Cambridge University Press
978-0-521-65007-6 - Mathematical Modeling: Case Studies from Industry
Edited by Ellis Cumberbatch and Alistair Fitt
Excerpt
More information

Introduction

"Technology transfer" has become one of the most well-used phrases of the end of the millennium. The realisation that the worlds of industry and academia cannot fruitfully progress separately has inspired both communities to build strong and mutually beneficial relationships. Often this has meant industry hiring individual professors as consultants, or industry supporting post-doctoral fellows (common in chemistry). An alternative structure has been the utilisation of a quasi-governmental organisation as a go-between, such as NACA/NASA in the US and the Aeronautical Research Council in the UK for work in aeronautics. **Direct** contact between **industry** and the **mathematics community** is more recent, achieving recognition via degree programmes, math-in-industry conferences and journals, all now in a global context. Half of the chapters in this book are products of Study Groups with Industry and Math Clinics, and to some extent the other half derive from similar direct interactions with industry prompted by the successes of these two initiatives.

Study Groups with Industry started in Oxford in 1968 when a small group of applied mathematicians (led by Alan Tayler and Leslie Fox) spent a week problem solving at Oxford University in conjunction with invited representatives from industry. A similar meeting has been held every year since. What *has* changed in recent years is the global nature of this industrial/academic collaboration. In the first 20 or so years Study Groups with Industry only happened in Oxford, and only once a year: by 1999 meetings running on similar lines to the Oxford model had also taken place in Australia, Canada, Denmark, Holland, Indonesia, Mexico, Norway, USA and other locations, in the same year. In many of these ventures domestic applied mathematics societies have encouraged and supported this expansion: pivotal roles have been played in Europe by ECMI (the European Consortium for Mathematics in Industry), in the USA and Canada by SIAM (Society for Industrial and Applied Mathematics) and PIMS (Pacific Institute for the Mathematical Sciences), in Mexico by SMM (Sociedad Matematica Mexicana) and in

Australia by ANZIAM (Australian and New Zealand Industrial and Applied Mathematics).

The *modus operandi* of Study Groups with Industry is now well established: ahead of the meeting there is solicitation from industry for the submission of problems, and probably some discussion regarding appropriateness for a week of brain storming. On the first day of the meeting, the academics, their graduate students, and representatives from industry gather. After each industrialist has described their specific unsolved problem, a room is allotted to each industry, and informal groups are formed. The next three days are spent working intensively on each problem, guided and assisted by the representative from industry. Fierce debate often rages: theories and countertheories come and go, and blackboards are filled with equations. Because of the time constraints that apply, the evenings are often used to work on the problems. Normally, a consensus is eventually reached. On the final day of the meeting, progress on each of the problems is summarised and subsequently a technical problem report is published.

Math Clinics originated at the Claremont Colleges in 1973. The structure here has a longer time frame: problems are solicited over the summer, and work on them is contracted for an academic year. Each problem is addressed over this time period by a team consisting of a faculty supervisor and 4–6 graduate and undergraduate students (one of whom is team leader, a managerial position). The team has the time to do basic research, and during the year makes a number of oral and written reports. Students are assigned grades as for a traditional course.

In both the Study Group and Clinic operations, the involvement of students is important. How problems originate, how industrialists expect them to be modelled, and what they expect from an answer, is valuable experience for students aiming for industrial employment, or starting on research in applied mathematics. Additionally students working on these problems learn the dynamics of group work, and the importance of good oral and writing skills.

What type of problems appear at Study Groups and Clinics? This question is difficult to answer, for the range is huge. Not only do the problems vary enormously in physical and mathematical terms, but the reasons that the industrialist has for bringing a specific problem may vary. Typical "industry questions" all begin "we have a process. . . " but may have a range of endings such as

- ". . . which is well-understood and has worked well for years, but we suspect that it could be made more efficient: how can we optimise it?"

