NPL
National Physical Laboratory

Report

NPL REPORT
DEM-ES-026

Software Support for Metrology – Good Practice Guide 15

Continuous Modelling

T J Esward and L Wright

March 2007

dti
FUNDED BY THE DTI

National Physical Laboratory | Hampton Road | Teddington | Middlesex | United Kingdom | TW11 0LW

Switchboard 020 8977 3222 | **NPL Helpline** 020 8943 6880 | **Fax** 020 8943 6458 | **www.npl.co.uk**

Software Support for Metrology
Good Practice Guide 15

Continuous Modelling

T J Esward and L Wright
Mathematics and Scientific Computing Group

March 2007

ABSTRACT

This good practice guide builds on work carried out during the second Software Support for Metrology (SS*f*M) programme. The guide continues the work carried out for project 1.4, "Continuous Modelling In Metrology", and relates it to the structure used for describing the modelling process in the best practice guide (BPG) "Software Support for Metrology Best Practice Guide 4: Discrete Modelling" to create a guide to good practice in continuous modelling. The guide is aimed at new and experienced users of continuous modelling techniques who wish to apply good practice to their modelling projects. It may also be of interest to those who want a better understanding of the similarities and differences between different numerical methods for solving models.

Please note that the majority of figures in this report are designed to be seen in colour, hence black and white printed copies may be less easy to interpret than the on-screen version.

ISSN 1744–0475

National Physical Laboratory
Hampton Road, Teddington, Middlesex, United Kingdom. TW11 0LW

We gratefully acknowledge the financial support of the UK Department
of Trade and Industry (National Measurement System Directorate)

Approved on behalf of the Managing Director, NPL
by Jonathan Williams, Knowledge Leader of the Electrical and Software team

Contents

1. Introduction

This good practice guide builds on work carried out during the second Software Support for Metrology (SS*f*M) programme. The guide continues the work carried out for project 1.4, "Continuous Modelling In Metrology" [1,2,3], and relates it to the structure used for describing the modelling process in the best practice guide (BPG) "Software Support for Metrology Best Practice Guide 4: Discrete Modelling" [4] to create a guide to good practice in continuous modelling.

In addition, this guide illustrates the similarities and differences between discrete and continuous modelling. In particular, the use of the same terminology as the discrete modelling BPG [4] shows that the main steps in the continuous modelling process are identical or analogous to those in the discrete modelling process.

The guide is aimed at new and experienced users of continuous modelling techniques who wish to apply good practice to their modelling projects. It may also be of interest to those who want a better understanding of the similarities and differences between different numerical methods for solving models. The scope and structure of the guide are described in section 1.2, and a summary of good practice is given in section 7.

Please note that the majority of figures in this report are designed to be seen in colour, hence black and white printed copies may be less easy to interpret than the on-screen version.

1.1 Mathematical modelling in metrology

Modelling can be regarded as a simplification process. A complex physical situation is turned into a solvable model by making simplifying assumptions about the factors that most influence the behaviour of the quantity of interest.

The vast majority of models assume that the physical world is a continuum and do not describe the interaction of individual atomic particles. Whilst this is a reasonable assumption for most metrology problems, in recent years there has been a growing interest in modelling micro- and nano-scale problems. The increased level of interest has lead to the question "at how small a scale do bulk models cease to be accurate?". This question, and other problems such as how to bridge the gap between molecular models and bulk models, is explored in the forthcoming NPL report "Mathematical Modelling for Metrology at the Nanoscale" [5], an output of the SS*f*M-3 programme.

In the discrete modelling BPG [4], the modelling process was seen as an evolutionary process with four components. The same approach will be taken in this guide, which leads to a new description of some aspects of continuous modelling. In particular, the model solving component suggests a new way of looking at the discretisation process that many continuous modelling techniques require, as will be explained in section 4.

The four components are:

Model building: Developing a mathematical model of the experimental system in terms of mathematical equations involving parameters that describe all the relevant aspects of the system. For continuous models, at least one of the equations will be either in differential or integral form, and the model will also require a solution domain, boundary conditions, and possibly initial conditions.

Model solving: Determining estimates of the model parameters from the measured data by solving the mathematical equations constructed as part of the model. The model solving steps of many continuous modelling techniques involve making assumptions about the local

behaviour of the solution within a small region to construct an approximate solution to the model.

Software implementation of solvers: Solution of continuous models often requires the solution of large systems of equations, making software implementation of the solver an important part of the process. A large number of proprietary packages exist that enable the user to build and solve the model within a single environment, often without the user needing to have any knowledge of the underlying mathematics.

Model validation: Determining whether the results produced are consistent with the input data, theoretical results, reference data, expert knowledge, etc. The results of an unvalidated model are little better than educated guesswork.

These descriptions use terms that need careful definition. Throughout this guide, the **parameters** of a model are the quantities that are to be estimated, and the **variables** are the quantities that are considered to be known. The definitions used here have been chosen to be consistent with the discrete modelling BPG [4]. These terms are discussed further in section 2.2.

1.2 Scope, structure and usage of this guide

This guide describes model building (section 3), model solving (section 4), and software implementation (section 5) using continuous models. Section 2 provides some general definitions of terms used throughout the guide and gives some motivation for the guide's contents and structure.

It considers the model in isolation, without relating the model to the problem that the metrologist wishes to solve, so optimisation methods for continuous modelling are not described. The issues generated by this approach are discussed further in section 2.2. A large amount of material on optimisation already exists, including some useful websites [6,7], and software resources such as the NAG library [8] and Matlab [9] include optimisation routines.

No guidance is given here on seeking analytical solutions to problems. Very few continuous modelling problems have analytical solutions, and the problems that do have them tend to involve simplifications and assumptions that make the difference between the model and reality large. However, some problems that have analytical solutions can be useful for model validation purposes, so such solutions can be valuable when they are available.

This guide does not describe model validation methods in detail, because an extensive report on model validation in continuous modelling has been published as part of a previous SS*f*M programme [1]. The main conclusions of the validation report are described in section 6.

A summary of good practice is provided in section 7, and illustrative case studies are shown in section 8. Sections 9 and 10 describe some of the models and solution techniques that are most common in metrology. Users wishing to understand good practice in continuous modelling without reading the details should read sections 1, 2, and 7 of this report. Sections 3, 4, 5, and 6 explain the reasons for the recommendations that are made in section 7, and sections 9 and 10 can be used as a reference guide to the standard model building and solving techniques for continuous models in metrology.

1.3 Acknowledgments

This guide was produced under the Software Support for Metrology (SS*f*M) programme, which is managed behalf of the DTI by the National Physical Laboratory. For more information on the SS*f*M programme visit the website at http://www.npl.co.uk/ssfm or contact the programme manager Mr Bernard Chorley (phone +44 20 8943 7050, email: ssfm@npl.co.uk, National Physical Laboratory, Hampton Road, Teddington, Middlesex, UK TW11 0LW).

The authors would like to thank Graham Machin and Dave Lowe of the Thermal Measurement Team, Division of Engineering and Process Control at NPL for their contributions to the first case study, and Stephen Robinson of the Acoustics Team, Division of Quality of Life at NPL for his contributions to the second case study.

1.4 General notation

Unless otherwise specified in the text, the following notation is used throughout.

$\mathbf{x}$	position vector in Cartesian (x, y, z), cylindrical polar (r, θ, z) or spherical polar (R, θ, φ) coordinates
t	time
Ω	model domain
$\partial\Omega$	domain boundary
V	three-dimensional volume model domain
S	closed surface model domain (often the boundary of V)
$\mathbf{n}$	normal vector to a surface
A	two-dimensional area model domain
$f^{(n)}(x)$	nth derivative of the function f with respect to x
$f_x(x, y)$	partial derivative of f with respect to x

2. Motivation and general definitions

This section discusses some of the reasons for modelling a measurement process, and suggests ways in which consideration of these reasons can help to improve clarity and communication during the modelling process.

The section also gives definitions of many of the terms that are used in this guide. Many definitions that are used in discrete modelling are directly applicable to continuous modelling. The discrete modelling BPG [4] gives some of the following, but they will be repeated for completeness. Section 3 includes further discussion of some of the ideas introduced here.

2.1 Why model a measurement process?

All models start from a question that cannot be answered by measurement alone. Examples include:

- What is the best design for this experiment/How can I achieve a uniform state with this equipment? If a new experiment is being developed, a model can help identify the best design for the equipment before measurements are possible. Similarly, if a particular state (field strength, temperature uniformity, stress state, etc.) is required during an experiment, a model can suggest alterations to the equipment that can achieve the requirements.
- What is happening inside this object? In general, measurements of quantities such as stress and temperature inside a test sample are difficult to achieve. A model can predict the values of useful quantities in places where they cannot be measured.
- How does that happen? Many models are made because a better understanding of a measurement process is required. The process may be producing unexpected results, or an event (e.g. damage or failure) with an unclear cause may have occurred during the measurement process. A model can help to explain what has caused the unexpected occurrence.
- Are my assumptions right/What happens if the assumptions I made are wrong? Some experiments require their results to be processed to obtain the value that is of interest. Often the data processing is based on assumptions such as a uniform stress state, constant temperature, an ideal uniform material, etc. A model makes it possible to test these assumptions and to investigate the effects of false assumptions in a controlled way.
- How can I generate the results I'm interested in from the measurements I have made? In some cases (see case study 2 for an example) the measured values require extensive processing via a continuous model in order to produce the results of interest.
- How sensitive is the measured value to temperature/pressure/frequency/location? This is a key question for measurement processes as the answer affects the uncertainty of the measurement results. Some form of model of the measurement process is the only way to investigate the sensitivity of the results to external factors.

Some modellers find the definition of this question the most difficult stage of the modelling process, and it is usually the case that the answer to one question leads to a new set of questions. For instance, knowing the temperature or stress distribution inside a material often leads to asking why those results have occurred and how they can be altered or improved. However, it is important to have a clear idea of the question to be answered because it affects everything about the model. The question affects which results need to be obtained, it affects

which features of the process must be included in the mode, and it affects the degree to which the model needs to be validated.

In many cases, the person carrying out the measurement process and the person modelling that process are not the same person and so communication can become a problem. This problem can be avoided by using a "user specification" so that the requirements of all parties are understood. Even if the experimentalist is carrying out the modelling, a user specification can help to clarify the aims and limitations of the model and thus speed up the model development process.

The user specification should be written in terms of the experiment rather than in mathematical terms, because the aim is to capture the needs of the experimentalist rather than to define the model in full. The specification needs to include:

- Definition of the question to be answered, i.e. the purpose of the modelling activity.
- Definition of the model results that will be used to answer that question, including their nature (e.g. numerical values, visualisations, etc).
- Definition of who will use the model: a model that is being run by its developer to generate results may differ from a model that is part of a tool to be run by other people.
- A description of the assumptions made about the experiment.
- A description of how the model will be validated (even if this is only a check that the results look right).
- A list of the information that the experimentalist will supply to the modeller (usually dimensions of the experimental apparatus, measured values, and material definitions).

These points ensure that the modeller and experimentalist understand what the model will do, what the modeller requires in order to get the results, and what the limitations of the model are. The user specification need not be a formal document, but a record of the points listed above should be kept.

2.2 Types of model and quantity

A **physical** model is one in which there is a theory that defines how the variables and parameters depend on each other. An **empirical** model is one in which a relationship between the variables is expected or observed but there is no supporting theory.

Many models have both physical and empirical components. The differential equations used in continuous modelling are almost always physical models, but often the equations used to describe the behaviour of some of their variables are empirical. A good example of this is the use of empirical models of temperature-dependent material properties in the heat equation, a physical model.

The model **domain** is the region of space and/or time within which a solution is required. At the limits of the domain, which may be infinite, a set of **boundary conditions** will be defined that describe the behaviour of some of the variables and parameters at the boundary. If the domain includes a time dimension, a set of **initial conditions** at time $t = 0$ will also be required.

An **explicit** model is one in which one or more of the variables is given as a directly computable function of the remaining variables and parameters (variables and parameters are defined in section 1.2). We write $y = \varphi(\mathbf{x}, \mathbf{a})$ to show that y is a function of the model variables $\mathbf{x}$ and parameters $\mathbf{a}$. If $\mathbf{x}$ and $\mathbf{a}$ are known, then the corresponding value for y can be calculated. An **implicit** model is one in which the variables and parameters are linked through

a set of equations. We write, for example, $f(\mathbf{x}, \mathbf{a}) = 0$ to show that the components of $\mathbf{x}$ are related implicitly. In general, these definitions are most helpful for discretised models rather than models involving differential equations.

It may be possible to rewrite an implicit model as an explicit model, but it may not be helpful to do so since often the explicit form is less amenable to solution than the implicit form. For example, in its time-dependent form, the heat equation for constant thermal properties is usually written

$$\frac{\partial T(\mathbf{x},t)}{\partial t} = \alpha \nabla^2 T(\mathbf{x},t), \quad \mathbf{x} \in \Omega, t \in [0, t_F], \tag{2.1}$$

where Ω is the domain, $\mathbf{x}$ is a position vector, T is temperature, t is time, and the boundary and initial conditions are written

$$a(\mathbf{x},t)T(\mathbf{x},t) + \mathbf{b}(\mathbf{x},t).\nabla T(\mathbf{x},t) = c(\mathbf{x},t), \quad \mathbf{x} \in \partial\Omega$$
$$T(\mathbf{x},0) = T_0(\mathbf{x})$$

where $\partial\Omega$ is the boundary of the domain. If the equation (2.1) is integrated with respect to time between 0 and t^*, it can be rewritten

$$T(\mathbf{x},t^*) = T_0(\mathbf{x}) + \int_0^{t^*} \alpha \nabla^2 T(\mathbf{x},t) dt, \quad \mathbf{x} \in \Omega, t^* \in [0, t_F]$$

This form is not necessarily helpful, since the quantities on the right hand side of the equation may not be easy to compute, but there are circumstances under which an explicit form may be useful (see section 2.2).

The terms implicit and explicit are commonly used in the solution of time-dependent models, since the solution methods used for such problems are generally divided into implicit and explicit types. The way that the terms are used in time-dependent applications can be regarded as a special example of the definition given above. This usage is explained further in section 4.2.5.

A **linear** model is one in which all of the parameters appear linearly in the governing equations and the boundary conditions. A **nonlinear** model is one in which one or more of the parameters appears nonlinearly in the governing equations or the boundary conditions. A model that is nonlinear in its parameters can easily be linear in its variables, so it is useful to state what quantities the linearity or nonlinearity refers to.

Coordinate quantities identify a point in the domain. They are usually variables, and consist of a position and possibly a time. They can be parameters if the problem to be solved is the location of the maximum value of some quantity. **State quantities** characterise the state of the system at a point. Examples are temperature, stress, pressure, and field strength. The quantities being differentiated in the governing equations are state quantities.

A **field quantity** is a state quantity that varies over a model domain. If the model is to have a non-trivial solution, the quantity that is differentiated in the model equation must be a field quantity. Field quantities can be scalars (e.g. temperature), vectors (e.g. displacement) or tensors (e.g. stress).

A **transient** model is a model that includes time as one of the coordinate quantities. A **static** or **steady-state** model is one that does not. Many transient models converge to a steady state as time increases. Steady-state models are easier to solve because they have fewer coordinate variables, and because they do not require the same stability requirements when approximated numerically.

2.3 Problems, models, parameters and variables

In section 2.1, parameters were defined as unknown quantities that are to be estimated, and variables as quantities regarded as known from measurement or assumptions. It is important to note that a problem to be solved using a model has two sets of variables and parameters: the **problem** variables and parameters, and the **model** variables and parameters.

The problem variables are the quantities that are known in reality, and the problem parameters are the quantities that need to be estimated. The model variables are the quantities to which values must be assigned for the model to be solvable (the domain geometry, the domain properties, the boundary conditions, etc.), and the model parameters are the values produced by the solution process. For a given physical system, the model variables and parameters will be the same no matter what questions are to be asked about the system, but the problem variables and parameters are defined by the question that is asked.

As an example, consider a bar of constant and uniform material properties and uniform square cross-section that is held at a steady temperature at one end, is insulated along its long faces, and is losing heat by convection to an environment of varying temperature at the other end. The governing equation, initial conditions, and boundary conditions for this situation will be

$$\frac{\partial T(x,y,z,t)}{\partial t} = \alpha \nabla^2 T(x,y,z,t), \qquad 0 \le x \le L, 0 \le y,z \le b,$$

$$T(x,y,z,0) = T_0(x,y,z), \qquad 0 \le x \le L, 0 \le y,z \le b,$$

$$T = T_f, \quad x = 0, 0 \le y,z \le b,$$

$$\lambda \frac{\partial T}{\partial y} = h(T)(T - T_{amb}(t)), \quad x = L, 0 \le y,z \le b,$$

$$\frac{\partial T}{\partial y} = 0, \quad 0 \le x \le L, y = \pm b, 0 \le z \le b,$$

$$\frac{\partial T}{\partial z} = 0, \quad 0 \le x \le L, 0 \le y \le b, z = \pm b,$$

where T is the temperature, T_0 is the initial temperature distribution, T_f is the fixed temperature at one end, $h(T)$ is the heat transfer coefficient, α is the thermal diffusivity, λ is the thermal conductivity, $T_{amb}(t)$ is the temperature of the environment, L is the length of the bar, and b is the length of the square cross-section.

In order for this model to be solved, the following quantities and functions must be known: T_0, T_f, $h(T)$, α, λ, $T_{amb}(t)$, L, and b. These are the model variables. The model parameters are the values of T that are not defined by the initial or boundary conditions.

Suppose that three problems are of particular interest for the metrologist:

1. How does the temperature of the free end of the bar vary with time?
2. When is the maximum temperature of the free end reached?
3. Given a set of measurements of the temperature at the centre of the free end of the bar, what is the best estimate of the thermal conductivity of the bar?

The first problem can be written "what is $T(L, y, z, t)$?", and so the problem parameter is one of the model parameters, a state quantity which can be obtained directly from the model results.

The second problem can be written "find t^* such that $T(L, y, z, t^*) \ge T(L, y, z, t)$ for all $t \ge 0$", and so the problem parameter is a coordinate quantity that can be found from the model parameters.

The third problem can be rewritten as a minimisation problem. Suppose that the temperature measurements at the centre of the bar occur at times t_i, $i = 1, 2, \ldots, m$, and that the measured temperature at time t_i is T_i. Then the minimisation problem could be written as

$$\underset{\alpha}{\text{minimise}} \quad \sum_{i=1}^{m} \left[T\left(L, b/2, b/2, t_i\right) - T_i\right]^2.$$

In this case the problem parameter is one of the model variables. The model cannot be solved without a value for α, but the value of α is not known.

If any of the problem parameters is a model variable, the solution to the problem can only be found using an optimisation technique. Such problems are particularly common in metrology. Many applications of modelling in metrology are simulations of experiments that are used to determine a physical property of a material sample, so one or more of the domain properties may be unknown.

Optimisation methods can iteratively alter the parameters in order to make some set of model results agree with a corresponding set of experimental data. For instance, in the example above, an optimisation method could alter the thermal diffusivity of the bar to make the model predictions as close as possible to the measurements. The optimisation method attempts to minimise some function of the differences between the variables and the equivalent measured values, called the **objective function**. A commonly used objective function is the sum of squares of the differences, as was used in the example given above. Optimisation techniques are often necessary in discrete modelling, because the solution procedure used to generate parameter estimates often involves minimisation of an error norm.

If an explicit model for the variables exists, so that $y_i = f(\mathbf{x}; \mathbf{a})$, the objective function is an expression whose dependence on the parameters is given. If the dependence of the objective function on the parameters is known, the optimisation is considerably more straightforward than the case where only an implicit model is available. The partial derivatives of the objective function with respect to each parameter are found and set to zero, producing a set of equations that can be solved to find the best possible values of the parameters.

If the model is implicit, optimisation is still possible but the methods that can be used are often less efficient and require more computational effort. Continuous models do not generally have an explicit analytic solution, so the objective function for a continuous model is usually implicit.

The application of optimisation methods to continuous models can be problematic, particularly when commercial proprietary software is used. The unavailability of derivatives means that some of the most efficient techniques cannot be used, the need for repeated evaluations of the original model can increase computational effort to an unacceptable degree, and non-uniqueness of the solution can be difficult to disprove. However, such problems are not impossible, and several examples of optimisation using finite element analysis to generate the objective function exist [10, 11].

If the problem is to be expressed as an optimisation, it is important to be sure that the routines that generate the objective function are correct. This requirement means that the model development and validation processes are identical whether optimisation is required or not. If optimisation is required, the initial validation will treat variables of the optimised model as parameters and vice versa, and further validation will be required to ensure that the optimisation routines are working correctly.

3. Model building

A continuous model is usually defined by four entities: the governing equations, the problem domain, the domain properties, and the boundary conditions. These four entities can depend on one another: in particular, it is fairly common for the boundary conditions to depend on the material properties and in many cases the domain depends on the governing equations. In some cases the domain properties create further governing equations. Extra entities may be required in some cases, for example time-dependent models require definition of a set of initial conditions, and in some cases one or more of the entities may be unnecessary. In general, however, these four entities uniquely define a model. Each of the four entities will be discussed in more depth in the remainder of this section.

3.1 Governing equations

The main governing equations of a continuous model are either ordinary differential equations (odes), partial differential equations (pdes), integral equations, integro-differential equations, or some combination of these types. Each type of equation will be discussed in more detail in the following sections. Issues of existence and uniqueness of solutions to the governing equations will not be discussed here, but it should be remembered that such issues may exist. It should be noted that not every model that can be built has a solution, and that non-linear models in particular may not have a solution.

The choice of governing equations can involve approximation. In many models, phenomena that would affect the model results are neglected because their effects on the results are insignificant compared to the effects of other factors, and they add unnecessary complication to the model. A common example of this is the assumption that a process is isothermal: many physical processes that generate heat locally, such as contact involving friction, are regarded as isothermal because the temperature variations that are produced are sufficiently small that they have an insignificant effect on the model results of interest.

In some cases, a governing equation can feature more than one unknown quantity. The most common example is that of determining resonant frequencies. Resonant frequencies occur in many applications, including dynamic analysis of structures, acoustics problems, and electromagnetics problems. They are problems of the form

$$L(\mathbf{u}) = \lambda \mathbf{u}$$

where L is a differential operator, usually a linear differential operator, $\mathbf{u}$ are the unknown eigenfunctions, and λ are the unknown eigenvalues. The eigenvalues and eigenfunctions are generally infinite sets. These problems are often generated by making assumptions about the behaviour of the solutions of other governing equations.

It is becoming more common for continuous modelling software packages to describe their governing equations in physical terms rather than mathematical ones, that is defining the problem and not the model. The majority of finite element packages do not state the pdes being solved in their user manuals, probably because the packages are aimed at engineers and designers from less mathematical backgrounds. Instead there is often a separate theory manual that covers the equation definitions. The posing of problems in physical terms may mean that the user cannot clearly state which pde is of interest, but since the model solution method has already been determined by the choice of software this need not be a problem.

It is worth remembering that alternative solution methods may be available and could be better, so that an awareness of the underlying equations is generally advantageous. In addition, understanding the underlying governing equations can identify when approaches successfully applied in one area of physics can be used in another area.

3.1.1 Ordinary differential equations

Ordinary differential equations are equations that involve a function of a single quantity x and the derivatives of that function. Their generic form is

$$F\left(x, f(x), f'(x), f''(x), ..., f^{(N)}(x)\right) = 0,$$

where F could be a non-linear function of any of its arguments, and the functions f and F could be vectors. This is an Nth order system of odes. It is possible to transform an Nth order ode into a system of N first-order odes by defining

$$f_0(x) = f(x), \quad f_1(x) = f'(x), \quad f_2(x) = f''(x), ..., f_{N-1}(x) = f^{(N-1)}(x),$$

so that the equations can be written

$$\begin{aligned} &F(x, f_0(x), f_1(x), f_2(x), ... f_{N-1}(x), f'_{N-1}(x)) = 0, \\ &f'_0(x) = f_1(x), \\ &f'_1(x) = f_2(x), \\ &\vdots \\ &f'_{N-2}(x) = f_{N-1}(x). \end{aligned}$$

Whilst this transformation offers little advantage when seeking an analytical solution, it means that numerical approximation methods that are only suitable for first-order problems can be applied to the problem, which may be a computational advantage.

The most general form of an Nth order ode has N independent solutions, and if the equation is linear then any linear combination of these solutions will be a solution of the ode. The boundary conditions specify what combination of the possible solutions is required.

3.1.2 Partial differential equations

Partial differential equations occur when the function f depends on more than one quantity, and the equation involves the partial derivatives of f with respect to its arguments. If the function of interest if $f(\mathbf{x})$ and $\mathbf{x} = \{x_i: i = 1, 2, \ldots, p\}$ then the general form for an Nth order pde is

$$F\left(\mathbf{x}, f(\mathbf{x}), \left\{\frac{\partial^j f}{\partial x_1^{i_1} \partial x_2^{i_2} ... \partial x_p^{i_p}} : i_1 + i_2 + ... + i_p = j, j = 1,2,..., N\right\}\right) = 0.$$

Most pdes that occur in metrology are second order or less. Second-order pdes can be written

$$\sum_{i=1}^{p}\sum_{j=1}^{p} C_{ij}\left(\mathbf{x}, f, \left\{\frac{\partial f}{\partial x_k} : k = 1,2,..., p\right\}\right)\frac{\partial^2 f}{\partial x_i \partial x_j} + \sum_{k=1}^{p} D_k(\mathbf{x}, f)\frac{\partial f}{\partial x_k} + E(\mathbf{x})f(\mathbf{x}) = g(\mathbf{x}),$$

where the quantities $\mathbf{x}$ could include time. Second-order pdes can be classified by the value of the determinant of the matrix

$$\mathbf{C} = [C_{ij}]$$

If $\det(\mathbf{C}) > 0$, the equation is elliptic and is probably derived from a physical equilibrium problem. If $\det(\mathbf{C}) < 0$, the problem is hyperbolic, and if $\det(\mathbf{C}) = 0$, the problem is parabolic. Hyperbolic and parabolic problems are generally derived from physical propagation problems. It should be noted that if the coefficients C_{ij} are not constant, the classification of the equation can change. Typical examples of each class of pde are:

Elliptic: the steady-state heat equation $\nabla^2 T = 0$

Hyperbolic: the wave equation $\nabla^2 u = c^2 u_{tt}$

Parabolic: the diffusion equation $\kappa \nabla^2 f = f_t$

3.1.3 Integral and integro-differential equations

Integral equations involve the integral of a field quantity, and are less common in continuous modelling than differential equations. Integro-differential equations are even less common than integral equations and involve both integral and differential terms, sometimes with the differential under the integral sign.

Many of the integral and integro-differential equations that occur in physics can be reformulated as differential equations. For example, if the domain Ω is the exterior to a closed volume V, then the Helmholtz equation (an equation commonly occurring in acoustic models) can be written as

$$\nabla^2 \varphi(\mathbf{x}) + k^2 \varphi(\mathbf{x}) = 0, \qquad \mathbf{x} \in \Omega \tag{3.1}$$

or

$$\int_{\partial V} \left(G_k(\mathbf{x}, \mathbf{x}') \nabla \varphi(\mathbf{x}') - \varphi(\mathbf{x}') \nabla' G_k(\mathbf{x}, \mathbf{x}') \right) . d\mathbf{A}' + \varphi(\mathbf{x}) = 0, \qquad \mathbf{x} \in \Omega, \tag{3.2}$$

where k is the constant wavenumber, φ is the complex-valued pressure, and G_k is the Green's function for equation (3.1). The integral equation form (3.2) is commonly solved using boundary element techniques, whereas the differential form (3.1) is more difficult to solve over an infinite exterior domain Ω.

It is often more useful to use an integral formulation when the domain is extremely large in comparison to the region that is really of interest. In the example above, the domain is the whole of space outside of Ω, but the region of interest to the metrologist may only be a few points within that infinite region. Solution of the differential equation (3.1) would give φ at all points in Ω, and calculation of such a solution would be computationally intensive and much less efficient than using the integral form.

3.2 Model domain

Two questions need to be considered when defining a model domain: which results are of interest, and what affects those results. The domain needs to encompass the locations of the required results, and enough of the rest of the universe to simulate the physical situation accurately.

The most common problem in definition of a model domain is knowing what not to include. The definition needs to be sufficiently detailed that the model results will be close to measured values, but it needs to be sufficiently simple that the model can be solved using a reasonable amount of computational effort. The complexity of a domain can sometimes be reduced by choosing appropriate boundary conditions. Three good ways of controlling the complexity of the domain are reducing the amount of detail, using symmetry or planarity, and limiting the extent of the domain.

Care needs to be taken when simplifying the domain to ensure that the simplified version still produces good results. If possible, model results from models with and without the simplification should be compared in order to assess the effect of the change.

An example of this is shown in figure 3.1. The left-hand model has a rounded inner corner and the right-hand model does not. The contour plots show the maximum principal stress. The maximum stress in the rounded model is 11.4 MPa, and the maximum in the unrounded

sample is 14.2 MPa, an increase of about 25%. Clearly the flattening of the corner, which has simplified the geometry, has affected the results to a significant degree.

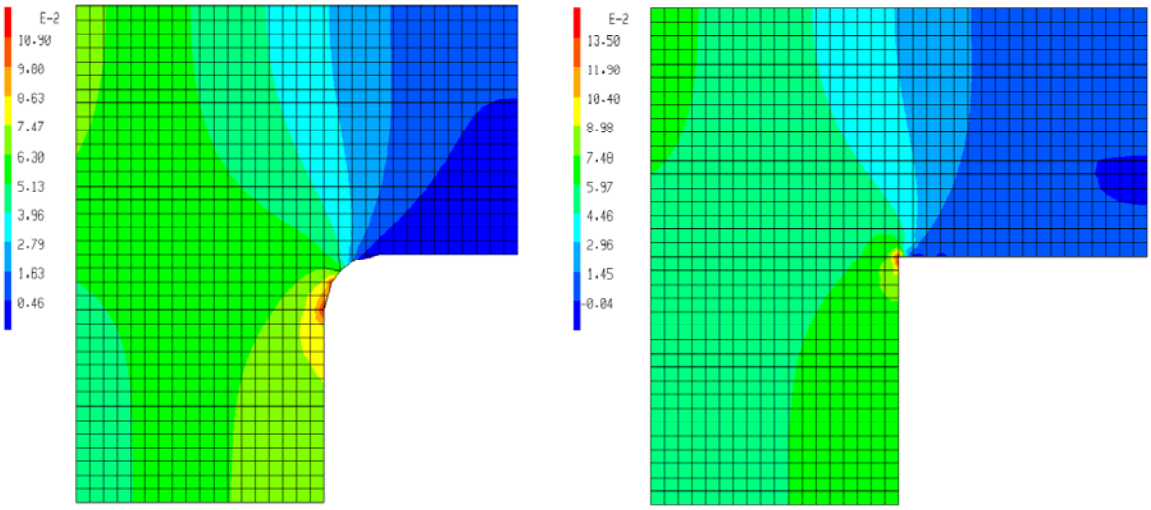

Figure 3.1: Comparison of results from a model with and without a rounded corner.

3.2.1 Reduction of detail

If the model is simulating an experimental procedure, there may be parts of the equipment and the environment of the experiment that do not need to be modelled in detail. Often it is sufficient to model the effect that the equipment has on the process of interest without modelling the details of the equipment accurately.

As an example, consider the equipment shown in figure 3.2. The equipment uses a hemispherical ram to load a circular sample of material on a support. The ram and the support are made of much stiffer material than the sample. The results of interest are the stresses, strains, and deformation of the sample. The main effects of the equipment on the sample are the support holding the plate up, and the ram pushing it down.

Since the ram and the support are much stiffer than the sample material, their deformation can be neglected. Figure 3.3 shows the chosen model domain. The entire model is treated as axisymmetric, and only the material sample is modelled in detail. The use of axisymmetry leads to an altered form of the governing equations that can be solved more quickly. The ram and the support are modelled as rigid surfaces and the rest of the equipment is not modelled at all.

This simplification process is often applied to thermal models. It is often assumed that the air temperature around a hot object does not vary spatially, even though in reality this will not be the case. Similarly, there are often geometric simplifications that can be made. Models of vehicle impacts often represent large complicated parts such as the engine as lumped masses, since their mass affects the deformation of the frame of the car but they do not themselves deform.

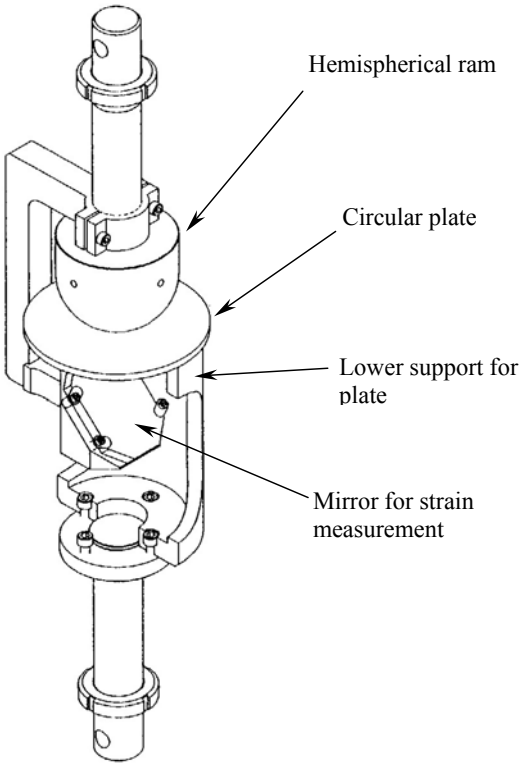

Figure 3.2: Equipment for centrally loading a circular plate.

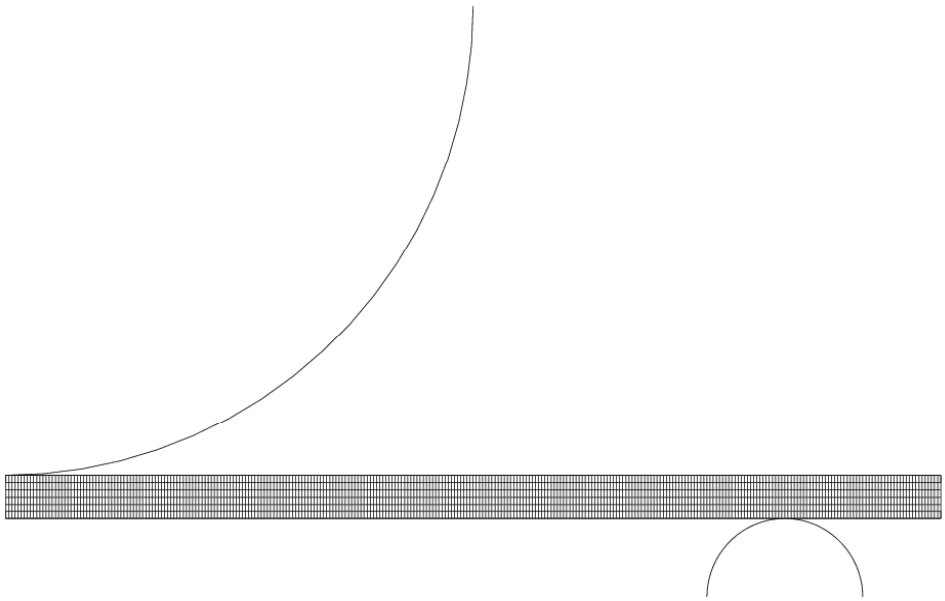

Figure 3.3: Model simulation of the equipment in figure 3.2.

3.2.2 Symmetry and planarity

Symmetry, periodicity and planarity are used to reduce the size and sometimes dimensionality of a domain. Symmetry can be reflective or rotational, and means that only a fraction of the full domain needs to be modelled. Planarity means that the parameters and variables are not expected to vary in one direction, so the problem can be modelled as a two-dimensional slice.

In general, use of symmetry requires the problem to have symmetric boundary conditions as well as a symmetric domain.

