UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

FORD MOTOR COMPANY

Petitioner,

v.

PAICE LLC & ABELL FOUNDATION, INC.

Patent Owner.

DECLARATION OF PAULINA LUBACZ

I, Paulina Lubacz, hereby declare as follows:

1. I am presently employed as the Chief Operating Officer at Durham University. Durham University is a collegiate research university located in the Durham, England. I have personal knowledge of the matters stated below. I am over 18 years of age, and I am competent to testify regarding the following.

2. It is the normal course of business for the library services to index and catalogue doctoral thesis papers that are submitted by students at Durham University.

3. Attached as **Exhibit A** to my declaration is a true and accurate copy of a doctoral thesis titled "Some drive train control problems in hybrid i.c. engine/battery electric vehicles" that was authored by Philip Wilson Masding.

4. The first page of **Exhibit A** includes an imprint of Durham University Main Library's property stamp together with a stamped date of "2 November 1989." This property stamp and date would have been placed on the thesis at the time it was being processed by the library services at Durham University.

5. As is the normal practice at Durham University, the property stamp is also repeated on page 101 of the doctoral thesis attached as **Exhibit A**.

6. I have been informed by library services that between the processing date of November 2, 1989 and November 2011, if no embargo on access was requested by the author, there would have existed no time between the processing date of November 2, 1989 and November 2011 when the doctoral thesis attached as **Exhibit A** was not publically available at the library of Durham University. The library has no record of any embargo having been requested for the doctoral thesis attached as **Exhibit A**.

7. The doctoral thesis attached as **Exhibit A** would therefore have been indexed and searchable by the general public since around November 2, 1989.

8. The doctoral thesis attached as **Exhibit A** would have been indexed and searchable by the general public well before September 1997.

9. Durham University is currently in the process of digitizing hard-bound doctoral thesis papers.

10. The hard-bound version of the doctoral thesis attached as **Exhibit A** was removed from the library's shelves in November 2011 for digitizing.

11. The doctoral thesis attached as **Exhibit A** was digitized on December 10, 2012 and deposited with the Durham University e-Thesis repository on February 8, 2013.

12. The doctoral thesis attached as **Exhibit A** is now available for download by the public through Durham University's e-Thesis repository at the following weblink: http://etheses.dur.ac.uk/6408/.

I declare under the penalty of perjury that the foregoing is true and accurate to the best of my knowledge.

6. 11. 2014
Date

Paulina Lubacz

Some Drive Train Control Problems In Hybrid i.c Engine/Battery Electric Vehicles

by Philip Wilson Masding B.Sc.

A Thesis Submitted for the Degree of Doctor of Philosophy

School of Engineering and Applied Science
University of Durham

1988

Declaration

None of the work contained in this thesis has been previously submitted for a degree in this or any other university. It is not part of a joint research project.

Copyright

Acknowledgements

I would like to acknowledge the tremendous support and encouragement I have had from my supervisor, Dr. J.R. Bumby, throughout this research. I am also grateful to Mr N. Herron and Mr K. McGee for their assistance in maintaining and building the rig and to Mr N. Herron in particular who designed the mechanical system for the gearbox. Finally I would like to thank Mr C. Dart for his help with electrical work on the rig.

I am also grateful to the Ford Motor Company, Lucas Chloride E.V. Systems Ltd., and Shell Research Ltd. for the provision of equipment.

Abstract

This thesis describes the development of a microprocessor based control system for a parallel hybrid petrol/electric vehicle. All the fundamental systems needed to produce an operational vehicle have been developed and tested using a full sized experimental rig in the laboratory.

The work begins with a review of the history of hybrid vehicles, placing emphasis on the ability of the petrol electric design to considerably reduce the consumption of oil based fuels, by transferring some of the load to the broad base of fuels used to generate electricity.

Efficient operation of a hybrid depends on the correct scheduling of load between engine and motor, and correct choice of gear ratio. To make this possible torque control systems using indirect measurements provided by cheap sensors, have been developed. Design of the control systems is based on a theoretical analysis of both the engine and the motor. Prior to final controller design, using the pole placement method, the transfer functions arising from the theory are identified using a digital model reference technique. The resulting closed loop systems exhibit well tuned behaviour which agrees well with simulation.

To complete the component control structure, a pneumatic actuation system was added to a 'manual gearbox' bringing it under complete computer control. All aspects of component control have been brought together so that an operator can drive the system through simulated cycles. Transitions between modes of operation during a cycle are presently based on speed, but the software is structured so that efficiency based strategies may be readily incorporated in future.

Consistent control over cycles has been ensured by the development of a computer speed controller, which takes the place of an operator. This system demonstrates satisfactory transition between all operating modes.

CONTENTS

List of Symbols

a	Zero of $g_c(w')$.
C_d	Coefficient of drag
C_r	Coefficient of rolling resistance
e_a	Back e.m.f. voltage, V
f_c	Counter value from flywheel speed probe
f_t	Number of teeth on the flywheel speed probe gear
g	Gain of $g_c(w')$.
$G_a(s)$	Armature closed loop transfer function
$g_c(w')$	P+I Controller in w'-plane form.
$g_c(z)$	Bilinear discretisation of $g_c(w')$.
g_f	Final drive ratio
$G_f(s)$	Field closed loop transfer function
i_a	Armature current, A
i_f	Field current, A
$\bar{i}_f$	Mean field current during a transient, A
J	Flywheel inertia, kgm^2
J_{eq}	Equivalent inertia of vehicle mass M, kgm^2
K	Constant relating dynamometer speed to load
K_c	Scaling factor, flywheel count to roadspeed
K_e	Back e.m.f. constant ($= \frac{2\pi K_T}{60}$)
k_i	$gaT_s/2$
K_f	Manifold filling delay gain
K_p	Gain of armature controller
K_{qt}	Power stroke delay gain
K_T	Motor torque constant
K_2	Gain of engine torque/speed transfer function
$k_\theta(\theta, p_m)$	Manifold pressure/throttle gain

L_a	Armature inductance, H
L_f	Field inductance, H
$\dot{m}_a$	Mass air flow into inlet manifold, kg/s
$\dot{m}_c$	Mass charge flowing out of inlet manifold, kg/s
M	Vehicle mass, kg
m_m	Mass of gaseous mixture in the inlet manifold, kg
N	Engine or motor speed, r.p.m.
N_0	Speed at which linearisation takes place, r.p.m.
P	Adaptive gain matrix
P_{ic}	Engine power, kW
p_m	Inlet manifold depression, mbar
$p_{m(0)}(N)$	Manifold depression at no load and N r.p.m., mbar
R_a	Armature resistance, Ω
R_f	Field resistance, Ω
r_w	Vehicle wheel radius, m
T_{em}	Motor torque, Nm
T_f	Torque in gearbox output shaft, Nm
T_{ic}	Engine torque, Nm
T_l	Vehicle aerodynamic and rolling resistance loss torque, Nm
t_r	Controller design criteria: rise time, sec
T_s	Control system sampling period, sec
v_a	Armature voltage, V
v_f	Field voltage, V
v_j	Adaptive error at j^{th} sampling point
$\alpha(\theta_m)$	Field current set-point function
β_j	Plant model coefficient vector
$\beta(\theta_m)$	Armature current set-point function
Δ	Signifies small change in variable

ξ	Controller design criteria damping factor
θ	Engine throttle position, 0.9° steps
θ_d	Demand throttle position, 0.9° steps
θ_m	Electric motor accelerator/ brake demand
$\bar{\theta}_m$	Mean demand value during transient
τ_f	Inlet manifold filling time constant, sec
τ_2	Engine inertia time constant, sec
ϕ_f	Field flux, Wb
ϕ_j	Plant model input/output vector
y	Controller torque demand

List of Figures

CHAPTER 1

INTRODUCTION

Conventional internal combustion (i.c.) engine vehicles, based on the Otto and Diesel cycles, have dominated road vehicle transport for the greater part of this century. Among their many advantages are unlimited range, ease of refuelling, good power characteristics and the low cost materials required for their construction [JPL, 1975 pp 3-2][Unnewehr and Nasar, 1982 pp 214-217] Despite these facts conventional vehicles bring with them problems of air pollution and inefficient use of dwindling fuel resources, which at the present rates may be depleted early in the twenty-first century [Foley, 1976]. Although diminishing energy resources, and in particular depletion of fossil fuels, poses a problem to many areas of technology it is particularly pronounced in the transport sector where the choice of fuels is most restricted. This dependence on one type of fuel is illustrated by examining figures for total energy consumption in Britain over recent years. Energy consumption in Britain peaked in about 1973 [Bumby and Clarke, 1982] subsequently falling by 11% in the years up to 1981 as a result of price rises, conservation measures and recession. Price rises were particularly great for oil based fuels and many non transport users switched to gas or coal over this time period. As a result of this switch over amongst more flexible users, the transport sector has become responsible for using a progressively greater proportion of the total amount of oil used in the U.K. Over the period 1973-1981 this proportion rose dramatically from 29% to 47%. Further analysis of the transport sector itself shows that most fuel is used by passenger cars. In 1978 such vehicles accounted for 63% of all fuel used in vehicles, whilst the public transport sector for example, including taxis, coaches and buses used only 4.4% of the total [Dept. of Energy, 1978]. Faced with this heavy dependence on oil, inspite of its likely exhaustion in the relatively near term, there has been

increasing interest in possible alternative vehicle propulsion systems over the last twenty years. Amongst the many possible alternative power plants are the Stirling engine, the Brayton engine (gas turbine) the all electric vehicle and the hybrid vehicle [JPL, 1975]. Of these alternatives the hybrid power plant is the subject of the present study. A hybrid power plant maybe defined as one that is powered by two or more energy sources and as such there are a multitude of possible combinations. Examples of some which have been investigated are:

Heat engine/battery electric

Flywheel/battery electric

Heat engine/flywheel/battery electric

Pneumatic/battery electric

Battery/battery

The purpose of this study is to consider the control problems relating to the heat engine/battery electric power train in particular. As many authors have pointed out [Mitcham and Bumby, 1977][Unnewehr and Nassar, 1982] using two power plants adds considerably to the complexity and cost of the whole vehicle and hence the power train must offer other compensating factors such as lower running costs, reduced pollution and less noise in sensitive areas. Of the alternatives to conventional i.c. vehicles mentioned earlier, the electric vehicle would seem to offer all these advantages without the added complexity of a duel power source. Unfortunately, due to current battery technology, electric vehicles are basically 'energy deficient' [Unnewehr and Nassar, 1982]. Despite the fact that many types of advanced batteries have been the subject of much research and development, many people still feel that for the present, at least, the only practical storage battery for use in electric vehicles is the lead/acid type [van Niekerk et al, 1980][Dell, 1984], which has an energy density of 40 Wh/kg as opposed to the energy

density of petroleum which is 12300 Wh/kg [Unnewehr and Nassar, 1982]. In addition the energy density of lead/acid batteries falls at high discharge rates which always occur in vehicle applications during acceleration. As a consequence of these factors even advanced electric vehicles such as the ETV-1, developed for the U.S. Department of Energy by Chrysler and General Electric, suffer from severe range limitations varying in this case from 160 km at 60 km/h down to 95 km at 90 km/h even under favourable steady cruise conditions [Kurtz et al, 1981]. One way of reducing the severe range restrictions of the straightforward all electric vehicle is to combine it with some secondary energy source in a hybrid design. As can be seen in the list of hybrid drives above, all incorporate an electrical system. Of these systems the flywheel/battery electric, the pneumatic battery/electric and the battery/battery can be considered as primarily electric since the secondary energy storage device has very limited capacity. Nevertheless this secondary device can make a substantial improvement to the vehicle range simply by providing energy for short bursts of acceleration. For example flywheel systems, despite their modest energy storage capacity, can easily provide enough energy to accelerate a vehicle up to 30 mph once, thus protecting the batteries from the damaging high discharge rates which normally occur during such operation. In one such study simulations showed that electric vehicle range could be extended by 80% by using a flywheel of about 0.5 kWh capacity [Burrows et al, 1980].

On the other hand the hybrid heat engine/ battery electric is by no means primarily electric, with the power and stored energy capacity of the engine typically exceeding that of the electrical system. In this case the motor takes on a secondary role in shielding the engine from operating conditions which are particularly unfavourable to it. Typically this means that such a vehicle will operate electrically when the total power requirement is low.

Since low loading is typical of urban driving this vehicle has the potential to reduce pollution and noise in the urban environment whilst at the same time still offering the extended range capability characteristic of the engine. In general this system seeks to emphasise the particular strengths of each power source whilst playing down the disadvantages.

When considering any type of hybrid electric vehicle two basic mechanical configurations form the basis of nearly all designs. These are the series and parallel configurations, which are illustrated by figures 1.1 and 1.2 respectively. In the parallel hybrid both the engine and the motor are connected directly to the roadwheels. Provision is usually made for them to operate either independently or jointly in powering the vehicle. In the series hybrid however, only the motor is connected to the road wheels and must therefore be sized to meet all vehicle acceleration and top speed power requirements. In this case the engine supplies power indirectly via the generator battery combination. Although the series system has been considered for large bus applications [Brusaglino, 1982][Bader, 1981][Roan, 1976], on the grounds that it gives greater flexibilty in positioning components and ease of controlling the i.c. engine, the multiple energy conversions invariably lead to a low overall efficiency. Consequently almost all designs for passenger car applications have tended to favour the parallel configuration. As seen previously this latter application is most important because passenger cars offer the greatest potential fuel savings in the country as a whole. Clearly if hybrid vehicles could offer significant fuel savings in this transport sector, the benefit to society would be substantial. Before considering what potential the hybrid heat engine/electric has for making fuel savings in this sector it is important to consider what it is competing against in terms of likely improvements to conventional vehicles and other alternative vehicle power systems.

1.1 Possible Improvements or Alternatives to I.C. Engine Vehicles

It is important to recognise that any alternative to i.c. engine vehicles is competing against a mature, established technology backed up by a large industrial research base. As a result conventional vehicles have shown a steady improvement over the years and seem likely to continue to do so. It is therefore pointless to compare possible economic and environmental advantages of a new vehicle technology against the i.c. engine vehicles which are being built today. A more realistic comparison is to judge the new technology against an estimate of what conventional vehicles will achieve 10-15 years from now. This 10-15 year period will then allow a reasonable time span for the new technology to be put into production on a worthwhile scale. Some of the likely improvements in i.c. engine design which might be made over the next few years are discussed by Blackmore and Thomas [Blackmore and Thomas, 1979]. This review suggests that major economy improvements could be achieved by using higher compression ratios coupled with leaner mixtures thus saving 20% of fuel use over all operating conditions. Lean mixtures must be used in conjunction with the higher compression ratios to avoid problems with knocking. Other savings can be achieved by improvements to cold starting (saving 25% on short trips), power modulation by variable mixture strength as opposed to the inefficient throttling method (saving 20% on average) and finally up to 10% can be saved by improving fuel and lubricants. Apart from changes to the engine itself further improvements can be made by altering vehicle aerodynamics and designing better tyres and gearboxes.

Electronic engine control systems also offer substantial improvements in fuel economy. One major application for electronics is in optimising ignition timing. Standard mechanical distributors produce a crude variation in spark timing on the basis of engine speed, as sensed by a centrifugal weight

system, and engine load, as determined indirectly by a vacuum diaphragm lever system connected to the inlet manifold. Load and speed are in fact the correct indicators needed to optimise ignition timing, but the mechanical system can never re-produce the complex function relating them to optimum spark advance. In contrast, in an electronic system, it is possible to store the best spark advance as a function of manifold depression and speed in a two dimensional look-up table. As described by Main [Main, 1986] Ford have produced a microprocessor system which not only controls spark timing but also controls a continuously variable transmission system (CVT).

Unlike some proposed vehicle systems the hybrid heat engine/battery electric can incorporate all the changes mentioned above in its own engine and so its potential advantages remain undiminished.

1.1.1 Novel Engines

Apart from all electric and hybrid vehicles which may generally be considered as the electric alternative, whether partial or complete, there are two other major possibilities; alternative fuels and novel engines. Leading examples of alternative fuels are those based on agricultural products such as alcohol made from sugar. This has been successfully applied in Brazil with 95% alcohol by volume being used in cars with modified carburetters. According to the Jet Propulsion Laboratory (JPL) [JPL, 1975] lesser concentrations, up to 28%, maybe used in ordinary cars without any such modifications. A second alternative is to produce synthetic petroleum from coal as has already been done in S. Africa. Neither of these alternatives are without drawbacks though, alcohol for example is particularly prone to absorb water in storage and the petrol from coal process is undesirable on both environmental and efficiency grounds [Mitcham and Bumby, 1977].

Amongst the most likely novel engines are the gas turbine, based on

the Joule Brayton cycle, and the Stirling engine. Gas turbines are already highly successful as aero-engines but need much development for small scale use in cars. In addition they use more expensive materials in order to withstand the high temperatures associated with a continuous combustion process [Bumby and Clarke, 1982]. External combustion engines such as the Stirling design promise a broader choice of fuels and the ability to optimise combustion conditions for better efficiency and reduced emissions. Despite improvements to heat exchangers however the Stirling engine still suffers from large thermal inertia.

In their study of alternative vehicle power systems, JPL concluded that the Stirling engine and the gas turbine might achieve gains of 20% and 30% respectively when compared with the conventional vehicles of the time (1975). When such improvements are viewed in the light of the predicted improvements to the conventional car itself, these figures are not outstanding, particularly when it is considered that in all cases the alternative engine costs more to produce.

1.1.2 The Electric Vehicle

Although of somewhat limited range the all electric vehicle must not be dismissed prematurely as a likely replacement means of transport, particularly in certain applications. Fleet trials of electric vehicles have shown them to be very useful in applications with a well defined use pattern and limited range requirement, such as the small delivery vehicle. Volkswagen [Altendorf et al, 1982] undertook one such trial with 40 electric vehicles beginning in 1978 and found them to be reliable and acceptable to their operators. Unfortunately this success will not be repeated in the car market until improved battery technology is widely available. Although many battery types such as nickel/zinc and nickel/iron offer considerable improvements in

energy storage capacity over lead-acids, they still do not give an electrical vehicle comparable performance with a conventional vehicle. One possible battery technology which does promise large vehicle range, up to 200 km, is sodium/sulphur. Here the main problem is that battery operating temperature must be above 300°C.

Assuming that some battery technology such as sodium/sulphur could provide vehicle performance that was acceptable to the consumer, a whole new problem is created when the infrastructure needed for refuelling is considered. To assess this problem Watson [Watson et al, 1986] carried out a computer simulation study to determine the cost of the vehicle refuelling infrastructure that would be needed for an all electric car fleet of 14 million vehicles; representing Britain's entire fleet in 1983. It was assumed that this fleet would follow the usage pattern established for British drivers by the 1978-79 national travel survey. Several likely methods of charging were considered, each providing different proportions of the total refuelling need. Most important was recharging at the vehicle owner's premises, providing 80-85% of vehicle energy needs, followed by 10% at other charging locations and finally 5-10% at battery exchange stations. As a result the total cost to the nation was calculated to be between 205 and 478 million pounds depending on the predominant battery type used in the vehicles. By far the greatest part of this cost is accounted for by the battery exchange stations. This fact explains the great dependence of the total cost on battery type, since battery exchange stations would have to be spaced according to the range of the electric vehicles. In this respect the hybrid has a definite advantage since the need for battery exchange stations might well be zero, due to the range provided by the engine and the ability to recharge from the engine itself should battery charge become dangerously low. As far as engine refuelling is concerned the hybrid vehicle could make use of the massive petroleum

distribution infrastructure already in existence.

Any one of the above systems may one day play a role in road vehicle transport, however the main emphasis in this introduction must lie with the possible application of the hybrid heat engine/ battery electric power train. Such vehicles are by no means a new idea, and in the following section their history will be traced from its earliest beginnings to the present day.

1.2 The History of Hybrid I.C. Engine/Electric Passenger Cars

Hybrid i.c. engine/electric vehicles have a long history dating back to the early part of this century. At this time i.c. engine vehicles were underpowered and the hybrid, originally patented by Pieper in 1909 [Pieper, 1909] was seen as one way of of achieving better performance. At the same time the all electric vehicle enjoyed a brief period of popularity since it offered greater reliability than the i.c engine. The peak period for electric vehicles came in about 1915 when 4500 vehicles were produced annually in the U.S. [Kordesch, 1977]. However i.c. engine production had already passed the half million mark at this stage and its rapidly improving performance, coupled with cheap petrol, effectively stopped all interest in electric and hybrid cars by 1930. Specialist electric vehicles have however remained popular for certain applications, most notably as milk floats in the U.K. Away from such limited applications the use of i.c engine vehicles remained unquestioned until the late 1960's when the staggering exhaust emission problems of some North American cities began to arouse public concern, forcing the U.S. government to enact strict emissions legislation. The scale of this pollution problem is illustrated by 1966 U.S. government statistics which showed that total air pollution from all sources amounted to 141 million tons per year, of which 61% came from motor vehicles [U.S. Dept. of Commerce, 1967]. It was for this reason that the U.S. Environmental Protection Agency (EPA) sponsored extensive studies

of hybrid vehicles of all types as a possible means of meeting the new emissions legislation which was due to come into force in 1976. Several organisations received contracts to carry out the studies amongst them Aerospace and TRW evaluated specifically heat engine/battery hybrids [Sampson and Killian, 1972][TRW, 1972]. Both these studies concluded that any improvements in emissions or fuel economy would not justify the added complexity of the hybrid system. Vehicle manufacturers themselves showed renewed interest in hybrids and other alternative vehicle drive trains. This interest arouse from the manufacturers initial technical inability to meet the required emission standards; a fact they admitted to publicly. As one example Ford sponsored the design study carried out by the JPL [JPL, 1975], which covered a wide range of automobile power systems. Other manufacturers and research organisations went on to produce operational hybrid vehicles at this time [Minicars, 1972][Agarwal et al, 1969][EPA Office of Air and Waste Pollution Programs, 1975][Fersen, 1974][Wouk, 1976]. These vehicles were built mainly to reduce exhaust emissions, with little consideration paid to the fuel saving potential of hybrid vehicles. On the whole though car manufacturers were not seriously committed to hybrid vehicles and eventually they successfully concentrated their efforts towards meeting emissions regulations by modifying conventional i.c. engine vehicles. These efforts resulted in crude, inefficient, systems which absorbed a great deal of engine power. To maintain acceptable driveability and performance the manufacturers simply installed larger engines causing average fuel economy in the U.S. to reach a record low in 1975 [Roan, 1984].

As mentioned previously the EPA study was concerned with the potential of hybrid vehicles of all types to reduce exhaust emissions. Most of the study groups involved concluded that hybrids could produce good results in this respect, however in the light of the cheap fuel prices at the time they

were not regarded as a cost effective solution. Almost as soon as these results were announced however the whole situation was changed by the Arab Oil Embargo and the subsequent steep rise in energy prices. New emphasis was now placed on fuel economy, which meant that manufacturers were faced with the duel problem of maintaining emission standards whilst at the same time reducing fuel consumption. As has been previously mentioned improvements in emissions had previously been obtained at the expense of fuel consumption. Such difficulties prompted Ford to re-examine the potentials of hybrid vehicles using a combination of computer studies and practical dynamometer trials [Unnewehr et al, 1976]. Although the dynamometer testing was rather limited in scope, it did establish the practicality of some important fuel saving measures useful for hybrid vehicles, in particular fuel of at idle was successfully used. In addition the researchers found that results from the dynamometer agreed very well with simulations of the same experiments on a computer. Encouraged by these consistent results the research went on to compare the simulated performance of experimental hybrid drive trains and control strategies with results obtained from production vehicles. All substitute drive trains were of parallel design and were sized to give approximate equivalent acceleration performance as the conventional system they replaced. Operating strategy for the hybrid called for battery state of charge to be maintained in a narrow band around 75%, thus all energy ultimately came from the fuel tank, with no use of mains electricity. Large fuel economy gains were predicted for the hybrid when compared with two vans produced by Ford at the time, these gains ranged from 30-70%. Unfortunately these results are not quite as encouraging as they appear since the comparison vehicles had, by modern European standards, massively overrated engines. For example the larger of the two vans had a 142 kW (190 BHP) engine and an automatic transmission. This design was compared with a replacement combination

of a 67 kW (90 BHP) engine and a 352 Nm (260 ft/lb) motor, which together achieved the largest quoted comparitive fuel saving over the Federal CVS-H cycle. Much less improvement was noted for the second van which had originally a 86 kW (115 BHP) engine. Despite the fact that these measurements of fuel saving are not applicable in a present day context, this work does establish some important principles of hybrid vehicle design. Firstly the hybrid drive outperforms its conventional counterpart by restricting engine operation to the high efficiency region of its characteristic. Further improvements accrue, particularly during urban operation, from the fuel off at idle strategy and energy reclaimed by regenerative braking. It was also shown that losses involved in turning the engine with fuel off could be reduced by collapsing the valves as an alternative to completely declutching the engine. Of equal importance to these results was the demonstration that computer aided parameter studies could optimise component ratings for hybrid vehicle applications. For example varying the final drive ratio was shown to effect both emissions and fuel economy.

Despite these efforts by Ford and considerable further work carried out in Universities [Roan, 1978][Beachley et al, 1978] major advances in hybrid vehicle technology really had to wait until the 'Electric Vehicle Research and Developments Act' of 1975 in which the U.S. government provided $180M for research into electric and hybrid vehicles. This act resulted in a major project aiming to develop an operational hybrid vehicle which had the potential to be mass produced in the near term. Management of the project was undertaken by JPL on behalf of the U.S. Department of Energy. In phase 1, beginning in 1978 JPL, gave sub-contracts to four companies to perform initial design studies. All four companies reported back to JPL by 1979 with their design proposals [Burke and Somuah, 1980][Fiat, 1979][Schwarz, 1980][Minicars, 1980]. Interestingly all four design teams recommended the parallel hybrid design.

After considering each proposal JPL awarded the second phase of the project to General Electric, who were to proceed with building the vehicle which became known as the HTV-1. In their initial study General Electric had used two computer packages. A first stage investigation was carried out using a package called HYVELD which examined the economic potential of various driveline configurations and component ratings. Inputs to the HYVELD program included annual vehicle usage pattern, engine and motor efficiencies and various economic factors effecting the vehicle owner. From these inputs the program quickly established that the series hybrid design could not compete with the conventional vehicle in producing an acceptable power to weight ratio at a reasonable cost. This was essential to the second phase of the project because the hybrid vehicle had to match the performance and comfort of a reference vehicle, specifically a General Motors A-Body Malibu Century [Trummel and Burke, 1983]. Important vehicle requirements were 5 adult capacity fuelled by wall plug electricity and petroleum fuel, cruise speed of 56 mph and acceleration of 0-56 mph in 15 seconds or less, gradeabilty of 3% at 56 mph for 0.6 miles and compliance with 1981 emission standards. General Electric formed a team with three other contractors each with responsibility for different vehicle sub-systems. General Electric themselves provided the motor whilst Globe Batteries provided lead acid batteries, Triad Services were responsible for the vehicle structure and finally Volkswagen produced a modified engine. An important new development in the HTV-1 was the use of a sophisticated microprocessor control system, which JPL considered to be essential for both efficient vehicle operation and to provide acceptable vehicle driveability [Trummel and Burke, 1984]. Of all the hybrid vehicle projects to date the HTV-1 control system has the most similarities with the work described in this thesis. Consequently specific aspects of the control will be discussed in context in later chapters, however

at this stage it is relevant to mention the overall results obtained from the HTV-1. Vehicle fuel economy was assessed under both urban and highway driving conditions. In both cases economy declined with distance travelled due to falling battery state of charge. Over five cycles of the 12 km FUDS cycle, battery state of charge (SOC) fell from 100% to 19%, indicating fairly limited urban range. Average battery losses of 17 Wh/km over the cycle caused this drop. Over the same five cycles the fuel economy fell from 54.9 mpg (23 km/l) to 17.7 mpg (7.5 km/l) [Trummel et al, 1984]. It is worth noting that the high power load from accessories, typical of American vehicles, was responsible for using 63 Wh/km and had this been reduced by only 25% the vehicle would have achieved unlimited urban range. A similar result could have been achieved by using a more advanced body design than the HTV-1 which had poor aerodynamics ($C_d = 0.43$). On the highway cycle the net fuel economy ranged from 34.6 mpg down to 26.2 mpg after 450 miles when battery SOC forced the use of heat engine recharge. Using a combination of these results the annual fuel consumption of the vehicle was predicted for typical usage patterns. The pattern chosen was for the 50th percentile all purpose 4-5 passenger vehicle in 1985. Under these conditions the HTV-1 achieved a combined urban and highway economy of 39 mpg which compared to the 1985 reference vehicle represented a petroleum saving of 31%. Although the performance of the i.c. engine reference vehicle is not particularly impressive by present standards, it does not detract from the achievement of this vehicle in demonstrating the considerable potential for substituting petroleum use with the more broadly based electricity.

Following the development of the HTV-1, American interest in hybrid vehicles has been reduced to further paper studies. JPL have carried out an assessment of advanced vehicle technology to establish what they think might be worthy of development in the 1990's [Hardy and Roan, 1980]. Computer

simulations were used extensively in this study to evaluate various vehicle drive trains. Each vehicle was designed to provide the same performance as an i.c engine equivalent vehicle in its category. These vehicle categories were based on range and passenger carrying capability and covered a broad spectrum of vehicles from a two passenger commuter car to a full size, general purpose car designed to carry five passengers. All electric vehicles proved to be an attractive option for the small commuter vehicle but no battery technology was capable of replacing the general purpose vehicle with adequate range at reasonable cost. In contrast the hybrid vehicle was shown to be ideally suited to this type of application. Its main advantage is that the presence of the petrol engine means that only limited range is needed from the batteries; 80 km was considered perfectly adequate in this case. Choice of battery technology is therefore unrestricted and so the cheapest available type can be used. This study recommended the bipolar lead acid as the most suitable design. Overall the JPL study concluded that hybrid vehicles might save 70-75% of liquid fuel.

The Electric Power Research Institute in America has also carried out a survey of hybrid vehicle technology fairly recently [Renner, 1986]. From the point of view of the electricity generating utilities, represented by this organisation, electric and hybrid vehicles represent a large potential market for off-peak electricity since such vehicles would most likely be charged overnight. In their study they identify the limited energy storage capacity of batteries and the complexity of the vehicle as a whole as the main barriers to the success of hybrids.

Development of actual prototype vehicles is now continuing outside America and in particular Volkswagen are building on experience gained during their involvement with the HTV-1 project. Volkswagen's continued interest in hybrid vehicles stems from the belief stated by Kalberlah [Kalberlah, 1986],

that the 'flexibility in the hybrid drive system makes it into a universal drive system not predestined to be restricted to the second car' (unlike the electric vehicle). In addition he notes that even if some advanced battery technology such as sodium sulphur does extend the range of electric vehicles to a respectable 250 km, the stored energy requirement of 50 kWh can not be charged overnight from a common power socket. To demonstrate the potential of hybrids Volkswagen have recently developed a parallel hybrid version of the Golf [Miersch et al, 1987]. This vehicle has very limited electrical capability, the 3 phase asynchron dynamotor can only provide 6 kW, which corresponds to a cruise speed of 50 km/h. This small motor follows from Volkswagen's belief that the cost of the electrical components must be kept to an absolute minimum if hybrids are to be an economic proposition. Control philosophy for the vehicle is relatively simple, the driver can select one of three operating modes; all electric, all i.c. engine or hybrid. In the hybrid mode the engine provides all power requirements above 6 kW, being shut down if such extra power is not needed. In assessing the potential of this system in the German car market Kalberlah notes that due to its unlimited range it has an application potential equal to that of the conventional Golf. A large potential market is seen to be vitally important because the vehicle would have to take full advantage of the economies inherent in high volume production to be competitive. Taking into account German usage patterns a typical mix of urban and highway driving would allow the vehicle to achieve 67% petroleum substitution. Total energy consumption by the vehicle over the ECE15 cycle is 3.3 litres/100 km plus 16 kWh/100 km. In monetary terms, at 1985 prices, the electricity used costs the equivalent of an extra 2.8 litres of petrol to the German consumer. Taking this into account a substantial cost saving results when compared with the baseline performance of the Golf which uses 10.11 litres/100 km over the ECE15 cycle.

1.3 The Context of the Present Work

Having brought this survey of hybrid vehicle technology up to date with the discussion of the latest Volkswagen results, the relevance of the present work and the computer studies which lead up to it, can now be established. Computer studies of hybrid vehicles have been carried out at the University of Durham using a general purpose road vehicle simulation package called Janus [Bumby et al, 1985]. This program, developed over a number of years in the School of Engineering and Applied Science, is capable of predicting the energy use of a variety of power train configurations. Computer simulation studies, using programs such as Janus, have formed an important part of hybrid vehicle research since the mid 1970's. Their great strength lies in their ability to systematically assess the fundamental design options available in hybrids such as component configuration, component layout and component ratings. In 1977 when JPL undertook a survey of electric and hybrid vehicle simulation software, it reported 111 operational programs [JPL, 1978], many of which had been written as a result of the 'Electric Vehicles Research and Development Act'. Obviously these programs performed many functions but two main categories can be identified [Wolfson and Gower, 1983]. Least concerned with actual vehicle design is the first category, which attempt to define what vehicle performance and range is necessary to meet the requirements of a given percentage of vehicle users or vehicle missions. Such studies are usually based on national travel statistics which are produced in both Britain and the U.S. [Department of Transport, 1979][U.S. Dept. of Transport 1972]. Often included at this stage is an assessment of the economic implications of a particular hybrid or electric design both in cost to the consumer and on required infrastructure such as electricity generating utilities. One such study was carried out by General Electric of America [Burke and Smith, 1983] to establish what performance would be required

from a hybrid vehicle.

Janus belongs to the second category which use models of individual drive train components to calculate the energy use of a vehicle when driven over a simulated driving cycle. In the HTV-1 project the equivalent program was called HEAVY, (Hybrid and Electric Advanced Vehicles Systems Simulation) [Hammond and Beach, 1981] developed by Boeing Computer Services in conjunction with NASA. This program stores a number of standard drive train configurations and drive train components together in a library module, allowing standard and novel drive trains to be investigated. In addition the user may define a drive train component that the program knows about by entering a list of parameters required to define that component. Having specified a design the program will put the resulting vehicle through a driving cycle. As the cycle proceeds the program models the drag and rolling resistance forces and thus produces the vehicle energy requirement at the road wheels which can be reflected back through the transmission, taking due account of losses, to arrive finally at the amount of petrol or battery charge consumed. Once the simulated cycle is complete the user has at his disposal a breakdown of energy requirements on an individual component basis. Using a method similar to this Janus has been used to thoroughly investigate the economic potential of parallel i.c. engine/electric hybrids [Bumby and Forster, 1987]. Although the results show that a hybrid can save considerable energy overall, when compared with present day vehicles, this ability reduces dramatically as likely improvements are incorporated in the conventional vehicle used for comparison. Neverthless they are attractive because they can shift energy use away from petroleum towards the broader base of fuels used to generate electricity. This petroleum substitution potential is also sensitive to the conventional vehicle technology used for comparison, although to a lesser extent than overall energy saving. Placing a precise

figure on the percentage of petroleum which might be saved by the hybrid is complicated by the vehicle use pattern [Forster and Bumby, 1988]. In the past overall energy consumption for vehicles have been based on an assumption that 40% of vehicle driving is under urban driving conditions, 50% at 90 km/h cruise and 10% at 120 km/h cruise. When applying this mix to the hybrid design it can reasonably be assumed that urban driving will be electrically powered and that the 120 km/h cruise will be powered by the engine. The 90 km/h cruise however might well occur in either mode, with petroleum substitution depending on the relative use of the engine and the motor. By assuming equal use of engine and motor at 90 km/h Bumby and Forster [Forster and Bumby, 1988] calculated that a parallel hybrid could save 50% of petrol when compared with an advanced conventional vehicle featuring an efficient continuously variable transmission (CVT). Clearly such a vehicle represents a much more formidable target performance than vehicles considered in American studies.

When considering the parallel hybrid drive train Janus predicted the effects of component ratings on overall performance in much the same way as the General Electric computer study did for phase 1 of the HTV-1 project. In the case of Janus though, the drive train performance was designed to match that of a modern mid-size European car. Such a vehicle represents the 1.6-2.0 litre category, and since over 40% of petrol in Britain is consumed by such vehicles they offer the greatest potential for saving fuel. Differences in the General Electric and Janus studies also appear in the way that they arrived at a detailed control strategy for operating the hybrid vehicle. General Electric considered the operating conditions which gave the best efficiencies for the engine and motor and then postulated a control strategy which they thought would best achieve the design goals of maximising petroleum substitution whilst still retaining acceptable performance. These broad objectives remain

valid in the Janus study, however the control strategy designed to meet them came out of a formal optimisation study. The optimisation process defined an energy cost function, which allows arbitrary values to be placed on petrol from the tank, E_1, and electric charge from the battery, E_2, as follows:

$$F = \lambda_1 E_1 + \lambda_2 E_2 \tag{1.1}$$

By varying the two weighting factors, λ_1 and λ_2, the cost function can be made to take into account the prevailing monetary value of the two fuels and at the same time other factors such as a penalty for the environmental damage associated with use of petrol can be considered. Having decided on the relative values of the two primary energy sources Janus then puts the experimental vehicle drive train through a driving cycle. At every one second interval the torque requirement of the vehicle at the wheels, is calculated to meet the instantaneous speed and acceleration defined by the cycle speed profile. This torque requirement is then reflected back through the drive train taking into account successive losses in each component, such as the final drive and gearbox. Once the gearbox is reached however there are two important decisions to be made, concerning which gear to use and how to split the total torque demand between the engine and motor. It is the gear ratio which defines what the total torque demand will be at the input shaft to the gearbox but this total demand can be met in any number of ways between the two extremes of all electric operation and all i.c. engine operation. Choice of gear ratio and torque split is in fact crucial to energy economy, since the efficiencies of both the engine and the motor are functions of torque and speed. This fact is clearly illustrated for the engine by figure 1.3 which is a typical efficiency map for a 50 kW engine. In order to calculate which gear and torque split to use, Janus evaluates the cost function for all gear ratios and a representative selection of torque splits. Thus the cost function

is determined by two independent variables, with the optimal solution being the minimum value. Having chosen the optimum solution for gear ratio and torque split, the energy use is calculated by assuming steady state conditions over the one second time increment. At this point the whole calculation is repeated.

The net result of the optimisation process ultimately depends on the weighting factors as illustrated by figure 1.4. This result shows range and liquid fuel consumption as a function of weighting factor for an urban driving cycle. If the ratio λ_1/λ_2 is greater than about 1.5 then the penalty for using petrol is so high that Janus will operate the vehicle almost entirely electrically, resulting in inevitable range limitations but effectively infinite liquid fuel economy. In contrast when the ratio is less than about 0.3 the preffered strategy includes so much operation by the engine that range is unlimited but fuel economy, measured in terms of litre/100 km, is low. As the weighting factor ratio is varied two differing operating philosiphies present themselves. At about $\lambda_1/\lambda_2 = 0.25$ raw energy use is minimised, suggesting an energy saving strategy. On the other hand at $\lambda_1/\lambda_2 = 0.6$ the engine is restricted to relatively efficient operation but use of electricity is small enough for sufficient urban range to satisfy nearly all such journeys. This type of operation could be considered as aiming for maximum petroleum substitution without compromising vehicle utility. Of these two philosophies petroleum substitution has perhaps the greatest potential. As mentioned previously there is likely to be considerable improvement in i.c. engine vehicles over the next few years and this means that although the energy saving hybrid makes economic sense at the moment it may not do so in the future. In contrast the petroleum substitution strategy will make increasing sense in future as liquid fuel prices increase due to the rising scarcity of the resource.

Whenever the weighting factors do allow use of the engine the op-

timisation process favours operation near to the engine maximum efficiency point which lies within the 32% contour on figure 1.3. In fact as the 'cost' of petrol is increased the optimisation tends to restrict use of the engine to ever smaller regions around this point. As an example figure 1.5 shows the use made of the engine over a typical urban driving cycle with the weighting ratio $\lambda_1/\lambda_2 = 0.6$. Each figure on the diagram represents the percentage of the total cycle time that the engine spent operating at that point. Although a good deal of the cycle was powered electrically the figures for the engine are clearly clustered round the maximum efficiency contour.

Carrying out the optimisation process in full involves quite complex calculation, particularly to determine losses in the prime movers. Consequently although it would be ideal to include it in an operational vehicle there simply is too much work involved for real time computation. By considering the usage patterns for both the engine and the motor over an optimally controlled cycle however, it is possible to devise a sub-optimal control algorithm, based on a number of simple rules, which produces virtually the same economy as the fully optimal case [Bumby and Forster, 1987]. It is the practical component control problems raised by this sub-optimal control strategy which provided the motivation for the work described in this thesis.

The job of the sub-optimal controller can be seen as that of selecting one of the possible operating modes that are available with a parallel hybrid drive train. All of the operating modes are described in Table 1.1. With maximum petroleum substitution as the goal each mode tends to be suited to a particular type of vehicle operation

Electric Mode. Used for journeys within the electric range of the vehicle.

Primary i.c. Engine Mode. Used for long distance high speed travel. Electric power is used to supplement the engine during hill climbing or acceleration.

Hybrid Mode. Used for long range urban or sub-urban travel beyond the capability of the motor alone. In this case the motor starts the vehicle from rest, with the engine taking over when efficiency considerations allow.

Accelerator 'kick-down'. Overrides other modes to give the driver full power in a panic situation.

Battery Charge Mode. Used to prevent battery charge from becoming dangerously low.

Regenerative Braking. Used whenever possible to recoup vehicle kinetic energy.

The inputs to the mode controller are the driver's demand signals represented by the position of the accelerator and brake pedals. Not all of the modes listed in Table 1.1 will actually be selected by the sub-optimal controller working on purely efficiency grounds. In particular it would never select the battery recharge mode because the multiple energy conversions would always be too inefficient. To overcome this problem the purely efficiency based strategy must be overriden by battery state of charge considerations. Accelerator 'kick-down' is another mode where efficiency considerations are overriden, this time by the needs of the driver. Normally the engine would be limited to 90% of its maximum power since above this the carburetter system would reduce efficiency, however in a panic situation the driver must be given full power regardless of the cost. Another useful input to the mode controller might be the journey type specified by the driver. If maximum petroleum substitution is the design goal then for example on a short urban journey it would be advantageous to restrict the vehicle to all electric operation since there is no possibility of dangerously depleting the batteries. Similarly on long journeys use of the motor should be kept to a minimum. Combining all these considerations produces the vehicle control hierarchy illustrated by figure 1.6. At the top of the hierarchy is the driver who makes inputs to

the system in two ways. Firstly his unique knowledge of intended destination allows him to select one of the three possible journey types. Secondly he communicates his power demand through use of the brake and accelerator. Journey type has a large influence on which operating strategy is best for the vehicle and hence which operating modes should be used. Restricting the use of some of these modes however are the battery charge constraints which seek to keep the battery SOC above 20% at all times. After processing through these constraints it is then up to the mode controller, which embodies the sub-optimal strategy to operate the vehicle in the best way to meet the power demands coming from the driver.

Even if the precise form of sub-optimal strategy mentioned above were not employed, any other mode controller would still pose the same set of component control problems. As it responds to the driver's ever changing input the mode controller produces only three outputs; the torque demand for the motor, the torque demand for the engine and the required gear ratio. Component control then involves matching the torque output of engine and motor to their respective demands and controlling the transmission. If the theoretical economic potential of the vehicle is to be realised in practice these demands must be met as quickly and accurately as possible. A further complication arises in the case of the engine, in that when it is not needed it must be shut down and remain stationary. Once torque is again required from the engine a control subsystem is needed to start it and match its speed with the rest of the drive train so that torque is reapplied smoothly.

Microprocessor based control systems capable of meeting all of these control needs have been developed on a full-sized laboratory test bed and are described over subsequent chapters. Using a test bed represents a logical progression from the computer simulations previously carried out at Durham. Unlike an operational vehicle problems such as component packaging do not

have to be tackled during the development of the rig. This leaves the way clear to test the fundamental feasability of the control strategies suggested by the simulations. In this context it should be mentioned that only the control of a warm engine has been considered on the grounds that if satisfactory engine control could not be achieved under such conditions then additional work would be pointless. Furthermore the vehicle manufacturers are currently undertaking the development of engine management and cold start systems [Meyer et al, 1983].

The individual component controllers developed for the rig have been assembled into a fully integrated drive train system capable of responding to the outputs of the mode controller. Final testing of this complete system is made easier by the development of a computer based driving cycle controller which mimics the action of a driver in producing the necessary accelerator and brake signals to follow the cycle. Using the computer based cycle speed controller leads to greater repeatability in the cycle tests. In addition the cycle controller is sufficiently flexible for any cycle to be quickly defined, and then driven, provided it is within the power limitations of the engine and motor on the rig. Although this system is intended for use with the sub-optimal mode controller described earlier, its operation is independent of the mode control strategy and is a necessary part of any hybrid drive train controller.

In the next chapter a description of the laboratory test system is given thus providing the necessary background information on the operating environment of each of the component controllers. Subsequently chapters 3 and 4 examine the relevant theory to allow indirect methods of torque measurement to be developed for both the engine and the motor. Chapter 5 then proceeds to design torque control systems on the basis of these measurements and a speed control system used during engine starting. The

dynamic performance of both torque and speed controllers is tested against simulations. The final control sub-system is a pneumatic actuator for a 'manual' gearbox described in chapter 6. In chapter 7 torque controllers, engine starting system and automated gearbox are combined into the fully operational system.

