from Figure 2.13, and were then converted into torques which have been plotted with crosses in Figure 2.14.

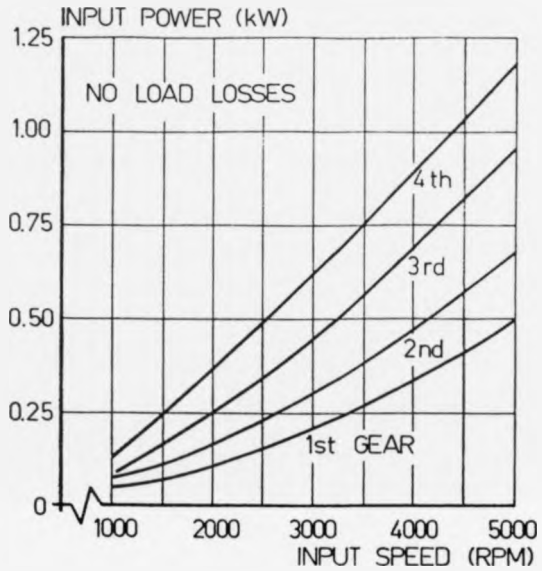

Figure 2.13: No load power losses in manual transmission.

The lines on Figure 2.14 are least squared best fits of functions of the form of Equation 2.22, and the coefficients of these lines are, therefore, the values of c_1 and c_2 for each gear ratio. In general, the crosses on Figure 2.14 lie close to the best fit straight lines. Better fits would obviously have been achieved using higher order polynomials, however, since the points for the graphs were obtained by taking measurements by hand, the use of more complex functions could not be justified.

In order to relate transmission no load torque loss to the gear ratio, it would be sensible to relate the coefficients obtained for Equation 2.22 to the gear ratios in the transmission. The coefficient values, and ratio values, are given in Table 2.2, which suggests that the reciprocal of gear ratio might be approximated by a linear relationship to the coefficient values.

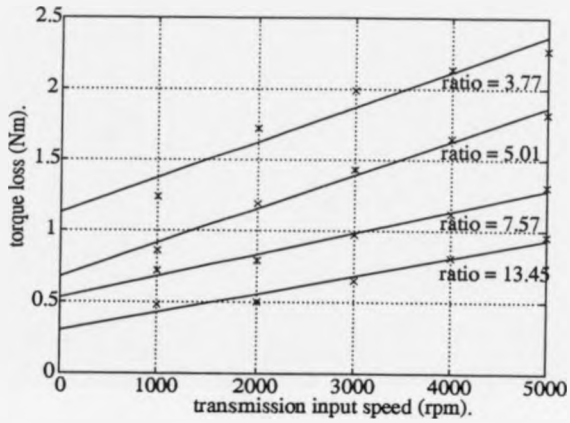

Figure 2.14: No load torque losses for manual transmission using data taken from Figure 2.11.

gear ratio	c_2 value (Nm/rpm)	c_1 value (Nm)
13.45	0.127×10^{-3}	0.299
7.57	0.151×10^{-3}	0.529
5.01	0.238×10^{-3}	0.676
3.77	0.248×10^{-3}	1.128

Table 2.2:

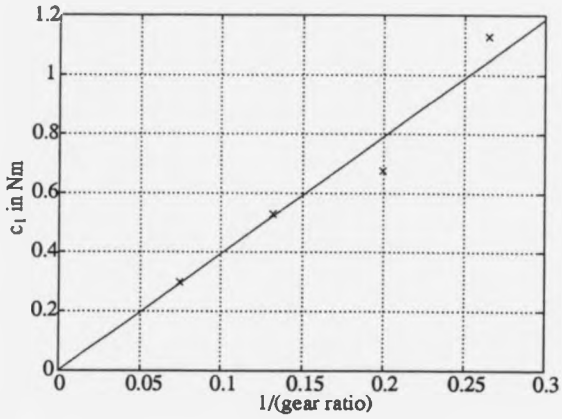

Figure 2.15: c_1 against reciprocal of gear ratio.

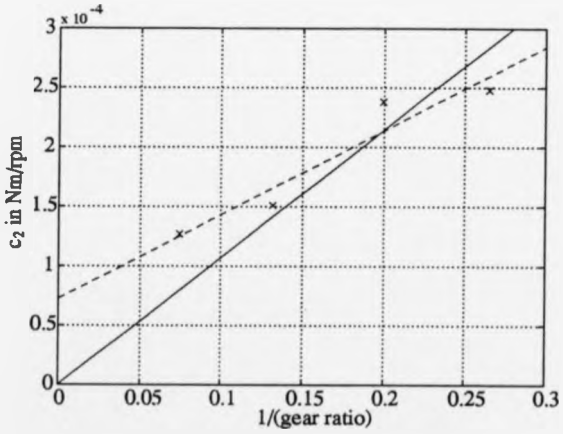

Figure 2.16: c_2 against reciprocal of gear ratio.

Figures 2.15 and 2.16 show the values of c_1 and c_2 plotted against the reciprocal of gear ratio. They also include best least squared straight line fits constrained to pass through the origin for each set of data, shown by the solid lines. This allows the coefficient values to be related to the reciprocal of the gear ratio by the gradient of the graph. In the case of Figure 2.15, the fit is really quite good and it is obvious that an unconstrained best line fit would pass close to the origin. The solid line fit obtained in Figure 2.16 is not as good, and an unconstrained fit, shown by the dashed line, actually passes some distance from the origin. However, the physical consequence of the gear ratio approaching infinity (and its reciprocal approaching zero), is that as the transmission input shaft turns the output shaft remains still. It therefore seems reasonable that the term relating lost torque to input shaft speed should be zero at this value of gear ratio, since, as the input shaft speed increases, the output shaft speed remains zero. From the Figure 2.15 the following relationship is found:

$$c_1 = \frac{3.95}{gear\ ratio} \tag{2.23}$$

and using the solid line in Figure 2.16:

$$c_2 = \frac{1.071 \times 10^{-3}\,\omega}{gear\ ratio} \tag{2.24}$$

A final expression for the lost torque is then obtained by combining Equations 2.22, 2.23 and 2.24 to get:

$$T_l = \frac{(1.071 \times 10^{-3}\,\omega + 3.95)}{gear\ ratio} \tag{2.25}$$

The values that this function predicts are compared with the measured data in Figure 2.17, from which it is seen that the predictions of Equation 2.25 agree very well with the measured data for the lower three gear ratios. For the ratio of 13.45, the reason for the poor agreement is found in Figure 2.16 where the point for c_2 for the this gear ratio does not lie close to the line. Note, that in Figure 2.15 the

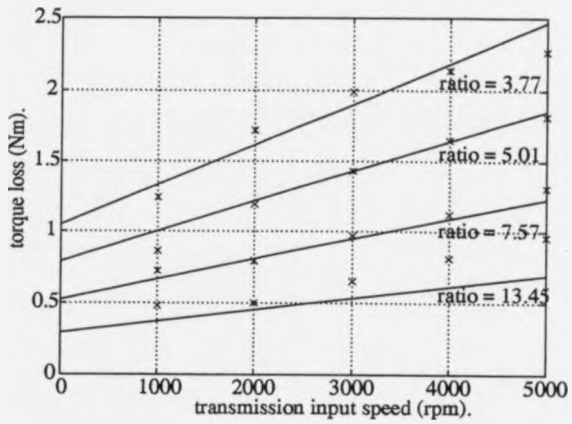

Figure 2.17: Comparison of the measured torque loss values and the values predicted by the torque loss model.

value for c_1 is close to the line. The no load torque loss model is derived from experimental data which has been subsequently measured by hand, and there are good physical reasons why the line in Figure 2.16 should pass through the origin. The error that was introduced by constraining the line in this way was, therefore, felt to be acceptable.

The no load torque loss modelling work has obtained expressions for c_1 and c_2. Some physical insight is gained by noting that c_1 relates the gear ratio to the torque losses in bearings and seals, which will be largely invariant with input speed. Also, c_2 relates the gear ratio to the torque losses due to oil churning, whose power will be proportional to a power of the input speed. Further work in this area might be able to relate these values to the internal design of the transmission.

Modelling the gear friction torque loss.

Figure 2.18 shows the measured efficiency characteristics for the manual transmission, again reproduced from [68]. The frictional loss in the gear teeth was investigated using these plots and the expression derived in the previous section for no load torque loss. The output torque T_o can be written:

$$T_o = T_i \, \eta_o \, g_r \tag{2.26}$$

where η_o is the overall transmission efficiency, T_i is the input torque and g_r is the gear ratio. An imaginary torque, T_{im}, that is equal to the input torque, less the no load lost torque, could be considered to act on the input shaft of the transmission. Any losses between the point at which this imaginary torque acts, and the transmission output shaft, should then be due to friction in the gear teeth. This allows an estimate of the teeth efficiency, η_t, to be calculated:

$$\eta_t = \frac{T_o}{T_{im} \, g_r} \tag{2.27}$$

or:

$$\eta_t = \frac{T_o}{(T_i - T_l) \, g_r} \tag{2.28}$$

Using the transmission efficiency curves in the graphs presented in Figure 2.18, it is possible to calculate a transmission output torque, for any input speed and torque, and any gear ratio. The torque loss model can then be used with Equation 2.28 to calculate an estimate for η_t. This calculation was carried out for each gear ratio, and each speed, at input torque values taken every 12.5 Nm, from 12.5 Nm to 87.5 Nm. It was found that the values obtained at 12.5 Nm could not be trusted (in one case the value of η_t was greater than unity). This was because values of η_o could not be measured accurately from the graphs due to the shape of the curve in this region. Having discarded the points obtained at 12.5 Nm, 84 points remained at which the value of η_t had been estimated. The average value

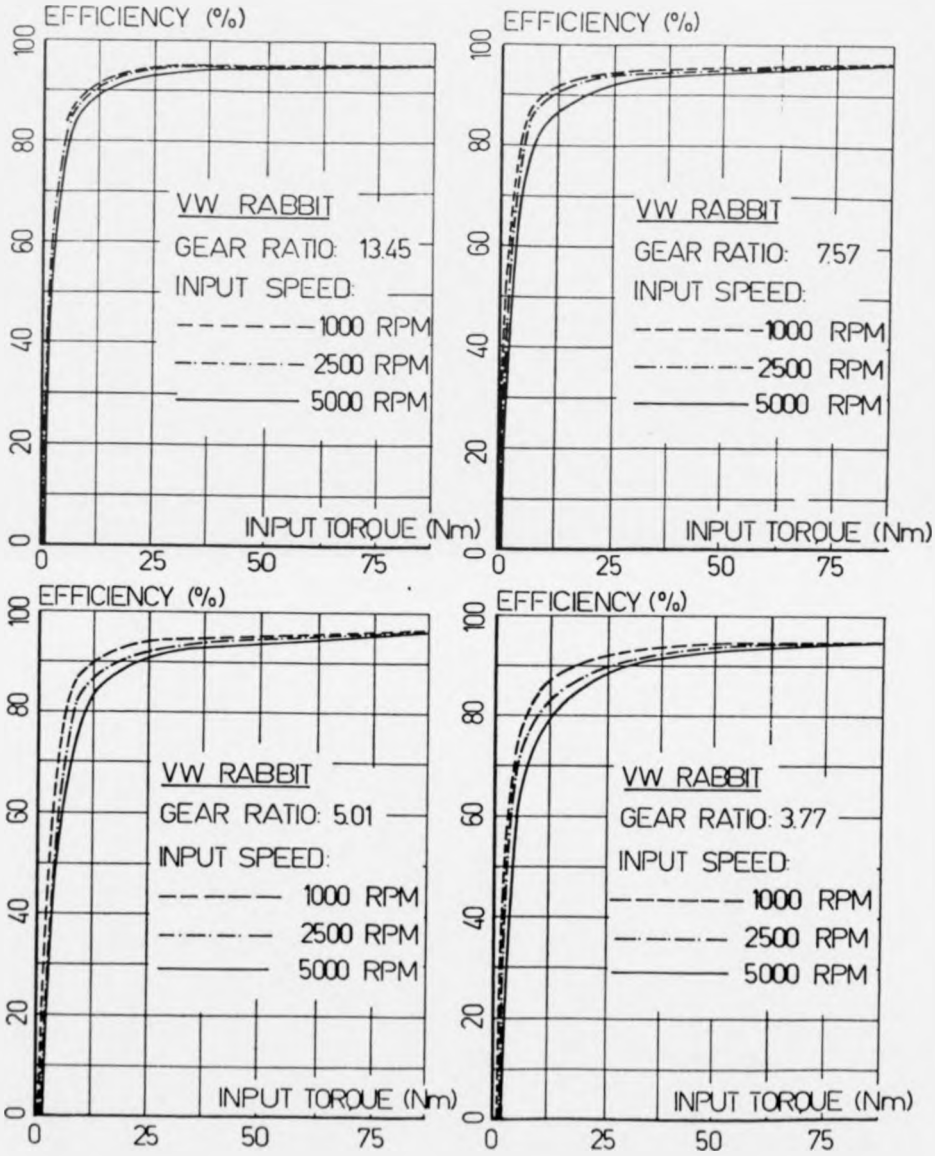

Figure 2.18: Manual transmission efficiency.

for each of the twelve traces was then found, and these averages lay in the range 0.957 to 0.977, with all but three values lying in the range 0.960 to 0.972. Given the nature of the technique used to obtain these values, a result quoted to more than two significant figures cannot be justified. A value of 0.96 was chosen as the tooth efficiency because of the clustering of the results around this value.

An expression for the output torque of the transmission can then be written:

$$T_o = (T_i - T_l)\,\eta_t\,g_r \tag{2.29}$$

2.7 Vehicle dynamics.

The purpose of modelling the vehicle dynamics is to be able to predict the acceleration of the vehicle. Clearly, the forward acceleration of the vehicle will be equal to the overall forward force on the vehicle divided by the mass of the vehicle. The forward force on the vehicle is equal to the tyre forward force, less the drag forces due to tyre rolling resistance and aerodynamic drag. The tyre rolling resistance, F_{rr}, is given by:

$$F_{rr} = mgc_r \tag{2.30}$$

where c_r, the coefficient of tyre rolling resistance, relates the drag force caused by the tyre to the vertical load on the tyre. The coefficient of rolling resistance is a function of many variables such as tyre inflation pressure, tyre temperature, vehicle speed and the nature of the road surface, which are largely impossible to model. Generally, acceptable results are obtained without the need to consider these effects. Typical values for tyre drag are between 0.015 and 0.020, though the Metrolec vehicle had special low rolling resistance tyres that have a c_r value of about 0.010. The aerodynamic drag is given by the well known equation:

$$F_{ae} = \frac{1}{2}\rho c_d A v^2 \tag{2.31}$$

where F_{ae} is the force due to aerodynamic drag, ρ is the density of air, c_d is the drag coefficient of the vehicle, A is its frontal area and v is its velocity. The acceleration of the vehicle is then:

$$\dot{v} = \frac{F_f - F_{rr} - F_{ae}}{m_{ef}} \tag{2.32}$$

where m_{ef} is the effective mass of the vehicle and F_f is the forward force in the tyre. The effective mass of the vehicle is greater than the actual mass of the vehicle, because of the rotational inertias of the engine and the drivetrain.

Considering the conventional vehicle, the acceleration of the engine is obtained from the acceleration of the vehicle, $\dot{v}$, since the engine and the wheels of the vehicle are assumed to be rigidly coupled. The following Equation of motion for the engine acceleration can then be written:

$$T_e = \frac{J_e \dot{v} g_r}{r_r} + T_i \tag{2.33}$$

where J_e is the engine inertia, $\dot{v}$ is the vehicle acceleration, g_r is the gear ratio as defined in Section 2.6, r_r is the tyre rolling radius (which is slightly larger than the loaded radius) and T_i is the transmission input torque. The transmission output torque, T_o, can be related to the input torque by:

$$T_i = \frac{T_o}{g_r \eta_{tr}} \tag{2.34}$$

where η_{tr} is an approximate transmission efficiency, used only for the purposes of calculating m_{ef}, and has a value of 0.96. The transmission output torque passes via driveshafts to the wheel and, equating torques at the wheel gives:

$$T_o = \frac{J_d \dot{v}}{r_r} + F_f r_l \tag{2.35}$$

where J_d is the rotational inertia of the wheels, tyres and driveshafts and r_l is the rolling radius of the tyre. The overall acceleration of the vehicle is then found:

$$F_f - F_{rr} - F_{ae} = m\dot{v} \tag{2.36}$$

where m is the mass of the vehicle. Combining Equations 2.33, 2.34, 2.35 and 2.36 and rearranging gives:

$$\frac{T_e g_r \eta_g}{r_l} - F_{rr} - F_{ae} = \left(\frac{J_e g_r^2 \eta_g}{r_r r_l} + \frac{J_d}{r_r r_l} + m \right) \dot{v} \tag{2.37}$$

Equation 2.37 is now in the familiar, $F = ma$, form in which the force on the left hand side is the transmission output torque divided by the loaded radius of the tyre less the drag forces, and the right hand side is a mass multiplied by the vehicle acceleration. The effective mass of the vehicle is then given:

$$m_{ef} = \frac{J_e g_r^2 \eta_g}{r_r r_l} + \frac{J_d}{r_r r_l} + m \tag{2.38}$$

and can be used in Equation 2.32. The additional effective mass of the rotational inertias is surprisingly important, amounting to about 250-300 kg in first gear which, on a vehicle whose mass is around 1000 kg, is very significant. Equation 2.38 has been developed around the conventional vehicle, for the hybrid vehicle the inertia of the motor would have to be added to that of the engine. If there are any gear ratios between the motor and the engine then the motor inertia would have to be multiplied by the square of the gear ratio before being added to the engine inertia.

2.8 Completed vehicle model.

Figure 2.2 is a schematic diagram of the completed hybrid vehicle model. The engine, motor, transmission and vehicle dynamics blocks all implement the models described above. The driver and controller blocks are described below. The components are all assumed to be rigidly coupled to the wheels of the vehicle from which their speeds are derived. It is assumed that the wheels do not slip and their speed is derived from the vehicle speed.

In the hybrid vehicle that was initially considered, the motor and the engine had very similar maximum speeds (though the engine has an effective minimum

speed of around 1000 r/min while the motor has no minimum speed). The engine and the motor both turn at the same speed, their output torques being added together at the input to the transmission.

2.8.1 Controller block.

The controller block is used to interpret the signals from the driver block and generate demand signals for the engine and the motor. The signal for the motor is demanded armature current and the signal for the engine is demanded throttle angle. It is assumed in this work that the engine has an electronic throttle controller whose dynamics are sufficiently fast to be neglected. There is, therefore, no mechanical link between the accelerator pedal and the engine or the motor. The purpose of this thesis is to develop control strategies that can be used in this controller.

2.8.2 Driver block.

The purpose of the driver block is not to model the actions of a human driver but, simply, to apply the required pedal values such that the vehicle follows a desired speed-time profile. Driver modelling has been done using fuzzy logic, [70], and using traditional control theoretic techniques, [71]. Generally, the emphasis in this sort of work is in steering control and the effect of the driver on the overall directional stability of the vehicle, though [70] includes some consideration of speed control.

The driver actions used in this work were implemented using a simple proportional plus integral controller which takes, as its input, the error between the vehicle speed in m/s and a demanded reference speed in m/s. Originally, the driver block was developed around the conventional vehicle and was first used in the validation work described in Section 2.9.1. Suitable proportional, and inte-

gral, gains were found to be 15.0, and 8.0. These gains were decided upon by comparing the simulated vehicle speed profiles over drive cycles with real vehicle speed responses over drive cycles using human drivers. It was found that the same gains could be used in all the different gears, though a different gain was needed to implement braking. Braking control was implemented by taking negative values from the driver output as meaning zero throttle angle, and multiplying the output by a further gain of 120.0 to get the braking torque at the wheel in Nm.

The reference data describing the speed profile that the vehicle has to follow also defines which gear should be selected in a manual transmission. This gear is then passed to the transmission model.

The integration step size of the simulations carried out using this vehicle model was 0.1 second and was determined empirically by conducting simulations during the development of the vehicle model.

2.8.3 Representing the action of the clutch.

In a conventional vehicle with a manual transmission, the engine is coupled to the transmission input by means of a friction clutch, allowing its speed to differ from that of the transmission input at low vehicle speeds and, also, allowing the engine to be decoupled when gear selection takes place. The actions of a driver when controlling the clutch of a vehicle are extremely difficult to model, the approach taken in this work is as follows. When the vehicle is stationary, the throttle angle is zero, the engine is at idle and the engine speed and fuel flow rate are as described in Section 2.4. To simulate the vehicle moving away from rest the following condition is detected. The vehicle speed is lower than the reference speed, first gear is selected and the transmission input speed is lower than 1200 r/min. Under these circumstances, the engine speed is set to 1200 r/min and the throttle angle is taken from the PI controller as usual. When the vehicle

has accelerated to such a speed that the transmission input speed reaches 1200 r/min, the engine speed is then set equal to the transmission input speed. When the vehicle is slowing down, and the transmission input speed is lower than 1200 r/min, the engine is declutched and at idle.

The modelling of manual gear changes in performance predictions is extremely difficult, since the effectiveness of a gear change is very dependant upon driver technique. The approach taken in this work is to instantaneously change the gear ratio and this, instantaneously, changes the engine speed. In reality, sudden changes in engine speed of this type produce uncomfortable jerks, see [72], but, as the function of the simulation is to predict fuel consumption rather than accurately simulate transients, this approach is acceptable.

2.8.4 Drive cycles.

The reference speed profiles, that the vehicle is made to follow by the driver block, define the speed of the vehicle and the gear selected over the duration of the cycle. The cycles chosen can be arbitrary, however, certain defined emissions and fuel consumption test cycles are used regularly. The most important of these, for this work, is the European ECE-15 test cycle, part of which is shown in Figure 2.19. Each of the speed excursions shown in Figure 2.19 is, somewhat incorrectly, known as a "hill", although obviously no gradient is involved. The full ECE-15 cycle consists of the three hills performed four times.

First gear is, obviously, used to start to ascend each hill, the crosses indicate where the cycle specifies that the driver should change from first to second gear, and the circle indicates where third gear should be used. The gears are not specified for descending the hills, although, third gear should be used until the top of the final descent on the third hill.

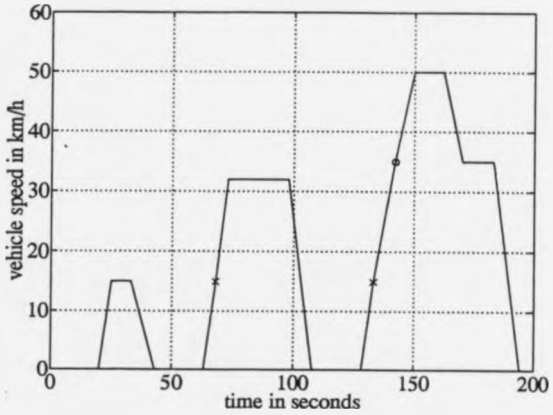

Figure 2.19: Speed profile for the first three "hills" of the ECE-15 driving cycle.

2.9 Model validation.

As stated in Section 2.3.3, an advantage of the use of real vehicles as a base for modelling work is that the model predictions can be compared with the performance of the vehicles. The validation work for each type of vehicle is described in the next two sections.

2.9.1 Petrol engine vehicle model validation.

Performance evaluation.

An obvious performance figure that is quoted for production vehicles is the 0-60 mph acceleration time. However, this performance value is strongly affected by the way in which the vehicle accelerates from rest and, then, the way in which the gear changes are performed, both of which are only modelled approximately for reasons explained in Section 2.8.3. Performance prediction using in-gear acceleration times is a much better means of investigating the model validity in this case.

The parameters of the in-gear acceleration simulations are shown in Table 2.3. The engine inertia includes the inertia of the flywheel, ancilliaries and transmission components that rotate at the speed of the engine. The wheel inertia includes the inertia of the wheels and tyres, braking components, driveshafts and CV joints, and transmission components that rotate at the same speed as the wheels.

The simulation results, and comparison with the official performance figures for the Metro GTi, are shown in Table 2.4. It is seen that the model predicts the acceleration times in fourth gear for all the tests, except the fastest one, to within an accuracy that is acceptable. The last test result could be inaccurate because of inaccuracies in the engine model, since this is the only test that uses engine speeds above 5900 r/min (the fastest fifth gear test takes the engine speed to 5300 r/min). There is no immediately obvious reason why the results for the simulations using fifth gear are consistently too fast by about the same amount. Since the overall gearing in the simulation exactly matches the correct value, an error could be present in the transmission losses. The gear ratio of fifth gear is outside the range of the gear ratios used in the development of the model and, also, the transmission model is based on a four speed transmission, rather than a five speed transmission as used on the vehicle. The combination of these factors could explain why the errors are generally about 3% more negative in fifth gear than in fourth gear.

Fuel consumption prediction.

A further validation of the conventional vehicle models is fuel consumption prediction over standard drive cycles. This is really a more meaningful validation, since it is closer to the purpose for which the models were developed, and exercises them over a wide range of conditions. The validation chosen for this purpose was the prediction of the fuel consumption of the Metro GTi over the ECE-15 fuel

Parameter.	Value.
Vehicle mass	1021 kg
Engine inertia	0.15 kgm^2
Wheel inertia	3.2 kgm^2
Tyre coefficient of rolling resistance	0.015
Aerodynamic drag coefficient	0.36
Frontal area	1.78 m^2
Tyre rolling radius	0.259 m
Tyre loaded radius	0.245 m

Table 2.3:

economy cycle. The results of this simulation are shown in Figure 2.20.

The top plot in Figure 2.20 is the vehicle speed over the simulation, the reference speed is shown dotted. The second plot shows the pedal values that the PI controller in the driver block applied to cause the vehicle to follow the cycle, the positve values are throttle angle, the negative values are the braking torque signal before the gain of 120 is applied. The lowest plot shows the engine speed and illustrates how the engine speed is held at 1200 r/min until the transmission input turns synchronously.

The fuel consumption predicted by this simulation is 32.1 miles per gallon, whereas the vehicle achieves 33.4 mpg over the test, a difference of 3.9% between simulation and test. In reality, this result is not quite as good as it might at first seem, since the engine map used in the model is fully warm, whereas the vehicle starts at 20°C in the test. The vehicle will, therefore, run with an air fuel ratio that is richer than usual until it is warm. The engine is not fully warm until about 100 seconds into this test, and recalling that the three hills are performed four times in the real test, the effect of using a fully warm map over the entire simulation is low. Using a fully warm engine map throughout the simulation should result in the fuel consumption prediction being about 2% too low. This is complicated

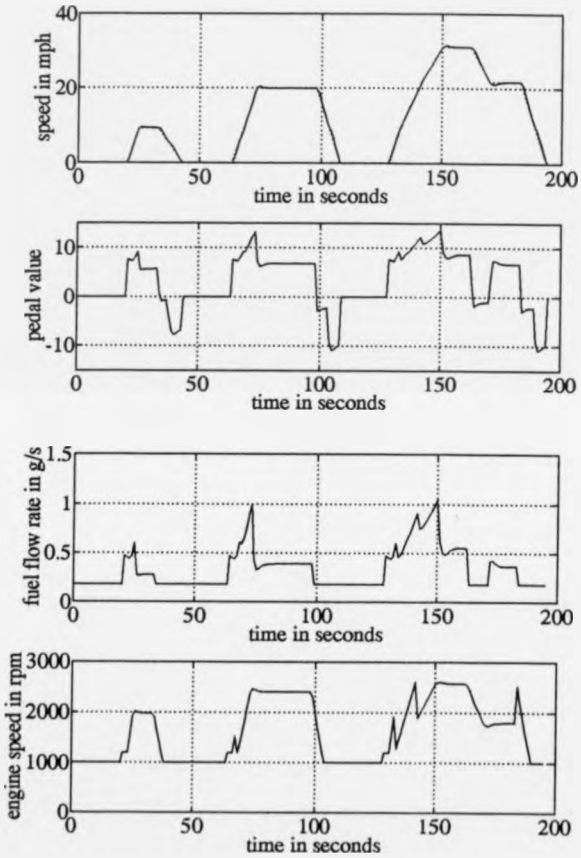

Figure 2.20: Simulation variables obtained during conventional vehicle fuel consumption validation.

Speed range (mph)	Gear	Simulation time	Test time	% difference
30 - 50	4	6.70	7.0	-4.3
40 - 60	4	6.93	7.2	-3.7
50 - 70	4	7.38	7.5	-1.6
60 - 80	4	7.94	8.2	-3.2
70 - 90	4	9.12	9.1	+0.2
80 - 100	4	12.68	11.9	+6.6
30 - 50	5	8.89	9.6	-7.4
40 - 60	5	9.14	9.8	-6.7
50 - 70	5	9.84	10.5	-6.3
60 - 80	5	11.06	12.0	-7.8
70 - 90	5	12.75	13.6	-6.3
80 - 100	5	15.44	16.5	-6.4

Table 2.4:

by the fact that, during cold running the engine tends to produce slightly more power due to the low air fuel ratio and this effect is completely unmodelled.

This validation indicates that the models can be used to predict fuel consumption to within around 5% over ECE-15, and similar, cycles.

2.9.2 Electric vehicle model validation.

An electronic instrumentation system was developed to validate the electric vehicle models by measuring the intermediate variables such as driveshaft torque, motor currents, and so on, as the vehicle was driven. Unfortunately, when the instrumentation was completed the vehicle was very unreliable for a considerable time and then became unusable due to faults in the high temperature batteries. This means that the only performance validation that could be carried out was the simulation of some in-gear acceleration tests that were done on a test track, whose results are shown in Figure 2.21.

The first gear test seems to indicate reasonable accuracy after the first two seconds of simulation. At the time at which the test work was done, the rate of

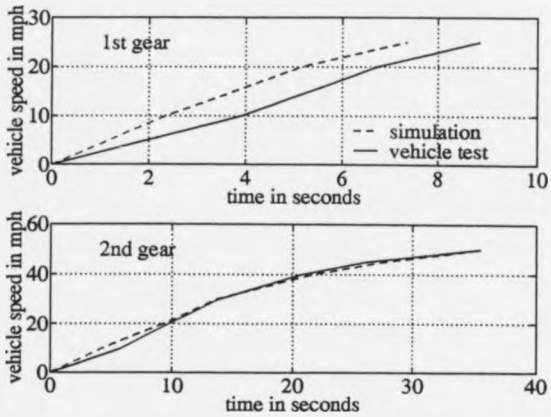

Figure 2.21: Comparison of Metrolec acceleration tests and simulations.

change of current was limited by the motor controller software. The maximum armature current of 160A was reached after around 0.5 seconds, additionally, other lags, such as the delays involved in causing currents to rise in the motor inductances, and the delays of the motor control loops that have not been modelled, will tend to make the car slower than the simulation as it moves from rest.

The results obtained in second gear agree well with the test. However, tests performed in higher gears, at higher speeds, did not agree as well with the simulations. A possible reason for this is that the tests were not performed on a flat track (normally acceleration tests are repeated in opposite directions on a reasonably flat track). Even very small gradients have a very large effect on the performance of an electric vehicle at speeds approaching its maximum speed, therefore errors in the high speed tests are to be expected.

The maximum speed of the Metrolec was calculated by plotting the driveshaft torque, at the maximum armature current, against road speed on the same plot as a plot of the driveshaft torque required to maintain constant roadspeed. The point

at which the two lines meet is the maximum speed prediction. This technique produces a maximum speed estimate of 64.02 mph and, in use, the maximum speed of the vehicle was found to be between 65 and 70 mph. The mass used in these calculations was 1127 kg, and the tyre coefficient of rolling resistance was 0.011.

Current consumption tests, leading to battery state of charge predictions, could not be carried out due to problems with the vehicle. Such predictions would be extremely difficult because the vehicle has electrical loads which are significant, but essentially unpredictable, such as battery heaters and coolers and other normal vehicle electrical loads, such as windscreen wipers and demisters.

2.10 Conclusions.

This chapter has stated the reasons for the development of the hybrid vehicle powertrain models and defined a hypothetical vehicle on which to base the modelling studies. This vehicle is based on an existing prototype electric vehicle and a production petrol engined car.

Two different approaches to powertrain modelling, forward and reverse dynamic modelling, are described and it is found that forward dynamic modelling is the most appropriate for this application. Models are then developed for a petrol engine, an electric motor, a vehicle transmission and the vehicle dynamics. The petrol engine model uses static engine maps of torque and fuel flow rate. The electric motor model is derived considering the circuit formed by the motor, its controller and the batteries of an electric vehicle. The transmission model is derived from the numerical analysis of experimental results from transmission tests. The vehicle dynamics are modelled by considering the balance of the tractive and drag forces on the vehicle and by calculating the total inertia of the vehicle, and its rotating inertias.

Validation exercises are described in which the predictions of the models are compared with the performance of the real vehicles. It is seen that a model of a production petrol engine vehicle predicts its in-gear accelerations to within 5%, in the areas of interest, and the fuel consumption over a defined emissions and fuel consumption cycle also to within 5%. A complete validation of the electric vehicle models was not possible, however, the reliable tests that were done indicated that the simulation results modelled the vehicle performance acceptably.

Having developed a set of hybrid vehicle component models the performance of a hybrid vehicle cannot yet be simulated, since, this requires a powertrain controller, the form of which is currently unspecified. The development of such controller is the task of the rest of this thesis.

Chapter 3

An initial application of fuzzy logic in the control of a hybrid vehicle powertrain.

3.1 Introduction.

The role of the powertrain controller in a hybrid vehicle has been described in detail in Section 2.2. Previous work in the area of hybrid powertrain control will be described and considered in the context of automotive and control engineering in the early 1990s. The engineering issues of the specific hybrid vehicle powertrain control task addressed in this work are then considered and this leads to a problem formulation. It will be seen that fuzzy logic is ideally suited to this problem whereas other, more conventional, control methods are less well suited.

The way in which the fuzzy logic is used is then considered in some detail. The information that the fuzzy logic relates is defined, and suitable units of measurement for the variables used to describe the fuzzy information are obtained. The detailed operation of the fuzzy logic is then considered, and the techniques that were used in early work in the field of fuzzy control are applied and found to have drawbacks. In the light of these drawbacks, the fuzzy implication method, the defuzzification method and the set shapes used are then revised, these revisions

being justified on philosophical grounds. It is found that, having been justified in this manner, the revised methods result in a controller with a lower computational burden.

These revisions allow some conclusions about the role of fuzzy logic in control engineering to be drawn and these are briefly considered. The results of simulation experiments using the fuzzy controllers derived previously are then presented and discussed. Some individual results are considered in detail and the general form of the overall results obtained is also discussed. The simulation experiments highlight some important hybrid vehicle operation points which are not specific to the way in which the controller operates and these are described. Some aspects of this work have been published elsewhere, see [73].

3.2 A review of previous work in hybrid vehicle powertrain control.

The powertrain controller in a hybrid vehicle is not only fundamental to the operation of the vehicle but also, to a large extent, determines how well the vehicle performs, in terms of the way in which the vehicle uses energy and its driveability. It is, therefore, surprising that relatively little work in the area of hybrid vehicle powertrain control has been reported. Possible reasons for this are that much of the work reported has used simulation studies and it is possible to neglect the powertrain control issues in a simulation study more than in a vehicle, where the practical aspects are immediately obvious. Also, as stated in the introduction, relatively few prototype hybrid vehicles have been produced and the form of the powertrain controllers used is very commercially sensitive information. Interestingly, [7] recognised the problems of component control in hybrid powertrains almost 90 years ago!

3.2.1 Hybrid powertrain control strategies using on/off engine operation or simple functions of pedal value.

Most of the hybrid vehicle powertrain control studies identify regions in the vehicle speed/pedal value space and define control actions for those regions, based on considering the efficiency maps of the components and the current controller objectives. The work that uses this approach will be described by first considering controllers implemented on vehicles and then controllers that have been investigated using computer simulation studies.

Reference [74] describes practical hybrid vehicle controllers for four vehicles. In each of these examples the actions taken by the controller are relatively straightforward. In the first example, applied to a city taxi, the controller uses only the electric motor of the vehicle when in towns, but starts the engine outside towns. Where the engine power exceeds the instantaneous demands of the vehicle, the motor absorbs the power to recharge the batteries. The powertrain is designed with a single transmission gear ratio. In another application described in [74], the engine is normally used to power the vehicle. However, when the accelerator pedal is released, if it is not depressed again within 0.5 seconds, the engine is switched off. When the accelerator pedal is pressed again, the electric motor is used. If the demands of the driver are not met by the motor, the engine is restarted and used to supplement the motor. There are two possible drawbacks to such an approach. The first is that if the engine is not used reasonably frequently then the emissions controlling catalyst will cool down and stop converting the pollutants. This would have a great effect on the emissions of the vehicle. From the point of view of catalyst operation, the engine should either be run continuously or not at all. The second problem is restarting the engine. If the engine does not fire immediately, unburnt fuel is emmited which would have a disastrous effect on vehicle emissions. Also, there is the problem of the driveability effects of restarting the

engine and controlling its speed to be synchronous with the vehicle speed before it can be used. A loss of driveability is, necessarily, implied in the operation of this powertrain control strategy.

A third example is a very elegantly constructed hybrid vehicle whose motor is on the end of the engine, in the place that would normally be occupied by the engine flywheel. Some slight mechanical rearrangement is required because the motor, whilst very short, is longer than the thickness of the flywheel. Again, this hybrid has a very simple control policy which results in poor driveability, though, having a diesel engine, the other disadvantages of on/off engine operation are lessened, see [75]. The last example reported in [74] is a vehicle adapted from a conventional four wheel drive vehicle in which the layout of the drivetrain allows the rear wheels of the vehicle to be driven by an electric motor and the front wheels to be driven by the petrol engine. No details of the control actions are given for this vehicle.

Reference [12] describes five hybrid vehicle concepts and gives details on the control strategies used. The first vehicle described uses two motors, an engine, a separate flywheel and an epicyclic gearbox. This high component count and cost would make the vehicle unlikely to be produced commercially. The control strategy that is used is very dependant on the unusual vehicle configuration and involves varying the ratio of the gearbox and the speed of the flywheel. The second example from this work is a delivery van with a very small diesel engine and an electric motor in a simple hybrid configuration. Where the diesel engine has insufficient power to meet the requirements of the driver, the motor is used, in addition, to power the vehicle. The reduction in the battery state of charge caused by this motor effort is replaced, to a certain extent, by regenerative braking in which the motor is used, as well as the mechanical brakes, to stop the vehicle, thereby increasing the charge in the battery. This simple idea achieved consider-

able improvements in fuel consumption if the range of the vehicle was relatively small, about 100 km. The third example involved simulating a passenger car in which the engine torque passed through an electric motor before driving a conventional transmission. The control strategy used in this case is to make the motor assist the engine at high throttle angles but to work against the engine at low throttle angles. The reasons for this approach, which would appear to cause very poor driveability, are not clear. The other vehicles which are simulated in this work are a high performance passenger car and a 10 tonne bus, both of which use rather unusual components and have control strategies that specifically make use of them.

A very early report on the control of a hybrid vehicle in which the engine torque passes through the motor is given in [76]. This vehicle has only a single gear ratio in its transmission and uses the motor when the car accelerates from rest. When the vehicle speed reaches 20 km/h the engine is started and is used to propel the vehicle. The powertrain controller then uses the electric machine as a generator for low pedal values and as a motor for high pedal values. Whilst the performance of the vehicle in terms of the maximum speeds and accelerations was quite comparable to the original Ford Escort estate on which the vehicle was based, the part-load driveability must have been worse because of the way in which the motor drives the vehicle from rest and acts against the engine at low pedal values.

A hybrid powertrain test facility is the subject of references [77] and [78] which describe the data acquisition equipment and the low level control actions that are required in order to control the powertrain in real time. The control actions taken at a higher level on the test bed are described in [79] (discussed below) but, disappointingly, the implementation details are not described in the practical work reported in [77] and [78]. Reference [80] is also from the same body of work and

describes the control of a gearbox, engine and motor to allow the transmission on the powertrain test facility to be automatically controlled.

Reference [79] describes a simulation study in which a reverse dynamic model of a parallel hybrid vehicle is used to investigate the correct control strategy. As explained previously, the reverse dynamic modelling approach involves finding the torque required in the vehicle driveshaft to meet the speed/acceleration operating point of the vehicle. Emissions drive cycles define the gear that should be used at each instant in the cycle, but in this work this constraint was relaxed and the objective of the work was to determine the values of gear ratio and torque split that minimised a cost function:

$$F = \lambda_1 E_1 + \lambda_2 E_2 \tag{3.1}$$

where E_1 and E_2 are the amounts of electrical and fuel energy used, and λ_1 and λ_2 are weights used to influence the relative amounts of energy used. At each time instant, the requirements of driveshaft torque and speed are defined and a direct search technique in the gear ratio/torque split space is used to determine the values of gear ratio and torque split that minimise F. However, this search would be very difficult for a vehicle to carry out on an instant-by-instant basis so a suboptimal technique is used that could be implemented in real time. This technique defines a box in the transmission input torque/engine speed space which is in the region of the most efficient operation of the engine. When the operating point of the powertrain falls inside the box the engine alone is used. Below the minimum torque line the electric motor only is used (and it is implied that the engine is switched off). Above the maximum torque line the engine and the motor are used together to meet the high torque demand. The speed limits on the box are the minimum speed of the engine and a disallowed high speed operation region. Where more than one gear can be used, the gear that places the engine nearest its most efficient operating point is used.

One disadvantage of this method is that the box that is used will vary as different values of λ_1 and λ_2 are used, though this in itself is not a large problem since any controller will have to vary its policy as its performance requirements change. What is more worrying is that the controller switches the engine off and this will have poor effects on the emissions performance of the vehicle and also on its driveability. A further problem is that there might be frequent changing of the gear ratio selected since the method does not seem to have any hysteresis around the switching lines. However, the method has the very great advantage of being motivated by the desire to look for the optimal control point in the control space.

It should be noted that it was possible to derive this control policy because of the vehicle modelling approach it would have been extremely difficult to do this using a forward dynamic model because the transmission input torque required value is not known. However, the use of reverse dynamic models has resulted in a control policy that has aspects that would be difficult to implement on a practical vehicle. The extent of of these difficulties could only be determined by a practical implementation of the controller.

Reference [81] considers the operating cost of a parallel hybrid vehicle. The method of operating the hybrid vehicle is that the engine is not used beneath 3km/h, above this speed two different control strategies are used, dependant upon the battery state of charge. These strategies are defined by considering the power split between the engine and the motor. The real difference between these two strategies is that, at constant vehicle speeds, for high battery states of charge the motor provides all the power, whilst, at low battery states of charge the engine provides all the power. During acceleration, the motor and the engine work together. The study varies the power split during acceleration, the speed threshold that defines when the engine will be used and the capacity of the vehicle battery. The effect that these parameters have on the vehicle operating cost is

then calculated. Again, these control strategies have the effect of introducing sudden bursts of engine use which would cause problems on a practical vehicle.

Another study that is driven by the economics of car ownership is presented in [82], in which the powertrain control policy is not a central feature. The details that are given of the control policy include on/off engine operation, use of the electric motor only at low speeds, sharing of power requirement at high power demands and cruising under the power of the heat engine. This powertrain control strategy is very similar to a strategy described in [83] and to the other examples.

3.2.2 Hybrid powertrain control strategies using constant power engine operation.

Reference [84] describes the control strategy used to control the operation of a very small parallel hybrid vehicle. In this work the accelerator pedal is taken to be a vehicle speed demand. Knowing the power required to meet the speed demand, the engine throttle is set to generate this power from the engine. Power for transient operation and for speed regulation on gradients is drawn from the electric motor. A serious problem associated with this approach is determining the correct power required to meet a given road speed, since the mass of a production vehicle in every day use can vary by around 30%, due to varying the number of passengers and other loads. The controller used in this work is unusual in that it made use of conventional linear control methods to control the speed of the vehicle and to provide the correct motor currents. A similar approach is described in [55] in which the performance of a variety of types of vehicles, including serial hybrid vehicles and buses, are simulated. It is found in this work that the parallel hybrid vehicle is upto 40% more efficient than an equivalent serial hybrid vehicle.

The motivating principle of this approach is the low fuel consumption and low rate of pollutant production of the heat engine when used at constant power. However, by making the control policy so prescriptive, the ratio of use of fuel

and electric energy cannot be changed and no real account is taken of efficient operation of the electric motor.

3.2.3 Serial hybrid powertrain control strategies.

Reference [85] contains a very theoretical technique to determine the constant operating power required of the heat engine in a passenger bus with a serial hybrid powertrain. The method is very complex and contains algorithms that are not necessarily convergent. It also has the advantages and disadvantages of constant power operation.

In [86] a serial hybrid vehicle using a natural gas engine is considered. The most efficient way of operating the engine and the generator is established and a control policy based on this is used. Since the paper concentrates on the interface between the engine, the generator and its associated control electronics, and this interface does not exist in a parallel hybrid vehicle, the work is not useful in the parallel hybrid vehicle arrangement.

3.2.4 Discussion of previous work.

The major drawback associated with the majority of the previous work is that the controllers that were used almost always switched off the engine reasonably frequently. This was probably not considered a great problem in previous work because the engines used were not fitted with a catalyst, or were diesels which do not have such large difference in operation between cold and warm running. However, the catalyst in a modern car will cool to a point where it stops converting, or "goes out" if the engine is merely left at idle for too long. Clearly, the catalyst will not work at all well when the engine is only being used intermittently.

The other major disadvantage associated with these controllers was that they did not consider the effects of the control actions on the driveability of the vehicle.

In the cases where the controllers were developed on real vehicles, the control actions have, in some cases, been shown to cause driveability problems. Where simulations have been used, the effect of driveability has tended to be ignored. A possible reason for this is that the simulations performed to date have tended to assume that the torque can be accurately set for any of the components without considering the control actions needed to schedule engine or motor torque. If there are regions in which the majority of the tractive effort swings rapidly from one prime mover to the other over small regions of pedal value, then the throttle angles and motor currents will have to change rapidly as the pedal value changes. In a practical vehicle, maintaining a smoothly changing output torque could be difficult where the components are not well behaved computer models but have unmodelled effects that will vary from one unit to the next.

3.3 Problem formulation and contrast with conventional control.

Figure 2.1 shows the role of the powertrain controller and it is seen that no aspects of feedback control, disturbance rejection or reference tracking are involved. In effect, this is not a conventional control application, though clearly, the hybrid powertrain is being controlled.

Referring to Figure 2.2 it is seen that the prime movers, whose actions are dictated by the powertrain controller, each have a torque output, the engine has a fuel flow rate output and the motor has a battery current output. On a production vehicle, shaft torques are not measured. This is partly because the measurement task is difficult and partly because measuring torques does not provide very much useful information on production vehicles. In this work, it is assumed that the engine output torque, the motor output torque and the driveshaft torques cannot be measured. Conversely, the battery current and the engine fuel flow rate are

very easily measured.

The controller has, therefore, to provide control actions whose effects cannot be completely measured or observed. The actions of the engine vary non-linearly, with both speed and throttle angle and this, the lack of observable outputs and the essential problem formulation have made the application of conventional control methods very rare in hybrid vehicle powertrain control. In the reported cases where conventional methods have been applied, they are only used in the low level dynamic control actions and are not part of the high level energy management strategy that is used.

The lack of applicability of traditional control methods is reflected in the high level powertrain control strategies that are used. In the practical examples these strategies were probably obtained by using engineering knowledge of the operating characteristics of the components. They have a form which is generally expressed by simple line diagrams or linguistic statements.

The hybrid powertrain controller could be made to work by the use of maps of the engine performance which would relate the generation of torques, the use of fuel and the battery current to the controlling inputs. This would allow the powertrain controller to choose the best control action at any instant, since, in theory, it would be fully aware of all the consequences of its actions. However, this would involve the storage of large amounts of data and would not take into account variations between individual components and variations due to ageing.

The nature of control strategies described in the previous section reflects the understanding of the engineers designing the vehicles. This is borne out by the ways in which these strategies are described. This sort of control system design approach in which the basic understanding of an engineer is embodied in a conceptually straightforward manner which is easily implemented in a vehicle is very common in the motor industry. There is no finer example of this than the design

of engine management systems, but other examples would include the design of transmission controllers and chassis control systems.

The nature of the hybrid vehicle powertrain control problem, past experience in hybrid powertrain control and the current working practices of the motor industry all indicate that the use of fuzzy logic in this application would be beneficial.

3.4 Application of fuzzy logic to hybrid vehicle powertrain control.

The central idea that lies at the origin of fuzzy logic is the usefulness of the vague way that humans think. In 1965 Lotfi Zadeh proposed the fuzzy set, [87], which was intended to represent the vague ways in which human beings handle concepts, particularly quantitative concepts such as "big", "fast" and "furry". In [88] and [89] the idea is extended to a method of reasoning that makes use of fuzzy sets and it is this idea that is the basis of fuzzy logic and fuzzy control.

Readers unfamiliar with fuzzy logic are referred to [90] and [91] for an introduction to the subject. For the complete novice [92] provides a very useful introduction to fuzzy control and an excellent coverage of a great deal of more advanced material.

3.4.1 Form of the fuzzy controller.

Figure 2.2 shows the inputs that the powertrain controller uses and the outputs that it will generate. The function of the controller is to relate the demanded throttle angle and demanded armature current to the pedal value. The real time control of the throttle angle and the armature current are achieved by lower level controllers that simply track the values generated by the fuzzy logic.

The fuzzy rules therefore have the following form:

IF pedal value is V THEN throttle angle is Θ

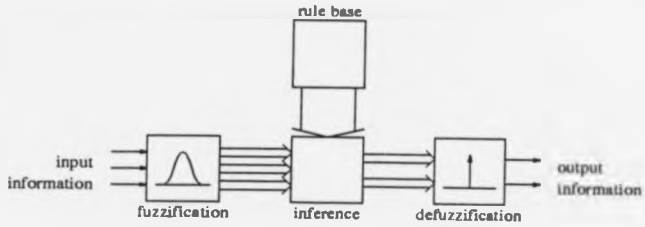

Figure 3.1: Fuzzy controller block diagram.

AND armature current is I.

where V, Θ and I are fuzzy sets defined on the pedal value, throttle angle and armature current universes of discourse. The rules, therefore, have a one input, two output form.

3.4.2 Initial inference method, defuzzification method and fuzzy set shapes.

When constructing a fuzzy controller, having determined the information on which the controller is to act and the information that is to be generated, the method of obtaining the output information from the input information has then to be defined.

The processes that are traditionally assumed to go on inside a fuzzy controller are shown in Figure 3.1.

The process of fuzzification involves taking a number from the outside world and making it fuzzy. Inference, as the term is used here, is the method of obtaining the inferred output fuzzy sets from the input fuzzy sets and the fuzzy rules. Defuzzification is the process of obtaining a crisp output value (a number) from the output fuzzy sets that were inferred by the fuzzy rules.

The first controllers that will be considered use the max-min method of implication in inferring the output fuzzy sets and take the centroid of the fuzzy union

of the output sets as the defuzzified output value.

The max-min method of implication used in the compositional rule of inference, as defined by Zadeh in [89], was used in the early applications of fuzzy logic for control. These early studies used a mean of maxima method of defuzzification, in which the defuzzified output is the value of the output universe of discourse which has the highest grade of membership in the union of the inferred output fuzzy sets. Where more than one output has the same maximum grade of membership a mean of these values is used. Examples of work using max-min implication and mean of maxima defuzzification include [93, 94, 95, 96, 97, 98, 99, 100, 101, 102].

The motivation for the mean of maxima defuzzification method is that it chooses the value in which the fuzzy logic most strongly believes. A serious disadvantage of this method of defuzzification is that it tends to produce output values that switch instantaneously between one output value and another value elsewhere. The reason for this is that normally, when fuzzy logic is being used, more than one rule will have non-zero grades of membership of its input variables in its antecedent fuzzy sets. Therefore, more than one output fuzzy set will be inferred over the output universe of discourse. As the input variables change, their grades of membership in the input fuzzy sets change, and the maximum grades of membership of the inferred, corresponding output fuzzy sets change accordingly. If, for example, two rules are inferred, then as the input values change, first one output fuzzy set and then the other will have the highest grade of membership. At the point at which the grades of membership are equal, the output value of the fuzzy logic swings immediately between the points with the maximum grade of membership of the output fuzzy sets. This discontinuity in the input-output relation of a set of fuzzy rules can be advantageous in applications where an actuator is being used which has only discrete output values, see [103], but is otherwise undesirable.

The defuzzification method that will be considered in the following sections is the centroid of the union of the output fuzzy sets implied by each of the rules. The principle advantage of this defuzzification method over the mean of maxima is that the output value of the fuzzy logic varies continuously as the input variables change. An early example of this method of defuzzification is [104], after which its use grew steadily. It is now used in the majority of fuzzy control applications. Reference [105] includes a detailed comparison of the two main defuzzification methods, finding the centroid method more advantageous.

Practical implementation.

This section will consider the issues that arise when implementing max-min inference and the centroid defuzzification method in a practical, automotive controller. The constraints of automotive control will be borne in mind. This means that the controller would have to be implemented on a conventional automotive microcontroller so that the performance aspects of the code in terms of time and memory requirements must be considered, see Section 1.5.

