
1

Introduction

Hyperbolic systems of partial differential equations can be used to model a wide variety of phenomena that involve wave motion or the advective transport of substances. This chapter contains a brief introduction to some of the fundamental concepts and an overview of the primary issues discussed in this book.

The problems we consider are generally time-dependent, so that the solution depends on time as well as one or more spatial variables. In one space dimension, a homogeneous first-order system of partial differential equations in x and t has the form

$$q_t(x, t) + Aq_x(x, t) = 0 \tag{1.1}$$

in the simplest constant-coefficient linear case. Here $q : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^m$ is a vector with m components representing the unknown functions (pressure, velocity, etc.) we wish to determine, and A is a constant $m \times m$ real matrix. In order for this problem to be *hyperbolic*, the matrix must satisfy certain properties discussed below. Note that subscripts are used to denote partial derivatives with respect to t and x .

The simplest case is the constant-coefficient *scalar* problem, in which $m = 1$ and the matrix A reduces to a scalar value. This problem is hyperbolic provided the scalar A is real. Already this simple equation can model either advective transport or wave motion, depending on the context.

Advective transport refers to a substance being carried along with fluid motion. For example, consider a contaminant being advected downstream with some fluid flowing through a one-dimensional pipe at constant velocity $\bar{u}$. Then the concentration or density $q(x, t)$ of the contaminant satisfies a scalar advection equation of the form

$$q_t(x, t) + \bar{u}q_x(x, t) = 0, \tag{1.2}$$

as derived in Chapter 2. It is easy to verify that this equation admits solutions of the form

$$q(x, t) = \tilde{q}(x - \bar{u}t) \tag{1.3}$$

for any function $\tilde{q}(\xi)$. The concentration profile (or waveform) specified by $\tilde{q}$ simply propagates with constant speed $\bar{u}$ and unchanged shape. In this context the equation (1.2) is generally called the *advection equation*.

The phenomenon of *wave motion* is observed in its most basic form if we model a sound wave traveling down a tube of gas or through an elastic solid. In this case the molecules of

the gas or solid barely move, and yet a distinct wave can propagate through the material with its shape essentially unchanged over long distances, and at a speed c (the speed of sound in the material) that is much larger than the velocity of material particles. We will see in Chapter 2 that a sound wave propagating in one direction (to the right with speed $c > 0$) can be modeled by the equation

$$w_t(x, t) + cw_x(x, t) = 0, \quad (1.4)$$

where $w(x, t)$ is an appropriate combination of the pressure and particle velocity. This again has the form of a scalar first-order hyperbolic equation. In this context the equation (1.4) is sometimes called the *one-way wave equation* because it models waves propagating in one particular direction.

Mathematically the advection equation (1.2) and the one-way wave equation (1.4) are identical, which suggests that advective transport and wave phenomena can be handled by similar mathematical and numerical techniques.

To model acoustic waves propagating in both directions along a one-dimensional medium, we must consider the full acoustic equations derived in Chapter 2,

$$\begin{aligned} p_t(x, t) + Ku_x(x, t) &= 0, \\ u_t(x, t) + (1/\rho)p_x(x, t) &= 0, \end{aligned} \quad (1.5)$$

where $p(x, t)$ is the pressure (or more properly the perturbation from some background constant pressure), and $u(x, t)$ is the particle velocity. These are the unknown functions to be determined. The material is described by the constants K (the bulk modulus of compressibility) and ρ (the density). The system (1.5) can be written as the first-order system $q_t + Aq_x = 0$, where

$$q = \begin{bmatrix} p \\ u \end{bmatrix}, \quad A = \begin{bmatrix} 0 & K \\ 1/\rho & 0 \end{bmatrix}. \quad (1.6)$$

To connect this with the one-way wave equation (1.4), let

$$w^1(x, t) = p(x, t) + \rho cu(x, t),$$

where $c = \sqrt{K/\rho}$. Then it is easy to check that $w^1(x, t)$ satisfies the equation

$$w_t^1 + cw_x^1 = 0$$

and so we see that c can be identified as the speed of sound. On the other hand, the function

$$w^2(x, t) = p(x, t) - \rho cu(x, t)$$

satisfies the equation

$$w_t^2 - cw_x^2 = 0.$$

This is also a one-way wave equation, but with propagation speed $-c$. This equation has solutions of the form $q^2(x, t) = \tilde{q}(x + ct)$ and models acoustic waves propagating to the left at the speed of sound, rather than to the right.