Mode	Description
Electric Mode	All propulsion power supplied by the electric traction system
I.C. Engine Mode	All propulsion power supplied by the i.c. engine
Primary Electric Mode	The electric traction system provides the principle torque, but when necessary its maximum torque is augmented by the engine
Primary I.C Engine Mode	The i.c. engine provides the principle torque, but when necessary its maximum torque is augmented by the motor
Hybrid Mode	Both the i.c. engine and the electric traction system provide torque split between them in some way.
Battery Charge Mode	The i.c. engine provides both the propulsion power and power to charge the batteries, with the traction motor acting as a generator.
Regenerative Braking	During braking the vehicle kinetic energy is returned to the battery with the traction motor acting as a generator.
Accelerator 'kick-down'	Essentially a primary i.c. engine mode when full engine torque is allowed to give maximum acceleration

Table 1.1 Possible Operating Modes for the Parallel Hybrid Vehicle

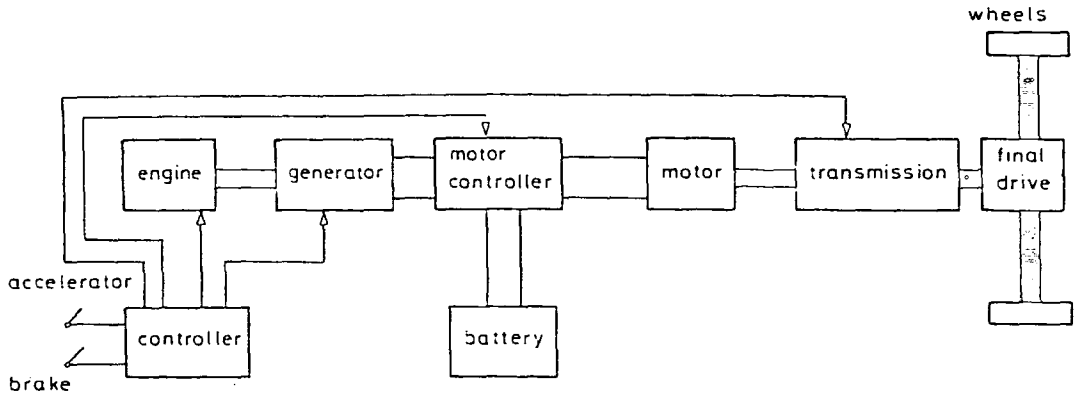

Fig. 1.1 Series Hybrid Electric Vehicle Drive Train

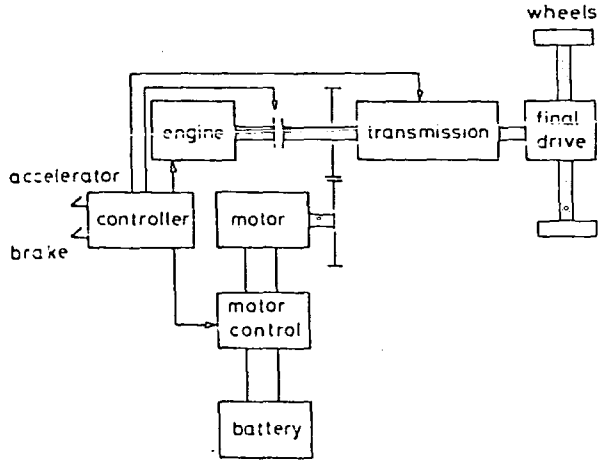

Fig. 1.2 Parallel Hybrid Electric Vehicle Drive Train

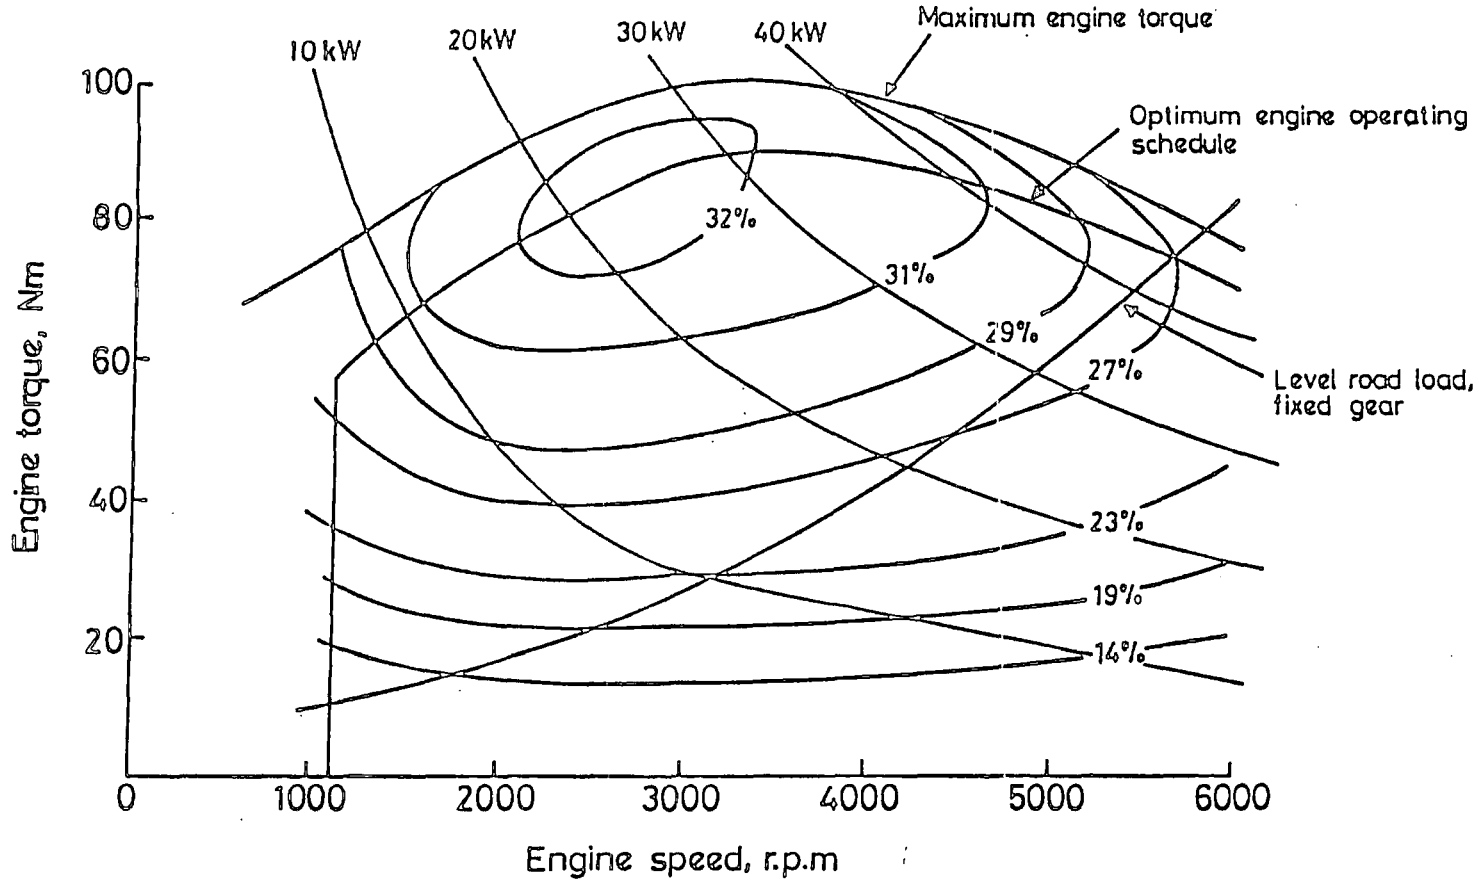

Fig. 1.3 Typical Efficiency Map for a 50 kW I.C. Engine

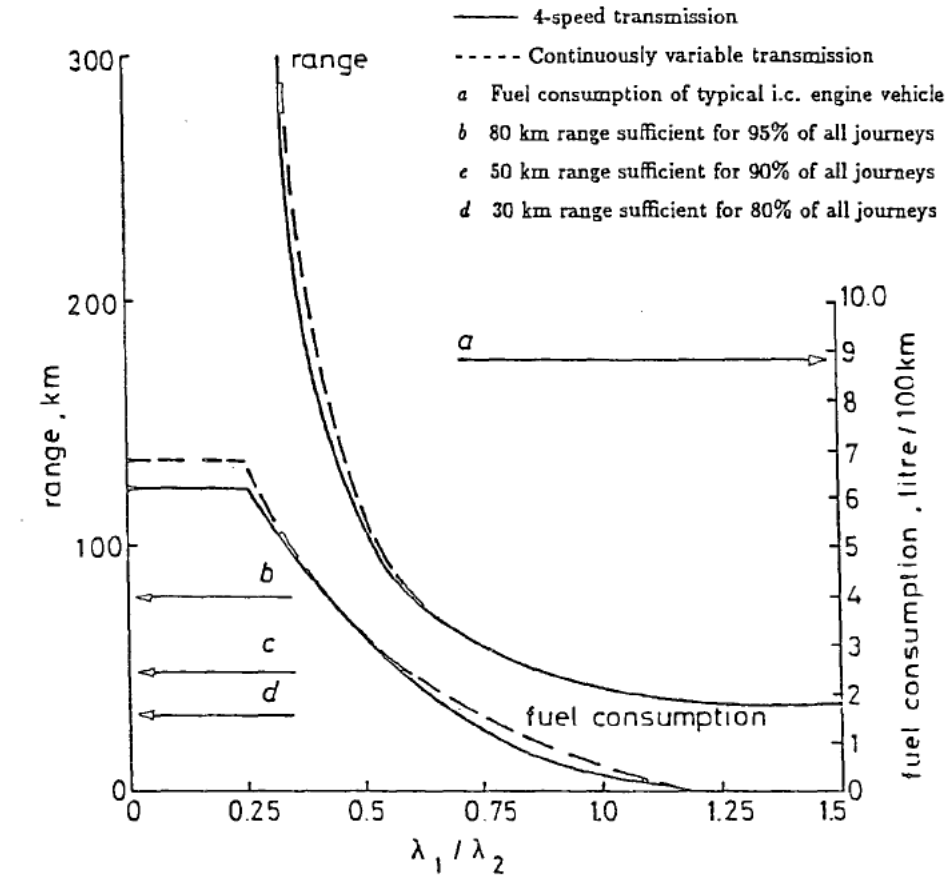

Fig. 1.4 Influence of Weighting Factor on the Performance of the Hybrid Vehicle

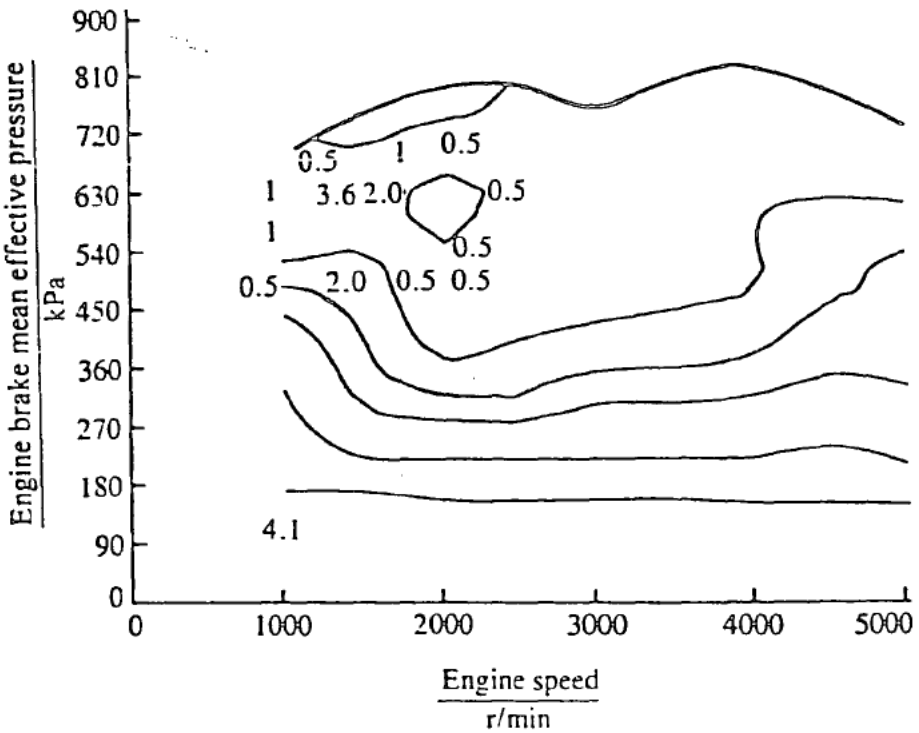

Fig. 1.5 Use of the Engine Over an Urban Driving Cycle for an Optimally Controlled Hybrid Vehicle

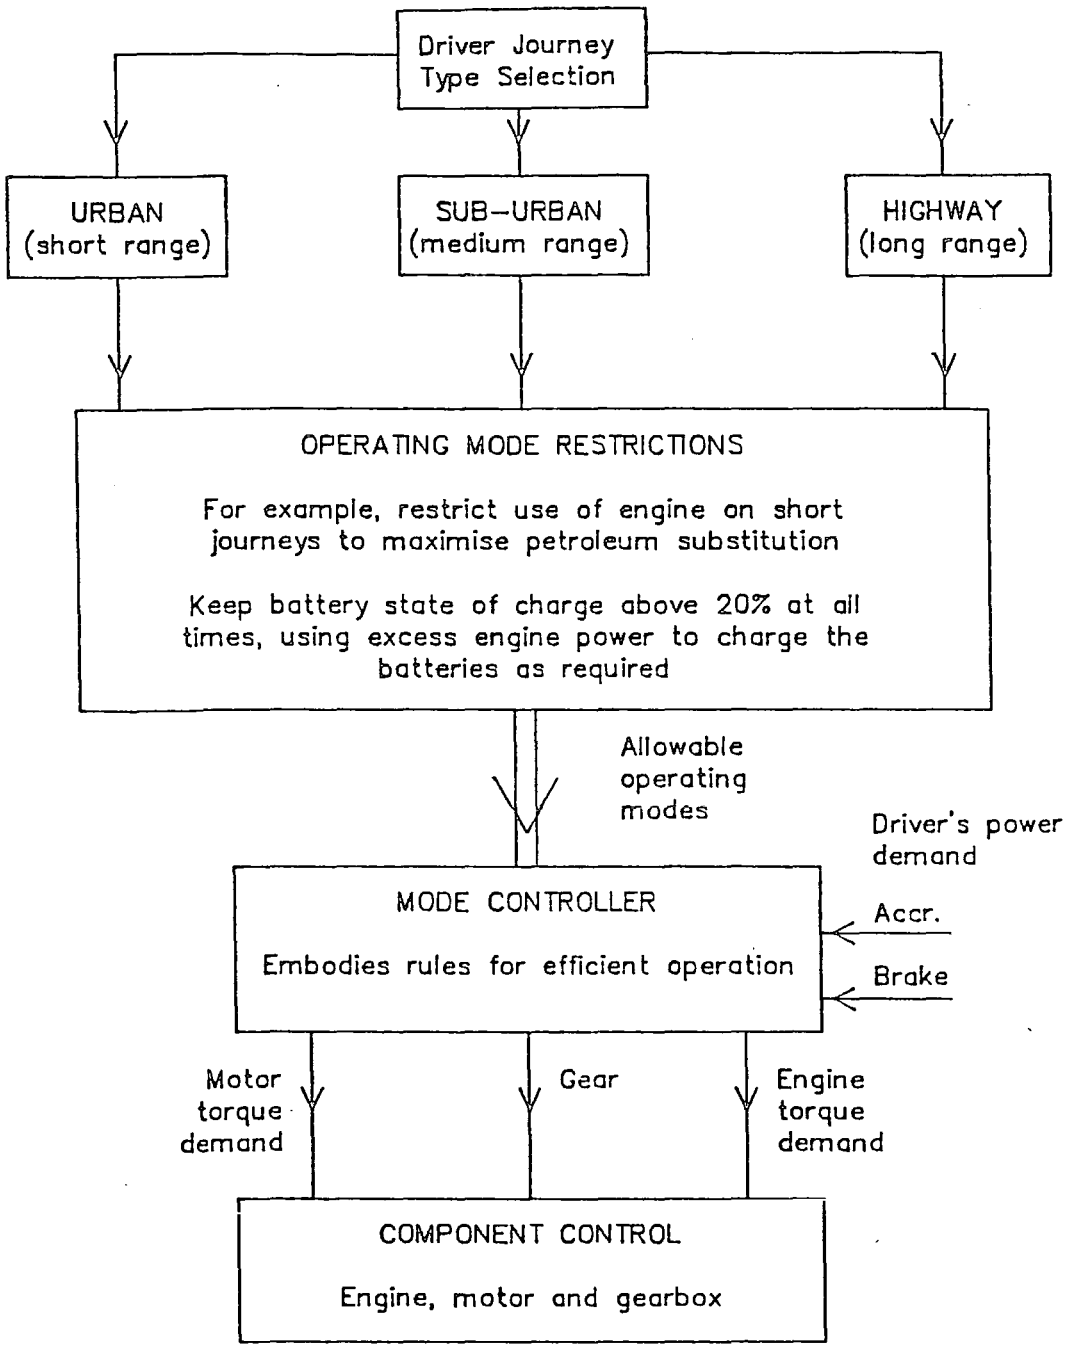

Fig. 1.6 Possible Hybrid Vehicle Control System

CHAPTER 2

THE LABORATORY TEST SYSTEM

The laboratory test system is the hardware realisation of the parallel hybrid drive shown in figure 1.2. The component layout in the rig system is shown by figure 2.1. Overall the rig consists of three subsystems; firstly the flywheel and dynamometer representing vehicle inertia and road load, secondly the hybrid drive train itself and finally the computers and signal monitoring equipment [Bumby and Masding, 1988]. Each of these subsystems are now described in detail.

2.1 Load Emulation

On the test bed layout of figure 2.1, all the components forward of, and including, the transmission represent the drive train components present in an actual vehicle. To the rear of the transmission the dynamometer and flywheel are included to represent the losses associated with a moving vehicle. Vehicle mass is simulated by the variable inertia flywheel which has an available inertia range of 2.02-15.47 kgm^2. In simulating any real vehicle of mass M, the necessary equivalent flywheel inertia J_{eq} is

$$J_{eq} = M \left(\frac{r_w}{g_f} \right)^2 \tag{2.1}$$

Where r_w is the vehicle wheel radius and g_f is the gear ratio of the vehicle final drive system. A nominal flywheel inertia of 6.54 kgm^2 was used throughout the experimental work. Independent calibration experiments presented in chapter 7 confirm the accuracy of this value.

Using typical values for a mid-range passenger car of $r_w = 0.3$ m and $g_f = 3$ the inertia of 6.54 kgm^2 represents a vehicle weighing 650 kg. Typically hybrid vehicle curb weights would be much higher than this, for example the HTV-1 [Trummel and Burke, 1983] weighed 2032 kg and a

hybrid vehicle built by Lucas Chloride weighed 2100 kg [Harding et al, 1983]. Simulating an unrealistically low weight on the rig does not however invalidate any of the control systems developed in subsequent chapters; indeed in most cases it is more difficult to obtain satisfactory smooth operation with low vehicle weight. Keeping flywheel inertia low is also an advantage on safety grounds, since it reduces the stored energy in the system. The dynamometer is of eddy current design and is water cooled. Control electronics for the dynamometer provide three load settings; constant speed, constant torque and power law. Most important for the hybrid vehicle work is the power law setting which ideally reproduces the effects of aerodynamic drag and rolling resistance which act on a real vehicle. Road wheel torque needed to overcome these two factors is given by [Bumby et al, 1985].

$$T_d = \frac{1}{2}\rho C_d A V r_w \qquad \text{Nm} \tag{2.1}$$

$$T_r = C_r M g r_w \qquad \text{Nm} \tag{2.2}$$

Where

$\rho =$ Air density (1.226 kg/m^3 @ 15°C, 10^5 Pa)

$C_d =$ Coefficient of drag

$C_r =$ Coefficient of rolling resistance

$A =$ Vehicle frontal area m^2

$V =$ Vehicle velocity m/s

$r_w =$ Wheel radius m

$M =$ Vehicle mass kg

$g =$ gravitational constant

In fact the dynamometer control system is not complex enough to reproduce the combined effects of rolling resistance and drag. Instead the dynamometer power law setting produces a load which varies according to

speed squared, but the curve is made of piecewise linear sections. Since the dynamometer system is typically operated at fairly low speeds, it rarely reaches the second linear sub-section of its loading characteristic and hence in most experiments loading on the rig is simply proportional to speed. Consequently neither drag nor rolling resistance are accurately modelled. In order to assess the seriousness of this defficiency figure 2.2 shows the dynamometer loading curve and the equivalent requirement predicted from equations (2.1) and (2.2) using likely parameters for a parallel hybrid as suggested by Bumby and Forster [Bumby and Forster, 1987]. Their values for such a car are $C_d = 0.35$, $C_r = 0.01$, $A = 1.95$ m^2, $g_f = 3.5$, $M = 1880$ kg and $r_w = 0.3$ m. Rotational speeds for the flywheel are converted to km/h by assuming the same values of g_f and r_w. Essentially this is the same as saying that the speed of the flywheel represents the prop-shaft speed in a real vehicle. As can be seen the dynamometer loading is at least realistic up to speeds of 50 km/h, which is the maximum encountered in the European ECE15 Urban driving cycle adopted for many of the rig tests. To operate at speeds above this would require modifications to the control electronics governing the dynamometer.

2.2 The Hybrid Drive System

The power sources in the experimental hybrid drive train are a standard Ford 1100cc petrol engine and a Lucas Chloride separately excited d.c. traction motor. The traction motor has been used in operational electric vehicles, specifically the Bedford CF and Freight Rover Sherpa 1 tonne panel vans [Manghan and Edwards, 1983]. Both the engine and the traction motor are connected to a common drive point at the input to the gearbox by a toothed belt drive giving a 1:1 speed ratio. Smooth performance is achieved by accurately balancing the belt and associated pulley wheels. The belt

itself has a trapezoidal tooth section and ground back. The i.c. engine is connected to the common drive point through a friction clutch, short drive shaft and torque transducer. During normal operation the friction clutch is not needed, since the rig is invariably started from rest by the electric motor. It is retained however for experiments requiring complete isolation of the engine, such as no-load speed testing. In addition this device also makes it possible to start the flywheel from rest on the engine alone, should the motor not be operational for some reason. An additional one way clutch or free wheel, allows the engine to remain stationary whilst the rest of the drive train is in motion. In common with the engine, the electric traction system is connected to the common drive point via a flexible coupling and torque transducer. The flexible coupling is designed to cope with the motoring and regenerative braking torques which are encountered in the motor drive shaft. In this instance a tyre type coupling is used which has proved ideal, both because it is easy to connect and has favourable mechanical properties such as zero backlash and ability to marginally damp high frequency torques. Such torques appear because of the nature of the chopper control system, which is a Lucas Chloride Mark III design and connects the motor to the traction batteries. It employs a thyristor armature chopper and transistor field chopper. A summary of all the drive train ratings is given in Table 2.1. In total the combined power output of the motor and engine can match that of a medium sized conventional passenger car which typically has a maximum engine power of 55 kW and weighs about 1 tonne [Bumby and Forster, 1985]. As such the engine and motor on the rig are representative of what might be installed in an equivalent hybrid vehicle. In contrast the 1 tonne lead acid battery pack used on the rig would preferably be replaced by 400 kg of similar batteries (or better still 200 kg of Ni/Zn batteries [Bumby and Forster, 1987]). The greater size of the battery pack used with the rig is

simply a consequence of its original application in an all electric van.

2.3 Computer Control and Signal Monitoring

Control of the rig is provided by a Motorola M68000 based microcomputer system. In carrying out its control function the M68000 has responsibilty for receiving and processing data from the transducers round the rig. Operator inputs and data display are handled by a second computer which is a Duet 16 personal computer (PC). Communication between the two computers is achieved via an RS232 serial data line operating at 9600 baud. A schematic diagram of the two computers and instrumentation interface is shown in figure 2.3.

2.3.1 Motorola M68000 system

At the heart of the drive train control system is the M68000 microprocessor unit which has two main tasks to perform. Firstly it must operate the engine motor and transmission as required to meet the demands of test driving cycles and secondly it must act as a data logger. Both control and data logging depend on the extensive range of interface hardware available to the M68000. In total this hardware comprises four 6522 versatile interface adaptor (VIA) chips with interrupt facility; a 16 channel, 12 bit, analogue to digital convertor (ADC) also with interrupt facility; eight 12 bit digital to analogue converters (DAC) and finally two standard timer/counter (STC) chips each with five, 16 bit, timer counters. The M68000 itself operates at 8MHz and has 192k of random access memory (RAM).

Output control signals from the M68000 are summarised in Table 2.2. The i.c engine output is controlled by a stepper motor servo-system on the throttle valve, with the stepper motor pulses and direction signal being generated by the M68000 using two of its VIA chips. On initial start

up the throttle control system is initialised by closing the throttle onto a micro-switch. Pulse rate for the stepper motor is software controllable but is consistently set to 1 step per 5 ms, representing a practical lower limit which ensures that no steps are lost when moving from rest under load. With each step being 0.9°, zero to full throttle requires 85 steps and is thus achieved in 425 ms. Ideally a slightly faster response is desirable and other drive by wire systems have achieved zero to full throttle in 150 ms [Collonia, 1979]. When describing the throttle servo for an automated gear change system requiring engine speed control, Ford [Main et al, 1987] state that to be transparent to the driver full to zero travel time should be less than 200 ms and positional accuracy should be 0.5° or better. The system described in this reference used an analogue throttle servo which offers greater resolution than the discrete motion of the stepper system. Originally such a device was employed on the rig, however the system proved extremely sensitive to noise from the engine ignition system causing unsatisfactory hunting around the command position leading to its replacement with the current stepper system.

The electric motor is controlled via the M68000 DAC interface which feeds analogue voltages to the Lucas Chloride power electronic control unit. One signal is used to control acceleration and another to control regenerative braking. Both software and hardware interlocks are used to prevent brake and motor signals from being applied simultaneously. In the event of a software failure a watchdog routine has the ability to disable the output from both DACs with a single signal from a VIA; with the result that no further erroneous controller action will be effective.

Software for the M68000 is written mostly in the 'C' language, with small amounts written in M68000 assembler. This assembler is necessary for manipulation of registers within the VIA and STC chips. The 'C' programming language is particularly useful for real time control applications

as it is sufficiently high level to allow complex control structures to be easily written, and yet compiles efficiently to produce fast execution times. 'C' also interfaces easily with M68000 assembler routines with a simple system of parameter passing between the two. All program development is carried out on an M68020 Unix based system and down line loaded into the M68000.

2.3.2 The Duet 16 Personal Computer

This computer does not carry out any direct control of the drive train. Its purpose is to interact with the operator for overall rig control, to display graphically in real time selected test results and, when a test is complete, to accept data from the M68000 system and display selected data after suitable analysis. All the P.C. software is written in BASIC. Although BASIC suffers from slow running speed and lack of structure its interactive nature and ease of use in graph plotting has proven extremely useful for displaying results. The Duet has a useful range of peripheral devices including a dual disk drive, colour graphics monitor and dot matrix printer which has been used to obtain most of the graphs in this thesis.

2.4 Instrumentation

Instrumentation on the rig is used to gather three categories of data; vital control values, safety measurements and visual checks. Control signals passed to the M68000 are shown in Table 2.3. In each case the transducer type and input signal type are given. In all cases these signals first pass through an interface electronics unit, housed under the M68000 itself, prior to connection to the appropriate input chip. Close proximity of the interface and microprocessor ensure minimal contamination by noise in the intervening wiring. Least accurate of the measuring devices listed in Table 2.3 is the engine fuel flow transducer. This device consists of a small turbine which

generates pulses as it rotates. Unfortunately at low flow rates measurement accuracy is poor and so this method cannot be used for measuring the idling fuel consumption of the engine. An alternative manual method uses a digital balance to accurately weigh the amount of fuel used and can easily measure idle fuel flow rates. A simple tap arrangement allows fuel to be drawn either from the main tank or a container on top of the balance.

Electronic signal processing provides the necessary filtering and amplification appropriate to each signal type and input chip combination. All interface circuits are based on the 741 operational amplifier (op-amp) and the filters use a second order Butterworth design [Hilburn, 1973]. Due to the presence of the op-amps a dual power supply is required, this is provided by a number of dc-dc converters which are intended to demonstrate that the whole system could function off the single 12v engine battery.

Speed signals are vital to rig safety and so duplicate systems provide speed readings to both the M68000 and hard wired overspeed trip circuitry. Three speed measurements are provided covering the engine, motor and flywheel. Reference to the rig layout diagram shows that this set of measurements completely defines the state of the rig. All speed signals come from magnetic probes activated by toothed wheels. Those providing data to the M68000 include integral circuitry to generate a pulse train, in contrast the safety probes produce unconditioned analogue output. Signals from the safety probes are fed to the operators display panel where, in addition to passing through the trip circuitry, they are used to drive panel speed meters.

In figure 2.1 the operator's control panel is labelled safety trips, instrumentation and test bed control. This unit includes the electronics for the torque transducers which process the raw signals prior to passing to the M68000. In addition these devices provide panel meters displaying torque, speed and power. Other displays on the panel show engine cooling water

temperature, oil pressure and battery state of charge.

Throughout the whole rig temperature sensitive paper is used liberally on components liable to overheat such as bearing mountings. Regular checks of these indicators show the maximum temperature which has been reached and thus provide an additional safety precaution. Finally two emergency stop buttons are provided, one on the operator's control panel and one near the gearbox, these cut the engine and motor ignition and apply an air brake which is fitted to the flywheel.

2.5 Software System

Computer software for the rig system provides real time data gathering and control, with simultaneous real time displays. Data analysis programs are also available for later analysis of results. Such varying demands are met using the software structure shown in figure 2.4. During an active test the M68000 and Duet run parallel programs which proceed in three broad phases regardless of the test being carried out. As can be seen from figure 2.4, the three software stages are initialisation, active rig control and data gathering, and finally data transfer and disk storage. Because of this common core structure programs designed to carry out widely differing tasks can be quickly produced by modifying existing programs.

2.5.1 System Initialisation

During the initialisation phase, system parameters are relayed from the Duet to the M68000. Maximum flexibility is achieved by passing nearly all variables in this way, however many seldom change between tests. To avoid the need to enter such variables every time a program is run, they can be selected as a default option by the user. A most important example of a variable of this type is the base sampling period which will apply to all digital

control algorithms and data sampling. This sampling period is generated by interrupts from a timer on the ADC board. Selecting a sampling period for a digital system is quite a complex process, however in this case it represents a compromise between allowing for the fast response times characteristic of the torque control systems and still leaving sufficient time to carry out all necessary control computations at each interrupt. As a result of this trade off the standard sampling period used in all experiments is 20 ms. Should a key variable such as this, ever need to be altered for some reason this means altering a line of BASIC which is a far quicker process than recompiling the M68000's 'C' program.

Ultimately the interrupt period is based on a 0.8 MHz clock on the ADC board which is then subdivided by countdown timers to provide the 20 ms period. Interrupts on the M68000 can have any one of 8 priority levels, with level 7 the highest and level 0 the lowest. Only levels 3 and 4 are used on the hybrid vehicle rig processor, the data sampling and control interrupt generated by the ADC is level 3 and the VIA chips can generate at level 4, either when one of their timers times out or in response to an external signal. Both these possibilities are represented by the three interrupt applications carried out by the VIA's on the rig system. One timer is used to count pulses output to the stepper motor. When all necessary pulses have been sent, the counter reaches zero and pulses are cut off by the interrupt routine. The second timer application is a watchdog system included in the software as a safety measure. During normal operation the watchdog timer is reset at each sampling interrupt and consequently it never reaches zero. In the event of a software failure causing the sampling interrupt to become inactive, the watchdog timer times out and generates a level 4 interrupt. When this happens the interrupt routine shuts the rig down by cutting the engine ignition and disabling motor accelerator and brake signals. Finally the

only external signal generating an interrupt comes from a gear lever position sensor as described in chapter 6.

Another important initialisation is concerned with data logging. Before an experiment begins the Duet presents the user with a full menu of available data sources. He is then able to select any combination of data sources which may be relevant to the particular test. Furthermore the user can request that these sources be recorded either after every 20 ms interrupt period, or after any multiple of it. Thus although data is sampled every 20 ms for control purposes, it need only be recorded as often as the user considers appropriate. Both these features of source selection and extended sampling period encourage best use of available memory space and reduce the time needed for data transfer to the Duet and its subsequent retrieval from disk.

An essential feature of the software in most applications is the cycle following capability. A test cycle commonly refers to a speed/time profile but in addition the engine throttle or motor accelerator can be made to follow a cycle of positions or settings. It is also possible to define a torque/time profile when testing the engine and motor torque control systems. Regardless of the cycle type, cycle parameters are transferred from the Duet to the M68000 during the initialisation phase. The operator can select any cycle which has been previously defined and stored on disk. A small BASIC program is used for cycle definition. When this program is run, the user is prompted firstly to input the number of data points followed by the time and demand level for each one. When the M68000 interprets the cycle it interpolates between each data point so that cycles consist of a number of straight line sections.

Having received all the basic system variables the M68000 carries out hardware initialisation, this consists of moving the throttle-servo to zero and placing the gearbox in neutral. After this stage further variables are transferred to the M68000. In this case they are mostly controller gains,

precisely what values are needed depends on the particular test, with the full set only being necessary when the completed hybrid control system is operated over a cycle. Once this stage is finished the initialisation phase of the software comes to an end.

2.5.2 Active Rig Control

With initialisation complete the main phase of both programs begins. In the M68000 the software divides into time critical tasks which are executed every 20 ms in the interrupt routine and non time critical tasks which run in the background whenever time is available. Immediately after initialisation the main programs give the operator at the Duet full manual control of the rig. By pressing appropriate command keys, as detailed on a menu, the user can operate the engine ignition, starter motor or change gear. A further Duet command key toggles the action of a small hand held potentiometer between control of engine throttle or motor accelerator demand. The signal from this potentiometer is fed back to the M68000's ADC board. All command keys cause a single character to be passed from the Duet to the M68000 which continually scans for such inputs in its background routine. A summary of the command keys and their functions is given in Table 2.4. Although processing a user request maybe delayed by an interrupt routine, the delay is so small that as far as the user is concerned his commands are obeyed immediately.

As indicated in Table 2.4 one of the user commands instructs the M68000 to commence the test cycle. Once this option has been selected all the manual control commands are overridden. At each interrupt throughout the cycle, the M68000 calculates cycle time by incrementing a counter on each execution. This time value is used by the interpolation routines which calculate the required setpoint at each point in the cycle. In a simple test

this may provide the demand signal to the throttle servo-system but there are other applications as mentioned previously. As the test proceeds relevant parameters are passed to the Duet for use in a real time display which may be graphical or numerical. These parameters are not passed on any strictly regular basis but depend on the processing time needed for the Duet to display each one, subject to any delays due to interrupts in the M68000. For example during a driving cycle test values of speed in km/h are passed to the Duet. These are displayed numerically and they are plotted over the cycle profile to give an immediate visual check on the accuracy of cycle following. Over a typical cycle time period the rate of on line data transfer is more than sufficient to fully utilise the resolution of the graphics screen. At any time during the cycle the operator has the option to abort it by pressing the appropriate key. If he takes up this option the rig is shut down by cutting the engine ignition and applying the regenerative brakes until the rig is stationary. At this stage the cycle may be restarted from the beginning with all control and data logging reset.

2.5.3 Data Logging

As soon as active rig control commences the M68000 begins gathering data from all sources at each interrupt. Thus data is always available for control as soon as the cycle starts or in some cases for user displays during manual control. Scanning the data sources gets first priority as soon as the interrupt occurs. Since this scanning process only consists of reading chip registers and transferring the measured values to an array, execution time is very small. Only raw transducer values are recorded so that no unecessary time is used in computation. During this initial data retrieval routine the watchdog timer is also reset.

Component control routines follow the first data logging routine ac-

cording to the individual program application. In the most complex case, this comprises fully integrated control of the engine motor and transmission to meet the demands of the test driving cycle. Due to the compact nature of the initial data logging the component controllers will have suffered minimal delay in receiving the necessary data and need only process those data sources which they require.

Before each interrupt routine ends, a final data logging routine is executed. This routine transfers those data sources selected by the user during initialisation to a large array which builds up in memory until the test is complete. With extended data logging this only happens after the appropriate number of interrupts have occured. Having final data logging after component control is advantageous in that control is not unecessarily delayed and in addition parameters derived during the control process may be recorded. A final refinement to the data logging software allows recording to begin at a specified point in a cycle as input from the Duet during initialisation. As a result it is possible to sample a small part of a long cycle at maximum sampling rate without building up an unecessary amount of data. This is particularly useful during transient experiments when often there is a large preliminary phase of the cycle, designed to allow steady state conditions to be reached, before the transient input is applied.

2.5.4 Data Transfer

The final phase of the software on both the M68000 and the Duet is concerned with data transfer. To start this process the user presses the appropriate command key and inputs the name of a disk file to be used by the Duet. At this point the M68000 scales all the raw transducer readings that have accumulated in its data array into meaningful values, for example the raw ADC reading from the manifold pressure transducer is converted to

mbar. During the transfer the Duet prompts the M68000 for each data value and immediately writes it to the disk file. Before the disk file is closed the user can input any number of comments concerning experimental conditions. Following this action the M68000 and Duet programs both end. It is possible to end both programs without going through the data transfer phase.

2.5.5 Data Analysis

Any disk file may be retrieved at a later date and then displayed graphically on the screen of the Duet using a BASIC graph plotting program. The nature of BASIC made it easy to refine the graphics display into a pleasing format. This feature has been exploited to produce a graph drawing program capable of coping with all the different data types present in experiments on the rig. The graphics routines begin by retrieving the sampling period and data code word from the disk file. This provides the information necessary to label the y-axes and also specifies what increment to use between points on the time-axis.

With two or more data sources the program can use two independently scaled y-axes. The program allows the operator to specify which data sets from the file are to be plotted on the same axis. It then retrieves all data points from the disk file and searches for maximum and minimum values to automatically scale the axes and therefore make best use of the screen size. As an alternative to plotting results against time, an additional program allows any pair of data sources to be plotted against each other. Apart from graphical display several data analysis utilities have been written, one of the most useful of these will produce a regression line for any pair of data sources.

2.6 Discussion

Foremost of the considerations taken into account when designing the rig has been retaining flexibility in operation. For this reason manual control has been retained as a possibility in almost all cases, as an alternative to computer operation. For example it is still possible to select a gear by hand when the rig is stationary, despite the fact that this meant designing an oversize shift mechanism. Clearly this is an advantage since many experiments are carried out in a fixed gear and there is then no need to operate the compressed air system merely to initially engage that gear. The only exception to this manual option rule is the engine throttle, which must be operated by the computer controlled stepper servo-system. Nevertheless no major modifications have been made to the carburetter and it would be straightforward to install the conventional manual linkage to the butterfly valve once again should the need arise.

As well as considering freedom of operation in the mechanical design, the software too has been structured to allow as much experimental flexibility as possible. As a result of this approach the M68000 software consists of a hard core of control and data logging routines which provide the basic framework for individual experiments. These core routines provide an interrupt programming area which carries out data logging of all available sources and a background routine which allows easy transfer of user commands from the Duet. To build up a program for a particular experimental application only requires that the necessary component control be added to the interrupt programming area. This is made easier by a number of basic component control routines such as those for stepper-servo control and gear changing. By adopting this approach there is no need to write a fundamentally different program whether the experiment is concerned with simply observing the effects of transient throttle movements or carrying out fully integrated control of all

the drive train components.

The microprocessor and electronics system on the rig represent what would be a vital part of any operational road vehicle based on the design. Only ten years ago this microprocessor control system and associated instrumentation, would have seemed exceedingly sophisticated when compared with conventional vehicles of the time, however this is almost certainly not the case today. As Westbrook pointed out in a review of automotive electronics [Westbrook, 1986], there has been a rapid increase in the electronics content of conventional vehicles over the last few years and this trend seems set to continue. In all complex control systems applied to vehicles to date, the main problem facing manufacturers has been achieving high reliability and low cost in the relatively harsh vehicle environment. Despite this difficulty the incentive of achieving better economy and reduced emissions has been sufficient to allow microprocessor based engine systems to be considered for use in moderately priced vehicles [Kawamura et al, 1985][Main, 1986]. These examples illustrate that the control system needed by a hybrid vehicle should not prove an insurmountable obstacle to its acceptance. Furthermore these systems also illustrate that the high volume production associated with the passenger car market even make it possible to design custom chips for automotive applications. Often this comes about as a result of close collaboration between vehicle and semiconductor manufacturers, as is the case with Nippondenso and Toshiba [Kawamura et al, 1985].

In the software field the attention of vehicle manufacturers has focussed on reliability and efficient use of memory space [Srodwa et al, 1984]. When the low cost of electronic memory is considered this latter consideration might seem surprising, but it is merely indicative of the highly competitive state of the vehicle market. In the above paper, Ford discuss the advantages of writing small engine control programs in assembler rather than using the usual

combination of high level language plus compiler. They state that the added difficulty in producing the coding is worthwhile even if the resultant program is only 25% smaller than the high level language version. It is however unlikely that this advantage would remain as control structure becomes more complex, as it would be in a hybrid vehicle. Since conventional vehicles are almost certain to incorporate an equivalent level of sophistication in their control systems the hybrid should not suffer any disadvantage in this respect.

Transducers are another area which are vital to the operation of the rig, since without them the complex control strategies necessary for optimum control of the hybrid drive train would be impossible. The reliability and cost of these devices has an important bearing on whether the hybrid vehicle control concept, investigated on the rig, could ever be put into commercial production. As far as cost is concerned the transducers are often far more important than the controlling electronics; typically they account for 65% of the cost in an engine management system [Westbrook, 1986]. Fortunately all the engine transducers used on the hybrid vehicle rig have already been successfully employed in conventional production vehicles. A summary of the state of engine sensor technology in 1985 was carried out by Wolber and Ebaugh [Wolber and Ebaugh, 1985]. As this paper shows both magnetic speed sensors and the strain gauge manifold pressure transducers have passed the twin tests of low production costs and reliability in the engine operating environment. When it is realised any component operating in the engine compartment must survive temperatures ranging from $-40°C$ to $120°C$, as well as the effects of salt spray, it is apparent what a stiff test of transducer, and electronic technology, this represents. In the case of the motor the only transducers required by the hybrid drive system are the Hall effect current probes. These devices have long been an essential part of electric traction control systems and indeed one is included in the Lucas Chloride power

electronics control unit on the rig.

The purpose of the preceding discussions has been to show that the microprocessor based control system envisaged for a hybrid vehicle would not pose any extreme difficulties to vehicle manufacturers. Possibly a far more important economic consideration is the viability of having two power sources in the drive train. However this is an exceedingly difficult question to answer and lies beyond the scope of the work described here. In this case the primary goal is to establish the technical feasability of blending the various operating modes possible with a hybrid drive train. The rig system described in this chapter provides the necessary test environment to achieve this, however it represents only a beginning. In the next chapter the design of the control system begins in earnest with an analysis of the relevant physics of the engine and motor.

2.7 Software Documentation

For reasons of commercial confidentiality software listings are not included in this thesis. However full listings and supporting documentation are available within the School of Engineering and Applied Science. This documentation includes an operating guide for the rig and circuit diagrams for all the transducer circuitry.

Component	Description
Traction Motor	Lucas Chloride separately excited d.c. motor, Type MT286 rated 37 kW (1/2 hr)
Motor Control	Lucas Chloride Type Mk. IIIB current controlled SCR Armature Chopper and transistor field chopper
Batteries	Lucas Chloride Type EV5C, 216v, 184 Ah (5hr rate)
Engine	Ford 1100cc petrol engine 32 kw at 5,500 r.p.m. 71 Nm at 3,000 r.p.m.
Transmission	1st 3.656:1 2nd 2.000:1 3rd 1.425:1 4th 1.000:1
Flywheel	Variable inertia 2.02 to 15.57 kgm^2
Dynamometer	Froude Consine EC38TA water cooled dynamometer Max. torque 475 Nm

Table 2.1 Test Bed Component Ratings

Signal Name	Interface Hardware	Signal Type
Engine Throttle Command	2 × 6522 VIA Chips	Pulse Train to Stepper Motor drive card
Motor Accelerator Command Motor Brake Command	12 bit DAC "	Analogue signal to power electronics controller "
Starter Motor	6522 VIA o/p port	1=ON 0=OFF
Gear shift signals	6522 VIA o/p port	ON/OFF signals to solenoid valves
Engine ignition	6522 VIA o/p port	1=ON 0=OFF
Engine Throttle direction signal	6522 VIA o/p port	1=Move to Full 0=Move to Zero
Motor Accelerator/ Brake enable/disable	6522 VIA o/p port	1=Enable 0=Disable

Table 2.2 M68000 Output Control Signals

Signal Name	Interface Hardware	Signal Type
Engine Speed	VIA Counter	Pulse Train from 135 tooth engine starter ring
Motor Speed	STC Counter	Pulse Train from 90 tooth wheel
Flywheel speed	STC Counter	Pulse Train from 60 tooth wheel
I.C. engine fuel flow	STC Counter	Pulse Train from turbine device
Inlet manifold depression	ADC	Integrated strain gauge bridge 0-100mV=0-300 Nm
Motor and Engine Torque	ADC	Strain gauge Transducer
Motor Field and Armature current	ADC	Hall effect current transducers 0-10V=0-200A-turns
Motor Field and armature voltage	ADC	Voltage transposers Transfer ratio 25:1
Gear lever Sensor switches	VIA Input Port	1=Active 0=Inactive
Manual accelerator signal	ADC	0-5V from hand held pot.

Table 2.3 M68000 Input Data Signals

Command Letter	M68000 Response
D	Commence Cycle
A	Abort cycle:- Engine throttle set to zero, ignition off, regenerative braking until flywheel is stationary
I	Toggle engine ignition
U	Record data sources requested during initialisation
Q	Stop data recording:- Engine throttle set to zero, ignition off, stop motor, disable motor brake and accr., stop ADC interrupts, disable stepper motor, exit program
R	As Q but transfer data to Duet prior to program exit
M	Toggle manual pot. between motor accr. and engine throttle control
S	Toggle engine starter motor
0-4	Move into neutral or appropriate gear