Appendix B.1 shows a pseudo-code implementation of max-min inference and centroid of the union of the output fuzzy sets defuzzification on a single input, single output (SISO) fuzzy controller. The upper loop obtains the grades of membership of the input variable in each of the input fuzzy sets and the two nested lower loops simultaneously generate and defuzzify the output fuzzy set. The input and output universes of discourse are defined by arrays that have points at all values that fuzzy sets are defined to have values. This limits the available fuzzy sets to piecewise linear sets, but, since alternative fuzzy set definitions tend to use exponents, see [90] and [106] for examples, or polynomials, see [92], [107] and [108], which are less easily implemented on real time control hardware, this imposes no real restriction in automotive control. Generally, functions will produce

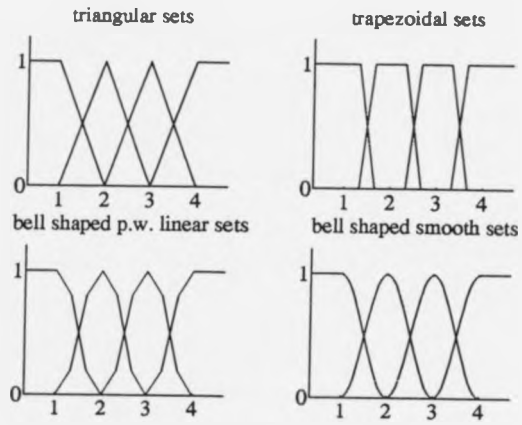

Figure 3.2: Four commonly used fuzzy set shapes.

a smoother output surface and in practice, a balance has to be struck between smoothness and numerical expense for each individual application. Figure 3.2 shows three examples of piecewise linear set shapes and one example of smoother fuzzy sets.

A modification that would allow the code to use less memory can be made if the fuzzy sets used over a particular universe of discourse all have the same basic shape. This allows the fuzzy set to be stored only once, but requires the use of offsets that define the position of the set in the universe of discourse. This procedure is used in [109]. The fuzzy sets (and hence the fuzzy rules) can then be stored as a series of offsets. In the case of the SISO controller, two single dimensioned arrays would then describe the controller. This modification would, however, make the code run slightly more slowly.

Linear interpolation is used within the arrays defining the input fuzzy set to obtain the exact grade of membership of an arbitrary input value in the set. Before the interpolation can be carried out, the location of the points that bracket the

input value in the universe of discourse array must be found, in the pseudo-code this is represented by the subroutine `getind`. The array can be searched from one end taking an average of $n/2$ iterations to find the bracketing indices, where n is the length of the array. Alternatively, for large arrays the vector can be searched using a variant of the subroutine `locate` in [49], taking rather less iterations, on average $\log_2 n$. A further refinement would be to use a search that started from the last element used and searched out from that point. This is because, in practice, the current value of an input to a fuzzy controller should move relatively slowly compared with the sampling interval and is, therefore, strongly correlated with the last value. These refinements would be used dependant upon whether the particular application justified the additional complexity, bearing in mind that each algorithm has overheads that are associated with deciding upon the next point for the search. Further discussion of a microprocessor implementation of a fuzzy controller is given in [109] and [110].

Having defined the code that will perform the processes of inference and defuzzification, it is necessary to define the fuzzy sets that will be used. When piecewise linear fuzzy sets are used in reports of practical applications of fuzzy control, the sets are generally either triangular, trapezoidal or bell-shaped. An important feature that determines the operation of the rules is the overlap between the fuzzy sets and the nature of the fuzzy sets in the region of the overlap. It has become common practice to use input fuzzy sets that overlap with their nearest as far as the point at which the nearest neighbour has its highest grade of membership, see Figure 3.2. The effect of this pattern of overlap for the triangular sets and bell-shaped sets, as drawn in Figure 3.2, is that for any value of input two rules are normally implied.

Trapezoidal sets, when used as shown in Figure 3.2, have the property that, in the region of the plateau, the grade of membership of the set does not vary as the

input varies. This means that the output sets will not vary and neither, therefore, will the output. In control terms, the effect of this is that, for these regions of input, controllability is lost. For this reason, unless a deadband is specifically desired, trapezoidal sets should be avoided. A more sensible way of maintaining the output at the same value, should this be desired, is to have two adjacent rules with the same output fuzzy set.

Each of the plots in Figure 3.2 shows lowest and highest fuzzy sets that have grades of membership that extend outwards at a value of 1. The highest and lowest sets should always extend far enough over the input universe of discourse to cover all the values of the input variable that will occur. Where the sets extend at ends of the universe of discourse, as shown in Figure 3.2, only one rule will be inferred. However, this rule will normally have the effect of, for example, demanding an extreme value from an actuator, which would be its maximum or minimum value.

The fuzzy sets that were initially used in the control of the hybrid vehicle powertrain are shown in Figures 3.3 and 3.4 together with the linguistic labels given to each set. The engine throttle angle can very between 0° and 90° so the universe of discourse was defined accordingly. The motor armature current can take values from -200 A to 200 A so the demanded armature current universe of discourse was defined to cover this range. The above discussion indicates that only piece-wise linear, non-trapezoidal sets should be used and accordingly bell-shaped sets were used initially.

There are no convenient, meaningful units of pedal value. The physical displacement of the pedal could be used, but, this in itself is not a particularly intuitive measure of pedal movement. Alternatively, on a hybrid vehicle, the pedal value will operate a potentiometer, rather than directly operate the throttle on the engine and the voltage on the potentiometer will ultimately be the input to the fuzzy controller. Initially, there is nothing very intuitive about the voltage on

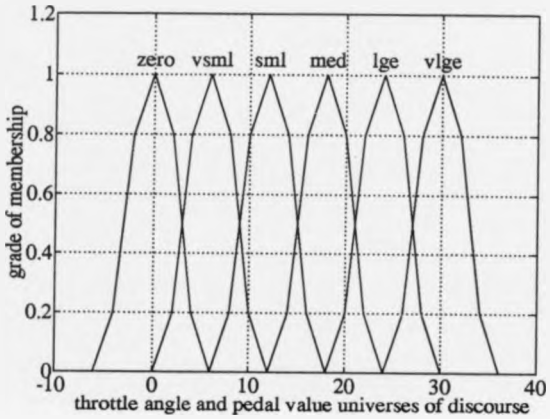

Figure 3.3: Fuzzy sets defined on the throttle angle and pedal value universes of discourse.

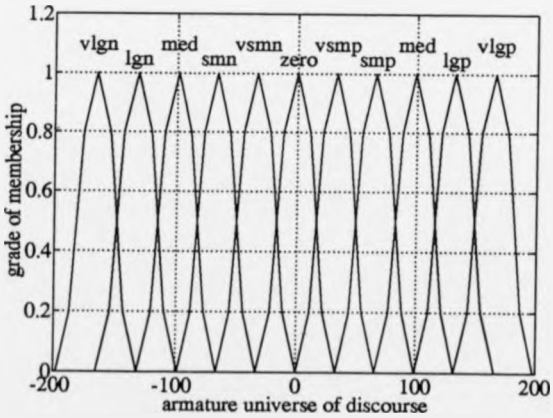

Figure 3.4: Fuzzy sets defined on the armature current universe of discourse.

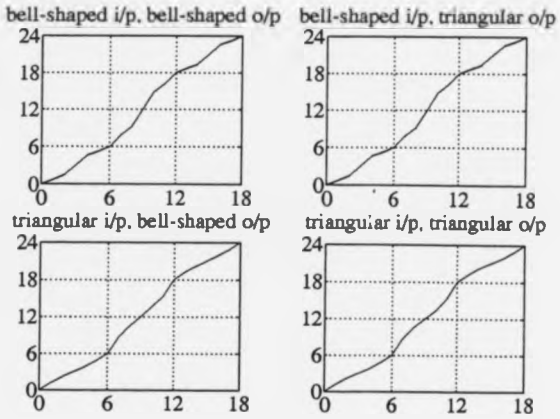

Figure 3.5: Plots of throttle angle against pedal value for an example hybrid powertrain fuzzy controller. The plots show the controller responses for different shaped input (i/p) and output (o/p) fuzzy sets.

this potentiometer either.

The units of pedal value that are used are "original degrees of throttle angle", which again, appear to have little immediate validity. The advantage of using this measure is that it allows a baseline controller to be defined whose performance is then readily compared with other controllers. The baseline controller demands a throttle angle equal to the pedal value and demands zero armature current. This has the effect of making the vehicle operate as a conventional heat engine vehicle. Pedal values from other controllers are then readily and intuitively compared with the pedal values for the baseline vehicle.

Figure 3.5 shows a comparison of the controller pedal value–throttle angle responses for the controller with the following four rules.

IF pedal = ZERO THEN throttle = ZERO AND current = ZERO
IF pedal = VSML THEN throttle = VSML AND current = ZERO
IF pedal = SML THEN throttle = MED AND current = VSMN

IF pedal = MED THEN throttle = LGE AND current = VSMN

The triangular sets used in the generation of Figure 3.5 have the same centres and degree of overlap as the bell shaped sets shown in Figures 3.3 and 3.4. The input set centres are at 0, 6, 12, and 18 and the output set centres are at 0, 6, 18 and 24. The four plots show clearly how fuzzy logic, when implemented in this manner, is really only a method of interpolation between the input and output set centres. The variations in the resulting responses indicate the arbitrary nature of the interpolation function that results from the use of differently shaped sets. The shape of the output sets used has relatively little effect on the controller response, when compared with the effect of the shape of the input sets. The bell-shaped input sets have the effect of smoothing the response as it passes through the set centres at the expense of introducing kinks elsewhere.

When triangular shaped input sets are used, the degree of inference of one rule or another changes linearly as the input point moves between the input set centres of the rules. Any non-linearities in the controller response are, therefore, introduced by the defuzzification process. If the centroid of the fuzzy union defuzzification method is considered geometrically, for output sets that do not overlap, it is readily shown to produce a non-linear output as the input grades of membership vary. Furthermore, when the output sets overlap, the form of this nonlinearity will change.

The conclusion of the above discussion is that when the max-min method of inference is used with the centroid of the fuzzy union of the output sets defuzzification technique, the choice of the set shapes that should be used is somewhat arbitrary. Different shaped sets generate different responses, however, the interpolative effects of fuzzy logic remain the same. Since varying the set shapes has no global effect on the controller response, improvements in the system performance when fuzzy control is used, which are achieved by different set shapes, must be

due to the localised effects of a system operating in a particular region of the control space for extended periods.

When the fuzzy rules have more than one input, as will be the case in subsequent sections of this work, the method used to determine the extent to which a particular rule is inferred, when max-min inference is used, is to take the minimum of the grades of membership of the inputs in their respective fuzzy input sets as the extent to which the rule is inferred. This means that changes in the less significant input do not affect the output of the rule being considered (though they may affect the output of other rules). This means that there is a loss of controllability for the less significant input for any particular rule, though not necessarily for the complete set of rules. A further, well known, consequence of taking the minimum is that the output control surface is not smooth, see [111].

3.4.3 Revised inference method, defuzzification method and fuzzy set shapes.

The methods of inference and defuzzification described in the previous section are far from the only ones that have been considered for use in fuzzy logic. In fact, in [112] 72 different versions of the strongly related implication operator are considered, [113] is similar. In addition to the minimum operator, engineering applications have generally restricted themselves to one other inference operator, the product operator, in which the grades of membership of the inferred output fuzzy set are multiplied by the grade of membership of the input in the input fuzzy set.

An advantage of this inference method is that the shape of the output fuzzy set is maintained. Intuitively, for the fuzzy logic purist, the inferred fuzzy set contains the same information as the original fuzzy set, but conveys this information less strongly. Comparing the product operator with the minimum operator, since the minimum operator changes the shape of the fuzzy set by truncating it, different

information is contained in the set. Whilst this is a rather conceptual distinction it does have a practical implication, if the mean of maxima method of defuzzification is to be used, because the output of the logic has to take discrete values.

The most significant advantage of product inference is found when using rules with more than one input. In this case, the accepted method of inferring the output fuzzy set, for any given rule, is to multiply the grades of membership of the input variables in the corresponding input fuzzy sets for that rule, and use the value obtained to infer the output set using the product operator. This method has the advantage of giving the inputs that have a lower grade of membership some influence over the consequences of the rule. This means that, in general, controllability will not be lost for the less important inputs and the output surface of the fuzzy controller will be smoother.

Another revision to the fuzzy methods outlined above is to use a modified defuzzification method in which the area of each inferred fuzzy set is used in the defuzzification procedure rather than the area of the union. Again, intuitively, there is an advantage in doing this because, in defuzzifying the sets shown in Figure 3.6, the shaded portion of the diagram will be counted twice, having been inferred twice. Additionally, in forming the fuzzy union of output sets, it is possible that a small set could be inferred that lay completely within the union of the other fuzzy sets. In this case, the information from this set would be lost by simply using the centroid of the output set union. This modified method of defuzzification will be called the moment of area method of defuzzification and is equivalent to forming the union using the summation operator.

This method of defuzzification can be stated mathematically as:

$$O = \sum_{i=1}^{n} o_i a_i \alpha_i / \sum_{i=1}^{n} a_i \alpha_i \tag{3.2}$$

where O is the output of the defuzzification, there are n rules, o_i is the centroid of the i-th fuzzy output set, a_i is the area of the i-th fuzzy output set and α_i is

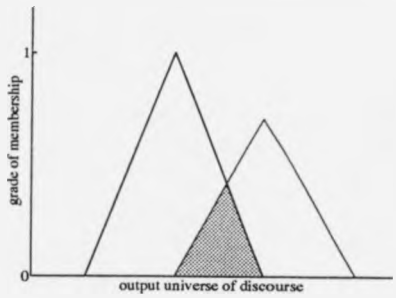

Figure 3.6: Shaded area is counted twice in modified centre of area method of defuzzification.

the extent to which the i-th rule is inferred. If the output fuzzy sets all have the same area, then Equation 3.2 reduces to:

$$O = \sum_{i=1}^{n} o_i \alpha_i / \sum_{i=1}^{n} \alpha_i \tag{3.3}$$

At this point it should be noted that the output of the defuzzification procedure is independant of the shape of the output fuzzy sets. If triangular input fuzzy sets that extend to the centres of the adjacent fuzzy sets are used, $\sum \alpha_i$ will always equal 1 and for any inferred rule, i, α_i will vary linearly with the input value. The defuzzification output value is, therefore, a linearly weighted sum of the output set centres and the complete process of inference and defuzzification reduces to linear interpolation.

The price paid for the use of triangular input sets is a reduction in the smoothness of the controller output as the input passes through the input set centres. Figure 3.7 shows the input-output responses of two controllers to two different sets of rules. The continuous bell-shaped sets that were used were generated by the use of a quadratic polynomial. It is seen that, in this case triangular input sets cause the controller to generate more sensible values than continuous bell-shaped sets. Both responses for the bell-shaped sets are continuous in both value and deriva-

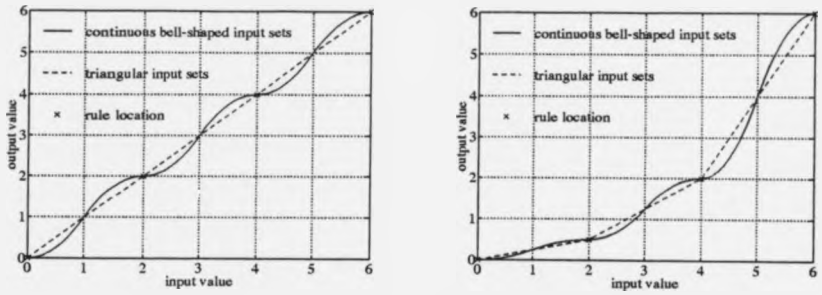

Figure 3.7: Comparison of controllers using continuous bell-shaped fuzzy input sets and triangular input sets with revised methods of inference and defuzzification.

tive, but this is achieved by enforcing a gradient of zero at the location of the rule. Note, it is impossible for controllers of this form to really closely approximate the curve that rule locations define since this would require the interpolation result to be influenced by points outside the current interpolation region. This could be achieved by the use of fuzzy sets that overlap to a greater extent so that more than two rules are inferred at any instant. However, this would interfere with the interpolative nature of fuzzy logic and it would not then be possible to guarantee that, in general, the input output relationship of a set of rules would pass through the points defined by the highest grades of membership of the input and output fuzzy sets. From the point of view of product liability this is a considerable disadvantage. If the fuzzy rules used have weights as suggested in [92], then it is possible for the combined rule set to have an output that is manipulated by the weights as well as the output set locations. This would allow the input-output relation to pass smoothly through the output set centres, at the considerable expense of making the controller very much more difficult to adjust. At this stage, the use of such rules was not considered.

The above discussion indicates that, in terms of input-output response, there is

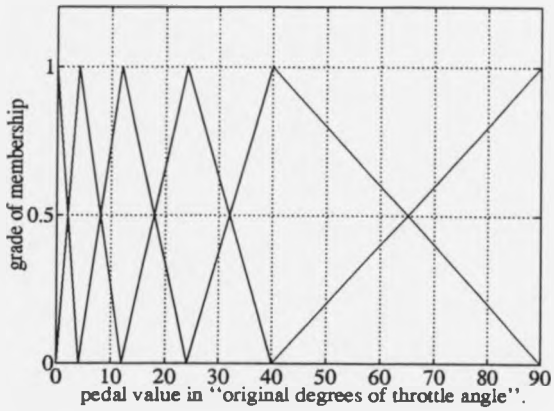

Figure 3.8: Input fuzzy sets used on the pedal value universe of discourse.

no clear advantage or disadvantage in using triangular input fuzzy sets. However, the use of triangular input sets allows the complete process of product implication and moment of area defuzzification to be reduced to n-dimensional linear interpolation, (where n is the number of inputs), see [114]. This gives these sets a very considerable computational advantage in a practical, particularly a real time, automotive, application.

Accordingly, the fuzzy input sets for pedal value that were used in the remainder of the work described in this chapter were the triangular sets shown in Figure 3.8. As previously stated, the shape of the output sets is immaterial and they are most straightforwardly defined as fuzzy singletons. For the throttle angle, these sets are located at the peaks in the pedal value fuzzy sets and, for the armature current, are located at the peaks of the sets shown in Figure 3.4.

The peaks in the input sets are at 0, 4, 12, 24, 40 and 90. The sets are nonlinearly distributed in the input space because virtually all the torque from the engine being modelled was available at about 40° of throttle angle. This meant

that small pedal values would be required to drive a conventional vehicle with this engine. Since one of the aims of the controller would be to make the vehicle feel consistent in use, accurate control in this region was needed. Early simulation experiments indicated that around six sets were required to get reasonable control of the powertrain and that, in order to get the vehicle model to behave well when moving from rest, particular attention had to be paid to very low pedal values.

3.5 Simulation results.

This section describes the results obtained using controllers of the type outlined in the previous section in the modelling environment described in the previous chapter. The parameters of the vehicle model are as given in Table 2.3, with the exception that the vehicle mass used is 1200 kg. There is an additional inertia for the hybrid vehicle since the inertia of the motor has to be added to that of the engine. The motor inertia used was 0.025 kgm^2. For the most part, the simulations were carried out using the European ECE-15 cycle as the vehicle speed reference data. Other simulations were carried out using Japanese drive cycles for comparison. The simulations continued to use an integration step size of 0.1 seconds and this was also chosen as the sampling frequency of the hybrid powertrain fuzzy controller. This time was chosen for the sampling interval because a practical vehicle controller could easily achieve an update time of 100ms and, if the sampling rate were to be reduced, the driver of the vehicle might start to notice the sampling interval.

3.5.1 Baseline controller vehicle simulation.

As stated in Section 3.4.2, the units of pedal value, namely "original degrees of throttle angle" are chosen because they allow a baseline controller to be defined, in which the throttle angle is equal to the pedal value and the armature current is

zero. The form of the fuzzy sets used allows any controller to be defined by three vectors of numbers. The baseline controller can be described by the vector triple:

$$\begin{matrix} [0 & 4 & 12 & 24 & 40 & 90] \\ [0 & 4 & 12 & 24 & 40 & 90] \\ [0 & 0 & 0 & 0 & 0 & 0] \end{matrix}$$

in which the first vector represents the location in the input universe of discourse of the point at which input fuzzy set has a grade of membership of 1, see Figure 3.8. The second vector represents the locations of the corresponding fuzzy output set singletons in the throttle angle universe of discourse and the third vector is the corresponding vector for the demanded armature current universe of discourse.

Since the same input sets are always used the controller can be represented more concisely, using just the output sets, by the vector pair:

$$\begin{matrix} [0 & 4 & 12 & 24 & 40 & 90] \\ [0 & 0 & 0 & 0 & 0 & 0] \end{matrix}$$

Using vectors to describe the controller removes the linguistic side of fuzzy control, however, the vector values soon become as meaningful as the linguistic terms and are a much more concise form of notation.

Figure 3.9 shows the values of five variables plotted against time for a simulation of the vehicle using the baseline controller (which causes the vehicle to behave as a conventional vehicle) over the ECE-15 drive cycle. The fuel consumption achieved by the vehicle over this cycle was 9.09 l/100km or 31.05 mpg. The negative values of pedal value shown in the second plot occur where the PI controller, that causes the vehicle to follow the reference speed (representing the driver of the vehicle in the simulation) applies the brakes to make the vehicle slow down. Positive values of pedal value are used as input to the fuzzy controller to generate the throttle value and the motor current demand. The action of the baseline controller is demonstrated in the third and fifth plots which show the throttle angle equal to the positive values of the pedal value and the demanded armature current equal to zero.

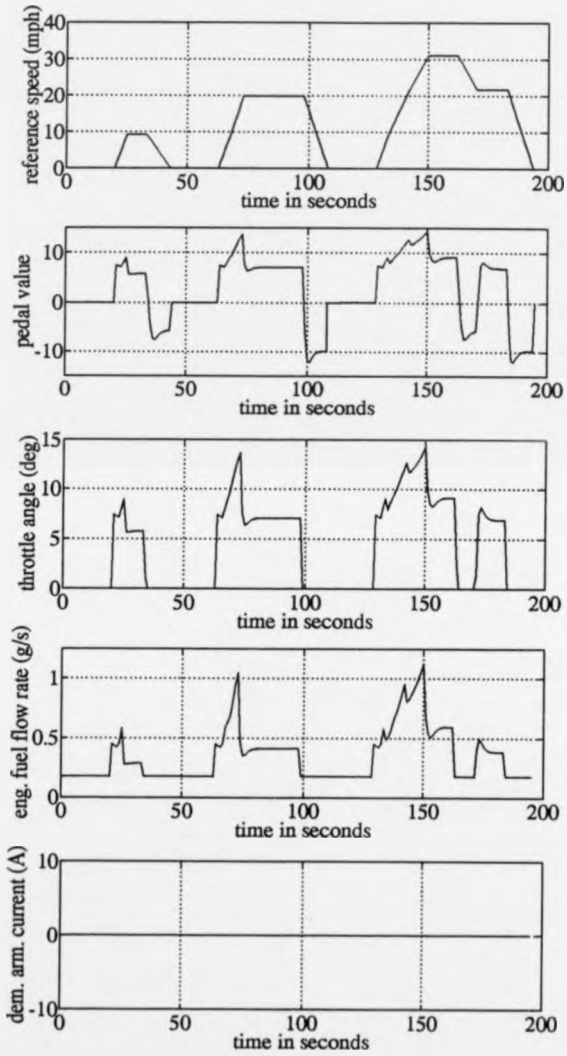

Figure 3.9: Parameter values against time for vehicle with baseline controller over ECE-15 fuel consumption cycle.

The fourth plot shows the fuel flow rate against time for the cycle showing the base value of the idle fuel flow rate (0.178 g/s), the peaks occurring as the vehicle accelerates and the steady flow rates as the vehicle maintains a constant speed.

3.5.2 Controller design procedure.

Having defined the baseline controller, many different controller combinations were used. The values of fuel consumption and change in battery state of charge obtained using these controllers over ECE-15 are given in Figure 3.11 and in Appendix C. In total around 40 different controller designs were investigated over 4 different cycles. The controllers were designed by considering the relative amounts of engine and motor activity that would cause the vehicle to have a desired fuel consumption and battery state of charge depletion. Simulation time histories were examined to determine modifications in controller design. Essentially this procedure is a fuzzy controller being designed and developed by a system expert. Figure 3.10 shows the simulation time histories of a typical controller whose fuzzy output sets are defined by the vectors:

$$\begin{matrix} [0 & 4 & 4 & 12 & 24 & 90] \\ [0 & 0 & 33 & 99 & 132 & 165] \end{matrix}$$

This controller uses the motor to assist the engine in powering the vehicle. Consequently, lower throttle angles are used than in the baseline controller, giving rise to lower fuel flow rates, and significant demanded armature currents cause positive battery currents to flow.

It was noted, in performing these simulations, that the maximum engine speed was just over 2600 r/min. The motor operates with low efficiency at low speeds such as this. It was, therefore, decided that a gear ratio between the engine and the motor, called m2egr, should be introduced. This gear ratio took the values 1 or 2. If m2egr is 2, the motor turns twice as fast as the engine and its output torque is multiplied by two before being added to the engine torque to give the

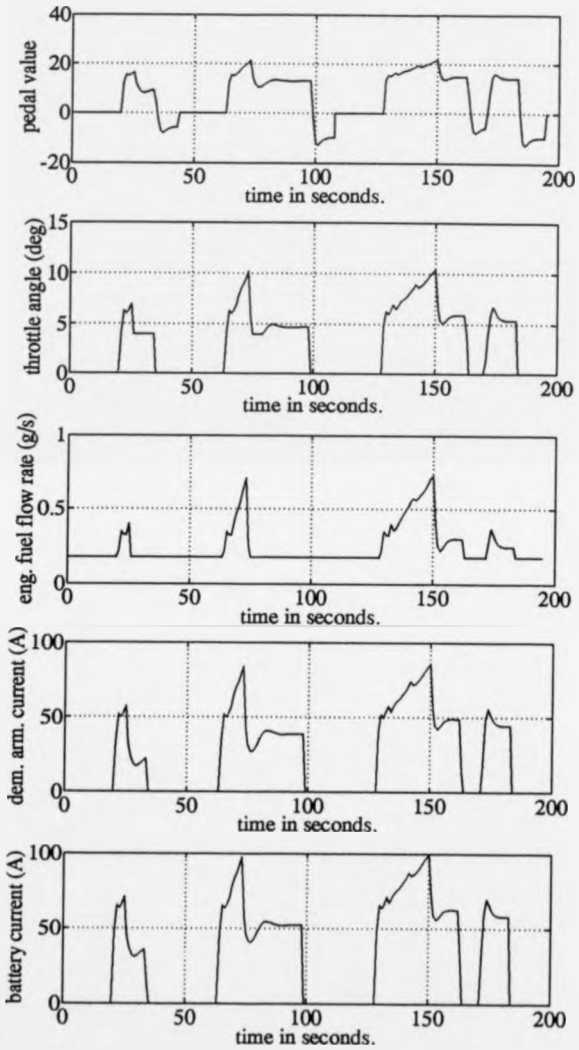

Figure 3.10: Parameter values against time for vehicle with a typical fuzzy controller over ECE-15 fuel consumption cycle.

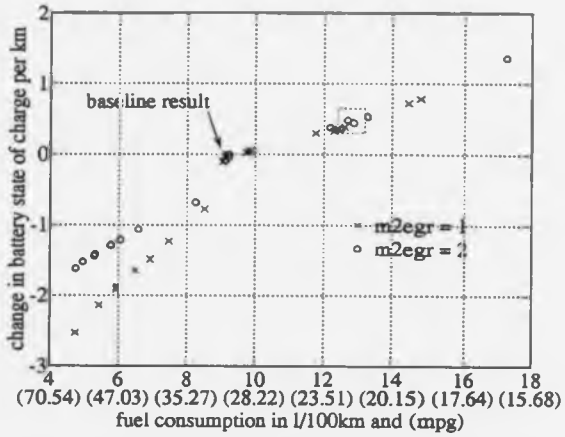

Figure 3.11: Simulation results obtained using different controllers over the ECE-15 drive cycle.

transmission input torque.

In general operation, it is desirable to use as little fuel as possible and to deplete the battery state of charge as little as possible, therefore, simulation results in the top left hand corner of Figure 3.11 are more desirable than results in the bottom right hand corner. In order to illustrate the points of interest arising from the simulation experiments, the two different types of results, those in which the motor assists the engine and those in which the motor works against the engine, are discussed in the next two sections.

3.5.3 Controllers that decrease the battery state of charge.

The results in which the battery state of charge is depleted by the controller (results to the left of the baseline result in Figure 3.11) are considered below.

The position of the baseline result with respect to the top left hand corner of the plot, would indicate that its performance from the standpoint of energy utilisation was the best. The poorer performance of the results to the left of the

baseline result can be explained by considering the efficiencies of the motor and the engine. When the demanded motor armature currents are low, the motor efficiency is low because a field current is required which will often be the nominal field current of 13.5 A. If the demanded armature current is in the order of a few tens of Amps, then the field current makes up a very significant proportion of the battery current and consequently the motor draws a high current but generates very little torque. This explains why, on the left hand side of the plot, the results move away from the baseline result in a direction parallel to the vertical axis, indicating a depletion in battery state of charge before a significant reduction in fuel consumption is achieved.

A further reason why the overall efficiency of the left hand results is low is that, when the motor assists the engine, there is a reduction in the fuel flow rate. Since the engine runs continuously, the lowest value that the fuel flow rate would fall to is the idle fuel flow rate. The two leftmost results in Figure 3.11 are achieved with the engine at idle throughout the cycle and the vehicle using electric energy alone for propulsion (in fact the m2egr = 1 result does not actually quite meet the demands of the cycle). These electric only results indicate the limitations imposed by the continuous operation of the engine. The form of the points shown in Figure 3.11 indicates that the combined effect of the field current and the idle fuel consumption is to cause a depletion in the battery state of charge that is never outweighed by the reduction in the fuel consumption.

Figure 3.11 serves only to qualitatively illustrate the effects of different controllers, since an absolute figure of merit for each controller is difficult to obtain because the usefulness of any given controller would depend upon the circumstances under which the vehicle operated. This would include the cost of fuel and electricity, the presence of emission regulations, the length of the daily journey and so on.

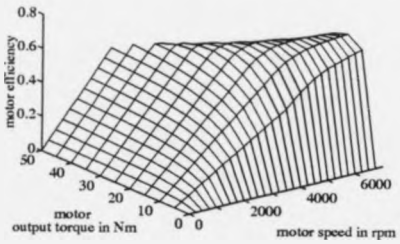

Figure 3.12: Motor efficiency against motor speed and output torque.

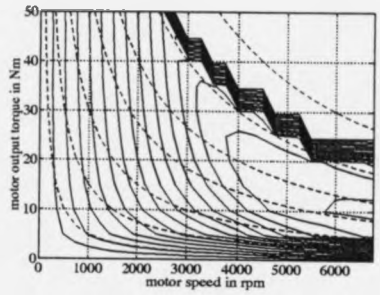

Figure 3.13: Contour plot of motor efficiency (solid lines) and motor power output (dashed lines).

One effect that is clearly indicated in Figure 3.11 is the value of using a motor to engine gear ratio of two rather than one. The reason for this is shown in Figures 3.12 and 3.13 which show the efficiency of the motor (including losses in the batteries) against its speed and output torque. The cascade of contours in Figure 3.13 is a result of the efficiency of the motor being calculated and stored in an array before the contours were calculated. The contours are close together because field weakening prevents the motor from operating in the region to the top right of the plot and the efficiency in this area was set at zero. The contours of constant power and efficiency in Figure 3.13 indicate that the motor will operate much more efficiently at higher speeds than lower ones, an improvement of 25-30% being possible for increasing the motor speed by a factor of two.

For any given controller, using m2egr = 2 results in twice as much torque going from the motor to the gearbox and this means that smaller pedal values are required to meet any given operating point. The smaller pedal values give rise to lower armature currents but also lower throttle angles and, therefore, the fuel consumption is reduced.

The battery state of charge is better for m2egr = 2 than for m2egr = 1 because

the system efficiency is a stronger function of motor speed than armature current and points of equivalent output power operate at higher efficiencies at higher speeds, see Figure 3.13.

The disadvantage of using m2egr = 2 is that at high engine speeds the motor could be damaged by exceeding its maximum speed. The main conclusion to be drawn from these results is that, since no controller clearly gives better results than other controllers giving similar results, the "best" controller will be determined as much by the requirements of a particular user as by considerations of the efficient operation of the powertrain.

3.5.4 Controllers that increase the battery state of charge.

The results of simulating controllers that increase the state of charge of the battery are shown to the right of the baseline result in Figure 3.11.

The results that are close to the baseline value move away from it in a direction that is parallel to the fuel consumption axis, indicating that there is an increase in fuel consumption before any useful increase in the battery state of charge. Again, this is because before any armature current can be made to flow into the battery, a field current has to applied by the battery. Where a negative current is demanded but a positive current flows, the model of the motor and motor controller assumes that the motor controller would not apply any currents, a consequence of which is that no torque is generated by the motor.

In contrast to the results which reduce the state of charge of the battery, there is no great improvement in the performance of the controllers for m2egr = 2 as opposed to m2egr = 1. The reasons for this are shown in Figures 3.14 and 3.15, in which it is seen that the contours of constant power, in regeneration, follow the contours of constant motor efficiency. Consequently, little advantage is gained by the use of higher motor speeds.

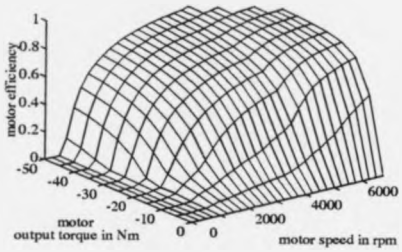

Figure 3.14: Motor efficiency against motor speed and output torque for positive armature currents.

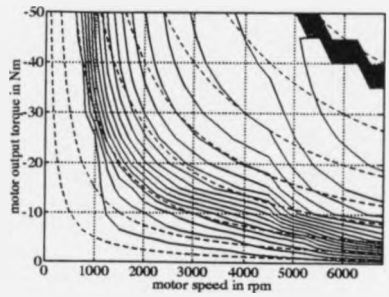

Figure 3.15: Contour plot of motor efficiency (solid lines) and motor power output (dashed lines) for positive armature currents.

Possibly the most interesting results are the two marked with circles in the dotted box in Figure 3.11, in which one result has both very slightly better fuel consumption and a considerably better battery state of charge than the other. This is one of only two instances of both performance measures improving simultaneously and the parameter values for the simulation are shown in Figure 3.16. The controller whose performance is shown with the solid lines (controller A) has the following output throttle angle and demanded armature current sets:

$$\begin{matrix} [0 & 4 & 12 & 24 & 40 & 90] \\ [0 & 0 & -33 & -99 & 0 & 165] \end{matrix}$$

The other controller, (controller B) is described by the vectors:

$$\begin{matrix} [0 & 12 & 24 & 24 & 40 & 90] \\ [0 & -33 & -66 & 0 & 0 & 165] \end{matrix}$$

The first point to note from Figure 3.16 is that the pedal values required to cause the vehicle to follow the reference speed are very different and they are also very different from the pedal values required for the baseline vehicle. This implies that the vehicles would feel extremely different to drive even though they actually have the same ultimate performance. The large difference from the pedal values used

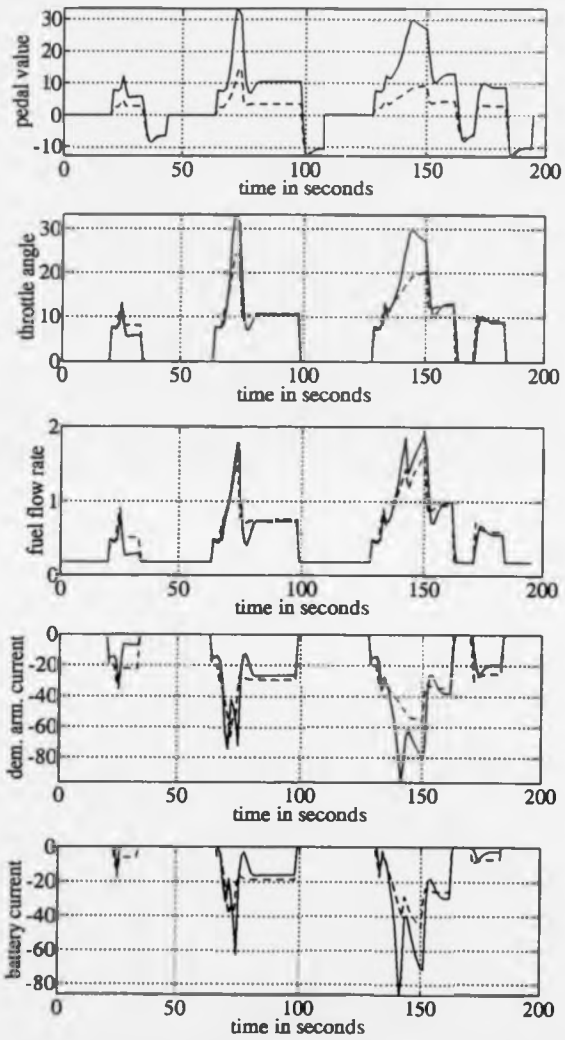

Figure 3.16: Parameter values against time for vehicles with different controllers (controller A shown with solid lines, controller B shown with dashed lines).

by the baseline controller would make them unacceptable from this point of view. The variation in the required pedal values is caused by the limitation imposed by restricting the available output sets to predefined values. The resulting control actions cannot then be finely adjusted to give an improved system performance.

Some early simulation experiments were carried out using fuzzy controllers of the type discussed in Section 3.4.2. Some of the controllers that were used in these experiments resulted in regions in which the controller caused the driveshaft torque to fall as the pedal value increased, or even, in extreme cases, generated negative driveshaft torques for positive pedal values. A disadvantage of applying fuzzy control in the manner described here is that there is no way that the driveability of the vehicle can be maintained when hand tuning the controllers.

The reasons for the improved performance of the controller with the solid line are also shown in Figure 3.16. On the third hill controller A demands a more negative current than controller B and, because the motor is operating relatively efficiently at this point, generates a large negative battery current at the expense of a small increase in the fuel flow rate. Conversely, on the first hill controller B demands a higher armature current but, because the motor is operating relatively inefficiently at this stage, has to increase the fuel flow rate a disproportionate amount for the improvement in the battery state of charge that is gained. By making efficient use of the motor in this way, controller A is able to increase the battery state of charge 10% more than controller B whilst using less fuel. In the other two points which display a similar improvement in both performance measures the improvement occurs for similar reasons.

It should be noted that it was not possible to predict which controllers would do well and which would do badly. To some extent, this is due to the effects of the fuzzy interpolation method since, occasionally, a controller might happen to be very efficient in a region which was used frequently by the cycle. Where the

cycles have large periods of steady state operation this is particularly likely. A further consequence of this is that controllers that did well on one cycle might not do well on another, the next section briefly looks at results obtained on other cycles.

The main reason why some controllers are very much better than others for m2egr = 2 but not for m2egr = 1 is that, as seen in Figure 3.15, motor efficiency is much more dependant on motor output torque (and therefore armature current) at high motor speeds than at low motor speeds. Since the controllers vary the armature current demand, this accounts for the variation in the efficient utilisation of the energy resources of the vehicle.

3.5.5 Other remarks.

In order to investigate whether the performance of the controllers was dependant upon the nature of the drive cycle used, the simulations whose results are shown in Figure 3.11 were repeated for three Japanese drive cycles also represented by straight line speed-time profiles. It was found that, for one of the cycles, the same two instances of both performance measures for one controller being an improvement over another occurred. For the other two Japanese cycles this did not occur because, at the pedal values that demand the correct armature currents for the potentially increased efficiency, the motor was not operating at a speed at which it has markedly different efficiencies for the different controllers. This indicates that different results will be obtained over different drive cycles and also that in order to improve the efficiency of the combined powertrain operation it might be useful to have a second fuzzy input of powertrain speed. However, if this second input were included, the tedious task of hand tuning the controller would be further complicated.

Some work was done in which the powertrain controller acted on both positive

and negative pedal values to allow the vehicle to make use of regenerative braking to increase the battery state of charge by using the motor in addition to the mechanical brakes. In order to make this work the brake torque had to be brought under control as well and so a further fuzzy output, demanded braking torque, was created. The creation of this output assumes that in a practical vehicle braking torque could be controlled and this is current practice in the automotive industry in anti-lock braking systems and slip-reduction or traction control systems. By using regenerative action in this way it was possible to improve the battery state of charge by around 0.17 Ah/km. The control of the regenerative action of the motor during braking represents an interesting sub-problem, but is less complex than the combined operation of the engine and the motor. The regenerative operation of the motor during braking could hide some effects of the actions of the controllers during accelerations and cruises and, in order to focus attention on the engine-motor interaction, regenerative braking is not used in the remainder of this work. It should be borne in mind that regenerative braking is the only instance of "something for nothing" in the hybrid vehicle because the energy regained by regenerative braking would otherwise be lost as heat in the braking system of the vehicle.

Finally, a revision was made to the model, in that the actions of a one-way clutch on the output to the engine were simulated. Without this clutch, if the fuzzy controller demanded 0° throttle angle and positive armature currents at the motor, the motor has to supply the power for propelling the vehicle along the road and turning the engine. The need for this clutch was realised early in the work and the results discussed above were obtained using a vehicle model including a one-way clutch.

3.6 Comments on the role of fuzzy logic in control engineering.

The recent surge in interest in fuzzy control has lead to the production of various VLSI devices which implement the inference and defuzzification processes in analogue parallel hardware, see [115, 116, 117]. Some of the fuzzy hardware was developed from digital signal processing devices which are discussed in [118]. The use of fuzzy hardware would allow some of the disadvantages described in Section 3.4.2, particularly the slow operation of the code, to be overcome when the methods described in that section are used. Similarly, some simulation tools now offer fuzzy control simulation facilities, see [119]. As has been shown in the previous section, no particular disadvantages are attached to a straightforward implementation of fuzzy control that would not require specific hardware or simulation tools which will, inevitably, be relatively expensive.

It has been shown in the previous sections that fuzzy logic can really be thought of as an interpolation method between fuzzy sets. In fact, this interpolative property was made use of in [103], when crisp rules for a discrete system where extended using fuzzy logic to control a continuous system. This interpolative nature has important consequences when considering commonly held beliefs about fuzzy logic. There is no reason why, when used in feedback reference tracking control, a fuzzy controller should be robust, unless the rules in the controller have been specifically designed to be tolerant to plant variations and sensor noise. Even when account has been taken of these factors in the design of the rule base, the fuzzy approach would give no indication of the degree of robustness. Fuzzy control makes use of expert understanding, and generally, experts will have a very good qualitative understanding of the operation of their system. However, when considering issues such as robustness and stability an accurate *quantitative* knowledge of the system is required and, for applications suited to fuzzy control, this knowledge is

unlikely to be available. These comments apply only to fuzzy control that makes use of expert knowledge to design and tune the rule base, they do not apply to the various forms of adaptive or learning fuzzy control that are discussed in the next chapter.

Once the limitations and advantages of fuzzy control are understood, it is sensible to try to define the sorts of control problems that would benefit from its application. The most obvious class of problems that would benefit from the application of fuzzy control are those for which there are no other control solutions, a good example of which is the hybrid powertrain control problem.

A second group of potential fuzzy control problems are problems in which the relevant input information is not obvious if conventional control methods are applied. Examining a control problem from the standpoint of a human being, rather than a controller, may reveal that relevant control information is obscure and clarify its nature. This is borne out by the number of novel, and apparently ideal, applications of fuzzy control that involve the use of new types of sensors. Having been identified, and then instrumented, this information can then be used, with benefit, in other control methods. A further advantage of fuzzy control is that, in some applications, it may generate a hardware controller very quickly and without the involvement of a specialist control engineer. This is because fuzzy control can be viewed as simply a method of allowing a component specialist to express control ideas to a software engineer in a concise manner. Fuzzy logic, when applied in this manner, fosters involvement, team working and shared responsibility, which are increasingly important in competitive industrial engineering. It is this aspect of fuzzy control that may be its greatest single advantage in commercial applications.

3.7 Conclusions.

This chapter began by reviewing previous work in the area of automotive hybrid vehicle powertrain control and found that the control methods adopted were developed largely with the powertrain efficiency in mind. In the simulation studies that were reviewed it was felt that some difficulty might be found in implementing the work reported in a practical vehicle. In the work that was conducted on vehicle prototypes some driveability problems were encountered because of the on-off operation of the engine and the motor. In the light of petrol engine emission control methods on-off engine operation is no longer feasible and other control methods that take this into account should be found.

The powertrain control problem was then considered in detail and it was pointed out that the lack of applicability of conventional control methods and the tendency of automotive engineers to favour techniques which have a practical orientation and make use of their experience makes this problem a good candidate for fuzzy control. The exact form of the fuzzy logic controller was then considered and it was found that inference methods, defuzzification methods and fuzzy set shapes that could be justified by considering the nature of their resulting input-output relations could be implemented in compact, fast code that could be used in a real-time automotive application.

The results of simulations which used the models described in the second chapter and the fuzzy controllers developed in this chapter were then presented and discussed. The objectives of the powertrain controller are set out in Section 2.2. The first requirement, was that the controller should make use of the energy resources of the vehicle in the correct ratio. This requirement was partially met. In practice, when deciding to use a particular controller, the author had some idea of the possible effects of the controller but was unable to make more than vague predictions about the relative amounts of fuel or electrical energy that would be

used.

The second requirement, that the demands of the driver be met in an energetically efficient manner was occasionally achieved, though more by luck than judgement. It was found that efficient use of the powertrain depended not only upon the demanded armature currents and throttle angles, but also on the powertrain speed. It would seem sensible to take account of the speed in the control actions, however, hand tuning such a controller would be impractical.

The last requirement, that of ensuring a consistent response to changes in pedal value was not met at all. These comparisons of the performance of the controller with its requirements indicate that some method of automatically tuning the controllers so that they maintain vehicle driveability, use the energy resources of the vehicle in the correct ratio and result in efficient operation of the powertrain should be used. Also, in use there are variations in the performance of individual engines and also variations in the way in which the motor will perform as the battery discharges and it would be desirable for controllers to adapt to these changes as they occur. Lastly, different users have widely different patterns of vehicle usage, and the controller that is correct for one user may well be inappropriate for another.

In order to develop a controller that is able to adjust itself to better meet its performance requirements, some familiarity with self-organising, adaptive and learning fuzzy control methods was felt to be important and, to that end, the next chapter reviews work that has been done in this field.

Chapter 4

Methods of fuzzy controller tuning, adaptation and learning.

4.1 Introduction.

The previous chapter described the implementation of a fuzzy controller used in the management of a hybrid vehicle powertrain in simulations of a hybrid vehicle over drive cycles. This controller was tuned by hand and some representative simulation results have been presented. It was seen that, whilst fuzzy logic provided an excellent tool for tackling the control problem, the three requirements of the controller were difficult to meet when the controllers were manually tuned. It was also pointed out that a more effective controller would include the powertrain speed as an input and that, whilst hand-tuning the single input controller had been difficult and only partially successful, hand-tuning the two input controller would be completely impractical.

These observations motivate the use of some automatic controller modification scheme to both assist in the design of the fuzzy controller and to modify it to improve its performance in use. This chapter will survey work in the field of fuzzy control that addresses the issue of automating the modification of fuzzy controllers. This is not intended to be an exhaustive review of this area of work as the production of such a document would be an almost never ending task. The

object of this survey is to cover the major techniques and methods that have been used and to discuss their relevance in the control of a hybrid vehicle powertrain.

4.1.1 Terminology.

The definitions of the terms used in the title of this chapter have been the subject of considerable attention. The large number of different definitions, used to cover the same term, was discussed as early as 1977 in [120], which emphasized differences in opinion at a time when the subject of adaptive control (for want of a more general term) was not very mature. As recently as 1989, Åstrom and Wittenmark, [121], still felt that formal definitions were absent and adopted a pragmatic definition for adaptive control. It would seem sensible that, when using terms for which no universally accepted definition is available, definitions appropriate to the context of the discussion at hand should be stated and used.

The term "tuning", as it is used here, means the adjustment of a controller whose basic method of operation has been established (fuzzy, PI, etc.) in order to improve its performance in a particular application. The adjustments would be necessary because of time invariant and unmodelled uncertainties in the system to be controlled. Tuning would, therefore, not include an automated continual process of on-line controller modification.

The terms "adaptive" controller and "learning" controller are more difficult to define but, for the purposes of this discussion, would both include controllers which are automatically modified in use to improve their performance in specific circumstances for particular applications. The systems being controlled by these controllers would have uncertainties which were time variant or non-linear structures that would require a different controller in different regions of plant operation. As stated previously, attempts have been made to define rigorously the terms "adaptive" and "learning", but none have really gained universal ac-

ceptance. Possibly, a reason for this is that the exact definition of such terms is not actually *useful*. Perhaps, in this case, it is useful to make a distinction between them based upon dictionary definitions for adapt: "to adjust (someone or something, esp oneself) to different conditions, a new environment, etc." and for learn: "to gain knowledge of (something), to acquire skill in (some art or practice), ... to commit to memory", see [122]. The main difference between the two definitions being that learning involves knowledge acquisition and remembering, but adapting really only means changing. The distinction that can, therefore, be made between adaptive controllers and learning controllers is that an adaptive controller will change its actions from one circumstance to another but, being incapable of benefitting from experience, would have to use the same process to go back to the first situation. On the other hand a learning controller would make use of previous experience to determine, as it moves from one situation to another, why the control actions need to be modified. In this way, it is more likely to perform well in situations which it has not previously encountered. For a detailed consideration of this distinction see [123].

When looked at in this manner, tuning could be considered to be a subset of adaptation, adaptation a subset of learning and learning a subset of true intelligent control, although, given the vagueness with which these terms are understood and used, they are clearly fuzzy subsets!

4.2 The Self-Organising Fuzzy Logic Controller.

Early reports on fuzzy control [93, 98, 99, 100] indicated that, having designed a fuzzy controller, it was normally necessary to tune it by some means to improve the performance of the original controller. This need to modify the controller was indicated as early as 1975, by Mamdani and Baaklini, [94], which describes a "rule-modification algorithm".

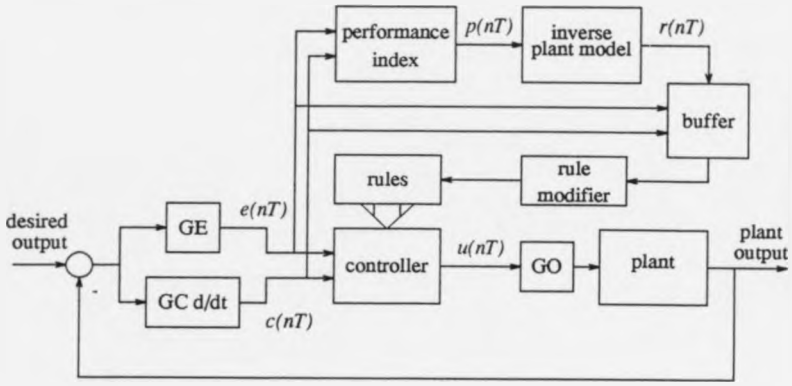

Figure 4.1: Structure of the single input single output self organising fuzzy logic controller.

The first widely adopted method of automatically modifying a fuzzy controller is given by Procyk in [124], the important details of which are summarised in [102] and [125]. This method will be known as the Self Organising Fuzzy Logic Controller (SOFLC).

4.2.1 Self Organising Fuzzy Logic Controller structure and operation.

The structure of the SOFLC, for single-input single-output (SISO) plants is shown in Figure 4.1. The fuzzy controller itself is defined to have inputs of plant output error, and change in output error, and to generate an incremental output which is the change in desired plant controlling input. It operates in discrete time, and on discrete universes of discourse, using three gain coefficients which scale the two inputs and the output, and which are akin to the gain coefficients in a linear PI controller. Using an incremental controller output has the effect of integrating the controller input values, making the controller analogous to the linear PI controller.

In order for a controller to improve its performance, it has to be able to measure its current performance to determine if it is improving. The SOFLC designer defines a decision table called the performance index and, in doing so, effectively defines the closed loop performance that will be obtained from the controller and the plant when in use. Since, clearly, the closed loop system cannot have instantaneous transients and perfect tracking, the performance table will have to incorporate some expectations of the achievable performance of the system. The performance table is an instantaneous view of the performance of the system. That is to say, if the output is away from the setpoint, but the change in error of the output is such that the error is being reduced, the performance table should indicate adequate performance. It is, therefore, possible to coarsely define the damping, overshoot and so on using this performance measure.

A further feature of the SOFLC is the inverse plant model. In the case of SISO plants, this model is simply a number that inverts the "gain" from the plant input to the plant output. Having determined the desired plant output from the performance measure when particular values of error and change in error occur, the inverse plant model can then be used to determine a plant input value that should be used for those values of error and change in error. Since the controller relates error and change in error to plant input, the effect of the performance measure and the inverse model is to suggest "correct" controller output values. Should the controller output values differ from the suggested values, the suggested values can then form the basis of a controller modification.

The knowledge of the control engineer is also needed in making some assessment of the dynamics of the plant. When the controller determines that the current closed loop system performance is poor, there is a need to relate the control actions that caused the current poor performance to the modifications suggested by the performance measure and the inverse plant model. Since the plant has

lags and delays, the control actions that lead to poor performance will have been taken some time in the past and should be modified so that, when similar circumstances are encountered in the future, more appropriate control actions will be taken. The method by which this is achieved is termed credit assignment. The fuzzy control action that should be modified is determined by a delay (called the delay in reward). The delay in reward is the number of samples in the past that the most significant incorrect control action is judged to have occurred. Workers extending Procyk's ideas subsequently applied the modifications to more than one incorrect control action in the past. The function of the buffer in Figure 4.1 is to store the rules upto the delay in reward for later modification.

If the delay in reward is m, T is the sampling time and the current time is nT, then the past control action:

$$e((n-m)T) \rightarrow c((n-m)T) \rightarrow u((n-m)T) \tag{4.1}$$

should be replaced by the control action:

$$e((n-m)T) \rightarrow c((n-m)T) \rightarrow u((n-m)T) + r(nT) \tag{4.2}$$

where $e(nT)$ is the plant error, $c(nT)$ is the plant change in error, $u(nT)$ is the plant input, $r(nT)$ is the output from the performance index and the inverse plant model (known as the reinforcement), and m is the delay in reward. The symbol $\rightarrow$ indicates the process of implication.

