Design of a Magnetorheological Brake System

by

Luís Falcão da Luz
Licenciado in Aerospace Engineering, Instituto Superior Técnico, Lisboa, 2002

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
MASTER OF APPLIED SCIENCE
in the
Department of Mechanical Engineering.

We accept this thesis as conforming
to the required standard

University of Victoria

Supervisors: Dr. Afzal Suleman, Dr. Edward Park

Abstract

The hydraulic brakes presently used in cars exhibit several important limitations. These include the slow response to the driver's command; difficulties in control due to the hydraulic nature of the system; and a large number of components spread throughout the car with critical components such as the disk surface, the brake pads and the fluid pipings vulnerable to damage from gravel or other external sources. To overcome these problems, intrinsic to the concept of hydraulic brakes, a new system must be devised. Solutions are sought in the use of smart materials, including the application of piezoelectric or electrostrictive materials and electrorheological or magnetorheological fluids to car brakes. A detailed study of each material is carried out, in terms of their possibilities and limitations. It is seen that present piezoelectric and electrostrictive materials are unable to meet the performance requirements needed for application to car brakes and that electrorheological fluids are less suitable than magnetorheolical fluids for this application. Consequently, an innovative car braking system is designed using magnetorheological fluids.

The design procedure comprises the study of theoretical models for the performance of a magnetorheological brake and, given the absence of closed-form solutions for the braking torque of an arbitrary brake system, finite element models are built to provide a means to analyse the performance of the magnetorheological brake system. The formulation of these models (including the definition of the geometry, material properties, boundary conditions and meshing process, as well as necessary assumptions) are described. The results obtained with the finite element models are presented and analysed.

In order to obtain an optimum design, i.e. one with high braking power and low weight, an optimisation procedure is carried out, centred on the finite element

analysis. Three different optimisation methods are used (subproblem approximation, first order and simulated annealing). Their performance and relative methods are compared.

Based on the results of the optimisation problem, a final design is proposed, taking into account manufacturing constraints and a study of its longevity and reliability is carried out. A scaled-down prototype is also proposed to serve as a proof of concept.

Finally, the strengths and weaknesses of magnetorheological brakes and the expected evolution of this techonology are discussed, as well as conclusions regarding the use of piezoelectric or electrostrictive materials for brake actuators.

Examiners:

Table of Contents

List of Tables

List of Figures

Nomenclature

Conventions

All values in this thesis are in SI units, except where explicitly noted otherwise.

Variable Definitions

Below is a list of the variables used throughout the thesis and the meaning of each.

a	Acceleration
A	Area
B	Magnetic flux density
C_D	Aerodynamic drag coefficient
C_p	Specific heat at constant pressure
D	Aerodynamic drag
E	Energy; Young's modulus
E_c	Kinetic energy
f_r	Rolling friction coefficient
F	Force
F_{max}	Blocked force of a piezoelectric or electrostrictive element
F_r	Rolling friction
g	Gravitational acceleration
G	Shear modulus
h	Convection coefficient
H	Magnetic field intensity
I	Current intensity
J	Polar moment of inertia
k	Thermal conductivity

L	Inductance
m	Mass
N	Number of disk surfaces; Number of wire turns in the coil
Nu	Nusselt number
P	Power
Pr	Prandtl number
$\dot{Q}$	Heat transfer rate
R	Electric resistance
Re	Reynolds number
t	Time
T	Torque; Temperature
T_B	Torque due to the magnetic field
T_μ	Torque due to the plastic viscosity
T_∞	Far-field temperature
U	Voltage (electric potential); Fluid velocity
v	Velocity
V	Volume
W	Work; Weight
Δx	Displacement
Δx_{free}	Free displacement of a piezoelectric or electrostrictive element
$\dot{\gamma}$	Strain rate
κ	Thermal diffusivity
μ	Magnetic permeability
μ_p	Plastic dynamic viscosity
ν	Kinematic viscosity; Poisson's ratio
Φ	Magnetix flux; Angular displacement
ρ	Density
τ	Shear stress
τ_e	Electro shear stress (due to the magnetic field)
τ_p	Plastic shear stress
ω	Angular velocity

Acronyms

While an effort was made to introduce the meaning of each acronym upon its first reference in the text, the most significant ones are included below for easy reference.

3D	Three Dimensions/Three-Dimensional
ABS	Anti-lock Braking System
AWG	American Wire Gauge
CFD	Computational Fluid Dynamics
CNG	Compressed Natural Gas
EPA	[United States] Environmental Protection Agency
ER	Electrorheology/Electrorheological
FE	Finite Element
FEM	Finite Element Model/Modelling
FTP	[United States] Federal Test Procedure
IUT	In-Use Thickening [of magnetorheological fluids]
ISA	International Standard Atmosphere
LPG	Liquified Petrol Gas
MR	Magnetorheology/Magnetorheological
SUV	Sport Utility Vehicle

Acknowledgements

A university degree is a long commitment involving many challenges and its accomplishment involves the help of many people. While it would be impossible to acknowledge them all in here, I wish to mention those more closely attached to this project: my supervisors, Dr. Afzal Suleman and Dr. Edward Park, for sharing their time and knowledge and, more importantly, for creating such a fantastic workgroup spirit; Stan Burns for the finite element models of the electrostrictive inchworm actuator; Steve Ferguson for all the knowledge of cars and for the invaluable help with the construction of the hydraulic brake setup; Rodney Katz and Ken Begley for the interest and time dedicated to this project and all the practical advice given; Dilian Stoikov for extending the possibilities of this project through the development of control algorithms; everyone in the workgroup for the friendship and the discussions that always stimulated further work; and, most importantly, those who made these 2 years of my life unforgettable: David, Diogo, Ahmad, Joana, Ana and Pedro at home, and Gonçalo, Marc, Sandra and Scott in the office; last, but certainly not least, the family and friends in Portugal and elsewhere who encouraged me and who, in one way or another, were always close in spite of the geographical distance. A particular appreciation goes to my mother and to Rita.

To my parents, who taught me everything

Chapter 1

Introduction

This year's surge in oil prices was one more episode in a long series of crises since the 1970s. This continued instability has prompted major developments in the use of alternative energies for industrial production, home use and transportation. Focusing on transportation, it is interesting to notice that indeed the first cars to have low fuel consumption as a key asset (such as the Renault 5, Fiat Panda and Citröen AX, the latter's body using composite panels to reduce weight) were introduced in the late 1970s and early 1980s. This era also witnessed major developments in Diesel engines for cars, due to their lower fuel consumption for comparable performances. These developments include the direct injection turbo-Diesels first seen on the Fiat Croma in 1988 and the groundwork for the common rail technology that would become commercially available with the Alfa Romeo 156 JTD in 1997.

But the research interest spurred by the 1970's oil crisis was not limited to improving gasoline and Diesel engines. Much attention was given to electric motors as an alternative to the classical thermal engines for the propulsion of cars, resulting in the availability of electric versions of several car models (such as the Peugeot 106 since November 1995 and the Citroën Saxo since July 1997). However, the size and

weight associated with the batteries seriously limits the energy capacity and hence the travel range of these cars is currently restricted to between 80 and 90 km [1]. This can be identified as the biggest obstacle to a greater success of electric cars. Further growth of electric car sales is also hampered by the existence of non-electric components in the car (among which the brakes), thus requiring heavy electromechanical components to be added, increasing the weight of the car and further reducing the weight available to batteries.

One way to overcome the limitations of electric cars due to the limited energy supply of the batteries was found in hybrid propulsion (already commercially available in e.g. the Toyota Prius, Honda Insight and Honda Civic, with more cars expected to be available by the end of 2004 including the Ford Escape and other models by Chevrolet and Mercedes-Benz [2]). These possess both an electric motor and a thermal engine. When driving for long periods of time at constant speed (e.g. on a highway) the thermal engine is used both to power the car and charge the batteries. When driving in traffic in urban areas, the thermal engine is switched off and only the electric motor is used as it is much more efficient in these conditions, given that unlike a thermal engine it does not waste energy when the car is stopped.

Simultaneously, recent years have witnessed a significant increase in environmental awareness, prompting a variety of changes in all areas of engineering. Vehicles, due to their impact in global energy consumption and pollutant emissions, are among the systems undergoing more pronounced improvements. Given the pollutant emissions of all thermal engines and the increasing pressure towards cleaner air, electric motors are the subject of increased research interest and expected to grow beyond their current niche position towards becoming a viable option. Since the major obstacle to a greater competitiveness of electric cars lies in the limitations of current batteries, one of the most active areas of research is that of energy supply for electric cars (with

a particular focus on the use of fuel cells), both in terms of technical solutions and the economic aspects of their implementation [3].

As a result of all the development carried out in the past decades, recent years have finally witnessed widespread interest in alternative (more environment-friendly) energy supplies for road vehicles:

- In individual consumers: rise in the number of cars running on LPG (liquified gasoline gas) in Europe; importance of cars powered by biofuel (combination of gasoline and alcohol) in Brazil, which accounted for 17% of the sales in the first months of 2004 [2]; increasing interest in electric vehicles - in France, for example, the cumulative number of new electric vehicles registered grew from 296 in 1993 to 5608 in 1999 [1].

- In companies: to name but a few examples, British Columbia's transit operator BC Transit participates in the development of fuel-cell technology with three vehicles received in 1999 [4]; French mail service La Poste conducted numerous experiments throughout the 1990s with electric cars [5] operating a total of 700 vehicles as of October 2002 [6]; EDF (Electricité de France) operated 1500[1] electric vehicles and expected to acquire a further 1500 [6]; the European Community sponsors different research projects in the area of fuel cell vehicles, most notably:

  - project CUTE which involves 9 cities in the European Union (Amsterdam, Barcelona, Hamburg, London, Luxembourg, Madrid, Porto, Stockholm, Stuttgart) in conjunction with projects in Iceland and Western Australia, based on Mercedes-Benz Citaro buses [7];

[1] Also as of October 2002

- under project Joule/Thermie, a study involving various companies and institutions and the transit operators of Berlin, København and Lisboa, centred on a new bus developed by MAN [8];

However, the fact remains that whatever type of engine is used, the energy cannot come entirely from "clean" (environment-friendly) sources. Even electricity is largely obtained from thermal powerplants burning fuel. Hence, it became clear that it was fundamental to devise ways to achieve optimum efficiency of the engines (that is, to increase the ratio between the work produced by the engine and the energy required). This has been addressed by research on electronic control of the engine, which resulted in a major improvement: whereas gasoline engines were until the early 1990s fed by a carburettor (a purely mechanical device that supplies the engine an amount of fuel proportional to the accelerator position), modern engines have electronic fuel injection. This system's core is a highly sophisticated electronic unit which decides the ideal amount of fuel to be supplied to the engine at each time, based not only on the accelerator position but also on parameters such as the engine speed and temperature. It can even provide a short power boost when required (e.g. when the driver presses the accelerator to its fullest to overtake another car) by temporarily disabling the air conditioning. This integrated control of various systems in order to achieve optimum performance would not be possible with purely mechanical components and highlights the advantage of electronic systems. Therefore, it comes as no surprise that more and more of a car's traditional systems are being replaced by electronic components. For example, the power steering systems that provide some of the force required to steer the car's wheels (therefore relieving the driver) have traditionally relied on hydraulics. Recently, however, electric power steering systems have been developped[2].

To gain a better perspective of the growing influence of electric and electronic

[2]e.g. in the new-model Fiat Punto introduced in 1999

systems in cars, it suffices to consider that the length of electrical wiring in cars has doubled over the past thirty years [9]. This massive growth of electrical functions and associated wiring was only expected to increase, which prompted car manufacturers to tackle this issue. As a result, the Peugeot 607 became the first mass-built car equipped with multiplexing. This consists of replacing the bundles of wires (one for each function, such as warning lights, air conditioning, windscreen wiper, and so on) currently found in most cars with only a handful of wires[3] and having electronic processors distribute the various signals among this reduced number of wires[4]. This paves the way to the all-electric car: indeed, all seems ready for every function of the car to be controlled by electric signals! In fact, this trend has been best summed up as the "transition from mechanical cars with electronics, to electronic cars with mechanical portions" [11].

In order to complete this transition, electric solutions are sought to replace the remaining mechanical components. The various advantages of such a change in the braking system, which is the focus of this thesis, are discussed in the following section.

1.1 Motivation

As it was mentioned in the previous section, many traditionally mechanical systems in cars are being replaced with electrical components. Whereas the initial impetus for this change stemmed from economical and environmental concerns, it was soon realised that the use of electric and electronic components could bring major improvements in performance. One area where this is particularly clear is the braking system. While a driver may not be aware of how to achieve the optimal braking and

[3]Instead of the traditional copper wires, optic fibres may be used [10]
[4]In the Peugeot 607, the number of connections was reduced four-fold [9]

may sometimes just slam on the brake pedals causing the wheels to lock and thus losing steering control, a system comprising electronic sensors and processors may detect when a wheel has locked and adjust the brake pressure accordingly - this is the principle of operation of an ABS (anti-lock braking system). However, as long as hydraulic brakes are used, the ABS must rely on electromechanical interfaces such as valves to adjust the braking pressure. This slows down the response of the system, compromising the braking distance, and brings extra weight and cost to the system due to the inclusion of extra components. Recall from the previous section that this extra weight limits the battery capacity (and hence, the autonomy) of electric cars.

Another inherent limitation of hydraulic brakes lies in the fact that between the moment the driver presses the pedal to its fullest and the moment full pressure is transmitted to the brake pads, a time delay of between 200 and 300 ms occurs as the pressure propagates from the master cylinder to the callipers throughout the brake fluid circuit. This brings a significant increase to the braking distance as it means that at 120 km/h the car travels over 6.5 metres before even starting to fully brake after the driver acted. Even at a lower speed (e.g., 100 km/h), the increase in braking distance due to this time delay will be of 5.5 metres, which is more than the length of one car and constitutes a penalty of 12% of a typical braking distance (about 45 metres).

A further drawback of conventional hydraulic brakes is the vulnerability of the disk which, because of its open configuration, is exposed to external elements which may seriously degrade braking performance (e.g., oil or greases) or permanently damage the disk's surface (e.g., gravel). Also, the very nature of hydraulic brakes with pipes and hoses throughout the car makes the system more vulnerable to catastrophic failures due to fluid leaks or to premature ageing of the brake fluid[5].

[5]If the brake fluid is in contact with air it will absorb water vapour, causing a reduction in the

So far, the developments in car braking systems have been centred in minor improvements and fixes to hydraulic brakes. As indicated above, these exhibit intrinsic performance limitations and require electromechanical parts that prevent optimum operation of electronic control systems. In order to address these issues, electric brake systems must be devised. While some concepts for electric brake systems have already been presented, recent work in materials science has resulted in smart materials suitable for application to brakes. Smart materials are those whose properties can be changed by the user, such as piezoelectric crystals which exhibit a deformation proportional to the applied electric field, or shape memory alloys whose form depends on the temperature.

Given the great interest surrounding electric components for cars that has been mentioned above, some research has already been conducted in the area of electric brakes. These developments are the object of the next section.

1.2 State of the Art

In 2002 Delphi has introduced an electric calliper for disc brakes [13]. The concept of this brake system is the same as that behind today's brakes but with the traditional hydraulically-actuated calliper replaced with one whose clamping force is provided by an electric motor. Given that the electric motor produces rotational movement, gears are needed to translate that rotation into the linear movement necessary to push the brake pads against the disk and pull them away. Difficulties associated with the transmission of forces from the pads to the motor were addressed by a system involving a set of planetary gears, subject of a patent submitted by Delphi [14].

boiling temperature of the fluid (normally between 230 and 290 °C) of 60 to 80 °C per year [12]. This lower boiling temperature causes the brakes to fail at lower temperature and hence reduces the system's endurance when subject to demanding braking for long periods of time (e.g., long descents)

Another system whose application to car brakes has been the object of recent research involves the use of eddy currents [15, 16]. Eddy current brakes have been used for a long time in trains in addition to the standard pneumatic system. They consist of having a disk of a conductive (but not magnetically permeable) material rotating within a varying magnetic field. This varying magnetic field induces swirl-like currents (the eddy currents) in the conductor and in turn a magnetic field opposed to the applied one is generated. This interaction generates a force on the conductor that slows it down. The main advantage of this system is the possibility of transforming the kinetic energy of the car into electrical energy instead of dissipating it as heat (thus wasting it) as happens in friction brakes. The main disadvantage is the very poor performance of eddy current brakes at low speeds which has, until now, meant than an auxiliary pneumatic or hydraulic braking system must be used to provide sufficient braking torque at low speeds.

1.3 Objectives

Considering all that has been said above, the goal of the present project can be summarised as follows: to conduct a study of smart materials suitable for application to car brakes and to develop an innovative system based on the most promising materials. The intended advantages of this new system over conventional hydraulic brakes are:

- Improved performance (braking distance);
- Better "packaging" (reduction in the number of components spread throughout the car);
- All-electric control (eliminating the mechanical components).

1.4 Thesis Outline

The thesis presents the design process of the proposed magnetorheological (MR) brake and therefore the structure follows the sequence of tasks carried out from the definition of the target specifications to the final design.

Chapter 2 presents an overview of the operation of today's hydraulic brakes to identify the requirements and constraints faced by a new system, this information and experimental data of common braking performance is then used to establish precise performance targets that must be met by the new brake system. A literature review of suitable smart materials for application to car brakes (piezoelectric and electrostrictive materials, electrorheological and magnetorheological fluids) is then conducted, particularly in terms of their properties and performance capabilities, as well as existing successful applications and a detailed study of the potential for application to car brakes is performed.

Once the required performance of the brake system has been determined and the most adequate smart materials have been selected, the concept behind the proposed braking system is presented in Chapter 3, and various configurations are devised and their relative merits are studied. This leads to the selection of the most promising configurations for detailed analysis. The criteria leading to the selection of the materials to be used in components such as the brake disk, shaft and casing as well as the type of coil wire that will be used to produce the magnetic field are also presented in this chapter, along with a comparison of the properties of the various materials. The procedure to be used in the analysis and optimisation of the design is also introduced at this point.

In order to obtain a complete understanding of the operation of the MR brakes, finite element models were built to simulate the magnetostatics behaviour, the fluid

flow and heat transfer within the system, as well as its structural resistance. These, along with the assumptions and simplifications that were made, are presented in Chapter 4, together with considerations regarding the transient behaviour of the system under time-varying braking demands and the influence of phenomena of mechanical behaviour of materials (such as fatigue and creep) that may have an impact in the life of the components. The results obtained are presented and discussed, as well as the design changes motivated by the insight gained from those results.

After the various concepts have been updated to better address all the issues faced by car brakes (reflecting the knowledge gathered from the initial finite element analyses) and finite element models have been built to describe the response of these MR brake concepts, an optimisation procedure may be carried out so as to obtain the best possible design (configuration, dimensions, MR fluid type) in terms of weight and performance (braking torque). This is described in Chapter 5, which first introduces the fundamentals of optimisation procedures, discusses the application of three different methods to the present problem and presents the optimum designs returned by each method. These results are analysed to compare the relative merits of each optimisation method.

Having reached the optimum configuration for the MR brake, several considerations regarding manufacturing issues and the sealing of the fluid are described in Chapter 6 and taken into account to yield the detailed design. Longer term matters associated with the implementation of MR brakes in common vehicles, such as the expected longevity and reliability are also introduced.

Finally, the conclusions of the project are presented in the last chapter, where the limitations of this design are discussed and form the basis for the layout of future work that may help improving this technology.

Chapter 2

Background

2.1 Hydraulic Brakes

Since it is the goal of this thesis to present an innovative design for car brakes able to replace the current systems, it was important to obtain an in-depth knowledge of today's brakes, so that performance targets could be established for the new concept ensuring its competitiveness and also in order to gain the maximum possible insight into the challenges involved in the design and operation of car brakes. To this end, a typical brake system has been assembled in the laboratory. It was important to use all the components from the same car model and version to ensure that this experimental setup represented a real configuration. Also, it was decided to model a car without power-assisted brakes which would greatly increase the complexity of the setup (requiring an engine to drive the vacuum pump) without any change in the braking performance (brake assist systems are only intended to reduce the force which the driver must exert on the pedal). With these two considerations in mind (no power-assistance and easy availability of all required components from the same version), the choice was to use the brakes of a 1983 Chevrolet Chevette. The setup

of this experimental system was done in conjunction with colleagues Steve Ferguson and David Cruz in the period from May to August 2003. A picture of the overall system is shown in Fig. 2.1.

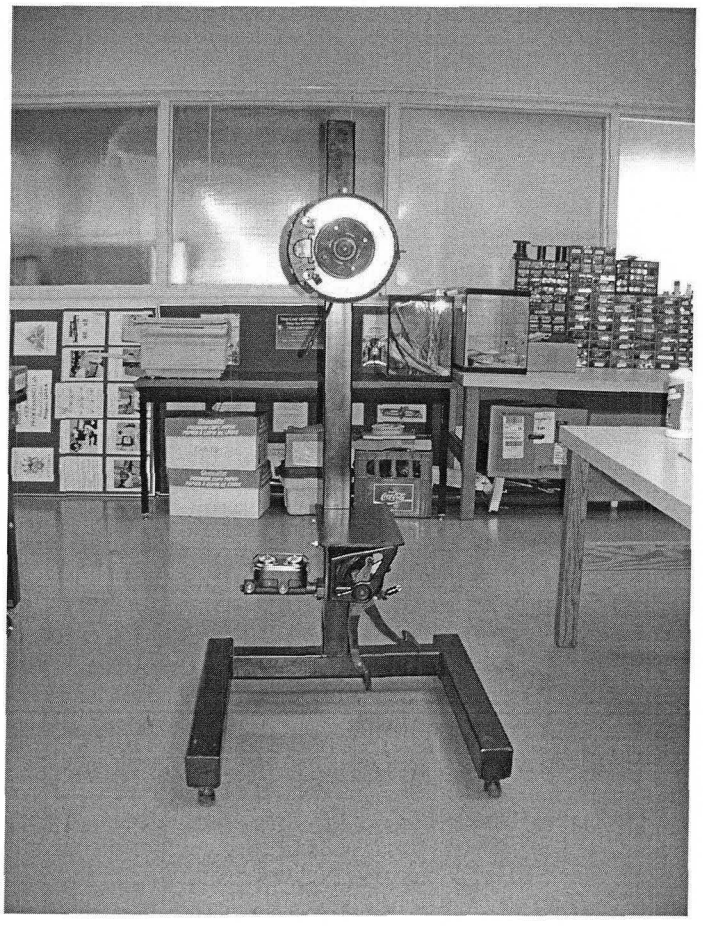

Figure 2.1: Laboratory setup of a hydraulic brake system

The individual components are pictured in greater detail in Figs. 2.2 (the master cylinder on the left and the pedal on the right), 2.3 (the disk and wheel hub on the centre and the calliper on the left edge of the disk) and 2.4 (the calliper). With reference to these pictures, the operation of a hydraulic disk brake system can be

summarised as follows: when the driver presses the pedal, a piston inside the master cylinder compresses the brake fluid. This increased pressure in the brake fluid propagates through the brake lines (pipes that extend from the master cylinder until each of the wheels) all the way to the calliper at each wheel. The calliper is located around the rotating disk and the increased fluid pressure pushes brake pads (made of a high friction material) against the surface of the disk. The resulting friction slows down the disk (and hence the wheel, given that they move together).

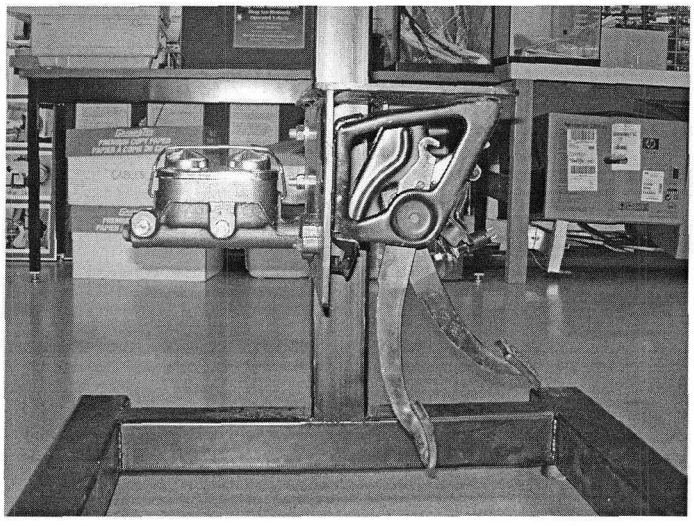

Figure 2.2: Detail of the brake pedal and master cylinder

Fig. 2.5 illustrates one of the disadvantages of hydraulic brakes that have been mentioned earlier: the need to have pipes (the "brake lines") taking the fluid from the master cylinder to each of the wheels. These add extra weight to the car and, more importantly, make the system more vulnerable: leaks can occur if the pipes (or the flexible hoses connecting the pipes to the callipers) are damaged or in the connections between the various components. In the lab setup, for example, excessive tightening of a bolt in a junction between two brake lines led to leakage of brake fluid.

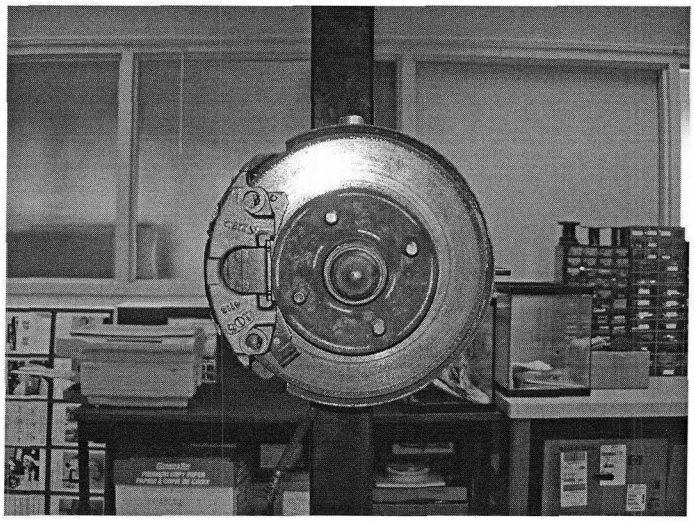

Figure 2.3: Detail of the disk and wheel hub

The hydraulic brake system was used to estimate the force required from the actuator pressing the pads. This was achieved by fitting a pressure gauge to the brake line, thus knowing the fluid pressure. The linear force F exerted by the fluid on the pads against the disk is then obtained from:

$$F = n \cdot p \cdot A \tag{2.1}$$

where n is the number of pistons (1 in the 1983 Chevrolet Chevette) p is the fluid pressure and A is the cross-sectional area of the calliper piston[1]. Fully pressing the brake pedal originates a fluid pressure of approximately 300 psi [2]. Hence, considering a fluid pressure of 300 psi and a piston area of 1 in^2 leads to an estimate of 300

[1] Approximately 1 in^2 for the 1983 Chevrolet Chevette

[2] Note that this value was obtained by applying normal pressure on the pedal (just as when driving a car). In the laboratory, it was possible to reach higher fluid pressures (in excess of 600 psi) by applying the maximum possible load on the pedal. Note that even if it was possible for a driver inside the car to press the pedal with such strenght, the braking performance would likely remain unchanged as once the wheels have locked further pressure does not make a difference.

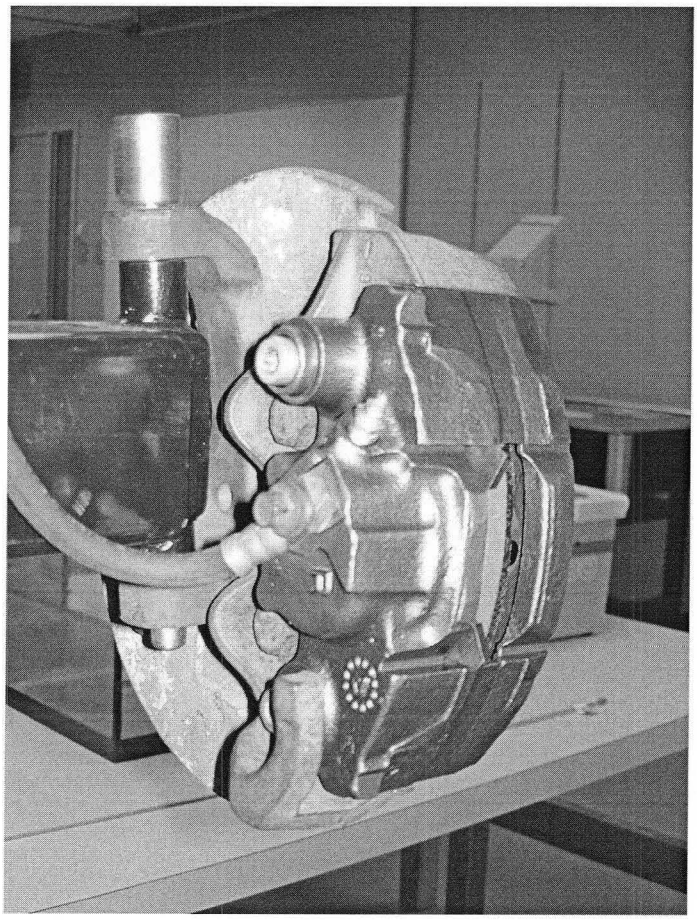

Figure 2.4: Close-up of the calliper

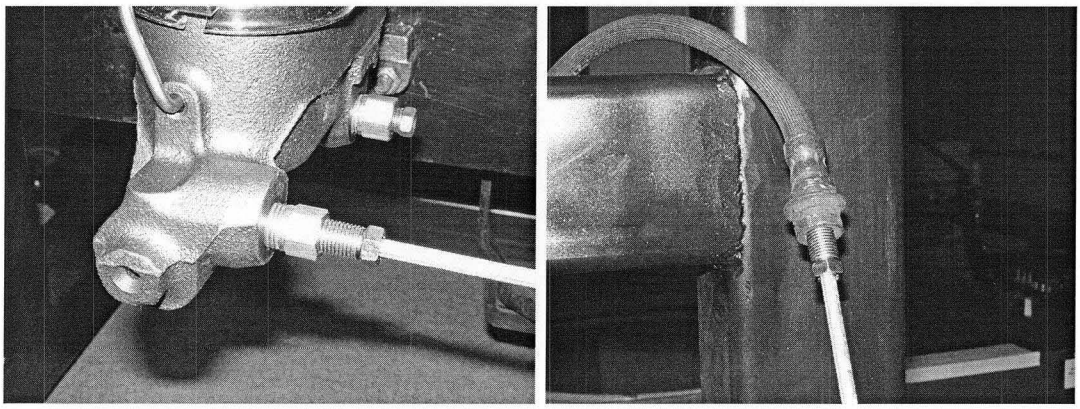

Figure 2.5: Vulnerability of the hydraulic brake lines

lb (approximately 1334 N) for the brake actuator force. The actuator must also be capable of overcoming the distance between the brake pads and the disk. This may be less than 1 mm with new pads but increases as the pads wear out. The thickness of the friction material on new brake pads for the Chevrolet Chevette system was measured to be 7 mm. The minimum allowable thickness of the brake pads depends on the model but generally varies between 0.3 and 0.5 mm. This means that in the present case the distance that the brake pads must travel to touch the disk increases by $7 - 0.3 = 6.7$ mm from new condition to maximum wear. Hence, the maximum travel that may be required of each brake pad is equal to $1 + 6.7 = 7.7$ mm, assuming a 1 mm travel for new brake pads.

2.2 Definition of Performance Targets

Cars were once considered to possess good braking if the stopping distance in metres did not exceed the square of one tenth of the initial speed in kilometres per hour [17]. For a car stopping from an initial velocity of 100 km/h, this would mean a braking

distance of $(100/10)^2 = 100$ m. Thankfully, the improvements in brake technology and tyre-road adherence have reduced this distance to under a half that value in the last decades. This requires important braking torques and the purpose of this section is exactly the determination of the torque associated with today's brake performance.

The first estimate of the required braking power can be obtained from the equations of motion. For a body subject to a constant acceleration, these can be written for the position x and velocity v at time t as:

$$x = x_0 + v_0 t + \frac{1}{2} a t^2 \tag{2.2}$$

$$v = v_0 + at \tag{2.3}$$

In the problem of determining the required braking power, knowing (from road tests) the braking performance of a car (i.e. the distance $\Delta x \equiv x - x_0$ travelled while going from a known initial velocity v_0 to a stop $v = 0$), the unknowns are the braking time t and acceleration a. To obtain these, the velocity equation can be rewritten as:

$$a = \frac{-v_0}{t} \tag{2.4}$$

Introducing this in the position equation yields:

$$\Delta x = v_0 t + \frac{1}{2} \frac{(-v_0)}{t} t^2 \tag{2.5}$$

and hence:

$$t = \frac{2\Delta x}{v_0} \tag{2.6}$$

This equation is used to obtain the braking time, after which Equation 2.4 gives the braking deceleration. Once the deceleration has been determined, the braking force

F can be obtained from Newton's second law which, for a body whose mass m does not change with time, becomes the well-known $F = ma$.

The value of F thus obtained is the sum of all external forces applied on the car. This includes not only the force produced by the brakes but also the contributions of other sources of resistance to the motion of the car, such as the aerodynamic drag, the friction between the tyres and the road and the resistance from the powertrain.

One way to know the fraction of this force coming from the actual brakes is to resort to experimental data. Another possibility is to carry out an analytical study of the contribution of each of the other sources of drag.

Lee [18] presents a graph with the percent contribution of the brakes to the total braking force, at a speed of 50 km/h, as a function of the deceleration rate. It is seen that the brakes are responsible for between 80 and 90% of the total braking force for decelerations greater than 0.2g. For less pronounced decelerations, the influence of the brakes is lower as the other sources of drag nearly suffice by themselves for slowing down the car.

However, that graph is only valid for a speed of 50 km/h. In order to determine the influence of the brakes on the total deceleration at different speeds, it is necessary to resort to the equations describing the other sources of drag. Starting with the aerodynamic drag D, which is given by [19]:

$$D = C_D \frac{1}{2} \rho v^2 S \tag{2.7}$$

where C_D is the drag coefficient (which depends on the shape of the car and can be obtained from wind tunnel tests), ρ is the density of air (1.225 kg/m^3 at 15 °C), v is the velocity of the car and S a reference surface (the frontal area of the car). Typical values for C_D and S for a relatively aerodynamic car are respectively 0.3 and 2.5 m^2.

At 50 km/h, this gives for the aerodynamic drag: $D = 0.3\frac{1}{2}1.225\left(\frac{50}{3.6}\right)^2 2.5 = 88.61$ N

ITI [20] suggests the following expression for the drag associated with the friction between the tyres and the road:

$$F_r = f_r \cdot m \cdot g \tag{2.8}$$

where f_r is the rolling friction coefficient, equal to 0.01 for an asphalt road, according to the same reference. Note that this is a somewhat simplistic approximation in that the tyre friction drag is assumed to be completely independent of the velocity. More sophisticated models are available. Namely, Volvo Powertrain Corp. proposes the set of expressions in Table 2.1 for the value of the rolling friction coefficient.

Table 2.1: Rolling friction coefficient

Type of tire	f_r
Radial-ply passenger car tyre	$0.0136 + 0.04 \times 10^{-6} \times (v \times 3.6)^2$
Bias-ply passenger car tyre	$0.0169 + 0.19 \times 10^{-6} \times (v \times 3.6)^2$
Radial-ply truck tyre	$0.006 + 0.23 \times 10^{-6} \times (v \times 3.6)^2$
Bias-ply truck tyre	$0.007 + 0.45 \times 10^{-6} \times (v \times 3.6)^2$

(The velocity v is in m/s)
Source: Volvo Powertrain Corp., as cited in [21]

Since we are looking at the application of MR brakes to passenger cars, the equations of interest are the first two in Table 2.1, and particularly the first one, given that virtually all cars currently produced are equipped with radial-ply tyres. It is interesting to note that any of the expressions for passenger cars give higher values for the rolling friction coefficient than the one suggested by ITI. This can be due to ITI having considered heavy vehicles or slippery roads (there is no mention to the vehicle type and the surface is only referred to as "asphalt road"). Due to the existence of such uncertainties in the value proposed by ITI and its simplistic nature, the first expression proposed by Volvo Powetrain Corp. will be used in the present work to

model the friction between the tyres and the road. This leads to the values of f_r (the rolling friction coefficient) indicated in Table 2.2 for different velocities. Analysis of Table 2.2 demonstrates that the tyre friction coefficient varies with speed albeit not too significantly.