Symmetry is imposed through the use of appropriate boundary conditions. For example, consider the shape shown in figure 3.4, and assume that the material is uniform and elastic. The shape is being loaded by equal and opposite forces on the faces at $x = \pm L$ and $y = \pm H$, and the forces in both cases are acting in a direction normal to the loaded face. Since the forces on $x = \pm L$ are equal and opposite and the geometry has reflectional symmetry about $x = 0$, the points along the plane $x = 0$ are not loaded in the x direction, and so they will not move in that direction. Similarly, the points along the plane $y = 0$ are not loaded in the y direction and will not move. These symmetries mean that it is only necessary to model the region $0 \leq x \leq L$, $0 \leq y \leq H$, and the displacements of the rest of the shape can be deduced from the symmetry.

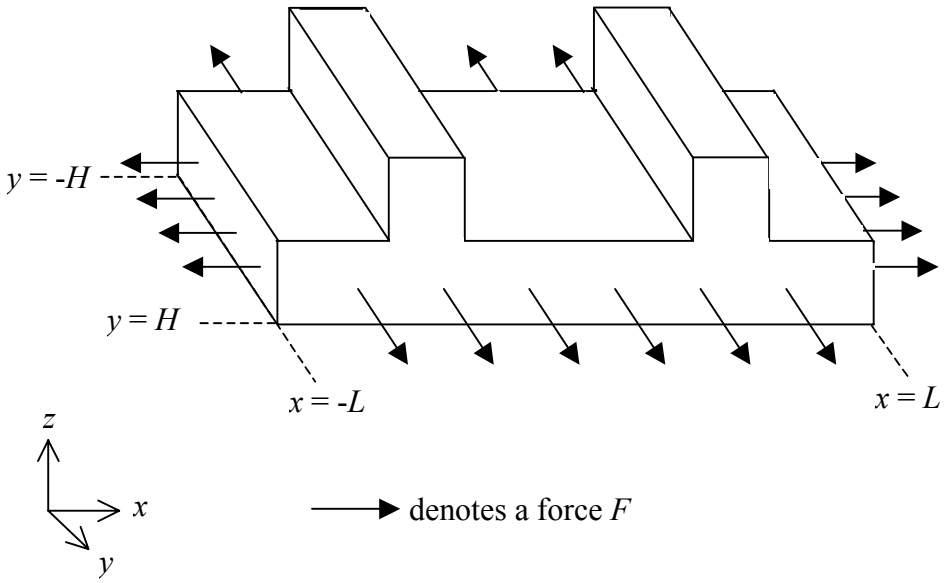

Figure 3.4: An example of a symmetric problem.

Some model domains have periodic symmetry, meaning that they are made up of a smaller unit repeated a number of times. It may be possible to reduce the domain size by only considering a single representative unit. An axisymmetric model is an extreme version of this where the repeated unit is an infinitely thin slice of the full domain. An example of an axisymmetric domain is shown in figures 3.2 and 3.3 above. The boundary conditions and domain are fully rotationally symmetric, so a two-dimensional slice of the geometry can be considered.

It may be possibly to use periodic symmetry by changing the way that the boundary conditions are expressed. An example of this technique is used in one of the case studies (see section 8.2), where rotational symmetry is imposed by expressing the boundary conditions as a Fourier series, leading to a set of problems with various orders of periodic symmetry.

Similarly, planarity could be regarded as an infinitely thin slice in one of the coordinate directions. Planarity can be imposed when the results are not expected to vary significantly over one of the dimensions of a three-dimensional domain. For example, consider a thin, wide beam, fixed at one end and loaded in the vertical direction at the other, as shown in figure 3.5 where L and b are much larger than h.

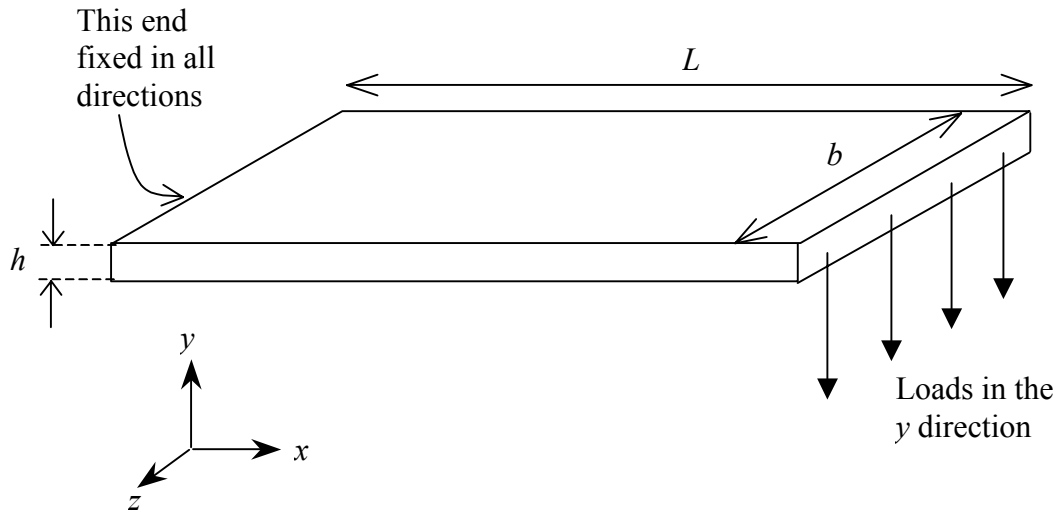

Figure 3.5: Thin wide beam that can be treated as plane strain.

Since there are no forces acting directly in the *z* direction, it is likely that there will be no displacement in the *z* direction and that the *x* and *y* displacements will not depend on *z*. If these assumptions are true, the model can be reduced to a *plane strain* model and treated as a two-dimensional case. Similarly, if it can be shown that the normal stress in the *z* direction is zero (for instance if the beam had been narrow in the *z* direction), then a *plane stress* model can be used. Both of these two assumptions reduce a three-dimensional model to a two-dimensional one.

These assumptions are effectively semi-analytic methods: they make an assumption about the behaviour of the solution and simplify the governing equations and model domain in line with the assumption. Semi-analytic methods are described in more detail in section 10.5.

3.2.3 Limiting the extent of a domain

Some models, particularly models of wave propagation problems, use infinite domains. In general such problems apply a boundary condition at infinity such as steady state conditions or no incoming waves. Such problems can be modelled over a more manageable domain if "infinity" is defined as "sufficiently far away from the areas of rapid change". Even when the boundary condition is to be applied at a finite distance, assumptions about how far away the boundary needs to be can significantly reduce the computational effort required to obtain a good solution.

For example, consider a small object in a large container of gas, such as is shown in figure 3.6. Experiments like this are used to calculate gas density by measuring the effective mass difference between artefacts of differing volumes but the same true mass and surface area. If the large container is hotter at the bottom than at the top, the change in density of the gas with temperature will cause the gas to flow upwards, as shown in figure 3.6. The details of the gas flow around the small object will then affect the subsequent density calculation, and so the metrologist is interested in the flow around the small object. A full model of this situation would require a model of the fluid flow and heat flow including simulation of convection, and would simulate the whole of the large container of gas.

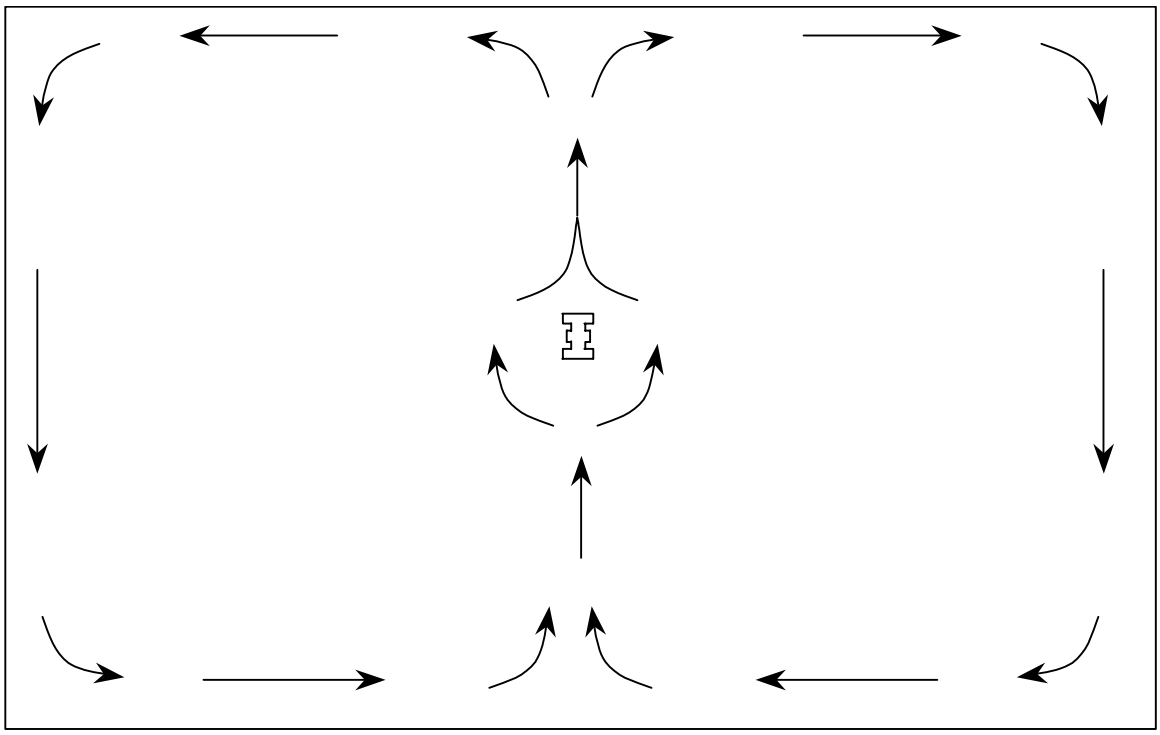

Figure 3.6: Example of the gas flow expected in an experiment to calculate air density.

The fluid flow in most of the container will be uniform and of constant velocity, driven by the temperature difference between the top and bottom of the container. The two regions where the flow will not be uniform are the edges of the container and the area around the small object. Of these, only the area around the object is of interest to the metrologist. Hence instead of using a full model, a smaller region was chosen. The borders of this region were chosen far enough away from the small object and the edges of the container that the uniform flow assumptions were reasonable, but close enough to the small object that the domain size was much smaller than that of the full problem. A typical set of results are shown in figure 3.7. The results nearest to the artificial wall would be ignored as they are not of interest and are unlikely to affect the flow around the object.

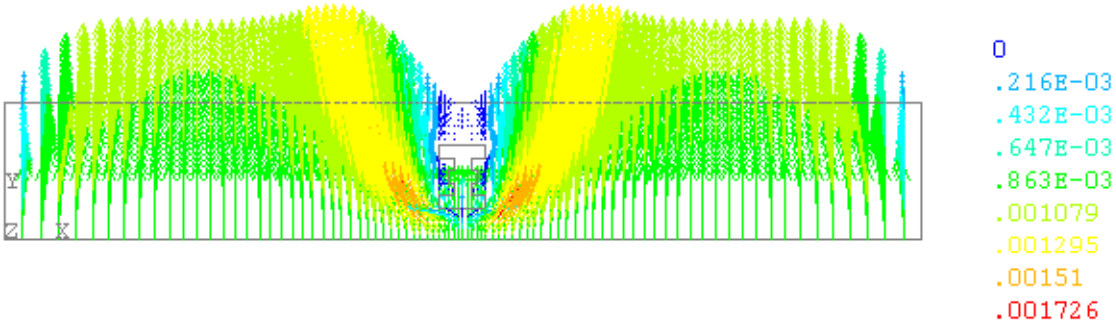

Figure 3.7: A sample set of velocity results from the model on a reduced domain.

For many models (particularly high-frequency electromagnetic and acoustic problems), defining infinity as “sufficiently far away” still leads to an excessively large computational domain. Alternative techniques are needed for the boundaries of such models. In some cases, reformulation of the governing equations can remove the need to consider the boundary at infinity. This technique is commonly used to rewrite the Helmholtz equation as a surface integral equation rather than a pde. Some problems can be limited to a domain lying inside a reflectionless artificial boundary, which can be interpreted numerically in a number of ways [12, 13]. Some techniques, such as infinite elements, are applied during the discretisation

process. Such techniques generally involve making an assumption about the behaviour of the field variable as the position tends towards infinity.

3.3 Domain properties

Domain properties are usually the physical properties of the materials that make up the domain of the model. In many cases, these quantities are the problem parameters, but since the model cannot be solved without them, they are model variables. The most common decision that needs to be made when defining domain properties is whether to regard them as constant or use a more complicated model.

The most common assumption used in modelling domain properties is that the domain being modelled has constant, uniform, isotropic properties. This assumption usually leads to a simpler form of the differential equation and can produce a linear governing equation instead of a non-linear one. However, sometimes constant isotropic properties do not result in accurate predictions, so thought should be given to the alternatives.

Domain properties can depend on a variety of different quantities, depending on the property and the governing equation. Some common examples are:

- dependence on location,
- temperature dependence (most material properties display this to some degree but in some cases it can be neglected),
- pressure dependence (e.g. density of a gas depends on temperature and pressure),
- dependence on strain and strain rate (plasticity properties of materials undergoing deformation are often strain or strain rate dependent)
- frequency dependence (particularly in electromagnetic and acoustic applications)

Some domain properties are written as time-dependent quantities. Care should be taken when using such formulations to ensure that the model uses elapsed time, which requires that a definition of $t = 0$ is stated clearly. If possible, the model should be reformulated so that the properties are dependent only on state variables rather than coordinate variables (see section 2.1 for definitions). For example, one of the most common models used for metal creep is

$$\frac{d\varepsilon_{cr}}{dt} = A\sigma^n t^m,$$

where t is time, ε_{cr} is the creep strain, σ is some property of the stress tensor (usually the second invariant), and A, n, and m are model parameters. This model can be rewritten as

$$\frac{d\varepsilon_{cr}}{dt} = \left\{A\sigma^n\left[(m+1)\varepsilon_{cr}\right]^m\right\}^{1/m+1},$$

which is a preferable form as it defines the creep strain in terms of state variables and so is independent of a frame of reference and coordinate system. McCartney [14] gives a deeper explanation of why this is important.

In some cases, the domain properties may be expressed in a form that makes them difficult to substitute into the governing equations. The properties may be in implicit form or may be differential equations themselves. In such cases it is better to consider the governing equations and the domain properties as forming a system of equations rather than attempting to form a single more complicated equation.

Empirical models are commonly used for modelling domain properties, because models of domain properties are usually derived from measurement rather than theory. Several important points should be remembered when using such models.

The first point is that empirical models should not generally be used outside of the range of the data from which they are derived. For instance, if a model of the dependence of thermal conductivity on temperature is derived from measurements made at 293 K, 323 K, and 373K, it should not be used to predict the thermal conductivity at 673 K unless there is an excellent reason for believing the model to be reliable at that temperature. It is better to use a constant value taken from measurement than the predictions of an empirical model based on data collected at unsuitable points. This point is particularly important if the model contains functions such as negative powers or logarithms that become singular as the variable tends towards zero.

The second point is that it may be possible to simplify an empirical model if the data on which the model is based was collected over a much wider range of variables than would occur in the model. For instance, suppose that a sixth order polynomial is being used to describe the behaviour of some quantity $k(x)$ for $0 \leq x \leq 100$, but that in the model x only varies between 12 and 13. Then it may be possible to use a linear model for k, which would make the model easier to solve.

Finally, if the domain properties are defined as a function of coordinate variables (i.e. position or time), an origin for the coordinates must be clearly stated. Often an origin is specified when the domain is defined, but in some cases it may be more useful to specify a local origin for definition of properties since this can result in a more numerically stable expression.

Anisotropic properties have become more common as the use of laminate and fibre-reinforced materials has become more prevalent. Such materials are usually designed to have specific properties that make them suitable for particular applications, so their properties are often sufficiently well-characterised that a reliable model is available. Similarly, piezoelectric materials often have well-characterised properties that are uniform in a plane but vary in the out-of-plane direction (this is called transversely isotropic behaviour).

Some materials cannot easily have their properties modelled in anything but a statistical sense. Examples of such materials include solids with embedded particles (e.g. rubber-toughened materials), particulate mixtures (such as gravel and concrete), biological tissue (which has too complex a structure to model in detail), and other multi-phase mixtures. These materials are usually characterised using averaged properties, but it is important to remember that the local properties of such materials can be very different from the averaged properties, and so the detailed model predictions may not describe a real situation well.

It should be noted that material properties can often be difficult to obtain, particularly if the material in question is not being used in its unprocessed form. For example, the physical properties of a bulk material, a thin film of that material, and fibres or wires made from that material are likely to vary widely. Similarly, the properties of a thin film usually depend on how the film was created, since the creation method affects the microstructure and the microstructure affects the physical properties.

This difficulty should be taken into account during model building. Care should be taken to obtain material properties that are appropriate to the application, where possible. Sensitivity tests (such as response surface analyses) should be carried out to explore the importance of the relevant properties, if computationally feasible. If neither of these options are possible, the difficulty should be noted and reported as an assumption made during the model building process so that any subsequent decisions based on the model results take the limitation into account.

3.4 Boundary conditions

The boundary conditions of a model identify the unique solution to the particular problem of interest out of all of the possible solutions to the governing equation, given the domain properties. The number of boundary conditions that must be specified for a model to be solvable are related to the order of the differential equation and the number of space and time

dimensions. Boundary conditions are usually generated either from measurement data or from making assumptions about the model.

The most common forms of boundary condition used in pde problems are:

$$u(\mathbf{x},t) = f(\mathbf{x},t) \quad \text{on } \partial\Omega \quad (\text{Dirichlet conditions}),$$
$$\mathbf{n}.\nabla u(\mathbf{x},t) = g(\mathbf{x},t) \quad \text{on } \partial\Omega \quad (\text{Neumann conditions}),$$
$$a(\mathbf{x},t)u(\mathbf{x},t) + b(\mathbf{x},t)\mathbf{n}.\nabla u(\mathbf{x},t) = h(\mathbf{x},t) \quad \text{on } \partial\Omega \quad (\text{Robin conditions}),$$

where $\partial\Omega$ is the boundary of the domain, $\mathbf{n}$ is the normal to $\partial\Omega$, $\mathbf{x}$ denotes position, t is time, a, b, f, g, and h are all functions of position and time, and u is the unknown in the governing equation. Dirichlet and Neumann conditions are special examples of the more general Robin conditions. Another way of writing the Neumann conditions is

$$\frac{\partial u}{\partial n} = g(\mathbf{x},t) \quad \text{on } \partial\Omega,$$

where $\partial u/\partial n$ is the "normal derivative" of u.

Boundary conditions that are derived from measurements are most likely to be Dirichlet conditions, although in some case (such as acoustics) the measured quantity can be a derivative of the unknown of the governing equation. Boundary conditions that are derived from assumptions about the model are commonly either Dirichlet conditions or Neumann conditions. For instance, an assumed "no flow" boundary condition for a thermal problem would be $\partial T / \partial n = 0$, a Neumann condition, and the symmetry assumptions mentioned in section 3.2.2 specified the displacements on the boundaries, a Dirichlet condition.

Some boundary conditions do not fit into either category, such as the friction condition used in contact analyses, which has

$$\mu\|\mathbf{F}.\mathbf{n}\| = \|\mathbf{F} - (\mathbf{F}.\mathbf{n})\mathbf{n}\|$$

where $\mathbf{F}$ is a force, $\mathbf{n}$ is the normal to the surface and μ is the coefficient of friction. Whilst this expression specifies the ratio of the normal and tangential forces at the boundary, it does not specify their values. Another common example is the Stefan-Boltzmann model for radiative heat transfer, which has

$$\lambda \frac{\partial T}{\partial n} = A\varepsilon\sigma\left(T^4 - T_{\text{amb}}^4\right),$$

where T is the temperature of the radiating body in K, T_{amb} is the ambient temperature in K, A is the surface area, ε is the emissivity, σ is the Stefan-Boltzmann constant, and λ is the thermal conductivity.

Sometimes boundary conditions are given in the model that it is difficult or impossible to implement in a numerical formulation. A common example of this phenomenon is the assumption of stress-free boundaries. The normal stress on any surfaces of a body undergoing deformation that are not under direct load or constraint is generally assumed to be zero. Whilst this assumption can lead to useful analytic solutions, it is difficult to implement in most finite element packages, where boundary conditions are prescribed as displacements or forces rather than stresses.

Ode models can be divided into two broad classes according to how their boundary conditions are specified. Suppose the governing equation and domain for an ode of order N are

$$f^{(N)}(x) = F\left(\left\{f^{(j)}(x), j = 0,1,\ldots,N-1\right\}, x\right), \quad a \le x \le b,$$

where f and F could be vectors (as in section 3.1). Then an initial value problem (IVP) specifies all of $f, f^{(1)}, f^{(2)}, \ldots, f^{(N-1)}$ at either a or b, whereas a boundary value problem specifies

some conditions at *a* and some conditions at *b*. For example, if the governing equation and domain were

$$f^{(2)}(x)+f(x)=0,\quad 0\le x\le \pi/2,$$

then an IVP with the solution sin(*x*) could have boundary conditions

$$f(0)=0,\quad f^{(1)}(0)=1,$$

and a BVP with the same solution could have boundary conditions

$$f(0)=0,\quad f(\pi/2)=1.$$

In general, BVPs are more difficult to solve than IVPs. It should be noted that not all BVPs can be solved: in particular, nonlinear systems may not have a solution.

One of the difficulties encountered when defining boundary conditions in more than one dimension is knowing how to define conditions at the point where two different parts of the boundary meet. For example, consider a box filled with fluid with a lid that is moving with constant velocity, as shown in figure 3.8.

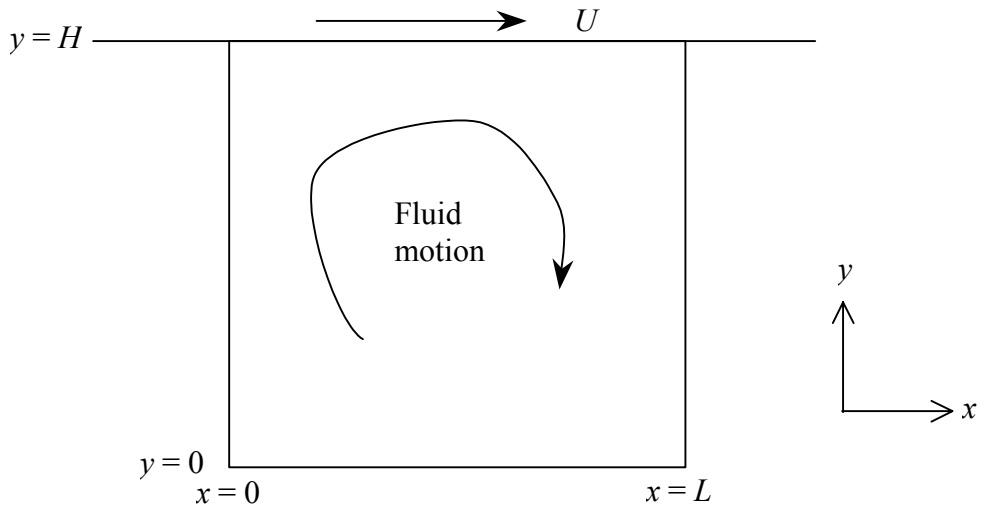

Figure 3.8: A fluid-filled box with a lid moving at constant horizontal velocity *U*.

There is no flow through the walls of the box, so if u is the horizontal component of fluid velocity, $u = 0$ on $x = 0$ and $x = L$, $0 \le y \le H$. The lid is moving with constant velocity, so if the fluid is moving at the same velocity as the wall (a "no slip" condition), $u = U$ on $y = H$, $0 \le x \le L$. This means that the horizontal velocities at $x = 0$, $y = H$ and $x = L$, $y = H$ are ill-defined.

Physical considerations show that the horizontal fluid velocities should be zero at both of these points, or there will be fluid entering and leaving the box as opposed to purely circulatory flow. In general, if there are two possible contradictory boundary conditions that could be applied at a point, physical considerations should be used to choose between them.

3.5 Model parameters and variables

Section 2.2 explained that model variables are the quantities that need to be known for the governing equation to be solvable, and that the parameters are the solution of the governing equation at some set of points in the domain. A good choice of parameters can lead to significant benefits in terms of model stability and conditioning. These issues are discussed further in section 4.3.

The main decision that needs to be made regarding these quantities when building the model of a problem is the choice of units. In general, the choice of units for a model is driven by common sense: there is no point modelling the deformation of a structure whose dimensions are in millimetres using length units of metres. Some choices are slightly less clear-cut: if a transient thermal model is entirely linear (including its domain properties and boundary conditions) then it may be worth using the temperature difference from some average or initial value as the model parameter, particularly if the expected values of the temperature difference are only small.

Care needs to be taken when transforming between sets of units, particularly for empirical models. For example, suppose that the yield stress of a material, σ_Y, is related to the strain rate r through the expression

$$\sigma_Y = A + B\ln(r),$$

where A and B are constants. Suppose that experiments have been carried out to obtain values for A and B that measured σ_Y in MPa and r in s^{-1}. If the model uses units of mm, ms, and g, then the values must be changed, to A^* and B^* say. It can be shown that MPa are the appropriate unit of stress for this choice of base units, so only the change of units for the strain rate needs to be considered. Let the strain rate in ms^{-1} be r^*. Then $r^* = 10^{-3}r$, and so

$$\begin{aligned}\sigma_Y &= A + B\ln(r)\\ &= A + B\ln\left(10^3 r^*\right)\\ &= A + B\ln\left(10^3\right) + B\ln\left(r^*\right)\\ &= A + 3B\ln(10) + B\ln\left(r^*\right)\\ &= A^* + B^*\ln\left(r^*\right),\end{aligned}$$

where $A^* = A + 3B\ln(10)$ and $B^* = B$. This is not an intuitively obvious transformation.

Some models are best defined in terms of a non-Cartesian coordinate system, or in terms of more than one local coordinate system. Polar coordinate systems are useful where the domain geometry has rotational symmetry, but care must be taken with points around the origin because some differential operators become singular as $r \to 0$. Local coordinate directions can be useful for definition of anisotropic materials or for accurate definition of curved domains. A well chosen coordinate system can give a better description of the physical problem and can improve the conditioning of the model.

3.5.1 Variable uncertainties

All measurements have associated uncertainties. In particular, the measurements of model variables will have uncertainties associated with them. These uncertainties will produce corresponding uncertainties associated with the model parameters. This is true of both discrete and continuous models, and the sections in the discrete modelling BPG [4] on error structure and statistical models for measurement error are relevant to model variables for both types of model, so their advice will not be repeated here.

Calculation of the uncertainties associated with the parameters of continuous modelling problems is a complicated subject that will not be addressed in depth here. The main complications are:

- the model parameters and the problem parameters are not always the same (see section 2.2),
- the standard approach as detailed in the Guide to Uncertainty in Measurement (GUM) [15] is often not applicable,

- the computational effort required to solve many continuous models is sufficiently large to make sampling methods such as Monte Carlo simulation unsuitable.

These complications have led to the development of a range of techniques for calculating uncertainties associated with continuous model parameters efficiently. A report was published as part of the second SS*f*M programme [2] that explains most of the main issues and methods, and readers are referred to that report for more detail.

The main conclusions of the report [2] are:

- the GUM approach [15] is suitable for some models, and can be a good way of providing rough estimates from simplified models prior to using more complicated models requiring more advanced methods,
- sampling methods are readily applicable to continuous models, but the computational expense of some models can mean that it is preferable to use methods other than MCS to produce uncertainty estimates,
- some analytical techniques (e.g. stochastic differential equations, response surface methodologies, series expansion methods, etc.) can be useful, but often they require high-level mathematics to implement them,
- some problem-specific methods are available, for instance statistical energy analysis, but these methods cannot always be applied to every type of problem.

The point mentioned above concerning computational expense of sampling methods is gradually becoming less of an obstruction. Sampling methods generally involve calculating several thousand solutions of the model independently using randomly chosen input values, so the time-consuming part of the calculation involves a large number of independent repetitions of the same solution process. This independence of the repeated step makes sampling methods ideal for solution using a distributed computing network, and some work [16] has been carried out using parallel processing for the repetitive runs.

The first case study in section 8.1 of this report uses a response surface and several sampling methods to investigate the sensitivity of the model results to the domain properties. See this section for more detail on sampling methods.

4. Model solving

In the discrete modelling BPG [4], model solving is defined as

"the process of determining from data gathered from a measurement system, a particular model that adequately represents the system behaviour. Constructing the model space [i.e. model building] is concerned with defining where we should look to define the behaviour; model solving is concerned with selecting the best candidate from the options defined by the model space".

The same statement is true of the model solving process as defined here for continuous models. A schematic of the processes is shown in figure 4.1. The model building process creates a general definition of the model. Variable values, known from measurement or other prior knowledge, are used to define a specific model. The specific model is then solved to give parameter values.

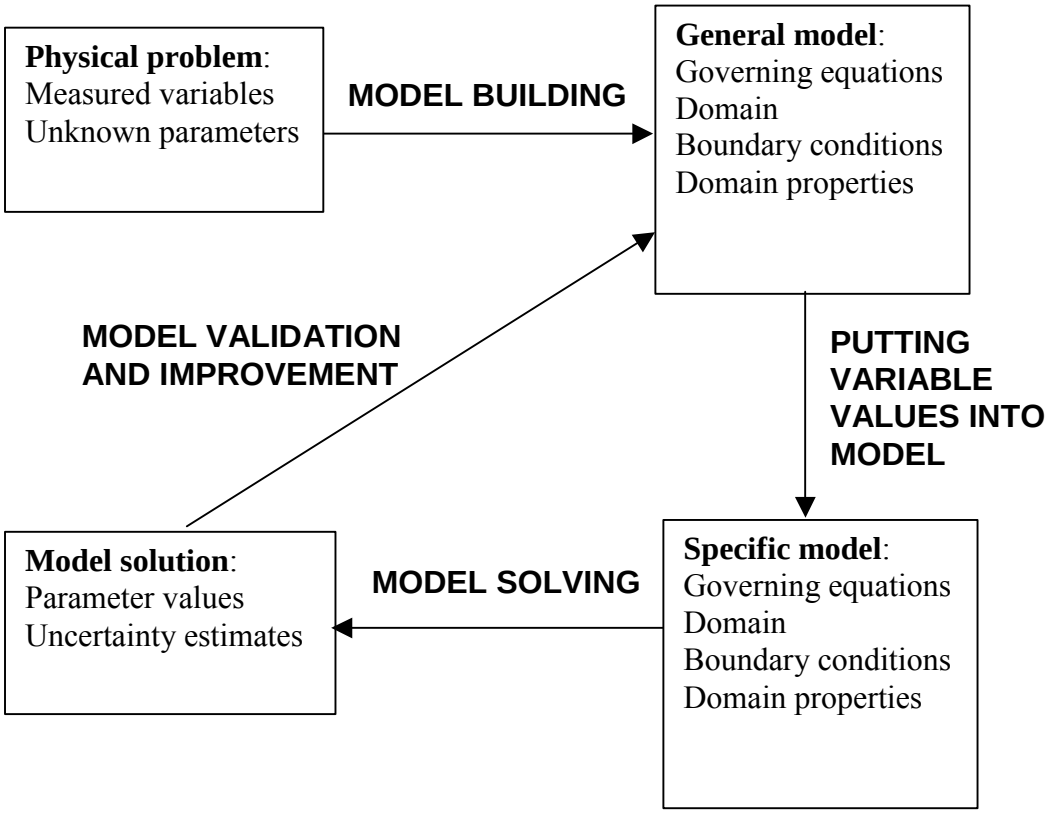

Figure 4.1: Schematic of the model building and solving process.

The most general strategy for solving a continuous model is to split the domain into small regions and make assumptions about the behaviour of the solution within each small region. These assumptions will rely on a set of local parameters, generally the solution values at some set of points within the small region. The subdivision of the domain is usually called the **mesh**, and the points at which the solution is evaluated are called the **mesh points**. If the governing equation is an ode, the distance between two adjacent mesh points is called the **step size**.

When the assumptions are applied to all of the small regions, a set of simultaneous equations whose unknowns are the set of all the local parameters will be produced. This process is called discretisation. This set of equations is generally of the form $\mathbf{Ax} = \mathbf{b}$, or sometimes $\mathbf{Ax} = \lambda\mathbf{x}$, although other forms exist. Often if the equations are non-linear they are locally linearised so that they are in the form $\mathbf{Ax} = \mathbf{b}$. The matrix of equations can then be solved using standard linear algebra techniques.

The most important properties of a good solution method are discussed in section 4.1. The assumptions that underlie some of the most common solution methods used in continuous modelling are described in section 4.2. The methods themselves will be described more fully in section 10.

Semi-analytical solutions are discussed in section 10. These solution methods make an assumption about the dependence of the solution on one or more of the coordinate variables, and derive a simplified form of the governing equation as a result. Such methods can be good way of imposing a known behaviour on a solution.

4.1 Solution method properties

The main property that a good solution method must have is that of **convergence**. Suppose that the governing equation is of the form

$$Du(\mathbf{x},t) = f(\mathbf{x},t), \quad \forall \mathbf{x} \in \Omega, 0 \le t \le t_F,$$
$$Lu(\mathbf{x},t) = g(\mathbf{x},t), \quad \forall \mathbf{x} \in \partial\Omega, 0 \le t \le t_F,$$
$$u(\mathbf{x},0) = u_0(\mathbf{x}),$$

where D and L are differential operators representing the governing equations and the boundary conditions respectively, possibly involving material properties and other parameters, and f, g, and u_0 are known functions. Then a solution method for this problem could produce an approximate solution $U = \{U_j^n : j = 1, 2,\ldots, J, n = 0, 1,\ldots, N\}$ at a set of mesh points $\{\mathbf{x}_j, t_n\}$ by

$$\mathbf{AU} = \mathbf{F},$$
$$\mathbf{BU}_B = \mathbf{G},$$
$$U_j^0 = u_0(\mathbf{x}_j),$$

where $\mathbf{A}$, $\mathbf{B}$, $\mathbf{F}$, and $\mathbf{G}$ are matrices and $\mathbf{U}_B = \{U_j^n : \mathbf{x}_j \text{ lies on } \partial\Omega\}$. This method will be **convergent** if

$$\left\| u(\mathbf{x}_j, t_n) - U_j^n \right\| \to 0, \quad j = 1,2,...,J, n = 0,1,...,N$$

as J and N increase, where $t_N = t_F$, and the $\mathbf{x}_j$ are distinct points in Ω. This property ensures that the calculated solution converges to the true solution as more mesh points are used.

Some valid models contain singularities, for instance models describing cracked materials include a stress singularity at the crack tip. If a singularity is present, it may not be possible to achieve convergence as defined here. It may be possible to consider the convergence of a quantity that avoids the singularity instead. Results in the immediate vicinity of the singularity could be ignored if they are unlikely to be of interest. It may be possible to examine the convergence of a function of the result such as a surface integral, since these quantities are often bounded even though the integrand values are infinite.

Two other properties are useful, and are easier to prove than convergence: **consistency** and **stability**. For well-posed linear problems, a solution method that is consistent and stable will also be convergent.

Stability is a property of solution methods used for time-dependent problems and initial value problems in odes, and it ensures that if a model's variables are perturbed, the resulting perturbations in the model solution will be bounded. Formally, if two problems have initial conditions $\mathbf{V}_0 = \{V^0_j, j = 1, 2, \ldots, J\}$ and $\mathbf{W}_0 = \{W^0_j, j = 1, 2, \ldots, J\}$ then a zero-stable method has

$$\left\| \mathbf{V}^n - \mathbf{W}^n \right\| \le K \left\| \mathbf{V}^0 - \mathbf{W}^0 \right\|, \quad t_n \le t_F,$$

where K is a constant independent of the choice of $\mathbf{x}_j$ and t_n. This is an important property for a solution method because finite precision arithmetic will always cause perturbations during the solution procedure, and if these perturbations can grow without bound then the errors will swamp the true solution. In general, a method is only stable for certain choices of time increment, usually related to the domain properties and the choice of $\mathbf{x}_j$. There are a number of other kinds of stability that define conditions on the properties of a method that will lead to zero-stability, and generally these conditions are easier to prove than proving zero-stability directly.