- "... which has worked well for years, but we don't have much idea why. Now we're trying to extend and change it. How does it work?"
- "... which normally works pretty well. Sometimes, though, something goes wrong. Under what circumstances might we expect this to happen, what should we measure to give us warning signs and how might we cure the problem?"
- "... in mind which is new and very promising. Before running expensive and/or dangerous experiments, we'd like to get an idea of how things might turn out."
- "... which works well and we know is safe. But we have to satisfy an external regulatory body to *prove* that it is safe. Can you model the 'worst case' circumstances for us?"
- "... which we have simulated (ourselves, or bought software for). The data and the simulations do not agree. What is going wrong?"

As far as the representatives from industry are concerned, finding the answers to questions of this sort is usually the main priority. There are many other possible advantages of the interaction. The opportunity to spend time discussing their problem with a group of experts is attractive, as is the chance to see what sort of problems other industries have and how they cope with them. Frequently the discussion widens, and other problems may be considered. Finally, there is always the chance to build more longer lasting professional relationships and to recruit promising students.

For the process to be mutually beneficial, everybody concerned must have a good reason for wanting to take part. Some of the benefits for students have been referred to above. What about faculty? It turns out that there are many possible reasons for participating. These include:

- A constant supply of interesting and novel problems which often lead to publications in leading journals
- The chance to form closer relationships with companies which may lead to joint studentships and research contracts
- An opportunity to broaden the range of new problem areas leading to a freshening of the teaching syllabus
- The transfer of mathematics across applications
- A chance to work as part of a team – frequently a novelty for mathematicians.

How much does all of this cost? Considering the possible benefits to both industry and the academics involved, the sums involved are surprisingly small. British Study Groups cost between 10,000–15,000 pounds to run; the money is provided by the participating industries (each problem presenter is asked to contribute about 1000 pounds), and various other organisations such as the London Mathematical Society and (until recently) EPSRC also contribute much-needed revenue. As far as the industrial participants are concerned, the biggest cost is normally that of having their expert in a particular field away from the office for a whole week. For clinics, the financial details are rather more large-scale; each clinic project costs the industry concerned a sum between 35,000 and 45,000 dollars. Although this might sound a lot, for their money the client gets a year's work from a dedicated group of experts: the productivity often equals that of an engineer-year which costs the employer five times as much. Do they consider this money well spent? This is a question that can only truly be answered by the industrial scientists themselves, but it is significant that a large number of projects have been brought to the Claremont clinics by "repeat customers".

Thirteen different problems are considered in this book; all originate from real collaborations with industry. Although some of them are a little difficult to classify, four of the problems are recognisably elliptic, six are parabolic, and two are hyperbolic. This distribution may be said to fairly represent the frequency of each type of problem that is normally encountered at Study Groups.

The order chosen for the articles is related to ordinary and partial differential equation classification. Chapters 1–4 concern models based on essentially elliptic partial differential equations. Chapters 5–10 all involve some form of diffusion, with nonlinear effects, convection and reaction, and finally radiation being successively introduced. The classification of the underlying equations in chapter 11 is less clear; the equations contain both parabolic and hyperbolic features. The concluding two chapters, 12 and 13, address two hyperbolic problems.

CAMBRIDGE

Cambridge University Press
978-0-521-65007-6 - Mathematical Modeling: Case Studies from Industry
Edited by Ellis Cumberbatch and Alistair Fitt
Excerpt
More information

0

Mathematical Preliminaries

0.1 The Continuum Model

Most of the modelling introduced in the following chapters uses the continuum approach. In this introductory chapter, we list the commonly used equations: those describing diffusion, convection, radiation, and fluid and solid mechanics. We do not attempt to give an even partially rigorous derivation of any of these equations; our purpose is to provide a ready resource, and to indicate source books of wider scope. Above all, this section should be seen as answering the question "why did they start from *those* equations?".