Table 2.2: Rolling friction at different velocities

Velocity [km/h]	Rolling friction coefficient
0	0.0136
50	0.0137
100	0.0140
130	0.0143

The powertrain drag comprises mostly the engine braking but also some minor sources of friction throughout the transmission. Most data on engine braking is relative to the retarding torque achieved by the engines of heavy vehicles, where it is a fundamental contribution to maintain moderate speeds in long descents without overheating the conventional brakes[3]. A quantification of the engine braking deceleration independent of the vehicle is provided by [22]: the deceleration due to engine braking is comprised between 0.2 and 0.7 g at 50 km/h and between 0.4 and 1 g at 100 km/h, varying linearly for intermediate speeds, and hence it is seen that the maximum engine braking force is $F_e = 0.7 \times 9.8 \times m$ at 50 km/h and $F_e = 1.0 \times 9.8 \times m$ at 100 km/h.

The braking force contributed by the brakes (F_b) is then given by $F_b = F - D - F_r - F_e$. Once this has been determined, it is neccessary to know how much of this force is done at each wheel. A simplistic approach would be to assume that the braking force is divided equally among the four wheels and hence each brake would be responsible

[3]which would lead to failure due to brake fade

for one quarter of the total force required from the brakes. However, this is seldom true as it would likely result in locking of the rear wheels. Front wheel drive cars are typically designed to have a weight distribution of 2/3 in the front axle and 1/3 in the rear axle [23]. This means that applying the same maximum braking power to both axles would result in the rear wheels locking first due to the lower adherence of the rear tyres (caused by a lower loading). This is an extremely dangerous situation that may entrain loss of control of the car. This phenomenon is actually aggravated by the fact that braking induces a twisting moment that increases the loading on the front axle and reduces the loading on the rear axle[4]. This shift in the load distribution during braking means that even on rear wheel drive cars (which have a more balanced weight distribution between both axles) the braking power to the rear wheels must be lower. Therefore, the problem becomes that of determining what is the ratio of braking at each axle, so that the maximum braking effort required from each of the front wheels (those that contribute most to the braking) is determined.

Until recently, at maximum decelerations approximately 90% of the braking came from the front wheels [18, 24]. This impressive percentage is due to the fact that most cars have had drum brakes mounted on the rear wheels. Since these are significantly less progressive than disk brakes and therefore more prone to premature locking, it has been necessary to fit most car models with extremely conservative proportioning valves that keep the brake fluid pressure to the rear wheels at a slight value at all times [25]. Without this constraint, the contribution of the rear brakes could be significantly higher. Lee [18] presents a plot of the ideal braking distribution: it is seen that the ideal value for the rear braking ratio r_b (the percentage of the total braking effort contributed by the rear brakes) decreases linearly from approximately 45% at very low decelerations ($a \approx 0$) to approximately 33% at decelerations of 0.85g.

[4]That's why on heavy braking the front of the car plunges noticeably

Therefore, the front axle is responsible for up to 67% of the force which means that each of the front brakes must be able to develop 33% of the total force required from the brakes.

Finally, once the braking force F at each wheel has been determined, the torque is obtained from:

$$T = F_{each\ wheel} \times r \tag{2.9}$$

where r is the radius of the wheel and tyre assembly.

We now have all the necessary expressions to determine the braking torque required to produce the deceleration associated with typical braking distances. These can be obtained from car tests. Results from such tests are compiled in [24], mostly in the form of the braking distances from an initial speed of 130 km/h. While this information is valuable, it is not the most useful in that the determination of the residual sources of drag associated with this speed involves a greater uncertainty than would be achieved for an initial speed of 100 km/h. Also, the statistical treatment of the data is hampered by the high deviation from one car model to another: values range from 63 to 91 metres and a meaningful analysis would require computing the braking force associated with each case taking into account the weight of each of the 500 cars featured in the study. Ideally, a case study should be analysed (preferably for an initial speed of 100 km/h) - and this is presented in the same article [24], with an in-depth study of one of the cars with best braking performance - the Citroën Xantia 2.1 TD. The main characteristics of this car are presented in Table 2.3.

Based on the car data in Table 2.3 and on the braking test results (distance of 46 m to go from 100 km/h to a full stop [24]), it is possible to determine the performance of the braking system using the equations introduced throughout the present section. The results are summarised in Table 2.4. Based on the calculations summarised in Table 2.4, a target of 1010 Nm was set for the braking torque of the present project.

Table 2.3: Characteristics of the case study car

Car model and version	Citroën Xantia 2.1 TD (1998)
Kerb weight [a] [b]	1381 kg
Tyres [b]	205/60VR15 [c]
Wheel+tyre radius	313.5 mm [d]
Aerodynamic drag coefficient (Cx) [e]	0.34
Frontal area (S) [e]	2.07 m^2
S·Cx [e]	0.69 m^2

[a]Note that since 1996, the definition of kerb weight used in Europe includes the empty weight of the car plus 68 kg for the driver and 7 kg for luggage [26]

[b]Source: [27]

[c]Hence, the width of the tyres is 205 mm, the height is 60% of the width and the wheel diameter is 15 in

[d]$205mm \times 60\% + \frac{1}{2}15in \times 25.4mm/in$

[e]Source: [28]

Table 2.4: Brake performance summary

Braking distance	46 m
Braking time	3.3 s
Deceleration[a]	8.4 m/s^2
Total braking force	11600 N
Braking force contributed by the brakes	9744 N
Braking force contributed by each front brake	3216 N
Braking torque contributed by each front brake	1008 Nm

[a]The braking distance indicated above was calculated with the help of a device measuring the deceleration [24] and hence it is the distance between the moment the pressure is available at the brakes and the moment the car comes to a full stop.

Another quantity of interest is the braking power, i.e. the power dissipated by all the sources of drag in the car and which equals the rate of change of the car's kinetic energy (E_c):

$$P = \frac{dE_c}{dt} \tag{2.10}$$

For braking from 100 km/h to a full stop in 3.3 s (v. Table 2.4), this leads to a braking power of 161.5 kW.

Finally, it was important to compare the values estimated above with experimental data. To this end, the author had the opportunity to witness brake testing conducted in a Mercedes-Benz 190E at CIMA in Oeiras, Portugal, which indicated a braking force of 2.5 kN in each of the front wheels and 1.8 kN in each of the rear wheels. It is seen that in this case the front wheels are responsible for approximately 58% of the braking, reasonably less than the 66% that were mentioned above. One contributing factor to this difference may lie in the fact that this is a rear wheel drive car, with a greater weight supported by the rear axle and hence with more braking power in that axle. Nevertheless, this means that the assumption that 66% of the braking is brought by the front wheels is a conservative one in that it led to the design of brakes with a greater braking force (so that each accounts for half of 66% of the total brake contribution) than is actually necessary (half of 58%). Naturally, this distribution of brake forces among the two axles is only possible in a car with rear wheel disk brakes (as is the case of the Mercedes-Benz 190E), as it would lead to premature locking of the brakes in a car fitted with drum brakes in the rear wheels.

2.2.1 Driving Patterns

So far, the study has focused on the maximum braking torque, i.e. on the worst case scenario. That was necessary in order to determine the extreme performance

required of the new system. However, in terms of determining the actual demands placed on the brakes and to plan their resistance and longevity, it becomes necessary to somehow take into account its real use. While it is naturally impossible to model the full driving history experienced throughout the life of the brakes, it is possible to devise a representative circuit that includes the various types of demands imposed on the brakes in typical proportions and thus models real life conditions. Such circuits have been created in the past following studies of everyday driving habits of common drivers. The most popular drive cycle in use today is the FTP 75 (Federal Test Procedure) based on the earlier FTP 72. Their significance is asserted by the fact that they have been chosen by other countries as the basis for their own test procedures: A10 or CVS (Constant Volume Sampler) in Sweden, ADR (Australian Design Rules) 27 and ADR 37 [29].

The FTP 75 cycle has recently been complemented by two other cycles [30]: US06 to model a more agressive driving style and SC03 to simulate the effects of air conditioning on engine loads, an issue of particular interest for the problem of modelling pollutant emissions, but which does not affect the braking patterns.

For the analysis of the demands placed on the braking systems, the most relevant cycles are the FTP 75 and US 06. The key figures of both are presented in Table 2.5. The speed vs. time curves for both cycles are included in Appendix A.

Table 2.5: Main characteristics of the FTP75 and US06 driving cycles

	FTP75	US06
Description	Urban driving	Aggressive driving
Distance travelled	11.04 mi (17.77 km)	8.01 mi (12.8 km)
Duration	1874 s	596 s
Average speed	21.2 mph (34.1 km/h)	48.4 mph (77.9 km/h)

Source: [29]

2.3 Smart Materials

Section 1.1 explained the shortcomings of current braking systems and how they could be addressed through the use of smart materials. It was therefore fundamental to conduct a literature review of the various candidate materials for application to car brakes in order to gain a complete knowledge of the properties and limitations of each material. This literature review and the material choices that emerged from it are presented in the following sections.

2.3.1 Piezoelectric and Electrostrictive Materials

The first approach considered involved the use of piezoelectric or electrostrictive materials to actuate the brake pads. These materials exhibit a deformation when a voltage is applied[8] and thus constitute a simple way of pressing the brake pedals against the disk based solely on an electric signal. They are also extremely fast, which is particularly important for the present application.

Three different solutions involving piezoelectric or electrostrictive materials were considered:

- Direct actuation;
- Inchworm linear actuation;
- Rotary actuation.

The first consists of simply having a piezolectric or electrostrictive stack pushing the brake pads against the disk. When voltage is applied to the stacks, the material

[8] And vice-versa, making them suitable not only for actuators but also for sensors as a voltage is produced when a deformation is applied.

expands thus pushing the brake pads; as soon as the voltage is removed the stack returns to its original shape, pulling the pads away from the disk. This is the simplest concept but the least flexible one, being directly dependent on the performance of the stack.

Given that piezoelectric and electrostrictive materials are only capable of very small elongations each time the voltage is applied but exhibit very quick response, concepts have been developped in the past to augment the total displacement by applying many voltage steps in sequence so that the total displacement is the sum of all the individual elongations. Hence, the inability of the material stack to stretch as much as desired in one move is bypassed by having it scrambling rapidly. This concept mimics the motion of an inchworm and thus such systems became known as "inchworm actuators". A schematic illustration of the operation of an inchworm actuator can be found in [31].

Another concept that takes advantage of the fast response of these materials is the rotary actuator: it is possible to have an arrangement with three actuators whereby one produces the actual movement, one clamps the shaft to the deflection mechanism and another clamps the shaft to a static component. This mechanism is presented in [32].

Given that the goal of the actuator now being design is to produce linear movement (pushing the brake pads against the disk), the first two options were initially considered.

Direct Actuation

The first concept that was considered consists of using a stack of piezoelectric or electrostrictive materials to push and pull the brake pads. This is the simpler design

as no extra components are required and no moving parts are involved, but the performance is limited by the capabilities of the material being used.

It must be noted that the performance capabilities of piezoelectric and electrostrictive actuators are represented by the maximum free displacement (Δx_{free}) and the maximum blocked force (F_{max}). The former is the maximum displacement achieved by the actuator without any force applied; the latter is the maximum force exerted by the actuator without displacement. Therefore, the actuator is capable of a displacement Δx_{free} at zero force or a force F_{max} at zero displacement. It is impossible to achieve the maximum force and the maximum displacement simultaneously. Rather, the relation between the output force and displacement of a piezoelectric or electrostrictive actuator is given by a linear relation, as exemplified in Fig. 2.6.

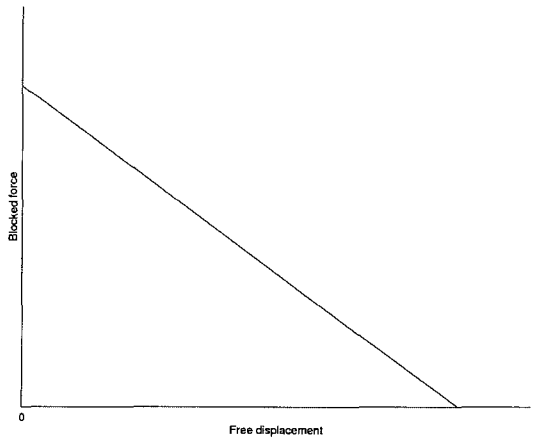

Figure 2.6: Typical piezoelectric actuator force vs. displacement performance

A study of suitable piezoelectric actuators currently available was carried out and the most promising ones are summarised in Table 2.6. Note that these are only intended to provide a quick overview of the capabilities of today's materials. Given that these actuators have different dimensions and operating voltages, their performances

are not directly comparable - for example, the actuator of Physik Instrumente (PI) is seen to have a much greater force than any of the other actuators but this is obtained for 1000 V unlike the other materials, that require maximum voltages of 150 or 300 V.

Table 2.6: Overview of piezoelectric and electrostrictive materials performance

Material	Blocked force [N]	Free displacement [μm]
Adaptronics APA230L [a]	1350	236
Cedrat Technologies DPA80 [b]	3500	80
PI P-247 [c]	30000	120
Sensortech ASM-10 [d]	10000	40

[a]Source: [33]
[b]Source: [34]
[c]Source: [35]
[d]Source: [36]

Recall from section 2.1 that it is estimated that an actuation force of 1334 N is required to produce a force comparable to that of today's hydraulic brakes and that each brake pad must travel up to 7.7 mm (and hence a total displacement of up to $2 \times 7.7 = 15.4$ mm, as the disk brake system comprises one brake pad on each side of the disk). Clearly this displacement is not immediately available from any of the actuators in Table 2.6. The possibility of using a lever system to increase the displacement (albeit reducing the force) was therefore studied. However, it is important to keep in mind that due to conservation of energy, the total work (force times displacement) output from the lever is always limited by the input work. That is, in order to achieve a higher displacement, the force produced by the actuator would have to be sacrificed. Also, the size and cost of these high-force and high displacement actuators constitutes a major obstacle to their implementation in common cars. This leads to the conclusion that direct actuation with existing piezoelectric or electrostrictive materials is not possible.

In order to overcome the limitations associated with the direct use of piezolectric or electrostrictive stacks as actuators, different concepts have ben studied over time that take advantage of the very fast response of these materials. Two such configurations are addressed in the following sections.

Inchworm Actuator

One of the possible solutions in order to achieve higher displacements and forces is the inchworm concept, already described. Burns [37] has proposed and built an inchworm actuator based on electrostrictive elements (Sensortech BM600). Recall that the main advantage of the inchworm concept lies in its dynamic behaviour, i.e. repeating individual steps in order to obtain a considerable displacement. Hence, it is important to determine its dynamic response, which is controlled by the frequency of the electrical signal supplied to the piezoelectric stacks. The minimum time for charging and discharging each of the electrostrictive stacks in Burns' configuration is 0.075 s. Given that each inchworm step comprises a sequence of 6 clamp or release operations, the minimum time per cycle is of $6 \times 0.075 = 0.45$ s. With a maximum displacement of 389.1 μm per cycle, this leads to a maximum velocity of $389.1/0.45 = 864.7$ μm/s. Given that a total distance of up to 15.4 mm must be travelled by the brake pads before they are in contact with the disk[9] (and only then exerting a braking force), the proposed actuator must be capable of travelling such a distance in the shortest amount of time - if a maximum acceptable actuation time of 100 ms is defined[10], this leads to a required speed of the actuator equal to $15.4/0.1 = 154$ mm/s. Unfortunately, this required velocity is 2 orders of magnitude above that achievable with currently existing materials. In order to bridge the gap between the

[9] v. section 2.3.1

[10] Recall that hydraulic brakes exhibit a lag of 200 to 300 ms before full braking power is available and the present system is intended to present a much faster response

current actuator speed of 864.7 μm/s and the minimum required value of 154 mm/s, materials with faster response times and greater displacements must be devised.

Rotary Actuator

While the objective of the desired brake actuator is to press the pads against the disk (linear motion), nothing precludes the use of a rotary actuator connected to a rotary-linear converter[11]. In fact, this is the concept used by Delphi in the electric actuator mentioned in section 1.2.

A study of the piezoelectric rotary actuator proposed and built by Gursan [32] was therefore conducted. However, once more the performance limitations of current piezoelectric materials are apparent. The torque of the system reaches a maximum of 13 Ncm for a speed of approximately 0.22 rpm. In turn, the maximum speed is of 2 rpm with a torque of 1 Ncm. This performance is far from suitable for the desired brake actuator and highlights that the advantage of piezoelectric and electrostrictive actuators lies in the precision of actuation rather than on brute force capabilities.

General Considerations

It was seen in the previous sections that present actuators based on piezoelectric or electrostrictive materials lack the performance required for the present application. Also, it must me noted that piezoelectric and electrostrictive materials will exhibit a deformation only if the applied electric potential has a frequency within the operating range of the particular material being used. Static applications, such as that of the brake actuators where it is intended to keep the material in its excited condition for up to several seconds, are not ideal.

[11] Such as a rack and pinion system

With the above considerations in mind, it must be concluded that neither piezoelectric nor electrostrictive materials currently constitute a viable option to actuate the brake pads. The potential for application of such materials to car brakes in the future presents one more motivation for the development of new, improved materials. However, the difficulty of these materials in dealing with constant voltage and displacement seems to be intrinsic to their nature and may be more difficult to overcome.

In short, the study described in this section led to the conclusion that piezoelectric or electrostrictive materials may become suitable actuators to car brakes once their performance (force and displacement) and behaviour at very low frequencies (close to 0 Hz) are improved. Presently, other options must be sought for the electric brakes and a study of another interesting family of smart materials is conducted in the following section.

2.3.2 Electrorheological and Magnetorheological Fluids

Electrorheological and magnetorheological fluids are liquids whose viscosity can be controlled by applying an electric or magnetic field, respectively. Electrorheological fluids are obtained by introducing semiconducting solid particles (such as corn starch) in a dielectric carrier liquid (such as silicone oil) [38]. Magnetorheological fluids are in turn obtained by adding ferrous (magnetically permeable) particles (typically 3 to 5 microns) to a non-magnetic carrier liquid [39].

In order to evaluate the possibility of using either of these materials, it is important to know the characteristics they possess. Table 2.7 [40] presents an overview of the properties of both types of materials, providing a basis for comparison of the relative merits of each.

Table 2.7: Magnetorheological versus electrorheological fluids

	MR	ER
Max. Yield Stress	50-100 kPa	2-5 kPa
Max. Field	250 kA/m (limited by saturation)	4 kV/mm (limited by breakdown)
Viscosity	0.1 - 1.0 Pa.s	0.1 - 1.0 Pa.s
Operable Temp. Range	-40 to +150 °C (limited by carrier fluid)	+10 to +90 °C (ionic,DC) -25 to +125 °C (non-ionic,AC)
Stability	Unaffected by most impurities	Cannot tolerate impurities
Response Time	< milliseconds	< milliseconds
Density	3-4 g/cm^3	1-2 g/cm^3
$\mu_p/\tau^2_{y(field)}$	10^{-10} to 10^{-11} s/Pa	10^{-7} to 10^{-8} s/Pa
Max. Energy Density	0.1 Joule/cm^3	0.001 Joule/cm^3
Power Supply (typical)	2-25 V @ 1-2 A (2-50 watts)	2-5 kV @ 1-10 mA (2-50 watts)

Source: [40]

From this comparison between ER and MR fluids, it can be seen that the latter present several important advantages regarding the application to car brakes: higher yield stress, greater energy dissipation per unit volume, wider operating temperature range. The lower density of the ER fluids does not bring an important advantage, given that the fluid volume to be used is reduced and hence the weight difference is negligible. In reality, the lower energy density of ER fluids implies that for the same application a greater volume is required than would be using MR fluids and this increase in volume more than offsets the gain associated with the lower density.

For all these reasons, it is not surprising that MR fluids have enjoyed a greater popularity and already have numerous applications, presented in detail below. This brings an important advantage in terms of existing know-how regarding the use of MR fluids under real-life conditions. It will be seen in Section 6.4.1, when the longevity of MR fluids is discussed, that it was indeed the knowledge resulting from life-cycle testing of products based on MR fluids that identified and permitted to overcome serious longevity limitations of early fluids. Due to the much lesser extent of study on ER fluids, it is possible that similar problems exist but have not yet been identified, which might result in real-life experiments of an hypothetical ER brake demonstrating the inability of existing fluids to achieve the desired life.

Presently, magnetorheological fluids are commercially available from two different suppliers:

- Lord Corporation, from Cary (NC, USA) has an extensive experience in the manufacture of MR fluids and applications;
- Liquids Research, from Bangor (Wales, UK) has also been synthesising MR fluids for several years and has recently partnered with BASF, one of the worldwide leaders in chemical engineering and magnetic products.

Existing applications

Altough ER and MR fluids have been known for several decades[12], actual applications of these technologies were until very recently hampered by limitations in the formulation of the fluids. On ER fluids these still subsist in part (which explains the far lower number of applications) whereas on MR fluids these were mostly related with stability and longevity issues. These difficulties are described in detail in Section 6.4.1, along with the solutions found to overcome them.

This means that altough systems based on ER and MR fluids were first designed in the 1940s ([42, 43] and [44, 45] respectively), only in recent years have there been successful commercial applications (and for the time being, only of MR fluids). Altough the research on ER fluids is not as intense as that in their MR counterparts, it nevertheless includes study and applications in a variety of areas from vehicle suspension dampers [46] to haptic instrument controls[13] [47].

It is however clear that in the last decade MR fluids were the subject of a far greater interest, which translated into a higher number of industrial applications. A good overview of implementations of MR fluids in the 1990s is found in [48]. Since then, the areas covered by MR technology vastly expanded and the most promising uses are described in the following paragraphs.

One application of particular interest for the present work is a small MR brake[14] developped by Lord Corporation for selected industrial uses [49, 50, 51]. MR brakes have also been applied to static bicyles by Nautilus Corp. [52, 53]. This past work on small MR brakes constitutes a good departure point for the present development of

[12] the description of an early MR fluid made of mercury, steel dust and graphite can be found in [41]

[13] These are systems designed to improve the man-machine interface by providing the operator with some feedback

[14] Diameter: 96.6 mm; Maximum torque: 5.6 Nm

an MR brake for use in cars, in that concepts have already been proposed and some mathematical models formulated. However, it will be seen in the next chapter that these models cannot be applied to a car MR brake, due to the greater complexity associated with the need to obtain a very high braking torque from a small system. Furthermore, the application to car brakes brings different difficulties associated with the need to provide a high braking torque in a small space and volume with a limited power supply. The steps necessary to overcome these challenges will be presented in the following chapters, where it is seen that the design of an MR brake for cars cannot be obtained simply by scaling-up the small Lord MR brake.

Another application of magnetorheological fluids to vehicle systems is in dampers. The first commercial application happenned in January 2002 in the Cadillac Seville STS [54]. Furthermore, MR dampers are not limited to car suspensions: smaller systems have been used to reduce vibrations on the driver's seat of heavy vehicles and on household washing machines [55]. On the other hand, larger damper units are finding their way to civil engineering applications, namely in seismic response control [56].

Magnetorheological fluids have also been successfully applied to clutches in different configurations: radial double-plate system [57]; input plate enclosed within a housing serving as output plate [58].

The tremendous possibilities of control of MR fluid-based devices, the high yield stresses that these can exhibit and the low voltage required make this technology very attractive for application to bio-mechanics, such as in prostethics [59, 60], rehabilitation devices [61] and cancer treatment [62].

Finally, the very fast response and good precision of magnetorheological fluids have made them suitable for use in valves as a means to control pressure in hydraulic circuits [63].

MR fluid theory

In order to properly characterise a system based on MR fluids, it is necessary to understand their behaviour. Whereas other fluids of interest in engineering - including air and water - exhibit a Newtonian behaviour, i.e., the shear stress (τ) is a linear function of the shear strain rate $\dot{\gamma}$ expressed as $\tau = \mu\dot{\gamma}$ where μ is the viscosity of the fluid, MR fluids are non-Newtonian in that their viscosity is not a constant. Instead, the shear stress of the MR fluid is the sum of two effects: an electro-stress (τ_e, due to the applied magnetic field) which does not vary with the strain rate and a contribution similar to that of a Newtonian flow ($\tau_p = \mu_p\dot{\gamma}$ where μ_p is the plastic viscosity, assumed equal to the no-field viscosity obtained from experimental tests of the fluid with no magnetic field applied). This viscoplastic model, comprising a strain rate-independent stress and a linear function of the strain rate, is known as the Bingham model. Experimental results support the use of the Bingham model to describe the behaviour of MR fluids [64]. Hesselbach and Abel-Keilhack [65] have suggested that the Herschel-Bulkley model is more adequate to describe the behaviour of MR fluids particularly for systems where there is a large interval of shear rate values. However, this conclusion is based on the comparison of the Bingham and Herschel-Bulkley values with a highly non-linear experimental curve of shear stress vs. shear rate for Lord Corporation's MRF-132LD fluid. For the fluids now being considered (Lord Corporation's MRF-132AD and MRF-241ES) the experimental values of shear stress vs. shear rate form a nearly linear curve, well approximated by the Bingham model.

Table 2.8 summarises the properties of the two fluids most suitable for this particular application.

The viscosity values are those without a magnetic field applied. It is important to note that they cannot be directly compared between these two materials since

Table 2.8: Fluid properties for Lord's MRF-132AD and MRF-241ES

Fluid	MRF-132AD	MRF-241ES
Base Fluid	Hydrocarbon	Water
Operating Temperature	-40 to +130 °C	-10 to +70 °C
Density	3.09 g/cc	3.86 g/cc
Specific Heat @ 25°C	0.80 J/g°C	0.94 J/g°C
Thermal Conductivity @ 25°C [a]	0.25-1.06 W/m°C	0.85-3.77 W/m°C
Flash Point	>150°C	>93°C
Viscosity	0.09 ± 0.02 Pa.s [b]	10.8 ± 1.5 Pa.s [c] 2.2 ± 0.4 Pa.s [d]

Source: [66, 67]

[a]Values were calculated with and without magnetic fields applied. Thermal conductivity of MR fluids is not strongly dependent on temperature from -30°C to 100°C.

[b]Calculated for slope between 800 1/s and 500 1/s at 40°C

[c]@ 10 1/s Shear Rate

[d]@ 50 1/s Shear Rate

they were calculated at different shear rates. If the shear stress vs. shear rate curve is known, the viscosity at any shear rate is given by the slope of the curve at that particular point. This can be done for the two fluids (the product bulletins [66, 67] include this curve) and the viscosity values thus obtained are very close for both materials. Recall that the Newtonian component of the shear stress is given by $\tau_p = \mu_p \dot{\gamma}$. This means that at low speeds ($\dot{\gamma} \approx 0$), the Newtonian contribution[15] to the shear stress is low. However, at higher speeds this contribution can become significant if μ_p (the no-field viscosity) is high. For this reason, the value of μ_p at high shear rates is of particular importance to the design of the MR brake. From the shear stress vs. shear rate curves indicated above, it is seen that $\mu_p \approx 0.09 Pa \cdot s$ for shear rates greater than 100 s^{-1}.

[15]Which we want to keep low since it is always exerted regardless of whether or not the brake pedal is pressed

Chapter 3

Preliminary Design

3.1 MR Brake Concept

As it was seen in section 2.3.2, electrorheological and magnetorheological fluids possess the property of changing their viscosity as an electric or magnetic field is applied. This changing viscosity is in turn proportional to the friction exerted on any body moving within the fluid. Hence, the application to brakes consists of having a disk immersed in MR fluid. The braking torque exerted on the disk is then controlled by the magnetic field applied to the fluid. This requires the inclusion of an electromagnet, a permanent magnet or a combination of both. However, the magnetic field is to vary according to the driver's pressure on the brake pedal and while it would be possible to achieve that using only a permanent magnet, that solution would require moving parts and result in a complex system. Hence, an electromagnet will be used to produce the magnetic field.

Bydoń [68] noted that the friction torque T of an MR brake is influenced by two distinct contributions: one due to the viscosity of the carrier liquid (T_μ) and another

due to the magnetic field effect (T_B) and proposed the following relations:

$$T = T_\mu + T_B \tag{3.1}$$

$$T_\mu = \frac{N\mu\pi f}{2h}\left(R_z + R_w\right)^2\left(R_z^2 - R_w^2\right) \tag{3.2}$$

$$T_B = \frac{N\pi H^\beta}{2h}\left(R_z + R_w\right)\left(R_z^2 - R_w^2\right) \tag{3.3}$$

where R_z and R_w are the outer and inner radii of the brake disk, respectively; N is the number of surfaces in contact with the MR fluid (2 for 1 disk with fluid on both sides, 4 for 2 disks, etc.); f is the frequency of rotation of the shaft; H is the magnetic field strength; β is a constant of the particular MR fluid being used, that relates the yield stress with the magnetic field; h is the thickness of the fluid gap; and μ is the dynamic viscosity (equal to the product of the density ρ by the kinematic viscosity ν).

Several problems can be identified in this mathematical model of an MR brake:

- The expression for T_μ uses a constant value of H which is only possible if the magnetic field is nearly uniform throughout the MR fluid. While this may be an acceptable approximation in Lord Corporation's MR brake studied by Bydoń, which has reduced dimensions and low torque, it may introduce a significant error in the present problem of much larger dimensions with parts operating close to magnetic saturation.

- The term H^β represents the shear stress due to the magnetic field effect on the MR fluid (τ_e). The value of β is chosen to perform a best fit of the τ_e vs. H curve for the MR fluid being used. However, this expression (H^β) does not allow accurate approximations of the τ_e curve for all MR fluids. Using an approximation of the form kH^β (where the parameter k has units of Pa·m/A

and is determined in conjunction with β from the τ_e vs. H curve of the fluid) produces much better approximations[1].

To address these difficulties, a new mathematical model for MR brakes is required. To obtain this new, more general model, it is necessary to depart from the basic equations describing the phenomena that occur in an MR brake. Consider an elemental area dA on the lateral surface of the disk (which is subject to shear stress τ). The elemental force on this area due to the shear stress is given by:

$$dF = \tau dA \tag{3.4}$$

The corresponding torque is then given by:

$$dT = \tau r dA \tag{3.5}$$

where r is the distance from the centre of rotation to the elemental area dA. Since we are analysing a disk, it is convenient to change to polar coordinates. Recall that the transformation from Cartesian to polar coordinates is described by $dA \equiv dxdy = rdrd\theta$ where the inclusion of r results from the Jacobian of the coordinate transformation. Hence:

$$dT = \tau r^2 drd\theta \tag{3.6}$$

The total torque T exerted on the disk's surface is the sum of all the elemental contributions dT:

$$T = \int_0^{2\pi} \int_{r_w}^{r_z} \tau r^2 drd\theta \tag{3.7}$$

This expression gives the torque exerted by the friction of the MR fluid on each disk

[1]This alternative expression also addresses an inconsistency of units in Bydoń's model

surface. Therefore, the total braking torque of the MR brake is:

$$T = N \int_0^{2\pi} \int_{r_w}^{r_z} \tau r^2 dr d\theta \tag{3.8}$$

where N is the number of disk surfaces in contact with the fluid. For an axisymmetric problem, this can be simplified to:

$$T = 2\pi N \int_{r_w}^{r_z} \tau r^2 dr \tag{3.9}$$

In the above expression for the torque, r_w and r_z are known geometrical characteristics of the disk (the same as in Bydoń's model introduced earlier) and τ is given by the Bingham model (refer to section 2.3.2):

$$\tau = \tau_e + \tau_p \tag{3.10}$$

Hence:

$$T = 2\pi N \int_{r_w}^{r_z} (\tau_e + \tau_p) r^2 dr \tag{3.11}$$

τ_e can be related to the magnetic field by an approximation of the form $\tau_e = kH^\beta$ as discussed earlier in this section. As for τ_p, it has been seen in 2.3.2 that it is a function of the no-field viscosity of the fluid (μ_p) and the shear rate $\dot{\gamma}$: $\tau_p = \mu_p \dot{\gamma}$, where $\dot{\gamma}$ is the velocity gradient near the wall of the disk, i.e., $\dot{\gamma} = dv/dn$ with n the direction normal to the wall. For a narrow fluid gap, it is possible to consider a linear velocity distribution and hence:

$$\dot{\gamma} = \frac{dv_z}{dy} \approx \frac{v_{z,rotor} - v_{z,stator}}{h} = \frac{v_{z,rotor}}{h} \tag{3.12}$$

where $v_{z,rotor}$ is the tangential velocity of the rotor and h is the thickness of the fluid

gap. This leads to the following expression for the braking torque:

$$T = 2\pi N \int_{r_w}^{r_z} \left(kH^{\beta} + \mu_p \frac{v_{z,rotor}}{h} \right) r^2 dr \tag{3.13}$$

Finally, the linear velocity can be replaced by the angular velocity by means of the relation $v = \omega r$, yielding:

$$T = 2\pi N \int_{r_w}^{r_z} \left(kH^{\beta} + \mu_p \frac{\omega r}{h} \right) r^2 dr \tag{3.14}$$

We now have an expression that gives the torque as a function of the magnetic field H and the angular velocity of the wheels ω. All the other values are known characteristics of the MR brake geometry (N, r_w, r_z, h) or of the fluid (k, β, μ_p). Furthermore, this equation can be divided into the two components T_μ and T_B associated with the plastic viscosity and the magnetic field, respectively, using the relation (3.1). This leads to:

$$T_\mu = 2\pi N \int_{r_w}^{r_z} \mu_p \frac{\omega r}{h} r^2 dr \tag{3.15}$$

$$T_B = 2\pi N \int_{r_w}^{r_z} kH^{\beta} r^2 dr \tag{3.16}$$

An interesting exercise consists in manipulating this expression so that it can be readily compared with the model proposed by Bydoń. Considering first the contribution of the no-field viscosity, we have:

$$T_\mu = 2\pi N \int_{r_w}^{r_z} \mu_p \frac{\omega r}{h} r^2 dr \tag{3.17}$$

The quantities μ_p, ω and h are independent of the radius, which leads to:

$$T_\mu = 2\pi N \frac{\mu_p \omega}{h} \int_{r_w}^{r_z} r^3 dr = 2\pi N \frac{\mu_p \omega}{h} \frac{r_z^4 - r_w^4}{4} = \frac{N\pi\mu_p\omega}{2h} \left(r_z^4 - r_w^4 \right) \tag{3.18}$$

Finally, substituting $2\pi f$ for ω results in:

$$T_\mu = \frac{N\pi^2\mu_p f}{h}\left(r_z^4 - r_w^4\right) \tag{3.19}$$

Similar manipulation can be carried on the equation for the contribution of the magnetic field on the torque:

$$T_B = 2\pi N \int_{r_w}^{r_z} kH^\beta r^2 dr \tag{3.20}$$

This is the equation that exhibits the most significant differences given the approach described earlier (the shear stress is approximated by kH^β instead of H^β and H is allowed to change over the radius). If H was assumed to be constant, the following expression would be obtained:

$$T_B = \frac{2\pi N k H^\beta}{3}\left(r_z^3 - r_w^3\right) \tag{3.21}$$

Both the equations proposed by Bydoń and those now obtained are summarised in Table 3.1 for quick comparison.

It has been seen that MR fluids present a nearly linear relation between the yield stress and the magnetic field for a very large interval of magnetic field intensities. Hence $\beta = 1$ may be used for both fluids being considered. The value of k is simply the slope of the yield stress vs. magnetic field curve and is equal to approximately 0.269 Pa·m/A (35kPa @ 130kA/m) for MRF-132AD and 0.476 Pa·m/A (50kPa @ 105kA/m) for MRF-241ES.

