


ASTR 670 • Hydrodynamics

by Benedikt Diemer (University of Maryland)

Updated May 1, 2025



Contents

0	About these notes	4
0.1	Acknowledgments and commonly used references	4
0.2	Notation	4
1	Hydrodynamics: from particles to fluids	7
1.1	What is a fluid?	7
1.2	Astrophysical fluids	8
1.3	Averaging over particles	8
1.4	Basic fluid quantities	9
1.5	Towards the fluid equations: a schematic outline	11
1.6	Temperature and the equation of state	12
2	The Euler equations	15
2.1	Eulerian and Lagrangian derivatives	15
2.2	The continuity (density) equation	16
2.3	The velocity equation	17
2.4	The internal energy equation	18
2.5	Summary of the Euler equations	19
3	Equilibrium and steady flows	21
3.1	Hydrostatic equilibrium	21
3.2	Adiabatic flows	22
3.3	Barotropic flows and Bernoulli's principle	23
3.4	The limits of inviscid fluid dynamics: d'Alembert's paradox	26
4	Sound waves	28
4.1	The linearized Euler equations	28
4.2	The properties of sound waves	29
5	Computational hydro I: Theoretical background	31
5.1	The self-similarity of hydrodynamics	31
5.2	The conservation-law form of the equations	33
5.3	The advection equation	34
5.4	A very basic taxonomy of partial differential equations	35
5.5	Computing in space and time: finite difference vs. finite volume	36
5.6	The Courant-Friedrichs-Lewy condition	37
5.7	Computational error terms	37
6	Computational hydro II: Finite-differencing schemes	38
6.1	First-order schemes for the advection equation	38
6.2	Stability analysis	40
6.3	Understanding numerical diffusion	42
6.4	Higher-order finite-difference schemes	43
7	Shocks	45
7.1	Why shocks arise: Burgers' equation	45
7.2	What is a shock?	46
7.3	The Rankine-Hugoniot conditions	48
7.4	Ideal gas shocks	49

7.5	Isothermal shocks	51
7.6	Supernova blast waves	52
8	Computational hydro III: Finite-volume schemes	56
8.1	Godunov schemes	56
8.2	Riemann solvers	58
8.3	The shocktube problem	62
8.4	Higher-order schemes: Reconstruction and slope limiters	63
8.5	Higher-order schemes: Time integration	66
8.6	Multiple dimensions	68
8.7	Popular hydrodynamics codes in astrophysics	68
9	Fluid instabilities	70
9.1	The Jeans instability	70
9.2	Perturbations of a general two-fluid interface	72
9.3	Surface gravity waves	73
9.4	The Rayleigh-Taylor instability	74
9.5	The Kelvin-Helmholtz instability	74
10	Magnetohydrodynamics	76
10.1	The interaction between fluids and electromagnetism	76
10.2	The two-fluid approach	77
10.3	Eliminating electric fields and currents	78
10.4	The equations of ideal MHD	80
10.5	Basic MHD dynamics: flux freezing and dynamos	81
10.6	MHD waves	83
10.7	Observing magnetic fields via the rotation measure	83
11	Turbulence	85
11.1	Self-similarity and the turbulent cascade	85
11.2	The Kolmogorov-Obukhov law	86
A	Mathematical Background	89
A.1	Vector operators	89
A.2	Index notation and tensors	89
A.3	Vector identities	90
A.4	Eigenvalues and eigenvectors	91
B	Derivations	92
B.1	Conservation-law form of the Euler equations	92
B.2	The matrix form of the fluid equations and the eigenvalue perspective	93
B.3	Dispersion relation of perturbations at a two-fluid interface	95
C	ULULA: a lightweight hydro code in Python	98
	References	99

0 About these notes

These notes serve as a guide to ASTR 670, a graduate class in fluid dynamics and the interstellar medium (ISM) at the University of Maryland. Since hydrodynamics accounts for only half the course (about 12 lectures), we have to omit numerous important topics and interesting details. The goal of the class is to develop an intuitive, physical understanding for fluid dynamics rather than to give an exhaustive mathematical treatment. In particular, we will skip lengthy derivations and refer to textbooks instead. In the spirit of learning how to tackle the kinds of hydrodynamics problems that arise in the practice of astrophysics, a significant portion of the course is devoted to modern computational methods. The numerical chapters are deliberately interspersed with physics-focused ones, but the topics could also be taught in a different order.

0.1 Acknowledgments and commonly used references

My treatment of hydrodynamics is, of course, based on the much more substantial works of many colleagues. The most commonly cited text is the excellent book by Clarke & Carswell (2014, hereafter [CC](#)), while the classic text of Shu (1992, hereafter [Shu](#)) provides a comprehensive mathematical background. I also extensively refer to the outstanding lecture notes by Frank van den Bosch (2020, hereafter [vdB](#); you can find the notes on his [website](#)). For the computational aspects of the course, I have built on the innovative computational hydrodynamics book by Michael Zingale (2021, hereafter [Zingale](#); the book is publicly available on [github](#)). Toro (2009, hereafter [Toro](#)) provides a comprehensive reference for the mathematical underpinnings of the numerical methods discussed.

Finally, I am grateful to my colleague Alberto Bolatto who taught this course before me, and whose notes provided a basis for my version of the class. I owe a great debt to the late Alexei Khokhlov and to Andrey Kravtsov, whose graduate courses at the University of Chicago introduced me to hydrodynamics and numerical astrophysics.

The image on the front page shows a simulation of the Kelvin-Helmholtz instability created with the moving-mesh code AREPO (Springel 2010), using initial conditions and a plotting script provided by Philip Mocz.

0.2 Notation

It is virtually impossible to find two texts on hydrodynamics that agree on the same set of symbols. Is pressure denoted p or P ? Is m the mass of a particle or the mass in a fluid element? Is the fluid velocity denoted \mathbf{v} or \mathbf{u} ? There are no correct answers, of course, but we will use the symbols listed in Tables 1 and 2 consistently. Throughout the notes, bold symbols denote 3-vectors, e.g., $\mathbf{u} = (u_x, u_y, u_z)$; we use row and column notation interchangeably. We will try to use vector notation throughout, but sometimes the index summation notation is clearer (see §A.2).

Sym.	Units	Meaning	§
α_{cfl}	—	CFL number	5.6
\mathbf{B}	$\text{G} = \text{g}^{1/2} \text{cm}^{-1/2} \text{s}^{-1}$	Magnetic field (the name of the unit is Gauss)	10.1
c_s	cm/s	Sound speed	4.2
E	erg/cm ³	Total energy per unit volume in a fluid element	1.4
\mathbf{E}	$\text{g}^{1/2} \text{cm}^{-1/2} \text{s}^{-1}$	Electric field (the units are statV/cm)	10.1
ε	erg/g = cm ² /s ²	Thermal (“internal”) energy per unit mass in a fluid el.	1.4
η	g/cm/s	Dynamical coefficient of viscosity	11.1
f	s ⁻¹	Frequency	4.2
\mathcal{F}	(multiple)	Vector of flux terms	5.2
Φ	erg/g = cm ² /s ²	Gravitational potential	1.4
γ	—	Ideal gas equation of state parameter	1.6
Γ	erg/cm ³ /s	Heating rate per unit volume	2.4
H	erg/g = cm ² /s ²	Bernoulli constant	3.3
\mathbf{I}	—	Identity matrix (usually 3 × 3)	—
\mathbf{J}	$\text{g}^{1/2} \text{cm}^{-1/2} \text{s}^{-2}$	Electric current density (the units are statA/cm ²)	10.1
\mathbf{k}	cm ⁻¹	Wave vector (with amplitude $k = 2\pi/\lambda$)	4.2
K	$\text{g}^{1-\gamma} \text{cm}^{3\gamma-1} \text{s}^{-2}$	Constant in barotropic equation of state	3.2
κ	cm ² /s	Diffusivity	5.4
λ	cm	Wavelength	4.2
λ_k	(depends)	Eigenvalues of a matrix	A.4
Λ	erg/cm ³ /s	Cooling rate per unit volume	2.4
λ_{mfp}	cm	Mean free path of particles	1.3
m	g	Mass in a fluid element	1.4
m_{ptl}	g	Mass of individual particles	1.4
\mathcal{M}	—	Mach number	7.2
μ	—	Mean mass of particles in units of proton mass m_p	1.4
n	cm ⁻³	Number density	1.3
N	—	Number of some species in a fluid element	1.3
ω	s ⁻¹	Angular wave frequency ($\omega = 2\pi f$)	4.2
$\boldsymbol{\omega}$	s ⁻¹	Vorticity, the curl of the bulk velocity	3.3
p	g cm/s	Momentum of particles	1.3
P	dyne/cm ² = erg/cm ³	Pressure in a fluid element	1.4
Π	erg/cm ³	Momentum flux density tensor	B.1
q	$\text{g}^{1/2} \text{cm}^{-3/2} \text{s}^{-1}$	Net charge density of fluid (the units are statC/cm ³)	10.1
\mathbf{q}	(depends)	Position in Lagrangian space	2.1
Q	(unknown)	Unspecified scalar quantity	—
\mathbf{Q}	(unknown)	Unspecified vector quantity	—
r	cm	Radial coordinate in Eulerian space	—
Re	—	Reynolds number (relative importance of viscosity)	11.1
ρ	g/cm ³	Mass density	1.4
S	erg/K/g	Entropy per unit mass in a fluid element	3.2

Table 1: Continued on next page.

Sym.	Units	Meaning	§
\mathbf{S}	(multiple)	Vector of source terms	5.2
σ	cm ²	Cross section of particles	1.3
σ_e	s ⁻¹	Electrical conductivity	10.3
t	s	Time coordinate	—
T	K	Temperature in a fluid element	1.4
\mathbf{u}	cm/s	Bulk (average) velocity of particles	1.4
\mathbf{U}	(multiple)	Vector of conserved fluid quantities	5.2
\mathbf{v}	cm/s	Total velocity vector of particles, $\mathbf{v} = \mathbf{w} + \mathbf{u}$	1.4
V	cm ³	Volume of a fluid element (or other volumes)	1.3
\mathbf{V}	(multiple)	Vector of primitive fluid quantities	5.2
\mathbf{w}	cm/s	Random velocity vector of particles	1.4
\mathbf{x}	cm	Coordinate position vector in Eulerian space	—

Table 1: Definition of symbols used throughout the text. The section given in the right column typically refers to the section where a quantity is first defined. We use the CGS unit system throughout.

Sym.	Value	Units	Meaning
G	6.67×10^{-8}	cm ³ /g/s ²	Gravitational constant
m_p	1.67×10^{-24}	g	Proton mass
m_e	9.11×10^{-28}	g	Electron mass
N_A	6.02×10^{23}	mol ⁻¹	Avogadro constant (atomic mass units per mole)
k_B	1.38×10^{-16}	erg/K	Boltzmann constant
\mathcal{R}	8.31×10^7	erg/K/mol	Gas constant ($\mathcal{R} = N_A k_B$)
c	3.00×10^{10}	cm/s	Speed of light
q_e	4.80×10^{-10}	g ^{1/2} cm ^{3/2} s ⁻¹	Electron charge (the units are statC)
g	9.81×10^2	cm/s ²	Gravitational acceleration on Earth
P_{air}	1.01×10^6	dyne/cm ²	Atmospheric air pressure on Earth
ρ_{air}	1.23×10^{-3}	g/cm ³	Density of air at standard pressure and temperature
μ_{air}	29.0	—	Mean particle weight in air in proton masses
AU	1.50×10^{13}	cm	Astronomical unit
pc	3.09×10^{18}	cm	Parsec
kpc	3.09×10^{21}	cm	Kiloparsec
Mpc	3.09×10^{24}	cm	Megaparsec
yr	3.16×10^7	s	Year in seconds
M_{\odot}	1.99×10^{33}	g	Solar mass
M_{\oplus}	5.97×10^{27}	g	Earth mass
R_{\odot}	6.96×10^{10}	cm	Solar radius
R_{\oplus}	6.38×10^8	cm	Earth radius

Table 2: Physical and astronomical constants and unit conversions in CGS units as used throughout the text.

1 Hydrodynamics: from particles to fluids

Broadly speaking, hydrodynamics is the study of the bulk dynamics of fluids, as opposed to their microscopic particle properties. In this section, we define what we mean by the term “fluid” and consider situations where fluids are important in astrophysics.

1.1 What is a fluid?

Simply put, a fluid is a material that can flow, meaning that it deforms under stresses. The term “stress” means forces that act differently on different locations. While a solid reacts elastically to such forces and returns to its original form once they are removed, a fluid does not. However, the definition can be a little ambiguous since it depends on the timescale considered. For example, rocks deform over geological timescales. On a particle level, fluidity corresponds to an ability of particles to move more or less freely past each other. Fluids can be further classified by a number of properties (see [vdB §1](#) for more details):

- **Collisionality** means that the constituent particles of a fluid (frequently) collide with each other. For example, the molecules in air do, but the stars in a galaxy or dark matter particles do not. We might nevertheless consider them a fluid because the particles exchange energy via gravitational interaction and thus represent a statistical ensemble. For the purposes of this course, however, we will consider only collisional fluids. This distinction will become important when we relate pressure and density.
- **Compressibility** means that a fluid changes its density under pressure. This property typically distinguishes gases from liquids. While there are technically no truly incompressible liquids, substances such as water are so resilient to pressure that we can think of them as incompressible (at least under the conditions encountered on Earth). By contrast, all gases are compressible.
- **Viscosity** means that the collisions in a fluid smooth out velocity differentials. Motions in viscous fluids (such as honey) will decay faster than motions in fluids with low viscosity (such as water). An “inviscid” fluid has no viscosity at all.
- **Conductivity** is similar to viscosity but for heat, meaning that a highly conductive fluid smooths out temperature gradients. If a fluid is inviscid and does not conduct, we call it a “perfect” or “ideal” fluid. Even though such fluids technically do not exist, ideal gases are a common approximation in astrophysical systems.
- **Ionization** means that a significant fraction of the atoms in a fluid are ionized. A partially or fully ionized gas is called a *plasma* and presents a physically different system where we need to take electromagnetic forces into account. We will return to plasmas when we study magnetohydrodynamics (MHD) in §10. For the time being, we assume all fluids to be neutral.

On Earth, the most common fluids (e.g., air and water) are collisional, neutral, and exhibit a range of viscosities and compressibilities. Many problems of interest involve solids (e.g., the aerodynamics of an airplane). In astrophysics, there are almost no solid-fluid interfaces to worry about and viscosities tend to be very low. However, plasmas are relatively common, e.g., in the intracluster medium and parts of the interstellar medium.

While gases are compressible, their *flows* are often non-compressing in practice. On Earth, we typically use solids to compress gases (think airplane cabins, pistons in engines, balloons, and so

on). Without being contained, the gas moves out of the way rather than being compressed (e.g., when you exhale). Given that there are no solid containers in the Universe, astrophysical flows are often incompressible. However, gas is being compressed by gravity, e.g., in stars and galaxies.

1.2 Astrophysical fluids

Hydrodynamics acts almost everywhere in the Universe. Here is an incomplete list of hydrodynamical astrophysical objects and systems (cf. [vdB §1](#)):

- **Stars** and **gas planets** are balls of gas in hydrostatic equilibrium between their gravity and pressure.
- **Planet atmospheres** are also in hydrostatic equilibrium but around a solid planet.
- **White dwarfs** and **Neutron stars** can be described as fluids, albeit with unusual equations of state (§1.6) that take into account the enormous pressure due to the quantum degeneracy of the respective particles.
- **Supernovae** are hydrodynamical systems (either stars or white dwarfs) where the dynamics is influenced by neutrinos and nuclear burning.
- **Accretion disks** and **proto-planetary disks** arise because of the angular momentum of gas.
- The **interstellar medium (ISM)** is the gas between the stars in galaxies. Roughly speaking, it consists of hot, warm, and cold phases.
- The **circum-galactic medium (CGM)** is diffuse gas that resides outside galaxies but within their dark matter halos.
- The **intra-cluster medium (ICM)** is the hot plasma in galaxy clusters (outside of their galaxies).
- The **Intergalactic medium (IGM)** is the gas well outside galaxies that forms part of cosmic structure, which can be observed via the Lyman- α forest.
- The **photon-baryon fluid** in the early Universe is a hydrodynamical system, with radiation playing a major role.

Note that gravity plays an important role in all of these systems. There are also various collisionless “fluid” systems such as the stars in galaxies and the particles in dark matter halos. The Jeans equation that describes these systems is closely related to the equations of hydrodynamics, but we will focus on the latter in this course.

1.3 Averaging over particles

We often speak of quantities such as pressure or temperature without thinking about how they relate to the fundamental physics of matter. Before we write down the equations that govern the evolution of fluids, we need to establish the connection between the microscopic particles that make up fluids and the macroscopic, averaged description that we will focus on in this course.

In principle, a full description of the physics of fluids would depend on the motions of all their constituent particles. However, such a description is not only unattainable in practice, it is also of no interest since most of the particle motions are random. Instead, we are interested in the bulk motion of particles. In hydrodynamics, we thus average over particle motions within small, imaginary volumes in space.

How small or large can these imaginary “fluid elements” be? Let us envision a cubic element with side length δx and volume $V = \delta x^3$, although the actual shape of the element is unimportant. The first criterion arises from the requirement that particles must collide frequently, which is equivalent to the mean free path being much shorter than the size of our fluid element,

$$\delta x \gg \lambda_{\text{mfp}} = \frac{1}{n\sigma}, \quad (1.1)$$

where $n = N/V$ is the number density of particles and σ their cross-section. We can estimate the cross-section as πr^2 , where $r \approx 10^{-8}$ cm is the size of an atom. The number density of an ideal gas (such as air) at room temperature is about $n_{\text{air}} \approx 2 \times 10^{19}/\text{cm}^3$, so $\lambda_{\text{mfp}} \approx 2 \times 10^{-4}$ cm. This is a small distance, but not truly microscopic. Thus, the fluid approximation holds within the kinds of volumes we typically consider but not in arbitrarily small “fluid pixels.” A second criterion arises from our need to average over enough particles to avoid significant shot noise, implying

$$\delta x^3 n \gg 1. \quad (1.2)$$

To be a little more specific, we could require that our averages have a $1/\sqrt{N}$ error of less than 1%, which would mean averaging over at least 10,000 particles per element. For air, that translates to $\delta x \gtrsim 10^{-5}$ cm, which is a little less restrictive than the mean free path criterion. Finally, we will treat the quantities in our fluid element as locally constant, meaning that the element must be small enough so that any quantity Q does not vary significantly within it,

$$\delta x \ll \frac{Q}{|\nabla Q|}. \quad (1.3)$$

This criterion depends on the problem we are trying to solve as well as the properties of the fluid. In astrophysics, we typically need to imagine much larger fluid elements than in terrestrial applications. For example, for the warm neutral medium phase of the ISM, $n \approx 1/\text{cm}^3$, implying $\lambda_{\text{mfp}} \approx 10^{15}$ cm ≈ 200 AU. In such a region, we find $\approx 10^{46}$ particles, so we are averaging over large numbers. These examples demonstrate that we should not blindly expect the fluid approximation to be valid at any scale!

1.4 Basic fluid quantities

We are now ready to define averaged fluid quantities. The first is **density**, $\rho = m/V$, where m is the amount of mass in the fluid element and V its volume, $V = \delta x^3$. Alternatively, we can define the **number density** $n = N/V$, where N is the number of some species (e.g., hydrogen atoms) in the fluid element.

Second, we define the **fluid velocity**, \mathbf{u} , a vector that has three components u_x , u_y , and u_z . This velocity represents the average over the total particle velocities \mathbf{v} , such that $\mathbf{v} = \mathbf{u} + \mathbf{w}$, where \mathbf{w} are the random velocities of particles.

Third, **pressure** measures the mechanical force exerted by the particles in a fluid. We can understand this effect by considering a box as shown in Figure 1. A pressure is, by definition, the force on an area; for now, we consider the left side of the box in the x -direction. Imagine particles with mass m_{ptl} bouncing back and forth in the box due to their random velocities (we assume no bulk velocity for now). A particle will hit the left wall once per $\Delta t = 2\delta x/w_x$, where δx is the size of the box (or fluid element). The particle imparts a momentum of $\Delta p = 2m_{\text{ptl}}w_x$ per collision because it reverses direction. The same happens for all N particles in the box, which gives a total force on the left wall of

$$F_{\text{left}} = N \frac{\Delta p}{\Delta t} = N \frac{m_{\text{ptl}} \langle w_x^2 \rangle}{\delta x}, \quad (1.4)$$




Figure 1: Pressure can be understood as the “banging” of particles against the walls of a container. During each reflection, a particle imparts twice its momentum in the direction of the wall. Image from Wikipedia.

where $\langle w_x^2 \rangle$ is the average square-velocity of the particles in the x -direction. We can relate this velocity to the overall velocity because no direction is special, and thus

$$\langle |\mathbf{w}|^2 \rangle = \langle w_x^2 \rangle + \langle w_y^2 \rangle + \langle w_z^2 \rangle = 3\langle w_x^2 \rangle. \quad (1.5)$$

The isotropy of the problem also implies that all walls feel the same force. We plug $\langle w_x^2 \rangle = \langle |\mathbf{w}|^2 \rangle / 3$ into Equation 1.4 and write down the definition of pressure,

$$P \equiv \frac{\text{force}}{\text{area}} = \frac{F}{\delta x^2} = \frac{Nm_{\text{ptl}}\langle |\mathbf{w}|^2 \rangle}{3\delta x^3} = \frac{m\langle |\mathbf{w}|^2 \rangle}{3V} \quad (1.6)$$

which simplifies to

$$P = \frac{1}{3}\rho\langle |\mathbf{w}|^2 \rangle \quad (1.7)$$

This formula makes it clear that pressure has units of an energy density: the more kinetic particle energy there is contained in a certain volume, the more force those particles exert. Pressure is thus intimately related to the thermal energy per unit mass of fluid,

$$\varepsilon_{\text{th}} \equiv \frac{1}{2}\langle |\mathbf{w}|^2 \rangle = \frac{3}{2}\frac{P}{\rho}. \quad (1.8)$$

However, we will typically use the more general total **internal energy per unit mass**, ε . To understand the distinction between ε_{th} and ε , we need to first take a detour into statistical mechanics (§1.5) and thermodynamics (§1.6). For now, we define the **total energy per unit volume** (not per unit mass!) as the sum of the kinetic energy, the internal energy, and external potentials,

$$E \equiv \rho \left(\frac{|\mathbf{u}|^2}{2} + \varepsilon + \Phi \right) \quad (1.9)$$

where we have specifically included a **gravitational potential** Φ because gravity is by far the most relevant external force on astrophysical fluids. Note that E and P now have the same units of energy per volume.

One final definition for this chapter is to write the particle mass in terms of the mean particle weight in proton masses such that

$$m_{\text{ptl}} \equiv \mu m_{\text{p}}. \quad (1.10)$$

When considering “pristine” gas with the Big Bang composition of $\approx 76\%$ hydrogen and $\approx 24\%$ Helium, we can assume a mean particle mass of $1/(0.76 + 0.24/4) \approx 1.22$ (in units of m_{p}). However, this number depends on whether the gas is atomic or molecular, and whether it contains higher elements (“metals”).

1.5 Towards the fluid equations: a schematic outline

Our goal is to obtain differential equations that tell us how the averaged fluid quantities evolve in space and time. In principle, these *fluid equations* can be derived via statistical mechanics from the full 6D phase-space distribution of particles, $f(\mathbf{x}, \mathbf{v}, t)$. This function tells us how many particles there at time t at a given coordinate and with a given velocity. With lots of math (see vdB §6), it can be shown that the evolution of the distribution function is described by the **Boltzmann equation**,

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} + \frac{\partial \mathbf{v}}{\partial t} \cdot \frac{\partial f}{\partial \mathbf{v}} = \left(\frac{df}{dt} \right)_{\text{collisions}}, \quad (1.11)$$

where the right-hand side contains all changes in f due to the collisions of particles. We are generally interested in the fluid quantities as a function of space and time, but we do not care about the microscopic particle velocities. Thus, we can obtain useful relations by integrating over velocity. For example, we can compute the fluid momentum by integrating over the total momentum of the individual particles in the distribution function,

$$[\rho \mathbf{u}](\mathbf{x}, t) = \int m_{\text{ptl}} \mathbf{v} \times f(\mathbf{x}, \mathbf{v}, t) d^3v. \quad (1.12)$$

This is called a “moment” of the Boltzmann equation. We can compute similar moments for the particle mass and kinetic energy, which lead to the conserved quantities of density and internal energy (Shu §2). The key insight is that if a quantity is conserved in individual particle collisions (as are particle mass, momentum, and kinetic energy), collisions also cannot change the conserved moments over time. Thus, we can derive a lowest-order set of fluid equations by dropping the collision term.