Table 2.4 Operator Command Keys

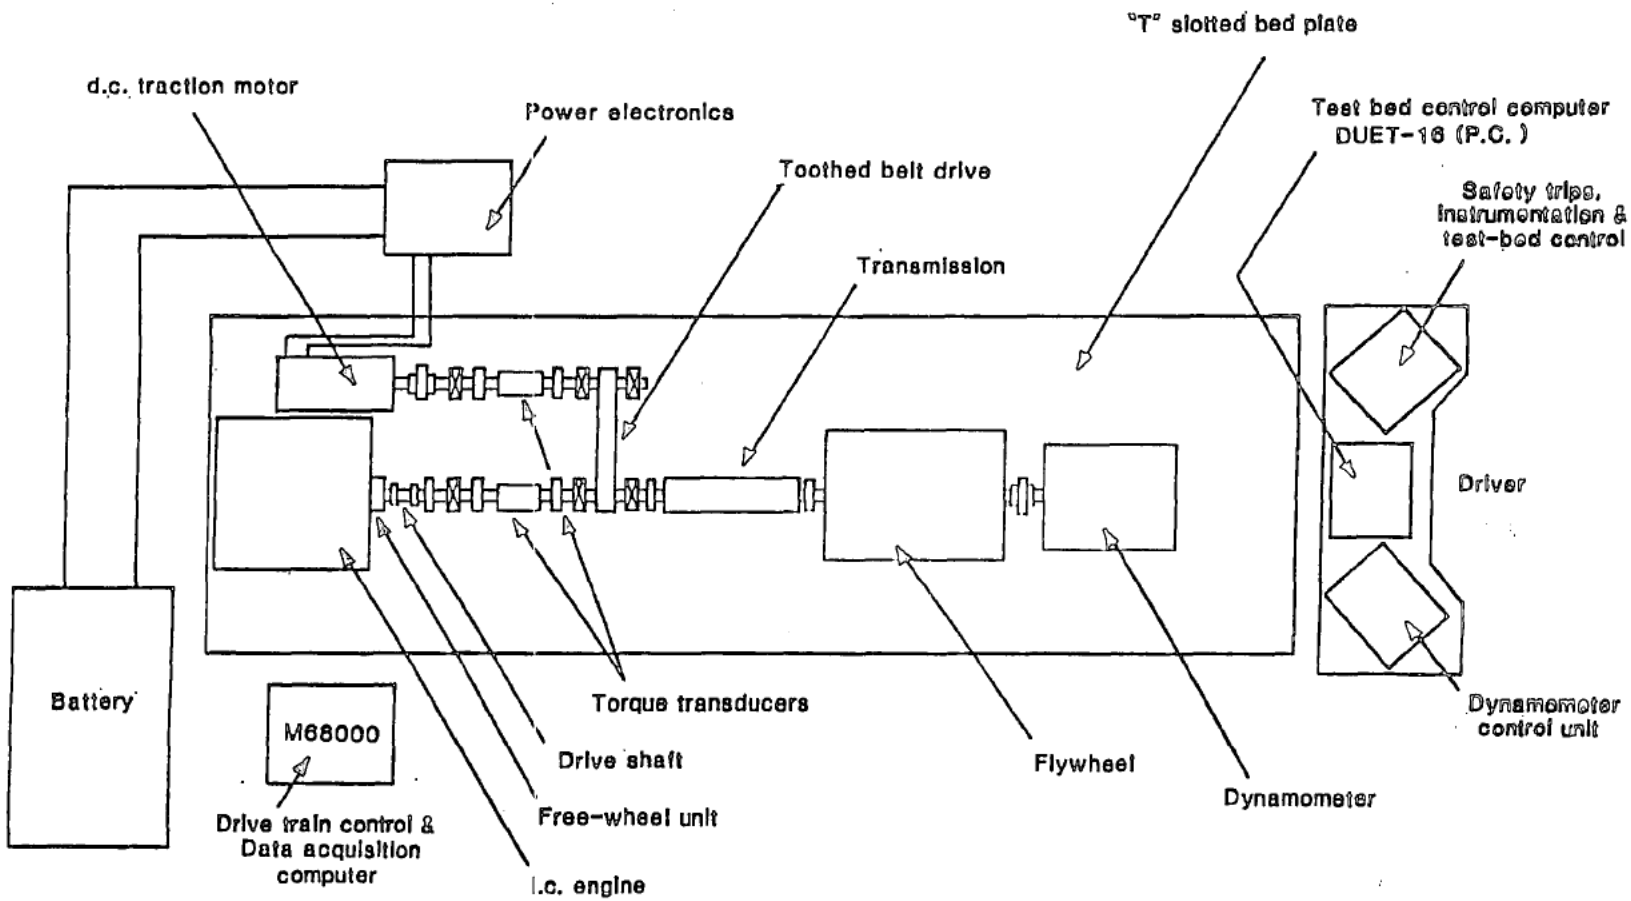

Fig. 2.1 Test Bed Layout

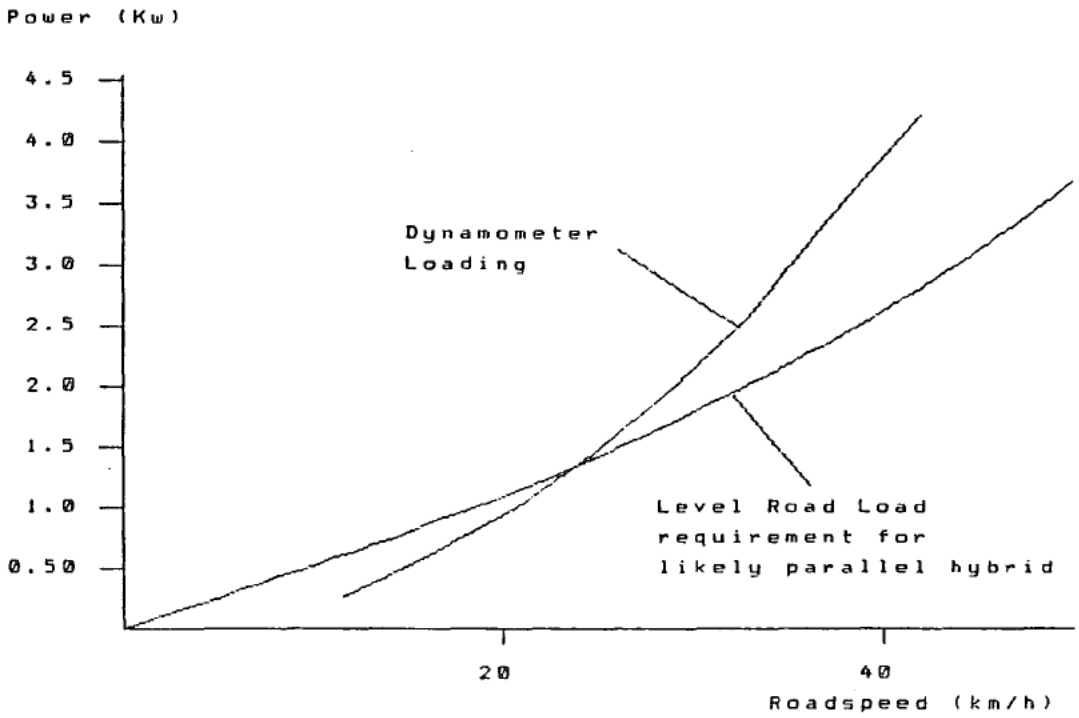

Fig. 2.2 Power Loading Provided by the Dynamometer

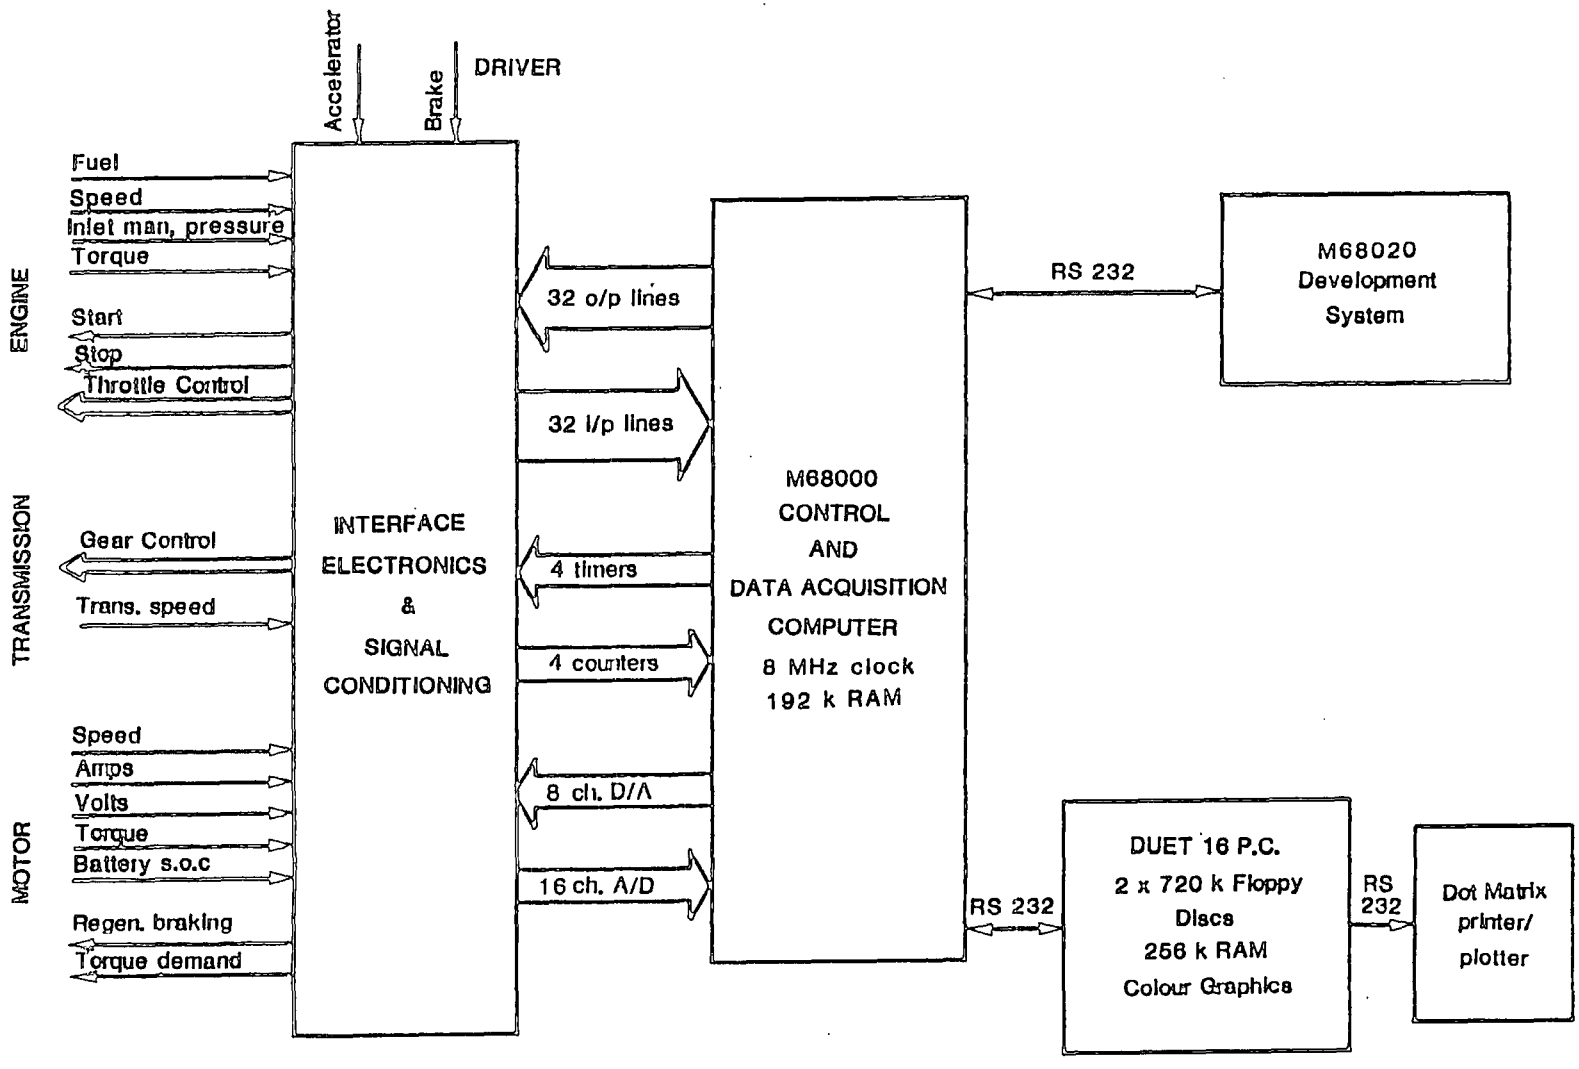

Fig. 2.3 Test Bed Computing and Interface Equipment

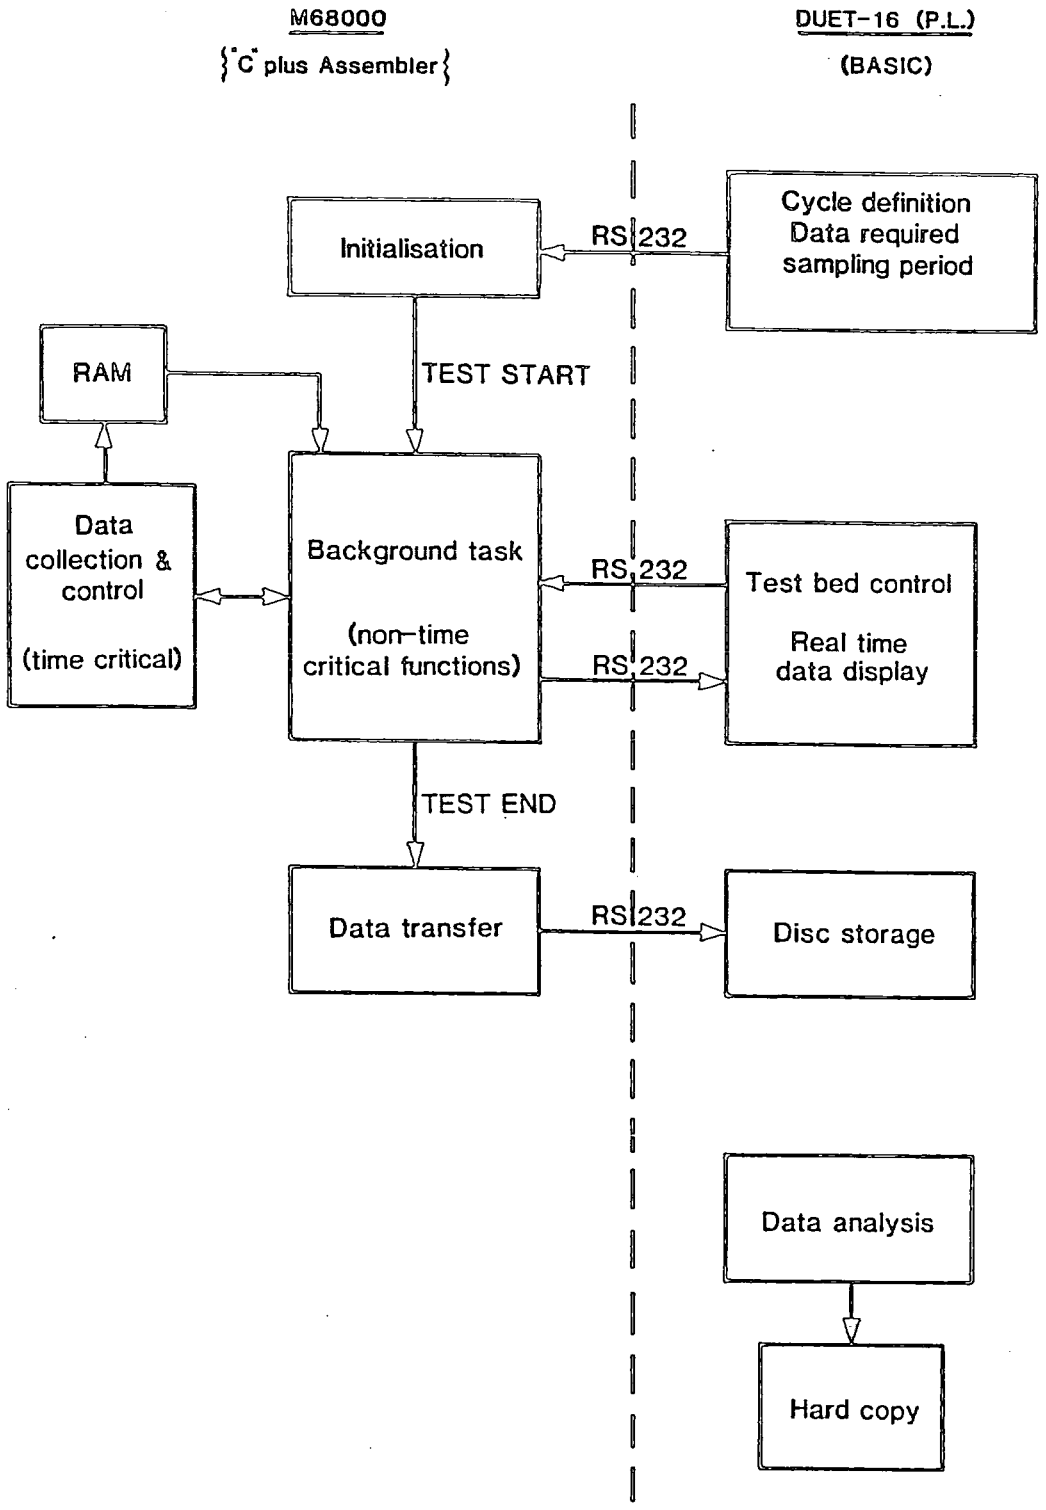

Fig. 2.4 Software Structure

CHAPTER 3

PHYSICAL ANALYSIS OF THE ENGINE AND MOTOR

In the introduction the concept of the various possible operating modes of the parallel hybrid drive system was put forward. Once a given mode has been selected the necessary torque and gear shift control must be undertaken to actually bring about the chosen operating state. To achieve this goal the primary requirements are engine and motor torque control and transmission shifting. A further control sub-system is needed to start the engine when required.

The simplest way of providing torque control for either the engine or the motor is to monitor the signals from the torque transducers and then adjust the demand to either the engine throttle servo system or the motor power electronics, until measured output matches demand. Unfortunately torque transducers, such as those used on the rig, are not practical in a production vehicle due to their prohibitive cost. This problem would be likely to remain even if the devices were manufactured in the large numbers applicable to the mass car market. As a result, torque measurements must be made which are based on cheap, robust sensors. Once these alternative methods are devised and calibrated they can be used as the feedback signals to closed loop torque control systems. On the hybrid vehicle rig the indirect torque measurement method for the engine is based on inlet manifold depression and speed, [Masding and Bumby, 1988 (b)] whilst the equivalent system for the motor uses values of armature and field currents [Masding and Bumby, 1988 (c)].

Before any control algorithms can be designed, mathematical models of both the i.c engine and the motor must be developed. These mathematical models will describe the dynamics of the relationship between the controlled variables, which are the calculated torque outputs of the engine and motor,

and the adjustable inputs which are throttle setting and demand voltage to the power electronics unit. All the control is carried out in the M68000 using discrete control algorithms, suggesting that the most appropriate form for the mathematical models is discrete transfer functions. Identification techniques described in the next chapter make it possible to arbitrarily assume what order of transfer function might be appropriate to say, the relationship between motor torque and accelerator demand and then determine the necessary coefficients from experimental data. Using this technique however brings with it the danger of underestimating the order of a key transfer function and hence eventually to poor controller design. Less seriously overestimating the order of a transfer function needlessly increases the identification time. Consequently an important first step towards digital transfer function identification, is an analysis of the physics of both the engine and motor. Such an analysis is best carried out in the continuous s-plane which more easily relates to most of the systems involved. Having established, through this analysis, what order of transfer function is appropriate to each plant the equivalent digital systems have a sound analytical foundation. In the next two sections analysis of first the motor and then the engine will be brought to the stage where an appropriate digital transfer function can be written down to relate the individual control signal to the resulting torque output as calculated through the indirect method. In addition a similar result is obtained which relates engine speed under no load conditions to throttle setting. At the conclusion of this process the coefficients of these digital transfer functions are unknown however in chapter 4 they are successfully determined on the basis of experimental input/output data.

3.1 Analysis of the Electric Traction System

The main elements which make up the electric traction system are shown in figure 3.1. The motor field current is controlled by a transistor field chopper while the armature current is controlled using a thyristor (SCR) chopper unit. Both systems operate at a nominal battery voltage of 216v. Connected in the armature circuit is a large inductance for regenerative braking purposes and to ensure continuous current flow. The M68000 microprocessor control unit interfaces with the power electronic controller by outputting an accelerator or brake demand signal, θ_m, to the power electronic controller via two DAC's. All identification and analysis is carried out in terms of this variable θ_m, which is the number set in the DAC register by the M68000. Any relationship involving θ_m can be converted to volts by taking into account that numbers in the DAC registers range from 0-4095 and this corresponds to 0-5 volts. Raw output from the DAC's is modified by buffer amplifiers which have a gain of 0.68, in the case of the accelerator signal, and 0.56 in the case of the brake. According to the original design for the electric traction system the brake and accelerator signals originate from potentiometers moved by the driver's pedals. Supply voltage for this potentiometer comes from the power electronics and does not exceed 3.5 volts in the case of the accelerator or 3 volts in the case of the brake. Consequently not only do the buffer amplifiers isolate the computer DAC from direct contact with the power electronics, but in addition the attenuation they provide ensures that even the full DAC voltage can not exceed the design limits for the controller.

3.1.1 D.C. Machine Analysis

The basic equations governing the behaviour of any d.c. machine are

$$v_a = (R_a + L_a s) i_a + e_a \tag{3.1}$$

$$v_f = (R_f + L_f s) i_f \tag{3.2}$$

$$e_a \propto \phi_f N \tag{3.3}$$

$$T_{em} \propto \phi_f i_a \tag{3.4}$$

Equations (3.1) and (3.2) can be expressed in transfer function form as:

$$\frac{i_a}{v_a - e_a} = A(s) = \frac{1}{R_a + L_a s} \tag{3.5}$$

$$\frac{i_f}{v_f} = F(s) = \frac{1}{R_f + L_f s} \tag{3.6}$$

whilst ϕ_f, the field flux, can be expressed as a function of the field current, i_f, i.e.

$$\phi_f = f(i_f) \tag{3.7}$$

If no saturation of the magnetic circuit is present, then

$$\phi_f \propto i_f \tag{3.8}$$

In road vehicle applications the traction motor must provide the necessary propulsion torque to overcome vehicle inertia effects aerodynamic drag and vehicle rolling resistance such that

$$T_{em} = T_l + J s \omega \tag{3.9}$$

where T_l is the vehicle aerodynamic and rolling resistance loss torque and J is the vehicle inertia both referred to the motor shaft.

In his paper on electric machines and power systems, Barton [Barton, 1987] elegantly represents these equations in the block diagram form of figure 3.2. Unfortunately this system is non-linear because of the presence of the two multiplier blocks and the saturation function which relates field current

to field flux. Examination of figure 3.2 suggests that two basic methods are available for controlling the torque output, and hence speed, of the electric traction system. Firstly, field voltage can be maintained constant and all control carried out on the armature via the armature chopper. With field current constant the motor speed is essentially proportional to the applied voltage so that, assuming constant current flow, as the applied armature voltage is varied so too is the motor speed. Consequently with full field current applied to the field winding motor speed can be varied from zero to break speed by controlling the voltage applied to armature terminals via the armature chopper. When full armature voltage is applied the motor will run at break speed. In this condition the multiplier blocks reduce to a simple gain which produces a linear transfer function relating both motor torque and motor speed to the controlled armature voltage. In the second basic control mode armature voltage can be maintained constant while the field current is varied. Because motor speed is inversely proportional to field current, equation (3.3), this field weakening mode is usually employed to obtain motor speeds in excess of the break speed. To provide smooth control in this field weakening, high speed, region some degree of armature control may also be necessary.

In some instances a third mode of operation maybe used at speeds below break speed in which both the field current and the armature current are varied together, in some way, so as to improve the overall response of the motor/controller system.

3.1.2 Motor Torque Equation

In any D.C. motor air-gap torque is related to armature current and field flux by equation (3.4) which, if saturation is not present, reduces to

$$T_{em} = K_T i_f i_a \tag{3.10}$$

As machine saturation is mainly dependent on the value of field current, a more accurate statement of equation (3.10) would be

$$T_{em} = K_T(i_f) i_a i_f \tag{3.11}$$

3.1.3 Effect of the Power Electronic Controller

Initial tests on the power electronic control unit showed that all three modes of control described in section 3.1.1 are used to control the motor over its complete speed range. Each of the three modes of operation used can be characterised by the behaviour of the field current as shown in figure 3.3. Over the range $\theta_m = 0 - 250$ both field and armature current are directly controlled. This mode, termed the field boost mode, is characterised by the rapidly increasing field current. At speeds below 2500 r.p.m. the field boost mode ends when i_f reaches 8.5A the motor then moving into a full field mode where field current is constant and all control takes place in the armature. The motor remains in this full field mode for all higher values of θ_m provided that the speed is no greater than about 2000 r.p.m. At higher speeds the field weakening mode, characterised by gently falling field current, is entered. The precise value of θ_m at which this happens is speed dependent, for example at 2300 r.p.m. $\theta_m = 500$ while at speeds above 2500 r.p.m., the motor break speed, full field mode is not used and the motor switches directly from field boost to field weakening at an initial field current value that depends on speed.

In the power electronic control unit both field and armature current are controlled. Examination of the circuit diagram for the power electronics shows that a Hall effect device is used to derive the feedback signal for control

of armature current, whilst a current transformer is used for control of field current. Each closed loop system uses a set point based on the accelerator demand. Combining this information with the governing equations, (3.1)-(3.6), produces the control block diagram of figure 3.4.

The set point in both the field and armature current control loops is calculated from the accelerator input, θ_m, by the set point functions $\alpha(\theta_m)$ and $\beta(\theta_m)$ respectively both of which depend on the mode of operation.

Analysis of this block diagram allows the relationship between θ_m and motor torque, T_{em}, to be determined. Consider first the field current control loop. Here a compensator, $C_f(s)$, outputs the voltage, $v_f(s)$, to the field windings on the basis of the current error which it measures. This action results in a certain closed loop current response for the compensator plus field windings which may be defined as $G_f(s)$, where:

$$G_f(s) = \frac{C_f(s)F(s)}{1 + C_f(s)F(s)} \tag{3.12}$$

hence the response of field current to the accelerator setting is

$$i_f(s) = \alpha(\theta_m(s))G_f(s) \tag{3.13}$$

and for small changes in demand, $\Delta\theta_m$,

$$\Delta i_f(s) = \Delta\alpha(\theta_m(s))G_f(s) \tag{3.14}$$

Similarly for the armature loop, the set point current is $\beta(\theta_m(s))$, and the closed loop current response for the armature in the absence of $e_a(s)$ is

$$G_a(s) = \frac{i_a(s)}{\beta(\theta_m(s))} = \frac{C_a(s)A(s)}{1 + C_a(s)A(s)} \tag{3.15}$$

giving the full armature loop response in the presence of $e_a(s)$ as:

$$i_a(s) = G_a(s)\left(\beta(\theta_m(s)) - \frac{K_e N i_f(s)}{C_a(s)}\right) \tag{3.16}$$

To evaluate equation (3.16) it is necessary to know the transfer function of the current controller, $C_a(s)$. First order op-amp compensators are used by the power electronics in both the field and armature current control loops so that the form of the armature controller is:

$$C_a(s) = \frac{K_p}{1 + \tau_p s} \tag{3.17}$$

Hence by substituting for the controller and also for the field current, $i_f(s)$, as defined by equation (3.13) the complete armature current response is obtained as

$$i_a(s) = G_a(s)\{\beta(\theta_m(s)) - (1 + \tau_p s)N\alpha^*(\theta_m(s))G_f(s)\} \tag{3.18}$$

where $\alpha^*(\theta_m(s)) = \frac{K_e}{K_p}\alpha(\theta_m(s))$

During normal operation any change in motor speed occurs only slowly due to the large vehicle inertia. In contrast current changes occur rapidly so that during current control motor speed can be regarded as constant at the initial value N_0, so that, for small changes in demand

$$\Delta i_a(s) = G_a(s)\{\Delta\beta(\theta_m(s)) - N_0(1 + \tau_p s)\Delta\alpha^*(\theta_m(s))G_f(s)\} \tag{3.19}$$

The importance of equations (3.13) and (3.18) lies in the way in which they can be used to relate the accelerator demand to the motor currents and hence to the output torque via the torque equation, (3.11). Particularly important is the linearised form of this equation which is given by

$$\Delta T_{em}(s) = K_T(i_f)(i_{a(0)}\Delta i_f(s) + i_{f(0)}\Delta i_a(s)) \tag{3.20}$$

Substituting for $\Delta i_f(s)$ and $\Delta i_a(s)$ from equations (3.14) and (3.19) gives:

$$\Delta T_{em}(s) = K_T(\bar{i}_f)[G_a(s)\{\Delta\beta(\theta_m(s)) - (1+\tau_p s)N_0\Delta\alpha^*(\theta_m(s))G_f(s)\}i_{f(0)} + G_f(s)\Delta\alpha(\theta_m(s))i_{a(0)}] \tag{3.21}$$

Provided that the setpoint functions, $\Delta\alpha(\theta_m(s))$ and $\Delta\beta(\theta_m(s))$, may be written as $K\Delta\theta_m(s)$, where K is some constant, then equation (3.21) describes a linear transfer function relating $\Delta\theta_m(s)$ to $\Delta T_{em}(s)$. Such a result is the key to the control of motor torque since it relates the output to be controlled to the only available input, $\Delta\theta_m$. To evaluate the expression for $\Delta T_{em}(s)$ two channels of identification are required.

A. Steady state experiments to determine the set point functions $\Delta\alpha(\theta_m)$, $\Delta\alpha^*(\theta_m)$, $\Delta\beta(\theta_m)$ and to calibrate the torque relationship $K_T(\bar{i}_f)$.

B. Transient experiments to determine the unknown dynamic elements $G_a(s)$, $G_f(s)$ and $(1+\tau_p s)$.

Both the steady state and transient identification experiments are presented in chapter 4. The transient experiments are used to obtain digital transfer functions for the dynamic elements. Before such identification can begin however it is necessary to transform the s-domain transfer functions into the correct z-domain form.

Equation(3.15) defines $G_a(s)$ in the abscence of $e_a(s)$. During identification experiments $e_a(s)$ can be made constant so that when difference values are considered the same expression applies.

$$G_a(s) = \frac{\Delta i_a(s)}{\Delta\beta(\theta_m(s))} = \frac{C_a(s)A(s)}{1 + C_a(s)A(s)} \tag{3.22}$$

Substituting for the first order compensator included in the power electronics, equation (3.17), reveals the form of the s-plane transfer function for $G_a(s)$

$$G_a(s) = \frac{\Delta i_a(s)}{\Delta\beta(\theta_m(s))} = \frac{K_p/R_a}{(1 + L_a/R_a s)(1 + \tau_p s) + K_p/R_a} \tag{3.23}$$

A zero order hold is included in series with the transfer function $G_a(s)$ to model the digital to analogue converter to give the pulse transfer function

$$G_a(z) = (1 - z^{-1})Z\left(\frac{G_a(s)}{s}\right) \tag{3.24}$$

with result:

$$G_a(z) = \frac{b_1 z + b_2}{z^2 - a_1 z - a_2} \tag{3.25}$$

A similar equation to equation (3.25) holds for $G_f(z)$.

3.2 Physical Analysis of the I.C. Engine

3.2.1 Inlet Manifold Pressure Variations

As previously mentioned the engine torque may be calculated from speed and inlet manifold depression. This method was successfully applied for the HTV-1 [Somuah et al, 1983]. The analysis of the engine must therefore investigate the physical processes governing air and fuel flow through the inlet manifold. All the relevant physical processes are shown in the block diagram of figure 3.5. An air/fuel mixture flows through the inlet manifold into the cylinders with the volume of the inlet manifold acting as a plenum chamber smoothing out pressure changes as the flow rate varies. When the throttle suddenly moves the manifold pressure will change over a number of engine cycles. Essentially the rate of flow of the air/fuel mixture into the inlet manifold depends on the throttle opening and the inlet manifold depression [Morris et al, 1981]. As the air fuel ratio is of the order of 15:1 this flow is dominated by the mass air flow $\dot{m}_a$. The mass charge flowing out of the inlet

manifold, $\dot{m}_c$, depends on manifold pressure and engine speed. Balancing mass flows in and out of the inlet manifold gives:

$$\frac{d}{dt}m_m = \dot{m}_a - \dot{m}_c \tag{3.26}$$

Where m_m is the mass of the gaseous mixture in the inlet manifold. Assuming ideal gas conditions then the mass, m_m, can be related to the volume and temperature of the inlet manifold by:

$$m_m = \left(\frac{V_m}{R_m T_m}\right) p_m \tag{3.27}$$

giving

$$\dot{m}_a - \dot{m}_c = \left(\frac{V_m}{R_m T_m}\right) \frac{dp_m}{dt} \tag{3.28}$$

As explained above, the mass charge flowing into the cylinders is dependent on the manifold pressure and the engine speed and for small changes can be expressed as [Morris et al, 1981],

$$\Delta\dot{m}_c = k_N(N, p_m)\Delta N + k_{pc}(N, p_m)\Delta p_m \tag{3.29}$$

while the mass airflow into the inlet manifold is dependent on throttle opening and manifold pressure, which for small changes, gives:

$$\Delta\dot{m}_a = k_\theta(\theta, p_m)\Delta\theta - k_{pa}(\theta, p_m)\Delta p_m \tag{3.30}$$

Combining equations (3.27)-(3.30) gives the change in inlet manifold pressure as

$$\Delta p_m(s) = \frac{K_f}{\tau_f s + 1}(k_\theta(\theta, p_m)\Delta\theta(s) - k_N(N, p_m)\Delta N(s)) \tag{3.31}$$

Where the manifold filling time constant

$$\tau_f = \frac{V_m}{R_m T_m (k_{pc}(N, p_m) + k_{pa}(\theta, p_m))} \tag{3.32}$$

and the associated gain

$$K_f = \frac{1}{k_{pc}(N, p_m) + k_{pa}(\theta, p_m)} \tag{3.33}$$

both of which are dependent on engine speed, manifold pressure, and throttle opening. Finally, because of the nature of the Otto cycle, a power stroke delay is present due to the time between filling a cylinder and combustion torque being produced. This power stroke delay amounts to one revolution of the engine, as represented by the term $e^{-\frac{60s}{N}}$ on figure 3.5. Physically this means that following any change in manifold depression, engine torque is delayed by an amount inversely related to speed ranging from 15 ms at 4000 r.p.m. to 60 ms at 1000 r.p.m. In their analysis Morris et al, defined the sampling period for the digital engine model as the time for one engine revolution, arguing that this was the fundamental time period of the engine. As a result the power stroke delay could be simply modelled by the unit delay operator z^{-1}. Unfortunately this step did not mean that the other terms in the digital model were independent of speed. In the context of the present system modelling the power stroke delay poses a difficult problem in that a fixed sampling period must be used regardless of engine speed. Not to do so would complicate software design and give a sample period irrelevant to the traction motor. Consequently the power stroke delay is dropped in all subsequent analysis of the engine and it is assumed that torque appears simultaneously with variations in manifold depression. Even under the worst case assumption of low speed engine operation the delay amounts to only 3 system sampling periods. Such a small discrepancy between the torque that

the controller assumes that the engine is producing and that which it actually is, would be hardly noticeable to the driver. The power stroke gain, K_{qt} relates changes in the engine torque to the change in the manifold pressure and, as the experimental results in chapter 4 will show, it is dependent on engine speed.

3.2.2 Transfer Function for Engine Torque

The objective of the engine analysis is to relate the combustion torque to throttle position thus allowing a torque control system to be designed. Combining the result of equation (3.31) with the fact apparent from figure 3.5 that

$$\Delta T_{ic} = K_{qt}(N)\Delta p_m \tag{3.34}$$

achieves this goal, in the form of the following equation:

$$\Delta T_{ic}(s) = \frac{K_f}{1 + s\tau_f}(k_\theta(\theta, p_m)\Delta\theta(s) - k_N(N, p_m)\Delta N(s))K_{qt}(N) \tag{3.35}$$

The above expression can be considerably simplified by noting that over the time scale of manifold pressure and torque variations, the large vehicle inertia will maintain the speed constant at some value N_0. As a result it is possible to directly relate torque to throttle variations by a simple transfer function.

$$\frac{\Delta T_{ic}(s)}{\Delta\theta(s)} = \frac{K_f}{1 + s\tau_f}k_\theta(\theta, p_m)K_{qt}(N_0) \tag{3.36}$$

The gain function $K_{qt}(N_0)$ is best measured under steady state conditions as part of the calibration of the whole manifold pressure, speed, torque relationship. Omitting this gain leaves the first order lag element, called the manifold filling delay, defining the dynamic relationship between two easily measured variables.

$$M_{fd}(s) = \frac{\Delta p_m(s)}{\Delta\theta(s)} = \frac{K_f k_\theta(\theta, p_m)}{1 + s\tau_f} \tag{3.37}$$

In practice manifold pressure changes from this element are not observed directly, they must be filtered to avoid aliasing problems at the sample rate of 50Hz. This is achieved by an active second order Butterworth filter. As the response of the filter is much faster than the manifold filling delay, equation (3.37) remains valid for the filtered pressure values, $\Delta\hat{p}_m$, measured by the computer.

$$\frac{\Delta\hat{p}_m(s)}{\Delta\theta(s)} = M_{fd}(s) \tag{3.38}$$

To aid in the identification process a discrete version of equation (3.37) is necessary which, in conjunction with the throttle transfer function, developed in section 3.2.3, will allow changes in inlet manifold pressure to be related to demand throttle opening. However the transfer function identified for the throttle unit only relates discrete values of throttle position at each sample interval whereas the throttle position applied to the engine is continuously changing during the sample period. Consequently some method of reconstructing a continuous signal from the discrete values produced by the throttle pulse transfer function must be included in the discretisation of the engine transfer function. This is best achieved by incorporating a triangular hold in series with $M_{fd}(s)$, such that

$$M_{fd}(z) = \frac{(z-1)^2}{T_s z} Z\left\{\frac{M_{fd}(s)}{s^2}\right\} \tag{3.39}$$

where T_s is the sampling period. Evaluating the z-transforms gives:

$$\frac{\Delta\hat{p}_m(z)}{\Delta\theta(z)} = \frac{b_0 + b_1 z^{-1}}{1 - a_1 z^{-1}} \tag{3.40}$$

Once $K_{qt}(N)$ is determined variations in engine torque can be predicted from throttle movements.

3.2.3 Throttle Servo-System

As explained in chapter 2 the stepper motor controlling the engine throttle moves 0.9 degrees per step at a rate of 4 steps/sample period (20ms). Consequently should a throttle step demand of less than four steps be requested the throttle will be seen to reach the new position at the next sample instant. Due to the inherent discrete action of the stepper motor system this simply translates to a pulse transfer function without the intermediate s-plane analysis:

$$G_T(z) = \frac{\Delta\theta(z)}{\Delta\theta_d(z)} = \frac{1}{z} \tag{3.41}$$

With a change in excess of four steps per sample however, the throttle will not have reached its final destination after one sample and the new, updated position will simply be four steps in advance of the old. This process results in the block diagram for the throttle system as shown in figure 3.6. Full throttle opening occcurs at 85 steps and appears as an output limit block. For step changes of less than or equal to four steps/sample the transfer function reduces to that of equation (3.41) and is that used by Morris et.al. [Morris et al, 1981]

3.2.4 Engine Speed on No-Load

The i.c engine is only started when the drive system is rotating at speeds above 1000 r.p.m. Consequently to bring the stationary engine on line it is necessary to automatically start the engine and match its speed with the moving drive train. Design of a controller to achieve this requires a model of the engine which relates engine speed to throttle position when the engine

runs on no load. Examination of figure 3.5 suggests that, ignoring the small power stroke delay, a second order transfer function

$$\frac{\Delta N(s)}{\Delta\theta(s)} = \frac{K_f K_2 K_{qt} k_\theta(\theta, p_m)}{(1 + s\tau_f)(1 + s\tau_2) + K_f K_2 K_{qt} k_N(N, p_m)} \tag{3.42}$$

will relate throttle angle to speed. This model can not be expected to give entirely accurate results since the manifold filling delay terms, τ_f and K_f, are themselves speed dependent as are $k_\theta(\theta, p_m)$, $k_N(N, p_m)$ and the power stroke gain, K_{qt}. Of these, variations in the gains K_f and K_{qt} will be the most significant because changes in the time constant, τ_f, will always be swamped by the delay, τ_2, due to the engine inertia. Indeed it is the engine inertia which dominates the speed response, and by its slowness further justifies the omission of the power stroke delay in this case.

3.2.5 Transfer Function for Engine No-Load Speed

Using a triangular hold and discretising the second order transfer function of equation (3.42) gives:

$$\frac{\Delta N(z)}{\Delta\theta(z)} = \frac{b_0 + b_1 z^{-1} + b_2 z^{-2}}{1 - a_1 z^{-1} - a_2 z^{-2}} \tag{3.43}$$

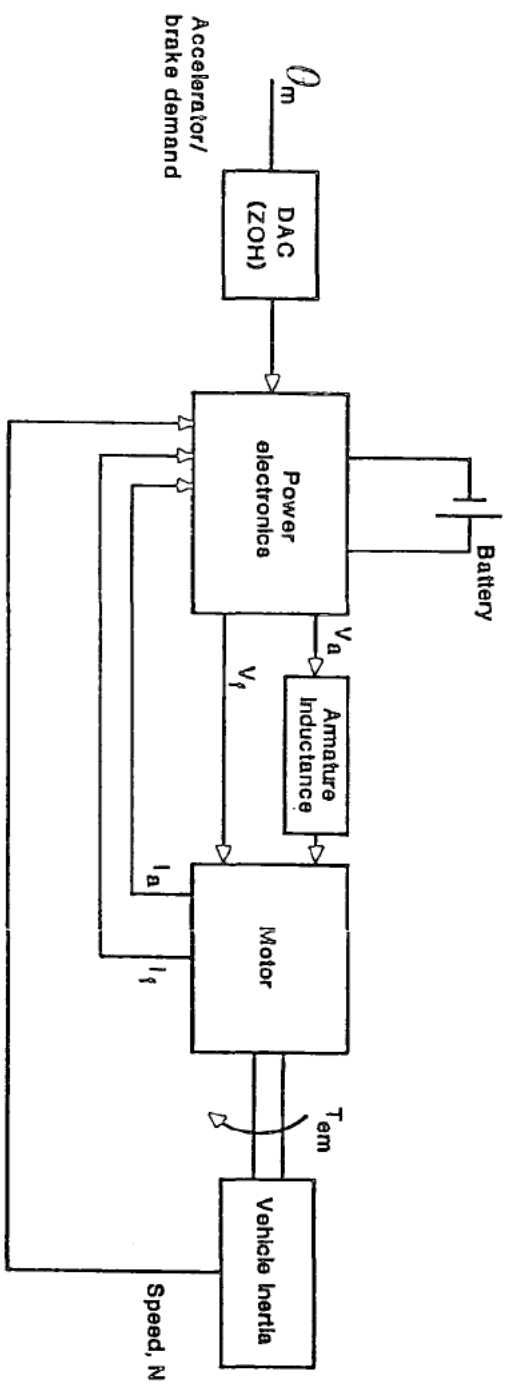

Fig. 3.1 Elements of the Electric Traction Drive

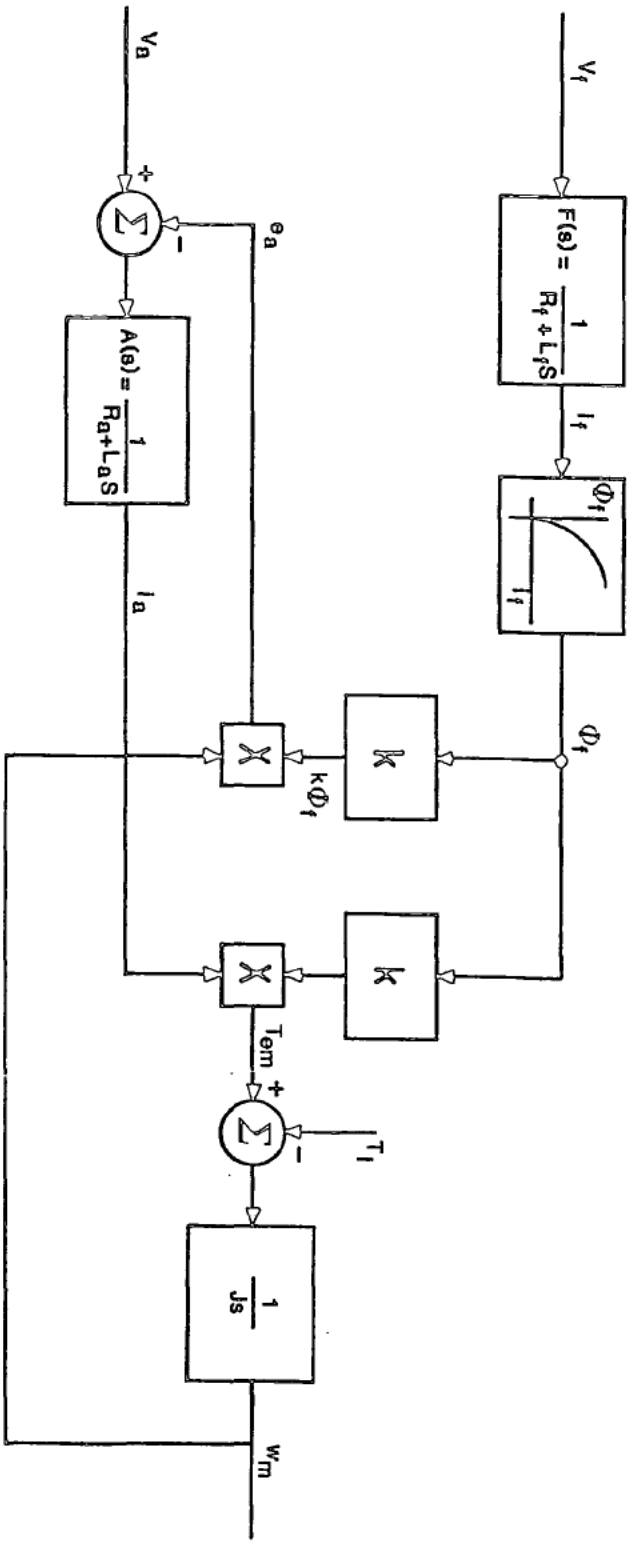

Fig. 3.2 Block Diagram of Electric Motor Equations

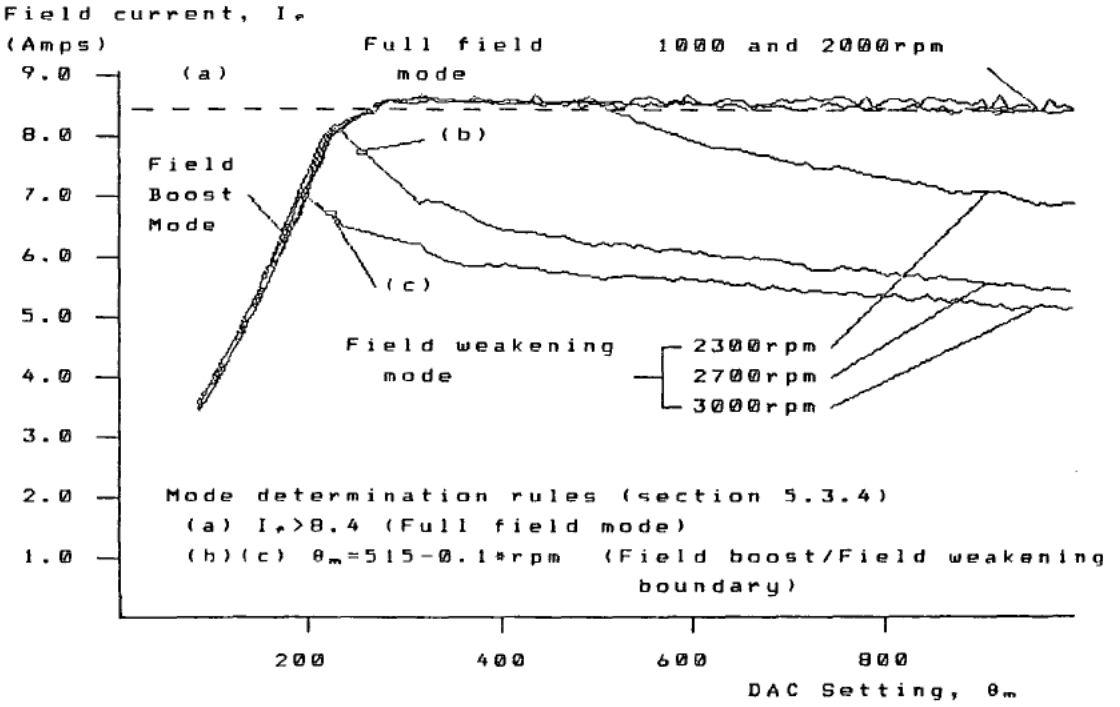

Fig. 3.3 Variation of Field Current With Operating Mode

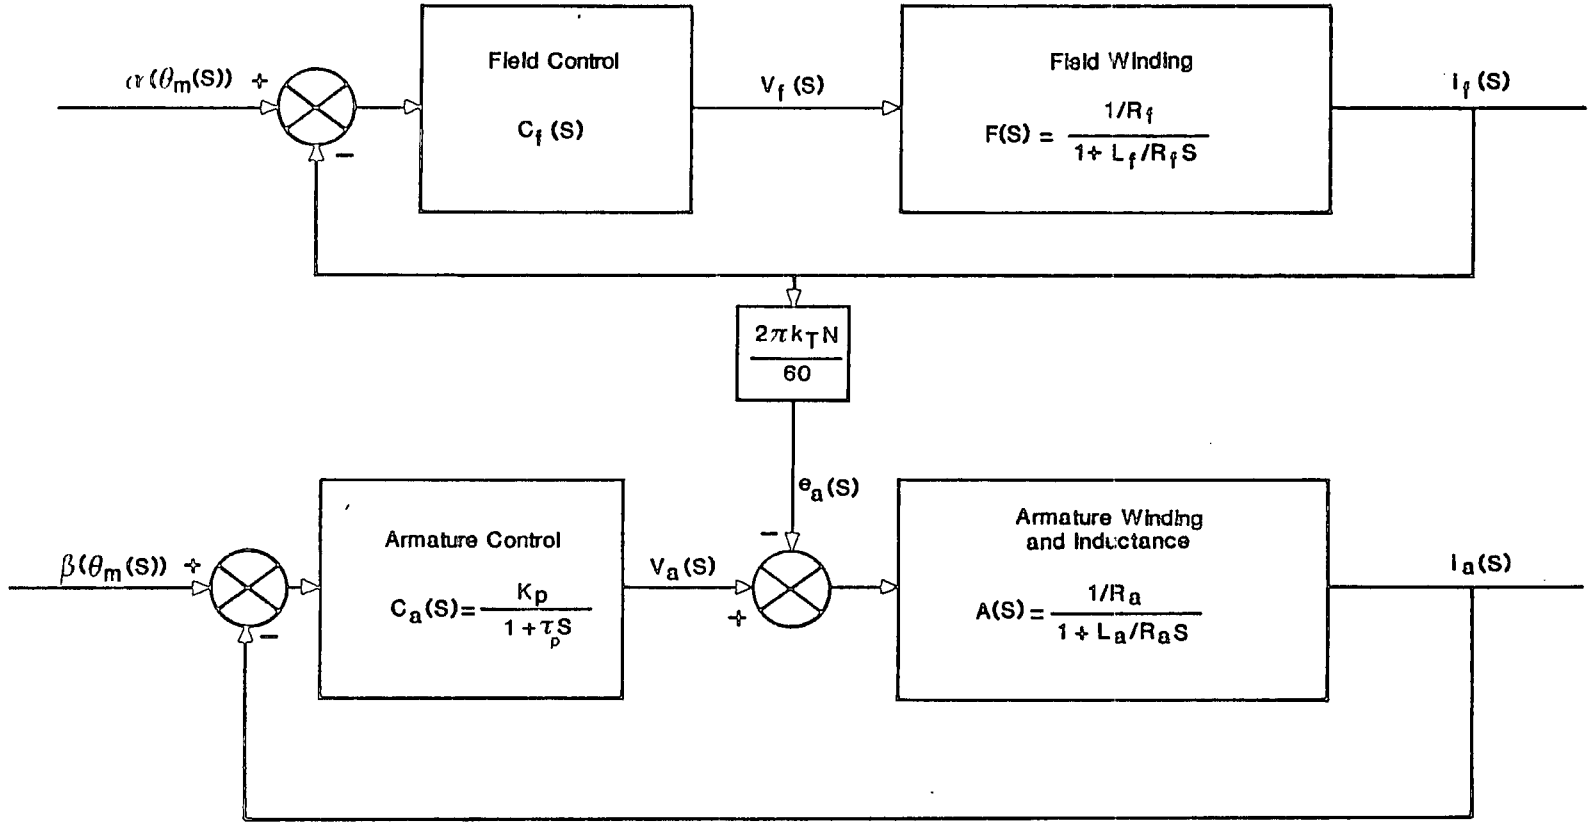

Fig. 3.4 Block Diagram for Motor Current Control

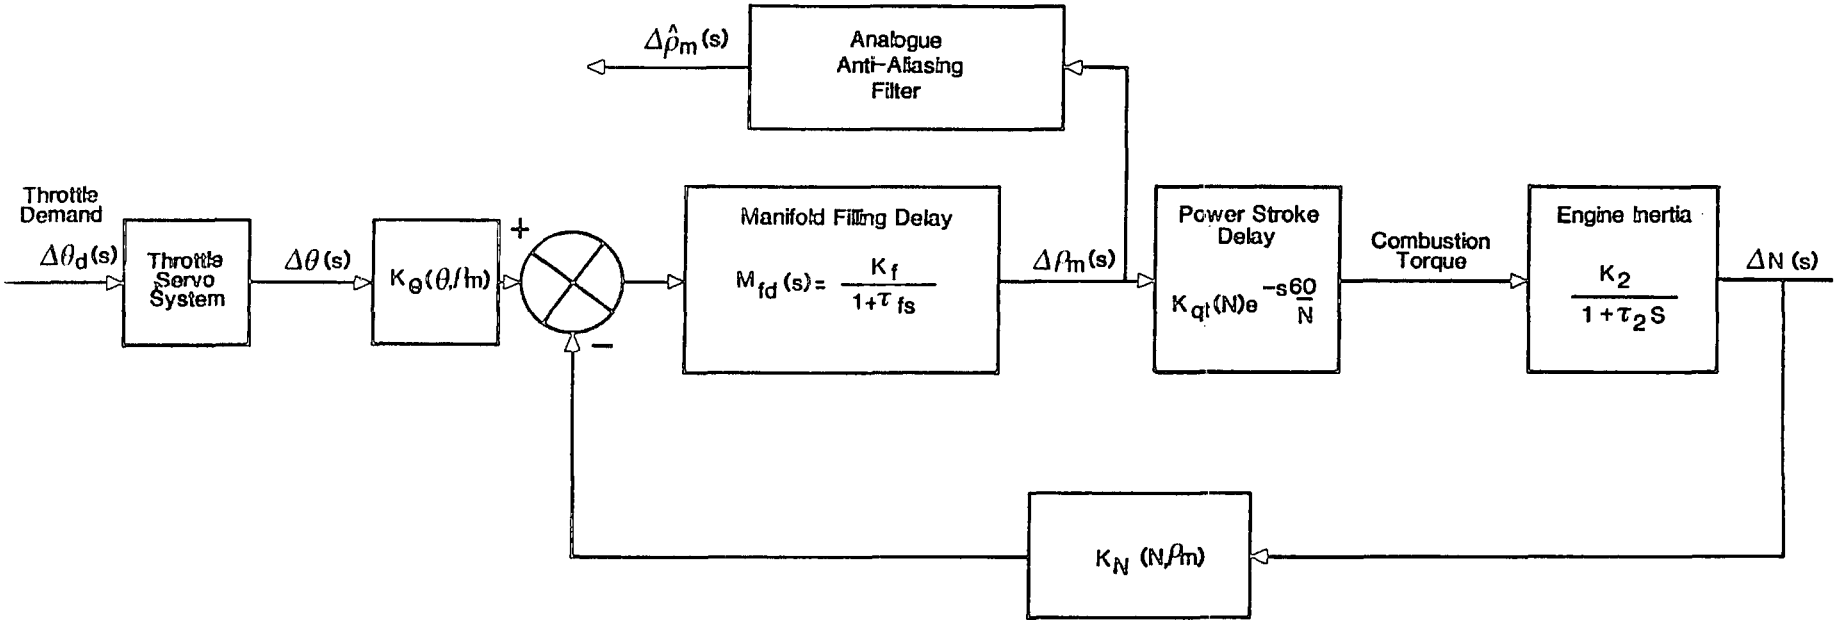

Fig. 3.5 Engine Physical Processes

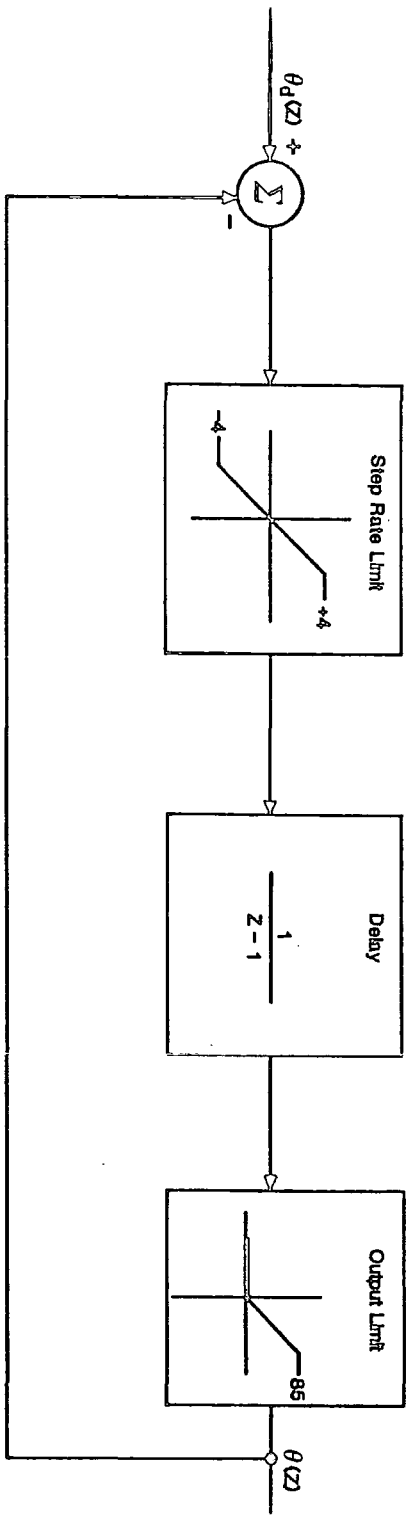

Fig. 3.6 Block Diagram of the Throttle Servo-System

CHAPTER 4

EXPERIMENTAL IDENTIFICATION RESULTS

In the previous chapter measurement techniques were proposed which allow the torque output of the engine and motor to be determined from other related variables rather than by direct measurement. Two experimental investigations are required before these methods can be applied directly to the control of torque.