The previous equations give the braking torque as a function of the magnetic field H experienced by the fluid. However, the user does not directly control H_f but rather the current I fed to the coil, which in turn induces the magnetic flux in the circuit.

Table 3.1: Comparison of different models for the components of the torque produced by an MR brake

Term	Bydoń[a]	Proposed new model
T_μ	$\frac{N\mu\pi f}{2h}\left(R_z+R_w\right)^2\left(R_z^2-R_w^2\right)$	$\frac{N\mu\pi^2 f}{h}\left(r_z^4-r_w^4\right)$
T_B	$\frac{N\pi H^\beta}{2h}\left(R_z+R_w\right)\left(R_z^2-R_w^2\right)$	$\frac{2\pi N k H^\beta}{3}\left(r_z^3-r_w^3\right)$ [b] $2\pi N\int_{r_w}^{r_z} kH^\beta r^2 dr$ [c]

[a]Source: [68]
[b]For H independent of the radius
[c]For H function of the radius

The question then arises as to whether it is possible to obtain an expression relating the magnetic field H_f with the current I. The search for such a solution must depart from the equations of magnetostatics [69]:

$$N \cdot I = \oint H dl = H_f L_f + H_s L_s \tag{3.22}$$

$$\Phi_f = \Phi_s \Rightarrow B_f A_f' = B_s A_s \tag{3.23}$$

where the subscripts $_f$ and $_s$ designate values in the fluid and in the steel, respectively. N is the number of coil turns, L is the length of the path travelled by the magnetic flux in each material, H is the magnetic field intensity, Φ is the magnetic flux, B is the magnetic flux density and A is the cross-sectional area "seen" by the magnetic flux. While for the steel this is simply A_s (the actual cross-sectional area of the steel element), for the MR fluid a different effective value (A_f') is used to account for the phenomenon of magnetic fringing[2] [69].

[2]This is the tendency of the magnetic flux to disperse through a greater area when crossing a material of low permeability, such as air, aluminium or MR fluid

There are 4 unknowns (H_f, H_s, B_f and B_s) in these 2 equations. Hence, two more equations are required and these are the constitutive relations for each of the materials (steel and MR fluid):

$$B_s = f_s(H_s) \tag{3.24}$$

$$B_f = f_f(H_f) \tag{3.25}$$

The relation between the magnetic flux (B) and the magnetic field (H) is generally non-linear (and that is the case for the two materials present in this system, steel and MR fluid). For this reason, the usual way to represent that relation for each material is by means of a B-H curve. The functions f_s and f_f cannot therefore be expressed analitically but may instead be obtained by approximation of the B-H curve for each material. Hence, a closed-form solution to the equations of magnetostatics cannot be obtained for systems comprising non-linear materials.

3.2 Coil Wire

The choice of the wire to use in the construction of the coil is related to the definition of the current at which the MR brake will be operating and the number of turns in the coil. Thicker wires are capable of conducting greater currents but take more space and hence a smaller number of turns can be wound in the same area. It can be seen from Table 3.2 - which presents the relevant properties of copper wires of different AWG (American Wire Gauge) sizes - that the change in current carrying ability is inversely proportional to the number of turns per unit area. This means that if a wire is capable of carrying twice the current of another, it will only allow half the number of turns and hence NI, the product of the number of turns by the current, is the same for the various wire sizes. Given that the magnetic flux produced by the

coil is proportional to NI, the choice of a given wire dimension will not influence it. However, altough any combination of N and I giving the same value of NI would produce the same magnetic field, not all values for N or I are equally acceptable. In particular, high values of either variable are undesirable: given that the electric power is $P = VI$ where V is the voltage of the power supply, a high value of I would result in an undesired high power consumption; on the other hand, a coil with a very high number of turns may be unpractical to build. With this in mind, it is seen that the wire to be chosen for the coil of the MR brake must be such that a balance between the values of N and I exists.

It is important to distinguish two different phenomena that may limit the maximum current that can be carried by a given wire: voltage drop and heat buildup. The former is associated with the fact that a current I flowing through a conductor with resistance R originates a drop in voltage equal to $\Delta V = R \cdot I$ (Ohm's law) - hence, if a wire with too high a resistance is used for the electrical circuit in a house, an excessive voltage drop will occur and instead of 110 V, the voltage available at the wall electrical outlets becomes $110 - \Delta V$. The need to keep the voltage drop low in wires used for power transmission results in the maxium current carrying ability expressed by the lower limit of the interval on the column "Maximum current" of Table 3.2. If the wires are not meant to be used to supply power to appliances or other devices, voltage drop is no longer a problem and the current limit of a given wire gauge is higher as it is only limited by heat buildup[3] - this phenomenon is related to the fact that when current flows through a resisting conductor, some electrical energy is dissipated as heat at a rate $H_{gen} = V \cdot I = R \cdot I^2$. Higher currents will result in an important amount of heat generation, leading to a rise in the conductor's temperature, ultimately causing the conductor to fuse. Naturally, the allowable current

[3]When only heat buildup is to be taken into account, the current carrying ability of each wire is given by the higher of the values on the column "Maximum current"

for each wire gauge is much less than that which would lead to fusion of the wire (indicated on the table's third column, "Fusing current"). This is due to the fact that the temperature must not only be kept to a value below the fusing limit of the wire but to much lower values that prevent damage to the wire insulation and fire hazards.

Table 3.2: AWG copper wire properties

AWG number	Diameter[a] [mm]	Fusing current[a][A]	Maximum current[b] [A]	Ω/1000ft[c]	Wires/in[d]
34	.16	9.04	0.057 - 0.33	260.9	21003
33	.1798	10.2	0.072 - 0.43	206.9	16436
32	.2019	11.5	0.090 - 0.53	164.1	12913
31	.2268	12.8	0.114 - 0.7	130.1	10628
30	.2545	14.1	0.144 - 0.86	103.2	8416
29	.2860	15.5	0.181 - 1.2	81.83	6610
28	.3211	18.4	0.228 - 1.4	64.90	5328
27	.3604	21.5	0.288 - 1.7	51.47	4272
26	.4049	24.7	0.363 - 2.2	40.81	3460
25	.4547	29	0.458 - 2.7	32.37	2770
24	.5105	33.4	0.577 - 3.5	25.67	2204
23	.5733	38.1	0.728 - 4.7	20.36	1743
22	.6139	48	0.918 - 7	16.14	1403
21	.7229	58.6	1.16 - 9	12.80	1115
20	.8118	69.9	1.46 - 11	10.15	888.4
19	.9116	81.5	1.84 - 14	8.051	711.1
18	1.024	107	2.32 - 16	6.385	566.9
17	1.150	132	2.93 - 19	5.064	450.8
16	1.290	166	3.69 - 22	4.016	360.7
15	1.449	198	4.65 - 28	3.184	287.3
14	1.628	232	5.87 - 32	2.525	228.5
13	1.828	286	7.40 - 35	2.003	183
12	2.003	344	9.33 - 41	1.588	146

[a]Source: [70]
[b]Sources: [71, 72]
[c]Source: [72]
[d]Source: [73]

It can be seen that the values for the number of wires per square inch indicated in Table 3.2 are not the same as would be obtained by dividing one square inch by the cross-sectional area of each wire:

$$N_{in^2} = \frac{1}{\pi \left(\frac{d}{2}\right)^2} \tag{3.26}$$

with the diameter d of the wire expressed in inches. The difference between the value obtained with this expression and that indicated in the AWG table is explained by the fact that the wires have circular cross-section and hence when winding the wire, there will be interstices between the various turns. This suggests the use of rectangular cross-section wire. Such wires do exist and are generally designated by the AWG number corresponding to their area. That is, a rectangular cross-section AWG 20 wire has such dimensions that it has the exact same cross-sectional area of a standard AWG 20 wire and hence the same current-carrying ability. This means that once a given AWG wire number has been selected according to the criteria indicated above, the choice between circular or rectangular cross-section wire can be made easily without requiring further calculations or conversions. This choice will be dictated by the needs faced by the system: if the demand for a greater number of wire turns in a limited area is pressing, the choice for the rectangular cross-section is inevitable; otherwise, the lower cost and much better availability of circular cross-section wire justify its selection.

3.3 Metallic Materials Selection

As it was seen in Section 3.1, the MR brake comprises not only the MR fluid and the coil that produces the magnetic field but rotating components such as the shaft and disk brake and static components such as the casing and the ring between the

fluid and the coil. The materials used in the construction of these components play an important role in the performance of the MR brake, as well as in its dimensions, weight and longevity. This section will describe the requirements for each of these components, a survey of appropriate materials and the criteria behind each choice.

The first aspect to consider in the choice of materials for each component is the desired magnetic behaviour: if the component is to be part of the magnetic flux path, then a magnetically permeable material must be chosen. If no flux is intended to occur across a component, then it must be of a non-permeable material, i.e., a non-ferrous material.

The choice of which magnetically permeable material to use is determined by the needs of each particular problem. In the case of the MR brake components, it is important to have a high saturation flux density (given that the flux is constant troughout the magnetic path, a high flux to the MR fluid will only be available if the ferrous parts are able to sustain high fluxes) and a high magnetic permeability, so that the magnetic field intensity on the ferrous parts is kept to a low value, thus reducing energy losses. Recall from section 3.1 that the required product of current and turns in the coil is equal to $H_f L_f + H_s L_s$. A high magnetic permeability in the ferrous components will keep H_s low and thus minimise the coil energy requirements.

The materials that best meet these criteria are low-carbon steels and supermendur (a Fe-2%V-49%Co alloy). The maximum relative permeability and the saturation flux (B_{max}) of these materials are presented in Table 3.3.

Table 3.3: Soft magnetic material properties

Material	Max. Relative Permeability	B_{max} [T]
Low-carbon steel	4000	2.14
Supermendur	100,000	2.30

Source: [74]

Altough supermendur is technically superior, the better availability of low-carbon steels determined its choice. Among these, Lord Corporation [69] recommends the use of AISI-12L14, AISI-1008, AISI-1010 or AISI-1018. For the remainder of this thesis, AISI-1010 will be assumed, given the availability of exact specifications for finite element modelling[5]. This is a 0.1% carbon steel, also referred to as SAE 1010. Its mechanical properties are listed in Table 3.4, together with those of candidate materials for the non-magnetic parts. For these, the considerations involved in the selection of the most adequate material involve the density, structural strength, cost and availability and heat conduction. Suitable candidates include aluminium, magnesium and titanium.

Table 3.4: Structural properties of materials

Material	Density [kg/m^3]	Yield Strength [MPa]	Young's modulus [GPa]
AISI 1010 Steel[a]	7860	276-414	179-310
Aluminium[b]	2630-2800	100-500	70-75
Magnesium[b]	1800	250	45
Titanium[b]	4730	830	115

[a]Sources: [75, 76]
[b]Source: [75]

It is seen from Table 3.4 that magnesium is the lighest of the 3 non-magnetic materials and hence a strong candidate. The advantages of aluminium and titanium in terms of strength and Young's modulus may not be relevant as the structural load is not expected to be significant. However, the lower cost and greater availability of aluminium may favour it. The next section addresses another aspect that may influence the materials selection: the corrosion behaviour.

[5]The ANSYS software material library contains the B-H curve for this particular steel

3.3.1 Corrosion Considerations

One aspect to take into account in the selection of the materials for the MR brake is the likelihood of occurrence of corrosion. Corrosion is due to an oxidation reaction, whereby atoms of a metal (the anode) give up some (n) electrons and form ions, corroding that metal:

$$M \rightarrow M^{n+} + ne^{-} \tag{3.27}$$

A reduction reaction in which an element (cathode) is reduced in valence charge must occur simultaneously at the same overall rate to prevent buildup of electric charge in the metal [77]. For an hydrogen cathode, this is written as:

$$2H^{+} + 2e^{-} \rightarrow H_2 \tag{3.28}$$

The likelihood of corrosion of each material depends on the environment surrounding it, being most common in water and the atmosphere. For this reason, the tendency of a material to corrode is generally obtained from the electrode potential of that material in an aqueous solution. Materials with an electrode potential greater than that of hydrogen have a lower tendency to corrode, whereas those with an electrode potential lower than that of hydrogen are most likely to corrode. The electrode potentials at 25°C are presented in Table 3.5 [77].

It is seen that noble metals such as gold, platinum and silver have very high electrode potentials and hence do not corrode easily, whereas materials such as aluminium and magnesium have very low electrode potentials. However, aluminium reacts easily with oxygen, forming a compound that prevents further corrosion. That is why corrosion of aluminium is uncommon. Hence, the choice of aluminium for the non-magnetically permeable parts of the MR brake does not pose a problem, but should magnesium be chosen, care will have to be taken to prevent corrosion. This

Table 3.5: Standard electrode potentials at 25°C

Element	Electrode potential, E° [V][a]
Au	+1.498
$2H_2O$	+1.229
Pt	+1.200
Ag	+0.799
2Hg	+0.788
Fe^{2+}	+0.771
$4(OH)^-$	+0.401
Cu	+0.337
Sn^{2+}	+0.150
H_2	0.000
Pb	-0.126
Sn	-0.136
Ni	-0.250
Co	-0.277
Cd	-0.403
Fe	-0.440
Cr	-0.744
Zn	-0.763
Al	-1.662
Mg	-2.363
Na	-2.714

Source: [77]

[a]using the standard hydrogen electrode as reference, i.e., 0 volts

is generally achieved by means of protective coatings on the surface of the material.

The same action may be necessary for the steel parts. Recall that the parts which are to conduct the magnetic flux will be made of a low-carbon steel. Unlike stainless steels, these are prone to corrosion and may require a protective coating to avoid corrosion. Given the need to ensure good cooling of the MR brake, this coating must be such that the heat transfer is not adversely affected.

Galvanic Corrosion

The choice of materials must also consider the possibility of development of galvanic corrosion. Galvanic (or bimetallic) corrosion may occur when different metallic elements are in contact in the presence of an electrolyte (in this case the MR fluid). In this case, the electrons formed at the anode (recall the equations in the previous section) flow through the electrical contact between the two materials, whereas the ions will move through the electrolyte.

Given that it is the anode that corrodes (simultaneously protecting the cathode from corrosion), it is fundamental to determine which of the two materials in contact is the anode and which is the cathode. This depends on the electrode potential of each metal in the environment in which they operate, the anode being the most electronegative material of the two. The galvanic series is a ranking of different materials according to their anodic or cathodic tendency when immersed in a certain electrolyte. Altough various parameters (temperature, velocity, biocide treatments) influence the potential of metals, the relative ranking of the different materials is not significantly affected by these factors or even by the electrolyte, and the galvanic series for seawater is commonly used to assess the corrosion behaviour of metals in most media [78].

The galvanic series for seawater is presented in Table 3.6. Depending on the particular testing conditions the order of the materials may not be exactly the same and slightly different galvanic series will be obtained from other references (eg [74, 79]), but the changes are minor and localised thus confirming that in spite of the complexity of the phenomenon of corrosion, the galvanic series is a good indication of the anodic/cathodic tendency of different metals.

It is clear from Table 3.6 that aluminium will be the most anodic when paired with steel or iron, and thus the most likely to corrode. However, recall from the previous section that in most media aluminium develops a layer that protects it from corrosion. The same phenomenon occurs with titanium, avoiding its corrosion. It is important to note that this phenomenon is far less important in titanium than in aluminium as it can be seen from the galvanic series that the former would rarely be the anodic element. Hence, when combining titanium and most other alloys, the greatest risk of corrosion is faced by the other element. Study of the phenomenon of galvanic corrosion has therefore resulted in important considerations regarding the choice of material for the non-magnetic parts: magnesium (otherwise a strong candidate due to its excellent strength to weigth ratio) is a poor choice due to the high chance that it will suffer dramatically from corrosion; titanium should also not be used as its combination with steel may cause accelerated corrosion in the latter; finally, aluminium has an electrode potential close to that of steel and even though it is the most anodic in that combination, corrosion will be avoided by the layer formed after the initial reaction with oxygen.

Table 3.6: Simplified galvanic series in seawater

Electrode Potential	Alloy	
	Graphite	
⇑	Platinum	
Electro	Gold	
Positive	High Alloy Stainless Steels	{ Super Austenitic }
(Cathodic)		{ Super Duplex }
	Titanium	
	Nickel Chrome Molybdenum Alloys	{ 625; C-276 }
	Low Alloy Stainless Steels (eg 316)	(PASSIVE)[a]
	Alloy 400/Alloy K-500	
	Silver	
	Nickel Aluminium Bronze	
	Copper nickel (70/30; 90/10)	
	Gunmetals/Tin Bronzes	
	Brasses	
	Tin	
	Lead	
	Austenitic Cast Iron	
	Low Alloy Stainless Steels (eg 316)	(ACTIVE)[a]
Electro	Cast Iron	
Negative	Carbon Steel	
(Anodic)	Aluminium Alloys	
⇓	Zinc	
	Magnesium	

Source: [78]

[a]Stainless steels and some nickel-based alloys have a thin film that while intact protects the material from corrosion. The material is then said to be passive. If this film fails, the material is more prone to corrosion and said to be active. [78, 74]

3.4 Design Methodology

Engineering studies fall into one of two categories: analysis problems, where the goal is to understand the operation and performance of existing systems; and design problems, where the goal is to conceive a system capable of achieving a predefined performance. The former are direct problems in that they depart from the characteristics of the system (dimensions, materials, etc.) to obtain its output whereas the latter are inverse problems where the question is that of determining which characteristics produce a given output. Direct problems are generally simpler to solve in that they only require the evaluation of known functions (describing the behaviour of the system being analysed) for the characteristics of the problem at hand; inverse problems, on the other hand, require the solution of the corresponding set of design equations.

In the particular case of an MR brake, the analysis problem would simply consist of running the finite element models of a given brake configuration (i.e., with known dimensions and materials). Clearly, the solution to the design problem of finding the ideal configuration of the MR brake is not so direct. Without closed-form solutions of the equations describing its dynamics, no immediate approach exists. In such cases, the design problem generally consists of an iterative process involving a sequence of analyses problems. Based on the results of each analysis some changes are made to the design and the performance of the resulting configuration is compared with that of the previous one, continuing this process until no further improvements in the design are achieved. Naturally, this kind of problems has been studied thouroughly in the past and a variety of optimisation techniques now exist to obtain an "ideal"[9]

[9]What exactly is an ideal design depends on the priorities of the designer and on the problem at hand. It is the responsibility of the engineer to define a priori the goals to achieve (such as lowest cost or best performance).

design using the information gathered from the analysis of candidate configurations.

Hence, in order to design a suitable MR brake, two different tasks are necessary: a model capable of analysing the performance of a given brake design and an optimisation tool capable of using the results of such analyses to produce improved designs. These two main tasks in the design of the MR brake will be the subject of Chapters 4 and 5 respectively.

The finite element analysis process can in turn be subdivided into the following steps: since the friction that originates the braking torque depends on the viscosity of the MR fluid and this in turn depends on the magnetic field, determination of the magnetic field distribution must be the first step in the study of the MR brake. When the magnetic field is known, the corresponding shear stress is obtained from the fluids's specification and a CFD (Computational Fluid Dynamics) analysis may then be conducted to determine the friction exerted on the disks' surfaces and the temperature distribution.

3.5 Definition of Candidate Configurations

As seen in the previous section, the design procedure will consist of an iterative analysis process: given an initial design, its behaviour is analysed, modifications are carried out, the resulting design is analysed, new modifications carried out and this sequence continued until no further improvements can be achieved. Therefore, this procedure requires the specification of an initial design. That is the subject of the present section.

It has been seen in the previous sections that an MR brake is formed by a disk rotating within magnetorheological fluid enclosed in a static casing. However, this

general concept leaves room for a variety of different configurations depending on the choices made regarding the following aspects:

- Number of disks;
- Method of generation of the magnetic field: choice between permanent magnet, electromagnet or both, and the definition of their location;
- Dimensions of the components;
- MR fluid to be used;

The proposed configuration of the MR brake is illustrated in Fig. 3.1, where a cut has been made to highlight the cross-section that will be considered from here on.

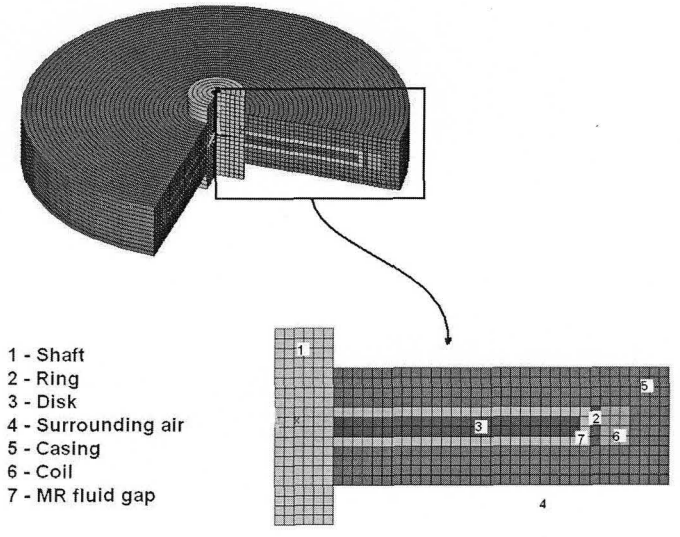

Figure 3.1: Proposed configuration of the MR brake and detail of the cross-section to be modelled

When defining the thickness of the MR fluid gap, several aspects must be considered. On one hand, a narrower gap will have the MR fluid subject to a higher magnetic

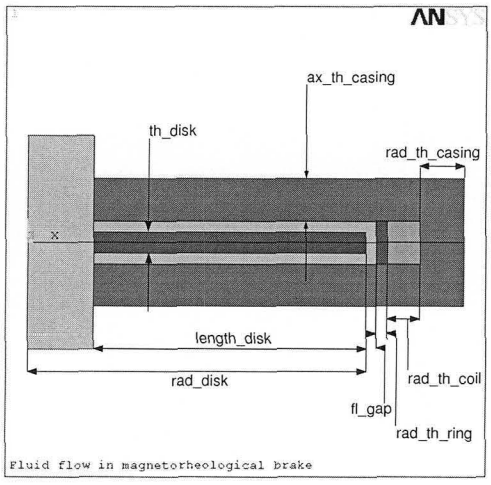

Figure 3.2: Illustration of the different variables used

field intensity and will reduce the volume of MR fluid (hence, the weight and cost). On the other hand, too narrow a fluid gap will increase the production cost as it will require tighter tolerances in the construction and assembly of the parts. Also, recall from the equations of the braking torque introduced in the previous section that the component T_μ associated with the viscosity of the carrier liquid is inversely proportional to the thickness of the fluid gap (h). Given that this component of the torque is present regardless of whether the brake pedal is pressed or not, it is desirable to reduce it, which implies increasing h. With all the above in mind, an initial value of 1 mm was chosen for the thickness of the MR fluid gap.

In consideration of the dimensions of the various components of the MR brake, the first factor to take into account is the existence of physical limitations. For example, if the brake is to be placed within the wheel rim (where today's disk or drum brakes are located), the overall diameter must be such that it will fit within that area: it is recommended that a minimum clearance of 3 mm exists between the brake and wheel

rim and spokes [80]. The maximum acceptable diameter for the MR brake is then, for a 16" wheel[10], $16in \times 25.4mm/in - 2 \times 3mm = 400.4mm \approx 40$ cm, resulting in a maximum acceptable radius of 20 cm. Note that this dimension is valid for wheels with a perfectly cylindrical rim. Many wheels have curved rims, which may require either a smaller MR brake or moving it closer to the suspension and away from the outside of the car (where the rim diameter may be smaller).

Recall from Section 3.1 that the braking torque due to the magnetic field is given by $T_B = 2\pi N \int_{r_w}^{r_z} \tau_e r^2 dr$ (Equation 3.16) and hence proportional to the third power of the radius. For this reason, it is important to maximise the radius of the disk. With the above restriction of the overall radius of the brake (20 cm) in mind and considering that the brake comprises other components outside the disk (such as the coil and casing), a value of 17 cm is tentatively assigned to the disk radius.

Considering that the shaft has a radius of approximately 3 cm (the value of the shaft of the Chevette used in the hydraulic brake setup described in 2.1) and that components such as bearings and seals will likely be mounted between the shaft and the casing, the latter's inner radius has initially been given a value of 5 cm[11].

For the purpose of determining the starting configuration of the optimisation procedure, it is acceptable to resort to the assumption of a constant magnetic field distribution (which, as seen in Section 3.1 is not exact). Equating the value of T_B given by (3.16) to 1010 N·m with $N = 2$ (one disk, two surfaces) and r_w and r_z equal to 5 and 17 cm, respectively, gives a value of 50359 Pa for the electro shear stress. From the specifications of the two candidate MR fluids, it is seen that in the case of MRF-241 this corresponds to a magnetic field intensity of approximately

[10]Larger wheels do exist (sizes up to 18" are normal, with some models - e.g., Momo's Storm [81] - available in dimensions up to 24") but are not too common other than in sports cars or SUVs.

[11]That is, the length of the disk, defined as the difference between its outer radius and its inner radius, is equal to $17 - 5 = 12$ cm

110 kA/m, whereas MRF-132 is not capable of achieving such a shear stress. In other words, either the disk radius or the number of disk surfaces in contact with the fluid must be increased in order to obtain a design that meets our prespecified torque target using MRF-132 fluid. In what concerns the system based on MRF-241 fluid, the knowledge of the required magnetic field intensity can be used to obtain the necessary NI, the product of the number of turns and the current intensity. Again using a relatively simple approach, it is possible to assume that in this initial design the steel components will be operating far from saturation and hence experiencing a very low magnetic field intensity (i.e., $H_s \approx 0$). If this is the case, the relation $NI = H_f L_f + H_s L_s$ reduces to $NI = H_f L_f$ and for this particular case (with H_f equal to 110 kA/m and L_f, the total thickness of MR fluid, equal to 2 times[12] the individual fluid gap of 1 mm), $NI = 220$ Ampère-turns.

For the initial design, a current of 1.5 A can be used (recall from Table 2.7 that MR fluids typicall operate at currents between 1 and 2 A). The number of turns is then $N = 220/1.5 = 147$. Using the most conservative criteria for 1.5 A from Table 3.2, AWG20 wire is chosen. From the same table, it is seen that this wire can be wound to approximately 888.4 turns per square inch. Hence, the necessary 147 turns correspond to a coil cross-sectional area of 0.165 in^2 or 1.07 cm^2. If the disk is chosen to have a thickness of 1 cm (the same as in the Chevette brake), the axial length of the coil (equal to the disk thickness plus the total thickness of the fluid gaps) is 1.2 cm. Hence, the radial length of the coil must be $w_{coil} = 1.07/1.2 \approx 0.9$ cm.

We have seen that the maximum radius of the MR brake is 20 cm and have set the disk radius to 17 cm and the width of the coil to 0.9 cm. In addition to these two quantities, the radius of the brake must also accommodate the fluid gap (0.1

[12]There are two MR fluid gaps, one between each side of the disk and the adjacent surface of the casing

cm) adjacent to the edge of the disk, the aluminium ring that will provide separation between the MR fluid and the coil and the casing. If the aluminium ring is chosen to have 5 mm[13], the maximum possible thickness for the end part of the casing is $20 - 17 - 0.1 - 0.5 - 0.9 = 1.5$ cm. This is the value that shall be used in the initial design.

The only remaining independent dimension is the thickness of the lateral walls of the casing. For the purpose of the initial design, it will be given the same value of the end part of the casing, i.e., 1.5 cm.

Other than the dimensions, there is one additional parameter that influences the performance of the MR brake: the number of disk surfaces. Whereas having a simple disk gives the simplest design and would be ideal, it was seen in Section 3.1 that the braking torque varies linearly with the number of surfaces and hence the addition of more disks may be necessary to generate more torque. It was therefore decided to study two different configurations: a single disk configuration, and a multiple disk configuration. For the latter, it was decided to use 3 disks, which represents a good balance between ease of manufacture (not too many disks) and torque improvement (50% more than in a 2 disk geometry).

[13]Since it is only a barrier between the MR fluid and the coil, it can be as small as desired. However, making it too thin would introduce structural problems and unnecessarily complicate the manufacturing and assembly process.

Chapter 4

Finite Element Modelling

Given the multiple disciplines involved in the operation of a magnetorheological brake (magnetostatics, fluid flow, heat transfer, solid mechanics), the presence of non-linearities (magnetic saturation and non-newtonian fluid behaviour) and the absence of closed-form solutions, the analysis of MR brakes is carried out using finite element modelling. This is an analysis technique that replaces the exact differential equations of the problem being studied with a system of simultaneous algebraic equations [82]. This is done by using a modified mathematical formulation of simpler solution (weak formulation) for the problem. This modified mathematical formulation is valid on average within the problem domain, unlike the original differential equations (strong formulation), which hold for every point in the domain. Hence, in order to obtain meaningful solutions with the weak formulation, it is necessary to divide the original problem domain (the system being analysed) into many small domains (the elements), so that the "average" introduced by the weak formulation does not introduce significant errors in the solution.

Many finite element analysis programs are commercially available, each with its own strengths and weaknesses. Among these are ANSYS and ModuleF, both cur-

rently in use at the University of Victoria's Department of Mechanical Engineering. While ModuleF has an extremely high adaptability given that it is an "open source" program that may be freely modified by the end user, it does not allow the same degree of control of the analysis as ANSYS (the latter's input files can include some computer logic and programming constructs such as *if* cycles and *do* loops while a ModuleF analysis requires the execution of several different programs, each reading input data from a different file) and is less complete than ANSYS in the disciplines modelled by the default elements. Particularly, any magnetostatics capabilities would imply the programming of new element types. ANSYS, on the other hand, has very broad-ranging capabilities including the analysis of structures, electric and magnetic fields, heat transfer and fluid flow[1]. Another important advantage of ANSYS for this project is the fact that it possesses an optimisation module, comprising two different methods. This means that, once the finite element models for the MR brake have been completed they can be used to determine the best possible configuration for the system. Furthermore, ANSYS is an extremelly popular program, thouroughly tested and validated.

Furthermore, ANSYS possesses strong capabilities when dealing with coupled-field problems. These are problems involving various engineering disciplines, with inter-dependence between them, i.e., the results of each discipline influence at least one of the other disciplines. That is not the case in the present problem, since the magnetic field distribution controls the fluid viscosity which influences the fluid flow and heat transfer but these do not in turn influence the magnetic field distribution. The problem is then said to be uncoupled and can be solved by running first the mag-

[1]The fluid flow module of ANSYS is its least achieved area, requiring a reasonable amount of time and fine-tuning of the mesh to converge. However, this problem is not specific to ANSYS but results from a difficulty of application of finite element techniques to fluid dynamics, which explains why programs dealing exclusively with fluid flow analysis often use other methods such as finite difference discretisation of the problem's equations.

netostatics analysis to determine the magnetic field and hence the MR fluid viscosity and then the computational fluid dynamics (CFD) analysis to determine the friction exerted on the disks and the temperature distribution.

4.1 Magnetic Field Model

Since the MR fluid's viscosity is controlled by the magnetic field, the first step in the analysis of an MR brake system is a magnetostatics analysis that determines the magnetic field distribution.

The first choice to be made when conducting finite element analysis is that of how to approximate the real system so that the model built to describe it is simple yet accurate. Finite element models exist for one, two and three-dimensional problems and it is therefore necessary to determine the minimum number of dimensions required to faithfully describe the behaviour of the system. Although every object in the universe is three-dimensional, many phenomena can be studied using one or two-dimensional models. For example, in a structural element subject only to axial loading, the displacement varies only along the axial direction and a one-dimensional model is sufficient. In the case of a plate subject to normal loads, the resulting displacement is only a function of the two in-plane coordinates and thus a two-dimensional model may be used.

A first glance at the MR brake suggests that the magnetic potential varies along the three coordinates and hence a three-dimensional model may be required. However, it is also clear that the geometry of the brake is axisymmetric. A question arises: can this be used somehow to reduce the dimensionality of the analysis? Indeed, two-dimensional models can be used to describe the behaviour of an axisymmetric problem, where not only the geometry has to be axisymmetric, but so are the material

properties and the boundary conditions. In the present problem, all these conditions are met, as the material properties do not vary in the tangential direction, and neither does the applied load in the coil.

Geometry

The construction of the finite element model begins with the definition of the geometry of the system to be analysed. This task is generally the most time-consuming in finite element analyses due to the challenge posed by the need to devise a model that accurately represents the real system while keeping computational demands to a minimum.

An important distinction must be made between geometric model and mesh. The former is the description of the form of the system being analysed while the latter is the set of elements and nodes[2] that will be used in the solution of the equations. Hence, when studying a given problem, the geometrical model must first be built and later this geometry is divided into smaller elements. In other words, while the geometrical model is characteristic of the system being analysed, the mesh is associated with the mathematical formulation of finite elements. The present section deals with the construction of the geometry. The meshing process can only be done once the geometry is defined and will be discussed in a subsequent section.

The geometry of the cross-section of the MR brake to be modelled is shown in Fig. 4.1.

For the analysis of the magnetic field distribution, it is prudent to also model the air around the component being studied so that any leakage of magnetic flux that

[2]The nodes are the points where the quantities of interest (displacements, stresses, magnetic fluxes, etc.) are computed. The location of the nodes on each element depends on the element type - e.g. some triangular elements have 3 nodes (one on each vertex of the triangle), whereas other have 6 nodes (one on each vertex plus one mid-point on each side).

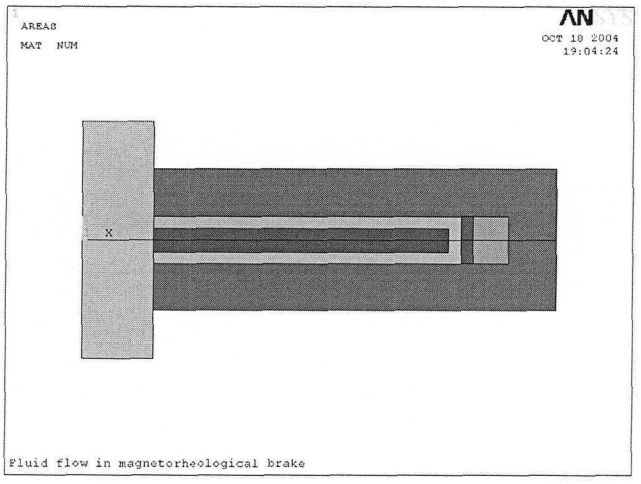

Figure 4.1: Geometry of the MR brake cross-section

may occur is taken into account. Altough no significant leakage is expected in the MR brake (given the use of steel for the casing, with a magnetic permeability 4000 times greater than that of air), a layer of air was included to provide more accurate results and verify that no leakage occurs. The resulting geometry of the MR brake is shown in Fig. 4.2 and the various quantities of interest are made indicated in Fig. 4.2.