However, we first need some closed form of the velocity distribution to integrate over. It turns out that such a distribution can be derived in the special case of an **ideal gas**, meaning an ensemble of particles whose isotropic random motions are determined only by its **temperature**, T . Physically, this case corresponds to a gas where the particles collide often enough to establish equilibrium and fulfill the conditions for the fluid approximation from §1.3, but not so often that they transfer momentum or heat between fluid elements, which would lead to viscosity and conductivity. While this scenario may sound contrived, ideal gases are usually a very good approximation in astrophysics. Most gas has extremely low density compared to terrestrial applications, and thus very low viscosity. For an ideal gas, it turns out that the random particle velocities follow the **Maxwell-Boltzmann distribution**,

$$f_{\text{MB,3D}}(\mathbf{w}) = n \left(\frac{m_{\text{ptl}}}{2\pi k_{\text{B}} T} \right)^{3/2} \exp \left(-\frac{m_{\text{ptl}} |\mathbf{w}|^2}{2k_{\text{B}} T} \right), \quad (1.13)$$

where k_{B} is the Boltzmann constant. This distribution seems to make the counter-intuitive prediction that the most likely state is a velocity of $\mathbf{w} = 0$. However, the probability is per velocity

vector, and there are way fewer possible short than long vectors. Since $f_{\text{MB},3\text{D}}$ is isotropic, we can reduce the triple integral over velocity space to a single integral,

$$\int f_{\text{MB},3\text{D}}(\mathbf{w})d^3w = \int 4\pi w^2 f_{\text{MB},3\text{D}}(\mathbf{w})dw = \int f_{\text{MB},1\text{D}}(w)dw \quad (1.14)$$

where

$$f_{\text{MB},1\text{D}}(w) = n \left(\frac{2}{\pi}\right)^{1/2} \left(\frac{m_{\text{ptl}}}{k_{\text{B}}T}\right)^{3/2} w^2 \exp\left(-\frac{m_{\text{ptl}}w^2}{2k_{\text{B}}T}\right). \quad (1.15)$$

This distribution asymptotes to zero at zero velocity, which makes much more sense. Since it is independent of the fluid's bulk velocity \mathbf{u} , the integrals over \mathbf{v} that appear in the moment equations cleanly split into bulk and random velocity, allowing us to obtain a set of differential equations for ρ , \mathbf{u} , and ε (Shu §2). At the lowest, effectively collision-free order, these equations are called the *Euler equations*, which we study in detail in §2. Physically, the fact that the Maxwell-Boltzmann distribution is isotropic tells us that the random particle velocities contribute equally in all directions, regardless of the fluid's bulk velocity. This assumption is broken if we allow for significant collision terms, which lead to viscosity and conductivity. The corresponding, more complex equations are called the *Navier-Stokes equations*. In this course, we skip this more complicated case in the interest of brevity and because viscosity is negligible in many (though by no means all) astrophysical problems.

This section has presented a highly simplistic and incomplete summary of the procedure for deriving the fluid equations from particle dynamics. The full derivation is beyond the scope of this course but can be found in §2 and §3 of Shu or §6 and §7 of vdB.

1.6 Temperature and the equation of state

In the previous section we introduced the concept of temperature, which is perhaps a more intuitive way to think about the thermal energy of a fluid. We can connect temperature to the thermal energy as defined in §1.4 by integrating the Maxwell-Boltzmann distribution to find the average square random velocity in an ideal gas. The result is

$$\langle |\mathbf{w}|^2 \rangle = \frac{3k_{\text{B}}T}{m_{\text{ptl}}}. \quad (1.16)$$

We rewrite this expression to get the average thermal (or kinetic) energy per particle in an ideal gas,

$$\varepsilon_{\text{th}}m_{\text{ptl}} = \frac{1}{2}m_{\text{ptl}}\langle |\mathbf{w}|^2 \rangle = \frac{3}{2}k_{\text{B}}T, \quad (1.17)$$

or, rearranged,

$$\varepsilon_{\text{th}} = \frac{3}{2}\frac{k_{\text{B}}T}{m_{\text{ptl}}} \implies T = \frac{2\varepsilon_{\text{th}}m_{\text{ptl}}}{3k_{\text{B}}}. \quad (1.18)$$

These expressions are, in some sense, as close to a definition of temperature as we will get. A non-ideal fluid would feature a different velocity distribution, and our intuition for temperature would break down. For example, there is no guarantee that our skin would perceive “how hot” a fluid is in the same way.

In §1.4, we already saw that pressure and thermal energy are intimately related. Our main goal for this section is to further define the relation between pressure and temperature in a so-called **equation of state** (EoS). An EoS encodes the microphysical properties of the particles in a given fluid, which takes us into the domain of **thermodynamics**. In a nutshell, thermodynamics

deals with the physics inside a fluid element, whereas hydrodynamics deals with the motion and evolution of fluid systems as a whole. For a more formal derivation of the results in this section, see CC §4.2 or any textbook on thermodynamics. We begin with the pressure-temperature relation, $P = P(\rho, T)$.¹ We combine the definition of pressure (Equation 1.7) with Equation 1.16 to find the equation of state for an ideal gas,

$$P = nk_{\text{B}}T = \frac{\rho k_{\text{B}}T}{\mu m_{\text{p}}} \quad (1.19)$$

While this equation connects pressure to temperature, we have yet to make a connection to the total internal energy ε and thus to the total energy E . The distinction between ε_{th} and ε arises because fluids can have non-thermal ways to store internal energy, known as **degrees of freedom**. Statistical mechanics tells us that the internal energy is in **equipartition**, meaning that each degree of freedom stores the same amount of energy. Let us call the number of degrees of freedom N_{dof} . The kinetic, thermal energy provides $N_{\text{dof}} = 3$, one for each dimension. In diatomic molecular gases, the dumbbell-shaped molecules can also rotate around two axes, leading to a total of $N_{\text{dof}} = 5$, and possibly vibrate, leading to $N_{\text{dof}} = 7$. Similarly, chemical potentials can add to the internal energy of a fluid. To make it even more complicated, whether certain degrees of freedom are active can depend on the temperature! Some degrees may be quantum-mechanically “frozen out” at low temperatures, reducing N_{dof} . Regardless of these complexities, we assume that we can somehow count the degrees of freedom and define the dimensionless **adiabatic index** of an ideal gas,

$$\gamma \equiv \frac{N_{\text{dof}} + 2}{N_{\text{dof}}}, \quad (1.20)$$

which reduces to $\gamma = 5/3$ for a monoatomic gas and $\gamma = 7/5$ for a diatomic gas with rotational degrees of freedom (e.g., air). Given that the thermal energy corresponds to three degrees of freedom, we can write the total internal energy as

$$\varepsilon = \frac{N_{\text{dof}}}{3} \varepsilon_{\text{th}} = \frac{2}{3(\gamma - 1)} \varepsilon_{\text{th}} = \frac{1}{\gamma - 1} \frac{k_{\text{B}}T}{\mu m_{\text{p}}} = \frac{P}{\rho(\gamma - 1)}. \quad (1.21)$$

We have succeeded in connecting the total internal energy to temperature and pressure, and we rearrange to get the desired ideal gas equation of state written as $P = P(\rho, \varepsilon)$,

$$P = \rho\varepsilon(\gamma - 1) \quad (1.22)$$

In the $\gamma = 5/3$ case where there is only thermal internal energy, we recover the relation $\varepsilon = \varepsilon_{\text{th}} = 3P/(2\rho)$ from Equation 1.8. The vast majority of astrophysical gases are atomic (e.g., H I, H II, and He) or molecular and very cold (e.g., H₂ in molecular clouds). By default, we will thus assume an ideal gas EoS with $\gamma = 5/3$ unless otherwise stated.

There are, however, some noteworthy exceptions. For example, the interior of gas planets is thought to obey a poorly understood, non-ideal EoS due to the higher densities and pressures. In white dwarfs, the Fermion degeneracy from electrons leads to a strongly increased pressure. The EoS of neutron stars is a highly active research topic. The ideal gas law is valid for relativistic

¹The ideal gas EoS is sometimes given in units of the gas constant $\mathcal{R} = N_{\text{A}}k_{\text{B}} \approx k_{\text{B}}/m_{\text{p}}$. The latter equality is approximate because N_{A} is defined as the number of atomic mass units (amu) per gram. The amu is defined as 1/12 of a Carbon atom rather than the proton mass; the two differ by about 0.7%. In these notes, we define μ as the mean particle mass in units of m_{p} , which is just as self-consistent as using \mathcal{R} instead of k_{B} and defining μ as the mean weight in amu. In astrophysics, the k_{B} notation seems to be somewhat more common.

gases, but the relation between kinetic energy and momentum changes from $p^2/2m_{\text{ptl}}$ to pc , which changes the P - ε relation from $P = (2/3)\rho\varepsilon$ to $P = (1/3)\rho\varepsilon$, corresponding to $\gamma = 4/3$. In general, an EoS where pressure increases weakly with density is called “soft,” whereas one where pressure increases strongly is called “stiff.”

Finally, we note a few particularly simple cases. If the temperature is constant across a fluid system, we have an **isothermal** EoS, which leads to a constant ε and thus $P \propto \rho$. An isothermal EoS is an example of a **barotropic** EoS, where pressure is a function of only density but not temperature, $P = P(\rho)$.

2 The Euler equations

In this chapter, we write down the equations that determine how the fluid quantities defined in the previous chapter evolve in space and time. There are many levels of “difficulty” of these equations, depending on how general we allow the properties of our fluids to be (e.g., compressibility, conduction, and viscosity) and which additional physics we include (e.g., gravity, heating, cooling, and electromagnetism). We begin with a relatively simple set called the *Euler equations*, which were first presented in three papers by Leonhard Euler in 1757.²

The Euler equations describe a **compressible, inviscid, non-conductive, and electrically neutral fluid**. We will include the effects of gravity, heating, and cooling because they tend to be relevant in astrophysical applications. Instead of the full derivation procedure outlined in §1.5, we motivate the Euler equations with physical reasoning. In a nutshell, the equations arise from the conservation of mass, momentum, and energy in a small fluid volume δV .

2.1 Eulerian and Lagrangian derivatives

Before we consider specific fluid quantities, we must pause to contemplate what we mean when we say that a fluid quantity changes in space and time. There are two fundamentally different ways to think about such changes: at fixed position and at a fixed fluid element.

So far, we have considered quantities at a fixed spatial position \mathbf{x} , i.e., $Q = Q(\mathbf{x}, t)$. In this *Eulerian* framework,³ we denote the change of Q over time as $\partial Q/\partial t$. The partial derivatives indicate that we understand the change to refer to a fixed location. There is, however, an entirely different way to think about the fluid: from the perspective of each fluid element. In this *Lagrangian* perspective,⁴ we are moving with the fluid element, $Q = Q(\mathbf{q}, t)$, where \mathbf{q} is some coordinate that labels the fluid elements. The position of the element is then $\mathbf{x}(\mathbf{q}, t)$.

The Lagrangian time derivative, denoted DQ/Dt , is also called the *substantial* or *total* derivative because it contains all changes experienced by a fluid element as it moves through space. We can translate Eulerian to Lagrangian derivatives by considering the total change in Q as the fluid element moves,

$$DQ = \left. \frac{\partial Q}{\partial t} \right|_{\mathbf{x}} Dt + \left. \frac{\partial Q}{\partial x} \right|_{t,y,z} Dx + \left. \frac{\partial Q}{\partial y} \right|_{t,x,z} Dy + \left. \frac{\partial Q}{\partial z} \right|_{t,x,y} Dz, \quad (2.1)$$

where we have made explicit the meaning of the partial derivatives with the other space and time coordinates being held fixed. We divide by Dt to get

$$\frac{DQ}{Dt} = \frac{\partial Q}{\partial t} + \frac{\partial Q}{\partial x} u_x + \frac{\partial Q}{\partial y} u_y + \frac{\partial Q}{\partial z} u_z, \quad (2.2)$$

where we have used that $Dx/Dt = u_x$ and so on. By the definition of the gradient and dot product

²The three papers were originally published in French under the names *Principes généraux de l'état d'équilibre des fluides* (General principles concerning the state of equilibrium of fluids), *Principes généraux du mouvement des fluides* (General principles concerning the motion of fluids), and *Continuation des recherches sur la théorie du mouvement des fluides* (Continuation of the research on the theory of the motion of fluids). The equations in these papers are lengthy, not least because vector notation had not yet been invented!

³The name is not specific to the Euler equations but rather refers to the general way of taking derivatives at a fixed location. It is not clear (to me) why it was named after Euler.

⁴Named after Joseph-Louis Lagrange, presumably because he introduced generalized coordinate systems to solve for the trajectories of mechanical bodies as part of his “Lagrangian mechanics” framework.




Figure 2: Visualization of streamlines and streaklines. Left: steady flow around a wing, illustrating the Bernoulli principle (§3.3). Here, pathlines, streamlines, and streaklines are coincident. Right: streaklines in a wing vortex, a non-steady flow. Images from van Dyke (1982).

operators (§A.1), the derivatives are then related as

$$\boxed{\frac{DQ}{Dt} = \frac{\partial Q}{\partial t} + \mathbf{u} \cdot \nabla Q} \quad (2.3)$$

Physically, this equation says that the infinitesimal change experienced by a fluid element is composed of the change at the Eulerian location of the fluid element plus gradients in Q that the fluid element moves through. The dot product ensures that only gradients in the direction of motion contribute. Equation 2.3 allows us to translate between physics that affects a moving fluid element and physics that happens at a particular location. Moreover, the Eulerian and Lagrangian viewpoints give rise to entirely different numerical methods: we can choose to compute the fluid state in spatially fixed (Eulerian) cells or to follow Lagrangian fluid elements (“particles”). The Lagrangian perspective also underlies certain *kinematical* concepts, namely so-called pathlines, streamlines, and streaklines.

- A **pathline** is the track of a Lagrangian fluid element.
- A **streamline** is tangent to the velocity of the flow.
- A **streakline** connects all fluid elements that have passed, and will pass, through the current location of the fluid element. Such a line could be created by continuously adding a dye at the origin. Compared to pathlines, streaklines move with the flow.

For a steady flow that does not change with time, all three lines are equivalent. Figure 2 shows visual examples; see CC §1.5 for more details.

2.2 The continuity (density) equation

As advertised, we will base our understanding of fluid evolution on the conservation of certain quantities, beginning with mass. In the fluid picture, mass conservation means that any change in mass (or equivalently, density of a fixed-volume element) must be balanced by inflows and outflows. Mathematically, we can express the change of mass as a volume integral,

$$m = \int_V \rho dV \implies \frac{\partial m}{\partial t} = \frac{\partial}{\partial t} \int_V \rho dV. \quad (2.4)$$

The flux of mass into/out of our volume element is density times velocity, but only velocity components perpendicular to the surface contribute. We convert the resulting surface integral to a volume integral using the divergence theorem (§A.3),

$$-\int_S \rho \mathbf{u} \cdot d\mathbf{S} = -\int_V \nabla \cdot (\rho \mathbf{u}) dV. \quad (2.5)$$

The minus sign enters because a positive divergence corresponds to an outflow, and thus to a negative change in the mass in the fluid element. We now set the right hand sides of Equations 2.4 and 2.5 equal and drop the volume integrals to obtain the **Eulerian continuity equation**,

$$\boxed{\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0} \quad (2.6)$$

This equation is general to any conserved quantity that moves along with the fluid velocity. It merely states that changes in that quantity are caused by the flow of the fluid into and out of the fluid element. What does this look like from the fluid element's perspective? We can compute the Lagrangian time derivative of ρ using Equation 2.3,

$$\frac{D\rho}{Dt} = \frac{\partial \rho}{\partial t} + \mathbf{u} \cdot \nabla \rho = -\nabla \cdot (\rho \mathbf{u}) + \mathbf{u} \cdot \nabla \rho = -\rho \nabla \cdot \mathbf{u} - \mathbf{u} \cdot \nabla \rho + \mathbf{u} \cdot \nabla \rho = -\rho \nabla \cdot \mathbf{u} \quad (2.7)$$

where we have substituted the Eulerian derivative from Equation 2.6 and used a vector identity to expand the divergence of $\rho \mathbf{u}$. We thus have the **Lagrangian continuity equation**,

$$\boxed{\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{u}} \quad (2.8)$$

There is also another, perhaps more intuitive way to see the meaning of the Lagrangian formulation, which follows the general derivation in §2.1. By construction, the Lagrangian time derivative takes into account both changes in $\rho(\mathbf{x})$ at fixed t and changes in $\rho(t)$ at fixed \mathbf{x} ,

$$D\rho = \left. \frac{\partial \rho}{\partial t} \right|_{\mathbf{x}} Dt + \left. \frac{\partial \rho}{\partial x_i} \right|_t Dx_i = \left. \frac{\partial \rho}{\partial t} \right|_{\mathbf{x}} Dt + D\mathbf{x} \cdot \nabla \rho|_t. \quad (2.9)$$

Dividing this equation by Dt and using $Dx_i/Dt = u_i$, we find

$$\frac{D\rho}{Dt} = \left. \frac{\partial \rho}{\partial t} \right|_{\mathbf{x}} + \mathbf{u} \cdot \nabla \rho, \quad (2.10)$$

which is the same as Equation 2.3 for $Q = \rho$ and the starting point of Equation 2.7. Either way, the Lagrangian continuity equation tells us that any change in density in the fluid element must be accompanied by a divergence in the velocity field. If the flow is incompressible (density does not change), it must also be divergence-free.

2.3 The velocity equation

We now use the conservation of momentum to derive an equation for changes in the fluid velocity. The derivation is easiest to understand from the Lagrangian point of view because it relies on accelerations (changes in velocity). We apply Newton's second law to a fluid element, $\mathbf{F} = m\mathbf{a}$, or rather $D\mathbf{u}/Dt = \mathbf{F}/m$ where \mathbf{F} is a vector describing the forces acting on the fluid element. The first of these forces is due to **pressure gradients**. Imagine our usual fluid element with volume

δx^3 , with slightly different pressures P_L and P_R at the left and right interfaces. As we showed in Equation 1.7, pressure equals force per area, meaning that the force differential across the fluid element is

$$\delta F = F_R - F_L = -(P_R - P_L) \delta x^2 = -\delta P \delta x^2. \quad (2.11)$$

The minus sign enters because a higher pressure on the right would mean a net force to the left and vice versa. The same equation holds in all three dimensions; the acceleration from the pressure gradient becomes

$$\left(\frac{\delta u_i}{\delta t}\right)_{\text{prs}} = -\frac{\delta P_i \delta x^2}{m} = -\frac{1}{\rho} \frac{\delta P_i}{\delta x} \implies \left(\frac{D\mathbf{u}}{Dt}\right)_{\text{prs}} = -\frac{\nabla P}{\rho}. \quad (2.12)$$

We can also add other accelerations to the right-hand side of this equation. Most importantly, gravity causes an acceleration $\mathbf{g} = -\nabla\Phi$, where Φ is the gravitational potential, which obeys the Poisson equation,

$$\nabla^2\Phi = 4\pi G\rho. \quad (2.13)$$

We combine these effects to write down the **Lagrangian velocity equation**,

$$\boxed{\frac{D\mathbf{u}}{Dt} = -\frac{\nabla P}{\rho} - \nabla\Phi} \quad (2.14)$$

We convert to the **Eulerian velocity equation** using Equation 2.3,

$$\boxed{\frac{\partial\mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla\mathbf{u} = -\frac{\nabla P}{\rho} - \nabla\Phi} \quad (2.15)$$

The second term looks a little counter-intuitive. It means that, at a fixed location in space, the velocity also changes if the fluid elements move along a gradient in velocity. Note the meaning of the $\nabla\mathbf{u}$ term: the ∇ operator is not the gradient of a scalar field here but applied to a vector field. In index notation,

$$(\mathbf{u} \cdot \nabla)\mathbf{u} = u_i \frac{\partial u_j}{\partial x_i} = u_x \frac{\partial \mathbf{u}}{\partial x} + u_y \frac{\partial \mathbf{u}}{\partial y} + u_z \frac{\partial \mathbf{u}}{\partial z}, \quad (2.16)$$

which is a vector quantity (as it must be to enter Equation 2.15).

2.4 The internal energy equation

The final conserved quantity is the energy contained in a fluid element, which we again consider from the Lagrangian viewpoint. The first law of thermodynamics states that the internal energy of a fluid can be changed by two mechanisms: cooling and/or heating processes, and the mechanical work done by the expansion or compression of the fluid, W . We can write this law as

$$\frac{D\varepsilon}{Dt} = -\frac{DW}{Dt} + \left(\frac{D\varepsilon}{Dt}\right)_{\text{heating}} - \left(\frac{D\varepsilon}{Dt}\right)_{\text{cooling}}. \quad (2.17)$$

The minus sign means that we consider the work done by the fluid element, rather than the work done to the fluid element. To understand the W term, imagine one wall of our cubic fluid element being shifted outwards by the force of pressure from within. According to Newtonian mechanics, the work done is

$$m \times dW = Fdx = P\delta x^2 dx = PdV, \quad (2.18)$$

which is why this mechanical work is often called “PdV work.” The added factor of m highlights that Equation 2.17 is in units of energy per unit mass. The volume dV is an imaginary construct, so we convert to physically meaningful quantities by taking the time derivative of Equation 2.18,

$$\frac{DW}{Dt} = \frac{P}{m} \frac{DV}{Dt} = P \frac{D(1/\rho)}{Dt} = -\frac{P}{\rho^2} \frac{D\rho}{Dt} = \frac{P}{\rho} \nabla \cdot \mathbf{u}, \quad (2.19)$$

where we could pull m into the derivative $V/m = 1/\rho$ because the mass of a Lagrangian fluid element is constant. In the final step, we used the Lagrangian continuity equation (2.8). We further define the heating and cooling rates

$$\Gamma \equiv \rho \left(\frac{D\varepsilon}{Dt} \right)_{\text{heating}} \quad \text{and} \quad \Lambda \equiv \rho \left(\frac{D\varepsilon}{Dt} \right)_{\text{cooling}}, \quad (2.20)$$

which are now cooling rates per unit volume. We substitute Equations 2.19 and 2.20 into Equation 2.17 to obtain the **Lagrangian internal energy equation**,

$$\boxed{\frac{D\varepsilon}{Dt} = -\frac{P}{\rho} \nabla \cdot \mathbf{u} + \frac{\Gamma}{\rho} - \frac{\Lambda}{\rho}} \quad (2.21)$$

The **Eulerian internal energy equation** follows from Equation 2.3,

$$\boxed{\frac{\partial \varepsilon}{\partial t} + \mathbf{u} \cdot \nabla \varepsilon = -\frac{P}{\rho} \nabla \cdot \mathbf{u} + \frac{\Gamma}{\rho} - \frac{\Lambda}{\rho}} \quad (2.22)$$

where the second term once again arises due to the motion of fluid elements down a gradient in internal energy (as seen from a fixed, Eulerian position). If there is no heating and cooling ($\Gamma = \Lambda = 0$), we call the flow **adiabatic**.