Firstly the indirect measurement methods must be calibrated under steady state conditions, for the particular engine and motor in use on the rig. In the first part of this chapter the results of these calibration experiments are presented and then the resulting torque relationships are tested under normal operating conditions. Once these torque signals are available and proven to be accurate, then the second stage is to determine the dynamic torque response of the engine and motor to input signals from the two accelerator systems. Analysis in chapter 3 has already shown what form of transfer functions ought to relate indirect torque measurements to throttle setting, in the case of the engine, and DAC setting in the case of the motor. For the engine there is only one dynamic element governing this relationship known as the manifold filling delay. It was shown that this element, describing the flow of air through the inlet manifold, may be modelled by a simple first order digital transfer function. All that remains to obtain a complete picture of the engine is to identify the numerical values of the transfer function coefficients at a range of different operating speeds.

The situation with the motor is similar but now three dynamic elements, describing the behaviour of field and armature current, remain to be quantified. In this instance however further steady state calibration must precede the identification of the dynamic elements. This first stage measures the accelerator gain functions as defined in chapter 3 thus removing the

non-linear effect of the power electronics controller. The experiments which allowed these gain functions to be determined follow on from the calibration and testing of the two torque models.

Once the gain function problem is solved the way is clear to identify all the necessary transfer functions. This is achieved using a model reference adaptive identification algorithm, thus completing the picture of the torque control environments relating to both engine and motor. To confirm the accuracy of the resulting transfer functions each one is tested against a completely independent set of data. These independent tests confirm that the engine and motor both behave consistently and so subsequent controller design can be tackled with confidence.

4.1 Calibration of Indirect Torque Measurements

4.1.1 Engine Torque Model Calibration

Calibrating the engine torque model involves measuring the relationship between speed, inlet manifold depression and torque output. In the steady state engine combustion power can be measured on the engine torque transducer and when plotted against inlet manifold depression at different speeds gives the variation shown in figure 4.1. These variations suggest that engine power can be expressed as:

$$P_{ic} = f_1(N) + f_2(N)\hat{p}_m \tag{4.1}$$

At any speed $f_1(N)$ and $f_2(N)$ can be quantified by fitting a least squares regression line to the data of figure 4.1. The values of $f_1(N)$ and $f_2(N)$ are themselves expressed as a function of speed in figure 4.2. This figure shows that $f_1(N)$ and $f_2(N)$ can in turn be computed at any speed by means of a regression analysis to give:

$$f_1(N) = -0.22112169 + 5.587 \times 10^{-3}N \tag{4.2}$$

$$f_2(N) = -6.152 \times 10^{-4} - 8.843 \times 10^{-6}N \tag{4.3}$$

which when substituted into equation 4.1 and converted to combustion torque gives:

$$T_{ic} = \frac{-2112.457}{N} + 53.359 - \left(\frac{5.875}{N} + 0.084472\right)\hat{p}_m \quad \text{Nm} \tag{4.4}$$

Normally the rise time associated with torque and inlet manifold depression is significantly faster than that associated with engine speed, because the large vehicle inertia restricts the rate at which engine speed can change. Assuming the engine speed to be constant, then changes in engine torque can be related to manifold pressure changes by linearising equation (4.4) about the initial operating speed to give

$$\frac{\Delta T_{ic}}{\Delta \hat{p}_m} = -\left(\frac{5.875}{N_0} + 0.084472\right) = K_{qt}(N_0) \tag{4.5}$$

where K_{qt} is the gain of the power stroke delay shown in the block diagram of figure 3.5.

4.1.2 Engine Torque Model Testing

Although the global engine torque model described by equation (4.4) has been calibrated over a large engine operating range it is important to carry out independent verification. Only then can the effect of inaccuracies in the linear regression analysis be properly assessed. To achieve this the engine was operated on load over a wide speed range. Throttle changes were made manually during the experiments. All adjustments were made slowly to ensure pseudo steady state conditions and so allow the model and

torque transducer readings to be compared directly without concern for their differing dynamic responses. In figure 4.3 comparison is made between the two readings by plotting the ratio transducer torque over model torque against speed. The resulting graph shows the indirect model giving acceptable results at high speeds, whilst at speeds below 1200 r.p.m. the model begins to break down and predicts torques higher than those measured by the transducer. In an attempt to compensate for this defficiency in the model the scaling factor $\frac{N}{1200}$ is applied whenever the engine is operating below 1200 r.p.m.

In the hybrid arrangement engine torque control at such low speeds is unusual, since low speed operation is mostly all-electric. If greater accuracy is required then a possible solution would be to map torque as a function of manifold pressure and speed into a two dimensional array or look-up table. One disadvantage of this method is the greater computational expense involved in interpolating between values.

A separate experiment plotting the prediction of the global torque model against time reveals an interesting effect. As shown by figure 4.4., the model gives a torque prediction that is approximately 3 Nm higher than the torque transducer measurement whenever the radiator cooling fan is not operating. During engine characterisation the cooling fan was always on and consuming some engine torque via the alternator. When the fan is off this torque is available at the output shaft and is registered by the torque transducer. As a consequence of this the torque prediction algorithm shows greatest accuracy in figure 4.4 during the indicated 10 second period when the radiator cooling fan was on.

4.1.3 Motor Torque Model Calibration and Testing

During the constant speed experiments used to determine the motor operating modes, as shown in figure 3.3, the M68000 was recording all data sources associated with the motor. After the experiments the torque

transducer and current readings were processed to show how the the motor torque constant, K_T, varies with field current by applying equation (3.10). As figure 4.5 shows these results demonstrate a slight degree of saturation at high field currents the effect of which is quantified by the regression line which is fitted to the data.

$$K_T(i_f) = 0.109 - 2.52 \times 10^{-3} i_f \tag{4.6}$$

and hence

$$T_{em} = (0.109 - 2.52 \times 10^{-3} i_f) i_f i_a \qquad \text{Nm} \tag{4.7}$$

Further investigation of these results showed no variation of the torque constant with armature current. A comparison of predicted torque using equation (4.7) with that measured from the torque transducer is shown in figure 4.6 for both motoring and regenerative braking operation. Agreement is at all times within $\pm$ 3 Nm.

4.2 Gain Functions for the Three Operating Modes of the Motor

Open loop gain for the field and armature current control loops is given by $C_f(s)F(s)$ and $C_a(s)A(s)$ respectively. Assuming that in both cases the gain is substantially greater than one, then the corresponding closed loop transfer functions $G_a(s)$ and $G_f(s)$ will have unity gain as can be seen from equations (3.12) and (3.15). Once this simplifying assumption is made the steady state expression for field and armature current can be predicted on the basis of equations (3.13) and (3.18) respectively

$$i_f = \alpha(\theta_m) \tag{4.8}$$

$$i_a = \beta(\theta_m) - N\alpha^*(\theta_m) \tag{4.9}$$

and in linearised form at constant speed N_0

$$\Delta i_f = \Delta\alpha(\theta_m) \tag{4.10}$$

$$\Delta i_a = \Delta\beta(\theta_m) - N_0\Delta\alpha^*(\theta_m) \tag{4.11}$$

Steady state experiments in each of the three operating modes should therefore reveal the form of the gain functions.

4.2.1 Field Boost Mode

In the field boost mode the experimental results of figure 4.7(a) indicate that the set-point function $\alpha(\theta_m)$ may be accurately represented by a second order polynomial.

$$\alpha(\theta_m) = a + b\theta_m + c\theta_m^2 \tag{4.13}$$

By defining the mean value of θ_m as:

$$\bar{\theta}_m = \theta_{m(0)} + \frac{\Delta\theta_m}{2} \tag{4.14}$$

and expanding equation (4.13) as a second order Taylor series gives

$$\Delta\alpha(\theta_m) = (b + 2c\bar{\theta}_m)\Delta\theta_m \tag{4.15}$$

A least squares analysis of the experimental data shown in figure 4.7(a) gives the coefficients for equation (4.15) as $a = 1.97, b = 9.683 \times 10^{-3}$ and $c = 7.472 \times 10^{-5}$.

From the definition $\alpha^*(\theta_m) = \frac{K_T}{K_p} \times \alpha(\theta_m)$ it follows that $\alpha^*(\theta_m)$ should have the same form as equation (4.13). As the second order term in $\alpha(\theta_m)$ is small however, reduced further by the gain modifier, $\frac{K_T}{K_p}$, a linear variation for $\alpha^*(\theta_m)$ is adequate i.e.

$$\alpha^*(\theta_m) = d + e\theta_m \tag{4.16}$$

Experimental evidence similarly suggests that a linear variation for the armature set-point function $\beta(\theta_m)$ can be used i.e.

$$\beta(\theta_m) = f + g\theta_m \tag{4.17}$$

Experimental results expressing armature current as a function of motor speed for constant θ_m are shown in figure 4.7(b) and, as suggested by equation (4.9), are straight lines of intercept $\beta(\theta_m)$ and slope $\alpha^*(\theta_m)$. Regressing the gradients and intercept values gives the constants d, e, f and g in equations (4.16) and (4.17) as

$$\alpha^*(\theta_m) = 4.914 \times 10^{-3} - 6.6 \times 10^{-5}\theta_m \tag{4.18}$$

and

$$\beta(\theta_m) = 26.618 + 0.2104\theta_m \tag{4.19}$$

which for small changes implies

$$\Delta\alpha^*(\theta_m) = -6.6 \times 10^{-5}\Delta\theta_m \tag{4.20}$$

$$\Delta\beta(\theta_m) = 0.2104\Delta\theta_m \tag{4.21}$$

These results are summarised in Table 4.1.

4.2.2 Full Field Mode

In the full field mode field current is constant which means that $\Delta\alpha(\theta_m(s)) = 0$. All control is carried out on the armature which, in the absence of $\Delta\alpha^*(\theta_m(s))$, must vary according to

$$\Delta i_a = \Delta\beta(\theta_m) \tag{4.22}$$

Slow ramping of the demand input, θ_m, at a number of constant speeds, figure 4.8, shows the field current constant at 8.5 A and reveals that in this mode the set point function, $\beta(\theta_m)$, depends only on the initial armature current and the demand value θ_m as

$$\beta(\theta_m) = I_a(N) + 0.137\theta_m \tag{4.23}$$

Where $I_a(N)$ is the initial armature current at motor speed N, giving the linearised results

$$\Delta\beta(\theta_m) = 0.137\Delta\theta_m \tag{4.24}$$

$$i_{f(0)} = 8.5\text{A} \tag{4.25}$$

$$\Delta\alpha^*(\theta_m) = \Delta\alpha(\theta_m) = 0 \tag{4.26}$$

As a consequence of these results for the set point gain functions, and the simplification of the motor block diagram in the armature control mode described in section 3.1.1, this mode is the simplest of the three to analyse, since it has constant gain and dynamics regardless of the initial conditions.

The result for this set point function is summarised in Table 4.1.

4.2.3 Field Weakening Mode

When operating in the field weakening mode, experiments show that the field current reduces with both speed and accelerator demand. The field and armature currents are again given by equations (4.8) and (4.9) respectively. The set-point functions $\alpha(\theta_m)$ and $\beta(\theta_m)$ can be evaluated by slowly ramping θ_m at a number of constant speeds and recording the variation in both field current and armature current.

For the field current the results of figure 4.9(a) suggest that $\alpha(\theta_m)$ has the form

$$\alpha(\theta_m) = a(N) + b(N)\theta_m \tag{4.27}$$

Regressing the slopes and intercepts gives

$$\alpha(\theta_m) = (19.957 - 0.004611N) + (1.605 \times 10^{-6}N - 6.134 \times 10^{-3})\theta_m \tag{4.28}$$

With speed constant and taking small variations gives

$$\Delta\alpha(\theta_m) = (1.605 \times 10^{-6}N_0 - 6.134 \times 10^{-3})\Delta\theta_m \tag{4.29}$$

Armature current is given by equation (4.8) but unlike the field boost mode, $\alpha^*(\theta_m)$ has negligible effect on i_a since typical values of $\alpha(\theta_m)$ are ten times smaller. The experimental results of figure 4.9(b) show that the armature current set-point function, $\beta(\theta_m)$, behaves in a similar manner to that in the full field mode, equation (4.23), but with a slightly different gain, namely

$$\beta(\theta_m) = I_a(N) + 0.147\theta_m \tag{4.30}$$

The important relationships for the field weakening mode are summarised in Table 4.1.

	Field Boost	Full Field	Field Weakening
$\alpha(\theta_m)$	$1.97+9.683\times10^{-3}\theta_m$ $+7.472\times10^{-5}\theta_m^2$	–	$19.957-4.611\times10^{-3}N$ $+(1.605\times10^{-6}N-6.134\times10^{-3})\theta_m$
$\Delta\alpha(\theta_m)$	$(9.683\times10^{-3}+1.494$ $\times10^{-4}\bar{\theta}_m)\Delta\theta_m$	0	$(1.605\times10^{-6}N_0-6.134\times10^{-3})\Delta\theta_m$
$\beta(\theta_m)$	$26.618+0.2104\theta_m$	$I_a(N_0)+0.137\theta_m$	$I_a(N_0)+0.147\theta_m$
$\Delta\beta(\theta_m)$	$0.2104\Delta\theta_m$	$0.137\Delta\theta_m$	$0.147\Delta\theta_m$
$\alpha^*(\theta_m)$	$4.914\times10^{-3}-6.600\times10^{-5}\theta_m$	–	–
$\Delta\alpha^*(\theta_m)$	$-6.600\times10^{-5}\Delta\theta_m$	0	–

Table 4.1 Set Point Functions

4.3 Transfer Function Identification

Both the analysis of the engine and motor began in the s-plane since it more readily relates to the continuous time processes involved. As a final stage discrete equivalents of the resultant transfer functions were produced since ultimately a digital control system was envisaged. Before either form of transfer function can be of practical use in controller design however the numerical coefficient values must be determined from experimental data.

Despite the discrete nature of the final control system, identification can be carried out in two ways. Either the s-plane model can be evaluated with subsequent discretisation via the appropriate hold algorithm or alternatively a digital model maybe identified directly. Eventually a particular version of the latter method was successfully applied in all cases on the rig, but in one instance the alternative s-plane method was tried for comparison. In comparing various methods of identification Whitfield [Whitfield, 1986] mentions the linear least squares technique due to Levy [Levy, 1959], which involves fitting a curve to frequency response data in the classical Nyquist form of gain and phase shift. This method was applied to the transfer function relating throttle position to engine speed on no load, by applying a sinusoidal input to the throttle via the stepper motor and measuring the corresponding speed variations. A comparison of the results from this method and those from the direct digital identification is given in section 4.6.

It was felt however that success with this method was likely to be fairly limited, both because of experimental difficulties in applying the requisite sinusoidal inputs (particularly in the faster acting torque transfer functions) and since as Whitfield points out [Whitfield, 1986] the method produces biased parameter methods and locally poor fits to the experimental data. Consequently the method was abandoned in favour of a method which is both unbiassed and leads directly to a digital model, which in the present

context seems a more natural choice. The identification technique used is Landau's Model Reference Adaptive Identification algorithm [Landau, 1976]. Many methods are available for off line identification of digital models and some are discussed by Warwick [Warwick, 1986]. He notes that Landau's method is strongly immune to disturbances affecting the output, an important advantage in a system liable to noise. The principle of Landau's identification procedure is to apply a common input to the actual system and an adjustable model, as illustrated schematically by figure 4.10. The output of the system and the model is then compared to produce an adaptive error, v_j, which is processed by an adjustment mechanism to provide an update to the model, β_j. Applying the technique assumes that the plant can be described by the general discrete transfer function:

$$\frac{y(z)}{u(z)} = \frac{b_0 + b_1 z^{-1} + \ \ldots \ + b_m z^{-m}}{1 - a_1 z^{-1} - \ \ldots \ - a_n z^{-n}} \tag{4.31}$$

The output of the model at the j^{th} sampling interval is calculated from a difference equation, which can be expressed in vector form as:

$$y_j = \beta^T \phi_j \tag{4.32}$$

Where β^T is the plant model coefficient vector:

$$\beta^T = [\ a_1 \quad a_2 \quad \ldots \quad a_n \quad b_0 \quad b_1 \quad \ldots \quad b_m\ ] \tag{4.33}$$

Similarly, ϕ_j is the model input/output vector:

$$\phi_j^T = [\ y_{j-1} \ \ldots \ y_{j-n} \ u_j \ \ldots \ u_{j-m}\ ] \tag{4.34}$$

The identification algorithm seeks to find approximate values of a *limited number* of terms from the plant model coefficient vector. In the case of a first order model the vector $\beta^T = [\ a_1 \ b_0 \ b_1\ ]$ will be identified. The

algorithm works with a set of experimental inputs, u, and plant outputs, y_p. At each step it applies the same input u_j to a reference model and calculates the model output given by:

$$y_j = \beta_{j-1}^T \phi_j \tag{4.35}$$

Where β_{j-1}^T is the best estimate of the plant model parameters from the previous $j-1^{th}$ step. An adjustment mechanism is now applied to calculate an updated model parameter set β_j

$$\beta_j = \beta_{j-1} + \frac{P_{j-1}\phi_j v_j}{1 + \phi_j^T P_{j-1}\phi_j} \tag{4.36}$$

Two new terms appear in equation (4.36), the adaptive gain matrix P_{j-1} and the error term v_j. Matrix P is square with dimension m+n+1 and, at the beginning of the identification it may be defined arbitrarily but must be non zero. In the present work diagonal elements of P are initialised to 0.1 with other elements left at zero. As the identification proceeds P is updated, after the update of β, according to the equation

$$P_j = P_{j-1} + \frac{P_{j-1}\phi_j \phi_j^T P_{j-1}}{1 + \phi_j^T P_{j-1}\phi_j} \tag{4.37}$$

The error term, v_j, is a measure of the difference between the plant output and the model output and can be considered as an adaptive error. Calculation of v_j, by equation (4.38), is the final stage for each data point.

$$v_j = y_{p(j)} - y_j + \sum_{i=1}^{n} g_i e_{j-i} \tag{4.38}$$

On the right hand side of this expression the difference between the plant and model output appears as well as the weighted sum of the n previous *a posteriori* errors, $e_{(j-i)}$, *i=1 to n* . The model output, y_j, is calculated from

the previous parameter set β_{j-1}^T so that the difference, $(y_{p(j)} - y_j)$, constitutes the *a priori error* .

Once the parameter set, β_j^T, has been updated the new *a posteriori* error is calculated from

$$e_j = y_{p(j)} - y_j^0 \tag{4.39}$$

With y_j^0 being the *a posteriori* model output:

$$y_j^0 = \beta_j^T \phi_j \tag{4.40}$$

At each step the previous n *a posteriori* errors are used and each is multiplied by a fixed gain g_i before summation as in equation (4.38). Values of g_i are important to the convergence of the whole algorithm as explained by Landau [Landau, 1976]. In practice it is sufficient to set $g_i = -a_{i(0)}$; where $a_{i(0)}$ is a reasonable initial estimate of the corresponding plant denominator coefficient. Suitable initial estimates can be obtained from the s-domain transfer function of the plant. In the case of the engine and the motor, an appropriate form for this transfer function can be obtained from an analysis of the physics of the plant, with an approximate evaluation of parameters being obtained from the open loop step response. After suitable discretisation, an initial estimate of the discrete transfer function is obtained.

A slight adjustment to experimental data is needed when identifying a transfer function with reduced order numerator as was found to be appropriate for the motor (equation 3.25). In this case it is required that the coefficient b_0 be zero. To achieve this it is not sufficient to set $m = n - 1$ since the identification algorithm then produces a transfer function with b_n equal to zero rather than b_0. If in addition however the experimental data set is shifted one step back in time prior to the identification the problem is solved,

because then the coefficient given by the algorithm as b_0 is in fact b_1. Finally it should be noted that for identification purposes both input and output data have their offset at zero time removed.

4.4 Transfer Function Identification for the Motor

4.4.1 The Closed Loop Transfer Functions for Current

In section 3.1 a complete model relating motor torque to accelerator demand was developed. To complete this model the transfer functions $G_a(z)$ and $G_f(z)$ are needed. To carry out the necessary identifications, short accelerator transients are applied to the system and the resulting variations in field and armature current are recorded. Particular behaviour in different operating modes makes them more or less suitable for identifying a given transfer function. In the field boost mode the large variations of field current provide ideal conditions for identifying $G_f(z)$, while a full field mode experiment, conducted at constant speed, ensures that the motor back e.m.f. is constant, an important condition for identifying $G_a(z)$. Once the experimental input/output data has been gathered the identification algorithm can be applied. Although the algorithm can be applied indefinitely there comes a point when no further worthwhile improvement can be made to the fit between model and experimental data. This point must be judged arbitrarily by the operator but to aid in this decision the identification program, written for the M68000, calculates the mean difference between experimental and model output over each pass through the data. Figure 4.11 shows how well the identified transfer function $G_a(z)$ fits the experimental data.

It is not easy to identify the time constant in the armature torque current controller directly, however once $G_a(z)$ and $G_f(z)$ are known an indirect method is possible. Initially these pulse transfer functions are transformed to their s-plane equivalents, $G_a(s)$ and $G_f(s)$, using the inverse bilinear transform.

Then by comparing the actual response of i_a in the field boost mode with the prediction of equation (3.19), τ_p can be adjusted until a similar response is obtained. The identified transfer functions and time constant, τ_p, are tabulated in Table 4.2.

Transfer Function	
$G_a(z)$	$\frac{0.1304z+0.08034}{z^2-0.8300z+0.04187}$
$G_a(s)$	$\frac{0.02676(s+420.75)(s-100)}{s^2-102.371s+1131.866}$
$G_f(z)$	$\frac{0.01536z+0.1008}{z^2-1.183z+0.2973}$
$G_f(s)$	$\frac{0.0344(s-135.97)(s-100)}{s^2+56.657s+460.63}$
τ_p	0.085

Table 4.2 Transfer Functions Identification

4.4.2 Direct Identification of the Motor Torque Transfer Function

With the identification of the three dynamic elements in Table 4.2, the current based model relating torque to accelerator setting is complete and could be used as a basis for controller design. There is however a gain in accuracy if a model is identified directly between accelerator setting and torque. Such a model could have been identified at the outset with no knowledge of the intervening system, in this case the motor and power electronics combination would be treated as merely a 'black box'. A serious drawback of such an approach is that the resultant transfer function would have arbitrary order, possibly resulting in important dynamic properties being overlooked. With the current model available however, a direct identification can be constructed from a logical foundation. As an example the following paragraphs describe how the direct identification of the motor torque transfer function for the field boost mode follows from the current model.

When the dynamic elements, identified in section 4.4.1, are transformed into the s-plane and substituted into equation (3.21), they produce a fourth

order torque response transfer function for the field boost mode. Experimental evidence and analysis of this fourth order model however clearly shows that the system is predominantly second order. For example, a particular field boost transient at 1500 r.p.m. had the initial data tabulated in the left hand column of Table 4.3.

	Field Boost	Full Field	Field Weakening
Speed N_0 r.p.m.	1497	1500	3030
$\bar{\theta}_m$	165	500	550
$i_{a(0)}$ (Amps)	47.07	81.81	64.82
$i_{f(0)}$ (Amps)	4.10	8.50	5.63
$\bar{i}_f$ (Amps)	5.76	8.50	5.57

Table 4.3 Initial Conditions for a Transient in Each Operating Mode

Combining this data with the mode dependent gains defined in Table 4.1 and adding all the s-plane dynamic elements defined in Table 4.2 provides the complete fourth order torque model with poles and zeros as shown in Table 4.4.

Poles	Zeros
-9.389	-11.637
-12.610	-52.316
-46.818	-1528.622
-89.761	369.23
	100.00

Table 4.4 Poles and Zeros for the Field Boost Transfer Function Based on Current Relationships

Several of these poles and zeros do not have a significant effect on the

transient performance of the model. For example, after a suitable modification to the system gain, the pole at -89.76 and the zeros at -1528, 100 and 369 may be dismissed since they are, at least, an order of magnitude larger than the dominant poles and zeros. Of the remainder, the pole-zero pair at around $s = -12$ cancel and so contribute little to the system which leaves the second order result:

$$\frac{\Delta T_{em}(s)}{\Delta \theta_m(s)} = \frac{1.7398(s + 52.316)}{(s + 9.839)(s + 46.818)} \tag{4.41}$$

Adding the necessary zero-order hold and discretising the result using equation (3.24) gives

$$\frac{\Delta T_{em}(z)}{\Delta \theta_m(z)} = \frac{0.03292z - 0.01147}{z^2 - 1.2134z + 0.3220} \tag{4.42}$$

This equation serves as a starting point for the adaptive identification procedure to identify a second order transfer function which more closely matches the experimental data. For comparison a fourth order identification was also carried out but no increase in accuracy was achieved. The results of the different identification procedures are tabulated in Table 4.5.

	$\frac{\Delta T_{em}(z)}{\Delta \theta_m(z)}$
FromCurrentModel (ReducedOrder)	$\frac{0.03292z-0.01147}{z^2-1.2134z+0.3220}$
DirectIdentification (4thOrder)	$\frac{0.01168z^3-0.009281z^2+}{z^4-1.282z^3+0.4149z^2}$ $\frac{0.001506z+0.001028}{0.00537z-0.01728}$
DirectIdentification ReducedOrder	$\frac{0.01171z-0.008569}{z^2-1.3184z+0.422}$

Table 4.5 Comparison of Torque Transfer Functions For the Field Boost Mode

These results not only vindicate the use of the reduced order model

but also confirm the accuracy of the transfer function derived directly from the current model. In figure 4.12 a comparison of the experimental torque response and the simulated torque response predicted from the transfer functions listed in Table 4.5 is made. All three have similar gains and dynamics.

For any given set of initial conditions a similar process is possible for the field weakening and full field modes, conveniently carried out in five stages:

Stage 1. Transform $G_a(z)$ and $G_f(z)$ into the equivalent s-domain transfer functions using some appropriate inverse mapping procedure e.g. the inverse bilinear transform.

Stage 2. Substitute the mode dependent gains and dynamics from Tables 4.1 and 4.2 into equation (3.21) namely $G_a(s), G_f(s), (1+\tau_p s), \Delta\beta(\theta_m), \Delta\alpha(\theta_m), \Delta\alpha^*(\theta_m)$.

Stage 3. Determine the parameters specific to the individual transient $K_T(\bar{i}_f)$, i_{f0}, i_{a0}, $\bar{\theta}_m$ and N_0.

Stage 4. Examine the full current model and if possible obtain a reduced order equivalent.

Stage 5. Combine the resulting transfer function with a zero-order hold and discretise using equation (3.24). Carry out a direct identification between $\Delta\theta_m$ and ΔT_{em} using the reduced order model as a starting point.

Results from the final stage of this process are set out for all three modes in Table 4.6.

	Field Boost	Full Field	Field Weakening
Reduced Order Current Model	$\frac{0.0329z-0.0115}{z^2-1.213z+0.322}$	$\frac{0.0125z-0.00637}{z^2-0.943z+0.129}$	$\frac{0.00866z-0.00442}{z^2-0.943z+0.129}$
Direct Identification	$\frac{0.0117z-0.00857}{z^2-1.318z+0.422}$	$\frac{0.0152z-0.00310}{z^2-1.308z+0.427}$	$\frac{0.000359z-0.00543}{z^2-1.276z+0.355}$

Table 4.6 Identified Torque Transfer Functions for All Three Operating Modes

4.5 Manifold Filling Delay and Engine Torque Variations

Identification of the transfer functions relating inlet manifold depression to throttle position, equation (3.40), using the model reference equations of section 4.3 were carried out at engine operating speeds of 1000, 1500, 2000, 2500 and 3000 r.p.m. In each case data for identification was obtained by measuring manifold pressure changes as the throttle was moved about an average setting for that speed. Numerical results relating to equation (3.40) for all the identification experiments are given in Table 4.7. Included in this table are the steady state gains of each transfer function derived from the final value theorem

$$K_f k_\theta(\theta, p_m) = \lim_{z \to 1}(z-1)H(z)G(z) \tag{4.43}$$

where $H(z)$ is the input to the plant $G(z)$. If $H(z)$ is the unit step $= z/(z-1)$; then the gain is

$$K_f k_\theta(\theta, p_m) = \lim_{z \to 1} M_{fd}(z) = \frac{b_0 + b_1}{1 - a_1} \tag{4.44}$$

Speed r.p.m.	a_1	b_0	b_1	System Gain
1000	0.766	0.184	-10.26	-40.45
1500	0.648	0.177	-10.964	-31.65
2000	0.640	0.158	-11.360	-31.11
2500	0.644	-0.185	-9.985	-28.31
3000	0.536	-0.608	-11.14	-25.31

Table 4.7 Identified Transfer Function Coefficients and Gain for the Engine Manifold Filling Delay

Results obtained under steady state conditions are compared with

these transfer function gains in figure 4.13. Lines representing these gains have been added to the steady state data at speeds of 1000, 2000, and 3000 r.p.m. Apart from the consistent slopes, the form of the steady state results confirms the theoretical prediction of equation (3.31) for the steady state case i.e.

$$p_m = p_{m(0)}(N) - K_f k_\theta(\theta, p_m)\theta \tag{4.45}$$

Before being used by the torque prediction algorithm, equation (4.4), an additional, digital, filter is included in the manifold pressure signal to remove the transitory effects of inlet valves opening, particularly at low speeds. Acceptable results are obtained by using a first order filter of the form:

$$\frac{\Delta\hat{p}'_m(z)}{\Delta\hat{p}_m(z)} = \frac{b_0^f}{1 - a_1^f z^{-1}} \tag{4.46}$$

Combining the results of equations (3.40),(3.41),(4.5) and (4.46) allows a complete linear model to be produced which relates variations in engine torque to changes in throttle demand as shown in the block diagram of figure 4.14. For simulation purposes this model can conveniently be expressed in state space notation by defining the state space vector $X = [\ X_1 \quad X_2 \quad X_3\ ]$ so that:

$$\begin{aligned} X_{k+1} &= AX_k + BU_k \\ Y_k &= CX_k \end{aligned} \tag{4.47}$$

where:

$$A = \begin{bmatrix} 0 & 0 & 0 \\ b_1 + a_1 b_0 & a_1 & 0 \\ a_1^f b_0^f b_0 & a_1^f b_0^f & a_1^s \end{bmatrix} \tag{4.48}$$

$$B^T = [\ 1 \quad 0 \quad 0\ ] \tag{4.49}$$

$$U_k = \Delta\theta_d \tag{4.50}$$

$$C = \begin{bmatrix} 1 & 0 & 0 \\ b_0 & 1 & 0 \\ b_0^f b_0 & b_0^f & 1 \end{bmatrix} \tag{4.51}$$

$$Y^T = [\quad \Delta\theta \quad \Delta\hat{p}_m \quad \Delta\hat{p}'_m \quad] \tag{4.52}$$

with changes in engine torque, ΔT_{ic}, being obtained from equation (4.5).

To check the accuracy of the engine model a common throttle signal was applied to the model and the real engine, for a number of different operating speeds. In each case the test signal demanded that the throttle be opened and closed twice, first moving by 5 steps and then by 3, finally returning to the average setting for that speed.

At the low engine speed of 1000 r.p.m. (figure 4.15) the engine does not operate very smoothly causing a noisy manifold pressure signal, and slightly non-linear behaviour, both of which reduce model accuracy. This uneven performance at low speed is the main reason for the additional digital filtering of the manifold pressure signal.

At 3000 r.p.m. engine performance is greatly improved as shown by figure 4.16. At this speed the manifold filling effect is both more linear and less noisy in its response than at 1000 r.p.m., allowing greater accuracy from the simple low order model used.

4.6 Engine No-Load Speed Transfer Function

Applying a simple step input to the engine, allowed the transfer function given by equation (3.43) to be quantified by the identification equations:

$$\frac{\Delta N(z)}{\Delta\theta(z)} = \frac{0.838 - 1.510z^{-1} + 1.922z^{-2}}{1 - 1.790z^{-1} + 0.795z^{-2}} \tag{4.53}$$

To verify the model a signal comprising two step changes in throttle over 3.5 seconds was applied. In figure 4.17 the resulting variations in engine speed are compared with not only the digital model given above, but also with a third order transfer function identified by Levy's method:

$$\frac{\Delta N(s)}{\Delta \theta(s)} = \frac{-0.0236s^3 + 0.2522s^2 - 1.3479s + 14.8812}{0.0182s^3 + 0.0188s^2 + 1.0747s + 1} \tag{4.54}$$

As can be seen the digital model achieves greater accuracy and does not exhibit the non-minimum phase behaviour displayed by the Levy model. It should be mentioned that the simulation for the Levy model did not attempt to model the lag caused by the throttle, however figure 4.17 suggests that its inclusion would have further reduced the predicted variation in speed.

4.7 Discussion

Previous sections in this chapter have demonstrated that there are two ways of identifying transfer functions relating motor torque to accelerator input.

The first, and most rigorous, involves identifying the behaviour of the current control loops within the power electronics. When considered in isolation the response of armature and field current to a given transient input is highly dependent on the initial operating conditions. This non-linear effect is cancelled by investigating set point functions which describe how current varies with demand under steady state conditions, over the whole operating range of the motor. Once these set point functions are obtained it is relatively easy to calculate the change in current which will result from a given change in demand, $\Delta\theta_m$, and a given set of initial conditions. Such changes in current do not appear instantaneously but are governed by the dynamics of the closed loop current control systems $G_a(z)$ and $G_f(z)$. The real advantage of using the set point functions is that these dynamic responses are now seen

to be independent of operating point.

Coupling both dynamic elements and set point functions together through equation (3.21) completely defines the torque response of the motor for any initial condition. Before considering the simplified version of this model, the accuracy of the global description of the motor will be demonstrated for three example transients. Table 4.3 indicates the three sets of initial conditions, one for each of the operating modes. In figures 4.18-4.20 results for both torque and current, calculated from equations (3.14), (3.19) and (3.21), are compared with experimental data.

By using the complete model as a basis a second method can produce a more compact form of torque transfer function. In this case a second order model was fitted directly between θ_m and torque, T_{em}, for a specific set of data obtained in each mode. If controller design could be based on such transfer functions then then the task would be considerably simplified, however they are strictly speaking only valid for one operating point. Before extending their use to controller design further tests are needed to see how their accuracy deteriorates at the extremes of a mode. As illustrated by figure 4.21, the direct model for the field boost mode is extremely accurate at original conditions, however at an entirely different point figure 4.21 also illustrates that this model is still capable of providing acceptable results. To exploit this fact direct models were also obtained for the full field and field weakening modes. Figure 4.22 demonstrates that the full field model is accurate at any point within its mode as expected from the theory. Furthermore the field weakening model is still adequate away from its original identification condition, as shown by figure 4.23. This has important implications for controller design as it removes the requirement to produce a controller design which adapts to widely varying gains and thus remains well tuned at all times.

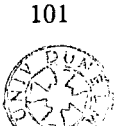

Modelling of the engine has proven relatively straightforward with only one first order dynamic element governing the torque response loop. When designing a controller it must be remembered that that this element is speed dependant so that design performance may well deteriorate away from design speed.

In both systems the indirect torque measurements have proven accurate under steady state conditions, however a difficulty arises when trying to assess the accuracy of these measurements under transient conditions. In this case the problem lies with the torque transducer signal which required heavy filtering to produce a meaningful result, thus vastly reducing its usefulness when transient accelerator variations are being considered. Consequently, with the present instrumentation available on the rig, it is difficult to ascertain whether transient behaviour of true torque differs sufficiently from model torque to affect vehicle driveability. Even if improved rig design did solve this problem, questions of driveability are more naturally approached using an operational road vehicle.

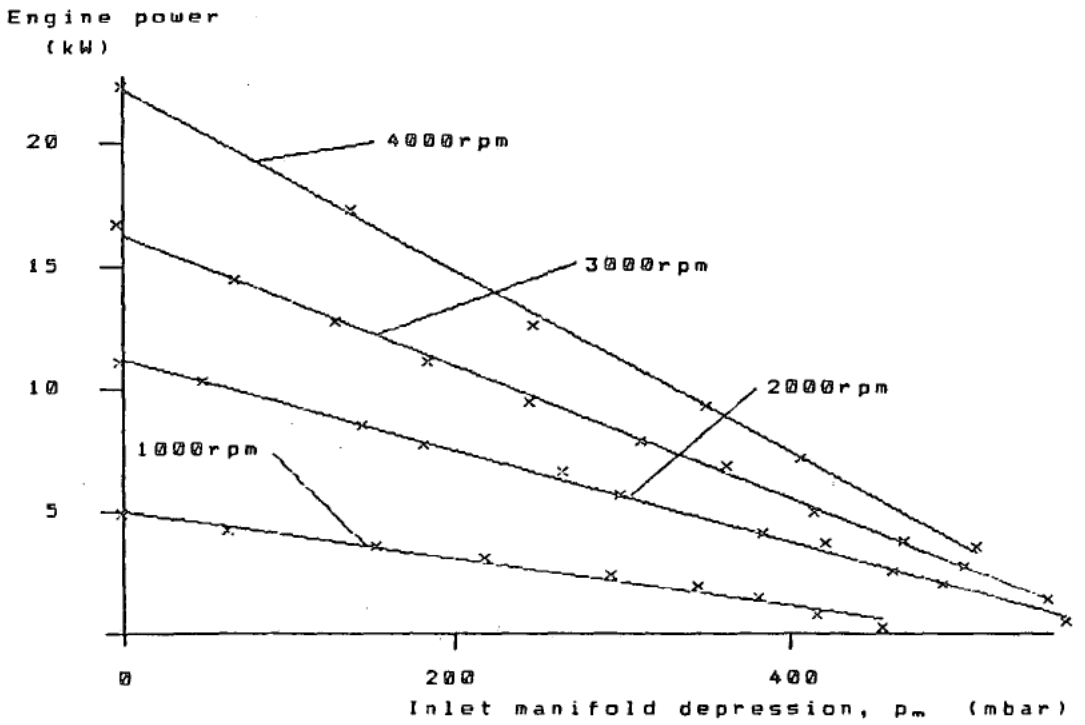

Fig. 4.1 I.C. Engine Power Characteristics

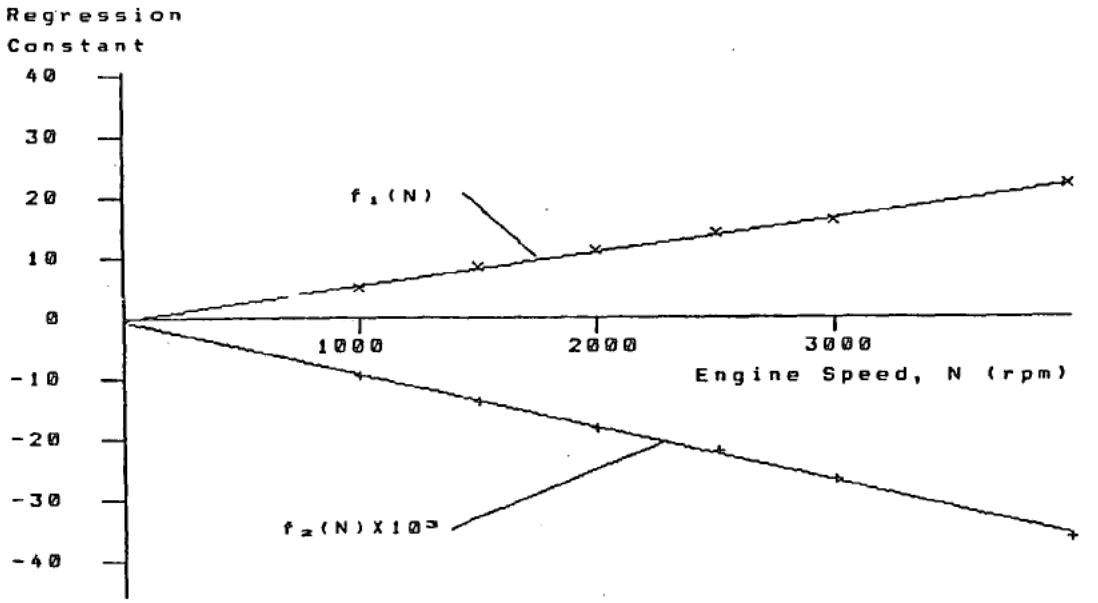

Fig. 4.2 I.C. Engine Power Characteristics: Variation of Regression Constants with Speed

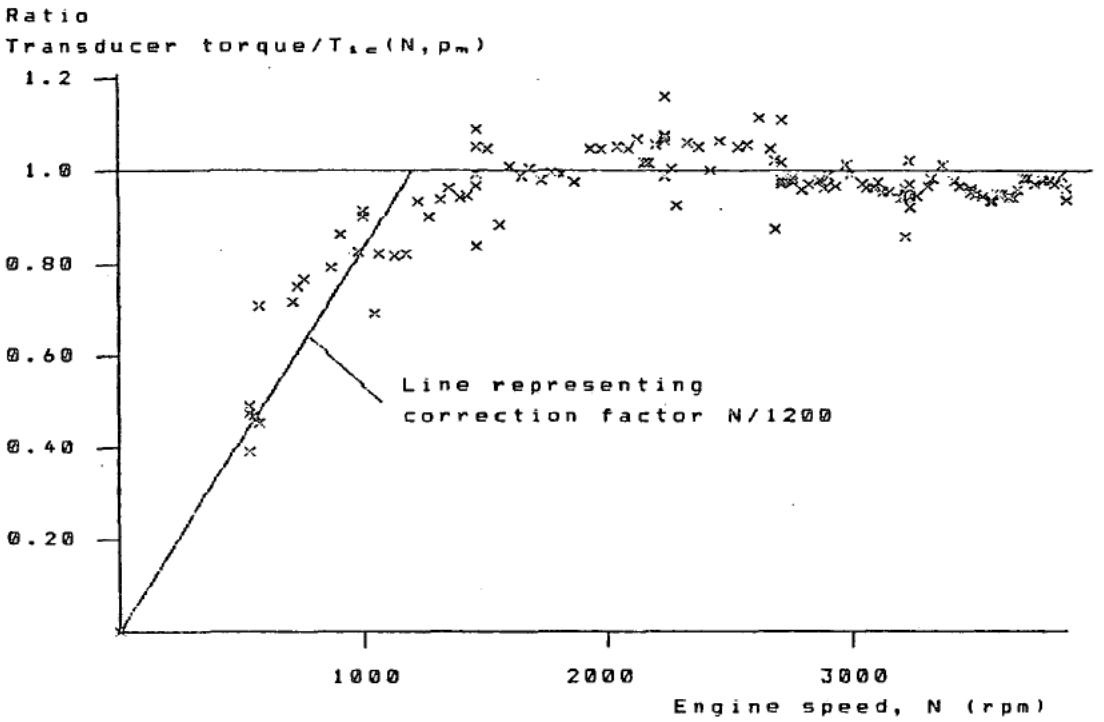

Fig. 4.3 Comparison of Torque Transducer Measurements with I.C. Engine Torque Prediction