Material Properties

The relation between the load and the primary variable is described by the stiffness matrix[3]. In order for the finite element program to be able to compute it, the relevant material properties must be specified. For the case of magnetostatics, this is the

[3]The name comes from the analogy with structural problems where the relation between force and displacement is related to the stiffness of the elements

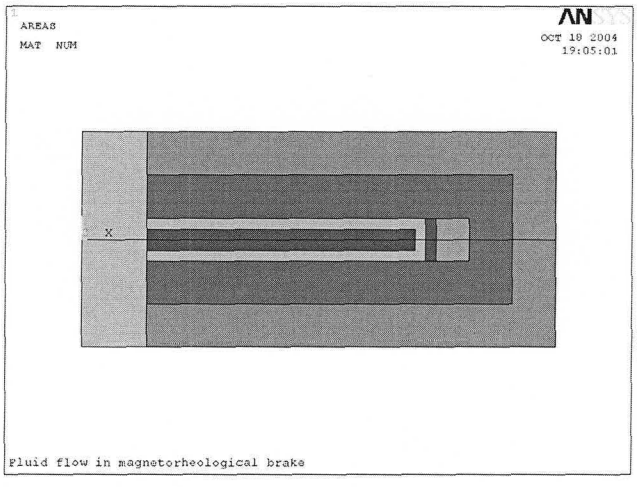

Figure 4.2: Geometry of the MR brake cross-section including the surrounding air elements

magnetic permeability that relates the magnetic flux B and the magnetic field H:

$$\mu = \frac{\partial B}{\partial H} \tag{4.1}$$

The simplest case is that of a material possessing constant magnetic permeability for any value of the magnetic field. Unfortunately actual engineering materials depart from this ideal case in two aspects:

- First, the relation between B and H is generally non-linear, i.e., the value of the permeability μ is not a constant for the entire domain of operation of the material but a function of the magnetic field;
- Second, actual materials are not capable of carrying infinite magnetic fluxes. Instead, the material saturates at some point, i.e., for values of the magnetic field intensity greater than the limit of the material H_{sat} the magnetic flux no

longer increases.

While in the case of steel or iron the errors associated with the first aspect are moderate (the behaviour of these materials does not depart too much from a linear relation between the flux density and the field intensity), the consequences of the second aspect are dramatic. Hence, the magnetic behaviour of most materials must be described by the full B vs. H curve rather than by a constant scalar (μ), as would be preferable from a computational point of view.

It is interesting to note that Walid [83] used a constant magnetic permeability value for all metallic materials when modelling an MR damper. The results (maximum flux density of approximately 1T) indicate that saturation does not occur in that system and hence the errors associated with this assumption are probably small. In the present project, however, it is important to reduce the weight of the steel casing that conducts the magnetic flux, while simultaneously allowing the latter to be as high as possible. This (high magnetic flux in reduced cross-section) will result in a high magnetic flux density and hence it is expected that the steel casing will operate close to saturation, which must then be taken into account in the finite element model. Therefore, rather than defining a constant magnetic permeability for steel, the actual B-H curve for SAE 1010 steel supplied with ANSYS' material library was used.

The MR fluid's B-H curve departs even more from the ideal linear model than that of steel and hence the full curve must also be used to describe the magnetic behaviour of this material. This curve is available from the technical specifications supplied by the manufacturer [66, 67].

As for aluminium, it exhibits a negligible magnetic permeability and its value is generally taken to be the same as that of free space ($4\pi \times 10^{-7}$T·m/A). A constant

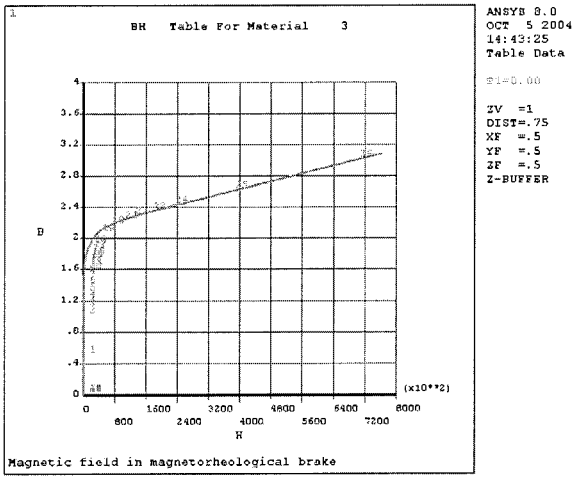

Figure 4.3: B-H curve for SAE 1010 steel
Source: [84]

value for μ can be used in this case since it does not vary with the magnetic field and that, given its very low value, saturation will not occur[4].

Boundary Conditions

The response of a system depends not only on its geometry and the properties of the materials it is made of but also on the loads that are applied and on any contraints imposed to it. These are collectivelly referred to as the boundary conditions of the problem.

The load in the present problem is the current flowing through the coil, responsible for the magnetic flux. It is input to the finite element model as a current density *jdens*

[4]Recall that for a material with constant magnetic permeability, the magnetic flux density is given by $B = \mu H$ and therefore for the values of H that can be obtained from practical coils, the magnetic flux will always be low

defined as:

$$jdens = \frac{NI}{A} \tag{4.2}$$

where N is the number of coil turns, I is the current and A is the area of the cross-section of the coil.

It is also necessary to define boundary conditions for the magnetic potential. Several possibilities are available and are described in [85]. The most adequate for the problem at hand is "flux-parallel", i.e., assuming that the magnetic flux at the boundaries is parallel to the boundary lines or, in other words, that there is no magnetic flux leakage to the outside (which would imply a component of the magnetic flux normal to the boundary). As indicated above, the present finite element model includes elements to model the air around the casing. This means that leakage from the casing to the air around it will be taken into account and the "flux-parallel" condition is only negleting magnetic flux leakage from this layer of air to the far-field atmosphere, which is approximately zero[5].

Meshing

As mentioned earlier, one of the central aspects of the finite element model is the division of the domain into finite elements, the process known as meshing.

The first step in the meshing process is the selection of the finite elements to use, among those that meet the needs of the problem at hand (in this case, two-dimensional axisymmetric elements with application to magnetostatics). Generally, several elements with different geometries and number of nodes are available to solve

[5]Actually, since the permeability of the steel used in the casing is several orders of magnitude greater than that of air and given that it will be operated below saturation, even the occurence of leakage from the casing to the modelled surrounding air is extremelly unlikely. It will therefore be interesting to look at the results and confirm if indeed the leakage from the casing to the surrounding air is minimal.

the desired set of equations. For irregular geometries, triangular elements are more adequate as it is easier to divide a domain comprising narrow angles in triangular areas than in quadrangular ones. For domains bound mostly by lines or areas at right angles, quadrangular elements may be used and this option has the advantage of requiring half the elements that would be necessary using triangles. It is also possible to use elements with a different number of nodes. For example, some elements have midside nodes, which means that triangular elements with 6 nodes instead of 3 are available, as are quadrangular elements with 8 nodes instead of 4. Elements with midside nodes are particularly suited to problems with curved boundaries, as these are much better approximated, albeit at a higher computational expense. Considering all the advantages and disadvantages associated with each type of element, a 4 node quadrilateral element was chosen for the present problem (given that the geometry is a simple one, with straight boundaries at right angles). The ANSYS 4 node quadrilateral element for magnetostatics is named PLANE13.

Once the element type to be used has been chosen, the actual discretisation of the domain may proceed. At this point, a distinction can be made between two different types of meshes: structured and non-structured. In the former, some nodes are defined by the user (generally in the boundary) and the remaining nodes are obtained as a function of those pre-defined nodes, using a geometric relation (also specified by the user). This results in a mesh whose node distribution forms a pattern. On a non-structured mesh, on the other hand, no nodes are pre-specified and all nodes are placed in the manner the meshing program determines to be most adequate. Looking at a non-structured mesh, generally no pattern in the node distribution can be identified and the nodes will appear to be randomly placed throughout the domain.

For the present problem, element divisions were specified along the edge lines of the model, but some freedom was given to ANSYS' smart mesher to find the best

mesh, resulting in a slightly unstructured mesh.

Whichever type of mesh is selected, the size of the elements must be chosen. This is a matter of striking the right balance, given that large elements produce a coarser mesh with lower computational requirements and faster solution time but may entrain important errors in the determination of the unknowns or even prevent convergence of the solution altogether, whereas a refined mesh made of very small elements will imply an exaggerate solution time without noticeable improvements in the accuracy of the solution[6]. The definition of the fineness of the mesh is therefore a delicate problem of compromise between conflicting interests. It is often a trial and error process until the ideal dimensions for a given problem are found. Analyses may be conducted with different meshes to determine the ideal element size (the largest element dimension that does not depart significantly from the results obtained with a very fine mesh). It is also common to use different element sizes at different stages of the design process: in the initial phases, when it is important to conduct many analyses of different concepts to study the effects of possible design changes and the goal is only to obtain a rough comparison of the various designs, a coarser mesh is used; later, when proceeding to detailed analyses of only a few configurations, a finer mesh may be used to obtain more accurate solutions.

A useful feature of ANSYS and many other finite element programs is the ability to locally refine the mesh. That is, after a mesh has been created for the entire domain using elements of a given dimension, it is possible to change only a part of the mesh so that smaller elements are used in that area. This is useful if a problem has one area where the gradients of the unknowns are more pronounced, as it allows faster convergence and greater accuracy without overloading the model with

[6]Too fine a mesh may also cause difficulties in the convergence of the solution, due to numerical problems.

a large number of elements and nodes in areas where they are not necessary. Local refinement of the mesh was not necessary in the present problem given the absence of significant gradients and the relative simplicity of the equations of magnetostatics (when compared with those of fluid dynamics, for example).

Solution

After all the preceding steps (definition of the geometry, material properties and boundary conditions and creation of the finite element mesh) have been completed, the problem at hand is fully characterised and the solution process can be executed. This is an automated task, during which the finite element analysis program carries out the following operations [86]:

- Computation of the individual stifness matrices for each element, Gaussian integration and matrix multiplication;
- Assembly of the element matrices into the global stifness matrix and triangularisation;
- Solution of the global system of equations, to obtain the values of the degrees of freedom (the magnetic potential, in the case of magnetostatics) at each node;
- Computation of derived quantities (such as the magnetic flux, in the present problem).

While these various steps are automated and executed by the program alone, the user is responsible for specifying some parameters that control how the solution will be carried out. For example, since in the present problem non-linear magnetic permeabilities were specified (for both the steel and MR fluid), the problem becomes

non-linear thus requiring an iterative approach. It is recommended that instead of applying the full load (coil current at once), the problem be divided into various substeps and the current progressively increased, as this accelerates convergence of the problem. For this particular problem, 5 substeps were chosen and given the quick convergence of the solution procedure found to be enough.

Post-processing

Given that it is important to know the weight of each MR brake configuration, the first post-processing task is its determination. This is done by selecting all elements of each material, obtaining from ANSYS their volume and multiplying this by the density of the material.

Once the magnetic field distribution is known, an estimate of the braking torque can be obtained. To achieve this, the following operations were necessary:

- Creation of a routine in the ANSYS input file to select the nodes along the disk/fluid boundaries, obtain the value of the magnetic field on the fluid adjacent to it[7] and save both the coordinates and the magnetic field intensity to an external file;
- Writing of a post-processing program (using Fortran) to calculate the braking torque, carrying out the following steps:
  - Acquisition of the results saved by the ANSYS routine described above;

[7]Since these nodes are on the boundary between the disk and the fluid, which have different magnetic permeabilities, the magnetic field is discontinuous. The nodal result saved by ANSYS is the average of the magnetic field on the disk and that on the fluid. Given its much higer magnetic permeability, the field on the disk is close to 0, which means that the nodal solution given by ANSYS is $H = \frac{H_{disk}+H_{fluid}}{2} \approx \frac{0+H_{fluid}}{2} = \frac{H_{fluid}}{2}$ and hence the magnetic field intensity experienced by the fluid can be obtained using $H_{fluid} = 2H$.

- Determination of the electro shear stress (τ_e) corresponding to the magnetic field at each point, by interpolation of the τ_e vs. H curve included in the fluid properties information.
- Numerical integration along the boundary of each disk of the product $\tau_e r^2$. Given the smooth variation of $\tau_e r^2$ along r, a simple zero-order integration was carried out as given by the following formula:

$$\int_{x_1}^{x_2} f(x)dx = \frac{f(x_1) + f(x_2)}{2}(x_2 - x_1) \tag{4.3}$$

It is seen that this Fortran routine estimates the braking torque taking into account solely the electro-stress contribution τ_e. Given the impossibility of quantifying the contribution of the term τ_p from the magnetostatics analysis[8] and the fact that it represents only a small fraction of the total friction in the MR brakes at high deceleration, this term is neglected. It must be noted that this is a conservative approach in that the braking torque (our target performance) is being underestimated.

Results

Fig. 4.4 presents the distribution of the magnetic flux density for the initial 1 disk design defined in Section 3.5. It is seen that the magnetic flux follows the intended path around the coil, crossing the MR fluid gaps and disk and going along the casing. However, it is seen that on the lateral walls of the casing the flux exceeds 1.7 T, a value outside the linear B-H relation of SAE 1010 steel[9]. This is in line with what was expected and confirms that the magnetic behaviour of steel could not have been defined simply by a constant magnetic permeability, as discussed earlier under the

[8]Recall that $\tau_p = \mu_p \dot{\gamma}$ and the determination of $\dot{\gamma}$ can only be achieved with knowledge of the velocity distribution.

[9]Valid for flux densities up to approximately 1.6 T (refer to Fig. 4.3)

heading "Material Properties". This also means that in the regions where the flux reaches such values, a high intensity of the magnetic field will be present. Recall that the coil current is related to the magnetic field intensity by $NI = H_f L_f + H_s L_s$ (Equation 3.22) and hence, if H_s (the magnetic field experienced by the steel) is not close to zero everywhere, it inevitably follows that for the same number of turns and current (NI), the magnetic field experienced by the MR fluid will be lower than desired. This scenario is confirmed from the analysis of Fig. 4.5 which shows the magnetic field distribution in the same MR brake. Indeed the maximum field intensity is about 42 kA/m, less than half the intended 110 kA/m, which would provide the desired 1010 Nm braking torque. The postprocessor routine described earlier[10] gave a value of 249.77 Nm for the braking torque in this system. Another important conclusion that can be drawn from the magnetic flux distribution is that, as expected, there is no leakage of magnetic flux to the surrounding air: the value of the magnetic flux density is approximately zero everywhere outside the disk. This confirms the validity of the choice of the "flux-parallel" boundary condition.

Consider now a version of the same MR brake with a thicker casing (5 cm instead of 1.5). This should cause the magnetic flux density (magnetic flux per unit area) to decrease, thus bringing the magnetic field intensity everywhere in the casing to very low values and ensuring that the magnetic field is maximum in the MR fluid. Figs. 4.6 and 4.7 plot the distributions of magnetic flux density and magnetic field intensity, respectively, in this new geometry. In the former, it can be seen that the maximum flux density is now below 1.6 T which means that the magnetic field is now very small throughout the steel casing and consequently the magnetic field in the MR fluid gaps increases to over 94 kA/m (twice the value of the original geometry). The torque for this configuration is of 848.17 Nm, still below the expected value of 1010

[10]Under the heading "Post-processing" in the present section

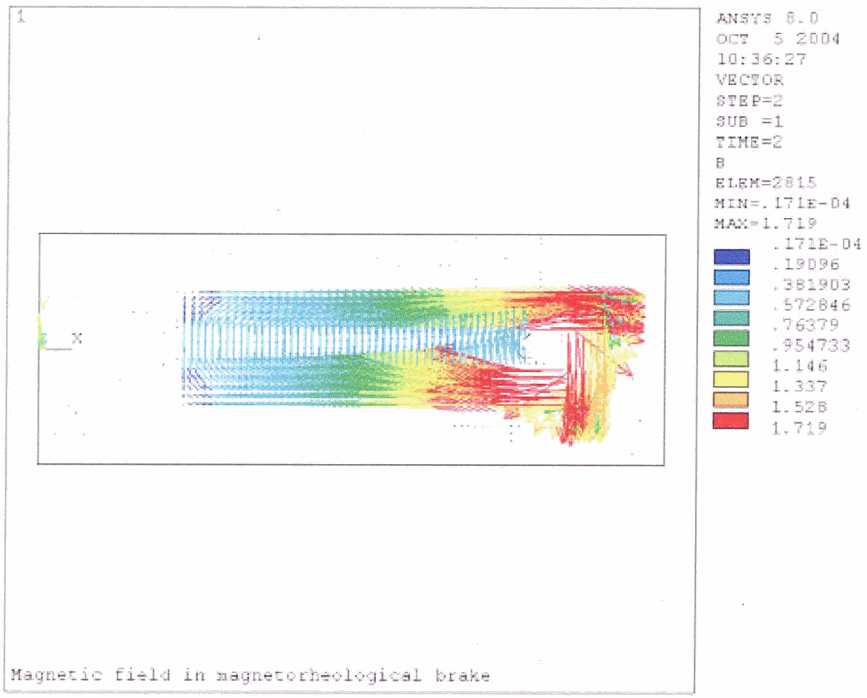

Figure 4.4: Magnetic flux density distribution in the initial 1 disk geometry using MRF-241 fluid

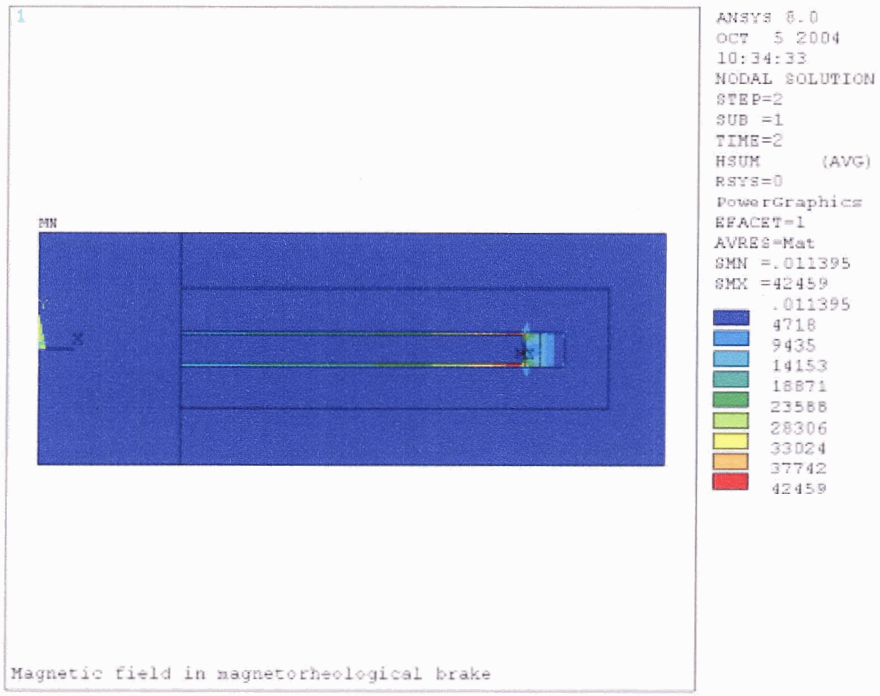

Figure 4.5: Magnetic field distribution in the initial 1 disk geometry using MRF-241 fluid

Nm. This deviation is due to the fact that the maximum magnetic field intensity in the MR fluid gap is still below the desired value of 110 kA/m (which would be the maximum limit attainable with 147 turns and 1.5 A if the magnetic field was zero everywhere in the steel) given that the steel parts also experience some magnetic field, albeit small, that was neglected. Furthermore, the 1010 Nm braking torque would be reached based on the assumption of a constant magnetic field throughout the MR fluid gap. It can be seen that this is not the case, as the portions of fluid closer to the shaft are subject to lower magnetic field intensities. This fact backs the remark made in Section 3.1 that Bydoń's model of the MR brake assuming a constant H would not be accurate for a large MR brake.

It is clear that using a thicker casing addressed the excessive flux density in the steel found in the initial design but it is important to note that this geometry is not a feasible one - it was obtained by increasing the thickness of the casing which means that it no longer respects the maximum radius allowable for the MR brake in a 16" wheel.

Figs. 4.8 and 4.9 present the results for an MR brake with the initial geometry (that of Figs. 4.4 and 4.5) but using instead MRF-132 fluid.

The magnetic field distribution in this configuration (using MRF-132 fluid) is more intense than in the system with the same geometry but using MRF-241 fluid. However, this does not indicate that the braking torque is now higher - indeed, recall that the torque depends on the shear stress exhibited by the fluid and MRF-132 presents lower shear stresses than MRF-241 for identical magnetic field intensities.

Table 4.1 summarises the braking torques of the various configurations as estimated by the postprocessing routine described earlier. As expected (recall that the torque is proportional to the number of surfaces in contact with the MR fluid), the systems with 3 disks exhibit the highest braking torque. Note also that in order to

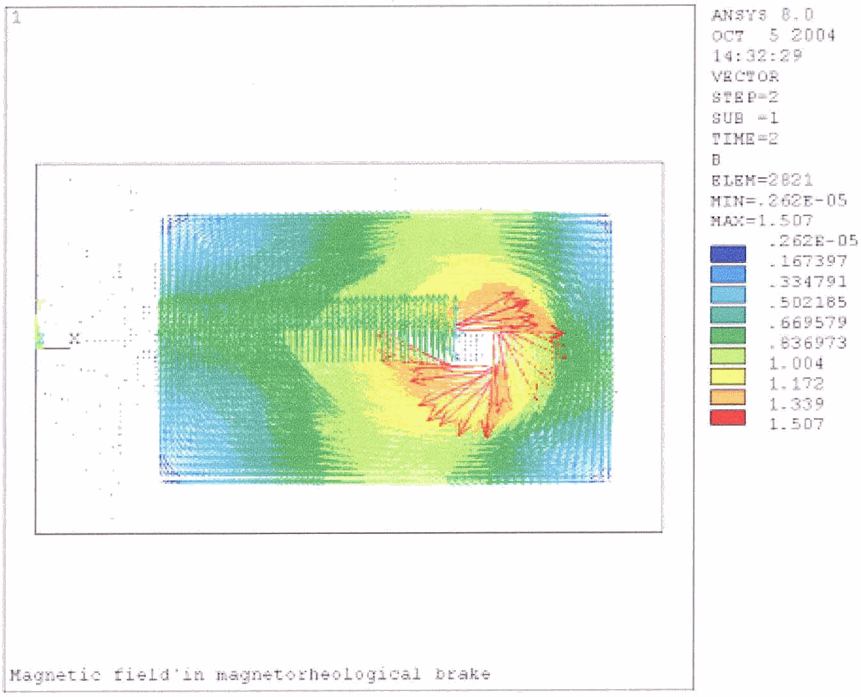

Figure 4.6: Magnetic flux density distribution in the 1 disk geometry using MRF-241 fluid with thicker casing

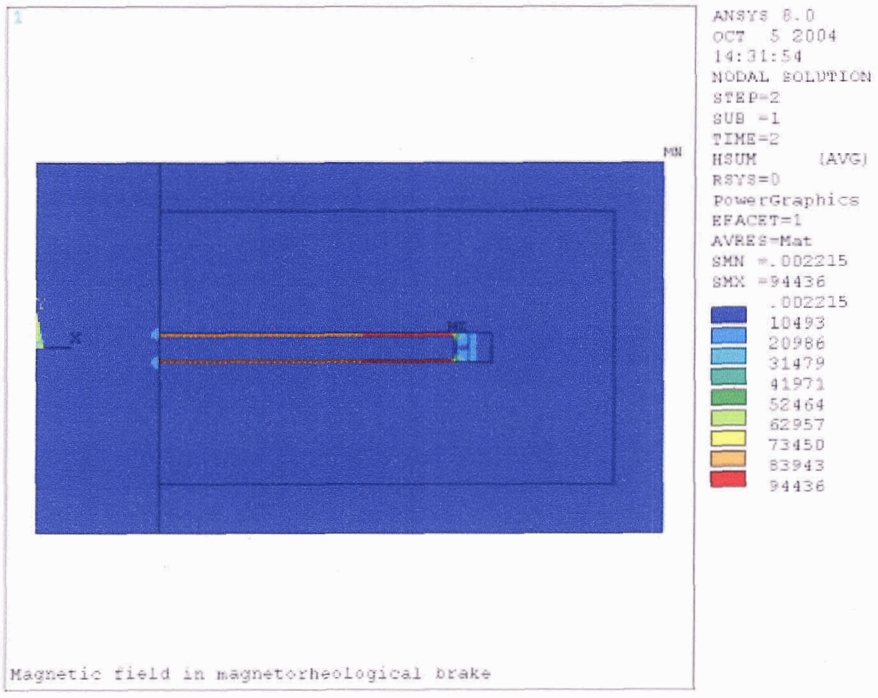

Figure 4.7: Magnetic field distribution in the 1 disk geometry using MRF-241 fluid with thicker casing

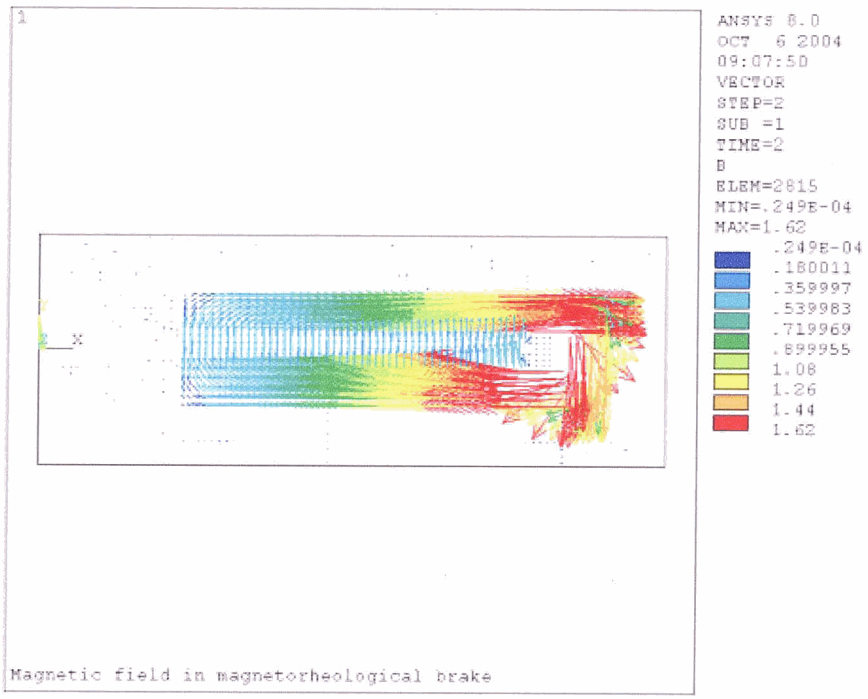

Figure 4.8: Magnetic flux density distribution in the 1 disk geometry using MRF-132 fluid

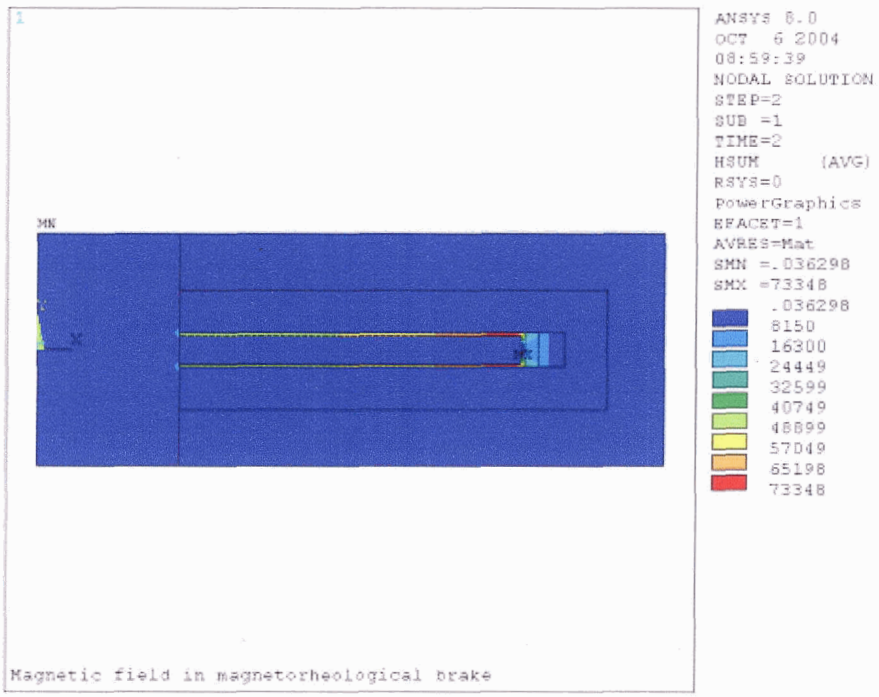

Figure 4.9: Magnetic field distribution in the 1 disk geometry using MRF-132 fluid

take advantage of the higher shear stresses achieved by the MRF-241 MR fluid, a thicker casing is necessary in order to prevent saturation in the steel path. Hence, on the configuration with thick casing (where saturation of the steel will not occur), the best results (highest torque) are achieved with the fluid capable of higher shear stresses (MRF-241) but on the other configurations magnetic saturation limits the flux and hence does not allow MRF-241 to match the torque of MRF-132 (which presents a lower shear stress but requires a much lower magnetic flux).

Table 4.1: Estimated braking torque for different configurations

Configuration	Braking torque [Nm]
1 disk, MRF-132	359.85
1 disk, MRF-241	249.77
1 disk (thick casing), MRF-132	569.96
1 disk (thick casing), MRF-241	848.17
3 disks, MRF-132	1320.2
3 disks, MRF-241	909.91

4.1.1 Transient Effects

The model presented above describes the magnetostatic behaviour of the MR brake, i.e., the magnetic field produced by a constant electric current flowing through the coil, but does not take into account the effects of a time-varying current. Changing the current induces an electromotive force (emf), which in turn influences the current, as stated by Lenz's law: the emf induced by a change in magnetic flux (due to the change in current) is oriented so as to oppose the variation in magnetic flux (and current) [87]. The evolution of the current with time is therefore exponential and given by [88]:

$$I(t) = \frac{U}{R}\left(1 - e^{-\frac{R}{L}t}\right) \tag{4.4}$$

where U is the voltage, R the resistance of the circuit and L its inductance. The term R/L, in particular, is the inverse of τ, which is the time constant of the circuit. This is the time that it will take until the intensity reaches $(1-e^{-1}) = 0.632$ of its nominal value U/R.

Given that one of the motivations for the MR brake is the faster response compared to a conventional system, it is important to ensure that I reaches a value close to V/R shortly after the pedal is pressed, i.e., in a few dozen milliseconds at most. The design of the MR brake is such that a moderate current originates the high magnetic flux required to produce the desired shear stress in the fluid and this implies that the inductance of the circuit will necessarily be high, which may slow down the response of the system. In order to determine how long it will take for the current to reach its nominal value, the exact values of R and L must be determined.

Recalling that the inductance is given by [89]:

$$L = \frac{N^2}{l/\mu A} \tag{4.5}$$

it is seen that reducing the number of turns greatly reduces the inductance, thereby producing a faster responce from the MR brake. Given that the magnetic field intensity is proportional to NI, the current must be increased by the same proportion as the number of turns is decreased. Given the practical limitations in the current (too high a value would lead to exaggerate power consumption), a value of 12 A was chosen. This is 8 times the initial value of 1.5 A, which means that the new number of turns is 1/8 of the initial value.

For values of R and L that do not permit the fast response that is desired, a higher voltage can be used. Since the battery supplies DC (direct current), it is impossible to simply use a traditional transformer as these require AC (alternate current) to

work. However two different possibilities exist:

- An adapter (the ignition coil) already exists in the car to increase the battery supplied voltage to a value of several kV required by the spark plugs. Altough this value is too high and should not be used to supply the MR brake, an adapter similar to the ignition coil but with a much more moderate output voltage can be used in parallel with it.
- The 12V output of the battery can be directly increased to the desired value, using converters. Altough these are more complex than traditional AC transformers, they are nevertheless light and unexpensive.

It is interesting to note that Delphi's electric calliper (v. Section 1.2) is offered to work with 12V and 42V power supplies, given that the latter are becoming increasingly common. If a 42V power supply is available, the transient response of the proposed MR brake is excellent without requiring any modification or voltage adapters. This means that in the future, as 42V becomes a standard value in power supplies, the use of MR brakes will be even more straightforward.

4.2 Fluid Flow and Heat Transfer Model

Once the magnetic field distribution is known, the shear stress can be obtained from the fluid's specifications and, together with the no-field viscosity, describes the behaviour of the fluid flow. A CFD model can then be built to obtain the velocity and wall shear stress distribution (the latter is the friction on the disks that causes the braking torque). Depending on the CFD program used, the temperature distribution may also be obtained. The ANSYS FLOTRAN[11] solver uses the momentum and

[11]The ANSYS' CFD module is named FLOTRAN

mass conservation equations to determine the values of the velocity components and pressure and may optionally solve the energy equation to obtain the temperature distribution. Since this is an important variable in the present problem (it determines whether the MR fluid's operating conditions are met), the energy equation was also solved.

Two elements are available for FLOTRAN analyses: FLUID141 (two-dimensional) and FLUID142 (three-dimensional). Since the present problem is an axisymmetric one and FLUID141 has the option to model axisymmetric systems, it was chosen for the analysis.

It must be stressed that the axisymmetry of the problem means that the geometry, material properties and loads are independent of the tangential coordinate (θ) and thus that all quantities are independent of θ. It does not mean that quantities along the tangential component are zero - in particular, the velocity component along that direction (v_θ) is not zero, only it does not vary along θ.

For problems in which v_θ is not zero throughout the domain, the "swirl" option of FLOTRAN must be activated. This forces the algorithm to solve the 3rd momentum equation (along θ), which it would not normally solve in a two-dimensional problem. It is interesting to note that using FLUID141 with the "swirl" option is not the same as using FLUID142 (the full three-dimensional model). Altough the number of equations that are solved is the same (conservation of mass, 3 components of momentum and energy), the FLUID141 + "swirl" model takes advantage of axisymmetry to cancel all derivatives in the tangential direction (terms $d(...)/d\theta$)

For elements modelling solid materials, there is only one unknowkn (the temperature) and only the energy equation is solved.

Geometry

As it was seen in Section 4.1, the geometrical model is a function of only the form of the system being studied. Hence, the same geometry defined for the magnetostatics analysis is used for the fluid flow and heat transfer study. Note, however, that for the present analysis it is no longer necessary to model the air surrounding the brake whose only purpose was the study of magnetic leakage.

Material Properties

Section 2.3.2 introduced the Bingham model that describes the behaviour of MR fluids. One of the advantages of this model is that it has been extensively studied and is included with ANSYS. However, it must be noted that the ANSYS Bingham model [85] differs slightly from the ideal model. The shear stress in the ideal Bingham model is the sum of that due to the magnetic field (τ_e, independent of the strain rate) and the plastic contribution, similar to the behaviour of Newtonian fluids ($\tau_p = \mu_p \dot{\gamma}$). That is, at zero strain rate the shear stress is positive and equal to τ_e, the magnetic contribution. This poses numerical problems and led to the use in ANSYS of a slightly modified model (the biviscosity model), consisting of two different sections:

- The shear stress is now zero at zero shear strain rate but increases rapidly[12] with the strain rate until the shear strain rate reaches $\frac{\tau_e}{\mu_r - \mu_p}$;
- After this point the shear stress is then given by the usual Bingham equation as $\tau = \tau_e + \mu_p \dot{\gamma}$.

This model overcomes the numerical difficulties associated with the implementation of the ideal Bingham model whilst replicating its behaviour, with differences existing

[12] for a generic fluid, ANSYS [85] suggests for this slope (μ_r) a value equal to 100 times the plastic viscosity μ_p; for the particular case of an ER fluid, Bullough [64] recommends a value of 10^4

only for very low strain rates (up to $\frac{\tau_e}{\mu_r - \mu_p}$). Note that in the biviscosity model the stress at zero strain rate is equal to zero instead of τ_e as would be the case in the Bingham model. The difference between both models is illustrated in Fig. 4.10, where the solid line represents the biviscosity model, combining the pure Bingham model for most strain rates with the Newtonian model for very low strain rates.