Consistency ensures that the local error caused by approximating a continuous quantity by a set of discrete values tends towards zero as more and more values are used. A consistent method will have

$$\left(\mathbf{Au}-\mathbf{F}\right)_j^n \to 0$$

as J and N increase, where $\mathbf{u}$ is the matrix of the exact solution evaluated at the points ($\mathbf{x}_j$, t_n). This is equivalent to the property of consistency as defined for estimators in discrete modelling [4].

In general it is not necessary to prove the convergence of a method such as finite element analysis or finite difference schemes, as such methods are well-established and have already been thoroughly analysed numerically. However, there are problems that feature sufficient nonlinearity that non-convergence can cause a problem so it is important to be aware of the possible issues.

4.2 Derivation of methods

This section briefly describes the assumptions that underlie some of the most common methods of continuous modelling. Note that there are circumstances under which the assumptions made can produce the same set of linear equations derived using different methods. The methods are discussed in more detail in section 10, and an alternative explanation of how these assumptions are arrived at is given in section 4.4.

4.2.1 Assumptions of finite element and boundary element methods

Finite element and boundary element methods make an assumption about the behaviour of the model parameters within a small region of the domain. It is assumed that the behaviour of the parameter within the small region can be written as a function (usually a polynomial) written in terms of the values of the parameter at fixed points in the region. In the simplest case, it is assumed that the parameter is constant within each element (this is more common in boundary element methods than finite element methods). For further explanation, see sections 10.1 and 10.2.

The assumption breaks down if the parameter does not vary in the manner assumed over the small region. This is particularly likely in areas of rapid change of parameters, where it is often necessary to use small elements or high-order approximation to describe the behaviour correctly

4.2.2 Assumptions of finite volume methods

Finite volume techniques write the rate of change of some quantity within each small region in terms of the fluxes of the quantity across the boundaries of the small region. This technique assumes that the quantity is constant within each region and that the fluxes are constant within each boundary. The most commonly used quantities are mass, momentum, and energy. For further explanation, see section 10.3.

The assumption breaks down when the quantity or its fluxes vary within the small region. If this occurs, regions can be subdivided but such problems quickly become computationally intensive.

4.2.3 Assumptions of finite difference methods

Finite difference methods are based on the assumption the solution can be expanded as a local Taylor series. This assumption means that the value of a parameter Q at a point $x + \Delta x$ can be written in terms of its value at a point x, so that

$$Q(x+\Delta x) = Q(x) + \Delta x \frac{dQ}{dx} + \frac{\Delta x^2}{2!}\frac{d^2Q}{dx^2} + \frac{\Delta x^3}{3!}\frac{d^3Q}{dx^3} + ...$$

Similar expressions using partial derivatives apply for functions of more than one variable. Expressions of this form are then combined to obtain approximations to the derivatives of the parameter Q. The assumption that the Taylor series is valid means that the higher-order terms in Δx must be sufficiently small, or the series will not converge. The majority of methods used for ordinary differential equations also rely on this assumption. For further explanation, see sections 10.4 and 10.6.

The assumption breaks down when the higher-order derivatives of the function become large, or when the derivatives do not exist. For instance, step-change functions cannot be expanded as a Taylor series in the region where the change occurs since the derivative is infinite there.

4.2.4 Assumptions of semi-analytic methods

Semi-analytic methods usually make assumptions about the dependence of the solution on one of the coordinate quantities, and then use the assumed form to produce a simplified version of the governing equation. Some of the most common assumptions are discussed further in section 10.5.

4.2.5 Time dependence

Time can be treated in the same way as other coordinate variables, so for instance a transient job could be modelled using finite elements in the time dimension as well as the spatial dimensions. However, this is not a common approach since it can lead to unnecessarily computationally intensive problems. The most common approaches to modelling time-dependent problems are semi-analytic methods and finite difference methods.

It is often only necessary to consider periodic solutions to a time-dependent problem, or those that decay to zero as time passes. If that is the case, then a semi-analytic solution based on a sum of the form

$$\sum_{n=0}^{N} e^{ik_n t} f_n(\mathbf{x}), \quad \text{or} \quad \sum_{m=0}^{M} e^{-l_m t} g_m(\mathbf{x})$$

where the k_n and l_m are constants determined from knowledge of the periodicity or decay time constant, or from any time-dependence of the boundary conditions, and f_n and g_m are the functions to be determined. Substitution of these forms into the governing equation and using the independence of the terms in the sum gives a set of N or M new governing equations. The functions f_n and g_m can then be determined using any method suitable for solution of the new governing equations.

The forms are the truncated discrete equivalents of the Fourier and Laplace transforms that can be used to look for analytic solutions. The use of the periodic solution is particularly useful for wave transmission problems, since the frequency for such problems is usually

known. The case study in section 8.2 uses a periodic solution to examine transmission of underwater sound waves.

Finite difference methods are particularly useful for problems with many mesh points that do not require the solution to be known at many intermediate times. Approximation of a second derivative using finite differences requires three values, and approximation of a first derivative requires two. It is unusual for derivatives of above second order with respect to time to appear in physical models, although they do occur in some applications such as nonlinear acoustics. Hence the majority of time-dependent models can be solved using the solutions at the two previous time steps to calculate the solution at the next time step. This means that the solution does not need to be stored at every time step, which is a significant advantage in terms of computer memory. It is common for commercial finite element packages to do all of their time-dependent calculations using finite differences.

Finite difference methods can be divided into two broad classes: **implicit** and **explicit**. Explicit methods express the solution at some mesh point $\mathbf{x}_j$ at a time t_n in terms of the solutions at all of the mesh points at times $0 = t_0, t_1, \ldots, t_{n-1}$, usually via an iterative definition so that only the previous few steps are required to calculate the new value. Implicit methods express the solution at some mesh point at a time t_n in terms of the solutions at all of the mesh points at times $0 = t_0, t_1, \ldots, t_{n-1}, t_n$. These definitions mean that solving a model using an explicit method for the time dependence does not require inversion of a matrix, since all of the unknown parameter values are expressed in terms of parameter values that have already been calculated. In contrast, the use of an implicit method will require a matrix inversion, because the unknown parameter values are inter-dependent. However, the use of an explicit method often imposes stability restrictions on the maximum time step that can be used. This restriction is discussed further in section 4.3.1.

4.3 Accuracy of results

There are two definitions of accuracy that are relevant to continuous modelling. One measures how well the model describes reality, and one measures how well the solution to the discretised model approximates the solution to the problem posed at the end of the model building process. The first type of accuracy will not be addressed here. It is a model validation issue rather than a model solving issue. Model validation is discussed in section 6.

It was stated in section 4.1 that the solution obtained using a convergent method gets closer to the solution of the original differential equation as the number of mesh points used in the approximation increases. In theory, this means that the more points are used, the better the solution will be. In real applications, this is not always the case. The choice of mesh points has an effect on the accuracy of the solution, and in some cases the problem may be such that a method only produces a convergent solution for an unacceptably small range of meshes. These problems are usually connected with stability and occur because the governing equation is a stiff system. Stiffness will be defined and discussed below.

In general, a mesh with a good choice of points has a large number of points in areas where the gradient of the solution is changing rapidly, the elements defined by the points are of more or less the same size, and the elements are close to the “ideal” shape for their type.

The solution at any given mesh point is dependent on the solution at the mesh points to which it is connected. Areas of rapidly changing gradients require a lot of points because the solution gradient is more or less constant between two adjacent points, so for regions of rapid change to be described accurately there must be a lot of points close together.

If some adjacent mesh points are very close together and some are very far apart, the resulting matrix of equations is likely to contain some very large and some very small elements. It is likely that the small matrix elements will either be neglected during the solution procedure, or that they will become swamped by the rounding errors, which can lead to an inaccurate solution. Hence all of the elements in a good mesh should be approximately the same size.

Finite elements are usually required not to deviate too far from canonical shapes such as right-angled triangles, squares, and cubes. Deviating from these recommendations often leads to an ill-conditioned matrix or other stability problems, caused by ill-conditioning of the coordinate transform described in section 10.1. Some methods require the mesh points to be distributed in a particular way. Finite difference methods and finite volume methods both work best when the grid is uniformly spaced in each direction, although multi-grid techniques can be used if extra detail is required in some regions. Finite difference methods and finite volume methods are usually only applied to problems that can be described by evenly-spaced orthogonal grids.

4.3.1 Well-posedness, conditioning and stiffness

Well-posedness is a property of the continuous model. It is the equivalent to stability of a solution method as defined in section 4.1. If a model is well-posed, perturbing its variables will lead to bounded perturbations in its solution. Well-posedness is equivalent to model conditioning as defined in the discrete modelling BPG [4].

It was stated in section 4 that most continuous modelling methods lead to a set of simultaneous linear equations. Hence the numerical accuracy of the solution will depend on the conditioning of this linear algebra problem. This conditioning can be analysed using the methods described in the discrete modelling BPG [4]. In summary, if the matrix associated with the problem has singular values $s_1 \geq s_2 \geq \ldots \geq s_n$, then the larger the value of s_1/s_n, the worse the conditioning of the matrix.

Ill-conditioning of a matrix can be linked to the idea of **stiffness** of a system of odes. Stiffness is a phenomenon of the differential equations rather than of the method used to solve them. It cannot be defined in strict mathematical terms, but its main effect is that a solution method applied to a stiff system is only stable for an excessively restrictive range of step sizes. Excessively restrictive in this case means that the step size is much smaller than the smoothness of the solution being approximated.

As an example, consider the following system of equations and boundary conditions:

$$y_1'(x) = -\frac{k+1}{2} y_1(x) - \frac{k-1}{2} y_2(x), \quad 0 \leq x \leq 10, \quad y_1(0) = C,$$

$$y_2'(x) = -\frac{k-1}{2} y_1(x) - \frac{k+1}{2} y_2(x), \quad 0 \leq x \leq 10, \quad y_2(0) = D,$$

where C and D are arbitrary constants and k is a positive constant with $k > 1$. This problem has the analytic solution

$$y_1(x) = \frac{(C-D)}{2} e^{-x} + \frac{(C+D)}{2} e^{-kx},$$

$$y_2(x) = -\frac{(C-D)}{2} e^{-x} + \frac{(C+D)}{2} e^{-kx},$$

and so for large values of k there will be a short initial period where both exponential terms affect the solution, but as x increases both solutions behave like e^{-x}. The simplest finite difference model for solving a problem like this is

$$y_1^{n+1} = y_1^n - h\left(\frac{k+1}{2} y_1^n + \frac{k-1}{2} y_2^n\right), \quad y_1^0 = C,$$

$$y_2^{n+1} = y_2^n - h\left(\frac{k-1}{2} y_1^n + \frac{k+1}{2} y_2^n\right), \quad y_2^0 = D,$$

where h is the step size, and y_i^n is the approximation to y_i at $x_n = nh$. This method could be written as

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix}^{n+1} = \begin{pmatrix} 1-h\frac{k+1}{2} & -h\frac{k-1}{2} \\ -h\frac{k-1}{2} & 1-h\frac{k+1}{2} \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}^n$$

or $\mathbf{y}^{n+1} = A\mathbf{y}^n$,

and so if the method is applied repeatedly,

$$\mathbf{y}^{n+1} = A^{n+1}\mathbf{y}^0.$$

A can be written as $A = Q^{-1}SQ$, where Q is an orthogonal matrix and S is diagonal, so that $A^n = Q^{-1}S^nQ$. For the matrix given above,

$$Q = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}, \quad S = \begin{bmatrix} 1-hk & 0 \\ 0 & 1-h \end{bmatrix},$$

and so the boundedness of A^n as n increases depends on the size of $1 - hk$ and $1 - h$. Since $k > 1$, the first of these is the greater restriction. Stability requires that $|1 - hk| < 1$, so the restriction on the step size h is

$$0 < h < 2/k,$$

which is a severe restriction for large values of k, particularly for large values of x where the e^{-kx} term no longer has any significant effect on the solution.

Going back to the original ode system, it can be written as

$$\begin{pmatrix} y_1' \\ y_2' \end{pmatrix} = \frac{1}{2}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} -1 & 0 \\ 0 & -k \end{pmatrix}\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}\begin{pmatrix} y_1 \\ y_2 \end{pmatrix},$$

a singular value decomposition. The key feature of a decomposition like this is the ratio of the largest singular value to the smallest, which in this case has the value k. As with the definition of ill-conditioning of a matrix, the larger this ratio is, the worse the stiffness problem is. If all of the eigenvalues of the matrix defining a linear system of odes have negative real parts and the ratio of the largest value to the smallest is large, the system will be stiff.

This is the simplest possible example of a stiff system. A better choice of solution method would have led to a less restricted choice of step sizes, but there are equations for which most explicit methods impose strong restrictions on the step size. A more thorough discussion of such systems, including stiffness in nonlinear systems, can be found in "*Numerical Methods for Ordinary Differential Equations*" [17].

4.4 Discretisation methods and estimators

The key concept in solution of discrete models, defined in the discrete modelling BPG [4], is that of the estimator. An estimator is defined as a scheme for producing an estimate of the model parameters from the model variables. In general, estimators are written as the minimisation of an error function that gives a measure of how well the model results match the measured behaviour. Different estimators use different error functions.

The key concept in solution of continuous models is the creation of a discretised system of equations. This creation process is, in some ways, very similar to the choice and application of an estimator. Suppose that the model is of the form

$$A(\mathbf{u}(\mathbf{x})) = 0, \quad \mathbf{x} \in \Omega,$$
$$B(\mathbf{u}(\mathbf{x})) = 0, \quad \mathbf{x} \in \partial\Omega,$$

where A and B are differential operators, Ω is a closed domain with border $\partial\Omega$, $\mathbf{u}$ is the unknown quantity, and $\mathbf{x}$ are the coordinate variables. Then if $\mathbf{v}(\mathbf{x})$ and $\mathbf{v}_B(\mathbf{x})$are any functions defined on Ω and $\partial\Omega$ respectively, it must be true that

$$\int_{\Omega} \mathbf{v}(\mathbf{x}) A(\mathbf{u}(\mathbf{x})) dV + \int_{\partial\Omega} \mathbf{v}_B(\mathbf{x}) B(\mathbf{u}(\mathbf{x})) dS = 0, \tag{4.1}$$

because all of the terms in the integrals must be zero. The discretisation process writes

$$\mathbf{u}(\mathbf{x}) \approx \mathbf{U}(\mathbf{x}) = \sum_{j=1}^{N} \mathbf{a}_j \varphi_j(\mathbf{x}), \tag{4.2}$$

where the $\mathbf{a}_j$ are parameters and the φ_j are basis functions, usually defined over a subdivision of the domain. This approximation means that the equation (4.1) above can no longer be true for all functions $\mathbf{v}$ and $\mathbf{v}_B$. Hence in order to seek an estimate of the parameters $\mathbf{a}_j$, a version of (4.1) where $\mathbf{v}$ and $\mathbf{v}_B$ are restricted to a chosen class of functions is developed. Writing

$$\mathbf{v}(\mathbf{x}) \approx \sum_{j=1}^{N} \mathbf{b}_j w_j(\mathbf{x}), \quad \mathbf{v}_B(\mathbf{x}) \approx \sum_{j=1}^{N} \mathbf{b}_j w_j^B(\mathbf{x}) \tag{4.3}$$

where the $\mathbf{b}_j$ are arbitrary parameters and the w_j and w_j^B are weighting functions, and substituting (4.2) and (4.3) into (4.1),

$$\sum_{j=1}^{N} \mathbf{b}_j \left\{ \int_{\Omega} w_j(\mathbf{x}) A(\mathbf{U}(\mathbf{x})) dV + \int_{\partial\Omega} w_j^B(\mathbf{x}) B(\mathbf{U}(\mathbf{x})) dS \right\} = 0.$$

Since the $\mathbf{b}_j$ are arbitrary constants, the bracketed term must be zero for each j, so

$$\int_{\Omega} w_j(\mathbf{x}) A(\mathbf{U}(\mathbf{x})) dV + \int_{\partial\Omega} w_j^B(\mathbf{x}) B(\mathbf{U}(\mathbf{x})) dS = 0, \quad j = 1,2,...,N, \tag{4.4}$$

which is a set of N equations that can be used to determine the parameters $\mathbf{a}_j$ in (4.2) above. It is common, but not necessary, to choose the w_j and w^B_j to be the same set of functions. Different discretisation methods use different choices of basis and weighting functions to generate different systems of equations for the same problem. The choice of functions is equivalent to the assumptions made about the solution behaviour in the small mesh subdivisions as described in section 4.2.

Zienkiewicz and Taylor [18] point out that since $A(\mathbf{U})$ is the residual error produced by using an approximation to $\mathbf{u}$ in the governing equation and $B(\mathbf{U})$ is the residual error produced by using an approximation to $\mathbf{u}$ in the boundary conditions, the expression (4.4) is a weighted sum of residuals, a measure of how well the model matches the true solution.

The process of solving the model is the process of choosing a set of basis functions and weighting functions, and minimising the weighted sum of (4.4) to generate a set of equations that can be solved for the unknown model parameters. This process is analogous to the choice and minimisation of an error function that generates the equations to be solved during the solution of discrete models.

5. Software implementation

This section discusses the issues surrounding software implementation of continuous mathematical models. The main choice that must be made is whether to use a pre-written package for the modelling process, or to write a piece of software specifically for the problem. The factors affecting this choice are discussed in section 5.1 and 5.2. Section 5.3 discusses some points of relevance to the application of both types of software to metrology problems.

5.1 Generic software

Generic software falls into two broad classes: software that solves problems defined in physical terms, and software that solves problems defined in mathematical terms.

The majority of commercial finite element packages define problems in physical terms. Such packages enable the user to define the domain properties and boundary conditions in physical terms, and do not require the user to know the governing equations that are being solved. Often these packages also include pre- and post-processing routines that help the user to define the problem and the mesh and explore the model results via a graphical interface.

The main advantage of packages using physically-defined problems is their usability, particularly for people from a non-mathematical background. In general the full documentation supplied with such packages includes a theory manual giving a mathematical definition of the governing equations and the models used to define the boundary conditions and domain properties, but an understanding of these definitions is not always necessary for use of the software. However, the lack of this understanding can make the identification and improvement of inaccurate results difficult. An understanding of the mathematical definition of a model usually leads to an understanding of the most likely causes of inaccuracy and instability.

Packages that define problems mathematically tend to be part of a larger more general package such as the NAG library [8] or Matlab [9], or standalone routines without a graphical interface such as those listed in Netlib [19], although this is not always the case. For example, the finite element package COMSOL Multiphysics (formerly Femlab) [20] allows the user to express the problem in physical or mathematical terms, including a weak formulation option, but has a graphical user interface and an automatic mesh generator.

Software that defines problems in mathematical terms is often considerably more flexible than that defined in physical terms. Equations can be defined more generally in mathematical terms than physical ones. This may be of limited benefit in many cases, since the majority of problems in metrology can be defined in physical terms, but there are exceptions. In some cases, simplification or transformation of a common physical model can lead to a new differential equation that requires software defined in mathematical terms to solve it. Similarly, if a new model has been developed for a domain property it may only be possible to solve using software that defines problems mathematically.

Some software packages that define problems physically can accept user-supplied software routines that define domain properties or boundary conditions. For instance, the Abaqus finite element package [21] has an option for user-defined materials that requires the user to supply a Fortran routine to calculate the stress increment for a given strain increment, enabling new material models to be created. Such options enhance the flexibility of the software, but they often require a high level of technical expertise.

The benefits of using generic software are that the software is usually robust, easy to use, and it generally uses best-practice algorithms for solution of large matrix problems. In addition, the support services for such software are often excellent, including user groups as well as

support by the software suppliers, and in most cases extensive testing of the software has been carried out.

The main disadvantages of generic software are that the software is usually a black box, so no detailed knowledge of the algorithm is available, and in general few details of testing procedures are available, so fitness for purpose cannot be guaranteed. This issue is discussed further in section 5.3. The test and example problems usually supplied with software can give some idea of the range of problems it is designed to solve, but sometimes these problems are too simple to show fitness for purpose definitively. Another disadvantage is that some properties of the results that are useful to the mathematician more than the engineer, such as derivatives and Jacobians, are not easy to obtain.

5.2 Custom-made software

The development of high quality multi-purpose software for solution of continuous models is a time-consuming and difficult process, and in most cases generic software is the better choice. However, custom-made software is very useful for three types of problem:

- problems with a simple solution procedure that are to be run repeatedly,
- complicated problems that are beyond the scope of generic software, and
- problems requiring more than parameter estimates from the software.

Problems involving multiple solutions using a simple procedure can often be easier to automate if they are solved using custom-made software. Common examples include Monte Carlo simulations, variable sensitivity analyses, and optimisation problems such as those produced when model parameters and problem parameters are not the same. Custom-writing the software generally means that the creation of inputs for each model run and the accumulation of the desired results are more straightforward to carry out than the equivalent tasks using generic software.

Some problems are beyond the scope of generic software and hence custom-made software is the only possible choice. Even though generic software exists that can solve a wide variety of problems, there are still some problems that need custom-made software. Such problems commonly occur when semi-analytic methods are used. It is often more simple to create a single piece of software to take in the variables and derive and solve the simplified version of the model from the assumptions of the semi-analytic method, than it is to derive the simplified model and its variable values by hand, enter them into some generic software for solution, and post-process the results to obtain the model parameter values.

If more information than just parameter estimates is required by the user, custom-written software may be the best choice. Custom-written software is more likely to have source code access and so it is more straightforward to obtain other results such as error estimates, condition numbers, or sensitivity coefficients. Similarly, if the user wishes to have high degree of control over the solution procedure (for instance if a solution of known accuracy is required), then custom-written software may be preferable.

Custom-written software can benefit greatly from library software, particularly for the solution of large matrix problems. The solution of the matrix of equations is probably the step in the solution where choice of algorithm is most crucial for obtaining an accurate solution. Many library routines exist that provide good solutions and estimates of condition numbers, and it is advisable to use them where possible.

The main disadvantage of custom-written software is the time it can take to write and test it. In some cases this time is repaid when the software is used, but it is always worth considering whether generic software is available that could do the job. Another possible disadvantage is the lack of pre- and post-processing software. Packages such as Matlab, Excel, and IRIS

Explorer can be used to generate input variables and post-process results if the input and output file formats are chosen carefully, so this need not be a problem.

5.3 Issues common to both types of software

The most important issue affecting good practice in continuous modelling software implementation is testing. All the software used in model development, model solution, and interpretation of the model results needs to be tested as far as possible, including any visualisation software. Guidance has been developed during previous SS*f*M programmes on good practice for testing software [22, 23]. In particular, a methodology has been developed that tests software for fitness of purpose in an objective manner. The methodology relies on treating software as a black box and using reference results and a well-chosen metric to express the accuracy of the test results.

Pre- and post-processing software is particularly important for continuous modelling software. An accurate solution to a continuous model requires a good choice of mesh, and so a good pre-processing package is very important. A typical continuous model produces results at a large number of points in the domain, and these results can be difficult to interpret without the visualisations produced by post-processing software.

5.3.1 Pre- and post- processing software

Many commercial packages include automatic mesh generators that will create a mesh of triangular or tetrahedral elements within a domain. Whilst these are extremely useful tools, particularly for the novice user, the quality of the mesh is strongly dependent on the quality of the domain geometry defined by the user. If the domain contains many small, detailed features the mesh may contain very small elements, resulting in an unnecessarily large matrix. An automatically-generated mesh should always be checked visually to ensure that there are no obvious flaws. Some packages provide mesh statistics or a measure of mesh quality which should also be checked. Additionally, it is worth checking that the mesh is most detailed in the areas where the parameters are changing most rapidly, for instance in areas likely to experience large deformations or high temperature gradients.

If automatic mesh generation is not available, it can be useful to sketch the domain as a collection of simple geometric objects, then use tools such as Excel and Matlab to produce a mesh by subdividing these shapes. These tools enable the user to create mesh points and visualise them in two dimensions to check the mesh quality. Alternatively, the nodes could be generated using a custom-made program then visualised in one of these packages.

Post-processing software is often specific to a solution package: many software packages produce output files that can only be read by their own proprietary post-processing software. This can limit the choice of packages available, although some software houses release their file formats to producers of visualisation software so that users are not so limited.

It is important to remember that visualisation packages are as susceptible to error as every other kind of software. The high quality of the graphical output of post-processing packages can lead users to forget that it is important to examine the numerical values the software has produced without them being processed further.

A good practice guide to visualisation has been produced as part of a previous SS*f*M programme [24], containing advice on how to visualise a wide variety of quantities and including case studies and examples. Another report [25] has dealt specifically with visualisation and continuous modelling and has identified a number of issues that are relevant to interpretation of results that are visualised in post-processing software.

6. Model validation

As stated in section 1.2, model validation in continuous modelling has been covered in an earlier SS*f*M report [1]. This section describes some of the main conclusions in the report.

The report advocated a three-stage approach to the validation of continuous models, answering a question at each stage:

- Is the model description adequate and mathematically correct?
- Is the discretisation method adequate?
- How do the results compare with other sources of information, including experimental results?

The answers to these questions are produced by carrying out (respectively) internal, inter-model, and external consistency checks.

In order to validate as many aspects of the model as possible, many different validation methods should be used. The report described a number of methods for carrying out each type of check, illustrated with examples. The methods ranged from a visual inspection of the results, using the experimentalist's expert knowledge to validate the model, to more mathematically involved methods that supply error bounds for results of finite element analyses.

Advice was also given about general validation strategy. The key recommendations were:

- The choice of validation methods depends strongly on the aim of the modelling process. Is one concerned simply to gain a qualitative understanding of the phenomenon of interest, or are the results to be used to quantify sources of uncertainty in a measurement process? In the second case the aim may be to develop the most comprehensive descriptive model possible, which implies an increased level of validation compared with the first case.
- Validation strategies should be part of the model development process, an aspect that is particularly important when a complex model is being constructed from simpler models or sub-models.
- Validation by comparison with experimental measurements can be problematic. Measurements may be available only for a limited part of the domain being modelled, and measurements may be unavailable or impossible to obtain in key parts of the model.
- Non-unique solutions may exist. For example, the physical problem may have a unique solution, but the mathematical formulation of the problem may have multiple solutions. Similarly, multiple physical solutions may co-exist. The most common case is hysteresis, where the response of a system to its load depends on its previous loading history as well as its current state. Although experimental repeatability may indicate that a unique stable solution exists, it is still possible that more than one stable solution can exist, but that the existence of other solutions depends on an aspect of the experiment that may be too subtle to model.

Validation for industrial applications

An industrial seminar held in 2004 on "FEA versus test: comparison of modern-day computer simulation with real-life testing" featured various case studies where finite element analysis was used to simulate various problems that had occurred in industrial applications. A number of observations made at this meeting were relevant to model validation, particularly regarding the validation standards applied to industrial models.

The majority of applications that were described were problems in the sense of "situations that produce an undesirable effect that needs to be eliminated", with the undesirable effect often being a qualitative event such as "the sheet which should be flat to within a visual inspection was bent" rather than a measured event such as "the sheet was deflected by 12.1 cm". This meant that the aim of the model was to predict the occurrence of the undesirable effect, and that the magnitude of the effect was less important. Modified versions of the model were subsequently used to explore ways of removing the effect, with the most successful design being implemented as the physical solution when possible.

The majority of the models discussed at the seminar were validated by visual inspection of the results or comparison of actual and predicted locations of damage or failure. The majority of applications that were described involved physical situations in which measurement and data collection were difficult if not impossible. In many cases, this is why mathematical modelling is used in the first place: the problem could not be understood by taking measurements and analysing data, and so a model was used instead. Where measurements were available, it was unusual for more than one set of test results to exist due to the expense of testing and data collection. This was particularly true of any tests that involved testing until failure.

These observations illustrate that it is often only possible to validate a model to a fairly basic level, but this may be acceptable if the model is only required for qualitative results. The drawback of this low level of validation is that it reduces confidence in the predictions of altered models used to explore solutions. In order to compensate, the models used often involve a high level of detail to improve the modeller's confidence that the likely causes of the problem have been included. Often engineers can combine their experience with model predictions to improve their confidence in a designed solution.

In general, it would seem that validation of models used to solve industrial problems is restricted by the nature of the situation being modelled, the paucity of measurement data, and the importance of obtaining a viable design solution rapidly. Metrology applications often have more data available for validation, but in cases where the model is simulating some form of destructive test, sparse data may be a problem.

7. Summary of good practice

This section summarises the key points that have been described in the preceding chapters. It provides a summary of good practice in continuous modelling. Further illustration of the points made here is given in the case studies that form section 8.

All models start from a questions that cannot be answered by measurement alone. In order to clarify the purpose of the model, a "user specification" should be created that includes:

- Definition of the question to be answered.
- Definition of the model results that will be used to answer that question.
- Definition of who will use the model.
- A description of the assumptions made about the experiment.
- A description of how the model will be validated.
- A list of the information that the experimentalist will supply to the modeller.

Problem parameters are the quantities for which values are required. **Problem variables** are the quantities that are known or measured. **Model parameters** are the quantities for which the model determines values. **Model variables** are the quantities that are required to uniquely define the model. If a problem parameter is a model variable, then it will be necessary to use optimisation techniques to determine the problem parameter value. See section 2 for more detail.

Modelling is an iterative, evolutionary process. Models should start simply and increase in complexity. At each stage of its development the model needs to be validated to ensure that the latest version of the model is correct.

Model building in continuous modelling consists of the definition of

- a governing equation or system of equations, usually in differential or integral form,
- a problem domain, which is the region of space and time over which the problem is to be solved,
- domain properties, generally given as material properties, and
- boundary conditions on the limits of the domain, which may include initial conditions if the model is time-dependent.

The choice of these entities affects how well the model describes reality and how rapidly results can be obtained. A good model will describe reality sufficiently well to produce accurate results, but will not include unnecessary detail. In general the user has little choice over the governing equation, but the other entities often have more than one choice available. See section 3 for more on model building.

Good practice in model building

Start with the simplest model possible. To simplify a model:

- simplify the domain as much as possible by
  - omitting unnecessary physical details,
  - limiting the extent of the domain to areas that affect the results of interest,
  - using symmetry, axisymmetry, and periodicity to reduce the size of the problem,
- use constant, uniform, isotropic domain properties,

- linearise all boundary conditions, and
- choose units so that models are in as simple and numerically stable a form as possible.

When adding complexity to a model, only add one new feature at a time and validate after each change.

Model solving is the process of obtaining values for the model parameters given the model definition and values of the model variables. The majority of model solving techniques for continuous models subdivide the model domain into smaller regions (called a mesh) and use assumptions about the local behaviour of the solution to derive a set of discrete equations (called a discretisation process). It can be shown that the derivation of the discrete equations is equivalent to an error minimisation process over all solutions exhibiting the chosen local behaviour. See section 4 for more on model solving.

Model solving techniques must be **convergent** for the model equation so that the approximate solution will converge to the "true" solution as the number of model parameters is increased in a controlled way. The accuracy of the solution of a model is determined in part by the conditioning and stability of the discretised version of the model. Both qualities ensure that small perturbations in the input quantities produce bounded perturbations in the results, so that use of finite precision arithmetic will not render the results useless and the technique will be convergent.

Good practice in model solving

In the vast majority of cases, only convergent methods should be used. Exceptions may have to be made if the model is known to include a singularity and no alternative formulation for the problem exists.

A good mesh should have sufficient mesh points to describe the solution in all areas, well-shaped mesh elements, mesh elements whose areas lie within a reasonable range of sizes, and sufficiently few mesh points that the problem does not become too computationally expensive.

The mesh chosen for the model should be validated, preferably by comparing results to a reference solution or by checking the convergence properties of results generated using different meshes.

If a problem is ill-conditioned, it may be possible to find an alternative formulation that improves the conditioning, for instance by considering an alternative coordinate system, another choice of units, or a simplified model.

Software implementation takes the discretised version of the problem and solves it. There is a wide range of generic packages for solving continuous modelling problems, with finite element analysis being probably the most popular technique. However, there is a lot to be gained from using custom-written software, particularly for complicated models or models that are to be run repeatedly. See section 5 for more on software implementation issues.

Good practice in software implementation

All software, including pre- and post-processing software, should be tested. Whilst it is reasonable to expect that commercially available software has been tested, there are usually few details of testing procedures available to software users.

Generic software is often well-tested and easy to use, but may not be sufficiently flexible to use for problems with unusual features.

Custom-written software is usually designed with a particular problem in mind. This makes the software less flexible in terms of applications, but means that the algorithms used in the software can be chosen to take advantages of special features of the problem such as special

matrix structure. Such software can benefit greatly from use of library routines such as those for solution of large matrix problems.

Model validation checks that the model solution is consistent with the model definition, and that the model solution is consistent with all measurement data available and all other knowledge of the behaviour of the physical system being modelled. Validation methods range from visual inspection of the model results to more complicated a posteriori error estimation methods. Good practice in model validation for continuous models is described in detail in an earlier SS*f*M report [1]. See section 6 for more on model validation.

8. Case Studies

These case studies have been chosen to illustrate good practice in continuous modelling. Each one includes model building, model solving, and model validation stages, and at each stage the decisions that shaped the model are explained. In both cases, the initial studies have raised new questions and have indicated possible new areas for further development.

As well as indicating good practice, these case studies illustrate the remarks in the introduction to this report regarding mathematical modelling. It is a simplifying process involving assumptions, compromises, and approximations. As the model is developed, the effects of these assumptions can be assessed and accounted for.

8.1 Case Study 1: High-temperature thermometry

8.1.1 The physical problem

Radiation thermometers are calibrated at high temperatures (typically 1100 to 2500 °C) using high temperature furnaces. A new method that could lead to up to a factor 10 improvement in scale reduction uncertainty is being investigated by NPL and others. In these new fixed points, a metal-carbon eutectic is placed in a cylindrical graphite crucible with an internal graphite black body, as shown in figure 8.1, the crucible is placed in a furnace and the temperature is raised to near the melting point of the metal-carbon eutectic. Once the temperature has stabilised, the furnace is stepped above the melting point of the metal-carbon eutectic, and a radiation thermometer is used to measure the temperature of the bottom of the black body (marked in red in figure 8.1). The thermometer would then be calibrated under the assumption that the temperature on the measured surface is the same as the melting point of the metal-carbon eutectic.

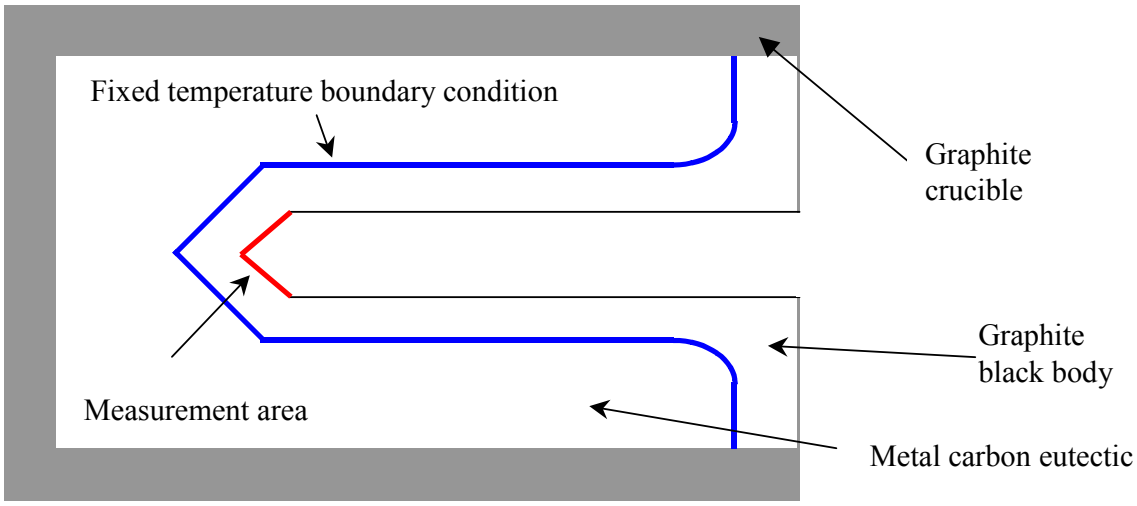

Figure 8.1: Cross-section of the graphite crucible and black body. The full shape is a cylinder.