2.5 Summary of the Euler equations

With all these equations and symbols, it is easy to lose the overview. The following boxes show a summary of the Eulerian fluid equations (2.6, 2.15, 2.22) and their Lagrangian counterparts (2.8, 2.14, 2.21):

Eulerian fluid equations	Lagrangian fluid equations
$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0$	$\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{u}$
$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} - \nabla \Phi$	$\frac{D\mathbf{u}}{Dt} = -\frac{\nabla P}{\rho} - \nabla \Phi$
$\frac{\partial \varepsilon}{\partial t} + \mathbf{u} \cdot \nabla \varepsilon = -\frac{P}{\rho} \nabla \cdot \mathbf{u} + \frac{\Gamma}{\rho} - \frac{\Lambda}{\rho}$	$\frac{D\varepsilon}{Dt} = -\frac{P}{\rho} \nabla \cdot \mathbf{u} + \frac{\Gamma}{\rho} - \frac{\Lambda}{\rho}$

Figure 3 shows a visual summary of the Euler equations. It highlights that the Eulerian derivatives (red and orange) and Lagrangian derivatives (yellow) are related in a uniform way for all quantities. The Eulerian continuity equation is generally written in the shorter form where the $\mathbf{u} \cdot \nabla \rho$ term is absorbed into $\nabla \cdot (\rho \mathbf{u})$, but it still follows the same fundamental form. The green/cyan fields mark

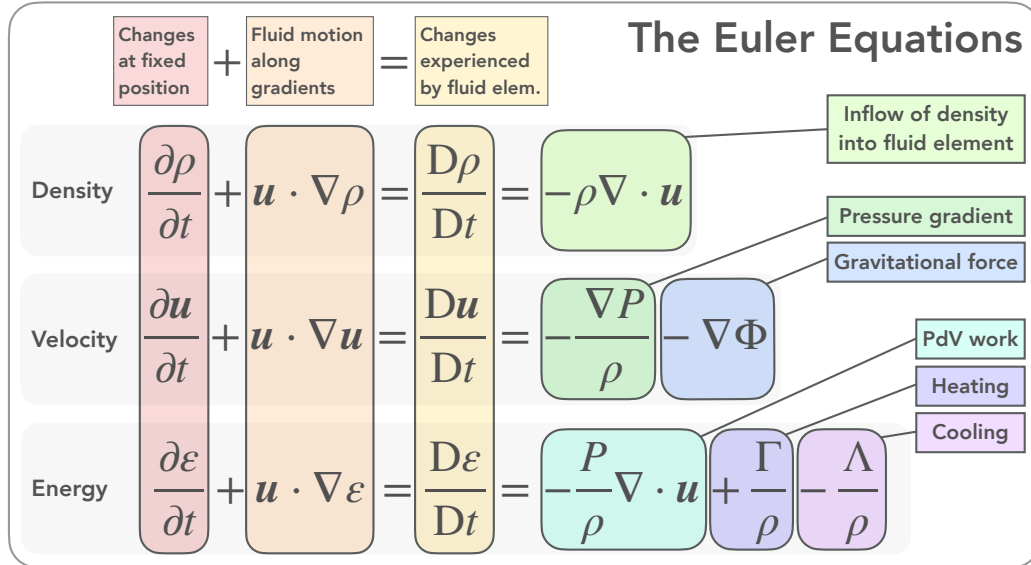


Figure 3: Visual summary of the Eulerian and Lagrangian Euler equations for inviscid fluids. For consistency with the other equations, the Eulerian continuity equation is separated into the gradient and divergence terms.

terms that cause changes in the fluid state due to other fluid variables. The blue/purple terms arise due to non-fluid physics such as gravity, heating, and cooling, and are often called *source terms*.

At this point, we might worry that the Euler equations are underspecified: they contain 6 unknowns (ρ , u_x , u_y , u_z , P , and ε) but we have only five equations.⁵ However, we recall from §1.6 that ρ , P , and ε are connected through the equation of state. This highlights a key point though: the Euler equations can only be solved by invoking knowledge about the microphysics of the gas as summarized by the EoS! Particularly simple solutions arise for barotropic or isothermal equations of state, where $P = P(\rho)$. In this case, pressure, temperature, and energy follow directly from density, and we thus do not need to solve the energy equation at all.

Combined with the equation of state and other closures such as the Poisson equation, we now have a solvable system of equations that predicts the evolution of the averaged fluid quantities. For example, we could imagine a numerical solution where we take a timestep to evolve the Eulerian equations by Δt and use the solution in ρ and ε to compute the ideal gas pressure $P = \rho\varepsilon(\gamma - 1)$. With all variables thus defined, we could move on to the next timestep. In practice, we will transform the Euler equations into a different form for numerical calculations (§5.2), but the system depicted in Figure 3 contains all the necessary information.

⁵The velocity equation has three dimensions. We do not count the additional variables Φ , Γ , and Λ because they need to be imposed externally, e.g., by the Poisson equation or based on atomic processes.

3 Equilibrium and steady flows

In this chapter, we consider situations where the Euler equations are simplified by certain assumptions such as static (equilibrium) solutions and time-invariant flows.

3.1 Hydrostatic equilibrium

The simplest solution to the Euler equations is one without motion, $\mathbf{u} = 0$. In this case, the continuity equation is trivially satisfied with $\partial\rho/\partial t = 0$. If the EoS is barotropic, pressure depends only on density and we can also forego the internal energy equation. In the remaining velocity equation, all time derivatives and all terms with \mathbf{u} cancel, so that we are left with the **hydrostatic condition**,

$$\frac{\nabla P}{\rho} = -\nabla\Phi. \quad (3.1)$$

We will most commonly apply this equation to a mass distribution that obeys either a spherical or planar symmetry. In the former case, we have

$$\nabla P = \frac{\partial P}{\partial r} = -\rho \frac{\partial\Phi}{\partial r} = -\frac{GM(r)\rho(r)}{r^2}, \quad (3.2)$$

where $M(r)$ is the mass enclosed within radius r , or the integral over $\rho(r)$. This form of the hydrostatic condition allows us to derive the density structure of gas balls such as stars (CC §5.4–5.6) or galaxy clusters. Another common application is a planar geometry, where the fluid evolves only in the vertical z -direction,

$$\nabla P = \frac{\partial P}{\partial z} = -\rho \frac{\partial\Phi}{\partial z}. \quad (3.3)$$

This equation describes, for example, the Earth's atmosphere to a good approximation. Here, we set $\partial\Phi/\partial z = g$, ignoring the very slight decrease of Earth's gravity with altitude since $z \ll R_\oplus$. Since we assumed a barotropic EoS, this equation is closed as soon as we provide the functional form of $P = P(\rho)$. To obtain a simple, analytical solution, we make the somewhat unrealistic assumption that the atmosphere is isothermal; in reality, the temperature drops significantly at high altitudes. Nevertheless, we press on and write

$$\frac{1}{\rho} \frac{\partial P}{\partial z} = \frac{1}{\rho} \frac{\partial \rho}{\partial z} \frac{k_B T}{\mu m_p} = \frac{\partial \ln(\rho)}{\partial z} \frac{k_B T}{\mu m_p} = -g, \quad (3.4)$$

where we have used the ideal gas law (Equation 1.19) with a constant T . We integrate over z ,

$$\ln(\rho) = -\frac{\mu m_p g}{k_B T} z + \text{const} \quad (3.5)$$

and exponentiate to get

$$\rho = \rho_0 \exp\left(-\frac{\mu m_p g}{k_B T} z\right) \equiv \rho_0 e^{-z/h_0} \quad (3.6)$$

where we have defined the density at $z = 0$, ρ_0 , and the exponential scale height, $h_0 = k_B T / \mu m_p g$. In air, $\mu \approx 29$ and we assume $T \approx 300$ K, which gives $h_0 \approx 9$ km. That is quite exactly the height of Mount Everest, which is why climbing it without oxygen is so hard! Figure 4 shows the results of numerical calculations assuming a fixed and varying temperature.

In astrophysics, hydrostatic equilibrium is most prominently applied in stars and galaxy clusters. Both systems as far from isothermal, meaning that one has to consider both the density and temperature profiles to fully understand the problem.




Figure 4: Approximate density, pressure, and temperature profiles of Earth’s atmosphere in hydrostatic equilibrium. The light-blue dashed line shows the prediction of Equation 3.6 for a temperature of 270 K. The dark-blue line shows a more realistic calculation with a varying temperature profile (taken from Wikipedia). The slightly lower temperature allows the gas to sink a little lower and have slightly lower pressure at high altitudes, but the isothermal approximation is quite realistic. One caveat is that the atmosphere is cut off at 30 km for numerical reasons. The calculations were performed with the ULULA code (§C).

3.2 Adiabatic flows

The EoS for an ideal gas relates the pressure to the density and internal energy, $P = \rho\varepsilon(\gamma - 1)$. Can we perhaps entirely eliminate the internal energy from the fluid equations by finding a direct relation between pressure and density? In the previous section, we used such a “barotropic” EoS for an isothermal gas, but that is a very constraining condition that will not apply in general.

It turns out we can find a barotropic EoS even if temperature varies, albeit by making another, weaker assumption: that there is no heating or cooling acting on the fluid. This situation is called **adiabatic**, meaning that no heat is transferred between a system and its surroundings. Adiabaticity does not mean that the internal energy of the fluid cannot change: there is still the PdV work that can “adiabatically” heat or cool a gas via compression or expansion. When we impose the adiabatic condition, $\Gamma = \Lambda = 0$, the Lagrangian internal energy equation (2.21) reduces to

$$\frac{D\varepsilon}{Dt} = -\frac{P}{\rho} \nabla \cdot \mathbf{u}. \quad (3.7)$$

We invert the ideal gas EoS to get an expression for the internal energy, $\varepsilon = P/[\rho(\gamma - 1)]$, and convert the RHS using the Lagrangian continuity equation (2.8), $D\rho/Dt = -\rho\nabla \cdot \mathbf{u}$,

$$\begin{aligned} \frac{1}{\gamma - 1} \frac{D(P/\rho)}{Dt} &= \frac{P}{\rho^2} \frac{D\rho}{Dt} \\ \implies \frac{1}{\gamma - 1} \left[\frac{1}{\rho} \frac{DP}{Dt} + P \frac{D(1/\rho)}{Dt} \right] &= \frac{P}{\rho^2} \frac{D\rho}{Dt} \\ \implies \left[\frac{1}{\rho} \frac{DP}{Dt} - \frac{P}{\rho^2} \frac{D\rho}{Dt} \right] &= (\gamma - 1) \frac{P}{\rho^2} \frac{D\rho}{Dt} \\ \implies \frac{1}{\rho} \frac{DP}{Dt} &= \frac{\gamma P}{\rho^2} \frac{D\rho}{Dt} \\ \implies \frac{1}{P} \frac{DP}{Dt} &= \frac{\gamma}{\rho} \frac{D\rho}{Dt} \\ \implies \frac{D \ln P}{Dt} &= \frac{D \ln(\rho^\gamma)}{Dt} \end{aligned} \quad (3.8)$$

which we can integrate to obtain the **equation of state for an adiabatic process**,

$$\boxed{P = K\rho^\gamma} \quad (3.9)$$

where K is some constant (for a given fluid element; it can vary across the fluid). We call a process adiabatic if K stays constant within each fluid element, even if P and ρ change. Given this condition, we could solve the continuity and velocity equations directly without considering the internal energy, which reduces to $\varepsilon_{\text{adiabatic}} = K\rho^{\gamma-1}/(\gamma-1)$.

Physically, Equation 3.9 means that we can think of pressure as having a one-to-one relation with density: the more we compress a fluid, the more it will resist further compression. Since pressure depends on the random motions of particles (§1.4), it depends on the fraction of the compression work that goes into thermal motions. This fraction is highest when there are no other degrees of freedom into which the mechanical work can be absorbed, i.e., when $\gamma = 5/3$. Since the Lagrangian energy equation was a re-statement of the first law of thermodynamics, Equation 3.9 can also be derived from thermodynamical considerations (CC §4.2).

What about entropy? It is tempting to assume that no heat being added also means that the entropy cannot increase. However, the second law of thermodynamics states that $dS \geq dQ/T$ (where dQ stands for an infinitesimal amount of heat), meaning that the entropy, S , can increase even if the heat does not! The equality holds only if the process is **reversible**, in which case we call the process **isentropic**. In other words, an adiabatic but irreversible process can still generate entropy.

While Equation 3.9 gives intuitive insight into adiabatic processes, it does not allow us to solve for the evolution of a general fluid because Lagrangian fluid elements move around and because K varies across the fluid. The exception to the second statement is an **isentropic fluid**, where K is the same everywhere in the fluid (corresponding to a constant entropy per unit mass).

On Earth, there are plenty of examples of approximately adiabatic processes. The pistons in a Diesel engine compress an air-fuel mixture until it ignites; this works only because the metal in the engine cannot dissipate the heat fast enough to prevent the temperature from rising. We will also encounter numerous adiabatic processes in astrophysics.

3.3 Barotropic flows and Bernoulli's principle

As we saw in the previous section, a barotropic EoS has the powerful advantage that it eliminates the internal energy equation, a major simplification of the Euler equations. In this section, we explore the consequences of an additional assumption: that the flow is steady, $\partial/\partial t = 0$ (but $\mathbf{u} \neq 0$). We recall that for time-invariant flows, pathlines, streamlines, and streaklines are identical. Thus, we can try to find a quantity Q that is conserved for each Lagrangian element, which would manifest as

$$\frac{DQ}{Dt} = \frac{\partial Q}{\partial t} + \mathbf{u} \cdot \nabla Q = 0 \implies \mathbf{u} \cdot \nabla Q = 0, \quad (3.10)$$

since $\partial Q/\partial t$ is zero for all quantities in a time-invariant flow. We thus try to express the Eulerian velocity equation in terms of a gradient. Most terms are already in this form. Since $\partial \mathbf{u}/\partial t = 0$, we can rewrite Equation 2.15 as

$$\mathbf{u} \cdot \nabla \mathbf{u} + \frac{\nabla P}{\rho} + \nabla \Phi = 0. \quad (3.11)$$

For the first term, we use the vector identity in Equation A.19 to write

$$\nabla \left(\frac{1}{2} u^2 \right) - \mathbf{u} \times \boldsymbol{\omega} + \frac{\nabla P}{\rho} + \nabla \Phi = 0, \quad (3.12)$$




Figure 5: Illustrations of the Bernoulli principle. Left: a Venturi meter measures the pressure at two pipe segments with different circumferences, finding a lower pressure where the fluid velocity is faster (image source unknown). Right: The stream falling from a faucet narrows as the water speeds up.

where we have defined the **vorticity**,

$$\boldsymbol{\omega} \equiv \nabla \times \mathbf{u}. \quad (3.13)$$

We will deal with the vorticity term in a second, but first we recognize that the barotropic condition, $P = P(\rho)$, allows us to convert the third term into a pure gradient field,

$$\frac{\nabla P}{\rho} = \nabla \int \frac{dP}{\rho}. \quad (3.14)$$

To obtain the condition in Equation 3.10, namely that Q be constant for a fluid element, we take the dot product of Equation 3.12 with the velocity,

$$\mathbf{u} \cdot \nabla \left[\frac{u^2}{2} + \int \frac{dP}{\rho} + \Phi \right] - \mathbf{u} \cdot (\mathbf{u} \times \boldsymbol{\omega}) = 0. \quad (3.15)$$

The last term cancels because the dot-product of a cross-product involving the same vector is always zero (§A.3). That means we have succeeded in defining a quantity that is conserved along streamlines in a steady, barotropic flow. It is called the **Bernoulli constant**,

$$H \equiv \frac{u^2}{2} + \int \frac{dP}{\rho} + \Phi \quad (3.16)$$

For an adiabatic, ideal fluid, we can solve the integral term. We have $P = K\rho^\gamma$ and thus $dP = K\gamma\rho^{\gamma-1}d\rho$, which lets us write

$$\int \frac{dP}{\rho} = \int \frac{K\gamma\rho^{\gamma-1}}{\rho} d\rho = \frac{\gamma K\rho^{\gamma-1}}{\gamma-1} = \frac{\gamma}{\gamma-1} \frac{P}{\rho} = \varepsilon + \frac{P}{\rho}. \quad (3.17)$$

We insert this solution to obtain the **Bernoulli constant for an adiabatic, ideal fluid**,

$$H = \frac{u^2}{2} + \varepsilon + \frac{P}{\rho} + \Phi \quad (3.18)$$

We can interpret H as the sum of kinetic energy, internal energy, work done by pressure gradients, and gravitational potential. Alternatively, the integral term can be expressed as internal energy

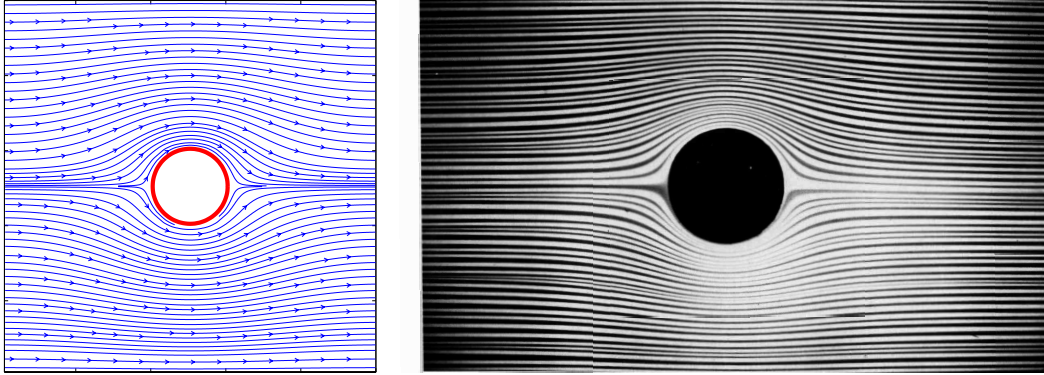


Figure 6: Potential flow around a cylinder in theory (left, from Zhang et al. 2007) and experiment (right, from van Dyke 1982). In the latter, potential flow is mimicked by a very viscous, slow flow between two glass plates; counter-intuitively, this happens to give very similar streamlines to the opposite case of zero viscosity. This type of flow, where parallel layers that do not mix, is also called a “laminar” flow.

plus P/ρ . In other words, the kinetic energy changes because the fluid element is accelerated by pressure gradients or by gravity. This can be shown more specifically assuming an **isentropic flow** (where entropy is constant) and that the gravitational potential varies slowly. In that case,

$$\mathbf{u} \cdot \frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho} \frac{DP}{Dt} \quad (3.19)$$

which is known as **Bernoulli’s principle**. It shows that an increase in fluid speed goes along with a decrease in pressure. For a formal derivation, see e.g. vdB §9. One of the most commonly cited consequences of Bernoulli’s principle is the flow around an airplane wing. The idea is that wings are designed so that the top streamlines are slightly longer than the bottom ones, leading to a pressure gradient and thus to lift (Figure 2), but there is significant debate as to the validity of this explanation.

For incompressible flows, we can visualize a Bernoulli flow as a tube with a varying cross-section. A smaller circumference means the fluid has to flow faster, and thus leads to lower pressure (Figure 5 shows some examples). We can intuitively understand this principle even without a physical tube. For example, the water stream falling from a faucet holds together without being confined. Does it narrow as it accelerates under gravity? Say the stream starts out at a height z_0 with zero velocity. The gravitational potential on Earth is $\Phi(z) = gz$. Water is very nearly incompressible under everyday conditions, implying $\partial\rho/\partial t = 0$. In Equation 2.5, we saw that this means that no momentum flux, $\rho\mathbf{u}$, traverses the surface area of our fluid elements. In the faucet scenario, this translates to the total flow through the stream being constant, and thus $\rho u_z A = \text{const}$. Here, $A = \pi r^2$ is the stream’s cross-section, which then depends on the velocity as $A \propto 1/u_z$. We now invoke

$$H = \frac{u_z^2}{2} + gz = \text{const}, \quad (3.20)$$

where we neglect the pressure term in Equation 3.16 because the conversion from gravitational to kinetic energy is by far the dominant process in this case. Thus, $u_z = -\sqrt{2g(z_0 - z)}$ and the radius of the stream evolves as

$$r \propto \sqrt{A} \propto u_z^{-1/2} \propto (z_0 - z)^{-1/4}, \quad (3.21)$$

which means that the stream narrows as it falls. You can prove this result experimentally at home by turning on a faucet (Figure 5)!




Figure 7: Examples of boundary layer separation in the potential flow around solid objects. The top row shows photographs of the average streamlines in low-viscosity flows around a sphere and a cylinder. The turbulent wake behind the objects is clearly visible. The bottom images show simulations of the flow around golf ball without (left) and with (right) dimples. In this case, the dimples lead to a more complex flow pattern which ends up creating a smaller low-pressure wake behind the ball, and thus a smaller drag force. Images from van Dyke 1982 (top) and cradle-cfg.com (bottom).

3.4 The limits of inviscid fluid dynamics: d’Alembert’s paradox

The situations we considered in the previous section seem to generate patterns where fluid flows in different layers that do not mix (also called a “laminar” flow). This impression can be formalized for a steady, adiabatic flow of an ideal fluid: if we start off with no vorticity, $\boldsymbol{\omega} = 0$, the flow stays vorticity-free. This conservation of vorticity is known as **Kelvin’s circulation theorem** (CC §9.1). A vorticity-free flow is called a **potential flow** because the velocity can be expressed as the gradient of a “velocity potential” ψ_u ,

$$\boldsymbol{\omega} = \nabla \times \mathbf{u} = \nabla \times (-\nabla \psi_u) = 0. \quad (3.22)$$

Figure 6 shows an example, namely the laminar flow around a cylinder. While potential flow solutions seem elegant, they run into a serious problem. Let’s consider the drag force on a sphere in a potential flow. There should be two pressure forces acting on the sphere: thermal pressure and so-called ram pressure. If the temperature and density of the fluid are constant (isothermal and incompressible flow), the thermodynamic pressure is isotropic and cancels out. Ram pressure is essentially a different term for momentum flux. The momentum per unit volume is $\rho \mathbf{u}$ and its flux is $P_{\text{ram}} = \rho \mathbf{u} \cdot \mathbf{u} = \rho u^2$. Since pressure is force per area, the drag force should be roughly $F_{\text{drag}} \approx \rho u^2 \pi r^2$, where r is the radius of the sphere.

This solution looks reasonable enough until we question what happens when we reverse the direction of the flow. The force is still the same! This is not specific to the drag flow problem: the velocity enters into Bernoulli’s constant only as u^2 , meaning that the streamlines are independent

of the flow direction. The only way to reconcile the solutions for two opposing flow directions is that the **drag force must be zero!** But we know from everyday experience that that is far from the truth: think about biking in a headwind, for example. This conundrum is called **d’Alembert’s paradox**. Its realization in 1752 discredited the field of fluid mechanics as a whole. After all, the “reversibility problem” traces right back to the Euler equations: their predictions for ρ and P do not change when $\mathbf{u} \rightarrow -\mathbf{u}$.

So what is the solution to the paradox? In one word, viscosity. The exchange of momentum between adjacent fluid elements becomes more important on small scales. Thus, we can always find a scale small enough for the inviscid large-scale behavior of a fluid to break down. At the edges of solid bodies, a so-called “boundary layer” forms between fluid that is “stuck” to the object and the rest of the flow. The large velocity gradient across this thin layer generates vorticity (because Kelvin’s circulation theorem is no longer valid for viscous flows), which causes the pressure distribution around the body to become anisotropic. In particular, a low-pressure region behind the body leads to a net drag force against the direction of the motion. In practice, the flow solutions get very complicated (see the example of a golf ball in Figure 7).

Given d’Alembert’s paradox, how can we justify using inviscid fluid dynamics at all in astrophysics? Luckily, there are virtually no fluid-solid interfaces we need to consider, meaning that the drag force on solid bodies is generally of little interest. In many astrophysical scenarios, the viscosity is so low that the Euler equations are a good approximation.

4 Sound waves

From everyday experience, we know that sound waves can travel long distances in air, which indicates that the Euler equations admit some sort of wave solution. In this section, we derive the corresponding wave equation and some properties of sound waves. Compared to the previous section, where we simplified the Euler equations by dropping various terms, we now simplify them by assuming that the variations due to waves are small.