The method by which the replacement is effected is by modification of the relation matrix. Procyk's original idea was to use a linguistic statement that would effectively remove the incorrect rule from the matrix and insert the new rule using fuzzy set operations such as union and intersection. In practice, this involved the manipulation of large amounts of data so that instead of modifying the relation matrix the fuzzy rules themselves were modified using De Morgan's Theorem. A

further saving was achieved by only modifying the rules if a sufficiently different rule was created.

4.2.2 Adjusting the performance of the self-organising fuzzy controller.

Experimental work performed on the above structure indicated that there are some fundamental compromises in this approach to self-organising fuzzy control. The overall performance obtained after a period of learning is determined by the performance index, the plant and by the gain factors GE, GC and GO. In particular, the gain factors are important as the most easily "tunable" parameters of the system as a whole.

The effect of the gain parameters is most easily understood in terms of the effect that they have on the signals reaching the performance index table. As the scale factors GE and GC change, the same values of error and change in error will be mapped into different areas of the performance table and will, consequently, cause different values of reinforcement to be generated. In turn, this will affect the controller modifications to more closely reflect the effect of the scale factors on the performance index.

For example, by changing the gain factors the "tolerance band", as it is known in [102], is adjusted in size, relative to the external error signals before GE and GC are applied. The tolerance band is the area of the performance index in the region of zero error, in which the control actions taken are modest since the plant is in the region of the set-point. Outside the tolerance band, the the control outputs are chosen such that the error is driven towards zero. Consequently, by changing the scale factors, different values of the error before scaling will be mapped into the tolerance band. This has the effect of modifying the ability of the controller to cause the plant to track a reference point and to approach the reference point from areas outside the tolerance band.

The non-linear nature of the performance table should, effectively, allow the gains of the closed loop system to be scheduled in order to achieve a rapid transient response and good set-point tracking. However, in many examples of reported work which use this basic method of self-organising fuzzy control [102, 126, 127, 128, 129, 130, 131] the same compromise occurs. Transient performance and set point tracking are traded off against one another in adjusting the gain factors. This is possibly due to the fact that simple linear gains where used to map the real error variables into the performance index, whereas a non-linear mapping might have allowed a more optimal compromise between tracking and response.

Generalising to multivariable SOFLC.

The structure outlined above is easily generalised to multivariable systems by the use of an inverse of the Jacobian of the plant model. This is illustrated in [102, 124, 126, 127] and also in work presented by Linkens and Hasnain [132].

Additions and extensions.

In addition to the original work, there have been additions and extensions as other workers have tried to improve the SOFLC. The important improvements are discussed briefly below.

The original work of Procyk was extended by Yamazaki and Sugiyama in their PhD theses [105] and [133]. Yamazaki investigated the use of different methods of implication and defuzzification, finding that product implication and the fuzzy centroid method of defuzzification have advantages in smoothness of output response and ability to make decisions in the absence of complete information. Previously, Procyk had used max-min implication and mean of maxima defuzzification. Yamazaki also attempted to optimise the performance of the controlled transients. Possibly the most important addition came with the use of the multi-

delay in reward, in which the modifications to the controller are applied to more than one rule that was executed in the past. This had the effect of reducing the sensitivity of the system to noise.

Sugiyama did further work on the implication method and extended the technique to consider the use of continuous inputs. The PI controller was extended by the use of an additional non-fuzzy term to become a PID controller, allowing the method to be more convergent for plants of high order. Possibly the most interesting feature of the work reported in [133] is the adaptation of the delay in reward, and GO scaling factor, as variations in the plant time delay and gain occur. These adaptive actions were achieved by the use of higher level fuzzy rules, representing the use of human intelligence in observing the state of the system. This leads naturally to the subject of hierarchical control with which Sugiyama concludes his thesis.

The work of Daley and Gill, reported in [126] and [127], is an application of the self-organising fuzzy controller on a plant with considerable complexity. In [127] they use a method of switching the scale factors in order to get the benefit of good transient performance and also good set point control. This method was later adopted by Linkens and Abbod in [130]. Both [127] and [130] report a considerable improvement in performance.

In [129] Shao "compiles" the relation matrix to produce a decision table or map. The map is then adjusted in order to avoid recomputing the entire relation matrix. The purpose of this extension is to facilitate easier implementation in hardware. Further, in order to improve the set point control, the error signals pass through a non-linear process to give the system more sensitivity for low values of error.

Farbrother, Stacey and Sutton and Sutton and Jess report the use of over-rules in [131] and [134], which assist in the convergence of the controller and, whilst in

[134] they are expressed in terms of the specific problem, they are generically applicable. One of these over-rules enforces the generation of a symmetric controller, allowing the controller to converge more quickly. Sutton and Jess also made use of the continuous variables proposed in [133].

4.2.3 Comments on the self-organising fuzzy logic controller.

SOFLC in feedback control.

The SOFLC is now a well understood control method and, as has been seen above, has been the subject of considerable development. Furthermore, SOFLC has now been used in hardware, [128, 129, 130, 135], demonstrating its value in practical situations. The high numerical burden of SOFLC has caused some problems and, in one case, use was made of transputers in order to obtain sufficient processing speed for a real time application, see [130].

A considerable disadvantage of the SOFLC is that it only operates on error signals rather than absolute signals and, therefore, is unable to distinguish between one region of the plant operating space and another. When operating around a particular set-point, it has been shown, see [133], that the SOFLC will adapt to the non-linear features of the plant operation in the region of the set-point. However, if the set-point is changed, then the SOFLC will have to adapt again in the new region of operation. Additionally, if the plant is returned to a set-point that has been previously used, the same process of adaptation must be undergone before the control performance previously achieved can be obtained once more. This reveals the adaptive rather than learning nature of the SOFLC.

SOFLC in the management of a hybrid vehicle powertrain.

SOFLC is not readily applied to the control of a hybrid vehicle powertrain as it is considered here, because hybrid powertrain control does not belong to the class of feedback control problems. In this respect, SOFLC is no more applicable than conventional linear control methods. It is possible that it might find application in the lower level control actions in a hybrid powertrain, or, in other feedback control applications in the automotive industry. Unfortunately, the numerical burden and memory requirements of SOFLC, see [92], would be very prohibitive in any automotive application. In this respect, the work by Shao, [129], in which the compiled decision table, rather than the relation matrix, is modified, achieving savings in the amount of memory and numerical expense of the method, is useful.

4.3 Adaptive and learning indirect fuzzy control using model inversion.

In this method of fuzzy adaptive or learning control the controller carries out a process of identification on the system to be controlled. The identified model is then used to control the system. This approach is described by Harris with co-workers Moore, Brown and Fraser [92, 136, 137], Graham and Newell [138], and Togia and Wang [139, 140].

In the work presented in [136] the comparison is drawn between fuzzy and linear methods of adaptive control. The adaptive fuzzy control method described in Section 4.2 is said to be analogous to the model reference adaptive controller as found in [121] whilst the indirect method, see [136], described in this section is analogous to the self-tuning regulator from the same text. The term indirect is used because it makes use of a model which is separately identified and then used to adapt the controller. Having obtained a plant model, this model can then be causally inverted and used in the control of the plant. Two methods of fuzzy

model causal inversion are explained in the following sections which are based on [136] and adopt the notation used in that paper.

4.3.1 Relation matrix causality inversion in closed loop adaptive control.

Relation matrix causality inversion.

An n-dimensional relation matrix is used to relate n variables which range over n discrete universes of discourse. Each dimension of the relation matrix will have a length l_i, where l_i is the number of discrete elements in the i-th universe of discourse. Each location in the relation matrix, therefore, corresponds to a particular combination of variable values. The number stored at each location is the grade of membership of the information represented by the location in the n-dimensional fuzzy relation that the relation matrix represents.

When using the relation matrix, $n-1$ fuzzy sets representing the input information are combined with the relation matrix using the compositional rule of inference. This generates a fuzzy output set which can then be defuzzified to generate the output information.

The relation matrix is simply an n-dimensional matrix of numbers, relating n universes of discourse with no specific dimension having any particular significance. This means that, no matter what physical causal relationship exists in the data from the system being identified, the relation matrix can be used in any direction. If the relation matrix representing the input-output relation of some system is known, then it can be used to infer the value of the input which would cause a given output to occur, even if the data in the matrix were obtained by applying input values and measuring output values. Put more simply, given an output variable value, the relation matrix can generate the input variable value required to cause the output to occur. This is known as relation causality inversion

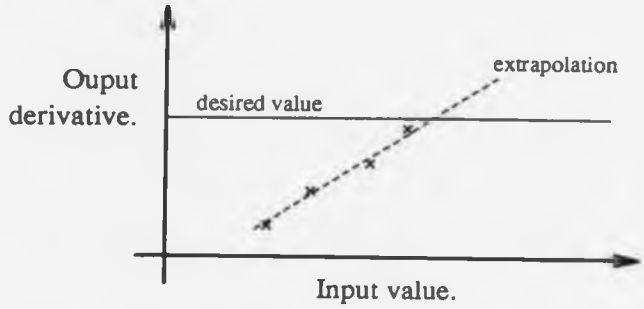

Figure 4.2: Relation causality inversion by prediction.

and takes place by simply reading the matrix in different directions.

When taking this view, the relation matrix can be thought of as an autoassociative memory [136, 92], or fuzzy associative memory [141].

If the relation matrix is stored in such a manner that makes causality inversion by this method difficult, an alternative method of inversion, called causality inversion by prediction, see [136, 138], can be used and is shown diagrammatically in Figure 4.2. In this case, the relation matrix relates an output derivative to an input value. Using other information, see below, the controller is able to determine a desired output derivative and requires the input value which causes the output to have the desired derivative. This is accomplished by the controller guessing a few input values and using interpolation, or extrapolation, from those guesses to get the value of input that causes the desired value of output to occur. Clearly, for this method of causality inversion, the plant model need not be a fuzzy relation matrix and any method of identification (be it fuzzy or not) can be used. However, given the nature of the method, a technique which is reasonably numerically inexpensive would be sensible.

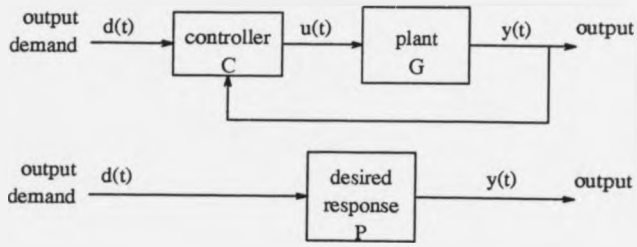

Figure 4.3: Closed loop controller and desired closed loop performance.

The use of causality inversion in closed loop control.

Figure 4.3 shows a plant being controlled by a fuzzy controller and, also, the desired closed loop response of the controlled system. The capital letters refer to fuzzy relations embodied as relation matrices. In this case, the description uses discrete time with sampling time T. The following analysis is, again, taken from [92, 136]. For the plant itself, the following fuzzy relation holds:

$$Y(t+T) = G \circ \{U(t) \times Y(t)\} \tag{4.3}$$

The $\circ$ operator denotes the compositional rule of inference.The fuzzy relation expressing the controller is as follows:

$$U(t) = C \circ \{D(t) \times Y(t)\} \tag{4.4}$$

The desired closed loop performance relation is expressed as:

$$Y(t+T) = P \circ \{D(t) \times Y(t)\} \tag{4.5}$$

The combined effect of the controller and the plant may be written:

$$Y(t+T) = G \circ \{Y(t) \times C \circ \{D(t) \times Y(t)\}\} \tag{4.6}$$

Combining Equations 4.5 and 4.6 gives:

$$P \circ \{D(t) \times Y(t)\} = G \circ \{Y(t) \times C \circ \{D(t) \times Y(t)\}\} \tag{4.7}$$

from which the controller relation matrix is obtained:

$$C \circ \{D(t) \times Y(t)\} = G^{-1} \circ \{Y(t) \times P \circ \{D(t) \times Y(t)\}\} \tag{4.8}$$

It can therefore be seen that the controller relation matrix is obtained from the inversion of the identified system model and the closed loop performance description.

The advantages of this method are, firstly, the controller contains a description of the system and the required closed loop performance in forms which are easily understood. Secondly, the controller formed by this method uses absolute values of input and output and will, therefore, behave differently in regions of the space where the system behaves differently. Controllers that use error signals have no indication of the region in the control space in which they are operating and so have to adapt, losing information as the system moves from one region to another. The crucial advantage of using absolute signals is that the controller becomes, to a certain extent, a learning controller. Obviously, the controller cannot "learn" about things for which it has no information, for example, variables which are not seen, formally, as plant inputs but which have an effect on the plant operation. The advantage of a learning approach over an adaptive approach will vary depending upon the situation, but, the comparison of the SOFLC and the relation causality inversion controller serves to illustrate the difference.

A drawback to the causality inversion approach is the computational expense of performing on-line identification of the relation matrix, inverting it and then combining it with the performance index for every time step in real-time. Also, the relation matrix has to store information for the entire control space but the SOFLC only stores information for the region in which the controller is currently operating. In order to obtain the same control performance in the region of a particular set-point, the relation matrix will have to store much more information than the SOFLC for the same degree of non-linearity.

4.3.2 Related Work.

Similarities to the approach outlined above can be found in other reports. In [106], Batur and Kasparian use an adaptive fuzzy technique that uses a model of the process which is not actually inverted but which predicts the next error value of the process using the current controller and the process model. If the controller has confidence in the model predictions, but finds that the predicted error is not small enough, then the controller is modified. This method departs from those outlined above, since the process model is not then used to determine the correct controller modification which is carried out by means of varying the membership values of the output fuzzy sets. The method of process identification is also different from that used in [136], using coefficients to relate the current output to past values of output and input, the coefficients being adjusted by means of a least square fit.

In [138], Graham and Newell describe adaptive control of a first order process which uses two methods of fuzzy process identification. The first method is as defined by Pedrycz [142], and uses the relation matrix in almost exactly the same way as in [136], with the addition of a confidence in relation matrix modification that can be used to determine whether modifications should take place. The second method is a rule-based method, developed by Tong, in which the rules are either introduced, if no similar rule exists, or modified, if a similar rule is present. The adaptive control method that is used simultaneously identifies and controls the process. The process is controlled by making use of the model in a similar way to that used in [136]. Successive input values are tried on the model and the one that results in the smallest subsequent process output error is chosen and applied.

Jager, Verbruggen, Bruijin and Krijgsman report, in [143], the use of an expert system that controls a process using feedback on the error between the reference signal, and the output signal and the change in error. A line in the phase plane formed by the error, and change in error, is drawn between the current plant

position and the origin (the origin being the desired next position of the process). A line is also drawn through the last position of the process and the current position of the process representing a prediction of the process behaviour. Clearly, the process should move along the line towards the origin and two lines are drawn either side of the line towards the origin indicating the fastest and slowest tolerable movements toward the origin. Control actions are then considered which move the predicted process behaviour within the tolerable band.

4.3.3 Fuzzy system identification.

The method of fuzzy control described above requires a relation matrix model of the system if relation matrix inversion is to be performed by an associative memory method. A means has, therefore, to be found of obtaining this relation matrix for the system being controlled. In [136], the fuzzy union of the current system relation matrix and fuzzy sets centred on the values of the current data is formed at each instant to give a new system relation matrix. This allows the matrix to acquire new information provided that this information lies "outside" the information learnt so far.

Clearly, if the current data is already contained in the relation matrix, then the relation matrix will not be modified by the formation of the union. However, if the data in the relation matrix is no longer correct because the system being identified is changing, then taking the union will not remove the old, incorrect, information. In order to do this, at each stage, the relation matrix is multiplied by a forgetting factor causing all the old information to be reduced in importance. A disadvantage of this technique for system identification is that, if the system does not enter a certain region of operation very frequently, then that area of the relation matrix can become entirely forgotten. There are methods that are used to prevent this from occurring, details of which are given by Brown, Fraser, Harris

and Moore in [137].

Other work on relation matrix identification and control has been done by Pedrycz [142], and by Czogała and Pedrycz [144, 145], in which the concepts of quality and optimality were considered. Xu and Lu, [146], developed an iterative method in which referential sets are used to produce an optimal model.

Other methods of fuzzy identification that do not involve manipulation of the relation matrix have also been considered. Takagi and Sugeno [147], demonstrate a method of fuzzy identification in which the rules infer a function that links the output variable to the input variable, rather than inferring the output variable itself. This work is significant in that it not only defines a further identification method, but also a new form for fuzzy rules. Kandel [148], develops fuzzy differential equations and uses them to describe fuzzy systems. In [149], Shen, Ding and Mukaidono use a set of fuzzy rules to describe the system being identified. The method of fuzzy identification modifies either the input sets, the output sets or the fuzzy rules themselves, dependant upon the error in the model prediction. These other techniques and methods that do not involve the relation matrix could be used for inversion by prediction, along with many other general identification methods.

4.3.4 Application of the inverse causality relation matrix method to the management of hybrid vehicle powertrains.

As stated above, this method of control is numerically expensive which will tend to make its use unattractive in automotive applications. It is also a closed loop setpoint control technique which, again, makes it difficult to apply, in hybrid vehicle powertrain control, for the reasons given in Section 4.2.3. However, an extremely attractive aspect of this method is the modularity that is inherent in its approach.

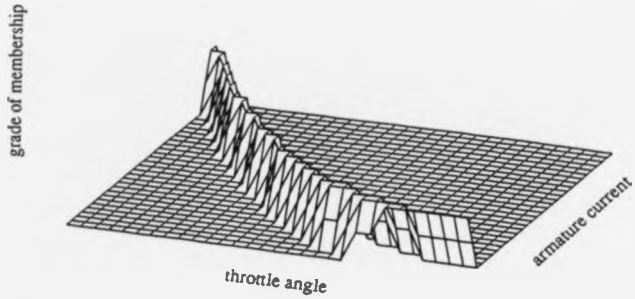

Figure 4.4: Fuzzy relation matrix showing those values of armature current and throttle angle that develop the same combined output torque.

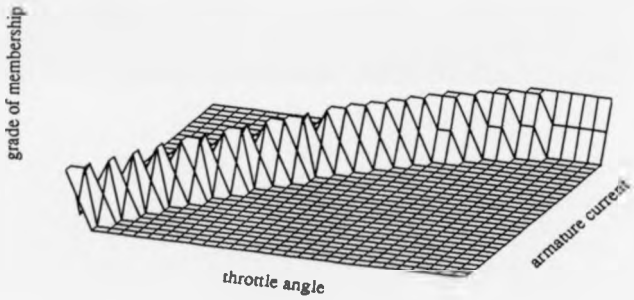

Figure 4.5: Fuzzy relation matrix showing those values of armature current and throttle angle that meet some combined energy use objective.

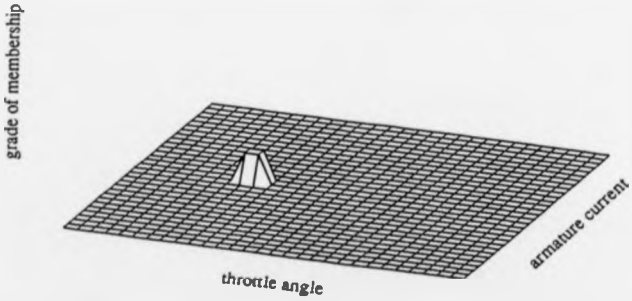

Figure 4.6: Intersection of torque and energy use objectives.

Figure 4.4 shows two dimensions (or one layer) of a three dimensional fuzzy relation matrix. The third dimension is transmission input torque (combined engine and motor output torque) which can be imagined to be constant over the layer shown. The complete three dimensional relation matrix would be composed of layers similar to Figure 4.4 vertically above one another, each corresponding to a different transmission input torque, giving a vertical axis of transmission input torque. Clearly, in this case, grade of membership would be stored in each location and could not be visually interpreted, as this would require a fourth dimension. The "hill" in the grade of membership shows those values of armature current and throttle angle that produce the value of torque that corresponds to this layer of the relation matrix.

There is at least one other dimension to the fuzzy relation describing the hybrid powertrain, since the relation matrix described above represents the powertrain at a particular engine speed. In order to completely represent the powertrain, a further dimension along which engine speed varies, would be required. This would result, effectively, in a vector of relation matrices such as the one described in the previous paragraph, each one obtained at a different engine speed.

Figure 4.5 shows a layer from a similar relation matrix to that shown in Figure 4.4 but, in this case, the vertical axis, that is not shown, is some control objective that can be varied. This control objective relates the relative amounts of throttle and armature current, and, therefore, the relative amounts of fuel and electrical energy that are used. For example, at low battery states of charge the controller might chose to apply regenerative currents. This would manifest itself by the controller picking a value for the objective axis that caused a bias, in the shape of the grade of membership ridge, towards negative armature currents and large throttle angles. The appropriate control action would then be represented by the fuzzy intersection of the relation matrix for the correct torque value and the

objective value, with the pedal value now demanding transmission input torque. This intersection would be as shown in Figure 4.6.

This approach is similar to that taken by Bellman and Zadeh in [150], in which the constraints and goals of an optimisation procedure are described by fuzzy sets and the solution that best meets the optimal requirements is given by the intersection of the fuzzy goals and constraints.

If the powertrain elements are assumed to be unchanging, an invariant 4-dimensional relation matrix, obtained by measurements, can be used to predict transmission input torque. Alternatively, if some means can be found of measuring or observing transmission input torque, then the relation matrix predicting this quantity can be identified on-line.

Clearly, the idea outlined above has a number of difficulties to overcome, including the numerical expense of the method, the amount of data to be stored and the dimensionality of the problem. If each of the two 4-dimensional relation matrices is defined on a $10\times10\times10\times10$ grid, and each number can be stored using one 8-bit byte, just under 20kbytes will be required to store the relation matrices. This is about one order of magnitude higher than the amount of memory which is generally taken up by the data for current automotive control applications. If the number of indices goes up to $15\times15\times15\times15$, the amount of memory goes up to about 50k. Also, the program would need, in automotive terms, a fair amount of "scratchpad", in order to perform implication using the current powertrain speed value, current controller objective and required transmission input torque, and then defuzzify the result onto two universes of discourse. A further drawback to this procedure is that most of the memory that is apparently needed by the associative memory would be unused, since as illustrated in Figures 4.4, 4.5 and 4.6, relation matrices are very sparse.

The practical side of the problem would be considerably worsened by attempt-

ing to identify the engine and motor on-line as this would require random access memory, having a far higher cost than either a masked permanent memory or a one time program memory. Also, computationally, the control problem would become very much harder.

In the light of the problems outlined above, this method would be unlikely to be adopted for practical hybrid powertrain control. However, the development of genuine multidimensional associative memories might make the approach feasible, if the devices were produced at reasonable cost and were capable of meeting an automotive temperature specification. The development of such devices would be a considerable breakthrough, since the above manipulation of relation matrices indicates the general applicability of the associative memory concept.

4.4 Neurofuzzy control.

The actions of human beings in performing tasks such as control and decision making were the inspiration for the development of fuzzy logic. One of the early objectives of fuzzy control was to replicate the conscious, reasoned actions of humans when performing control tasks such as driving and chemical plant control. The linguistic form of fuzzy rules makes this obvious. Over a similar time period, the parallel science of neural networks has been developed, in which the original objective of the network was to model the instinctive, subconscious acts of humans such as recognition of faces and limb movement.

In fuzzy logic, the central concept of the fuzzy set and the methods of approximate reasoning formed the framework of the decision making environment, the link with human beings being purely conceptual. In the case of neural networks, the form of the network was inspired by observations of the structure of the human brain, making the link with biology more direct. However, ultimately, both fuzzy logic and neural networks simply relate an input space to an output

space in a non-linear way (recursive neural networks could be thought of as an exception to this, since the output side of the network is fed back into the input side of the network, but the mapping between the input and the output is still simply a non-linear mapping).

Under some circumstances, certain classes of fuzzy logic rule bases and neural networks can be shown to be equivalent, see [151], and the close relationship between the two forms of network is exploited by Nie and Linkens, in [152], and Chien and Teng, in [153]. Two more examples of the application of neural methods in fuzzy networks are considered in some detail below. The chief attraction of combining neural networks and fuzzy logic is that a network whose topology combines aspects of neural networks and fuzzy logic can be initialised using fuzzy rules, to have some basic capability, and then subsequently trained using training techniques developed for neural networks.

Comments on the applicability of these controller modification methods will be made at the end of this section.

4.4.1 Neurofuzzy control using the CMAC.

The Cerebellar Model Articulation Controller (CMAC) was first suggested by Albus [154, 155], in 1975, as a neural network topology that was suitable for controlling the joints in a robotic manipulator. One of its most attractive properties is that the network globalises locally, which has two advantages. Firstly, similar inputs produce similar outputs, and secondly, training the network is relatively straightforward. For any particular input, only part of the network is used because of the localisation, therefore, only those output weights being used in that particular region of the input space need to be adjusted. In other neural topologies such as the multi-layer perceptron, the parameters that are used in computing the output value for any input, are not localised. Therefore, adjusting the output

weights for a better output match for a particular input will affect the output match for other inputs and, as a consequence, training is rather slow. The radial basis function neural network is very similar to the CMAC and is often used in a similar way in neurofuzzy control and identification, see [156].

Another attractive feature of the CMAC is the compact way in which the weights are stored. By the use of a generalisation factor ρ, and a pseudo-random hashing technique, the computational cost of an n-dimensional network does not increase exponentially with the dimension, but only linearly. Each input in a discrete CMAC is mapped, by means of the CMAC mapping algorithm, to ρ cells in a virtual association cell space. The nearest neighbour inputs will share $\rho - 1$ cells with the first input and inputs more than a distance $\rho - 1$ apart will share no cells, leading to local generalisation. Originally, the CMAC association cells had only a binary output, giving a piecewise constant network output, but higher order basis functions, see [92], or kernal functions, see [157], can be used to give a smoother output. When higher order functions are used, the CMAC becomes equivalent to a fuzzy network.

Two methods, the Least Mean Square method and the Normalised Least Mean Square method, are generally used to train a CMAC network, on-line, using instantaneous information. Using sets of data, allows batch training to be used and further gradient descent training methods are available, see [111], for details.

Reported applications of the use of the CMAC in control include [158, 159], with Majors, Stori and Cho describing the use of a CMAC in automotive engine control in [160]. These three papers were all presented at a recent conference, showing the current interest in combining the CMAC and fuzzy control.

4.4.2 Learning Vector Quantisation.

Learning vector quantisation (LVQ), was first introduced by Kohonen [161, 162], and is a technique used in pattern recognition and classification feedforward neural networks whose neurons have sigmoidal activation functions. Each first layer neuron has an n-dimensional vector of synaptic links with the n inputs. At each stage of learning, the synaptic vector which has the closest Euclidean distance with the input vector is adjusted, to more closely represent the input vector, by use of the formula:

$$\mathbf{m}_j(t+1) = \mathbf{m}_j(t) + c_t(\mathbf{x}(t) - \mathbf{m}_j(t)) \tag{4.9}$$

where $\mathbf{m}_j(t)$ is the vector of synapse values for the first layer neuron j, at time t, $\mathbf{x}(t)$ is the input at time t and c_t is a monotonically decreasing scalar gain. This is known as unsupervised competitive learning and will tend to make an individual neuron fire if the input vector is close to the pattern represented by that neuron's synapse vector. An extension to this is to apply the learning procedure to the N nearest synaptic vector neighbours of the input vector.

The network is trained using samples representative of the different fuzzy classifications that are to be distinguished. After a period of training, the synapse vectors will cluster around the distinct patterns that are present in the training data. This first layer then corresponds to the "IF" side of fuzzy rules, since the first layer neurons will fire to differing degrees dependant upon the extent to which the input matches the pattern classification that they represent. Since the neuron activation function is sigmoidal, it is important to note that the output from each neuron is only the extent to which a particular fuzzy set is inferred. The probability density function of the class into which inputs are clustered, effectively, forms the fuzzy set in this approach, rather than the activation function of the neuron itself, see [141]. This is fundamentally different from the way in which the radial

basis function and CMAC neural networks are used with fuzzy logic.

The subsequent layers of the network effectively perform the fuzzy process of inferring the output and, therefore, represent the "THEN" side of fuzzy rules. Accordingly, unsupervised competitive learning will only affect the "IF" side of fuzzy rules.

The process of inferring the output is often achieved by the use of various forms of associative memories, see [141]. Examples of this method include work by Yamaguchi, Takagi and Mita [163], by Yamaguchi, Tanabe, Kuriyama and Mita in [164], and by Pacini and Kosko [165].

4.4.3 Application of the use of neurofuzzy control to the management of a hybrid vehicle powertrain.

Neurofuzzy control cannot really be classified as either an adaptive method or a learning method, since this will depend on the nature of the input information and the subsequent use to which it is put. It is generally fair to say that methods which deal with error information are less likely to involve learning behaviour than methods which use absolute signals.

The principle problem of using neurofuzzy control in the management of a hybrid vehicle powertrain is that of training. Neural methods rely either upon the existence of a set of input-output training data, that is representative of the data that the network will encounter in use, or on-line network training by comparing the network output with a desired output.

In the hybrid vehicle powertrain problem, correct network outputs are not known and it is desirable that the method of adaptation, or learning, determines the correct outputs. Network modification will have to be based on observations of the performance of the network in controlling the hybrid vehicle. It was noted earlier, that an essential feature of the powertrain control method was the maintenance of a constant driveshaft torque, which would ensure that the vehi-

cle performed in a consistent manner. If some measure of "correctness" could be found, allowing on-line, unsupervised learning to be carried out, there is no guarantee that, as the controller evolved, it would maintain vehicle driveability. Since it is hoped that the controller might be able to adapt to individual vehicles, it must allow them to be continuously driven.

For these two reasons, it is unlikely that neurofuzzy control methods will prove useful in the hybrid vehicle powertrain control problem as it has been formulated in this work.

4.5 Fuzzy modification of linear control methods.

The work discussed in this section uses some means other than fuzzy sets and inference to generate the controller output value. Use is made of fuzzy sets in the modification of the lower level controller in order to improve its performance. The interesting observation, here, is the hierarchy that is applied by allowing the fuzzy controller to modify a lower level controller. If it is accepted that the adaptive, or learning level, would generally involve more "intelligent" actions than the lower level controller, then it is perhaps appropriate that a fuzzy controller, designed using expert knowledge, should be at the higher level.

4.5.1 Fuzzy sliding mode control.

In sliding mode control, as defined by Utkin [166], the state space of the system being controlled is partitioned into two spaces by a sliding surface which passes through the origin, which is assumed to be the point to which the system should be controlled. Different control actions are then defined in each of the two state space regions, each having the effect of driving the system towards the sliding surface. The phase trajectories of the controlled system in the region of the sliding surface

must be such that the system is moving towards the origin as it approaches the sliding surface. The combined action of both controllers will then be to drive the system towards the sliding surface and then along it, if, as the system states pass through the sliding surface, the control action is switched to the control demanded by the other part of the state space. Whilst this approach is not limited to the control of linear systems, linear systems are sometimes controlled in this way.

As the system slides along the sliding surface, the control action is continuously switched to keep the system on the surface. This results in rapidly changing control action that is undesirable. In order to prevent this rapid switching of the control action, a boundary layer may be defined along the switching surface, in which the actions are smoothly blended rather than abruptly switched. Palm [167], describes an implementation of the boundary layer and goes on to analyse the actions of the blending controller, as a subcontroller keeping the system on the surface, using a method originally proposed by Slotine and Sastry [168]. Palm is then able to show that the controller that brings the system towards the sliding surface, and then keeps it within the boundary layer, can be implemented using fuzzy logic. The use of fuzzy logic increases the speed with which the system approaches the sliding surface, and its smoothness within the boundary layer. These benefits are analogous to the improvements that fuzzy logic often brings when applied to conventional setpoint control. Vepa and Nowe [169, 170], extend this idea to include identifying a model of the system in the region of the sliding surface, using a neural network. This network is subsequently inverted and used to form a fuzzy controller whose objective is to keep the system on the sliding surface. Nam, Lee and Yoo [171], apply fuzzy sliding mode control to the regulation of air-fuel ratio in an automotive engine. The air-fuel ratio measurements have an unavoidable time delay, and predictive fuzzy control is used to keep the system on the sliding surface, to minimise the effects of the time delay.

4.5.2 Fuzzy modification of linear controller data.

In this section, controllers that are essentially linear, but whose inputs are modified by a higher level fuzzy controller, are considered. A particularly hierarchical applied approach to this was given by Saridis and Stephanou in [172], whilst Saridis and Lee [173] provide a more theoretical and generic viewpoint.

In [172], which considers the problem of the control of a prosthetic arm, various methods, including grammars, parsing algorithms and a fuzzy learning automaton, see [174], are used to learn the correct means of mapping the signals from the nerves of the amputee to the subcontroller inputs. A fuzzy reinforcement learning method is used to reward those mappings which result in improved control. The role of the learning fuzzy automaton in this work is the interpretation of imprecise control objectives to generate more precise commands for the lower level controllers. This approach is very numerically expensive and specific to the particular control problem considered.

Kim, Park and Lee [175], considered the control of a system using an actuator with a dead zone. Originally, PD control was used, leaving a steady state error. To remove this error, a PID controller was used, however, this controller had poor transient response. A fuzzy precompensator was then used, with the PD controller, to allow the steady state error to be reduced and to achieve good transient response.

4.5.3 Combined fuzzy and linear controllers.

Yeung and Sum [176], address the problem of controlling the tension in yarn as it is wound onto a bobbin. It was found that, in the area of the setpoint, adaptive linear controllers could not adapt quickly enough to avoid a process control failure resulting in the yarn breaking. For large process disturbances, a fuzzy algorithm is used to bring the process quickly towards the setpoint. As was stated in [133], the

control action when far away from the setpoint is not critical, since the controller should apply a large input to move the output quickly towards the setpoint, with more accurate control being required to control the output around the set point. In the region of the setpoint, a self-tuning regulator can take over the control of the process and adapt to system changes with sufficient speed to prevent a failure of process control.

The attractive aspect of this work is that linear methods are used, where their operation is adequate and are supplemented, by fuzzy control actions, where fuzzy logic has something to offer.

4.5.4 Fuzzy tuning of a PID controller.

Ollero and Garcia-Cerezo [177] describe two methods in which the gains of a PID controller are adjusted by an external fuzzy controller. In the first method, the controller varies the gains according to the current position of the controller in the input space. In the second method, the controller gains are adjusted as an external variable, which is known to affect the plant operating characteristics, changes. It can be argued that this is like gain scheduling but it is likely to be more effective than straightforward gain scheduling since it makes use of the intelligence of the designer.

In [178] Litt also applies an expert system to the tuning of a PID controller. The expert system applies modifications to the controller gains based upon the Ziegler-Nichols time domain adjustment procedure. Whilst the method used here is an expert system, and the paper refers to around 10 other applications of expert systems in similar roles, it also states that this is an excellent application for fuzzy logic.

4.5.5 Application of mixed fuzzy and linear controllers to the hybrid powertrain control problem.

The major impediment to the use of the above methods in the control of a hybrid powertrain is that they are essentially feedback control methods that operate on the error between an output and a reference signal and, as previously stated, the hybrid powertrain controller does not easily fit into this framework. Since the methods described above tend to involve the use of error signals, the fuzzy logic tends to be used in an adaptive, rather than a learning manner. A possible exception to this might be [172], in which some learning could be said to take place. Some of the methods used, such as [90], could really be called tuning since they automate the human process of controller design.

4.6 Miscellaneous adaptive fuzzy methods.

This section will cover several reported pieces of work that were not discussed in earlier sections.

4.6.1 Heuristic modification of fuzzy controllers.

This section will discuss various, seemingly *ad-hoc*, methods of adapting fuzzy controllers in heuristic ways which work for the controllers and processes concerned.

Gain coefficient tuning.

One method of gain coefficient tuning is described by Mallampati and Shenoi, in [179]. A simulated annealing technique is used which is an enhancement to the Monte Carlo method and involves the variation of parameters that influence some objective function. Simulated annealing involves always accepting modifications to the parameters that improve the objective function and also accepting,

with a decreasing probability, some parameter changes that worsen the objective function. This reduces the possibility of being trapped in local minima.

Another approach to fuzzy gain tuning is given by Bare, Mullholland and Sofer, in [180]. In this method, a fuzzy two input two output control problem is decomposed into two SISO problems by, effectively, decoupling them in time. Adjustments are made to one loop and the results are observed, and the other loop is adjusted less frequently. The method of adjustment is heuristic but phrased, in a crisp numeric manner, using the ratio between successive errors to determine whether the process and controller combination required "speeding up" or "slowing down", allowing the designer to use knowledge of the process operation in modifying the controller. The methods chosen are derived from human controller tuning which has traditionally been used.

Heuristic adaption followed by performance evaluation.

Bartolini and others [107], used an automated method of heuristically modifying the output sets in a fuzzy controller by means of five adaptation actions. The action to be taken was determined by the evaluation of six performance measures, which were functions of the error in the control of the level of liquid metal in a tundish which was part of a casting process. Two of the actions were associated with either the level being too high, or too low. Two more actions were associated with improving the controller sensitivity, and the last was to leave the controller unmodified.

Reference [181] describes how Murayama, and others, carried out an optimal fuel calibration exercise of a marine diesel engine in an attempt to minmise its use of fuel. This is complicated by an extremely noisy fuel flow rate signal. Previously, the optimisation had been carried out by a stochastic method known as evolutional operation. The method used here is essentially a gradient descent technique that

involves use of the designers common sense. Using this method, the engine was calibrated in one thirtieth of the time that had previously been taken using non-fuzzy methods.

4.6.2 Genetic Algorithms.

Genetic algorithms (GAs) emulate the process of genetic evolution in the natural world. Genetic algorithms are optimisation procedures that have the advantage that there is no need to directly relate the cost function of the optimisation problem to the parameters of the system that are varied in the optimisation.

In formulating an optimisation problem for the use of a GA, the parameters that are varied must be represented in a binary manner. The various binary parameter values are then concatenated to form a single long binary string, known as a chromosome. The task to be optimised is then carried out for all individuals in a "population", evaluating the cost function for each individual. Individuals performing well are used to produce the next population by a process of combination, known as crossover, which causes some of the offspring to inherit good properties from the parents and others to inherit bad properties. To add a degree of randomness to the process, that effectively stops it from being trapped in local minima, bits are inverted in the strings with a low probability. This is termed "mutation". Having established a new population of individuals, the process is repeated until little variation is seen in the population and the cost function value has reached the global minimum.

This procedure, and many enhancements to it, are explained by Goldberg in [182]. To evaluate a given population, the task must be performed for each individual, and, since GAs generally take thousands of iterations to converge, computer simulations are used to gain repeatability and speed. In the case of the hybrid vehicle, the control adaptation strategy must be made to work on an

actual vehicle in real time whilst driving. Whilst recent developments such as micro-GAs, see [183], have improved performance, GAs are not likely to be useful for hybrid vehicle control because of their slow rates of convergence. Also, the random nature of the algorithm might generate individuals that made the vehicle undriveable. Examples of fuzzy controller modification using GAs include work by Karr [183], and Foslien and Samad [184], whilst Tascillo and others combined the use of neural nets, fuzzy logic and GAs in [185].

4.6.3 Fuzzy cell-to-cell mapping.

The cell state space concept was developed by Hsu [186], to control non-linear systems in discrete time, using cells that are defined over discrete state spaces. The effect of applying control actions will be to move the system from one cell to another. Since the number of control actions is limited, only certain cells can be reached from any individual cell and a set of cell-to-cell mappings can then be defined. A cost function is then used to determine, for a given control objective, which mapping, and therefore which control action, is most appropriate at any instant in time.

Smith and Corner [187], extended Hsu's ideas by using state space cells, similar to those used by Chen and Tsao [188], to define an optimal control action. A fuzzy controller then uses these optimal control actions to train a set of fuzzy rules, of the Takagi and Sugeno form, using a method derived from Widrow-Hoff least square learning, see [189]. This gradient descent based method is also used by Nomura, Hayashi and Wakami [190], in the adaptation of a controller for a mobile industrial robot that is able to avoid obstacles.

Vepa and Nowe also combine the use of fuzzy logic and state space cells in [191]. Their technique is based upon a cell state space partition on which a reinforcement learning method, called Q-learning, is used, as originally defined

in [192]. Use of fuzzy logic is made by replacing the cells with fuzzy sets and defining nine output fuzzy sets for each of three input fuzzy sets. Each of the output fuzzy sets is modified by using a modification of the reinforcement method that takes into account the "safety" of the system control actions. Once again, these techniques are useful for feedback control but would be difficult to apply in hybrid vehicle powertrain control.

4.6.4 Adaptive fuzzy representation of information.

Two very theoretical papers by Lakov [193], and Blishun [194], investigate the representation of controller input information by the use of fuzzy sets. The adaptive aspect of the work reported is that the sets used to represent the information are modified, the controllers which subsequently make use of the information are static.

4.6.5 Other methods.

This section discusses other reports of adaptive fuzzy control which do not fall into any of the previous groups.

Handelman and Stengle [195], describe the adaptive control of an aircraft by means of an expert system that uses a series of blank template rules that it "recruits". The information in these rules is represented by fuzzy sets.

In [196], Borges de Silva and Oliver consider the control of a power converter. At the lowest level, four rules of the type proposed in [147] are used, which are adapted by two higher levels of fuzzy automata. Although the work discussed contains considerable detail on the converter itself, the fuzzy part is generally applicable.

Adaptive control is not, itself, the subject of the work reported by Andersen and Nielsen [197], which considers the representation of fuzzy rules. Their repre-

sentation of the rules is very similar to that used in Chapter 3 of this thesis and is claimed to be very efficient in terms of memory. This representation is useful for adaptive control where the adaptive action may cause a large increase in the number of rules.

Peng Xian-Tu constructs a fuzzy controller, in [198], by means of a series of input and output sets that are related to one another by mappings. In the MIMO case considered in this paper, the mappings take the form of matrices of scale factors which are then adjusted to meet some conditions for optimality by the use of the simplex algorithm.

4.7 Discussion of adaptive and learning fuzzy control survey.

The work described in the early sections of this chapter is not immediately applicable, because the control tasks considered were feedback reference tracking control problems and the problem addressed in this thesis is clearly not in this form. In some respects, this concentration on conventional control problems has been unfortunate for fuzzy control, since it has merely added a further competing method in an area in which there are already well established techniques. One of the principle attractions of fuzzy control must surely be its breadth of applicability. It should be noted that some of the most commercially successful applications of fuzzy logic are not in feedback control, but are in areas such as automotive transmission control, airport management and the operation of consumer products such as cameras and washing machines.

Despite the preoccupation of fuzzy control with feedback, reference following control, some of the work surveyed in this chapter contains ideas which can be made use of in the adaptive control required in this work. These ideas will be briefly discussed below.

The work by Shao, [129], in which a fuzzy controller is "compiled" to form a decision table and the adaptation process acts directly on the decision table makes a considerable contribution. When fuzzy logic is used in the way described in Section 3.4.3, the decision table is the same matrix as the rule output set offsets. Also, it should be remembered that automotive control engineers frequently make use of decision tables in the form of maps, which they are accustomed to heuristically modifying during the process of controller development. An important consequence of this is that an adaptive method that simply modifies maps will be readily accepted by automotive development engineers.

The greatest contribution that this chapter makes comes from Section 4.6, in which numerous different strategies are used to adaptively control various processes. Section 4.6 contains many different strategies, because each of the methods has been specifically designed to work on a particular plant. Examples of this are the work of Bare, Mullholland and Sofer [180], Bartolini, et al [107], and the work of Murayama, et al [181], in which expert knowledge in diesel engine calibration is directly implemented in rules which generate an optimal engine calibration. These examples are encouraging because they indicate that fuzzy logic provides a useful framework, in which many methods of adaptive control may be applied.

The work of Bellman and Zadeh [150], might, at first, seem to make a significant contribution to the work undertaken here, and the concept of the optimal decision being the union of the goals and constraints is obviously a valuable idea. However, [150] considers the task of taking a process from a particular starting point, to a particular finishing point. Since there are no specific start and end points for a hybrid powertrain control task (unless one considers the optimal use of energy over a specific journey) this aspect of the work is not particularly helpful.

4.8 Comments on adaptive and learning methods of fuzzy control.

The creation of the fuzzy set was motivated by the observation that vagueness and imprecision are useful concepts in complex control or decision making tasks. In certain respects, this has been borne out by the recent proliferation of products in which fuzzy logic is used. The commercial success of these products can be attributed to fuzzy logic, however, to what extent fuzzy logic gives these products a performance, rather than a marketing advantage, is not always clear.

A vast number of reported research applications of fuzzy control have drawn the conclusion that the original controller design had to be modified in order to perform satisfactorily. This is hardly surprising, when it is borne in mind that humans very rarely get anything right first time. The task of modifying a fuzzy controller has spawned the various adaptive and learning techniques that have been surveyed in this chapter. It can be seen that some of the methods discussed above actively embody human insight, experience and intuition in the controller modification process, [107, 167, 176, 181] are examples taken from Sections 4.5.4 and 4.6.

In contrast, in the work considered in Section 4.3, in which the fuzzy relation matrix is identified and then inverted as an autoassociative memory, the fuzzy sets are merely providing a general non-linear functional representation. Other than, perhaps, being able to short cut the identification procedure a little, by initialising relation matrices, no explicit use is made of the ability of fuzzy sets to represent the vagueness of human thinking.

The work discussed in Section 4.2 is somewhere between these two approaches, in that the performance index relation matrix is used to specify the closed loop system performance. However, the fuzzy operations, at a lower level, that involve modifying the relation matrix do so, again, without explicitly using experience

and insight.

The appropriate conclusion to be drawn from this short discussion is that the fuzzy set is no longer simply a means of representing vague, ill-defined information (although its utility in this role is, clearly, still of great value), but has also become a general method of non-linearly relating variables. This continuum of roles of fuzzy sets and fuzzy logic will ensure that fuzzy techniques will be used at various levels for the forseeable future in control applications.

Chapter 5

An adaptive fuzzy hybrid vehicle powertrain controller.

5.1 Introduction.

In Chapter 3 it was seen that the hand-tuned fuzzy controllers used did not meet the requirements set out in Section 2.2. It was found that hand-tuning the controllers was a laborious task that would possibly only be applicable to the specific task considered, the simulation of a hybrid vehicle over a particular drive cycle. Some method of automatically generating the correct control values for the powertrain controller was shown to be required. This method of generating the controller could be considered on various levels.

Firstly, an automated method of changing the controller values would be useful during the development of a hybrid vehicle which was to be sold with a fixed controller. This would allow the vehicle to be developed quickly and should be transferable from one engineering project to another.

Secondly, the controller might be made to change to meet the requirements of individual users, such that different patterns of usage resulted in the controller learning control actions that were most appropriate for any particular user.

Thirdly, the method could be extended to modifying the controller in use, allowing the characteristics of the individual vehicle to influence the control actions

that are taken. This would account for variations between vehicles and the effects of ageing.

In Chapter 4 adaptive and learning fuzzy control methods were surveyed and it was seen that none of the methods discussed were suitable for the modification of a hybrid vehicle powertrain controller. Accordingly, this chapter will describe a method of controller modification, specific to the hybrid vehicle control problem, that attempts to improve the performance of the vehicle with regard to the requirements set out in Section 2.2. The following sections will describe the structure of an adaptive hybrid vehicle powertrain controller and will then present early results arising from its use.

The hybrid vehicle powertrain controller can either use the motor to assist the engine or oppose the engine. In Chapter 3 it was seen that the potential for improving the efficiency of the powertrain operation was greater when the motor was operated against the engine. Also, the vehicle driveability is more difficult to maintain when the motor operates against the engine. Accordingly the adaptive hybrid powertrain controller was developed for the mode of operation in which the motor works against the engine.

5.2 The development of an adaptive hybrid powertrain control concept.

It will be recalled from Section 3.5 that the controller can be described in terms of three vectors. Figure 3.8 shows the form of the input sets used on the pedal value universe of discourse. Each input fuzzy set overlaps with its nearest neighbour as far as its nearest neighbour's maximum grade of membership. This allows the input fuzzy sets to be defined as a vector whose values are the locations of the maximum grades of membership of each fuzzy set. The corresponding output fuzzy sets on the throttle angle universe of discourse are all fuzzy singletons. The

vector of pedal value fuzzy set maximum grades of membership, Θ, the vector of throttle angle fuzzy singletons, Φ and the vector of demanded armature current fuzzy singletons, A, are defined as follows:

$$\begin{aligned} \Theta &= [\ 0 \quad 4 \quad 12 \quad 24 \quad 40 \quad 90\ ] \\ \Phi &= [\ \phi_1 \quad \phi_2 \quad \phi_3 \quad \phi_4 \quad \phi_5 \quad \phi_6\ ] \\ A &= [\ a_1 \quad a_2 \quad a_3 \quad a_4 \quad a_5 \quad a_6\ ] \end{aligned}$$

In Section 3.4.3 it was seen that the inference method and the defuzzification method used in this work reduce to linear interpolation, within the vectors, Θ, Φ and A, when the input sets are as shown in Figure 3.8.

In Section 2.2 it was seen that the controller had essentially three requirements. It had to use the energy resources of the vehicle in the correct ratio and as efficiently as possible. It also had to maintain a suitable relationship between the pedal value and the driveshaft torque. The adaptive controller will attempt to meet the first two requirements by observing the performance of the vehicle. The powertrain controller has no knowledge of the relationship between its control actions, the fuel flow rate and the battery current. There is the implicit assumption that battery current increases monotonically as demanded armature current increases and that the fuel flow rate increases monotonically as the throttle angle increases. The third requirement is met by a constraint in the rule modification procedure, in which the vectors Φ and A are modified, whilst Θ remains constant.

References [199, 200] are examples of the modification of fuelling control maps in engine controllers. Since use of these maps involves interpolation, their automatic modification is, at first sight, a similar problem to the modification of the hybrid vehicle powertrain controller. However, as a result of a separate feedback control action, the fuelling controller is able to determine the "correct" values for the fuelling maps. In the case of the hybrid powertrain controller, the adaptation process does not know the correct value to adapt to but has to find it out.

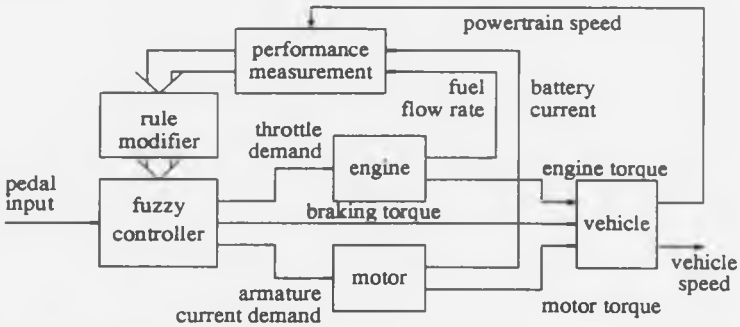

Figure 5.1: Schematic of an adaptive fuzzy hybrid powertrain controller.

5.2.1 An adaptive fuzzy hybrid vehicle powertrain controller structure.

Figure 5.1 shows the manner in which an adaptive hybrid powertrain fuzzy controller would work.

In order to observe the performance of the vehicle, two performance measures are used which operate on the values of the fuel flow rate and the battery current. The method by which these performance measures are evaluated is explained in the next section. The controller essentially functions by estimating the unknown performance measures over a period of time called the adaptation interval, T_a, and then making modifications to the controller based upon the performance measure values obtained over this interval.

Performance measurement.

As stated above, the modification procedure will address the problem of using the energy resources of the vehicle in the correct proportions and in an efficient manner. Two performance measure functions $\mathcal{F}(I_{bat}, \dot{m}_f)$ and $\mathcal{G}(I_{bat}, \dot{m}_f)$, which operate on the instantaneous values of the battery current and the fuel flow rate are used to determine the performance of the vehicle with respect to the two

objectives of using the energetic resources of the vehicle in the correct ratio and using them efficiently. $\mathcal{F}$ measures the relative use of the two sources of energy and $\mathcal{G}$ measures the total amount of energy used. They will be considered in more detail in subsequent sections.

The values of I_{bat} and $\dot{m}_f$ will vary as the powertrain operates at different speeds and different pedal values and also as the rules change. To determine the performance of any given set of controller rules in different speed-load regions of powertrain operation, the functions $\mathcal{F}(I_{bat}, \dot{m}_f)$ and $\mathcal{G}(I_{bat}, \dot{m}_f)$ should, therefore, be distinctly evaluated at various different operating speeds and pedal values. The vectors $\mathbf{P} = [p_1, p_2, \cdots, p_n]$ of pedal values and $\mathbf{\Omega} = [\omega_1, \omega_2, \cdots, \omega_m]$ of powertrain speed values, are used to partition the powertrain speed-pedal value space into an $n \times m$ grid, as shown in Figure 5.2. Each point in the grid formed by $\mathbf{P}$ and $\mathbf{\Omega}$ represents the performance in the region $\mathcal{S}_{(p_i,\omega_j)}$, shown inside the dashed box in Figure 5.2.

The elements of two matrices $\mathbf{F}$ and $\mathbf{G}$, of size $n \times m$, store estimates of the values of $\mathcal{F}(I_{bat}, \dot{m}_f)$ and $\mathcal{G}(I_{bat}, \dot{m}_f)$ under different powertrain operating conditions. Each element, $f_{i,j}$, in $\mathbf{F}$ stores an estimate of the value of $\mathcal{F}(I_{bat}, \dot{m}_f)$ in the region $\mathcal{S}_{(p_i,\omega_j)}$.