The system (1.5) of two equations can thus be decomposed into two scalar equations modeling the two distinct acoustic waves moving in different directions. This is a fundamental theme of hyperbolic equations and crucial to the methods developed in this book. We will see that this type of decomposition is possible more generally for hyperbolic systems, and in fact the definition of “hyperbolic” is directly connected to this. We say that the constant-coefficient system (1.1) is *hyperbolic* if the matrix A has real eigenvalues and a corresponding set of m linearly independent eigenvectors. This means that any vector in $\mathbb{R}^m$ can be uniquely decomposed as a linear combination of these eigenvectors. As we will see in Chapter 3, this provides the decomposition into distinct waves. The corresponding eigenvalues of A give the wave speeds at which each wave propagates. For example, the acoustics matrix A of (1.6) has eigenvalues $-c$ and $+c$, the speeds at which acoustic waves can travel in this one-dimensional medium.

For simple acoustic waves, some readers may be more familiar with the *second-order wave equation*

$$p_{tt} = c^2 p_{xx}. \quad (1.7)$$

This equation for the pressure can be obtained from the system (1.5) by differentiating the first equation with respect to t and the second with respect to x , and then eliminating the u_{xt} terms. The equation (1.7) is also called a *hyperbolic* equation according to the standard classification of second-order linear equations into hyperbolic, parabolic, and elliptic equations (see [234], for example). In this book we only consider first-order hyperbolic systems as described above. This form is more fundamental physically than the derived second-order equation, and is more amenable to the development of high-resolution finite volume methods.

In practical problems there is often a coupling of advective transport and wave motion. For example, we will see that the speed of sound in a gas generally depends on the density and pressure of the gas. If these properties of the gas vary in space and the gas is flowing, then these variations will be advected with the flow. This will have an effect on any sound waves propagating through the gas. Moreover, these variations will typically cause acceleration of the gas and have a direct effect on the fluid motion itself, which can also be modeled as wave-propagation phenomena. This coupling leads to *nonlinearity* in the equations.

1.1 Conservation Laws

Much of this book is concerned with an important class of homogeneous hyperbolic equations called *conservation laws*. The simplest example of a one-dimensional conservation law is the partial differential equation (PDE)

$$q_t(x, t) + f(q(x, t))_x = 0, \quad (1.8)$$

where $f(q)$ is the *flux function*. Rewriting this in the *quasilinear form*

$$q_t + f'(q)q_x = 0 \quad (1.9)$$

suggests that the equation is hyperbolic if the flux Jacobian matrix $f'(q)$ satisfies the conditions previously given for the matrix A . In fact the linear problem (1.1) is a conservation

law with the linear flux function $f(q) = Aq$. Many physical problems give rise to *nonlinear conservation laws* in which $f(q)$ is a nonlinear function of q , a vector of *conserved quantities*.

1.1.1 Integral Form

Conservation laws typically arise most naturally from physical laws in an integral form as developed in Chapter 2, stating that for any two points x_1 and x_2 ,

$$\frac{d}{dt} \int_{x_1}^{x_2} q(x, t) dx = f(q(x_1, t)) - f(q(x_2, t)). \quad (1.10)$$

Each component of q measures the density of some conserved quantity, and the equation (1.10) simply states that the “total mass” of this quantity between any two points can change only due to the flux past the endpoints. Such conservation laws naturally hold for many fundamental physical quantities. For example, the advection equation (1.2) for the density of a contaminant is derived from the fact that the total mass of the contaminant is conserved as it flows down the pipe and the flux function is $f(q) = \bar{u}q$. If the total mass of contaminant is not conserved, because of chemical reactions taking place, for example, then the conservation law must also contain *source terms* as described in Section 2.5, Chapter 17, and elsewhere.