4.1 The linearized Euler equations

In general, waves occur when a perturbation to some physical system leads to an oscillating motion around a constant state. For example, the sound of a violin is generated by small deviations from the default straight-line position of a string. For sound waves, the physical picture is that pressure and density oscillate around a constant fluid state. For the system to produce a stable wave, the perturbation cannot permanently alter this background state. Thus, we assume that the perturbation is small enough that it does not heat the medium significantly, i.e., that the system is adiabatic. We know that the adiabatic EoS is barotropic, so that we can neglect the energy equation (§3.2). Mathematically speaking, we express the perturbation as a small change $(\rho_1, \mathbf{u}_1, P_1)$ to an unperturbed, stationary fluid state $(\rho_0, 0, P_0)$,

$$\begin{aligned}\rho &= \rho_0 + \rho_1 \\ P &= P_0 + P_1 \\ \mathbf{u} &= \mathbf{u}_1.\end{aligned}\tag{4.1}$$

The continuity and velocity equations now read

$$\frac{\partial(\rho_0 + \rho_1)}{\partial t} = -\nabla \cdot ([\rho_0 + \rho_1]\mathbf{u}_1) = -\mathbf{u}_1 \cdot \nabla(\rho_0 + \rho_1) - (\rho_0 + \rho_1)\nabla \cdot \mathbf{u}_1\tag{4.2}$$

and

$$\frac{\partial \mathbf{u}_1}{\partial t} + \mathbf{u}_1 \cdot \nabla \mathbf{u}_1 = -\frac{\nabla(P_0 + P_1)}{\rho_0 + \rho_1}.\tag{4.3}$$

We cancel $\nabla \rho_0 = \nabla P_0 = \partial \rho_0 / \partial t = 0$ and **linearize** the equations, meaning that we throw out all terms where more than one perturbing quantity (ρ_1 , P_1 , or \mathbf{u}_1) are multiplied (since they will be sub-dominant). We are left with

$$\frac{\partial \rho_1}{\partial t} + \rho_0 \nabla \cdot \mathbf{u}_1 = 0\tag{4.4}$$

$$\frac{\partial \mathbf{u}_1}{\partial t} + \frac{\nabla P_1}{\rho_0} = 0,\tag{4.5}$$

where we have used that

$$\frac{\nabla P_1}{\rho_0 + \rho_1} \approx \frac{\nabla P_1}{\rho_0} \left(1 - \frac{\rho_1}{\rho_0}\right) = \frac{\nabla P_1}{\rho_0} + \mathcal{O}(\rho_1^2).\tag{4.6}$$

We can solve this system because the EoS is barotropic, allowing us to convert P to ρ so that we have only two variables for two equations. In particular, we expand the perturbation in pressure,

$$P(\rho_0 + \rho_1) = P(\rho_0) + \left(\frac{\partial P}{\partial \rho}\right)_0 \rho_1 + \mathcal{O}(\rho_1^2),\tag{4.7}$$

which we insert into the velocity equation to find

$$\frac{\partial \mathbf{u}_1}{\partial t} + \left(\frac{\partial P}{\partial \rho} \right)_0 \frac{\nabla \rho_1}{\rho_0} = 0. \quad (4.8)$$

We can now solve for either ρ_1 or \mathbf{u}_1 , for example by taking the time derivative of Equation 4.4,

$$\frac{\partial^2 \rho_1}{\partial t^2} = -\rho_0 \nabla \cdot \left(\frac{\partial \mathbf{u}_1}{\partial t} \right) = \left(\frac{\partial P}{\partial \rho} \right)_0 \nabla \cdot \nabla \rho_1, \quad (4.9)$$

which we recognize as a **wave equation**,

$$\boxed{\frac{\partial^2 \rho_1}{\partial t^2} - \left(\frac{\partial P}{\partial \rho} \right)_0 \nabla^2 \rho_1 = 0} \quad (4.10)$$

4.2 The properties of sound waves

So far, we have not assumed that the solutions to Equation 4.10 are necessarily waves. We show that by inserting the standard solution,

$$\rho_1 = \rho_s e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)} \quad \text{and} \quad \mathbf{u}_1 = \mathbf{u}_s e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}, \quad (4.11)$$

with amplitudes ρ_s and \mathbf{u}_s , a wave vector $\mathbf{k} = k \hat{\mathbf{k}} = 2\pi/\lambda \hat{\mathbf{k}}$, and an angular frequency $\omega = 2\pi f$ (where f is the frequency in Hz). This functional form allows for waves, but also for exponentially increasing or decreasing solutions (depending on whether \mathbf{k} and ω are real or complex). We find out which type of solution is present by computing the ratio ω/k , which is also the wave speed, $c_s = \lambda f = \omega/k$. By inserting the waveform and taking derivatives, we find

$$c_s = \frac{\omega}{k} = \sqrt{\left(\frac{\partial P}{\partial \rho} \right)_0}. \quad (4.12)$$

Most importantly, ω/k is real-valued, meaning that our fluid system supports the stable propagation of sound waves. If it was complex, the perturbation would exponentially grow or be damped. Another important property of Equation 4.12 is that c_s does not depend on ω , meaning that all frequencies propagate at the same speed. This “non-dispersive” nature of sound waves is good news for sound as a communication medium! Otherwise, we would perceive the different frequencies from far-away sound sources spread out over time. Physically, the expression for c_s also fits into our picture of waves, whose speed is generally

$$\text{wave speed} = \sqrt{\frac{\text{restoring force}}{\text{inertia}}}. \quad (4.13)$$

For example, for waves on a string, we have tension as a restoring force and density per unit length as inertia. We can thus interpret pressure as the force that restores the initial fluid state when perturbed and density as the inertia of the fluid. The sound speed depends on how strongly pressure responds to a change in density.

So far, we derived the wave properties from the excess density ρ_1 , but how does the fluid actually move? We can also solve Equations 4.4 and 4.5 for \mathbf{u}_1 to find

$$\mathbf{u}_1 \cdot \hat{\mathbf{k}} = \frac{\omega}{k} \frac{\rho_1}{\rho_0} = c_s \frac{\rho_1}{\rho_0}. \quad (4.14)$$

The fact that the ratio between density and velocity is rational means that they are in phase, i.e., that a density peak also corresponds to a velocity peak. Moreover, the motion of the fluid elements is parallel to the wave, implying that **sound waves are longitudinal** (as opposed to transverse waves on a string, for example). This behavior is encoded in the Euler equations because fluid elements are accelerated in the direction of pressure gradients. Finally, Equation 4.14 tells us that the fluid elements move with a small fraction of the sound speed, since our derivation rested on the assumption that $\rho_1 \ll \rho_0$. When the wave overdensity and excess pressure become comparable to the background density and pressure, the wave treatment breaks down and we are dealing with shocks (§7).

Given these findings, how should we think of the sources of sound? For a continuous wave, we clearly need an oscillating object such as a voice chord or string in an instrument, although our ears can also perceive less regular sonic events. Equation 4.14 implies that the velocity of the oscillating object (which is the same as that of the fluid elements) is much smaller than the speed of sound.

The last remaining question is what the sound speed actually is for realistic scenarios. We have left Equation 4.12 general, but in the majority of cases we will deal with ideal gases. Moreover, we assumed that the wave propagation is adiabatic, so $P = K\rho^\gamma$, and the sound speed becomes

$$c_{s,\text{ideal}} = \sqrt{\frac{\gamma P}{\rho}} \quad (4.15)$$

Here, the word “adiabatic” implies that the fluid cannot exchange heat quickly enough for the perturbations to even out in temperature, which is increased and decreased by the PdV work due to the oscillating pressure. This hypothesis is testable because we can also work out the sound speed for an isothermal sound wave, where $P \propto \rho$ and thus

$$c_{s,\text{isothermal}} = \sqrt{\frac{P}{\rho}} = \frac{c_{s,\text{ideal}}}{\sqrt{\gamma}}. \quad (4.16)$$

In air, we have $P = 10^6$ dyne/cm², $\rho \approx 1.2 \times 10^{-3}$ g/cm³, and $\gamma = 7/5$ for a gas composed of diatomic molecules, which gives the well known sound speed of $c_s \approx 3.4 \times 10^4$ cm/s. The isothermal sound speed would be about $1/\sqrt{7/5} \approx 15\%$ slower than the adiabatic sound speed. The adiabatic speed matches experimental data better, which means that the conduction of heat in air is slow enough to render wave propagation an adiabatic process.

5 Computational hydro I: Theoretical background

In this chapter, we consider some theoretical aspects that will become important when we solve hydrodynamics numerically. We will bring the Eulerian equations into a dimensionless, conservation-law form suitable for computing, briefly consider the mathematical properties of those equations, and study their basic numerical properties. As in any computational approach, we need to decompose the problem into a finite number of discrete numerical elements. We could adopt the Lagrangian viewpoint and consider finite mass elements (virtual “particles”), an approach that leads to algorithms such as Smoothed Particle Hydrodynamics (SPH). In this course, however, we will focus on Eulerian methods, where we divide space into an imaginary grid of cells.

5.1 The self-similarity of hydrodynamics

So far, we have been working in CGS units without much consideration for the reasoning behind this decision. If we want to solve astrophysical fluid systems, this choice may be unwise: after all, many astrophysical quantities correspond to large numbers in CGS. For example, expressing the Gpc scales of cosmological volumes in cm might cause computational issues and lead to values that are hard to interpret without unit conversions.

More generally, we might ask why there has not been a need to think about units. We have encountered a few dimensional quantities such as k_B and m_p , but they appeared exclusively in the EoS. In fact, the Euler equations do not contain any physical constants at all! What, then, is the best unit system for these equations?

Let us try to generalize our units by introducing a characteristic length, l_0 , a characteristic mass, m_0 , and a characteristic time scale, t_0 . These numbers could roughly correspond to the size of the system under consideration, to the mass of the system, and to the length of time for which the fluid processes act. We can also construct derived units, namely,

$$u_0 \equiv \frac{l_0}{t_0} \quad \rho_0 \equiv \frac{m_0}{l_0^3} \quad P_0 = E_0 \equiv \frac{m_0}{t_0^2} \frac{l_0}{l_0} \quad \varepsilon_0 \equiv \frac{l_0^2}{t_0^2}. \quad (5.1)$$

Note that we needed only three characteristic scales to define all derived scales. We can now transform our usual variables into re-scaled, dimensionless fluid quantities,

$$\tilde{\mathbf{u}} \equiv \frac{\mathbf{u}}{u_0} \quad \tilde{\rho} \equiv \frac{\rho}{\rho_0} \quad \tilde{P} \equiv \frac{P}{P_0} \quad \tilde{\varepsilon} \equiv \frac{\varepsilon}{\varepsilon_0} \quad \tilde{\Phi} \equiv \frac{\Phi}{\varepsilon_0} \quad \tilde{\Gamma} \equiv \frac{\Gamma}{E_0/t_0} \quad \tilde{\Lambda} \equiv \frac{\Lambda}{E_0/t_0} \quad (5.2)$$

and so on. The differential operators become

$$\frac{\partial}{\partial t} \rightarrow \frac{1}{t_0} \frac{\partial}{\partial \tilde{t}} \quad \nabla \rightarrow \frac{1}{l_0} \tilde{\nabla} \quad (5.3)$$

where $\tilde{\nabla} = \partial/\partial(x_i/l_0)$. We can now transform the Eulerian continuity equation 2.6 to

$$\frac{\partial \tilde{\rho}}{\partial \tilde{t}} \frac{\rho_0}{t_0} + \frac{1}{l_0} \tilde{\nabla} \cdot (\tilde{\rho} \tilde{\mathbf{u}}) \rho_0 u_0 = 0 \implies \frac{\partial \tilde{\rho}}{\partial \tilde{t}} + \tilde{\nabla} \cdot (\tilde{\rho} \tilde{\mathbf{u}}) = 0 \quad (5.4)$$

and proceed similarly for the velocity equation,

$$\frac{\partial \tilde{\mathbf{u}}}{\partial \tilde{t}} \frac{u_0}{t_0} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\mathbf{u}} \frac{u_0^2}{l_0} = -\frac{\tilde{\nabla} \tilde{P}}{\tilde{\rho}} \frac{P_0}{\rho_0 l_0} - \tilde{\nabla} \tilde{\Phi} \frac{\varepsilon_0}{l_0} \implies \frac{\partial \tilde{\mathbf{u}}}{\partial \tilde{t}} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\mathbf{u}} = -\frac{\tilde{\nabla} \tilde{P}}{\tilde{\rho}} - \tilde{\nabla} \tilde{\Phi} \quad (5.5)$$

and for the internal energy equation,

$$\frac{\partial \tilde{\varepsilon}}{\partial \tilde{t}} \frac{\varepsilon_0}{t_0} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\varepsilon} \frac{u_0 \varepsilon_0}{l_0} = -\frac{\tilde{P}}{\tilde{\rho}} \tilde{\nabla} \cdot \tilde{\mathbf{u}} \frac{P_0 u_0}{\rho_0 l_0} + \left(\frac{\tilde{\Gamma}}{\tilde{\rho}} - \frac{\tilde{\Lambda}}{\tilde{\rho}} \right) \frac{E_0}{t_0 \rho_0} \implies \frac{\partial \tilde{\varepsilon}}{\partial \tilde{t}} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\varepsilon} = -\frac{\tilde{P}}{\tilde{\rho}} \tilde{\nabla} \cdot \tilde{\mathbf{u}} + \frac{\tilde{\Gamma}}{\tilde{\rho}} - \frac{\tilde{\Lambda}}{\tilde{\rho}}. \quad (5.6)$$

In all cases, we got back the same Euler equations but for the dimensionless quantities! This is a profound insight: **any solution of the Euler equations represents an entire family of equal solutions** for problems with different l_0 , m_0 , and t_0 — as long as all other quantities are scaled accordingly. There is nothing special about length, mass, and time. For example, we could have set u_0 instead of t_0 and derived $t_0 = l_0/u_0$. However, we can only set three fundamental scales, or the remaining scales will be over-specified.

Let's visualize the power of rescaling a system with an example. Imagine a 1 cm^3 cell filled with 1 gram of an ideal gas with $T = 300 \text{ K}$ in which we have initialized some flow at a velocity of 1 cm/s . We could express this problem conveniently by choosing $l_0 = 1 \text{ cm}$, $t_0 = 1 \text{ s}$, and $m_0 = 1 \text{ g}$. Once we have found a solution for this situation, we could rescale it to an equivalent astrophysical problem, say, $m_1 \equiv 1 M_\odot$ of gas in a $l_1 \equiv 1 \text{ pc}$ cube over $t_1 \equiv 10^9$ years. These decisions would then fix the rest of the problem: the density is $\rho_1 = 1 M_\odot/\text{pc}^3 \approx 6.8 \times 10^{-23} \text{ g/cm}^3 \approx 4 m_p/\text{cm}^3$ and the velocity of the flow is $u_1 = 1 \text{ pc/Gyr} \approx 98 \text{ cm/s}$. We would also need to rescale the internal energy, pressure, or temperature. From the ideal gas EoS, we find the new temperature, $T' = T(l_1/l_0)^2(t_0/t_1)^2 \approx 2.9 \times 10^6 \text{ K}$. Moreover, Equation 5.2 tells us how the gravitational potential and heating/cooling terms would need to change to obtain the same problem. Of course, we could also go the other way and set a particular density or temperature that we are interested in, and work back to the size and timescale of the problem.

In hindsight, is it obvious that the Euler equations are fundamentally scale-free? After all, the units of each equation are consistent, meaning that the dimensionless scalings must cancel out of each equation. However, that does not mean that the equations cannot contain terms that would depend on l_0 , m_0 , and t_0 . The real reason for the scale-free nature of the Euler equations is that they do not contain any physical (dimensional) constants or fluid properties. Imagine, for example, that the proton mass were to appear in the equations: we would end up with a term like m_p/m_0 in the dimensionless equations, which would manifestly depend on our chosen unit system. One example where additional physics adds a scale is viscosity. In its simplest form, “dynamic” viscosity adds a term

$$+\frac{\eta}{\rho} \nabla^2 \mathbf{u} \quad (5.7)$$

on the RHS of the velocity equation, where η parametrizes the strength of viscosity. Since the units of the equation are those of $\partial \mathbf{u}/\partial t$ or $\text{g/cm}^2/\text{s}^2$ in CGS, the units of η must be g/cm/s (which is called a Poise). For example, water at room temperature has a viscosity of about 0.01 Poise, but is that a little or a lot? The answer depends on the problem in question, and it is best found by returning to the non-dimensional Euler equations. During the procedure laid out above, we multiply the entire velocity equation by t_0/u_0 (see Equation 5.5), meaning that the new term becomes

$$\frac{t_0}{u_0} \frac{\eta}{\rho_0} \frac{u_0}{l_0^2} \times \frac{1}{\tilde{\rho}} \tilde{\nabla}^2 \tilde{\mathbf{u}} = \frac{\eta}{\rho_0 u_0 l_0} \times \frac{1}{\tilde{\rho}} \tilde{\nabla}^2 \tilde{\mathbf{u}} \equiv \frac{1}{\text{Re}} \frac{1}{\tilde{\rho}} \tilde{\nabla}^2 \tilde{\mathbf{u}}, \quad (5.8)$$

where we have defined the dimensionless **Reynolds number**, $\text{Re} \equiv \rho_0 u_0 l_0/\eta$. While η tells us how important viscosity is in absolute terms (how much velocity gradients are smoothed out by particle collisions), the Reynolds number tells us how important viscosity is in a problem with given length, density, and velocity scales (with larger Reynolds number indicating a lower impact of viscosity, counter-intuitively). Staying with the example of water where $\rho_0 = 1 \text{ g/cm}^3$, we would consider

an ocean with $l_0 \sim 10^8$ cm, $u_0 \sim 100$ cm/s, and thus $\text{Re} \sim 10^8$, meaning that the viscosity term is strongly suppressed. Conversely, if we consider water under a microscope where, say, $l_0 \sim 0.1$ cm and $u_0 \sim 1$ cm/s, we have $\text{Re} \sim 10$ and thus a significant impact of viscosity.

The key takeaway is that while the basic Euler equations are scale-free, any additional physics likely is not and breaks the self-similarity. For the version of the equations that we consider in this course, this is true also for gravitational accelerations (e.g., g on Earth) or cooling (which depends on the physical temperature and density). However, we can always express constants such as g in units of l_0 , m_0 , and t_0 , and thus introduce dimensionless numbers into the non-dimensional Euler equations. The result is a huge set of such numbers. For example, the *Froude number* quantifies the importance of gravity in a problem, while the *Prandtl number* quantifies the relative importance of thermal diffusion and momentum diffusion (viscosity).

In terms of numerics, the scale-free nature of the Euler equations is an invaluable asset because we get to choose a set of “code units” (l_0 , m_0 , and t_0) for a given problem. The code then solves the equations in dimensionless units, and we translate back into physical units when analyzing the solution. This procedure not only ensures that the code can be almost entirely ignorant of the unit system, it also means that we can choose a system where the dimensionless fluid quantities remain within a reasonable range to avoid round-off errors (§5.7).

5.2 The conservation-law form of the equations

In §2, we derived the Euler equations from the conservation of mass, momentum, and energy. However, two of the three equations were not phrased in terms of the conserved quantities, since we wrote the momentum equation in terms of velocity and the energy equation in terms of internal energy rather than the conserved total energy. For computational hydrodynamics, it will be useful to recast the equations in terms of the conserved quantities ρ , $\rho\mathbf{u}$, and E .

The continuity equation is already in the form of a conservation law: the change of mass with time (per unit volume) is balanced by the divergence of momentum flux, which corresponds to the inflow/outflow of mass. We can generalize this form as a **conservation law** for any conserved quantity Q ,

$$\frac{\partial}{\partial t}(\text{density of } Q) + \nabla \cdot (\text{flux of } Q) = 0 \quad (5.9)$$

Purely based on dimensional analysis, the flux must have units of $Q \times \text{velocity}$, as exemplified by the mass conservation (continuity) equation where $Q = m$, the density is ρ , and the flux is $\rho\mathbf{u}$. For the other conserved quantities, we should be able to come up with similar expressions, given that a change of a conserved quantity with time must be balanced by an inflow or outflow. Indeed, we can recast the equations for \mathbf{u} and ε as conservation laws of $\rho\mathbf{u}$ and E . The derivation is a little tedious and can be found in §B.1. It brings our equations into a coherent form:

Fluid equations in conservation-law form

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\mathbf{u}) &= 0 \\ \frac{\partial(\rho\mathbf{u})}{\partial t} + \nabla \cdot (\rho\mathbf{u} \otimes \mathbf{u} + \mathbf{IP}) &= -\rho\nabla\Phi \\ \frac{\partial E}{\partial t} + \nabla \cdot ([E + P]\mathbf{u}) &= \rho\frac{\partial\Phi}{\partial t} + \Gamma - \Lambda \end{aligned}$$

We have introduced the momentum flux density tensor $\Pi_{ij} \equiv \rho \mathbf{u} \otimes \mathbf{u} + \mathbf{I}P = \rho u_i u_j + \delta_{ij}P$, where \mathbf{I} is the 3×3 identity matrix and δ_{ij} the Kronecker delta (§A.2). The total energy equation states something fairly profound: the flux of total energy is energy plus pressure times velocity. The RHS of the equations is not necessarily zero because we allow for so-called **source terms**, that is, changes due to fields other than the fluid quantities themselves (namely, gravity, heating, and cooling). We can make the notation even more compact by collecting the fluid quantities, fluxes, and source terms into vectors,

$$\boxed{\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathcal{F}(\mathbf{U}) = \mathbf{S}} \quad (5.10)$$

where we have defined the vector of conserved fluid quantities, \mathbf{U} , the vector of flux terms, $\mathcal{F}(\mathbf{U})$, and the vector of source terms, \mathbf{S} ,

$$\mathbf{U} \equiv \begin{pmatrix} \rho \\ \rho u_x \\ \rho u_y \\ \rho u_z \\ E \end{pmatrix} \quad \mathcal{F}(\mathbf{U}) \equiv \begin{pmatrix} \rho \mathbf{u} \\ \rho u_x \mathbf{u} + \delta_{xj}P \\ \rho u_y \mathbf{u} + \delta_{yj}P \\ \rho u_z \mathbf{u} + \delta_{zj}P \\ (E + P)\mathbf{u} \end{pmatrix} \quad \mathbf{S} \equiv \begin{pmatrix} 0 \\ -\rho \partial\Phi/\partial x \\ -\rho \partial\Phi/\partial y \\ -\rho \partial\Phi/\partial z \\ \rho \partial\Phi/\partial t + \Gamma - \Lambda \end{pmatrix}$$

where δ_{xj} means that P is only added to the x component of the flux vector, and so on. We have thus reduced the fluid equations to an orderly system of hyperbolic partial differential equations (PDEs, §5.4). Many numerical techniques have been developed to solve such systems. As we consider these algorithms in §6 and §8, we will often need to switch between the conserved vector \mathbf{U} and the standard, non-conservative quantities in the Euler equations. We call the latter the **primitive fluid variables** and also write them as a vector,

$$\mathbf{V} \equiv (\rho, u_x, u_y, u_z, P). \quad (5.11)$$

We can easily convert between \mathbf{V} and \mathbf{U} using the definition of $E = \rho(|\mathbf{u}|^2/2 + \varepsilon + \Phi)$ and the equation of state.

5.3 The advection equation

To get basic insights into how computational techniques work, it is helpful to reduce the conservation law of Equation 5.10 to its bare bones. What does it say, fundamentally? If we work in one dimension (x), the equation for an arbitrary, conserved quantity Q reads

$$\frac{\partial Q}{\partial t} + \frac{\partial}{\partial x} \mathcal{F}(Q) = S(Q). \quad (5.12)$$

Let's imagine a simple situation where there are no source terms ($S = 0$) and where pressure plays no role in the fluxes (e.g., because it is constant everywhere). The fluxes of density, momentum, and energy are then simply that quantity times the velocity, $\mathcal{F} = Qu$, where $u = u_x$ is the velocity in the one dimension of our problem. The conservation law reduces to

$$\frac{\partial Q}{\partial t} + \frac{\partial (Qu)}{\partial x} = \frac{\partial Q}{\partial t} + u \frac{\partial Q}{\partial x} + Q \frac{\partial u}{\partial x} = 0. \quad (5.13)$$

To further simplify the problem, let's assume that the velocity is uniform everywhere, $\partial u/\partial x = 0$, which results in the **advection equation**,

$$\boxed{\frac{\partial Q}{\partial t} + u \frac{\partial Q}{\partial x} = 0} \quad (5.14)$$

In this system, the quantity Q is simply shifted to the positive x -direction with a velocity u ; we say that Q is “advected” with the fluid. If this property of the equation is not obvious, you can insert the general solution $Q = Q(y)$ with $y \equiv x - ut$, which obeys the equation:

$$\frac{\partial Q}{\partial y} \frac{\partial y}{\partial t} + u \frac{\partial Q}{\partial y} \frac{\partial y}{\partial x} = 0 \quad \implies \quad \frac{\partial y}{\partial t} + u \frac{\partial y}{\partial x} = 0 \quad \implies \quad -u + u = 0. \quad (5.15)$$

The lines along which the solution is constant (because $x - ut$ is constant) are called **characteristics**. The advection equation presents an excellent test case for any numerical scheme because it is similar to the (more complicated) equations we eventually wish to solve. To see this more generally, we can write the Euler equations in matrix form, for example by expressing $\mathcal{F}(\mathbf{U})$ as a linear operator. The equations can then be reduced to an eigenvalue problem that has the form of the advection equation, with the eigenvalues providing the characteristic speeds of advection. This framework is worked out in §B.2.