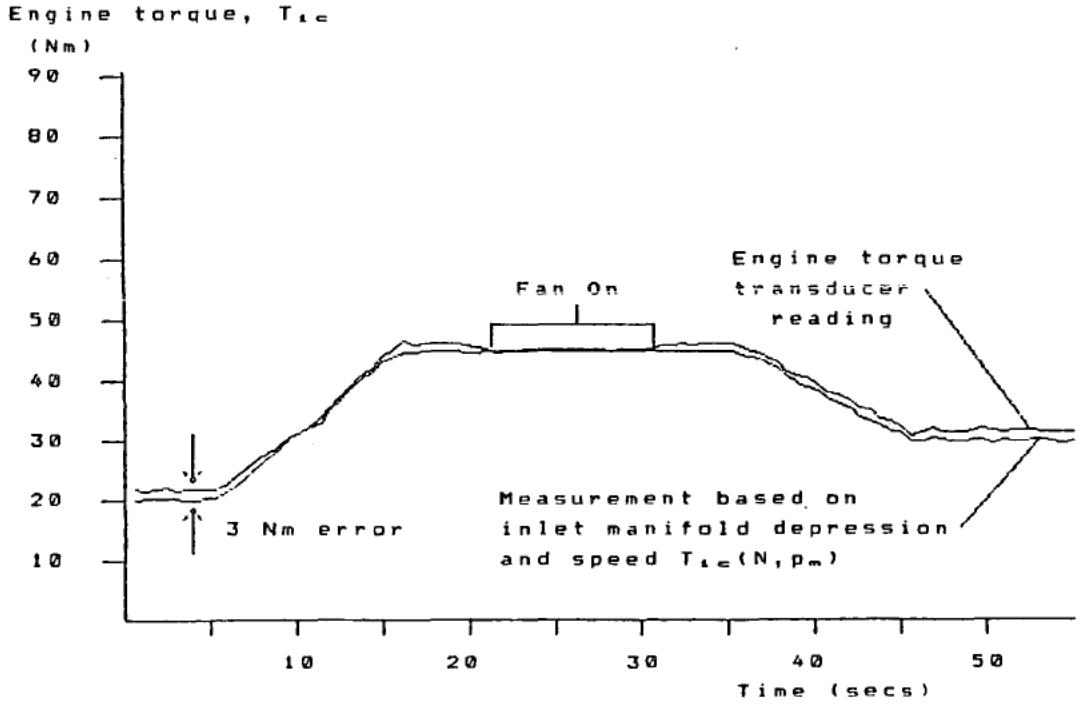

Fig. 4.4 The Effect of the Engine Cooling Fan on Torque Measurements

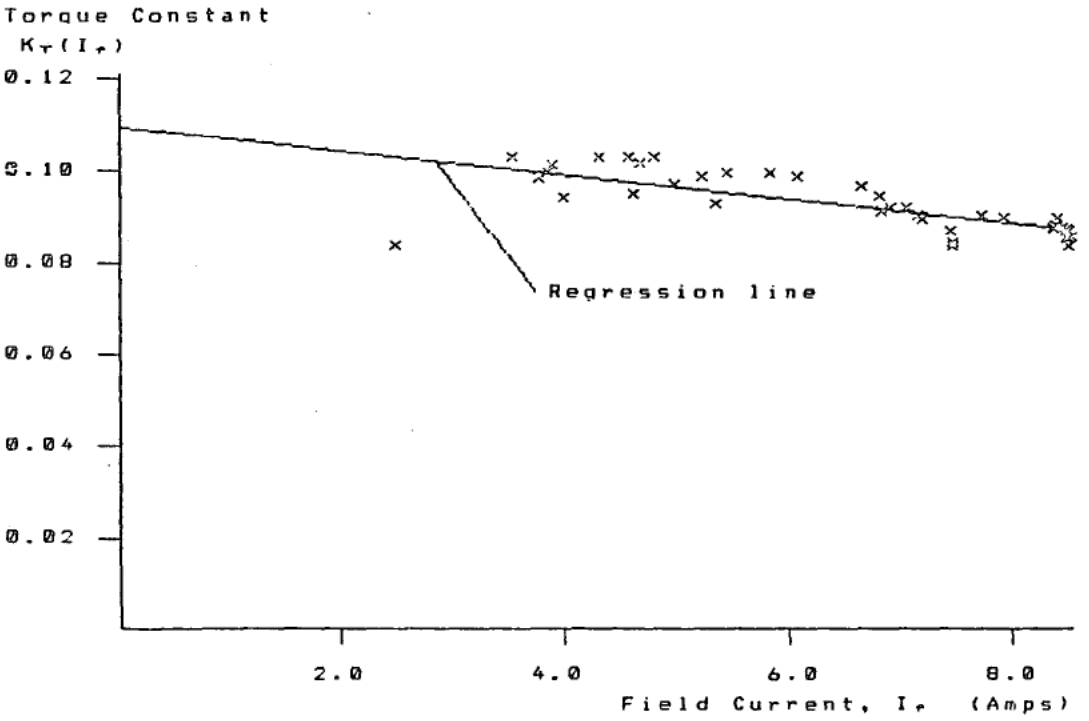

Fig. 4.5 Variation of Motor Torque Constant with Field Current

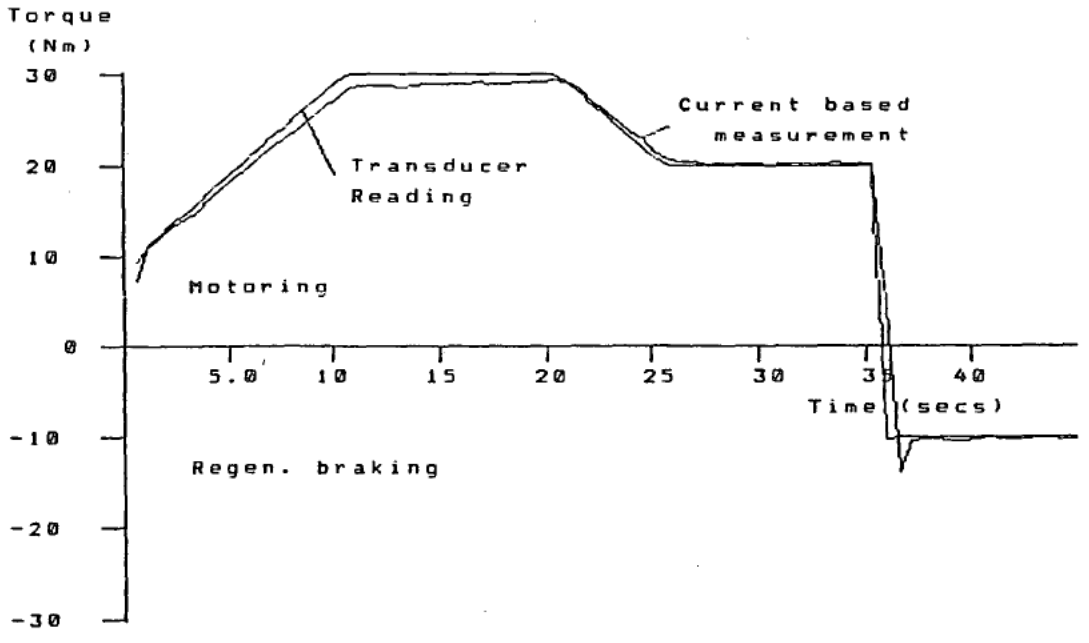

Fig. 4.6 Verification of Indirect Torque Measurements for the Motor

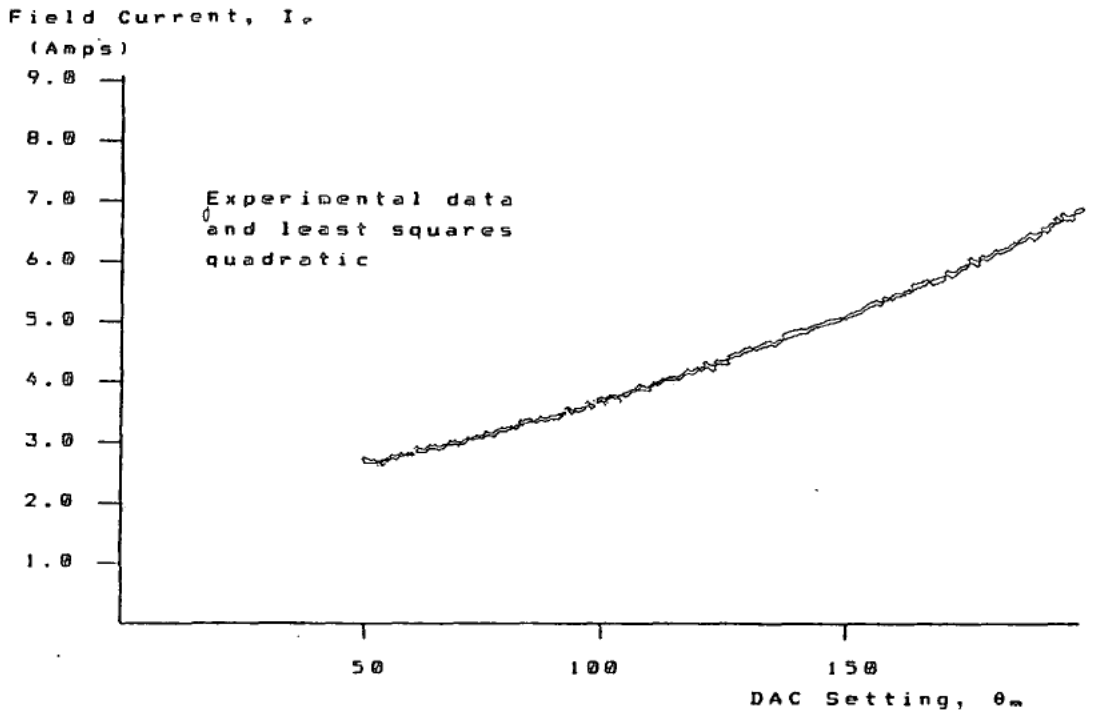

Fig. 4.7(a) Variation of Field Current with Input Demand in the Field Boost Mode

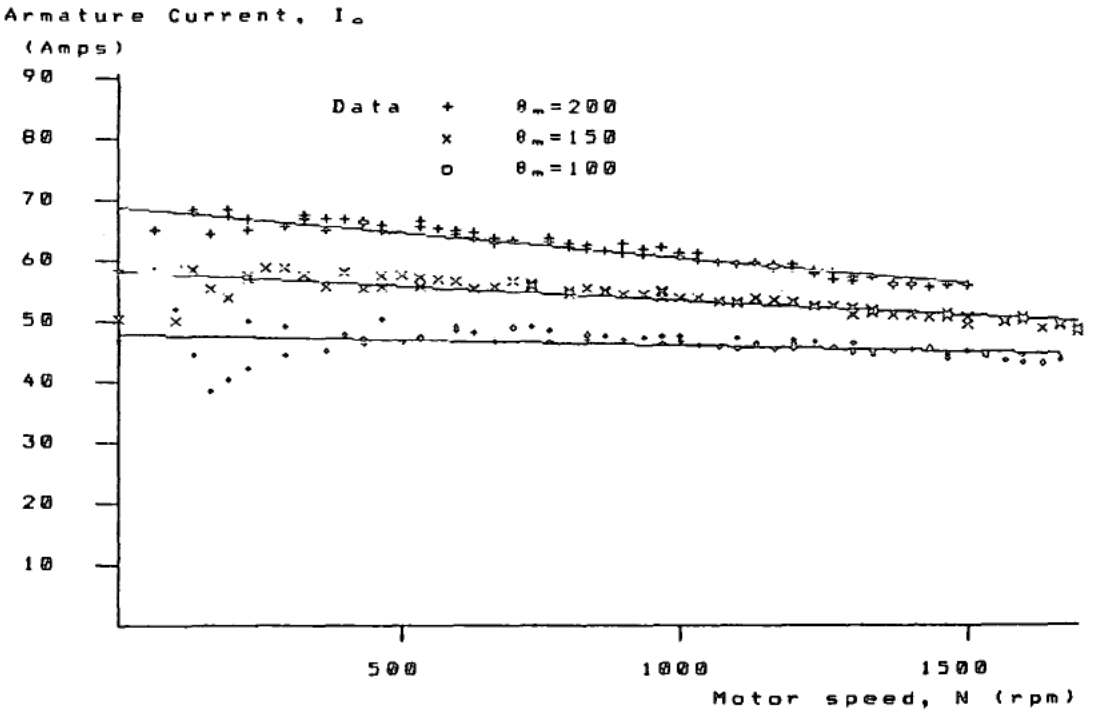

Fig. 4.7(b) Variation of Armature Current with Input Demand in the Field Boost Mode

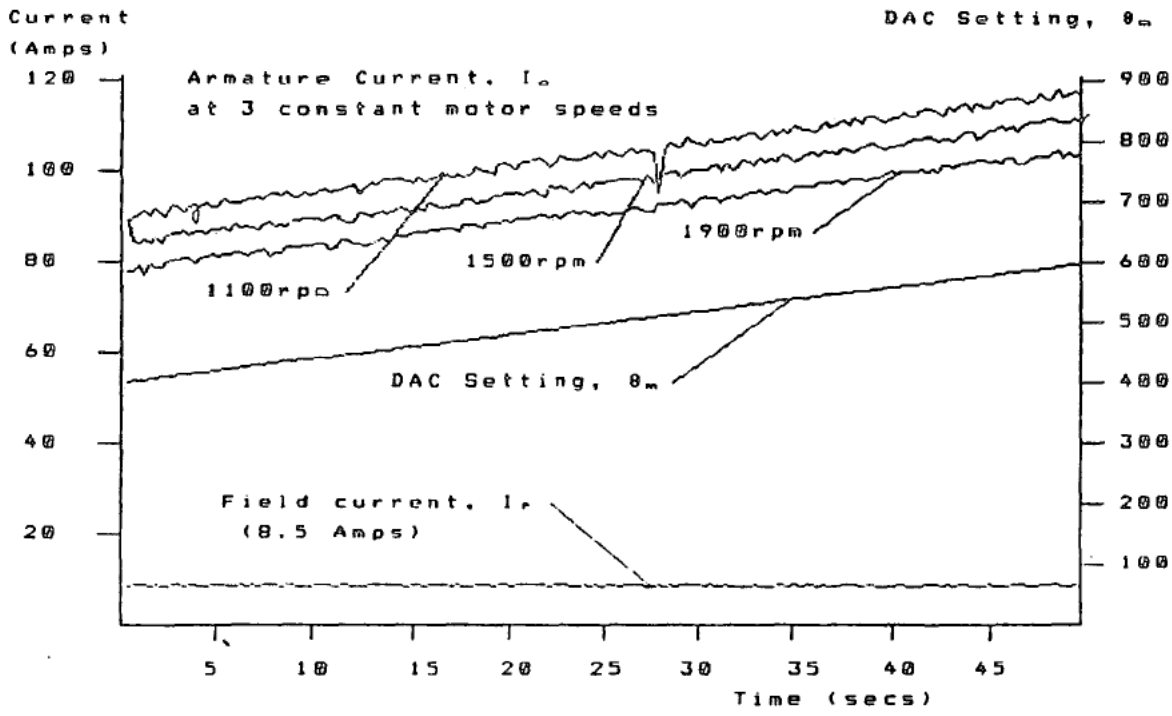

Fig. 4.8 Variation of Current with Input Demand in the Full Field Mode

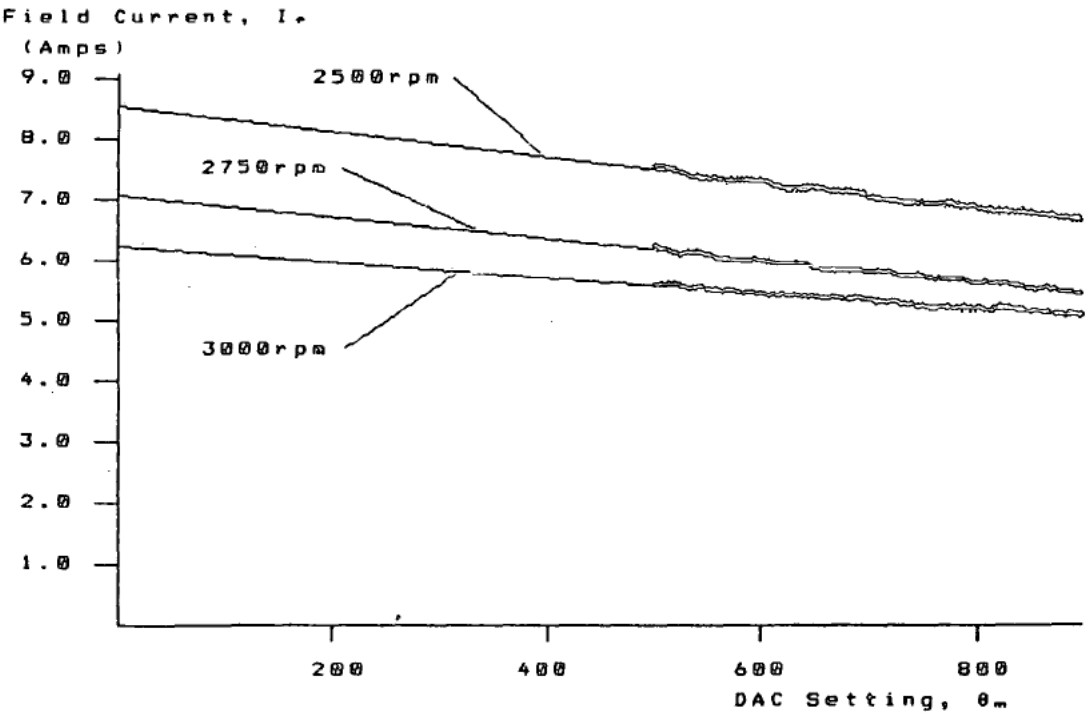

Fig. 4.9(a) Variation of Field Current with Input Demand in the Field Weakening Mode

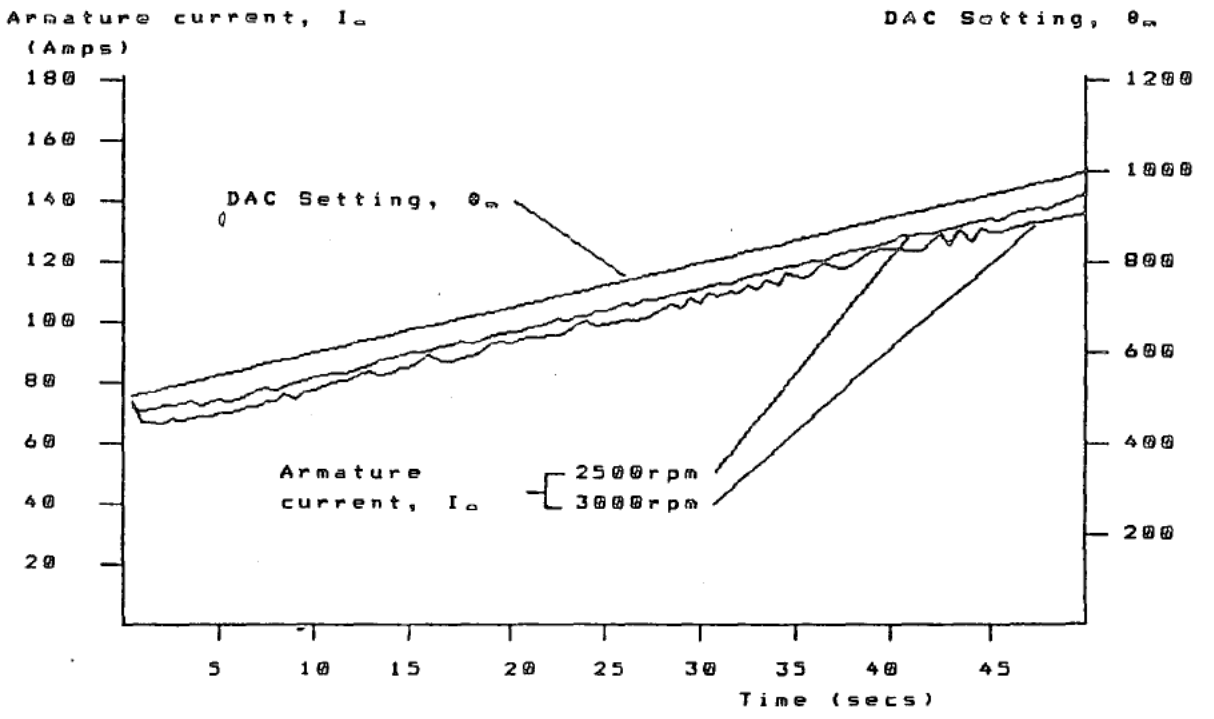

Fig. 4.9(b) Variation of Armature Current with Input Demand in the Field Weakening Mode

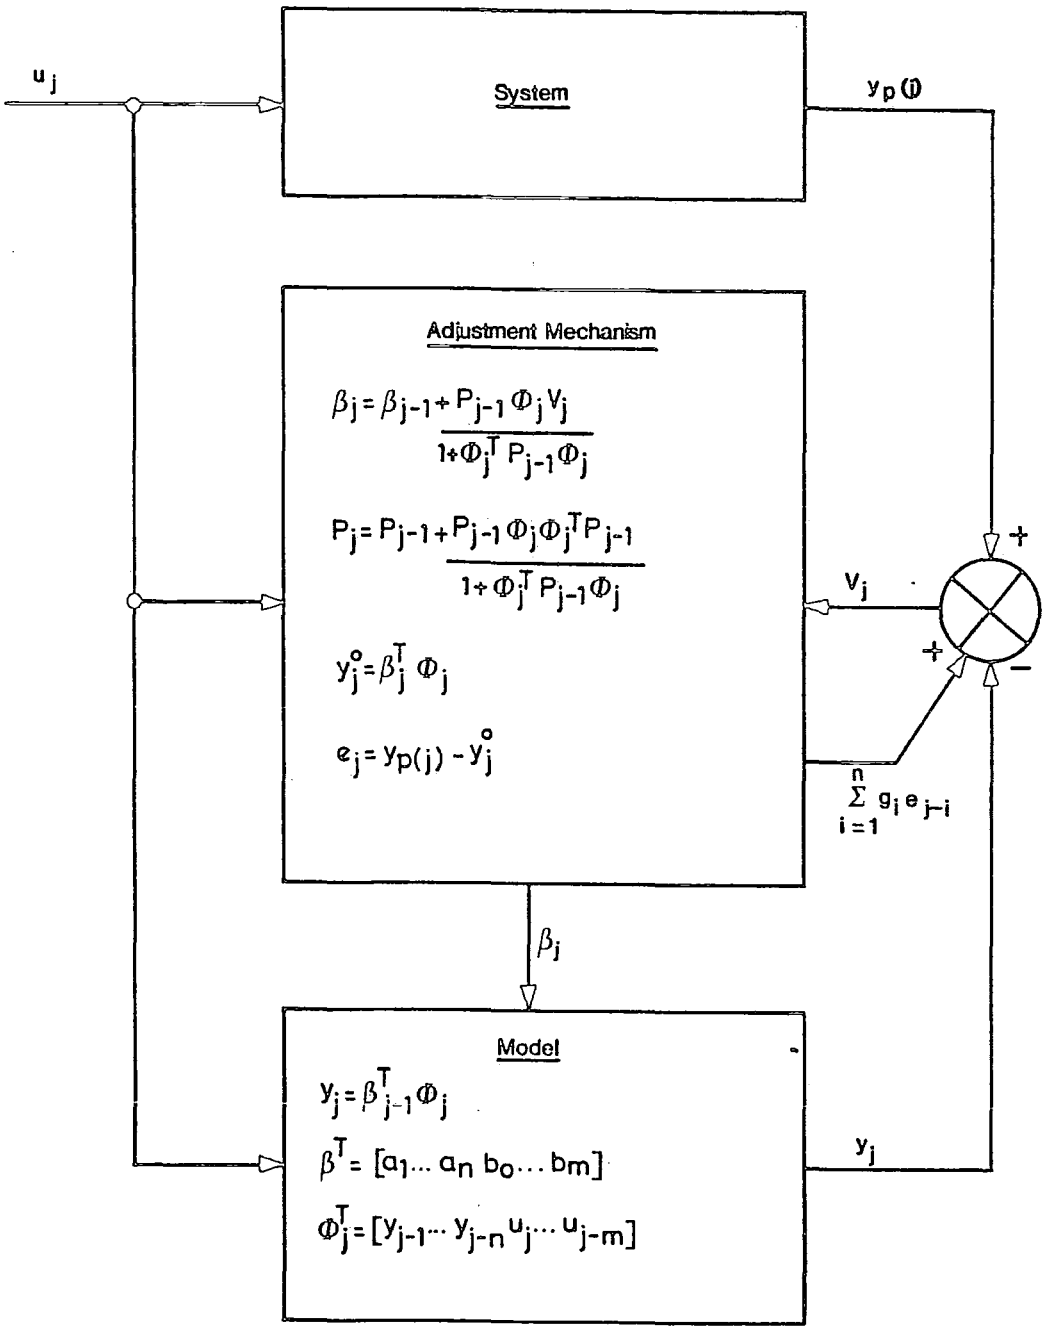

Fig. 4.10 Adaptive Identification System

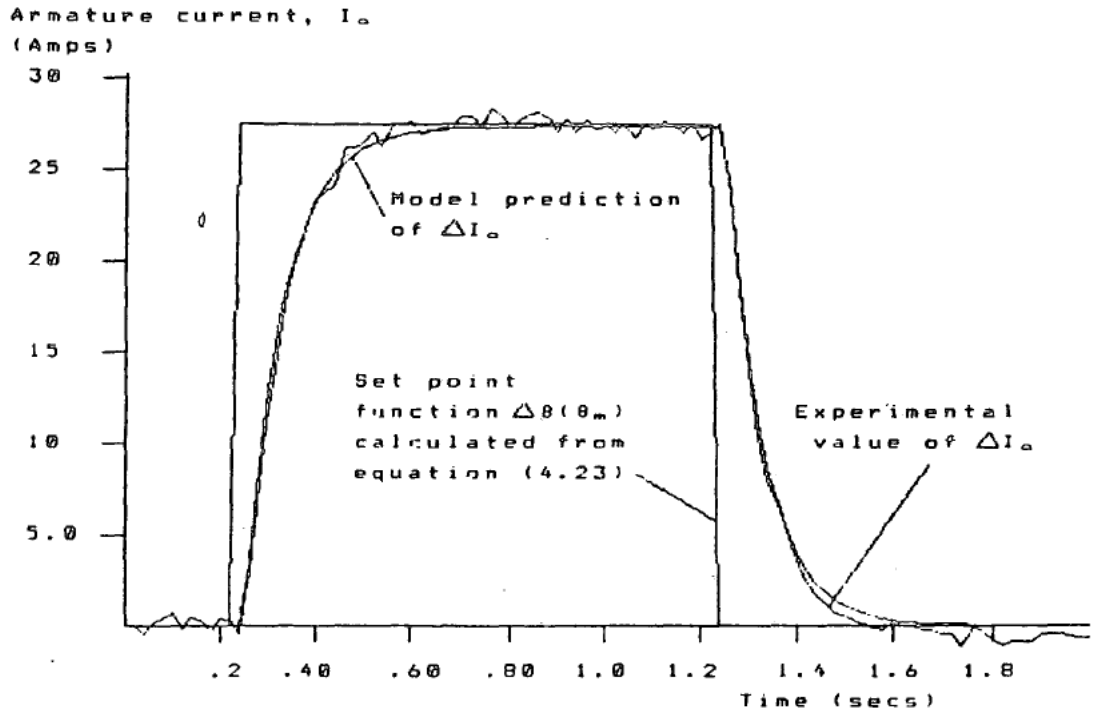

Fig. 4.11 Model Identification for the Closed Loop Armature Current Control Response $G_a(z)$

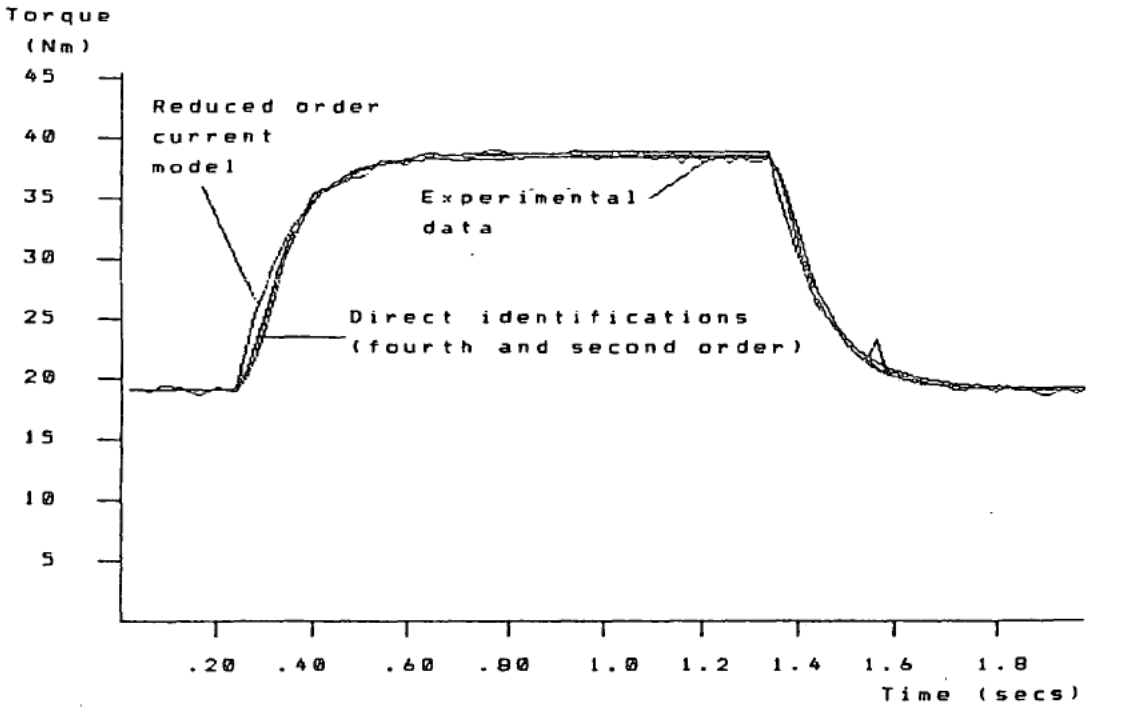

Fig. 4.12 Comparison of Torque Transfer Functions with Experimental Data for the Field Boost Mode

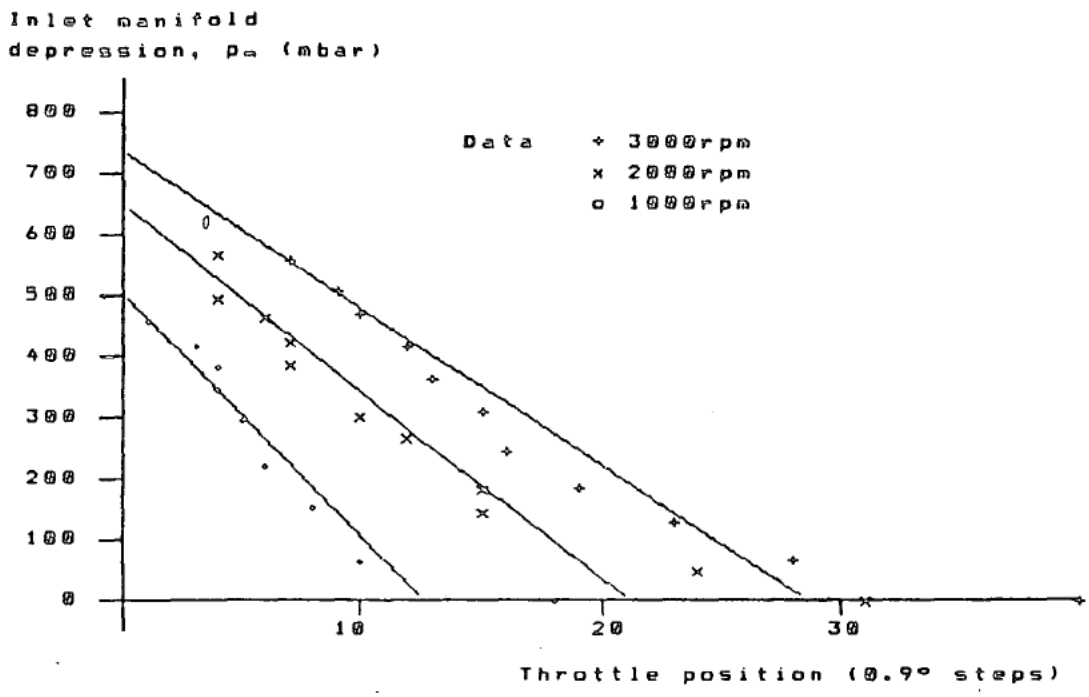

Fig. 4.13 Comparison of Gains for the Manifold Filling Delay and Steady State Data Relating Manifold Depression to Throttle Opening

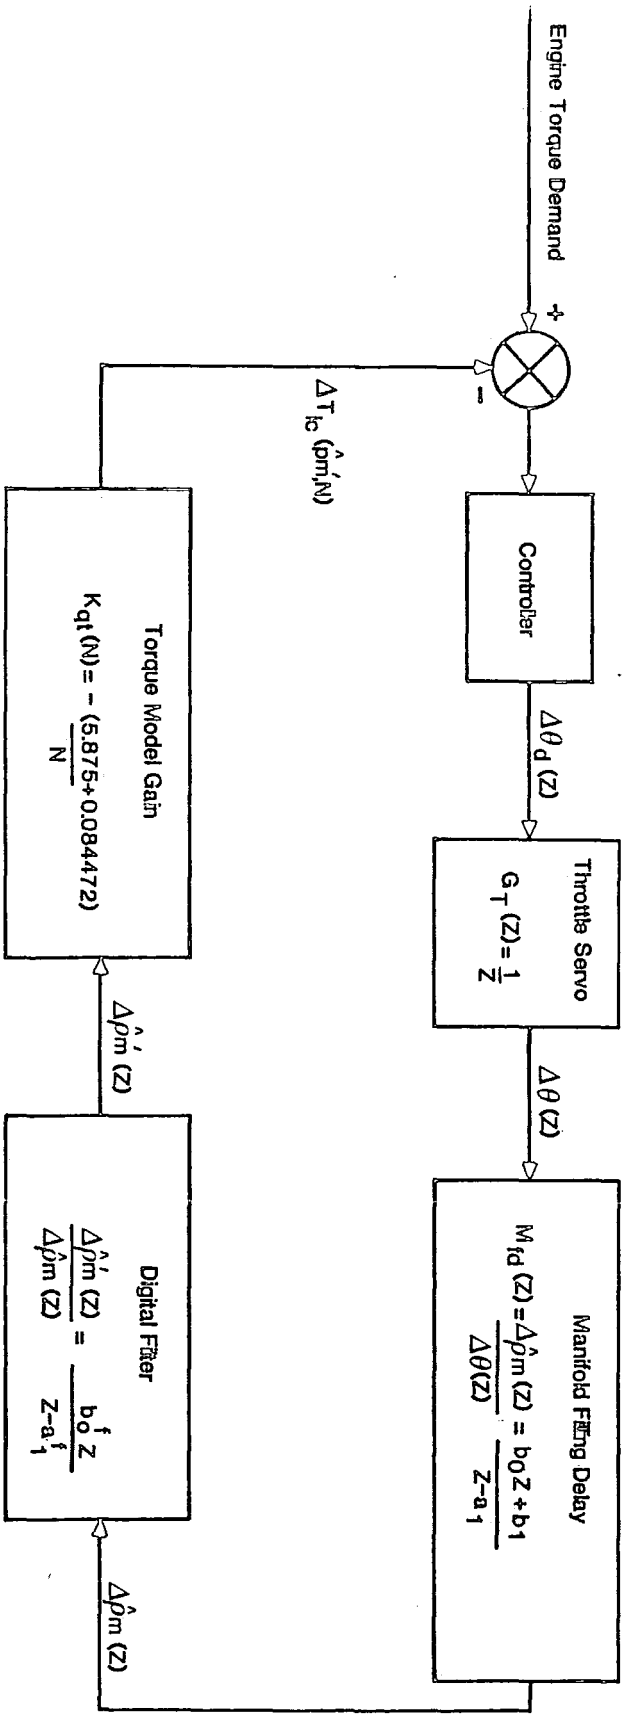

Fig. 4.14 Block Diagram of the Linearised Engine Model

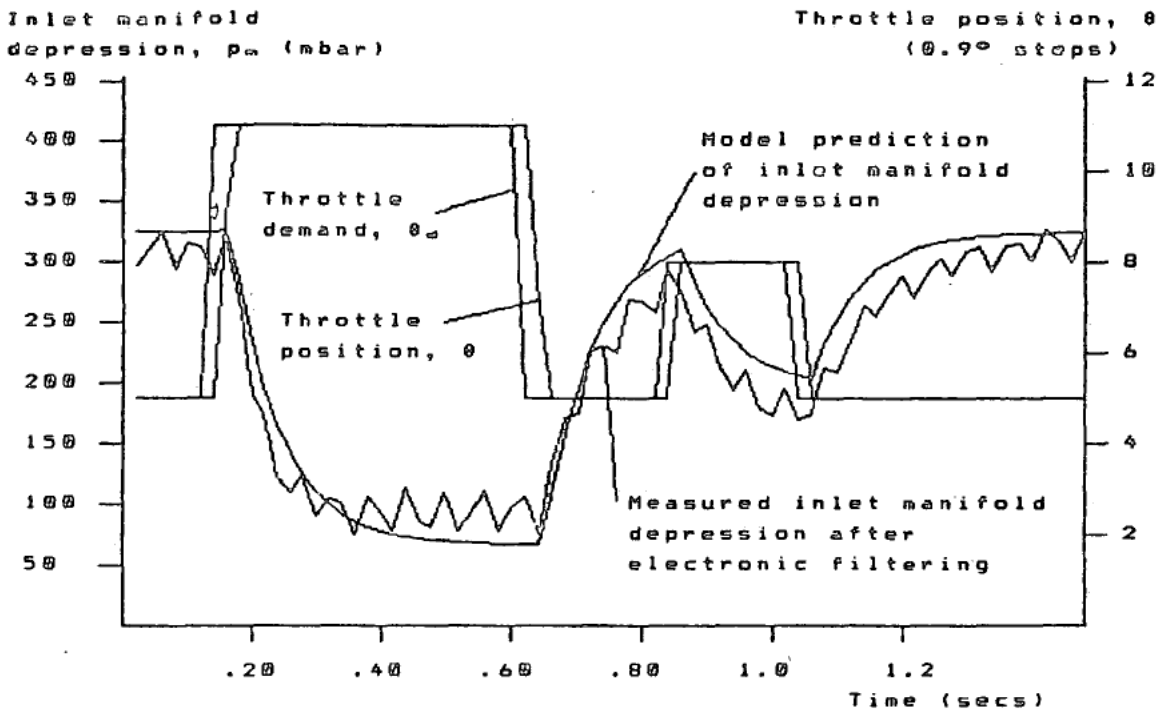

Fig. 4.15 Engine Behaviour at 1000 r.p.m.

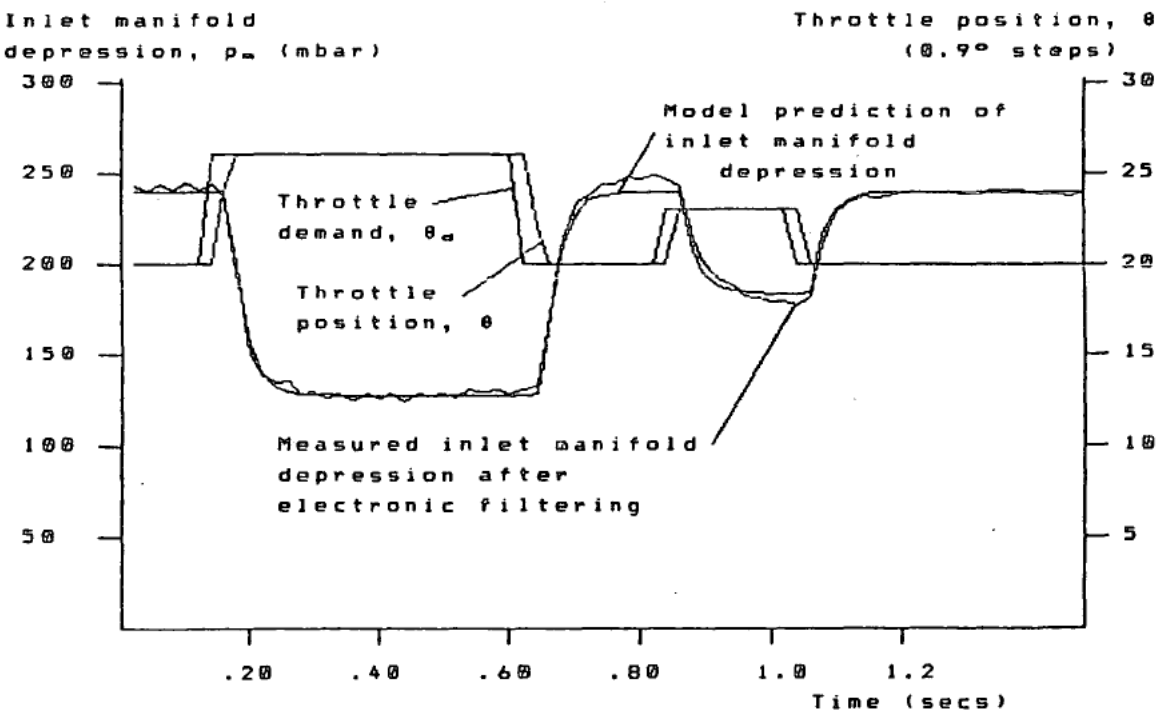

Fig. 4.16 Engine Behaviour at 3000 r.p.m.

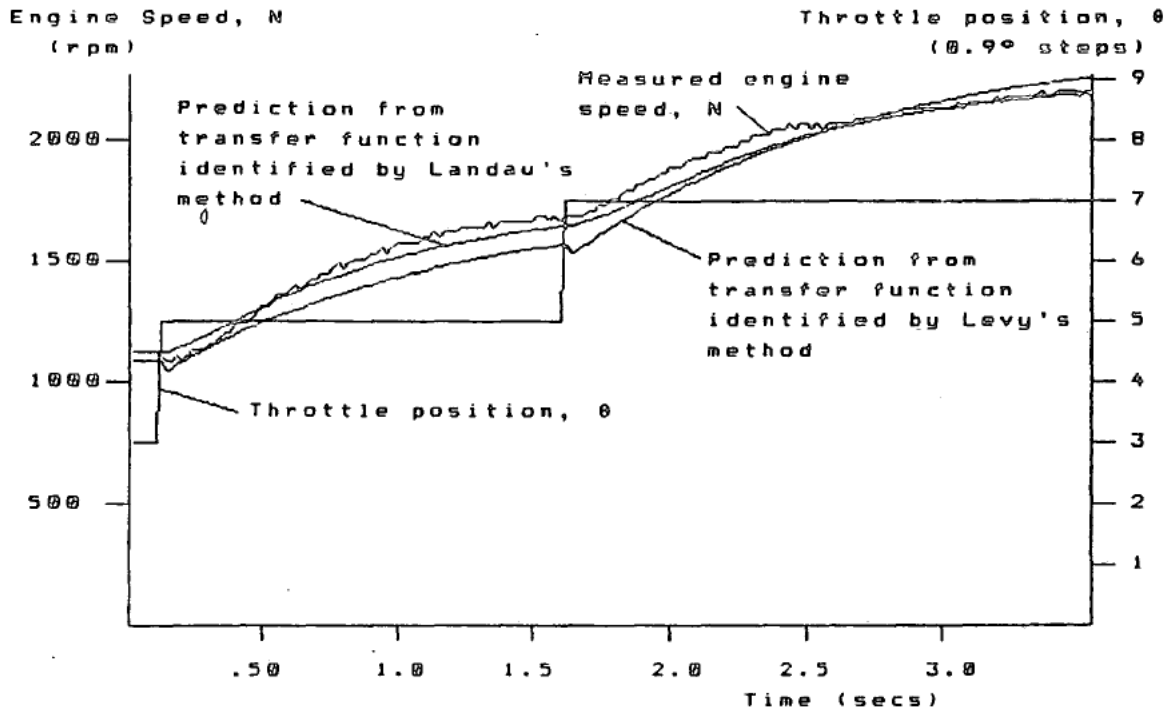

Fig. 4.17 Verification of Transfer Functions for Engine Speed on No-Load

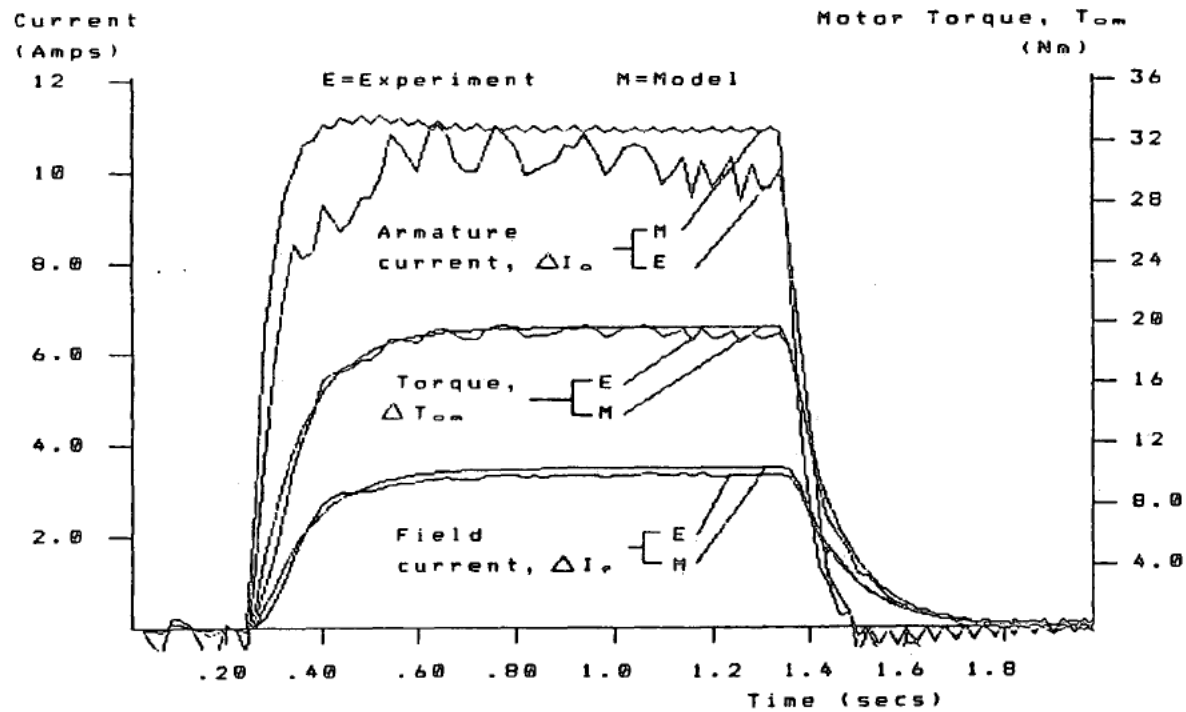

Fig. 4.18 Predicted and Measured Torque and Current Responses for a Transient in the Field Boost Mode

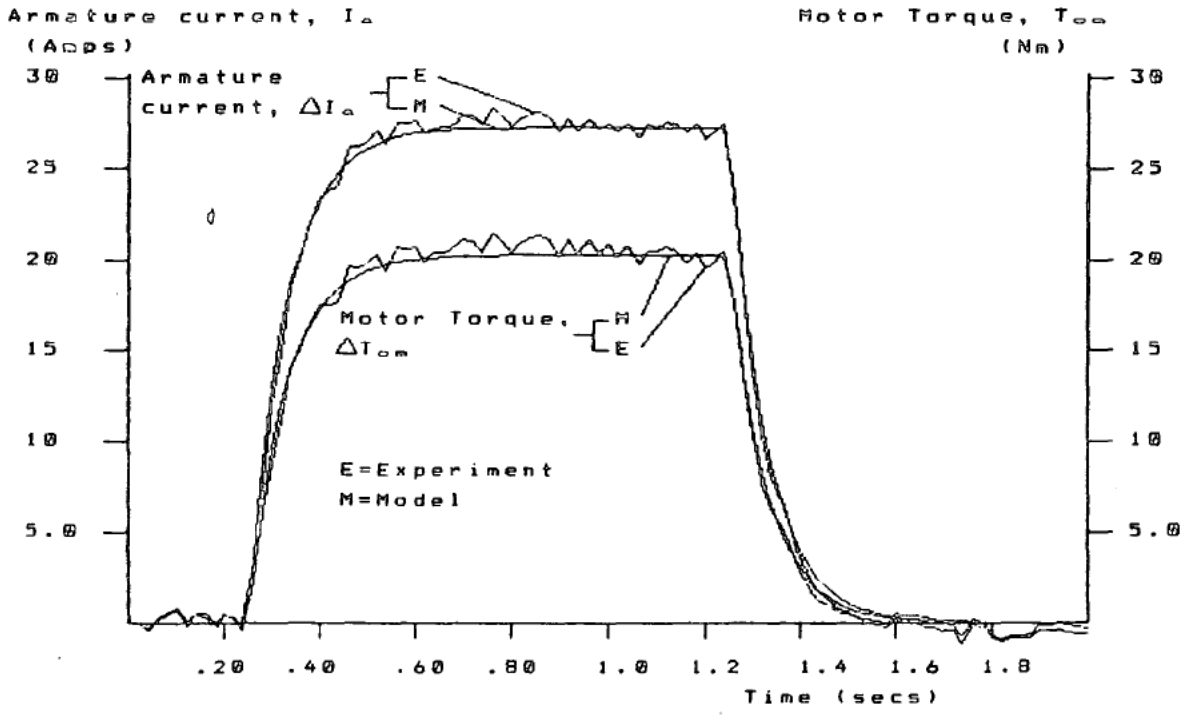

Fig. 4.19 Predicted and Measured Torque and Current Responses for a Transient in the Full Field Mode

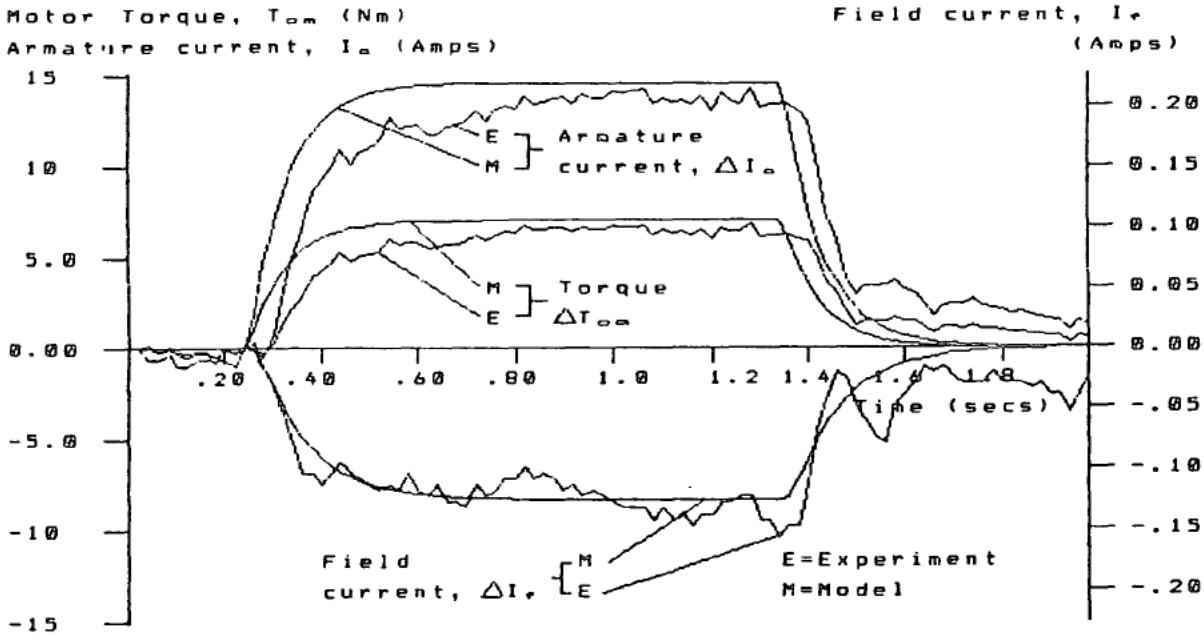

Fig. 4.20 Predicted and Measured Torque and Current Responses for a Transient in the Field Weakening Mode

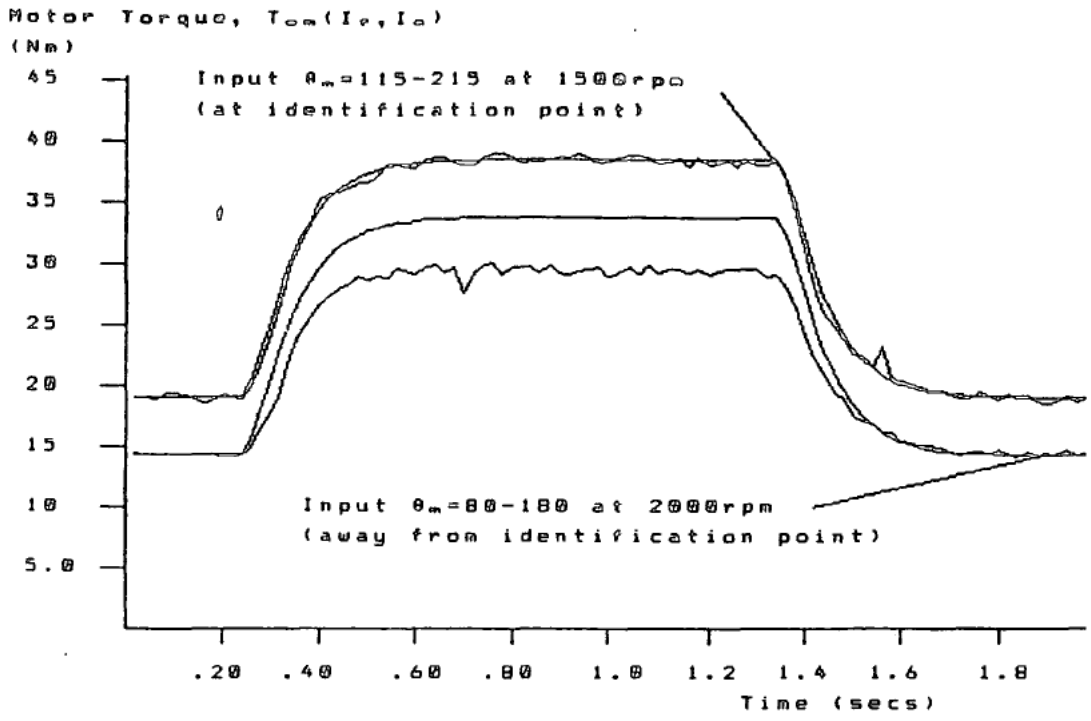

Fig. 4.21 Motor Torque Model: Test for the Direct Identification in the Field Boost Mode

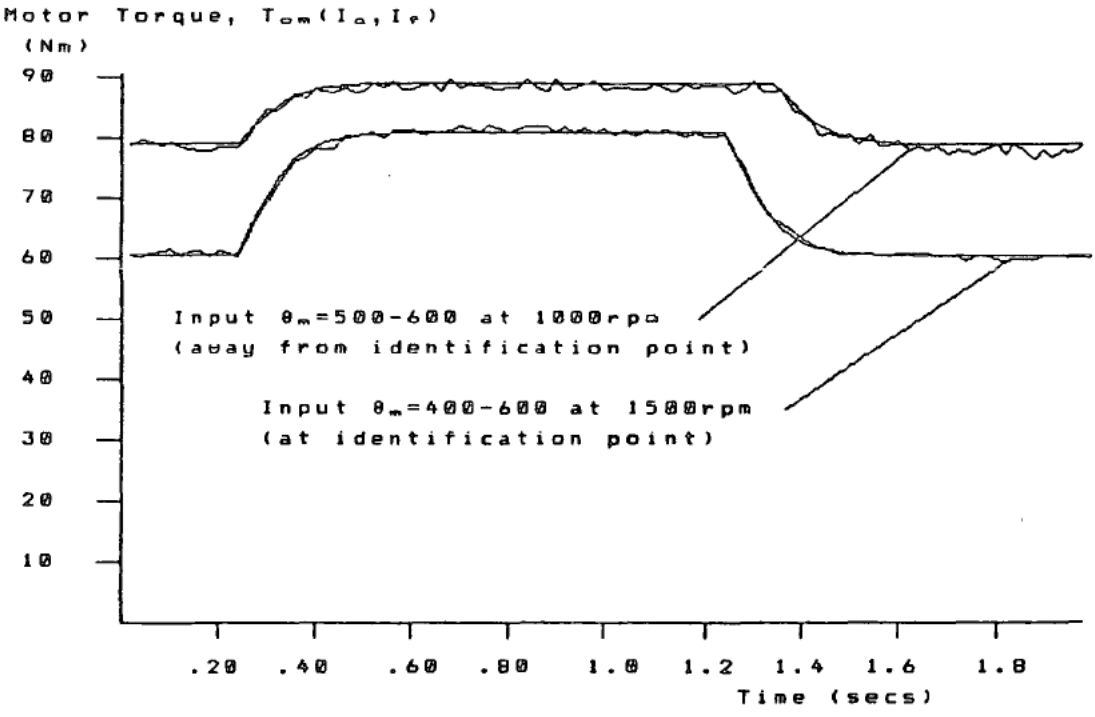

Fig. 4.22 Motor Torque Model: Test for the Direct Identification in the Full Field Mode

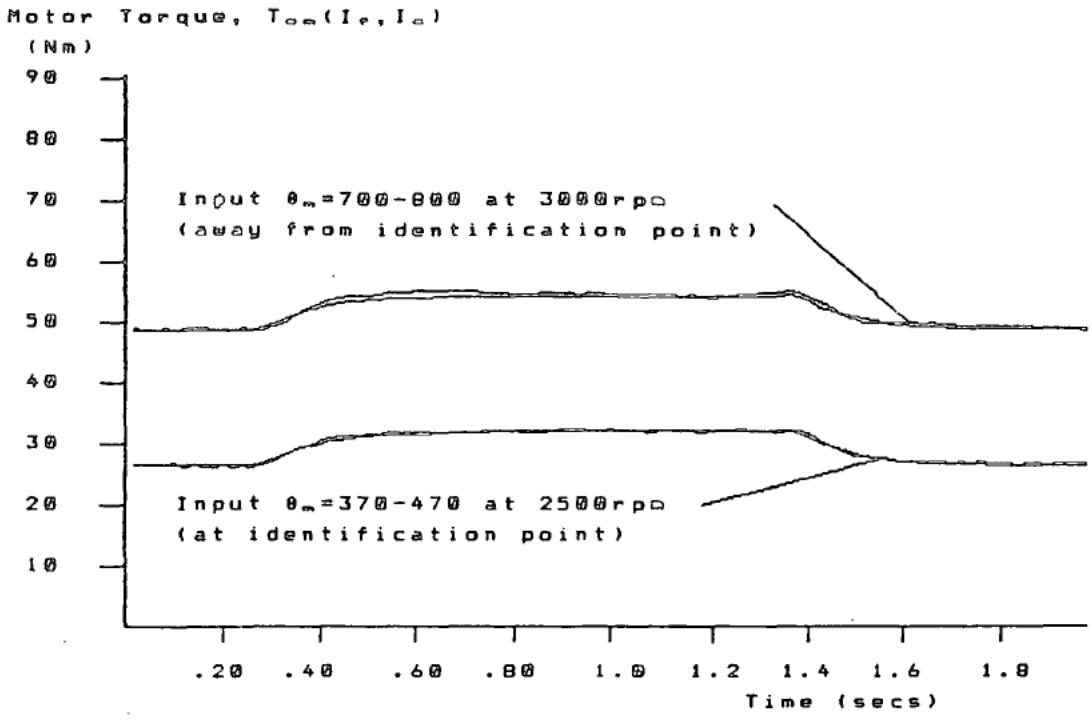

Fig. 4.23 Motor Torque Model: Test for the Direct Identification in the Field Weakening Mode

CHAPTER 5

CONTROLLER DESIGN

This chapter deals with the design of single input single output control systems based on the discrete models of the engine and motor identified from the theoretical analyses and experimental data. There are four separate closed control loop systems in the hybrid drive train as set out below.