Let us approach the continuum model by considering the example of diffusion. Diffusion is a molecular process. Consider the diffusion of heat: the diffusion happens because a hot region of a material has molecules of higher energy than those in cooler parts. Energy equalisation therefore takes place by molecular interaction – and the heat is said to "diffuse". To enable us to view this at the continuum level, local averages (for example, over many molecules) are taken: the molecular picture is smeared. The concept of heat as a variable having a value only at molecular sites is replaced by a framework in which heat is regarded as a variable that has continuous values. The laws governing changes in the continuous functions to describe heat transfer are treated "phenomenologically". Models are created at both levels; their usefulness, success, and relevance depends on the application.

0.1.1 Conservation Laws

Physical phenomena expressed in the continuum paradigm are usually phrased in terms of conservation laws. Let us introduce this in a generic fashion: we consider a substance, A, distributed continuously. For a subset V of the continuum,

$$\int_V A \, dV$$

Cambridge University Press
978-0-521-65007-6 - Mathematical Modeling: Case Studies from Industry
Edited by Ellis Cumberbatch and Alistair Fitt
Excerpt
More information

is the amount of our substance contained in V, so that $A(\boldsymbol{x})$ measures the local amount of A per unit volume at location $\boldsymbol{x}$ at time t. We anticipate that A is changing with time and with location, within (and outside) V, by its own motion relative to the continuum and by the motion of the continuum carrying it around. Let $\boldsymbol{Q}$ denote the flux of A, defined so that the amount of our substance flowing per unit time across the surface S surrounding V is

$$\int_S \boldsymbol{Q} \cdot \boldsymbol{n}\, dS,$$

where $\boldsymbol{n}$ is the unit outward normal to S. This definition means that, locally, for a small, flat element of area dS and normal $\boldsymbol{n}$, the amount of substance flowing across dS per unit time is $\boldsymbol{Q} \cdot \boldsymbol{n}\ dS$. We now assert that the rate of change of the amount of substance within V is due only to the amount flowing across its boundary. (That is, we assume no sources or sinks of A within V.) This gives

$$\frac{\partial}{\partial t} \int_V \mathrm{A}\, dV = -\int_S \boldsymbol{Q} \cdot \boldsymbol{n}\, dS. \qquad (0.1)$$

The right-hand side is transformed to a volume integral by use of the divergence theorem. Also, for a volume V fixed in space, the time derivative may be taken inside the integral. Thus (0.1) becomes

$$\int_V \left(\frac{\partial \mathrm{A}}{\partial t} + \nabla \cdot \boldsymbol{Q} \right) dV = 0. \qquad (0.2)$$

Since V is an arbitrary volume, it follows that the integrand is identically zero throughout the continuum. This argument relies on the assumption that the integrand is continuous, and it fails in its absence. (The latter occurs when shock waves are present.) This gives the *local* equation

$$\frac{\partial A}{\partial t} + \nabla \cdot \boldsymbol{Q} = 0. \qquad (0.3)$$

This equation is evident in various forms as we quote equations of continuum mechanics below.

0.2 Diffusion

The phenomenon of diffusion is all around us. We drink hot tea, and its heat, concentrated at first, is diffused throughout the body. The process is helped by warmed blood flowing to the colder extremities – this is advection. Similar processes are happening everywhere – concentrations of heat, moisture, salt, smoke, and other pollutants, are being dispersed in the atmosphere and oceans by the processes of diffusion and advection.

Let us consider pure diffusion first, and use the diffusion of heat as an illustrative example. The principle of the conservation of heat induces us to define the quantity A in (0.1) to be

$$A = \rho c_p T \tag{0.4}$$

where ρ is the mass density of the material in which the heat resides, c_p is the specific heat of the material, and T is its temperature. The flux of heat across a surface is taken as

$$\boldsymbol{Q} = -k\nabla T \tag{0.5}$$

where the factor k is called the thermal conductivity. This model, which asserts that the rate of heat flow is proportional to the temperature gradient (heat flows from high temperatures to colder ones) seems a natural one – and it works wonderfully well! With the substitutions of A and $\boldsymbol{Q}$ from (0.4) and (0.5) into the general conservation law (0.3), we find that

$$\frac{\partial T}{\partial t} = D\nabla^2 T \tag{0.6}$$

where

$$D = \frac{k}{\rho c_p} \tag{0.7}$$

is called the thermal diffusivity. In obtaining this equation, we have assumed that the component factors that make up D are all constant. This assumption may not hold in various applications, in which case T is governed by a more complicated equation than (0.6).