If the partitioning vector $\mathbf{P}$ is set equal to Θ, making the point $(p_i, \omega_j) = (\theta_i, \omega_j)$, the powertrain speed-pedal value space is partitioned by the vector of pedal values which have grades of membership of 1 in the input sets and the vector Ω. The pedal value fuzzy sets are then as shown at the top of Figure 5.2. It is seen from Figure 5.2 that the performance in each region, $\mathcal{S}_{(p_i,\omega_j)}$, which is now $\mathcal{S}_{(\theta_i,\omega_j)}$, will be strongly influenced by the rule whose fuzzy input set has its maximum grade of membership at θ_i. Therefore, the matrices $\mathbf{F}$ and $\mathbf{G}$ will contain several estimates of the values of the functions $\mathcal{F}(I_{bat}, \dot{m}_f)$ and $\mathcal{G}(I_{bat}, \dot{m}_f)$ for each pedal input rule, each of these estimates being taken from data close to

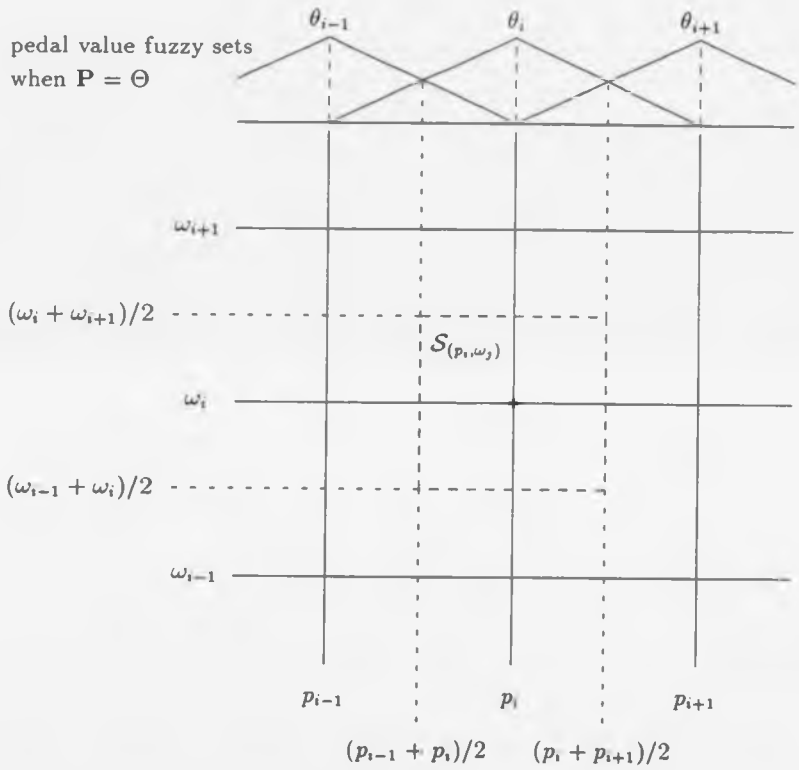

Figure 5.2: Method of clustering performance measure values.

the speed values, ω_j.

Consider the performance measure that attempts to measure the overall efficiency of the powertrain at any instant, $\mathcal{G}(I_{bat}, \dot{m}_f)$. This would involve measuring the total use of energy at any instant by a function such as:

$$\mathcal{G} = \alpha I_{bat} + \beta \dot{m}_f \tag{5.1}$$

The purpose of the performance measure matrices **F** and **G** is to store estimates of the performance of the vehicle in the regions $\mathcal{S}_{(\theta_i,\omega_j)}$. However, it is readily seen that, within each region, $\mathcal{S}_{(\theta_i,\omega_j)}$ as shown in Figure 5.2, different amounts of power are developed by the powertrain, more power being developed in the top right hand corner than the bottom left hand corner. As the vehicle is used, the values of pedal value and powertrain speed will vary, causing the operating point to enter numerous different regions, $\mathcal{S}_{(\theta_i,\omega_j)}$. If, during a period of operation, such as an adaptation interval, for a particular region, $\mathcal{S}_{(\theta_i,\omega_j)}$, the operating point is frequently in the higher power area of the region, then high values of the example performance measure, given in Equation 5.1, will result. Conversely, the opposite could occur, and the performance measure values would then be low. Before the modifications are made to the controllers, the average of the performance measure values obtained over the adaptation interval is required. Clearly, in the two cases described above, the average value obtained would not be representative of the performance that would have been obtained had the operating points been more uniformly spread over the entire region $\mathcal{S}_{(\theta_i,\omega_j)}$. To overcome this problem, before the fuel flow rate and the battery current are used, they are adjusted by referring them to the point (θ_i, ω_j), allowing (θ_i, ω_j) to be representative of the region, $\mathcal{S}_{(\theta_i,\omega_j)}$.The method by which the values of the fuel flow rate and the battery current are referred to the point (θ_i, ω_j) is now explained.

Two reference maps are used, the first relates the engine fuel flow rate to the throttle angle and the engine speed, and the second relates the battery current

to the demanded armature current and the engine speed. In Chapter 3 it was found that the motor was far more efficient if it turned at twice the speed of the engine, than at the same speed as the engine. At high engine speeds, turning the motor twice as fast as the engine would exceed the maximum speed of the motor. Therefore, the motor to engine gear ratio was set to be 2 at low engine speeds (below 3250 r/min) and 1 at higher engine speeds. The cost of a mechanism to implement this on a vehicle was felt to be justified by the considerable improvement in the resulting vehicle performance. The battery current reference map is defined over engine speed rather than motor speed because, in use the motor speed is determined by the engine speed. These reference maps are used to account for variations in individual components and were obtained by a process described in Section 5.2.2.

As explained in Section 3.5, the powertrain controller has a sampling interval of 0.1 seconds and the performance measure values are evaluated at each sampling interval. The method by which operating points are referred or "clustered" to the points (θ_i, ω_j) is shown graphically in Figure 5.3. The method used assumes that, whilst there will be variations in the magnitude of, for example, the engine fuel flow rate from one engine to another, the rate of change of the fuel flow rate with respect to the throttle angle and the engine speed will be similar for different engines in any region of operation.

The method by which points in the region $\mathcal{S}_{(\theta_i,\omega_j)}$ are clustered to the point (θ_i, ω_j) is now explained for the engine fuel flow rate (an analogous procedure is used for the motor).

1. Take the current pedal value and, using the current controller rules, determine the throttle angle that was inferred.

2. Using this throttle angle and the current engine speed, determine the reference fuel flow rate, $\dot{m}_{fr1}$, at the current operating point using bilinear

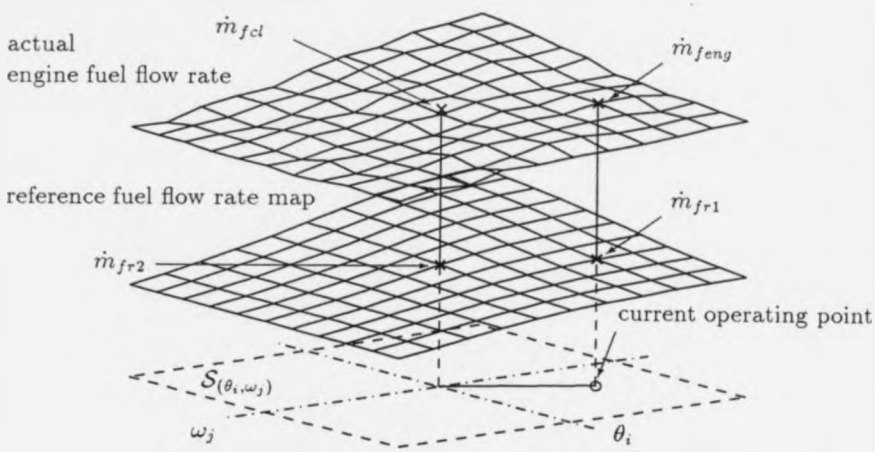

Figure 5.3: Diagram indicating the performance measure clustering procedure.

interpolation in the reference fuel flow rate map.

3. Determine the closest point, (θ_i, ω_j), in the performance measure space to which the fuel flow rate will be clustered.

4. Using the values θ_i and ω_j, the current controller rules and the reference fuel flow rate map, determine a second reference fuel flow rate, $\dot{m}_{fr2}$, at the point (θ_i, ω_j).

5. Add to the second reference fuel flow rate the difference between the first engine fuel flow rate and the actual engine fuel flow rate, $\dot{m}_{feng}$, calling the result $\dot{m}_{fcl}$.

6. Output the clustered fuel flow rate, $\dot{m}_{fcl}$, and the indices i, j of the point in the performance measurement grid.

Having obtained clustered values for both the engine fuel flow rate and the battery current, the performance measure values stored in the **F** and **G** matrices

can then be updated. On each occasion that the operating point is found in the space $\mathcal{S}_{(\theta_i,\omega_j)}$ the element $c_{i,j}$ of an $n \times m$ matrix, $\mathbf{C}$, of counters is incremented. The value of the performance measures, $\mathcal{F}(I_{batcl}, \dot{m}_{fcl})$ and $\mathcal{G}(I_{batcl}, \dot{m}_{fcl})$, are then calculated using the clustered values of battery current and fuel flow rate. The element $f_{i,j}$ of the matrix $\mathbf{F}$ at time NT is then updated using the formula:

$$f_{i,j}^{NT} = f_{i,j}^{(N-1)T} \frac{(c_{i,j} - 1)}{c_{i,j}} + \frac{\mathcal{F}(I_{batcl}, \dot{m}_{fcl})}{c_{i,j}} \tag{5.2}$$

which has the great advantage of producing an exact, unbiased average of any number of data points, without the need to store each individual value. The same procedure is carried out on the element $g_{i,j}$ in the matrix $\mathbf{G}$ using the function $\mathcal{G}(I_{batcl}, \dot{m}_{clust})$. Each element $f_{i,j}^{NT}$ and $g_{i,j}^{NT}$, is, therefore, an average value of the performance of the vehicle, as measured by the functions $\mathcal{F}(I_{bat}, \dot{m}_f)$ and $\mathcal{G}(I_{bat}, \dot{m}_f)$, in the region $\mathcal{S}_{(\theta_i,\omega_j)}$.

The controller is then modified using the matrices $\mathbf{F}$ and $\mathbf{G}$ after the adaptation interval has passed. The matrices $\mathbf{F}$, $\mathbf{G}$ and $\mathbf{C}$ are then cleared and the process is repeated.

Rule modification.

The rule modification procedure is responsible for modifying the fuzzy rules to improve the vehicle performance. As mentioned in the previous section, the input fuzzy sets remain fixed and the output fuzzy sets are moved. This movement could cause variations in the relationship between the pedal value and the driveshaft torque, so a method of moving the output fuzzy sets that maintains this relationship is required. It is assumed in this work that the driveshaft torque cannot be measured and its effects cannot be observed. In the absence of any information, some assumptions have to be made about the generation of engine and motor torque. The approach taken here is to use an estimate of the way in

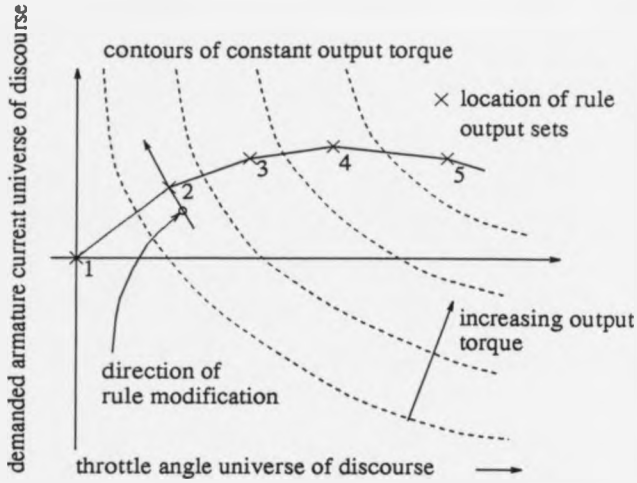

Figure 5.4: Controller output space, defined by throttle angle universe of discourse and demanded armature current universe of discourse.

which the engine torque changes, with respect to the throttle angle, and the motor torque changes, with respect to the demanded armature current. The reason for this is that the magnitude of the engine torque will vary between individual engines, but the rate of change of engine torque with respect to throttle angle should remain similar in any particular region of engine operation.

The performance of the electric motor is less sensitive to very small dimensional variations than the engine and so, in this sense, is more straightforward to control. However, the performance of the motor will be dependant upon variations in the circuit in the vehicle, particularly variations in the battery. Since the battery voltage can be measured directly, some of the variations can be directly accounted for. The rates of change of motor torque with respect to the demanded armature current could be varied as the battery voltage falls. At this initial, investigative stage, this was not done.

Since the rule input sets are fixed, the performance of the rules is determined

by the positions of the output sets. Considering the output space formed by the throttle angle and demanded armature current universes of discourse is helpful. A diagram of this output space is included in Figure 5.4.

Each of the crosses in Figure 5.4 represents the location of the output sets of a fuzzy rule, (ϕ_i, a_i), for which there is a corresponding input fuzzy set whose maximum grade of membership is at θ_i. Since, the method of fuzzy inference and defuzzification can be reduced to linear interpolation, the throttle value and demanded armature current will move along the line linking the crosses as the pedal value increases. Also shown in Figure 5.4, are the contours of constant combined engine and motor output torque, obtained at a particular engine speed. Considering this engine speed alone, in order to maintain a constant driveshaft torque, the output set whose locations are marked 2 in Figure 5.4, could be moved in the direction shown by the arrow.

When the rule is modified it moves to a location $(\phi_i + \Delta\phi_i, a_i + \Delta a_i)$. For the relationship between the pedal value and the driveshaft torque to remain constant the following expression should hold:

$$\Delta\phi_i \frac{\partial T_e}{\partial thr} + \Delta a_i \frac{\partial T_g}{\partial Iad} = 0 \tag{5.3}$$

where $\frac{\partial T_e}{\partial thr}$ is the rate of change of engine torque with respect to the throttle angle and $\frac{\partial T_g}{\partial Iad}$ is the rate of change of the motor output torque with respect to the demanded armature current. To move the rule being considered a distance, Δ_{mv}, the following constraint is applied:

$$\Delta\phi_i^2 + \Delta a_i^2 = \Delta_{mv}^2 \tag{5.4}$$

The rule offsets should then be moved as follows:

$$\Delta\phi_i = -\Delta_{mv}/\sqrt{1 + 1/\sigma^2} \tag{5.5}$$

$$\Delta a_i = \Delta_{mv}/\sqrt{1 + \sigma^2} \tag{5.6}$$

where:

$$\sigma = \frac{\partial T_g}{\partial Iad} / \frac{\partial T_e}{\partial thr} \tag{5.7}$$

The value that is used for Δ_{mv} is determined by the performance measurement values and the desired performance of the vehicle. It will readily be seen that this approach moves the locations of the output sets by simple Euler integration. The methods used to determine Δ_{mv} will be covered in subsequent sections that describe the simulation experiments that were carried out using the method.

The values $\frac{\partial T_g}{\partial Iad}$ and $\frac{\partial T_e}{\partial thr}$ are obtained by interpolation within maps which store the values of the partial derivatives at locations in a grid of points. The values of $\frac{\partial T_e}{\partial thr}$ are stored in a matrix, $\mathbf{J}_e$, using vertices in a grid in the throttle angle-engine speed space, and the values of $\frac{\partial T_g}{\partial Iad}$ are stored in a similar matrix, $\mathbf{J}_m$, at vertices in a grid in the demanded armature current space. Once again, bilinear interpolation is used to determine values of the partial derivatives at the location of the rule. A problem arises when evaluating these partial derivatives, since the value of, for example, $\frac{\partial T_e}{\partial thr}$ will vary as the engine speed changes. The problem of selecting the correct value will be addressed in Section 5.3.

The generation of the matrices $\mathbf{J}_e$ and $\mathbf{J}_m$ is covered in the next section.

5.2.2 Generation of the reference maps.

The generation of the reference maps is required for use in the clustering algorithm during performance measurement and also to create the matrices $\mathbf{J}_e$ and $\mathbf{J}_m$. A requirement of the clustering algorithm is that the reference maps are smooth, preventing local irregularities in the reference maps from influencing the clustering algorithm. A requirement of the rule modification algorithm is that the partial derivatives are smooth. Since these partial derivatives will be obtained by numerical differentiation of the reference maps the reference maps must be smooth in both value and derivative.

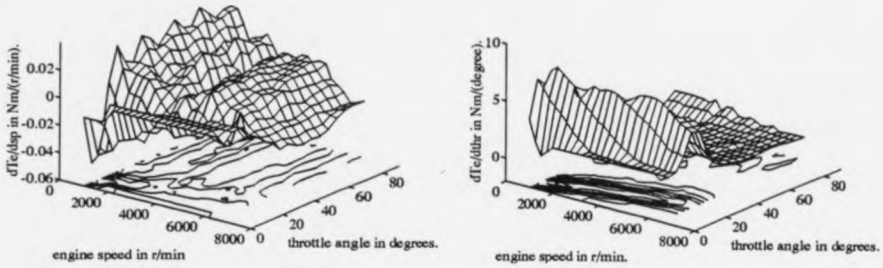

Figure 5.5: Mesh and contour plots of $\frac{\partial T_e}{\partial sp}$ and $\frac{\partial T_e}{\partial thr}$ obtained by numerical differentiation of raw engine data.

The generation of the maps for the engine will be considered first.

Generation of engine data reference maps.

Maps of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_e}{\partial sp}$, the partial derivative of engine torque with respect to the engine speed, obtained by numerical differentiation of the raw engine data, are shown in Figure 5.5.

It can be seen that the process of numerical differentiation has resulted in an extremely uneven surface, particularly for $\frac{\partial T_e}{\partial sp}$. Use of such surfaces in the rule modification algorithm would result in a rather erratic progression of the rules across the rule output space. Also, there are regions in the $\frac{\partial T_e}{\partial thr}$ map which have negative values, indicating that the torque would decrease as the throttle value increased. This would cause problems in the rule modification strategy and also would, clearly, not be representative of the actions of the engine. This situation might be improved by smoothing the surface before differentiating it. Consider a portion of the engine map:

$$\begin{matrix} c_1 & b_1 & c_2 \\ b_2 & a & b_3 \\ c_3 & b_4 & c_4 \end{matrix}$$

in which it is desired to adjust the value a, to obtain a "smoother" new value, a',

by application of the formula:

$$a' = \lambda a + (1 - \lambda) \left[\mu \sum_{i=1}^{4} b_i/4 + (1 - \mu) \sum_{i=1}^{4} c_i/4 \right] \tag{5.8}$$

where:

$$\mu = \frac{\sqrt{2}}{1 + \sqrt{2}} \tag{5.9}$$

and is used to weight the b_i values more than the c_i values in the estimate of the new value, a'. The value μ is obtained by considering the differing distances between the elements b_i and a and the elements c_i and a. The coefficient λ takes values between 0 and 1 and is a measure of the extent to which the value a is being adjusted, high values of λ leaving a largely unchanged. To avoid the effects due to starting in a particular location, the new values a' are placed in a separate matrix.

This algorithm was applied to the engine torque map before differentiation and the resulting partial derivative maps are shown in Figures 5.6 and 5.7 for $\lambda = 0.7$ and $\lambda = 0.3$.

The smoothing actions in Figure 5.6 with $\lambda = 0.7$ are insufficient. However, the effect of smoothing the engine map using lower values of λ such as $\lambda = 0.3$ is to change the shape of the derivative map. Also, there are still negative values in some regions in the $\frac{\partial T_e}{\partial thr}$ map. Smoothing the derivative maps after differentiating the raw engine data was also attempted, generating the maps shown in Figures 5.8 and 5.9.

Whilst the $\frac{\partial T_e}{\partial sp}$ map in Figure 5.9 is quite reasonable, comparing the $\frac{\partial T_e}{\partial thr}$ map with Figure 5.5 indicates that this map has changed shape significantly. Also, it still contains negative values. The lack of success in using smoothing methods to modify the maps obtained by simple numerical differentiation indicates that a more sophisticated method of obtaining the $\mathbf{J}_e$ matrix is required.

Linear regression is technique that can be used in such a way that the production of maps that can be guaranteed to be continuous in both value and derivative

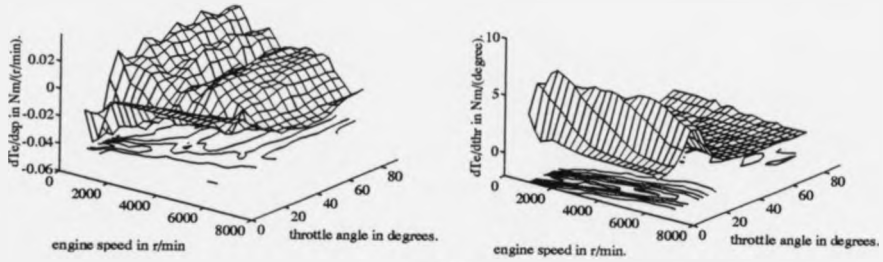

Figure 5.6: Mesh and contour plots of $\frac{\partial T_e}{\partial sp}$ and $\frac{\partial T_e}{\partial thr}$ obtained by numerical differentiation of smoothed engine data ($\lambda = 0.7$).

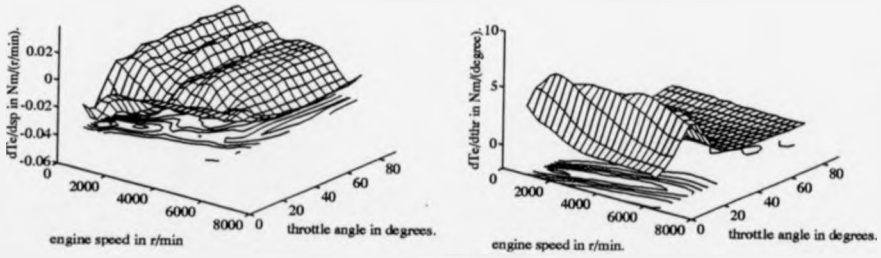

Figure 5.7: Mesh and contour plots of $\frac{\partial T_e}{\partial sp}$ and $\frac{\partial T_e}{\partial thr}$ obtained by numerical differentiation of smoothed engine data ($\lambda = 0.3$).

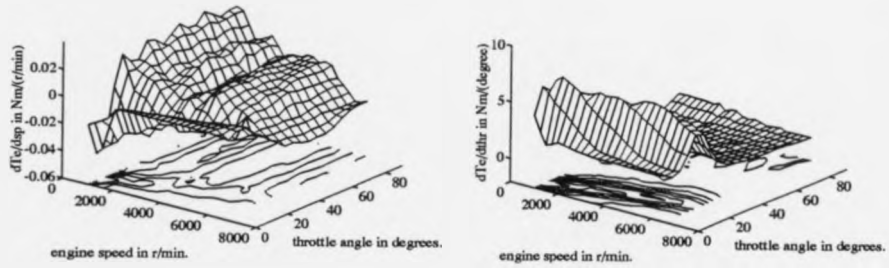

Figure 5.8: Smoothed ($\lambda = 0.7$) mesh and contour plots of $\frac{\partial T_e}{\partial sp}$ and $\frac{\partial T_e}{\partial thr}$ obtained by numerical differentiation of smoothed engine data.

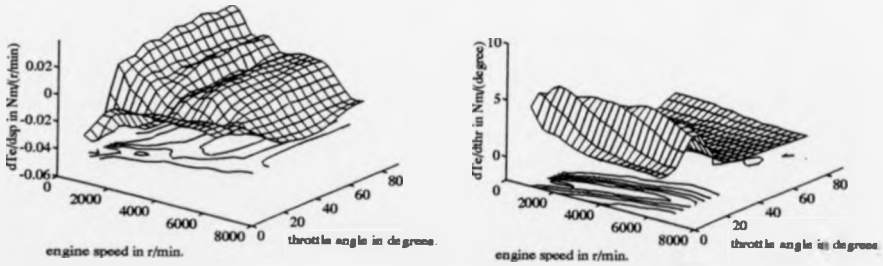

Figure 5.9: Smoothed ($\lambda = 0.3$) mesh and contour plots of $\frac{\partial T_e}{\partial sp}$ and $\frac{\partial T_e}{\partial thr}$ obtained by numerical differentiation of smoothed engine data.

is possible. Reference [201] gives a comprehensive introduction to the subject. Linear regression uses a linear combination of functions to approximate an unknown function and can be used to obtain "clean" values of a function in the presence of noise or to investigate the form of, for example, experimental test results.

Suppose it is desired to approximate a function, $\mathcal{H}(x_1, \ldots, x_n)$. Consider the following expression:

$$\mathcal{H}(x_1, \ldots, x_n) \approx \alpha_1 \mathcal{H}_1(x_1, \ldots, x_n) + \cdots + \alpha_m \mathcal{H}_m(x_1, \ldots, x_n) \tag{5.10}$$

The linear regression technique is used to obtain the values $\alpha_1, \ldots, \alpha_m$ that make the values generated by the right hand side a least mean square error approximation of the function $\mathcal{H}(x_1, \ldots, x_n)$. The success of a linear regression exercise is determined by the careful selection of functions that are used to approximate the values on the left hand side.

Both the engine torque and fuel flow rate maps were regressed in this manner, torque map regression will be considered first. The engine torque map is defined to be a function of two variables, engine speed and engine throttle angle. The values of engine torque are stored at throttle angle intervals of 5°, between 0° and 90° and engine speed intervals of 250 r/min, from 1000 r/min to 6500 r/min. Vectors Ψ and R, with elements ψ_i and r_j, of lengths 19 and 23, are used to refer to the throttle angles and engine speeds that define the points at which the torque values were measured.

Estimating functions of two variables that will satisfactorily approximate a two dimensional surface is far from easy. One method of approaching this problem is to consider how the data to be regressed vary with just one variable whilst the other is held fixed. It was felt that engine torque was likely to be a more complicated function of throttle angle than engine speed and accordingly, the engine speed was fixed and the torque data was regressed against the throttle angle.

Figure 5.10 shows plots of engine torque against throttle angle at engine speeds

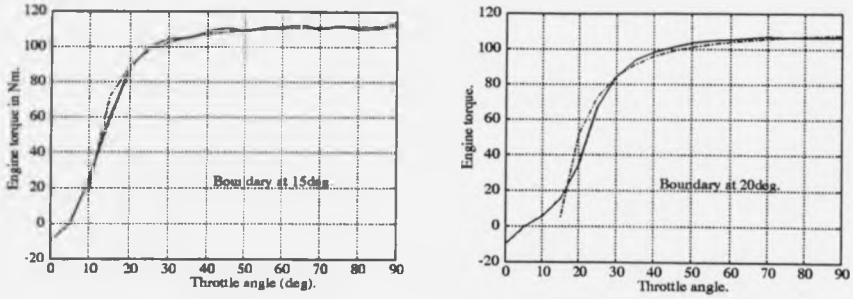

Figure 5.10: Plots of engine torque against throttle angle at engine speeds of 2500 r/min and 6000 r/min. Dashed lines show regression function fits for low throttle angles and dash-dot lines show regression fits for high throttle angles.

of 2500 r/min and 6000 r/min. No single function could be found that would give a reasonable fit to the throttle angle-engine torque relationship for all values of throttle angle. It was decided that different regression functions would be used for low throttle angles and high throttle angles and that these functions would then be blended. The relationship between the throttle angle and the engine torque was plotted at each engine speed at which engine torque data was available and regression fits were obtained for low throttle angles using the cubic:

$$T_e = \alpha_{1,r_j} + \alpha_{2,r_j}\psi + \alpha_{3,r_j}\psi^2 + \alpha_{4,r_j}\psi^3 \tag{5.11}$$

For high throttle angles the following function was used:

$$T_e = \alpha_{4,r_j} + \alpha_{5,r_j}/\psi^2 \tag{5.12}$$

where ψ is the general variable representing throttle angle. The values, α_{1,r_j}, etc., are returned by the linear regression method, for the particular engine speed, r_j, being considered. The distinction between high and low throttle angles was made by defining an index, b, and an element, ψ_b, in the throttle angle vector, Ψ, which formed a boundary between the two regions at each engine speed. The low throttle angle regression was then performed using the throttle values $[\psi_1, \ldots, \psi_{b+1}]$ and

the corresponding engine torque values, the high throttle angle regression was performed using the throttle values $[\psi_{b-1}, \ldots, \psi_{19}]$ and corresponding engine torque data. There is, therefore, an overlap between the two regions, $[\psi_{b-1}, \psi_b, \psi_{b+1}]$, which influences the regression values, α, returned. The region, $[\psi_{b-1}, \psi_b, \psi_{b+1}]$, is, however, the region which will be blended rather than regressed, the use of the overlap tends to make the blending action more straightforward.

The value, ψ_b, for each engine speed value was determined by examining the regression fits for candidate values, ψ_b, to determine which ψ_b gave the best regressed fit. In general, this was not a difficult task since there were only one or two candidate values. The regressed data are shown in Figure 5.10 by the dashed and the dash-dotted lines. The plots also show the value of ψ_b. In general, the value ψ_b, increased as the engine speed increased. Figure 5.11 shows the values of engine speed and throttle value in which the engine operates, the straight, solid black line segments showing the values that were chosen for ψ_b for different engine speeds.

At this stage, the regression exercise had related the engine torque to the throttle angle for all engine speeds considered. The objective of the regression exercise was to find analytic functions of throttle angle and engine speed to represent the engine torque. By considering the way in which the values $\alpha_{1,r_j}, \alpha_{2,r_j}$, etc varied, as the engine speed varied, candidate functions of throttle angle and engine speed were obtained. This was carried out by plotting the values $\alpha_{1,r_j}, \alpha_{2,r_j}$, etc against the values r_j, as shown in Figures 5.12 and 5.13 for the two regression coefficients, α_{4,r_j} and α_{5,r_j}. By examining the shape of the plots, functions of engine speed could then be used to approximate the values of $\alpha_{1,r_j}, \alpha_{2,r_j}$, etc. In Figures 5.12 and 5.13 the values of α_{4,r_j} and α_{5,r_j} are shown with solid lines and the regression function approximations are shown with the dotted lines. The functions chosen

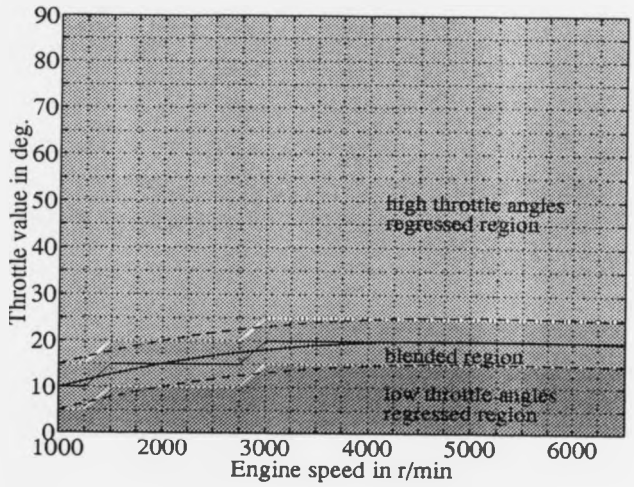

Figure 5.11: Plot showing the different regions used in the first stage of the engine torque map regression.

were:

$$\alpha_{4,r_j} \approx \beta_1 + \beta_2 r_j + \beta_3 r_j^2 + \beta_4 r_j^3 \tag{5.13}$$

$$\alpha_{5,r_j} \approx \beta_5 + \beta_6 r_j + \beta_7 r_j^2 + \beta_6 r_j^3 \tag{5.14}$$

This procedure was carried out for all the α values to obtain functions of both throttle angle and engine speed that could be used to approximate the engine torque at low throttle angles and at high throttle angles. By considering the variation of α_{1,r_j}, α_{2,r_j} and α_{3,r_j} as r_j varied, the functions obtained at low throttle angles were:

$$T_e \approx \gamma_1 + \gamma_2/r^2 + \gamma_3\psi + \gamma_4\psi/\sqrt{r} + \gamma_5\psi^2 + \gamma_6\psi^2/r + \gamma_7\psi^3 + \gamma_8\psi^3/r \tag{5.15}$$

By considering the variation of α_{4,r_j}, and α_{5,r_j} as r_j varied, the functions obtained for the high throttle angles were:

$$T_e \approx \gamma_9 + \gamma_{10}r + \gamma_{11}r^2 + \gamma_{12}r^3 + \gamma_{13}/\psi^2 + \gamma_{14}r/\psi^2 + \gamma_{15}r^2\psi^2 + \gamma_{16}r^3/\psi^2 \tag{5.16}$$

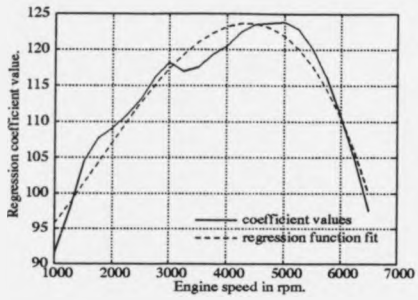

Figure 5.12: Plot of the value of α_{4,r_j} against the engine speed values, r_j.

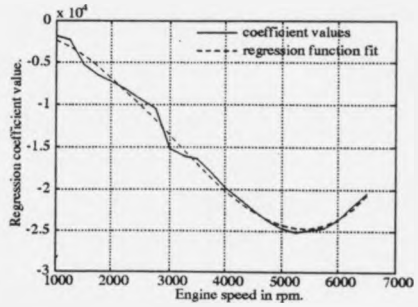

Figure 5.13: Plot of the value of α_{5,r_j} against the engine speed values, r_j.

where ψ and r are the general variables representing throttle angle and engine speed. Having obtained the γ values, functions relating the engine torque to the throttle angle and engine speed were obtained for the low and high throttle angle regions. Their combined effect is shown in Figure 5.14.

The next task is to blend the two regions. In order to do this, for each speed value, a cubic spline relating the engine torque to the throttle angle was calculated. The use of the cubic spline is advantageous because it can be guaranteed that the engine torque-throttle angle relationship will be continuous in value and derivative through the low throttle value region, the blending region and into the high throttle value region. In order to ensure that the engine torque-speed relationship is continuous in value and derivative, additional steps must be taken. The cubic splines are evaluated by using the values and derivatives of the regression functions at the edge of the blending region. However, the edges of the blending region are not smooth, therefore, the coefficients of the blending splines will not vary smoothly with the engine speed and the resulting engine torque-speed relationship will not be smooth in value and derivative. In order to overcome this difficulty, the blending splines are evaluated along the two dotted lines in Figure

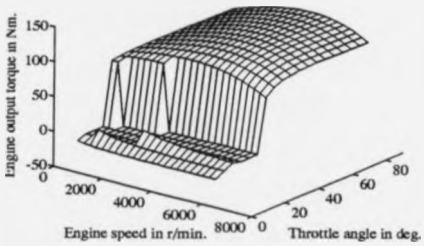

Figure 5.14: Plot of the surface obtained by regressing the engine torque data using different functions of throttle angle and engine speed for low and high throttle values.

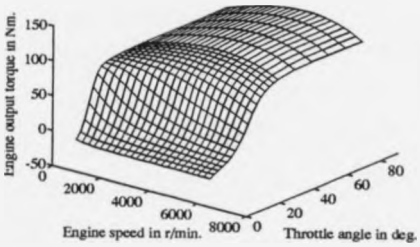

Figure 5.15: Plot of the regressed data with the blending function. Note that the values of throttle angle at which the new torque data is stored have been redistributed.

5.11. The values of throttle angle along these dotted lines vary smoothly as the engine speed varies. By evaluating the blending spline coefficients along these lines, the coefficient values will also vary smoothly as the engine speed varies. This will cause the engine torque-speed relationship to be smooth in value and derivative.

In order to obtain the dotted lines, the curved solid line was first obtained by fitting another cubic polynomial to the solid line with straight sections. The dotted lines are placed 5° of throttle angle either side of the smooth curved line. The last point to be made is that all the engine torque values within the dotted lines should be obtained using the blending splines and all the engine torque points outside the dotted lines should be obtained using the regression functions. In some cases this involved replacing regressed data with blended data. At this stage, an alternative vector of throttle angles, Ψ', whose values were spaced every 2.5° of throttle angle upto 40° and every 10° of throttle angle above this was used. The use of this new vector of throttle values allowed the important parts of the engine map to be represented in more detail, without storing any more information. The

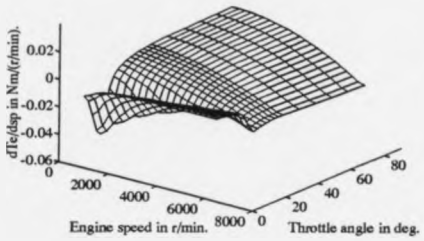

Figure 5.16: Plot of the map of $\frac{\partial T_e}{\partial sp}$ obtained by numerically differentiating the data shown in Figure 5.15.

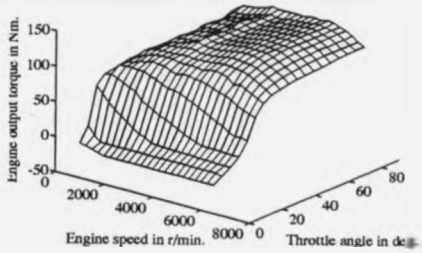

Figure 5.17: Plot of the original engine torque data.

regressed and blended data was then evaluated at these new throttle angles using the regressed and blending functions. The engine torque map obtained by this process is shown in Figure 5.15.

The map of $\frac{\partial T_e}{\partial sp}$ shown in Figure 5.16 was obtained by numerically differentiating the map shown in Figure 5.15. Comparing this map with that shown in Figure 5.5, it is seen that the objective of producing a smooth derivative map has been achieved, however, some of the detail in the shape of the surface has been lost. This is due to a slight ripple, probably due to inlet manifold "pulsing", that can be seen in the engine torque map shown again in Figure 5.17.

The second "hump" in this ripple is larger than the first and it was found that a function of the form:

$$r\cos((r-5000)\pi/1125) \tag{5.17}$$

was able to introduce the ripple into the regressed data. To restrict the action of the function shown in Equation 5.17, and the upper part of Figure 5.18, to the regions of the engine map in which the ripple is present, it was multiplied by the bump function shown in the lower part of Figure 5.18. Repeating the regression process produced the map shown in Figure 5.19, from which the derivative maps

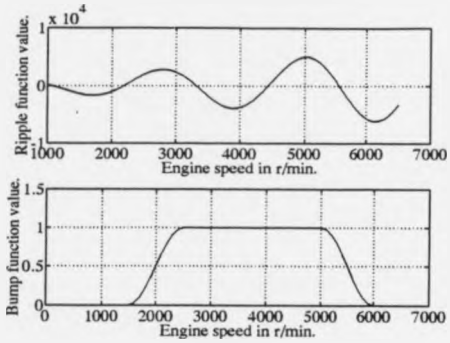

Figure 5.18: Upper plot shows trigonometric function used to include the ripple in the torque map, lower plot shows the bump function used to restrict the ripple function application.

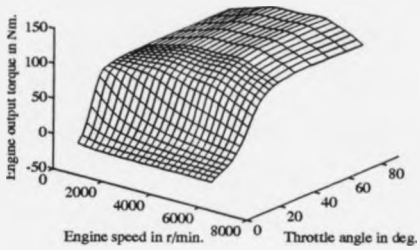

Figure 5.19: Plot of the regressed engine torque data including the ripple function.

shown in Figures 5.20 and 5.21 were obtained. There are still negative values in the $\frac{\partial T_e}{\partial th}$ map. However, they are around an order of magnitude smaller than the negative values in the map shown in Figure 5.5, obtained from the raw engine data. These values arise during the regression involving two independant variables and constraining the regression to remove this problem was not felt to be worthwhile.

The engine fuel flow rate map also has irregularities due to measurement noise associated with testing real engines. The partial derivatives of the engine fuel flow rate with respect to the engine speed and throttle angle are not required, since the fuel flow rate map is used in the clustering algorithm and not in the rule modification algorithm. However, the clustering algorithm relies upon the reference map being locally smooth and the partial derivatives of the engine fuel flow rate will give a good indication of the local smoothness of the map. The partial derivative maps obtained from numerical differentiation of the fuel flow rate maps are shown in Figures 5.22 and 5.23.

A very similar procedure to that employed in regressing the engine torque

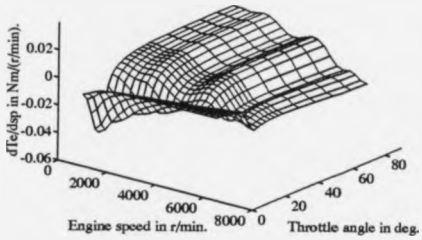

Figure 5.20: Plot of the $\frac{\partial T_e}{\partial sp}$ map obtained by numerically differentiating the regressed engine torque map including the ripple function.

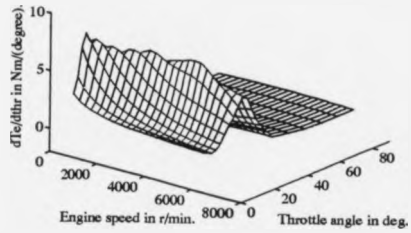

Figure 5.21: Plot of the $\frac{\partial T_e}{\partial sp}$ map obtained by numerically differentiating the regressed engine torque map including the ripple function.

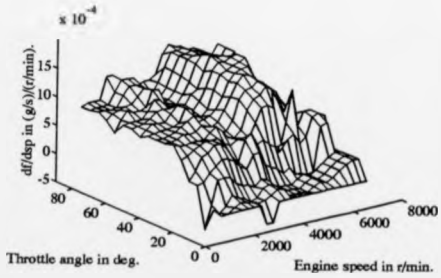

Figure 5.22: Plot of the $\frac{\partial \dot{m}_f}{\partial sp}$ map obtained by numerically differentiating the measured engine fuel flow rate map.

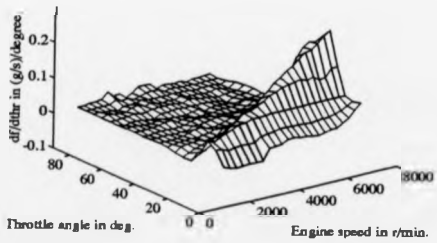

Figure 5.23: Plot of the $\frac{\partial \dot{m}_f}{\partial th}$ map obtained by numerically differentiating the measured engine fuel flow rate map.

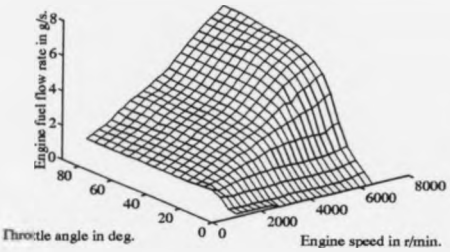

Figure 5.24: Plot of the original engine fuel flow rate map.

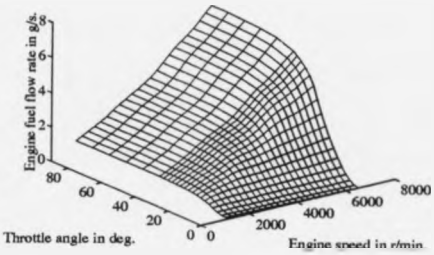

Figure 5.25: Plot of the regressed engine fuel flow rate map.

maps was used in regressing the engine fuel flow rate maps. Again, the map was regressed using two separate regions of throttle angle and subsequently blended. Similar numbers of functions were used and a cubic spline was used to blend the two regions. The blending region was a slightly different shape for the fuel rate map than for the torque map. There is no great significance in this, since the shape of the blending region will depend upon the functions used to regress the data. The engine fuel flow rate map that was developed in this manner is shown in Figure 5.25. Figure 5.24 shows the original fuel flow rate map for the purpose of comparison. The partial derivative maps are shown in Figures 5.26 and 5.27. The final modification that was made to the maps was to set the minimum negative torque developed by the engine to be 0. This is because, in the model of the powertrain used is this work, as described in Section 3.5.5, the output of the engine is passed through a one-way clutch. This aspect of the engine operation has to be accounted for in the engine reference maps, by setting the minimum torque value to be 0 Nm. This makes very little difference to the derivative maps.

The data stored in the engine reference maps is included in Appendix D.

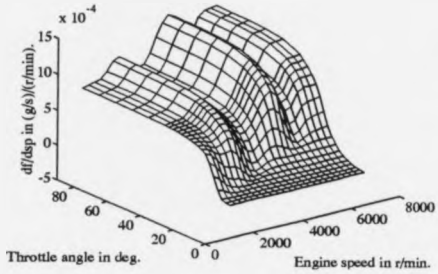

Figure 5.26: Plot of the $\frac{\partial \dot{m}_f}{\partial sp}$ map obtained by numerically differentiating the regressed engine fuel flow rate map.

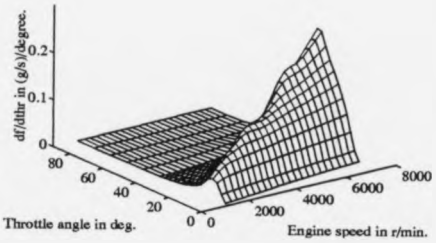

Figure 5.27: Plot of the $\frac{\partial \dot{m}_f}{\partial th}$ map obtained by numerically differentiating the regressed engine fuel flow rate map.

Generation of motor data reference maps.

In Chapter 2, a model of the traction motor was developed by considering the actions of the circuit formed by the motor, its controller and the batteries. Test data for the reference maps was not available, so the reference maps and partial derivative maps were generated by using the model to predict the performance of the motor over the range of conditions of interest. The motor maps stored motor output torque, and battery current, over a grid of values of engine speed and demanded armature current. It will be recalled that the motor speed is twice the engine speed for engine speeds below 3250 r/min and equal to the engine speed at engine speeds above this value. Engine speed is used to define the regions in the clustering algorithm and, for this reason, engine speed, rather than motor speed, is used to define the motor performance maps.

A map of the motor output torque against motor speed is shown in Figure 5.28 and a map of the motor output torque against the engine speed is shown in Figure 5.29.

Partial derivative maps, $\frac{\partial T_g}{\partial sp}$ and $\frac{\partial T_g}{\partial I_a}$ are shown in Figures 5.30 and 5.31.

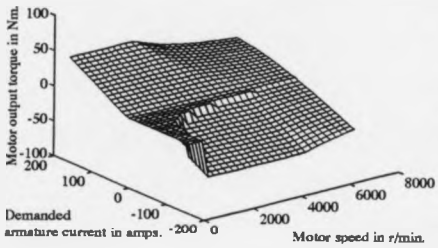

Figure 5.28: Plot of the motor output torque against demanded armature current and motor speed.

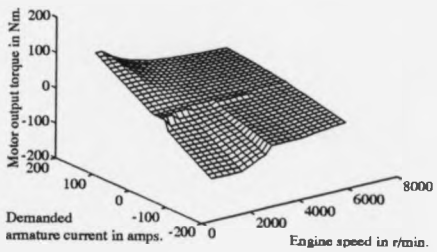

Figure 5.29: Plot of the motor output torque against demanded armature current and engine speed.

Clearly, the derivative maps are not very smooth. The irregularities on these maps are due to the decision taken when first modelling the motor, that if a negative armature current is demanded, but a positive battery current flows, then the motor controller will apply no currents at all. In Figure 5.28, it is seen that, as the demanded armature current becomes increasingly negative, a point is reached at which the battery current becomes negative. At this point, a significant negative torque is generated, causing the step in the motor torque map at this point. This causes the discontinuities in the derivative maps. It should also be noted that by turning the motor twice as fast as the engine this effect is reduced, since the engine has a minimum operational speed of 1000 r/min. The discontinuities were left in the motor derivative maps because they represent an actual, acute non-linearity that the adaptation process would have to take into account.

The battery current maps are shown in Figures 5.32 and 5.33. The corresponding derivative maps are shown in Figures 5.34 and 5.35. The data stored in the motor reference maps is also included in Appendix D.

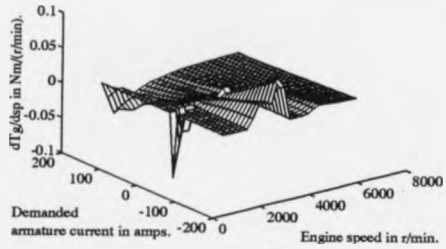

Figure 5.30: Plot of $\frac{\partial T_g}{\partial sp}$ against demanded armature current and engine speed.

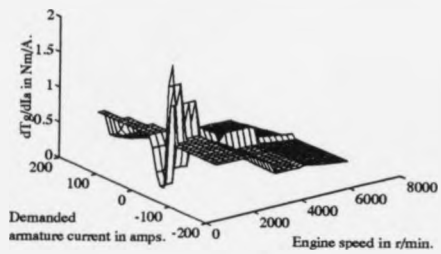

Figure 5.31: Plot of $\frac{\partial T_g}{\partial I_a}$ against demanded armature current and engine speed.

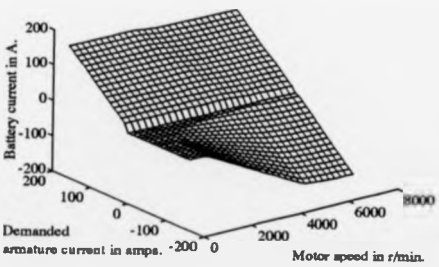

Figure 5.32: Plot of the battery current against demanded armature current and motor speed.

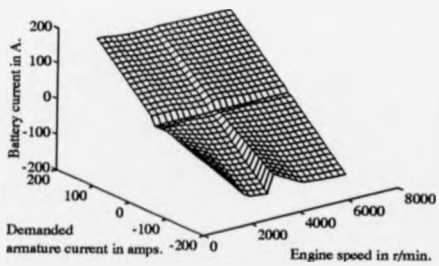

Figure 5.33: Plot of the battery current against demanded armature current and engine speed.

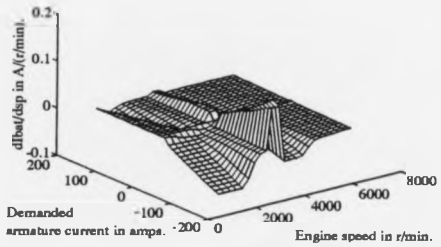

Figure 5.34: Plot of $\frac{\partial I_{bat}}{\partial sp}$ against demanded armature current and engine speed.

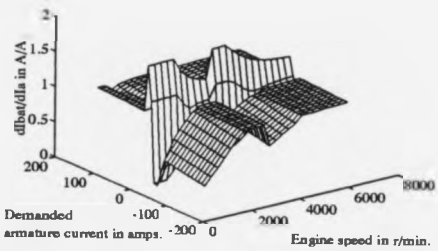

Figure 5.35: Plot of $\frac{\partial I_{bat}}{\partial I_a}$ against demanded armature current and engine speed.

5.3 Initial adaptive fuzzy hybrid vehicle powertrain controller.

The adaptive control methods developed here are applicable over the entire range of powertrain operation, however, they will be developed using controllers in which the motor works against the engine. This is because the experimental work described in Chapter 3 indicated that there was a greater opportunity for improving the efficiency of the powertrain in this region of operation. Also, variations in the torque response of the powertrain are more difficult to minimise when one of the prime movers develops negative torque and this will be a greater test of the constant torque rule modification procedure.

5.3.1 Determining the correct size and direction of rule modification.

The rule modification procedure effectively moves the rules along contours of constant output torque. This should cause controllers to be developed which attempt to maintain the driveshaft torque-pedal value relationship. The other aspect of the powertrain performance to be addressed in this initial adaptive

controller is that of using the stored fuel of the vehicle and its stored electrical energy in a ratio that can be controlled.

The performance measure used to assess the ratio of the use of electrical energy to fuel is

$$\mathcal{F} = I_{bat}/(k_e \dot{m}_f) \tag{5.18}$$

where k_e is a constant that defines the equivalence of fuel and electrical energy. By examining the shapes of the results presented in earlier work, see Figure 3.11, the value of k_e was set at 120 A/(g/s). The fuel flow rate appears in the denominator of the performance measure because its clustered value can never be zero, whereas I_{bat} can be zero. The first adaptive fuzzy controllers used the difference between a desired value, $\mathcal{F}_d$, and the performance measure value evaluated over a learning interval, to give the rule modification step size, Δ_{mv_i}, for the i'th rule

$$\Delta_{mv_i} = k_p(\mathcal{F}_d - \mathcal{F}_i) \tag{5.19}$$

where $\mathcal{F}_i$ is the value of the performance measure for the i'th rule evaluated over the last adaptation interval, T_a, and k_p is a constant that determines the rate at which the controller adapts. The value of $\mathcal{F}_i$ for any particular rule is the average of the elements $f_{i,j}$ taken over the i'th row of the matrix $\mathbf{F}$. The values, $c_{i,j}$, which count the number of times that an operating point enters a clustering pedal value-engine speed region, stored in the i'th row of the matrix, $\mathbf{C}$, are used to weight the average value of the performance measure. In forming the average, $\overline{f}_i$, over the row the following formula is used

$$\overline{f}_i = \frac{\sum_{j=1}^{m} f_{i,j} c_{i,j}}{\sum_{j=1}^{m} c_{i,j}} \tag{5.20}$$

$\mathcal{F}_i$ is then set equal to $\overline{f}_i$.

A restriction is placed on the values of Δ_{mv_i} which is that, if a rule has not been fired over the learning interval, then $\Delta_{mv_i} = 0$. This limits the rules being modified

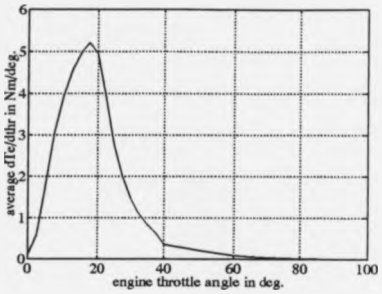

Figure 5.36: Plot of average value of $\frac{\partial T_e}{\partial thr}$ over all engine speeds (including action of one-way clutch).