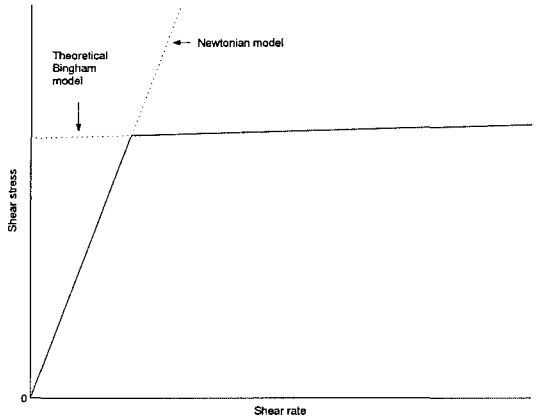

Figure 4.10: Comparison of the Bingham and biviscosity models

Recall from the comments on Bydoń's mathematical model of an MR brake in Section 3.1 that in the present problem using a constant value of the magnetic field is not accurate. At this point, this remark can indeed be confirmed from the magnetic field distributions obtained from the magnetostatics finite element models, where it is seen that the magnetic field intensity varies throughout the fluid. However, ANSYS requires that constant materials properties be specified and therefore a value based on the average magnetic field intensity was chosen for the electro shear stress of the MR fluid (τ_e). This will introduce errors in the values obtained for the wall shear stress (the friction exerted by the fluid on the disk surfaces). Nevertheless, the CFD analysis is important to study the MR fluid behaviour and obtain a good qualitative analysis of the flow within an MR brake, and to determine the temperature distribution.

Boundary Conditions

For the present analysis, boundary conditions are required in terms of the velocity on the fluid boundaries and the temperature in the boundary of the whole MR brake.

The evolution of the temperature throughout the MR brake is determined by the following contributions:

- Generation of heat by friction between the fluid and the disks and casing surfaces;
- Generation of heat due to the flow of electric current in the coil;
- Cooling by convection associated with the flow of air around the system.

The amount of heat generated by friction is computed by ANSYS from the MR fluid flow, and the total for the four wheels equals $\zeta_b \Delta E$, where ζ_b is the fraction of power dissipated by the brakes and ΔE is the change in kinetic energy of the car.

The heat formed in the coil is equal to VI, where V is the voltage drop in the coil and I is the current flowing through it. This contribution is modelled through the inclusion of a heat source (body load HGEN in ANSYS). The maximum continuous electrical power is 144W (for V=12V and I=12A). Compared with the power dissipated by friction (which accounts for most of the total braking power of 54000W [13] at each of the front wheels for maximum deceleration from 100km/h), this will generally be very low and thus not expected to seriously influence the temperature distribution during short stops. However, its influence on the temperature becomes more important for moderate continuous braking and, given that modelling this heat

[13] One third of the braking power determined using Eq. 2.10

source does not increase the solution time or complexity, it was taken into account in all analyses.

The convection cooling is included in the ANSYS solution through application of the surface load CONV, accompanied by the specification of the parameters that characterise convection: the convection coefficient h and the far-field temperature of the cooling fluid T_∞ which in this case is the air temperature. The heat transferred by convection is given by:

$$\dot{Q}_{conv} = h \cdot A \cdot (T_\infty - T) \tag{4.6}$$

where A is the area subject to cooling and T is the temperature at the surface of the component. The temperature value of the International Standard Atmosphere (ISA) at sea-level (15°C) was used for T_∞. As for h, the physical meaning of this parameter becomes clear if the above equation is rearranged into the following form:

$$h = \frac{\dot{Q}_{conv}}{A \cdot (T_\infty - T)} \tag{4.7}$$

The convection coefficient h is nothing more than the amount of heat transferred per unit area and unit temperature difference between the surface of the component and the far-field convective fluid.

The determination of h is the central problem in the study of convection. Equation 4.7 cannot be used to find h when analysing a new system since $\dot{Q}_{conv}$ is unknown. However, it is possible to experimentally measure the heat transfer on a known system and thus obtain h. Unfortunately, due to the complexity of the phenomenon of convection, such a value for h will only be valid for systems in the exact same conditions for which the measurement was made. This means that without a way to extrapolate the results from one configuration to another, the only way to determine h (which is

necessary to calculate the convective heat transfer) would be to experimentally measure the actual convective heat transfer. It would therefore be impossible to predict theoretically the value of $\dot{Q}_{conv}$. To overcome this obstacle, empirical relations for h have been develloped in the past, combining the theoretical knowledge of the process of convection with experimental values.

Many different expressions for the convection coefficient exist, given that it depends on a variety of factors, such as:

- Whether the flow is internal (e.g., water inside a pipe) or external (e.g., flow of air around a car);
- Whether there is forced convection (i.e., external forces such as a fan promote the convective flow) or free convection (i.e., the fluid is not set in motion by any external elements and any movement of fluid particles is due only to buoyancy associated with temperature differences);
- The geometry of the component;
- The flow orientation (parallel to the component's surface or normal to it);
- The nature of the flow (laminar, turbulent or mixed);
- The fluid properties.

The cooling of the MR brake is a problem of external flow with forced convection (as the air flows past the brake at a speed equal to that of the car). For these conditions and the geometry of the MR brake casing (flat walls subject to parallel flow) the most appropriate correlation to obtain h is that for forced parallel flow over a flat plate. For this particular case, different correlations exist depending on

the nature of the flow. This can be determined based on the Reynolds number, a dimensionless parameter of fluid flow defined as:

$$Re = \frac{V \cdot L}{\nu} \tag{4.8}$$

where V is the flow velocity, L is the reference length of the geometry and ν is the kinematic viscosity of the fluid. This is obtained from the dynamic viscosity μ as $\nu = \mu/\rho$ where ρ is the fluid density.

Another important dimensionless parameter is the Prandtl number, which relates the velocity and temperature boundary layers. It is expressed as:

$$Pr = \frac{\nu}{\kappa} \tag{4.9}$$

where κ is the thermal diffusivity of the fluid. This, in turn, can be obtained from the thermal conductivity (k), density (ρ) and specific heat at constant pressure (C_p):

$$\kappa = \frac{k}{\rho C_p} \tag{4.10}$$

Hence, the Prandtl number can be expressed as:

$$Pr = \frac{\nu \rho C_p}{k} \tag{4.11}$$

and for air at 15°C this equals 0.71657.

Once in possession of these parameters that describe the flow and fluid characteristics, the following equations [90] can be used to determine the local (h) and average ($\overline{h}$) values of the convection coefficient in laminar flow. They are based on empirical

correlations based on another dimensionless parameter, the Nusselt number (Nu):

$$Nu_x \equiv \frac{h_x x}{k} = 0.332 Re_x^{1/2} Pr^{1/3} \quad Pr \geq 0.6 \tag{4.12}$$

$$\overline{Nu}_x \equiv \frac{\overline{h}_x x}{k} = 0.664 Re_x^{1/2} Pr^{1/3} \quad Pr \geq 0.6 \tag{4.13}$$

In turbulent flow, the correlation is [90]:

$$Nu_x \equiv \frac{h_x x}{k} = 0.0296 Re_x^{4/5} Pr^{1/3} \quad 0.6 < Pr < 60 \tag{4.14}$$

Incropera does not explicitly include the equation for the average value, but this can be obtained from the following expression for $\overline{h}_x$:

$$\overline{h}_x = \frac{1}{x}\int_0^x h_x dx = 0.0296 \frac{k}{x} Pr^{1/3} \left(\frac{u_\infty}{\nu}\right)^{4/5} \int_0^x \frac{1}{x^{1/5}} dx \tag{4.15}$$

And hence:

$$\overline{Nu}_x \equiv \frac{\overline{h}_x x}{k} = 0.037 Re_x^{4/5} Pr^{1/3} \quad 0.6 < Pr < 60 \tag{4.16}$$

Finally, for the case of a mixed boundary layer (comprising both a laminar and a turbulent part), h can be obtained from the following equation [90]:

$$\begin{aligned} &\overline{Nu}_x \equiv \tfrac{\overline{h}_x x}{k} \\ &= \left[0.664 Re_{x,c}^{1/2} + 0.037\left(Re_L^{4/5} - Re_{x,c}^{4/5}\right)\right] Pr^{1/3} \end{aligned} \quad \left\{ \begin{array}{c} 0.6 < Pr < 60 \\ 5 \times 10^5 < Re_L \leq 10^8 \end{array} \right\} \tag{4.17}$$

where $Re_{x,c}$ is the transition Reynolds number. If this is assumed to be equal to

5×10^5, the above relation becomes:

$$\overline{Nu_x} \equiv \frac{\bar{h}_x x}{k} = \left(0.037 Re_L^{4/5} - 871\right) Pr^{1/3} \quad \left\{ \begin{array}{c} 0.6 < Pr < 60 \\ 5 \times 10^5 < Re_L \leq 10^8 \end{array} \right\} \tag{4.18}$$

It is clear from the above relations that the convection coefficient is higher under turbulent flow. Given that we intend to have as high a convection coefficient as possible in order to maximise the cooling of the brake, it is desirable to ensure that the flow of outside air around the casing is turbulent. The most common way to force transition to turbulent flow in low-speed aerodynamics is to use a transition wire [91]. This acts as a cylinder normal to the flow and its influence on the boundary layer velocity profile promotes transition of the flow to turbulent. When implementing this technique, one question arises: how big should the transition wire be? If the wire is too thick, it will bring an exagerate penalty to the overall flow; if, on the other hand, the wire is made too thin, it may be unable to promote transition to turbulent flow. It is therefore important to know the ideal dimension of the transition wire and this can be obtained from the Gibbings criterion [92]:

$$Re_{wire} = \frac{U \cdot d_{wire}}{\nu} \geq 826 \tag{4.19}$$

where d_{wire} is the diameter of the wire.

Meshing

While meshing is generally a simple step for structural problems (as well as for other scientific areas, such as the example of magnetostatics seen above), it becomes critical

for CFD problems. Altough some general suggestions exist on how to create a good[14] CFD mesh (such as refining the mesh close to the walls, where the gradients are higher), the ideal mesh for each problem can only be obtained from trial and error and after a good understanding of the problem has been obtained.

Results

The tangential velocity throughout the MR fluid is shown in Fig. 4.11. As expected, the no-slip condition causes the portion of fluid adjacent to the rotor to move at the same velocity as the rotor and the portion of fluid adjacent to the stator to present no motion. It is this velocity gradient that is responsible for the no-field component of the MR fluid's shear stress. The sum of this component with the contribution from the magnetic field gives the total wall shear stress, represented in Fig. 4.12 and in greater detail in Fig. 4.13.

Fig. 4.14 presents the distribution of the wall shear stress for the 3 disks geometry. Careful analysis of this figure reveals an important problem: the wall shear stress exhibits high (desired) values on the fluid gaps between each of the end disks and the stator close to it (the topmost and bottommost fluid gaps) but not on the gaps between the various disks (the two central fluid gaps). That is, of the 6 disk surfaces in contact with MR fluid only 2 (those across from the stator walls) are experiencing significant friction. The remaining 4 surfaces (those facing another disk) suffer minimum braking shear, due to the inability of such narrow fluid channels to withstand considerable shear. The solution to this problem lies in having each disk surface facing a static surface. This is achieved by fitting a stator between every two disks. The wall shear stress distribution for such a geometry (having two disks with a stator between them attached to the static ring between the fluid and coil) is

[14] A good mesh is one that leads to quick convergence and accurate results

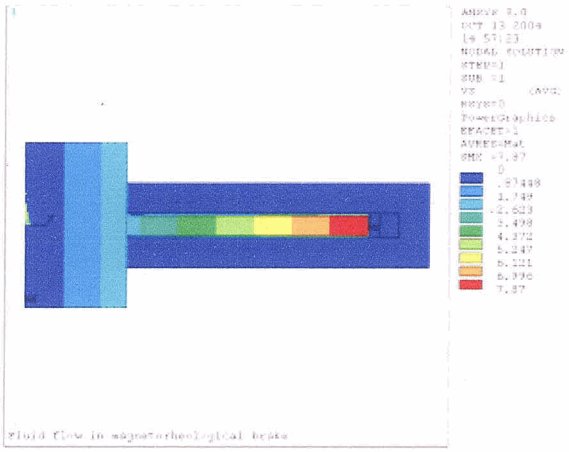

Figure 4.11: Tangential velocity distribution in the 1 disk geometry using MRF-132 fluid

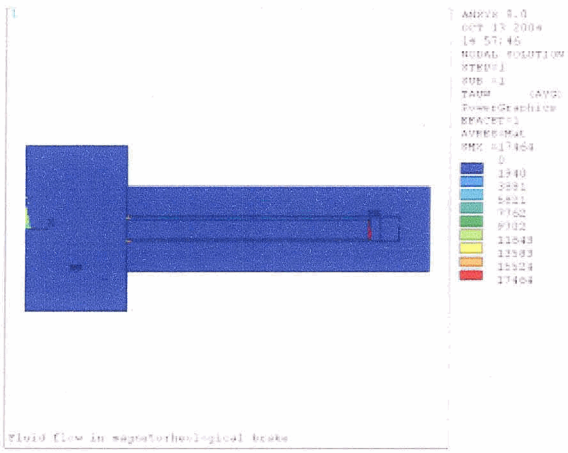

Figure 4.12: Maximum wall shear stress distribution in the 1 disk geometry using MRF-132 fluid

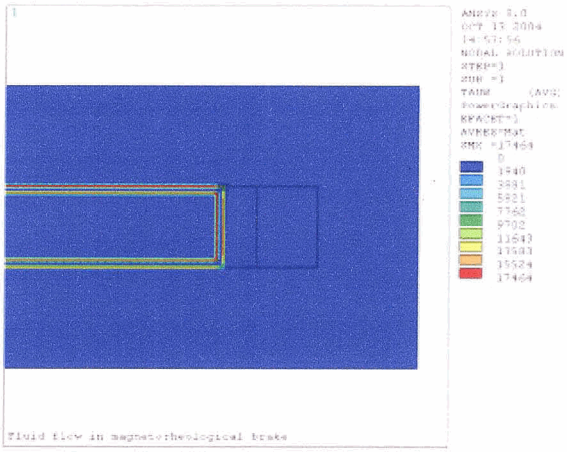

Figure 4.13: Detail of the maximum wall shear stress distribution in the 1 disk geometry using MRF-132 fluid

presented in Fig. 4.15, where it is seen that the desired maximum shear stress is now present in the entirety of the surfaces of all disks.

The steady-state temperature associated with constant braking at a moderate deceleration (0.05g, corresponding to a long downhill run) is presented in Fig. 4.16. While this value is of some interest, it is more useful to know the temperature after a series of alternate braking and accelerating periods. This will be done later when the transient response of the system is analysed.

Given that the MR fluid has some viscosity even without a magnetic field applied, some retarding torque will always be present, even when the brake pedal is not pressed. Naturally, it is desired that this torque be as small as possible, as it is an extra source of drag and hence increases the fuel consumption. Also, the permanent presence of this small friction also contributes to heat buildup and it is important to analyse the resulting temperature distribution. Given that the only contribution to

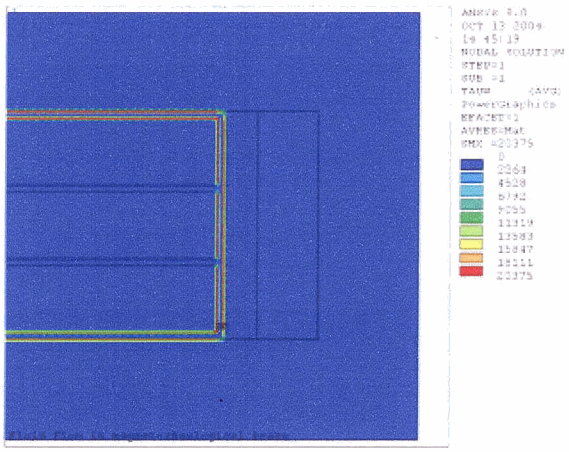

Figure 4.14: Details of the maximum wall shear stress distribution in the 3 disks geometry using MRF-132 fluid

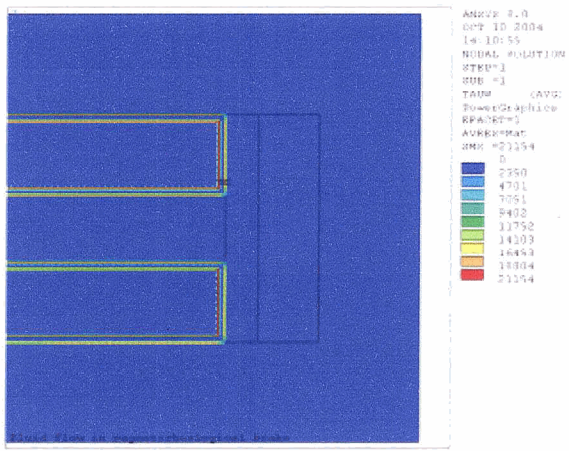

Figure 4.15: Detail of the maximum wall shear stress distribution in the 2 disks + stator geometry using MRF-132 fluid

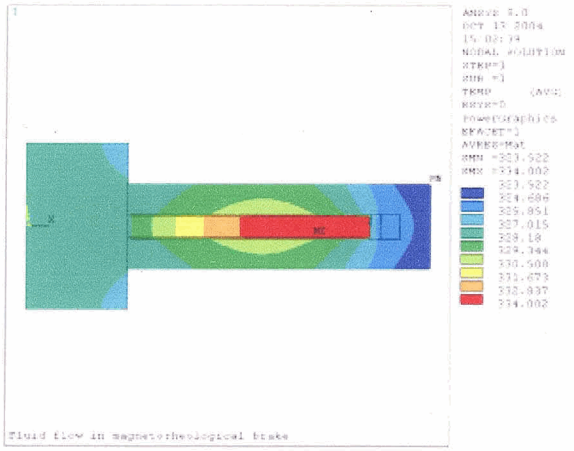

Figure 4.16: Steady-state temperature distribution in the 1 disk geometry using MRF-132 fluid

the friction without a magnetic field applied is the Newtonian one, the dependence on the velocity is now important[15]. For this reason, the no-field behaviour of the MR brake was studied at two different speeds: 50 and 150 km/h. The resulting distributions of wall shear stress and steady-state temperature are shown in Figs. 4.17 through 4.20. Regarding the wall shear stress, it is seen that as expected its dependency on the car's velocity is linear[16]: at 50 km/h, the maximum wall shear stress is equal to 754 Nm versus 2125 (approximately three times more) at 150 km/h. Analysis of the temperature distributions, on the other hand, reveals higher than desirable values: at 150 km/h, the temperature reaches 100 °C. While this is well within the fluid's operating limits (130 °C for MRF-132) and is only reached after a

[15]Recall that in the case of strong decelerations the velocity was not important as the component of the shear stress due to the magnetic field is much more important than the Newtonian one and does not depend on the velocity.

[16]Recall from Section 3.1 that without a magnetic field applied the only contribution to the shear stress is that due to the Newtonian behaviour of the carrier fluid, equal to $\mu_p \frac{\omega r}{h}$.

suficient time running at 150 km/h, it is a high temperature considering that it is solely due to the car's motion without any braking. This suggests that the design of the brakes may have to be modified in order to increase cooling. This can be achieved by adding fins to the casing so that the area exchanging heat with the surrounding air is increased or by making the shaft hollow in order to allow flow of cooling air within. Another possibility (commonly used in car racing) is to use a converging duct[17] - this ensures that the air used to cool the brakes is flowing faster than the outside air (and hence a higher convection coefficient) and at a lower temperature.

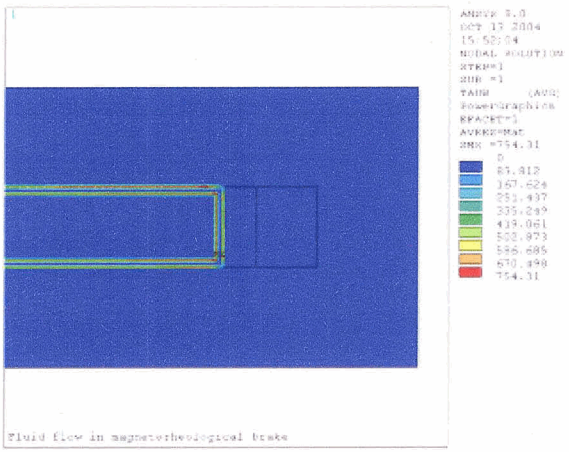

Figure 4.17: Detail of the no-field wall shear stress distribution in the 1 disk geometry using MRF-132 fluid, at 50 km/h

Furthermore, it is important that some braking torque be exerted even at zero speed, so that the car can be kept stopped on steep hills. The quantification of the

[17] As the name implies, this is simply a duct whose cross-section narrows between the inlet and the outlet. By conservation of mass, this causes a higher velocity at the outlet, which in turn leads to a lower pressure and temperature.

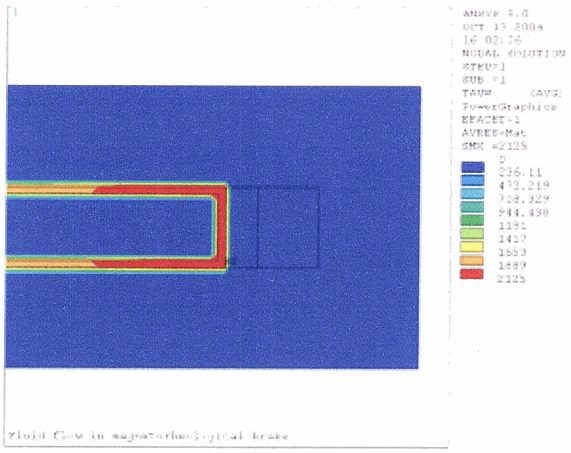

Figure 4.18: Detail of the no-field wall shear stress distribution in the 1 disk geometry using MRF-132 fluid, at 150 km/h

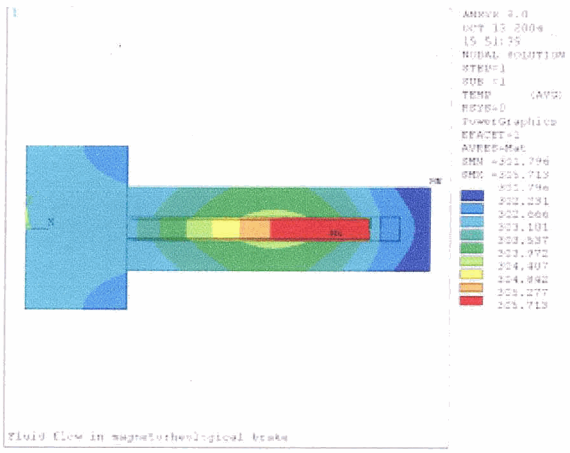

Figure 4.19: Steady-state temperature distribution in the 1 disk geometry using MRF-132 fluid, at 50 km/h

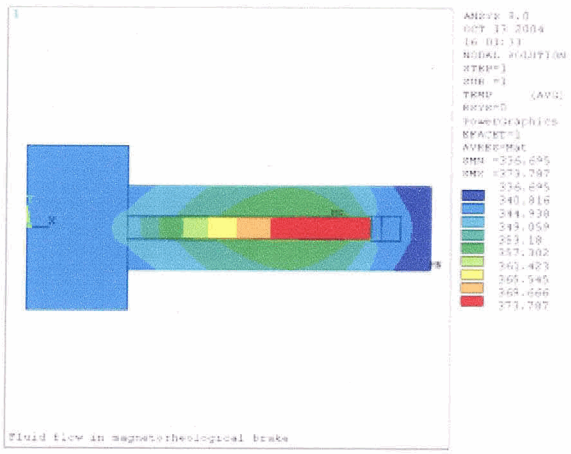

Figure 4.20: Steady-state temperature distribution in the 1 disk geometry using MRF-132 fluid, at 150 km/h

torque at zero speed is presented in Fig. 4.21, and is seen to be negligible. This is due to the biviscosity model used by ANSYS, which results in zero shear stress when no strain rate is present, contrary to the true Bingham model.

Figs. 4.22 through 4.23 show the distribution of the wall shear stress at low velocity (5 km/h) with different values for the Newtonian viscosity. It was seen above that in the theoretical model, the Newtonian viscosity would be equal to infinity. Lower values of the Newtonian viscosity make the convergence of the problem easier but may underestimate the shear stress as can be seen in Fig. 4.22. Therefore, higher values should be used unless convergence problems are encountered.

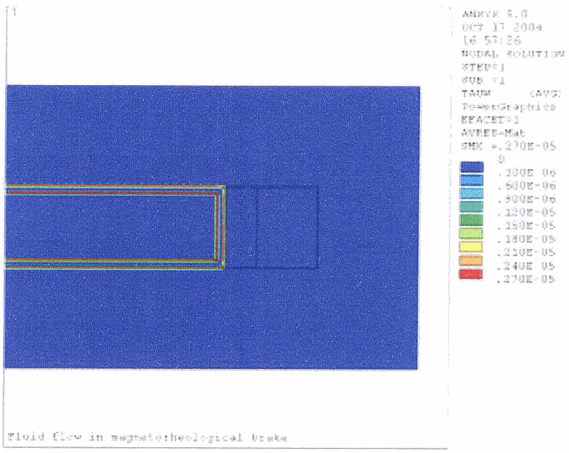

Figure 4.21: Detail of the wall shear stress distribution in the 1 disk geometry using MRF-132 fluid, at zero speed

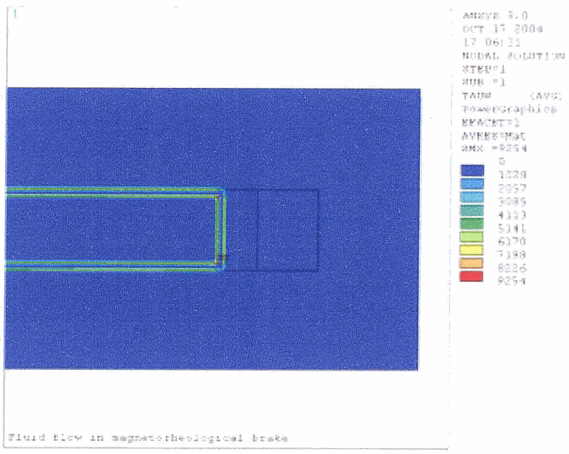

Figure 4.22: Detail of the wall shear stress distribution in the 1 disk geometry using MRF-132 fluid, at 5 km/h, with the Newtonian viscosity equal to 100 times the plastic viscosity

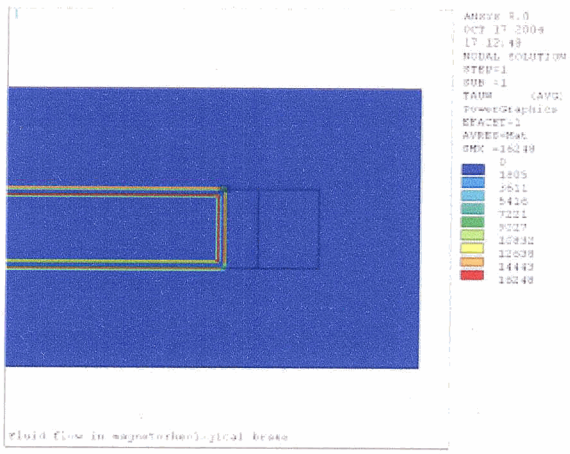

Figure 4.23: Detail of the wall shear stress distribution in the 1 disk geometry using MRF-132 fluid, at 5 km/h, with the Newtonian viscosity equal to 1000 times the plastic viscosity

4.2.1 Transient Model

Once again, we are interested not only on the steady-state solution, but also on the response of the system to a changing command. A very important problem in this category is the calculation of the temperature distribution following repeated brake-acceleration cycles. This corresponds more closely to the everyday operation of car brakes than the long, constant, moderate braking modelled by the steady-state solution.

However, whereas the closed-form solution for the change in current over time was sufficient for the study of the transient behaviour of the magnetic field (which is proportional to the current), it is not enough to obtain the change in temperature distribution with time, which results not only from the heat generated by friction but also from heat buildup in the system and the cooling on the outer surfaces (which does not depend on the current being input to the system). Hence, it is necessary to perform a transient finite element analysis of heat transfer to obtain the temperature distribution associated with a changing braking torque.

ANSYS can run transient problem analyses and the procedure is similar to that of steady-state solutions, with the difference that loads and boundary conditions can vary with time and the solution will be performed for each time step. The material properties, however, may not change over time. This is a limitation common to finite element programs, as allowing time-varying material properties would imply the recalculation of the stifness matrix at every time step. This limitation poses a problem to the analysis of the MR brake, where the fluid viscosity changes with time, as the magnetic field induced by the current in the coil varies. Hence, a different approach had to be sought in order to model the effect of alternate braking and non-braking periods. This was achieved by varying the boundary conditions instead of the material properties: on the accelerate steps, the velocity of the disk(s) and shaft

was set to zero, so that no friction occurs (v. Fig. 4.21).

Another detail specific to transient analyses is the need to specify initial conditions in addition to the boundary conditions. This is physically intuitive in that when determining the behaviour of the system over time it is necessary to know the state it had in the beginning - clearly, turning on the heater on a room initially at 10 degrees will result in a different temperature evolution than the one obtained when the same heater is turned on in the same room if it is initially at 0 degrees. The initial temperatures in each of these examples are the initial conditions of each problem. Note that this temperature must be specified for every point in the room. That is, the initial condition takes the form $T(x, y, z, t = 0) = T_i(x, y, z)$, where T_i is the known initial temperature at each point. For the axisymmetric case of the MR brake, this is expressed as:

$$T(r, z, t = 0) = T_i(r, z) \tag{4.20}$$

and similarly for the velocity components. At this point, a contrast can be highlighted between the form of the initial conditions and that of the boundary conditions: the former are defined for every point at $t = 0$ whereas the latter must be defined at each instant t but only for the points along the boundary of the system. Finally, the mathematical meaning of the initial conditions is also obvious: when solving a transient problem, the equations describing it become also differential in t as terms of the form $\frac{\partial}{\partial t}$ are included. In order to eliminate the integration constants associated with these terms and obtain a unique solution to the problem, the value of the unknown(s) at a given time t must be specified.

For the problem at hand the initial temperature was set to T_∞ (the far-field reference temperature) throughout the domain. The physical meaning of this assumption is that the finite element analysis starts with the car at rest.

Finally, the exact brake-accelerate schedule must be defined. This can be done

with reference to existing studies of hydraulic brakes. The same schedule presented in [93] was adopted, consisting of 15 cycles each comprising a 28 s acceleration period followed by a 0.3 g deceleration from 120 km/h to a full stop.

Post-processing

The ANSYS results file contain the values of all variables at every time-step. In order to analyse the evolution of the maximum temperature in the MR brake over time, it proved useful to create a post-processing tool that automatically gathers this information. Its operation consists of reading the results for each time-step, determining the maximum temperature and writing its value to a file. This file is later read in Excel, Matlab or another suitable program to create the temperature vs. time graphs.

Results

The evolution of the temperature with time for the 1 disk configuration based on MRF-132 fluid is presented in Fig. 4.24.

4.3 Structural Model

Whereas the performance of the brake system is determined from its magnetic and fluid flow behaviour and the MR fluid's capacity to operate at a certain viscosity depends on the temperature distribution, the resistance and longevity of the parts depend on their structural demands. It is therefore important to know the stresses present in the various components and ensure that they satisfy each material's limits.

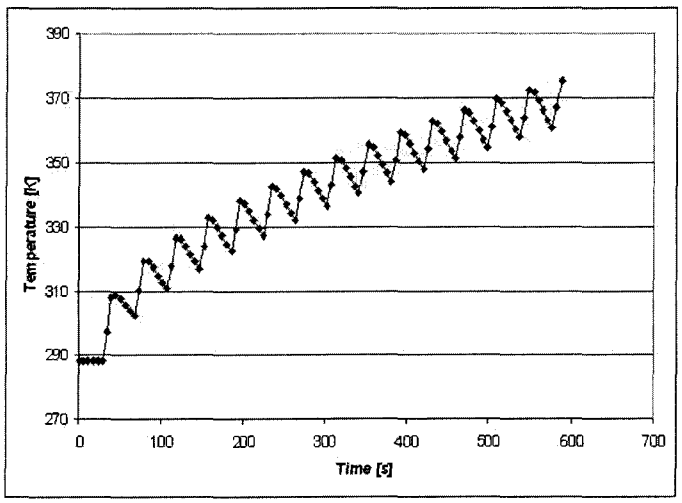

Figure 4.24: Evolution of the maximum temperature in the 1 disk geometry using MRF-132 fluid, subject to repeated brake-release cycles

The most delicate situation is that of the aluminium ring between the fluid gap and the coil in the 2 disk geometry (where it holds the stator), given its small thickness and the loads it is subject to (the weight of the stator and mostly the torque resulting from the friction exerted on the stator's walls during braking).

Components such as the shaft or brake disk are not expected to pose any problem, as their dimensions are identical to those used in conventional disk brakes and so are the loads.

Given the symmetry of the ring, it is possible to analyse only half its length. The advantage of this approach is that the half-ring is not statically indeterminate, unlike the full ring (which is clamped at both ends). Studying a statically determinate system greatly simplifies the calculations. Still, it was decided to validate this approach by running the FE models of the ring for both cases (full length and half length). The exact same results were obtained, confirming the validity of the approach.

For a shaft fixed at one end and subject to a torque on the other end, the displacement angle Φ and the maximum shear stress τ_{max} are:

$$\Phi = \frac{TL}{JG} \tag{4.21}$$

$$\tau_{max} = \frac{Tr_o}{J} \tag{4.22}$$

where T is the applied torque, L is the length of the shaft, J is the polar moment of inertia, G is the shear modulus and r_o is the outer radius of the shaft. For a shaft of hollow circular cross-section, the polar moment of inertia is given by:

$$J = \frac{1}{2}\pi \left(r_o^4 - r_i^4\right) \tag{4.23}$$

with r_i being the inner radius of the shaft.

When conducting the structural analysis of a hollow shaft in ANSYS, the element PIPE16 was chosen. This is a uniaxial element with tension-compression, torsion, and bending capabilities. Being a uniaxial element, the computed values are a function of one coordinate only (the length or axial coordinate). This is all that is required for the present analysis, as knowledge of the distributions of displacement and stress along a cross-section of the component is not required but only the determination of the maximum value at each point along its length.

The ring of the 2 disk + stator configuration was selected for analysis, given that it is the component under most adverse conditions: it must sustain the braking torque on the stator, which is equal to half the total braking torque.

Geometry

The definition of the geometry consists simply of specifying the coordinates of the endpoints of the ring, its outer radius and its thickness.

Material properties

The required material properties for this analysis are the Young's modulus (E) and Poisson's ratio (ν) for aluminium. Note that this is a pure torsion problem and hence the material property present in the equations is the shear modulus G. However, ANSYS computes this internally based on the values of E and ν using the expression:

$$G = \frac{E}{2(1+\nu)} \tag{4.24}$$

Boundary conditions

Both endpoints of the ring are clamped to the casing, thus implying a no-displacement and no-rotation boundary condition at both ends.

As for the applied load, it was seen above that since the ring is holding the stator, it will be subject to half the total braking torque. In the worst-case scenario, this will mean a braking torque of 1010/2 Nm.

Results

The maximum displacement and stress are shown in Table 4.2. It is seen that the models exhibit an excellent agreement. This validates the use of the analytical model presented above to the present problem. In what regards the actual numerical values, it is seen that even at maximum braking torque, both the displacement and the stress

are very low. The latter, of particular importance for the structural resistance of the ring, is 3 orders of magnitude lower than the yield stress of aluminium (v. Table 3.4).

Table 4.2: Comparison of the deformation and maximum stress associated with the braking torque, according to different models

	Analytical model	Finite element model
Deformation angle [radians]	9.52×10^{-7}	9.52×10^{-7}
Max. stress [Pa]	270794	270792

4.3.1 Fatigue

Materials subject to repetitive stresses over a large number of cycles will progressively lose strength and may eventually fail as cracks form in stress concentration points and later propagate throughout the component. This process of crack nucleation and propagation as a result of repeated loads is called fatigue. Smith [76] indicates that according to some estimates approximately 80% of the failures in machines are directly associated with the action of fatigue fractures.

The fatigue behaviour of a material is obtained from experimental testing of a part subject to alternating loads until it fails. This procedure is repeated at different stresses and the results are plotted in the S-N curve[18], which indicates the number of cycles (N) until the part failed at an applied stress S. The resulting S-N curve can then be used to determine the number of cycles a part will be able to withstand when subjected to a known stress. Conversely, if the life of the component is defined a priori, the maximum stress can be obtained from the S-N curve.