1. Control of motor torque
2. Control of engine torque
3. Control of engine speed on no load
4. Control of motor speed on no load

An additional system is developed in chapter 7 to control the speed of the flywheel in response to the demands of a test driving cycle. Unlike the systems mentioned above however, this control loop is not strictly part of the vehicle control system since it replaces a driver. Of the four remaining systems, the control of motor speed on no load is unique, in that no formal design method was used. Instead a simple method of tuning the control parameters on line was adopted because of the peculiar operating conditions of this system, within the gear changing process. In the case of the other three systems a single design technique was adopted and proved successful in all cases [Masding and Bumby, 1988 (d)].

5.1 Controller Design Method

In discussing methods for producing discrete control systems for continuous plant, Katz [Katz, 1981] discusses three methods. In the first, known as the analytical method, the desired system response is translated into a closed loop transfer function. Knowledge of the plant allows an exact solution to be obtained for the controller. Unfortunately in this case the designer has no control over the form of the resultant controller, which may be

of higher order than is strictly necessary to achieve satisfactory response. Such higher order is undesirable since it increases computation time and exaggerates numerical inaccuracies. A second technique uses root locus methods for pole-placement on the z-plane. Desirable closed loop pole locations are usually based on second order criteria such as damping factor and rise time. Loci of constant values of such criteria have rather more complex form in the z-plane than they do in the s-plane, making it more difficult to assess the effect of not precisely obtaining the calculated closed loop poles. Finally Katz proposes transforming a discrete version of the plant to the w-plane followed by controller design by frequency response methods.

Although all these methods have their merits, Katz notes that the w-plane method is very widely used and has resulted in many successful discrete designs. A slightly modified version of this technique has been used to design the three key closed loop control systems for the hybrid vehicle rig. The w-plane method usually begins by discretising the continuous plant and then transforming to the w-plane. Since this plane is very similar to the s-plane almost any analogue design technique may be used. In this case pole placement on the root locus diagram was chosen in preference to frequency response methods mentioned above. Essential to the design method is a knowledge of the discrete plant based on the continuous plant transfer function plus an appropriate hold device. In this application the correct form for the continuous plant was established for the engine and the motor through the physical analysis presented in chapter 3. Following on from this it is possible to add the appropriate hold device and determine the correct form for the z-transfer function before attempting to identify it. Consequently the z-transfer functions, which were finally identified from the experimental data, are ideally suited to the w-plane design method.

In it simplest form the w-plane transform involves a large distortion

of the frequency response which must be taken in to account in the design. A slight adjustment defines the w' transform which largely corrects this fault. Using this adjustment, z-transfer functions are mapped into the w'-plane by the transformation pair,

$$w' = \frac{2}{T_s}\left(\frac{z-1}{z+1}\right) \tag{5.1}$$

$$z = \frac{w' + 2/T_s}{w' - 2/T_s} \tag{5.2}$$

where T_s is the sampling period of the digital system. The complex variable w' is written $w' = \sigma \pm j\nu'$, where ν' is a fictitious frequency. Distortion of the frequency response is governed by the relationship between ν' and true frequency ω.

$$\nu' = \frac{2}{T_s}\tan\frac{\omega T_s}{2} \tag{5.3}$$

If the sampling rate is sufficiently high, then distortion will be negligible over a wide range of frequencies. In the case of the rig control system, with a sampling period of $T_s = 0.02$ seconds, equation 5.3 shows that distortion is under 5% for frequencies up to 38 radians/second (240Hz). In physical terms this frequency is much faster than the response of any of the electrical or mechanical systems to be controlled on the rig. Consequently even when considering the relatively fast acting motor torque control system, the part of the root locus containing the closed loop poles lies well below the 38 rad/sec mark on the imaginary axis. Once the plant is available in the w'-plane it remains to select the required pole-locations appropriate to the chosen performance criteria. Pole locations are chosen to give specified values of rise time, t_r, and damping factor, ξ. For a simple second order system these values are related to the closed loop pole locations by the following equations:

$$w' = \sigma \pm j\omega_d \tag{5.4}$$

$$\omega_d = \frac{\cos^{-1}(-\xi)}{t_r} \tag{5.5}$$

$$\sigma = \frac{\omega_d}{\tan(\cos^{-1}(-\xi))} \tag{5.6}$$

If the real system is dominated by one complex pole pair, then pole locations calculated from these equations will provide a good initial target for controller design. In practice the combination of hybrid vehicle plant plus the controller usually produces an additional closed loop pole on the real axis which has a significant effect on the transient response. Although this does cause a departure from the calculated performance, the effect is partly reduced by placing the controller zero near to the additional system pole. As a result the pole zero pair virtually cancel and so the the second order complex pair are the dominant factor determining the transient response [D'Azzo, 1975].

5.2 Control Algorithm

Although the M68000 microprocessor would be quite capable of implementing sophisticated control algorithms, experience has shown that satisfactory results can be obtained in all the hybrid vehicle applications by using a simple proportional plus integral (P+I) controller. Using such a low order controller brings several advantages which are important to the hybrid vehicle. Not least of these is the extreme robustness of the controller, even in the face of adverse system behaviour such as that occasionally encountered from a cold engine. This advantage has long been recognised by industrial users, who also capitalise on the fact that such control can be applied to a plant which has not been fully identified [Krikelis and Fassios, 1984]. A second advantage is that the P+I controller gives very fast execution times, thus allowing the M68000 time to carry out a multitude of other control

tasks within the interrupt period. Typically the control algorithm takes only 0.8 ms to execute in integer arithmetic. In an ideal digital system data sampling is followed instantaneously with output of the controller response at the sampling point. In practice this clearly can not be the case since the control algorithm must take a finite time to execute, however by using this fast algorithm the M68000 is able to carry out all data sampling and output the last of its four controller responses within 6 ms of each interrupt occurring. Such a delay although not ideal, is quite insignificant compared with the response times of all the systems involved.

Due to the similarity between the s and w'-planes, the proportional plus integral controller retains its usual form:

$$g_c(w') = g\left(\frac{w' + a}{w'}\right) \tag{5.7}$$

Phase one of the design process is concerned with determining g and a in order to best achieve the pole locations calculated from equations (5.4)-(5.6). In all cases the design criteria are $\xi = 0.707$ (critical damping) and a rise time which represents a realistically achievable minimum for the system concerned. Once g and a have been selected, equation (5.7) can be transformed in to the z-plane by the reverse mapping of equation (5.1) to give:

$$g_c(z) = \frac{(g + k_i)z + (k_i - g)}{z - 1} \tag{5.8}$$

where $k_i = gaT_s/2$. From this equation comes the direct realisation for the controller output u_k.

$$u_k = u_{k-1} + (g + k_i)e_k + (k_i - g)e_{k-1} \tag{5.9}$$

This control algorithm is then tested by examining the step response of the completed system. Since the assumption, made in the design, that the

resultant closed loop system would be second order is not true, the response of the system may be found to be unsatisfactory in which case the parameters g and k_i are fine tuned until that is no longer the case. Final versions are then implemented on the M68000 using equation (5.9) for testing in actual operation.

5.3 Design of Individual Controllers

5.3.1 Engine Torque

The engine torque controller receives its demand signal from the mode controller, and must match the indirectly derived torque measurement to this demand. Apart from the P+I control algorithm itself, the other key elements in the engine torque control loop are the throttle servo, the manifold filling delay, a digital filter and the torque model all of which are illustrated in figure 4.14. In drawing this diagram it was assumed that, due to the vehicle inertia, engine speed is essentially constant over the time scale of manifold pressure changes and hence, in the present context, speed is assumed constant over the time scale of engine torque control. In addition to being valid only for constant speed, a further limitation of this diagram is that the simple linear delay element used to model the throttle servo-system is only completely accurate when the controller demand does not exceed a movement of four steps per sample period. With these limitations taken in to account, the design aims for the system are critical damping and a rise time of 200 ms. Substituting these values in to equations (5.4)-(5.6) gives the pole locations $w' = -11.78 \pm j11.78$. Due to the speed dependent elements in the block diagram, it would be necessary to have a whole series of controller designs if these performance criteria were to be met at all times. Clearly this is not a practical option, and in any case variations in engine gain and dynamics are not great, making it reasonable to assume that one design

for a mid-range speed may well be sufficient. To test this theory a torque controller was designed for the engine operating at 2000 r.p.m., taking the relevant parameters for the manifold filling delay from Table 4.7. Prior to compensation, the system locus for 2000 r.p.m. appears as in figure 5.1(a) with two open loop poles, one due to the manifold filling delay and the other, nearer to the origin, due to the filter. After slight adjustment to give a satisfactory response the controller

$$g_c(w') = 0.4\left(\frac{w'+7}{w'}\right) \tag{5.10}$$

was selected which produces the modified system locus of figure 5.1(b). As shown the final pole locations are not too far removed from those calculated for guidance however the combination of engine plus controller has produced a third, real, closed loop pole. After applying the inverse transform equation (5.10) becomes

$$g_c(z) = \frac{0.428z - 0.372}{z - 1} \tag{5.11}$$

Live control system tests were carried out initially at design speed, by applying a step increase of torque demand amounting to 10 Nm. Results from the test are illustrated by figure 5.2(a). As the test proceeded the M68000 simultaneously carried out a complete simulation of the system using the state space equations developed in section 4.5. Three results from the simulation are plotted together with their experimental counterparts in figure 5.2(a). Of these, experimentally measured manifold pressure prior to digital filtering, shows considerable noise but the simulation accurately represents the general trend. Torque traces match almost exactly, since by this stage manifold pressure has been digitally filtered and noise effects are virtually eliminated. This slight difference in the torque repsonse may be partly

explained by the real pole. D'Azzo predicts that such a pole positioned between the complex pair and the origin, as is true in this case, will slow the system response and reduce overshoot. Close examination of figure 5.2(a) shows that the experimental response does not exhibit the slight overshoot predicted by the simulation.

Good system performance is to be expected at original design speed, however the engine must be accurately controlled at all speeds. To assess what effect the changing gain and dynamics has on system performance a second test was carried out at 3000 r.p.m. As figure 5.2(b) shows, acceptable performance is still obtained with the original control parameters chosen for the 2000 r.p.m. model. In this case the reduced manifold filling delay gain increases system rise time. By incorporating the correct model of the manifold filling delay in the simulation this effect is accurately predicted.

A small change in demand of 10 Nm was chosen for these initial tests so that the non-linear effects of the throttle step rate limitation would not affect the system. Under normal operating conditions the control system might well experience far greater changes in demand. For example large step changes in demand are likely when the vehicle operator wants to change from acceleration to cruise. Figure 5.3 shows how the control system reacted to a 25 Nm drop in demand, which is typical of such operating conditions. As this diagram demonstrates there is no significant deterioration in system performance. Also included on this diagram is the trace from the torque transducer which provides further confirmation of the accuracy of the indirect torque measurement.

5.3.2 Electric Motor Torque Control

Engine torque control must be carried out in the face of a system which has continuous speed dependent variations in both gain and dynamics. In the case of the motor similar variations in system transfer function occur dependent on the particular operating point. In addition motor operation is divided in to three distinct operating modes due to the action of the power electronics control unit. As defined in chapter 4 these three modes are named field boost, full field and field weakening.

Of all the three modes only full field operation allows the relationship between the controlled output, T_{em}, and the controllable input, θ_m, to be accurately described by a single linear transfer function. In the field weakening and field boost modes, this relationship, vital to controller design, can only be described by a complex fourth order model with non-linear gain functions. Nevertheless the analysis took due account of the fact that it would be extremely difficult to produce a single controller which could continually compensate for such variations in the system and thus stay optimally tuned at all times. Consequently the model was successfully reduced to a single linear second order result which is valid for one particular set of operating conditions. Furthermore tests showed that despite this apparent limitation one such transfer function identified at a representative point for each mode, could model motor performance across the whole of the mode with acceptable accuracy.

Building on this result it is possible to design a single fixed gain controller for each mode which maintains satisfactory performance at all times. One additional task which arises from the adoption of this strategy is accurate determination of the mode based on available measurements. Only by achieving this can the correct set of controller parameters be selected. To start with however three sets of controller parameters were produced using

the w' root locus design technique. The common design criteria used for all three modes was rise time of 150 ms and critical damping.

Figure 5.4 illustrates the block diagram for electric motor torque control. In this control system the plant is given by those discrete transfer functions listed in Table 4.6, which directly describe the relationship between accelerator demand, θ_m, and the indirect torque measurement, $T_{em}(i_a, i_f)$. As such the plant includes the power electronics, motor and electronic current filters. No additional filtering is necessary to obtain a satisfactory torque signal from i_a and i_f. Part of the controller design process is illustrated by figures 5.5(a) and 5.5(b) which show the uncompensated and compensated locus respectively for the field boost mode. Final results are given in Table 5.1 which lists a set of controller gains designed on the basis of the models chosen as representative of mid-band conditions for their mode, as listed in Table 4.6. Controller gains are given both in w'-plane form and in z-plane form ready for implementation on the M68000.

Parameters	Field Boost	Full Field	Field Weakening
w'-plane g	3.0	9.0	15.0
a	15.0	15.0	9.0
z-plane g	3.0	9.0	15.0
k_i	0.45	1.35	1.35

Table 5.1 Electric Motor Torque Control Parameters for All Three Operating Modes

5.3.3 Motor Torque Control Test Results

As with the engine, step tests of 10 Nm are ideal for testing torque controller performance. There are two reasons for carrying out these tests on the motor, firstly to verify system response at design conditions. Secondly, and more importantly, to check that performance does not deteriorate unacceptably

under extreme conditions, well removed from those of the design. Figure 5.6(a) shows the field boost controller giving excellent results under design conditions. Furthermore there is almost exact agreement between experimental data and the simulation carried out by the M68000. In a second test performed away from design conditions for the field boost mode, falling plant gain causes system rise time to increase as illustrated by figure 5.6(b).

As expected from the theory the full field controller produces precisely the same result regardless of the operating conditions. One representative experiment for this mode is shown in figure 5.7.

Like the field boost mode, the field weakening mode is expected to produce variations in controller performance. Under initial design conditions a good system repsonse is obtained which closely matches the simulation despite a slight change in system gain (figure 5.8(a)). Away from design conditions, figure 5.8(b) shows that changing plant dynamics have caused a slightly oscillatory response from the system.

There is obviously no absolute law which can decide whether or not any of these changes in performance are unacceptable, but in no instance is system stability even remotely called into question neither is there a large departure from the designed rise time.

5.3.4 Mode Determination

Isolated step tests carried out entirely within one mode are not representative of working conditions for the torque controllers. Under normal conditions the motor will regularly pass between modes and so in response, the torque controller must switch between parameters. There is no definite signal from the power electronics to indicate the operating mode of the motor, consequently a mode determination algorithm has been written into the software using measurable signals. As shown on the block diagram

of figure 5.4, the necessary inputs are speed, field current and accelerator demand. Reference to figure 3.3 shows that these measurements do define the operating mode. The mode determination rules are based on a simple set of inequalities:

1. If $i_f > 8.4$ then Mode = Full Field
2. else if $\theta_m < 515 - 0.1 \times$r.p.m. then Mode = Field Boost
3. else Mode = Field Weakening

To avoid rapid switching between controller gains a time constraint only allows a new mode to be selected after every 10 interrupts (200 ms). The mode determination rules were tested by instructing the motor to follow an arbitrary torque profile and allowing natural speed variations to occur. As can be seen from figure 5.9 the experiment succeeded in making the motor pass through all of its operating modes, but more importantly smooth control was achieved at the transition points.

With reference to rule 2 above, representative mode transition points have been added to figure 3.3. Bearing in mind that all areas to the left of the marked transition points will be treated as either full field or field boost it is apparent that on occasion these latter modes will be selected when in reality the motor is operating in the field weakening mode. This is deliberately chosen to be the case so that the relatively high field weakening gains are never used erroneously in the field boost mode when plant gain too is relatively high.

Tests show that should this condition arise the system is liable to go unstable whereas should the reverse mistake be made the use of low field boost control gains merely results in an extended rise time for the field weakening mode. In figure 5.10 the field weakening gains were deliberately used under operating conditions similar to the design state for the field boost controller. Although this particular step test does not illustrate unstable

behaviour, overshoot is unacceptable and there is evidence of amplification of noise in the controlled signal. Since this experiment took place near the design conditions for the field boost mode, it should be possible to predict these oscillations by using the field boost model in a simulation. When the field boost model is combined with the field weakening controller gains, considerable oscillation is predicted as shown by figure 5.10, but it is much less than that observed in experiment. Although this might in part be caused by operating conditions not precisely matching those applicable to the model, a more likely cause is noise and quantisation errors becoming a problem with the unsuitably high gains.

5.3.5 Engine Speed Synchronisation

Whenever the hybrid vehicle is operating in all electric mode or is stationary, the i.c. engine can be uncoupled from the drive train by means of the one-way clutch. Since in either of these situations the engine is not required to provide torque, the most obvious strategy is to shut it down entirely in order to conserve petroleum fuel. Adopting this strategy means that the next time the engine is needed it must be started and synchronised with the moving, and possibly accelerating drive train, before it can replace or augment the torque supplied by the electric traction system. Consequently a starting system is needed that has fast response and no tendency to overshoot the prevailing drive train speed, thus avoiding a shock torque in the drive shaft as the one-way clutch is engaged. Design of such a control system uses the transfer function relating throttle position to speed identified in chapter 4. When this is connected to the required control algorithm and throttle servo-system, the block diagram of figure 5.11 is produced.

For small variations in throttle demand, $G_T(z)$, reduces to $1/z$, producing a completely linear system which can be transformed to the w'-plane

for controller design. In order to produce an acceptably short synchronisation time for the engine a system rise time of $t_r = 0.5$ and critical damping is chosen as the design criteria. By equations (5.4-5.6) this suggests closed loop poles $w' = -4.71 \pm j4.71$. Figure 5.12(a) shows the uncompensated root locus for the system and figure 5.12(b) the system locus with the controller

$$g_c(w') = 0.012\left(\frac{w' + 1.1}{w'}\right) \tag{5.12}$$

With this controller the presence of the closed loop pole on the real axis modifies the system response so the performance criteria are not achieved with exactly the calculated imaginary pole locations given above. The simulated and experimental closed loop response of the system being shown in figure 5.13. On this diagram the experimental throttle trace shows the step rate non-linearity which was not included in the design. This is a consequence of the high gain needed to meet the fast system response requirement, and the large errors present at the beginning of the step demand. Actual engine response is delayed by the throttle step rate limitation causing it to lag behind the simulation, which does not allow for this effect. Despite this departure from design performance, the experimental result still exhibits satisfactory damping.

5.3.6 Engine Starting and Load Transfer

When required the warm engine will fire in typically 250 ms using the conventional electric starter motor, but there is a further delay whilst the engine accelerates up to the drive train speed. Inertia starting used in the HTV-1 project [Trummel and Burke, 1983] allowed the engine to be completely coupled into the drive train in 300 ms but the cost was the need for an additional clutch between the engine and the engine flywheel. A time analysis of the starting process on the rig is shown in figure 5.14. In this

experiment the motor was initially accelerating under load, as illustrated by the motor speed and torque traces. At time t=0.45 however, the computer receives the start command, immediately it turns on the ignition and engages the starter motor. At the same time the throttle is opened 9° and the computer then waits for the engine to fire. This is adjudged to happen when the engine speed passes 490 r.p.m. Above this speed the starter motor is turned off and the speed control algorithm is entered to run the engine up to the drive train speed. Synchronisation is deemed complete when the engine speed is within 45 r.p.m. of the drive train speed which in this case is achieved within 0.7 seconds of the original command to start. At this stage torque control is transferred to the engine which continues to accelerate the load. In figure 5.14 the slow rise time shown by engine torque is in fact false, since the trace represents the output of the highly filtered torque transducer. Total times for starting, speed synchronisation and transfer of load are consistently about 1 second as demonstrated by figure 5.14. Starting a cold and perhaps damp engine is still an unreliable feature of modern cars, consequently the software must be ready to cope with failure to start. In the event of the engine failing to start after five seconds the starter motor is disengaged, to allow battery recovery, before a second attempt is made. With a very hot engine better starting is often achieved with full throttle opening and this might be a useful strategy for the computer to adopt on the second attempt at starting if the engine had been operated very recently. Failure to start after perhaps five attempts would have to be treated as an error condition requiring the attention of the driver.

5.4 Quantisation Errors and Noise

Several effects cause errors in the digital control schemes described above. Taking a specific example of the engine torque control loop, six sources of error can be recognised:

1. Transducer noise
2. Measurement noise.
3. A/D quantisation and truncation
4. Parameter round off in the controller
5. Arithmetic errors in the controller
6. Quantisation of controller demand by the stepper motor throttle servo system.

Estimates of the magnitudes of all these error sources are included in Table 5.2, perhaps the most uncertain of these are the figures for measurement noise. Both the measured engine parameters, inlet manifold depression and speed, suffer from measurement noise but in disimilar ways.

Manifold pressure is measured by an analogue transducer, remote from the computer. As such the signal inevitably picks up noise, primarily from the power electronics, as it crosses the rig. Before reaching the ADC however, the manifold pressure signal is filtered electronically to reduce unwanted high frequency components and noise. Despite this filter it is still likely that measurable noise effects will reach the M68000. In order to assess the effect of noise from the traction system on the pressure transducer signal, two sets of measurements were taken, once during all electric operation with a stationary engine, and once whilst the rig was entirely inactive. As Table 5.2 shows, the effect of the motor is to virtually double the rms noise signal measured from the inlet manifold depression transducer. No attempt was made to isolate noise effects from the engine ignition system, which is likely to be a second major source on the rig, because with an operational engine it is difficult to

separate manifold pressure variations due valve openings from those due to noise.

In contrast the speed signal from the engine is digital consisting of a stream of pulses generated by the magnetic pick up mounted on the engine starter ring. These pulses are counted by a VIA chip and then read and the count reset at each interrupt. Noise becomes a problem in this system only if it is severe enough to cross the logic threshold voltages of the VIA, thus causing an erroneous pulse to be recorded. Some problems were encountered during the early development of the rig from erroneous pulses created in this way by noise from the power electronics. Subsequent liberal use of capacitors to decouple the power supplies to the magnetic probes has drastically reduced this problem, almost to zero, as indicated in Table 5.2.

Fairly comprehensive theory has been developed to assess the effects of random measurement noise on control systems. Nevertheless once digital transfer functions become involved it is particularly difficult to predict how noise will propagate through the system and according to Katz [Katz, 1981] no general theory exists to quantify this.

Even if the control system were to receive perfect measurements it has its own sources of internal error, largely concerned with the use of all integer arithmetic. On the M68000 all controller parameters use long integer type variables. As defined in the 'C' programming language, this variable type uses 32 bits for data storage. All calculations for control purposes are carried out to three decimal places by premultiplying the controller parameters and measurements by 1000. Taking into account the use of one bit for the sign, the 32 bit total word results in 2^{31} or about 2×10^9 being the largest number that can be stored in the system. Hence by preserving three decimal places this means that the largest decimal number that can be represented in any control calculation is about 2×10^6. Great care had to be taken

in designing all control routines to make sure that overflow errors could not occur as a result of this limitation. Given the truncation of all parameters to three decimal places, the precise error generated in the control algorithm depends on the way it is realised. There are various methods of realising any given digital controller or filter on a computer, and for this system the direct, series method has been adopted as represented by equation (5.9). Although this method result in a simple algorithm, it is more sensitive to arithmetic errors than some others, particularly when a high order controller is realised. In the hybrid vehicle control software, the fact that the controllers are all of low order vastly reduces the problem and makes the direct-series realisation a logical choice on the grounds of fast execution time.

Arithmetic errors were investigated experimentally by applying a random input of 1000 data pairs representing inlet manifold depression and speed to the engine torque control system. The output of the integer system was then compared with that from a full double precision floating point realisation of the same system. Although the floating point system is not immune to arithmetic errors itself, they will be much smaller than those from the integer system and can be ignored. Result show that in no case did the error exceed 0.5 in the final calculated demand to the throttle-servo. The importance of this last result is that it establishes quantisation by the stepper motor as by far the largest source of error in engine control systems. As a consequence of this fact there is no advantage to be gained by using any more accurate methods of either arithmetic or controller realisation. When the engine torque control system is operating the effect of stepper quantisation can be seen in the way the throttle position cycles between two values during steady state conditions, as clearly illustrated by figure 5.2(a).

	Error Type	Magnitude
Inlet Manifold depression	Transducer inaccuracy non-linearity and hysteresis	0.5% f.s.
	Temperature dependence	0.1% f.s.
	ADC Quantisation	1 L.S.B.=4.9 mV (0.45 mbar)
	Measurement noise	Motor ON 2.32 mbar (rms) Motor OFF 1.24 mbar (rms)
Speed	Quantisation	1 count=22.22 r.p.m.
	Measurement noise	Negligible
Stepper servo	Quantisation of demand signal	Resolution 0.9°
Arithmetic errors	Parameter round-off	All controller gains round off to 3 d.p.
	Algorithm calculation	Truncation to 3 d.p.
	Overflow	Maximum result 2×10^6

Table 5.2 Error Sources in the Engine Torque Control System

5.5 Model Reference Controller Design

Although the design process described in the previous sections does produce a fast well tuned control system, it has several disadvantages in the hybrid vehicle application. Foremost amongst these is that the design process is laborious, with possibly several cycles of parameter adjustment and simulation before a satisfactory result is obtained. In a production environment this would render the above method impractical if controller designs had to be tuned to the dynamic characteristics of individual vehicles. Furthermore at the end of a manual design process there is no guarantee that the chosen parameters are those which most closely achieve the specifications of damping

factor and rise time. One way of ensuring that the final parameters are the most suitable, and automating the whole design process, is to use a model reference technique. There are many advanced on-line model reference adaptive control techniques available, but in this section a simple off-line process using hill climbing optimisation methods will be applied to the design of the motor torque controller.

The principle of the model reference technique is to compare the output of the controlled plant, y_p, with that of a reference model, y_m, representing the desired closed loop system. For motor torque control this system must therefore have a rise time, t_r = 150 ms and damping factor $\xi = 0.707$. As before equations (5.4)-(5.6) provide the necessary closed loop pole locations for the equivalent second order system; now to be used as the reference model.

$$G_m(s) = \frac{\sigma^2 + \omega_d^2}{s^2 + 2\sigma s + (\sigma^2 + \omega_d^2)} \tag{5.13}$$

To quantify how closely the plant response matches that of the model a performance index, J, maybe defined:

$$J = \int_0^\infty |y_p - y_m| dt \tag{5.14}$$

J is conveniently calculated after a unit step has been applied to both systems. Clearly J cannot be calculated in a digital system but must be approximated by summing the difference term at each sampling interval:

$$S = \sum_{i=1}^{n} |y_{p(i)} - y_{m(i)}| \tag{5.15}$$

By considering only the first n terms it is assumed that the integral converges, which will be the case if the controlled plant has zero steady state error. Before the digital approach can be adopted, the reference model of equation

(5.13) must be discretised. Any discretisation method which preserves the two design specifications in the digital version is appropriate, and zero order hold equivalence is successful in this respect.

Having defined the performance index, S, and obtained the reference model, the remaining problem is to adjust the controller parameters until it is minimised. A simple way of doing this is described by Lin Luo [Lin Luo and Hill, 1986] and has been adapted to suit the present application.

The P+I controller as described by equation (5.9), has two adjustable parameters, namely g and k_i. This means that S is a function of g and k_i as illustrated by figure 5.15. Suppose initially S is evaluated at some arbitrary point A and at four surrounding points, labelled N, S, E and W on figure 5.15. One of these four points may well give a smaller value of S than A itself. If this is the case then that point is chosen as the new operating point and the process is repeated. Continuing in this way will eventually lead to the point B, which gives a smaller value of S than any point around it. Hence B is a minimum of the function $S(g, k_i)$ and should represent the control parameters giving the best possible match with the reference model. Before employing the method there are several practical points to consider. For example unless the function S has a single global minimum the parameter search may stop at a local minimum, giving perhaps dreadful controller performance. One possible way around this problem is to chose a number of well spaced starting points and to check that they all converge to the same solution. A second point is that the start point may not give a stable system, although Lin Luo and Hill counter this by stating that a stable operating point will always be reached.

5.5.1 Application of the Model Reference Technique to the Motor

One possible way of carrying out the optimisation process described above is to apply the control law directly to the motor, subsequently varying the parameters in real time in response to changing motor performance. However a problem which arises in applying the technique to the motor or engine, is how to apply the necessary step demands in the context of an operational vehicle. Furthermore control carried out at the start of the search might be exceedingly poor, hence there is a possibility of physical damage. An alternative solution is to use a model of the plant and adjust the the control law on that before applying it to the real system. Since this avoids both of the practical objections raised above, and since a model of the motor is available, this method is perfect for the present application. Using the direct second order model of the motor in the full field mode from Table 4.6 as an example, the design process outlined above gave the results displayed by figures 5.16 and 5.17. Figure 5.16 shows the locus of g and k_i as they converge on a single solution for four arbitrary start points. This single solution is encouraging since it indicates that a global rather than local minimum has been found. To achieve the result the step lengths are $\Delta g = 0.5$ and $\Delta k_i = 0.2$. In each case the optimised solution is given by $g = 8.5$ and $k_i = 1.4$ which agrees very closely with the parameter set found earlier by manual adjustment (see Table 5.1). Following the design, the response of the model and a simulation of the motor plus appropriate controller were compared as in figure 5.17. As might be expected this graph shows that it is impossible to precisely match the reference model by simply using a P+I controller; nevertheless a perfectly adequate design has been produced.

5.6 Discussion

All three controllers designed in this chapter have to cope with non-linear time varying systems. Gain and dynamics are a function of operating point in both the engine and motor torque control loops. Nevertheless satisfactory control has been achieved without recourse to sophisticated control algorithms or large numbers of parameters for different operating conditions.

Engine torque control achieves peak performance at the design speed of 2000 r.p.m., with system gain decreasing above that speed and increasing below it. At the higher speeds the falling gain causes an increase in rise time, whereas at lower speeds the increased gain causes a slight overshoot. Nevertheless even in this latter case system performance is not so seriously degraded as to be unacceptable. More importantly examination of the root locus based on the manifold filling delay for 1000 r.p.m. shows that system stability is in no doubt down to this speed; representing the lowest operational speed for the engine. To counter this problem a slight modification to the engine torque control loop would be to make controller gain speed dependent, thus largely cancelling out variations in plant gain and maintaining optimum system performance across the whole speed range. Such a philosophy was adopted for torque control in the HTV-1 [Somuah et al, 1983]. This vehicle also based its torque control on the an indirect measurement based on inlet manifold depression and speed, but included a variable gain in the control loop. As explained in the reference the controller was a microprocessor based implementation of a lead-lag network.

Turning to the motor, a single control algorithm is possible, despite the three operating modes. In this case the algorithm developed for the field boost mode can operate the motor across its entire speed range and all operating modes. Unlike the engine however, one controller cannot really cope with the wide range of system gains, and inevitably the response under

certain conditions is too slow to be acceptable. In particular the field boost controller will produce a very long rise time in the field weakening mode due to the very low plant gains it encounters. Consequently it is worthwhile adopting the three sets of controller parameters as described in the design.

Of the three control systems designed in this chapter, the engine speed control loop showed the greatest discrepancy between between measured performance and theoretical simulation. Apart from the enhanced effect of the step rate non-linearity in the throttle servo-system, the analysis in chapter 3 clearly showed that the transfer function relating throttle angle to speed was likely to be unreliable, since all coefficients are themselves functions of speed. This means that the controller is working against a continually varying system as it accelerates the engine. Once this fact is taken in to account the relatively small difference between experiment and the simple linear simulation is quite surprising.

The hill climbing technique introduced in section 5.5 offers an alternative method of applying the pole-placement technique in all three control systems dealt with earlier in this chapter. It would have clear advantages if it proved necessary to produce an individual controller for each hybrid vehicle in a fleet. Although this latter scenario is perhaps unlikely, it all depends on whether or not mass produced engines and motors have sufficiently consistent dynamics to make one controller design generally applicable. Whether or not the hill climb method is used for initial controller design, it nevertheless opens up the possibility of re-tuning controller parameters to changing drive line characteristics over the life of a vehicle. A possible approach to this problem is discussed in chapter 8.

Obviously the main incentive for shutting down the engine is to save fuel otherwise wasted during idling. Although this strategy does undoubtedly save fuel, the cost is some delay in the availability of engine torque however

small. In addition there must be some fuel penalty associated with starting the engine from rest and accelerating it up to the prevailing drive train speed before it can supply additional torque. This fuel loss is needed to replace the kinetic energy of the engine lost when it is shut down. To try to calculate the likely amount of fuel lost in this way, Beachley and Frank [Beachley and Frank, 1981] quote a typical engine inertia of 0.068 kgm^2, implying an energy loss of 0.933 Wh when such an engine is shut down from 3000 r.p.m. If fuel used is based on this figure alone the result is an unrealistically low estimate since some of the kinetic energy is replaced by the starter motor the next time the engine is used. This energy will later have to be made up by the engine via the inefficient alternator battery combination. As previously mentioned an operational hybrid might well crank the engine using a clutch and flywheel combination on the grounds that it provides a faster start. This method uses about 0.104 Wh according to Beachley and Frank, a figure slightly higher than that simply needed to replace engine kinetic energy, because of friction losses in the clutch. As a result of these losses it is uneconomical to shut the engine down for a very short space of time. Nevertheless Volkswagen considered potential savings due to fuel off at idle and overrun sufficient to include the feature in their Formel E range of cars [Schmidt, 1981]. Results stated in this paper suggest that stopping the engine for periods of less than 5 seconds is uneconomic, however fuel savings in normal urban driving of 30% are reported for the otherwise conventional i.c. engine vehicle.

Experiments using the balance to measure fuel consumption show that for an idling speed of about 800 r.p.m., the 1100 cc Ford engine on the rig uses 0.1206 grams of petrol per second. Since the measured density of the petrol was 0.4432 g/ml this equates to a total consumption of 979 ml/hour. Assuming that the Volkswagen engines have similar fuel consumption to the rig engine, then 1.36 ml of fuel is used in 5 seconds which, if they took no

other factors in to account, must represent Volkswagen's assessment of the total fuel penalty per start. Such a small amount of fuel is very difficult to measure on the rig, not because the balance does not have the necessary sensitivity, but because of the unevenness of fuel flow out of the tank. Further complications arise because there is no way of ensuring that the total amount of fuel in the supply line, fuel pump and carburetter is the same at the start and end of any calibration experiment. Despite anticipating these difficulties an experiment was carried on the rig to try to measure fuel use per start. The experimental method was to operate the engine on a 50% duty cycle for 128 seconds. During the first test the engine spent 64 seconds stationary then started, accelerated to 2000 r.p.m. and maintained that speed for the remaining 64 seconds. Subsequent tests repeatedly doubled the number of starts, until the practical limit of 16 starts was reached. This corresponded to each on/off cycle lasting 8 seconds. Using the digital balance to evaluate fuel use gave the results in Table 5.3.

No. of Starts	Fuel Used (g)
1	14.5
2	14.8
4	13.3
8	17.0
16	15.0

Table 5.3 Engine Start Tests

If it is reasonable to equate Volkswagen's 5 second economy rule directly to fuel used, then 16 starts would use 9 g of fuel which would have been easily detected by this experiment, had it occurred. Clearly some of

the apparent discrepancy might be explained by loss of battery energy during these experiments which the engine had yet to make up, however it seems unlikely that this is entirely responsible.

Whatever the fuel penalty associated with starting the engine is in reality, there can be no doubt that shutting down an inactive engine in the hybrid application is essential. Use of the motor means that inactive periods for the engine are much greater than in the conventional vehicle. Nevertheless on/off engine operation does pose a number of additional questions such as adverse effect on engine life and increased emissions. Initial test results from the HTV-1 [Trummel and Burke, 1983], suggested that engine life was not reduced by such action but that further work was needed to prove it beyond doubt. Emissions were in fact found to be a problem with the Volkswagen starting system in the HTV-1 project, possibly because, to ensure good starting, Volkswagen arranged for fuel to be sprayed in to the inlet manifold, resulting in a rich mixture. Eventually the system was disabled in the HTV-1 application as soon as engine cooling water temperature reached 95°F. No mention is made by Schmidt as to whether this spray system was retained by Volkswagen for the Formel E application; if it was retained then this might go some way towards explaining why the 5 second economy rule could not be confirmed on the rig. If the starting system used on the rig were adopted consideration would have to be given not only to the increased wear on the engine but the reduced lifetime of the starter motor.

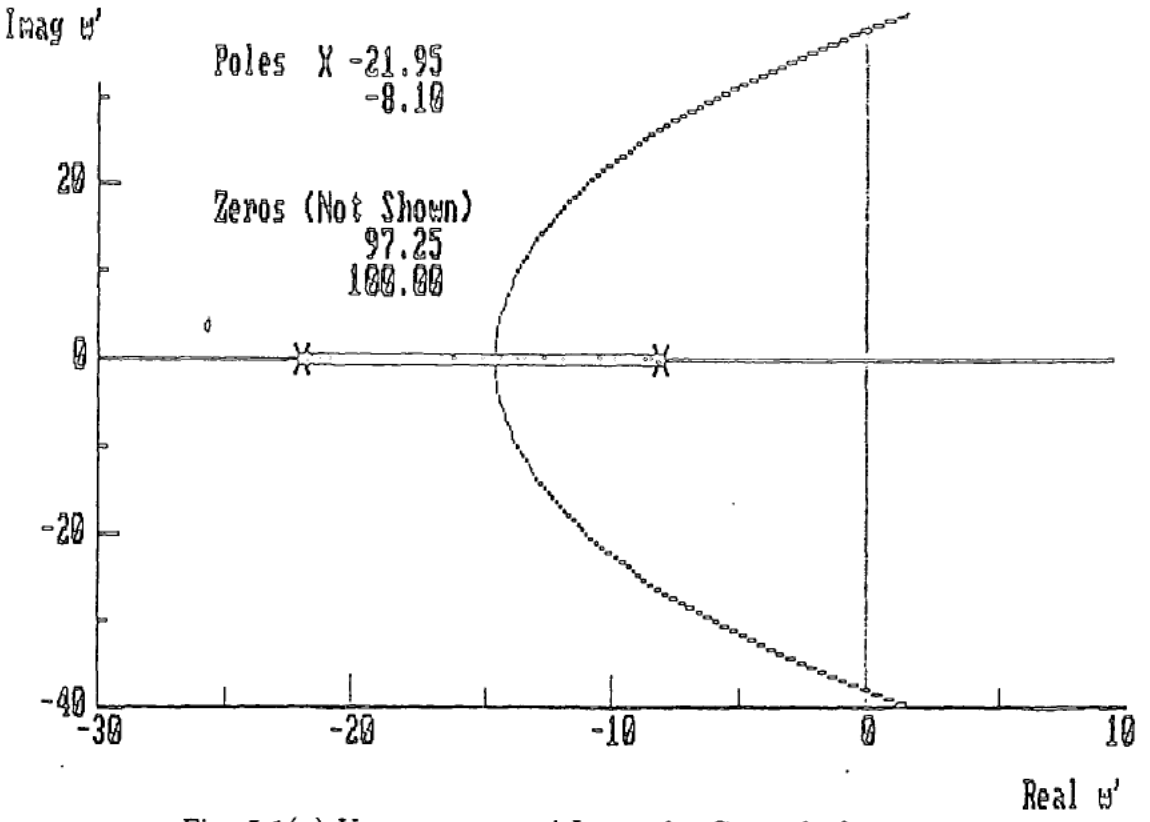

Fig. 5.1(a) Uncompensated Locus for Control of Engine Torque at 2000 r.p.m.

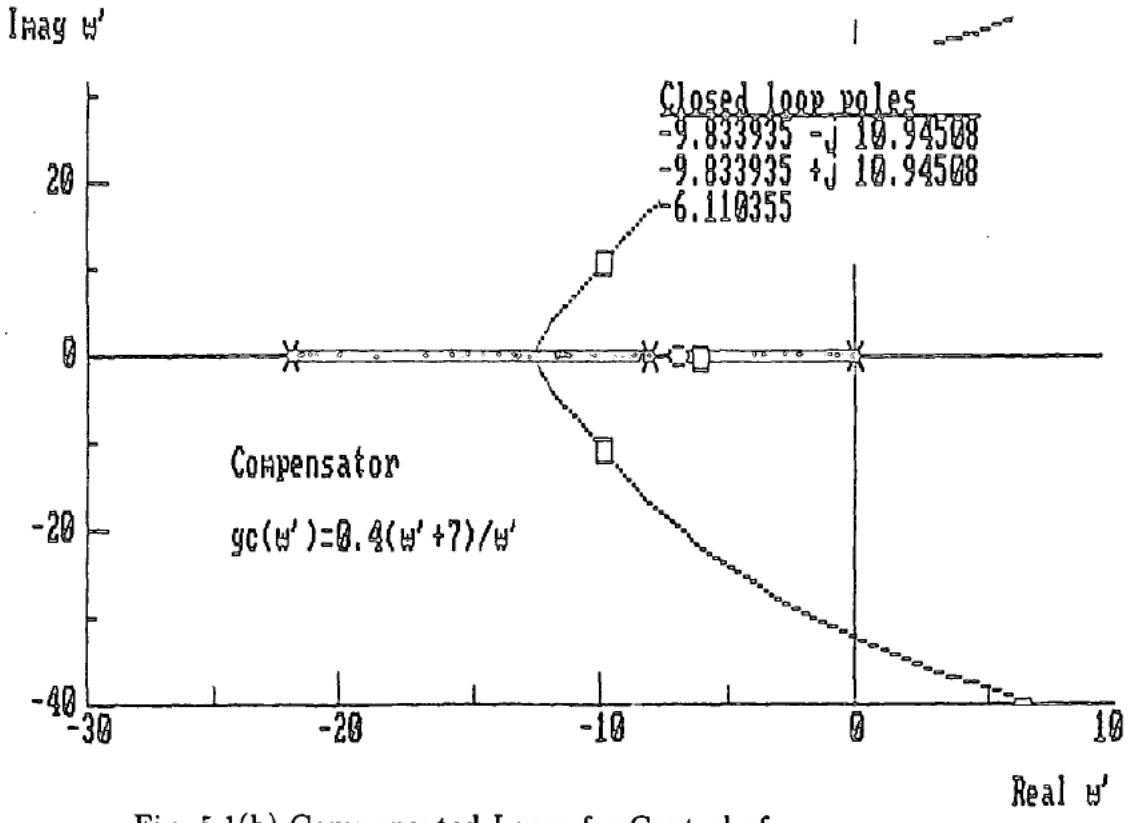

Fig. 5.1(b) Compensated Locus for Control of Engine Torque at 2000 r.p.m.