Whilst this assumption is reasonable for low accuracy measurements, if thermometers are to be calibrated to less than 1 K uncertainty then the assumption needs to be tested. Testing this assumption experimentally would require an alternative method of measuring the surface temperature of the black body, but at high temperatures no practical alternative to radiation thermometry exists. Hence the assumption needs to be tested using a model. There are a number of other questions about the experiment that the thermometrist would like to explore, and these can also be addressed using a model.

The questions to be asked are:

1. What is the average temperature difference between the inner surface of the black body and the fixed point temperature when the metal-carbon eutectic has melted and the system is in a steady state?
2. How is this value affected by the material properties (thermal conductivity and emissivity) of the graphite that the crucible and black body are made of?
3. How does the uniformity of the furnace affect the temperature difference at the end of the melting process?
4. How do all of these factors affect the melting process?

There are six different fixed points, corresponding to six different metal-carbon eutectics, that are of interest to the thermometrist, and the questions above are to be investigated for each fixed point. The process of developing a model to answer the fourth question will be described, but as yet the work has not been carried out so no results will be reported. The long-term aim of this work is to assist the thermometrist in designing equipment for this experiment.

8.1.2 Mathematical formulation

8.1.2.1 Governing equations

The model required is a high-temperature thermal model. The most general form of the governing equation of such a model is

$$\frac{\partial}{\partial t}\left(\rho c_p T\right) = \nabla.\left(\lambda \nabla T\right) + Q, \quad \mathbf{x} \in \Omega,$$

where $T(\mathbf{x}, t)$ is the temperature, $Q(\mathbf{x}, t)$ is the (known) power per unit volume generated within the body, and all of the material properties ρ (density), λ (thermal conductivity), and c_p (specific heat capacity) could all be functions of time, temperature, and position.

The different questions that are to be addressed lead to different simplifications of this equation. In all cases it is assumed that the material properties are independent of position within a given material, since there is no reason to assume otherwise.

The first three questions apply to the system in a steady-state condition, so all time dependence in the equation above can be ignored. Additionally, the source term Q can be neglected, since no internal heat sources or sinks exist.

It can be assumed that the material properties are independent of temperature within each model, but that they will vary for models of different temperatures. There is no reason to assume that the emissivity will not be constant, and whilst the conductivity of graphite is temperature dependent, the relationship is such that a change in temperature of 1 K will lead to a change in the thermal conductivity value of less than 0.05%. Such a change will not affect the temperature results of a steady-state model to any significant degree.

These simplifications mean that the first three questions can be answered using the governing equation

$$\nabla^2 T = 0, \quad \mathbf{x} \in \Omega.$$

It is important to note that the results will still depend on any of the material properties that appear in the boundary conditions, so for instance the thermal conductivity still affects the results even though it does not appear in the governing equation.

The last question requires a time-dependent model to answer in full, since the melting process is to be described. Simulation of a melting process requires the latent heat of melting to be taken into account. This leads to a modification of the time-dependent term in the heat

equation above. In its ideal form, this modification assumes that all melting takes place instantaneously at a single temperature, so that the energy used in melting, per unit volume, is $\rho L \delta(T - T_M)$ where T_M is the melting point. Hence the heat equation becomes

$$\frac{\partial}{\partial t}\left(\rho\left[c_p T + L\delta\left(T - T_M\right)\right]\right) = \nabla.\left(\lambda \nabla T\right) + Q, \quad \mathbf{x} \in \Omega,$$

In practice, this formulation is virtually impossible to implement and is not always a good description of real materials, so an altered form is used. It is assumed that the melting takes place over a small non-zero temperature range, $T_L \leq T \leq T_S$ where T_L and T_S are the liquid and solid temperatures respectively and between the two temperatures the substance is in a mixture of the two states. The time dependent term is the written as

$$\frac{\partial}{\partial t}\left(\rho\left[c_p T + L'\right]\right) = \nabla.\left(\lambda \nabla T\right) + Q, \quad \mathbf{x} \in \Omega,$$

$$L' = \begin{cases} LT/\left(T_L - T_S\right), & T_S \leq T \leq T_L, \\ 0, & \text{otherwise.} \end{cases}$$

The smaller $T_L - T_S$ is, the better the approximation to the original form.

In addition to this alteration, the model of the melting process will require the usage of temperature-dependent material properties. The temperature dependence of the properties of the graphite will need to be considered, and the metal-carbon eutectic will have different properties when it is a liquid.

8.1.2.2 Domain

From figure 8.1, it is clear that the domain geometry and its boundary conditions are radially symmetric and so a significant amount of computational effort can be saved by using an axisymmetric domain. It is assumed that the inner radius of the black body is sufficiently small that the amount of gas (argon is used) inside it can be neglected. This assumption neglects heat loss by convection, but at high temperatures radiation will be the dominant heat transfer mechanism so the assumption should be valid.

The domain also depends strongly on the question to be answered using the model. The first three questions given in section 8.1.1 are concerned with the temperature distribution of the surface of the graphite in the steady state at the end of the melting process. If the melting process is complete and the furnace is being held at a fixed temperature, then the liquid metal-carbon eutectic will be uniformly at the temperature of the furnace. It is reasonable to assume that the liquid metal forms a perfect thermal bond with the graphite black body, and so the whole of the surface of the black body in contact with the metal can be regarded as being at the furnace temperature.

Hence the only parts of the system that need to be modelled explicitly are the black body and the objects that it exchanges heat with. In this case, the main exchanges of heat occur between the furnace and the black body, and the “outside world” and the black body. The furnace and the outside world require less detail in the modelling as their temperatures are regarded as fixed. The effects of the liquid metal can be modelled as a boundary condition along one surface of the black body. An example of the model domain for the first three questions is shown in figure 8.2.

The final question looks at the melting process, and so the metal-carbon eutectic needs to be modelled in detail. In addition, the model will need to have a suitable time domain defined.

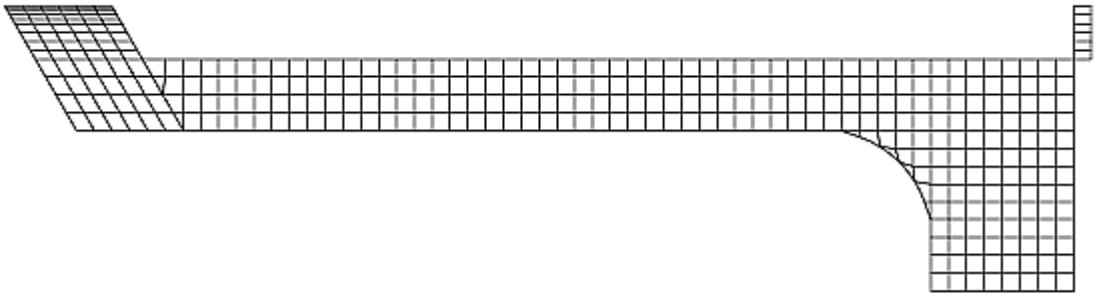

Figure 8.2: Model domain and mesh for the first two questions. The small volume at the right-hand side represents the outside world for radiative exchange and is at a fixed temperature.

8.1.2.3 Boundary conditions

The boundary conditions for this problem fall into two classes; fixed temperature conditions and radiative heat exchange conditions.

The fixed temperature conditions are applied to regions where the temperature is known. For the steady-state model, these regions are the outer boundary of the black body (where the black body is in contact with the liquid metal-carbon eutectic: the blue region in figure 8.1), the furnace, and the outside world. For the transient model, the fixed temperatures are the outside of the crucible nearest to the furnace, the furnace itself, and the outside world.

The radiative heat exchange condition between two surfaces can be written

$$\lambda \nabla T_1(\mathbf{x},t).\mathbf{n} = \left(\frac{1-\varepsilon_1}{\varepsilon_1} + \frac{1-\varepsilon_2}{\varepsilon_2}\frac{A_1}{A_2} + \frac{1}{F_{12}} \right)^{-1} A_1 \sigma \left(T_1^4 - T_2^4 \right) \qquad \mathbf{x} \in \partial\Omega,$$

where F_{12} is the view factor (a measure of how well the two surfaces can "see" one another: see section 9.2.3 for more detail), σ is the Stefan-Boltzmann constant, T_1 and T_2 are the temperatures of surfaces 1 and 2 respectively measured in K (absolute temperatures must be used), and A_1 and A_2 and ε_1 and ε_2 are the surface areas and emissivities of surfaces 1 and 2 respectively.

8.1.2.4 Material properties

As was mentioned above, a given steady state model can assume that the material properties of the graphite are constant. However, question two requires the material properties of the graphite to be treated as quantities that can vary. Since the question is essentially requesting a sensitivity analysis, it is helpful to define the distribution functions of the quantities. The emissivity of the graphite was taken to vary uniformly between 0.81 and 0.91 (the literature value as used in the other models was 0.86). The thermal conductivity was assumed to be normally distributed with a standard deviation of 0.097 times the mean value, so that 99% of the distribution lies between 75% and 125% of the mean. The mean value at each temperature was calculated by using the approximation

$$\lambda = \lambda_0 (T/T_0)^{-0.7},$$

where λ_0 = 60 W m^{-1} K^{-1} and T_0 = 1357.77 K (the copper fixed point). This approximation was taken from literature [26]. Three types of sampling method were used to choose the values used in the individual trials. These methods will be discussed further in section 8.1.5.

The material properties for the final question need to be temperature dependent, and will be dependent on position in as much as the properties in the region of the model representing the

graphite will be different from the properties in the region representing the metal-carbon eutectic.

In addition to the properties of the graphite and the metal-carbon eutectic, the emissivities of the outside world and the furnace must be defined. It is assumed that both of these are perfect black bodies and absorb all energy radiated to them, so that their emissivities are both 1.0.

It is assumed that any two materials that are in contact have a perfect thermal bond, so that only properties of well-defined materials need to be considered and artificial gaps need not be added.

8.1.3 Solution method

The problem features non-linear effects, including radiation and latent heat, and an axisymmetric geometry that can best be described using a mixture of squares, triangles, and parallelograms. These features make the finite element method the preferred solution method. The non-linear features mentioned are all offered as modelling options within the Abaqus [21] package. Other packages were considered, but Abaqus offered good surface-to-surface radiation algorithms, including view factor calculations, and good control over mesh density. In addition, Abaqus can carry out parametric studies through a scripting language, including running the jobs and collating the results, which makes efficient execution of uncertainty simulations possible.

8.1.4 Model validation

Initially, a simple version of the model was developed. Previous work [27] had involved simulation of an identical crucible using Ansys [28], another finite element software package. Correspondence with the author ensured that the same domain, boundary conditions, and material properties were used in the new work as in his work. This checking meant that the two sets of results could be compared for validation purposes.

Care had to be taken when obtaining the results of interest. The required result from a given model was the temperature averaged with respect to area over the red area in figure 8.1. The calculation of this area was complicated by the model being axisymmetric, and by the need to consider the approximation functions used in the finite element approximation when calculating the temperature distribution over the area. The correct weights were calculated and were checked by considering the cases of a uniform temperature and a linear distribution.

The first set of models used for validation simulated the crucible at the end of the melting process, and neglected the effects of the furnace. It was assumed that the only heat exchange occurring was between the inner surface of the crucible and the outside world, assumed to be held at a fixed temperature of 25 °C. Initially, the outside world was simulated using a single finite element, but this led to incorrect calculation of the view factors. A set of tests with increasingly refined meshes was run until a suitable mesh giving stable results was identified.

Once a suitable mesh had been found, three different tests were run at each of the six temperatures. The tests varied the thermal conductivity of the graphite, λ_{graph}, by ± 25%. The results of the tests are shown in table 8.1. The table shows the temperature of the inner surface of the crucible averaged over its area (the red area in figure 8.1) for all fixed points and all tests, in mK. The difference between the old and new results is at most about 5% of the old test result. This difference is likely to be due to differences between the meshes and software packages used.

The original study [27] also investigated the effects of the presence of the furnace. A second model was created using a 100 mm long furnace of radius 8 mm (the same radius as the outer edge of the black body), and it was assumed that the whole of this furnace was at the fixed point of the metal-carbon eutectic. The furnace was assumed to be exchanging heat with the inner surface of the crucible. The temperature averaged over the inner surface of the black

body was calculated, and the results are shown in table 8.2. The agreement is less good for this set of results. The Abaqus results are nearly uniformly 88% of the Ansys results. It is not clear why this is the case, although the uniformity of the difference almost implies a difference between the model boundary conditions. Other possible causes include software differences and inadequacies in the mesh.

	Temperature drops from results of original tests with Ansys (mK)			Temperature drops from results of repeated tests with Abaqus (mK)		
Metal	1.25 λ_{graph}	λ_{graph}	0.75 λ_{graph}	1.25 λ_{graph}	λ_{graph}	0.75 λ_{graph}
Co	10.5	13.0	17.0	10.5	13.2	17.5
Pd	16.5	21.0	27.0	16.8	21.0	27.8
Pt	29.0	38.0	50.0	30.9	38.5	51.2
Ru	49.0	60.0	80.0	49.5	61.8	82.1
Ir	96.0	113.0	150.0	95.8	119.2	158.1
Re	130.0	156.0	204.0	132.0	164.4	217.8

Table 8.1: Comparison between the results of equivalent models run in Ansys and Abaqus. Agreement is generally to within 5%, and results are shown in mK.

Metal	Co	Pd	Pt	Ru	Ir	Re
Ansys	10.5	16.8	30.5	49	94	129
Abaqus	9.3	14.8	27	43.3	82.8	113.6

Table 8.2: Results of tests simulating the presence of a uniform furnace. Results are shown in mK.

8.1.5 Simulation

Once the simple model had been validated, it was used to investigate the sensitivity of the results to the material properties of the graphite. Several sets of tests were run using different values of the thermal conductivity and the emissivity. Values were sampled from the distributions of these material properties described in section 8.1.2.4 using three different methods: sampling at regular intervals, Monte Carlo sampling, and Latin Hypercube sampling. Each method was used to select N^2 sets of values, where $N = 2$, 5, 10 and 32.

The regular sampling method set limits on each of the material properties and then took N evenly-spaced samples, including both limits, from each sample. The limits on the thermal conductivity were taken to be 75% and 125% of the mean value, which is a 99% confidence interval, and the limits on the emissivity were taken to be the limits of the interval over which it is distributed. So for example, the values chosen for the emissivity for $N = 5$ were 0.81, 0.835, 0.86, 0.885, and 0.91. The N values of each quantity were then combined to generate every possible pair of one thermal conductivity with one emissivity. These pairs of values were then used as input values for the model. The regular sampling method ensures that the whole of the ranges of the input quantities are sampled, but it completely ignores the distributions of the input quantities and so it does not give any information about the distribution of the results. As N tends towards infinity, this method becomes equivalent to the other two methods described here, but for practically applicable values of N it is not equivalent.

Monte Carlo sampling [29] samples randomly from the distribution function of the quantity. In order to generate samples, $2N^2$ random numbers between 0 and 1 were generated, and these values were used as cumulative probability levels to generate material property values by inverting each property's cumulative distribution function. This sampling method generates

good approximations to the probability distribution of the results when N^2 is large, but the method is not guaranteed to cover the full range of the input quantities and the quality of the approximation decreases as N decreases.

Latin Hypercube sampling [30] ensures that the full ranges of the input quantities are sampled, by combining regular sampling with Monte Carlo sampling. The ranges of the input quantities are divided into N^2 sub-ranges of equal probability, and a random value is chosen within each of the sub-ranges. The two sets of N^2 values are then paired up at random to create the input values for the model. It can be shown [30] that for independent input quantities, Latin Hypercube sampling is an unbiased sampling method.

Note that the random sampling methods can create values outside of the limits for the regular sampling method, because the thermal conductivity had artificial limits put on its range during the regular sampling, and it is expected that some values outside of this range will be created when N^2 is large.

Figure 8.3 shows an illustration of the values produced by the different sampling methods for $N^2 = 25$. The grid pattern produced by the regular sampling method is clear. The vertical axis is split into 25 equally-sized intervals, each representing a region of equal probability for the emissivity. It is difficult to see, but there is exactly one triangle, representing the Latin Hypercube sampling method, in each interval. The intervals of equal probability on the horizontal axis are also shown. They are not of equal sizes as the normal distribution is used, but again exactly one triangle is in each sub-interval. The Monte Carlo samples (circles) are in random places, and it is clear that most of the values are concentrated near to the centre of the thermal conductivity axis, as would be expected.

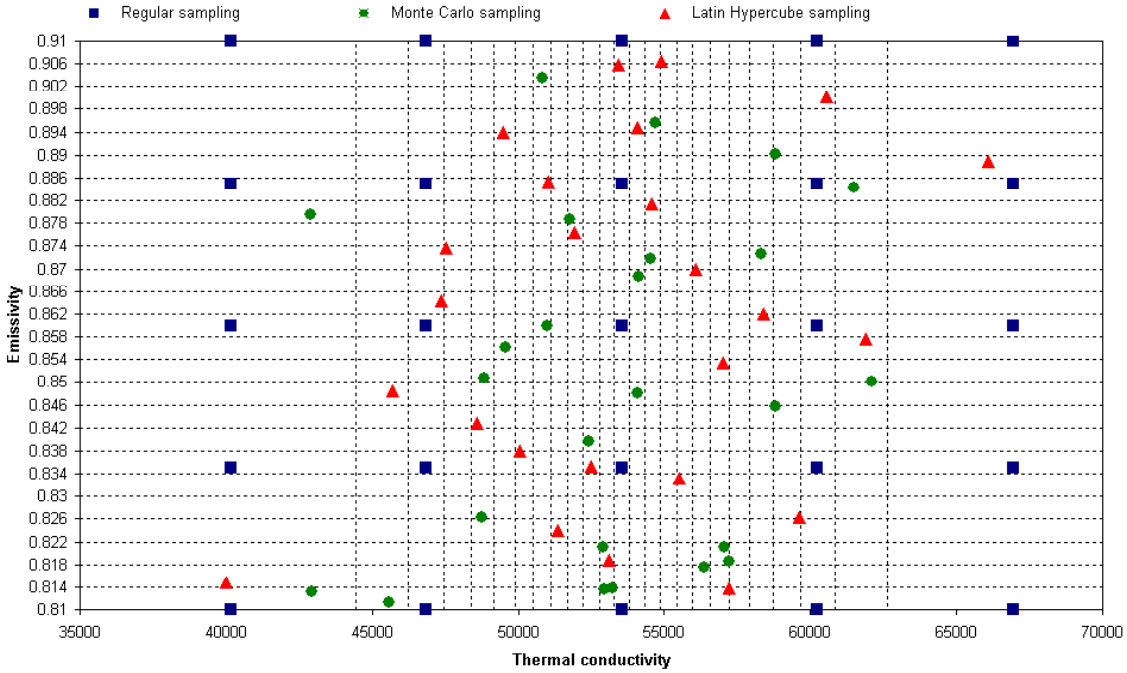

Figure 8.3: Sampled values for the three different methods from the space of possible input quantities.

The results shown in this section are from tests run at a single temperature. As yet the behaviour for other temperatures has not been tested, but it is likely to be similar.

Whilst the results of the regular sampling tests could not be used to determine the probabilistic behaviour of the results, they did provide useful information about the relationship between the input values and the results. Analysis of the regular sampling results for $N = 32$ led to a relationship of the form

$$\Delta T = \left(\frac{A}{\lambda} + B\right)(C\varepsilon + D)$$

being developed, where ΔT is the areally averaged difference between the fixed-point temperature and the measurement surface temperature. This relationship was tested using the results of the randomly-sampled tests, and relative accuracy was found to be within 0.1% for all values. This approach is similar to the response surface methodology commonly used in failure testing [31]. The form of the relationship is not entirely unexpected: the higher the thermal conductivity, the better heat will be conducted from the molten metal, and the lower the emissivity the less heat will be lost to the atmosphere.

The development of this simple relationship made it possible to generate a distribution function for the model results, by generating 10000 random sets of input values using Monte Carlo sampling, and using the simple relationship to generate the corresponding result value. This distribution function can be considered to be the "correct" distribution in order to assess the accuracy of the distributions generated by other methods. In particular, the result associated with a given level of cumulative probability can be calculated and compared to the value obtained from the Monte Carlo and Latin Hypercube sampling methods. For interest, this analytic model and the results of the 1024 regular sampling trials are shown in figure 8.4. Note that the regular trials are a bad approximation to the true distribution because the input values were not selected in a probabilistic way.

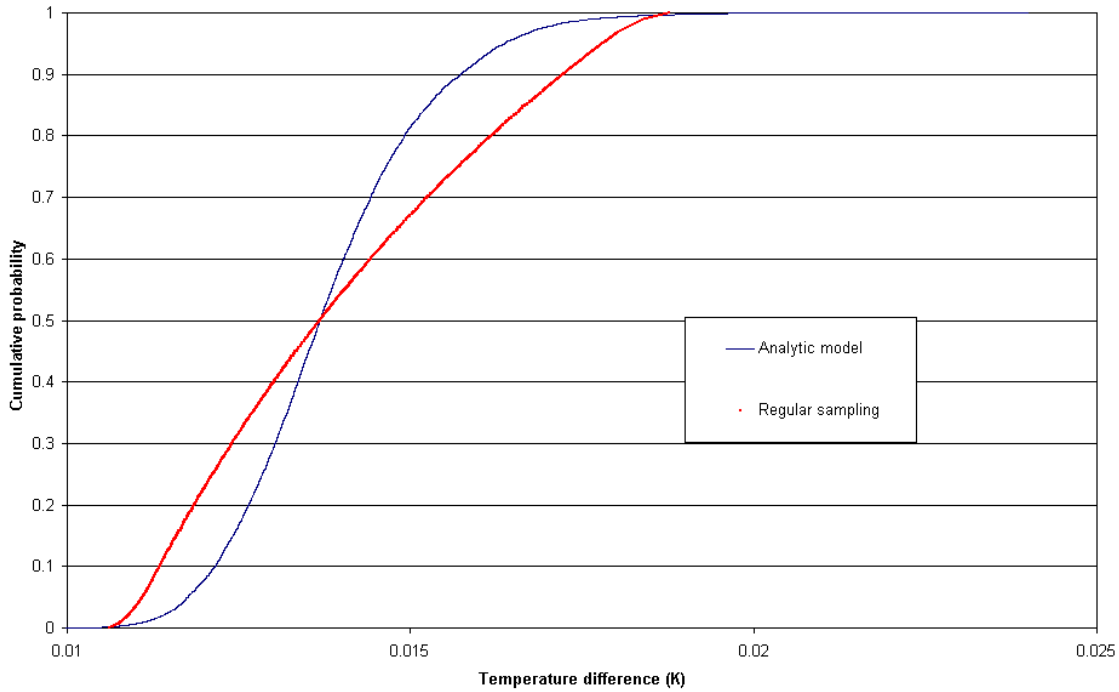

Figure 8.4: Analytic model and results of the regular sampling.

Figure 8.5 shows the results from the trials with $N^2 = 1024$. The "true" distribution cannot be identified clearly because it is covered by the other two distributions, and particularly by the Latin Hypercube sampling results. Figures 8.6 and 8.7 show the results for $N^2 = 100$ and 25 respectively. These plots show the distributions more clearly. In particular, it looks as though the Latin Hypercube sampling results are a better approximation than the MC sampling results.

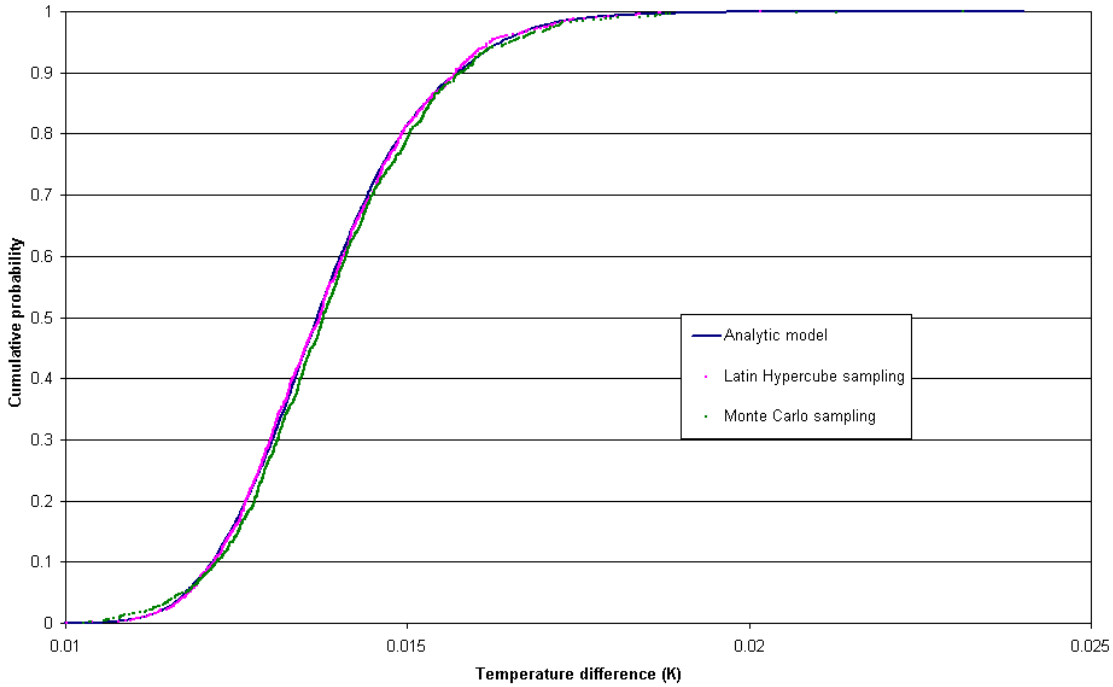

Figure 8.5: Test results for $N^2 = 1024$ samples.

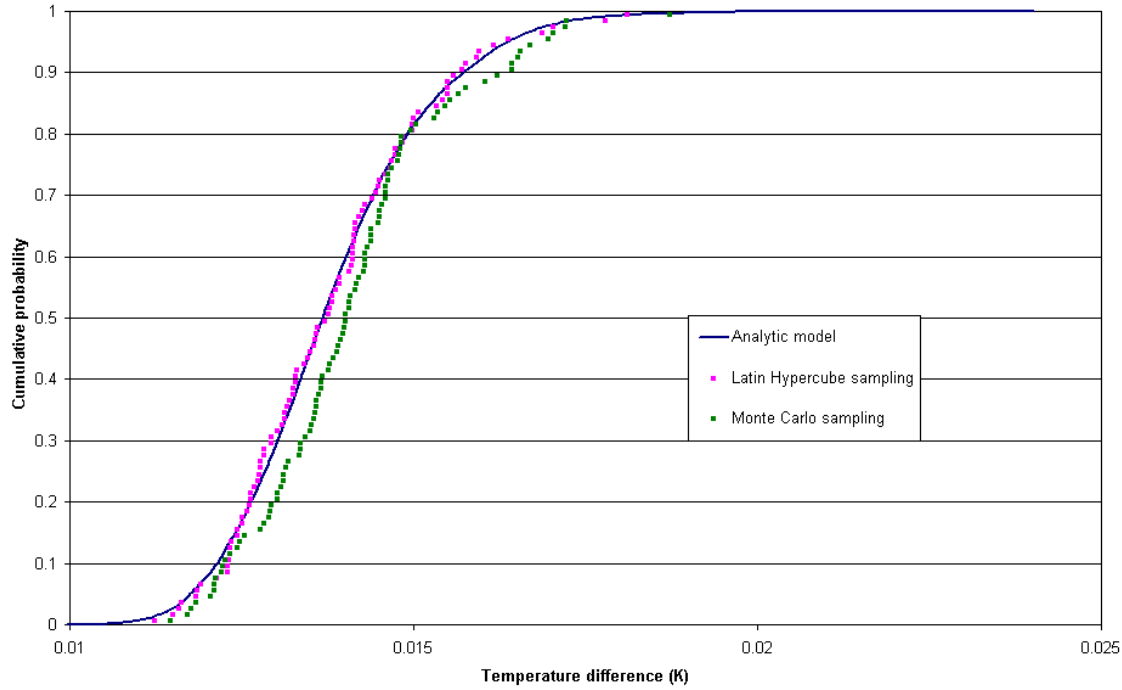

Figure 8.6: Test results for $N^2 = 100$ samples.

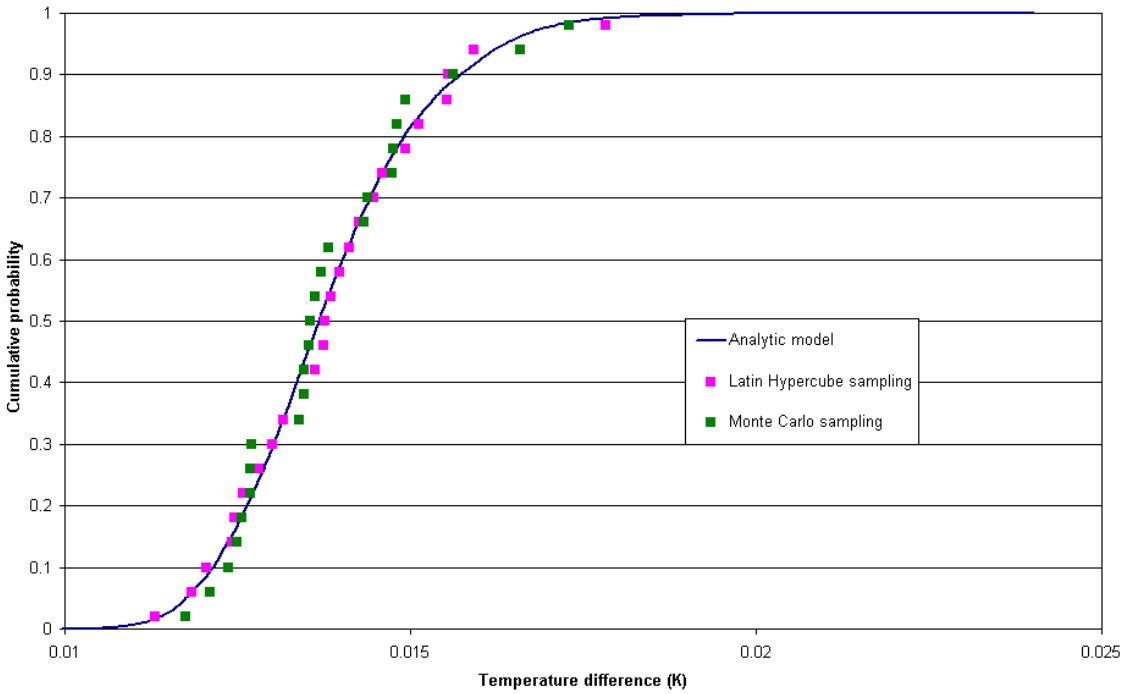

Figure 8.7: Test results for $N^2 = 25$ samples.

In order to check this hypothesis, the RMS relative error

$$e_{rel} = \sqrt{\frac{1}{N^2}\sum_{i=1}^{N^2}\left(1-\frac{\Delta T_s}{\Delta T_a}\right)^2}$$

was calculated, where ΔT_s is the result from the sampled tests and ΔT_a is the result from the analytic model. The results are shown in table 8.3.

The trend of e_{rel} with varying N for the Monte Carlo sampling is somewhat unexpected, since the result for N^2=25 is more accurate than that for $N^2 = 100$. The result is probably due to random chance, which illustrates the danger of drawing conclusions from a single set of results. Extra sets of 10 tests with $N = 100$ and $N = 25$ were run to obtain a more statistically valid basis for the conclusions, and average values of e_{rel} are also shown in table 8.3. These averaged values are more in line with expectations.

Number of samples	Single tests		Average over 10 tests	
	Monte Carlo	Latin Hypercube	Monte Carlo	Latin Hypercube
25	0.0156	0.0119	0.0352	0.0171
100	0.0229	0.0076	0.0129	0.0084
1024	0.0103	0.0037	-	-

Table 8.3: Relative RMS errors for Monte Carlo and Latin Hypercube sampling, for different values of N^2. Result of a single test and the average over 10 tests are shown.

8.1.6 Further results

After validation, further simulations of the NPL furnace and black body were carried out. These simulations aimed to answer the third question by examining the effects of furnace non-uniformity. Three different models of the black body and furnace were developed: one with a uniform furnace temperature, one with a non-uniform furnace temperature, and one with a non-uniform furnace temperature and extra insulation as shown in figure 8.8. The

insulation is present when the apparatus is in normal operation. The cavity was closed by a block of air at room temperature, as in previous models.

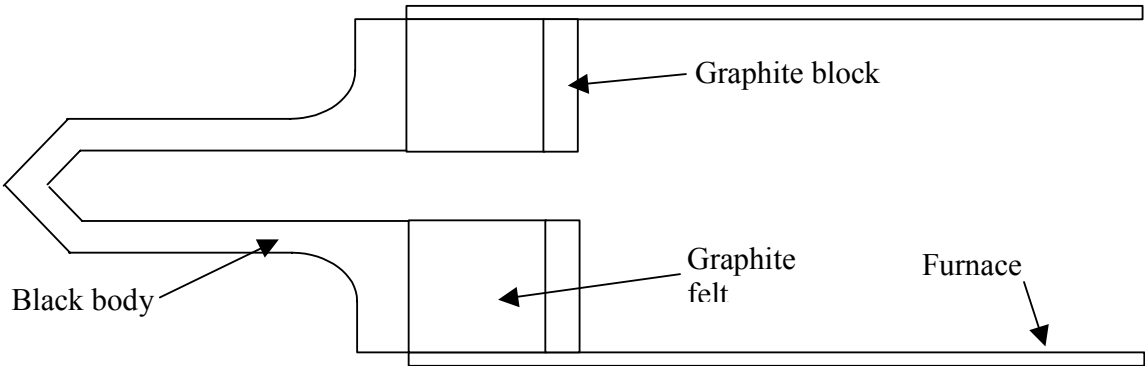

Figure 8.8: Diagram showing the furnace and extra insulation.

Five sets of measurements of the furnace temperature were available, each set consisting of temperatures measured at various points along the length of the furnace for a given target temperature in the absence of the black body and the insulation. Each set of measurements was interpolated linearly to obtain the boundary conditions for the model with a non-uniform furnace. The measurement data, plotted as temperature differences from the fixed-point temperature, are shown in figure 8.9. Note that the average furnace temperature is lower that the fixed point temperature, so it is expected that the non-uniform furnace will increase the temperature difference across the base of the black body.

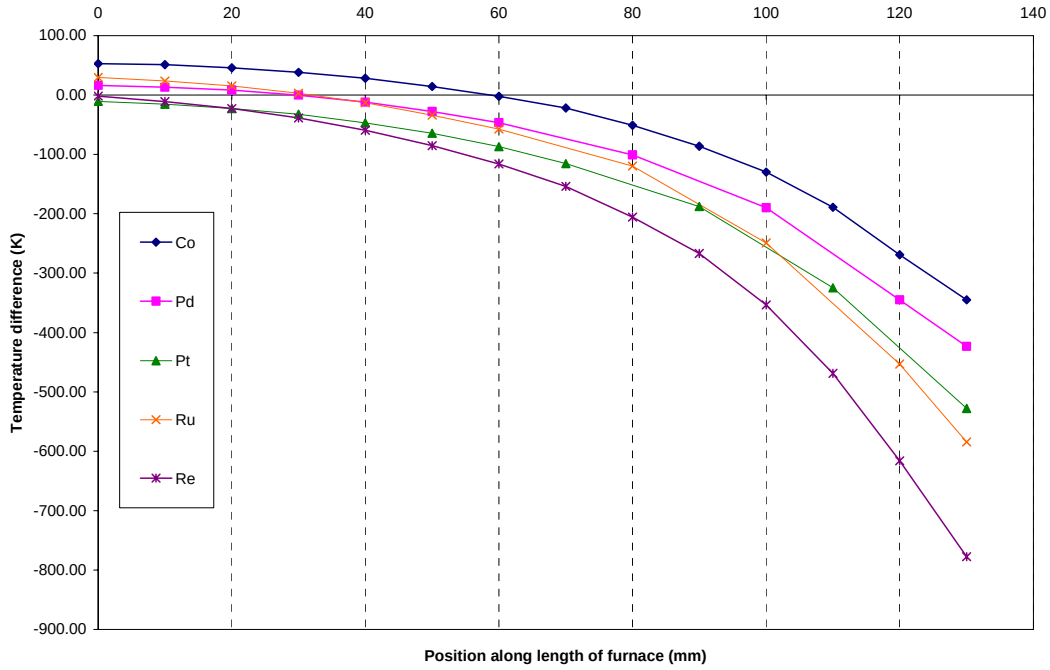

Figure 8.9: Measurement data from the non-uniform furnace, plotted as temperature difference from the fixed point against position along the length of the furnace (zero being the point at which the furnace meets the black body).