5.4 A very basic taxonomy of partial differential equations

Given the abundance of “ ∂ ” symbols in this text, it is probably clear by now that we are dealing with partial differential equations (PDEs). They come in a few flavors, namely, hyperbolic, parabolic, and elliptic. Since solving each type will demand different numerical methods, we should briefly contemplate the type of equation we are trying to tackle. Unfortunately, the exact mathematical definitions of PDE types are often less than helpful.⁶

The Euler equations are **hyperbolic** PDEs,⁷ which physically implies that the propagation speeds in the fluid system are finite. This is a profound property of the equations that our computational methods will need to respect. Another important characteristic of the Euler equations is that they are **non-linear**, which means that the properties of the initial conditions may not be conserved. For example, smooth initial conditions can generate discontinuous shocks (§7.1).

While the Euler equations are hyperbolic, additional terms (e.g., due to viscosity or conduction) can take on a dissipative nature. Mathematically, the dissipation arises because the temporal and spatial derivatives have different orders. One example of such a **parabolic** PDE is the heat equation,

$$\frac{\partial T}{\partial t} = \kappa \frac{\partial^2 T}{\partial x^2}, \quad (5.16)$$

where κ is the thermal diffusivity. If we inserted a wave solution like in §4.2, we would find an exponentially decaying term, meaning that a perturbation in temperature would be smoothed out by conduction rather than being advected. The larger κ , the faster a perturbation in $T(x, t)$ dissipates.

⁶For example, hyperbolic PDEs are those where “the Cauchy problem can be locally solved for arbitrary initial data along any non-characteristic hypersurfaces,” according to Wikipedia. A more useful taxonomy is to consider the coefficients of the second-order derivatives. For the purposes of hydrodynamics, the only terms that matter are $a \partial^2 Q / \partial t^2 + b \partial^2 Q / \partial x^2 + \text{lower order derivatives} = 0$. If $ab < 0$, the equation is hyperbolic; think, for example of the wave equation, where $a = 1$ and $b = -c_s^2$. If $ab > 0$, the equation is elliptic. This category includes the Poisson equation, $\nabla^2 Q = f$, where there are (positive) spatial derivatives in multiple dimensions. The derivatives do not need to be in space and time necessarily. Finally, parabolic equations arise for $ab = 0$, which happens if only one derivative appears in second order.

⁷At least if $\mathbf{A} = \partial \mathcal{F} / \partial \mathbf{U}$ is diagonalizable and all its eigenvalues are real. In §B.2, we show how this condition translates into finite propagation speeds of information.




Figure 8: Different ways to subdivide a computational space onto a grid. In the so-called finite-difference approach (top), we envision our grid to contain the values of fluid quantities at evenly spaced points $i \times \Delta x$. Alternatively, in the finite-volume approach (bottom), we imagine that each point represents a cell with walls at positions $i \pm 1/2$. The grid values now represent the average fluid quantities in each cell. Figure adapted from [Zingale](#).

5.5 Computing in space and time: finite difference vs. finite volume

When solving the Eulerian equations, we need to discretize space into a finite number of elements. We define a cell size and timestep such that

$$\begin{aligned} x_i &= x_0 + i \times \Delta x \\ t^n &= t_0 + n \times \Delta t \end{aligned} \quad (5.17)$$

which can, of course, be extended to multiple spatial dimensions with multiple indices. We use subscripts for position and superscripts for time so that we can combine them to write the fluid state at a particular cell and time as \mathbf{V}_i^n , one timestep into the future as \mathbf{V}_i^{n+1} , and one cell to the right as \mathbf{V}_{i+1}^n .

But what exactly do we mean by “the fluid state at some position?” We could simply imagine a grid of points in space that represent the fluid quantities at those points. This system is called a **finite-difference** approach (Figure 8), for reasons that will become apparent in §6. Here, we imagine that the true solution to the Euler equations could be interpolated between the discrete grid points.

However, we can also think about the division of space a little differently: what if the points x_i represent the solution in grid cells whose walls are at $i \pm 1/2$? In this **finite-volume** approach, we imagine that the values at x_i represent the average of the fluid quantities over the given cell, not the values at the mid-point (Figure 8). Instead of a smooth solution interpolated between grid points, we imagine a “histogram-like” solution with sharp jumps between cells. At first sight, this approach may seem counter-intuitive, but it will turn out to be extremely productive (§8).

5.6 The Courant-Friedrichs-Lewy condition

The Courant-Friedrichs-Lewy condition (or CFL condition, named for Courant et al. 1928) ensures that the numerical solution of advection-like hyperbolic PDEs cannot proceed faster than the physical speed of information transmission. In that sense, it ensures that **causality is obeyed**. Mathematically, the criterion can be written as

$$\alpha_{\text{cfl}} \equiv \left| c_{\text{max}} \frac{\Delta t}{\Delta x} \right| \leq 1 \quad (5.18)$$

where c_{max} is the maximum speed of any possible wave (§B.2). Computationally, the CFL condition limits the timestep Δt because it implies

$$\Delta t = \alpha_{\text{cfl}} \frac{\Delta x}{c_{\text{max}}} \leq \frac{\Delta x}{c_{\text{max}}}, \quad (5.19)$$

meaning that information cannot travel across a cell faster than the maximum wave speed. A Courant number smaller than unity provides a “buffer” preventing the numerical solution from reaching the fastest physically allowed wave speed. For the Euler equations, the maximum speed is $c_{\text{max}} = u + c_s$. When adding more terms and equations to the system, new, possibly faster waves arise (e.g., Alfvén waves in MHD, §10.6).

5.7 Computational error terms

Numerical solutions to any equation incur two fundamental types of errors. First, **round-off errors** occur due to the finite precision with which a number can be stored. These errors are ubiquitous but should not matter much as long as we use sufficiently large floating-point types. Moreover, we can reduce round-off errors by avoiding particularly susceptible operations such as adding large and small numbers.

In contrast, **truncation errors** are incurred “by design” due to the finite accuracy of an algorithm. For example, we will approximate derivatives as $\Delta Q/\Delta x$ over a finite distance Δx , which clearly neglects higher-order terms. The name “truncation error” refers to truncating a power series.

6 Computational hydro II: Finite-differencing schemes

In §5.3, we showed that a big part of what the Euler equations do is to move conserved fluid quantities through space. Thus, any numerical scheme must at least be able to solve the 1D advection equation (5.14), although success in this task does not automatically qualify the scheme for the full Euler equations. In this chapter, we explore various numerical schemes to solve the one-dimensional advection equation as a stand-in for general hyperbolic PDEs. See Zingale §4 and vdB §17-18 for a different presentation as well as additional material.

6.1 First-order schemes for the advection equation

We use V to denote a general, scalar fluid variable, although the same logic could apply to a vector of fluid variables, \mathbf{V} . We keep denoting the x -velocity as $u = u_x$ so that the advection equation reads

$$\frac{\partial V}{\partial t} + u \frac{\partial V}{\partial x} = 0. \quad (6.1)$$

What is the simplest numerical scheme we can think of? We can approximate any derivative using Taylor expansion, for example in space,

$$V_{i+1} = V_i + \left(\frac{\partial V}{\partial x}\right)_i \Delta x + \frac{1}{2} \left(\frac{\partial^2 V}{\partial x^2}\right)_i \Delta x^2 + \frac{1}{6} \left(\frac{\partial^3 V}{\partial x^3}\right)_i \Delta x^3 + \mathcal{O}(\Delta x^4) \quad (6.2)$$

$$V_{i-1} = V_i - \left(\frac{\partial V}{\partial x}\right)_i \Delta x + \frac{1}{2} \left(\frac{\partial^2 V}{\partial x^2}\right)_i \Delta x^2 - \frac{1}{6} \left(\frac{\partial^3 V}{\partial x^3}\right)_i \Delta x^3 + \mathcal{O}(\Delta x^4) \quad (6.3)$$

and similarly in time (note the alternating signs in the second equation). The simplest approximation to the derivatives is to keep only the first order of the expansion. Specifically, we take the forward derivatives in space and time,

$$\frac{V_i^{n+1} - V_i^n}{\Delta t} + u \frac{V_{i+1}^n - V_i^n}{\Delta x} = 0. \quad (6.4)$$

The forward derivative is visualized in Figure 9. Since the goal is to advance the solution V^n to the next timestep V^{n+1} , we write the **forward-time forward-space (FTFS) scheme** as

$$V_i^{n+1} = V_i^n - \alpha_{\text{cfl}}(V_{i+1}^n - V_i^n). \quad (6.5)$$

We have made the substitution $u\Delta t/\Delta x = \alpha_{\text{cfl}}$, which highlights the meaning of the Courant number: the higher α_{cfl} , the faster we are advancing the solution at each timestep. This shorthand will not be valid in general because the CFL number refers to the maximum speed with which any signal can propagate (§5.6). As long as we are studying the advection equation, that velocity is u , which is constant across the domain.

First-order schemes such as FTFS are called **explicit Euler schemes** — yes, another thing named after Euler! The term “explicit” means that the desired solution V^{n+1} appears only on the LHS but not on the RHS. The latter is also possible and leads to an **implicit scheme**, but that would mean having to solve the equation for V^{n+1} either analytically (before coding it up) or numerically.

Let’s test how our explicit Euler scheme performs in reality. We are free to choose any initial profile $V(x)$, which could represent a distribution of, say, density or temperature in a one-dimensional domain with periodic boundary conditions (think a tube that wraps around on itself). If all goes well, this initial profile should be transported to the right with speed u . We try two initial states:




Figure 9: Illustration of the backward, forward, and centered derivative approximations. The centered derivative is second-order accurate because the first-order errors from the backward and forward versions cancel. Figure from [Toro](#).

a smooth sine wave and a sharp tophat peak. Figure 10 shows the results: the FTFS scheme (left column) is highly unstable and diverges after few timesteps. We might suspect that our first-order derivatives are too inaccurate. Figure 9 suggests that the centered derivative is more accurate. Mathematically, this can be seen by subtracting Equation 6.3 from 6.2, which gives

$$V_{i+1} - V_{i-1} = 2\Delta x \left(\frac{\partial V}{\partial x} \right)_i + \frac{1}{3} \left(\frac{\partial^3 V}{\partial x^3} \right)_i \Delta x^3 + \mathcal{O}(\Delta x^4) \quad (6.6)$$

and thus

$$\left(\frac{\partial V}{\partial x} \right)_i = \frac{V_{i+1} - V_{i-1}}{2\Delta x} + \mathcal{O}(\Delta x^3). \quad (6.7)$$

Since the second-derivative terms cancel, we are left with third-order errors in V , and thus second-order errors in the derivative (where we divide by Δx). Using this new approximation for the finite difference between states, we obtain the **forward-time centered-space (FTCS) scheme**,

$$V_i^{n+1} = V_i^n - \frac{\alpha_{\text{cfl}}}{2} (V_{i+1}^n - V_{i-1}^n). \quad (6.8)$$

This scheme fares a little better in the sine-wave test (top center panel of Figure 10), but it still diverges. For completeness, we might also try the backward derivative, which yields the **forward-time backward-space (FTBS) scheme**,

$$V_i^{n+1} = V_i^n - \alpha_{\text{cfl}} (V_i^n - V_{i-1}^n). \quad (6.9)$$

Surprisingly, this scheme is much more stable! It passes both the sine-wave and tophat test (right column of Figure 10), although it diffuses the solution significantly. How is this possible? After all, FTBS is just the reverse of FTFS. However, our problem is inherently asymmetric because the solution is moving to the right with velocity u . Thus, the FTFS scheme uses values ahead of the solution (“downwind”) whereas the FTBS scheme uses values “upwind” of the cell in question.

This distinction is specific to the advection equation and will not hold in general problems where fluids can move in any direction. Nevertheless, the physics of this setup highlight a general problem with the flow of information, which is visualized in Figure 11. In the advection problem, the “domain of dependence” lies behind the solution (to the left), meaning that, physically speaking,

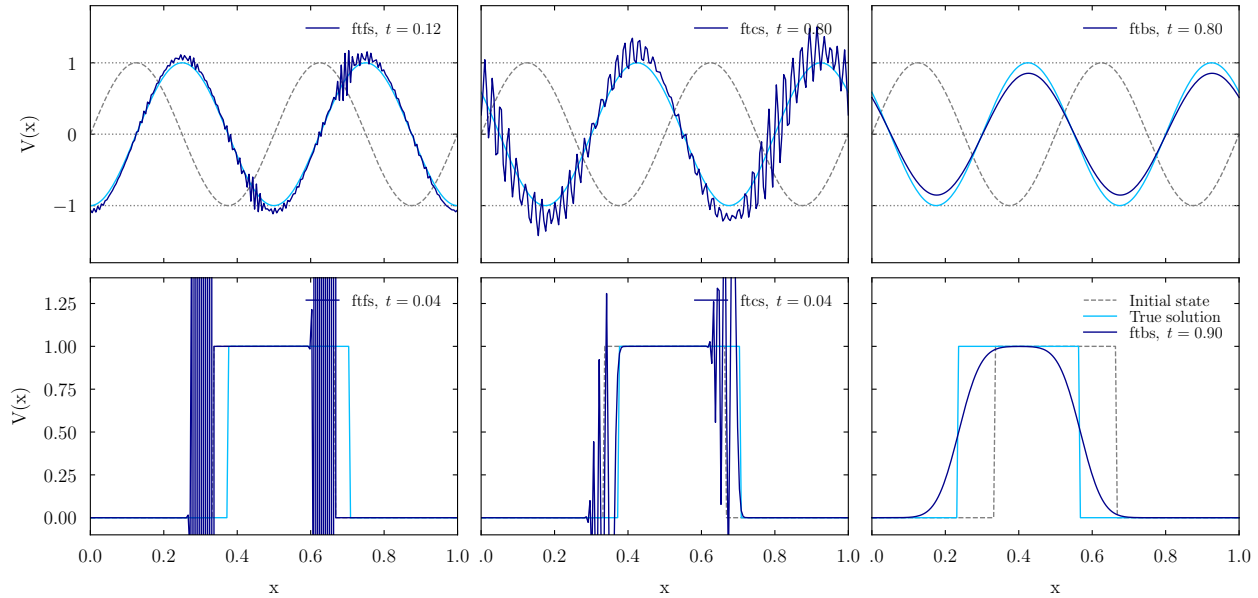


Figure 10: Advection test results for the sine wave (top) and tophat (bottom) setups. The dashed gray lines show the initial setup and the light-blue lines the correct, evolved solution. The dark blue lines show the numerical solution for the FTFS (left), FTCS (center), and FTBS (right) Euler schemes. All tests were run with $u = 1$ and $\alpha_{\text{eff}} = 0.5$. In the sine-wave test, the FTFS scheme (top left) diverges after a very short time. The FTCS scheme lasts longer (top center) but eventually also develops an instability that quickly explodes. The FTBS scheme, however, remains stable (top right). In the tophat test, we find similar results, but the unstable schemes diverge much faster due to the discontinuity. In the FTFS scheme (bottom left), the solution does not even move to the right; instead, oscillations rapidly spread to the left. The FTCS scheme shows similar oscillatory behavior (bottom center). The FTBS scheme (bottom right) runs stably to a much longer time, when the tophat has almost shifted through the domain once. However, the scheme adds numerical diffusion to both the sine and tophat shapes.

the information about the flow is moving left to right with a particular speed. The FTBS scheme performs best because its numerical domain of dependence overlaps with the physical one, whereas the FTCS and FTFS schemes rely on downwind information.

Finally, we might wonder what role the CFL number plays in our solution. Experimentally, we find that lower α_{eff} leads to more smoothing of the solution for the FTBS scheme. A curious case arises for $\alpha_{\text{eff}} = 1$: the scheme reproduces the known solution exactly! This only works for the advection equation though. A lower α_{eff} leads to a more dissipative solver because we are taking more timesteps for the same solution. The general takeaway is that decreasing α_{eff} leads to a more stable scheme, but also to more numerical dissipation.

6.2 Stability analysis

Let us now try to understand the experimental results from the previous section a little more systematically. First, we need to define a few important properties of numerical schemes:

- **Consistency:** A scheme is called consistent if all discrete, finite-difference operators converge to their continuous equivalents for infinitesimally small numerical steps. For example, the FTFS / FTCS / FTBS schemes (Equations 6.5, 6.8, and 6.9) are consistent because their derivative approximations tend to $\partial/\partial x$ as $\Delta x \rightarrow 0$ and to $\partial/\partial t$ as $\Delta t \rightarrow 0$.




Figure 11: Domains of dependence, meaning the region where the information for the next timestep originates from. The gray areas show the numerical domain of dependence: the cell itself and the cells to the left (FTBS, left), to the right (FTFS, right), or both (FTCS, center). The red area shows the physical domain of dependence for the advection equation: the information originates in an area of size $\alpha_{\text{cfl}}\Delta x$ to the left of the cell. Numerical schemes that use information from outside the physical domain are less stable. Based on a similar figure by Zingale.

- **Stability:** A scheme is called stable if noise (e.g., from the initial conditions or due to round-off error) does not grow.
- **Convergence:** A scheme is called convergent if its solution tends to the PDE's correct solution as $\Delta x \rightarrow 0$ and $\Delta t \rightarrow 0$.

The stability and convergence criteria seem related: if all derivative approximations tend to the correct values, the overall scheme should too. This is, indeed, the case, according to **Lax's equivalence theorem**: *for a well-posed initial value problem, the solution of a consistent scheme converges if and only if the scheme is stable*. Clearly, stability is key to predicting a scheme's success. But how can we know whether a scheme is stable, especially for general problems rather than for a single test case? One well-known technique to answer this question is **von Neumann stability analysis**. The idea is to insert a wave perturbation,⁸

$$V^n = A^n e^{ji\theta}, \quad (6.10)$$

and to check whether it grows in size. Note that $j \equiv \sqrt{-1}$ denotes the complex part of the solution in this case, since i already denotes the spatial index! We call the scheme stable if

$$\left| \frac{A^{n+1}}{A^n} \right| \leq 1. \quad (6.11)$$

Let's try this with our FTCS scheme. Inserting Equation 6.10 into 6.8, we get

$$\begin{aligned} A^{n+1} e^{ji\theta} &= A^n e^{ji\theta} - \frac{\alpha_{\text{cfl}}}{2} \left(A^n e^{j(i+1)\theta} - A^n e^{j(i-1)\theta} \right) \\ \implies \frac{A^{n+1}}{A^n} &= 1 - \frac{\alpha_{\text{cfl}}}{2} \left(e^{+j\theta} - e^{-j\theta} \right) = 1 - j \alpha_{\text{cfl}} \sin \theta. \end{aligned} \quad (6.12)$$

We are interested in whether this amplitude grows or decreases. There are different ways to write the result in the literature. For example, we can consider the square of the absolute value of the ratio, which for a complex number $a + jb$ is $\sqrt{a^2 + b^2}^2 = a^2 + b^2$, or in the case of FTCS,

$$\left| \frac{A^{n+1}}{A^n} \right|^2 = 1 + \alpha_{\text{cfl}}^2 \sin^2 \theta. \quad (6.13)$$

⁸We are cheating a little: we should insert the Fourier series over all possible perturbations with frequencies θ , $\Sigma_{\theta} A_{\theta}^n e^{ji\Delta x \theta}$. As long as the modes are independent, we can study a single mode with frequency θ to derive the correct behavior.

This result is bad news for the FTCS scheme! Regardless of the value of α_{cfl} , perturbations of (almost) any wavelength grow. When we apply the same method to the seemingly stable FTBS scheme, we get (Zingale §4.3 or vdB §18)

$$\left| \frac{A^{n+1}}{A^n} \right|^2 = 1 - 2\alpha_{\text{cfl}}(1 - \alpha_{\text{cfl}})(1 - \cos \theta), \quad (6.14)$$

which means that the scheme is stable if $0 \leq \alpha_{\text{cfl}} \leq 1$. This result is specific to the advection equation with positive u ; if we reversed the direction of u , the FTBS scheme would become an FTFS scheme. Mathematically, α_{cfl} would be negative in the equation above, meaning that all perturbations grow (as evidenced by our numerical experiments).

6.3 Understanding numerical diffusion

Having mathematically predicted which schemes are stable, can we go one step further and understand why the (stable) FTBS scheme diffuses the solution? We can gain insight from comparing the exact, Taylor-expanded solution (Equation 6.2) and our numerical approximation. The higher-order terms that we omitted in our numerical solution will constitute the truncation error incurred by the scheme. For the FTBS scheme (Equation 6.9), the approximations are

$$\frac{V_i^n - V_{i-1}^n}{\Delta x} = \frac{\partial V}{\partial x} - \frac{\Delta x}{2} \frac{\partial^2 V}{\partial x^2} + \mathcal{O}(\Delta x^2) \quad (6.15)$$

and, minding the signs in the Taylor expansion toward V_{i-1} (Equation 6.3),

$$\frac{V_i^{n+1} - V_i^n}{\Delta t} = \frac{\partial V}{\partial t} + \frac{\Delta t}{2} \frac{\partial^2 V}{\partial t^2} + \mathcal{O}(\Delta t^2). \quad (6.16)$$

Thus, instead of the advection equation, the FTBS scheme actually solves the equation

$$\frac{\partial V}{\partial t} + \frac{\Delta t}{2} \frac{\partial^2 V}{\partial t^2} + u \left[\frac{\partial V}{\partial x} - \frac{\Delta x}{2} \frac{\partial^2 V}{\partial x^2} \right] + \mathcal{O}(\Delta x^2) + \mathcal{O}(\Delta t^2) = 0. \quad (6.17)$$

Keeping terms to second order and using the advection equation to convert $\partial^2 V / \partial t^2 = u^2 \partial^2 V / \partial x^2$, the equation we are solving is

$$\begin{aligned} \frac{\partial V}{\partial t} + u \frac{\partial V}{\partial x} &= -\frac{\Delta t}{2} \frac{\partial^2 V}{\partial t^2} + \frac{u \Delta x}{2} \frac{\partial^2 V}{\partial x^2} \\ &= \left(-\frac{u^2 \Delta t}{2} + \frac{u \Delta x}{2} \right) \frac{\partial^2 V}{\partial x^2} \\ &= \frac{u \Delta x (1 - \alpha_{\text{cfl}})}{2} \frac{\partial^2 V}{\partial x^2}. \end{aligned} \quad (6.18)$$

The term on the right corresponds to a parabolic equation such as the heat equation (§5.4). Instead of the advection equation, we are actually solving a so-called **advection-diffusion equation** with a diffusion coefficient $\kappa = u \Delta x (1 - \alpha_{\text{cfl}}) / 2$. This finding immediately explains the experimentally found trends with α_{cfl} : when $\alpha_{\text{cfl}} = 1$, the second-order terms cancel, which is why the scheme becomes so accurate. The smaller α_{cfl} , the stronger the diffusion term becomes.

Our result also explains the instability of the FTFS scheme, which corresponds to $u \rightarrow -u$ (since the forward/backward differencing is with respect to the fluid motion). In the FTFS case, $\kappa \rightarrow -\kappa$, which is always negative and produces an unstable equation.




Figure 12: Same as Figure 10, but for higher-order schemes (with $\alpha_{\text{eff}} = 0.5$). While the Lax-Friedrichs scheme (left) is more diffusive, it is also more stable in the tophat test. Conversely, the Lax-Wendroff scheme (center, diffusion term suppressed by a factor of α_{eff}^2) performs very well in the sine-wave test but develops oscillations at sharp boundaries. The higher-order Fromm scheme (right) does slightly better at maintaining sharp edges but also exhibits numerical artifacts.