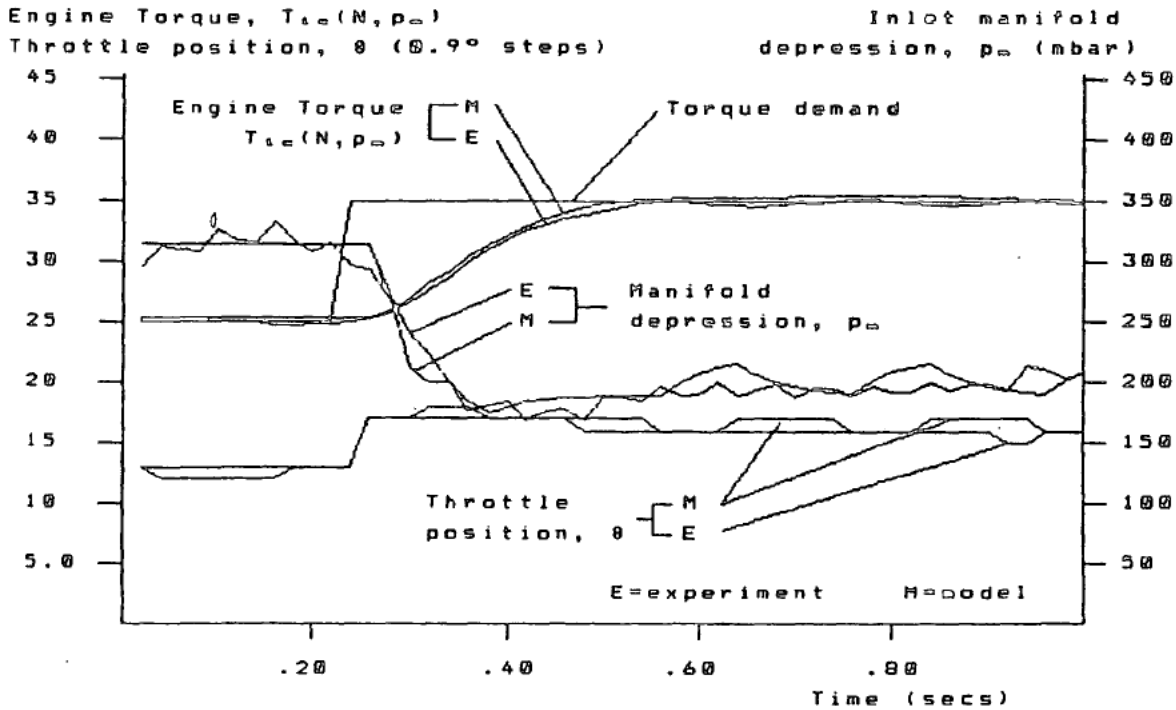

Fig. 5.2(a) Simulated and Experimental Performance of the Engine Torque Control System at 2000 r.p.m.

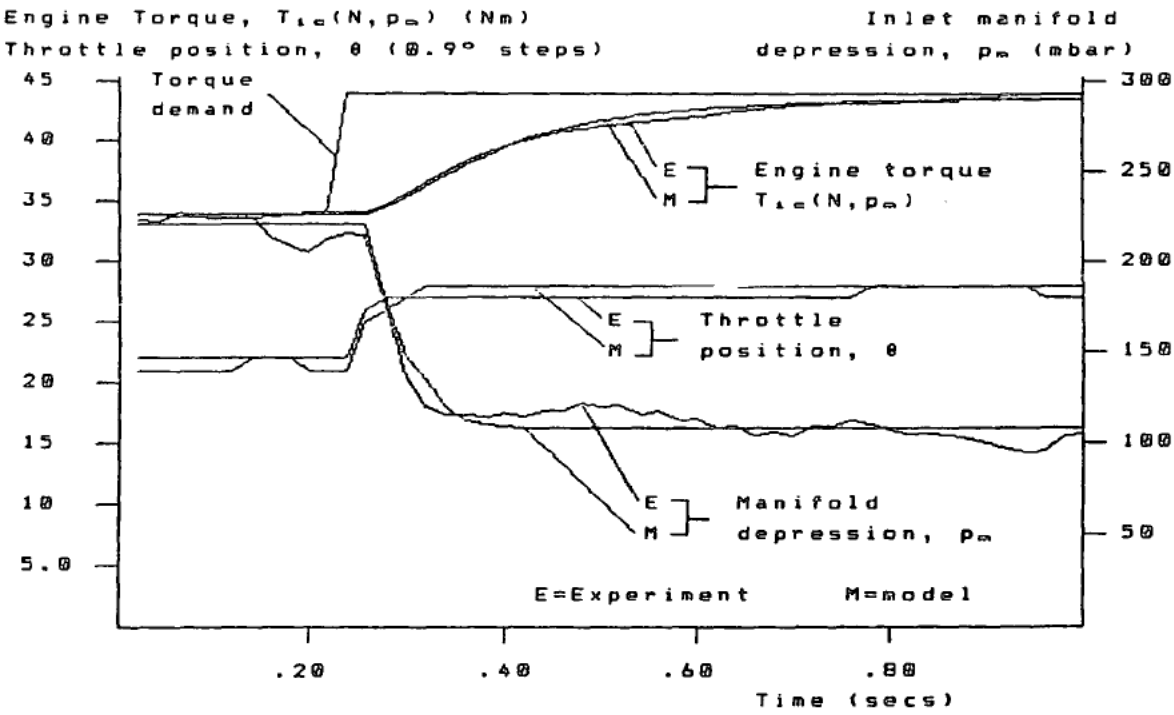

Fig. 5.2(b) Simulated and Experimental Performance of the Engine Torque Control System at 3000 r.p.m.

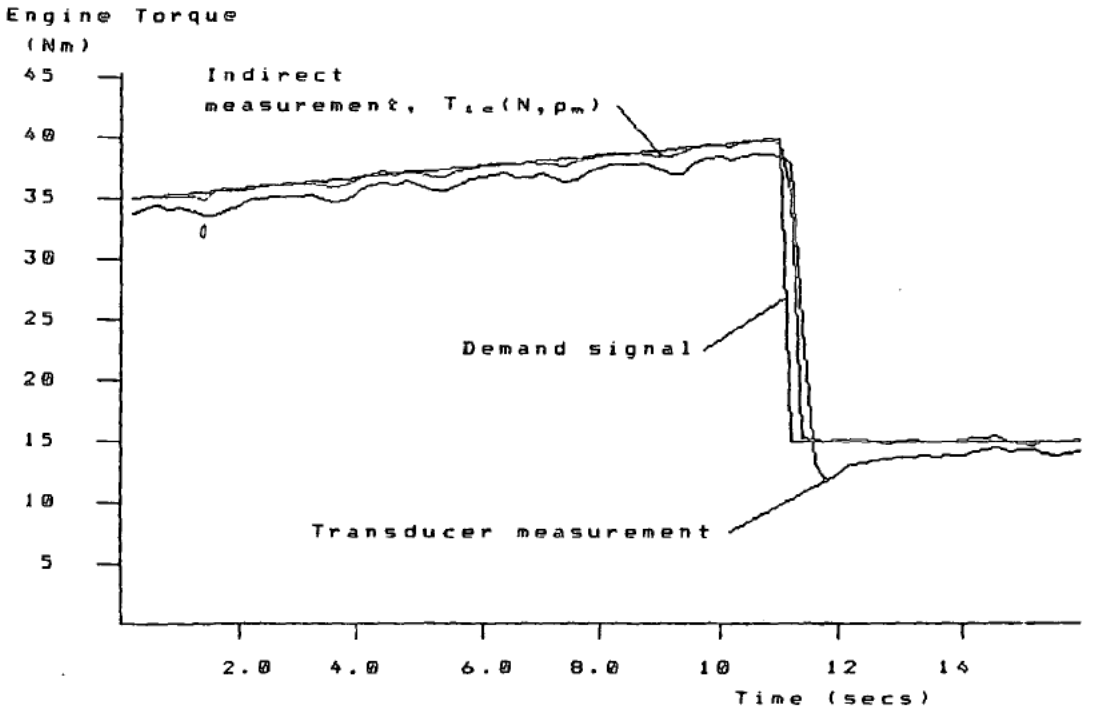

Fig. 5.3 Engine Torque Control System: Comparison of Indirect and Direct Torque Measurements During a Large Step Disturbance

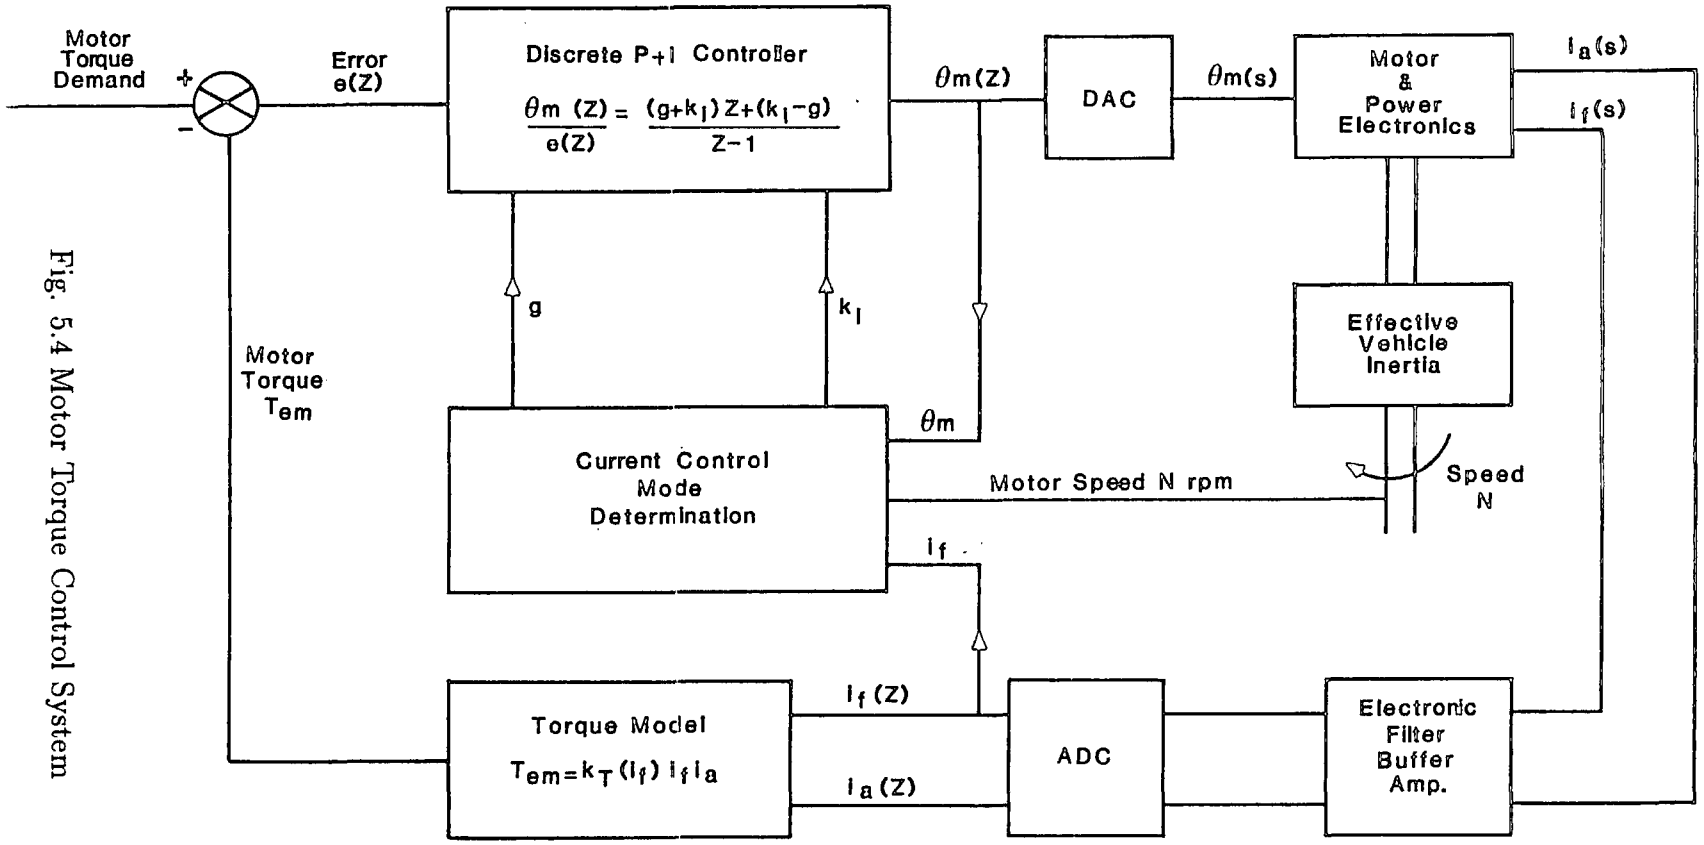

Fig. 5.4 Motor Torque Control System

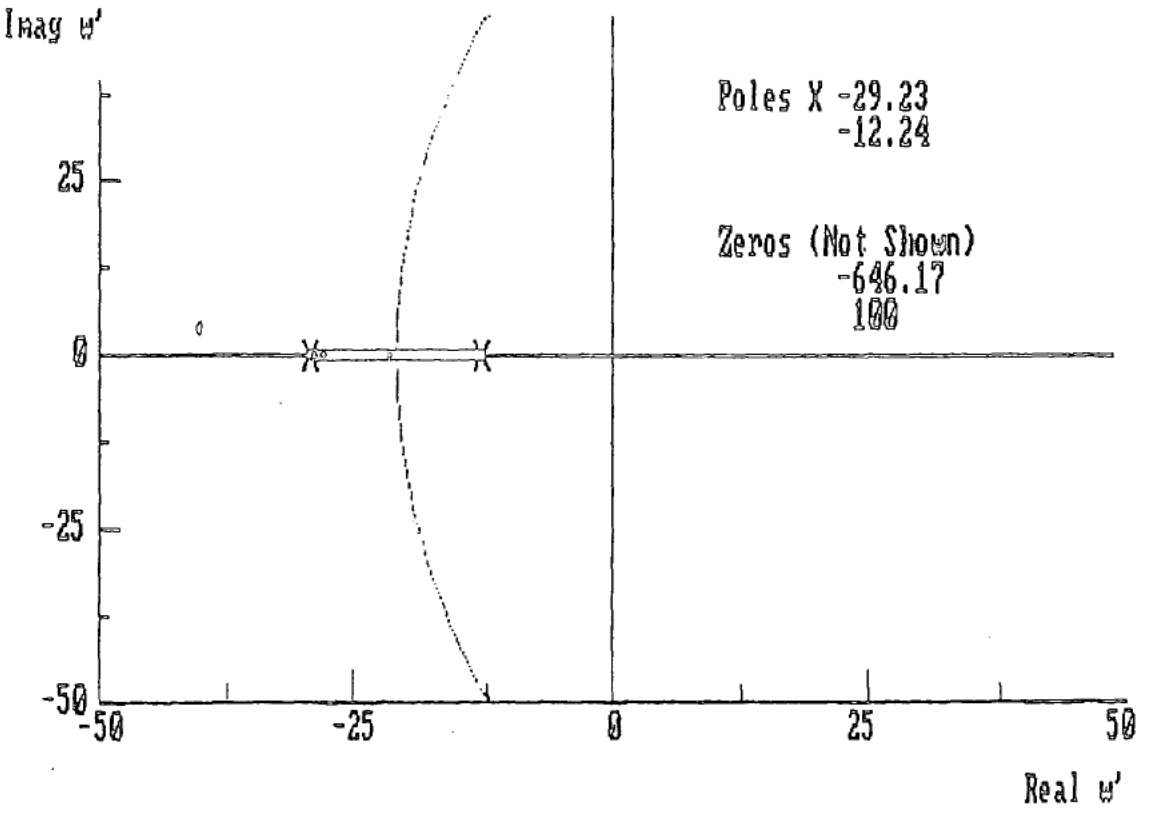

Fig. 5.5(a) Uncompensated Locus for Motor Torque Control in the Field Boost Mode

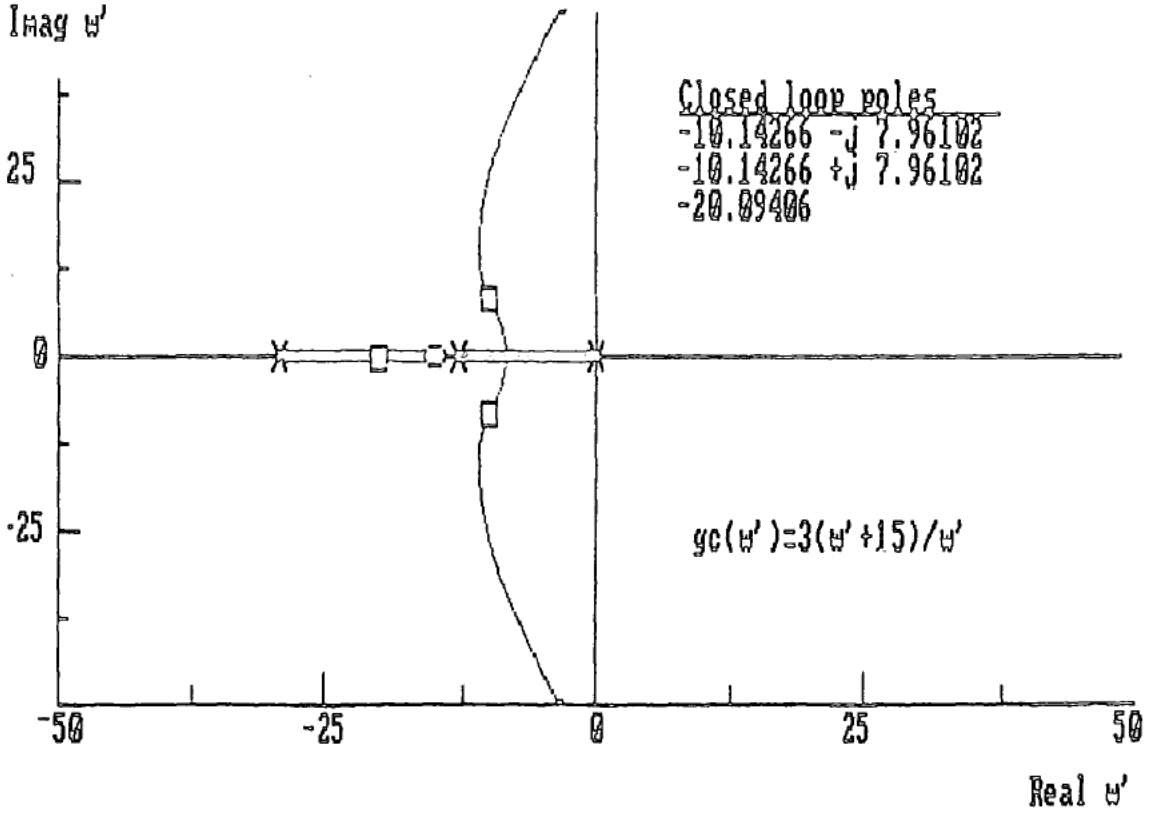

Fig. 5.5(b) Compensated Locus for Motor Torque Control in the Field Boost Mode

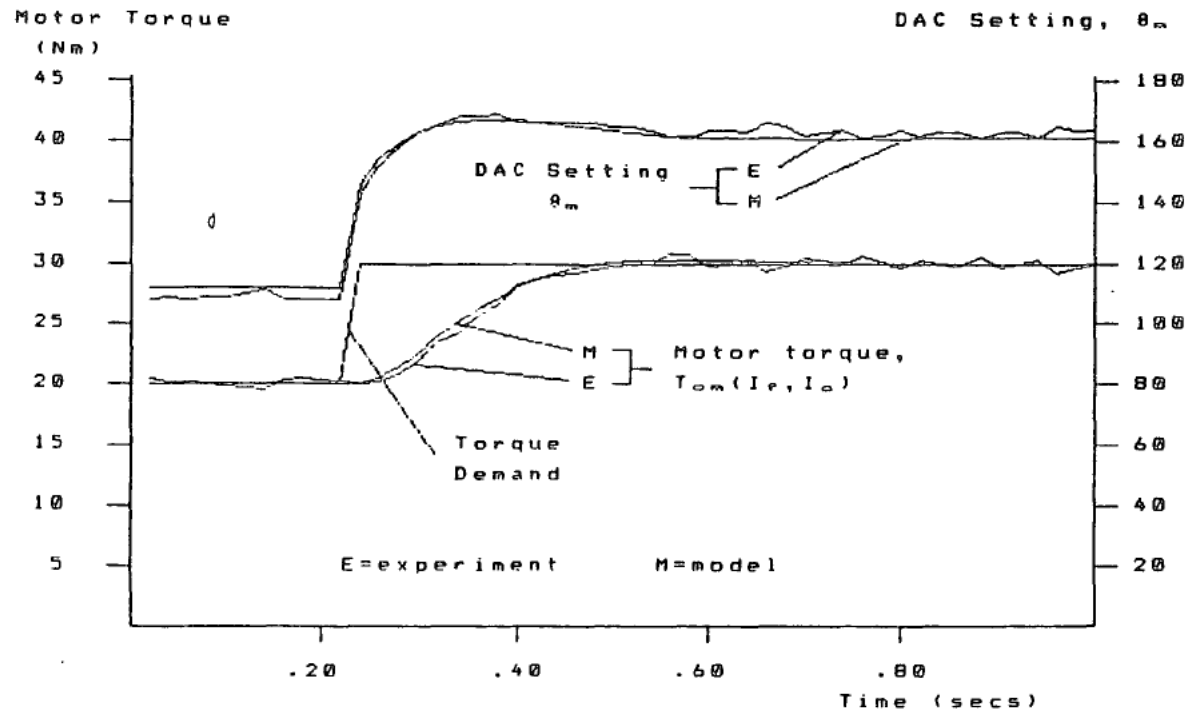

Fig. 5.6(a) Motor Torque Control Test in the Field Boost Mode at Design Conditions

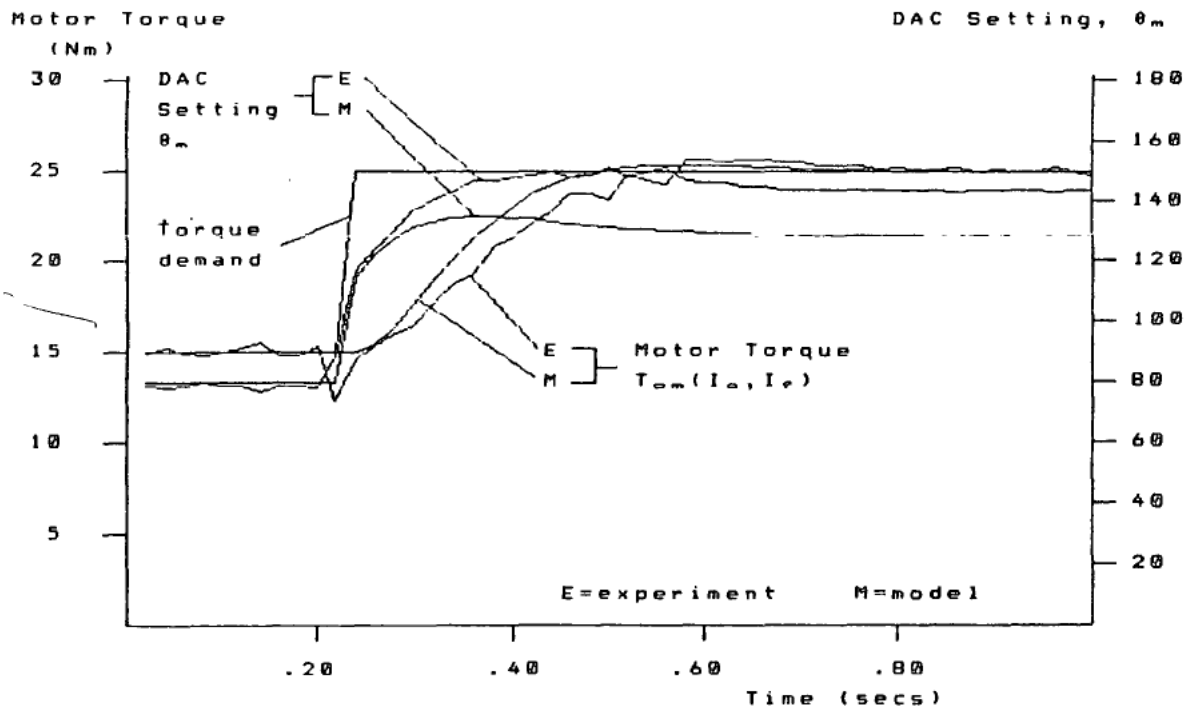

Fig. 5.6(b) Motor Torque Control Test in the Field Boost Mode Away From Design Conditions

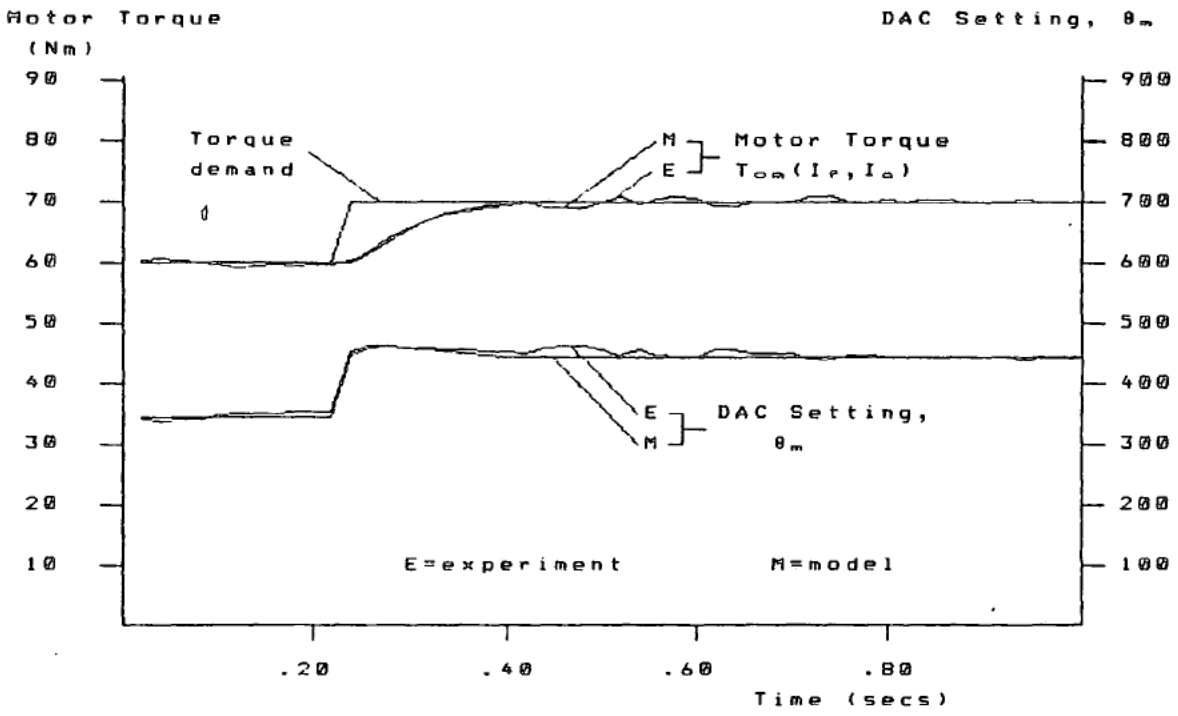

Fig. 5.7 Motor Torque Control Test in the Full Field Mode

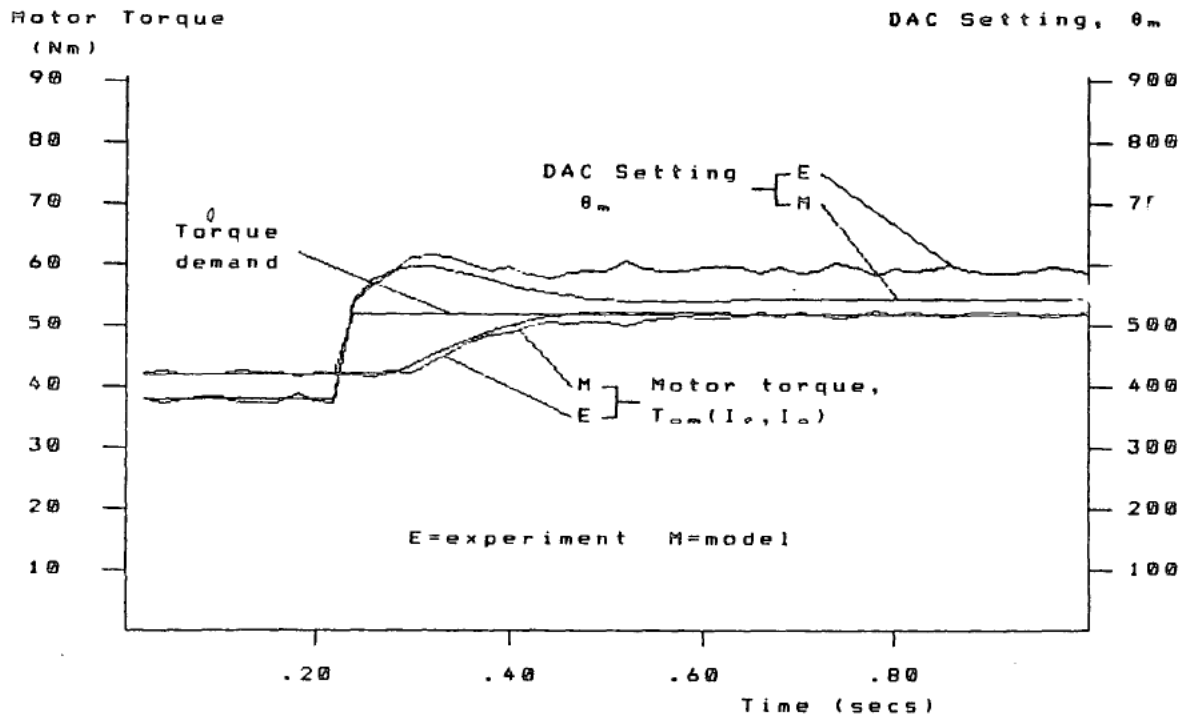

Fig. 5.8(a) Motor Torque Control Test in the Field Weakening Mode at Design Conditions

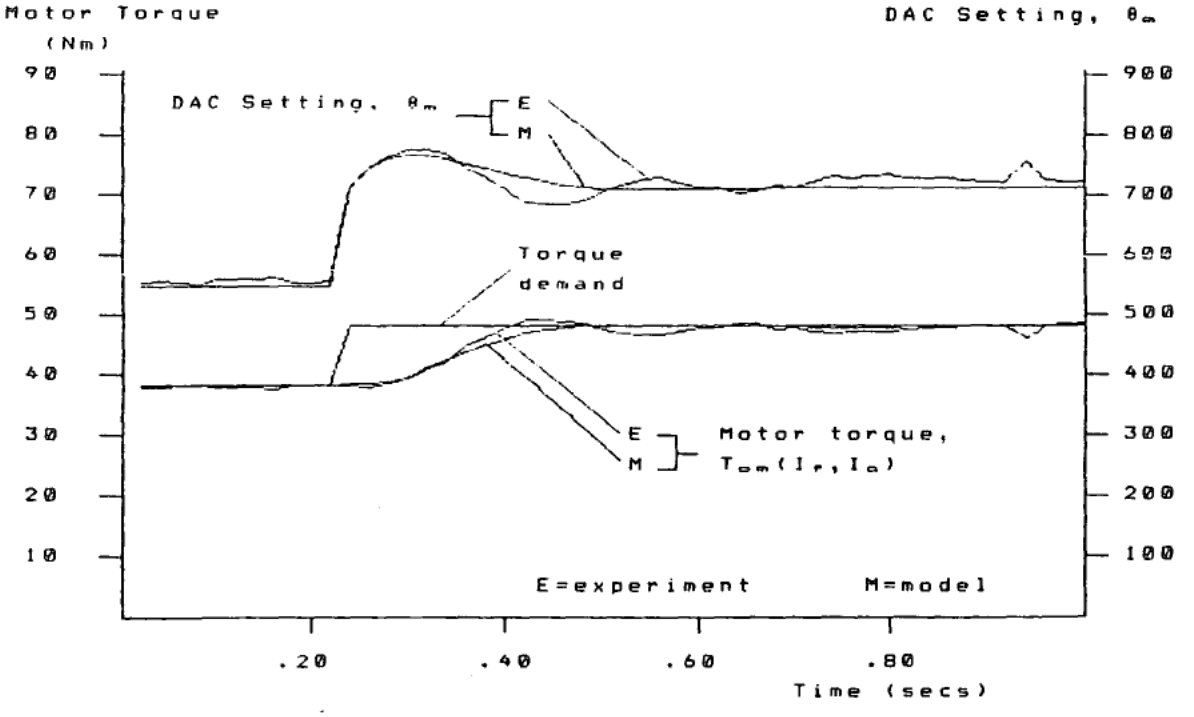

Fig. 5.8(b) Motor Torque Control Test in the Field Weakening Mode Away From Design Conditions

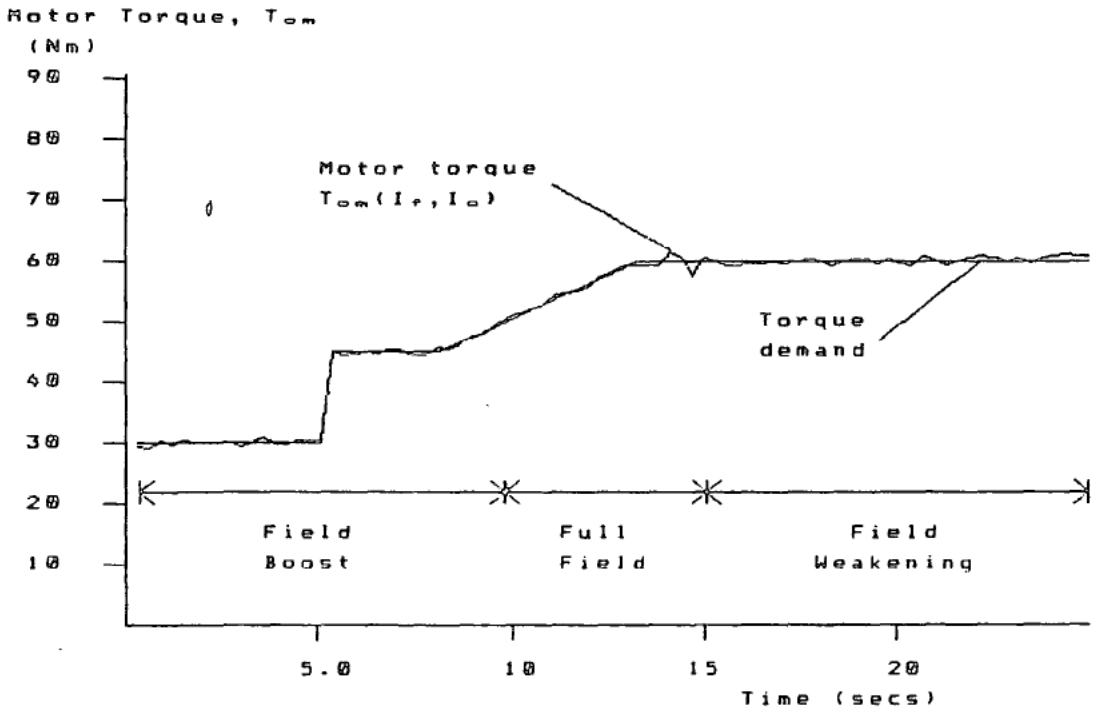

Fig. 5.9 Motor Torque Control Demonstrating Satisfactory Transition Between Operating Modes

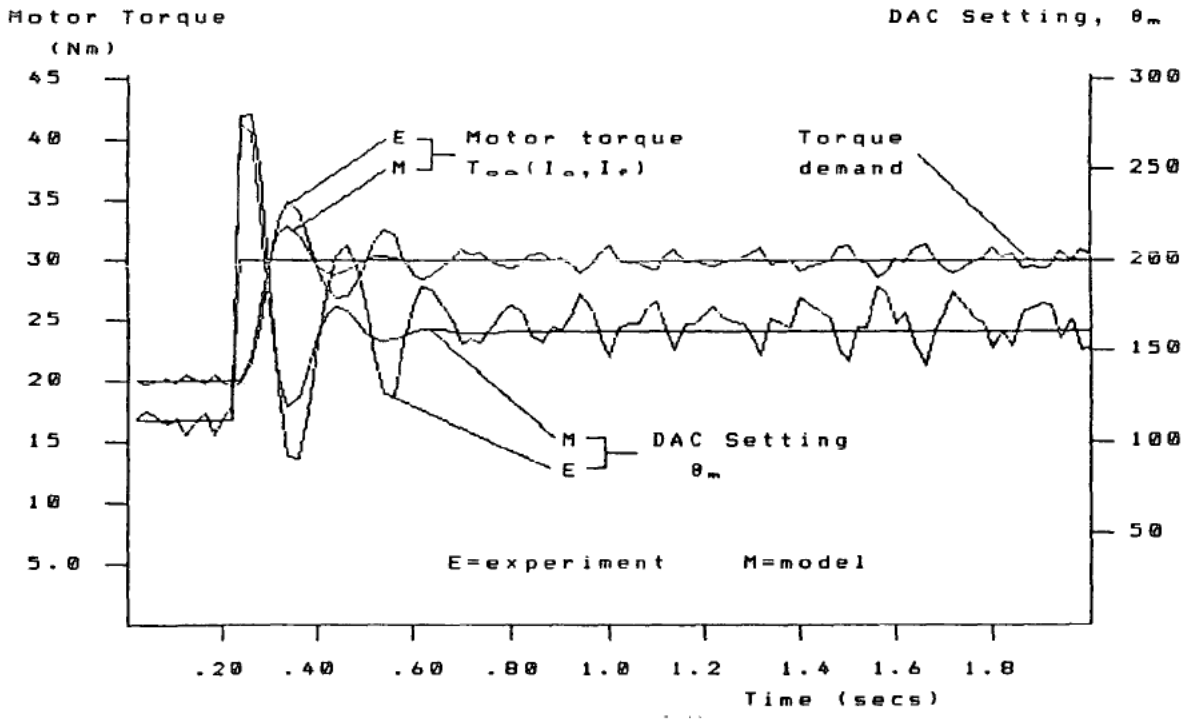

Fig. 5.10 Motor Torque Control: Effect of Using Field Weakening Gains in the Field Boost Mode