The model with a uniform furnace fixed the whole of the furnace to be at the melting point of the metal-carbon eutectic. In all cases, the furnace was modelled as a hollow cylinder of inner radius 12 mm radiating heat to the black body.

The graphite block was taken to be made from the same material as the black body. The graphite felt was taken to have temperature-dependent thermal conductivity of between 0.4 and 1.4 W m^{-1} K^{-1}. These values were taken from a model of another high-temperature

process involving graphite felt insulation [32], and agreed with other reference values. It is likely that only an order of magnitude value is required.

The results for the three models at five different temperatures are shown in table 8.4. The results show that the non-uniform furnace produces a higher temperature drop than the uniform furnace, as expected, but not hugely different (between 5% and 8% of the uniform furnace value).

Metal	Uniform furnace	Non-uniform furnace	Extra insulation
Co	11.4	12.0	3.3
Pd	18.3	19.4	5.8
Pt	33.4	36.0	11.2
Ru	53.5	57.2	16.5
Re	140.5	152.3	46.8

Table 8.4: Areally averaged temperature difference between the measurement area of the black body and the fixed point temperature.

It is also clear that the presence of the insulation has a large effect on the temperature drop. The area of the black body nearest to the outside air loses heat radiatively, and heat is conducted away from the measurement area to compensate. If an insulating material is placed between the surface that exchanges most heat with the outside air, the heat is not conducted away from the measurement area and so the temperature difference decreases. Additionally, the radiation strongly affects the view factor between the bottom of the black body and the outside world, which will also reduce the temperature difference. It is not clear at the moment which effect is most significant.

This idea is illustrated in figures 8.10, 8.11, and 8.12. Figures 8.10 and 8.11 show the temperature distribution within the black body without and with insulation respectively. There is a 65 K temperature drop within the insulation, meaning that little heat is conducted from the black body to the upper graphite block.

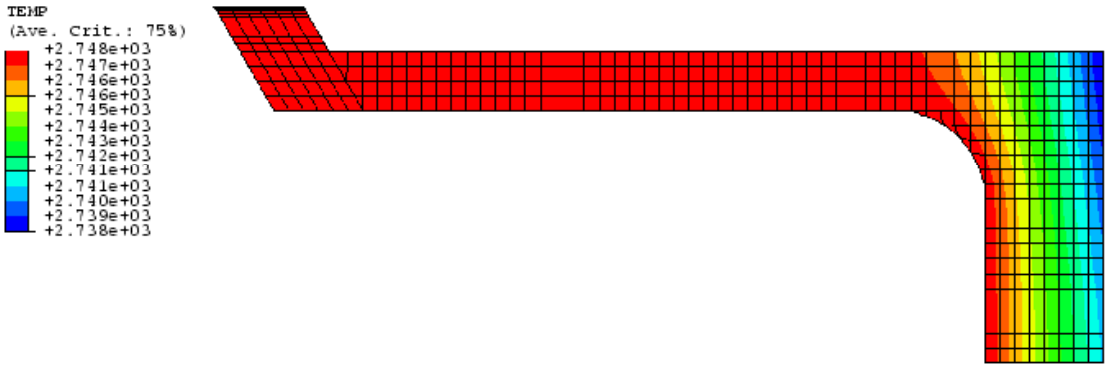

Figure 8.10: Contour plot of temperature within the black body for a model with a non-uniform furnace and no insulation. Contour values run from 2738 K to 2748 K. Note that one "unit" of colour is less than 1 K, so the scale labels may be misleading.

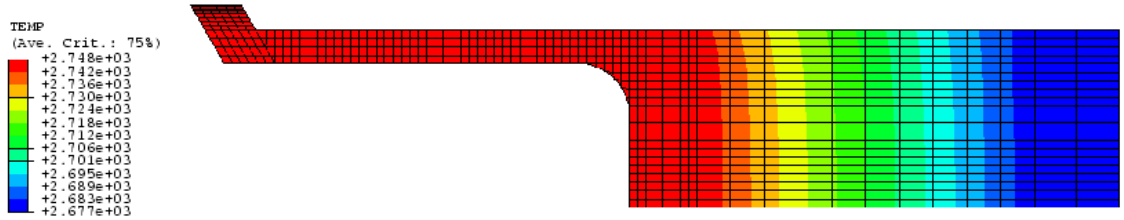

Figure 8.11: Contour plot of temperature within the black body and the insulation for a model with a non-uniform furnace and insulation. Contour values run from 2677 K to 2748 K, red area is 2742 – 2748 K.

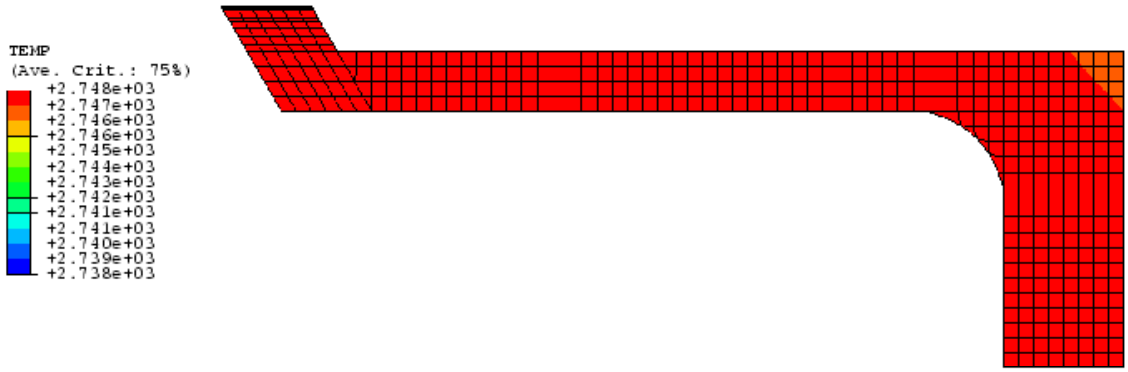

Figure 8.12: Contour plot of temperature within the black body for a model with a non-uniform furnace and insulation. Contour values run from 2738 K to 2748 K for easy comparison with figure 8.10. Note that one "unit" of colour is less than 1 K, so the scale labels may be misleading.

Figure 8.12 shows the temperature distribution within the black body in the presence of insulation to the same scale as figure 8.10. The insulation makes the temperature of the black body much more uniform (less than 2 K variation as opposed to approx 10 K). These contours were all taken from the Re models, which is the highest fixed point temperature, and hence they are likely to be worst-case scenarios.

8.1.7 Next developments

The next developments will be to model the time-dependent melting process as described in the preceding sections. This model will be validated against existing work and then extended to include the measured furnace non-uniformity.

8.2 Case Study 2: Near- to far-field underwater acoustics

Some of the work in this case study has been reported elsewhere [33, 34], and so full details of the derivation of some of the expressions and results quoted here are not given in this study.

8.2.1 The physical problem

The generation and detection of acoustic fields in water are typically undertaken using electroacoustic transducers, either deployed individually or in arrays. Such devices are most often characterised in terms of absolute sensitivity levels and device directivity. Measurement of the sensitivity of transducers and arrays may be carried out in a variety of types of facility, for example laboratory tanks or open-water sites. The sensitivities and directional response are usually required to be determined in the acoustic far-field.

However, if the array or transducer is physically large when measured in acoustic wavelengths (high *ka* value, where *k* is the wavenumber and *a* is the largest physical dimension), it may be impossible for far-field conditions to be achieved in a facility of finite size whilst maintaining free field environment and preserving steady state conditions. An open-water facility will in general enable a greater source-receiver separation to be used, but such facilities have the disadvantage that there is little or no control of environmental conditions and, in any case, even an open-water facility will place some limit on the maximum separation achievable. Making measurements at sea is an extremely expensive alternative and provides no environmental control at all. There is consequently a strong motivation to maximise the range of acoustic testing possible in laboratory tanks.

At NPL, measurements can be made in a laboratory tank of 5.5 metres diameter and 5 metres deep. To measure on a cylindrical surface surrounding the projecting transducer, a series of measurements is made on concentric circles of the same radius but offset in depth. Each circular data set is obtained by rotating the projector and not by scanning the hydrophone, as depicted in figure 8.13. In between the measurement of each successive circular data set, the depth of the projecting transducer and/or receiving hydrophone is adjusted so that the next data set is measured with the relative depth increased by a specified increment.

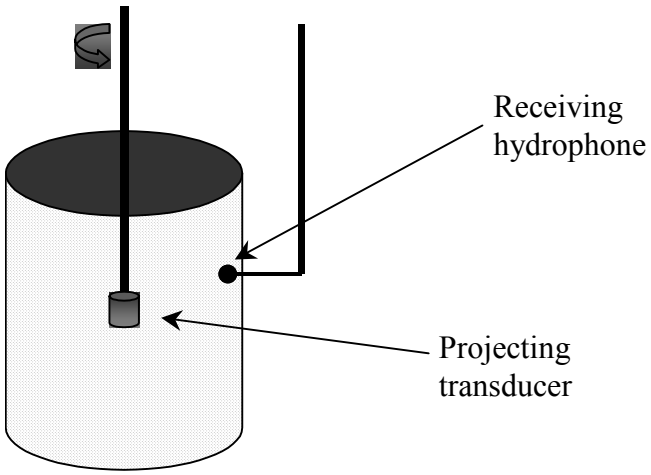

Figure 8.13 Schematic representation of a cylindrical scan.

One option to overcome the restrictions posed by finite-sized laboratory tanks is to undertake measurements in the acoustic near-field and then predict the acoustic far-field response from the near-field data. This case study describes one method for carrying out these predictions. The physical problem is to predict the pressure distribution caused by an underwater acoustic transducer at some point in space, given measurements of the pressure distribution on an open-ended cylindrical surface.

8.2.2 Mathematical formulation

8.2.2.1 Governing equation

The pressure distribution of a transducer in water is assumed to obey the time-dependent linear wave equation. This assumption, which neglects diffusivity and nonlinearity, is expected to be reasonable for the frequencies of interest. It is also assumed that the density and speed of sound in the medium are constant and uniform. The equation is

$$\frac{\partial^2 \Psi(\mathbf{x},t)}{\partial t^2} = c^2 \nabla^2 \Psi(\mathbf{x},t), \qquad \mathbf{x} \in \Omega,$$

where c is the wave speed and Ψ is the scalar time-dependent velocity potential. If it is assumed that the transducer signal is at a single frequency then this equation becomes

$$\nabla^2\varphi + k^2\varphi = 0, \qquad \mathbf{x} \in \Omega,$$

where $k = 2\pi f / c$ and f is the frequency in Hertz (see section 9.6.1 for how to obtain this equation). This is the Helmholtz equation. If the signal is likely to be made up of several frequencies, each can be modelled separately using the Helmholtz equation.

The boundaries of the real physical domain (e.g. the laboratory tank) are assumed to be far enough away that the measurements can be treated as though they are made in an infinite domain. This assumption produces a simpler model than that produced by considering interactions with the walls and free surface of the tank. The assumption is valid because the boundaries of the tank are a large number of acoustic wavelengths away from the measurement surface and gating techniques are used to ensure that the signal that is measured is not contaminated by echoes.

8.2.2.2 Domain

The results are to be determined in the region outside of the measurement cylinder, so the problem is an **exterior** one. The domain is taken to be the whole of free space outside of the measurement cylinder, an infinite domain. The main reason for this is that the metrologist is interested in the performance of the device in a domain, such as the sea or a lake, which is too big to be modelled in detail. An additional reason is that there is no justification for placing any other boundaries on the problem. If another detailed measurement scan had been taken on a cylinder of larger radius, then it would be justifiable to take the larger radius cylinder as an outer boundary, but this is not the case. The methods used to avoid the problems associated with infinite domains are discussed in section 8.2.3 below.

The measurement surface and the infinite domain are both axisymmetric. However, the measurement data is generally directional and so the domain cannot be regarded as two-dimensional if the measurement data is used in the form it is collected. Section 8.2.3 describes a method by which the three-dimensional problem can be reduced to a set of independent two-dimensional problems.

8.2.2.3 Material properties

The material properties within the medium and the frequency of the transducer are both regarded as constant throughout. The main material property required is the speed of sound in the medium. Additionally, density is required if it is necessary to convert from velocity potential to pressure.

The main two factors that will affect the variability of the material properties are the uniformity of the medium, the motion of the medium, and the temperature distribution within the medium. Under laboratory conditions, all three of these factors are under tight controls, and so the assumptions of a stationary uniform isothermal medium are valid and measurement data can be used to validate the model.

Under real operating conditions, it is very likely that the medium will be moving and that it will not be at a uniform temperature. Ideally, a sensitivity analysis would be used to see how the speed of sound in the medium affects the results, but as yet this has not been done.

8.2.2.4 Boundary conditions

There are two sets of boundary conditions that apply to this problem. The measurement data can be used to generate a fixed value boundary condition on the cylindrical boundary, and a condition governing the behaviour “at infinity” can be applied.

The usage of the measurement data creates a problem. Every domain needs to have closed boundaries. The measurement data are gathered on the curved surface of an open cylinder. No measurements are taken on the ends of the cylinder. Hence some sort of boundary conditions

must be created for the ends of the cylinder. Several possibilities were identified, based on the belief that the device is sufficiently directional that the pressure at the centre of the cylinder end will be zero. The whole of the cylinder ends could be taken to be at zero pressure, the pressure could be taken to decrease linearly to zero, or a more sophisticated approximation could be used.

Sets of tests were carried out applying these conditions to simulated data for which an analytic solution was known. It was found that the linearly decreasing conditions gave a reasonably accurate solution, and that the solution was more accurate when the measured pressures at the top and bottom of the cylinder were close to zero. This discovery can be used to determine how many measurement scan lines are required for a good solution.

The condition applied at infinity is the Sommerfeld radiation condition, which states that all pressure waves at infinity are outgoing. In three dimensions, this condition can be written as

$$\lim_{R \to \infty} R\left\{\frac{\partial \varphi(\mathbf{x})}{\partial R} - ik\varphi(\mathbf{x})\right\} = 0$$

where R is the radius of the point $\mathbf{x}$ in spherical polar coordinates.

8.2.3 Solution method

The first problem that the choice of solution method has to address is that of the infinite domain. Infinite domains are a mathematical idealisation of the real physical situation, and are computationally impossible. As mentioned in section 3.2.3, several methods exist to combat this problem, including introduction of artificial boundaries and use of infinite finite elements.

Since the solution is not required everywhere within the domain, and the Sommerfeld radiation condition makes the equation tractable to reformulation, the most efficient technique for this problem is to rewrite the Helmholtz equation as a surface integral and solve that. If the closed version of the measurement surface is S, and the domain external to the cylinder is Ω, then the reformulated problem can be written as

$$\int_{S(\mathbf{x}')} \left(G_k(\mathbf{x},\mathbf{x}')\nabla\varphi(\mathbf{x}') - \varphi(\mathbf{x}')\nabla G_k(\mathbf{x},\mathbf{x}')\right).d\mathbf{A}_{\mathbf{x}'} = \begin{cases} \varphi(\mathbf{x}), & \mathbf{x} \in \Omega, \\ \frac{1}{2}\varphi(\mathbf{x}), & \mathbf{x} \in S, \end{cases}$$

where

$$G_k(\mathbf{x},\mathbf{x}') = \frac{e^{ik|\mathbf{x}-\mathbf{x}'|}}{4\pi|\mathbf{x}-\mathbf{x}'|}$$

is the Green's function for the problem. The values of φ are known from measurement on the surface. The solution process solves the equation for $\mathbf{x}$ on S to obtain values of the gradient of φ, and then uses the values on S to calculate φ for $\mathbf{x}$ in Ω.

It can be shown that this formulation is non-unique for an infinite set of wavenumbers k, and that attempts to evaluate the integrals involved in the formulation are subject to numerical instability for values close to the non-unique wavenumbers. Several methods exist for avoiding this problem [35, 36] and two of the most popular were compared on some test problems (see section 8.2.4).

The boundary element method (BEM) was chosen to solve the surface integral problem. The surface S is split into a set of elements and the surface integral is approximated within each element, leading to a large set of linear equations that are solved to obtain the values on S and hence the values of φ at other points in the external domain. Other methods have also been tried, but they usually involve assumptions and approximations and so are not always suitable for real measurement data.

One problem with the boundary element method is the mesh size. For a device with a wavelength $\lambda = c/f$, at least three elements are required in every direction (vertical and circumferential) to produce good results, and it is often recommended that five elements are used per wavelength. Typically, the frequencies of interest are at least tens of kHz. For a frequency of 20 kHz in pure water (c = 1461 m.s^{-1} at 20 °C) and a measurement surface of radius 0.338 m and height 0.98 m, this results in 6 000 elements. Since the matrix of equations includes entries for every possible pair of interactions between elements, there will be 36 million terms to calculate and store, which leads to problems with computer memory. These memory problems can be removed if the model can be rewritten as a two-dimensional problem, as that would require around 100 elements for the same problem.

As was mentioned in section 8.2.2.2, the domain is axisymmetric but the boundary conditions as measured are not. However, if the boundary conditions on the measurement surface are written as a Fourier series,

$$\varphi(r,\theta,z) = b_0(r,z) + \sum_{k=1}^{\infty} \{a_k(r,z)\sin k\theta + b_k(r,z)\cos k\theta\},$$

then it can be shown that each component can be modelled separately as a two-dimensional problem, using a slight modification of the Green's function:

$$\int_{\zeta} c_m(\mathbf{x}') \frac{\partial}{\partial n_{\mathbf{x}'}} G_k^m(\mathbf{x},\mathbf{x}') - G_k^m(\mathbf{x},\mathbf{x}') \frac{\partial}{\partial n_{\mathbf{x}'}} c_m(\mathbf{x}') d\zeta_{\mathbf{x}'} = \begin{cases} \frac{1}{2} c_m(\mathbf{x}), & \mathbf{x} \in S, \\ c_m(\mathbf{x}), & \mathbf{x} \in \Omega, \end{cases}$$

where

$$G_k^m(r_x, z_x, r_{x'}, z_{x'}) = \int_{-\pi}^{\pi} \cos(m\alpha) G_k(r_x, 0, z_x, r_{x'}, \alpha, z_{x'}) d\alpha$$

c_m is any of the a_m and b_m, and ζ is the curve that is rotated around the axis to form the measurement surface (in this case, three straight lines that generate a cylinder). This reformulation makes it possible to parallelise the model and run the calculation for each term in the Fourier series separately. NPL has a distributed computing system that enables the entire calculation for all terms to be run in approximately the same time as a desktop PC takes to run the calculation for a single term.

The Fourier series given above involves summation of an infinite number of terms. Clearly this is impractical. Additionally, the series is obtained from a finite number of data points and so the number of terms that can be determined is limited. It is possible to calculate the goodness of fit of the approximation to the input data for a set of partial sums of the form

$$\varphi(r,\theta,z) = b_0(r,z) + \sum_{k=1}^{N} \{a_k(r,z)\sin k\theta + b_k(r,z)\cos k\theta\},$$

and truncate the series when the approximation is sufficiently accurate in some sense. The model truncated the series when the root mean square difference between the measured data and the Fourier series approximation was less than 5% of the mean of the measured data. The effect of the truncation on the results was tested (see section 8.2.5).

8.2.4 Model validation

The model was validated at each stage of its development. Two main sources of validation data were used: analytic solutions of the Helmholtz equation, and alternative methods of calculation.

Since the model domain is the exterior of a cylinder, any solution of the Helmholtz equation that is well-behaved outside of the cylinder can be used as an analytic source of validation data. In particular, the Green's function G_k which obeys

$$\nabla^2 G_k(\mathbf{x},\mathbf{x}') + k^2 G_k(\mathbf{x},\mathbf{x}') = \delta(\mathbf{x}-\mathbf{x}'), \qquad \mathbf{x},\mathbf{x}' \in \Omega,$$

can be used if $\mathbf{x}'$ is fixed to lie within the cylinder. G_k represents a point source or monopole at $\mathbf{x}'$. In general it is convenient to fix $\mathbf{x}'$ to be the origin, but in fact any linear sum of point sources will satisfy the Helmholtz equation. In particular it is common to consider two point sources 180° out of phase, separated by a small distance δ. The two point sources form a dipole, and as δ decreases to zero, the solution can be written as

$$G(\mathbf{x}) \approx \frac{e^{ik\sqrt{r^2+z^2}}}{4\pi\sqrt{r^2+z^2}} \cos\theta \left[ik + \frac{1}{\sqrt{r^2+z^2}} \right], \qquad \mathbf{x} \in \Omega,$$

where $\mathbf{x}$ is written in cylindrical polar coordinates. The monopole and dipole were used to test the model and software implementations throughout the process. An example of the results obtained using a dipole test case is shown in figure 8.14. Agreement is good.

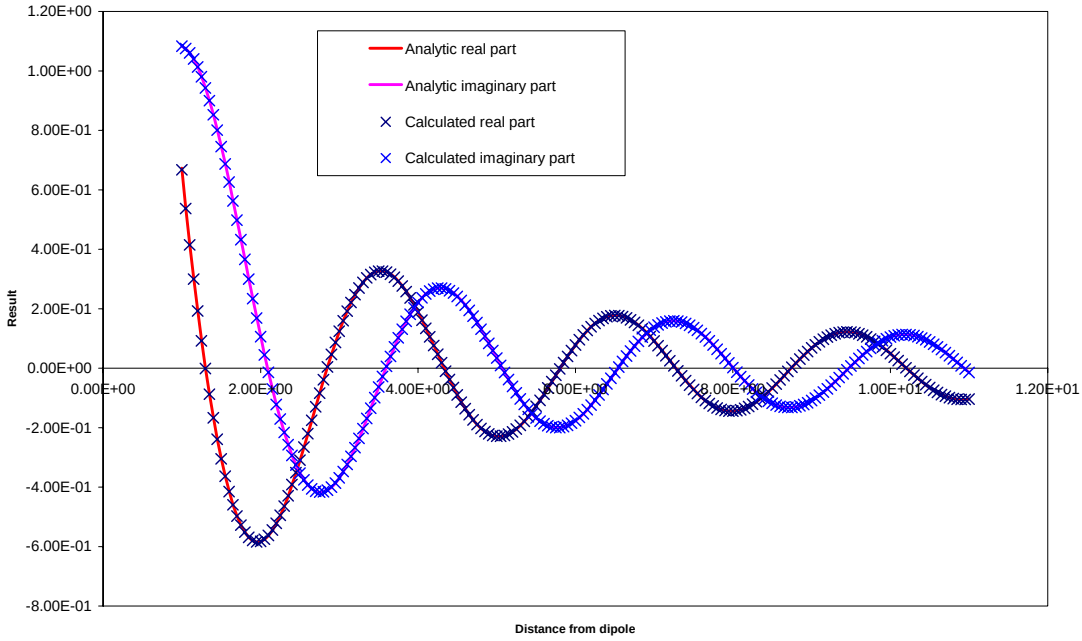

Figure 8.14: Real and imaginary parts of the distribution resulting from a dipole calculated along a straight line using an analytic solution (solid line) and boundary element method (crosses).

Several alternative methods exist for calculation. The Helmholtz equation can be solved using finite elements if suitable assumptions are made about the behaviour at infinity. An example of this technique was used [34] to test the software implementations of the two methods available for avoiding numerical singularity.

A well-understood problem was modelled using two different types of finite element, and the results of the finite element model were used as reference results for comparison with the results from the boundary element software. An example of the results is shown in figure 8.15. The problem is a two-dimensional scattering problem. The results are shown for two different input meshes and the reference solution. The agreement between the reference solution and the test results improved as the number of points in the mesh was increased.

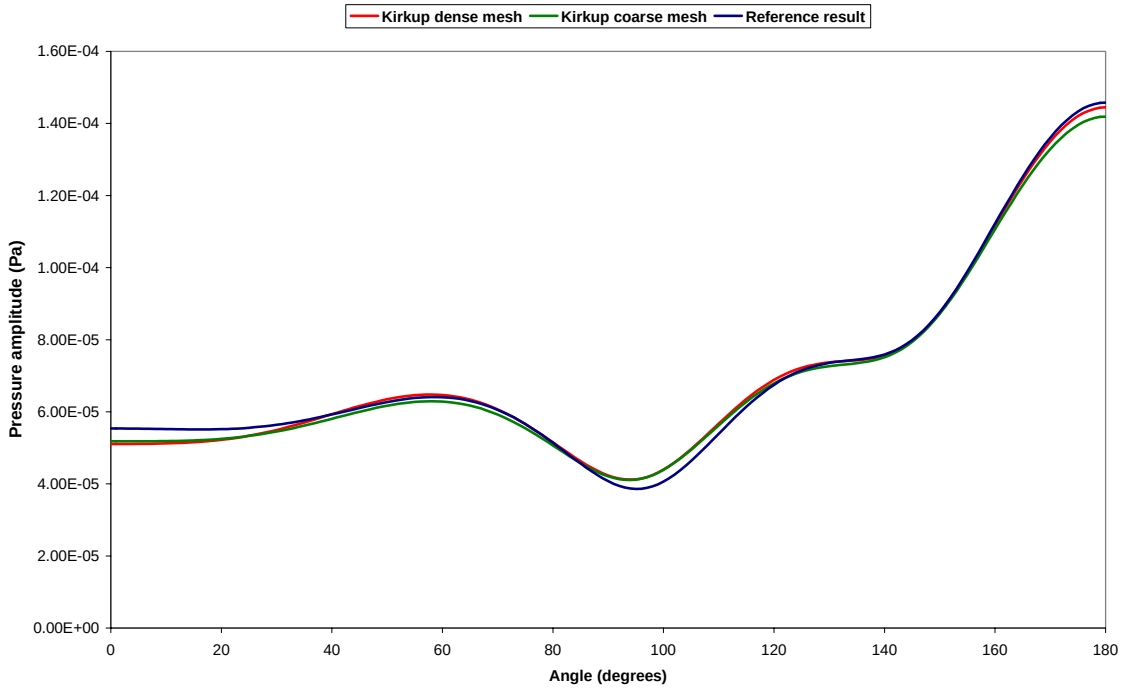

Figure 8.15: Test results from a scattering problem [34].

8.2.5 Results

Once the model had been validated, the software was applied to real measurement data sets. Four different input sets were used, gathered from two different transducers, each used at two different frequencies. One transducer was fairly uniform in the circumferential direction (i.e. for a fixed radius and height, the pressure was more or less the same all the way along the scan line, see figure 8.16 for an example), so the second transducer was designed to provide a more challenging data set.

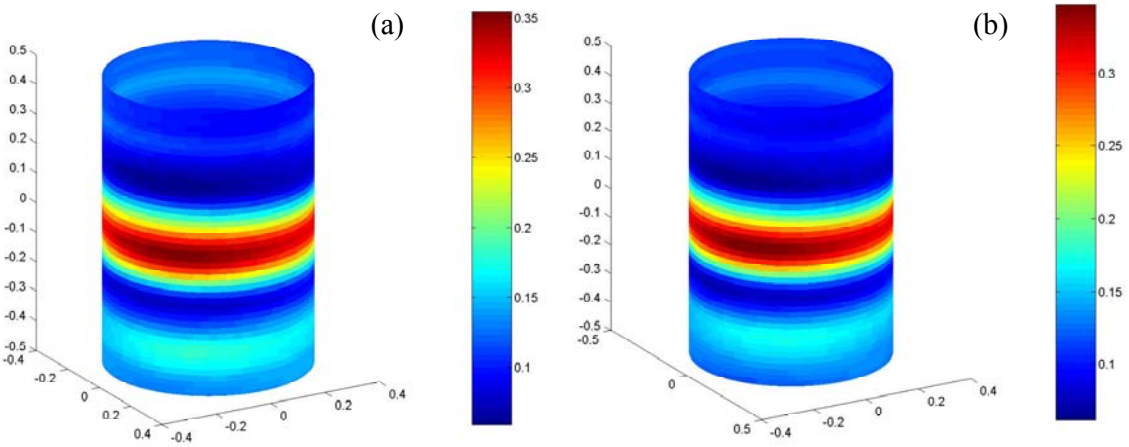

Figure 8.16 Measured data (a) and calculated results (b) in the near-field at a frequency of 13.9 kHz.

Typical results are shown in figures 8.16 to 8.19. All plots show the magnitude of the pressure. Results were calculated in the near-field as well as the far-field since the BEM does not require any assumptions about how far away the source is. A set of near-field results are shown in comparison with measured data in figure 8.16. For some cases it was possible to use two different sets of input data measured at different radii to calculate the same far-field results. An example of this is shown in figure 8.17. In general the agreement between measurement and calculation is quite good. Figures 8.18 and 8.19 show example results from

the more challenging transducer. It is clear that the agreement at the higher frequency is not as good.

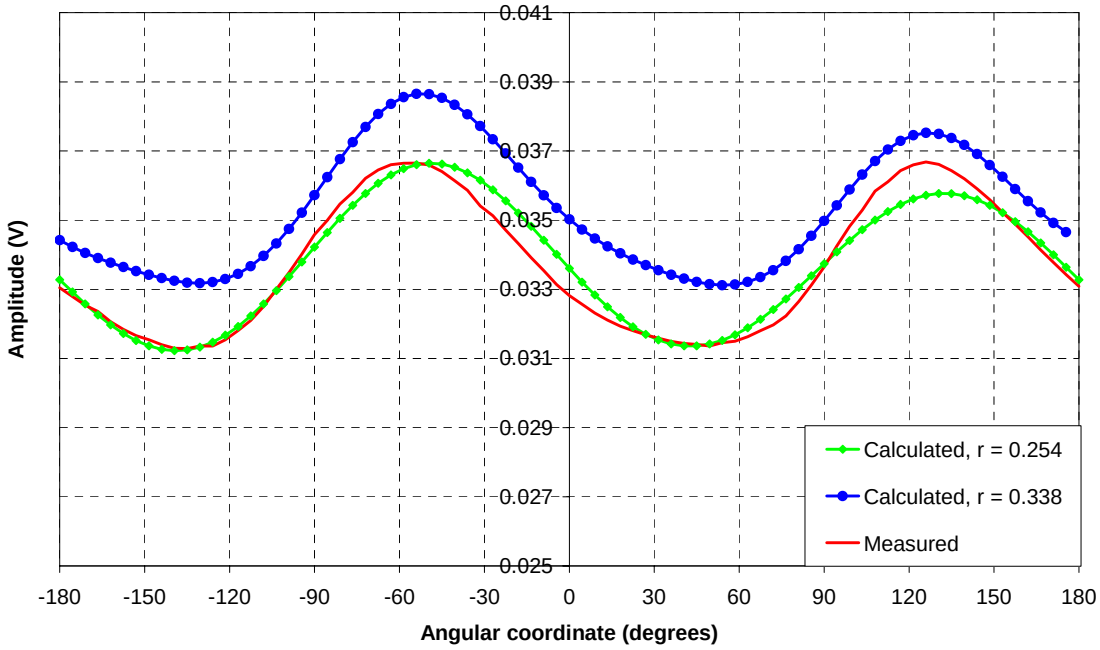

Figure 8.17: Measured far-field data along a single scan line compared with results calculated by propagating data from cylinders of radius 0.254 m and 0.338 m.

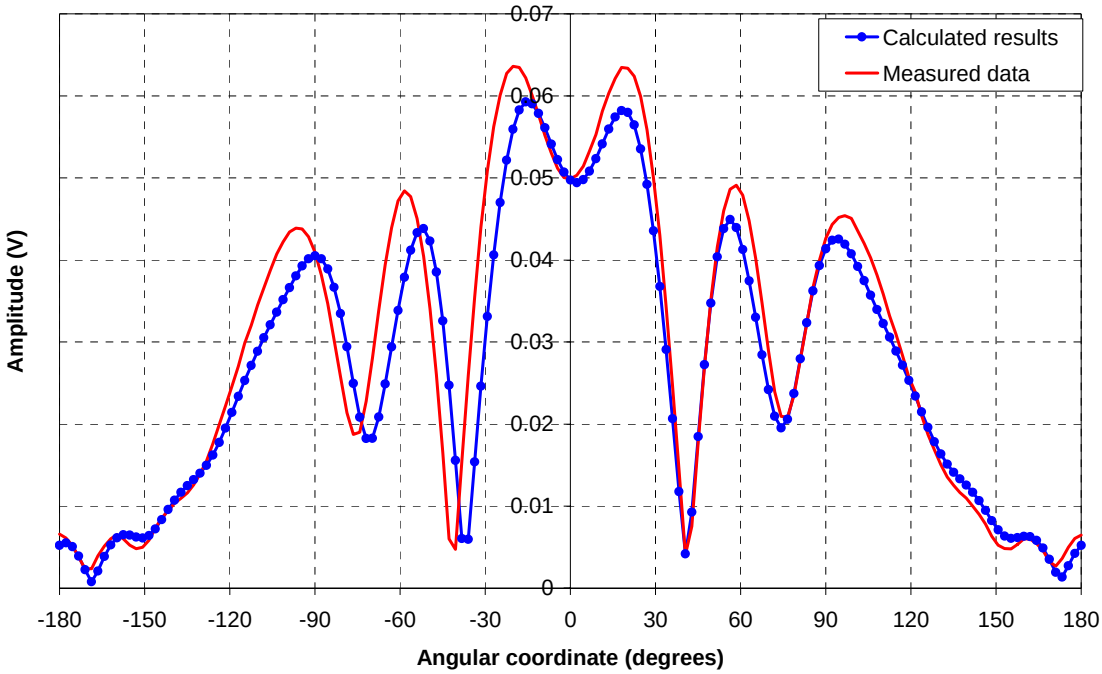

Figure 8.18: Far-field measurement data (red) and calculated results (blue) at 13.9 kHz at a radius of 3.5 m and a height of z = 0.

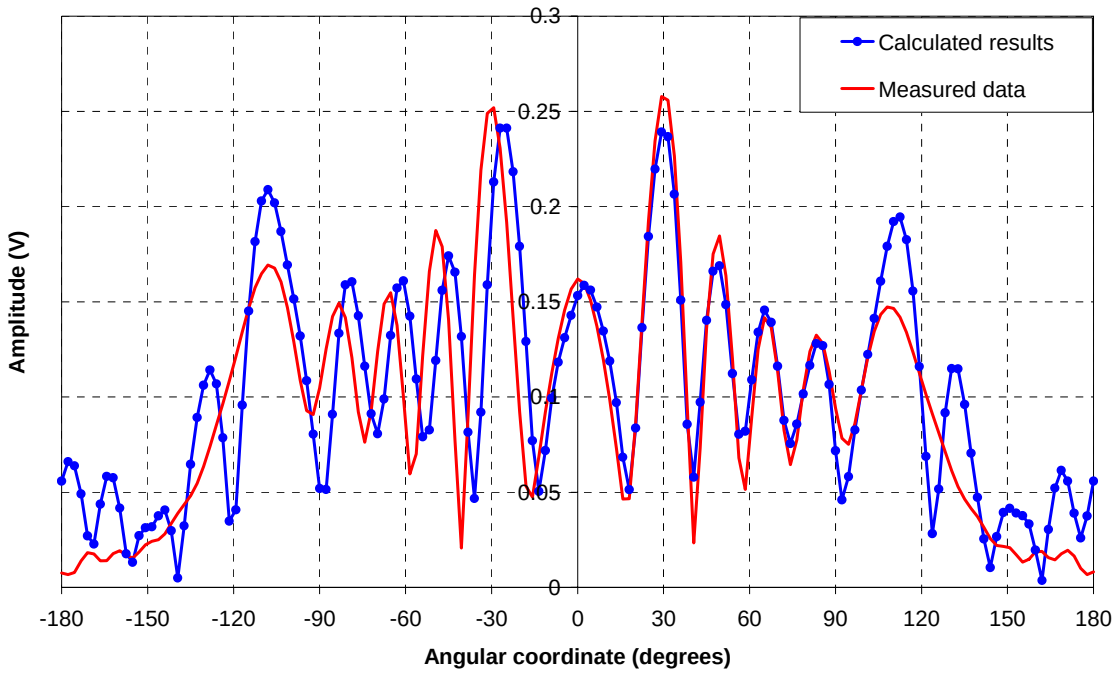

Figure 8.19: Far-field measurement data (red) and calculated results (blue) at 27.5 kHz at a radius of 3.5 m and a height of z = 0.