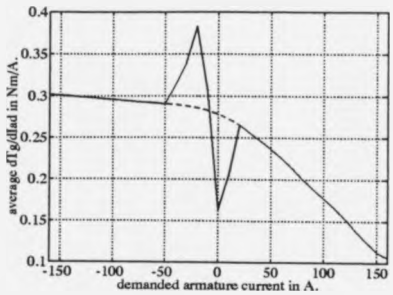

Figure 5.37: Plot of average value of $\frac{\partial T_g}{\partial Iad}$ over all engine speeds.

to those whose performance has been measured. It will be seen in Chapter 6 that there is a further large advantage to restricting the rule modification procedure in this way.

In Section 5.2.1, it was stated that the values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ are stored in matrices $\mathbf{J}_e$ and $\mathbf{J}_m$. Having decided upon the step size to be taken by the rule modification procedure, the values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ have to be obtained from $\mathbf{J}_e$ and $\mathbf{J}_m$ in order to calculate the changes in the rule output sets. The values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ at the values of throttle angle and demanded armature current of the rule being considered should be used. However, both $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ vary with engine speed. There remains the task of selecting single values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ at the particular throttle angle and demanded armature current values, from the range of values corresponding to the range of engine speeds.

Initially, the method of selecting the particular values was to pick the average value of the partial derivatives over all engine speeds. These average values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ are shown in Figures 5.36 and 5.37.

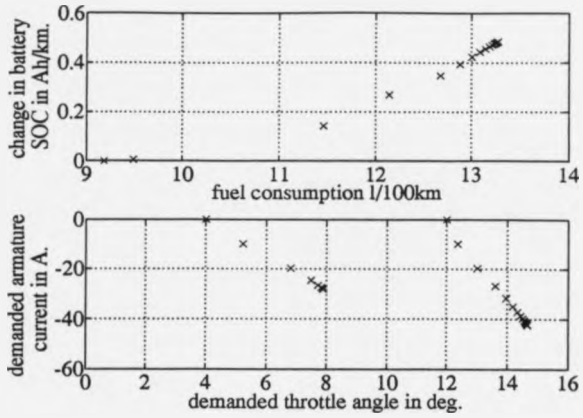

Figure 5.38: Plot of the results and controller values obtained by running successive simulations of the initial adaptive hybrid powertrain controller over the ECE-15 drive cycle.

5.3.2 Initial simulation results.

The performance of the adaptive powertrain controller was investigated by simulating successive ECE-15 cycles. At the end of each cycle the performance of the rules was measured, and the rule output set locations were then modified using the measured values of the rule performance over the last cycle. This procedure effectively sets the adaptation interval to be 196 seconds, which is the length of the portion of the ECE-15 cycle that is used in the simulations. By simulating whole cycles between between adaptations, the performance of the controller can be evaluated by the fuel consumption and change in battery state of charge obtained over each cycle. These values are shown in the upper plot of Figure 5.38 and the fuzzy output set locations that were generated by the adaptation process are shown in the lower plot of Figure 5.38 for the rules whose input sets have their maximum grades of membership at 4° and 12° of pedal value. The values of k_p

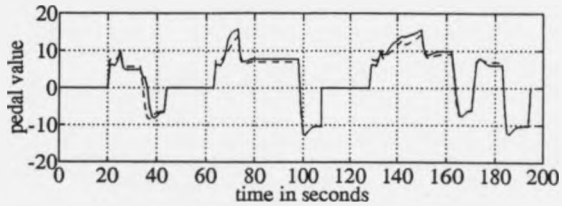

Figure 5.39: Pedal value time histories for the baseline controller (dashed line) and the adapted controller (solid line).

and of $\mathcal{F}_d$ were 50 and -0.2.

The ECE-15 cycle does not require large amounts of effort from the controller so that only the rules whose input sets have maximum grades of membership at 4° and 12° of pedal value are changed. The controller converges after around 13 ECE-15 cycles to the following values:

$$\begin{array}{lrrrrrr} \Phi & [0 & 7.91 & 14.67 & 24 & 40 & 90] \\ A & [0 & -28.15 & -42.28 & 0 & 0 & 0] \end{array}$$

Figure 5.39 shows a comparison of the pedal value time histories between the adapted controller and the baseline controller. The object of the adaptation strategy is to vary the relative use of the stored energy resources of the vehicle, whilst keeping the relationship between the pedal value and the output torque the same. Figure 5.38 illustrates how the relative use of the energy resources is varied by the adaptation strategy, and Figure 5.39, which is discussed in more detail below, gives an indication of how well the relationship between the pedal value and the driveshaft torque is maintained.

It can be seen from Figure 5.38 that the rule initially located at 4° of throttle angle converges before the rule that was originally located at 12°. This is because at 12° of throttle angle the fuel flow rates are higher than they are at 4° and, to get the same performance measure value, $\overline{f}_i$, the rule originally at 12° has to

move further. Since the rules move at a similar rate, the lower pedal value rule converges first. Comparing the results of the simulations presented in the upper plot of Figure 5.38, with those presented in Figure 3.11, it is seen that the adapted controllers use roughly similar amounts of fuel and electrical energy as the hand tuned controllers. In the sense of efficiency, the results are no better than the earlier results.

An interesting aspect of the clustering algorithm is revealed by considering a portion of the **F** matrix when the controller had converged. The top 8 rows of the **F** matrix together with the vector defining the pedal value input set locations and the speed values to which the points are to be clustered are shown below.

$$
\begin{array}{r}
\\
\text{engine speed values to which points are clustered } (\Omega)
\end{array}
\begin{array}{c}
\text{pedal value input set locations } (\Theta) \\
\begin{bmatrix} 0 & 4 & 12 & 24 & 40 & 90 \end{bmatrix} \\
\left\{ \begin{bmatrix} 1000 \\ 1250 \\ 1500 \\ 1750 \\ 2000 \\ 2250 \\ 2500 \\ 2750 \\ \vdots \end{bmatrix} \right.
\underbrace{\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
0 & -0.0549 & -0.0890 & 0 & 0 & 0 \\
0 & -0.1417 & -0.1299 & 0 & 0 & 0 \\
0.0160 & -0.2186 & -0.1563 & 0 & 0 & 0 \\
0 & -0.2639 & -0.1820 & 0 & 0 & 0 \\
0.0326 & -0.4421 & -0.2077 & 0 & 0 & 0 \\
0.0086 & -0.4200 & -0.2200 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots
\end{bmatrix}}_{\mathbf{F}}
\end{array}
$$

It can be seen from Figure 5.38 that the rules of the converged controller should not cause positive armature currents to flow. The positive values appear in the **F** matrix because of the actions of the clustering algorithm and are illustrated in Figure 5.40.

As explained previously, if low negative values of armature current are demanded, but the battery current is positive, no motor currents flow at all. The abrupt change in the motor action as the controller suddenly applies the currents will not be modelled in the battery current reference map, and the situation

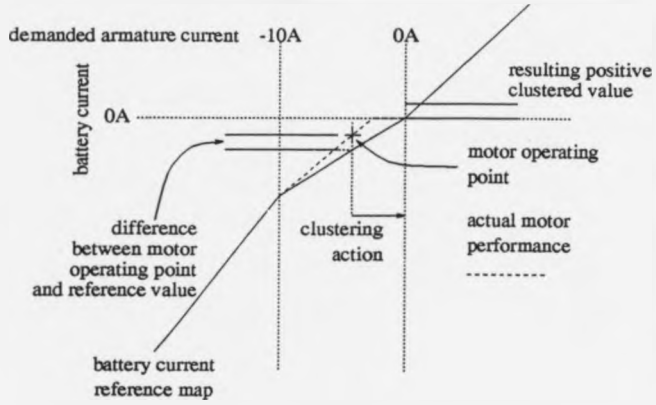

Figure 5.40: Example of clustering algorithm clustering a negative battery current value to the zero demanded armature current rule location and causing the clustered value to be positive.

shown in Figure 5.40 can occur, which results in the clustered value being positive. There is nothing particularly alarming or undesirable about this consequence of the clustering algorithm since it only occurs for the rule which is never moved.

Figure 5.39 shows the pedal value time histories of the baseline vehicle and the adapted controller. Compared with Figures 3.10 and 3.16 it is seen that the pedal value time history of the adapted controller is far more like the pedal value time history of the baseline controller than those of the hand tuned controllers. In this sense, the adaptive controller proves to be quite advantageous. The most obvious difference between the two time histories shown in Figure 5.39 is the dip in the adapted pedal value response at about 33 seconds. This dip is caused by the adaptive fuzzy controller failing to generate driveshaft torques similar to the torques generated by the engine at this instance.

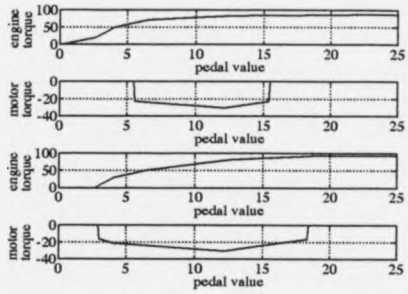

Figure 5.41: Plots of engine and motor output torque against pedal value for initial adaptive controller, (upper 2 plots at 1000 r/min, lower 2 plots at 1500 r/min).

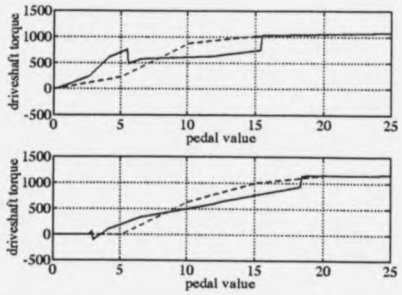

Figure 5.42: Plots of hybrid powertrain driveshaft output torque against pedal value, adaptive controller solid line, baseline vehicle dashed line, (upper plot at 1000 r/min, lower at 1500 r/min).

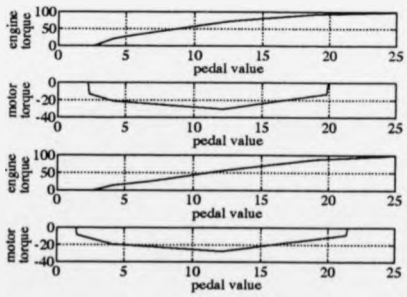

Figure 5.43: Plots of engine and motor output torque against pedal value for initial adaptive controller, (upper 2 plots at 2000 r/min, lower 2 plots at 2500 r/min).

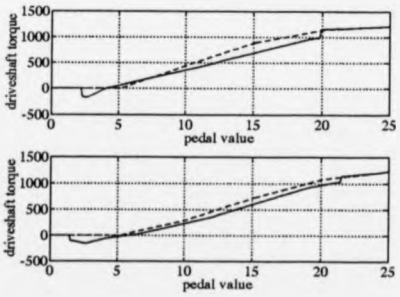

Figure 5.44: Plots of hybrid powertrain driveshaft output torque against pedal value, adaptive controller solid line, baseline vehicle dashed line, (upper plot at 2000 r/min, lower at 2500 r/min).

5.3.3 Initial adaptive controller driveability.

Figures 5.41, 5.42, 5.43 and 5.44 show the torque outputs for the engine and the motor, and the driveshaft torque, for the adapted controller. The driveshaft torque outputs for the baseline vehicle are also shown. The driveshaft torques give an indication of the vehicle driveability. In general, the driveshaft torques are maintained reasonably well. From the point of view of driveability, the rate of change of driveshaft torque with respect to the pedal value is probably as important as the magnitude of the pedal value itself. In this respect, the plots in Figure 5.44 are quite satisfactory. Plots in Figure 5.42 are less satisfactory in both value and derivative. The plot at 1000 r/min is by far the worst, but, from a practical point of view, the poor torque response at this speed is of little consequence because the powertrain never operates under load at this low speed.

The other major problem is illustrated by the motor suddenly starting to generate significant amounts of negative torque, causing a sharp dip in the driveshaft torque response, which would undoubtedly give poor driveability. In some respects, this appears worse than it actually is, since, on a conventional vehicle, the engine will develop significant negative torque at the low pedal values at which the dips occur. This negative engine torque is not shown in these plots because of the actions of the one way clutch. In use, it would be desirable for the powertrain to develop negative torque at zero throttle angle, since this would mimic the action of "engine braking" in a conventional vehicle. However, the rule at zero pedal value has to demand zero throttle angle and armature current, for those occasions when the vehicle is at rest with the engine at idle. This motivates the need to use an additional input of engine speed.

The other feature of the plots in Figures 5.42 and 5.44 is the sudden increase in torque as the motor stops operating, when the demanded armature current falls at a pedal value of about 20. This causes less of a problem in driveability because

the overall driveshaft torque is higher and the consequences of a step up in torque are not as great as those for a step down in torque. The real points of interest in these plots are the locations of the rules themselves. The rules are located at pedal values of 4 and 12. The driveshaft torques at the rule locations will only be different from the base vehicle driveshaft torques if the rules have been moved to incorrect locations. In between the rule locations, the driveshaft torque may differ from the base vehicle driveshaft torque due to sharp fluctuations in the motor or engine output torques as discussed above. The fact that the baseline driveshaft torque and the adapted controller driveshaft torque are not the same at these locations indicates that the rules have been moved to incorrect locations. This will be because the directions in which the rules are moved are calculated using partial derivative values that are average values taken over the complete range of engine speeds.

In an attempt to improve the performance of the controller adaptation in this respect, a different set of values of $\frac{\partial T_g}{\partial Iad}$ shown by the dotted line in Figure 5.37 was used. This modification had very little effect on the accuracy with which the adaptive controller torque outputs matched the baseline vehicle torque outputs. Final attempts at improving the torque matching involved forming different weighted average values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$. This was done by taking all the values, over the range of engine speeds, of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ for each rule and weighting their average by the number of times that a particular speed region was used by the vehicle. A further refinement to this was to weight the average for each rule by the number of times that the vehicle was operated in each speed region, whilst the rule being modified was used. These modifications had a slight improvement in the pedal values used over the cycle but very little effect on the torque response plots. The reason for the lack of improvement in the torque response plots was that the directions of the modifications were still being determined by a range of

engine speeds but were applied at specific engine speeds. The torque responses are a better way of determining the success of the rule modifications than the pedal value trace, because the latter type of graph simply shows how good the controller is over the ECE-15 cycle, whereas the torque response indicates how well the controller will perform in general. This is particularly important when it is considered that the ECE-15 cycle does not cover a large range of operating conditions.

Again, these observations motivate the use of engine speed as a second input, because each rule can then be moved using values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial lad}$ that are correct for the location of that rule.

5.3.4 Changing the adaptation interval.

Various simulation experiments were tried with the adaptation interval set at 20, 40, 60, 80, 120 seconds. Figure 5.45 shows the values of the throttle angle set location for the rule originally located at 4° of throttle angle when the learning interval was set at 20 seconds. It can be seen that the controller does not converge but the performance measure value for that rule oscillates around the desired performance measure causing cyclic changes to be made to the rule output set location. The data in Figure 5.46 is the same as that in Figure 5.45. The results obtained at 40, 60 and 80 second adaptation intervals were also oscillatory, but the causes of the oscillation are most easily understood by considering the results obtained with a 20 second adaptation interval.

Making reference to Table 5.1, the oscillatory behaviour is explained as follows. At modification number 15, the value $\overline{f}_2$ is more negative than the desired value, f_d and the rule is accordingly moved back towards 4°. Note that the values, $c_{2,1...7}$, show that during the last adaptation interval, when the second rule was used, the engine was operating at high speeds where more negative values of $\overline{f}_2$ are easily

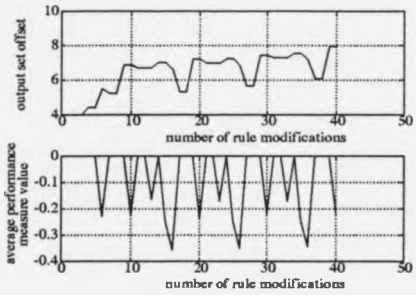

Figure 5.45: Upper plot shows the progression of the rule originally located at 4° of throttle angle, lower plot shows performance measure values for that rule.

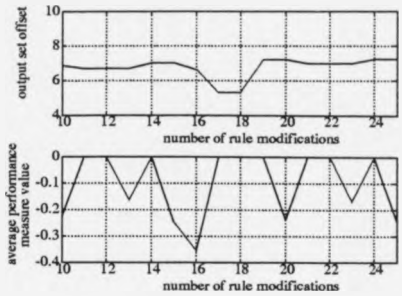

Figure 5.46: Detail from previous figure.

				Engine speeds in r/min						
mod. no.	time in cycle	$\overline{f}_2$		1000	1250	1500	1750	2000	2250	2500
15	80	-0.246	$c_{2,1\ldots7}$	0	33	6	0	0	0	59
			$f_{2,1\ldots7}$	0	-0.041	-0.114	0	0	0	-0.373
16	100	-0.356	$c_{2,1\ldots7}$	0	0	0	0	0	2	6
			$f_{2,1\ldots7}$	0	0	0	0	0	-0.383	-0.347
17	120	0	$c_{2,1\ldots7}$	0	0	0	0	0	0	0
			$f_{2,1\ldots7}$	0	0	0	0	0	0	0
18	140	0.022	$c_{2,1\ldots7}$	0	34	3	0	0	0	0
			$f_{2,1\ldots7}$	0	0.025	0.001	0	0	0	0
19	160	0	$c_{2,1\ldots7}$	0	0	0	0	0	0	0
			$f_{2,1\ldots7}$	0	0	0	0	0	0	0
20	180	-0.239	$c_{2,1\ldots7}$	0	0	0	76	0	0	13
			$f_{2,1\ldots7}$	0	0	0	-0.206	0	0	-0.433

Table 5.1: Performance measure values shown in Figures 5.45 and 5.46

achieved. At modification 16, the value, $\overline{f}_2$, is far too negative and the rule is, again, moved towards 4°. At modification 17, the rule has not been used and is not modified. At modification 18, very low engine speeds have been used when this rule fired and the value of $\overline{f}_2$ is actually positive due to the operation of the clustering algorithm as described above. This causes a large increase in the rule throttle angle and the cycle is subsequently repeated.

This oscillatory effect is most marked when very short adaptation intervals of 20 seconds are used. The effect disappears completely when 120 second adaptation intervals are used and would, in general, be less likely to occur on more random drive cycles. However, the existence of such oscillatory action, due to assessing powertrain performance over small ranges of engine speeds and then applying changes over all engine speeds, provides a further justification for the use of an additional input of engine speed to the powertrain controller.

5.4 Conclusions.

This chapter began by describing the basic functional objectives of an adaptive hybrid vehicle powertrain controller, motivated by the three performance requirements of hybrid powertrain control. An adaptive fuzzy hybrid powertrain controller structure was then described in terms of a block diagram, and the various tasks within the diagram were described in more detail. Two performance measure functions were considered, one assessing the performance of the powertrain with respect to its relative use of the two energy resources of the vehicle, the other measuring the overall efficiency of the powertrain.

In order to assess the merit of each individual rule, these performance measures were separately evaluated in the regions of operation in which each rule is dominant. Furthermore, in order to assess the merit of each rule at different engine speeds, the performance measure for each rule was separately evaluated in the

region of several different engine speeds. This performance measurement process required the use of a clustering algorithm to make the performance measured at individual points in a rule-engine speed region representative of the entire region.

The rule modification procedure was then described. In order to maintain the vehicle driveability, this procedure attempts to move rules along contours of constant driveshaft torque, starting from the rule locations of the baseline vehicle. Some information about the powertrain has to be used to move the rules and this is supplied by the use of reference maps of the derivatives of output torque with respect to the throttle angle and the demanded armature current for the engine and the motor.

The preparation of the engine reference maps involved the use of regression analysis in order to obtain smooth maps capable of being used to drive the rule modification procedure. The map generation for the motor being more straightforward because of the analytic nature of the motor model.

Results of early simulation studies were then presented. These studies show that, in principle, the adaptive controller is able to influence the relative use of fuel energy and electrical energy by the vehicle. When the pedal values for the adaptive controller are compared with the hand tuned controller, it appears that the adaptive controller maintains vehicle driveability very much better than the hand tuned controllers. However, a closer inspection of the driveshaft torques revealed that there was still a requirement for some improvement. Since only the performance measure which assesses the relative use of energy was used, no indication has yet been given of the suitability of the adaptive controller for improving the efficiency of the vehicle.

Three instances were found in which the operation of the adaptive controller would have been more effective, and straightforward, if a second control input of engine speed had been used. Additionally, since the controllers are now automat-

ically tuned, the daunting task of hand tuning a two input fuzzy hybrid vehicle powertrain controller has been removed.

The next chapter will, therefore, extend the ideas of this chapter into two input controllers and will investigate improving the powertrain efficiency.

Chapter 6

Refinements to the adaptive fuzzy hybrid vehicle powertrain controller.

6.1 Introduction.

The previous chapter described the basic method of operation of an adaptive fuzzy hybrid vehicle powertrain controller. It was seen that the significant differences between the hand tuned controllers of Chapter 3 and the adaptive controllers of Chapter 5 lay in performance measurement and in rule modification. Methods of carrying out these two tasks were developed and implemented generating an adaptive controller.

The adaptive controller successfully adapted each rule so that its performance in terms of the relative use of energy matched a demanded value. However, the controllers did not always maintain a constant powertrain output torque for any given pedal value. The major reason for this was that the controllers used only pedal value as their input.

In this chapter, an additional powertrain controller input of engine speed is introduced, allowing more accurate modification of the controller output set location. The disadvantage of the increased complexity in tuning a two input controller

is not important when the modification process is automated.

The research material covered in this chapter investigates the use of small fuzzy controllers that control various aspects of the rule modification process. This allows conclusions to be drawn regarding the relevance of fuzzy logic as a tool in adaptive and learning control applications.

Having established methods of using the extra input in the adaptation method, the rule modification technique goes on to use the overall performance of the vehicle, as well as the performance of each rule, for controller modification. The driveability of the resulting controllers is then improved and studies are carried out to investigate the possibility of improving the efficiency of the powertrain operation.

It was realised that the ECE-15 drive cycle was no longer an appropriate cycle to be assessing the performance of adaptation methods and, accordingly, data taken from an instrumented vehicle driven in an urban area, was used to exercise the vehicle model and the adaptive controllers described in this chapter.

6.2 A two input adaptive hybrid vehicle powertrain controller.

The addition of an extra input of engine speed in the hybrid powertrain controller is very straightforward when the controller is implemented as described in Section 3.4.3. When triangular fuzzy input sets, which overlap as far as their maximum grades of membership, are used with product implication, and the union of the output sets is formed using the summation operator, the overall input-output relation reduces to bilinear interpolation. The points within which the interpolation takes place are defined by the values in the input universes of discourse which have maximum grades of membership in the fuzzy input sets and by the centres of the output sets.

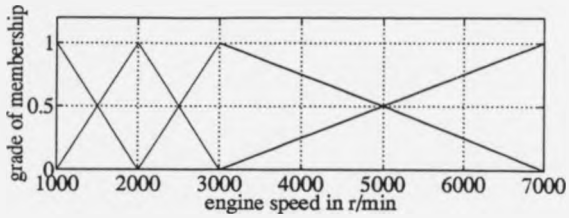

Figure 6.1: Fuzzy input sets initially used on engine speed universe of discourse.

In order to use a second input of engine speed, a series of input fuzzy sets must be defined on the engine speed universe of discourse. These sets were initially defined to have maximum grades of membership at the values 1000, 2000, 3000 and 7000 r/min and are shown in Figure 6.1. These set locations were decided upon by considering the range of speeds over which the engine usually operates.

The same pedal value sets were retained from the work described in the previous chapter. In the earlier work, the output fuzzy sets for the throttle angle and the demanded armature currents had been represented by a vector of values representing the locations of the fuzzy singleton output sets for each fuzzy rule. When generalising the method to include two input controllers, a matrix of output set singletons is required for each output. The original controller, from which the adaptation process started, was the baseline controller. It will be recalled from Section 3.5 that the baseline controller demands a throttle angle equal to the pedal value and no armature current. The matrix used for the initial locations of the output sets for the throttle angle was, therefore:

$$\begin{bmatrix} 0 & 4 & 12 & 24 & 40 & 90 \\ 0 & 4 & 12 & 24 & 40 & 90 \\ 0 & 4 & 12 & 24 & 40 & 90 \\ 0 & 4 & 12 & 24 & 40 & 90 \end{bmatrix}$$

The equivalent matrix for the demanded armature current output set locations was:

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The adaptation procedure is easily applied to the two input hybrid powertrain controller. In fact, in some respects, the implementation of the adaptation procedure is more elegant in the case of the two input controller. The vector $\mathbf{\Omega}$, used to partition the range of vehicle powertrain operating speeds can be set equal to the vector of locations with maximum grades of membership in the engine speed input sets. When the operation of the powertrain is clustered to the vertices of the performance measurement grid, each performance measure value in the performance measurement matrix then represents the performance of one rule in the fuzzy controller.

When the rules are modified, the derivatives of engine torque with respect to throttle angle, and motor output torque with respect to demanded armature current, can be evaluated at the powertrain operating speeds for each rule. Previously, average values obtained over all engine speeds had been used to determine the direction in which rules should be modified, leading to difficulties in maintaining a constant driveshaft torque as the controller adapted.

6.3 Alternative drive cycles.

The ECE-15 cycle is useful as a benchmark cycle over which to evaluate the fuel consumption and emissions performance of vehicles. However, it is not ideal as a cycle to exercise an adaptive control strategy, since the vehicle performance required to follow the cycle is not very large, few points of operation are visited on the cycle and much of the time is spent at idle. Essentially, the data contained in ECE-15 is not "rich" enough to exercise an adaptation strategy and the duration

of the cycle is rather short so that, during adaptation using successive simulation runs over the same cycle, effects due to repetition of the data occur.

6.3.1 Data acquisition.

In order to overcome these difficulties, the remainder of the adaptive controller work described here used data obtained by driving an instrumented vehicle over data acquisition test routes. During the data acquisition exercises, the engine speed, vehicle speed and time were stored using a sampling interval of about 1.1 seconds. The software and hardware used to carry out this task were not developed by the author, although use was made of previous work by the author, in the development of instrumentation for the Rover electric Metro.

The engine speed was measured by interrogating the vehicle engine management system via its serial communication channel. The vehicle speed was obtained by the use of a rotational transducer placed between the differential speedo cable drive and the speedo cable. This transducer emits 16 pulses per revolution, and the vehicle speed is calculated by measuring the time interval between each pulse.

Various drivers then drove the instrumented vehicle around a twenty minute urban journey in Coventry. An additional journey involving the use of dual carriage ways and extra urban roads was also used, but this data was not used in exercising the adaptive controllers, being less typical of general passenger car usage in Europe, see [202].

Figure 6.2 shows the measured engine speeds against the measured vehicle speeds. The theoretical "ingear" linear relationships between engine speed and vehicle speed have been shown by the straight lines. It can be seen that not all of the points lie on, or even in some cases, near the lines. These points occur during gear changes. It will be recalled, from Chapter 2, that the model requires the vehicle speed and gear selected as reference signals. Accordingly, the most

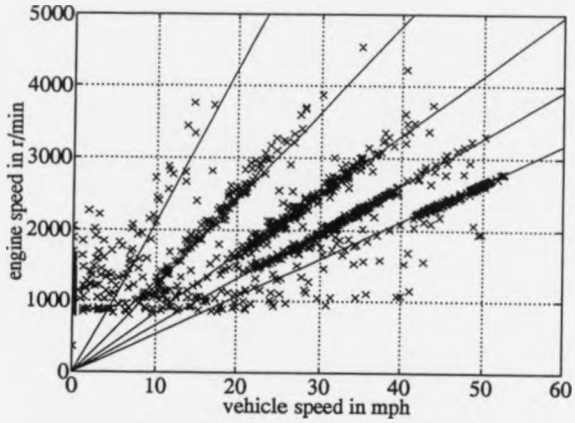

Figure 6.2: Engine speed against vehicle speed for the raw data from the measured cycle. Straight lines show the theoretical "ingear" relationship between the engine speed and the road speed.

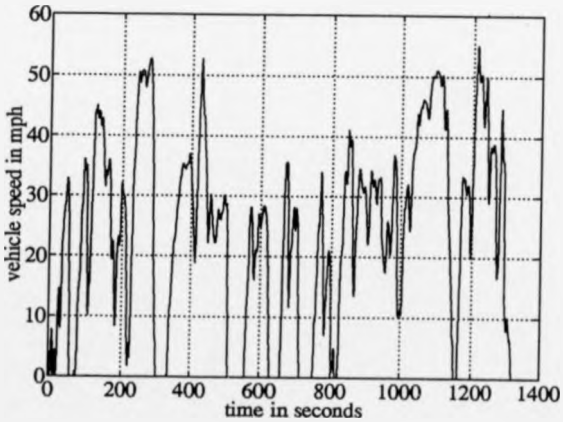

Figure 6.3: Reference speed against time for the measured urban drive cycle.

appropriate gear was calculated for each instant in the time history.

Figure 6.3 shows the vehicle speed against time for the cycle used to develop the adaptive controllers. Note the clustering of the vehicle speed values around the urban speed limits of 30 mph, 40 mph and 50 mph encountered on the journey.

6.4 Initial experiments with the adaptive two input hybrid powertrain controller.

6.4.1 Controllers using a single value of $\mathcal{F}_d$.

The migration of the fuzzy output singletons for a series of simulation runs representative of the early results is shown in Figure 6.4. The procedure for updating the rules was to use a modified version of Equation 5.19:

$$\Delta_{mv_{i,j}} = k_p(\mathcal{F}_d - \mathcal{F}_{i,j}) \tag{6.1}$$

where $\mathcal{F}_d$ = -0.08, k_p = 50 and the adaptation interval, T_a = 200 seconds. As in the adaptive single input controllers, if a rule has not been fired in the last adaptation interval, then it is not modified. Also, a mask matrix the same size as the matrices of rule output set locations consisting of the values 1 or 0 was included. The purpose of the mask was to fix the rules whose locations in the matrices of output singletons corresponded to the value 0 in the mask matrix. By this means, the rules at 0° and 90° of pedal value were fixed, whilst the other rules were allowed to be modified.

The simulations were run repeatedly over the cycle shown in Figure 6.3, the results shown in Figure 6.4, representing around 7000 seconds of vehicle usage. Four main points arise out of the work that Figure 6.4 represents:

- After a period of adaptation, many of the rule output set locations converge to final values where they remain. In some cases, some oscillatory behaviour occurred, for example, two distinct clusters of values are present

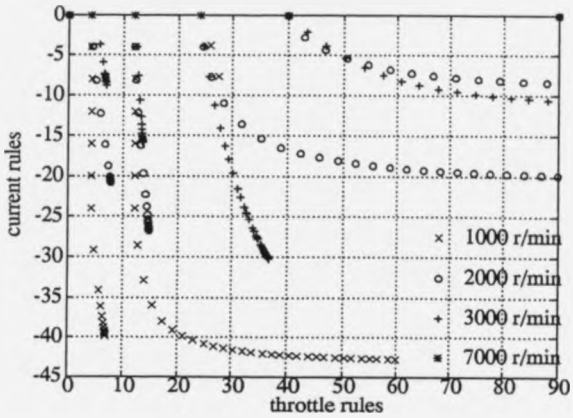

Figure 6.4: Progression of output set locations as the two input adaptive controller adapts.

in the locations of the rule whose output sets were originally located at 24° of throttle angle and 0A, at 3000 r/min. More sophisticated adaptation methods described later in this chapter were able to reduce the extent of such oscillations. However, when simulations over the same cycle are repeatedly used, in practice, it is extremely difficult to remove some form of cyclic variation. This is because there will, inevitably, be regions in any real cycle, in which individual rules perform well and other regions in which they perform badly.

- Some of the rules do not converge at all. The most obvious examples of this are the rules whose input sets are located at 3000 r/min, 12° of pedal value, 2000 r/min and 24° of pedal value, 2000 r/min and 40° of pedal value and 3000 r/min and 40° of pedal value. The reason why these rules do not converge is shown in Figures 6.5, 6.6 and 6.7, in which contours of constant powertrain output torque (calculated using the regressed engine map and

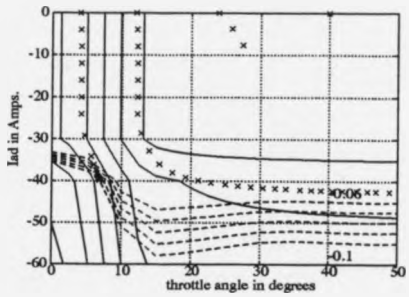

Figure 6.5: Contours of constant powertrain output torque in the throttle angle-demanded armature current space at a powertrain speed of 1000 r/min. Dashed lines are contours of constant performance measure value. Crosses show positions of fuzzy rule output sets as adaptation takes place.

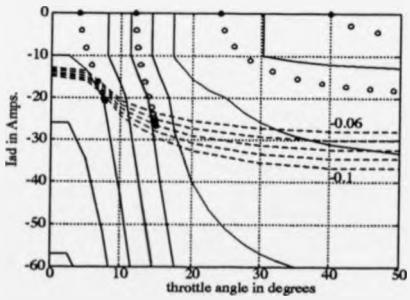

Figure 6.6: Contours of constant powertrain output torque at 2000 r/min.

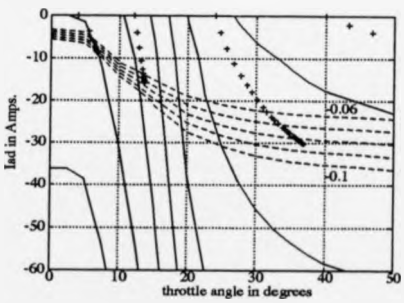

Figure 6.7: Contours of constant powertrain output torque at 3000 r/min.

the motor model) are shown as solid lines. The dashed lines are contours of performance measure ratio (calculated using the raw engine data and the motor model). The reason that some of the locations of the output sets do not converge is that the path along which these locations move does not cross the contour representing performance measure values equal to -0.08. Figure 6.7, also shows that the rate at which output set locations converge is strongly affected by the angle between the path along which the set locations move and the contours of constant performance measure.

- The output sets that were originally located at 4°, 0A do not converge exactly to a performance measure value of -0.08. The reason for this is that the engine changes most rapidly in this region of operation. It is, therefore, most difficult for the clustering algorithm to represent the action of the entire rule operation. Also, the contours of constant performance measure are most closely spaced in this area. In Figure 6.7 the locations of the output sets whose corresponding input sets are at 3000 r/min and 24° can be seen to oscillate either side of the performance measure = -0.08 contour.

- In Figure 6.4 the very different paths of the output sets originally at 24° and 2000 r/min and 3000 r/min are obvious. Figures 6.6 and 6.7 show that this is because the contours of constant powertrain output torque are quite divergent at these two powertrain speeds.

6.4.2 Controllers using different values of $\mathcal{F}_d$ for each rule.

In order to overcome the problem caused by the lack of convergence of the output set locations for some rules, it was decided to have a matrix, **H**, the same size as the matrices of controller output set offsets whose elements were the values $\mathcal{F}_d$ for each fuzzy rule. This allows much greater control over the converged rule output set locations. As before, the output sets were not moved if the rule had not been

fired or if they were fixed by the mask.

More negative values of $\mathcal{F}_d$ can be used for those rules with output set locations which can be moved a long way and still remain in a region in which the performance measure changes as the sets move. However, using a simple proportional control to adjust the location of the output sets will cause much larger steps to be taken when the performance of the rule is very different from the $\mathcal{F}_d$ value for that rule. When these larger steps were taken it became apparent that the simple Euler integration method of rule movement was inadequate and could cause the rules to move in quite the wrong direction. To overcome this problem, second order Runge-Kutta integration, as described in [49], was used.

The rule modification is performed as follows. The distance which the rule is to be moved is still calculated using Equation 6.1. Intermediate output set locations are first calculated, using Equations 5.5 and 5.6, setting $\Delta_{mv_{i,j}}$ equal to half the value given by Equation 6.1 and the values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ at the current location of the rule output sets. New values of $\frac{\partial T_e}{\partial thr}$ and $\frac{\partial T_g}{\partial Iad}$ are then calculated at the intermediate location of the output sets and these values, and the full value for $\Delta_{mv_{i,j}}$ given by Equation 6.1, are then used to move the output set locations from the current locations of the output sets (not the intermediate locations) to the next locations of the output sets. This procedure caused the paths taken by the output sets to more closely follow the contours of constant output torque, and also allowed the use of much larger step sizes in the adaptation process.

By adjusting the values stored in the **H** matrix, it is possible to have far greater control over the movement of the output sets during adaptation, and the values to which they converge. By plotting combined engine and motor output torque against pedal value, for both the baseline controller and the adapted controller, it is possible to study how well the adapted controllers maintain a constant powertrain output torque after adaptation.

It was quickly discovered that it is extremely difficult for an adapted controller to maintain the driveshaft torque of the baseline controller at 1000 r/min. This is because the operation of both the engine and the motor varies in an extreme way at such low speeds. At 1000 r/min the engine is very responsive to changes in the throttle angle. Also, the consequences of preventing motor currents from flowing, when the demanded armature current is negative but the battery current would be positive, are at their most severe at low speeds. It was, therefore, decided to fix the rules at 1000 r/min using the mask matrix and add an additional rule at 1500 r/min. This is a useful addition, since, from the point of view of the engine it is desirable that, as the engine comes off idle, it does not have the motor acting against it.

It was also found that the fuzzy rules whose input sets were located at 4° of pedal value caused significant problems in driveability. This is because of the sudden step of negative torque, which occurs as the conditions for motor currents to be applied are met. At higher engine speeds, this step of negative torque tended to occur as the engine itself was producing negative overrun torque, leading to some very negative powertrain output torque values. The contours of constant powertrain output torque are very closely spaced at low pedal values and this further hinders the maintenance of good driveability. To overcome these problems, the rules located at 4° were also fixed using the mask. Again, from a practical point of view, pedal values as low as 4° will only tend to be used as the engine comes off idle, and it is desirable to restrict the operation of the motor under these conditions.

During the controller development it was often found that the combined engine and motor output torque matched the torque output of the baseline controller very well, when the inputs were at the maximum grades of membership of the input fuzzy sets, but less well between the rules. This effect was reduced by adding two

more fuzzy input sets on the pedal universe of discourse. Figure 6.8 shows the progression of the output set locations, during adaptation, of the controller whose input sets and **H** matrix are shown below:

$$
\begin{array}{cc}
 & \text{pedal value input set locations} \\
 & \begin{bmatrix} 0 & 4 & 8 & 12 & 18 & 24 & 40 & 90 \end{bmatrix} \\
\begin{array}{r} \text{engine} \\ \text{speed} \\ \text{input} \\ \text{set lo-} \\ \text{cations} \end{array}
\begin{bmatrix} 1000 \\ 1500 \\ 2000 \\ 3000 \\ 7000 \end{bmatrix}
& \underbrace{\begin{bmatrix}
0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.00 & 0.00 & -0.25 & -0.15 & -0.05 & -0.02 & 0.00 & 0.00 \\
0.00 & 0.00 & -0.40 & -0.30 & -0.05 & -0.02 & 0.00 & 0.00 \\
0.00 & 0.00 & -0.60 & -0.50 & -0.25 & -0.08 & -0.01 & 0.00 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00
\end{bmatrix}}_{\mathbf{H}}
\end{array}
$$

The other simulation parameters that were used in obtaining the results shown in Figure 6.8 were $k_p = 50$ and $T_a = 50$ seconds. The simulation was stopped after 50 rule modifications. The output sets for the rules at 7000 r/min were only very slightly modified since they are fired very infrequently in the urban cycle used to exercise the controller. A feature to note from the upper two plots of Figure 6.8 is that few modifications were made to the output set locations of the rules whose input sets are at 1500 r/min and 24°, and 2000 r/min and 40°. This is because at low engine speeds, low pedal values are used predominantly and, since the rules are not fired, the adaptation scheme does not move the output sets. In use, the consequences of the rules overtaking one another are that, as the pedal value increases, the engine throttle angle decreases. However, this is not as much of a problem as it might at first seem, because, at the throttle angles used under these conditions, the rate of change of engine torque with respect to the throttle angle is very low. Also, at about this stage, the motor stops applying negative torque and the change in the motor torque is the most dominant effect. This is shown by the plots comparing baseline controller powertrain output torque and adapted controller powertrain output torque, in Figure 6.9.

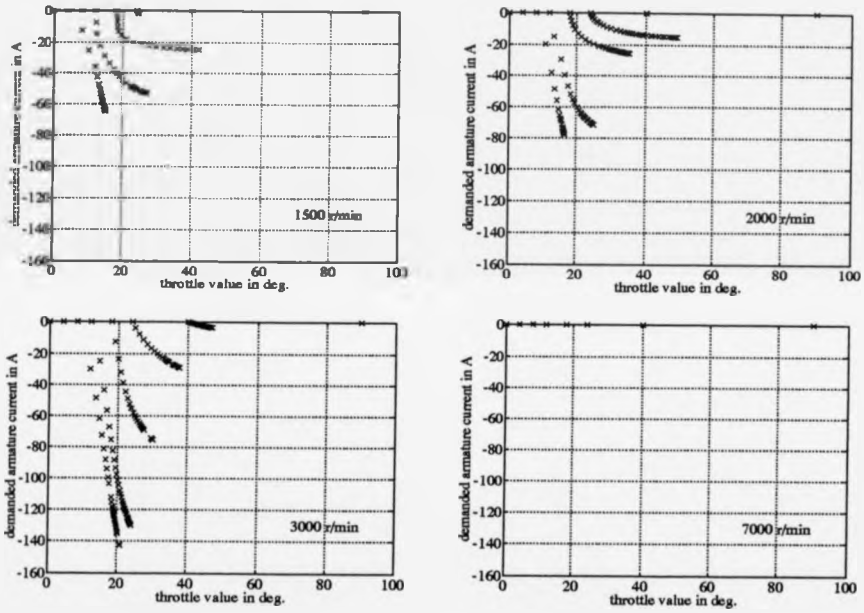

Figure 6.8: Progression of controller output set locations using two input adaptive hybrid vehicle powertrain controller.

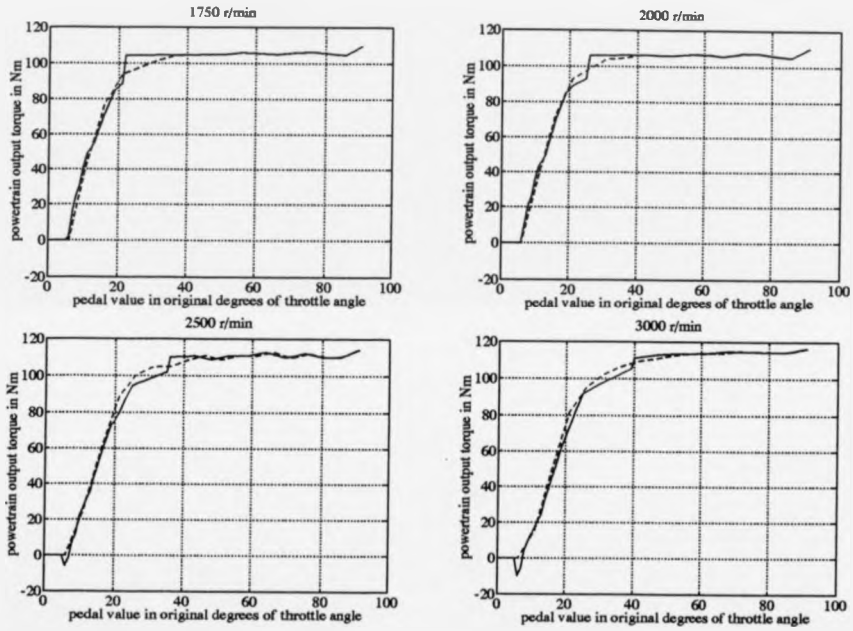

Figure 6.9: Plots comparing baseline controller powertrain output torque (dashed line) and adapted controller powertrain output torque (solid line) against pedal value.

In general, the torque responses of the adapted controller compare very well with the baseline controller torque responses, particularly at the points of maximum grade of membership of the input sets, as the lower two plots show. It should also be noted that the responses are also good between the engine speed input fuzzy sets, as the plots on the left of Figure 6.9 show.

The use of individual values of $\mathcal{F}_d$ allow the output sets to be placed by hand. However, the method has significant disadvantages. Firstly, the controllers have still not converged. This is shown by the fact that some output sets appear to have reached an appropriate position, but this was achieved by setting the $\mathcal{F}_d$ values such that, over the period of time that the adaptation took place, the sets reached these positions. Had a different number of rule modifications been made, they would have reached different positions.

The method also suffers from the disadvantage that, if the rule is not fired, its output sets are not modified, causing vastly different amounts of output set modification for different pedal values at the same speed. Additionally, this method of controller adaption requires a large number of values to be set, all of which need to be changed if more rules are added. Lastly, this method of setting the values gives no indication of how the vehicle will perform overall. The adaptation procedure described in the next section was developed to overcome these difficulties.

6.5 A controller adaptation procedure driven by the overall performance of the vehicle.

The controller adaptation procedure described in this section modifies the output sets of the rules of the controller by using the error between the performance of the vehicle, and the desired performance of the vehicle, in the controller adaptation. The modification of the output sets of each individual rule is then affected by the performance of the rule set rather than that individual rule.

6.5.1 Defining different types of fuzzy rules.

When the controllers were modified by setting different values in the **H** matrix it was realised that there were essentially two types of rule. The rules at relatively low throttle values tended to be able to affect the overall performance of the vehicle strongly because they could be moved to regions of operation where the performance measures had relatively negative values. Also, these rules were fired relatively often and so their performance has a strong effect on the overall effect on the vehicle. In contrast, it has been seen that the rules at higher pedal values are used relatively infrequently and, because of the path that the output sets take when being modified, do not achieve particularly negative values of the first performance measure. However, the rules with input fuzzy sets at higher pedal values need to be placed such that the controller driveability is maintained. This leads to a conceptual distinction between the two types of rules, those that should drive the rule modification process, and those that follow the modification process. The primary distinction between these two types is that the first type of rule is moving in a gradient of the performance measure and the second type of rule is not moving in a gradient of performance measure.

The controller modification strategy should then be to modify the rule output sets moving in a performance measure gradient, using the error between the vehicle performance and the desired vehicle performance. The rules whose output sets are not moving in a performance measure gradient should be modified according to the overall shape formed by the complete set of rules.

In order to implement this rule modification strategy, the gradient of the performance measure, in the direction in which the rule is moved, has to be evaluated. To do this, the performance measures for the controller rules are evaluated over an adaptation interval, T_a. The performance measure gradient for a particular rule is obtained by dividing the difference between two performance measure values,

for that rule, obtained over consecutive adaptation intervals, by the step size of the output set modification taken between the intervals.

The evaluation of the performance measures is affected by the actions of the clustering algorithm, which estimates the performance of the vehicle in the region of operation of each rule. The estimates provided by the clustering algorithm will have some statistical variation, because, in any adaptation interval, the operation of the powertrain is very unlikely to be spread uniformly over the region of operation of a particular rule. Generally, in any adaptation interval, the more time that the powertrain spends operating in the region of a particular rule, the more accurately the performance measure will reflect the true performance of that rule. Figure 6.10 shows the percentage accuracy of the performance measure estimates for a particular rule, plotted against the number of times that the powertrain operation was clustered to that rule. During the simulations used to generate these results, a fixed controller, which had been generated over earlier simulations, was used and was not allowed to adapt. The "correct" value of the performance measure, used to calculate the percentage error, was calculated by taking the average of all values of that performance measure during the simulation. Different values of T_a were used to obtain a range of maximum values of the count of the number of times that the performance was clustered towards a particular rule.

Figure 6.10 shows the results for the rule whose input sets had maximum grades of membership at 2000 r/min and 8° of pedal value, which is one of the more frequently used rules. If the region, in which the output sets of a rule in an adaptive fuzzy controller are moving, has a small gradient then, when the output sets are modified, the change in the true performance measure of the rule will quite small. When the controller is close to converging, small changes in the rule output sets will be made causing small changes in the true value of the performance measure. These small changes would be swamped by the statistical variations

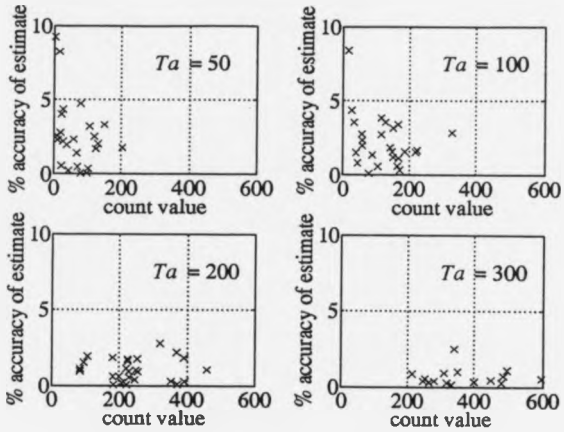

Figure 6.10: Plots of the performance measure values for the rule with fuzzy input sets at 2000 r/min and 8° of pedal value against the count of the number of times that the performance was clustered towards that rule.

between one performance measure estimate and the next, making the gradient estimate based upon the difference between successive values very inaccurate. As can be seen, the use of longer adaptation intervals tends to reduce this effect, but will cause the adaptation of the controller to be slow.

At this stage, the rules that drive the modification process were still only modified if they had been selected by the clustering algorithm during an adaptation interval. It was noticed that, for the most part, the same rules were fired as were moving in a gradient of performance measure. The calculation of the gradient of the performance measure involves the use of two additional matrices to store the step size and the gradient, and extra computations. Since the gradient calculated has an accuracy that can only be quantified at further computational and storage expense, a rule was treated as a driver, rather than a follower, if it had been selected by the clustering algorithm, at some stage in the adaptation interval, and the estimation of the gradient was dropped.

6.5.2 Driving rule and following rule modification schemes.

Having established that a particular rule should have the effect of driving the rule modification process, the location of the output sets is then modified by the use of a simple proportional law, similar to that used in the adaptive controllers described above. In this case, the error between the current performance of the vehicle and the desired performance of the vehicle is used. This raises the issue of measuring the current performance of the vehicle. The instantaneous performance of the vehicle, as measured by the first performance measure given in Equation 5.18, is shown in the upper plot of Figure 6.11 for a vehicle with a fixed controller being driven over the urban cycle. The lower plot shows two moving average estimates of the performance of the vehicle, calculated using the following formula:

$$\overline{\mathcal{F}}_n = \overline{\mathcal{F}}_{n-1}\frac{(\nu - 1)}{\nu} + \mathcal{F}\frac{1}{\nu} \tag{6.2}$$

where $\overline{\mathcal{F}}_n$ is the moving average value of the instantaneous performance measure, $\mathcal{F}$, at time index n. The parameter ν can be used to adjust the rate at which the average is influenced by new information. In the lower plot in Figure 6.11 the solid line represents a moving average obtained using $\nu = 5000$ and the dashed line represents a moving average obtained using $\nu = 2500$. The values of ν are very high because the moving average is updated every 100 ms and the instantaneous measure has such a large amount of variation. The performance measure value used to adjust the driving rules of the controller is calculated using Equation 6.2.

The modification of the rules that are following the adaptation process, rather than driving it, is carried out separately for each row of the matrices of output set locations. Within each row, corresponding to a particular powertrain speed, the relationships between the pedal value and the throttle angle, and the pedal value and the demanded armature current are then adjusted. The shape formed by the output set values of the rules following the adaptation process should be

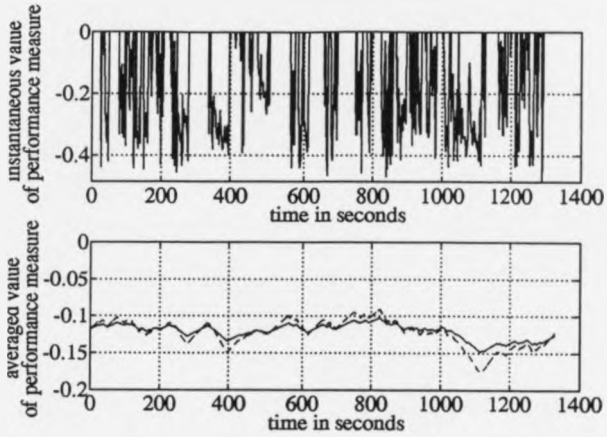

Figure 6.11: Upper plot shows an example plot of the instantaneous values of the first performance measure for a particular controller, whose output sets were fixed. Lower plot shows two moving average estimates of the same performance measure (solid line $\nu = 5000$, dashed line $\nu = 2500$). Note change in axis scales.

such that the powertrain output torque is maintained and the values of throttle, and demanded armature current, vary smoothly as the pedal value changes.

During any adaptation interval none, some, or all of the rules, at a particular engine speed, may be used. It is often the case that the rules are used over ranges or, more frequently, a range of pedal values which does not span the entire universe of discourse of pedal values. In these cases, the rules which have been used are adjusted by the proportional law used previously. The rules which have not been adjusted in this manner normally, also, lie over a range of pedal values.