In general, let us consider the variable A in (0.1) to be the concentration $c(\boldsymbol{x}, t)$ of a substance which is diffusing. The flux law for the diffusion of c similar to (0.5) is given by $\boldsymbol{Q} = -D\nabla C$. This is known as Fick's law. This alone would give (0.6) as the conservation law, with c replacing T. However, let us now add the process of advection to Fick's law. In addition to the concentration gradient ∇c transferring substance within the material, there is also material flow. If $\boldsymbol{u}(\boldsymbol{x}, t)$ denotes the velocity of the medium, the material crossing a surface element will contribute a term $c\boldsymbol{u} \cdot \boldsymbol{n}$ to the rate of transfer of c. Hence Fick's law when diffusion and advection are both present reads

$$\boldsymbol{Q} = -D\nabla c + c\boldsymbol{u} \tag{0.8}$$

Cambridge University Press
978-0-521-65007-6 - Mathematical Modeling: Case Studies from Industry
Edited by Ellis Cumberbatch and Alistair Fitt
Excerpt
More information

and (0.3) becomes

$$\frac{\partial c}{\partial t} + \nabla \cdot (c\boldsymbol{u}) = D\nabla^2 c. \tag{0.9}$$

Often the material velocity $\boldsymbol{u}$ may be taken as constant, in which case (0.9) simplifies. In general, however, $\boldsymbol{u}$ satisfies other equations which must be coupled with (0.9).

The equations of motion require boundary conditions. Let us now discuss these with relation to the heat diffusion equation (0.6); extension to the general diffusion equation (0.9) is straightforward. On a surface S at which the temperature is constant or is specified, we have

$$T = \text{constant} \qquad \text{or} \qquad T = F(\boldsymbol{x}, t) \tag{0.10}$$

for $\boldsymbol{x} \in S$.

On surfaces S that are insulated or across which the flux of heat is specified, we have

$$\frac{\partial T}{\partial n} = 0 \qquad \text{or} \qquad \frac{\partial T}{\partial n} = G(\boldsymbol{x}, t) \tag{0.11}$$

where $\boldsymbol{n}$ is the unit normal to the surface S and again $\boldsymbol{x} \in S$. There are situations for which neither (0.10) nor (0.11) apply, and a mixed boundary condition

$$\frac{\partial T}{\partial n} = \alpha(T - T_0) \tag{0.12}$$

is more appropriate for $\boldsymbol{x} \in S$. This boundary condition relates the flux of heat transferred across a boundary to the difference between the boundary temperature and the ambient temperature. It is often used at a solid–fluid interface where the fluid may be in motion. The coefficient α is an empirical constant, varying with the situation encountered. This boundary condition is often called the “radiation” boundary condition, though it is not related to radiative heat transfer.

The mathematical designation for (0.10), (0.11) and (0.12) are Dirichlet, Neumann, and Robin (or mixed) boundary conditions, respectively: for extensive discussion of the heat diffusion equation, its properties, and its engineering applications, the reader is referred to [4], [5], [11] and [22].

A boundary condition met in many applications concerns change of phase at an interface. In the case of the heat equation, the conditions are then known as the Stefan boundary conditions. In particular, consider modelling the melting of ice and, for illustrative purposes, take a plane geometry. Ice, at temperature $T = 0$°C, initially occupies the half-space $x > 0$. Beginning at $t = 0$ the ice is melted by applying a heat source at $x = 0$ which maintains a temperature $T = T_o > 0$ there. At time $t > 0$ there is water in $0 < x < s(t)$ and ice in $x > s(t)$, where $x = s(t)$ denotes the moving interface that separates ice from water. The Stefan boundary conditions at this interface are

$$T = 0 \quad \text{and} \quad -k\frac{\partial T}{\partial x} = L\rho\frac{ds}{dt} \quad \text{at} \quad x = s(t). \tag{0.13}$$