In order to try and identify why the 27.5 kHz results agreed less well with measurement than the 13.9 kHz results, the 27.5 kHz results were rerun using the maximum possible number of terms in the Fourier series (the original runs had used 51 terms, the maximum possible for the measured data is 159). It was found that the far-field results converged, and that the addition of extra terms beyond the 51st did not affect the far-field results to any significant degree. This is helpful because it demonstrates that the method of truncating the series is satisfactory, but it does not explain the discrepancy between the measured and calculated results.

8.2.6 Next developments

Additional work is in progress to improve the quality of the results calculated at 27.5 kHz. Examination of the input data set suggests that it does not have some properties that the far-field results do have, which may be affecting the quality of the results. In particular, the input data set is less symmetric about $\theta = 0$ than the far-field data, which may be affecting the calculated results.

It would be possible to apply this technique to other systems that obey the Helmholtz equation. A non-homogeneous vector version of the Helmholtz equation occurs in the solution of Maxwell's equations (see section 9.3) for propagating waves, and it may be possible to extend the technique to address such problems.

9. Continuous models in metrology

This section surveys some of the most common continuous models occurring in physics. For each model, the governing equation, domain properties, and common boundary conditions are given. Any special simplifications and solution methods are supplied.

Some simplification methods, such as use of symmetry or axisymmetry can be applied to any model and are not given in each section. Unless stated otherwise, it should be assumed that symmetry and axisymmetry can be applied to problems.

9.1 Stress analysis

9.1.1 Governing equations

The governing equations for a transient stress analysis are

$$\rho \frac{\partial^2 u_i}{\partial t^2} = \frac{\partial \sigma_{ij}(\varepsilon)}{\partial x_j}, \quad i = 1,2,3, \tag{9.1}$$

where the u_i are displacements, the x_i are coordinate directions, the sum over j is between 1 and 3, ρ is the material density, $\{\sigma_{ij}, i, j = 1, 2, 3\}$ is the (symmetric) stress tensor, and

$$\varepsilon_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$$

is the strain tensor. For static analysis, the time derivative vanishes. The domain properties define the model $\sigma_{ij}(\varepsilon)$, which is required for solution of the model. Some common models are given in the next section.

Equations governing wave motion in solid media and natural frequency vibration of structures can be derived from the equations (9.1). These will be discussed in section 9.5.

9.1.2 Domain properties

The domain property required for stress analysis is a model linking the stresses and the strains. The simplest model of domain properties is a simple elastic material. This gives the stresses in terms of the strains as

$$\sigma_{ij} = \frac{E}{(1+\nu)}\varepsilon_{ij} + \delta_{ij}\varepsilon_{kk}\frac{E\nu}{(1+\nu)(1-2\nu)}, \quad i, j = 1,2,3,$$

where E is the Young's modulus and ν is Poisson's ratio. This expression produces a set of governing equations of the form

$$\rho \frac{\partial^2 \mathbf{u}}{\partial t^2} = \frac{E}{2(1+\nu)}\left[\nabla^2 \mathbf{u} + \frac{1}{1-2\nu}\nabla(\nabla.\mathbf{u}) \right].$$

The model can be extended by considering the strains caused by thermal expansion, written as

$$\varepsilon_{ij}^{th} = \alpha_{ij}\Delta T$$

where ε^{th} is the thermal strain tensor, α_{ij} is the matrix of thermal expansion coefficients (usually diagonal) and ΔT is the temperature change relative to some "stress free" temperature, usually defined as either the temperature at which the domain was created or room temperature.

Most materials obey the elastic model when they are first loaded, but usually start to diverge from this as the strain increases. This divergence is usually modelled by splitting the strain into two parts: an elastic strain that is described by the model above, and an inelastic strain that is described by some other model. Elastic strains are recoverable, so that when the load is removed the domain returns to its original shape, whereas inelastic strains are not recoverable. Some materials, in particular rubbers, do not behave like this: instead they exhibit hyperelastic behaviour, where all deformations are recoverable. Models for such behaviour exist [37] but they will not be discussed here.

Almost all models of non-elastic behaviour are nonlinear and so increase the complexity of the problem. In addition, many of them are implicit, and most are only valid for small increments of strain and so the problem must be solved by calculating the cumulative effects of a series of small increments rather than solving in a single step. All of these factors increase the computational effort required to solve such problems.

There is an extensive range of material models for inelastic behaviour. Most apply to a specific type of material and have little application to more general materials. Many of the popular finite element packages offer such models, but only two models of inelastic behaviour will be discussed here: the von Mises' plasticity model and the power law creep model.

Plasticity models generally consist of three parts: a condition that defines when plastic deformation takes place, called the **yield surface**; a rule defining the resulting plastic strains, called the **flow rule**; and if necessary a model of how these entities evolve as the deformation continues, called the **hardening behaviour**. The von Mises' plasticity model is the simplest, and probably most common, model for plastic deformation. It is used to describe plasticity in metals, and has been the starting point for other more complicated plasticity models. Its yield surface and flow rule are

$$\sqrt{3}q - \sigma_Y = 0$$

$$\text{where} \quad q = \sqrt{\tfrac{1}{2}\left(\sigma_{ij} - \tfrac{1}{3}\delta_{ij}\sigma_{kk}\right)\left(\sigma_{ij} - \tfrac{1}{3}\delta_{ij}\sigma_{ll}\right)} = \sqrt{\tfrac{1}{2}s_{ij}s_{ij}}$$

$$d\varepsilon_{ij}^{pl} = d\lambda \frac{\sqrt{3}s_{ij}}{2q}$$

$$\text{where} \quad d\lambda = \frac{d\varepsilon_{ij}^{pl} s_{ij}}{\sigma_Y}$$

where σ_Y is the yield stress (a known value for a given material), $d\varepsilon^{pl}_{ij}$ are the plastic strain increments, s_{ij} is the deviatoric stress, $d\lambda$ is the plastic multiplier, and the summation convention is used for all repeated indices. The hardening behaviour can take a large number of forms, the most common being a yield stress that is dependent on the effective plastic strain.

Another important material model for domain properties is creep, or rate dependent plasticity. Creep models supply a strain rate for a given stress state, and so need to be integrated over time. One of the most commonly-used creep models is the power law creep model, which should be written

$$\dot{\bar{\varepsilon}}_{cr} = \left(Aq^n\left[(m+1)\bar{\varepsilon}_{cr}\right]^m\right)^{1/(m+1)}$$

where A, m and n are known values for the material, $\dot{\bar{\varepsilon}}_{cr}$ is the equivalent creep strain rate, and q is as defined in the plasticity model above. This model is sometimes written in terms of time, but this is not good practice, as explained in section 3.3.

9.1.3 Boundary conditions

Boundary conditions are commonly applied as either displacements, forces or pressures acting on part of the boundary. In general, such conditions are linear and do not increase the model's complexity. However, care needs to be taken when defining such conditions, particularly when forces and pressures are applied to individual points on the discretised boundary.

In general, forces and constraints applied at a single point are not physically realistic and should be avoided where possible. If point forces are the physically reasonable description (for instance if an object is supported on diamond tips), the solution in the immediate vicinity of the force or constraint is likely to be unreliable and should be ignored. In addition, care must be taken, when approximating pressures using point forces, to ensure that the pressure and the forces do the same amount of work when deforming the domain.

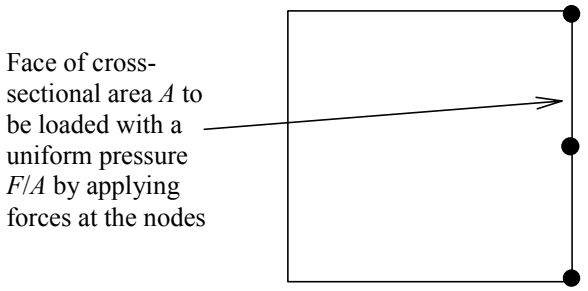

Figure 9.1: Quadratic element to be loaded with a uniform pressure on the boundary indicated.

For example, suppose that a uniform pressure $P = F/A$ is to be applied over the surface of a quadratic element as shown in figure 9.1. Then the nature of the approximation over this face is such that the applied forces need to be $F/6$ at the two corner nodes and $2F/3$ at the central node.

The boundary conditions for some problems do not specify displacements or stresses explicitly. The most common example is contact problems, where the boundary conditions are that two surfaces cannot interpenetrate and that if the two surfaces are in contact, equal and opposite forces act at the point of contact. This is a nonlinear boundary condition. It can be developed further to include models of friction or sticking, where the surfaces cannot break contact once they have touched.

9.1.4 Simplifications and solution methods

The main simplifications used in stress analysis are the assumptions of plane stress and plane strain. These assumptions treat either the stresses (plane stress) or the strains (plane strain) in one of the coordinate directions as being insignificant. These assumptions produce a simplified version of the governing equations, and a problem that can be treated as two-dimensional.

The majority of stress analysis problems are best solved using finite element methods. In particular, models with complicated geometries and nonlinear domain properties or boundary conditions can be considerably easier to build using a finite element approach. Many finite element packages have a wide range of material models, contact algorithms, and specialised elements for handling particular physical problems. Stress analysis problems and particularly engineering problems such as structural analysis were originally one of the main driving forces behind the development of finite element methods.

This does not mean that all stress analysis problems should be tackled using finite elements. Boundary element methods are particularly useful for stress problems involving damage and cracking, because they avoid the problems associated with the singularity of the stresses at the crack tip. Finite difference methods are not usually applied to stress analysis problems because most problems involve complicated geometries that cannot easily be described by an

evenly-spaced grid of points. Additionally, the equations are often non-linear second-order equations in three unknowns, which makes setting up and solving the discretised system of equations difficult.

Some problems can be tackled using semi-analytic methods. The application of such methods requires an understanding of the likely behaviour of the system, but it can be useful in situations where finite element analysis would be inefficient or inaccurate. In particular, semi-analytic methods can be useful for problems featuring long, thin layers that would require prohibitively large meshes [38]. Another example of the application of semi-analytic methods is the derivation of the model for the natural frequencies of a structure. This will be discussed further in section 9.5.

9.2 Heat transfer

9.2.1 Governing equations

Heat flow in a solid obeys the equation

$$\frac{\partial}{\partial t}\left(\rho c_p T\right) = \nabla.\left(\lambda \nabla T\right) + Q,$$

where T is the temperature, ρ is the density, c_p is the specific heat capacity, λ is the thermal conductivity, and Q is the heat generated per unit volume (often, but not always, zero). This equation models conduction. Heat transfer in fluids (liquids or gases) requires consideration of convection since matter is transported as well as heat. The equation governing conduction and convection in a fluid moving at velocity **u** (which may be a function of temperature) is

$$\frac{\partial}{\partial t}\left(\rho c_p T\right) + \mathbf{u}.\nabla\left(\rho c_p T\right) = \nabla.\left(\lambda \nabla T\right) + Q,$$

where the term involving **u** is an advection term. This equation must either have an expression for **u** supplied as one of the domain properties, or be solved in conjunction with the Navier-Stokes equations given in section 9.4.

Thermal models are often used in conjunction with physical models of other processes in order to give accurate domain properties for the other processes. In particular, models that involve melting, boiling, and other such phase changes will almost certainly require a thermal model to describe the phase change. In these models, the term Q can be used to model the energy that is used in the phase change process, or it can be simulated using a modified specific heat capacity (see section 8.1.2.1).

Note that the equation above, like most equations governing physical processes, is independent of the choice of units. This is not the case for all equations associated with thermal models. In particular, the equation describing radiation boundary conditions (which will be described in section 9.2.3) must have its temperature expressed in Kelvin. It is helpful to consider restrictions like this when designing a model, so that if temperatures in Kelvin are required for the boundary conditions then they can be used throughout.

9.2.2 Domain properties

The domain properties required for heat transfer problems are the density, the specific heat capacity, and the thermal conductivity. In general, these properties are temperature-dependent, but many problems only involve a small range of temperatures, and so assuming the properties to be constant may be a reasonable assumption. In cases where temperature dependency does affect the results, the advice given in section 3.3 should be considered. Temperature dependent material property models are empirical, and may need careful application.

The governing equation for conduction assumes perfect thermal contact throughout the domain. If the real problem involves several different objects that are adjacent but not tightly bonded, this assumption may not be valid. If this is the case, it may be worth introducing dummy layers of material that simulate the imperfect contact. The dummy layer should be of low density and specific heat capacity, so that it does not remove energy from the real physical problem, and should be of low thermal conductivity so that it simulates a poor thermal bond. "Low" means "low in comparison to the other materials used in the simulation, without being so low as to cause numerical problems". Values that are too low can result in severe restrictions on the range of time-steps that are stable.

The main use of models with dummy layers is to investigate how the quality of the thermal bond affects the results. As the thermal conductivity of the dummy gap is not known, values varying over different orders of magnitude can be tried. If the results vary widely, then the quality of the thermal bond is important.

9.2.3 Boundary conditions

Boundary conditions for heat transfer problems can be split into three types:

- Conduction
- Convection
- Radiation

Conduction conditions are mainly used at interfaces between two different regions with different thermal properties. These conditions ensure that the temperature and the heat flux are continuous across such interfaces, so that if the two regions are Ω_1 and Ω_2 and the points $\mathbf{x}_1$ and $\mathbf{x}_2$ lie in Ω_1 and Ω_2 respectively, then

$$\lim_{\mathbf{x}_1 \to \mathbf{x}} \{T(\mathbf{x}_1, t)\} = \lim_{\mathbf{x}_2 \to \mathbf{x}} \{T(\mathbf{x}_2, t)\},$$

$$\lim_{\mathbf{x}_1 \to \mathbf{x}} \{\lambda(\mathbf{x}_1)\nabla T(\mathbf{x}_1, t).\mathbf{n}\} = \lim_{\mathbf{x}_2 \to \mathbf{x}} \{\lambda(\mathbf{x}_2)\nabla T(\mathbf{x}_2, t).\mathbf{n}\},$$

for a point $\mathbf{x}$ lying on the interface between Ω_1 and Ω_2, and $\mathbf{n}$ the unit normal to the interface at the point $\mathbf{x}$. The main application is for regions that adjoin a perfect insulator, where it must be the case that

$$\lambda\nabla T.\mathbf{n} = 0,$$

because a perfect insulator has zero thermal conductivity. These conditions ensure that the governing equation is a valid statement throughout the domain.

Convection conditions are used to describe the heat transfer between a solid and a fluid, so they are useful for describing heat loss to an air environment. In many cases the fluid is not modelled in detail, and it is assumed that the fluid's temperature is known. The most general form of convection conditions is

$$\lambda\nabla T(\mathbf{x}, t).\mathbf{n} = h(T(\mathbf{x}, t), T_{fl}(\mathbf{x}, t))(T_{fl}(\mathbf{x}, t) - T(\mathbf{x}, t)), \qquad \mathbf{x} \in \partial\Omega,$$

where T_{fl} is the fluid temperature which may be a function of space and time, and h is a heat transfer coefficient. The coefficient h can be a function of a number of properties of the fluid flow, and it is often expressed as a function of the fluid velocity. Forced convection is convection in the presence of a fluid flow that would occur even if the temperature difference was not present, whereas free convection is convection involving a fluid flow driven solely by the temperature difference between the surface and the fluid. Forced convection generally produces faster heat transfer than free convection.

The convection coefficient h is usually determined experimentally. It is usual to assume that the coefficient depends on some power of the difference between T and T_{fl} so that

$$h\left(T(\mathbf{x},t),T_{fl}(\mathbf{x},t)\right)=h_0\left(T_{fl}(\mathbf{x},t)-T(\mathbf{x},t)\right)^{\alpha},\qquad \mathbf{x}\in\partial\Omega,$$

where h_0 may be dependent on characteristics of the fluid flow, and α may be zero. Some expressions exist for h_0, but these are usually for a specific geometry such as a flat wall. Where these expressions do exist they are usually given in terms of the Reynolds, Prandtl, and Nusselt numbers, which are properties of fluid flows, and they generally assume $\alpha = 0$. In general, if convection is a significant factor then an experimental determination of h_0 is preferable wherever possible.

Radiative transfer occurs either between the surface of an object and its surroundings, or between two surfaces. Radiation becomes a significant factor when the temperature difference between the radiating objects becomes large. The heat flux between two objects depends on their temperatures, the surface areas, the surface emissivities, and how much of one surface can be "seen" from the other (called the view factor F with $0 \le F \le 1$). The idea of view factors is illustrated in figure 9.2.

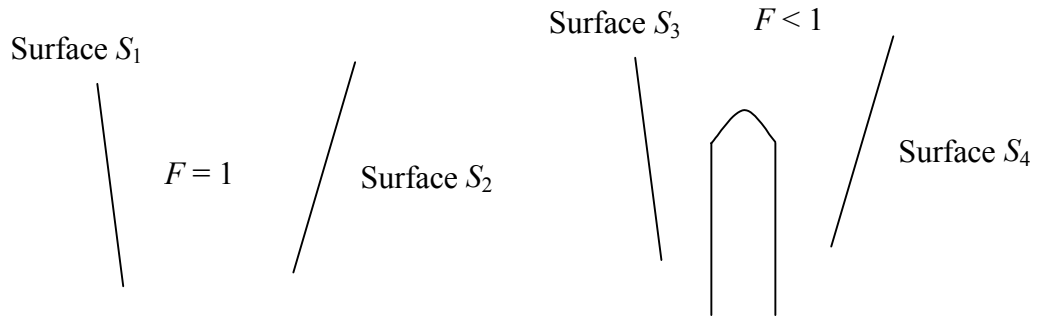

Figure 9.2: Surfaces S_1 and S_2 can see all of one another, so $F = 1$. Surfaces S_3 and S_4 cannot due to the intervening object, so $F < 1$.

For radiation from surface 1 to surface 2 with a view factor F_{12}, the boundary condition is given by

$$\lambda\nabla T_1(\mathbf{x},t).\mathbf{n}=\left(\frac{1-\varepsilon_1}{\varepsilon_1}+\frac{1-\varepsilon_2}{\varepsilon_2}\frac{A_1}{A_2}+\frac{1}{F_{12}}\right)^{-1}A_1\sigma\left(T_1^4-T_2^4\right)\qquad \mathbf{x}\in\partial\Omega,\qquad (9.2)$$

where T_i denotes the temperature of surface i, A_i is the surface area of surface i, ε_i is the emissivity of surface i, and σ is the Stefan-Boltzmann constant, which has a value of $5.67 \times 10\text{-}8$ W m^{-2} K^{-4}. If the surface is radiating to the atmosphere, it is assumed that $F_{12} = 1$ and that $\varepsilon_i = 1$ so that equation (9.2) above becomes

$$\lambda\nabla T_1(\mathbf{x},t).\mathbf{n}=\varepsilon_1 A_1\sigma\left(T_1^4-T_{amb}^4\right)\qquad \mathbf{x}\in\partial\Omega,$$

where T_{amb} is the ambient temperature.

9.2.4 Simplifications and solution methods

The most useful simplification is linearisation of the domain properties and of the boundary conditions. For small changes in temperature, it is likely that the majority of domain properties can be regarded as constant, and it is possible to linearise most forms of boundary conditions. For example, if we consider the radiation conditions

$$\begin{aligned}\lambda\nabla T_1(\mathbf{x},t).\mathbf{n} &= \varepsilon_1 A_1\sigma\left(T_1^4 - T_{amb}^4\right) \qquad \mathbf{x}\in\partial\Omega,\\ &= \varepsilon_1 A_1\sigma\left(T_1 - T_{amb}\right)\left(T_1^3 + T_1^2 T_{amb} + T_1 T_{amb}^2 + T_{amb}^3\right)\\ &\approx 4\varepsilon_1 A_1\sigma\left(T_1 - T_{amb}\right)T_{amb}^3,\end{aligned}$$

for $T_1 - T_{amb} << T_{amb}$. Linearisation of boundary conditions means that the matrix equations are linear and so can be solved more quickly.

If the equations are linear and the initial temperature distribution is known, it may be advantageous to reformulate the problem in terms of the temperature rise $\Delta T(\mathbf{x}, t) = T(\mathbf{x}, t) - T(\mathbf{x}, 0)$. The governing equation becomes

$$\frac{\partial}{\partial t}\left(\rho c_p \Delta T\right) = \nabla.\left(k\nabla\left[\Delta T\right]\right) + \tilde{Q}, \quad \mathbf{x}\in\Omega,$$

where

$$\tilde{Q} = Q + \nabla.\left(k\nabla\left[T(\mathbf{x},0)\right]\right),$$

and the boundary conditions are

$$\alpha(\mathbf{x})\Delta T + \beta(\mathbf{x})\frac{\partial \Delta T}{\partial n} = \gamma(\mathbf{x}) - \alpha(\mathbf{x})T(\mathbf{x},0) - \beta(\mathbf{x})\frac{\partial \Delta T(\mathbf{x},0)}{\partial n}.$$

The main advantage of this reformulation is that it is more suitable for modelling high-temperature problems whose results lie within a small range of values. The formulation as a problem in terms of temperature difference will produce more accurate results when numerical methods are applied.

Heat transfer problems can be solved using finite element analysis, finite difference methods, finite volume methods, or Laplace transforms. Laplace transforms are a semi-analytic integral transform. Given a function f, the Laplace transform is

$$L\left[f(t)\right](s) = F(s) = \int_0^\infty f(t)e^{-st}dt,$$

and its inverse is

$$f(t) = \frac{1}{2\pi i}\int_{\gamma-i\infty}^{\gamma+i\infty} F(s)e^{st}ds,$$

where γ is chosen so that all of the singularities of F are to the left of the line of integration in the complex plane.

This transform is particularly useful for problems where the equations and boundary conditions are linear, the material properties are constant, and the initial conditions are a uniform temperature distribution so that $T(\mathbf{x}, 0) = T_0$. These restrictions makes the transformation to a “temperature rise” problem as outlined above particularly simple:

$$\frac{\partial}{\partial t}(\Delta T) = \frac{k}{\rho c_p}\nabla^2\Delta T, \quad \mathbf{x}\in\Omega$$

$$\alpha(\mathbf{x})\Delta T + \beta(\mathbf{x})\frac{\partial \Delta T}{\partial n} = \gamma(\mathbf{x}) - \alpha(\mathbf{x})T_0, \quad \mathbf{x}\in\partial\Omega$$

$$\Delta T(\mathbf{x},0) = 0, \quad \forall \mathbf{x} \in \Omega.$$

The transformation of the problem then becomes

$$-sL[\Delta T] = \frac{k}{\rho c_p} \nabla^2 L[\Delta T], \quad \mathbf{x} \in \Omega,$$

$$\alpha(\mathbf{x}) L[\Delta T] + \beta(\mathbf{x}) \frac{\partial L[\Delta T]}{\partial n} = \frac{\gamma(\mathbf{x}) - \alpha(\mathbf{x}) T_0}{s}, \quad \mathbf{x} \in \partial\Omega, s > 0$$

which is easier to solve because the time dependency has been removed.

9.3 Electromagnetics

9.3.1 Governing equations

The most general form of the equations governing electromagnetic fields, called the Maxwell equations, is

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad \nabla \times \mathbf{H} = \left(\frac{\partial \mathbf{D}}{\partial t} + \mathbf{J} \right), \qquad (9.3)$$
$$\nabla . \mathbf{D} = \rho, \qquad \nabla . \mathbf{B} = 0,$$

where **E** is the electric field measured in Vm^{-1}, **B** is the magnetic field measured in T, $\mathbf{H} = \mathbf{B}/\mu$ where μ is the magnetic permeability measured in NA^{-2}, $\mathbf{D} = \varepsilon \mathbf{E}$ where ε is the electric permittivity, measured in Fm^{-1}, **J** is the vector current density measured in Am^{-2}, and ρ is the charge density in Cm^{-3}. The units are specified here because some choices of units eliminate the variables ε and μ. The equations are commonly written in integral form by applying the divergence theorem and Green's theorem to the integral of the above forms over a closed volume V or a closed surface A:

$$\oint_{\partial A} \mathbf{E} \cdot d\mathbf{s} = -\int_A \frac{\partial \mathbf{B}}{\partial t} \cdot d\mathbf{A} \quad \oint_{\partial A} \mathbf{H} \cdot d\mathbf{s} = \int_A \left(\frac{\partial \mathbf{D}}{\partial t} + \mathbf{J} \right) \cdot d\mathbf{A}$$
$$\oint_{\partial V} \mathbf{D} \cdot d\mathbf{A} = \int_V \rho \, dV \qquad \oint_{\partial V} \mathbf{B} \cdot d\mathbf{A} = 0.$$

The integral form can be more convenient for solution if finite volumes are to be used.

The equations are often simplified by assuming steady-state conditions (such problems are called electrostatic and/or magnetostatic). In the steady state,

$$\nabla \times \mathbf{E} = 0,$$

and so **E** can be written as

$$\mathbf{E} = -\nabla \varphi$$

where φ is the electric potential. Similarly, since

$$\nabla . \mathbf{B} = 0,$$

B can be written as

$$\mathbf{B} = -\nabla \times \mathbf{A}$$

where **A** is the magnetic vector potential. These simplifications mean that

$$\nabla.(\varepsilon\nabla\varphi) = -\rho, \quad \nabla\times\left(\frac{\nabla\times\mathbf{A}}{\mu}\right) = \mathbf{J}.$$

9.3.2 Domain properties

The domain properties that must be defined are the electric permittivity ε and the magnetic permeability μ. Both of these quantities can be complex-valued. The quantities are usually normalised and written

$$\mu = \mu_r\mu_0, \quad \varepsilon = \varepsilon_r\varepsilon_0,$$

where μ_r is the relative magnetic permeability and ε_r is the relative electric permittivity. The absolute values of μ_0 and ε_0 are $4\pi \times 10^{-7}$ N A^{-2} and $1/(c^2\,\mu_0)$ F m^{-1} respectively, where c is the speed of light *in vacuo*. The values of the relative quantities can be temperature dependent and frequency dependent, and the materials can be anisotropic.

In a conducting or semi-conducting material, it is common to define the conductivity, σ, such that

$$\mathbf{J} = \sigma\mathbf{E},$$

where σ is given in siemens per metre. This is a general form of Ohm's Law (simply written as resistance = electric potential difference / current). In some materials the relationship between **J** and **E** is more complicated.

9.3.3 Boundary conditions

Boundary conditions for most electromagnetic problems are developed by considering the continuity of the electric and magnetic fields across the boundary between two media, and then making assumptions based on the medium on the outside of the boundary.

At an interface between two media, the electric potential φ is always continuous (because otherwise the electric field would be infinite at the discontinuity). This means that if the potential just outside the boundary is known (e.g. due to the presence of an electrode) then φ can be set on the boundary.

This continuity also imposes conditions on the electric field:

$$(\mathbf{E}_1 - \mathbf{E}_2)\times\mathbf{n} = \mathbf{0},$$
$$(\mathbf{D}_1 - \mathbf{D}_2)\cdot\mathbf{n} = \sigma_f,$$

where **n** is the outward normal to the interface, subscripts 1 and 2 denote the two media with medium 1 being the domain and medium 2 being the outside, **D** is as defined in section 9.3.1, and σ_f is the free surface charge density on the interface. The equivalent conditions on the magnetic field are

$$(\mathbf{H}_1 - \mathbf{H}_2)\times\mathbf{n} = \mathbf{0},$$
$$(\mathbf{B}_1 - \mathbf{B}_2)\cdot\mathbf{n} = 0.$$

It can be shown that the electric field inside a perfect conductor is uniformly zero. Hence if the domain of interest is surrounded by a perfect conductor then the boundary conditions become

$$\mathbf{E}_1\times\mathbf{n} = \mathbf{0},$$
$$\mathbf{D}_1\cdot\mathbf{n} = \sigma_f.$$

For wave transmission problems in an infinite medium, some form of condition at infinity is required. It is often the case that wave transmission problems are described by the vector form of the Helmholtz equation, and a modified form of the Sommerfeld radiation condition can be applied.

9.3.4 Simplifications and solution methods

If it is assumed that the current density and charge density are zero, then a form of the wave equation can be derived for **E** or **B**. If it is then assumed that the solution is periodic in time, then a resonant frequency problem is obtained. These types of problem will be discussed further in section 9.5. Additionally, the steady-state form of the equations where the problem is reduced to solution for the electric potential and the magnetic vector potential was mentioned in section 9.3.1.

If the electric field can be written as a sum of terms of the form $\mathbf{E}_0 e^{i\omega t}$ for different values of ω, then it can be shown that

$$\nabla \times \left(\tfrac{1}{\mu} \nabla \times \mathbf{E}_0\right) - \omega^2 \mathbf{E}_0 = -i\omega \mathbf{J}$$

for each value of ω. This form is of most use when the boundary conditions for a problem can be written as the sum of a small number of terms. This type of problem is called a frequency domain problem, whereas the more general problems governed by the equations (9.3) are time domain problems.

The most suitable solution method for a given electromagnetic problem depends on its physical extent, and whether it is a time domain problem or a frequency domain problem. Useful advice regarding the most suitable method for different types of problem is given at [39].

Probably the most common solution methods for time domain problems are the Finite Difference Time Domain (FDTD) methods. These methods approximate the time differential by using a central difference method and use either finite volumes or finite differences in the physical domain. Some sources call the FDTD method using finite volumes in the physical domain FVTD. In each case the electric and magnetic fields are calculated on a pair of separate but interleaved grids.

Finite element methods are becoming more popular for electromagnetic problems now that computer power has made it possible for three-dimensional problems to be tackled. In general, electromagnetic problems are nonlinear and require solutions of vector equations, making FE more difficult to apply. Additionally there have been problems with existence of spurious solutions, but many of these issues have now been addressed. FE is particularly suited to problems in a finite domain where the fields at the boundary are well understood, and can be used in the time or frequency domain.

Another popular method is the method of moments, the most common example of which is the boundary element method. This method is more commonly used for problems with reasonably simple geometries and/or infinite domains.

In addition to these methods, there are other specialised methods that are suitable for particular classes of problem. Examples include the transmission line matrix method [40] and the generalized multipole technique [41], both of which can also be applied to acoustic problems. The transmission line matrix method solves problems in the time domain by linking mesh points with transmission lines and propagating the excitations of the fields along these lines. The generalized multipole technique solves problems in the frequency domain using basis functions derived from spherical wave fields corresponding to multipole sources away from the boundary.

9.4 Fluid flow

9.4.1 Governing equations

The most general form of the equations governing fluid flow are derived by considering conservation of mass, momentum and energy in a fluid whose pressure p, velocity $\mathbf{v} = \{v1, v2, v3\}$, temperature T and density ρ are varying. For a viscous compressible fluid, the equations are of the form

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i}(\rho v_i) = 0, \qquad (\text{cons}^{\text{n}} \text{ mass}),$$

$$\frac{\partial}{\partial t}(\rho v_i) + \rho v_j \frac{\partial v_i}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j} + \rho F_i, \qquad (\text{cons}^{\text{n}} \text{ momentum}),$$

$$\frac{\partial}{\partial t}(\rho E) + \frac{\partial}{\partial x_i}([\rho E + p]v_i) = \lambda \frac{\partial^2 T}{\partial x_i^2} + \frac{\partial}{\partial x_i}(\tau_{ij} v_j) + \rho F_i v_i, \qquad (\text{cons}^{\text{n}} \text{ energy}),$$

$$\tau_{ij} = \mu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i}\right) + \mu' \delta_{ij} \frac{\partial v_k}{\partial x_k},$$

$$E = c_p T + \tfrac{1}{2} v_j v_j - \frac{p}{\rho},$$

where the F_i represent the external accelerations that act on the fluid (e.g. gravity), μ and μ' are the viscosity coefficients (two are needed for a compressible material), λ is the thermal conductivity, and c_p is the specific heat capacity at constant pressure. The τ_{ij} are shear stresses, and E is the energy per unit mass.

For solution, these equations require a relationship between p and ρ. This is specified as a domain property and discussed in the next section. The most common assumption is that ρ is constant, which has a significant simplifying effect on the equations. If the density is constant, called an incompressible fluid, the equations become

$$\frac{\partial v_i}{\partial x_i} = 0, \qquad (\text{cons}^{\text{n}} \text{ mass}),$$

$$\frac{\partial v_i}{\partial t} + v_j \frac{\partial v_i}{\partial x_j} = -\frac{1}{\rho}\frac{\partial p}{\partial x_i} + \frac{\mu}{\rho}\frac{\partial^2 v_i}{\partial x_j^2} + F_i, \qquad (\text{cons}^{\text{n}} \text{ momentum}),$$

$$\frac{\partial}{\partial t}(\rho c_p T - p) + v_i \frac{\partial}{\partial x_i}(\rho c_p T - p) = \lambda \frac{\partial^2 T}{\partial x_i^2} + \tau_{ij} \frac{\partial v_j}{\partial x_i}, \qquad (\text{cons}^{\text{n}} \text{ energy}),$$

$$\tau_{ij} = \mu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i}\right),$$

which is a simpler form.

9.4.2 Domain properties

In general fluid properties can be pressure, velocity (strain rate) and temperature dependent. The main properties that must be defined for the most general fluid model are the density, the viscosity (kinematic or dynamic), the thermal conductivity and the specific heat capacity, but not all of these are needed if simplifying assumptions such as steady flow and a uniform temperature are made.

A particularly common form of flow is driven by the Boussinesq approximation, which models flow driven by gravity and the effects of temperature on density, writing $\rho = \rho_0(1 - \alpha\Delta T)$, and adding the resulting extra body force to the equation of conservation of momentum, giving

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}.\nabla)\mathbf{v} = -\frac{1}{\rho_0}\nabla p + \nu\nabla^2\mathbf{v} - \mathbf{g}\alpha\Delta T.$$

Turbulent flows occur between two layers with differing velocities, particularly where one layer is stationary (e.g. a wall). The resulting motion has features at a range of different scales and exhibits complicated rotational behaviour. Turbulent behaviour is commonly modelled using an averaged system of equations (averaging over time or mass) that feature an extra term in the stress tensors. This extra term, called a turbulent stress tensor and denoted τ_{ij}*, is then given as a domain property by writing it in terms of other physical quantities such as velocity and length. These models often lead to further unknown quantities requiring further equations (in particular the most common turbulence model, the K-ε model, requires equations for the unknowns K and ε).

9.4.3 Boundary conditions

Probably the most common boundary conditions at walls for fluids problems are the no-flow condition,

$$\mathbf{v}.\mathbf{n} = 0,$$

where **n** is the normal to the boundary, and the no-slip condition

$$\mathbf{v}.\mathbf{t} = 0,$$

where **t** is a tangent to the surface. These conditions define zero flow through the wall (no flow condition) and along the wall (no slip condition). Sometimes the no slip condition is replaced by a friction condition.

Inlets usually have a specified mass flow rate and density, and outlets often specify pressure. If there is only one outlet, the pressure there can usually be set to zero since pressure is a relative quantity in most fluid models.

More complicated boundary conditions occur when there is a free boundary, so that the level of the fluid becomes one of the unknown quantities. Such problems lead to a set of implicit equations that can be difficult to solve numerically.

9.4.4 Simplifications and solution methods

The most common method of simplifying fluid flow problems is to make assumptions about the fluid's behaviour that will simplify the governing equations. The simplest flows are those that are assumed to be inviscid, incompressible, and isothermal. This reduces the governing system to the Euler equations, which are considerably easier to solve.

Care needs to be taken when simplifying the domain of fluid flow problems. Some flows, particularly turbulent flows, involve three-dimensional phenomena that cannot be modelled adequately in a two-dimensional domain.

Many fluid flow problems are solved using finite volume methods, although finite element and finite difference techniques can be used. A good introduction to the use of the different techniques for different classes of problem is given in Chung [42].