The constant-coefficient linear acoustics equations (1.5) can be viewed as conservation laws for pressure and velocity. Physically, however, these are not conserved quantities except approximately in the case of very small amplitude disturbances in uniform media. In Section 2.7 the acoustics equations are derived from the *Euler equations* of gas dynamics, the nonlinear conservation laws that model more general disturbances in a compressible gas. These equations model the conservation of mass, momentum, and energy, and the laws of physics determine the flux functions. See Section 2.6 and Chapter 14 for these derivations. These equations have been intensively studied and used in countless computations because of their importance in aerodynamics and elsewhere.

There are many other systems of conservation laws that are important in various applications, and several are used in this book as examples. However, the Euler equations play a special role in the historical development of the techniques discussed in this book. Much of the mathematical theory of nonlinear conservation laws was developed with these equations in mind, and many numerical methods were developed specifically for this system. So, although the theory and methods are applicable much more widely, a good knowledge of the Euler equations is required in order to read much of the available literature and benefit from these developments. A brief introduction is given in Chapter 14. It is a good idea to become familiar with these equations even if your primary interest is far from gas dynamics.

1.1.2 Discontinuous Solutions

The differential equation (1.8) can be derived from the integral equation (1.10) by simple manipulations (see Chapter 2) *provided that q and $f(q)$ are sufficiently smooth*. This proviso is important because in practice many interesting solutions are not smooth, but contain discontinuities such as shock waves. A fundamental feature of nonlinear conservation laws

is that these discontinuities can easily develop spontaneously even from smooth initial data, and so they must be dealt with both mathematically and computationally.

At a discontinuity in q , the partial differential equation (1.8) does not hold in the classical sense and it is important to remember that the integral conservation law (1.10) is the more fundamental equation which does continue to hold. A rich mathematical theory of shock-wave solutions to conservation laws has been developed. This theory is introduced starting in Chapter 11.

1.2 Finite Volume Methods

Discontinuities lead to computational difficulties and the main subject of this book is the accurate approximation of such solutions. Classical finite difference methods, in which derivatives are approximated by finite differences, can be expected to break down near discontinuities in the solution where the differential equation does not hold. This book concerns finite volume methods, which are based on the integral form (1.10) instead of the differential equation. Rather than pointwise approximations at grid points, we break the domain into *grid cells* and approximate the total integral of q over each grid cell, or actually the *cell average* of q , which is this integral divided by the volume of the cell. These values are modified in each time step by the flux through the edges of the grid cells, and the primary problem is to determine good *numerical flux functions* that approximate the correct fluxes reasonably well, based on the approximate cell averages, the only data available. We will concentrate primarily on one class of *high-resolution* finite volume methods that have proved to be very effective for computing discontinuous solutions. See Section 6.3 for an introduction to the properties of these methods.

Other classes of methods have also been applied to hyperbolic equations, such as finite element methods and spectral methods. These are not discussed directly in this book, although much of the material presented here is good background for understanding high-resolution versions.

1.2.1 Riemann Problems

A fundamental tool in the development of finite volume methods is the *Riemann problem*, which is simply the hyperbolic equation together with special initial data. The data is piecewise constant with a single jump discontinuity at some point, say $x = 0$,

$$q(x, 0) = \begin{cases} q_l & \text{if } x < 0, \\ q_r & \text{if } x > 0. \end{cases} \quad (1.11)$$

If Q_{i-1} and Q_i are the cell averages in two neighboring grid cells on a finite volume grid, then by solving the Riemann problem with $q_l = Q_{i-1}$ and $q_r = Q_i$, we can obtain information that can be used to compute a numerical flux and update the cell averages over a time step. For hyperbolic problems the solution to the Riemann problem is typically a similarity solution, a function of x/t alone, and consists of a finite set of waves that propagate away from the origin with constant wave speeds. For linear hyperbolic systems the Riemann problem is easily solved in terms of the eigenvalues and eigenvectors of the matrix A , as

developed in Chapter 3. This simple structure also holds for nonlinear systems of equations and the exact solution (or arbitrarily good approximations) to the Riemann problem can be constructed even for nonlinear systems such as the Euler equations. The theory of nonlinear Riemann solutions for scalar problems is developed in Chapter 11 and extended to systems in Chapter 13.

Computationally, the exact Riemann solution is often too expensive to compute for nonlinear problems and *approximate Riemann solvers* are used in implementing numerical methods. These techniques are developed in Section 15.3.