The boundary condition (0.13) is to be used in finding the temperature profile in $0 < x < s(t)$. In (0.13) L is the coefficient of latent heat, that is, the amount of heat per unit mass required to change ice into water at $T = 0^o$C. Equation (0.13) gives two equations to be satisfied at $x = s(t)$, as distinct to the single condition in (0.10), (0.11), or (0.12) prescribed previously. Since the location of the ice–water interface is not known *a priori*, two conditions are necessary to solve for the temperature profile and the location $x = s(t)$ simultaneously. Problems of this type are known as free or moving boundary value problems; they are generally difficult to solve as they are nonlinear. For further details the reader is referred to [8] and [21].

0.3 Flow and Deformation of Solids and Fluids

In many of the problems in this book, a fluid or a solid deforms under the action of external forces; to solve the problem under consideration it is necessary to determine how the fluid or solid moves. We therefore need to consider how to propose general equations for the flow of such materials.

We distinguish between four states of matter: solid, fluid, gas, and plasma. Determining exactly which of these four states of matter best describes a given piece of matter is a highly nontrivial process, and we do not intend to pursue this in a detailed manner; rather we content ourselves with simple “coffee table” explanations. Consider a bucket of matter. If the bucket contains a solid, then the constituent particles of the solid are disinclined to move relative to each other. Moreover, if the bucket is “poured out”, then the solid will retain its shape. If the bucket contains a fluid, then particles can move relative to each other with much greater ease. Inversion of the bucket will now result in a differently shaped fluid. If the bucket contains a gas, then the gas will not retain its “shape” even *before* the bucket is poured out. Our approach is thus broadly as follows: we acknowledge that a plasma state of matter exists, but give no attention to modelling it, and we treat gases as “fluids” since the constituent particles of both gases and fluids can move relative to each other with ease.

0.3.1 Equations of Motion

To model the flow of solids and fluids, we shall assume that the rate of change of momentum is equal to the sum of the applied forces; this is Newton's second law.

The applied forces separate into two categories. There are body forces $\boldsymbol{b}$ acting on each continuum *volume* element: the force per unit volume is $\rho\boldsymbol{b}$, where ρ is the mass density. The simplest body force (and the only one, for most of non-electromagnetic physics) is that due to gravity. The second category of force for a continuum is that acting on *surface* elements. This is manifested by considering simple actions: (1) blowing up a balloon so that increasing the interior pressure expands the balloon volume until it equalises with the exterior ambient pressure. These pressures are expressed in terms of force per unit surface area, and elementary physics shows that in the static case the pressure acts in a direction normal to surface elements; (2) rotating a solid body, say a cylinder, immersed successively in two different fluids requires different torques to achieve the same angular velocity – the "stickier" (more viscous) the fluid the higher the torque required. It is evident here that there are surface forces acting tangential to the surface of the moving cylinder.

In general, the surface force acting on a surface with normal $\boldsymbol{n}$ has both normal and tangential components, and this force will depend on the orientation of $\boldsymbol{n}$. Again, elementary physical principles generate expressions for the force per unit area, $\boldsymbol{F}$, in terms of the normal $\boldsymbol{n}$: since $\boldsymbol{F}$ and $\boldsymbol{n}$ each have three components, the magnitude factor must have nine components, as in matrix multiplication. In terms of components,

$$F_i = \sum_{j=1}^{3} T_{ij} n_j$$

where

$$\boldsymbol{T} = \{T_{ij}\} \tag{0.14}$$

is the stress tensor or matrix.

A conservation law for linear momentum similar to (0.1) may now be invoked. Since the surface forces (0.14) may be expressed as a product of the stress and the normal vector, use of the divergence theorem allows the surface integral to be expressed in terms of a volume integral, as in (0.2). The local equation for conservation of linear momentum is therefore derived, similar to (0.3), as

$$\rho\boldsymbol{a} = \operatorname{div} \boldsymbol{T} + \rho\boldsymbol{b}. \tag{0.15}$$