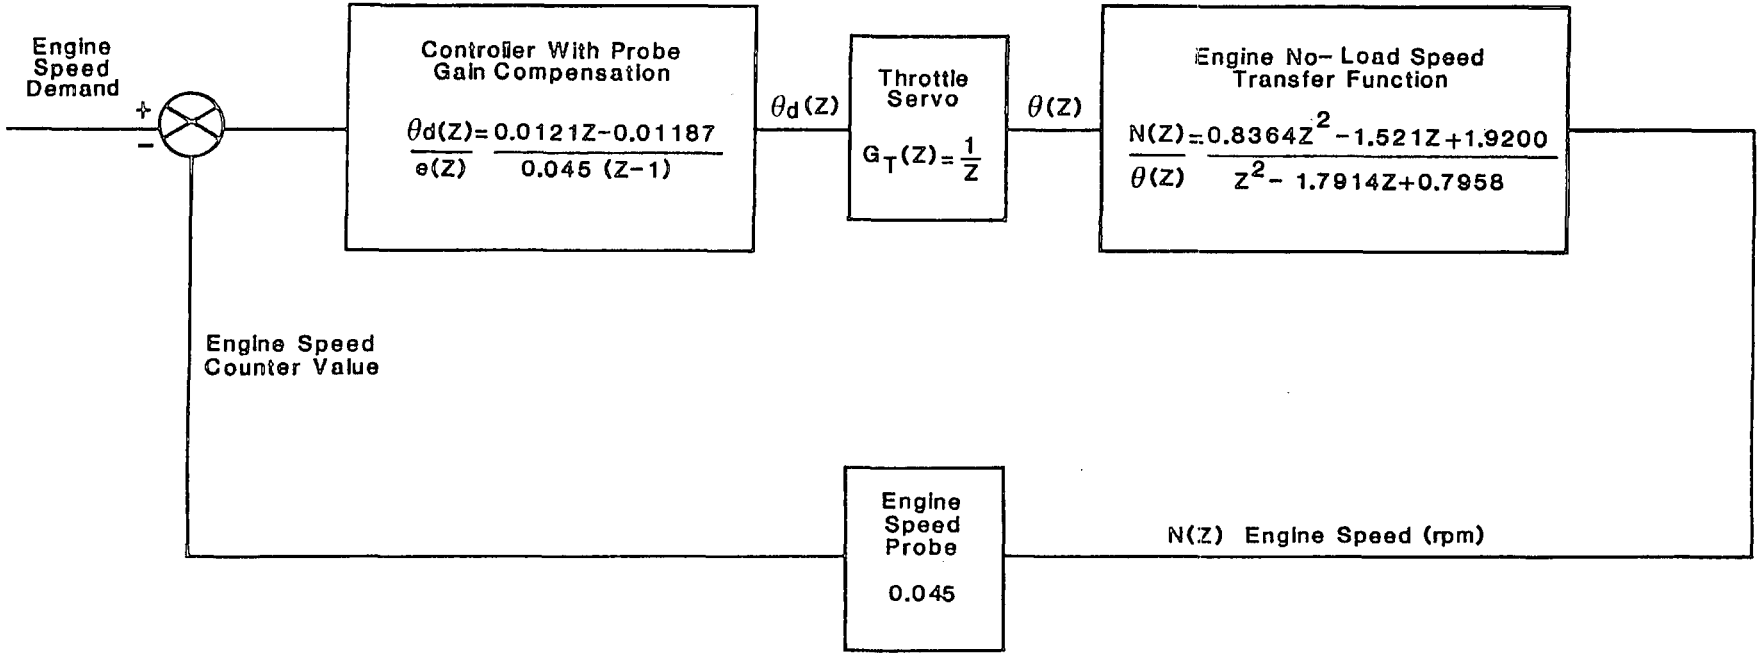

Fig. 5.11 Engine Speed Control Block Diagram

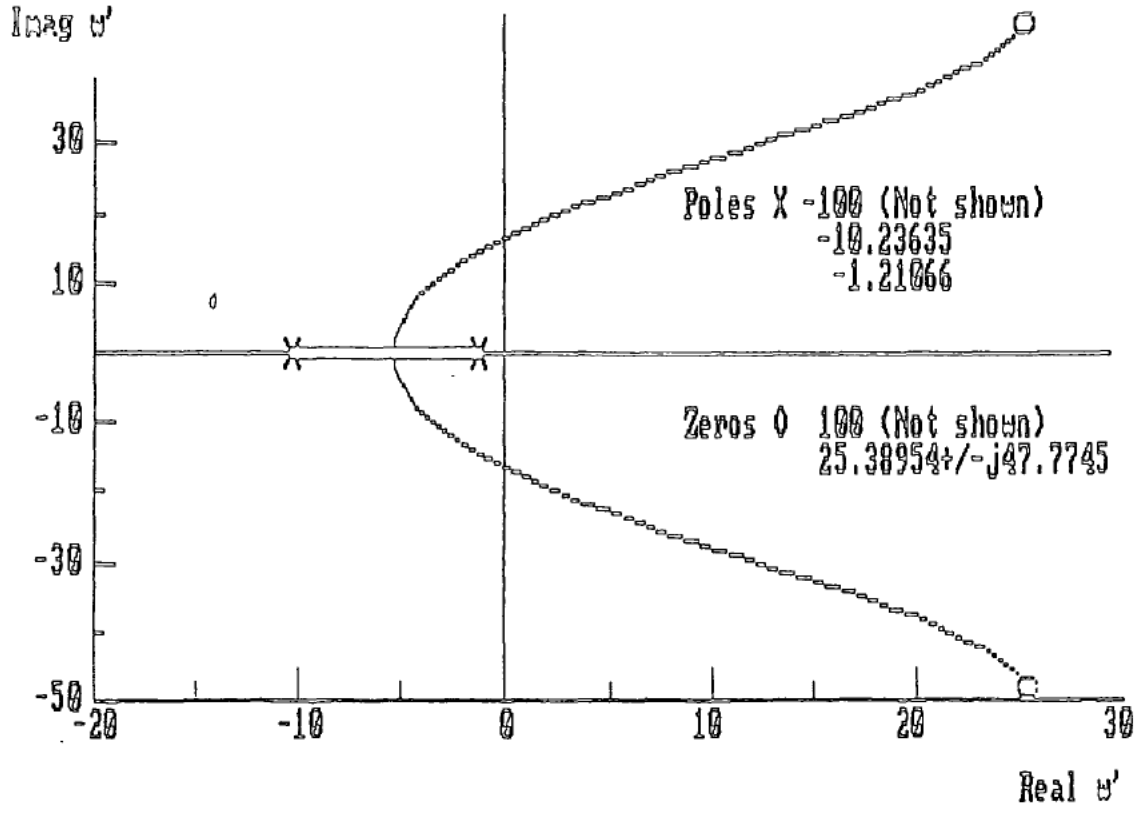

Fig. 5.12(a) Uncompensated Locus for Control of Engine Speed on No-Load

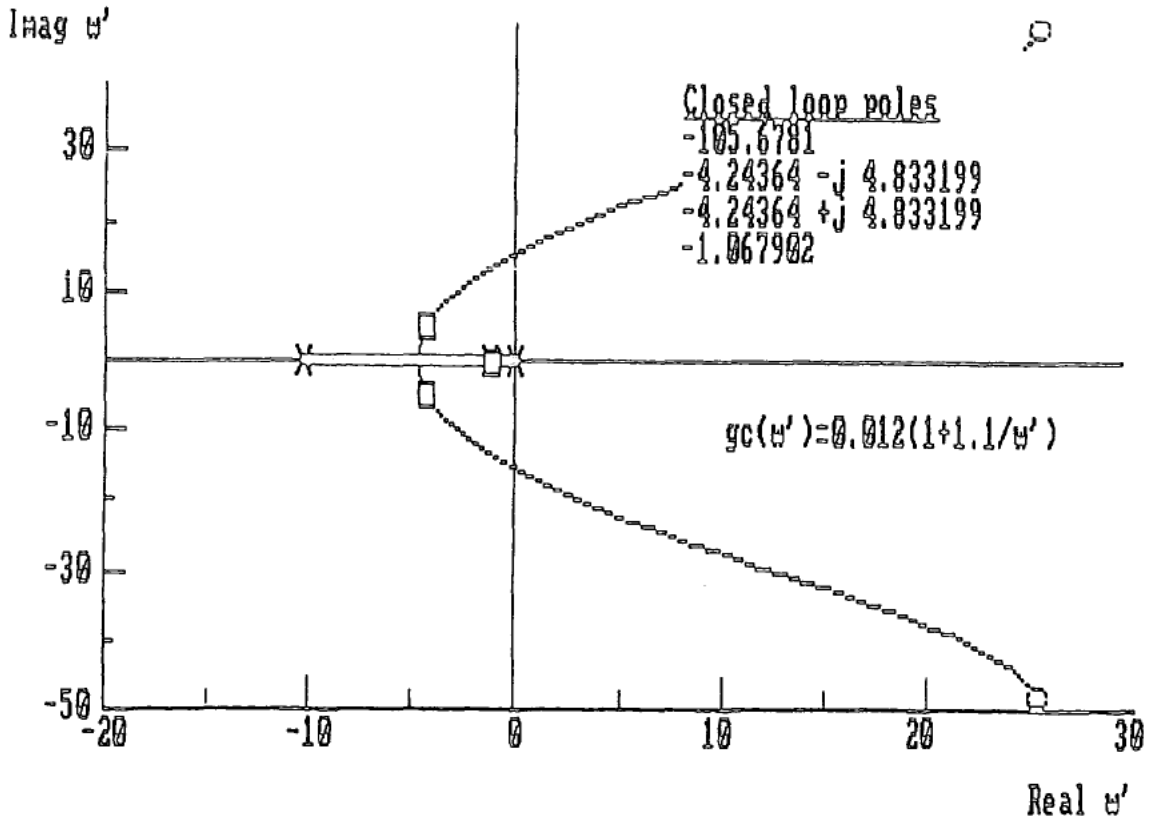

Fig. 5.12(b) Compensated Locus for Control of Engine Speed on No-Load

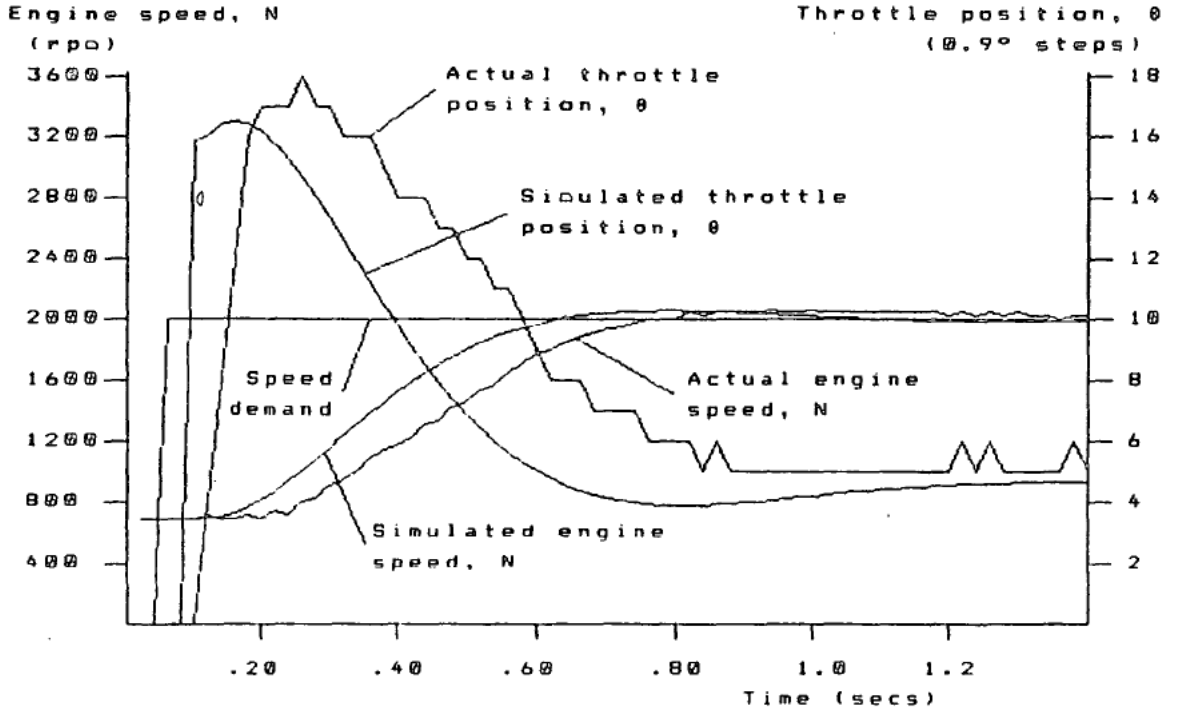

Fig. 5.13 Step Test for the Engine Speed Control System

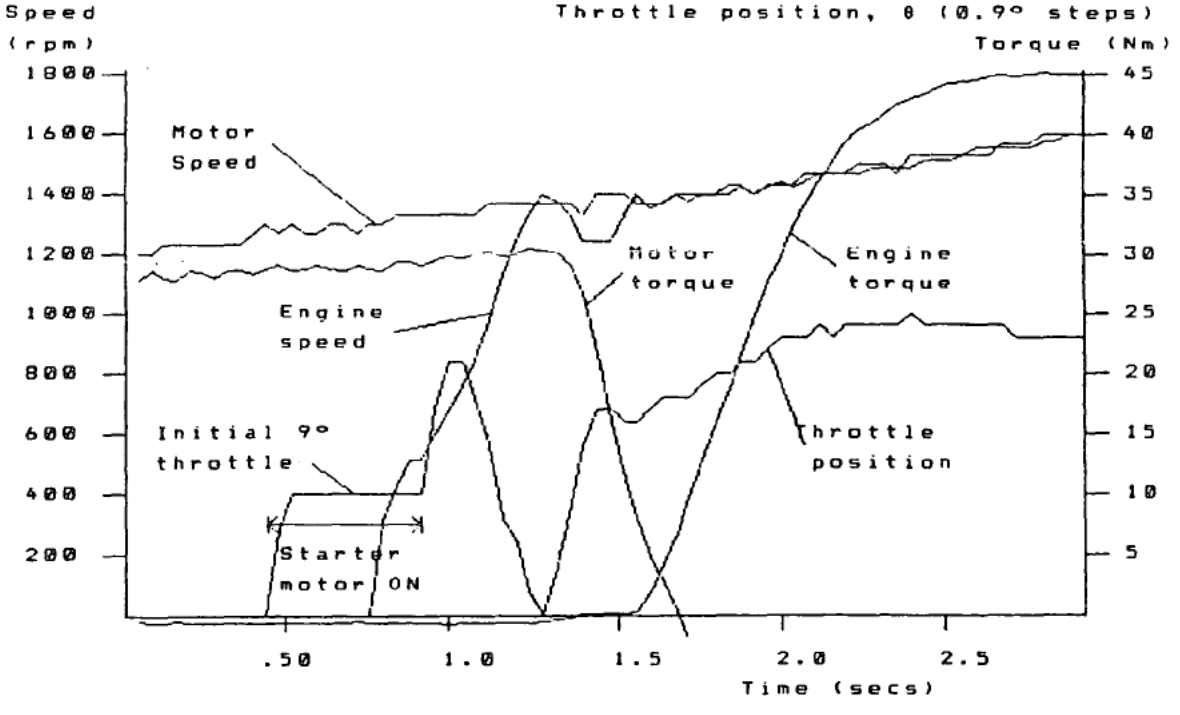

Fig. 5.14 Analysis of the Engine Starting and Load Transfer Process

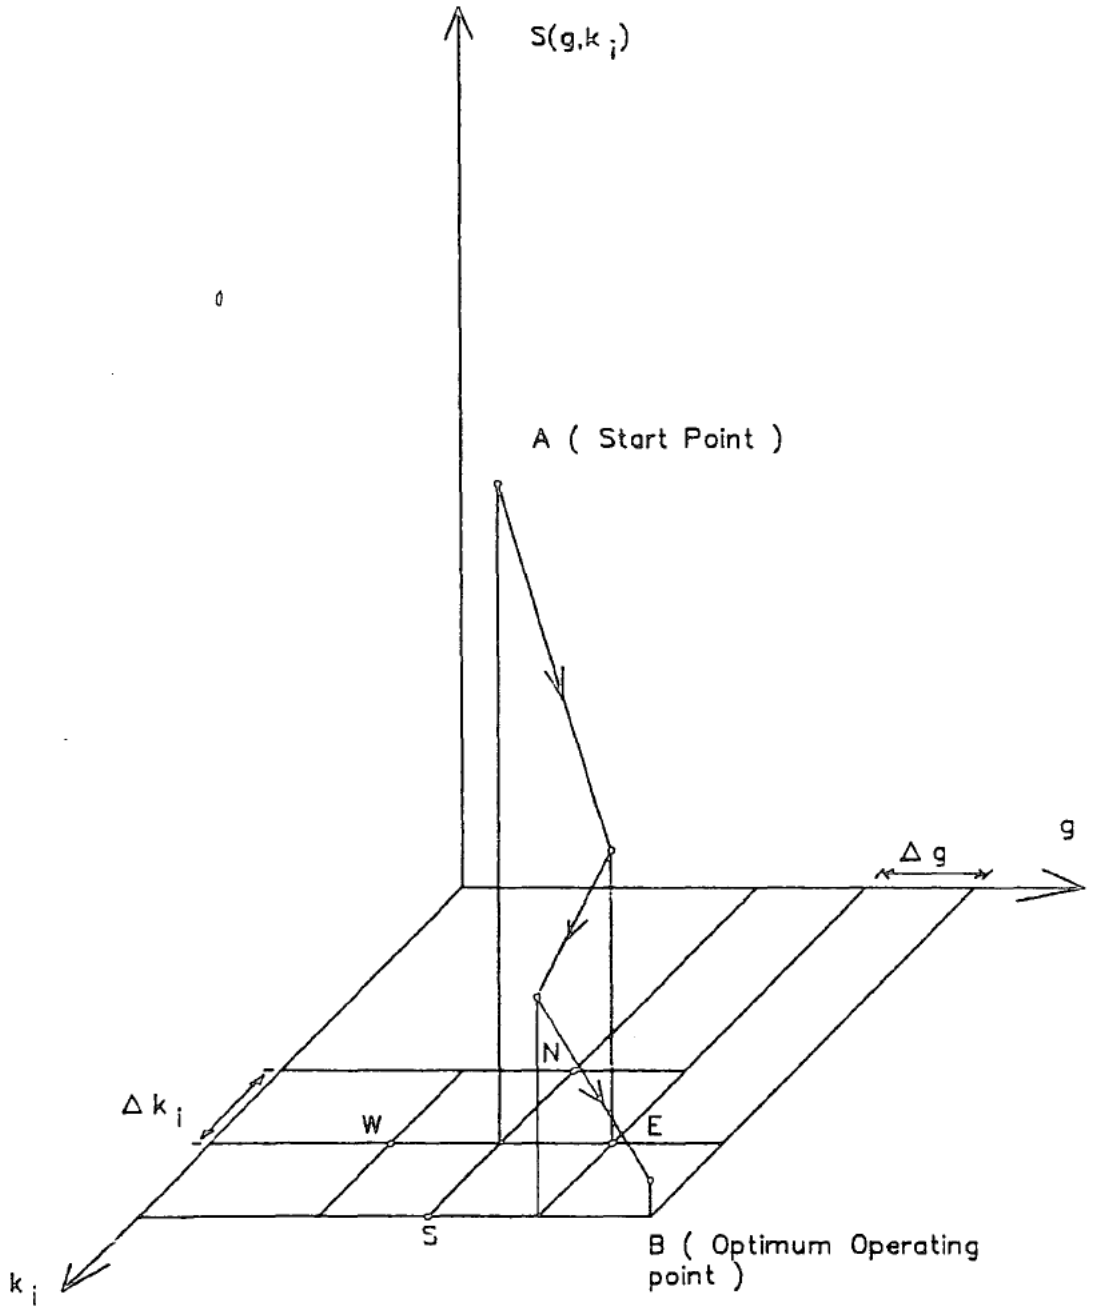

Fig. 5.15 Hill Climbing Control Design Process

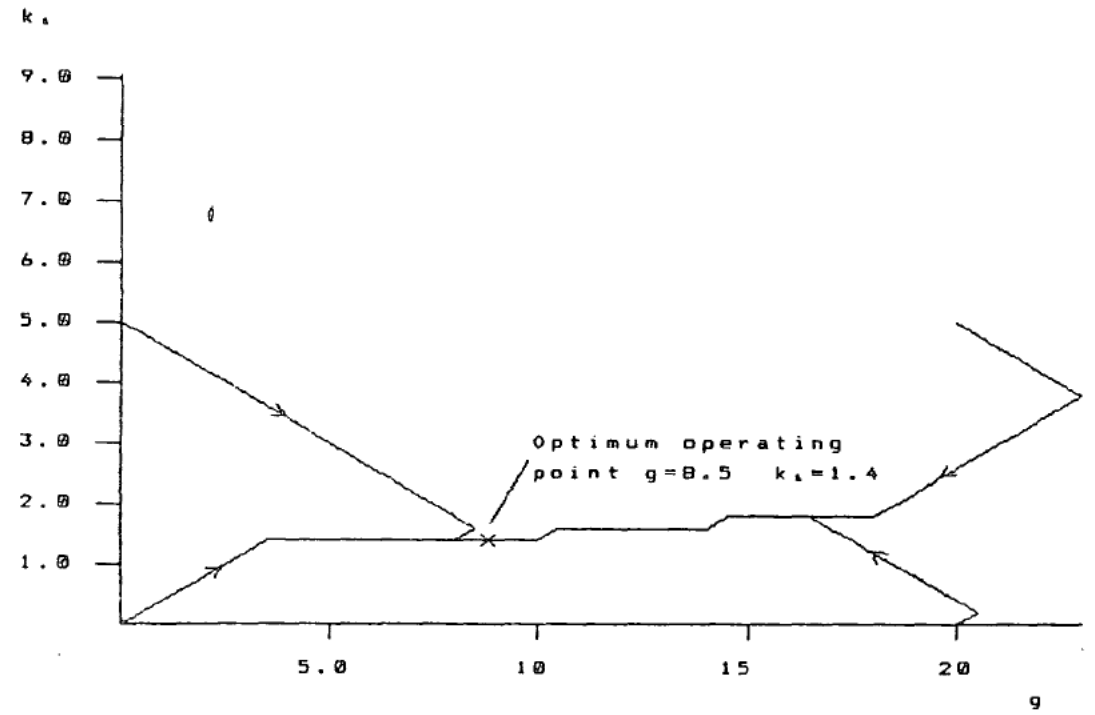

Fig. 5.16 Locus of Control Parameters During a Hill Climb Design Procedure

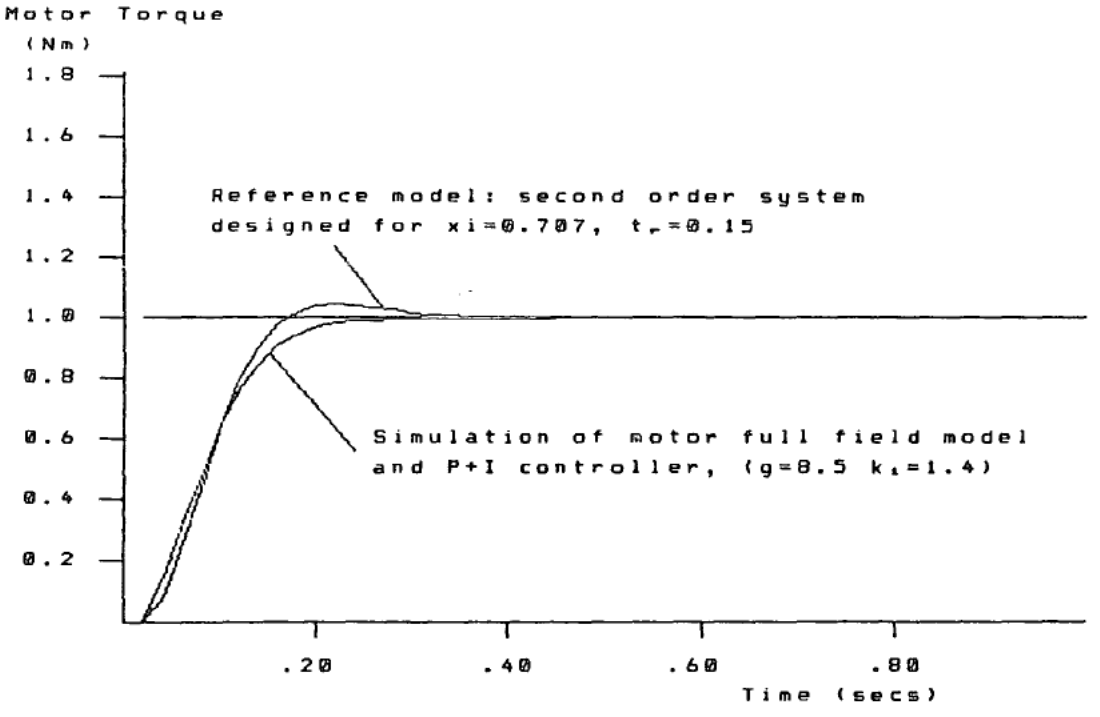

Fig. 5.17 Comparison of Reference Model Performance and Control System Designed by the Hill Climb Method

CHAPTER 6

THE AUTOMATED GEAR CHANGING SYSTEM

As mentioned in chapter 2, the hybrid vehicle rig includes a conventional 4 speed manual transmission between the engine/motor and the flywheel. For reasons of efficiency some form of variable transmission is essential to the overall vehicle control strategy. This need prompted the development of a pneumatic actuation system which allows the M68000 computer to change gear automatically as required [Masding and Bumby, 1988 (a)]. In this chapter the automatic gear changing system will be described and results presented to illustrate system performance. Before embarking on this description however, it is useful to consider why the hybrid vehicle requires a variable transmission at all, and what other transmission systems are available as possible alternatives to the conventional manual gearbox.

6.1 The Case for a Variable Transmission in Road Vehicles

The need for a variable transmission in a hybrid vehicle is a consequence of the power characteristics of both the engine and motor. In chapter 1 it was mentioned that optimum control of a hybrid would restrict use of the engine to the high efficiency region of its characteristic, which tends to be at relatively high load and low speed, as indicated by figure 1.3. Such operation can not be achieved if only a fixed ratio transmission is available. Figure 1.3 demonstrates this fact by showing the level road load seen by the engine when operating in a fixed gear. It is only at high speeds, and hence loads, that the engine operates at all efficiently. At low loads the operating point is well removed from the high efficiency (low specific fuel consumption) area. At a road load of 10 kW the engine operates at about 3000 r.p.m. and is relatively inefficient. By reducing the engine speed relative to the vehicle speed by a suitable change in gear ratio, the engine operating point

can be moved up, along the constant power line, towards the high efficiency region. As this operating point moves up this constant power line it would ultimately reach the engine optimum operating line, the locus of which links the maximum engine efficiency points at each speed. In order to follow this optimum engine operating schedule a continuously variable transmission (CVT) must be used, however as discussed later, a manual transmission can achieve comparable economy provided it is operated correctly.

Up to now the advantages of a transmission system have only been illustrated for the i.c. engine, nevertheless, the efficiency and performance of an electric vehicle can be similarly improved by using a variable transmission. A typical efficiency map for a separately excited d.c. traction motor is given by figure 6.1. In this case the high efficiency region occurs at relatively high speed and low torque in marked contrast with the engine. Efficient control of the motor therefore requires a different shifting logic from that of the engine. Another important consequence of the electric motor torque characteristic is the availability of torque at low speed, and in fact at zero speed. As a result neither a clutch nor a variable transmission are absolutely essential in any vehicle which has an electric traction system. This has lead many previous designs for all electric vehicles, such as the ETV-1 [Kurtz et al, 1981] and the Lucas Chloride van [Manghan and Edwards, 1983], to dispense with both clutch and variable transmission on the grounds that this gives a considerable weight saving and reduction in maintenance. Aside from the efficiency gains that a variable transmission offers to an electric vehicle, there are other advantages, such as improved performance and possibly simplified control electronics, to be considered when deciding whether or not to omit such a transmission.

As figure 6.1 illustrates, two distinct operating regions are required to cover the complete speed range of the motor. At about 2000 r.p.m. the

torque is seen to fall with speed. This transition point is termed the 'break speed', and is the operating speed for the motor with full field current and full armature voltage. Below the break speed it is necessary to control the armature voltage at full field current while speeds above the break speed are achieved by reducing the field current but at full armature voltage; a process known as field weakening. By including a variable transmission it is possible to make the motor break speed appear at relatively low road speeds, thus confining the need for armature control to very low vehicle speeds. Field weakening control can be extended still further by arranging for a battery switching system to halve the applied armature voltage. Once the minimum speed for field weakening is as low as 10 km/h, armature control electronics can be dispensed with altogether. Vehicle movement from rest is achieved by starting resistors in the armature circuit, as is the case in the Ford electric vehicle [Burba et al, 1986]. An electric system based on a separately excited motor, field chopper and starting resistors was identified as the most economic drive system available by General Electric during their component evaluation study for the HTV-1 project [Burke and Somuah, 1980].

Apart from the improvements in efficiency and reduction in electronic control requirements, the variable transmission also gives the electric motor improved performance as demonstrated by figure 6.2. In this diagram the motor maximum torque envelope of figure 6.1, has been converted to an equivalent traction force at the vehicle road wheels for three different transmission ratios. Also shown is a typical road load requirement at different gradients. With a fixed ratio of 5.65:1 the vehicle can cover the speed range 0-100 km/h but is incapable of starting on gradients greater than 15%. A typical performance specification for an i.c engine passenger car is to start from rest on 20% gradient and the ability to sustain 120 km/h on a 2% gradient. Obviously the fixed ratio electric vehicle has difficulty meeting either

requirement. By adding two additional gear ratios both low and high speed performances are improved.

In a hybrid petrol/electric vehicle the presence of the electric motor means that again a variable transmission is not absolutely essential as illustrated by the Lucas Hybrid [Harding et al, 1983] and an earlier hybrid built by Bosch [Fersen, 1974]. Nevertheless, if efficiency is an important design consideration one is almost certain to be included. With the differing efficiency maps of the engine and motor there is some conflict of shifting strategy in a hybrid and hence some improvement in vehicle economy can be made by using independent transmissions. Previous work at Durham [Bumby and Forster, 1987] suggests that overall an improvement of 10% might be expected, although it is debatable whether the increased complexity and weight would make this worth while; particularly when the hybrid design is already open to criticism for high component count when compared to conventional vehicles. A further argument against using two transmission systems for the present work is that the control strategies which have been designed to optimise energy consumption tend to operate the vehicle in either all electric or all i.c. engine modes thus allowing gear shifting strategy to be optimised for which ever unit is operational.

6.2 Variable Transmission Systems

At this stage the case for incorporating a variable transmission in an efficiency conscious hybrid vehicle has been clearly established. Once this decision has been made there are several possible transmission systems which might be adopted, with the CVT apparently representing the optimum solution. By far the largest category of CVT's are those based on the Van Doorne pulley [Steig and Spencer-Worley, 1981][Srinvasan et al, 1982], in which the power transmitting and power receiving pulleys vary in diameter

in opposing senses so that a constant length belt can be used to link them but still allow the speed ratio between their shafts to change. An alternative system is the Perbury, [Stubbs, 1981][Stubbs and Ironside, 1981] which makes use of discs with toroidal hollows in their faces separated by rollers with spherical rolling surfaces, these rollers being constrained in a cage system so that they can tilt and thus provide a variable drive ratio between the two discs. If these systems could match the efficiency, low cost, and reliability of a conventional manual gearbox and in addition were controlled to follow the engine optimum operating line they could achieve fuel savings of up to 30% relative to the conventional automatic and 15-20% relative to the conventional manual [Westbrook, 1986]. Unfortunately the belt and pulley CVT has a much lower efficiency than the 85-95% achieved by the manual and shows particularly poor performance under certain loading and speed conditions [Bonthron, 1985]. In contrast the Perbury CVT promises much higher efficiency but in this case it has been hampered by high manufacturing costs. For these reasons the only commercial application of CVT's up to 1984 was the Daf Variomatic belt system [Cuypers, 1984].

Due to the present low efficiency of practical CVT's there is a danger that any benefits accruing from improved engine utilisation can be lost in the transmission itself. Even with high efficiency the potential advantages of a CVT may be quite small when compared with an improved conventional system following an economic shifting strategy. A comprehensive study of the problem of matching engine and transmission carried out by Riccardo Engrs., [Thring, 1981], demonstrated that a discrete ratio transmission following an optimal shifting strategy can achieve virtually the same economy as a high efficiency CVT. The most important factor influencing whether or not high efficiency can be achieved with a stepped transmission is the span and not

the number of ratios. Specifically Thring points out that a gearbox with only 4 ratios can achieve very good fuel economy provided it has a span of at least 5. Gear boxes fitted to present production vehicles, such as the Ford gearbox fitted to the test rig, typically have a span of 3.5 with top gear giving 19mph/1000 r.p.m. An optimum design would provide about 35mph/1000 r.p.m. which together with a span of 5 would bring about a 30% improvement in fuel economy. Above these values a law of diminishing returns applies to both span and number of ratios, demonstrated by the fact that increasing the span to 8 yields only a further 2% gain in economy and a similar negligible improvement occurs if the number of ratios is doubled to 8.

Several manufacturers have recognised the fuel saving potential of using a conventional manual gearbox while automating the shifting process. Examples to date are those built by Ford [Main et al, 1987] and Isuzu of Japan [Isuzu, 1986]. Complicating the design of both these systems is the need to automate not only the operation of the gear shift lever but, in addition, the movement of the clutch. In practice the clutch has proven by far the more difficult item to automate successfully, with smooth and reliable operation being difficult to achieve. In their paper Ford suggest that the completed system would be most likely to find acceptance in light delivery vehicles and taxis rather than passenger cars. For this latter market Ford believe that conventional automatics designed for better fuel economy have the greatest potential because they offer better smoothness in shifting. Despite these difficulties considerable research effort continues into the automatic operation of the clutch system [Falzoni, 1983] [Tanaki, 1984]. In their study Fiat [Falzoni, 1983] reported major difficulties in obtaining smooth low speed operation, particularly in reverse.

The alternative option, mentioned above, of controlling a conventional

automatic so that it follows a more economic shifting schedule has also been investigated by Ford [Richardson et al, 1983] and Fiat [Busca et al, 1979]. In this case inevitable losses in the torque converter, even with lock up clutch, must limit the potential of such a system. A microprocessor controlled automatic, with torque converter removed, was chosen for the HTV-1 project [Somuah et al, 1983]. In this case there was no need to use a torque converter with the 3-speed General Motors automatic because the vehicle included two microprocessor controlled clutches capable of isolating the gearbox from both the engine and the motor during shifting. It is interesting to note that no problems with clutch control, even during low speed manoeuvres, are reported for the HTV-1.

From the above discussion it would seem that if a conventional manual gearbox can be successfully automated, it arguably represents the best choice in terms of fuel economy and production cost. Possibly the only barrier to its success in conventional i.c. engine vehicles is the need to control the clutch. This latter problem does not arise in a hybrid vehicle because the electric motor can be used to move the vehicle from rest and control the speed of the input shaft to the gearbox during shifting. As a result of this simplification it is believed that the system developed for the rig would be ideal for an operational hybrid.

Before moving on to consider the mechanics of the rig system, it is important to consider what characteristics constitute good driveability in an automated transmission. Firstly it is apparent that the manual gearbox cannot achieve a 'hot shift' and so it is important that the total shift time, during which torque is absent from the road wheels, must be as short as possible. In the Ford system [Main et al, 1987] total shift times amounted to 0.7 seconds which was considered to be acceptable. Secondly a key consideration is smooth action, with large torque transients in the prop-shaft being avoided

when power is restored.

6.3 Transmission System Hardware on the Hybrid Vehicle Rig

To achieve the automation of the gearbox on the rig a pneumatic actuation system has been mounted on the rear portion. This mechanism is illustrated by figure 6.3; overall dimensions are about 600 × 300 mm. In an operational system the design could be considerably reduced in size; in fact it would seem perfectly feasible that the final product would not be much larger than the original gearbox. Such scaling down has not been attempted on the rig because it would prevent gear shifting by hand, thus reducing experimental flexibilty.

The function of the gear-change mechanism is to move the gear lever in a similar manner to a human operator. Such freedom of movement is provided by the two pneumatic cylinders attached to the gear selection lever as shown in figure 6.3. Activating the longitudinal cylinder causes a plate to move backwards and forwards thus shifting the gear change lever between the 1/3 and 2/4 ends of the H gate. This plate is mounted between two sets of bearings rotating in the horizontal plane. Each bearing wheel has a grooved edge which locates with a bevel on the edge of the plate. A second cylinder is mounted on the plate and moves with it. This cylinder provides the necessary sideways movement of the gear selector between the 1/3 and 2/4 sides of the H gate.

The circuit diagram for the pneumatic system is shown in figure 6.4. Compressed air is supplied to the system from a reservoir which is recharged as necessary by an electric pump. During normal operation a regulator valve maintains the working pressure in the system at about 2 bar. Each of the two working cylinders is controlled by a piston valve. There are five ports on each piston valve with the flow path between them being controlled by

activating the appropriate solenoid. In the case of the longitudinal cylinder two additional components are used to help stop the piston in the central, neutral position. The first is a fast acting valve which cuts the air supply to the cylinder 12 ms after its solenoid has been energised. Secondly flow regulators are fitted to both supply lines feeding the cylinder. These allow air to flow freely into the cylinder but, when air is flowing out to the exhaust ports, the flow is restricted. As a result the speed of the piston's movement is decreased which reduces its tendency to overshoot the neutral region before the air cut-off valve can be activated.

Both cylinders have magnetic pistons which operate sensor switches attached to the outside of the casing, as illustrated in figure 6.3. A system of two switches on the transverse cylinder and four on the longitudinal cylinder completely define the the position of the gear lever. Each sensor switch can easily be adjusted by repositioning along the casing thus ensuring that it is activated when the gear lever reaches precisely the right position. The area of the H gate covered by each position sensor is shown diagramatically in figure 6.5.

The six sensors alone are sufficient to confirm when a particular gear is engaged and when the gearbox is in neutral. Logical OR of the two centre sensors provides a single signal which covers the whole of the neutral region. This signal is interfaced to a positive active edge interrupt input on one of the VIA chips. Whenever this signal generates an interrupt the air supply to the longitudinal valve is cut using the fast acting cut off valve. However, ambiguous positions exist, for example, the position reading obtained between first and neutral is the same as that obtained between neutral and second. A solution to this problem is provided by connecting the two centre sensors to a positive edge triggered D type flip-flop decoder circuit. The flip-flop circuit uses one centre sensor as clock and the other as D, relying on the fact that

the two centre sensors are displaced slightly from one another. As a result an output is produced which is high forward of neutral and low to the rear.

6.4 Interface Circuitry

In order to give the M68000 full control of the pneumatic system both position sensors and solenoids are interfaced to the VIA chips. Two chips are used, one assigned to input signals from the position sensors and decoder circuit and the other assigned to output signals for the control valves. Power for the position switches comes from the 12v engine starter battery, making the signals prone to noise from the ignition, in addition to any other noise picked up as the signal cables cross the rig. Consequently each signal passes through a dc input module, which includes some filter circuitry, before finally arriving at the M68000 VIA input port.

Solenoid valve control signals from the VIA passs through opto-isolator units before reaching the relevant valve. These devices not only provide noise protection for the M68000 but switch the necessary 12v battery supply to the valves.

6.5 Software Control

As described in chapter 2, the M68000 software is structured so that time critical control tasks run every 20 ms in interrupt routines whilst non-critical tasks such as user requests run in the background. Gear changing encompasses both types of software task; the fundamental control of the gearbox runs in interrupts whilst manual requests to change gear come from the background. When the complete drive system is running however, gear change requests are also generated in interrupts. Gear change logic is provided by a single 'C' sub-routine, which has the flow chart illustrated by figure 6.6. Direct access to the input and output registers of the VIA chips involved is

provided by two small assembler routines. All that is necessary to activate the main 'C' routine is to call it in every interrupt passing the required gear each time. On completion of the gear change the routine returns false. Each gear change breaks down into three stages as described in the following sections.

6.5.1 Stage 1: Shift into Neutral

At the start of any gear change the loading on the gearbox is removed and the gear lever is moved into neutral. During its first call the gear change routine activates the necessary solenoids to move into neutral by sending a bit pattern to the relevant VIA output register. A two dimensional array provides the program with all the bit patterns it needs based on the current and next gear. These codes activate the longitudinal cylinder as appropriate to get out of the current gear and simultaneously activate the transverse cylinder to move left or right ready for the next gear. When the gear lever moves into the central neutral region a high level interrupt signal (level 4) is generated. At this stage the air is supply to the longitudinal cylinder is immediately cut. Control now reverts to the main gear shifting program which is executed in the level 3 interrupt software every 20 ms. It would be slightly simpler to allow this routine to cut the supply to the longitudinal cylinder, however this would entail a delay in this action which might amount to 20 ms. It was felt that the need for prompt action in this respect justified the use of the higher level interrupt. Reaching neutral does not necessarily mean that the gear lever is on the correct side of the H gate, and consequently the main routine continues to check the position sensors until the appropriate position is reached. Checking the position consists of reading the VIA input register ten times and only accepting a given result if all ten give the same correct reading.

Shifting into neutral is made easier if no power is being transmitted by the gearbox, and so the program takes steps to remove the loading from both the motor and engine. Unfortunately in neither case is this simply a matter of setting an accelerator signal to zero. Complications in reducing the engine loading arise from the fact that it must pick up load smoothly again once the gear change is completed. If the engine throttle is simply set to zero, engine speed will fall far behind the rest of the drive train during the gear change. Consequently when the throttle is returned to its previous setting, the engine can accelerate very rapidly until it hits the rest of the drive train, possibly causing an unacceptable jolt in a real vehicle. To solve this problem closed loop speed control of the engine is provided to maintain the engine at about 88 r.p.m. (4 counts on the engine speed transducer) behind the drivetrain throughout the gearchange.*

Set points for the engine speed controller depend on whether the gear change is up or down. With an upshift the engine can immediately start to slow towards the speed it will need when the new gear is engaged. In contrast it is not physically possible to accelerate the engine to the speed that will be necessary after a downshift, until the gearbox is in neutral. Hence two setpoints are used; for a downshift the engine follows 88 r.p.m. behind the motor as it accelerates and for an upshift the engine is sent immediately towards the speed needed for the next gear.

When the electric motor has been providing power prior to the gear change particular attention must be paid to the accelerator and brake signals so that is simply rotates without loading the drive shaft. If both brake and accelerator signals are set to zero then the logic within the power electronics becomes undefined causing the contactors to switch continuously between the acceleration and braking modes. Application of a small signal to either input

solves this problem and causes the controller to switch positively into one or other mode, with acceleration taking priority if both signals are present. Accordingly the gear change routine applies a small signal to the accelerator in preparation for a down change and a brake signal for an up change.

6.5.2 Stage 2: Speed Matching

Once neutral has been obtained the next stage is to synchronise the speed of the input shaft of the gearbox to that of the output taking into account the ratio of the next gear to be engaged. This is achieved by sensing the speed of the flywheel and controlling the speed of the electric motor to a setpoint based on the flywheel speed times the gear ratio. Action of the synchronisation algorithm depends on the size and sign of the error in a way that has been designed to minimise the time taken to match the speeds. Intuitively mixed use of accelerator and brake, depending on the sign of the demand from say a P+I algorithm might seem an obvious form of control. However this is not the case because of the delays that arise when the power electronics have to switch between braking and accelerating. In practice these delays are caused by the contactors closing and the field current reversing, which amounts to an unacceptable 0.5 second delay. To avoid this problem, and achieve synchronisation in about 0.7 seconds the controller takes different courses of action depending on whether an up or down shift is needed. On a change down, having preset the contactors to accelerate prior to making the neutral move, speed control is achieved by a P+I algorithm using the accelerator alone. During a change down accelerator action is all that should be required, however if the system overshoots slightly natural slowing of the motor provides adequate correction without recourse to regenerative braking. One slight modification to the basic P+I algorithm is that the gains are reduced for small errors since the unloaded motor is highly

sensitive to variations in the input signal. On an up shift the system has been primed for braking and so a constant brake value is applied for as long as the overspeed error is greater than 200 r.p.m. Once the overspeed error has been reduced below this level the controller sends out a small accelerator signal in anticipation of the time it takes the electronics to make the necessary changes. Meanwhile the motor continues to slow naturally and eventually reaches the required speed, at this stage the P+I accelerator algorithm takes over to maintain this speed until the gear is successfully engaged.

6.5.3 Stage 3: Engaging the New Gear

As soon as the speed error is less than 66 r.p.m. the longitudinal cylinder is activated to engage the new gear. Any residual speed errors are partly corrected by the synchromesh in the gearbox, assisted by the continuing action of the P+I control algorithm. Once the engage move has been made the program waits for the sensors to confirm that the position for the new gear has been reached. Finally the program returns false to the calling routine thereby signalling that power may be restored to the drive shaft.

6.5.4 Error Handling

Several error conditions are feasible during the gear change operation, the main control routine has been designed to detect the occurence of three of them. Paramount consideration in handling error conditions is to protect the gearbox against the possibility of attempting to engage a gear when there is a large speed error, or when either prime mover is attempting to provide power. If this precaution is not fully observed, the result can be destruction of the gearbox.

As indicated above it is possible for the gear selector to stop in a position that, although in neutral is not covered by the centre sensor. If this

happens the main gear change routine would wait indefinitely for the neutral position to be reached. Accordingly 25 interrupts (0.5 seconds) are allowed for the neutral position to be reached and after this time it is assumed that an error has arisen. When this occurs control is passed to an error handling routine which attempts to nudge the gear selector back towards neutral. Interrogation of the output from the decoder circuit determines which direction the longitudinal cylinder must move to achieve this.

Another error condition occurs if an erroneous interrupt from the centre position sensor causes the air supply to be cut before the engage or neutral moves are complete. To guard against this problem the program allows three seconds overall to engage the gear and if this is not achieved the error handling code forces the main program to remake the necessary move; subject to successful speed matching.

Finally if the gear change is not completed after five seconds, the error handling code assumes that something is disastrously wrong and it abandons the whole gear change. In taking this action the gear change program cuts the engine ignition and disables both motor accelerator and brake action. It is therefore impossible for external routines to restore power to the drive shaft.

Should any of the above error conditions occur key control variables and position signals are stored by the software so that a later diagnosis can be carried out by the operator. It should be mentioned that with the present hardware and software configuration none of the above errors has occured, they came to light as possibilities during earlier development work. For example erroneous interrupts occured when the dc input modules were placed near the cylinders so that they did not filter the signal just prior to it reaching the computer.

6.5.5 Location on Power Up

On initial power up a gear change initialisation routine is activated that locates the gear selector lever in the neutral left position ready for the first gear change. This is done without use of the decoder circuit since the output from the flip-flop might erroneous until the neutral position is crossed for the first time. A simple series of logical movements based on the sensor readings makes up the structure of the location routine.

6.6 Results

Typical results for both a down-change and an up-change are shown in figures 6.7(a) and 6.7(b) respectively. Both changes occured when the drive train was operating in i.c. engine mode. As such the engine provides all the power before and after the gearchange. This type of operation provides a stiffer test of the gear changing system than a change in electric mode because of the delays in restoring torque inherent in the way engine speed lags slightly behind the gearbox input shaft during the change.

Figures 6.8(a) and 6.8(b) shows speed synchronisation in more detail and indicate how much of the total shift time is taken up by each stage of the shifting process. Speed matching is the most time consuming stage with considerable dead time before the motor speed responds at all to the setpoint. This set point speed becomes defined as soon as neutral is reached and is then related to the flywheel speed through the appropriate gear ratio. Slight decay in the flywheel speed can be seen during both gear changes as a result of the dynamometer loading. Variations in engine speed given by figure 6.8 illustrate how it immediately drops away towards the final setpoint value during an up change whereas in a down change it is tightly linked to the motor speed as soon as neutral is reached. Natural decay of speed for the engine is seen to be much slower than active braking for the motor.

6.7 Discussion

Before the gear changing system was described in detail two performance criteria were suggested; total shift time and smooth reapplication of torque. Experimental results show that typical shift times for the rig system are 1.16 seconds for an up change and 1.48 seconds for a down change. Greater adjustments in speed are required if the shift is between non-sequential gears which typically increases the total shift time by 200 ms. Figure 6.8 shows that of this total time 200-400 ms are required to reach neutral, speed synchronisation requires 600-800 ms and the engagement process requires a further 100-300ms. This wide range of times for the engagement move is a result of varying residual speed errors when engagement is first attempted. Although quite reasonable for development work on the rig, this total shift time is actually rather slow for acceptable driveability in a production vehicle, a more realistic target is the 700 ms figure quoted by Ford. Fortunately there are a number of modifications to the present system which make it quite likely that the total shift time could be reduced to a far better value.

A substantial improvement in shift time could result from the use of a more advanced electronic control unit now manufactured by Lucas [ETV World, 1988]. The Mark V power electronics control unit described in this article makes very limited use of mechanical contactors, thus making it possible to switch instantaneously between motoring and braking. As a consequence it would be possible to incorporate mixed use of regenerative braking and accelerator action in a single P+I control algorithm; furthermore dead time apparent in the present system before it responds to a control signal should be substantially reduced. Despite continued use of the older Mark III controller on the rig there is still some scope for improving shift times by altering the control algorithm. At present controller gains are fixed over the whole operating range of the motor. As was revealed in the analysis of the motor

in chapters 3 and 4, three operating modes exist each with differing dynamic characteristics. Consequently using one set of controller parameters for speed matching must reduce performance under certain conditions. A look up table of gains to be adopted over different speed bands is one possible solution to this problem. A further refinement worthy of consideration is to increase the bias brake/accelerate signal during the neutral move in an attempt to reduce dead time once full speed control commences. Any gains here though would have to be weighed against a possible increase in time required to reach neutral due to slight loading of the gearbox.

A second possible improvement may be obtainable by adopting a radically different control strategy. In a similar system, developed for an all electric vehicle, [van Niekerk et al, 1980] satisfactory engagement was achieved by simply sweeping the motor speed past the desired set point and attempting to engage as soon as a 10% error band is entered. Using this technique total shift times of 700 ms were achieved.

Finally it must be recognised that the mechanical system has been partly compromised from the ideal by the desire to retain a manual shift facility. If the longitudinal cylinder acted directly on the selector forks and the transverse cylinder were replaced by rotary action, not only would the resulting system be more compact, but the shifting action would be faster. Following this course of action would require higher operating pressures to compensate for the loss of the mechanical advantage offered by the gear lever. Many operational systems such as that built by Ford, use a hydraulic oil based actuation system working at much higher pressures than those used on the rig and thus allowing a more positive shift action.

Leaving aside operation time, perhaps one of the greatest advantages of the present transmission control system is the way that it fulfills the second design criteria, that of smooth reapplication of torque once the gear change

is complete. In figure 6.7, the engine torque traces shows no substantial osscilations on re-engagement, so that good driveability in this respect can be anticipated. Some indication of the amount of torque oscillations in the prop-shaft encountered when a typical driver starts a vehicle from rest is given by Lucas [Lucas and Mizon, 1978]. Their results, and those quoted by Ford in [Main et al, 1987] as representing a perfectly acceptable gear change, show far greater oscillations than figure 6.7, suggesting that the speed matching system exceeds requirements.

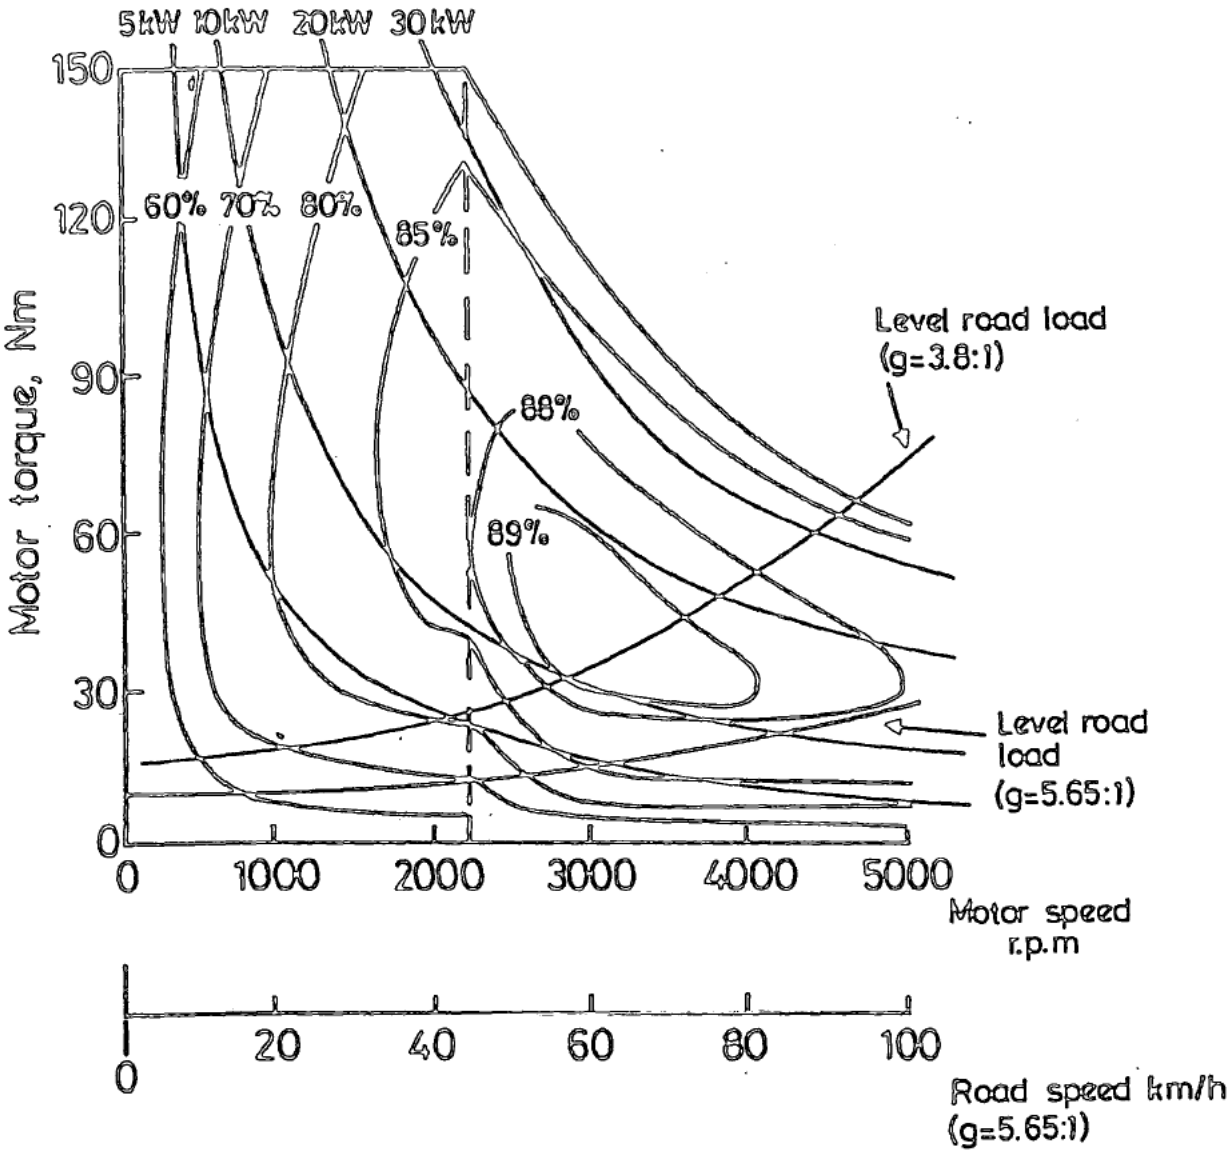

Fig. 6.1 Road Load and Motor Operating Curves
for All Electric Vehicles

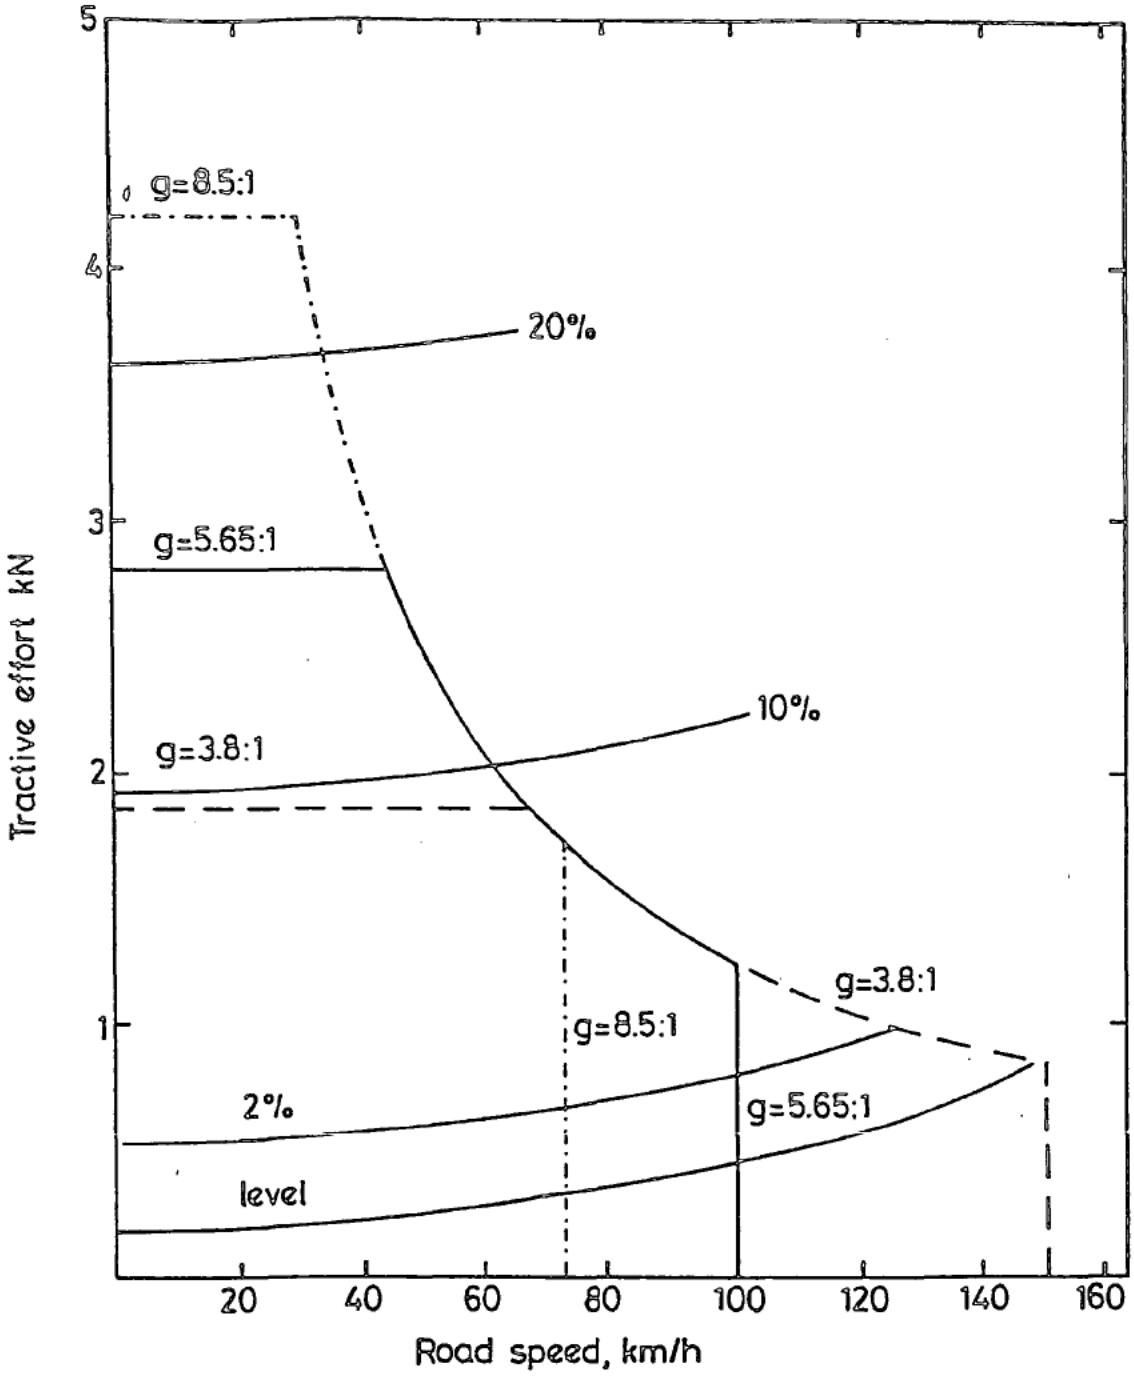

Fig. 6.2 Tractive Effort and Road Load Curves for All Electric Vehicles Showing the Effect of Gear Ratio