[18] also known as the Wöhler curve [74]

It has been seen that for some materials (primarily ferrous materials), applied stresses up to a certain value (named the fatigue endurance limit), will not lead to failure by fatigue, regardless of the number of cycles. Altough the fatigue endurance limit depends on a variety of factors such as the surface finish, the temperature, the medium in which the material is operating and others, it has been found that for ferrous materials a good approximation is to use a value equal to half the ultimate strenght. Most other materials do not possess such an absolute value, altough the slope on the S-N curve also decreasess significantly (in absolute value), i.e., as the applied stress decreases the number of cycles increases ever more rapidly. Qualitative S-N curves for steel (dashed line) and aluminium (dotted line) are plotted in Fig. 4.25 (note that the number of cycles is in a logarithmic scale).

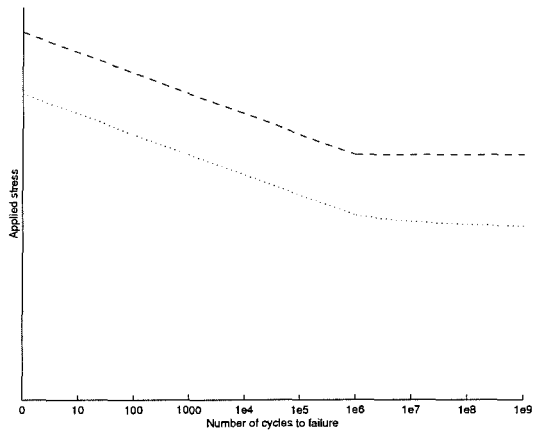

Figure 4.25: Typical S-N curves for ferrous and non-ferrous materials

It must be noted that due to a variety of factors (such as the exact dimensions of the test specimen, the surface finish and homogeneity of the material and the precision of the testing equipment), the number of cycles to failure at a certain stress may sometimes vary one order of magnitude from one specimen to another [94]. It is therefore important to keep in mind that the results expressed by the S-N curve are

statistical values and that as the applied stress approaches the value given by the S-N curve for a given number of cycles, the possibility of failure increases. Hence, when designing a component it is important to ensure not only that it barely complies with the S-N curve for the material in question, but that a good safety margin exists.

In order to determine whether there is risk of failure by fatigue, the stress on the ring (which was found to be equal to 271 kPa in the most adverse circumstances, v. Table 4.2) must be compared with the S-N curve values for aluminium. A comparison of the S-N curves for various aluminium alloys is found in [95], and it is seen that even in the worst case (2017-T4 alloy) the number of cycles tends to infinity as the applied stress approaches 120 MPa, a value 3 orders of magnitude above that present in the MR brake ring. It is therefore safe to conclude that the this component is not prone to failure by fatigue.

4.3.2 Creep

Fatigue is not the only phenomenon that lowers the resistance of a material. A constant load (even if the yield stress of the material is not exceeded) for an extended time at high temperatures may cause permanent deformation, and eventually lead to failure. This phenomenon is known as creep and is responsible for a large number of failures in components used at high temperatures [74].

A representation of the strain as a function of time, at constant temperature and stress, is known as the creep curve and consists of three different stages [94]:

- Primary creep, during which the change in strain decreases with time;
- Secondary creep, during which the change in strain remains constant (i.e., the strain progresses at a constant rate over time);

- Tertiary creep, during which the change in strain increases rapidly until fracture occurs.

A typical creep curve is shown in Fig. 4.26, with the three different stages marked.

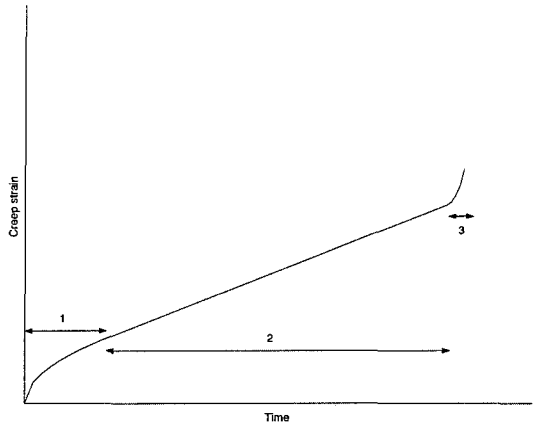

Figure 4.26: Typical creep curve

Whereas the phenomena involved in the growth of the elongation by creep is complex (v. e.g., [94] for a good description of the mechanisms associated with each of the three creep stages), its effects are relatively simple and the creep behaviour of a given material from an engineering point of view is generally represented by the time to failure at a certain applied stress and temperature. Hence the stress values determined above for the MR brake ring can be compared with existing experimental data [96, 97]. These include tests of specimens of various aluminium alloys for lifetimes of up to 100,000 hours and at temperatures comprised between 25°C and 300°C. Given that the temperature in the MR brake shall not exceed 130°C (due to the operating temperature envelope of the MR fluid), the data obtained from tests at 150°C (the closest temperature) was used. Under these conditions, it is seen that in the worst case rupture occurred for an applied stress of 26 MPa for 1100-O aluminium. This is

2 orders of magnitude above the stress faced by the ring in the MR brake, therefore confirming that it is not prone to creep rupture.

4.4 Summary

Using the models developped in the present chapter, it is possible to improve the design by trial-and-error, i.e., by changing this or that parameter in the design and observing the effect produced by that change. However, given the number of parameters involved (number of disks, type of fluid, dimensions) and the fact that they are coupled (changing one parameter alters the influence of other parameters), this non-scientific approach would hardly lead to a truly ideal system. For this reason, a scientific means of improving the design based on the results of the analyses models now built must be introduced. That is the purpose of the next chapter.

Chapter 5

Optimisation

An optimisation problem consists of finding the values of the design variables (e.g., the dimensions of a component) that minimise an objective function (e.g., the cost or the weight of the component) while satisfying certain requirements, known as the constraints of the problem (e.g., the load that the component must be able to support) and can be mathematically formulated as:

$$find\ the\ minimum\ of\ f(x)$$

$$subject\ to:\ g(x) \leq 0$$

$$with\ x_{i\min} \leq x_i \leq x_{i\max}\ for\ 1 \leq i \leq N$$

where $f(x)$ is the objective function, x is the vector containing the N design variables, $g(x)$ are the constraints[1] and $x_{i\min}$ and $x_{i\max}$ are the minimum and maximum values respectively for design variable i.

The goal of optimisation can then be summed up as getting the most results out

[1] defined in ANSYS as state variables of the problem

of the least means. It has been studied for long, but the massive demands of World War 2 in terms of military performance prompted the development of what is now known as "operations research": scientific tools to improve the decision-making in the areas of logistics and operations, and where optimisation techniques play a central role [98].

The discipline of optimisation comprises a tremendous variety of tools, each with specific strengths and weaknesses. These determine the choice of the optimisation method to be used to address each problem. A good overview of optimisation techniques is found in [99]. Before discussing in greater detail the selection of the method to be used, some important aspects of optimisation problems that influence the way in which they are addressed and solved will now be introduced.

At this point, a distinction can be made between constrained optimisation (where only those designs that respect some imposed conditions are accepted) and unconstrained optimisation, where no restrictions exist other than the interval of each design variable (the $x_{i\min}$ and $x_{i\max}$ in the formulation above). Unconstrained optimisation is a simpler problem and many techniques exist to address it, but relatively uncommon in actual engineering problems. It is however possible to adapt unconstrained optimisation algorithms to solve constrained problems, and this is described in Section 5.2.

Another important factor to take into account is the possibility of existence of several local minima. Whereas a polynomial of the second degree will only have one extreme, irregular functions can have many thus making it hard for an optimisation algorithm to know when the true global optimum has been found. This is illustrated in Fig. 5.1: the function on the left only has one minimum and regardless of where along that curve the optimisation procedure begins, it will converge to the minimum; the figure on the right, on the other hand, illustrates a function with several local minima

and it is easy to see how an optimisation procedure can have problems - suppose the initial design has $x = 0$, the algorithm will progress along the curve as the objective decreases but, as it reaches $x = 0.25$, no further reduction in the objective function can be achieved and the program terminates giving $x = 0.25$ as the design variable value that yields the best design. This is not true, since the actual optimum is found for $x = 0.75$. This example shows the difficulty associated with optimisation of more irregular functions. There are several techniques to overcome the obstacles posed by the existence of many local minima, such as starting the optimisation procedure various times from different initial designs or using a method that will accept uphill movements (i.e., movements that progress towards a worse value of the objective function), as is the case with random-search methods.

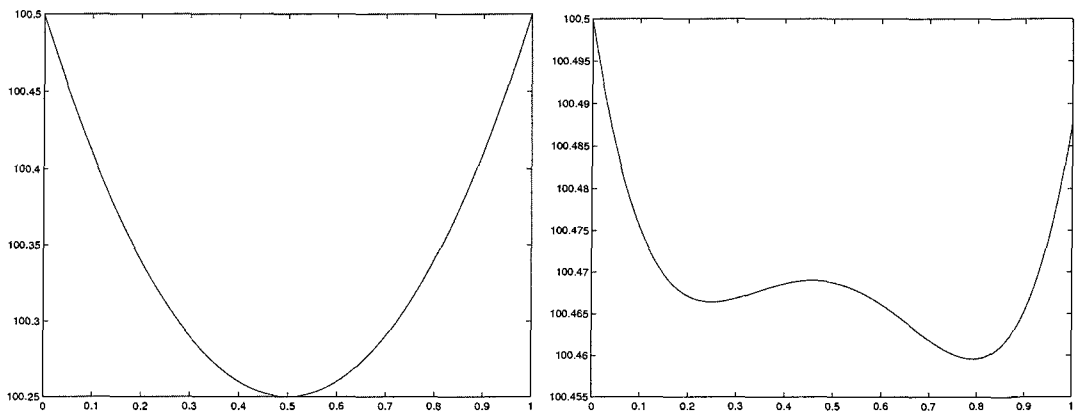

Figure 5.1: Comparison between a single minimum (left) and various local minima (right)

Once the study of the principles of optimisation was conducted, it was possible to proceed to the application of optimisation techniques to the problem at hand. This is described in the next section.

5.1 Mathematical Formulation of the MR Brake Optimisation Problem

The first step in the optimisation process is the definition of the actual problem, i.e., the formulation of the objective function and design and state variables. In the design of the MR brake, the following concerns may be identified:

The torque T must at least meet the minimum value specified in 2.2 (T_{min}). This suggests that a constraint of the form $T \geq T_{min}$ be included in the optimisation problem. While further improvements of the torque beyond the value T_{min} are positive, they do not bring a significant benefit to the system, in that for application to a conventional car T_{min} is an acceptable value. Hence, the objective function may take into account the torque as one of the variables to be optimised but this is not a priority.

Optimisation of the weight, on the other hand, plays a central role. For the MR brake system to be most competitive, it is important to have as low a weight as possible. Another important aspect is that a lower weight represents a lower cost in materials. These considerations lead to the choice of the weight as the most important variable in the objective function. Simultaneously, it is good practice to establish a maximum acceptable value for the weight, as a constraint of the form $W \leq W_{max}$. It is important that excessively restrictive constraints are not imposed or otherwise the design space will be excessively narrowed thereby greatly difficulting the operation of the optimisation procedure. For this reason, and considering that the weight is the major component of the objective function (hence it will necessarily be brought to very low values), a permissive value (65 kg) was specified for the maximum weight. This will allow the optimisation procedure to search a wider design space without being limited by tight feasibility boundaries.

It was chosen to consider only the weight of the steel components (disk(s) and casing) as this is by far the leading contribution and given that the weight of the other components (ring, shaft, coil, fluid) is nearly independent of the dimensions of the system (hence, no optimisation of the weight of these components would be achieved).

Also, it is important to keep in mind that the weight of an MR brake system cannot be directly compared with the on-wheel weight of a hydraulic disc or drum system, given the extra weight required by the hydraulic components that are no longer required in an all-electric system: the master cylinder, vacuum pump, brake lines, fluid and electromechanical control systems (such as the ABS) which in the MR brake can be replaced by software.

Finally, the dimensions must be compatible with the space physically available within the wheel and in particular the diameter must be lower than a pre-determined maximum (related to the diameter of the wheel ring). Therefore, the geometrical dimensions of the MR brake must be taken into account as constraints so that the design produced is a valid one. It was seen in Section 3.5 that the overall diameter of the MR brake is limited by the diameter of the wheel ring. Specifically, it was seen that a clearance of a least 3 mm must exist between the casing of the MR brake and the wheel. Therefore, the following constraint was added to the optimisation problem:

$$width \leq r_{wheel} - 0.003 \tag{5.1}$$

where $width$ is the outer radius of the MR brake casing and r_{wheel} is the inner radius of the wheel (set to 20 cm in Section 3.5).

The considerations presented above suggest the inclusion in the objective function of both the torque and the weight. However, it must be noted that the goal is to increase the torque and decrease the weight. Since in an optimisation problem the

objective function is minimised, dealing with the weight (which is to be minimised) is straightforward. As for the torque, it suffices to remark that maximising the function T (torque) is the same as minimising $-T$. Hence, the following form for the objective function $f(x)$ might be suggested:

$$f(x) = W - T \tag{5.2}$$

where W is the weight of the MR brake. Looking more thoroughly at this expression, one limitation becomes apparent: the objective function has been obtained by combining the weight and torque, however these two variables have very different values - the weight may be around 40 kg while the torque is greater than 1000 Nm for valid designs. The problem in the above expression for the objective function lies in the fact that this difference was not taken into account. An improvement of 1 kg in the weight produces the same change in the objective function as an improvement of 1 Nm in the torque. However, the former represents an improvement of 2.5% (1kg/40kg) whereas the latter represents an improvement of only 0.1% (1Nm/1000Nm). Clearly, the objective function must be changed to take into account the different orders of magnitude of each variable. This is accomplished by dividing each variable by a reference value, hence defining the objective function in terms of a dimensionless weight and a dimensionless torque:

$$f(x) = \frac{W}{W_{ref}} - \frac{T}{T_{ref}} \tag{5.3}$$

where W_{ref} and T_{ref} are the reference values of the weight and of the torque, respectively. The question that arises is that of which values to choose for these reference quantities. This is done by referring to the knowledge gathered from the finite element analyses conducted so far and taking into account the values that each variable

is expected to take. It has been seen that the torque must meet the specified performance target (1010 Nm) and if possible should be improved beyond that value. Hence, it is expected that the torque will take values somewhat greater than the minimum permitted and so a value of 1200 Nm was chosen for the reference torque. In what concerns the weight, it was seen from the finite element analyses conducted in the previous chapter that the 1 disk geometry requires a greater weight than the 2 disk geometry for the same torque. Considering the values observed in Chapter 4, the reference weight was set to 60 kg in the case of 1 disk configurations and 30 kg for 2 disk configurations.

Unfortunately, this expression does not yet meet our desired goal: to improve first and foremost the weight and only secondarily the torque. This shortcoming is due to the fact that the dimensionless weight and torque are considered equally in the objective function. In order to privilege improvements in one of the variables, ponderating factors may be introduced in the expression for $f(x)$:

$$f(x) = k_W \frac{W}{W_{ref}} - k_T \frac{T}{T_{ref}} \tag{5.4}$$

where k_W and k_T are the scalar factors that determine the importance of each variable. As discussed earlier, the major goal of the present optimisation is to reduce the weight. Improvements in the torque beyond the minimum acceptable value, while positive, are of little importance. With this in mind, k_W and k_T were set to 0.9 and 0.1 respectively, meaning that improvements in weight are considered 9 times as important as improvements in the torque or, in more practical terms, a new design is considered better if:

- It improves the weight and this percent improvement is greater than 1/9 of the percent penalty in torque or

- It improves the torque and this percent improvement is greater than 9 times the percent penalty in weight.

Finally, the expression for the objective function takes into account one more consideration imposed by the ANSYS optimisation algorithms [85]: these will exhibit errors if the objective function takes a value close to zero. For this reason, it is recommended to add a finite number of such a magnitude as to ensure that f will always be much greater than zero. Given that in the above expression for f the terms $k_W \frac{W}{W_{ref}}$ and $k_T \frac{T}{T_{ref}}$ will each normally be less than one (and so will the difference between both terms), adding 100 to the objective function guarantees that its value will never approach 0. This leads to the final expression for the objective function:

$$f(x) = 100 + k_W \frac{W}{W_{ref}} - k_T \frac{T}{T_{ref}} \tag{5.5}$$

The values for k_W, W_{ref}, k_T and T_{ref} are summarised in Table 5.1.

Table 5.1: Objective function parameters

Geometry	k_W	W_{ref} [kg]	k_T	T_{ref} [Nm]
1 disk	0.9	60	0.1	1200
2 disks	0.9	30	0.1	1200

When defining the design variables to be taken into account during the optimisation procedure, a trade-off must be reached between the greater control and design possibilities associated with a higher number of variables and the computational expense that too high a number of design variables would imply. Since we are therefore interested in using the lowest possible number of design variables while allowing the greatest possible variety of designs, it is important to refer to the analytical model of the MR brake obtained in Section 3.1 in order to select the appropriate set of

design variables. Recall from Eq. 3.14 that the torque is influenced by the number of disk surfaces, the inner and outer radius of the disks and the magnetic field intensity. The latter, in turn, is influenced by the product of the current (which was fixed to 12 A following the considerations on the transient response of the magnetic field in the previous chapter) by the number of turns (controlled by the width of the coil cross-section *rad_th_coil* and by its height which is a function of the disk thickness) and by the dimensions of the magnetic path (*ax_th_casing* and *rad_th_casing*).

The design variables (using the names presented in Fig. 3.2) and the interval of allowable values for each are shown in Table 5.2. It is seen that large design intervals were chosen so that as many different designs as possible are considered.

Table 5.2: Design space for each variable

Variable	Interval of allowable values [cm]
th_disk	1-5 (1 disk) ; 0.5-2.5 (2 disks)
rad_disk	13.0-18.5
rad_th_coil	0.25-2.5
rad_th_casing	1-5 (1 disk) ; 0.5-2.5 (2 disks)
ax_th_casing	0.25-2.5
length_disk	3-8

It was seen in the previous chapter that, while the CFD model is important to obtain a qualitative analysis of the fluid behaviour, the numerical results for the torque thus obtained are less accurate than those estimated based solely on the magnetic field distribution and the fluid properties (yield stress vs. magnetic field intensity) given that only this latter option can take into account the effect of non-constant magnetic field distributions. For this reason, the evaluation of the torque is done based on the magnetostatics model. Furthermore, this is advantageous in that the solution of the magnetostatics model alone requires only a fraction of the time that would be required to solve both the magnetostatics and CFD models. This plays

a major role in an optimisation procedure, where many finite element analyses are required.

Regarding the choice of a more powerful optimisation method than those supplied with ANSYS, while both genetic algorithms and simulated annealing are methods requiring considerable fine-tuning of parameters to obtain good results without incurring prohibitive computational requirements, this is probably easier in simulated annealing given the great number of selection strategies available in genetic algorithms. Indeed, it has been found that among implementations of both methods to the same problem, simulated annealing produces results only marginally worse than those obtained with genetic algorithms, but at a fraction (one or more orders of magnitude of difference) of the computational time [100].

5.2 ANSYS built-in tools

The Ansys software [84] includes an optimisation processor with two different methods: the subproblem approximation method (described below in Section 5.2.1) and the first order method (described in Section 5.2.2). These are fairly simple methods, capable of minimising relatively regular functions (use of these methods with very irregular objective functions can and probably will result in obtaining a local minimum instead of the global minimum).

Both are unconstrained optimisation methods that can be used to perform the optimisation of constrained problems via the inclusion of a penalty function. That is, whilst the method does not explicitly take into account the constraints, the objective function is modified so that constraint violations will greatly increase it. Hence, the minimisation of the objective function will first result in the elimination of the constraint violations and only then in the minimisation of the desired objective.

Because of this, the objective function used by the optimisation algorithm is not always the one defined by the user, as penalties may have been added for those constraints that are violated (and these will often change over time). It must be stressed that the objective function in the ANSYS results file is that used by the algorithm (i.e., including any penalties, and thus often not the objective function defined by the user). Care must be taken when comparing designs from different iterations.

The ANSYS built-in methods provide a good platform to conduct early optimisation efforts, given that they can be used as a black box requiring no more than the definition of the objective function and of the design and state variables and have been thouroughly tested and validated. This allows the user to focus on the actual problem being studied rather than on the development of the optimisation algorithm. Furthermore, they are not too computer-intensive and yield results quickly, which allows the user to run the optimisation several times with different parameters and thus gain a better understanding of the role each of them plays in the problem under study. Another advantage lies in the fact that these are thorougly tested routines that are useful to benchmark the simulated annealing procedure purpose-built later.

Once a good knowledge of the problem has been obtained from the execution of these built-in methods, more sophisticated and accurate algorithms can be implemented, taking advantage of all the data gathered in these optimisation runs.

5.2.1 Subproblem Approximation Method

The first and simpler of the two algorithms available with ANSYS is subproblem approximation. This is a zero-order method as it conducts the optimisation based solely on the values of the objective function and state variables, without taking into

account their derivatives. After performing function evaluations for different designs, it is possible to obtain a polynomial approximation to the constraints and objective function, as a function of the design variables. For example, the objective function is approximated by $\hat{f}$, expressed as [85]:

$$\hat{f} = a_0 + \sum_{i}^{n} a_i x_i + \sum_{i}^{n} \sum_{j}^{n} b_{ij} x_i x_j \tag{5.6}$$

where x_i is the design variable i $(1 \leq i \leq n)$ and a_0, a_i and b_{ij} are the coefficients of the approximation, obtained by performing a least-squares fit of the function evaluations. The same procedure is used for each constraint. After such approximations have been obtained for the objective function and all the constraints, the algorithm then performs an optimisation of these approximations. The advantage of this approach is that the optimisation procedure can now do a very large number of function evaluations since these only require the computation of the values of polynomial functions of the form of $\hat{f}$ given above. If the method did not construct those polynomial approximations to the objective and constraints, it would require one execution of the finite element model for each design configuration, resulting in a much higher solution time.

5.2.2 First Order Method

The second optimisation technique included with ANSYS is a first order method, requiring not only the objective function and state variables values but also their derivatives relative to the design variables. These derivatives are computed by ANSYS using forward finite differences, in which the gradient of a function f at a reference

point x^r relative to variable x_i is given by the following expression [85]:

$$\frac{\partial f(x^r)}{\partial x_i} = \frac{f(x^r + \Delta x_i e) - f(x^r)}{\Delta x_i} \tag{5.7}$$

where Δx_i is the step size of the forward difference in direction i and e is the unit vector in direction i (i.e., a vector whose ith component is equal to 1 and all others are equal to 0). The default value for Δx_i is 0.5% of the design space of variable x_i (that is, $\Delta x_i = \frac{0.5}{100}(x_{imax} - x_{imin})$.

5.3 Simulated Annealing

Simulated annealing is a very powerful optimisation method whose major drawback is the very high number of function evaluations that are required. It is a random-search method, i.e., the evolution from one design configuration to another is determined randomly. Several random-search methods exist and the difference between them is in the criteria used to compare different designs. The simplest of these is the Monte Carlo method, which consists of a large number of trials with random values of the design variables. If enough trials are conducted, it is expected that the global optimum will be found. This method has important merits such as the ability to find the global optimum, even on very irregular functions when many local minima exist, the ease of implementation (all that is necessary is to generate random values for the design variables and analyse the resulting design, repeating this procedure a large number of times) and a great robustness (the computational implementation is extremely simple and there are no computations involved, hence leaving no margin for failure or numerical errors). The usefulness of the Monte Carlo method however is hampered by the fact that in order to ensure the entire design space is analysed (i.e., all possible values of the design variables are tested) a massive number of function

evaluations is required, thus making it an unpractical option for most engineering problems.

With this in mind, improvements to the basic Monte Carlo procedure have been sought in the past, with the objective of reducing the number of necessary function evaluations while simultaneously retaining all the advantages it presents. One possibility is to replace the purely random nature of the Monte Carlo algorithm with a philosophy of "educated guesses" in which random designs are combined with a study of the problem in hand. This is the reasoning behind the Metropolis method introduced in 1953 [101], which will not only proceed to the new point if this move decreases the objective function but may also accept some moves that actually increase the value of the objective function, subject to an acceptance probability.

The difficulty associated with this algorithm lies in the fact that too high an acceptance probability will result in nearly all moves being accepted, regardless of whether the objective function is decreasing or increasing (and by how much), and too low a value for the acceptance probability will only accept downward moves, thus being equivalent to a simple first order method and not overcoming local minima.

A practical way to use this algorithm successfully in optimisation problems was devised independently by Kirkpatrick *et al* [102] and Černy [103] and simply consists of running the Metropolis algorithm sequentially for decreasing values of the acceptance probability, so that the method will start by searching the entire space for the region with the lowest objective function, and then search for the minimum in that region, accepting ever smaller upward moves and thus converging to the global minimum. The analogy of this procedure with the slow cooling of a material until it reaches its minimum energy solid configuration (referred to in condensed matter physics as "annealing") originated its name: simulated annealing. This procedure is schematically represented by the flowchart in Fig. 5.2.

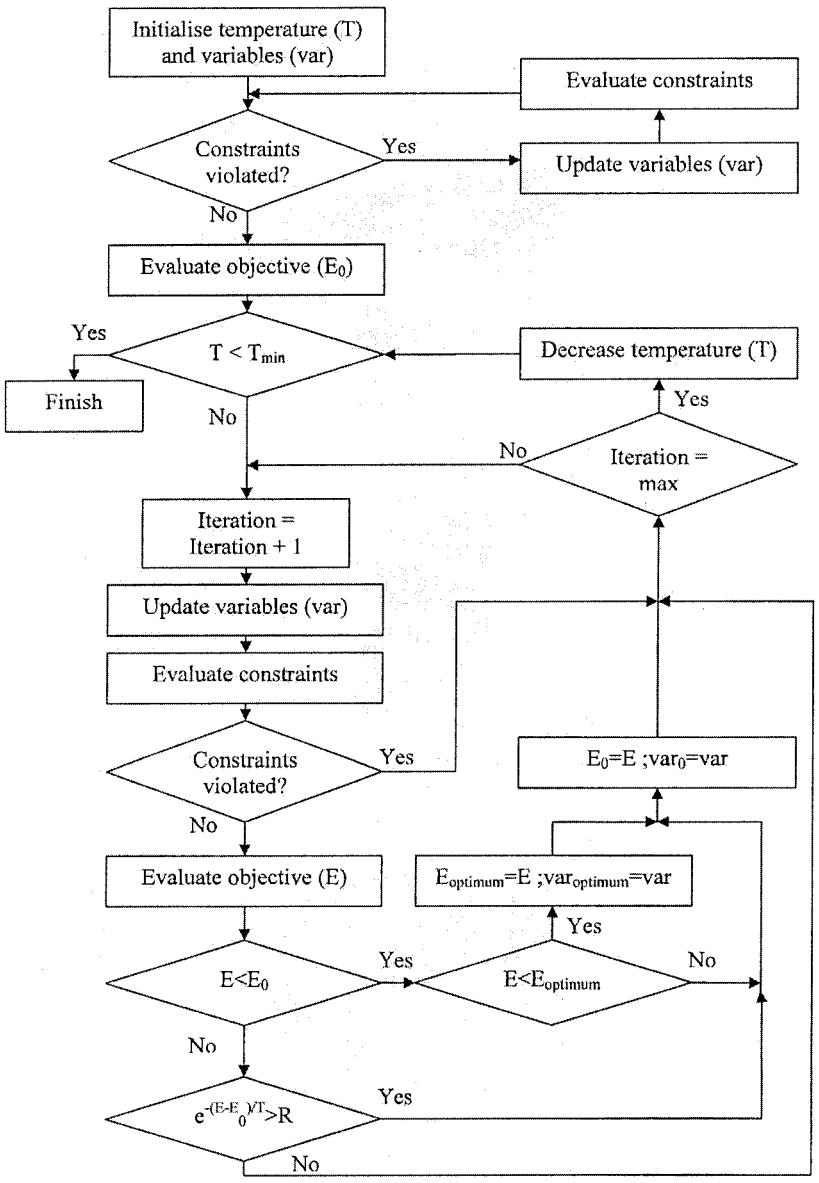

Figure 5.2: Schematic representation of the simulated annealing procedure

In the physical reality, a material is heated up to the liquid phase - at this stage, its molecules can move freely throughout the material and find a distribution that minimises the global energy. However, this is a highly unstable configuration, since the molecules move randomly (being in the liquid state). Suddenly freezing down the material would maintain the given arrangement of molecules, regardless of whether that corresponds to minimum energy. It is therefore necessary to slowly cool down the material, so that the intermolecular connections will become stronger ant the molecules will only move within an ever smaller neighbourhood towards a distribution of lower energy. This concept can be applied to the optimisation of any function simply by replacing energy by a cost or objective function and temperature by a parameter controlling the acceptance probability of a higher "energy" configuration.

At this point, one important question arises: how shall this cooling be conducted, i.e., how does the temperature change from one iteration to another? This is one of several parameters that control the execution of the simulated annealing procedure, along with others such as the number of iterations, the number of random moves per iteration and the maximum step size in the random moves. While a good understanding of the theory behind simulated annealing (v. e.g., [104]) is crucial for the correct choice of values for those parameters, it is not enough - each case is different and the ideal simulated annealing parameter values for each optimisation problem can only be obtained from experimenting with the particular problem at hand. This is why, altough the optimisation problem of the inchworm actuactor (described in Section 2.3.1) served as a good introduction to simulated annealing and provided some insight and experience regarding the role of each of the parameters, some trial and error with the MR brake was required in order to obtain values that ensured a good solution (with the entire design space searched) at an acceptable computational expense.

If the initial temperature of the Simulated Annealing algorithm is too high, all the moves will be successful, meaning that computational time is being wasted without progressing towards a solution. If, on the other hand, it is assigned too low a value, it is possible that too few (or none at all) uphill movements will be accepted and the procedure may get trapped in a local minimum. To overcome the difficulties posed by the choice of the initial temperature, it has been suggested and successfuly implemented (v. e.g., [105]) to set it to a value such that approximately 80% of the moves are accepted. A very simple implementation of this idea was carried out, consisting of two objective evaluations close to the initial point (whose objective function value had already been computed). Simply, starting from the initial point, the algorithm changes the design by adding $(maximum\ step)/2$ to the value of each design variable and computes the objective function for this new point. Again starting from the initial point, a new move is made in the opposite direction, subtracting $(maximum\ step)/2$ to the initial value of each variable and the objective function for this point is also obtained. The average of the differences between each of these objective values and the initial objective is then computed (ΔE) and substituted into the Metropolis Monte Carlo equation:

$$e^{-|\Delta E|/T_0} = 0.8 \tag{5.8}$$

to yield the initial temperature:

$$T_0 = -|\Delta E| / ln(0.8) \tag{5.9}$$

One question comes up when planning the simulated annealing algorithm for a constrained problem: how should constraint violations be addressed? Several possibilities exist:

- After the design is evaluated the algorithm can test it for feasibility and, should the design be unfeasible, simply discard it;
- The classic approach of incorporating penalty functions into the objective function to penalise constraint violations and minimise the modified objective can be used;
- The method can ignore contraint violations and proceed freely throughout the design space regardless of whether each design is unfeasible or not and at the end of each temperature step select the best feasible design.

The last option is ideal for problems with many constraints, where it may happen that to reach the global optimum the method will need to go through unfeasible areas. Since that is not the case in this problem, where only two constraints exist (torque and weight[2]) and they are simple inequalities, either of the former two methods is preferred since they will not waste computation time studying the unfeasible region. Among those two options, the use of penalty functions is a solution adopted when a model for the constraint violations is required. That is the case for deterministic[3] search methods, which move from one design configuration to another based on the knowledge of the objective funtion and constraint violations. In random search methods, on the other hand, such knowledge is not required since the algorithm will update the design randomly and can verify a posteriori whether the new design is a feasible one. Hence, penalty functions are common in deterministic methods but given the extra complexity and uncertainty associated with their use (in the definition of the function to use as well as in the choice between interior and exterior penalties[4])

[2]It will be described in a subsequent paragraph how it was possible to eliminate the width constraint while ensuring that only designs that fit in the wheel would be produced.

[3]i.e., non-random

[4]The former depart from a feasible point and do not allow the design to approach the constraint

not justified in random-search procedures. The first of the 3 candidate techniques to address the constraint violations was therefore chosen, consisting only of a test of the constraints after each function evaluation.

This implementation therefore requires that the initial design be a feasible one, which was not a problem when solving both systems with the 2 disks geometry or the 1 disk system based on the MRF-241ES fluid, for which feasible solutions had been found during the solution by the ANSYS built-in methods. However, for the 1 disk system using the MRF-132AD fluid no feasible solution was obtained by either the subproblem approximation or the first order method and thus a tecnhique to obtain a feasible solution had to be devised. The methodology that was followed consisted of running the simulated annealing procedure to minimise the function "sum of constraint violations". This was defined as:

$$violation = violation_{torque} + violation_{weight} \tag{5.10}$$

with

$$violation_{torque} = max\left(0, \frac{T_{min} - T}{T_{min}}\right) \tag{5.11}$$

$$violation_{weight} = max\left(0, \frac{W - W_{max}}{W_{max}}\right) \tag{5.12}$$

being the relative violation of each constraint, i.e., how much percent below the minimum required torque and how much percent above the maximum permitted weight is the current design. If either constraint is satisfied its violation is 0. Hence, if both are simultaneously satisfied the sum of constraint violations is 0 and a feasible solution has been found. When this happens, the procedure stops and this feasible design can then be used as the starting point for the simulated annealing minimisation

boundaries, whereas the latter minimise the constraint violations whenever the design is outside the feasible region

of the actual objective function.

Given that simulated annealing is a random-search method and therefore requires a very high number of function evaluations to ensure that the entire design space is analysed, reducing the number of design variables would greatly improve the computer requirements of the method and its accuracy by allowing a more thorough study of the remaining variables. Knowledge gathered from the execution of the two ANSYS built-in optimisation methods, together with information from the physics of the problem allows the design variables to be treated so that their number is reduced.

Looking at the optimum design sets produced by the ANSYS built-in algorithms, it was seen that in all cases but one the state variable *width* (the outer radius of the MR brake) is in the vicinity of the maximum allowed value of 20cm and the only design which does not follow this tendency produces poorer results than the other method in the same problem. This seems to indicate that the optimum relation between torque and weight for the MR brake occurs at the maximum width. This observation prompted an analysis of the dependency of the torque and weight on the radius to determine whether indeed the maximum width would always yield the best design. This dependency can be obtained from the equations for the torque and weight as follows:

$$T = 2\pi N \int \tau r^2 dr \propto r^3 \tag{5.13}$$

$$W = mg = V\rho g = r^2 h\rho g \propto r^2 \tag{5.14}$$

This implies that the torque to weight ratio $\frac{T}{W}$ is proportional to r. Hence, the optimum torque to weight is obtained for the highest possible r (the variable *width* in this problem), which is limited by the radius of the wheel rim. It is therefore possible to eliminate one of the design variables by forcing $width = maximum$ instead of

$width \leq maximum$, with $width$ defined as:

$$width = rad\_disk + fl\_gap + rad\_th\_ring + rad\_th\_coil + rad\_th\_casing \quad (5.15)$$

Hence:

$$rad\_disk + fl\_gap + rad\_th\_ring + rad\_th\_coil + rad\_th\_casing = width_{max} \quad (5.16)$$

This leads to the following expression for rad_disk:

$$rad\_disk = width_{max} - (fl\_gap + rad\_th\_ring + rad\_th\_coil + rad\_th\_casing) \quad (5.17)$$

The outer radius of the disk (rad_disk) is now a dependent variable and no longer a design variable. Furthermore, this modification brings about another significant improvement, the elimination of one of the constraints. This is an important contribution in that not only it eliminates the need for the optimisation algorithm to account for one more constraint, it eliminates the possibility of unfeasible designs due to excessive width, as this is now forced to assume an acceptable value.