Note, once again, that these conclusions hold only for the advection equation. Nevertheless, the general principle applies to virtually all finite-differencing schemes: some of the second-order truncation error terms correspond to diffusion-like physics, leading to a blurring of sharp boundaries (see vdB §18 and Zingale §4.3 for additional details).

6.4 Higher-order finite-difference schemes

While the success of the FTBS scheme is encouraging, it is clearly not generalizable: if the fluid flowed in the $-u$ direction, we would revert to the unstable FTFS scheme. It stands to reason that any generally applicable scheme needs to be symmetric in space (although not in time). Such schemes can be constructed by making the FTCS scheme more stable.

For example, the **Lax-Friedrichs** and **Lax-Wendroff** methods add an **artificial viscosity** term to the equations. This term is essentially the term on the RHS of the advection-diffusion equation (6.18), with the second derivative based on Taylor expansions:

$$V_i^{n+1} = V_i^n - \frac{\alpha_{\text{eff}}}{2}(V_{i+1}^n - V_{i-1}^n) + \frac{\beta}{2}(V_{i+1}^n - 2V_i^n + V_{i-1}^n), \quad (6.19)$$

where $\beta = 1$ gives the Lax-Friedrichs scheme and $\beta = \alpha_{\text{eff}}^2$ the Lax-Wendroff scheme. Their performance in the sine-wave and tophat tests is compared in Figure 12. Once again, we face a trade-off between diffusivity and stability: the diffusion term is larger in the Lax-Friedrichs scheme (because $\alpha_{\text{eff}} = 0.5$), leading to a less accurate sine wave but no instabilities in the tophat test. Conversely, the Lax-Wendroff scheme is not diffusive enough to suppress artifacts at sharp boundaries. We can further extend our finite-difference schemes to higher orders, with examples

including the **Beam-Warming scheme**,

$$V_i^{n+1} = V_i^n - \frac{\alpha_{\text{cfl}}}{2}(3V_i^n - 4V_{i-1}^n + V_{i-2}^n) + \frac{\alpha_{\text{cfl}}^2}{2}(V_i^n - 2V_{i-1}^n + V_{i-2}^n) \quad (6.20)$$

and the **Fromm scheme**,

$$V_i^{n+1} = V_i^n - \frac{\alpha_{\text{cfl}}}{4}(V_{i+1}^n + 3V_i^n - 5V_{i-1}^n + V_{i-2}^n) + \frac{\alpha_{\text{cfl}}^2}{4}(V_{i+1}^n - V_i^n - V_{i-1}^n + V_{i-2}^n). \quad (6.21)$$

Note that these schemes are asymmetric and thus not generally applicable. Moreover, it is not clear that they provide solutions that are qualitatively better than the Lax schemes (Figure 12).

We conclude that it is impossible to find a finite-differencing scheme that is both stable and avoids significant numerical diffusion. Moreover, Figure 12 highlights two other, fundamental issues: there is no guarantee that a positive solution stays positive, or that the total V in the domain is conserved! Both present a real problem if we are, for example, advecting mass (density) or momentum; negative values and non-conservation are clearly unphysical. In fact, our experiences with finite-differencing schemes are representative of a more fundamental, underlying result known as **Godunov's theorem**, which states that *there are no linear higher-order schemes for advection that retain a positive solution*. To get around this issue, we will try an entirely different computational approach in §8.

7 Shocks

In the previous chapter, we considered a sharp tophat perturbation without ever motivating why such discontinuities might be of interest in fluid dynamics. In this chapter, we introduce shocks, a particularly important type of discontinuity that is commonplace in astrophysical applications.

7.1 Why shocks arise: Burgers' equation

We motivated the advection equation based on the general form of conservation laws, but we did not specify the physical nature of the advected quantity Q . Let's stay in 1D where $u = u_x$ but return to a more physical system. We will simplify the Euler equations by assuming constant pressure everywhere, but not constant density or velocity. The continuity equation in 1D then reads

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x}(\rho u) = 0 \quad \implies \quad \frac{\partial \rho}{\partial t} + \rho \frac{\partial u}{\partial x} + u \frac{\partial \rho}{\partial x} = 0 \quad (7.1)$$

and the momentum equation (in conservation form) becomes

$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial}{\partial x}(\rho u^2 + P) = 0 \quad \implies \quad u \frac{\partial \rho}{\partial t} + \rho \frac{\partial u}{\partial t} + \rho(2u) \frac{\partial u}{\partial x} + u^2 \frac{\partial \rho}{\partial x} = 0. \quad (7.2)$$

The two equations now contain many of the same terms. We multiply the continuity equation 7.1 by u and subtract it from 7.2 to find **Burgers' equation**,

$$\boxed{\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = 0} \quad (7.3)$$

This equation is another form of the advection equation, except that the advected quantity is now the velocity itself! Burgers' equation also directly follows from the momentum equation with $\nabla P = \nabla \Phi = 0$. In this special case, we can solve the momentum equation without any reference to the continuity equation because changes in density have no impact on the velocity in the absence of pressure gradients. The second term in Burgers' equation corresponds to the $\mathbf{u} \cdot \nabla \mathbf{u}$ term in the momentum equation, which makes the equation **non-linear**, a fundamental feature of the Euler equations. When we considered sound waves, that term was assumed to be small for a small velocity perturbation, which kept the equations linear. Now, we no longer assume that u is small.

While Burgers' equation can be solved analytically using the so-called method of characteristics (Figure 14), it is easier to compute its solution numerically. We recognize that the equation can be written as a conservation law,

$$\frac{\partial u}{\partial t} + \frac{\partial(u^2/2)}{\partial x} = 0, \quad (7.4)$$

where the flux of velocity has simplified to $u^2/2$ because we dropped the pressure term (compared to Equation 5.11). We discretize this expression in a fashion similar to our treatment of the advection equation. By now, we know that we should use upwind derivatives, and specifically the backwards derivative if we again constrain ourselves to positive initial $u(x)$,

$$u_i^{n+1} = u_i^n - \frac{\Delta t}{2\Delta x} \left[(u_i^n)^2 - (u_{i-1}^n)^2 \right], \quad (7.5)$$

where we have used the spatial and temporal indices i and n introduced in §5.5.⁹ The solution is shown in Figure 13. Regardless of whether we use a smooth (Gaussian) or sharp (tophat) initial

⁹We could also try to discretize the equation using only linear derivatives, $u_i^{n+1} = u_i^n - \Delta t/\Delta x u_i^n (u_i^n - u_{i-1}^n)$, but this solution turns out to be less accurate than the conservation form.




Figure 13: Evolution of Burgers’ equation for a Gaussian initial velocity profile (left) and for a tophat (right). In either case, the initial profile steepens into a shock wave and leaves behind a rarefaction zone. The solution was computed using Equation 7.5 and a timestep limited by $\alpha_{\text{eff}} = 0.5$.




Figure 14: Solution of the Burgers’ equation by the method of characteristics. The red lines indicate so-called characteristics where the velocity remains at its initial value. Where multiple characteristics intersect, a shock forms; where they diverge, we get a rarefaction wave. Figure from vdB.

velocity profile, those parts of the solution that have higher velocity move to the right faster and pile up, forming a shock discontinuity. At the back side of the flow, we encounter a **rarefaction wave**, i.e., a region where the density decreases.

Figure 14 shows a schematic of the analytical solution using so-called characteristics, lines of a constant fluid state that move with velocity $u(x)$. Where the characteristics pile up, a shock forms; where they diverge, we get rarefaction. A similar solution also applies to strong sound waves: in the peaks, the fluid speed is higher than in the troughs, which eventually leads to the wave steepening into a shock.¹⁰ This effect is minor in many everyday situations because the sound waves are weak compared to the background fluid state, but this example, as well as the solutions to Burgers’ equation in general, highlight that shocks are a natural consequence of the Euler equations.

7.2 What is a shock?

In a nutshell, shocks are interfaces where supersonic flows (faster than sound) meet obstacles, which could be a solid boundary or subsonically flowing gas (slower than sound). As we saw in §5, the Euler equations do not allow information to propagate faster than the speed of sound, c_s . Thus, the information that an obstacle is in the way of a flow cannot propagate backwards (“upstream”) fast

¹⁰It is sometimes stated that the steepening of sound waves arises because the sound speed is higher in the peaks than in the troughs. While that is true for an ideal gas, it is not the point: the fluid speed itself is faster, meaning the peaks “run into” the slower troughs. This effect can be demonstrated for an isothermal equation of state where the sound speed is manifestly constant.




Figure 15: Example of subsonic and supersonic flow around airplanes. The flow around the subsonically moving plane (left) can adjust to form steady streamlines because information about the obstacle can propagate faster than the obstacle itself. In the supersonic example on the right, the fluid cannot adjust ahead of the plane and smashes into the air that is being dragged along with the obstacle. The shock slows down the air so that it can subsonically flow around the plane. Figure adapted from Shu.




Figure 16: A shock front moving to the right in the frame of the fluid, converting the fluid state \mathbf{V}_1 to \mathbf{V}_2 . When we zoom in (left), the shock has a finite thickness Δx that corresponds to a few mean free paths. Typically, this scale is extremely small compared to the size of the system L so that the shock appears as an infinitely sharp discontinuity when we zoom out (right). Figure from Shu.

enough to let the fluid know that it should find a path of less resistance. Instead, the supersonic fluid smashes into the obstacle without warning. This situation is illustrated with airplanes in Figure 15. We describe supersonic flows by their **Mach number**,

$$\mathcal{M} \equiv \frac{u}{c_s} \tag{7.6}$$

where u is the supersonic flow speed. Note that just accelerating fluid particles to a velocity $u > c_s$ is not sufficient to create a shock. In the frame of the fluid particles, they are at rest and nothing special happens. The supersonic nature of the flow only materializes when compared to the restframe of the obstacle. Thus, the defining characteristic of a shock is the Mach number in the rest frame of the shock.

The situation at the shock front clearly violates the fluid picture laid out in §1.3 because the region across which the fluid state changes becomes infinitesimally small, incompatible with the idea that changes should happen over a scale much larger than our imaginary fluid elements. In reality, a shock does have a thickness across which the “upstream” fluid state (which we will call \mathbf{V}_1) changes to the “downstream,” or “post-shock,” fluid state \mathbf{V}_2 (Figure 16). However, the width of the shock is only a few mean free paths, λ_{mfp} , just enough for the fluid particles to smash into each other. The size L of our fluid problem should, by construction, be much larger than the mean free path, $L \gg \lambda_{\text{mfp}}$ (§1.3). Thus, we can think of shock fronts as infinitesimally thin discontinuities (although not all discontinuities are shocks; we will discuss other types in §8.2).

Shocks represent large, non-linear perturbations in the fluid that we cannot describe with linear approximations as we did for sound waves. Besides violating the fluid picture, all shocks are viscous. Thus, we abandon any hope of describing the particle-level physics inside the shock front from first




Figure 17: Mach cones in theory and practice. As the flow velocity exceeds the sound speed, information propagates only behind the source of the perturbation (left). The subsequent wavefronts line up along two lines that deviate from the flow direction by the Mach angle, ϕ . The Schlieren image on the right shows a fighter jet flying supersonically, with a Mach cone originating from its tip. Figures adapted from Wikipedia (left) and NASA (right).

principles. Instead, we will write down conservation laws that are obeyed across the shock front regardless of the microphysics. One quantity that is not conserved, however, is entropy: shocks are **non-reversible** phenomena that increase the entropy of the fluid.

In Figure 16, we assumed that the velocity u is perpendicular to the shock front. This is not true in general: **oblique shocks** propagate at an angle compared to the fluid velocity. For example, the flow around the nose of a supersonic airplane creates oblique shocks whose angle marks a **Mach cone**, that is, the area to which disturbances from the flow can propagate at the sound speed. This principle is illustrated in Figure 17. By considering a triangle formed by the distance traveled by the source of the perturbation and by sound waves originating from it, we can see that $\sin \phi = c_s \Delta t / (u \Delta t)$, meaning that the **Mach angle** is

$$\phi = \sin^{-1} \left(\frac{1}{\mathcal{M}} \right) \quad (7.7)$$

This angle approaches 90° for $\mathcal{M} \rightarrow 1$, which would correspond to the flat front in the center image of the left panel of Figure 17. For higher \mathcal{M} , the angle sharpens. The relatively shallow angle of the Mach cone around the jet in the right panel of Figure 17 implies that the plane is flying at only mildly supersonic speed.

7.3 The Rankine-Hugoniot conditions

As discussed above, we will not even attempt to understand exactly what happens inside the shock front. Instead, we ask how the downstream (post-shock) fluid state \mathbf{V}_2 relates to the upstream (pre-shock) fluid state \mathbf{V}_1 given a particular Mach number. The answer to this question is contained in the conservation-law form of the fluid equations (§5.2). Basically, the conservative fluxes have to be the same on both sides of the shock to not violate mass, momentum, and energy conservation. We can prove this statement mathematically by integrating a 1D conservation law for a quantity Q across the shock front,

$$\frac{\partial Q}{\partial t} + \frac{\partial \mathcal{F}(Q)}{\partial x} = 0 \quad \implies \quad \frac{\partial}{\partial t} \int Q dx + \mathcal{F}(Q_2) - \mathcal{F}(Q_1) = 0. \quad (7.8)$$

Since Q does not accumulate in the shock over time, the first term is zero, and thus $\mathcal{F}(Q_2) = \mathcal{F}(Q_1)$. Here, we are working in the frame of the shock, meaning that u is the fluid velocity relative to the

shock front. Moreover, we are assuming that u is in the direction of the shock front. For oblique shocks, additional trigonometric terms appear, but they do not fundamentally alter the nature of the physics.

All that is left is to insert the conserved fluxes, namely, ρu for density, $\rho u^2 + P$ for momentum, and $(E + P)u$ for the total energy. We expand the definition of total energy, $E = \rho(u^2/2 + \varepsilon + \Phi)$, and omit the gravitational Φ term because any gravitational potential will be smooth across the shock. Moreover, we assume that the shock is adiabatic, i.e., that it passes too fast for significant heating or cooling to occur. We divide the energy condition by the constant ρu for simplicity and thus obtain the **Rankine-Hugoniot conditions**,

Rankine-Hugoniot conditions

$$\rho_1 u_1 = \rho_2 u_2$$

$$\rho_1 u_1^2 + P_1 = \rho_2 u_2^2 + P_2$$

$$\frac{u_1^2}{2} + \varepsilon_1 + \frac{P_1}{\rho_1} = \frac{u_2^2}{2} + \varepsilon_2 + \frac{P_2}{\rho_2}$$

What is the physical meaning of these equalities? One striking property is that they are, by construction, symmetric, and thus do not impose a direction on the shock. However, it can be shown that a shock that increases density also increases entropy, whereas a de-compressing “rarefaction” shock would decrease entropy. Thus, nature allows **only compressive shocks**, meaning $\rho_2 > \rho_1$ in all cases.

With this limitation in mind, the first equation enforces mass conservation and means that the fluid slows down in the shock, $u_2 < u_1$. This makes intuitive sense, given that a supersonic flow hits a stationary or slower obstacle. The momentum equation allows the conversion between ram pressure, ρu^2 , and thermal pressure. Since $\rho_2 > \rho_1$ and ρu is constant, the post-shock thermal pressure is always larger, $P_2 > P_1$. The slowing down of the fluid is achieved by increasing the average velocity of the particles. Finally, we note that the conserved total energy is identical to the Bernoulli constant (Equation 3.18), which is not surprising given that both are based on the concept of energy conservation.

7.4 Ideal gas shocks

While the Rankine-Hugoniot conditions are appealing in their simplicity, they are not terribly practical because they do not tell us how much the density, pressure, and temperature of the gas change across a shock. Assuming an ideal gas equation of state, $P = \rho\varepsilon(\gamma - 1)$, we can write the ratios of the pre- and post-shock quantities as a function of only the Mach number and the adiabatic index. Here, we define the Mach number of a shock as $\mathcal{M}_1 \equiv u_1/c_s$, where u_1 is the pre-shock fluid velocity in the frame of the shock front. The algebra is tedious and not particularly enlightening (e.g., §89 in Landau & Lifshitz 1987); we simply quote the result:




Figure 18: Evolution of the fluid properties across shocks in an ideal gas with $\gamma = 5/3$ (dark blue) and $\gamma = 7/5$ (light blue). The post-shock density (left) approaches a fixed compression ratio of 4 or 6 for the two γ values, whereas the pressure and temperature can rise by an arbitrary factor depending on the Mach number.

Ideal gas shock conditions

$$\frac{\rho_2}{\rho_1} = \frac{u_1}{u_2} = \frac{(\gamma + 1)\mathcal{M}_1^2}{(\gamma - 1)\mathcal{M}_1^2 + 2}$$

$$\frac{P_2}{P_1} = \frac{2\gamma\mathcal{M}_1^2 - \gamma + 1}{\gamma + 1}$$

$$\frac{T_2}{T_1} = \frac{(2\gamma\mathcal{M}_1^2 - \gamma + 1)([\gamma - 1]\mathcal{M}_1^2 + 2)}{(\gamma + 1)^2\mathcal{M}_1^2}$$

Figure 18 shows plots of these ratios for shocks in ideal fluids with $\gamma = 5/3$ and $7/5$. Given that $\mathcal{M}_1 > 1$ by the definition of a shock, we reaffirm that **shocks always decelerate, compress, pressurize, and heat** the fluid. As expected, all expressions reduce to unity for $\mathcal{M}_1 \rightarrow 1$. While the Rankine-Hugoniot conditions assumed that the shock is adiabatic, we once again emphasize that this does not mean that entropy is conserved! The constant K in the adiabatic equation of state, $P = K\rho^\gamma$, will not be conserved across shocks.

One unclear aspect of the post-shock flow is its Mach number. Intuitively, it must be smaller than unity because otherwise a second shock would form right away. From the expressions above, $u_2 < u_1$ and $c_{s,2} > c_{s,1}$ since $c_s \propto \sqrt{T}$ in an ideal gas, but that does not imply that $\mathcal{M}_2 < 1$. The full result that can be (tediously) derived from the ideal gas shock conditions is the post-shock Mach number in units of the pre-shock Mach number (CC §7.2),

$$\mathcal{M}_2 = \sqrt{\frac{(\gamma - 1)\mathcal{M}_1^2 + 2}{2\gamma\mathcal{M}_1^2 - \gamma + 1}}. \quad (7.9)$$

Though not obvious, the important conclusion from this expression is that $\mathcal{M}_2 < 1$ regardless of \mathcal{M}_1 , meaning that the **post-shock flow is always subsonic**.

The curves in Figure 18 seem to approach relatively simple shapes at high Mach numbers, meaning strong shocks. As $\mathcal{M}_1 \rightarrow \infty$, the post- and pre-shock properties approach ratios of

$$\frac{\rho_2}{\rho_1} \rightarrow \frac{\gamma + 1}{\gamma - 1} \quad \frac{P_2}{P_1} \rightarrow \frac{2\gamma\mathcal{M}_1^2}{\gamma + 1} \quad \frac{T_2}{T_1} \rightarrow \frac{2\gamma(\gamma - 1)\mathcal{M}_1^2}{(\gamma + 1)^2}. \quad (7.10)$$

For $\gamma = 5/3$, we get $\rho_2/\rho_1 \rightarrow 4$. Regardless of how strong a shock is, it cannot compress the gas arbitrarily! By contrast, the post-shock pressure rises as $P_2/P_1 \rightarrow 5/4\mathcal{M}_1^2$ for strong shocks, and the temperature as $T_2/T_1 \rightarrow 5/16\mathcal{M}_1^2$, albeit with a more complex behavior at moderate \mathcal{M}_1 . In this regime, we make an observation that is particularly relevant in astrophysics: for strong shocks, the post-shock temperature is a function of only the pre-shock velocity. This surprisingly simple behavior arises as follows. We note that the sound speed depends only on the temperature,

$$c_{s,1} = \sqrt{\gamma P_1/\rho_1} = \sqrt{\frac{\gamma k_B T_1}{\mu m_p}}. \quad (7.11)$$

This lets us convert the RHS of the high- \mathcal{M} temperature expression in Equation 7.10,

$$T_2 = \frac{5}{16} T_1 \mathcal{M}_1^2 = \frac{5}{16} T_1 \frac{u_1^2}{c_{s,1}^2} = \frac{5}{16} \frac{\mu m_p T_1 u_1^2}{\gamma k_B T_1} = \frac{3}{16} \frac{\mu m_p}{k_B} u_1^2 \approx 2800\text{K} \left(\frac{u_1}{10 \text{ km s}^{-1}} \right)^2, \quad (7.12)$$

where we have used $\mu \approx 1.22$ for a cosmological gas of hydrogen and helium (§1.4). Physically, this behavior arises because shocks convert parts of the kinetic energy of the incoming gas to thermal energy. In strong shocks, the kinetic energy of the incoming gas dominates, meaning that the thermal energy of the pre-shock gas is approximately determined by only u_1^2 .

Equation 7.12 offers a first-order estimate of the temperature of gas after passing an through an “accretion shock” around dark matter halos (a discontinuity where cold, infalling gas runs into the gas already in the halo). For example, the Milky Way halo has a circular velocity of about 200 km/s. Assuming that the gas reaches roughly this velocity at infall, we expect it to get shocked to $T_2 \approx 10^6$ K, which is close to the so-called “virial temperature” of the Milky Way halo.

7.5 Isothermal shocks

So far, we have treated shocks as adiabatic because a shock front is extremely thin, meaning that the brief amount of time it takes the fluid to travel through it leaves little time for cooling. After the shock, however, the gas is free to “cool away” some of the acquired thermal energy. Astrophysical gases can be very efficient at cooling, meaning that an isothermal shock might be a more realistic description. Here, we imagine that the gas travels for some “cooling length” after the shock before it settles back into the original temperature, $T_2 = T_1$.

In isothermal gas, we have $P \propto \rho$ and thus $c_s^2 = P/\rho = \text{const}$ (pre- and post-shock). Moreover, the first two Rankine-Hugoniot conditions are still true, while the third is not because it did not take cooling into account. We write $P = \rho c_s^2$ and convert the second condition to

$$\rho_1(u_1^2 + c_s^2) = \rho_2(u_2^2 + c_s^2) \quad (7.13)$$

and, using the first condition,

$$(u_2 - u_1)c_s^2 = (u_2 - u_1)u_1u_2 \quad \implies \quad c_s^2 = u_1u_2, \quad (7.14)$$

which tells us the much simpler **jump conditions for isothermal shocks**,

$$\boxed{\frac{\rho_2}{\rho_1} = \frac{P_2}{P_1} = \frac{u_1}{u_2} = \mathcal{M}_1^2} \quad (7.15)$$

Notably, there is no more limit on the compression due to the shock. By cooling away the internal energy, the gas can settle into a lower-pressure and higher-density post-shock state than it would have without cooling.

7.6 Supernova blast waves

One quintessential astrophysical example of shocks are the blast waves caused by exploding supernovae. These stellar explosions deposit about 10^{51} erg of energy almost instantaneously. In terms of kinetic energy, this number corresponds to one solar mass (the typical mass of the ejected material) moving at about 10,000 km/s! In reality, the energy is deposited as both kinetic and thermal energy, but that does not really matter for our purposes.

At first, the ejecta expands freely because its mass is greater than that of the surrounding material it has swept up. This **free expansion phase** ends roughly when the shell has swept up than its own mass, which slows down the shock front. The shock converts the kinetic energy into heat, leaving behind a hot, high-pressure, expanding supernova bubble. This is called the **Sedov-Taylor or adiabatic phase** because radiative cooling is too slow to change the progress of the shock significantly. Eventually, cooling does become important and the bubble radiates away its thermal energy. The remaining kinetic energy is still pushing more material outwards in the **snow plow phase**. At some point, the velocity of the shell becomes subsonic compared to the surrounding gas' sound speed. The shock dies out and the disturbance propagates like a sound wave. We can think of the three phases as regimes where we can neglect both the surrounding pressure and density (free expansion), the surrounding pressure but not density (Sedov-Taylor), and eventually neither (snow plow).

While the process may sound complicated, there is a beautiful solution for the Sedov-Taylor phase that depends only on the energy of the supernova, E_{SN} , the mass of the ejected material, M_e , and the density of the surrounding medium, ρ_1 (the notation highlights that this is the pre-shock density). Given the enormous input energy, the temperature (internal energy) of the surrounding gas can be neglected (along the lines of the discussion on strong shocks in §7.4).