1.2.2 Shock Capturing vs. Tracking

Since the PDEs continue to hold away from discontinuities, one possible approach is to combine a standard finite difference or finite volume method in smooth regions with some explicit procedure for tracking the location of discontinuities. This is the numerical analogue of the mathematical approach in which the PDEs are supplemented by jump conditions across discontinuities. This approach is often called *shock tracking* or front tracking. In more than one space dimension, discontinuities typically lie along curves (in two dimensions) or surfaces (in three dimensions), and such algorithms typically become quite complicated. Moreover, in realistic problems there may be many such surfaces that interact in complicated ways as time evolves. This approach will not be discussed further in this book. For some examples and discussion, see [41], [66], [103], [153], [154], [171], [207], [289], [290], [321], [322], [371], [372].

Instead we concentrate here on *shock-capturing* methods, where the goal is to capture discontinuities in the solution automatically, without explicitly tracking them. Discontinuities must then be smeared over one or more grid cells. Success requires that the method implicitly incorporate the correct jump conditions, reduce smearing to a minimum, and not introduce nonphysical oscillations near the discontinuities. High-resolution finite volume methods based on Riemann solutions often perform well and are much simpler to implement than shock-tracking methods.

1.3 Multidimensional Problems

The Riemann problem is inherently one-dimensional, but is extensively used also in the solution of multidimensional hyperbolic problems. A two-dimensional finite volume grid typically consists of polygonal grid cells; quadrilaterals or triangles are most commonly used. A Riemann problem normal to each edge of the cell can be solved in order to determine the flux across that edge. In three dimensions each face of a finite volume cell can be approximated by a plane, and a Riemann problem normal to this plane solved in order to compute the flux. Multidimensional problems are discussed in the Part III of the book, starting with an introduction to the mathematical theory in Chapter 18.

If the finite volume grid is rectangular, or at least logically rectangular, then the simplest way to extend one-dimensional high-resolution methods to more dimensions is to use *dimensional splitting*, a fractional-step approach in which one-dimensional problems along each coordinate direction are solved in turn. This approach, which is often surprisingly effective in practice, is discussed in Section 19.5. In some cases a more fully multidimensional

method is required, and one approach is developed starting in Chapter 20, which again relies heavily on our ability to solve one-dimensional Riemann problems.

1.4 Linear Waves and Discontinuous Media

High-resolution methods were originally developed for nonlinear problems in order to accurately capture discontinuous solutions such as shock waves. Linear hyperbolic equations often arise from studying small-amplitude waves, where the physical nonlinearities of the true equations can be safely ignored. Such waves are often smooth, since shock waves can only appear from nonlinear phenomena. The acoustic waves we are most familiar with arise from oscillations of materials at the molecular level and are typically well approximated by linear combinations of sinusoidal waves at various frequencies. Similarly, most familiar electromagnetic waves, such as visible light, are governed by the linear Maxwell equations (another hyperbolic system) and again consist of smooth sinusoidal oscillations.

For many problems in acoustics or optics the primary computational difficulty arises from the fact that the domain of interest is many orders of magnitude larger than the wavelengths of interest, and so it is important to use a method that can resolve smooth solutions with a very high order of accuracy in order to keep the number of grid points required manageable. For problems of this type, the methods developed in this book may not be appropriate. These finite volume high-resolution methods are typically at best second-order accurate, resulting in the need for many points per wavelength for good accuracy. Moreover they have a high cost per grid cell relative to simpler finite difference methods, because of the need to solve Riemann problems for each pair of grid cells every time step. The combination can be disastrous if we need to compute over a domain that spans thousands of wavelengths. Instead methods with a higher order of accuracy are typically used, e.g., fourth-order finite difference methods or spectral methods. For some problems it is hopeless to try to resolve individual wavelengths, and instead *ray-tracing* methods such as geometrical optics are used to determine how rays travel without discretizing the hyperbolic equations directly.