Non-dimensionalisation is commonly used in fluid flow problems: flows are often described by several non-dimensional parameters (Reynolds number, Prandtl number, Nusselt number, etc.) and a suitable solution method or simplification is chosen on the basis of this description.

9.5 Waves and resonant frequencies

Wave equations and resonant frequency problems occur in a number of different areas of science. As they are generic classes of equation, the domain properties and boundary condition information should be taken from the physical problem, on the assumption that the conditions are consistent with the assumptions made in deriving the governing equation. Many wave transmission problems occur in infinite domains. The advice in section 3.2.3 regarding infinite domains should be considered when solving wave transmission problems.

A system undergoing some form of excitation is most likely to respond at its resonant frequencies. They are the frequencies at which the system response can grow uncontrollably if the motion is not damped. This can be a great disadvantage in measurement equipment.

9.5.1 Governing equations

Wave equations are equations of the form

$$\frac{\partial^2 f(\mathbf{x},t)}{\partial t^2} = c^2 \nabla^2 f(\mathbf{x},t),$$

where c is the wave speed. In one spatial dimension, the solution can take the form $f(x, t) = g(x - ct)$, and in three dimensions if $R = |\mathbf{x}|$ a solution can be written $f(\mathbf{x}, t) = g(R - ct)/R$. It is common to seek solutions involving a term of the form $g(\mathbf{k}.\mathbf{x} - ct)$ where $\mathbf{k}$ is some vector: these are travelling wave solutions.

Wave equations can be derived from many of the equations mentioned in the preceding sections. Examples of wave equations include

$$\rho \frac{\partial^2}{\partial t^2}(\nabla.\mathbf{u}) = \frac{E(1-\nu)}{(1-2\nu)(1+\nu)} \nabla^2(\nabla.\mathbf{u}),$$

from stress modelling, where $\mathbf{u}$ is the displacement, E is the Young's modulus, υ is the Poisson's ratio, and ρ is the density,

$$\rho \frac{\partial^2 \mathbf{u}}{\partial t^2} = \nabla(K \nabla.\mathbf{u}),$$

from fluid motion, where $\mathbf{u}$ is the velocity, ρ is the density, and K is the bulk modulus of the fluid, and

$$\varepsilon_0 \mu_0 \frac{\partial^2 \mathbf{E}}{\partial t^2} = \nabla^2 \mathbf{E}$$

from electromagnetics in vacuo, where E is an electric field and ε_0 and μ_0 are the electric permittivity and magnetic permeability of free space respectively.

Resonant frequencies (also known as natural frequencies) are similar to eigenvalues of matrices (in fact, the discretised versions of such problems result in a matrix eigenvalue problem). Resonant frequency solutions (called modes) are the solutions to a governing equation that are periodic in time. If a problem involves excitation of an object a some random set of frequencies, after a period of initially chaotic behaviour the response of the object will be made up of a weighted superposition of its resonant frequency modes.

Resonant frequency equations are derived by assuming a solution of the form $f(\mathbf{x})e^{i\omega t}$ exists and substituting it into the governing equation. This produces a modified form of the governing equation featuring $f(\mathbf{x})$ and ω as unknowns. The boundary conditions are then used to derive an equation for ω. This solution will have an infinite number of solutions, but generally only the first few are required.

Examples of resonant frequency equations include:

From stress modelling,

$$\frac{E}{2(1+\nu)}\nabla^2\mathbf{u}+\frac{E}{2(1+\nu)(1-2\nu)}\nabla(\nabla.\mathbf{u})+\rho\omega^2\mathbf{u}=\mathbf{0},$$

where **u** is the displacement, E is the Young's modulus, υ is the Poisson's ratio, ρ is the density, and ω is the resonant frequency. The solution produces the frequencies ω and the displacements **u**.

From electromagnetics,

$$\nabla\times(\nabla\times\mathbf{B}_0)+\mu\varepsilon\omega^2\mathbf{B}_0=\mathbf{0},$$

where $\mathbf{B}_0$ is the magnetic field strength, μ is the magnetic permeability, ε is the electric permittivity, and ω is the resonant frequency. The solution produces the frequencies ω and the field strength $\mathbf{B}_0$.

9.6 Acoustics

9.6.1 Governing equations

This section of the guide is concerned with linear acoustic wave propagation in fluids. More general wave equations, including the equation for wave transmission in solids, were discussed in the previous section.

The simplest form of acoustic waves in homogeneous isotropic fluids obey the wave equation

$$\frac{\partial^2\Psi(\mathbf{x},t)}{\partial t^2}=c^2\nabla^2\Psi(\mathbf{x},t),$$

where c is the wave speed and Ψ is the scalar time-dependent velocity potential, such that if **V** is the particle velocity within the fluid, then

$$\mathbf{V}(\mathbf{x},t)=\nabla\Psi.$$

As well as this equation, there are more complicated forms of nonlinear wave equation that include diffusion and absorption by the acoustic medium. These will not be discussed in detail here, but are covered elsewhere [43].

The linear wave equation can then be simplified further by seeking periodic solutions of the form

$$\Psi=\sum_{n=1}^{N}\varphi_n(\mathbf{x})e^{-i\omega_n t}$$

where the $\omega_n = 2\pi f_n$ are known (angular) frequencies, where f_n is the frequency in Hertz. Then each of the φ_n obeys (dropping the subscripts)

$$\nabla^2\varphi+k^2\varphi=0,$$

where $k = 2\pi f / c$. This is the Helmholtz (reduced wave) equation, which can be obtained from any wave equation in the same fashion.

9.6.2 Domain properties

Acoustic problems generally require definition of two domain properties: the speed of sound and the density. The density is generally required to link the model results to a measurable

quantity such as pressure. Using the same notation as in section 9.6.1, if p is the sound pressure and ρ is the density, then

$$p(\mathbf{x}) = i\rho\omega\varphi(\mathbf{x}).$$

Both density and sound speed can be dependent on temperature and frequency, in which case the governing equation may be more complicated than has been stated.

Some sources use bulk modulus B and density to define the properties of fluid domains. The speed of sound is then given by

$$c = \sqrt{B/\rho}.$$

9.6.3 Boundary conditions

Acoustic problems fall into two classes: interior problems, where the pressure distribution inside some closed surface is sought, and exterior problems, which seek the pressure outside of a closed surface. Both of these problems require boundary conditions on the closed surface of the form

$$\alpha(\mathbf{x})\varphi + \beta(\mathbf{x})\frac{\partial\varphi(\mathbf{x})}{\partial n} = \gamma(\mathbf{x}),$$

where α, β, and γ are known functions. Frequently these conditions are derived from measurement and have either α or β being zero, but surfaces with different reflection characteristics may need a different expression.

In addition to these conditions, exterior problems need a condition at infinity, since one of their boundaries is effectively infinite. Physical considerations usually dictate that all radiated and scattered waves are outgoing, called the Sommerfeld radiation condition. The condition is written

$$\lim_{R\to\infty} R^{1/2}\left\{\frac{\partial\varphi}{\partial R} - ik\varphi\right\} = 0$$

in two dimensions and

$$\lim_{R\to\infty} R\left\{\frac{\partial\varphi}{\partial R} - ik\varphi\right\} = 0$$

in three dimensions, where R is $|\mathbf{x}|$ in both cases. Some sources use a + sign instead of a – sign in the condition. This choice usually means that the formulation has been obtained by writing

$$\Psi = \sum_{n=1}^{N} \varphi_n(\mathbf{x})e^{i\omega_n t}$$

instead of the form given in section 9.6.1.

9.6.4 Simplifications and solution methods

As was mentioned in section 9.6.3, exterior acoustic problems involve a boundary condition at infinity. This requirement can be problematic when applying numerical techniques to acoustic problems. A discussion of some methods that avoid the problem is given in section 3.2.3.

As with electromagnetic problems, the type of solution method available for acoustic problems strongly depends on the type of problem.

For wave transmission problems there are several different analytic and semi-analytic methods in common use, including the angular plane-wave spectrum method (decomposes a wave front into its plane wave spatial frequency components and then propagates these in the frequency domain), ray, image and beam tracing (commonly used in architectural acoustics), and high-frequency asymptotic methods (often used to study non-destructive testing problems). In general these methods rely on the problem domain having certain properties (e.g. planarity, axisymmetry, etc.). The boundary element method can solve a much more general class of problems, and includes the condition at infinity automatically (see section 10.2 for details).

For problems where fluid-structure interaction is a concern, finite element methods are a convenient way to simulate all parts of the domain. The methods can be used for frequency or time-dependent problems. Many common FE packages now have acoustics simulation options.

Problems that are solved by decomposing as a sum of periodic components may require a full transient solution, using either finite differences or finite elements. In particular, if the problem has been derived from a nonlinear wave equation then it is likely to require specially written software. Finite difference schemes are commonly used for such problems because it is often simpler to implement a finite difference solution to a nonlinear problem than it is to use a finite element method.

9.7 Ordinary differential equations

Ordinary differential equations (odes) do not commonly occur as the general forms of governing equations for physical situations. The main reason for this is that the majority of continuous models represent two- or three-dimensional problems, and the quantities of interest vary throughout the domain so the changes with respect to all coordinate directions (and sometimes time) must be considered.

Odes are more commonly produced when semi-analytic methods are used to simplify a partial differential equation. They also occur when techniques such as separation of variables are used to look for an analytic solution or a Green's function. Some problems can be simplified to odes by assuming that only one of the coordinate direction is significant. For instance if a thin beam is being bent, only the variation of deformation along its length needs to be considered.

This is not to say that there are no odes used to define continuous models. Chemistry uses odes to describe how concentrations of different substances vary with time. The equations of motion of a particulate system are a set of coupled odes. Quantum mechanics applications often involve odes. Other examples occur in population dynamics, molecular modelling, and electrical circuit analysis.

All odes have a linear domain, which may be infinite, and boundary conditions are point values of the functions, the derivatives or a combination of the two at the limits of the boundary. If the boundary is infinite, the boundary conditions may be given as a limiting value for the solution as the coordinate tends to infinity.

Only a few of the most common odes will be given here, but there are useful resources online that give fuller lists [44]. The following odes occur in a wide range of physical applications. The Airy equation,

$$y'' = xy,$$

occurs in wave propagation, Bessel's equation

$$x^2 y'' + xy' + \left(x^2 - \nu^2\right)y = 0,$$

the modified version

$$x^2 y'' + xy' - \left(x^2 + \nu^2\right)y = 0,$$

and the spherical version

$$x^2 y'' + 2xy' + \left(x^2 - \nu(\nu+1)\right)y = 0,$$

often occur when looking for separable solutions for equations in cylindrical and spherical coordinate systems, and the associated Legendre equation

$$\frac{d^2 y}{d\theta^2} + \frac{\cos\theta}{\sin\theta}\frac{dy}{d\theta} + \left(m(m+1) - \frac{n^2}{\sin^2\theta}\right)y = 0$$

occurs in separable solutions to Helmholtz's equation in spherical coordinate systems (common in acoustics and electromagnetics). Many initially unpromising equations can be transformed into these forms by suitable coordinate transformations.

These equations have solutions that are defined in terms of special functions that can either be written as infinite sums, or can be defined in terms of recurrence relations. Many resources on mathematical methods and special functions [45, 46] describe Bessel functions and Legendre functions more fully. Bessel functions are particularly useful since the solutions to many other common equations can be written as a sum of Bessel functions.

Application-specific odes include:

- $m_j \ddot{\mathbf{x}}_j = \sum_{i=1}^{N} \frac{F_{i,j}\left(\mathbf{x}_i - \mathbf{x}_j\right)}{\left|\mathbf{x}_i - \mathbf{x}_j\right|} + \mathbf{F}_{gen,j}, \quad j = 1,2,...,N,$ describing the motion of an N-particle system, where m_j and $\mathbf{x}_j$ are the mass and position of the j^{th} particle, $F_{i,j}$ is the magnitude of the force caused by the interaction between the i^{th} and j^{th} particles (assumed to act along the line connecting the particles and using an appropriate sign convention), and $\mathbf{F}_{gen,j}$ represents any other forces present acting on the j^{th} particle,
- $\frac{d\left[A_j\right]}{dt} = k_j \left[A_1\right]^{n_{1,j}} \left[A_2\right]^{n_{2,j}} ... \left[A_N\right]^{n_{N,j}}, \quad j = 1,2,...,N,$ describing the evolution of an N-system chemical reaction where the k_j and $n_{i,j}$ are known constants and $[A]$ denotes the concentration of chemical A,
- $\frac{d^2}{dx^2}\left[EI\frac{d^2 y}{dx^2}\right] = q(x),$ which describes the small vertical deformations y of a beam of moment of inertia I and Young's modulus E under a vertical load q per unit length, where x is the coordinate along the length of the beam, and
- $\frac{d^2 Q}{dt^2}\sum_{i=1}^{N_C}\frac{1}{C_i} + \frac{dQ}{dt}\sum_{i=1}^{N_R} R_i + Q\sum_{i=1}^{N_L} L_i = 0,$ which describes the charge Q within a circuit containing N_R resistors R_i, N_C capacitors C_i, and N_L inductors L_i in series.

Odes are solved using a variety of methods depending on whether the problem is a boundary value problem or an initial value problem (see section 3.4 for definition). The methods will be discussed more thoroughly in section 10.6.

10. Model solution methods

This section discusses some of the most common discretisation methods, as introduced in section 3.2. The advantages and disadvantages of each method are discussed.

It should be noted that much of the detail given here is not necessary for a user to begin using commercially available software to apply the techniques. However, the detail becomes more useful when the techniques are to be applied to complicated problems involving nonlinearity, and when the user wishes to develop custom-written software.

Throughout the following, a "structured" grid or mesh is a mesh made up of parallelograms (preferably rectangles) or parallelepipeds (preferably cuboids), such as is shown in figure 10.4. The nodes lie in patterns along sets of parallel and perpendicular lines. An "unstructured" mesh, such as that shown in figure 10.1, has no pattern to its nodal points.

10.1 Finite element techniques

Finite element analysis splits the domain into a set of non-overlapping elements that cover the entire domain. In one dimension, these elements are line segments, in two dimensions they are usually triangles or quadrilaterals and in three dimensions they are generally hexahedra, wedges, or tetrahedra. In all cases, the elements may have curved sides of they are second order or above. In addition to these standard elements, some developers have used special elements for modelling springs, point masses, joints, infinite boundaries, and so on. The mesh points are placed at the corners of the elements, and for higher-order approximations points are placed mid-way along the element sides and sometimes in the centre of the element. The order of the element defines its shape and the order of approximation used within the element.

An example of a two-dimensional mesh using first-order triangular elements is shown in figure 10.1. The circles denote mesh points, usually called **nodes**, and the numbered areas are elements. Note that the mesh points are not placed at regular intervals, and that the elements are not all the same size and shape.

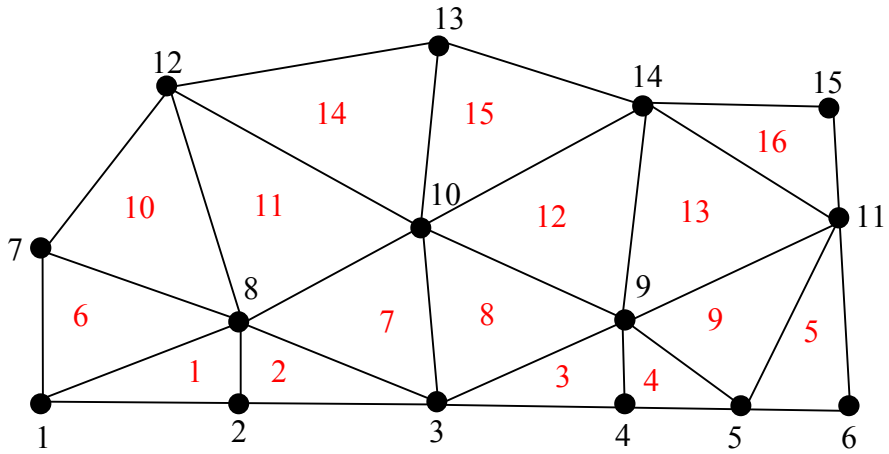

Figure 10.1: A typical two-dimensional finite element mesh. Numbers in red are element numbers, and numbers in black are node numbers.

If two adjacent but distinct regions are modelled, such as two different materials perfectly bonded together, then the two regions will share a line of nodes and the continuity of various of the field quantities will be assumed along the join. The boundary conditions are applied to the outermost nodes of the mesh, so in figure 10.1 all the nodes except nodes 8, 9, and 10 would have boundary conditions applied.

The assumption made about the behaviour of the solution is that the solution within an element can be written as a sum of functions as suggested in section 4.4, and parameterised by the solution values at the mesh points. The functions are usually polynomials in the appropriate number of spatial dimensions, and are generally chosen such that if the $\mathbf{x}_i$ are the mesh points and the N_i are the corresponding basis functions then

$$N_i(\mathbf{x}_k) = \delta_{ik}, \quad 1 \le i,k \le N.$$

In this formulation, the basis function N_i is zero within any element that does not include the i^{th} node. For example, on the mesh in figure 10.1, N_2 will be zero everywhere apart from within elements 1 and 2 whereas N_{10} will be non-zero over elements 7, 8, 11, 12, 14, and 15.

Not all basis functions are formulated in this way. Hierarchical basis functions use non-physical parameters for higher order approximations rather than adding extra nodes to the mesh. The main advantage of this technique is that the resulting matrix of equations is close to diagonal and so has improved conditioning properties. The main disadvantage is that the non-physical parameters can be confusing when interpreting results, and the element formulation can be tricky.

Once the basis functions have been chosen, the contributions of each basis function to the behaviour of the solution within each element can be calculated. For example, if a parameter Q is of interest and one element is the triangle bounded by $x = 0$, $y = 0$, and $x + y = 1$ (see figure 10.2) with vertices A, B and C, then the a first order (i.e. linear) finite element method would assume that

$$Q(x,y) = (1-x-y)Q_A + xQ_B + yQ_C$$

within the element shown. If a quadratic approximation was required, then the values at the mid-points of each side (D, E, and F) would be required and the resulting expression would be

$$\begin{aligned} Q(x,y) = {} & (1-x-y)(1-2x-2y)Q_A + x(2x-1)Q_B + y(2y-1)Q_C \\ & + 4x(1-x-y)Q_D + 4xyQ_E + 4y(1-x-y)Q_F. \end{aligned}$$

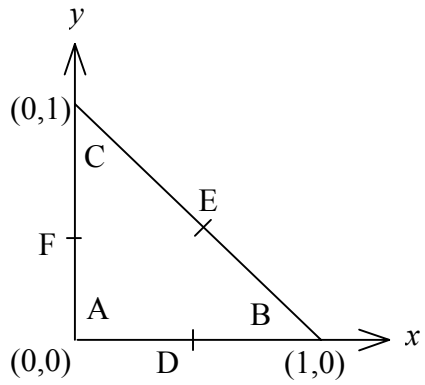

Figure 10.2: A small region of the domain, typical of those used in finite element analysis.

If the triangle in figure 10.2 did not have two of its sides lying along the x and y axes, this formula would not be so straightforward. However, any triangle can be mapped onto this triangle. Instead of x and y, suppose that the axes in figure 10.2 are ξ and η. Consider a general triangle with corners (x_1, y_1), (x_2, y_2), and (x_3, y_3). Then writing

$$x = x_1 + \xi(x_2 - x_1) + \eta(x_3 - x_1), \quad 0 \le \xi, \eta \le 1,$$
$$y = y_1 + \xi(y_2 - y_1) + \eta(y_3 - y_1), \quad 0 \le \xi, \eta \le 1,$$

so that

$$\xi = \frac{(x-x_1)(y_3-y_1)-(x_3-x_1)(y-y_1)}{(x_2-x_1)(y_3-y_1)-(x_3-x_1)(y_2-y_1)},$$
$$\eta = \frac{(x-x_1)(y_2-y_1)-(x_2-x_1)(y-y_1)}{(x_3-x_1)(y_2-y_1)-(x_2-x_1)(y_3-y_1)},$$

gives a suitable mapping. Similar expressions can be obtained for other element shapes. This mapping is useful during the assembling of the matrix coefficients.

The matrix equations are assembled using the weak formulation given in equation 4.4 of section 4.4. The basis functions φ_j and the weighting functions w_j are usually both chosen to be the N_j as defined above. Hence it is necessary to calculate expressions of the form

$$\int_{\Omega} N_j(\mathbf{x}) A\left(\sum_{k=1}^{N} U_k N_k(\mathbf{x})\right) dV + \int_{\partial\Omega} N_j(\mathbf{x}) B\left(\sum_{k=1}^{N} U_k N_k(\mathbf{x})\right) dS = 0, \quad j = 1,2,...,N.$$

These expressions can be simplified by recalling that Ω is the union of the set of elements $\{S_n, n = 1, 2, \ldots, N_e\}$ and by rewriting the integrals as a sum of integrations over each of these sub-intervals:

$$\sum_{n=1}^{N_e} \int_{S_n} N_j(\mathbf{x}) A\left(\sum_{k=1}^{N} U_k N_k(\mathbf{x})\right) dV + \sum_{n=1}^{N_e} \int_{\partial S_n} N_j(\mathbf{x}) B\left(\sum_{k=1}^{N} U_k N_k(\mathbf{x})\right) dS = 0, \quad j = 1,2,...,N.$$

Each basis function will only be non-zero over a small number of elements, considerably reducing the number of calculations required.

The calculations can be further simplified by using the mapping outlined above, by writing

$$\iint_{S_n} N_j(x,y) A\left(\sum_{k=1}^{N} U_k N_k(x,y)\right) dxdy = \iint_{S^*} N_j(\xi,\eta) A^*\left(\sum_{k=1}^{N} U_k N_k(\xi,\eta)\right) |J| d\xi d\eta,$$
$$J = \frac{\partial(x,y)}{\partial(\xi,\eta)},$$

where A^* has been altered to use the coordinates ξ and η, and S^* is the canonical triangle shown in figure 10.2. Whilst this form looks more complicated, it is easier to evaluate numerically.

The numerical evaluation of the expressions above is carried out using quadrature, because the operator A may involve non-linear terms, and often the software implementations of the finite element method are designed to solve for a wide range of operators so a generally applicable methodology is required. The accuracy of quadrature can be improved by using quadrature points that are not the corners and mid-points of the elements. Hence the unknowns U_k in the expressions above are generally the values at the quadrature points within each element rather than the values at the nodes. The results are often output as nodal values, which are calculated by extrapolation from the values at the quadrature points.

The matrices generated when finite element techniques are used are quite sparse, because the various unknown quantities are only linked to their nearest neighbours due to the use of very localised test functions.

Transient models are commonly solved in one of two ways using finite elements. If a time-dependent problem has conditions that only change slowly, is approximately linear, and does not feature complicated failure or contact conditions, then it can be regarded as "quasi-static" and an implicit time integration technique can be used (see section 4.2.5). Otherwise, an explicit integration method is used. Caution should be exercised when considering results

obtained using explicit solution methods since it is often easy to generate inaccurate results without realising it, through a poor choice of time step. This is particularly true for models involving complicated nonlinear behaviour and user-defined materials.

The main benefit of the finite element method is its flexibility. Since the model can be derived from its weak form, as explained in section 4.4, a wide variety of continuous models can be solved using the method. It can be applied in one, two or three dimensions and does not require a regularly-spaced grid of mesh points. The irregularity of the mesh means that complicated geometries can be modelled easily, and that a mesh can be designed to be detailed in areas of rapid change and sparse in areas of constant behaviour. The accuracy of the approximation can be increased by increasing the number of elements or by increasing the order of approximation used within the elements. The sparseness of the matrices used in finite elements often means that efficient storage and solution techniques can be used, making solution comparatively rapid even for problems with a large number of unknowns.

The main drawback of the finite element method is that it requires a mesh covering the whole of the domain. This restriction is a particular problem for wave transmission problems that use a condition at infinity as one of the boundary conditions. Infinite elements can be used to avoid this problem, but it is often more straightforward to use an alternative solution technique such as the boundary element method.

Another disadvantage of the finite element method relates to problems involving high levels of deformation (principally transient stress analysis and flow problems). As the computational domain moves, the nodes move and so the elements become distorted. If the elements distort too much, numerical problems occur leading to poor results. This problem can be alleviated by using automatic remeshing, but this technique is very computationally expensive.

Finite element methods are suitable for most types of continuous model, and much of the software that is available is of a good quality. Many proprietary packages include easy-to-use model development tools and include many common material models and boundary conditions as pre-defined options, making the model development stage easier.

10.2 Boundary element methods

Boundary element methods are formulated in a similar way to finite element methods. The domain is covered by a mesh of possibly irregular non-overlapping elements, and assumptions are made about the local behaviour within each element. The weak form of the equation is used to generate a set of linear equations linking the unknowns, and these equations are solved numerically.

The principal difference between the boundary element method and the finite element method is that the boundary element method expresses the weak form in terms of surface integrals rather than volume integrals, and that the weighting function as described in section 4.4 is the Green's function of A^T, the adjoint of the operator A with respect to the bilinear form

$$B(u,v) = \int_\Omega uv dV.$$

The adjoint is defined as the operator such that $B(Au, v) = B(u, A^T v)$ for all functions u and v. Many of the operators commonly occurring in continuous models are self-adjoint so that $A^T = A$.

As an example, consider the steady-state heat equation on a three-dimensional domain. The operator is self-adjoint, and its Green's function is

$$G(\mathbf{x},\mathbf{x}') = \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}.$$

Substituting this into the weak form of the heat equation gives

$$\int_\Omega \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla^2 T(\mathbf{x})dV(\mathbf{x}) = 0.$$

Applying the divergence theorem twice gives

$$\int_\Omega \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla^2 T(\mathbf{x})dV(\mathbf{x}) = \int_\Omega \nabla.\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x})\right)dV(\mathbf{x}) - \int_\Omega \nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right).\nabla T(\mathbf{x})dV(\mathbf{x})$$

$$= \int_{\partial\Omega}\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x})\right).\mathbf{n}dS(\mathbf{x}) - \left\{\int_\Omega \nabla.\left(\nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right)T(\mathbf{x})\right)dV(\mathbf{x}) - \int_\Omega \nabla^2\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right)T(\mathbf{x})dV(\mathbf{x})\right\}$$

$$= \int_{\partial\Omega}\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x}) - \nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right)T(\mathbf{x})\right).\mathbf{n}dS(\mathbf{x}) + \int_\Omega \delta(\mathbf{x}-\mathbf{x}')T(\mathbf{x})dV(\mathbf{x}),$$

where **n** is the outward normal to the boundary of the domain, and so

$$\int_{\partial\Omega}\left(T(\mathbf{x})\nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right) - \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x})\right).\mathbf{n}dS(\mathbf{x}) = \begin{cases} T(\mathbf{x}'), & \mathbf{x}' \in \Omega, \\ \tfrac{1}{2}T(\mathbf{x}'), & \mathbf{x}' \in \partial\Omega, \\ 0, & \text{otherwise.} \end{cases}$$

The solution of this equation for points in the domain requires knowledge of T and its gradient on the whole of the boundary, which is generally not available. Hence either the calculation must be repeated with an alternative form of the Green's function that vanishes on the boundary (which is only possible for certain domain geometries) or the problem must be solved in two stages. First the boundary element method must be applied to

$$\int_{\partial\Omega}\left(T(\mathbf{x})\nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right) - \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x})\right).\mathbf{n}dS(\mathbf{x}) = \tfrac{1}{2}T(\mathbf{x}'), \qquad \mathbf{x}' \in \partial\Omega, \tag{10.1}$$

to get the temperature and gradient distribution on the boundary, and then those results can be used in

$$\int_{\partial\Omega}\left(T(\mathbf{x})\nabla\left(\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\right) - \frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|}\nabla T(\mathbf{x})\right).\mathbf{n}dS(\mathbf{x}) = T(\mathbf{x}'), \qquad \mathbf{x}' \in \Omega, \tag{10.2}$$

to obtain the temperature at a general domain point. Each of these two equations is solved by discretising the surface $\partial\Omega$, approximating the unknowns as a sum of parameterised local functions and using quadrature techniques to evaluate the integrals in terms of the unknown parameters. The simplest function used assumes that the unknown is constant within the element.

The matrices generated by application of the boundary element method are generally dense, particularly the matrices associated with the solution of equation (10.1). The density occurs because the unknown value at a given point on the surface is related to the unknown values at every other point in the surface, so every entry in the matrix is non-zero. This density can lead to problems with efficient storage and solution procedures, and an associated reduction in the maximum possible size of problem that can be solved. In some cases, such as the presence of rotational symmetry, the matrices have special structure that can reduce the problem, but in general this is unlikely to be the case.

The main advantage of the boundary element method is its suitability for problems involving an infinite domain. Most infinite domain problems assume that some function of the solution tends towards zero at infinity such that the integral over the boundary vanishes at the infinite boundary and the remaining integral is over a finite boundary.

Another advantage is that as boundary element formulations use an integral equation over the outside of a domain, the formulation may avoid singular behaviour, making problems such as crack propagation more amenable to numerical solution.

The main drawback to boundary element methods is that the Green's functions used in their derivation are often singular, and so the evaluation of the various integrals required in the formulation becomes difficult. In some cases the singularity can be avoided and a normal quadrature scheme can be used, but in other cases extra manipulation of the integral terms is required to make the problem numerically stable. In addition to this problem, the identification of a suitable Green's function can be difficult for some problems, making the formulation step tricky.

Boundary element methods are particularly suitable for wave transmission and scattering problems in acoustics and electromagnetics, and for problems in stress analysis (particularly damage models). They are also used for fluid flow problems and thermal modelling.

10.3 Finite volume techniques

Finite volume methods are usually developed by considering the changes in some quantity (e.g. mass, energy, momentum, etc.) within a small volume of the domain. They assume that the field quantities are constant either throughout the volume or across one of its faces, and calculate the resulting fluxes in terms of field quantities and element dimensions.

As an example, consider a small fixed two-dimensional volume of space through which an incompressible inviscid fluid is flowing in the absence of body forces (e.g. gravity), as shown in figure 10.3, and suppose that the flow is steady-state and the fluid is at a uniform temperature. Assume that the fluid velocities (u, v) are uniform across each of the walls of the volume, and that the pressure p_i is uniform throughout the volume.

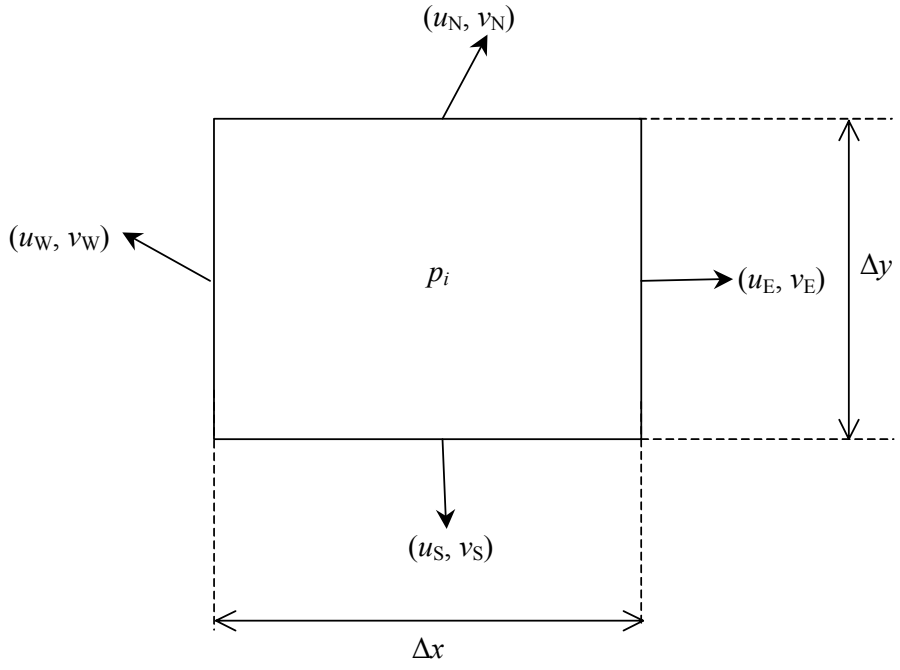

Figure 10.3: A fixed volume of space through which liquid is flowing.

Mass is conserved within the element. The only transport of mass is due to the fluid velocities, since the fluid is incompressible. The change in mass due to the flow is

$$\Delta m = \rho\{\Delta y(u_E - u_W) + \Delta x(v_N - v_S)\},$$

and so the discretised form of the equation for conservation of mass becomes

$$\frac{(u_E - u_W)}{\Delta x} + \frac{(v_N - v_S)}{\Delta y} = 0.$$

The conservation of momentum can be applied in two directions to generate similar equations:

$$\frac{u_E u_E - u_W u_W}{\Delta x} + \frac{u_N v_N - u_S v_S}{\Delta y} + \frac{p_E - p_W}{\rho \Delta x} = 0,$$

$$\frac{u_E v_E - u_W v_W}{\Delta x} + \frac{v_N v_N - v_S v_S}{\Delta y} + \frac{p_N - p_S}{\rho \Delta y} = 0.$$

where p_E, p_W, p_S and p_N are the pressures to the east, west, south and north of the volume shown. This over-simplistic formulation leads to numerical problems during solving, but the ideas are the same when deriving improved formulations.

Many problems commonly solved with finite volume methods involve more than one type of unknown quantity. Examples include Maxwell's equations and the Navier-Stokes equations (see sections 9.3 and 9.4 for further details of these equations). Such problems can benefit from the use of staggered grids, with a separate grid being used for each unknown (so for instance in the example above, pressure and volume grids are offset from one another.

The example above used a regularly-shaped orthogonal mesh. It is possible to generate non-uniform finite volume meshes, but the mathematics is considerably more involved. The methods can be extended to transient problems by using finite differences in the time direction (i.e. by considering the change in conserved quantities over a short time step).

In the general description of methods given in section 4.4, the finite volume method in its simplest (and most commonly-used) form uses the basis functions and weighting functions to be constant within the volume in question and zero outside of it. The matrices generated by the methods are quite sparse because each unknown value depends only on the values immediately around it.

One advantage of finite volume methods is that because they consider the volume mesh to be fixed and the amount of substance in the volume to be variable, flow problems and problems featuring large amounts of distortion can be easier to implement than they would be using finite elements or finite differences.

In addition, because they are derived directly from conservation laws when applied to steady-state problems, the results automatically satisfy those laws. Their derivation is reasonably straightforward conceptually, making them easy to explain. The methods can be extended to less structured grids than the one suggested above.

The main disadvantage of the methods is that naive implementations of the methods can have numerical problems: for example, in the formulation above, the pressure p_i does not appear in the equations for the volume considered. This leads to coupling between alternate cells which can lead to instability. Additionally, finite volume meshes are constructed from shapes with straight sides, and so finite volume meshes do not always give good approximations of curved geometries, although published work has addressed this drawback for some applications [48].

Finite volume techniques are most commonly used to solve fluid flow problems and electromagnetics problems (particularly in the time domain). They are less well-suited to problems involving complicated geometries.

10.4 Finite difference techniques

Finite difference techniques place mesh points at intervals on an orthogonal grid. The mesh points are usually placed at regularly-spaced intervals as this makes the calculation of

derivatives easier and can improve the conditioning. It is then assumed that the solution can be expanded locally as a Taylor series.