The formal derivation of the shock properties is somewhat involved (e.g., CC §8.1.1), but there is a shortcut: the problem is self-similar! As we recall from §5.1, self-similarity means that we can freely set the physical mass, length, and time scales of any solution to the Euler equations. The Sedov-Taylor problem falls into this category because we have specified a characteristic energy, mass, and density. From these three inputs, we should be able to construct characteristic length and time scales that can adjust to any values of the input variables. From Equation 5.1, we recall that $\rho_0 = m_0/l_0^3$ and $E_0 = m_0 l_0^2/t_0^2$ (since E_{SN} has units of energy rather than energy per unit volume). To find l_0 as a function of ρ_0 , E_0 , and t_0 , we need to cancel the m_0 factors, so we experiment:

$$\frac{E_0}{\rho_0} = \frac{m_0 l_0^2}{t_0^2} \frac{l_0^3}{m_0} = \frac{l_0^5}{t_0^2} \implies \frac{E_0 t_0^2}{\rho_0} = l_0^5 \implies l_0 = \left(\frac{E_0 t_0^2}{\rho_0} \right)^{1/5}. \quad (7.16)$$

Whichever way we combine the natural scales, there is only one possible combination of ρ_1 , E_{SN} , and t that will give us units of length! Thus, whatever the detailed physics of the problem may be, the shock radius must follow the **Sedov-Taylor solution**,

$$R_s \propto \left(\frac{E_{\text{SN}}}{\rho_1} \right)^{1/5} t^{2/5} \quad (7.17)$$

Admittedly, we have no idea why this solution comes out the way it does, but it is nifty that the self-similar nature of the problem lets us circumvent lots of complicated math. We now get to decide on a convenient “unit system” that connects our self-similar units to the three input parameters. It is tempting to set E_0 to the supernova energy and ρ_0 to the surrounding density, but there is a more elegant solution: to set l_0 and t_0 to the shock radius and time where the Sedov-Taylor solution




Figure 19: The Sedov-Taylor solution of the self-similar blast wave problem in density (left), pressure (center), and velocity (right). The light-blue dashed line shows the analytical solution, and the dark blue line shows the numerical solution according to the ULULA code (§C), which matches well. The density jump does not quite reach the maximum value of 4 times the surrounding density (for $\gamma = 5/3$) because the shock has a finite Mach number. While the density approaches zero at the center, the pressure approaches a constant value far behind the shock front.

starts to become valid. This happens roughly when the free-expansion phase ends, meaning that the ejecta has swept up a mass about equal to its own. We thus define the characteristic length via

$$\frac{4\pi}{3}\rho_1 l_0^3 = M_e \quad \Longrightarrow \quad l_0 = \left(\frac{3M_e}{4\pi\rho_1}\right)^{1/3}. \quad (7.18)$$

The ejecta reaches this radius with a velocity defined by its kinetic energy,

$$E_{\text{SN}} = \frac{1}{2}M_e v_e^2 \quad \Longrightarrow \quad v_e = \sqrt{\frac{2E_{\text{SN}}}{M_e}}, \quad (7.19)$$

and thus at our desired starting time for the Sedov-Taylor phase,

$$t_0 = \frac{l_0}{v_e} = \left(\frac{3M_e}{4\pi\rho_1}\right)^{1/3} \left(\frac{2E_{\text{SN}}}{M_e}\right)^{-1/2} = \left(\frac{3}{4\pi}\right)^{1/3} 2^{-1/2}\rho_1^{-1/3} E_{\text{SN}}^{-1/2} M_e^{5/6}. \quad (7.20)$$

Finally, we set m_0 to the ejecta mass, which fixes the rest of our natural unit system:

$$m_0 = M_e \quad v_0 = \frac{l_0}{t_0} = v_e \quad \rho_0 = \frac{m_0}{l_0^3} = \frac{4\pi}{3}\rho_1 \quad E_0 = \frac{m_0 l_0^2}{t_0^2} = 2E_{\text{SN}}. \quad (7.21)$$

We can now replace the proportionality in Equation 7.17 with the length and time scales where the Sedov-Taylor phase begins,

$$R_s = l_0 \left(\frac{t}{t_0}\right)^{2/5} \quad (7.22)$$

There are no additional factors because we chose l_0 to match the shock radius at time t_0 when this solution becomes valid. If we had chosen a different unit system, e.g., by setting E_0 and ρ_0 to the supernova energy and surrounding density, we would obtain a similar equation but with additional non-dimensional factors. Figure 19 shows the self-similar solution for the radial profiles of density, pressure, and velocity.

Let's put in some numbers to make things more concrete. We assume that E_{SN} is the conventionally assumed supernova energy of 10^{51} erg, that the mass of the ejected material is $M_e = 1M_{\odot}$, and that ρ_1 corresponds to a density of one hydrogen atom per cc, $\rho_1 = m_p/\text{cm}^3 = 1.67 \times 10^{-24}$ g/cm³. The ejecta velocity is then

$$v_0 = v_e = \sqrt{\frac{2E_{\text{SN}}}{M_e}} = \sqrt{\frac{2 \times 10^{51} \text{ g cm}^2}{2 \times 10^{33} \text{ g s}^2}} = 10^9 \frac{\text{cm}}{\text{s}} = 10^4 \frac{\text{km}}{\text{s}}. \quad (7.23)$$

This velocity is extremely fast, which justifies our assumption of a strong shock where the temperature of the surrounding medium does not matter much. The Sedov-Taylor phase begins once the shock has reached a radius of

$$l_0 = \left(\frac{3M_e}{4\pi\rho_1}\right)^{1/3} = \left(\frac{3 \times 2 \times 10^{33} \text{ g}}{4\pi \times 1.67 \times 10^{-24} \text{ g cm}^3}\right)^{1/3} \approx 2.2 \text{ pc}, \quad (7.24)$$

which happens after a free-expansion phase that lasts about

$$t_0 = \frac{l_0}{v_0} = \frac{6.6 \times 10^{18} \text{ cm}}{10^9 \text{ cm/s}} \approx 200 \text{ yr}. \quad (7.25)$$

In reality, the criterion is that the explosion should have swept up much more mass than M_e , so the initial radius and timescale are only approximate. For example, for Tycho's supernova remnant from an explosion in 1572, it is not clear that the system has entered the Sedov-Taylor phase yet.

The Sedov-Taylor phase ends when the shell of ejecta becomes sub-sonic with respect to the surrounding material, that is, when the speed of the shock wave falls below the sound speed outside the hot bubble. This is roughly the same time when the pressure behind the shock wave becomes comparable to the outside pressure, or when the total thermal energy in the bubble equals the explosion energy (meaning that the initial kinetic energy has been dissipated, see CC §8.4). We can estimate the duration of the Sedov-Taylor phase by equating the shock speed,

$$v_s = \frac{dR_s}{dt} = \frac{2}{5} \frac{l_0}{t_0} \left(\frac{t}{t_0}\right)^{-3/5} = \frac{2}{5} v_0 \left(\frac{t}{t_0}\right)^{-3/5}, \quad (7.26)$$

with the sound speed outside the bubble,

$$c_{s,1} = \sqrt{\frac{\gamma P_1}{\rho_1}} = \sqrt{\frac{\gamma k_B T_1}{\mu m_p}}, \quad (7.27)$$

to get

$$\frac{2v_0}{5} \left(\frac{t_{\text{end}}}{t_0}\right)^{-3/5} = \sqrt{\frac{\gamma k_B T_1}{\mu m_p}} \implies t_{\text{end}} = t_0 \left(\frac{25\gamma k_B T_1 M_e}{8\mu m_p E_{\text{SN}}}\right)^{-5/6}. \quad (7.28)$$

Using $\mu = 1.22$ and the previous values for supernova energy and ejecta mass, we only need T_1 to find the time t . This dependence makes sense, since we now need to know how much pressure the surrounding gas provides to stop the shock wave (or, equivalently, what its sound speed is). We assume $T_1 = 10^4$ K to find

$$t_{\text{end}} \approx t_0 \times 4.2 \times 10^7 (T_1/\text{K})^{-5/6} \approx 2 \times 10^4 t_0 \approx 4 \times 10^6 \text{ yr}. \quad (7.29)$$

From Equation 7.22, we get

$$l_{\text{end}} \approx l_0 \times (2 \times 10^4)^{2/5} \approx 50 l_0 \approx 100 \text{ pc}. \quad (7.30)$$

Our estimates for t_{end} and l_{end} are somewhat optimistic because radiative cooling starts to become significant before the Sedov-Taylor phase is over, meaning that some of the thermal energy inside the bubble is lost. Moreover, the final bubble size is of the same order as the typical scale height of the gas in disk galaxies. This means that some of the bubble will push out of the disk. Between radiative losses and this “blowout” effect, only a few percent of the initial supernova energy is really transferred to the surrounding ISM.

On the other hand, the final size of the bubble highlights how effective supernovae are at messing up the ISM. The supernova rate in the MW is thought to be about $10^{-7}/\text{Myr}/\text{pc}^3$ (CC). When we multiply this number with t_{end} and the final volume of our bubble, we reach the conclusion that any patch in the ISM is, on average, part of about one supernova bubble! This high “filling factor” means that we need to consider the effects of overlapping bubbles when predicting the impacts of supernova feedback on the ISM.

8 Computational hydro III: Finite-volume schemes

In §6, we tried our hand at numerically solving simplified, Euler-like systems such as the advection equation. Unfortunately, our finite-differencing schemes ran into a number of fundamental issues: they were either unstable or highly diffusive, did not conserve advected quantities such as mass, and allowed the solution to become negative even if the initial conditions were strictly positive (as predicted by Godunov’s theorem). Moreover, in §7 we saw that the Euler equations inevitably lead to shocks and other discontinuities, which present particular challenges to finite-differencing schemes (as evidenced by the tophat advection tests in §6.1). To solve these issues, we return to the drawing board and consider a finite-volume approach, which seems to give similar results at first but turns out to be much more flexible in practice.

Throughout the section, we use $\mathbf{U} = (\rho, \rho u_x, E)$ to denote the vector of conserved fluid quantities in 1D and $\mathbf{V} = (\rho, u_x, P)$ for the primitive quantities. In terms of a numerical scheme, they are equivalent because we can always convert one into the other via the equation of state. It is also understood that the flux vector $\mathbf{F} = \mathbf{F}(\mathbf{U}) = \mathbf{F}(\mathbf{V})$ even if the dependence is not explicitly specified.

Most of the algorithms described in this chapter are implemented in ULULA, a lightweight, publicly available python code that is briefly introduced in §C. The purpose of ULULA is to experiment with how well different schemes perform on simple 1D and 2D hydro problems.

8.1 Godunov schemes

Let’s return to the different ways of discretizing space introduced in Figure 8. So far, we have thought of grid points as representing a solution at discrete locations; but what about the finite-volume approach? There, we are thinking of each grid point as a fluid cell, conceptually similar to the “fluid elements” on which we based our understanding of the Euler equations. One way to enforce conservation is to take the conservation-law form of the equations seriously on a cell-by-cell level. After all, a 1D conservation law,

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}(\mathbf{U})}{\partial x} = 0, \quad (8.1)$$

tells us that a change in $\mathbf{U}(x, t)$ must be accompanied by an inflow or outflow, meaning a differential flux into and out of the cell. The basic idea of the **finite-volume formulation** is to compute the **flux across cell boundaries** and to add the same amount of \mathbf{U} to one cell that is subtracted from another. Such a scheme will, by construction, preserve the total amount of \mathbf{U} to machine precision. Mathematically, we can derive a general expression by integrating the conservation law over one cell and one timestep. We use $i - 1/2$ and $i + 1/2$ to denote the positions of the left and right walls of cell i (Figure 8),

$$\int_{x_{i-1/2}}^{x_{i+1/2}} dx \int_{t_n}^{t_{n+1}} dt \left[\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} \right] = 0. \quad (8.2)$$

The integral over space of $\partial/\partial x$ is just the difference between the values at the lower and upper limits, and the same for time, so

$$\int_{x_{i-1/2}}^{x_{i+1/2}} dx \left[\mathbf{U}(x, t^{n+1}) - \mathbf{U}(x, t^n) \right] + \int_{t_n}^{t_{n+1}} dt \left[\mathbf{F}(x_{i+1/2}, t) - \mathbf{F}(x_{i-1/2}, t) \right] = 0. \quad (8.3)$$

We identify the spatial terms on the left as the average values in the cell at times t^n and t^{n+1} ,

$$\mathbf{U}_i^n \equiv \frac{1}{\Delta x} \int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{U}(x, t^n) dx \quad \mathbf{U}_i^{n+1} \equiv \frac{1}{\Delta x} \int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{U}(x, t^{n+1}) dx. \quad (8.4)$$

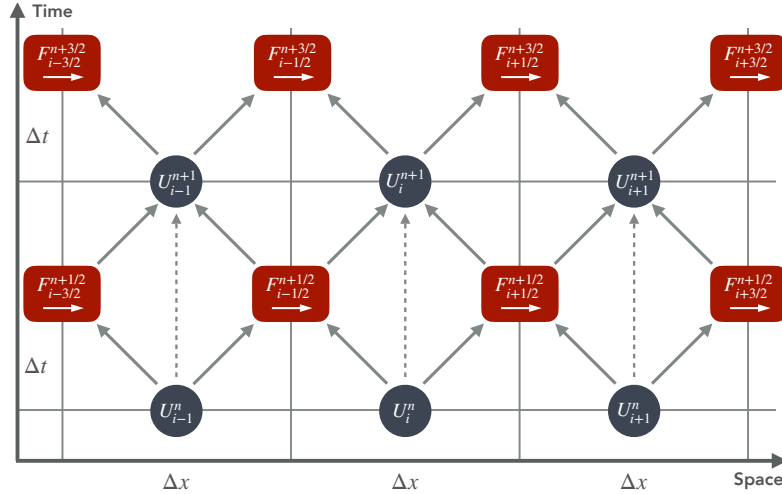


Figure 20: Schematic depiction of Godunov’s scheme (Equation 8.6). The vectors of conserved fluid variables are depicted as blue circles that “live” at cell centers i and timesteps n . From those values, we somehow compute fluxes across the cell interfaces, averaged over a timestep. Those fluxes, in turn, shift conserved quantities between cells, advancing the fluid state to timestep $n + 1$.

Staying with the notation of $\pm 1/2$ for the midpoints between two locations or timesteps, the integral over time would then quantify the time-averaged flux at half a timestep and at the left and right cell boundaries,

$$\mathcal{F}_{i-1/2}^{n+1/2} \equiv \frac{1}{\Delta t} \int_{t_n}^{t_{n+1}} \mathcal{F}(x_{i-1/2}, t) dt \quad \mathcal{F}_{i+1/2}^{n+1/2} \equiv \frac{1}{\Delta t} \int_{t_n}^{t_{n+1}} \mathcal{F}(x_{i+1/2}, t) dt. \quad (8.5)$$

When we insert these definitions into Equation 8.3 and rearrange, we obtain a **general scheme for conservative updates**,

$$\boxed{U_i^{n+1} = U_i^n + \frac{\Delta t}{\Delta x} \left(\mathcal{F}_{i-1/2}^{n+1/2} - \mathcal{F}_{i+1/2}^{n+1/2} \right)} \quad (8.6)$$

This scheme has a clear, physical interpretation: the conserved fluid quantities \mathbf{U} at the next timestep $n + 1$ are those at the previous timestep plus any flux that has entered from the left (the $i - 1/2$ side) minus any flux that has exited to the right (the $i + 1/2$ side). Figure 20 shows a schematic visualization of this process. Note that the left/right conventions rely on positive velocity going to the right.

Equation 8.6 is exact given the integral-based definitions in Equations 8.4 and 8.5. Any numerical (truncation) error we make when implementing the scheme will be in the averaged fluxes, but such errors will be inevitable because integrating over time exactly would correspond to taking an infinite number of timesteps. In summary, we have reduced the problem of solving the Euler equations to finding numerical approximations to the time-averaged fluxes across the “walls” between cells. The resulting algorithms are called **Godunov schemes** (after Godunov 1959).

We can see how Equation 8.6 reduces to finite-differencing schemes by applying it to the advection equation. The flux of an advected quantity Q is simply $\mathcal{F}(Q) = uQ$, if we assume that the velocity u is constant. As long as u is positive, the correct flux at the left cell wall is uQ_{i-1} and at the right wall uQ_i (Figure 21). At this point, we make no attempt to time-average the fluxes, i.e.,




Figure 21: Schematic depiction of Godunov’s scheme for the advection equation with the velocity u going to the right. The red shaded area show a conserved quantity Q , the yellow shaded area its fluxes across the interfaces during one timestep, which in this case are just $\Delta t \mathcal{F}_i = \Delta t u_{i-1} Q_{i-1}$. Here, the fluxes are determined by the left cell values because the fluid is going to the right. In general, we need to check at each interface in which direction the fluid is moving.

we assume $\mathcal{F}_i^{n+1/2} \approx \mathcal{F}_i^n$ to get

$$Q_i^{n+1} = Q_i^n + u \frac{\Delta t}{\Delta x} (Q_{i-1}^n - Q_i^n). \quad (8.7)$$

This is exactly our FTBS scheme from Equation 6.9! In this case, the scheme turns out to be exact (for $\alpha_{\text{eff}} = 1$) because the linear changes of the fluxes over a timestep fortuitously cancel. However, this will not be the case in general: just taking the value at one limit of the time integral (n instead of $n + 1/2$ in this case) will lead to first-order errors in the scheme. Time-averaging the fluxes will give us second-order (or higher) accuracy in time (§8.5).

8.2 Riemann solvers

For the advection equation, it was easy to guess the flux across each interface because we knew that the entire fluid is going in one direction. As a result, we picked the “right” scheme, namely, an upwind scheme (FTBS rather than FTCS or FTFS). But what should the flux be in the general situation, where we have two arbitrary, different fluid states at the left and right sides of an interface? This situation is called a **Riemann problem**, an initial-value problem with a discontinuity. Mathematically speaking, the problem is to find the solution $\mathbf{V}(x_0, t)$ to the initial conditions

$$\mathbf{V}(x, 0) = \begin{cases} \mathbf{V}_L & \forall x < x_0 \\ \mathbf{V}_R & \forall x > x_0 \end{cases} \quad (8.8)$$

where we will refer to \mathbf{V}_L and \mathbf{V}_R as the left and right states. Solving Riemann problems is the price we are paying for the finite-volume approach: the solution is not smooth at the cell walls because we are not interpolating between grid points. Thus, we need to solve one Riemann problem for each cell interface and timestep. This means that a good “Riemann solver” will be a key piece to implementing efficient, accurate Godunov schemes. We do not actually care about the full solution $\mathbf{V}(x_0, t)$ but only about the resulting flux across the discontinuity or cell interface, $\mathcal{F}_{i\pm 1/2}(\mathbf{V})$. Entire books have been written on the topic of Riemann solvers (most notably that of Toro). Here we will discuss the basic idea and one simple, approximate Riemann solver.

To solve a Riemann problem, we need to know about the physics of the underlying equations. This approach is fundamentally different from the finite-differencing schemes in §6, where we remained ignorant of the meaning of the PDEs we were trying to solve. Figure 22 shows the general




Figure 22: General structure of the Riemann problem. From the discontinuity at time 0, various waves, shocks, and discontinuities travel to the left and/or right at characteristic speeds λ_i . For a system of m equations, there will be m such speeds, which are ordered by their velocity from left to right. The Riemann state at the location of the discontinuity will be determined by the fluid state that these waves leave behind. Figure from Toro.

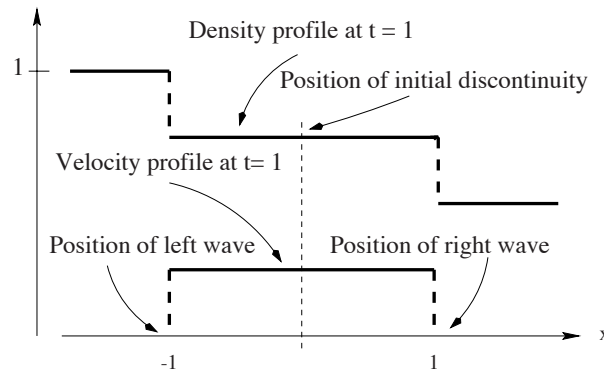


Figure 23: Solution of an example Riemann problem for the linearized fluid equations 8.9 and the initial conditions $\rho_L = 1$, $\rho_R = 1/2$, and $u_L = u_R = 0$ (a fluid at rest with an initial density jump). At $t = 1$, the fluid is moving to the right, creating a state of intermediate density that spreads in both directions at speed c_s . Figure adapted from Toro.

idea. We know that the Euler equations admit information-transmitting “waves” such as sound waves, shocks, and other discontinuities. The solution at the location of the initial discontinuity is determined by the various waves that travel left and right at certain speeds. This so-called “Riemann fan” of possible waves contains the physics of the problem, but that does not necessarily mean that it can be solved analytically.

The first step is to work out the speeds of the waves in our general fluid system and to compute them for the particular Riemann problem we are trying to solve (for particular left and right states). It turns out that the speeds of the characteristic waves can be found as the eigenvalues λ_k of the Jacobian matrix, $A_{ij} = \partial \mathcal{F}_i / \partial U_j$. This follows directly from the structure of the Euler equations, which can be turned into an advection equation for each of the eigenvector fluid states. These states then move with a speed equal to the respective eigenvalue (see §B.2 for a derivation). One simple example are the linearized fluid equations that we encountered when deriving sound waves




Figure 24: The Riemann problem for the 1D Euler equations. Since there are three equations (corresponding to mass, momentum, and energy conservation), there are always three waves in the problem with eigen-speeds u and $u \pm c_s$. However, the type of waves depends on the initial setup. The four panels on the right show examples of possible outcomes, including shocks (thick lines), contact discontinuities (dashed lines), and rarefaction waves (fans) going to the left and/or right. The central wave is always a contact discontinuity. Note that shocks can travel supersonically depending on the fluid velocity u . Figure adapted from [Toro](#).

in §4.1,

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \rho_0 \frac{\partial u}{\partial x} &= 0 \\ \frac{\partial u}{\partial t} + \frac{c_s^2}{\rho_0} \frac{\partial \rho}{\partial x} &= 0. \end{aligned} \quad (8.9)$$

For this system of two equations, we find two eigenvalues $\lambda_k = \pm c_s$ and corresponding fluid states with density ρ_0 and speed $u = \pm c_s$ (§B.2). Clearly, this system describes signals moving either left or right at the speed of sound. Due to its simplicity, we can analytically solve the general Riemann problem (e.g., [Toro](#) §2.3.4). For initial states $\mathbf{V}_L = (\rho_L, u_L)$ and $\mathbf{V}_R = (\rho_R, u_R)$, we get a solution $\mathbf{V}(x_0, t) = (\rho_*, u_*)$ with

$$\begin{aligned} \rho_* &= \frac{1}{2}(\rho_L + \rho_R) + \frac{\rho_0}{2c_s}(u_L - u_R) \\ u_* &= \frac{1}{2}(u_L + u_R) + \frac{c_s}{2\rho_0}(\rho_L - \rho_R). \end{aligned} \quad (8.10)$$

This result is illustrated for an initial density jump in [Figure 23](#). The situation gets more complicated when we consider the full 1D Euler equations. Since we now have three equations, we also have three possible waves with eigen-speeds u and $u \pm c_s$, meaning that information can travel with the sound speed added to the local fluid velocity (§B.2). The resulting Riemann fan has three waves ([Figure 24](#)). The middle wave is always a **contact discontinuity**, a discontinuity across which pressure and normal velocity are constant but density, temperature, and entropy are not. For example, we can have higher density and lower temperature on one side of the discontinuity, which just moves along with the fluid velocity u . The other discontinuities can be two shocks, two rarefaction waves, or one of each (see examples in [Figure 24](#)). In 2D or 3D, this structure does not change fundamentally because the additional eigenvalues are also $\lambda = u$. These “degenerate” eigenvalues are a reflection of the fact that the fluid can flow in any direction.