However, there are some situations in which high-resolution methods based on Riemann solutions may have distinct advantages even for linear problems. In many applications wave-propagation problems must be solved in materials that are not homogeneous and isotropic. The heterogeneity may be smoothly varying (e.g., acoustics in the ocean, where the sound speed varies with density, which may vary smoothly with changes in salinity, for example). In this case high-order methods may still be applicable. In many cases, however, there are sharp interfaces between different materials. If we wish to solve for acoustic or seismic waves in the earth, for example, the material parameters typically have jump discontinuities where soil meets rock or at the boundaries between different types of rock. Ultrasound waves in the human body also pass through many interfaces, between different organs or tissue and bone. Even in ocean acoustics there may be distinct layers of water with different salinity, and hence jump discontinuities in the sound speed, as well as the interface at the ocean floor where waves pass between water and earth. With wave-tracing methods it may be possible to use reflection and transmission coefficients and Snell's law to trace rays and reflected rays at interfaces, but for problems with many interfaces this can be unwieldy. If we wish to model the wave motion directly by solving the hyperbolic equations, many high-order methods can have difficulties near interfaces, where the solution is typically not smooth.

For these problems, high-resolution finite volume methods based on solving Riemann problems can be an attractive alternative. Finite volume methods are a natural choice for heterogeneous media, since each grid cell can be assigned different material properties via an appropriate averaging of the material parameters over the volume enclosed by the cell. The idea of a Riemann problem is easily extended to the case where there is a discontinuity in the medium at $x = 0$ as well as a discontinuity in the initial data. Solving the Riemann problem at the interface between two cells then gives a decomposition of the data into waves moving into each cell, including the effects of reflection and transmission as waves move between different materials. Indeed, the classical reflection and transmission coefficients for various problems are easily derived and understood in terms of particular Riemann solutions. Variable-coefficient linear problems are discussed in Chapter 9 and Section 21.5.

Hyperbolic equations with variable coefficients may not be in conservation form, and so the methods are developed here in a form that applies more generally. These *wave-propagation methods* are based directly on the waves arising from the solution of the Riemann problem rather than on numerical fluxes at cell interfaces. When applied to conservation laws, there is a natural connection between these methods and more standard flux-differencing methods, which will be elucidated as we go along. But many of the shock-capturing ideas that have been developed in the context of conservation laws are valuable more broadly, and one of my goals in writing this book is to present these methods in a more general framework than is available elsewhere, and with more attention to applications where they have not traditionally been applied in the past.

This book is organized in such a way that all of the ideas required to apply the methods on linear problems are introduced first, before discussing the more complicated nonlinear theory. Readers whose primary interest is in linear waves should be able to skip the nonlinear parts entirely by first studying Chapters 2 through 9 (on linear problems in one dimension) and then the preliminary parts of Chapters 18 through 23 (on multidimensional problems).

For readers whose primary interest is in nonlinear problems, I believe that this organization is still sensible, since many of the fundamental ideas (both mathematical and algorithmic) arise already with linear problems and are most easily understood in this context. Additional issues arise in the nonlinear case, but these are most easily understood if one already has a firm foundation in the linear theory.

1.5 CLAWPACK Software

The CLAWPACK software (“conservation-laws package”) implements the various wave-propagation methods discussed in this book (in Fortran). This software was originally developed as a teaching tool and is intended to be used in conjunction with this book. The use of this software is briefly described in Chapter 5, and additional documentation is available online, from the webpage

<http://www.amath.washington.edu/~claw>

Virtually all of the computational examples presented in the book were created using CLAWPACK, and the source code used is generally available via the website

<http://www.amath.washington.edu/~claw/book.html>

1.6 References

9

A parenthetical remark in the text or figure captions of the form

[claw/book/chapN/examplename]

is an indication that accompanying material is available at

<http://www.amath.washington.edu/~claw/book/chapN/examplename/www>

often including an animation of time-dependent solutions. From this webpage it is generally possible to download a CLAWPACK directory of the source code for the example. Downloading the tarfile and unpacking it in your claw directory results in a subdirectory called claw/book/chapN/examplename. (You must first obtain the basic CLAWPACK routines as described in Chapter 5.)

You are encouraged to use this software actively, both to develop an intuition for the behavior of solutions to hyperbolic equations and also to develop direct experience with these numerical methods. It should be easy to modify the examples to experiment with different parameters or initial conditions, or with the use of different methods on the same problem.