Another modification can be carried out as a result of the observation that the optimum designs are those in which the flux is the same throughout the steel path around the coil. This implies that $rad_t h_c asing$ and $ax_t h_c asing$ are not independent but rather that there is an ideal relation between them, such that the areas 1 and 2 in Fig. 5.3 are equal.

Equating the expressions for areas 1 and 2 gives:

$$2\pi r_1 h_1 = \pi \left(r_3^2 - r_2^2\right) \quad (5.18)$$

where $r_1 = rad\_disk$, $h_1 = ax\_th\_casing$, $r_3 = width$ and $r_2 = r_3 - rad\_th\_casing$.

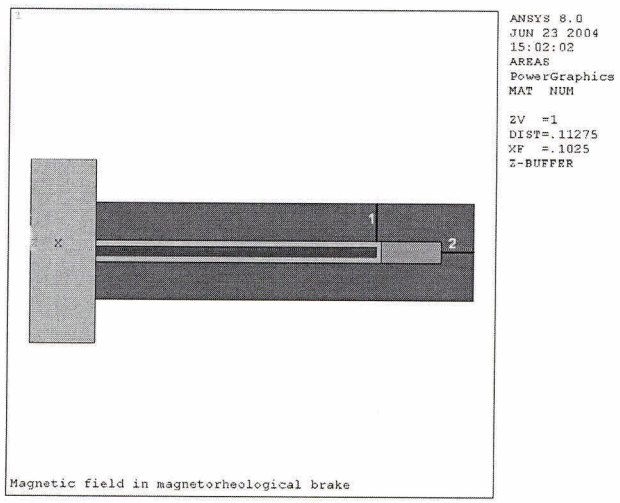

Figure 5.3: Illustration of the different cross-sections along the steel path

Solving for h_1:

$$h_1 \equiv ax\_th\_casing = \frac{(r_3^2 - r_2^2)}{2r_1} \tag{5.19}$$

It is now possible to determine *ax_th_casing* as a function of *rad_th_casing*, keeping only the latter as a design variable.

The definition of the maximum step size permitted for each variable (i.e., how much the value of that variable can change from one design to the next) must ensure that the whole design space is analysed, which requires a balance between the need to search the full interval of design values for each variable from the lower bound to the upper bound, while giving a good resolution (i.e., analysing values of the design variable close to each other). The chosen values are presented in Table 5.3, together with the other parameters controlling the simulated annealing procedure.

Table 5.3: Simulated annealing procedure parameters

Initial temperature	Such that 80% of the moves are accepted
Cooling schedule	$T_i = 0.78 \times T_{i-1}$ [5]
Number of temperatures	8 (1 disk) ; 12 (2 disks)
Number of moves per temperature	1200 (1disk) ; 1000 (2 disks)
Step size for each variable:	
th_disk	0.2 cm
rad_th_coil	0.225 cm
rad_th_casing	0.2 cm
length_disk	0.25 cm

5.4 Results

5.4.1 Overview

In order to compare the results of the different methods, it is important to ensure that the objective functions are defined in the same manner for all of them. It has been pointed out in Section 5.2 that the objective function values written by ANSYS to the results file may differ from that specified by the user and even vary from one iteration to another due to penalties for constraint violations. Therefore, the first step in the analysis of the results is to have an algorithm that reads the state variables (braking torque and weight) for all iterations from the ANSYS results file and then computes the objective function solely based on those values. This will provide a meaningful comparison between designs obtained with different methods and at different iterations as the same formula was used in the calculation of all objective function values.

Furthermore, it has been noted earlier that in the definition of the objective function both the torque and weight were divided by reference values to convert them to dimensionless parameters. In particular, the weight was divided by a reference value

of 60kg for the single disk geometry and a reference value of 30kg for the double disk geometry. Since it is useful to have a comparison between the results obtained with the various geometries, for the purpose of Table 5.4, a new objective function was defined with the same reference weight (30kg) used for all the geometries. Please note that all objective function values other than in Table 5.4 were obtained with the expressions defined earlier and therefore are not meant to form a basis for comparison between the different designs.

Table 5.4: Optimisation results: objective function comparison

Configuration	Subproblem Approximation	First Order Method	Simulated Annealing
1 disk, MRF-132	-[a]	-[a]	-[a]
1 disk, MRF-241	-[a]	101.7301[b]	101.7311[b]
2 disks, MRF-132	100.8975	100.9598	100.7905
2 disks, MRF-241	100.7144	100.5946	100.4547

[a]No feasible design found

[b]After the objective function was modified in order to allow a comparison between different geometries (as described in the text)

It is seen from Table 5.4 that none of the methods found a solution that met the specified targets of torque and weight for the 1 disk geometry using MRF-132 fluid. For the 1 disk geometry based on MRF-241 fluid, subproblem approximation was unable to reach a feasible solution, unlike the other two methods. It is also seen that this is the only configuration for which simulated annealing was outperformed (albeit only marginally). For the configurations with 2 disks and stator, simulated annealing always obtained the best results. Altough the exact values of torque and weight associated with the optimum design found by each method for each of the configurations will be presented in detail later, it is clear that the 2 disks geometry have much superior performance - the improvement in the torque and weight components

of the objective function was over 50%[10].

Another important factor in optimisation problems is the solution time. Approximate values for the solution time of each method are shown in Table 5.5. It is seen that the better performance of simulated annealing comes at a high cost: its solution time is 25 times greater than that of the first order method. This, in turn, requires 12 times more time than subproblem approximation. Given the good results (v. Table 5.4) obtained with the latter method whenever it was able to converge, this is seen to be a good introductory optimisation method.

Table 5.5: Approximate solution time of each method

Method	Solution Time
Subproblem approximation	20 minutes
First order method	4 hours
Simulated annealing	100 hours

The following sections discuss in greater detail the results obtained for each configuration of the MR brake.

5.4.2 1 Disk Geometry, MRF-132 Fluid

This was the most problematic of the 4 configurations studied, in that no feasible design was found. Nevertheless, the best design[11] obtained with each of the optimisation methods for this configuration is presented in Table 5.6.

[10]Recall that the objective function comprises a component due to the torque, one due to the weight and one constant (100) introduced solely for numerical reasons. The actual improvement in torque and weight may be assessed simply by subtracting this physically meaningless component (100) from the objective function value.

[11]When no feasible designs are found, the best one is considered to be that closest to the feasible region, i.e., that has the lowest constraint violations.

Table 5.6: Optimum values for the 1 disk geometry using MRF-132AD fluid

	Subproblem Approximation	First Order Method	Simulated Annealing
Objective function	[a]	[a]	[a]
Torque [N]	441.07	770.74	748.18
Weight [kg]	26.035	74.437	64.818
th_disk [cm]	1.7927	3.4773	3.3206
rad_disk [cm]	18.001	[b]	[b]
rad_th_coil [cm]	0.79981	0.63105	0.69970
rad_th_casing [cm]	1.0004	2.5266	2.2982
ax_th_casing [cm]	1.7000	[b]	[b]
length_disk [cm]	4.4984	8.9139	7.9970

[a]No feasible designs found. The other values in this column are for the design with lowest constraint violations.

[b]Not a design variable in this solution. Refer to the explanation in 5.3

Analysis of the results presented in Table 5.6 show considerable differences between the various methods. It is seen that subproblem approximation led to a design with a very good weight value (well within the feasible limit of 65 kg) but a torque far away from the required minimum (1010 Nm). The first order method, on the other hand, divided the constraint violations by both quantities, leading to a design that does not meet either constraint but in which none of the quantities is as away from the acceptable values as in subproblem approximation. Finally, simulated annealing seems to have performed the best of the 3 methods, in that its weight satisfies the imposed limit while the torque is only slightly worse than that of the first order method.

5.4.3 1 Disk Geometry, MRF-241 Fluid

The value of the objective function at every iteration for each of the methods is presented in Fig. 5.4. Note that only the objective function values corresponding to

feasible designs are shown. The absence of values for the subproblem approximation method indicates that no feasible solution was found to this configuration. Also note that 1 iteration is not the same as 1 function evaluation: simulated annealing has the least iterations (temperature steps) but a massive number of function evaluations per iteration. This explains that altough requiring the lowest number of iterations of the 3 methods, simulated annealing actually takes the longest time to obtain the solution (v. Table 5.5). As a curiosity, notice that the first order method only found a feasible design after a number of iterations greater than that required by the simulated annealing procedure for the entire analysis.

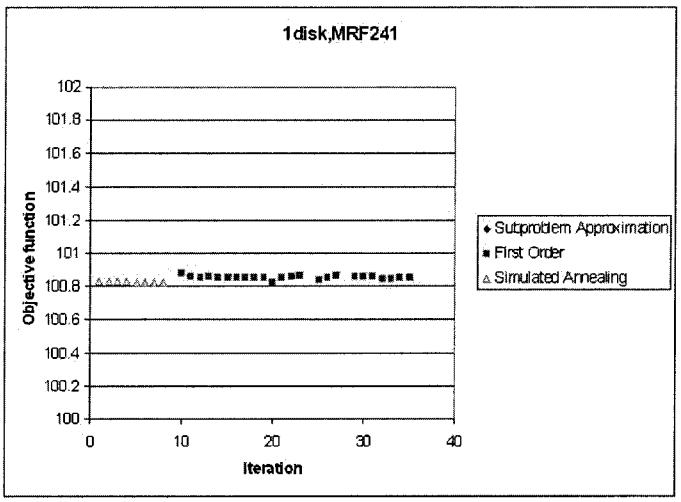

Figure 5.4: Evolution of the objective function for each optimisation method in the 1 disk geometry using MRF-241 fluid

The optimum values of the objective function, torque, weight and design variables obtained by each method are shown in Table 5.7.

For this configuration, only the first order method and simulated annealing obtained feasible designs. It must however be stressed that the results of the various methods are not too different. The subproblem approximation method was close to

Table 5.7: Optimum values for the 1 disk geometry using MRF-241ES fluid

	Subproblem Approximation	First Order Method	Simulated Annealing
Objective function	[a]	100.82	100.82
Torque [N]	1002.0	1010.3	1011.1
Weight [kg]	64.759	60.476	60.513
th_disk [cm]	4.5532	2.9910	2.8872
rad_disk [cm]	[b]	[b]	[b]
rad_th_coil [cm]	0.47442	0.62216	0.7705
rad_th_casing [cm]	2.4295	2.5495	2.5211
ax_th_casing [cm]	[b]	[b]	[b]
length_disk [cm]	5.7039	6.1859	6.3394

[a]No feasible designs found. The other values in this column are for the design with lowest constraint violations.

[b]Not a design variable in this solution. Refer to the explanation in 5.3

the feasible region (The torque is 8 Nm away from the minimum) whereas the 2 other methods do little better than barely meeting the acceptable targets. It is also seen that the designs obtained by the first order method and simulated annealing are identical - all variables have close values, which naturally leads to similar values for the torque and weight in the designs of both methods.

5.4.4 2 Disks Geometry, MRF-132 Fluid

The evolution of the objective function throughout the optimisation procedure, according to the various methods, is presented in Fig. 5.5. An immediate conclusion that can be taken is that for the present problem subproblem approximation outperformed the more sophisticated first order method, achieving a better design (lower value of the objective function) in fewer iterations. The fact that the first order objective function hovered around a constant value for so long a period (from iteration 21 until the end) instead of progressively decreasing towards the optimum (as would

be expected from a gradient-based method) suggests that it may have got trapped in a local minimum[22].

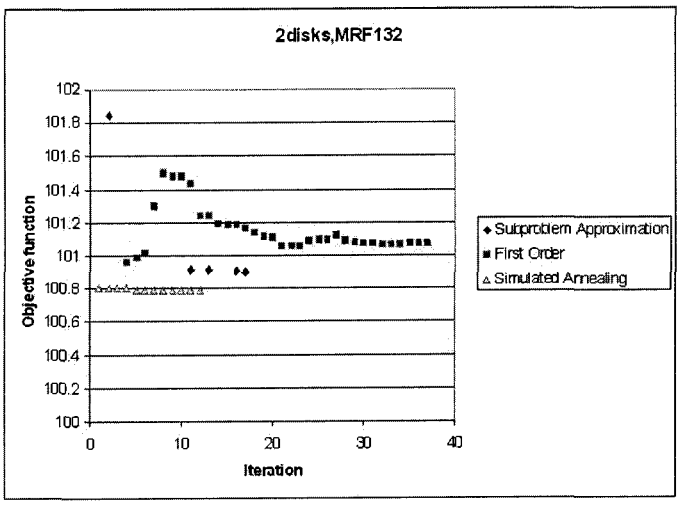

Figure 5.5: Evolution of the objective function for each optimisation method in the 2 disks geometry using MRF-132 fluid

The characteristics of the optimum designs obtained by the different methods are presented in Table 5.8. It is seen that indeed the first order method is quite off the global optimum reached by the other two methods with a disk thickness approximately 33% greater than that obtained by the other methods (and hence the higher weight).

5.4.5 2 Disks Geometry, MRF-241 Fluid

Fig. 5.6 presents the evolution of the objective function for the 2 disks configuration using MRF-241 fluid. Once more, simulated annealing reaches the best results of all three methods, now with the first order method outperforming the subproblem

[22]Note that the first order method is more likely to encounter a local minimum than the subproblem approximation method [84].

Table 5.8: Optimum values for the 2 disks geometry using MRF-132AD fluid

	Subproblem Approximation	First Order Method	Simulated Annealing
Objective function	100.90	100.94	100.79
Torque [N]	1016.6	1071.8	1012.7
Weight [kg]	32.741	34.172	29.162
th_disk [cm]	0.50895	0.66108	0.51130
rad_disk [cm]	16.778	17.185	[a]
rad_th_coil [cm]	0.97289	0.80453	1.1601
rad_th_casing [cm]	1.4982	1.5248	1.4413
ax_th_casing [cm]	1.7422	1.7050	[a]
length_disk [cm]	5.9936	5.6474	4.9575

[a]Not a design variable in this solution. Refer to the explanation in 5.3

approximation method. It is also seen that the qualitative evolution of the solution of each method is similar to that observed in the previous section: simulated annealing is able to find a very good design on the first iteration and thereafter progresses slowly, whereas the subproblem approximation method first obtains very poor values and is capable of achieving great improvements over a period of several iterations. And again, the first order method seems to progress towards a worse solution and only later converges to its optimum. The optimum values obtained by each method are in Table 5.9.

Given that, among the proposed configurations (1 and 2 disks, and both types of MR fluid), this is the one with more disk surfaces and using the fluid with a higher yield stress, it was expected to produce the best results. Indeed, that is the case. Comparing these values with those based on the same geometry using MRF-132 (v. Table 5.8), it is seen that in addition to a slightly improved braking torque (1025.4 Nm vs. 1012.7 Nm) a weight gain of over 38% (18.005 kg vs. 29.162 kg) has been achieved. This means that, as expected, this is the most promising configuration. Unfortunately, its suitability for application to car brakes is hampered

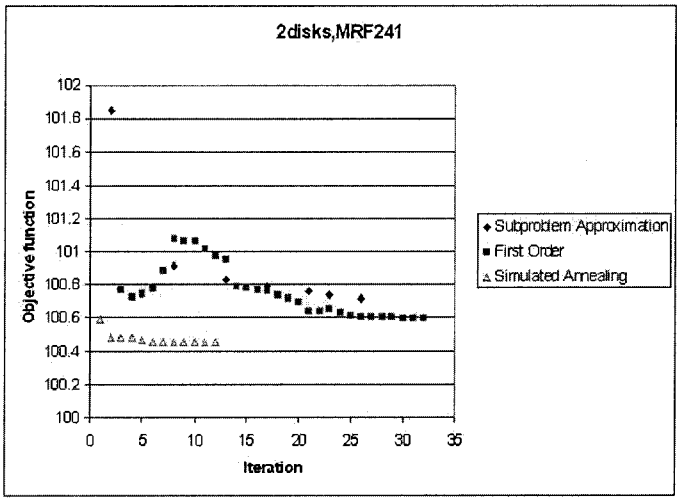

Figure 5.6: Evolution of the objective function for each optimisation method in the 2 disks geometry using MRF-241 fluid

Table 5.9: Optimum values for the 2 disks geometry using MRF-241ES fluid

	Subproblem Approximation	First Order Method	Simulated Annealing
Objective function	100.71	100.59	100.45
Torque [N]	1010.6	1040.6	1025.4
Weight [kg]	26.620	22.712	18.005
th_disk [cm]	1.5172	0.66309	0.50000
rad_disk [cm]	16.996	15.864	[a]
rad_th_coil [cm]	0.69963	1.1685	1.0560
rad_th_casing [cm]	1.6715	1.5024	1.1325
ax_th_casing [cm]	1.2101	1.6091	[a]
length_disk [cm]	2.2540	2.8038	3.0000

[a]Not a design variable in this solution. Refer to the explanation in 5.3

by the limitations of the MRF-241 fluid: lower operating temperature limit (given that it is water-based instead of oil-based) and poor compatibility with dynamic seals, which will be addressed in greater detail in Section 6.1.2.

Chapter 6

Implementation

6.1 Accessory components

So far, the design procedure has addressed all major components of the MR brake: the rotating shaft and disks, the static casing, the coil that induces the magnetic field and the aluminium ring between the fluid and the coil. In order to actually bring together all these parts, some auxiliary components must be included in the design. For instance, the combination of rotating and static components requires bearings to allow relative motion, and seals must be fitted to keep the MR fluid within the casing.

6.1.1 Bearings

Bearings are used to connect parts that must be allowed to rotate relative to each other, e.g., to support a rotating shaft on a static structure. Desired properties of the bearings are:

- Ability to transmit high loads;

- Low friction;
- Resistance to high temperatures and high velocities.

Bearings are necessarily used on cars to connect the wheel axles to the chassis, allowing the former to rotate freely. Hence, in traditional disk brakes, the disk is attached to the axis and the calliper (which must not rotate) is attached to the chassis. On the MR brake, two different possibilities exist:

- The current configuration can be maintained, with the casing attached to the chassis and hence static;
- Alternatively, bearings may be fitted within the MR brake (between the shaft and casing) with the chassis in turn fixed to the casing. This means that the weight of the car would be supported on the casings of all the brakes, hence requiring careful dimensioning of their thickness. This option would also increase the complexity of the brakes by adding more parts (and complicating the assembly process given the precise requirements of bearings) but might bring weight and space savings by effectively using the MR brake unit to connect the axis to the chassis instead of dedicated structures.

This suggests that to take full advantage of the geometry of MR brakes to realise an extra weight and space gain, the bearings may be incorporated in them. However, given that the goal of the present project is to devise a simple system capable of highlighting the principle of operation and the performance capabilities of MR brakes, ease of construction considerations dictated the option not to include bearings and to maintain the configuration of hydraulic disk brakes in this aspect.

6.1.2 Seals

It was seen in the previous section that bearings may be used to join parts while still allowing them to rotate independently. Because they may face high velocities and temperatures, lubrication is generally required. This originates one problem: how to retain the lubricating agent within the bearing unit? The solution to this problem is the use of seals that form a barrier around the bearings, efectively keeping the lubricant enclosed. In the particular case of the MR brake there is another, even more critical, requirement for seals: the necessity to maintain the MR fluid within the brake system, preventing any leaks to the exterior (as this would render the brake inoperative). The goal of the seal is therefore to prevent any leak of fluid, whilst simultaneously keeping friction to a minimum and being able to withstand the velocities and temperatures present in the MR brake. Many designs for seals have deen developped in the past, each with particular advantages and disadvantages. It is therefore necessary to choose the design most adequate to the demands of the current system: capacity to withstand moderate velocities and temperatures, excellent fluid retention[1], very good longevity. In addition to the great variety of available seal designs, many different materials are also used. The choice is once again dictated by the characteristics of the application but also by compatibility limitations of the MR fluids. For example, some MR fluids are not compatible with rubber, a material commonly used in seals, and hence particular care must be taken in the choice of the seal material, referring to the material compatibility indications supplied by the fluid manufacturer [106, 107]. It is also important to note that MRF-241ES is not recommended for long-term use with dynamic seals [106], which would require the

[1]Depending on the type of fluid (density, viscosity), different designs may be suitable. MRF-132AD is an oil-based MR fluid and hence seals for oil retention are the most adequate.

use of different sealing configurations, such as labyrinth seals[2]. However, the greater complexity and lower availability of such sealing systems dictates the option to use conventional dynamic seals in the present problem, thus leading to the selection of MRF-132AD.

Another aspect to take into account is the allowable deviations in the dimensions of the shaft and housing bore. This concern stems from the critical balance that the seal must strike: it must be tight enough that no fluid is allowed to leak but cannot press too hard against the surrounding surfaces (which would result in unacceptable friction). Therefore, seals are designed to be applied to parts of specific dimensions. Tables 6.1 and 6.2 list the tolerances that must be met in the manufacturing of the shaft and the housing bore, respectively [108].

These are relatively tight tolerances but nevertheless they can be obtained from most well-equipped machining facilities.

6.2 Manufacturing

One of the tasks of the engineer working on the design of new components is the definition of the manufacturing processes to be used in the creation of the actual parts. Study of manufacturing technologies is also very important to ensure that all components can be made as specified and within acceptable cost and duration. The goal of the present section is therefore to bridge the gap between the on-paper design of the MR brake and the process of manufacturing the actual unit.

[2]These are contactless configurations, which instead of having an element of rubber or other material forming a solid barrier in the gap between the shaft and the casing have a long intrincate path (the "labyrinth") of narrow gaps between the shaft and casing, with a form carefully designed so that the fluid is not allowed to escape through these gaps.

Table 6.1: Shaft tolerances for metric seals

Shaft diameter Nominal [mm]		Diameter tolerance Deviation [μ m]		Circularity tolerance Deviation [μ m]
over	incl.	high	low	max
6	10	0	-90	22
10	18	0	-110	27
18	30	0	-130	33
30	50	0	-160	39
50	80	0	-190	46
80	120	0	-220	54
120	180	0	-250	63
180	250	0	-290	72
250	315	0	-320	81
315	400	0	-360	89
400	500	0	-400	97
500	630	0	-440	110
630	800	0	-500	125
800	1000	0	-560	140
1000	1250	0	-660	165
1250	1600	0	-780	195
1600	2000	0	-920	230
2000	2500	0	-1100	280
2500	3150	0	-1350	330
3150	4000	0	-1650	410
4000	5000	0	-2000	500

Source: [108]

Table 6.2: Housing bore tolerances for metric seals

Nominal diameter [mm]		Diameter tolerance Deviation [μ m]		Chamfer [μ m]
over	incl.	high	low	max
10	18	+27	0	0,3
18	30	+33	0	0,3
30	50	+39	0	0,4
50	80	+46	0	0,4
80	120	+54	0	0,8
120	180	+63	0	0,8
180	250	+72	0	0,8
250	315	+81	0	0,8
315	400	+89	0	0,8
400	500	+97	0	0,8
500	630	+110	0	0,8
630	800	+125	0	0,8
800	1000	+140	0	0,8
1000	1250	+165	0	0,8
1250	1600	+195	0	0,8
1600	2000	+230	0	0,8
2000	2500	+280	0	0,8
2500	3150	+330	0	0,8
3150	4000	+410	0	0,8
4000	5000	+500	0	0,8

Source: [108]

The most pressing question is that of how to create the various parts (disks, shaft, casing) that form the MR brake. Each manufacturing process has its own strengths and weaknesses thus making it more suitable for different situations. For example, in the present problem it must noted that the construction of the prototype (one unit only, for experimental and demonstration purposes) does not possess the same characteristics as the intended series production of units for commercialisation. The latter is particularly adequate for foundry, a method that presents lowest costs for high production batches but is unpractical for the construction of just one unit, given the high costs involved in the construction of the moulds. Hence, it is expected that the components of the prototype will be obtained by cutting standard steel plates in order to obtain the desired shape. This method also allows better finishing quality[3] but the waste of material (all that is cut from the steel plates until the desired shape is obtained) is considerable - thus rendering it unreasonable for the manufacture of large batches.

Another aspect to consider is the ease of manufacturing associated with the selected configuration: components involving very small tolerances should be avoided as this greatly increases the production cost and time. It has been seen above that the seals require relatively tight tolerances, but that cannot be avoided given that they are fundamental components and this is an intrinsic limitation. The bearings also require precise finishing of the parts they will be in contact with and this is a variable to take into account when deciding whether or not to include the bearings in the MR brake.

[3]Parts obtained by foundry may present small defects such as air bubbles and cavities, in addition to a less perfect surface finish [109]

6.2.1 Assembly

Once the various parts have been manufactured, it is necessary to assemble them in order to obtain the MR brake. Also, it is important to wind the coil wire using the appropriate tension. This can be obtained from the wire manufacturer or from tables of common coil data, such as [73].

Another important step in the assembly process is to ensure that the disks are properly balanced, i.e., their centre of gravity is perfectly aligned with the axis of rotation. This is the same procedure that is carried out on the car wheels, whenever the tyres are fitted and ensures that the rotating part (in our case the disk) will not vibrate excessively when rotating at high speeds.

6.3 Control

Until recently, car brakes were essentially open-loop systems: the pressure of the brake fluid (and hence the braking power) were proportional to the travel of the brake pedal. The only exception was the pressure limiting device to avoid premature locking of the rear wheels. Over the last few years, however, sophisticated control systems have been added to improve the braking performance and the overall stability of the car. These determine the ideal braking power on each wheel at every moment, based of information given by a variety of sensors (car speed and acceleration, wheel slip, etc.) in order to provide maximum braking capabilities and optimum handling.

One of the advantages of the magnetorheological brake system is the existance of wider control possibilities, given that the braking power is determined only by the coil current. For this reason, advanced control systems can be designed for MR brakes. However, in order to achieve optimum performance, these control systems should take

into account the dynamics of MR brakes, rather than just replicating those found on hydraulic systems. Hence, the equations derived in 3.1 should form the model of the MR brake to be used by the MR brake controller to be developped in the future.

At this point, it is expected that the automatic control system for the MR brake will be implemented using dSpace. This is a very powerful suite for control applications, used in many different areas. One limitation associated with it is the low current supplied by most dSpace hardware versions, in the order of 50mA. Given that the MR brake requires a current of up to 12A, the work to be conducted in the future on the control area must include the definition of the amplifier to use.

6.4 Commercialisation

The car manufacturing industry is quite conservative in nature, but the public is even more so. This means that new technologies generally face an important challenge in overcoming scepticism around their advantages. In a safety-critical component, such as the brakes, this scepticism is likely to be much more pronounced. It is interesting to note that the public presentations[4] of the MR brake project have aroused much interest in the audiences but even amidst this enthusiasm the questions regarding the safety of this concept were a constant. This mirrors well the scepticism of the public regarding the application of a new technology to a safety-critical component.

Therefore, the study of the reliability and safety of an MR brake system is an important part of its design. The first step (presented in Section 6.4.1) involves the study of the effects of time and wear on the MR fluid's properties, after which an analysis of the reliability of the brake is performed in Section 6.4.2. The reliability

[4]at the ASI Exchange in Vancouver in March and at the Mechanical Engineering Graduate Research Colloquium in Victoria in May

(and safety) of the MR brake system depends not only on the deterioration of the fluid's properties but also on the performance of all the parts that make up the design.

6.4.1 MR Fluid Stability and Longevity

One factor of the utmost importance regarding the actual implementation of magnetorheological devices is the longevity and stability of the MR fluid. Until recently, these issues have been major obstacles to real-world applications of MR fluids. In early MR compounds, the ferrous particles tended to detach from the carrier fluid thus producing a non-homogeneous material without the desired properties. The use of additives has greatly overcome this problem and presently the iron particles in MR fluids will only separate from the fluid after long periods of inactivity (several months)[5] and, most importantly, moderate mixing is enough to return the MR fluid to its original properties.

This means that after a car equipped with MR brakes has been stopped for several weeks the iron particles may have separated but will remix with the carrier fluid after the car is in motion and the brakes have been applied. This means that the driver must be particularly careful in this situation as the brakes response in the first few brake uses may not be optimal. It must be noted that this situation presents absolutely no difference from the behaviour of today's hydraulic brakes, whose components (particularly the calliper) may be sluggish and react poorly (or not at all) after the car has been stopped for several weeks. The important safety advice for MR brakes after prolonged immobilisation is therefore exactly the same as for existing brake systems: particular care must be taken in the first minutes of driving.

[5] after 6 months storage at 25°C, a clear layer of 4 to 20% (by volume) forms on a 1 litre container, depending on the particular fluid [106]

When the stability problems were solved, prototypes of applications based on MR fluid for commercial application were built and tested and over the course of life-cycle experiments, serious degradation of the fluid after repeated use was noticed. In an early version of Lord Corporation's MotionMaster^{TM} RD-1005 truck seat damper, the off-state[6] force (which would ideally remain constant at the initial low value over the life of the product) increased by a factor of 2.5 after 600,000 on-state cycles, rendering it uneffective [110]. This increase in the off-state force is due to an increase in the off-state viscosity of the fluid, a phenomenon known as "In-Use Thickening" (IUT) and believed to be caused by spalling of the surface layer of the particles in the MR compound. After two years of fluid development, this problem too was overcome and presently MR fluids can endure long on-state periods without deterioration. The exact MR fluid life depends on a variety of factors, including the shear rate, temperature and duration. An estimate of the usable life of an MR fluid can be obtained from the following relation [110]:

$$LDE = \frac{1}{V}\int_0^{life} P dt \tag{6.1}$$

where V is the fluid volume (since LDE is generally given in J/cm^3, this should be expressed in cm^3) and P is the dissipated power. Hence, LDE is the energy dissipated per unit volume of the fluid throughout its operation and can be compared to the permissible value obtained from experimental data. According to [110], today's MR fluids can sustain approximately 10^7J/cm^3 before excessive thickening occurs.

If the braking power is assumed constant throughout the life of the brake, this can be rewritten as:

$$LDE = \frac{P}{V} life \tag{6.2}$$

As an example, this would mean that a brake system with 300 cm^3 of MR fluid would

[6] i.e., with no magnetic field applied

be capable of dissipating a constant 7.5 kW for:

$$life = \frac{LDE \cdot V}{P} = \frac{10^7 \cdot 180}{7500} = 240000\,\text{s} \tag{6.3}$$

However, such a number for the MR fluid life is of little interest in reality as it assumed a constant value for the braking power. In order to determine the actual life of the fluid, the actual use of the brake must be taken into account. This can be done by referring to a typical drive cycle, such as the FTP 75 and US 06 cycles introduced in Section 2.2.1.

To this end, a procedure was implemented to calculate the dissipated energy over the braking periods of the driving cycles. This starts by computing the deceleration a at each time-step t as a function of the velocity v:

$$a(t) = min\left(\frac{v(t+\Delta t) - v(t)}{\Delta t}; 0\right) \tag{6.4}$$

The *min* function ensures that for those time-steps in which the car is accelerating or cruising (i.e., $v(t+\Delta t) \geq v(t)$) the deceleration is set to 0. Once in possession of the deceleration at each time step, the associated braking torque can be computed. As it was seen in Section 2.2 the dependency of the torque on the deceleration is linear. Hence, a simple proportion can be used:

$$T = \frac{a}{a_{max}} T_{max} \tag{6.5}$$

where T_{max} is the maximum torque that produces the maximum deceleration a_{max}. Note that using this simple relation neglects the effect of different velocities on the residual sources of drag (aerodynamic, powertrain, and tyre friction). This is a conservative approach in that for higher speeds (those at which a greater amount of energy

is dissipated) these residual sources of drag are being underestimated and therefore the contribution from the braking torque is overestimated.

Recall that the power dissipated by the brakes is given by:

$$P = T\omega \tag{6.6}$$

Therefore, ω must be determined. Recalling that it is given by:

$$\omega = \frac{v}{r} \tag{6.7}$$

where r is the radius of the wheel+tyre assembly, the angular speed at each time-step is given by:

$$\omega(t) = \frac{\frac{v(t)+v(t+\Delta t)}{2}}{r} \tag{6.8}$$

Note that the function of the numerator is simply to compute the average speed over the time-step in question.

The dissipated power can now be obtained from the expression above, and from this the energy dissipated over one complete cycle can be calculated:

$$E_{cycle} = \int_{t_0}^{t_{max}} P \tag{6.9}$$

where t_0 and t_max are the start and end time of the driving cycle, respectively.

Recall that:

$$LDE = \frac{1}{V}\int_0^{life} Pdt = \frac{1}{V}E \tag{6.10}$$

and the energy dissipated throughout the life of the fluid can be written as:

$$E = n_{cycles}E_{cycle} \tag{6.11}$$

That is, the total energy dissipated by the fluid equals the amount of energy dissipated in each cycle (E_{cycle}) multiplied by the number of cycles (n_{cycles}). Substituting in the previous expression yields:

$$LDE = \frac{1}{V} n_{cycles} E_{cycle} \tag{6.12}$$

which can be rewritten for the number of cycles:

$$n_{cycles} = \frac{LDE \times V}{E_{cycle}} \tag{6.13}$$

Since LDE is a known characteristic of the MR fluid, V is known from the dimensions of the brake and E_{cycle} has been computed above, this expression allows the determination of the number of cycles the fluid will be able to withstand before important In-Use Thickening occurs.

Based on the number of cycles and the length of the cycle (L_{cycle}), the total distance d that the car may travel over the useable life of the MR fluid is simply given by:

$$d = n_{cycles} L_{cycle} \tag{6.14}$$

The longevity results obtained using these methodology for the FTP 75 and US 06 cycles are presented in Table 6.3.

Table 6.3: Longevity results for the MR fluid

	FTP 75	US 06
Energy dissipated throughout one cycle [J]	1,142,800	1,002,635
# of cycles the fluid can withstand	1575.079	1795.269
Total distance travelled during MR fluid life [km]	27989.15	22979.44

This is a satisfactory longevity for the MR fluid (comparable to that of brake pads in today's systems) and it can be expected that this figure may greatly increase

with further development of MR fluids. Comparing the values for the two cycles, it comes as no surprise that the life of the fluid is shorter when the driving style is represented by the US 06 cycle - this cycle models aggressive driving with more pronounced accelerations and decelerations. Also, it is interesting to note that as MR fluids are improved their life may no longer be affected by uses in which the power dissipated is small. That is, one can speculate about the possibility of MR fluid behaviour becoming similar to the phenomenon of fatigue in ferrous materials, in which as long as the stresses are below a certain threshold, the life is virtually infinite.

The actual lifespan of the material is not the only important consideration. Because the demands the fluid is subject to will vary from one car to another and over time depending on the driving pattern, it is important to have a way to monitor the fluid condition. Since In-Use Thickening manifests itself through an increase in the off-state viscosity, the simplest way to determine the fluid's performance is to measure the off-state braking torque (ie, the drag caused by the MR brake system without the brakes applied) - once this exceeds a pre-defined threshold, a warning light in the instrument panel can advise the driver of the need to replace the fluid soon. Alternatively, this measurement can be done in a few seconds during scheduled maintenance of the car with equipment that already exists in most garages. It is also important to note that, unlike conventional brakes where extreme wear of the brake pads will result in reduced braking torque, the In-Use Thickening phenomenon of MR brakes is not a safety critical problem given that its only consequence is the increase in the off-state braking torque and thus a greater fuel consumption.

Another noteworthy aspect is the ease of replacement once the operating life has been reached. On conventional hydraulic brakes, the calliper must be removed from around the disk, the pads are then replaced, and finally the calliper is once again

placed around the cylinder while keeping it open with a wrench or similar tool. On most cars, this operation has to be conducted without the wheel in place so that the calliper can be removed. In total, replacing the 4 sets of brake pads requires over 90 minutes of work. It may also be necessary to replace the brake fluid which requires emptying the existing fluid from the hydraulic circuit, adding the new fluid and then bleeding the air that entered the circuit. This requires approximately 30 more minutes. On the MR brake system, on the other hand, only the MR fluid in each wheel has to be replaced. Since the brake at each wheel has its own fluid and there is no hydraulic system around the car with points at different heights, there is no need to bleed the system. The replacement procedure is much simpler in that it consists only of opening a valve in the brake casing, draining the old fluid and then inserting the new fluid, once for each wheel. Furthermore, it may be possible to conduct this operation without removing the wheel.