For example, suppose that the solution to the model is a function $f(x, y)$ of the spatial coordinates x and y. Then the local Taylor expansions for the four points $f(x\pm\Delta x, y)$ and $f(x, y\pm\Delta y)$ will be

$$f(x\pm\Delta x,y)=f(x,y)\pm\Delta x\frac{\partial f}{\partial x}(x,y)+\frac{\Delta x^2}{2}\frac{\partial^2 f}{\partial x^2}(x,y)\pm\frac{\Delta x^3}{3!}\frac{\partial^3 f}{\partial x^3}(x,y)+\frac{\Delta x^4}{4!}\frac{\partial^4 f}{\partial x^4}(x,y)+\ldots,$$

$$f(x,y\pm\Delta y)=f(x,y)\pm\Delta y\frac{\partial f}{\partial y}(x,y)+\frac{\Delta y^2}{2}\frac{\partial^2 f}{\partial y^2}(x,y)\pm\frac{\Delta y^3}{3!}\frac{\partial^3 f}{\partial y^3}(x,y)+\frac{\Delta y^4}{4!}\frac{\partial^4 f}{\partial y^4}(x,y)+\ldots,$$

and so rearranging gives

$$\frac{\partial^2 f}{\partial x^2}(x,y)=\frac{1}{\Delta x^2}\{f(x+\Delta x,y)+f(x-\Delta x,y)-2f(x,y)\}-\frac{2\Delta x^2}{4!}\frac{\partial^4 f}{\partial x^4}(x,y)+\ldots,$$

$$\frac{\partial^2 f}{\partial y^2}(x,y)=\frac{1}{\Delta y^2}\{f(x,y+\Delta y)+f(x,y-\Delta y)-2f(x,y)\}-\frac{2\Delta y^2}{4!}\frac{\partial^4 f}{\partial y^4}(x,y)+\ldots,$$

so for instance

$$\nabla^2 f=\frac{\partial^2 f}{\partial x^2}(x,y)+\frac{\partial^2 f}{\partial y^2}(x,y)$$

$$=\frac{f(x+\Delta x,y)+f(x-\Delta x,y)-2f(x,y)}{\Delta x^2}+\frac{f(x,y+\Delta y)+f(x,y-\Delta y)-2f(x,y)}{\Delta y^2}+O(\Delta x^2,\Delta y^2).$$

Other similar approximations can be obtained for other derivatives. The size of the mesh determines the accuracy of the solution. Some transient pdes generate finite difference approximations that are only stable for a limited range of time steps.

The matrices generated by finite difference methods are generally sparse because the unknown value at each mesh point is related to its nearest neighbours. Using the terminology of section 4.4, finite difference methods have a delta function weighting, and a set of localised polynomial basis functions.

The form of the approximation means that the mesh points must lie on an orthogonal grid, or on a grid that can be mapped onto an orthogonal grid (for instance a set of concentric circles and radial lines). This requirement can have disadvantages. One is that it limits the range of geometries that can be described accurately. Another is that some boundary conditions can be difficult to implement, particularly if they require the normal derivative of a curved surface.

In general, a fine mesh is required where the solution is changing rapidly. This can mean that the mesh is unnecessarily fine in some areas. For instance, consider the mesh shown in figure 10.4, which has been designed for a problem that has very steep gradients close to the origin and is less variable away from the origin.

The mesh has more detail than is necessary for areas near to the axes but away from the origin. This unnecessary detail increases the solution time and matrix storage requirements, and may produce stability problems in some cases. In order to avoid such problems, multigrid methods have been developed. In their simplest form, these methods solve for the whole problem on a coarse mesh, then solve the area around the origin on a finer mesh, then use the solution on the fine mesh to update the solution on the coarse mesh and re-solve, and then re-solve for the fine mesh, and so on until convergence is achieved. The methods can provide converged solutions more rapidly with fewer numerical problems.

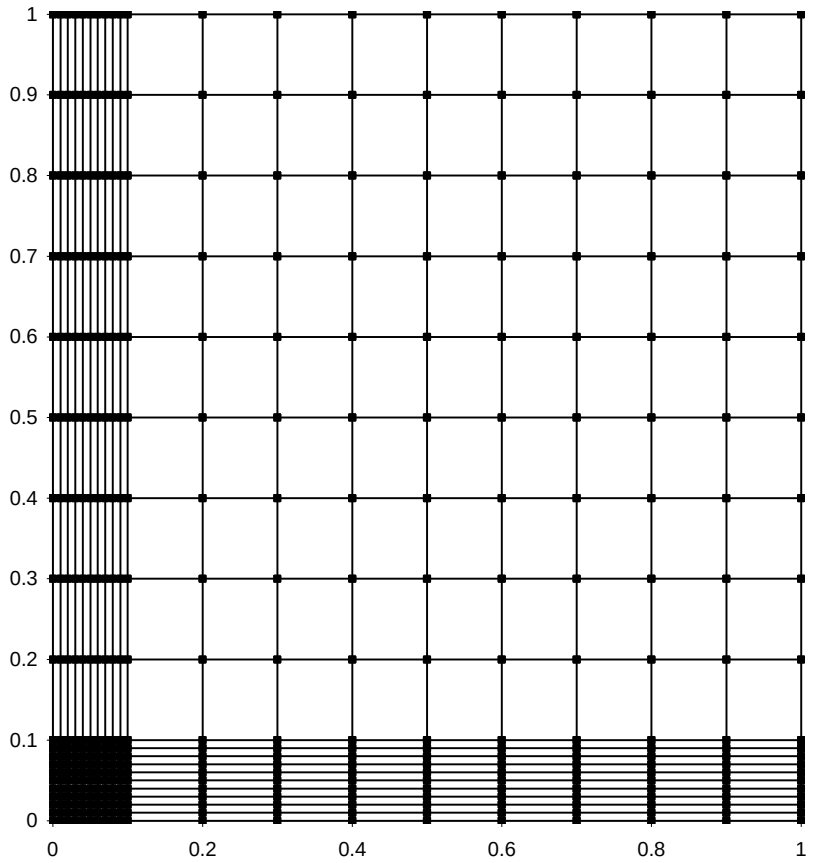

Figure 10.4: A finite difference mesh for a problem with sharp gradients around the origin.

Finite difference methods are commonly used for the time dimension of transient equations because the limitations of the methods (geometric limitations, approximations of boundary conditions) do not affect the time dimension.

The main advantage of finite difference methods is their ease of implementation. The straightforward nature of the formulae mean that they can be adapted to nonlinear problems and problems featuring spatial variations of properties reasonably easily. The methods have convergence properties and truncation errors that are derived from well-established mathematical analysis, meaning that the methods can be applied with confidence.

The main disadvantage of finite difference methods is that the requirements of the structured mesh make modelling of small geometric features and accurate approximation of some boundary conditions problematic.

10.5 Semi-analytic methods

Semi-analytic methods are generally used to reduce the dimensionality of a problem. The methods usually result in a simpler set of equations that can be solved more quickly than the original formulation.

Semi-analytic methods are particularly common for reducing transient problems to (often coupled) sets of static problems. This usage is because it is often easier to predict how a system will evolve over time in general terms: often the long-term behaviour of a system is relatively well understood and is the main result of interest.

The term "semi-analytic methods" covers too wide a range for every common approximation to be discussed here. Three of the most common methods are:

1. make an assumption about the behaviour of the solution in one dimension and generate a new equation on the basis of that assumption,
2. use an integral transform (Fourier transform, Laplace transform, etc.) to move the time-dependence of the problem into the frequency domain,
3. expand the solution as a sum of functions dependent on some parameter, usually known to be small or large.

An example of the first type of method is reduction of a three-dimensional stress analysis problem to a plane strain problem. The full set of equations for a static elastic stress analysis problem are

$$\frac{\partial \sigma_{ij}(\varepsilon)}{\partial x_j} = 0, \quad i = 1,2,3,$$

$$\varepsilon_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad i, j, = 1,2,3,$$

$$\sigma_{ij} = \frac{E}{(1+\nu)}\varepsilon_{ij} + \delta_{ij}\varepsilon_{kk}\frac{E\nu}{(1+\nu)(1-2\nu)}, \quad i, j = 1,2,3.$$

Assuming that the domain is much larger in the x_3 direction than the other two directions, and that no significant forces act in the x_3 direction, then it is reasonable to assume that $u_3 = 0$ everywhere, and that gradients in the x_3 direction can be ignored. These assumptions mean that the only three non-zero strains are ε_{11}, ε_{22}, and ε_{12}, and the non-zero stress terms are σ_{11}, σ_{22}, σ_{12}, and σ_{33}, so the equations become

$$\frac{E(1-\nu)}{(1+\nu)(1-2\nu)}\frac{\partial^2 u_1}{\partial x_1^2} + \frac{E}{2(1+\nu)}\frac{\partial^2 u_1}{\partial x_2^2} + \frac{E}{2(1+\nu)(1-2\nu)}\frac{\partial^2 u_2}{\partial x_1 \partial x_2} = 0,$$

$$\frac{E(1-\nu)}{(1+\nu)(1-2\nu)}\frac{\partial^2 u_2}{\partial x_2^2} + \frac{E}{2(1+\nu)}\frac{\partial^2 u_2}{\partial x_1^2} + \frac{E}{2(1+\nu)(1-2\nu)}\frac{\partial^2 u_1}{\partial x_1 \partial x_2} = 0,$$

or

$$\frac{\partial^2 u_1}{\partial x_1^2} + \frac{\partial^2 u_1}{\partial x_2^2} = \frac{1}{2(1-\nu)}\left(\frac{\partial^2 u_1}{\partial x_2^2} - \frac{\partial^2 u_2}{\partial x_1 \partial x_2} \right),$$

$$\frac{\partial^2 u_2}{\partial x_1^2} + \frac{\partial^2 u_2}{\partial x_2^2} = \frac{1}{2(1-\nu)}\left(\frac{\partial^2 u_2}{\partial x_1^2} - \frac{\partial^2 u_1}{\partial x_1 \partial x_2} \right).$$

These equations can generally be solved considerably more quickly than the equivalent three-dimensional version.

An example of the application of a Laplace transform to a thermal problem has been given in section 9.2.4. Fourier transforms are commonly used in electromagnetic and acoustic wave transmission problems.

The third type of method is commonly used in two ways: either a perturbation expansion or an asymptotic expansion. Both methods approximate an unknown function with a sum of terms, then compare terms of the same order to generate a new set of equations. The difference between the methods lies in their convergence properties. Suppose a function $f(x, \varepsilon)$, where ε is a small parameter, is approximated as

$$f(x,\varepsilon) \approx f_N(x,\varepsilon) = \sum_{n=1}^{N} g_n(x) h_n(\varepsilon).$$

Then a perturbation solution will have

$$f_N(x,\varepsilon) \xrightarrow[N\to\infty]{} f(x,\varepsilon), \quad \forall x, \text{fixed small } \varepsilon,$$

and an asymptotic expansion will have

$$f_N(x,\varepsilon) \xrightarrow[\varepsilon\to 0]{} f(x,\varepsilon), \quad \forall x, \text{fixed } N.$$

The asymptotic expansion does not necessarily converge as N increases: in fact, there is very likely to be an optimal choice for N. The most common choice for h_n is $h_n = \varepsilon^n$, but cases where singular behaviour exists can require other functions. The methods are commonly used to develop new pdes (often nonlinear) for the functions $g_n(x)$. Examples of perturbation methods are found in nonlinear acoustics [43]. Asymptotic expansions are commonly used for evaluation of integrals, and have applications in a wide range of physical areas, particularly wave transmission problems.

Semi-analytic methods can be applied to any problem in which the behaviour on one of the coordinate directions (including time) is believed to be understood. They are particularly common for transient problems.

10.6 Methods for ordinary differential equations

There is a vast range of methods for solution of ordinary differential equations. Different methods have been designed to give good performance for different types of equation. Boundary value problems and initial value problems (defined in section 3.4) need different types of method, and stiff problems (defined in section 4.3.1) require special consideration. This section will only give an outline of some of the most common methods for solution of odes, and the reader requiring more information is directed to the references [17, 47]. Additionally, most of these methods are implemented in Chapter D02 of the NAG software library [8].

The simplest methods for ordinary differential equations are derived directly from the Taylor expansion of the equation. Suppose we want to solve the equation for $\mathbf{y}(x)$

$$\mathbf{y}' = \mathbf{f}(x,\mathbf{y}), \qquad 0 \le x \le X,$$

subject to a set of initial conditions $\mathbf{y}(0) = \mathbf{y}_z$. The Taylor expansion for $\mathbf{y}$ in the region of some point x_0 is

$$\mathbf{y}(x) = \mathbf{y}(x_0) + (x - x_0)\mathbf{y}'(x_0) + \frac{(x-x_0)^2}{2!}\mathbf{y}''(x_0) + \frac{(x-x_0)^3}{3!}\mathbf{y}'''(x_0) + \cdots$$

and so if $|x - x_0|$ is small, neglecting higher order terms and substituting for the first derivative gives

$$\mathbf{y}(x) = \mathbf{y}(x_0) + (x - x_0)\mathbf{f}(\mathbf{y}(x_0), x_0) + O\left((x - x_0)^2\right)$$

and so if the value of $\mathbf{y}(x_0)$ is known, then $\mathbf{y}(x)$ can be calculated to first order. Hence if we have an initial value problem, by repeatedly carrying out the calculation

$$\mathbf{y}_{n+1} = \mathbf{y}_n + (x_{n+1} - x_n)\mathbf{f}(x_n, \mathbf{y}_n), \quad n = 1,2,\ldots,N,$$
$$\mathbf{y}_0 = \mathbf{y}_z$$
$$0 = x_0 < x_1 < x_2 < \cdots < x_{N-1} < x_N = X$$

we can calculate the function throughout the domain. This is Euler's method. The step size $x_{n+1} - x_n$ is often written h_n, or h if the step size is constant (this convention will be used in this section).

Similar rules can be developed that involve more values known from previous steps, provided the x_n are equally spaced. These are called **multistep** methods, and particularly common are linear multistep methods, which are of the form

$$\sum_{j=0}^{k} a_j \mathbf{y}_{n+j} = \sum_{j=0}^{k} b_j \mathbf{f}_{n+j},$$

where $\mathbf{f}_n = \mathbf{f}(x_n, \mathbf{y}_n)$ and the a_j and b_j are constants. In the example above, $k = 1$, $a_1 = 1$, $a_0 = -1$, $b_1 = 0$, and $b_0 = (x_{n+1} - x_n) = h$. Commonly-used multistep methods include the Adams methods (also known as Adams-Bashforth methods if explicit and Adams-Moulton methods if implicit). Multistep methods can require a few steps of a single-step method to generate enough initial values to start the method.

In general, implicit multistep methods have good stability properties but can take a long time to calculate as they require iterative solution of a set of implicit (possibly non-linear) equations. Explicit methods are easy to calculate but are often stable only for very limited step sizes. A common way of combining the good properties of the two types of method is to use the explicit method to generate the initial guess required for the calculation with the implicit method. Such methods can be shown to have good stability and convergence properties. These methods have been described in previous SS*f*M work on continuous modelling [1].

One of the main disadvantages of multistep methods is that the step size must be uniform throughout. This can be very computationally wasteful if, for instance, an initially varying system has reached steady state. **Runge-Kutta** (RK) methods can be used with a variable step size, and hence are much more computationally efficient for most problems. RK methods calculate intermediate values between xn and xn+1 in order to reduce the truncation error of the method. Probably the most commonly-used RK method is the fourth-order method

$$\mathbf{y}_{n+1} = \mathbf{y}_n + \frac{\mathbf{k}_1}{6} + \frac{\mathbf{k}_2}{3} + \frac{\mathbf{k}_3}{3} + \frac{\mathbf{k}_4}{6} + O\left(h^5\right),$$

$$\mathbf{k}_1 = h\mathbf{f}(x_n, \mathbf{y}_n), \qquad \mathbf{k}_2 = h\mathbf{f}\left(x_n + \frac{h}{2}, \mathbf{y}_n + \frac{\mathbf{k}_1}{2}\right),$$

$$\mathbf{k}_3 = h\mathbf{f}\left(x_n + \frac{h}{2}, \mathbf{y}_n + \frac{\mathbf{k}_2}{2}\right), \qquad \mathbf{k}_4 = h\mathbf{f}(x_n + h, \mathbf{y}_n + \mathbf{k}_3).$$

Several different techniques can be used to estimate the error produced by a single step of an RK method (for instance Richardson extrapolation, mentioned in [1], or embedded RK methods, see [47]). The error estimate can then provide guidance for step size alterations: if

the error is very small then the step size could be increased, and if it is too large the step calculation could be repeated with a smaller step size.

It can be shown that the stability properties of all explicit linear multistep methods and explicit RK methods make them unsuitable for solving stiff systems (see section 4.3.1 for a discussion of stability and stiffness): instead an implicit method must be used. There is a class of implicit multistep methods called **backward differentiation formulae** (BDF) that have good stability properties and are relatively straightforward to implement (although as with all implicit methods an iterative solver is required for nonlinear problems), and so are recommended for stiff systems. Similar implicit RK methods exist but these methods are less popular because they require more functional evaluations, and hence more iterative loops, than the BDF. The general form of a BDF is

$$\sum_{j=0}^{k} a_j \mathbf{y}_{n+j} = h b_k \mathbf{f}_{n+k},$$

with the simplest example being the backward Euler method, which has $k = 1$, $b_1 = 1$, $a_1 = 1$, and $a_0 = -1$. More complicated methods involve higher values of k, and have better truncation errors.

All of the methods described above require known values of all the unknowns to start the method, which means that they are only suitable for initial value problems. Boundary value problems require a different set of approaches. The generic approach is to generate a first guess at the solution, then use an iterative technique to improve the first estimate.

One commonly-used relatively robust class of methods is **shooting methods**, which estimate the unknown conditions at one end of the interval, solve the resulting initial value problem and then adjust the estimates of the unknown conditions according to the difference between the calculated values at the other end of the interval and the prescribed conditions there.

An alternative method (sometimes called a **relaxation method**) uses finite difference methods to approximate the ode on a set of points, then uses a Taylor expansion to generate equations for increments of each of the **y** at each of the points such that the new approximation will be closer to the desired solution.

The success of both types of method is partly governed by the quality of the initial estimates of the unknown quantities. in addition, the shooting method can encounter problems for stiff systems, and the relaxation methods may not be suitable for systems exhibiting oscillatory behaviour over small regions.

A class of method that can be used for both type of problem is collocation methods. In these methods, the solution to an ode is approximated using polynomials (often Chebyshev polynomials) and the coefficients of the polynomials are the quantities to be calculated. The method is suitable for linear problems, and can be applied to nonlinear problems if an iterative technique is used.

10.7 Summary

The benefits and drawbacks of each of the methods described in this chapter are:

Finite element methods

+ Unstructured meshes can be used so most geometries can be modelled
+ Exact implementation of differential boundary conditions
+ Wide range of software available
+ Has useful error minimisation properties

- Formulation from scratch requires mathematical understanding and rigor
- Can be very computationally expensive

Boundary element methods

+ Formulated as a surface problem so does not need a mesh over the entire domain
+ "Condition at infinity" for wave problems is automatically obeyed
+ Unstructured meshes can be used
- Formulation requires good mathematical understanding
- Can have stability and uniqueness problems
- Reformulation may not be suitable for every governing equation

Finite volume methods

+ Fairly easy to formulate
+ Physically intuitive
+ Easy to implement as software
+ Unstructured meshes can be used (with some difficulty)
+ Has useful conservation properties
- Can be very computationally expensive
- Differential boundary conditions approximated inexactly in most cases
- Based on flux calculations so may not be suitable for every governing equation

Finite difference methods

+ Mathematically easy to formulate
+ Easy to implement as software
+ Stability and consistency qualities well-established
- Structured mesh required
- Differential boundary conditions approximated inexactly
- Can be computationally expensive

Semi-analytic methods

+ Can enforce required behaviour on a solution
+ Lead to a simpler form of the governing equation
- Formulation requires imposing assumptions about the solution

Methods for ordinary differential equations

+ Usually developed for good stability & consistency
+ Higher-order derivatives easily handled by turning problem into a first-order system
- Can be hard to identify suitable method for stiff equations
- Boundary value problems can be difficult to solve

11. Further reading

This section suggests a range of books for further reading. Some of the titles have been cited in the main text, but they are repeated here for completeness.

The majority of these books require some degree of mathematical knowledge. This is an inevitable consequence of the mathematical nature of modelling measurement processes. Where possible, some introductory texts starting from the simplest possible point have been suggested. In addition, any books regarded as the "bible" of their field have been recommended as reference guides.

In addition to these books, many of the NAFEMS "Why do" and "How To" guides offer a broad introduction to topics, usually approached from a finite element perspective.

General introductions to mathematical modelling

N. D. Fowkes and J. J. Mahony, *An Introduction to Mathematical Modelling*, John Wiley, 1994, ISBN: 0471934224.

This book has been found to be useful in previous SS*f*M work on validation. The book outlines the skills and methods required to formulate and solve mathematical models. The book contains information on a wide variety of analytical methods, and shows examples of the methods not working (with reasons why) as well as successful examples. Indications are given of when numerical methods are required. Maple code is used throughout.

A. B. Tayler, *Mathematical Methods in Applied Mechanics*, Oxford University Press, 2002, ISBN: 0198515596.

This book covers the basic methods of model formulation and describes many analytic techniques for solution of simple models. These analytic techniques can often be used to obtain a first approximate solution and to understand the dependencies and behaviour of the solutions. The techniques are illustrated using real example problems drawn from physics, and indications of how the models could be extended and developed are given. The three main chapters cover elliptic, parabolic, and hyperbolic partial differential equations, with a final chapter on nonlinear equations. No specific descriptions of numerical methods are given, but references for suitable methods are supplied.

N. Gershenfeld, *The Nature of Mathematical Modelling*, Cambridge University Press, 1999, ISBN: 0521570956.

This book divides models into three classes, analytical, numerical, and observational, according to the most suitable solution method for the model. The material on analytical models largely covers model building and solution under simplifying assumptions. The section on numerical models acknowledges that very few differential equations can be solved analytically, and outlines some of the key numerical techniques used for solution. Observational models (the largest section) are obtained by manipulating measurement data (e.g. curve/surface fitting, digital signal processing, filtering, etc.), something that SS*f*M classes as discrete modelling. Key points are illustrated by examples throughout. Problems for the reader are given in each chapter with full solutions at the end of the book.

R. Aris, *Mathematical Modelling Techniques*, Dover Publications, 1995, ISBN: 0486681319. This book is a general introduction to the ideas and motivation behind mathematical modelling. It is written in a conversational style and is illustrated by some general examples.

N. Bellomo and L. Preziosi, *Modelling Methods and Scientific Computation*, CRC Press, 1995, ISBN: 0849383315. This book is a comprehensive review of the equations and methods of modelling in the physical sciences, including algorithms and code fragments.

E. A. Bender, *An Introduction to Mathematical Modelling*, Dover Publications, 2000, ISBN: 048641180X. This book introduces many of the basic ideas of mathematical modelling and features many examples. The majority of the examples are problems that can be solved analytically, and so few of the examples are drawn from physics. Partial differential equations are only mentioned briefly.

Stress Analysis

S. Timoshenko, J.N. Goodier, *Theory of Elasticity*, McGraw-Hill Publishing Co.,1970, ISBN: 0070858055.

S.P. Timoshenko, D.H. Young, W. Weaver, *Vibration Problems in Engineering*, John Wiley, 1990, ISBN: 0471632287.

These books derive analytical solutions to many simple problems in elastic stress and strain and vibration. Whilst such solutions are not directly applicable to many problems, they provide a good first approximation to many real-world problems. In addition, following the logic used to obtain the solution leads to an improved understanding of the theory underpinning the results, and the assumptions that have been made to generate them.

R. D. Cook, *Finite element modeling for stress analysis*, John Wiley, 1995, ISBN: 0471107743. As well as introducing the concepts and mathematical basis for FE, this book explains key points and pitfalls to be aware of when applying FE to real problems in stress analysis. It takes a practical approach and is suitable for engineers and physicists as well as mathematicians.

H. Kolsky, *Stress waves in solids*, Dover Publications, 2003, ISBN: 0486495345. This gives an introduction to the propagation of stress waves within solid materials. Geotechnical texts may be a good source of further material on this topic.

Heat Transfer

F. Incropera, D. DeWitt, *Fundamentals of Heat and Mass Transfer*, John Wiley and Sons (WIE), 1996, ISBN: 0471304603. This is a general introduction from an engineering perspective, including material on conduction, convection, and radiation. The derivation of equations and the physical phenomena are explained, and there is a large number of examples illustrating the points made. There is not much material on phase changes such as melting, although boiling is covered.

Electromagnetics

The following are general introductory texts to analytical electromagnetic modelling.

S. Ramo, J. R.Whinnery, and T. van Duzer, *Fields and Waves in Communication Electronics*, John Wiley, 1994, ISBN: 0471585513.

J. Kraus, *Electromagnetics* (5^{th} ed), McGraw-Hill, 1999, ISBN: 0071164294.

R. Harrington, *Time-harmonic electromagnetic fields*, John Wiley, 2001, ISBN: 047120806X.

The next three are regarded as standard references.

J. D. Jackson, *Classical electrodynamics* (3rd ed), John Wiley, 1999, ISBN: 047130932X. This is THE classic among physicists, but has more emphasis on fundamentals than applications and hence may be most suitable as a reference.

J. van Bladel, *Electromagnetic Fields*, Hemisphere, 1985 ISBN: 0891168192. This is very thorough and complete, and is regarded as a classic for electrical engineers.

C. Muller, *Foundations of the Mathematical Theory of Electromagnetic Waves,* Springer-Verlag, 1969 ISBN: 0387045066. This book is a classic for mathematicians.

Finally, these books have applications in specific areas.

M. Born, and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, Cambridge University Press, 1999, ISBN: 0521642221.

R. E. Collin, *Foundations for Microwave Engineering* (2nd ed), John Wiley, 2000, ISBN: 0780360311

All of the books above focus on analytical models. In general, books on numerical electromagnetic modelling tend to be covering a single specific method. References for some such books are given in section 9.3.

Waves and natural frequencies

Many books on specific application areas address waves and frequencies of vibration. For instance, most books on electromagnetism include material on electromagnetic waves. Please see the sections on application areas for such titles. One book that links the theories underpinning wave propagation in acoustics and electromagnetics is

D. S. Jones, *Acoustic and Electromagnetic Waves*, Clarendon Press, 1989, ISBN: 0198533802.

Acoustics

The following three texts are generally regarded as classics, largely focussing on derivation of equations and their analytical solution.

P. M. Morse & K. U. Ingard, *Theoretical Acoustics*, Princeton University Press, 1987, ISBN: 0691024014.

E. Skudrzyk, *Foundations of Acoustics, Basic Mathematics & Basic Acoustics*, Springer-Verlag, 1972, ISBN: 0387809880.

D. T. Blackstock, *Fundamentals of Physical Acoustics*, John Wiley, 2000, ISBN: 0471319791.

Numerical methods

The following books cover specific numerical solution methods. In general they are oriented towards mathematicians who wish to understand the background to the methods and wish to be able to check the convergence and stability of the methods.

O. C. Zienkiewicz and R. L. Taylor. *The Finite Element Method* (three volume set), McGraw-Hill, 2005, ISBN: 0750664312. This is a classic text, including theory, applications, and worked examples. As well as explaining the theory underpinning finite element analysis, it presents correlations and extensions, and links finite element methods to finite difference and boundary element methods via a common general formulation (as outlined briefly in section 4.4). The second volume focusses on solid and structural mechanics applications, and the third volume covers computational fluid dynamics.

K. W. Morton and D. F. Mayers, *Numerical Solution of Partial Differential Equations: an Introduction*, Cambridge University Press, 2005, ISBN: 0521607930. This book concentrates on the use of finite difference methods to solve partial differential equations. A wide range of methods and their strengths and weaknesses are illustrated with examples, and analysis is used to demonstrate why the instabilities and inaccuracies observed in the examples occur. In addition to the material on finite difference methods, the book outlines other methods and analysis techniques and gives references for further reading.

W. S. Ames, *Numerical methods for partial differential equations*, 3rd edition, Academic Press, 1992, ISBN: 012056761X. This book gives a thorough discussion of the application of all types of numerical method (including finite difference, finite element, and spectral methods) to the three main classes of partial differential equation. It includes worked examples and has a separate section on special topics that were of interest to researchers at the time of writing the book such as shock waves and singularities.

J. D. Lambert, *Numerical Methods for Ordinary Differential Systems*, John Wiley, 1991, ISBN. This book gives a thorough introduction to the range of methods that can be applied to the initial value problem of a system of ordinary differential equations. It includes a lengthy section on stiff problems, and recommendations of which methods are suitable for which types of problems.

The following books describe methods and algorithms that are needed for implementing most solution methods, including matrix methods and integration and quadrature methods.

G. H. Golub and C. F. van Loan, *Matrix Computations*, Johns Hopkins University Press, Third Edition, 1996, ISBN: 0801854148.

This book gives efficient, stable algorithms for carrying out a wide variety of matrix computations in finite precision arithmetic. Many of the algorithms are included in the free LAPACK routines. This type of computation is essential for anyone wishing to create their own continuous modelling software.

W. Press, S. A. Teukolsky, W. T. Vetterling, and B. Flannery, *Numerical recipes in C/C++/F77/F90*, CUP. ISBN: 0521431085 (C), 0521750334 (C++), 052143064X (F77), 0521574390 (F90).

This book contains descriptions of algorithms for solving a variety of common problems and the mathematics that underpins them. The code sections are not always ideal for the language in which they are written, but much of the mathematics is sound.

12. References

All reports marked * are or will be available for download from the SS*f*M website, http://www.npl.co.uk/ssfm/download/

[1] T. J. Esward, G. J. Lord, and L. Wright, *Model Validation in Continuous Modelling*. NPL Report No. CMSC 29/03, December 2003.*

[2] G. J. Lord and L. Wright, *Uncertainty Evaluation in Continuous Modelling*. NPL report No. CMSC 31/03, 2003.*

[3] T. J. Esward and L. Wright, *Guide to the use of finite element and finite difference software*. NPL report No. 30/03, 2003.*

[4] M. G. Cox, A. B. Forbes, and P. M. Harris, *Software Support for Metrology Best Practice Guide No. 4: Discrete Modelling and Experimental Data Analysis*. April 2004.*

[5] T. J. Esward, *Mathematical modelling for metrology at the nanoscale* NPL report due 2006.*

[6] http://www.mat.univie.ac.at/~neum/glopt.html

[7] http://www-fp.mcs.anl.gov/otc/Guide/

[8] http://www.nag.co.uk/numeric/numerical_libraries.asp

[9] http://www.mathworks.co.uk/products/matlab/

[10] P. Macey, *Sonar array optimization using a coupled finite/boundary element formulation* Presented at conference *Boundary Integral Methods: Theory and Applications*, in the workshop *Innovative Boundary Element Techniques in Computational acoustics and Electromagnetics* at University of Bath on 15th September 2000.

[11] T. J. Esward and L. Wright, *Finite Element Optimisation For Problems In Acoustics*. Proceedings of the National Measurement Conference 2001.

[12] D. Givoli, *Non-Reflecting Boundary Conditions*, J Comp Phys, v**94**, p1-29, 1991.

[13] J. P. Berenger, *A perfectly matched layer for the absorption of electromagnetic waves*, J Comp Phys, v**114**, p185-200, Oct 1994.

[14] L N McCartney, *No time – gentlemen please!* Philosophical Magazine, 1976, v**33**(4), p689-695.

[15] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, and OIML, *Guide to the Expression of Uncertainty in Measurement*. ISBN 92-67-10188-9, Second Edition.

[16] T. J. Esward, N. J. McCormick, K. M. Lawrence and M. J. Stevens, *Distributed computing for metrology applications*, forthcoming SS*f*M Good Practice Guide, due March 2006.*

[17] J. D. Lambert, *Numerical Methods for Ordinary Differential Systems: The Initial Value Problem*, John Wiley & Sons, 1991, ISBN 0471 92990 5.

[18] O.C. Zienkiewicz & R.L Taylor, *The Finite Element Method, Volume 1: The Basis*, Fifth Edition, Butterworth-Heinemann, 2002, ISBN 0750 65049 4

[19] http://www.netlib.org

[20] http://www.uk.comsol.com/products/multiphysics/

[21] http://www.abaqus.co.uk

[22] H. Cook, M. Cox, M. P. Dainton and P. M. Harris, *A Methodology for Testing Spreadsheets and Other Packages Used in Metrology*, NPL Report CISE 25/99, September 1999.*

[23] T. J. Esward, E. G. Johnson, and L. Wright, *Testing Continuous Modelling Software*, NPL Report CMSC 41/04, March 2004.*

[24] J. Gilby and J. Walton. *SSfM Good Practice Guide 13: Data Visualisation*, revised 2005.*

[25] L. Wright, *Visualisation, Uncertainties, and Continuous Modelling*, NPL Report No. CMSC 39/04, February 2004. *

[26] Y. S. Touloukian, R. W. Powell, C. Y. Ho and P. G. Klemens, *Thermal Conductivity: Nonmetallic Solids, Thermophysical Properties of Matter, The TPRC Data Series, 2*, New York, Plenum Press, 1978.

[27] P. Jimeno Largo, *Novel high temperature metal-carbon eutectic fixed points: influence of impurities and of the temperature distribution at the radiant cavity*, Doctoral thesis, Universidad de Valladolid, 2005.

[28] http://www.ansys.com

[29] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, and OIML, *Guide to the Expression of Uncertainty in Measurement, Supplement 1, Numerical Methods for the Propagation of Distributions*, To be published.

[30] A. J. Olsson, G. E. Sandberg, *Latin Hypercube Sampling for Stochastic Finite Element Analysis*, Journal of Engineering Mechanics, **128**, p121-125, 2002.

[31] A. Saltelli, K. Chan, and E. M. Scott (eds.) *Sensitivity Analysis*, Wiley, ISBN 0-47-199892-3, 2000.

[32] J. Geiser, O. Klein, and P. Philip, *WIAS-HITNIHS: A software tool for simulation in sublimation growth of SiC single crystals: Applications and methods*, International Congress of Nanotechnology 2004, San Francisco.

[33] L. Wright, S. Robinson, V. Humphrey, P. Harris and G. Hayman, *The application of boundary element methods to near-field acoustic measurements on cylindrical surfaces at NPL*, NPL Report DQL-AC 011, January 2005.

[34] T. J. Esward, K. Lees, D Sayers, and L. Wright, *Testing Continuous Modelling Software: Three Case Studies*, NPL Report CMSC 42/04. March 2004.

[35] A. J. Burton and G F. Miller, *The application of integral equation methods to the numerical solution of some exterior boundary value problems*, Proc. Royal Society, London Vol **A323**, p 201-210, 1971.

[36] H. A. Schenck, *Improved integral formulation for acoustic radiation problems*, J. Acoust. Soc. Am., Vol. **44**, p 41-58, 1967.

[37] B. Duncan and L. Crocker, *Measurement Good Practice Guide 45: Characterisation of Flexible Adhesives for Design* Appendix I (Hyperelastic Models), July 2001, available from http://www.npl.co.uk/adhesives/fepublications.html#mgpg

[38] L. N. McCartney and S. E. T. Saunderson, *A recursive method of calculating stress transfer in multiple-ply cross-ply laminates subject to biaxial loading. I: Development of model*, NPL report DMM(A)150, June 1995.

[39] http://www.emclab.umr.edu/EMoverview.html

[40] C. Christopoulos, *The Transmission-Line Modelling Method TLM*, John Wiley & Sons Inc ISBN: 0-78-031017-9, 1995.

[41] C. Hafner, *The Generalized Multipole Technique for Computational Electromagnetics*, Artech House, ISBN 0-89-006429-6, 1990.

[42] T. J. Chung, *Computational Fluid Dynamics*, Cambridge University Press, ISBN 0-52-159416-2, 2002.

[43] M. Hamilton and D. Blackstock, *Nonlinear Acoustics: Theory and Applications*, Academic Press, ISBN 0-12-321860-8, 1998.

[44] Eric W. Weisstein, *Ordinary Differential Equation* From MathWorld--A Wolfram Web Resource. http://mathworld.wolfram.com/OrdinaryDifferentialEquation.html

[45] M. Abramowitz and I. A. Stegun, (eds.). *Legendre Functions* and *Orthogonal Polynomials*, Chs. 8 and 22 in *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 9th printing. New York: Dover, pp. 331-339 and 771-802, 1972.

[46] C. M. Bender and S. A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers*, McGraw-Hill International Editions, ISBN 0-07-Y66173-1, 1978.

[47] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran: The Art of Scientific Computing*, Second Edition, Chapters 16 & 17, Cambridge University Press, ISBN 0 52 143064 X, 1992.

[48] B. Krietenstein, R. Schuhmann, P. Thoma and T. Weiland, *The Perfect Boundary Approximation technique facing the challenge of high precision field computation*, Proc. XIX International Linear Accerlerator Conference (LINAC '98), Chicago, USA. p860-862, 1998.