Figure 6.12 shows example rule output set locations after the driving rules have been modified using the proportional law. The output sets of the rules which have been modified in this way, or which have been fixed by the mask matrix, are shown with crosses. The circles indicate the locations of the output sets which have not yet been modified, which will follow the movement of the

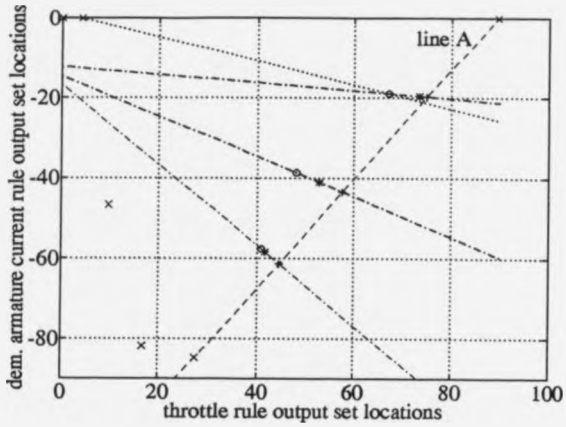

Figure 6.12: Diagram showing the modification of rule output set locations that follow, rather than drive, the rule modification process.

driving rules. The dash-dotted lines show the directions in which the unmodified output sets are constrained to move in order to maintain a constant powertrain output torque. An obvious method of modifying the output sets following the adaptation process would be to apply a modification based on the modification that was applied to the driving rule output sets. However, this would not take into account the differing directions in which the output sets have to move in order to maintain the output torque and the effect these different directions will have on the overall shape formed by the rules in this row. Clearly, if the output sets, at about 70° of throttle angle and -20A demanded armature current, were moved very much further, the throttle output set location would go past 90°, which is very undesirable.

The dashed line, marked line A, in Figure 6.12 links the output sets one rule to either side of a range of sets that have yet to be modified. Since the mask always fixes the sets at the ends of the pedal value universe of discourse this line

is always defined. The rules which follow the adaption process are adjusted along the dash-dot lines towards the points marked "+" where the dash-dot and dashed line intersect. The points to which the output sets are moved are calculated by a weighted average of the current locations of the output set and the intersection point. The weights used in the average are calculated by a small fuzzy controller. The linguistic statement that originally described the actions of the fuzzy controller was "If your confidence in the current locations, of the output sets of the rule being considered, is high, then leave the sets where they are, otherwise move them towards the intersection point".

This statement was used because occasionally the rules at a particular engine speed region of operation would be fired in an unusual manner. If, during a period of adaptation, the low pedal value rules are not used, but the high pedal value rules are used, then the line towards which the sets would be moved might, for example, be as shown by the dotted line in Figure 6.12. If the other sets were moved back towards this line, the effect on the overall shape of the output sets would be disastrous. The purpose of the above linguistic statement is to limit the movement of the rules to reflect the general operation of the powertrain, rather than just the operation over the last adaptation interval. For this reason, a measure of the confidence that might be placed in the location of a rule would be the length of time since that rule was last used. If the rule was recently used, the confidence in its current location should be high, whereas, if the rule has not been used for some time, the confidence in its location is low. This measure of confidence is appropriate because, if the rule has been used, then its output sets will have been adjusted by the proportional law and, consequently, its location would have, to a certain extent, been influenced by the recent performance of the vehicle.

Fuzzy sets were then defined over the universes of discourse "number of adapta-

tion intervals since last modification as a driving rule" and "confidence in current rule location". Given the simple nature of the controller, only two rules were used. Since the fuzzy controller was implemented in the same manner as the main powertrain controller, it was reduced to two vectors of length 2 within which linear interpolation was carried out. The first fuzzy input set has its maximum grade of membership at 1. The position of the maximum grade of membership of the second fuzzy input was seen as a tunable parameter, known as n_{ad}. The output set locations were at 1 and 0. The overall effect of this was that, if only one adaptation interval had passed since the rule was modified as a driver, the confidence value was 1 and the rule was not altered. If the number of adaptation intervals since the rule was last modified was greater than or equal to the position of the maximum grade of membership of the second fuzzy set, then the rule output sets were moved immediately to the intersection. In practice, this second case is very extreme, since $n_{ad} = 20$ adaptation intervals was eventually used. The main effect of the location of the second fuzzy set is to modify the behaviour of the adaptation scheme for rules that have not been modified as driving rules for perhaps five or six adaptation intervals. This illustrates the importance of considering the cumulative effect of the rule modification scheme, rather than the exact action that occurs at any individual stage in the modification process.

One further addition to the modification process was required at this stage. It is possible that the line linking the ends of a series of output sets to be modified can be almost parallel to the direction in which the rule output sets are moved. In this case, the location of the intersection, can be a great distance from the current locations of the output sets. In order to prevent rules from being moved too far towards this intersection, a limit was set on the maximum step size of any output set movements. This limit was normally set at 5.

Figure 6.13 shows the progression of the windowed average performance mea-

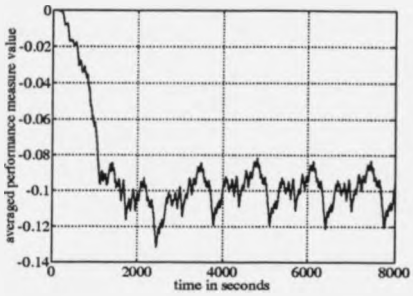

Figure 6.13: $\overline{\mathcal{F}}$ against time with $k_p = 50$, $\nu = 5000$ and $n_{ad} = 20$.

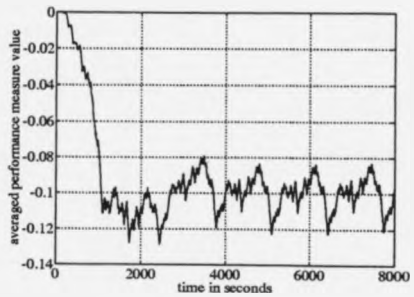

Figure 6.14: $\overline{\mathcal{F}}$ against time with $k_p = 100$, $\nu = 5000$ and $n_{ad} = 20$.

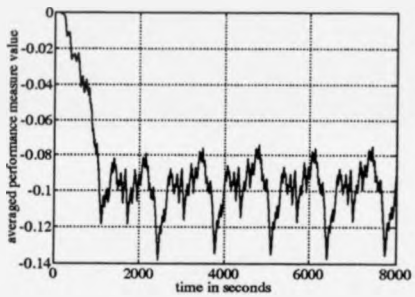

Figure 6.15: $\overline{\mathcal{F}}$ against time with $k_p = 50$, $\nu = 3000$ and $n_{ad} = 20$.

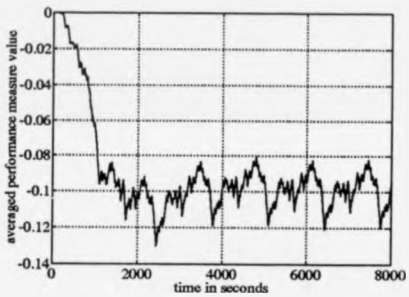

Figure 6.16: $\overline{\mathcal{F}}$ against time with $k_p = 50$, $\nu = 5000$ and $n_{ad} = 10$.

sure towards a demanded value of -0.1. For the purposes of comparison, Figures 6.14, 6.15 and 6.15 show similar responses in which the principle adaptation parameters have been changed. The simulation exercises were performed by repeatedly using the urban cycle shown in Figure 6.3.

Figure 6.17 shows the locations of the output set offsets of the controller as it was at the end of the simulation whose performance measure time history is shown in Figure 6.13. The controller output set locations have been linked with lines to emphasise the shape formed by the adaptation process.

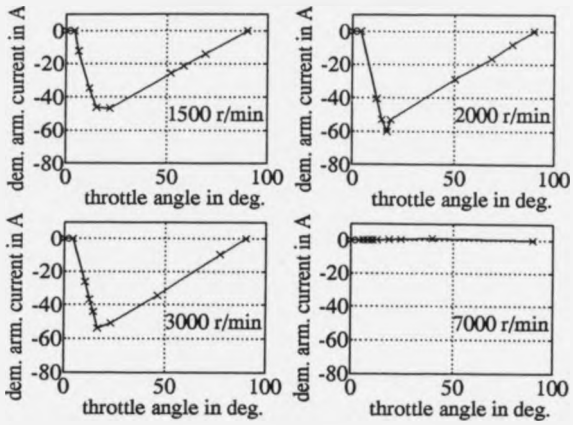

Figure 6.17: Output set locations for controller whose adaptation parameters were $k_p = 100$, $\nu = 5000$ and $n_{ad} = 20$.

A comparison of Figures 6.14, 6.15 and 6.16 with Figure 6.13 illustrates the effect that the adaptation parameters have on the performance measure progression. Firstly, it should be noted that, using values of ν in the order of a few thousand, a poor compromise is reached between controller responsiveness and reducing the periodic variations in the performance measure, caused by operating a vehicle over different regions of a cycle. It is seen that increasing k_p, which is effectively the gain on the adaptation process, increases the overshoot and cyclic variations. The cyclic variations are increased because of the lag between changes to the controller and resulting changes in the performance measure, due to the actions of the moving average. Clearly, the greater these variations are, the more the controller will be adjusted over the course of an individual cycle. Reducing ν causes the performance measure to become far more responsive and, again, this causes much higher cyclic variations. Changes in n_{ad} make relatively little difference since the primary function of this part of the adaptation procedure is to

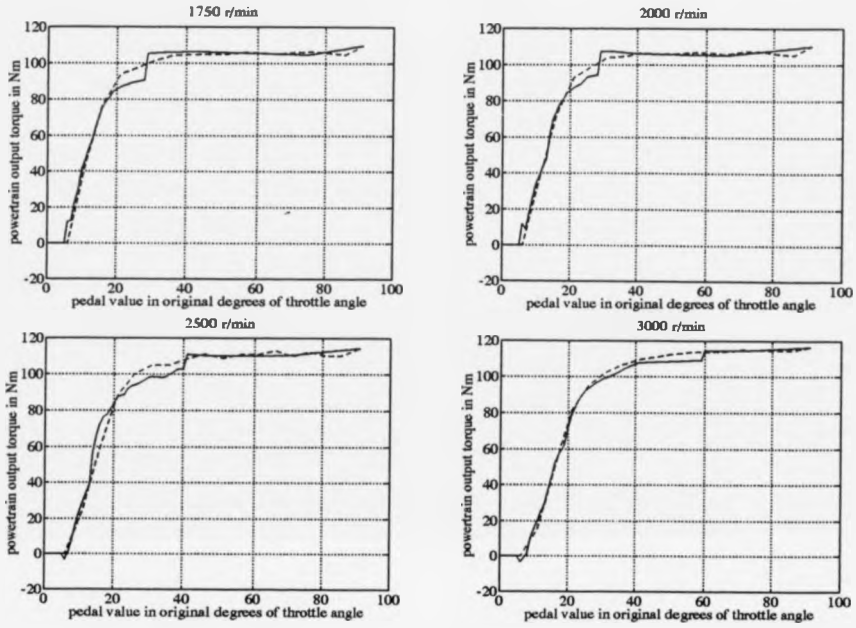

Figure 6.18: Plots of baseline controller output torque and adapted controller output torque against pedal value. The adaptation parameters were $k_p = 100$, $\nu = 5000$ and $n_{ad} = 20$.

prevent rare, and disastrous, rule modifications. In some circumstances, reducing n_{ad} can have a similar effect to increasing k_p, since it allows the output sets following the adaptation process to move more quickly.

The effects of the three principle adaptation parameters can be summarised as follows: ν is a measure of the ability of the controller to recognise new performance data, k_p is a measure of the controller's desire to improve its performance and n_{ad} is a measure of the reluctance of the controller to forget old controller values.

Figure 6.17 shows the locations of the output sets generated by the simulations whose performance measure values are shown in Figure 6.14. It can be seen that the adaptation process has moved the output sets for the rules whose input sets are

at 7000 r/min in the opposite direction to the rules at other speeds. This is simply because these rules are chosen by the clustering algorithm very infrequently. On the occasions when they are chosen, the performance of the vehicle is such that the rules are moved in the opposite direction from those at other speeds.

Figure 6.18 compares the torque responses of the baseline controller and the adaptive controller whose output set locations are shown in Figure 6.17. The responses are reasonable and would not, at first sight, appear to cause many problems in the operation of the vehicle. However, the output torque is not maintained particularly accurately and there are small spikes in the responses, most noticeably at powertrain speeds between the powertrain speed input fuzzy sets. Theses spikes are caused by the angular nature of the shapes formed by the fuzzy output sets. They do not appear to be very significant in these plots, because the match between the engine, and motor, maps and the component models is quite accurate. However, on real vehicles, they are potentially more significant due to variations between components due to ageing and manufacturing tolerances. Output sets causing smoother shapes would be less likely to cause problems in the torque responses.

The modifications to the adaptation process described in the next section attempt to reduce the cyclic variation of the performance measures, obtaining controllers which converge rather than exhibit cyclic variations. The adaptation process will also be modified so that output sets are placed at locations which form a smoother curve in the throttle angle-demanded armature current output space.

6.6 Varying the moving average measure of performance and producing smoother controllers.

6.6.1 Automatically controlling ν.

In order to improve the compromise between the responsiveness of the controller and its ability to converge, the value of ν was brought under the control of an algorithm. This algorithm set low values for ν, when the performance of the vehicle was some distance from the desired value, and much larger values, when the performance was closer to the desired value.

A series of five decreasing threshold values, $\epsilon_{1,\ldots,5}$ are defined. These values are used to represent proportions of the value $\mathcal{F}_d$. When the difference between the moving average of the performance measure, and demanded value of the performance measure, is less than these proportions of $\mathcal{F}_d$, different values of the parameter ν are used, these values being stored in a vector $\nu_{1,\ldots,5}$. As the threshold values become smaller, the corresponding ν values increase. The motivation behind this is that, as the performance measure value is close to the desired value, then the requirement for it to be responsive is reduced, but the requirement for it to represent the performance of the vehicle over longer time scales is increased. Let i be the current value of the indexing variable used to address the elements of the vectors $\epsilon_{1,\ldots,5}$ and $\nu_{1,\ldots,5}$. If the condition

$$|\mathcal{F}_d - \overline{\mathcal{F}}| \leq \epsilon_{i+1}\mathcal{F}_d \tag{6.3}$$

is met for a continuous period of time T_{av}, then the indexing variable, i is incremented. If the condition

$$|\mathcal{F}_d - \overline{\mathcal{F}}| > \epsilon_i\mathcal{F}_d \tag{6.4}$$

is ever met, the indexing variable is decremented. These changes are made within the restriction $1 \leq i \leq 5$ and are tested every control interval (0.1 sec).

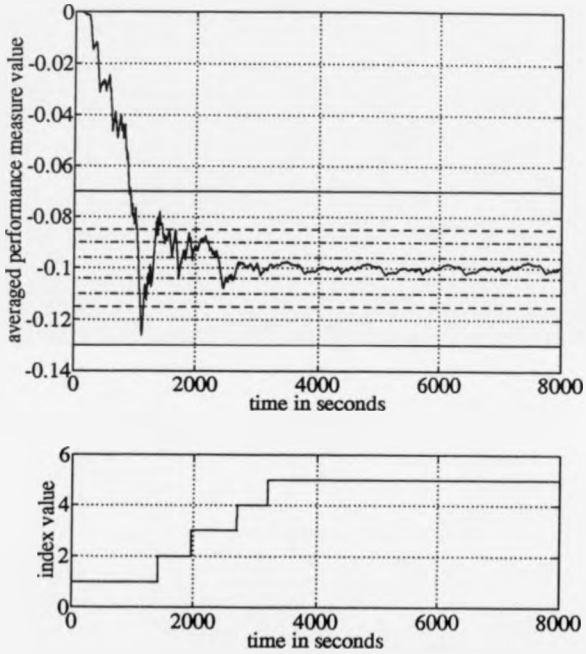

Figure 6.19: Upper plot shows the value of the moving average performance measure against time when the value of ν is varied throughout the simulation dependant upon the error between the actual vehicle performance and the desired performance. Lower plot shows the value of the indexing variable used to address elements in the vector $\nu_{1,\cdots,5}$ of ν values. $\mathcal{F}_d$ = -0.1.

The upper plot in Figure 6.19 shows the value of the moving average performance measure against time when the vectors $\epsilon_{1,\cdots,5}$ and $\nu_{1,\cdots,5}$ have the following values:

$$\begin{aligned} \nu_{1,\cdots,5} &= [\ 2500 \quad 5000 \quad 10000 \quad 20000 \quad 40000\ ] \\ \epsilon_{1,\cdots,5} &= [\ 1 \quad 0.3 \quad 0.15 \quad 0.1 \quad 0.04\ ] \end{aligned}$$

The parameter T_{av} was set to 500 seconds and $\mathcal{F}_d = -0.1$. The horizontal lines on the plots are drawn at the values, $(1 \pm \epsilon_i)\mathcal{F}_d$, of performance measure. The lower plot in Figure 6.19 shows the value of the indexing variable used to address the elements of the vectors $\epsilon_{1,\cdots,5}$ and $\nu_{1,\cdots,5}$ against time. By successively increasing the value of ν, the average performance measure can be made to follow the desired value with decreasing cyclic variation. This reduction in the variation of the performance measure has the desirable consequence of reducing the variation in the locations of the controller output sets, as shown in Figure 6.20.

A further example of the operation of the algorithm that adjusts the value of ν is shown in Figures 6.21 and 6.22, in which the same controller parameters were used, but with a value of $\mathcal{F}_d$ = -0.14.

Figure 6.23 shows the output set locations produced by the controller whose performance measure progression is shown in Figure 6.19. It is seen that the shapes formed by the set locations are similar to those without the adjustment of ν and are, therefore, still rather angular. Since the driving sets are chosen by the statistical variation in the use of the pedal value input sets, it was felt that increasing the number of pedal value input sets might improve the smoothness of the shape formed by the output set locations. In fact, one of the output sets was still used very much more than its nearest neighbours, but, because its neighbours were now closer to it, the overall shape formed by the output set locations was worse than the shape produced by the controller with fewer rules. The next task in the development of the adaptive controller was to generate smoother output set locations.

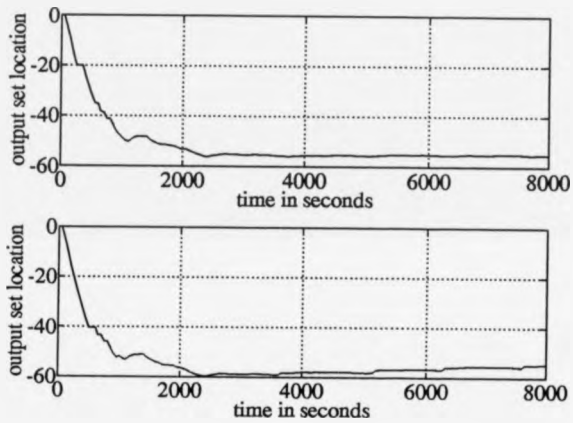

Figure 6.20: Upper plot shows the progression of the demanded armature current output fuzzy set over the duration of the simulation for the rule whose input sets have maximum grades of membership at 2000 r/min and 12° of pedal value. Lower plot shows the same information for the rule whose input sets are at 3000 r/min and 12°. $\mathcal{F}_d$ = -0.1.

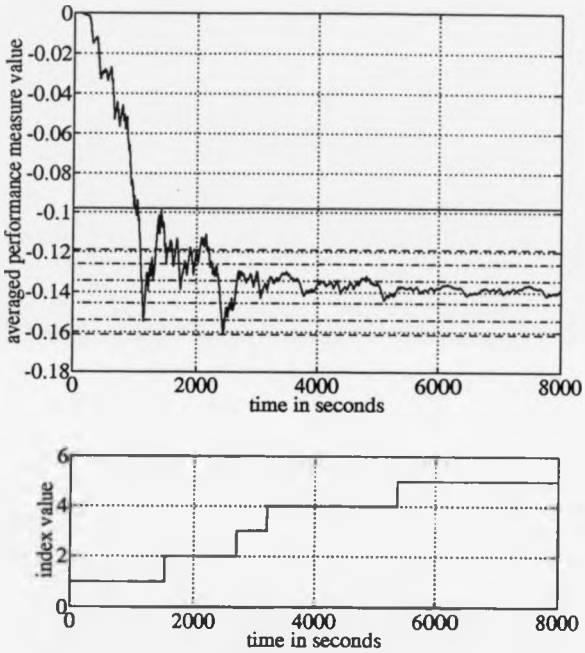

Figure 6.21: Upper plot shows the value of the moving average performance measure against time when the value of ν is varied throughout the simulation dependant upon the error between the actual vehicle performance and the desired performance. Lower plot shows the value of the indexing variable used to address elements in the vector $\nu_{1,\ldots,5}$ of ν values. $\mathcal{F}_d$ = -0.14.

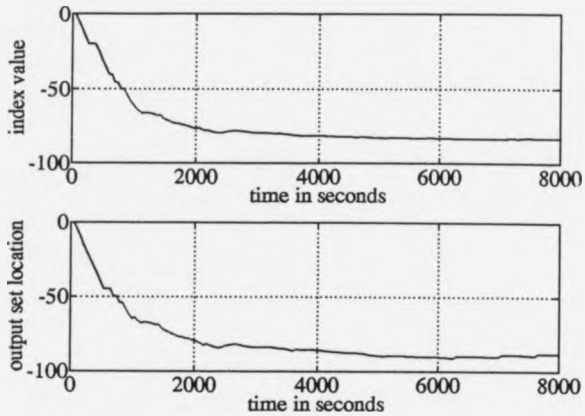

Figure 6.22: Upper plot shows the progression of the demanded armature current output fuzzy set over the duration of the simulation for the rule whose input sets have maximum grades of membership at 2000 r/min and 12° of pedal value. Lower plot shows the same information for the rule whose input sets are at 3000 r/min and 12°. $\mathcal{F}_d$ = -0.14.

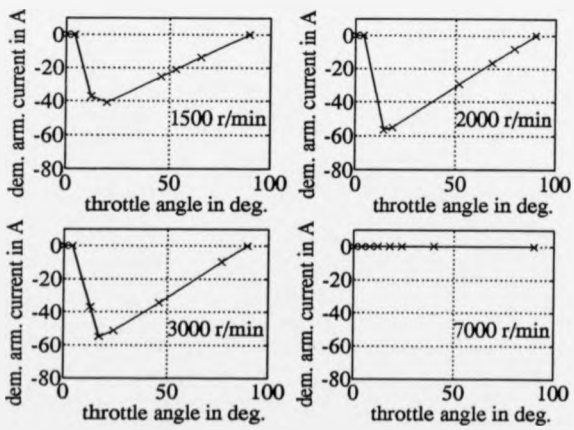

Figure 6.23: Final output set locations for the controller whose values of ν were varied throughout the simulation dependant upon the error between the actual vehicle performance and the desired performance.

6.6.2 Adjusting the movement of the driving rules to produce smoother output set locations.

The movement of the rules that drive the modification process is affected by two factors, the error between the vehicle performance and the desired vehicle performance, and the usage of the rules. It is the relative amounts of rule usage that determine the shape that the output set locations form, and it was seen, in the previous section, that the manner in which the driving rules were used did not cause the output set locations to form a smooth curve. It is, therefore, necessary to modify the manner in which the rules are moved when they are driving the rule modification process, to take account of the shape which they form.

In order to get controllers whose output sets form a smooth curve, some consideration of the angle caused by the current location of the rule output sets is needed. If the output sets in Figure 6.23 are considered, it is seen that the controller output shapes would be smoother if, where the angles at the output sets are relatively large, the output sets were moved outwards. Conversely, where the output sets form small angles, either they should be moved inwards or their outward movement should be restricted.

Consider the angle formed by the location of the output sets of a particular rule. Two vectors, **a** and **b**, can then be defined as shown in Figure 6.24. The cosine of the angle formed at the output sets of the rule being considered can then be found by the use of the scalar product of the the vectors **a** and **b**.

The value of $\cos\psi$ and the error value $\mathcal{F}_d - \overline{\mathcal{F}}$ can then be used as the inputs to a two input fuzzy controller whose output is then used in the modification step size calculation instead of the term $(\mathcal{F}_d - \overline{\mathcal{F}})$. Let the modified error output of the two input fuzzy controller, for the rule whose output sets are on row i, column j, of the output set matrices, be $e_{i,j}$. The output set modification step size is then

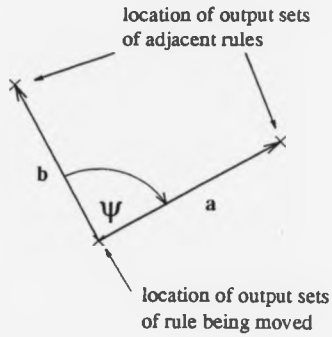

Figure 6.24: Calculation of the angle formed by the location of the output sets of a particular rule.

calculated:

$$\Delta_{mv_{i,j}} = k_p e_{i,j} \tag{6.5}$$

The smoothing fuzzy controller uses four input sets on the $\cos\psi$ universe of discourse and two input sets on the error universe of discourse. Using the notation and implementation adopted previously for fuzzy controllers in this thesis, this controller can be expressed:

$$\begin{array}{cc} & \overbrace{\begin{bmatrix} -1 & \kappa_n & \kappa_p & 1 \end{bmatrix}}^{\cos\psi \text{ input set locations}} \\ \begin{array}{c}(\mathcal{F}_d - \overline{\mathcal{F}}) \\ \text{input set} \\ \text{locations}\end{array} \left\{ \begin{bmatrix} -5 \\ 5 \end{bmatrix} \right. & \underbrace{\begin{bmatrix} \Delta_{mv_1} & -5 & -5 & \Delta_{mv_2} \\ \Delta_{mv_3} & 5 & 5 & \Delta_{mv_4} \end{bmatrix}}_{\text{modified error value, } e_{i,j}} \end{array}$$

The values κ_n, κ_p, Δ_{mv_1}, Δ_{mv_2}, Δ_{mv_3} and Δ_{mv_4} are tunable parameters whose functions will now be explained. If the value of $\cos\psi$ lies between κ_n and κ_p then, for values of $(\mathcal{F}_d - \overline{\mathcal{F}})$ which lie between -5 and 5, the output is $(\mathcal{F}_d - \overline{\mathcal{F}})$. For values of $(\mathcal{F}_d - \overline{\mathcal{F}})$ outside the range -5 to 5, the output is restricted to -5 or 5. The values κ_n and κ_p, therefore, define a range of values of $\cos\psi$ within which the adaptive actions are unaffected. The values Δ_{mv_1}, Δ_{mv_2}, Δ_{mv_3} and Δ_{mv_4} are used

to modify the adaptive actions taken by the controller when the angles are very large or very small (note the fuzzy statement). When the angle is very large ($\cos\psi$ is close to -1) the output sets of the rules should tend to be moved outwards more and inwards less. When the angle is very small the reverse should occur. The following conditions should, therefore, apply to the values of Δ_{mv_1}, Δ_{mv_2} Δ_{mv_3} and Δ_{mv_4}:

$$\begin{array}{rcccr} -5 & \leq & \Delta_{mv_1} & \leq & 0 \\ & & \Delta_{mv_2} & \leq & -5 \\ 5 & \leq & \Delta_{mv_3} & & \\ 0 & \leq & \Delta_{mv_4} & \leq & 5 \end{array}$$

The values of κ_n, κ_p, Δ_{mv_1}, Δ_{mv_2}, Δ_{mv_3} and Δ_{mv_4} were chosen by a process of tuning by observing the actions of the adaptation process. The values Δ_{mv_2} and Δ_{mv_4} were set to the relatively aggressive values of -10 and 0. These values were used because they would only occur for very small angles. The overall effect on these small angles was adjusted by the parameter κ_p. For the larger angles, the rules concerned are, generally, following rules and the purpose of adjusting the angle is to make the transition from the driving rules to the following rules smoother. Since large angles themselves have no adverse effect on driveability, the value of κ_n was set relatively close to -1 and the values of Δ_{mv_1}, and Δ_{mv_3}, were used to adjust the controller actions.

Two further conditions have to be applied to Δ_{mv_1}, Δ_{mv_2}, Δ_{mv_3} and Δ_{mv_4}:

$$\begin{array}{rcrcrcr} \Delta_{mv_1} & - & (-5) & = & \Delta_{mv_3} & - & 5 \\ -5 & - & \Delta_{mv_2} & = & 5 & - & \Delta_{mv_4} \end{array}$$

These are, essentially, symmetry conditions and are used to ensure that the smoothing action does not bias the controller towards moving inwards or, outwards.

Figure 6.25 shows the progression of the performance measure when the adaptation process was implemented using the smoothing fuzzy controller whose parameter values were as follows:

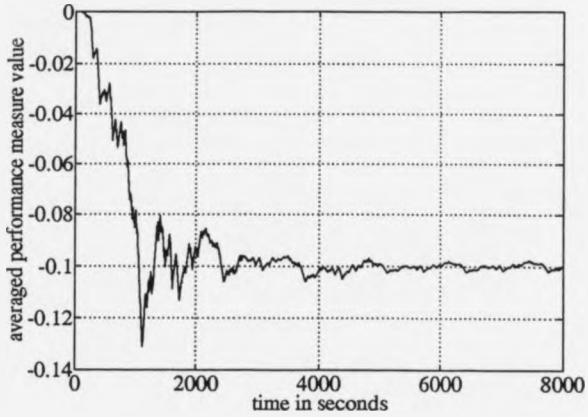

Figure 6.25: Averaged performance measure against time when the adaptation process includes the smoothing fuzzy controller. Powertrain fuzzy controller uses 8 fuzzy input sets on the pedal universe of discourse.

$$\begin{bmatrix} -5 \\ 5 \end{bmatrix} \begin{bmatrix} -1.00 & -0.95 & 0.30 & 1.00 \\ -4.2 & -5.0 & -5.0 & -10.0 \\ 5.8 & 5.0 & 5.0 & 0.0 \end{bmatrix}$$

The locations of the output sets generated by this adaptation process are shown in the four plots on the left hand side of Figure 6.26, in which it is seen that the smoothing fuzzy controller greatly reduces the angular nature of the controller output set locations. The output set locations at 7000 r/min have moved in the opposite direction. However, this is not significant since movements of the rules at 7000 r/min happen very infrequently and the direction of these movements will depend upon the sign of the error when the output sets are moved.

The right hand four plots of Figure 6.26 show the output set locations that were generated when two additional pedal value input fuzzy sets were added. These input sets had maximum grades of membership at 6°, and 10°, of pedal

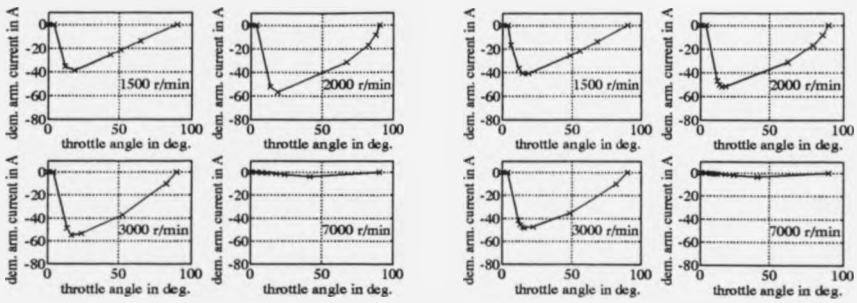

Figure 6.26: Output set locations after controller convergence using the smoothing fuzzy controller. Four plots on left hand side use 8 input sets on the pedal universe of discourse, four plots on right hand side use 10.

value and have been moved such that, in the area where the shape formed by the output set locations is most critical, they define a smooth curve.

6.7 Experiments with the adaptive hybrid powertrain controller.

6.7.1 Varying the demanded performance measure value.

In order to investigate the variation in the performance of the controller, in terms of use of fuel and electrical energy, when the value of $\mathcal{F}_d$ was varied, the adaptive controllers were allowed to adapt over six urban cycles (about 8000 seconds). They were then fixed, and the performance of the vehicle and adapted controller was measured over a single urban cycle. The results obtained by this method are shown in Figure 6.27 for a range of values of $\mathcal{F}_d$ from 0 to -0.16. $\mathcal{F}_d$ = -0.16 approaches the limit of operation of the adaptive controller, as values more negative than this demand armature currents close to the capacity of the motor.

When the controller is adapted in this manner, there does not seem to be any huge advantage, in terms of efficiency, in using any particular value of $\mathcal{F}_d$, though

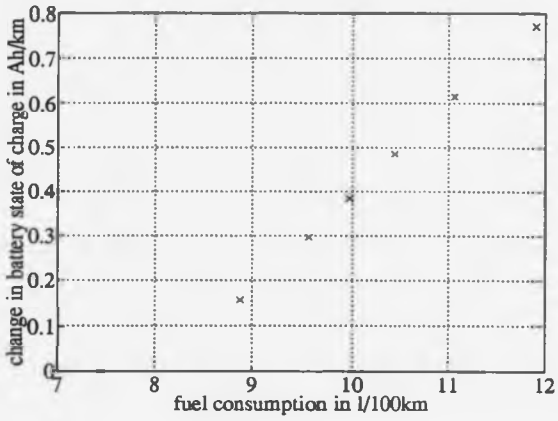

Figure 6.27: Change in battery state of charge against fuel consumption for $\mathcal{F}_d$ = 0, -0.04, -0.08, -0.10, -0.12, -0.14 and -0.16.

in the case of individual users, there may well be advantages in one value over another.

Also, after the adaptation process, for the fixed controllers, it was observed that :

$$\frac{\int I_{bat}}{k_e \int \dot{m}_f} \neq \mathcal{F}_d \tag{6.6}$$

The reason for this difference, which was generally small, is that the performance measure calculates the moving average version of

$$\overline{\left(\frac{I_{bat}}{k_e \dot{m}_f}\right)} \tag{6.7}$$

whereas

$$\frac{\int I_{bat}}{k_e \int \dot{m}_f} = \frac{\overline{I_{bat}}}{k_e \overline{\dot{m}_f}} \tag{6.8}$$

The difference between the demanded performance of the vehicle, and the performance of the vehicle using a fixed controller over a cycle, is very slight. It could be overcome by calculating moving averages of both the battery current, and the fuel

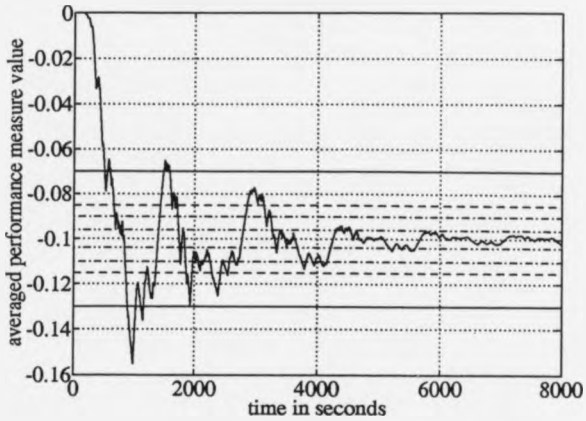

Figure 6.28: Averaged performance measure against time for a simulation using the data from Driver B.

flow rate and then using these values in the calculation of instantaneous values of the performance measure.

6.7.2 Using data from different drivers.

The route whose speed-time data is shown in Figure 6.3 was driven by two drivers. Let the data shown in Figure 6.3 represent the results obtained when Driver A drove the route. Figure 6.28 shows the progression of the performance measure value for a simulation of the vehicle being driven round the route by a second driver, Driver B. Figure 6.29 shows the progression of two of the output set locations during the adaption process for the same simulation, and Figure 6.30 shows converged output set locations for the controller after the simulation. Comparing Figure 6.28 with Figure 6.25, it is seen that the progression of the performance measure is far more oscillatory for Driver B, than for Driver A. Also, comparing Figure 6.30 with Figure 6.26, it is seen that the controller demanded armature current output sets have converged to values that are far lower for Driver B, than

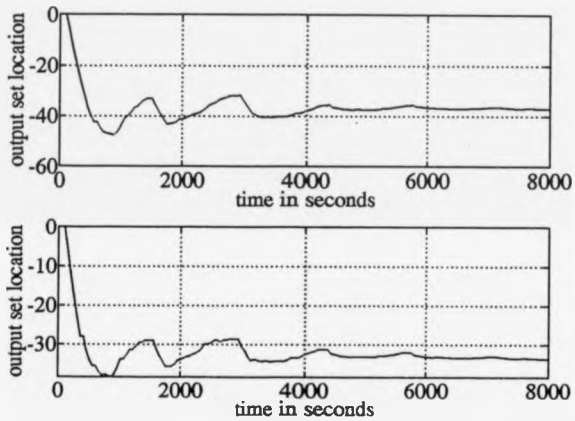

Figure 6.29: Upper plot shows the progression of the demanded armature current output fuzzy set over the duration of the simulation of the cycle driven by Driver B for the rule whose input sets have maximum grades of membership at 2000 r/min and 12° of pedal value. Lower plot shows the same information for the rule whose input sets are at 3000 r/min and 12°. $\mathcal{F}_d$ = -0.1.

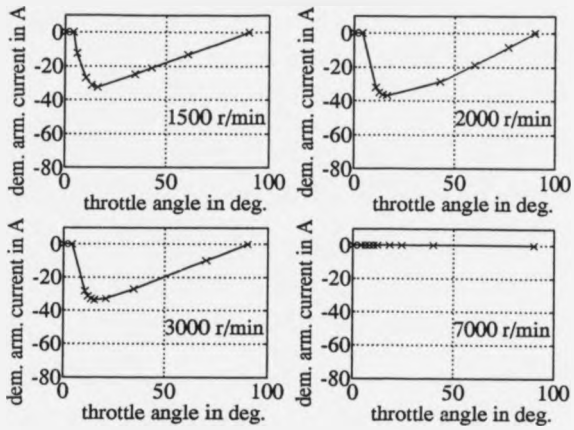

Figure 6.30: Output set locations for the controller having been adapted over the data taken from Driver B.

for Driver A. The reason for these surprisingly large differences is shown in Figure 6.31, which compares the engine speeds used by Drivers A, and B, over the same route. It should be borne in mind that Driver B takes 1452 seconds to do the route, whereas Driver B takes only 1361 seconds. The engine speeds used by Driver B are much more closely gathered around one particular value than for Driver A.

Figure 6.32 shows the same information as Figure 6.31 but uses the locations of the maximum grades of membership of the engine speed input fuzzy sets as the bins for the histogram. These plots, therefore, show the number of occasions that the operating points are clustered towards a particular engine speed input set. The relative number of times that the clustering algorithm selects the engine speed input set at 2000 r/min, as opposed to the other input sets, is far greater for Driver B, than for Driver A. The motor speed that is used when the set at 2000 r/min is selected by the clustering algorithm will always be relatively high, because the motor turns twice as fast as the engine at these speeds. As the motor turns

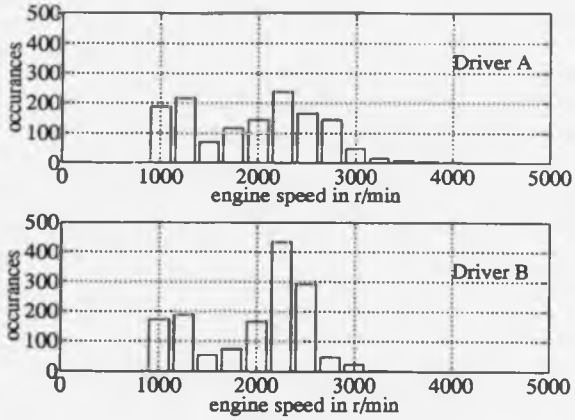

Figure 6.31: Comparison of the engine speed values used by Drivers A and B over the same urban route.

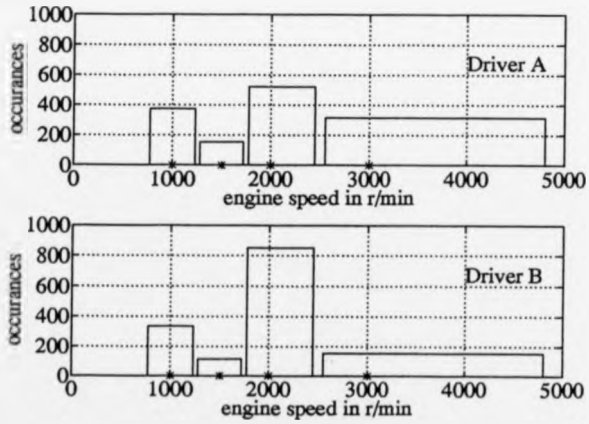

Figure 6.32: Comparison of the engine speed values used by Drivers A and B over the same urban route, using the locations of the maximum grades of membership of the engine speed input fuzzy sets as the bins in the histogram.

faster, it will generate larger regenerative currents for any demanded armature current. Conversely, less negative demanded armature currents are required to cause the same value of the performance measure $\mathcal{F}_d$, resulting in the adapted controller having to demand less negative values of armature current.

The frequent use of the engine speed set at 2000 r/min causes the corresponding output sets to be adjusted frequently, and, since these output sets are very influencial in determining the value of the averaged performance measure, the averaged performance measure value is very much more oscillatory. The frequent use of one particular engine speed by the clustering algorithm has the same effect as increasing the gain, k_p, of the adaptation algorithm.

Figure 6.33 shows the averaged performance measure value, and the value of the indexing variable used to select the value of ν, for a simulation in which the data from Driver B was used for 8000 seconds and then the data from Driver A was used for 8000 seconds. Figure 6.34 shows the movements of two of the demanded armature current output sets during the same simulation. These plots show how the controller, having converged for Driver B, readapts to give the correct performance for Driver A. It is therefore seen that the controller is able to adapt to different driving styles, of different users, which was one of the features that was felt to be desirable in an adaptive hybrid powertrain controller.

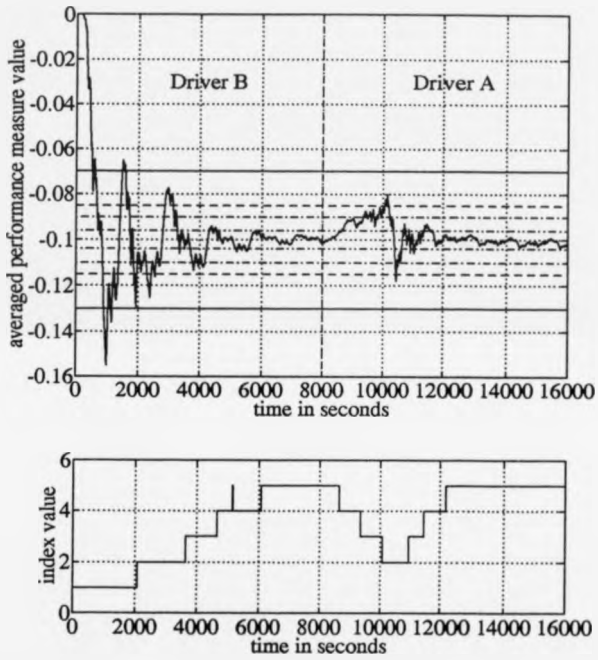

Figure 6.33: Upper plot shows the value of the moving average performance measure against time when the first 8000 seconds of the simulation use data from Driver B and the second 8000 seconds of the simulation use data from Driver A. Lower plot shows the value of the indexing variable used to address elements in the vector $\nu_{1,\ldots,5}$ of ν values. $\mathcal{F}_d$ = -0.1.

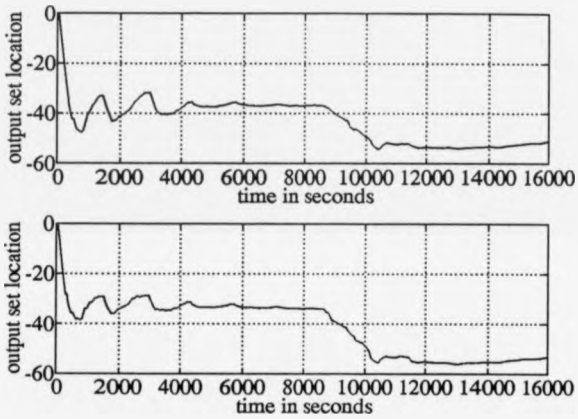

Figure 6.34: Upper plot shows the progression of the demanded armature current output fuzzy set over the duration of the simulation for the rule whose input sets have maximum grades of membership at 2000 r/min and 12° of pedal value. Lower plot shows the same information for the rule whose input sets are at 3000 r/min and 12°. Urban cycle data taken from Driver B for the first 8000 seconds and from Driver A for the second 8000 seconds. $\mathcal{F}_d$ = -0.1.

6.8 Investigating improving the efficiency of the adaptive hybrid powertrain controller.

The earlier material in this thesis has developed an adaptive hybrid powertrain controller that modifies the matrices of fuzzy output set locations, on a row-by-row basis, to achieve a desired overall performance measure, whilst maintaining the vehicle driveability. After the powertrain controller has adapted, the final converged controller is not unique. By moving one row of output set locations inwards, and another outwards, the same overall vehicle performance, in terms of the first performance measure, could be achieved. This section will investigate whether it is possible to improve the efficiency of the vehicle operation by varying the positions of the rows relative to one another.

By varying the positions of the rows of the matrices of output set locations, the same basic modification procedure can be used for each row, as has been used previously. The difference in this work will be the values that are used to drive the row-by-row modification procedure.

6.8.1 Making use of the second performance measure.

It will be recalled from Chapter 5, that the second performance measure, $\mathcal{G}$, is intended to be a measure of the efficiency of the powertrain operation. Equation 5.1 gives the basic form of this performance measure. More specifically, $\mathcal{G}$ is defined:

$$\mathcal{G} = I_{bat} + k_e \dot{m}_f \tag{6.9}$$

in which k_e is the equivalence of fuel and stored electrical energy, and takes the same value of 120 A/(g/s), as in Equation 5.18. When the rows are to be adjusted to improve the efficiency of the powertrain, the ratio of the use of electrical energy to fuel should be maintained. Therefore, the way that the rows should be moved, relative to one another, has to take account of the differing amounts that both

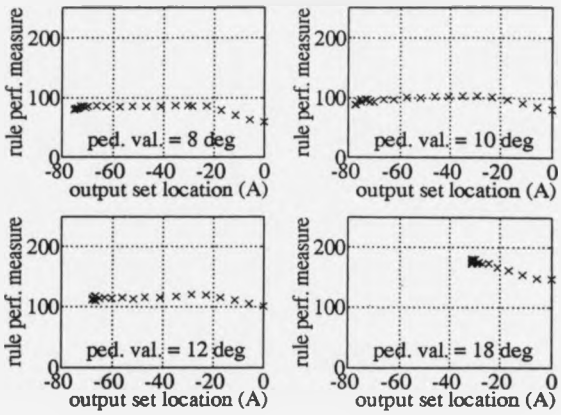

Figure 6.35: Values of $g_{i,j}$ plotted against demanded armature current output set location for four controller rules whose engine speed input sets have a maximum grade of membership of 2000 r/min. Maximum grades of membership of pedal value input sets are shown on the plots

performance measures will change as the rule output sets, in a particular row, are moved. This implies that some measure of the gradient of the performance measure for any given rule, with respect to the location of the rule, is important.

It will also be recalled from Chapter 5, that the matrix $\mathbf{G}$ is used to store the values $g_{i,j}$ of the performance measure $\mathcal{G}$, for each rule. Figures 6.35 and 6.36 show example values of $g_{i,j}$ plotted against the demanded armature current output set locations for eight rules, as the controller is modified by the adaptation algorithm.

The plots in Figure 6.36 display some variation in the values of $g_{i,j}$, because the rules whose engine speed input sets have a maximum grade of membership at 3000 r/min cluster operating points from engine speeds of 2500 r/min to 5000 r/min. If the clustering algorithm chooses a particular set when engine speeds towards the top of this region are used, and this set is chosen by the clustering algorithm only once or a few times, then some variation in the value of $g_{i,j}$ is seen.

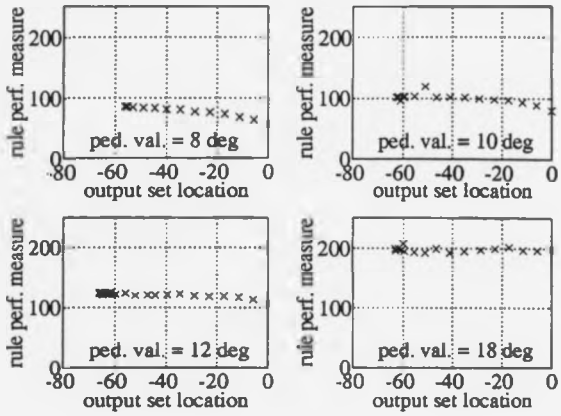

Figure 6.36: $g_{i,j}$ for four controller rules whose engine speed input sets have a maximum grade of membership of 3000 r/min. Maximum grades of membership of pedal value input sets are shown on the plots

Once the output set location has moved beyond a certain, fairly low, negative value, there is very little change in the trend followed by the values of $g_{i,j}$. However, the specific values of $g_{i,j}$ exhibit variations due to the operation of the clustering process. Establishing the value of a gradient in $g_{i,j}$ with respect to the output set location would, therefore, be extremely difficult using this data. Before developing a complex adaptive algorithm that attempts to improve the efficiency of the powertrain, it would seem sensible to attempt to quantify the improvements that such an algorithm could possibly achieve, and this is the subject of the next section.

6.8.2 Controlling the adaptation of each row of the controller individually.

The purpose of this section is to determine whether worthwhile improvements in the efficiency of the controller operation could possibly be gained by improvements

to the adaptive powertrain controller. One very straightforward way in which efficiency of the powertrain can be improved is by restricting the use of the motor at speeds where its efficiency is known to be low. To a certain extent, this has already been achieved by the use of the clutch, allowing the motor to spin twice as fast as the engine, at low engine speeds, or at the same speed as the engine, at higher engine speeds. However, a further improvement might be possible by preventing the controller from demanding armature currents for the fuzzy rules with engine speed fuzzy input sets at 1500 r/min. This is easily achieved by modifying the mask matrix to prevent the movement of the output sets for the fuzzy rules whose engine speed fuzzy input sets are at 1500 r/min.

The results obtained by the controllers adapted in this way are compared with the previous results in Figure 6.37, in which it is seen that a small improvement in the overall powertrain efficiency has been achieved by constraining the rules at 1500 r/min. In some places, the battery state of charge could be improved by around 5% for the same use of fuel.

It would seem, from these results, that a small improvement in the controller performance can be achieved straightforwardly. To investigate whether any greater improvement in the overall efficiency can be made, the adaptation procedure was modified to move each row of the controller output set location matrices according to some performance measure which was specific to that row. In order to do this, it was necessary to estimate the performance of each row of the controller. The **F** matrix can be used for this purpose, since it includes a performance estimate, $f_{i,j}$ for each rule. These performance estimates should be combined across the row to give an overall estimate for the performance of the row. This row performance measure is obtained by forming a mean of the performance measure values of all the rules across the row, and by weighting this mean by the frequency with which the rule is used, relative to the frequency with which the

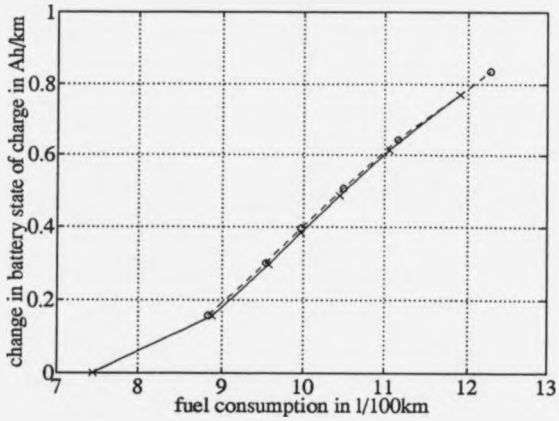

Figure 6.37: Comparison of results obtained by the adaptation process with the output sets for the rules at 1500 r/min unconstrained (shown with crosses) and fixed (shown with circles).

row is used. This mean is then used, in a moving average algorithm, to generate moving average performance measures for each row of the controller output set matrices, analogous to the overall moving average performance measure for the entire vehicle. The average performance measure value for the r'th row, $\overline{\mathcal{F}_r}$, is found:

$$\overline{\mathcal{F}_{rn}} = \overline{\mathcal{F}_{rn-1}}\frac{\nu_r - 1}{\nu_r} + \frac{1}{\nu_r}\frac{\sum_{j=1}^{n_p} f_{r,j}c_{r,j}}{\sum_{j=1}^{n_p} c_{r,j}} \tag{6.10}$$

The value of ν_r is selected from a vector of values $\nu_{r1,\cdots,5}$, in a similar way to the value of ν used in the overall moving average performance measure, and n_p is the number of pedal value input fuzzy sets.

The rule modification process is analogous to the modification process for the overall performance measure. For each row, r, a row demanded performance measure $\mathcal{F}_{d_r}$ is used. The difference $(\mathcal{F}_{d_r} - \overline{\mathcal{F}_r})$ is used as the input to the smoothing fuzzy controller that generates the modified error value $e_{r,j}$. This modified error

value is then used to calculate the change in the output set location using:

$$\Delta_{mv_{r,j}} = k_{p_r} e_{r,j} \tag{6.11}$$

The value k_{p_r} is chosen from a vector of values $k_{p_{1,\cdots,5}}$, addressed using the same indexing variable as is used to address $\nu_{r1,\cdots,5}$. Different values of k_p are required because, when large values of ν are used, the output can only change slowly. Consequently small errors between demanded and actual values of performance measure can exist for long periods. These small values would lead to large controller modifications if the value of k_p were not decreased as the value of ν were increased.

One of the rows is not modified using the error $(\mathcal{F}_{d_r} - \overline{\mathcal{F}_r})$, but uses the error $(\mathcal{F}_d - \overline{\mathcal{F}})$. This ensures that the controller will tend to give the vehicle a determined measure of overall performance. The manner in which this performance is achieved can be varied by the use of the various values $\mathcal{F}_{d_r}$.