These examples can also serve as templates for developing codes for other problems. In addition, many problems not discussed in this book have already been solved using CLAWPACK and are often available online. Some pointers can be found on the webpages for the book, and others are collected within the CLAWPACK software in the applications subdirectory; see

<http://www.amath.washington.edu/~claw/apps.html>

1.6 References

Some references for particular applications and methods are given in the text. There are thousands of papers on these topics, and I have not attempted to give an exhaustive survey of the literature by any means. The references cited have been chosen because they are particularly relevant to the discussion here or provide a good entrance point to the broader literature. Listed below are a few books that may be of general interest in understanding this material, again only a small subset of those available.

An earlier version of this book appeared as a set of lecture notes [281]. This contains a different presentation of some of the same material and may still be of interest. My contribution to [287] also has some overlap with this book, but is directed specifically towards astrophysical flows and also contains some description of hyperbolic problems arising in magnetohydrodynamics and relativistic flow, which are not discussed here.

The basic theory of hyperbolic equations can be found in many texts, for example John [229], Kevorkian [234]. The basic theory of nonlinear conservation laws is neatly presented in the monograph of Lax [263]. Introductions to this material can also be found in many other books, such as Liu [311], Whitham [486], or Chorin & Marsden [68]. The book of Courant & Friedrichs [92] deals almost entirely with gas dynamics and the Euler equations, but includes much of the general theory of conservation laws in this context and is very useful. The books by Bressan [46], Dafermos [98], Majda [319], Serre [402], Smoller [420], and Zhang & Hsiao [499] present many more details on the mathematical theory of nonlinear conservation laws.

For general background on numerical methods for PDEs, the books of Iserles [211], Morton & Mayers [333], Strikwerda [427], or Tveito & Winther [461] are recommended. The book of Gustafsson, Kreiss & Oliger [174] is aimed particularly at hyperbolic problems and contains more advanced material on well-posedness and stability of both initial- and initial-boundary-value problems. The classic book of Richtmyer & Morton [369] contains a good description of many of the mathematical techniques used to study numerical methods, particularly for linear equations. It also includes a large section on methods for nonlinear applications including fluid dynamics, but is out of date by now and does not discuss many of the methods we will study.

A number of books have appeared recently on numerical methods for conservation laws that cover some of the same techniques discussed here, e.g., Godlewski & Raviart [156], Kröner [245], and Toro [450]. Several other books on computational fluid dynamics are also useful supplements, including Durran [117], Fletcher [137], Hirsch [198], Laney [256], Oran & Boris [348], Peyret & Taylor [359], and Tannehill, Anderson & Pletcher [445]. These books discuss the fluid dynamics in more detail, generally with emphasis on specific applications.

For an excellent collection of photographs illustrating a wide variety of interesting fluid dynamics, including shock waves, Van Dyke's *Album of Fluid Motion* [463] is highly recommended.

Many more references on these topics can easily be found these days by searching on the web. In addition to using standard web search engines, there are preprint servers that contain collections of preprints on various topics. In the field of conservation laws, the *Norwegian preprint server* at

<http://www.math.ntnu.no/conservation/>

is of particular note. Online citation indices and bibliographic databases are extremely useful in searching the literature, and students should be encouraged to learn to use them. Some useful links can be found on the webpage [c1aw/book/chap1/].

1.7 Notation

Some nonstandard notation is used in this book that may require explanation. In general I use q to denote the solution to the partial differential equation under study. In the literature the symbol u is commonly used, so that a general one-dimensional conservation law has the form $u_t + f(u)_x = 0$, for example. However, most of the specific problems we will study involve a velocity (as in the acoustics equations (1.5)), and it is very convenient to use u for this quantity (or as the x -component of the velocity vector $\vec{u} = (u, v)$ in two dimensions).

The symbol Q_i^n (in one dimension) or Q_{ij}^n (in two dimensions) is used to denote the numerical approximation to the solution q . Subscripts on Q denote spatial locations (e.g., the i th grid cell), and superscript n denotes time level t_n . Often the temporal index is suppressed, since we primarily consider one-step methods where the solution at time t_{n+1} is determined entirely by data at time t_n . When Q or other numerical quantities lack a temporal superscript it is generally clear that the current time level t_n is intended.