However, the behaviour of the MR fluids under the very high shear rates expected in the proposed brake configuration remains to be seen. Whether these high shear rates will adversely influence the longevity of the fluid must be assessed from the testing to be carried out. Finally, the fact that the problem of In-Use Thickening was only discovered during life-cycle experiments underlines the great importance of such tests, which are the only way to uncover issues that are not apparent during the research and design stages.

6.4.2 Safety and Reliability

Since this is an entirely new system for the car brakes (a safety critical component), it is extremely important to study the reliability of the system and the safety of its operation in the event of a failure.

It has been seen in Section 6.4.1 that one of the possible deteriorations faced by an MR brake is due to the phenomenon of In-Use Thickening, and altough the braking torque is not adversely affected by this problem, the performance of the car may suffer greatly and hence IUT is an important factor to take into account in the dimensioning of a system for commercial application.

A more serious problem that would entrain total failure of the MR brake is the possibility of leakage of the fluid. Given that the braking torque is achieved by the friction exerted by the fluid on the disk surfaces, loss of fluid would imply the loss of braking power. To ensure that this does not happen, it is fundamental that the seals be capable of withstanding the friction exerted by the rotating shaft and the system temperatures without degradation of properties. Moreover, it is worth considering a design with redundant sealing: two seals may be used in each opening and a sensor may be fitted in the space between the two seals - if it detects the presence of fluid, it may issue a warning (e.g., by means of a light on the car's instrument panel) indicating that the first seal failed and must be replaced. While this replacement does not occur, the second seal will ensure that the MR fluid is not lost.

Another problem that may affect the operation of the MR brake is the stoppage of electric supply. Altough unlikely, this can occur in the event of major problems with the battery or wires. Because such a failure would also entrain total brake failure, provisions must be made for alternative energy supply. This can come from a small energy storage (a mini-battery or a capacitor) at each brake. Either option would provide electricity for the coil (and thus braking torque) for the limited amount of time required to safely bring the car to a stop. Another, more sophisticated, solution involves the use of permanent magnets [111]. The MR brake can be built such that it comprises both an electromagnet and a permanent magnet. When the brakes are not required (i.e., the driver is not pressing the pedal), the electromagnet produces a

magnetic flux contrary to that induced by the permanent magnet, thus cancelling it and ensuring that the net magnetic flux is zero. As the pedal is progressively pressed, the flux produced by the electromagnet (contrary to that of the permanent magnet) approaches zero (at which point the flux is entirely due to the permanent magnet). Further pressing the pedal causes the electromagnet to produce a flux in the same direction as that due to the permanent magnet thus adding both effects and producing a greater overall magnetic flux. This concept presents two important advantages and two serious disadvantages: on the positive side, it is a fail-safe system in that power failure eliminates the flux produced by the electromagnet to cancel the permanent magnet and hence automatically results in a moderate braking torque being applied; also, the electromagnet can now be made smaller, as it only contributes to half the maximum magnetic flux, the rest being assured by the permanent magnet. On the negative side, this design increases the complexity of the system (more parts and the need to ensure that without the brake pedal pressed the electromagnet efectivelly cancels the flux of the permanent magnet); moreover, the electric power consumption is increased, as this system requires current to be fed to the electromagnet whenever the car is moving.

When application to actual cars is considered, the inclusion of some magnetic shielding to reduce the influence of external magnetic fields may also be considered.

6.5 Detailed Design

With all the above considerations in mind, the design of the MR brake can be finalised. This section presents the frozen configuration, its characteristics and performance. The characteristics of the proposed MR brake are summarised in Table 6.4 and the results of the finite element analyses for this particular configuration are presented in

Figs. 6.1 through 6.3.

Table 6.4: Final configuration summary

Property	Value
MR Fluid	MRF-132AD
Number of disks	2 (with a stator between them)
Maximum current	12 A
Coil wire	AWG 14
Approximate number of coil wire turns	80
Fluid gap thickness	0.1 cm
Volume of fluid	180 cm^3 [a]
rad_th_ring	0.5 cm
th_disk	0.5 cm
rad_disk	16.8 cm
rad_th_coil	1.2 cm
rad_th_casing	1.4 cm
ax_th_casing	1.6 cm
length_disk	5.5 cm [b]
Suggested seal	SKF CR 186x226x16 HS8 R

[a] The actual quantity may be slightly greater than this, given the need to fill the seal height

[b] Recall from table 5.8 that the optimum design had a length of 5 cm. However that would result in an inner radius of $16.8 - 5.0 = 11.8$ cm for the casing and no seals are available for that value. For this reason, the length was slightly increased so that SKF's CR 186x226x16 HS8 R may be used.

It is seen that the maximum temperature in the repeated braking cycle exceeds the maximum operating temperature of the MR fluid (130°C=403 K). This means that successful application to actual cars will require the inclusion of some of the solutions indicated in Section 4.2.

Should the results from the experimental analysis of the MR brake prove unsatisfactory in terms of the braking torque at very low speeds (close to zero), it is suggested to give the disk(s) a rough finish, so that the contact area with the fluid is increased. Small holes may also be drilled on the disk(s) surfaces to increase the ability of the fluid to grasp the disk(s).

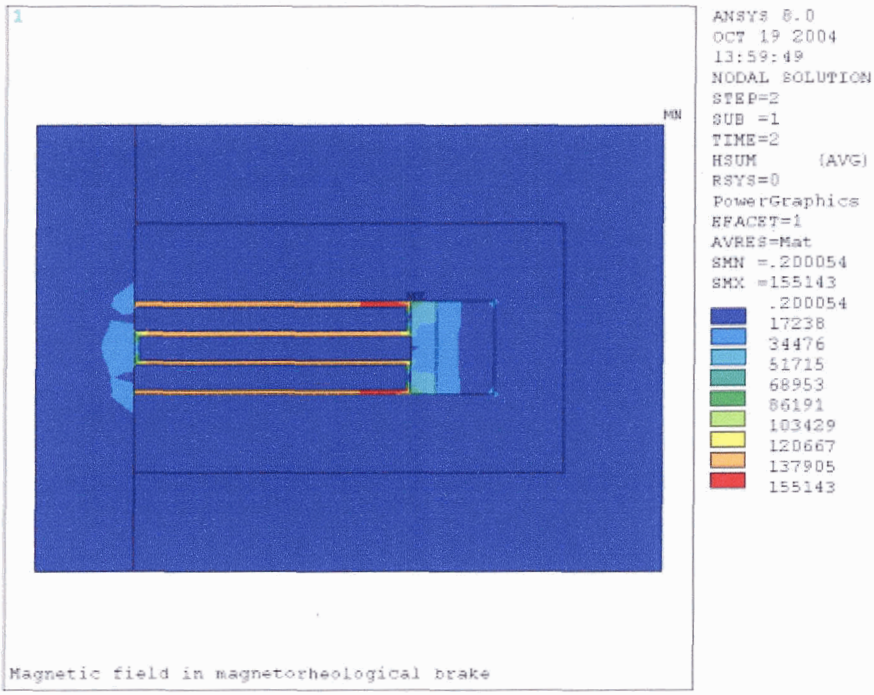

Figure 6.1: Magnetic field intensity in the final configuration

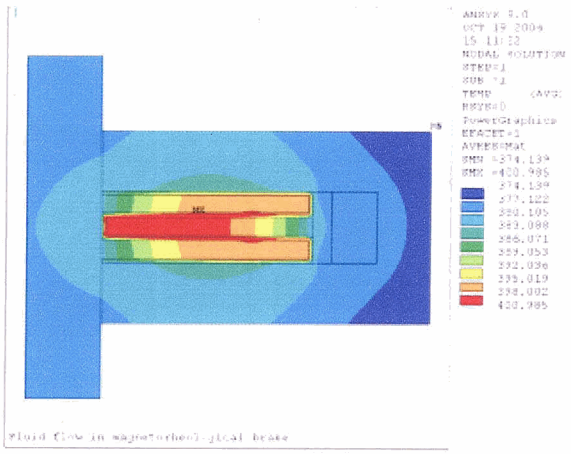

Figure 6.2: Steady-state temperature distribution in the final configuration

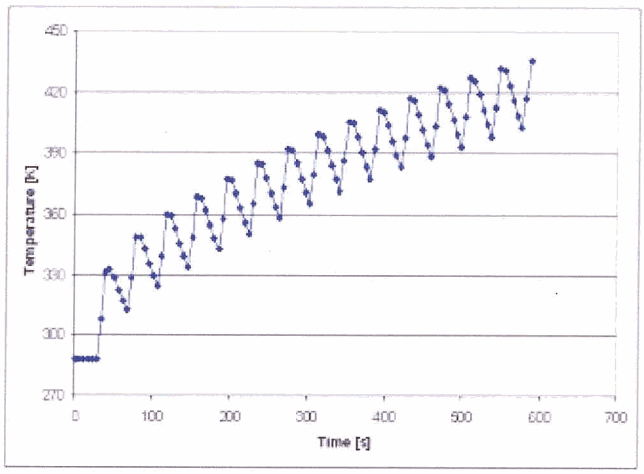

Figure 6.3: Evolution of the maximum temperature in the final configuration, subject to repeated brake-release cycles

Chapter 7

Conclusions and Future Work

It is unquestionable that a strong trend towards the totally electric car exists today. More and more mechanical, pneumatic and hydraulic systems are being replaced with electric counterparts. The brakes are among the candidates for such replacement, given the performance limitations of the traditional hydraulic design and the advantages in terms of integrated control of a totally electric brake system over today's electro-hydraulic configuration.

Solutions centred upon the use of piezoelectric or electrostrictive materials as actuators of the brake pads were initially pursued but were seen not to be a viable option with today's materials. Possible application to car brakes therefore constitutes one more motivation to the development of new piezoelectric and electrostrictive materials capable of producing greater forces and displacements.

A study of electrorheological or magnetorheological materials was then conducted and the latter were seen to be more adequate for an application to car brakes, given their power requirements, yield stress and energy dissipation and the operating temperature range. The existing analytical models of MR brakes (developped for small

systems producing low braking torques) were then analysed and seen to be inadequate for the present problem. This led to the deduction of new, more general equations of magnetorheological brakes. While no closed-form solution of these equations exists for the case of an MR brake of arbitrary geometry and comprising materials with non-linear magnetic permeability, they nevertheless gave an important insight into the role of the various design variables. This insight gained from the analytical models allowed the definition of preliminary MR brake configurations.

Finite element models were created and their results prompted important design changes to improve the overall performance of the MR brake. These analyses covered all disciplines involved in the operation of the brake: magnetostatics, fluid flow, heat trasnfer and structures. These studies were performed not only in terms of the performance capabilities of each design but also in terms of reliability.

In order to achieve the best designs, an optimisation procedure was carried out. Three different methods were used (two included with ANSYS and one custom-built). It was seen that the latter method (simulated annealing) gave the best results - although at a much greater computational expense. The improvements in performance brought by the optimisation methods were significant, with weight savings of over 50%. Designs meeting the initially proposed targets were obtained, with the desired improvement in response time over current hydraulic brakes and bringing a minor weight penalty compared with hydraulic brakes.

After the best design was obtained, final considerations regarding the manufacturing and assembly of the MR brakes, as well as their implementation were studied and led to minor modifications in the design (e.g., to meet the available seal sizes), leading to the final configurations for both a full-size model and a scaled-down proof-of-concept model.

Taking into account all the information from the design process, it is possible

to highlight the most important future work to be carried out in order to achieve commercial success for MR brakes:

- Extensive testing of the proposed prototype, particularly in terms of:
  - Confirming the results of the finite element models;
  - Evaluating the braking torque at very low speeds (less than 1 km/h);
  - Assessing the lifetime of the MR fluids and their behaviour under high shear stresses;
- Development of MR fluids with increased operating temperature ranges;
- Reduction in the off-state torque of MR fluids (e.g., through use of different carrier fluids). This is important not only to reduce the parasite drag which is present even when the brakes are not pressed (and which therefore brings a penalty in fuel consumption) but also to reduce the heat buildup whenever the car is not braking;
- Development of a control system for the MR brake, to fully exploit the possibilities presented by an all-electrical system.

Keep in mind that the configuration presented in the previous chapter was defined taking into account manufacturing constraints associated with the production of a small number o units, as well as the fact that the main goal of the present thesis is the creation of a new brake system and its proof of concept. Further improvements are possible and warranted once the commercialisation of MR brakes becomes a reality. For example, based on the results of the heat transfer analyses and the conclusion that cooling of MR brakes must be improved for increased safety and performance under extreme conditions, the cooling concepts introduced earlier may be incorporated.

Also, based on the results of the magnetostatics analyses, it is possible to determine that the region of the casing closest to the shaft is operating at very low magnetic flux densities. This suggests using a tapered casing, narrower closer to the disk, which would save weight without noticeable penalty to the performance of the system. Furthermore, this change would also increase the cooling area, thus lowering the temperature of the MR brake.

7.1 Original Contributions

The following original contributions of the present thesis are identified:

- Deduction of a new analytical model for the torque of an MR brake, more generic and comprehensive than the present state of the art;
- Creation of finite element models of an MR brake in the various disciplines of interest (magnetostatics; both steady-state and transient fluid flow and heat transfer; structural analysis), involving adaptations to commercially existing FEM software (ANSYS) to take into account phenomena of great importance in MR brakes (such as time-varying material properties);
- Development of a simulated annealing optimisation procedure for use in ANSYS, applicable to any problem regardless of the field(s) of study involved;
- Comparative study of different optimisation techniques for the particular problem of MR brakes, including the analysis of their performance and computational cost.
- Conceptual design of an MR brake suitable for application to a car, meeting all the goals presented in 1.3.

All in all, it is seen that the concept of magnetorheological brakes is promising and will become very competitive once better fluids (especially in terms of operating temperature, longevity and ability to withstand high shear rates) become available.

References

[1] Institut Français de l'Environment, Délégation à l'Aménagement du Territoire et à l'Action Régionale and Services du Ministère de l'Aménagement du Territoire et de l'Environnement. *Aménagement du territoire et environnement: Politiques et indicateurs.* Institut Français de l'Environment (IFEN), Orléans, France, July 2000.

[2] Paulo Santos. Hidrogénio - o próximo passo. *Visão*, (599):90–91, August 26 2004.

[3] S. Prince-Richard, M. Whale and N. Djilali. A techno-economic analysis of electrolytic hydrogen production for fuel cell vehicles. In *10th Canadian Hydrogen Conference*, pages 760–768, Quebec City, May 2000.

[4] Chris Pritchard. The little company that could. *The Peak*, 101(13), April 6 1999.

[5] Philippe B. de l'Arc. Histoire de la voiture electrique. Retrieved September 23, 2004, from http://perso.club-internet.fr/pboursin/pdgve4b.htm.

[6] Urs Schwegler. Electric vehicle symposium EVS 19: 19/23 Oktober 2002, Busan (Südkorea). Tagungsbericht, e'mobile - Schweizerischer Verband für elektrische und effiziente Strassenfahrzeuge, Fischingen, Switzerland, November 2002.

[7] Björn Hugosson, Kristina Haraldsson and Anders Folkesson. Demonstration of fuel cell buses under varying climate conditions: the Stockholm CUTE project. In *EVS 20 - The 20th International Electric Vehicle Symposium and Exposition*, Long Beach, CA, USA, November 15-19 2003.

[8] Tiago L. Farias. The fuel cell bus Berlin, Copenhagen and Lisbon: a sustainable and environmental solution for the urban passenger transportation sector. In *Transport and Air Pollution: 10th International Scientific Symposium*, Boulder, CO, USA, September 17-19 2001.

[9] The Peugeot 607. Sales and Marketing, Peugeot Motor Company plc, Coventry UK, July 2003.

[10] Abelardo López-Lagunas, Sek Meng Chai, J. Cross, B. Buchannan, Lawrence A. Carastro, S. Wang, D. Scott Wills, Nan Marie Jokerst, Martin A. Brooke and Mary Ann Ingram. Bidirectional single fiber low-cost optoelectronic interconnect for automotive applications. *IEEE Transactions on Vehicular Technology*, 49(1):281–287, January 2000.

[11] Martin Ward, editor. *Components News.* Ricardo Information Services, Shoreham-by-Sea, West Sussex, UK, July 2000.

[12] Heinrich Riedl. *Manual Práctico del Automóvil.* Ediciones CEAC, Barcelona, Spain, 1992.

[13] Brake modules - electric caliper. Delphi, Troy, MI, USA, 2002.

[14] David Bernard Drennen, Gary Chris Fulks and Robert John Disser. Electric brake caliper. US Patent 6,412,610 B1, United States Patent Office, July 2 2002.

[15] Kapjin Lee and Kyihwan Park. Analysis of an eddy-current brake considering finite radius and induced magnetic flux. *Journal of Applied Physics*, 92(9):5532–5538, 2002.

[16] David Cruz. Design and modeling of an eddy-current car brake system. In *MEGRC 2004 - Mechanical Engineering Graduate Research Colloquium*, University of Victoria, Victoria, BC, Canada, May 18 2004.

[17] Walquir Baptista de Moura, editor. *Freios e Sistemas Antibloqueio.* Grande Enciclopédia Pratica: Mecânica do Automóvel. Editora Século Futuro, Penha, R.J., Brasil, 1988.

[18] Kwangjin Lee. Numerical prediction of brake fluid temperature rise during braking and heat soaking. In *SAE International Congress and Exposition*, number 1999-01-0483 in SAE Technical Paper Series, Detroit, MI, USA, March 1-4 1999.

[19] John D. Anderson, Jr. *Fundamentals of Aerodynamics.* McGraw-Hill International Editions, second edition, 1991.

[20] ITI, Dresden, Germany. *SimulationX Tutorial 3: Automotive Powertrain and Vechicle Simulation*, 2002.

[21] E1102: Design fundamentals using advanced computer technologies - 4.1 Segway human transporter in simplified mechanics. Retrieved July 25, 2004, from http://gateway.seas.columbia.edu/e1102/week05/page2.html.

[22] Proposal for draft ammendment of TRANS/WP.29/GRRF/2003/25. 54 th GRRF, agenda item 1.3, Informal document 11, United Nations Economic Commission for Europe (UNECE) Working Party on Brakes and Running Gear (GRRF), Genève, Switzerland, 2003.

[23] Andrea de Adamich. *Corso di guida sicura.* Arnoldo Mondadori Editore, Segrate (Milano), Italy, 1991.

[24] Jean-Rémy Macchia. Dossier freinage - qui s'arrête le plus vite? *L'Automobile Magazine*, (625):117–127, July 1998.

[25] *Home Mechanic.* Gregory's Scientific Publications, Sydney, Australia, third edition, 1988.

[26] European Commission directive 95/48/EC. *Official Journal of the European Communities (English)*, L(233):73–85, September 30 1995. Office for Official Publications of the European Communities, L-2985 Luxembourg.

[27] Gérard Flocon, editor. *Toutes les voitures du monde 98/99.* L'Automobile Magazine, Issy-les-Moulineaux, France, 1998.

[28] Iranian automobiles - Citroen Xantia (English). Retrieved October 1, 2004, from http://www.farhangsara.com/cars/xantia.htm.

[29] Dieselnet website. Ecopoint Inc., 2004. Retrieved August 28, 2004, from http://www.dieselnet.com.

[30] Environmental Protection Agency. Motor vehicle emissions federal test procedure revisions; final regulations. *Federal Register*, 61(205):54851–54906, October 22 1996.

[31] David Newton, Ephrahim Garcia and Garnett C. Horner. A linear piezoelectric motor. *Smart Materials and Structures*, 6:295–304, 1997.

[32] Selcuk Gursan. Development of a continuous-motion piezoelectric rotary actuator for mechatronics and micropositioning applications. Master's thesis, University of Victoria, 1996.

[33] Amplified actuators L series. Adaptronics Inc., Troy, NY, USA. Retrieved October 16, 2003, from http://www.adaptronics.com/products /piezoelectric_actuators/amplified_piezo_actuators/APAL.html.

[34] Direct piezo actuators DPA. Cedrat Technologies, Meylan, France. Retrieved November 30, 2003, from http://www.cedrat.com/piezo/fiches/DPA.pdf.

[35] P-243 P-247, Preloaded open- & closed-loop extra-high-load HVPZT translators. Physik Instrumente (PI), Karlsruhe, Germany. Retrieved November 6, 2003, from http://www.physikinstrumente.de/pdf/P_243_247.pdf.

[36] Active strut member - ASM 10. Sensor Technology Limited, Collingwood, Ontario. Retrieved October 16, 2003, from http://www.sensortech.ca/asm10.html.

[37] Stan Burns. Design and modeling of an electrostrictive inchworm actuator. Master's thesis, University of Victoria, 2001.

[38] R. Stanway, J. L. Sproston and A. K. El-Wahed. Applications of electro-rheological fluids in vibration control: a survey. *Smart Materials and Structures*, 5:464–482, 1996.

[39] Mark R. Jolly, J. David Carlson and Beth C. Muñoz. A model of the behaviour of magnetorheological materials. *Smart Materials and Structures*, 5:607–614, 1996.

[40] What is the difference between MR and ER fluids? Lord Corporation Materials Division, Cary, NC, USA, March 2002.

[41] Ralph Lucas. Magnetizable fluid. US Patent 2,149,782, United States Patent Office, March 7 1939.

[42] Willis M. Winslow. Method and means for translating electrical impulses into mechanical force. US Patent 2,417,850, United States Patent Office, March 25 1947.

[43] Willis M. Winslow. Induced fibration of suspensions. *Journal of Applied Physics*, 20:1137–1140, 1949.

[44] Jacob Rabinow. The magnetic fluid clutch. *Transactions of the American Institute of Electrical Engineers*, 67:1308–1315, 1948.

[45] Jacob Rabinow. Magnetic fluid torque and force transmitting device. US Patent 2,575,360, United States Patent Office, November 20 1951.

[46] D.J. Peel, R. Stanway and W.A. Bullough. Dynamic modelling of an ER vibration damper for vehicle suspension applications. *Smart Materials and Structures*, 5:591–606, 1996.

[47] J. Melli-Huber, B. Weinberg, A. Fisch, J. Nikitczuk, C. Mavroidis and C. Wampler. Electro-rheological fluidic actuators for haptic vehicular instrument controls. In *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'03)*, pages 262–269, Los Angeles, California, USA, March 22-23 2003.

[48] J.D. Carlson, D.M. Catanzarite and K.A. St. Clair. Commercial magneto-rheological fluid devices. In W.A. Bullough, editor, *Proc. 5th Int. Conf. on ER Fluids, MR Fluids and Assoc. Tech., July 1995*, World Scientific, pages 20–28, Singapore, 1996.

[49] J. David Carlson, Douglas F. LeRoy, John C. Holzheimer, Donald R. Prindle and Robert H. Marjoram. Controllable brake. US Patent 5,842,547, United States Patent Office, December 1 1998.

[50] J. David Carlson. Magnetorheological brake with integrated flywheel. US Patent 6,186,290 B1, United States Patent Office, February 13 2001.

[51] MR brake MRB-2107-3 product bulletin. Lord Corporation Materials Division, Cary, NC, USA, 2003.

[52] Greg Paula. Exercise equipment puts on magnetorheological brakes. *Mechanical Engineering*, May 1997. Published by ASME - American Society of Mechanical Engineers, New York, NY, USA.

[53] Gregory M. Webb. Exercise apparatus and associated method including rheological fluid brake. US Patent 5,810,696, United States Patent Office, September 22 1998.

[54] Russ Gilbert and Mark Jackson. Magnetic ride control. *GM Tech Link*, 4(1):1–2, January 2002.

[55] J. David Carlson. Controlling vibration with magnetorheological fluid damping. *Sensors*, 19(2), February 2002.

[56] S. J. Dyke, B. F. Spencer Jr., M. K. Sain and J. D. Carlson. Modeling and control of magnetorheological dampers for seismic response reduction. *Smart Materials and Structures*, 5:565–575, 1996.

[57] B. M. Kavlicoglu, F. Gordaninejad, C. A. Evrensel, N. Cobanoglu, Y. Liu, A. Fuchs and G. Korol. A high-torque magneto-rheological fluid clutch. In *Proceedings of SPIE Conference on Smart Materials and Structures*, San Diego, CA, USA, March 2002.

[58] Swaminathan Gopalswamy, Samuel Miller Linzell and Gary Lee Jones. Magnetorheological fluid clutch with minimized reluctance. US Patent 5,845,752, United States Patent Office, December 8 1998.

[59] Stephanie v.L. Henkel. MR knee is a step up. *Sensors*, 18(1), January 2001.

[60] Hugh Herr. Prosthetic and orthotic limbs. *Journal of Rehabilitation Research and Development*, 39(3 (Supplement)):11–12, May-June 2002.

[61] Salvatore Liberatori. *Progetto costruttivo di una pedana riabilitativa a controllo magnetoreologico.* PhD thesis, Università degli Studi di Pisa, 2004.

[62] G. A. Flores and J. Liu. In-vitro blockage of a simulated vascular system using magnetorheological fluids as a cancer therapy. *European Cells and Materials*, 3(Supplement 2):9–11, 2002.

[63] Jin-Hyeong Yoo and Norman M Wereley. Design of a high-efficiency magnetorheological valve. *Journal of Intelligent Material Systems and Structures*, 13(10):679–685, October 2002.

[64] W. A. Bullough. Smart machine systems. In Afzal Suleman, editor, *Smart Structures: Applications and Related Technologies.* Springer Verlag, 2002.

[65] J. Hesselbach and C. Abel-Keilhack. Active hydrostatic bearing with magnetorheological fluid. *Journal of Applied Physics*, 93(10):8441–8443, 2003.

[66] Hydrocarbon-based MR fluid MRF-132AD product bulletin. Lord Corporation Materials Division, Cary, NC, USA, 2003.

[67] Water-based MR fluid MRF-241ES product bulletin. Lord Corporation Materials Division, Cary, NC, USA, 2002.

[68] Sławomir Bydoń. Simulation of induction motor shaft positioning system with magnetorheological brake. In *XXVIII ASR '2003 Seminar, Instruments and Control*, Ostrava, May 6 2003.

[69] Engineering note - magnetic circuit design. Lord Corporation Materials Division, Cary, NC, USA, 1999.

[70] Daniel Marques. *Memorandum do Electricista.* Abel d'Oliveira, Lda., Lisboa, Portugal, second edition, 1924.

[71] Wire gauge and current limits. PowerStream, Orem, UT, USA, 2004. Retrieved October 4, 2004, from http://www.powerstream.com/Wire_Size.htm.

[72] Donald Herrington and Stanley Meacham, editors. *Handbook of electronic tables & formulas.* Howard W. Sams & Co., Indianapolis, IN, USA, second edition, 1962.

[73] Copper magnet wire data - winding data. Weico Wire & Cable, Edgewood, NY, USA. Retrieved October 4, 2004, from http://www.weicowire.com/WindingData.htm.

[74] Donald R. Askeland and Pradeep P. Phulé. *The Science and Engineering of Materials.* Brooks/Cole, fourth edition, 2003.

[75] Ferdinand P. Beer and E. Russel Johnston, Jr. *Mechanics of Materials.* McGraw-Hill, second edition, 1992.

[76] William F. Smith. *Princípios de Ciência e Engenharia dos Materiais.* McGraw-Hill de Portugal, third edition, 1998.

[77] William F. Smith. *Principles of materials science and engineering.* McGraw-Hill International Editions, second edition, 1990.

[78] Bimetallic corrosion. National Physical Laboratory, Teddington, Middlesex, UK, 2000.

[79] J.R. Rossum. Fundamentals of metallic corrosion in fresh water. Roscoe Moss Company, Los Angeles, CA, USA. Retrieved July 22, 2004, from http://www.roscoemoss.com/metallic_corrosion.html.

[80] Brembo caliper clearance guide. Retrieved October 4, 2004, from http://www.buybrakes.com/brembo/caliper-clearance.html.

[81] Momo website - ruote. Retrieved October 4, 2004, from http://www.momo.it/Momo/ruote.html.

[82] Daryl L. Logan. *A First Course in the Finite Element Method.* Brooks/Cole, third edition, 2002.

[83] Walid H. El-Aouar. Finite element analysis based modeling of magneto rheological dampers. Master's thesis, Virginia Polytechnic Institute and State University, 2002.

[84] Ansys Inc., Canonsburg, PA, USA. *ANSYS Release 8.0 Documentation*, 2003.

[85] Ansys Inc., Canonsburg, PA, USA. *ANSYS Theory Reference*, 2003.

[86] Eliahu Zahavi. *The Finite Element Method in Machine Design.* Prentice-Hall, 1992.

[87] Arthur L. Cook and Clifford C. Carr. *elements of Electrical engineering.* John Wiley and Sons, fifth edition, 1947.

[88] Raymond A. Serway. *Physics for scientists and engineers.* Saunders College Publishing, fourth edition, 1996.

[89] P. C. Sen. *Principles of Electric Machines and Power Electronics.* John Wiley & Sons, second edition, 1997.

[90] Frank P. Incropera and David P. DeWitt. *Fundamentals of Heat and Mass Transfer.* John Wiley and Sons, fourth edition, 1996.

[91] Vasco de Brederode. *Fundamentos de Aerodinâmica Incompressível.* Edição do Autor, 1997.

[92] J.C. Gibbings. On boundary-layer transition wires. ARC CP 462, Aeronautical Research Council, 1959. As cited in [91].

[93] Thomas Valvano and Kwangjin Lee. An analytical method to predict thermal distortion of a brake rotor. In *SAE 2000 World Congress*, number 2000-01-0445 in SAE Technical Paper Series, Detroit, MI, USA, March 6-9 2000. Reprinted From: Brake Technology: ABS/TCS Systems, NVH, and Foundation Brakes (SP1537).

[94] Carlos A. G. de Moura Branco. *Mecânica dos Materiais.* Fundação Calouste Gulbenkian, third edition, 1998.

[95] B. B. Verma, J. D. Atkinson and M. Kumar. Study of fatigue behaviour of 7475 aluminium alloy. *Bulletin of Materials Science*, 24(2):231–236, April 2001.

[96] Typical mechanical properties of wrought aluminum alloys at various temperatures. Technical Report 06813G, ASM International, Materials Park, Ohio, USA.

[97] European Aluminium Association and Matter Project. aluSelect creep rupture properties. The University of Liverpool, Liverpool, UK, 2001. Retrieved October 18, 2004, from http://aluminium.matter.org.uk/aluselect/07_creep_browse.asp.

[98] Luís Valadares Tavares, Rui Carvalho Oliveira, Isabel Hall Themido and Francisco Nunes Correia. *Investigação Operacional.* McGraw-Hill de Portugal, 1996.

[99] Garret N. Vanderplaats. *Numerical Optimization Techniques for Engineering Design - With Applications.* McGraw-Hill, 1984.

[100] Francesco Pontrandolfo, Giuseppe Monno and Antonio E. Uva. Simulated annealing vs. genetic algorithms for linear spline approximation of 2d scattered data. In *XII ADM International Conference*, Rimini, Italy, September 5-7 2001.

[101] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.

[102] S. Kirkpatrick, C. D. Gellatt, Jr. and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 13 1983.

[103] V. Černy. Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45(1):41–51, 1985.

[104] Peter J.M. van Laarhoven and Emile H.L. Aarts. *Simulated Annealing: Theory and Applications.* D. Reidel Publishing Company, Dordrecht, Holland, 1987.

[105] Horacio Martínez Alfaro, Homero Valdez Peña and Jaime Ortega Consuegra. Uso de recocido simulado para minimizar costos en la industria del acero. *Transferencia de Posgrado, Investigación y Extensión en el Campus Monterrey*, 13(49), January 2000.

[106] Lord rheonetic^{TM} magnetically responsive technology MR fluid comparison sheet. Lord Corporation Materials Division, Cary, NC, USA, 2002.

[107] Lord rheonetic^{TM} magnetically responsive technology magneto-rheological fluid/material compatibility. Lord Corporation Materials Division, Cary, NC, USA, 2002.

[108] SKF website. SKF, Göteborg, Sweden. Retrieved October 4, 2004, from http://www.skf.com.

[109] Cláudio Roberto Losekann. *Processos de Fabricação I - Unidade I.* Universidade do Vale do Itajaí - Direção do Centro de Educação Superior, São José, Brasil, 2001.

[110] J. David Carlson. What makes a good MR fluid? In *Proceedings of the 8th International Conference on Electrorheological (ER) Fluids and Magnetorheological (MR) Suspensions*, Nice, France, July 9-13 2001.

[111] J. David Carlson. Permanent-electromagnet systems ("magnetic cancellation"). Lord Corporation Materials Division, Cary, NC, USA, May 2002.

[112] EPA - OTAQ - Federal test procedure revisions, 2004. Retrieved August 27, 2004, from http://www.epa.gov/otaq/sftp.htm.

Appendix A

Representative Driving Cycles

The speed versus time curves for the two representative driving cycles (FTP 75 and US 06) are presented in this appendix. These are based on the values from [112].

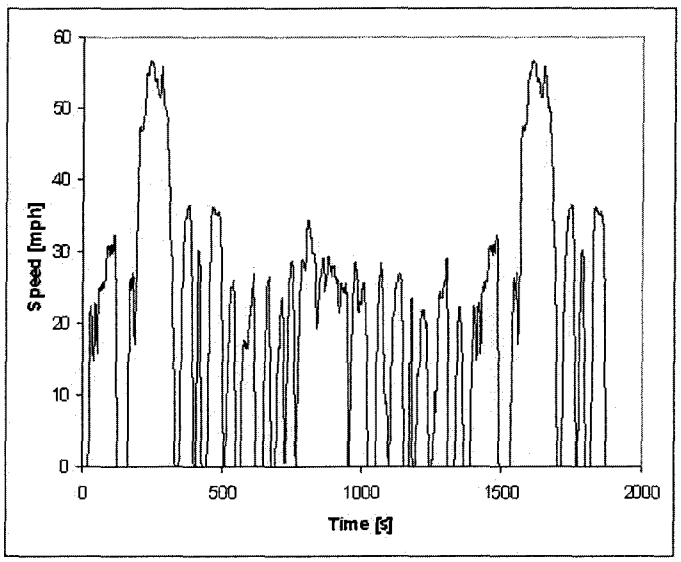

Figure A.1: Speed vs. time in the FTP 75 cycle
Source: [112]

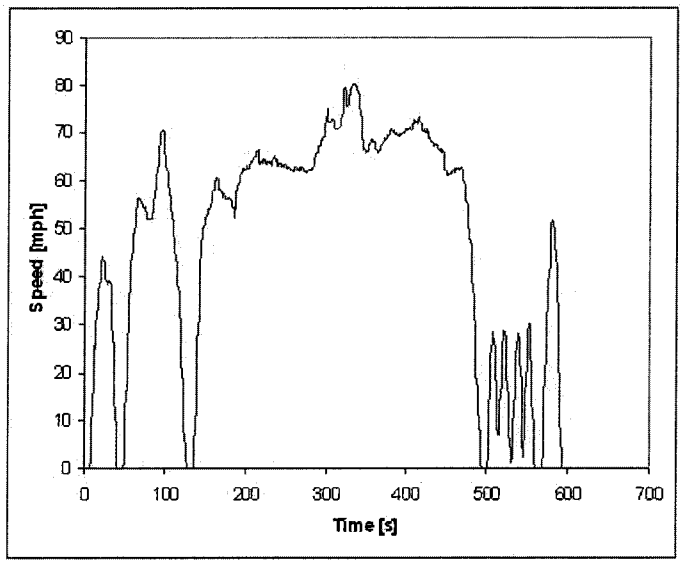

Figure A.2: Speed vs. time in the US 06 cycle
Source: [112]