Figures 6.38 and 6.39 show the results of two simulations in which separate values of $\mathcal{F}_{d_r}$ are used to adjust the output set locations. The values of $k_{p_{1,\cdots,5}}$ were 40, 40, 20, 20, and 10. In both cases the mask is used to prevent the output sets of the rules with engine speed input sets at 1000 r/min, and 1500 r/min, from moving. The sets for the rules at 3000 r/min are adjusted by the overall performance measure. In Figure 6.38, $\mathcal{F}_{d3} = -0.158$, which is the value of $\overline{\mathcal{F}_3}$ that the controller would arrive at when adapting purely on the overall performance measure value, which was set at -0.1. In Figure 6.39, $\mathcal{F}_{d3} = -0.2$ and it is seen that, in order to make the overall vehicle performance remain at -0.1, a much lower value of $\mathcal{F}_{d4}$ is required. In running these simulations, it was found necessary to increase the gain for the row of output set locations that is used to adjust the overall performance of the controller. This is because the performance of the vehicle, as a whole, is less responsive to adjustments made to just one row of output set locations, than to adjustments made to all the output set locations.

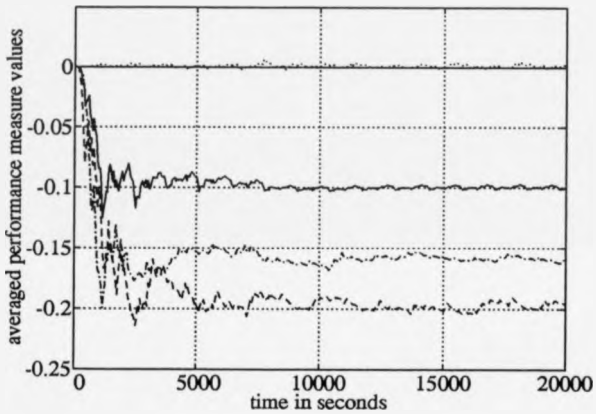

Figure 6.38: Plot of the overall vehicle averaged performance measure and row averaged performance measures for the controller whose output sets are modified by separate values of $\mathcal{F}_{dr}$. Solid curve is the overall vehicle performance measure, dashed curve is the performance of the rules at 3000 r/min, dash-dot is the performance of the rules at 2000 r/min and the dotted curve is the performance of the rules at 1500 r/min. The solid straight line is the performance of the rules at 1000 rmin. $\mathcal{F}_{d3} = -0.158$

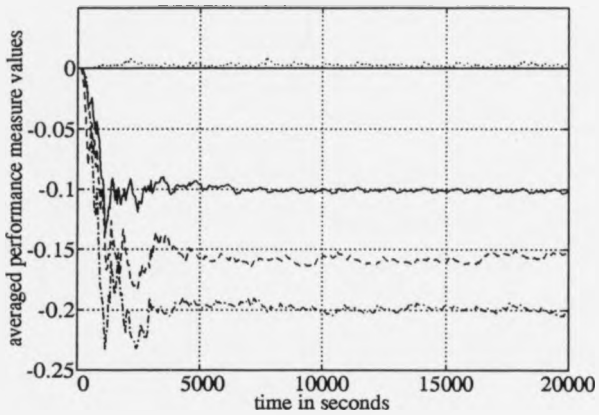

Figure 6.39: Plot of the overall vehicle averaged performance measure and row averaged performance measures for the controller whose output sets are modified by separate values of $\mathcal{F}_{d_r}$. Line types as previous plot. $\mathcal{F}_{d3} = -0.2$

In order to investigate whether improvements in the controller efficiency were possible by individually controlling the output set adjustments for each row of the controller, a sensitivity study whose results are shown in Figure 6.40 was conducted. In this study, the value of $\mathcal{F}_d$ was fixed and the controller adapted using various values of $\mathcal{F}_{d3}$, in a similar way to the two examples described above. Having been adapted the output set locations were then fixed and a further simulation was carried out to assess the performance of the static controller. The points plotted on the graph show the results of the sensitivity study and the dashed line is drawn between the results previously obtained with the output sets for rules at 1500 r/min fixed, shown in Figure 6.37.

From the results presented in Figure 6.40, it would seem that there is little justification for attempting to improve the powertrain operating efficiency using the adaptive controller. The results obtained in Chapter 3 had indicated that there was the possibility for some powertrain efficiency improvement. The reasons why

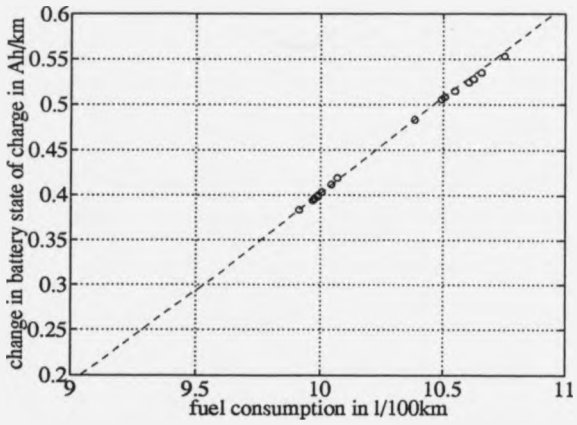

Figure 6.40: Results of sensitivity study investigating the variation of powertrain operating efficiency by the use of separate performance measure at different powertrain operating speeds.

no improvement has been possible concern the nature of the drive cycle used in Chapter 3, and the vehicle model used in the later work. The ECE-15 cycle is short and does not greatly exercise the vehicle. It is also composed of straight line segments that cause the vehicle to visit relatively few sites during a simulation.

In order to improve the apparent efficiency of operation over ECE-15, it is necessary to perform well over just a few sites. In order to improve the efficiency over the urban cycle acquired from a real vehicle, it is necessary to improve the operation of the powertrain over a much greater range of operating conditions and this seems to have been very difficult. The other reason for the difficulty in improving the efficiency of the vehicle is that the addition of the clutch allowing the motor speed to be set equal to the engine speed, or twice the engine speed, removed the less efficient regions of motor operation and made the potential for subsequent performance improvement considerably less. The conclusion of this sensitivity study is that, with the vehicle model including variable motor speed

operation, there is some justification in constraining the movement of the rules at low engine speeds as this does give a small performance benefit. However, at this stage, it is hard to justify further complicating the adaptive controller in the attempt to improve the vehicle efficiency.

6.9 Conclusions.

The basic adaptive fuzzy controller described in Chapter 5 was extended to include the use of an additional input of engine speed. It was seen that the inclusion of this input made the implementation of the adaptive fuzzy controller more elegant. Initially, the movement of the output set locations was controlled for each individual rule. However, this was laborious, caused problems in the convergence of the output set locations and gave little overall control of the vehicle performance.

The next stage in the development of the controller was to modify the rules of the controller for each engine speed input fuzzy set together, in response to changes in the overall vehicle performance. Each rule was modified as a rule modification driver, or a rule modification follower, depending upon the way in which that rule had been used. Originally, it was proposed to use rate of change of the performance measure of each rule, with respect to the location of the output set, to distinguish between driving and following rules. However, in practice, establishing this gradient proved difficult and the driving and following rules were distinguished simply by their use over the last adaptation interval.

The smooth, robust movement of the following rules required considerable development and was achieved by implementing parts of the rule modification process using small fuzzy controllers. The final adaptive controller was then used in a number of experiments.

The adaptive action of the controller could be used to control the relative amounts of fuel, and electrical energy, used very straightforwardly. Possibly, the

most significant aspect of the performance of the adaptive controller was its ability to adapt to the differing driving styles of more than one user, even when adapting using data taken from the same route.

Finally, a sensitivity study was undertaken in which the controller adaptation was controlled using performance measures for each engine speed as well as an overall vehicle performance measure. It was found that some improvements were possible by restricting the movements of rules at low engine speeds, but beyond this, developing the controller to achieve further improvements was unlikely to be worthwhile.

The conclusion from this sensitivity study is that there seems to be no requirement to measure the performance of individual fuzzy rules on a production vehicle, since the overall performance of the vehicle is used to modify the controller. However, the performance measures for each rule have been useful in the development of the adaptive controller and could also be useful in the development of prototype vehicles, and further hybrid vehicle research.

The largest contribution that this chapter makes in the more general field of fuzzy control is that of the use of small fuzzy controllers in the adaptation of other controllers. The smoothing fuzzy controller used in the adaptation process worked very quickly, and was easily implemented and tuned. It also proved robust in use, since the original controller parameters were used with varying numbers of output sets and different modification procedures. In this sense, the design method that fuzzy control forces is robust. This does not imply that fuzzy control is robust in the stricter sense that robustness has in the context of feedback control. These observations support the view that feedback control, whilst being an obvious area of application of fuzzy logic, is not, perhaps, its most natural area of application.

The linguistic statements that were used during the design of a particular fuzzy controller have been included because they illustrate the manner in which

the controllers were designed. The linguistic statement is useful in both stating the problem to be solved, and in identifying the information required to solve it. When the method of implementing fuzzy logic is straight forward, as it is in this thesis, a controller can then be quickly developed and implemented. In this sense, the ways in which fuzzy logic has been used in the adaptation procedure were motivated by the original philosophical beliefs about fuzzy control, but have been implemented in a manner suitable for use in a modern real time control application, using conventional hardware.

Chapter 7

Conclusions.

This thesis has considered the application of the theory of fuzzy logic to the practical automotive engineering problem of the management of an automotive hybrid vehicle powertrain. Conclusions and recommendations for further work will be drawn separately for hybrid vehicles and fuzzy logic.

7.1 Hybrid Vehicles.

Throughout the work described in this thesis, there have been three primary, customer related, requirements of the hybrid powertrain controller. These were to use the energetic resources of the vehicle in the correct amounts relative to one another, to use these resources efficiently and, also, to allow the vehicle to be driven in a consistent and sensible manner. Other, less directly customer related, requirements were that the work make use of existing automotive control technology and be technically feasible, from a cost and engineering development point of view.

In order to investigate the control of a hybrid powertrain, it was felt necessary to develop computer models of the components that make up such a vehicle. The modelling approach taken used forward dynamic models and attempted to embody some understanding of component operation. The use of efficiency maps

was avoided where possible, although the engine model used did make use of fuel flow rate and torque maps. In the case of the motor, the modelling approach was sufficiently general to allow the use of the same motor model in vehicles with different batteries, indeed this model has been used in other engineering projects within the Rover Group. Without doubt, there is a considerable advantage in the circuit based motor modelling approach taken in this work, over the efficiency map based models, because the motor model is more generally applicable and a far greater understanding of the vehicle operation is gained. For the fully warm engine, there is no analogous external factor that affects the performance in the same way as the battery affects the performance of the motor.

The difficulty in modelling various aspects of the engine enforced a constraint on the operation of the hybrid powertrain. Previous work in this field has not made use of an engine whose emissions are controlled by the use of a catalyst and this has meant that on-off engine operation has been feasible. In this work, the engine was left running at all times, because in a real vehicle, if the engine is left at idle, the catalyst stops converting, leading to disastrous emissions performance. To avoid this, on-off engine operation has not been used, and consequently, improving the efficiency of the powertrain operation has been much more difficult. It has also been apparent that the operation of the controller of a hybrid powertrain will be dependant upon the exact form of the powertrain. This is significant, since there are many ways in which the parallel hybrid vehicle powertrain can be implemented on a vehicle.

It was realised early in the work that the control of a hybrid vehicle powertrain was not a reference tracking feedback control problem, and in this sense, it may well not be what the control engineering community would call a "control" problem at all. As far as automotive engineering is concerned, there is no doubt that the problem is a "control" problem, and it was seen that none of the generally

accepted control methods were suitable because of the problem formulation. The lack of an appropriate conventional control technique, and the flexible nature of fuzzy logic, made it an obvious choice for the control of the hybrid powertrain. It should be borne in mind that the motivation for this work came from an industrial requirement and not from an interest in a particular control methodology. It has, therefore, been satisfying that the more work was done on the problem, the more fuzzy logic seemed to be an appropriate aid in its solution.

The first attempt at the use of fuzzy logic in the control of the powertrain was very much in line with the way in which fuzzy logic was originally intended for use. This exercise generated a great deal of understanding and experience in the operation of the powertrain, in the light of which, modifications were made to the vehicle model. These modifications included the inclusion of a one way clutch on the engine output to prevent the engine overrun torque from opposing the motor operation and, also, the inclusion of two discrete gears between the motor and the transmission input. These two gears allowed the motor to turn at either the same speed as the engine, or at twice the speed of the engine. This is an example of how the control strategy for a hybrid powertrain should reflect the nature of the components used, since, if a motor whose efficiency was less dependant upon speed were to be used, these gears would probably be unnecessary. Whilst this initial application of fuzzy logic generated considerable understanding, it was felt that there was a need to use adaptive methods to assist in the control of the powertrain of the vehicle.

In the introduction to Chapter 5, it was seen that an adaptive hybrid powertrain controller could be useful at three levels. The controller's adaptive capability could be used in the development of a vehicle, to generate a "production release tune" for a fixed controller. Alternatively, the controller might be able to adapt to the requirements of a particular user in service, and lastly, the controller might

be able to adapt to compensate for ageing of the vehicle components.

To assess the extent to which the adaptive controller is useful in these three requirements, it is sensible to consider the ways in which a hybrid vehicle would be used. Firstly, two major modes of operation would include the use of just the electric motor, because of its low cost of operation and perceived lack of environmental impact, and just the heat engine, for long journeys at high speeds where the electric side of the powertrain can clearly make very little contribution. Two other modes of operation also seem to exist. During general day-to-day use, the owner of a hybrid vehicle may wish to use as much electricity as possible, since this is a cheap source of energy. Many users will have performance, or range, requirements that will not be easily met by the use of purely the electric side of the powertrain and will therefore need to supplement the operation of the motor with power from the heat engine. The last mode of operation is a further combined mode of operation, in which the electric motor acts against the engine in order to raise the battery state of charge. This mode of operation would only be needed when approaching a region in which there was an emissions restriction that would require a battery state of charge higher than the current battery state of charge.

The adaptive controller was developed around the fourth mode described above since the engineering problems faced here are at their most severe. However, in practical vehicle operation, this mode will be used infrequently and, consequently, there will be little opportunity for adaptation in service. The adaptive controller would, therefore, only be useful for developing the production tune for this mode. Conversely, many users will spend considerable amounts of time using their vehicles in the other combined mode of operation in which the adaptation strategy would still work and might be implemented as follows. The user of the vehicle will clearly have to select one of the four basic modes of operation. If the combined

mode of operation in which the motor and the engine work together is selected, then a further control which indicates the relative split between the two power sources could also be adjusted. The vehicle user could set this control such that over the course of their daily motoring the most effective use is made of the electrical energy stored on the vehicle. The experiments reported in Chapter 6 indicate that the controller should be able to adapt to changes in the route or user easily within two or three days of motoring. There is also the possibility of a vehicle adapting to the requirements of more than one user, since modern methods of vehicle access allow the vehicle to distinguish easily between different drivers. The ability of the adaptive controller to adapt to the requirements of individual users is the largest contribution that this thesis makes to the field of hybrid vehicle engineering.

Further hybrid vehicle powertrain control work.

In the opinion of the author, little value will be added by further simulation studies using this vehicle model and controller approach. The work has reached a point where the simulation studies should be validated by experience with a vehicle. Additionally, some of the parts of the adaptation strategy require the evaluation of square roots, which is starting to be considered in engine control but which has not generally been done before. Engine control is characterised by the need to respond to events which occur every few hundred microseconds and whose real-time aspects are very severe. In contrast, the controller being considered here is only required to output a control value every 100 milliseconds and, consequently, will have large periods of time in which the adaptive calculations can be carried out. Experience with a real vehicle will be required to determine the degree of processing power that will be needed to implement the adaptation strategy, though some indication of its effectiveness in use has been gained through these

simulation studies.

Further work on the simulation of the hybrid vehicle powertrain considered here should only be carried out to investigate some of the thermal effects that have limited the control strategy to this date. The most important of these effects is the temperature of the catalyst and the engine. If these effects were realistically modelled, control strategies taking into account vehicle emissions could be developed. This would clearly be a very useful contribution, since any production vehicle, once developed will have to pass some kind of emissions test. Also, more general control strategies, utilising on-off engine operation, could be developed allowing a considerable reduction in the fuel used. These improvements would really come about as a result of improvements in the science of engine modelling, which is a separate, very active field of research.

The control strategy described here is able to respond to ageing effects of the engine and motor in their use of fuel and electrical energy. However, it is not able to respond to ageing effects in the generation of torque. In an engine with a modern closed-loop fuelling system, the torque generated and the fuel used should be related to one another. It might be possible to infer variations in the torque output of the engine, from easily observed variations in the engine fuel flow rate. This represents a further possible avenue of future research.

7.2 Fuzzy logic.

As stated in the previous section, the work described here was motivated by an industrial engineering problem, rather than by a desire to conduct research in the subject of fuzzy logic. It was soon realised that fuzzy logic provided a natural framework in which the subject of hybrid powertrain control could be investigated, this has allowed conclusions to be drawn about the use of fuzzy logic in modern engineering.

The experiences gained from the first application of fuzzy logic to the hybrid powertrain were, in some respects, typical of the experiences reported on many occasions previously. These would include the ease of setting up a solution to a relatively complex problem, and the subsequent requirement for "tuning" a controller whose performance was unsatisfactory.

In this early work, considerable attention was paid to the implementation of fuzzy logic from the point of view of the efficiency of the operation of the code. The interpolative effects of fuzzy logic, noted by many researchers, were exploited by using fuzzy logic in a particular way. The fuzzy sets used were triangular and overlapped with their nearest neighbours as far as their maximum grades of membership. Product implication was used and the union of the output sets was formed by the summation operator. This output set was defuzzified using the centre of area. Under these circumstances, fuzzy logic reduces to linear interpolation. This discovery came as something of a surprise and almost a disappointment to the author, who, being familiar with the extensive use of maps in automotive control, felt that there ought to be something very new, and somehow different, about fuzzy logic. However, it was soon realised that the constraints put on the implementation method are quite minor and that they can be justified on philosophical grounds. The sets, when drawn, still represent vaguely held beliefs and this type of fuzzy set can convey linguistic concepts. They also logically divide an input universe of discourse and, to some extent, formalise the development of the controller, as well as giving a fast, computationally compact fuzzy controller.

This method of implementing fuzzy logic has the further advantage that the output can be guaranteed to pass through the centres of area of the output sets. This is an extremely important result from the point of view of product liability which poses real limitations on many industries, with aerospace and automotive being obvious examples. A further advantage of this characteristic of the output

is that it allowed the rule modification process, developed later in the thesis, to ensure that the controller maintained a constant output torque. The major criticism of this method of implementing fuzzy logic is that the input-output relation is not continuous in gradient. In this application, the discontinuous gradient poses no great problem.

Having attempted a basic implementation of fuzzy logic it was found, at the end of Chapter 3, that there was a need to develop some adaptive method of placing the controller output sets. Accordingly, the work conducted to date in the field of adaptive fuzzy control was reviewed. It was found that very little of the work carried out previously was directly applicable. The most significant observation that came out of this section of the work was the very generic nature of the adaptive fuzzy problems tackled to date, and the very specific nature of the hybrid vehicle powertrain control problem. This observation highlights a very important issue concerning the original motivation of the subject of fuzzy control and areas in which research is currently being conducted. The original premise of fuzzy control, that of a human operator being able to successfully control plants with which little success had been enjoyed using automatic control, implies that there is something *specific* to the engineering task at hand that the human recognises. The majority of the adaptive fuzzy work done to date abandons this concept. This is illustrated by the fact that, in most adaptive fuzzy controllers, the fuzzy set does not convey a notion of some vaguely described quantity, but is part of a series of fuzzy sets, which when subjected to fuzzy operations, describe some function. It is really this function that is adapted rather than any concepts embodied in the fuzzy sets. This is as true of the basic hybrid powertrain fuzzy controller that is adapted in the work reported here, as it is of other adaptive fuzzy controllers.

The change in the use of fuzzy sets from representing vaguely held notions, to collectively representing functional relationships seems to be a consequence of

seeking for generic solutions to the generic class of feedback control problems. It seems, therefore, that the fuzzy set is useful as a means of representing functions in generic types of problems, and as a means of representing vaguely held beliefs in specific types of problems. It is interesting to note that the original motivations for the use of fuzzy logic are not generally used once the step is taken of adapting a controller whose basic operation is unsuitable.

In this work, it was felt that producing an adaptive controller whose adaptation process contained aspects of fuzzy logic would be an interesting task for a number of reasons. Firstly, there is a parallel with the reported work in which fuzzy controllers are used to tune PID controllers. This is a seemingly ideal application of fuzzy control, building on the vast human experience in tuning this type of controller. Secondly, implementing the adaptation process is likely to be successful if the premise that fuzzy logic really does describe the way humans think is true, since it is surely the adaptive and learning properties of human behaviour that are seen to be so attractive in adaptive and learning control.

To claim that the controller used in this work is adapted by a higher level fuzzy controller would be perhaps overstating the case, since the root of the adaptation controller is a simple proportional adjustment. However, two of the aspects of the adaptation procedure are implemented using small fuzzy controllers and the success of these methods should be considered in some detail.

Firstly, the two controllers were both very small. One had one input, one output and two rules, the other had two inputs with four fuzzy sets on one universe of discourse and two fuzzy sets on the other universe of discourse and, again, one output. This meant that the number of different situations that the controller could distinguish was limited to a small number that the human controller designer could also distinguish. This helps the human in suggesting appropriate output values.

The nature of the adaptation problem was relatively complex and reflected the ideas and understanding of the author and also the desired final performance of the adaptation algorithm. The controllers that were eventually developed carry out extremely specific tasks within the overall process of adaptation. It seems very unlikely that the specific controllers developed would have any use elsewhere in control engineering, but their success in this particular application recommends a similar approach in similar problems.

In order to determine the nature of the required adaptive action, linguistic statements were used. Initially, this was done to determine whether linguistics, which have not been used a great deal in recent work, had any real relevance. It was found that, for these small controllers, linguistic statements, which were not in the form of a set of fuzzy rules, but tended to be a simple statement, were extremely valuable. Their value lay in their ability to identify the correct information required for any control or decision making task. Once this information has been identified, effort was directed towards obtaining it and making use of it. A strong conclusion from this work is that, in some circumstances, linguistics are invaluable, not simply as an addition to fuzzy sets in the representation of vague concepts, but, more generally, as a vital stage in the controller development. It seems to the author that the more colloquial and natural the linguistic statement, the more likely it is to contain valuable insight into the control task at hand.

The method by which fuzzy logic was implemented, allowed these controllers to be developed and tuned extremely quickly, and once tuned, they were generally left unattended as successive layers of the adaptation algorithm were developed above them. In this sense, these small controllers were robust, though as stated previously, this in no way implies that fuzzy logic is robust in feedback control applications.

At an early stage it had been felt that the ability to measure the performance

of each rule in a set of fuzzy rules was important. As the research was conducted, it was realised that this aspect of the fuzzy controller was not really useful for production hybrid vehicles, though it had been useful in the research carried out in this thesis. This rule performance measurement technique may, perhaps, be useful in other applications of fuzzy logic.

Further work continuing from this thesis is not considered as a separate section for fuzzy control. This is because it is felt that the most significant contributions that this thesis makes to the field of fuzzy control, that of the hybrid powertrain controller itself, and the use of small fuzzy controllers in adaptive techniques, are inherently specific to the application considered. It is felt that the concepts that are embodied in this work will be of use in the consideration of any control engineering project which is new or about which there are no established techniques.

One of the major aims of this thesis, was to investigate the relevance of the original ideas motivating fuzzy control in modern control engineering, given the age of the subject and the nature of current fuzzy research activities. The growth of fuzzy logic in commercial products is an obvious indication of the relevance of fuzzy logic as a basic control technology. In many cases, it is possible to see the combination of new sensors and intelligent information processing, that are the hallmarks of good fuzzy engineering.

It is hoped that this thesis will provide evidence supporting the view that fuzzy logic is not just useful in small scale, almost trivial applications, but is capable, when applied in the correct manner, of making a considerable contribution in much larger applications of significant complexity. As the debate between the merits of conventional control techniques and intelligent control techniques continues, work such as this will hopefully encourage the use of a considered approach, in which control techniques are selected by their relevance for the application, rather than by intellectual dogma.

Bibliography

[1] L. Ridge. Environmental impact model - energy database. Rover Group, Internal Report., 1992.

[2] J. B. Heywood. *Internal Combustion Engine Fundamentals.* McGraw-Hill., 1989.

[3] Californian Air Resources Board. Resolution number 90-58., 1990.

[4] The Council of the European Communities. Draft proposal for a council directive on the setting of incentives to reduce the carbon dioxide per kilometre emissions from new passenger cars., 1993.

[5] C. Holman. Environmental impacts of urban electric vehicles. In *Proceedings of MIRA Electric Vehicle Technology Seminar, 29th April 1992, Leamington Spa, U.K.*, 1992.

[6] M. Cooper-Reade. Electric vehicles and the greenhouse effect. In *Proceedings of MIRA Electric Vehicle Technology Seminar, 29th April 1992, Leamington Spa, U.K.*, 1992.

[7] H. Pieper. Mixed drive for autovehicles. *United States Patent Office, Number 913,846.*, 1909.

[8] An electrical cab. *The Engineer.*, 20th August 1897.

[9] J. Thomason. The world's most unique spitfire? *The Courier.*, 159:30–31, September, 1993.

[10] M. Westbrook. Electric vehicles – developments and limitations. In *Proceedings of MIRA Electric Vehicle Technology Seminar, 29th April 1992, Leamington Spa, U.K.*, 1992.

[11] J. Barker. Electric shock. *Autocar and Motor.*, pages 50–53, 18th April 1990.

[12] K. J. Bullock. Design characteristics of hybrid vehicles with road test results from several prototypes. *SAE 852276*, 1985.

[13] A. Kalberlah. Hybrid drive systems for cars. *Automotive Engineering (SAE)*, Vol. 99, No. 7:pp 17–19, 1991.

[14] Robert Bosch GmbH. *Automotive Handbook.* 1986.

[15] Californian Air Resources Board. Mail-out #93-27., 1993.

[16] R. A Williams. Automotive active suspensions. In *Proceedings of IEE Colloquium on Active SUspension Technology for Automotive and Railway Applications.*, 1992.

[17] A. N. Costa and R. P. Jones. Motion management for automotive vehicles. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume Vol 2., pages pp 932 – 937, 1991.

[18] U. Kiencke. A view of automotive control systems. *IEEE Control Systems Magazine*, Vol. 8(No. 4):pp 11 – 19, August 1988.

[19] S. D. Farrall, A. S. Cherry, and R. P. Jones. Fuzzy control applied to automotive vehicles. In *IEE Colloquium "Two Decades of Fuzzy Control - Part I", London, U.K.*, 1993. Digest No. 1993/166.

[20] K. Howard. Wacky world of car tech No.16a. fuzzy logic. *Autocar and Motor.*, 196(2):68–69, 7th April 1993.

[21] L. Peters, K. Beck, and R. Camposano. Fuzzy logic controller with dynamic rule set. In *Proceedings of the 1992 IEEE International Symposium on Intelligent Control, Glasgow, UK.*, pages 216–219, 11th-13th August 1992,.

[22] Intel Corporation. *Intel Automotive Handbook.*

[23] U. Kiencke, S. Dais, and M. Litschel. Automotive serial controller area network. *SAE 860391*, 1989.

[24] Philips Semiconductors. *80C51-Based 8-Bit Microcontrollers, Data Handbook.*, 1993.

[25] S. D. Farrall. The development of a circuit based electric vehicle powertrain model. In *Proceedings of MIRA Electric Vehicle Technology Seminar, 29th April 1992, Leamington Spa, U.K.*, 1992.

[26] S. D. Farrall. The development and validation of hybrid vehicle powertain models for use in the design of high level hybrid vehicle powertrain controllers. In *Proceedings of the 11th International Electric Vehicle Symposium*, Florence, 1992.

[27] K. Lillie. Electrical propulsion using sodium sulphur batteries. In *Proceedings of Autotec Congress, Birmingham, U.K.*, pages 12–15, November, 1991. Seminar 20, Paper No. C427/20/211.

[28] A. Copus. DC traction motors for electric vehicles. In *Proceedings of International Conference on Electric Vehicles for Europe, Stoke-on-Trent, U.K.*, pages pp 63 – 80, May, 1991.

[29] J. T. Kummer and N. Weber. A sodium sulphur secondary battery. *SAE 670179*, 1967.

[30] J. W. Angelis, H. Birnbreier, and H. Haase. High-energy battery for electric road vehicles. Technical Report D HB 1253 87 E, Asea Brown Boveri., 1987.

[31] S. Preston. Volume production of sodium sulpur batteries. In *Proceedings of International Conference on Electric Vehicles for Europe, Stoke-on-Trent, U.K.*, pages pp 43 – 48, May, 1991.

[32] A. R. Tilley. Advanced batteries for vehicle applications. In *Proceedings of MIRA Electric Vehicle Technology Seminar, 29th April 1992, Leamington Spa, U.K.*, 1992.

[33] R. J. R. Johns. A unified method for calculating engine flows. *ASME Paper No. 84DGP18*, 1984.

[34] Sherman R.H. and Blumberg P.N. The influence of induction and exhaust processes on emissions and fuel consumption in the spark ignited engine. *SAE 770880*, 1977.

[35] Tabaczynski R.J. and Ferguson C.R. A turbulent entrainment model for spark ignition engine combustion. *SAE 770647*, 1977.

[36] Powell J.D. A review of ic engine models for control system design. In *Selected papers from the 10th Triennial World Congress of the International Federationn of Automatic Control, Munich, Germany.*, volume 3, pages 233-240. Pergammon Press, 27-31st July 1987.

[37] I. F. Kuriger. *Some Aspects of Modelling and Control of Automotive Power Systems.* PhD thesis, University of Warwick., 1985.

[38] J. E. Auiler, J. D. Zbrozek, and P. N. Blumberg. Optimisation of automotive engine calibration for better fuel economy - methods and applications. *SAE 770076*, 1977.

[39] J. F. Casssidy. A computerised on-line approach to calculating optimum engine calibrations. *SAE 770078*, 1977.

[40] A. R. Dohner. Optimum control solution of the automotive emission-constrained minimum fuel problem. *Automatica*, 17(3):441–458, 1981.

[41] N. F. Benninger and G. Plapp. Requirements and performance of engine managment systems under transient conditions. *SAE 910083*, 1991.

[42] H. Sindano. Phenomenological transient fuelling model for port injected engines. In *Proceedings of the Institute of Mechanical Engineers Seminar on Engine Transient Performance.*, 7th Novemner 1990. (not part of formal pagination).

[43] C. E. Baumgartner, C. H. Onder, and H. P. Geering. Multivariable control of lean-burn engines. In *Selected papers from the 10th Triennial World Congress of the International Ferderationn of Automatic Control, Munich, Germany.*, volume 3, pages 245–249. Pergammon Press, 27-31st July 1987.

[44] M. Nasu, A. Ohata, and T Meguro. Idle speed control by non-linear feedback. *JSAE Review*, 13(2):54–60, 1992.

[45] A. M. Foss, R. P. G. Heath, P. Heyworth, J. A. Cook, and J. McLean. Thermodynamic simulation of a turbocharged spark ignition engine for electronic control development. In *Proceedings of Institute of Mechanical Engineers Seventh International Conference. Automotive Electronics. London, Uk. 9-13 Oct. 1989.*, pages 195–202, 1989.

[46] P. W. Masding and J. R. Bumby. Identification and peformance simulation of a hybrid IC-engine/battery electric automotive powertrain. Part 1: The IC engine. *Transactions of the Institute of Measurement and Control.*, 12(1):27 – 39, 1990.

[47] J. R. Bumby, P. H. Clarke, and I. Forster. Computer modelling of the automotive energy requirements for internal combustion engine and battery electric-powered vehicles. *IEE Proceedings*, 132, Pt. A(5):265–279, 1985.

[48] D. B Gilmore. Fuel econnomy goals for vehicles with regenerative braking energy storage systems. *International Journal of Vehicle Design.*, 13(2):125–133, 1992.

[49] W. H. Press, B. P. Flannery, S. A. Tenkolsky, and W. T. Vetterling. *Numerical Recipes. The Art of Scientific Computing (FORTRAN Version).* Cambridge University Press., 1989.

[50] L. A. M. van Dongen, R. van der Graaf, and W. H. M. Visscher. Theoretical prediction of electric vehicle energy consumption and battery state-of-charge during arbitrary driving cycles. In *Proceedings of EVC Syposium VI*, 1981. Paper No. EVC 8115.

[51] Schmidt L. F. and Graf J. E. Electric vehicle performance predictions. In *Proceedings of EV Expo 80, Missouri, U.S.A.*, May 20-22, 1980. Paper No. EVC 8057.

[52] H. Shimizu et al. A simulation program for the design and evaluation of an electric vehicle. In *Proceedings of the 10th International Symposium on electric vehicles, EVS-10, Hong-Kong*, pages 345 – 350, 1990.

[53] A. F. Burke and G. H. Cole. Simulation of electric vehicles with hybrid

power systems. In *Proceedings of the 10th International Symposium on electric vehicles, EVS-10, Hong-Kong*, pages 692 – 707, 1990.

[54] A. F. Burke. Comparison of simulation and test for electric vehicles of recent design. *SAE 890818*, 1989.

[55] H. T. Sampson and H. J. Killian. Evaluation of powertrains for hybrid heat engine electric vehicles. *SAE 720194*, 1972.

[56] F. A. Fitz and P. B. Pires. A high torque, high efficiency CVT for electric vehicles. *SAE 910251*, 1991.

[57] Wolfson R.P. and Gower J.H. Modelling and simulation in electric and hybrid vehicle research and development. *IEEE Transactions on Vehicular Technology.*, VT-32(1):62 – 73, 1983.

[58] R. A. Hammond and R. F. Beach. Heavy - a flexible simulation for evaluating electric and hybrid vehicle performance. In *Proceedings of EVC Syposium VI.*, 1981. Paper No. EVC 8114.

[59] H Ishitani et al. A simulation program for the design and evaluation of an electric vehicle. In *Proceedings of the 10th International Symposium on electric vehicles, EVS-10, Hong-Kong*, pages 359 – 370, 1990.

[60] P. W. Masding and J. R. Bumby. Identification and peformance simulation of a hybrid IC-engine/battery electric automotive powertrain. Part 2: The electric traction system. *Transactions of the Institute of Measurement and Control.*, 12(2):97 – 112, 1990.

[61] M. G. Say and E. O. Taylor. *Direct Current Machines.* Pitman, 2nd edition, 1986.

[62] J. Weimer. Messungen an einer Gleichstrommaschine (Nelco). Technical Report 4/91, Universität Kaiserslautern., 28/2/91. (in German).

[63] *Carbon Brushes and Electrical Machines.* Morganite Electrical Carbon Ltd., 1978.

[64] M. E. Hayes. *Current collecting brushes in electrical machines.* Pitman, 1947.

[65] D. Hrovat and W. F. Powers. Powertain computer control systems. In *Selected papers from the 10th Triennial World Congress of the International Ferderationn of Automatic Control, Munich, Germany.*, volume 3, pages 213–219. Pergammon Press, 27-31st July 1987.

[66] Y. Hojo, K. Iwatsuki, H. Oba, and K. Ishikawa. Toyota five-speed automotic transmission with application of modern control theory. *SAE 920610*, 1992.

[67] R. P. Jones and others. Modelling and simulation of an automotive powertrain incorporating a perbury continuously variable transmission. In *Institute of Mechanical Engineers Conference on Integrated Engine/Transmission Systems.*, pages 41 – 49, Bath, U. K., 1986.

[68] Leo A. M. van Dongen. The efficiency characteristics of manual and automatic passenger car transaxles. *SAE 820741*, 1982.

[69] Y. Ko and K. Hosoi. Measurements of power losses in automobile drive train. *SAE 840054*, 1984.

[70] I. Kageyama and H. B. Pacejka. On a new driver model with fuzzy control. In *Proceedings of the 12th IAVSD Symposium on Dynamics of Vehicles on Roads from Tracks. Lyon, France, 26-30th August 1991.*, pages 314 – 324., 1991.

[71] R. A. Hess and A. Modjtahedzadeh. A control theoretic model of driver steering behaviour. *IEEE Control Systems Magazine.*, 10(5):3 – 8, 1990.

[72] J. C. Matthews and R. P. Jones. Optimal feedback control of clutch engagement in an automotive vehicle. In *Proceedings of the second European Control Conference, ECC'93, Groningen, The Netherlands.*, volume 2, pages 986–991, 1993.

[73] S. D. Farrall and R. P. Jones. Energy management in an automotive electric/heat engine hybrid vehicle powertrain using fuzzy decision making. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 463–468, 1993.

[74] A. Kalberlah. Electric hybrid drive systems for passenger cars and taxis. *SAE 910247*, 1991.

[75] R. Bremner. Voltswagen. *Car magazine.*, pages 38–41, June 1991. Special Supplement - The car in the future.

[76] Schneider. Fahrzeug mit hybridantrieb. Technical Report. K/EVE-2545., Robert Bosch., 1973. (in German).

[77] J. R. Bumby and P. W. Masding. A test-bed facility for hybrid ICE-engine/battery electric road vehicle drivetrains., journal = Transactions of the Institute of Measurement and Control., year = 1988, volume = 10, number = 2, pages = 87-97,.

[78] P. W. Masding and J. R. Bumby. Integrated microprocessor control of a hybrid IC engine/battery-electric automotive powertrain. *Transactions of the Institute of Measurement and Control.*, 12(3):128–146, 1990.

[79] J. R. Bumby and I. Forster. Optimisation and control of a hybrid electric car. *Proceedings of the Institution of Electrical Engineers.*, Vol. 134, Pt. D.(No. 6):373–387, 1987.

[80] P. W. Masding, J. R. Bumby, and N. Herron. A microprocessor controlled gearbox for use in electric and hybrid-electric vehicles. *Transactions of the Institute of Measurement and Control.*, 10(4):177–186, 1988.

[81] R. K. McGehee and R. A. Hammond. Comparison of energy management strategies for a heat engine/electric hybrid drive train. In *Proceedings of EVC Syposium VI.*, Baltimore., 1981. Paper No. EVC 8116.

[82] A. F. Burke. A family of hybrid (electric/ice) passenger cars. In *Proceedings of EVC Syposium VI.*, 1981. Paper No. EVC 8144.

[83] A. F. Burke and C. B. Samuah. Computer-aided design of electric and hybrid vehicles. *International Journal of Vehicle Design, Special Publication SP2.*, pages 61–81, 1982.

[84] Assbeck F., Bidan P., Marpinard J.C., and Salut G. Automatic control of a hybrid vehicle. In *Proceedings of the 8th Electric Vehicle Symposium.*, pages 228–234, Washington., 1986.

[85] A. Szumanowski. The model and simulation for determination of minimal energetic parameters of two source propulsion system. In *Proceedings of the 10th International Symposium on electric vehicles, EVS-10, Hong-Kong*, pages 708–721, 1990.

[86] Cambier C. and Gogue G. Engine management for hybrid electric vehicles (an advanced electric drivetrain for EVs). In *Proceedings of EVS 10, Hong Kong '90.*, pages 760 – 776, 1990.

[87] Lotfi Zadeh. Fuzzy sets. *Information and Control.*, 8:338–353, 1965.

[88] Lotfi Zadeh. Fuzzy algorithms. *Information and Control.*, 12:94–102, 1968.

[89] Lotfi Zadeh. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Trans. on Sys., Man and Cybernetics.*, SMC-3 No.1:28–44, 1973.

[90] C. C. Lee. Fuzzy logic in control systems: Fuzzy logic controller - Parts I & II. *IEEE Transactions on Systems, Man and Cybernetics.*, 20(2):404–453, 1990.

[91] M. Sugeno. An introductory survey of fuzzy control. *Information Sciences.*, 36:59–83, 1985.

[92] C. J. Harris, C. G. Moore, and M. Brown. *Intelligent Control: Aspects of Fuzzy Logic and Neural Nets.* World Scientific, 1993.

[93] E. H. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies.*, 7:1–13, 1975.

[94] E. H. Mamdani and N. Baaklini. Prescriptive method for deriving control policy in a fuzzy-logic controller. *Electronic Letters.*, 11:625–626, 1975.

[95] D. A. Rutherford and G. C. Bloore. The implementation of fuzzy algorithms for control. *Proceedings of the IEEE*, 64(4):572–573, April,1976.

[96] G. A. Carter and D. A. Rutherford. A heuristic adaptive controller for a sinter plant. In *Proceedings of 2nd IFAC Symposium on Automation in Mining, Mineral and Metal Processing.*, volume 8, pages 315–324, 1976.

[97] E. H. Mamdani. Advances in the linguistic synthesis of fuzzy controllers. *International Journal of Man-Machine Studies.*, 8:669–678, 1976.

[98] W. J. M. Kickert and H. R. Van Nauta Lemke. Application of a fuzzy controller in a warm water process. *Automatica*, 12:301–308, 1976.

[99] P. J. King and E. H. Mamdani. The application of fuzzy control systems to industrial processes. *Automatica*, 13:235–242, 1977.

[100] E. H. Mamdani and C. P. Pappis. A fuzzy logic controller for a traffc junction. *IEEE Trans. on Sys., Man and Cybernetics.*, SMC-7 No.10:707–717, 1977.

[101] W.J.M. Kickert and E.H. Mamdani. Analysis of a fuzzy logic controller. *Fuzzy sets and systems.*, 1:29–44, 1978.

[102] T. J. Procyk and E. H. Mamdani. A linguistic self-organising process controller. *Automatica.*, Vol. 15:15–30, 1979.

[103] A. S. Cherry. *An Investigation of Multibody System Modelling and Control Analysis Techniques for the Development of Advanced Suspension Systems in Passenger Cars.* PhD thesis, University of Warick., 1992.

[104] J.-J. Østergaard. Fuzzy logic control of a heat exchanger process. In M. M. Gupta, G. N. Saridis, and B. R. Gaines, editors, *Fuzzy Automata and Decision Processes.*, pages 285–320. North Holland, 1977.

[105] T. Yamazaki. *An Improved Algorithm for a Self-Organising Fuzzy Controller.* PhD thesis, Queen Mary College, University of London., 1982.

[106] C. Batur and V Kasparian. A real-time fuzzy self-tuning control. In *Proceedings of the IEEE 8th American Control Conference.*, volume 2, pages 1810–1815, 1989.

[107] G. Bartolini, G. Casalino, F. Davoli, M. Mastretta, R. Minciardi, and E. Morten. Development of performance adaptive fuzzy controllers with application to continuous casting plants. In M. Sugeno, editor, *Industrial Applications of Fuzzy Control.*, pages 73–86. Elsevier Science Publishers B. V., 1985.

[108] P. Albertos. Fuzzy controllers. In L. Boullart, A. Krijgsman, and R. A. Vingerhoeds, editors, *Application of Artificial Intelligence in Process Control.*, pages 343–367. Pergammon Press., 1992.

[109] N. Vijeh. *Microprocessor Engineering Aspects of a Self-organising Fuzzy Logic Controller.* PhD thesis, University of Warwick., 1988.

[110] T. Yamazaki and M. Sugeno. A microprocessor based fuzzy controller for industrial purposes. In M. Sugeno, editor, *Industrial applications of fuzzy control.*, pages 231 – 239. Elsevier Science Publishers B.V. (North-Holland)., 1985.

[111] M. Brown, D. J. Mills, and C. J. Harris. The representation of fuzzy algorithms used in adaptive modelling and control schemes. Technical report, Department of Aeronautics and Astronautics, University of Southampton., 1993.

[112] Z. Cao and A. Kandel. Applicability of some fuzzy implication operators. *Fuzzy sets and systems.*, 31(2):151–186, 1989.

[113] D. Schwarz. Fuzzy inference in a formal theory of semantic equivalence. *Fuzzy sets and systems.*, 31(2):205–216, 1989.

[114] P. Constancis and C. J. Harris. A comparison between fuzzy, PI and B-spline control. In *Proceedings of the 1992 IEEE International Symposium on Intelligent Control, Glasgow, UK.*, pages 372–373, 1992.

[115] L. Kourra and Y. Tanaka. Dedicated silicon solutions for fuzzy logic systems. In *IEE Colloquium "Two Decades of Fuzzy Control - Part 1"., London, U.K.*, 1993. Digest No. 1993/166.

[116] T. Yamakawa. Stabilization of an inverted pendulum by a high speed fuzzy logic controller hardware system. *Fuzzy Sets and Systems.*, 32(2):161–180, 1989.

[117] A. MacLeod. When maths fails fuzzy logic rules. *New Electronics.*, pages 19–22, April, 1992.

[118] I. Radivojevic, J. Herath, and W. S. Gray. High performance DSP architectures for intelligence and control applications. *IEEE Control Systems Magazine.*, 11(4):49 – 55, 1991.

[119] Integrated Systems Incorporated. An introduction to the knowledge based modules: RT/Expert and RT/Fuzzy. In *Seminar notes for "Systems Design, Simulation and Real Time Prototyping Using Graphics Modelling.", Coventry, U.K.* 1992.

[120] G. Saridis. *Self-Organising Control of Stochastic Systems.* Marcel Dekker, New York., 1977.

[121] K. J. Åstrom and B. Wittenmark. *Adaptive Control.* Addison-Wesley, 1989.

[122] P. Hanks, editor. *Collins Dictionary of the English Language.* Collins.

[123] J. A. Farrell and W. L. Baker. An introduction to learning control systems - workshop notes. In *The 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, 1993.

[124] T. J. Procyk. *A Self-Organising Controller for Dynamic Processes.* PhD thesis, Queen Mary College, University of London., 1977.

[125] E. H. Mamdani. Application of fuzzy logic to approximate reasoning using linguistic synthesis. *IEEE. Transactions on Computers.*, C-36(12):1182–1191., 1977.

[126] S. Daley and K. F. Gill. A design study of a self-organising fuzzy controller. *Proceedings of the Institution of Mechanical Engineers.*, 200(C1):59–69, 1986.

[127] S. Daley and K. F. Gill. Attitude control of a spacecraft using an extended self-organising fuzzy logic controller. *Proceedings of the Institution of Mechanical Engineers.*, 201(C2):97–106, 1987.

[128] R. Tanscheit and E. M. Scharf. Experiments with the use of a rule-based self-organising controller for robotics applications. *Fuzzy Sets and Systems.*, 26(2):195–214, 1988.

[129] S. Shao. Fuzzy self-organising controller and its application for dynamic processes. *Fuzzy Sets and Systems.*, 26(2):151 – 164, 1988.

[130] D. A. Linkens and M. F. Abbod. Self-organising fuzzy logic control for real time processes. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 2, pages 971–976, 1991.

[131] B. N. Farbrother, B. A. Stacey, and R. Sutton. Fuzzy self-organising control of a remotely operated submersible. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 1, pages 499–504, 1991.

[132] D. A. Linkens and S. B. Hasnain. Self-organising fuzzy logic control and application to muscle relaxant anaesthesia. *Proceedings of the IEE.*, 138(3, Pt. D):274–284, 1991.

[133] K. Sugiyama. *Analysis and synthesis of the rule based self-organising controller.* PhD thesis, Queen Mary College, London University., 1986.

[134] R. Sutton and I. M. Jess. Real-time application of a self-organising autopilot to warship yaw control. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 2, pages 827–831, 1991.

[135] E. M. Scharf, N. J. Mandic, and E. H. Mamdani. A self-organising algorithm for the control of a robot arm. In *Proceedings of the 23rd International Symposium on Mini and Micro Computers and their Applications, San Antonio, Texas.*, pages 8–13, 1983.

[136] C. G. Moore and C. J. Harris. Indirect adaptive fuzzy control. *International Journal of Control.*, 56(2):441–468, 1992.

[137] M. Brown, R. Fraser, C. J. Harris, and C. G. Moore. Intelligent self-organsing controllers for autonomous guided vehicles comparative aspects of fuzzy logic and neural nets. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 1, pages 134–139., 1991.

[138] B. P. Graham and R. B. Newell. Fuzzy adaptive control of a first-order process. *Fuzzy Sets and Systems.*, 31:47–65, 1989.

[139] M. Togai and P. P. Wang. Analysis of a fuzzy dynamic system and synthesis of its controller. *International Journal of Man-Machine Studies.*, 22:355–363, 1985.

[140] M Togai. Application of fuzzy inverse relations to synthesis of a fuzzy controller for dynamic systems. In *Proceeding of 23rd Conference on Decision and Control, Las Vegas*, pages 904–905, 1984.

[141] B. Kosko. *Neural Netwroks and Fuzzy Systems: A Dynamical Systems Approach to Machine Intelligence.* Prentice-Hall International Editions., 1992.

[142] W. Pedrycz. An identification algorithm in fuzzy relational systems. *Fuzzy Sets and Systems.*, 13:153–167, 1984.

[143] R. Jager, H. B. Verbruggen, P.M. Bruijin, and A. J. Krijgsman. Real-time fuzzy expert control. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 2, pages 966–970, 1991.

[144] E. Czogała and W. Pedrycz. Control problems in fuzzy systems. *Fuzzy Sets and Systems.*, 7:257–273, 1982.

[145] E. Czogała and W. Pedrycz. On identification in fuzzy systems and its applications in control problems. *Fuzzy Sets and Systems.*, 6:73–83, 1981.

[146] C-W. Xu and Y-Z. Lu. Fuzzy model identification and self learning for dynamic systems. *IEEE Transactions on Systems, Man and Cybernetics.*, SMC-17(4):683–689, 1987.

[147] T. Takagi and M. Sugeno. Fuzzy identification of systems and its applications to modelling and control. *IEEE Trans. on Sys., Man and Cybernetics.*, SMC-15(1):116–132, 1985.

[148] A. Kandel. On the control and evaluation of uncertain processes. *IEEE Transactions on Automatic Control.*, AC-25(6):1182–1187, 1980.

[149] Z. Shen, L. Ding, and M. Mukaidono. A self-learning approach for fuzzy simulating model. In *Proceedings of the International Conference on Fuzzy Logic and Neural Networks, Iizuka, Japan.*, pages 223–229, 1990.

[150] R. E. Bellman and L. A. Zadeh. Decision making in a fuzzy environment. *Management Science.*, 17:B-141–B-165, 1970.

[151] P. E. An, M. Brown, C. J. Harris, A. J. Lawrence, and C. G. Moore. Comparative aspects of associative memory networks for modelling. In *Proceedings of the second European Control Conference, ECC '93, Groningen, The Netherlands.*, volume 2, pages 454–459, 1993.

[152] J. Nie and D. A. Linkens. Neural network based approximate reasoning: Principles and implementation. *International Journal of Control.*, 56(2):399–413, 1992.

[153] Y-C. Chien and C-C. Teng. Fuzzy modelling using neural networks. In *Proceedings of the second European Control Conference, ECC '93, Groningen, The Netherlands.*, volume 2, pages 503–508, 1993.

[154] J. S. Albus. A new approach to manipulator control: The cerebellar model articulation controller (CMAC). *Transactions of the ASME, Journal of Dynamic Systems Measurement and Control.*, 97(3):220–227, 1975.

[155] J. S. Albus. Data storage in the the cerebellar model articulation controller (CMAC). *Transactions of the ASME, Journal of Dynamic Systems Measurement and Control.*, 97(3):228–233, 1975.

[156] S. Chen and S. A. Billings. Neural networks for nonlinear dynamics system modelling and identification. *International Journal of Control.*, 56(2):319–346, 1992.

[157] A. Krijgsman. Associative memories: The CMAC approach. In L. Boullart, A. Krijgsman, and R. A. Vingerhoeds, editors, *Application of Artificial Intelligence in Process Control.*, pages 403–421. Pergammon Press., 1992.

[158] Y. Jin, A. G. Pipe, and A. Winfield. Stable neural control of discrete systems. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 110–115, 1993.

[159] G. Calcev. A self-tuning neurofuzzy controller. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 577–581, 1993.

[160] M. Majors, J. Stori, and D. Cho. An adaptive neural network control method for automotive fuel injection systems. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 104–109, 1993.

[161] T. Kohonen. Learning vector quantisation for pattern recognition. Technical Report TKK-F-A601, Helsinki University of Technology., 1986.

[162] T. Kohonen. *Self-Organisation and Associative Memory.* Springer-Verlag., 1989.

[163] T. Yamaguchi, T Takagi, and T. Mita. Self-organising control using fuzzy neural networks. *International Journal of Control.*, 56(2):415–439, 1992.

[164] T. Yamaguchi, M. Tanabe, K. Kuriyama, and T. Mita. Fuzzy adaptive control with an associative memory system. In *Proceedings of the IEE International Conference Control '91, Edinburgh, U.K.*, volume 2, pages 944–948, 1991.

[165] J. Pacini and B. Kosko. Adaptive fuzzy systems for target tracking. *Intelligent Systems Engineering.*, 1:3–21, 1992.

[166] V. I. Utkin. Variable structure systems with sliding modes. *IEEE Transactions on Automatic Control.*, AC-22:212–222, 1977.

[167] R. Palm. Sliding mode fuzzy control. In *Proceedings of the IEEE '92 International Conference on Fuzzy Systems. San Diego, USA.*, pages 519–526, 1992.

[168] J. J. E. Slotine and S. S. Sastry. Tracking control non-linear system using sliding surfaces with application to robot manipulators. *International Journal of Control.*, 38(2):465–492, 1983.

[169] R. Vepa and A. Nowe. Desgin of fuzzy learning compensators and controllers for autonomous, redundant robots manipulators. In *Proceedings of the 1st International Workshop on Intelligent Autonomous Vehicles, IAV-'93, Southampton, UK.*, pages 312–317, 1993.

[170] R. Vepa. Fuzzy learning based on variable structure control systems theory. In *IEE Computing and Control Division Colloquium on "Two Decades of Fuzzy Control - Part 2".*, pages 5/1–5/4, 1992.

[171] S-K. Nam, M-H. Lee, and W-S. Yoo. Predictive sliding control with fuzzy-logic for fuel injected automotive engines. *Proceedings of the Institution of Mechanical Engineers, Journal of Systems and Control Engineering.*, 206:237–244, 1992.

[172] G. N. Saridis and H. E. Stephanou. A hierarchical approach to the control of a prosthetic arm. *IEEE Trans. on Sys., Man and Cybernetics.*, Vol-SMC7(No. 6):407 – 420, 1977.

[173] G. N. Saridis and C-S. G. Lee. An approximation theory of optimal control for trainable manipulators. *IEEE Trans. on Sys., Man and Cybernetics.*, Vol. SMC-9(No. 3.):152–159, 1979.

[174] W. G. Wee and K. S. Fu. A formulation of fuzzy automata and its application as a model of learning systems. *IEEE Transactions on Systems Science and Cybernetics.*, SSC-5:215–233, 1969.

[175] J-H. Kim, J-H. Park, and S-W. Lee. Control of systems with deadzones using PD controllers with fuzzy precompensation. In *Proceedings of the*

1993 IEEE International Symposium on Intelligent Control, Chicago, USA., pages 451–456, 1993.

[176] M. F. Yeung and K. W. Sum. An on-line intelligent control scheme for tension control. In *Proceedings of the IEE 1st Conference on Intelligent Systems Engineering.*, pages 244–249, 1992.

[177] A. Ollero and A. J. Garcia-Cerezo. Direct digital control, auto-tuning and supervision using fuzzy control. *Fuzzy Sets and Systems.*, Vol. 30.:135–153, 1989.

[178] J. Litt. An expert system to perform on-line controller tuning. *IEEE Control Systems Magazine.*, 11(3):18–23, 1991.

[179] D. Mallampati and S. Shenoi. Self-organising fuzzy logic control. In R. Sharda, editor, *Knowledge-Based Systems and Neural Networks.*, pages 271–282. Elsevier Science Publishing Co. Inc., 1991.

[180] W. H. Bare, R. J. Mullholland, and S. S. Sofer. Design of a self-tuning expert contoller for a gasoline refinery catalytic reformer. In *Proceedings of the American Control Conference.*, pages WA8 – 11:15, 1988.

[181] Y. Murayama, T. Terano, S. Masui, and N. Akiyama. Optimising control of a diesel engine. In M. Sugeno, editor, *Industrial Applications of Fuzzy Control.*, pages 63–71. Elsevier Science Publishers B. V., 1985.

[182] D. E. Goldberg. *Genetic algorithms in search, optimisation and machine learning.* Addison-Wesley, 1989.

[183] C. L. Karr. Design of an adaptive fuzzy logic controller using a genetic algorithm. In *Proceedings of 4th International Conference on Genetic Algorithms, San Diego, USA.*, pages 450–457, 1991.

[184] W. Foslien and T. Samad. Fuzzy controller synthesis with neural network process models. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 370–375, 1993.

[185] A. Tacillo, V. Skormin, J. Crisman, and N. Bourbakis. Intelligent control of a robotic hand with neural nets and fuzzy sets. In *Proceedings of the 1993 IEEE International Symposium on Intelligent Control, Chicago, USA.*, pages 232–237, 1993.

[186] C. S. Hsu. A discrete method of optimal control based upon the cell state space method. *Journal of Optimisation Theory and Applications.*, 46(4):547–569, 1985.

[187] S. M. Smith and D. J. Corner. Automated calibration of a fuzzy logic controller using a cell state space algorithm. *IEEE Control Systems Magazine.*, 11(5):18–28, 1991.

[188] Y. Y. Chen and T. C. Tsao. A description of the dynamical behaviour of fuzzy systems. *IEEE Transactions on Sys., Man and Cybernetics.*, SMC-19(4):745–755, 1989.

[189] B. Widrow and S. D. Stearns. *Adaptive Signal Processing.* Prentice-Hall, 1985.

[190] H. Nomura, I. Hayashi, and N. Wakami. A self-tuning method of fuzzy control by descent method. In *Proceedings of the IFSA Conference, Brussels, Belgium.*, pages 155–158, 1991.

[191] R. Vepa and A. Nowe. A reinforcement algorithm based on safety. In *Fuzzy Logic in artificial intelligence, FLAI'93, Linz, Austria.*, pages 47–58, 1993.

[192] C. Watkins. *Learning from delayed rewards.* PhD thesis, Cambridge University, U.K., 1989.

[193] D. Lakov. Adaptive robot under fuzzy control. *Fuzzy Sets and Systems.*, 17:1–8, 1985.

[194] A. F. Blishun. Fuzzy adaptive learning model of decision making process. *Fuzzy Sets and Systems.*, 18:273–282, 1986.

[195] D. A. Handelman and R. F. Stengel. Rule based mechanisms of learning for intelligent adaptive flight control. In *Proceedings American Control Conference.*, pages Paper No. WA7 – 10:45, 1988.

[196] I. Borges de Silva and G. Oliver. Real time fuzzy adaptive controller for an asymmetrical four quadrant power converter. In *Proceedings of the IEEE Industrial Applications Society Annual Meeting, Atlanta, USA.*, pages 872–878, 1987.

[197] T. R. Andersen and S. B. Nielsen. An efficient single output fuzzy control algorithm for adaptive applications. *Automatica.*, 21(5):539–545, 1985.

[198] Xian-Tu Peng. Generating rules for fuzzy logic controllers by functions. *Fuzzy Sets and Systems.*, 36:83–89, 1989.

[199] S. Nakaniwa, J. Furuya, and N. Tomisawa. Development of nest structured learning control system. *SAE 910084*, 1991.

[200] T. Manaka. Control apparatus for internal combustion engine. *European Patent Specification, Publication Number 0 281 962 B1*, 1991.

[201] M. S. Younger. *A Handbook for Linear Regression.* Duxbury Press., 1979.

[202] Department of Transport. *Transport Statistics Great Britain.* HMSO Publications., 1993.

Appendix A

Data used in engine model.

A.1 Engine torque data.

Throttle angle	Engine model output torque in Nm. Engine speed in r/min.											
	1000	1250	1500	1750	2000	2250	2500	2750	3000	3250	3500	3750
0°	0.0	-2.5	-5.0	-7.5	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0
5°	19.6	8.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10°	71.5	62.2	51.7	43.8	37.1	30.2	24.2	18.4	15.3	9.8	9.0	8.5
15°	82.8	84.9	81.3	76.9	72.5	68.8	58.9	54.2	49.1	44.0	41.5	36.3
20°	85.3	93.3	93.4	93.6	92.8	92.8	88.0	86.6	81.3	77.4	77.4	75.1
25°	87.7	92.2	93.8	97.5	98.6	102.1	100.1	99.7	95.1	92.5	95.4	94.7
30°	90.1	93.9	98.1	101.3	103.9	104.7	104.7	104.7	102.2	101.9	102.7	101.5
35°	90.8	98.7	102.4	104.1	104.6	104.7	105.0	106.1	106.5	104.9	105.5	105.9
40°	90.8	97.6	100.6	104.7	106.3	106.8	108.3	110.0	109.4	107.4	107.7	107.9
45°	90.8	95.5	101.3	104.8	106.1	108.2	111.0	114.0	110.8	109.9	110.0	110.8
50°	90.8	97.1	103.0	104.8	106.0	106.8	108.7	115.4	112.5	111.6	112.2	113.9
55°	90.7	99.0	102.7	106.0	106.8	107.5	110.9	114.7	113.6	111.6	112.4	113.2
60°	90.7	96.9	101.3	105.5	106.6	108.0	110.8	113.3	113.7	112.2	112.3	113.1
65°	90.6	96.5	102.6	105.0	105.6	107.5	112.9	112.7	113.9	112.9	112.1	114.1
70°	90.7	99.4	103.9	105.9	107.2	108.2	109.8	117.7	115.0	111.8	113.0	114.6
75°	90.8	97.9	101.8	106.3	107.3	107.9	112.3	115.3	114.8	112.3	112.7	114.7
80°	90.9	96.4	102.4	105.4	105.9	107.8	110.0	113.0	114.5	112.9	112.5	114.8
85°	90.9	97.4	103.4	104.6	105.1	107.8	110.3	112.4	114.4	112.4	112.3	114.9
90°	95.9	102.4	108.4	109.6	110.1	112.5	114.3	118.4	116.4	116.5	116.5	116.4
	4000	4250	4500	4750	5000	5250	5500	5750	6000	6250	6500	
0°	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	-10.0	
5°	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-2.0	
10°	8.2	7.9	7.6	7.3	7.0	6.7	6.4	6.1	5.8	5.5	5.2	
15°	31.2	27.1	23.1	19.4	17.2	16.5	16.5	16.0	15.6	15.3	15.0	
20°	69.7	67.0	62.3	57.3	54.3	48.3	44.6	38.0	34.8	35.1	32.6	
25°	92.6	92.4	89.5	86.1	83.2	80.8	76.4	71.2	67.1	63.8	60.8	
30°	100.1	103.0	104.0	102.2	99.4	90.4	93.8	89.6	84.4	80.7	76.0	
35°	105.9	107.8	108.5	107.7	106.5	105.3	102.1	98.2	93.7	88.1	82.5	
40°	109.2	111.1	112.5	111.5	110.7	109.3	107.1	103.6	98.3	92.9	85.6	
45°	112.5	114.3	115.9	115.7	114.9	113.1	110.3	106.1	101.1	95.5	88.6	
50°	114.2	116.1	116.4	116.3	116.2	115.8	112.2	108.0	103.5	97.8	90.0	
55°	114.2	115.9	117.2	117.0	116.6	115.5	114.6	109.9	104.9	99.2	91.0	
60°	115.0	116.5	117.9	117.4	116.5	115.8	115.3	110.7	105.4	99.9	92.0	
65°	115.8	117.3	117.3	117.4	117.9	117.7	115.2	111.2	106.2	100.6	92.6	
70°	115.8	116.8	116.8	117.4	117.9	118.0	115.0	111.0	107.0	101.0	93.7	
75°	115.6	116.2	116.6	117.4	117.9	117.6	114.8	111.8	106.5	101.0	95.0	
80°	115.3	115.9	116.6	117.4	118.1	117.2	114.4	111.9	106.5	100.8	95.7	
85°	115.1	115.8	116.6	117.4	118.6	116.8	115.1	112.0	106.7	100.6	95.8	
90°	116.4	118.1	118.0	117.4	119.1	118.1	114.6	112.0	108.7	100.6	95.8	

A.2 Engine fuel flow rate data.

Engine model fuel flow rate in g/s.												
Throttle angle	Engine speed in r/min.											
	1000	1250	1500	1750	2000	2250	2500	2750	3000	3250	3500	3750
0°	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.00	0.00	0.00	0.00	0.00
5°	0.25	0.22	0.20	0.21	0.22	0.22	0.22	0.20	0.21	0.22	0.27	0.27
10°	0.71	0.69	0.63	0.66	0.69	0.68	0.69	0.64	0.67	0.69	0.74	0.74
15°	0.95	0.83	0.91	1.00	1.09	1.17	1.19	1.24	1.25	1.27	1.34	1.34
20°	0.92	0.98	1.05	1.18	1.30	1.48	1.59	1.73	1.78	1.89	2.05	2.17
25°	0.88	0.94	1.08	1.25	1.41	1.64	1.80	1.98	2.05	2.17	2.42	2.59
30°	0.84	0.98	1.14	1.32	1.51	1.71	1.90	2.10	2.24	2.40	2.58	2.74
35°	0.83	1.02	1.19	1.37	1.55	1.73	1.91	2.14	2.36	2.54	2.69	2.85
40°	0.84	1.02	1.20	1.39	1.60	1.78	2.00	2.22	2.42	2.59	2.76	2.93
45°	0.85	1.03	1.21	1.40	1.60	1.82	2.06	2.31	2.46	2.64	2.82	3.03
50°	0.85	1.03	1.22	1.41	1.60	1.80	2.01	2.33	2.52	2.71	2.91	3.12
55°	0.85	1.01	1.22	1.42	1.62	1.82	2.08	2.34	2.55	2.76	2.94	3.13
60°	0.86	1.04	1.23	1.42	1.64	1.88	2.11	2.34	2.56	2.75	2.93	3.14
65°	0.86	1.04	1.23	1.41	1.60	1.84	2.12	2.34	2.56	2.75	2.93	3.16
70°	0.86	0.99	1.22	1.42	1.62	1.82	2.04	2.38	2.59	2.80	2.99	3.18
75°	0.86	1.03	1.24	1.43	1.66	1.86	2.14	2.40	2.64	2.82	3.00	3.21
80°	0.86	1.05	1.24	1.42	1.62	1.89	2.16	2.43	2.68	2.85	3.01	3.23
85°	0.86	1.04	1.23	1.41	1.60	1.88	2.16	2.44	2.71	2.87	3.02	3.25
90°	0.86	1.04	1.24	1.42	1.62	1.99	2.19	2.48	2.67	2.88	3.10	3.33
	4000	4250	4500	4750	5000	5250	5500	5750	6000	6250	6500	
0°	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
5°	0.25	0.25	0.23	0.22	0.24	0.26	0.26	0.31	0.36	0.45	0.53	
10°	0.72	0.72	0.70	0.69	0.71	0.73	0.73	0.78	0.83	0.92	1.00	
15°	1.32	1.32	1.30	1.29	1.31	1.33	1.33	1.38	1.43	1.52	1.60	
20°	2.21	2.29	2.31	2.33	2.39	2.34	2.39	2.32	2.38	2.70	2.78	
25°	2.73	2.82	3.01	3.16	3.26	3.44	3.58	3.66	3.85	4.09	4.30	
30°	2.91	3.19	3.48	3.77	4.02	3.91	4.41	4.57	4.75	5.02	5.20	
35°	3.06	3.38	3.67	4.02	4.37	4.68	4.92	5.12	5.32	5.52	5.71	
40°	3.18	3.52	3.84	4.19	4.55	4.87	5.18	5.45	5.66	5.87	5.96	
45°	3.30	3.61	3.99	4.36	4.72	5.04	5.34	5.60	5.83	6.06	6.20	
50°	3.36	3.67	4.03	4.40	4.78	5.16	5.45	5.71	5.96	6.22	6.29	
55°	3.36	3.71	4.06	4.43	4.81	5.20	5.60	5.81	6.04	6.29	6.36	
60°	3.39	3.70	4.09	4.46	4.84	5.24	5.63	5.81	6.06	6.31	6.41	
65°	3.42	3.68	4.08	4.47	4.84	5.23	5.58	5.86	6.08	6.34	6.46	
70°	3.44	3.73	4.13	4.49	4.86	5.22	5.53	5.81	6.10	6.36	6.54	
75°	3.45	3.78	4.15	4.50	4.87	5.22	5.53	5.85	6.10	6.35	6.59	
80°	3.46	3.81	4.16	4.51	4.88	5.21	5.53	5.86	6.10	6.33	6.64	
85°	3.48	3.83	4.18	4.53	4.89	5.21	5.54	5.86	6.11	6.32	6.63	
90°	3.55	3.82	4.20	4.54	4.89	5.22	5.52	5.86	6.11	6.33	6.64	

Appendix B

Pseudo-code implementation of fuzzy logic algorithms.

B.1 Max-min inference, centroid of fuzzy union defuzzification.

```
/*
 *there are ni rules
 *associated with each rule, i, there is an input and an
 *output fuzzy set
 *the 1-dimensional arrays input_uni and output_uni define
 * discrete universes of discourse over the input and output
 *variables, they have dimension nj and nk
 *an array, input_sets of dimension (ni,nj) defines the
 *grades of membership of the input fuzzy sets, the input
 *fuzzy set i is defined over the i'th row of the matrix
 *an array, output_sets of dimension (ni,nk) defines the
 *grades of membership of the output fuzzy sets
 *an array, inp_grds of dimension ni stores the grades of
 *membership of the input value in each of the input fuzzy
 *sets
 *temp is used as a scratch pad in the calculation of the
 *output fuzzy set value
 *output_fset stores the values in the output fuzzy set as
 *it is created, note there is no need to store the entire
 *set
 */

for i = 1, ni /*for each rule*/
    /*
```

```
     *obtain grade of membership of input in input fuzzy
     *set for current rule
     *find indices of elements of input_uni that bracket
     *current input value, store value in index
     */
    call getind(inp_val, input_uni, index)

    /*
     *linearly interpolate within area found by getind,
     *returning result in vector inp_grds
     */

    call linearinterp(inp_val,inp_uni(index),
                      input_sets(i,index), inp_grds(i))
end

/*
 *preset total area sum to zero
 */
area_sum = 0

/*
 *preset moment of area sum to zero
 */
momarea_sum = 0

for k = 1, nk /*for each element of the output fuzzy set*/
    /*
     *preset output fuzzy set at this element value to
     *zero
     */
    output_fset = 0

    for i = 1, ni /*for each rule*/
        /*
         *find fuzzy intersection of output set i, at
         *element k and grade of membership of input in
         *input fuzzy set
         */
        temp = min(inp_grds(i),output_sets(i,k))

        /*
         *find fuzzy union of output set i at element k,
```

```
        *and output sets from other rules
        */
       output_fset = max(temp, output_fset(k))
    end

    /*
     *increment area sum
     */
    area_sum = area_sum + output_fset
    /*
     *increment moment of area sun
     */
    momarea_sum = momarea_sum + output_fset*output_uni(k)
end

/*
 *finally defuzzify by dividing the moment of area by the
 *area
 */
output = momarea_sum/area_sum

end
```

Appendix C

Early simulation results.

Throttle offsets						Armature current offsets						m2egr	fuel cons.[1]	Δ Batt.[2]
0	4	12	24	40	90	0	0	0	0	0	0	1	9.111	0.000
0	4	12	24	40	90	0	0	0	33	66	165	1	9.110	-0.042
0	4	12	24	40	90	0	0	33	66	99	165	1	8.504	-0.773
0	4	12	24	40	90	0	33	66	99	132	165	1	6.927	-1.488
0	4	12	24	40	90	0	0	0	33	66	165	2	9.208	-0.039
0	4	12	24	40	90	0	0	33	66	99	165	2	8.239	-0.676
0	4	12	24	40	90	0	33	66	99	132	165	2	6.033	-1.210
0	4	12	12	40	90	0	0	0	33	66	165	1	9.040	-0.095
0	4	12	12	40	90	0	0	0	33	66	165	2	9.102	-0.086
0	4	12	12	40	90	0	0	0	66	99	165	1	9.018	-0.105
0	4	12	12	40	90	0	0	0	66	99	165	2	9.088	-0.088
0	4	4	12	40	90	0	0	33	99	132	165	1	6.483	-1.648
0	4	4	12	40	90	0	0	33	99	132	165	2	5.759	-1.287
0	4	4	12	40	90	0	0	66	99	132	165	1	5.949	-1.878
0	4	4	12	40	90	0	0	66	99	132	165	2	5.306	-1.415
0	4	4	12	40	90	0	0	99	132	132	165	1	5.430	-2.141
0	4	4	12	40	90	0	0	99	132	132	165	2	4.946	-1.521
0	0	0	0	0	0	0	33	66	99	132	165	1	4.735	-2.533
0	0	0	0	0	0	0	33	66	99	132	165	2	4.734	-1.618

[1] Fuel consumption in l/100km.
[2] Change in battery state of charge in Ah/km.

Throttle offsets						Armature current offsets						m2egr	fuel cons.	Δ Batt.
0	4	12	24	40	90	0	33	33	66	99	165	1	7.462	-1.232
0	4	12	24	40	90	0	33	33	66	99	165	2	6.576	-1.062
0	0	4	12	40	90	0	33	66	99	132	165	1	5.924	-1.912
0	0	4	12	40	90	0	33	66	99	132	165	2	5.284	-1.443
0	4	12	24	40	90	0	0	0	-33	0	165	1	9.111	0.000
0	4	12	24	40	90	0	0	0	-33	0	165	2	9.209	0.000
0	4	12	24	40	90	0	0	-33	-66	0	165	1	9.774	0.032
0	4	12	24	40	90	0	0	-33	-66	0	165	2	12.181	0.384
0	4	12	24	40	90	0	0	-33	-99	0	165	1	9.849	0.051
0	4	12	24	40	90	0	0	-33	-99	0	165	2	12.700	0.486
0	4	12	24	40	90	0	0	-66	-99	0	165	1	11.770	0.300
0	4	12	24	40	90	0	0	-66	-99	0	165	2	17.273	1.359
0	4	12	24	40	90	0	-33	-66	-99	0	165	1	12.624	0.388
0	12	24	24	40	90	0	-33	-99	-33	0	165	1	9.858	0.054
0	12	24	24	40	90	0	-33	-99	-33	0	165	2	13.283	0.538
0	12	24	24	40	90	0	-66	-99	-33	0	165	1	12.429	0.359
0	12	24	24	40	90	0	-99	-132	-66	0	165	1	14.747	0.784
0	12	24	24	40	90	0	-99	-99	-66	0	165	1	14.438	0.725
0	12	24	24	40	90	0	-33	-66	0	0	165	1	9.780	0.036
0	12	24	24	40	90	0	-33	-66	0	0	165	2	12.877	0.450
0	12	24	24	40	90	0	-66	-66	-33	0	165	1	12.296	0.329
0	12	24	24	40	90	0	-33	-33	0	0	165	1	9.697	0.023
0	12	24	24	40	90	0	-33	-33	0	0	165	2	12.505	0.363

Appendix D

Data used in reference maps.

D.1 Engine reference torque data.

Throttle angle	Engine reference output torque in Nm. Engine speed in r/min.											
	1000	1250	1500	1750	2000	2250	2500	2750	3000	3250	3500	3750
0.0°	-4.4	-6.5	-7.5	-8.2	-8.6	-8.9	-9.1	-9.3	-9.4	-9.5	-9.6	-9.6
2.5°	4.4	1.2	-0.6	-1.8	-2.7	-3.3	-3.8	-4.2	-4.5	-4.8	-5.0	-5.2
5.0°	19.8	13.8	10.0	7.3	5.4	4.0	2.8	1.9	1.1	0.5	-0.1	-0.5
7.5°	40.7	31.5	24.7	19.9	16.4	13.6	11.5	9.7	8.2	7.0	5.9	5.0
10.0°	61.6	51.6	43.1	36.3	30.6	26.2	22.6	19.7	17.3	15.3	13.5	12.0
12.5°	78.2	70.6	62.6	54.8	47.9	41.9	37.0	32.7	29.1	26.1	23.5	21.2
15.0°	85.7	83.4	78.8	72.9	66.5	60.4	54.5	49.0	44.1	39.9	36.5	33.4
17.5°	88.4	87.9	86.8	85.0	81.7	77.7	73.3	68.2	62.9	57.7	53.0	48.9
20.0°	90.1	90.6	90.6	90.0	89.1	88.2	87.2	84.7	80.9	76.3	71.5	67.1
22.5°	91.3	92.4	93.1	93.3	93.3	93.4	93.5	93.0	91.5	89.2	86.0	82.6
25.0°	92.1	93.6	94.8	95.7	96.2	97.0	97.8	98.0	97.3	95.7	93.5	91.3
27.5°	92.7	94.5	96.1	97.3	98.3	99.5	100.9	101.6	101.4	100.4	98.7	97.1
30.0°	93.2	95.2	97.0	98.6	99.8	101.4	103.1	104.2	104.5	103.8	102.6	101.4
32.5°	93.5	95.7	97.8	99.5	101.0	102.9	104.8	106.2	106.8	106.4	105.5	104.5
35.0°	93.8	96.1	98.3	100.2	101.9	104.0	106.2	107.8	108.5	108.4	107.7	107.0
37.5°	94.0	96.5	98.7	100.8	102.6	104.8	107.2	108.9	109.9	109.9	109.4	108.8
40.0°	94.2	96.7	99.1	101.2	103.2	105.5	108.0	109.9	110.9	111.1	110.7	110.3
50.0°	94.7	97.4	100.0	102.4	104.6	107.2	110.0	112.2	113.5	113.9	113.8	113.7
60.0°	94.9	97.7	100.5	103.0	105.3	108.0	110.9	113.2	114.7	115.3	115.3	115.2
70.0°	95.1	97.9	100.7	103.3	105.7	108.5	111.5	113.8	115.4	116.0	116.0	116.0
80.0°	95.2	98.1	100.9	103.5	106.0	108.8	111.8	114.2	115.7	116.4	116.4	116.4
90.0°	95.2	98.1	101.0	103.6	106.1	109.0	112.0	114.4	116.0	116.6	116.7	116.7
	4000	4250	4500	4750	5000	5250	5500	5750	6000	6250	6500	
0.0°	-9.7	-9.7	-9.8	-9.8	-9.8	-9.8	-9.8	-9.9	-9.9	-9.9	-9.9	
2.5°	-5.3	-5.5	-5.6	-5.7	-5.8	-5.9	-6.0	-6.0	-6.1	-6.2	-6.2	
5.0°	-0.9	-1.3	-1.6	-1.9	-2.2	-2.4	-2.6	-2.8	-3.0	-3.2	-3.3	
7.5°	4.2	3.4	2.8	2.2	1.7	1.3	0.8	0.5	0.1	-0.2	-0.5	
10.0°	10.7	9.5	8.4	7.5	6.7	5.9	5.2	4.6	4.0	3.4	2.9	
12.5°	19.3	17.5	15.9	14.6	13.3	12.2	11.1	10.2	9.3	8.5	7.8	
15.0°	30.6	28.2	26.0	24.1	22.4	20.8	19.3	18.0	16.8	15.7	14.7	
17.5°	45.3	42.2	39.4	36.9	34.6	32.4	30.5	28.7	26.9	25.0	22.9	
20.0°	63.3	60.2	57.4	54.8	52.1	49.0	45.7	42.8	39.9	36.5	32.5	
22.5°	79.7	77.4	75.5	73.4	70.7	66.7	62.0	57.9	54.0	49.5	44.4	
25.0°	89.5	88.3	87.4	86.2	84.0	80.0	75.1	70.9	67.0	62.5	57.8	
27.5°	95.9	95.2	94.9	94.2	92.4	88.9	84.4	80.5	76.9	72.8	68.2	
30.0°	100.5	100.3	100.3	100.0	98.6	95.5	91.3	87.7	84.3	80.4	76.0	
32.5°	104.0	104.1	104.4	104.4	103.2	100.3	96.4	93.0	89.8	86.0	81.7	
35.0°	106.7	106.9	107.5	107.6	106.7	104.0	100.1	96.9	93.8	90.1	85.9	
37.5°	108.7	109.1	109.8	110.1	109.3	106.7	103.0	99.8	96.8	93.1	89.0	
40.0°	110.3	110.8	111.6	112.0	111.3	108.8	105.1	102.0	99.1	95.4	91.2	
50.0°	113.9	114.7	115.7	116.3	115.7	113.3	109.7	106.6	103.6	99.9	95.5	
60.0°	115.5	116.4	117.4	118.0	117.4	115.0	111.3	108.1	105.0	101.1	96.5	
70.0°	116.3	117.2	118.2	118.8	118.1	115.6	111.9	108.6	105.3	101.2	96.5	
80.0°	116.7	117.6	118.6	119.1	118.4	115.8	112.0	108.6	105.2	101.0	96.1	
90.0°	117.0	117.8	118.8	119.3	118.5	115.9	112.0	108.5	105.0	100.6	95.6	

D.2 Engine reference fuel flow rate data.

Throttle angle	Engine reference fuel flow rate in g/s. Engine speed in r/min.											
	1000	1250	1500	1750	2000	2250	2500	2750	3000	3250	3500	3750
0.0°	0.14	0.13	0.12	0.12	0.11	0.11	0.10	0.09	0.09	0.08	0.08	0.07
2.5°	0.15	0.15	0.14	0.14	0.14	0.13	0.13	0.13	0.12	0.12	0.11	0.11
5.0°	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24
7.5°	0.36	0.39	0.40	0.41	0.41	0.42	0.42	0.42	0.43	0.43	0.44	0.44
10.0°	0.48	0.55	0.59	0.61	0.63	0.64	0.65	0.66	0.67	0.68	0.69	0.70
12.5°	0.58	0.69	0.77	0.82	0.86	0.89	0.92	0.93	0.94	0.96	0.98	1.00
15.0°	0.64	0.78	0.90	0.99	1.06	1.13	1.20	1.24	1.26	1.27	1.29	1.32
17.5°	0.69	0.84	0.98	1.10	1.21	1.33	1.45	1.53	1.57	1.59	1.62	1.66
20.0°	0.73	0.88	1.03	1.17	1.30	1.45	1.61	1.74	1.83	1.88	1.92	1.97
22.5°	0.76	0.92	1.08	1.23	1.37	1.54	1.72	1.88	2.01	2.10	2.17	2.26
25.0°	0.78	0.95	1.11	1.27	1.43	1.61	1.81	1.99	2.14	2.26	2.36	2.47
27.5°	0.80	0.97	1.14	1.30	1.47	1.66	1.88	2.07	2.24	2.38	2.50	2.64
30.0°	0.81	0.98	1.16	1.33	1.50	1.71	1.93	2.14	2.32	2.47	2.61	2.77
32.5°	0.82	1.00	1.17	1.35	1.53	1.74	1.97	2.19	2.38	2.55	2.70	2.87
35.0°	0.83	1.01	1.19	1.37	1.55	1.77	2.01	2.23	2.43	2.60	2.77	2.95
37.5°	0.84	1.01	1.20	1.38	1.57	1.79	2.03	2.26	2.47	2.65	2.82	3.01
40.0°	0.84	1.02	1.20	1.39	1.58	1.80	2.05	2.29	2.50	2.68	2.86	3.06
50.0°	0.85	1.03	1.22	1.41	1.61	1.84	2.09	2.34	2.56	2.76	2.95	3.16
60.0°	0.86	1.04	1.23	1.42	1.62	1.85	2.11	2.36	2.59	2.79	2.98	3.20
70.0°	0.86	1.04	1.23	1.42	1.62	1.85	2.12	2.37	2.59	2.80	3.00	3.22
80.0°	0.86	1.04	1.23	1.42	1.62	1.86	2.12	2.37	2.60	2.80	3.00	3.22
90.0°	0.86	1.04	1.23	1.42	1.62	1.86	2.12	2.37	2.60	2.80	3.00	3.22
	4000	4250	4500	4750	5000	5250	5500	5750	6000	6250	6500	
0.0°	0.06	0.06	0.05	0.04	0.04	0.03	0.03	0.02	0.01	0.01	0.00	
2.5°	0.11	0.10	0.10	0.10	0.09	0.09	0.09	0.08	0.08	0.07	0.07	
5.0°	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.23	0.23	0.23	
7.5°	0.44	0.45	0.45	0.46	0.46	0.46	0.47	0.47	0.48	0.48	0.48	
10.0°	0.71	0.72	0.73	0.74	0.75	0.76	0.77	0.78	0.79	0.80	0.81	
12.5°	1.02	1.04	1.05	1.07	1.09	1.11	1.13	1.15	1.17	1.19	1.21	
15.0°	1.35	1.39	1.42	1.45	1.48	1.51	1.55	1.58	1.61	1.64	1.67	
17.5°	1.71	1.76	1.80	1.85	1.90	1.95	2.00	2.04	2.09	2.14	2.19	
20.0°	2.05	2.14	2.23	2.30	2.36	2.41	2.48	2.54	2.61	2.68	2.75	
22.5°	2.37	2.51	2.65	2.78	2.87	2.94	3.01	3.10	3.19	3.29	3.39	
25.0°	2.63	2.82	3.03	3.22	3.37	3.46	3.54	3.66	3.80	3.94	4.09	
27.5°	2.82	3.04	3.30	3.55	3.76	3.90	4.01	4.17	4.36	4.54	4.74	
30.0°	2.97	3.21	3.49	3.77	4.02	4.20	4.34	4.54	4.78	5.01	5.24	
32.5°	3.08	3.34	3.64	3.94	4.21	4.41	4.58	4.80	5.06	5.32	5.57	
35.0°	3.17	3.45	3.76	4.07	4.36	4.57	4.76	5.00	5.28	5.56	5.84	
37.5°	3.24	3.53	3.85	4.18	4.47	4.70	4.90	5.16	5.46	5.75	6.04	
40.0°	3.30	3.59	3.92	4.26	4.56	4.80	5.01	5.28	5.59	5.89	6.20	
50.0°	3.42	3.73	4.08	4.44	4.76	5.02	5.26	5.55	5.89	6.22	6.56	
60.0°	3.47	3.78	4.14	4.50	4.84	5.11	5.35	5.65	6.00	6.34	6.69	
70.0°	3.48	3.80	4.16	4.53	4.86	5.14	5.38	5.69	6.04	6.38	6.74	
80.0°	3.49	3.81	4.17	4.53	4.87	5.15	5.40	5.70	6.05	6.40	6.75	
90.0°	3.49	3.81	4.17	4.54	4.88	5.15	5.40	5.71	6.06	6.41	6.76	

D.3 Motor reference torque data.

Motor reference torque in Nm.										
Demanded armature current.	Motor speed in r/min.									
	0	250	500	750	1000	1250	1500	1750	2000	2250
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10A	1.5	1.5	1.6	1.6	1.6	1.6	1.7	1.7	1.7	1.7
20A	4.7	4.8	4.8	4.8	4.9	4.9	4.9	4.9	4.9	4.9
30A	8.0	8.0	8.0	8.1	8.1	8.1	8.1	8.1	8.2	8.2
40A	11.2	11.2	11.3	11.3	11.3	11.4	11.4	11.4	11.4	11.4
50A	14.4	14.5	14.5	14.5	14.6	14.6	14.6	14.6	14.6	14.6
60A	17.7	17.7	17.8	17.8	17.8	17.8	17.9	17.9	17.9	17.9
70A	20.9	21.0	21.0	21.0	21.1	21.1	21.1	21.1	21.1	21.1
80A	24.2	24.2	24.2	24.3	24.3	24.3	24.3	24.3	24.4	24.4
90A	27.4	27.4	27.5	27.5	27.5	27.6	27.6	27.6	27.6	27.6
100A	30.6	30.7	30.7	30.7	30.8	30.8	30.8	30.8	30.8	30.8
110A	33.9	33.9	34.0	34.0	34.0	34.0	34.0	34.1	34.1	34.1
120A	37.1	37.2	37.2	37.2	37.2	37.3	37.3	37.3	37.3	37.3
130A	40.4	40.4	40.4	40.5	40.5	40.5	40.5	40.5	40.5	40.6
140A	43.6	43.6	43.7	43.7	43.7	43.7	43.8	43.8	43.8	43.8
150A	46.8	46.9	46.9	46.9	47.0	47.0	47.0	47.0	47.0	47.0
160A	50.1	50.1	50.1	50.2	50.2	50.2	50.2	50.3	50.3	50.3
	2500	2750	3000	3250	3500	3750	4000	4250	4500	4750
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10A	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.6	1.5	1.4
20A	4.9	4.9	4.9	4.9	4.9	4.9	4.9	4.7	4.4	4.2
30A	8.2	8.2	8.2	8.2	8.2	8.2	8.0	7.6	7.1	6.8
40A	11.4	11.4	11.4	11.4	11.4	11.4	10.9	10.3	9.8	9.3
50A	14.6	14.7	14.7	14.6	14.6	14.6	13.7	12.9	12.2	11.6
60A	17.9	17.9	17.9	17.9	17.9	17.4	16.3	15.4	14.6	13.8
70A	21.1	21.1	21.1	21.1	21.1	20.0	18.8	17.7	16.8	15.9
80A	24.4	24.4	24.4	24.4	23.9	22.4	21.1	19.9	18.8	17.8
90A	27.6	27.6	27.6	27.6	26.4	24.7	23.2	21.9	20.7	19.6
100A	30.8	30.8	30.8	30.7	28.7	26.8	25.2	23.8	22.5	21.3
110A	34.1	34.1	34.1	33.0	30.7	28.8	27.1	25.5	24.1	22.8
120A	37.3	37.3	37.3	35.0	32.7	30.6	28.7	27.1	25.6	24.2
130A	40.6	40.6	39.7	36.9	34.4	32.2	30.2	28.5	26.9	25.5
140A	43.8	43.8	41.6	38.6	36.0	33.7	31.6	29.7	28.1	26.6
150A	47.0	46.8	43.2	40.1	37.4	34.9	32.8	30.9	29.1	27.6
160A	50.3	48.4	44.6	41.4	38.5	36.0	33.8	31.8	30.1	28.5
	5000	5250	5500	5750	6000	6250	6500	6750		
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
10A	1.3	1.3	1.2	1.2	1.1	1.1	1.0	1.0		
20A	4.0	3.8	3.6	3.4	3.3	3.2	3.0	2.9		
30A	6.4	6.1	5.8	5.6	5.4	5.1	5.0	4.8		
40A	8.8	8.4	8.0	7.6	7.3	7.0	6.8	6.5		
50A	11.0	10.5	10.0	9.6	9.2	8.8	8.5	8.2		
60A	13.1	12.5	11.9	11.4	10.9	10.5	10.1	9.7		
70A	15.1	14.4	13.7	13.1	12.6	12.1	11.6	11.2		
80A	16.9	16.1	15.4	14.7	14.1	13.6	13.1	12.6		
90A	18.6	17.8	17.0	16.2	15.6	14.9	14.4	13.9		
100A	20.2	19.3	18.4	17.6	16.9	16.2	15.6	15.0		
110A	21.7	20.7	19.7	18.9	18.1	17.4	16.7	16.1		
120A	23.0	21.9	20.9	20.0	19.2	18.4	17.7	17.1		
130A	24.2	23.1	22.0	21.1	20.2	19.4	18.7	18.0		
140A	25.3	24.1	23.0	22.0	21.1	20.3	19.5	18.8		
150A	26.2	25.0	23.9	22.8	21.9	21.0	20.2	19.5		
160A	27.1	25.8	24.6	23.5	22.6	21.7	20.9	20.1		

Demanded armature current.	Motor reference torque in Nm. Motor speed in r/min.									
	0	250	500	750	1000	1250	1500	1750	2000	2250
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-10A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-20A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-30A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-11.3
-40A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-14.5	-14.5	-14.5
-50A	0.0	0.0	0.0	0.0	0.0	0.0	-17.8	-17.8	-17.8	-17.7
-60A	0.0	0.0	0.0	0.0	0.0	-21.0	-21.0	-21.0	-21.0	-21.0
-70A	0.0	0.0	0.0	0.0	-24.3	-24.3	-24.3	-24.2	-24.2	-24.2
-80A	0.0	0.0	0.0	0.0	-27.5	-27.5	-27.5	-27.5	-27.5	-27.5
-90A	0.0	0.0	0.0	0.0	-30.8	-30.7	-30.7	-30.7	-30.7	-30.7
-100A	0.0	0.0	0.0	0.0	-34.0	-34.0	-34.0	-34.0	-33.9	-33.9
-110A	0.0	0.0	0.0	-37.3	-37.2	-37.2	-37.2	-37.2	-37.2	-37.2
-120A	0.0	0.0	0.0	-40.5	-40.5	-40.5	-40.4	-40.4	-40.4	-40.4
-130A	0.0	0.0	0.0	-43.8	-43.7	-43.7	-43.7	-43.7	-43.7	-43.7
-140A	0.0	0.0	0.0	-47.0	-47.0	-46.9	-46.9	-46.9	-46.9	-46.9
-150A	0.0	0.0	0.0	-50.2	-50.2	-50.2	-50.2	-50.1	-50.1	-50.1
-160A	0.0	0.0	0.0	-53.5	-53.4	-53.4	-53.4	-53.4	-53.4	-53.4
	2500	2750	3000	3250	3500	3750	4000	4250	4500	4750
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-10A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-20A	0.0	0.0	0.0	-8.0	-8.0	-8.0	-8.0	-8.0	-8.0	-7.6
-30A	-11.3	-11.3	-11.3	-11.3	-11.3	-11.3	-11.3	-11.3	-11.2	-10.6
-40A	-14.5	-14.5	-14.5	-14.5	-14.5	-14.5	-14.5	-14.5	-14.5	-13.7
-50A	-17.7	-17.7	-17.7	-17.7	-17.7	-17.8	-17.8	-17.8	-17.8	-16.9
-60A	-21.0	-21.0	-21.0	-21.0	-21.0	-21.0	-21.0	-21.0	-21.0	-20.0
-70A	-24.2	-24.2	-24.2	-24.2	-24.2	-24.2	-24.2	-24.2	-24.3	-23.2
-80A	-27.5	-27.5	-27.5	-27.5	-27.5	-27.5	-27.5	-27.5	-27.5	-26.4
-90A	-30.7	-30.7	-30.7	-30.7	-30.7	-30.7	-30.7	-30.7	-30.7	-29.6
-100A	-33.9	-33.9	-33.9	-33.9	-33.9	-33.9	-34.0	-34.0	-34.0	-32.8
-110A	-37.2	-37.2	-37.2	-37.2	-37.2	-37.2	-37.2	-37.2	-37.2	-36.1
-120A	-40.4	-40.4	-40.4	-40.4	-40.4	-40.4	-40.4	-40.4	-40.4	-39.4
-130A	-43.7	-43.6	-43.6	-43.7	-43.7	-43.7	-43.7	-43.7	-43.7	-42.7
-140A	-46.9	-46.9	-46.9	-46.9	-46.9	-46.9	-46.9	-46.9	-46.9	-46.0
-150A	-50.1	-50.1	-50.1	-50.1	-50.1	-50.1	-50.1	-50.2	-50.2	-49.4
-160A	-53.4	-53.4	-53.4	-53.4	-53.4	-53.4	-53.4	-53.4	-53.4	-52.8
	5000	5250	5500	5750	6000	6250	6500	6750		
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
-10A	-4.3	-4.1	-3.9	-3.8	-3.6	-3.5	-3.3	-3.2		
-20A	-7.2	-6.9	-6.6	-6.3	-6.0	-5.8	-5.6	-5.3		
-30A	-10.1	-9.7	-9.2	-8.8	-8.5	-8.1	-7.8	-7.5		
-40A	-13.1	-12.5	-11.9	-11.4	-10.9	-10.5	-10.1	-9.7		
-50A	-16.0	-15.3	-14.6	-14.0	-13.4	-12.9	-12.4	-11.9		
-60A	-19.0	-18.1	-17.3	-16.6	-15.9	-15.2	-14.7	-14.1		
-70A	-22.0	-21.0	-20.1	-19.2	-18.4	-17.7	-17.0	-16.3		
-80A	-25.1	-23.9	-22.8	-21.8	-20.9	-20.1	-19.3	-18.6		
-90A	-28.1	-26.8	-25.6	-24.5	-23.5	-22.5	-21.7	-20.8		
-100A	-31.2	-29.7	-28.4	-27.2	-26.0	-25.0	-24.0	-23.1		
-110A	-34.3	-32.7	-31.2	-29.9	-28.6	-27.5	-26.4	-25.4		
-120A	-37.4	-35.7	-34.1	-32.6	-31.2	-30.0	-28.8	-27.7		
-130A	-40.6	-38.7	-36.9	-35.3	-33.8	-32.5	-31.2	-30.1		
-140A	-43.8	-41.7	-39.8	-38.1	-36.5	-35.0	-33.7	-32.4		
-150A	-47.0	-44.7	-42.7	-40.8	-39.1	-37.6	-36.1	-34.8		
-160A	-50.2	-47.8	-45.6	-43.6	-41.8	-40.2	-38.6	-37.2		

D.4 Motor reference battery current data.

Motor reference battery current in A.										
Demanded armature current.	Motor speed in r/min.									
	0	250	500	750	1000	1250	1500	1750	2000	2250
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10A	23.5	23.5	23.5	23.5	23.5	23.5	23.5	23.5	23.5	23.5
20A	33.5	33.5	33.5	33.5	33.5	33.5	33.5	33.5	33.5	33.5
30A	43.5	43.5	43.5	43.5	43.5	43.5	43.5	43.5	43.5	43.5
40A	53.5	53.5	53.5	53.5	53.5	53.5	53.5	53.5	53.5	53.5
50A	63.5	63.5	63.5	63.5	63.5	63.5	63.5	63.5	63.5	63.5
60A	73.5	73.5	73.5	73.5	73.5	73.5	73.5	73.5	73.5	73.5
70A	83.5	83.5	83.5	83.5	83.5	83.5	83.5	83.5	83.5	83.5
80A	93.5	93.5	93.5	93.5	93.5	93.5	93.5	93.5	93.5	93.5
90A	103.5	103.5	103.5	103.5	103.5	103.5	103.5	103.5	103.5	103.5
100A	113.5	113.5	113.5	113.5	113.5	113.5	113.5	113.5	113.5	113.5
110A	123.5	123.5	123.5	123.5	123.5	123.5	123.5	123.5	123.5	123.5
120A	133.5	133.5	133.5	133.5	133.5	133.5	133.5	133.5	133.5	133.5
130A	143.5	143.5	143.5	143.5	143.5	143.5	143.5	143.5	143.5	143.5
140A	153.5	153.5	153.5	153.5	153.5	153.5	153.5	153.5	153.5	153.5
150A	163.5	163.5	163.5	163.5	163.5	163.5	163.5	163.5	163.5	163.5
160A	173.5	173.5	173.5	173.5	173.5	173.5	173.5	173.5	173.5	173.5
	2500	2750	3000	3250	3500	3750	4000	4250	4500	4750
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10A	23.5	23.5	23.5	23.5	23.5	23.5	23.5	22.4	20.5	19.1
20A	33.5	33.5	33.5	33.5	33.5	33.5	33.5	31.5	29.9	28.5
30A	43.5	43.5	43.5	43.5	43.5	43.5	42.8	40.8	39.2	37.9
40A	53.5	53.5	53.5	53.5	53.5	53.5	51.8	50.1	48.6	47.4
50A	63.5	63.5	63.5	63.5	63.5	63.3	61.0	59.3	58.0	57.0
60A	73.5	73.5	73.5	73.5	73.5	72.2	70.3	68.7	67.5	66.7
70A	83.5	83.5	83.5	83.5	83.5	81.3	79.5	78.0	77.0	76.4
80A	93.5	93.5	93.5	93.5	92.7	90.5	88.8	87.5	86.7	86.1
90A	103.5	103.5	103.5	103.5	101.7	99.7	98.1	97.0	96.3	95.8
100A	113.5	113.5	113.5	113.3	110.8	108.9	107.5	106.6	106.0	105.6
110A	123.5	123.5	123.5	122.1	119.9	118.2	117.0	116.3	115.8	115.4
120A	133.5	133.5	133.5	131.1	129.0	127.5	126.6	126.0	125.5	125.2
130A	143.5	143.5	142.6	140.1	138.3	137.0	136.3	135.7	135.3	135.0
140A	153.5	153.5	151.4	149.2	147.6	146.6	145.9	145.5	145.1	144.8
150A	163.5	163.2	160.4	158.4	157.0	156.2	155.7	155.2	154.9	154.6
160A	173.5	171.9	169.4	167.6	166.5	165.9	165.4	165.0	164.7	164.4
	5000	5250	5500	5750	6000	6250	6500	6750		
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
10A	17.9	17.0	16.4	16.0	15.7	15.4	15.1	14.9		
20A	27.4	26.7	26.2	25.8	25.5	25.2	25.0	24.8		
30A	37.0	36.4	36.0	35.6	35.3	35.1	34.8	34.6		
40A	46.7	46.2	45.8	45.4	45.2	44.9	44.7	44.5		
50A	56.4	55.9	55.6	55.3	55.0	54.8	54.6	54.4		
60A	66.1	65.7	65.4	65.1	64.9	64.6	64.4	64.3		
70A	75.9	75.5	75.2	74.9	74.7	74.5	74.3	74.2		
80A	85.7	85.3	85.0	84.8	84.6	84.4	84.2	84.0		
90A	95.4	95.1	94.9	94.6	94.4	94.2	94.1	93.9		
100A	105.3	105.0	104.7	104.5	104.3	104.1	103.9	103.8		
110A	115.1	114.8	114.6	114.3	114.2	114.0	113.8	113.7		
120A	124.9	124.6	124.4	124.2	124.0	123.9	123.7	123.6		
130A	134.7	134.5	134.3	134.1	133.9	133.7	133.6	133.4		
140A	144.5	144.3	144.1	143.9	143.8	143.6	143.5	143.3		
150A	154.4	154.2	154.0	153.8	153.6	153.5	153.3	153.2		
160A	164.2	164.0	163.8	163.6	163.5	163.4	163.2	163.1		

Demanded armature current.	Motor reference battery current in A.									
	Motor speed in r/min.									
	0	250	500	750	1000	1250	1500	1750	2000	2250
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-10A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-20A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-30A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-1.4
-40A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-1.7	-4.0	-6.2
-50A	0.0	0.0	0.0	0.0	0.0	0.0	-2.5	-5.3	-8.2	-11.0
-60A	0.0	0.0	0.0	0.0	0.0	-2.1	-5.5	-8.9	-12.3	-15.6
-70A	0.0	0.0	0.0	0.0	-0.4	-4.4	-8.3	-12.3	-16.3	-20.2
-80A	0.0	0.0	0.0	0.0	-2.1	-6.6	-11.1	-15.7	-20.2	-24.7
-90A	0.0	0.0	0.0	0.0	-3.7	-8.8	-13.9	-19.0	-24.1	-29.2
-100A	0.0	0.0	0.0	0.0	-5.2	-10.9	-16.5	-22.2	-27.8	-33.5
-110A	0.0	0.0	0.0	-0.4	-6.6	-12.9	-19.1	-25.3	-31.5	-37.8
-120A	0.0	0.0	0.0	-1.2	-8.0	-14.8	-21.6	-28.4	-35.2	-42.0
-130A	0.0	0.0	0.0	-1.9	-9.3	-16.6	-24.0	-31.4	-38.7	-46.1
-140A	0.0	0.0	0.0	-2.6	-10.5	-18.4	-26.3	-34.3	-42.2	-50.1
-150A	0.0	0.0	0.0	-3.1	-11.6	-20.1	-28.6	-37.1	-45.6	-54.1
-160A	0.0	0.0	0.0	-3.6	-12.6	-21.7	-30.8	-39.8	-48.9	-57.9
	2500	2750	3000	3250	3500	3750	4000	4250	4500	4750
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-10A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-20A	0.0	0.0	0.0	-1.1	-2.2	-3.3	-4.5	-5.6	-7.0	-9.0
-30A	-3.1	-4.8	-6.5	-8.2	-9.9	-11.6	-13.3	-15.0	-16.8	-18.9
-40A	-8.5	-10.8	-13.0	-15.3	-17.6	-19.8	-22.1	-24.4	-26.6	-28.8
-50A	-13.8	-16.6	-19.5	-22.3	-25.1	-28.0	-30.8	-33.6	-36.5	-38.6
-60A	-19.0	-22.4	-25.8	-29.2	-32.6	-36.0	-39.4	-42.8	-46.2	-48.5
-70A	-24.2	-28.2	-32.1	-36.1	-40.0	-44.0	-48.0	-51.9	-55.9	-58.4
-80A	-29.3	-33.8	-38.3	-42.8	-47.4	-51.9	-56.4	-61.0	-65.5	-68.2
-90A	-34.3	-39.3	-44.4	-49.5	-54.6	-59.7	-64.8	-69.9	-75.0	-78.1
-100A	-39.2	-44.8	-50.5	-56.1	-61.8	-67.5	-73.1	-78.8	-84.4	-88.0
-110A	-44.0	-50.2	-56.4	-62.7	-68.9	-75.1	-81.4	-87.6	-93.8	-97.8
-120A	-48.8	-55.5	-62.3	-69.1	-75.9	-82.7	-89.5	-96.3	-103.1	-107.7
-130A	-53.4	-60.8	-68.1	-75.5	-82.9	-90.2	-97.6	-104.9	-112.3	-117.5
-140A	-58.0	-65.9	-73.9	-81.8	-89.7	-97.6	-105.6	-113.5	-121.4	-127.4
-150A	-62.5	-71.0	-79.5	-88.0	-96.5	-105.0	-113.5	-122.0	-130.5	-137.2
-160A	-67.0	-76.0	-85.1	-94.2	-103.2	-112.3	-121.3	-130.4	-139.4	-147.0
	5000	5250	5500	5750	6000	6250	6500	6750		
0A	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
-10A	-0.6	-1.9	-2.8	-3.4	-3.9	-4.2	-4.5	-4.8		
-20A	-10.5	-11.8	-12.7	-13.4	-13.8	-14.2	-14.5	-14.7		
-30A	-20.4	-21.7	-22.7	-23.3	-23.8	-24.2	-24.5	-24.7		
-40A	-30.3	-31.6	-32.6	-33.3	-33.8	-34.1	-34.4	-34.7		
-50A	-40.2	-41.5	-42.5	-43.2	-43.7	-44.1	-44.4	-44.7		
-60A	-50.1	-51.4	-52.4	-53.2	-53.7	-54.1	-54.4	-54.6		
-70A	-60.0	-61.3	-62.4	-63.1	-63.6	-64.0	-64.4	-64.6		
-80A	-69.9	-71.2	-72.3	-73.1	-73.6	-74.0	-74.3	-74.6		
-90A	-79.7	-81.1	-82.2	-83.0	-83.6	-84.0	-84.3	-84.6		
-100A	-89.6	-91.0	-92.1	-92.9	-93.5	-93.9	-94.3	-94.5		
-110A	-99.5	-100.9	-102.1	-102.9	-103.5	-103.9	-104.2	-104.5		
-120A	-109.4	-110.8	-112.0	-112.8	-113.4	-113.9	-114.2	-114.5		
-130A	-119.3	-120.7	-121.9	-122.8	-123.4	-123.8	-124.2	-124.5		
-140A	-129.2	-130.6	-131.8	-132.7	-133.3	-133.8	-134.1	-134.4		
-150A	-139.1	-140.5	-141.7	-142.6	-143.3	-143.7	-144.1	-144.4		
-160A	-148.9	-150.4	-151.6	-152.6	-153.2	-153.7	-154.1	-154.4		