The Riemann problem for the 1D Euler equations can be solved exactly, or at least be approximated numerically to arbitrary accuracy (e.g., [Colella & Glaz 1985](#)). However, approximate




Figure 25: Schematic depiction of Godunov schemes without and with reconstruction. The symbols have the same meaning as in Figure 20. The simplest Godunov scheme (left) is first-order in both space and time because we used the cell-centered states \mathbf{U}_i to compute the fluxes across interfaces. Moreover, those fluxes were taken at time n instead of being averaged over the timestep. With spatial reconstruction (right), we compute the left and right edge states in each cell (yellow circles) and estimate the fluxes based on those. However, the fluxes are still taken at time level n instead of $n + 1/2$, making the scheme second-order in space but not in time. We use the symbol \mathbf{U} to denote all fluid states even though the reconstruction is commonly performed on the primitive variables \mathbf{V} .

solutions tend to be computationally faster. We will restrict ourselves to one of the simplest approximations, the **HLL Riemann solver** (named after Harten, Lax, and van Leer; Harten et al. 1983). Here, we do not try to obtain the full solution of the Riemann problem but rather an averaged flux at the discontinuity (or cell interface),

$$\mathcal{F}_{i\pm 1/2}^{\text{HLL}} = \begin{cases} \mathcal{F}_L & \forall S_L \geq 0 \\ \frac{S_R \mathcal{F}_L - S_L \mathcal{F}_R + S_L S_R (\mathbf{U}_R - \mathbf{U}_L)}{S_R - S_L} & \text{otherwise} \\ \mathcal{F}_R & \forall S_R \leq 0 \end{cases} \quad (8.11)$$

We have adopted a common notation of L and R subscripts, which indicate the cells to the left and right of an interface (cells i and $i + 1$ for the $i + 1/2$ interface, for example). Similarly, \mathcal{F}_L is the cell-centered flux in the left cell, and so on. Moreover, S_L and S_R are the velocities of the fastest possible waves going left and right, which we approximate as $S_L = u_L - c_s$ and $S_R = u_R + c_s$. If $S_L > 0$, all waves are traveling to the right and we can use the flux corresponding to the left state (as in the advection equation with positive velocity). The opposite case is that $S_R < 0$, meaning that all waves travel to the left. In the majority of cases, we will end up in the intermediate “star region,” where we use the averaged HLL flux (Equation 8.11). Importantly, we have ignored the contact discontinuity that splits the central region in two (Figure 24). The more complicated HLLC (HLL + contact) Riemann solver fixes this issue by adding conditions to figure out whether the interface lies in the left or right star region (Toro et al. 1994). One can add further complications such as MHD to Riemann solvers, resulting in a rich, technical field of research that is beyond the scope of these notes (see Toro for a comprehensive treatment).

We have now assembled the necessary pieces for our first Godunov scheme: the conservative update formula (Equation 8.6) and the HLL Riemann solver to compute the fluxes across cell interfaces (Equation 8.11). The left panel of Figure 25 shows a schematic of this scheme, which is only first-order accurate in both space and time because we have not yet made any effort to average the fluid variables or fluxes.

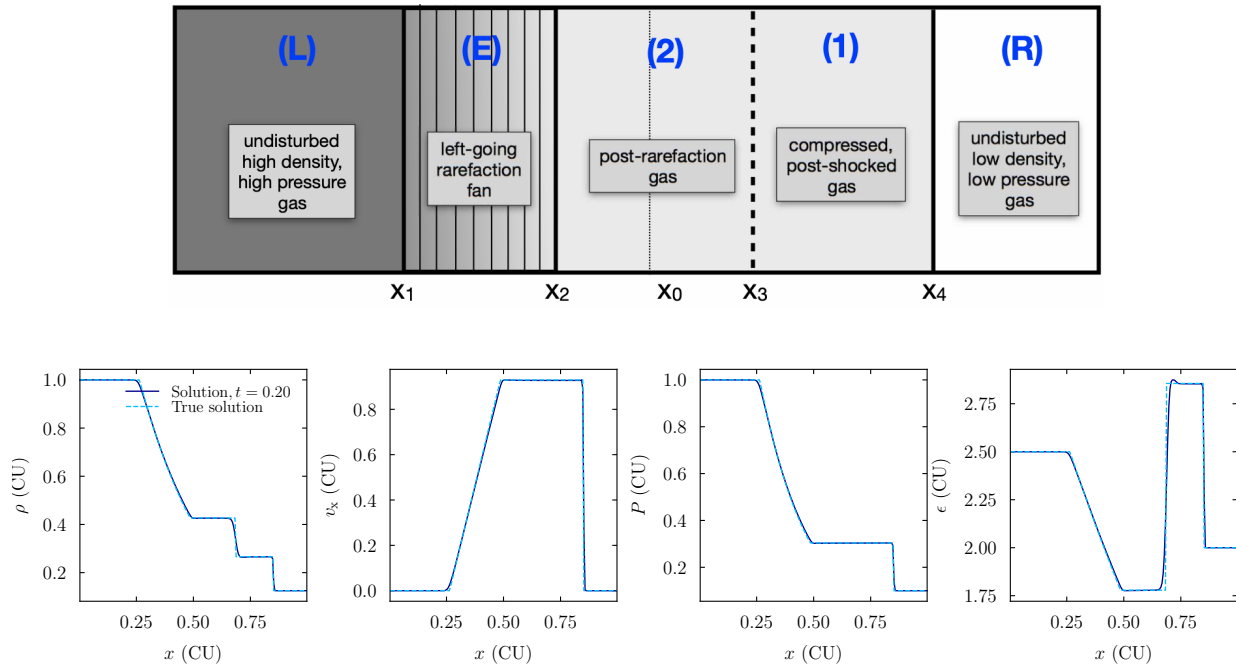


Figure 26: Schematic representation (top, from [vdB](#)) and simulation (bottom) of the Sod shocktube problem. The top panel shows the five separate fluid states that develop. The (L) and (R) states are the initial states to the left and right of position x_0 , which have not yet been influenced by any wave. At x_4 , the right-going shock is compressing and condensing the pre-shock (R) gas. At x_3 , a contact discontinuity separates the (1) state from (2), which has the same pressure and velocity but higher density (and thus lower internal energy and temperature). To the left of x_2 , a rarefaction wave connects the high-density state (L) and state (2). The bottom panel shows the numerical solution at time $t = 0.2$ (dark blue). The simulation was run with ULULA using 200 grid cells and the HLL Riemann solver. The dashed light-blue line shows the analytical solution (see [vdB](#) §21 for the mathematical expressions).

8.3 The shocktube problem

To test whether our Godunov algorithm gives the desired results, we need a more stringent test problem than the simple advection or Burgers' equations. In fact, the ideal test problem is one that produces the kinds of discontinuities that our Riemann solver should be able to deal with. This test is provided by the well-known **Sod shocktube problem** (named after Sod 1978), which basically represents a Riemann problem with an initial setup of

$$\mathbf{V}(x, 0) = \begin{cases} \mathbf{V}_L & \forall x < x_0 \\ \mathbf{V}_R & \forall x > x_0 \end{cases} \quad (8.12)$$

in a domain from $0 < x < 1$. One can choose different values for the (L) and (R) states; we use $\rho_L = 1$, $\rho_R = 1/8$, $P_L = 1$, $P_R = 1/10$, $u_L = u_R = 0$, $x_0 = 1/2$ and $\gamma = 7/5 = 1.4$. The Riemann problem that develops is shown in the top right of the four examples in Figure 23: a shock and a contact discontinuity moving to the right and a rarefaction wave moving to the left. The problem thus splits into five areas, namely, pre-shock (R), post-shock (1), post-contact discontinuity (2), rarefaction (E), and the initial left state (L). Figure 26 shows a schematic and the solution at $t = 0.2$. Note the constant pressure and velocity across the contact discontinuity. The dark blue lines show a numerical solution with the ULULA code (§C), using the HLL Riemann solver. As $\Delta x \rightarrow 0$, the scheme gets arbitrarily close to the analytical solution.

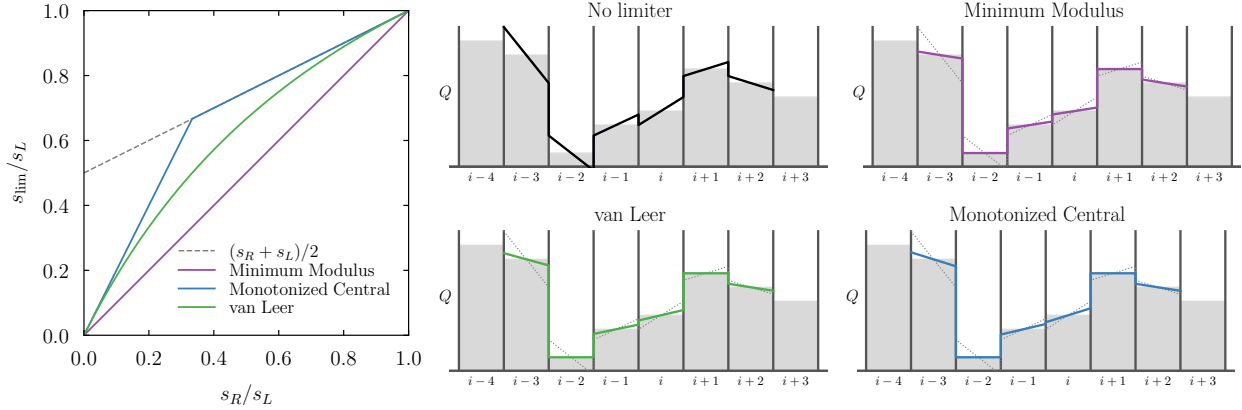


Figure 27: Slope limiters and their effect on the reconstruction of a fluid quantity $Q(x)$. The left panel shows the limited slope as a function of the ratio of the left and right slopes s_L and s_R according to a number of popular slope limiters. All limiters share a few fundamental properties, namely that the limited slope $s_{\text{lim}} = 0$ if s_R and s_L have opposite signs, that $s_{\text{lim}} = 0$ if either s_L or s_R is zero, and that they are symmetric in s_L and s_R . Thus, in this figure, we can assume that $s_R/s_L > 0$ and $s_R < s_L$, as the opposite situation would give the same results. The right panels show the slope limiters applied to an imaginary fluid quantity Q . The dotted gray lines show the same solution as the no-limiter panel to highlight the differences. The most conservative limiter is minmod (purple), which always sets the slope to the smaller of the two absolute values. The most aggressive is MC (monotonized central, blue), which maintains a slope closer to the mean of s_L and s_R as long as they are not too dissimilar.

8.4 Higher-order schemes: Reconstruction and slope limiters

The success of our Godunov scheme in the shocktube test is reassuring, but the scheme is formally only first-order accurate. Specifically, it is first-order in space because our discrete representation \mathbf{U}_i approaches the true solution $\mathbf{U}(x)$ with an accuracy proportional to Δx (see, e.g., the Taylor expansion of Equation 6.2). The scheme is also first-order in time because we made no effort to average our fluxes over the timestep, i.e., we approximated $\mathcal{F}^{n+1/2} \approx \mathcal{F}^n$ when we computed the HLL fluxes from \mathbf{U}_i^n (Figure 25). Thus, the error in the fluxes will scale as Δt .

To improve the situation, we replace the “piecewise-constant” fluid state in each cell with a linear interpolation.¹¹ For this purpose, we need a suitable approximation to the slope of the solution within each cell, which should reduce unphysically large jumps without smoothing out physically meaningful discontinuities. For convenience, we define the left, right, and central slope vectors,¹²

$$\mathbf{s}_L \equiv \frac{\mathbf{V}_i - \mathbf{V}_{i-1}}{\Delta x} \quad \mathbf{s}_R \equiv \frac{\mathbf{V}_{i+1} - \mathbf{V}_i}{\Delta x} \quad \mathbf{s}_C \equiv \frac{\mathbf{s}_L + \mathbf{s}_R}{2} = \frac{\mathbf{V}_{i+1} - \mathbf{V}_{i-1}}{2\Delta x}. \quad (8.13)$$

In Figure 9, we saw that the central derivative is a better approximation than the left and right derivatives, partly because the latter two have an unphysical preference for one direction. We use \mathbf{s}_C to represent the fluid state as linearly evolving between each cell’s walls,

$$\mathbf{V}(x) = \mathbf{V}_i + \mathbf{s}_C(x - x_i), \quad (8.14)$$

¹¹This does not mean that we linearly interpolate between cell centers! If we did that, the solution would become discontinuous at the cell centers, meaning that the interpolation within each cell would no longer be linear. Moreover, the values on the left and right of each cell interface would be equal by construction, and our Riemann solver would simply give $\mathcal{F}_{i+1/2} = \mathcal{F}([U_{i+1} - U_i]/2)$. We would not be solving any actual Riemann problems, negating the shock-capturing properties of the scheme.

¹²The reconstruction is usually performed in primitive variables \mathbf{V} , but in principle one can also use conserved quantities.

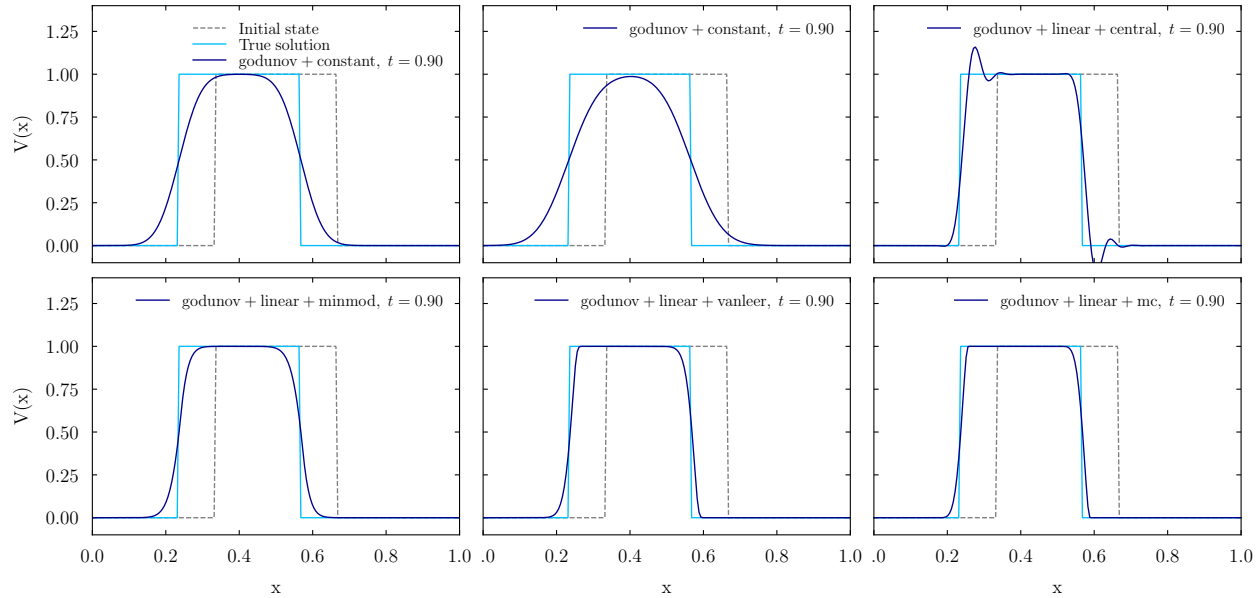


Figure 28: Advection test with Godunov schemes. As in Figures 10 and 12, the tophat test was run until $t = 0.9$ with 200 spatial cells. The first-order Godunov scheme with piecewise-constant cell values is identical to the FTBS finite-difference scheme, as shown in the top left (with $\alpha_{\text{cfl}} = 0.5$, see §8.1). For our limiter tests, we set $\alpha_{\text{cfl}} = 0.01$ to force the solver to take many timesteps, which adds diffusivity (top center). Adding reconstruction with central derivatives (without a limiter) leads to oscillations (top right). The bottom row shows the results with the minmod, van Leer, and MC slope limiters. All three result in stable schemes and produce much more accurate results than the piecewise-constant method. The more aggressive limiters (MC) are less diffusive than the more conservative ones (minmod).

which leads to states at the left and right interfaces of

$$\mathbf{V}_{i\mp 1/2} = \mathbf{V}_i \mp \frac{\Delta x}{2} \mathbf{s}_C. \quad (8.15)$$

This process is called **reconstruction**, a term that refers to recovering the states at the cell interfaces that were lost when we discretized the domain. The resulting interpolations are schematically shown with black lines in the top-center panel of Figure 27. The reconstruction looks smoother than the piecewise-constant representation in most areas, but the example also highlights a potential risk: the fluid state seems to overshoot or undershoot the neighboring cells at some interfaces. This will lead to unphysical, excessive gradients and thus instability. Most worryingly, the solution becomes negative at one interface.

Once again, we explore the effect of the reconstruction using the advection equation. We recall that, since u is constant everywhere, our Riemann solver is trivial: it simply takes the state on the upwind side of an interface and computes the corresponding flux. Thus, the advection test isolates the effects of the spatial reconstruction. Figure 28 shows the tophat shifted almost a full period (as in Figures 10 and 12, where we investigated finite-difference schemes). The top left panel shows the solution for piecewise-constant cell states as discussed in §8.1. Indeed, the result from the FTBS scheme is recovered exactly (compare to the bottom right panel of Figure 10).

For the following comparisons, we exaggerate the scheme’s diffusivity by setting $\alpha_{\text{cfl}} = 0.01$, which forces the solver to take a large number of timesteps (top center panel of Figure 28). The top right panel shows what happens when we add the linear reconstruction with centered slopes (Equation 8.15): the solution develops oscillations and eventually diverges! Clearly, the over- and




Figure 29: Linear (left) and parabolic (right) reconstruction of states at the cell interfaces. Dashed lines show the unlimited derivatives and parabolic fits, and solid lines show a limited version. Figure from Zingale.

undershoot near the tophat discontinuity present a problem.¹³

The solution is to impose a **slope limiter**, that is, an algorithm that limits the linear slope in those cells where it risks over- or undershooting the values in neighboring cells. Most limiter algorithms identify such cases purely based on s_L and s_R , meaning that they can be represented as a function

$$s_{\text{lim}} = \phi(s_L, s_R). \quad (8.16)$$

Here, we consider a single fluid quantity and write s instead of \mathbf{s} . This process is repeated for each individual quantity in the state vector \mathbf{V} . One property is shared between all slope limiters: if s_L and s_R have opposite signs, we set $s_{\text{lim}} = 0$ because the current cell is either a trough or peak compared to its neighbors (Figure 27). Similarly, the most aggressive (steep) slope we should ever choose is s_C , the average between s_L and s_R (Equation 8.13).

Figure 27 visually summarizes three popular slope limiters. Here, we assume that s_L and s_R have the same sign and that $|s_L| > |s_R|$; the figure would look the same in the opposite case since all slope limiters are symmetric. The dashed gray line shows the average slope s_C . Where the lines fall below s_C , the limiter chooses a more conservative (shallower) slope. The most extreme cases are the “minmod” limiter, which always chooses the shallower alternative, and “monotonized central” (MC), which chooses s_C unless the slopes differ significantly. The “van Leer” limiter represents an intermediate choice (van Leer 1974). The right panels of Figure 27 show the impact of these limiters on the example distribution of Q . All limiters avoid negative solutions and create a smoother interpolation, with some subtle differences between them.

The bottom three panels of Figure 28 show that the reconstruction scheme works very well with any of the limiters. The more conservative limiters lead to slightly more diffusion, which we can understand as follows: with a more conservative reconstruction, we end up with larger differences at the interface boundaries, and thus larger Riemann fluxes that cause larger diffusivity. Nevertheless, using reconstruction with any limiter drastically reduces the diffusion compared to the first-order scheme without reconstruction (top center panel).

In summary, we have finally succeeded in solving the advection equation with a stable, second-order (in space) scheme that is not terribly diffusive. This kind of scheme (Godunov with linear reconstruction and slope limiter) is sometimes called MUSCL, for “Monotonic Upstream-centered Scheme for Conservation Laws” (van Leer 1979). Of course, the advection test is relatively basic. MUSCL-type schemes can be tested on the full Euler equations with the ULULA code (§C).

Given the success of a linear reconstruction, we might be tempted to go to even higher order. For example, the “piecewise-parabolic” method replaces the linearly evolving state vector within each cell with a quadratic polynomial, making the scheme 3rd-order accurate in space (Colella & Woodward 1984, see example in Figure 29). However, higher-order schemes are also more costly

¹³It can be shown that a linear reconstruction without a limiter corresponds exactly to some of the higher-order finite-difference methods analyzed in §6.4, depending on whether we choose the left, central, or right derivative (leading to the Beam-Warming, Fromm, and Lax-Wendroff schemes, respectively; see vdB §21).




Figure 30: Schematic depiction of a MUSCL-Hancock scheme. The symbols have the same meaning as in Figures 20 and 25. As before, we reconstruct the cell-edge states, but now we also advance them by a half timestep (the so-called Hancock step). We feed the time-advanced edge states to the Riemann solver to obtain approximate time-averaged fluxes at time level $n + 1/2$. As in Figure 25, we use \mathbf{U} for all states, regardless of whether they are expressed in the primitive or conserved representation in practice.

computationally and add complexity to the code. As a result, second-order schemes are often the method of choice.

8.5 Higher-order schemes: Time integration

Our current Godunov scheme represents a first-order “Euler” integration method in time, where we take the local derivative in time (the flux difference) and multiply it by the timestep,

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x} \Delta \mathcal{F}, \quad (8.17)$$

where we have abbreviated our approximation to the differential fluxes across a cell’s left and right interfaces as $\Delta \mathcal{F}$. Unfortunately, first-order schemes are notoriously unstable in any type of numerical integration. Regardless of how we approximate $\Delta \mathcal{F}$, we can obtain higher-order time integration by applying the flux operation twice and averaging,

$$\begin{aligned} \mathbf{U}^* &= \mathbf{U}^n + \Delta t / \Delta x \times \Delta \mathcal{F}(\mathbf{U}^n) \\ \mathbf{U}^{**} &= \mathbf{U}^* + \Delta t / \Delta x \times \Delta \mathcal{F}(\mathbf{U}^*) \\ \mathbf{U}^{n+1} &= (\mathbf{U}^n + \mathbf{U}^{**}) / 2. \end{aligned} \quad (8.18)$$

This scheme is called a **second-order Runge-Kutta** integrator (RK2). It has the advantage of simplicity because both steps demand exactly the same operations. On the other hand, the RK2 scheme is computationally expensive because we need to perform two full timesteps instead of one, including the relatively expensive Riemann problems.

In practice, another idea turns out to be more fruitful. We recall that the goal is to time-average the fluxes to get $\mathcal{F}^{n+1/2}$ instead of the \mathcal{F}^n . One way to achieve a second-order estimate (without solving extra Riemann problems) is to advance the reconstructed states at a cell’s left and right edge by a half timestep, where we approximate the time evolution with the flux difference between the left and right sides of the cell. This flux difference occurs purely because of the reconstruction, which leads to different states at the left and right edges of each cell. Mathematically, the **Hancock**




Figure 31: 2D advection test. A sharp blob of dense gas is initially placed at the center and transported across the domain with speed $u_x = 0.5$ and $u_y = 0.4$. We analyze the results when the blob has roughly crossed the domain once, after about 550 time steps. An ideal scheme would perfectly preserve the blob (top left panel). The simplest scheme, Euler with no spatial interpolation and the HLL Riemann solver (top second panel), is extremely diffusive. Using an HLLC Riemann solver helps significantly in this particular test because of the contact discontinuity at the edge of the blob (top third panel). In this case, using Hancock time-interpolated states does not make a huge difference (top right panel). Adding spatial interpolation and the conservative MinMod slope limiter improves the solution significantly, although the blob shows deformations when using the Euler scheme (bottom left panel). Using Hancock timestepping is more diffusive but also reduces those artifacts (bottom second panel). Combining the more aggressive MC slope limiter with Euler timestepping leads to instability (bottom third panel). However, the Hancock time interpolation restores stability and leads to the best solution (bottom right panel).

step can be written as

$$\Delta U_i \equiv \frac{\Delta t}{\Delta x} \left[\mathcal{F}(U_{i-1/2}^n) - \mathcal{F}(U_{i+1/2}^n) \right] \quad \text{and} \quad U_{i\pm 1/2}^{n+1/2} = U_{i\pm 1/2}^n + \frac{\Delta U_i^n}{2}. \quad (8.19)$$

In other words, we compute an approximate time update for the state in cell i and apply it equally to the reconstructed states on the left and right. An equivalent (and often faster) way to compute the time-advanced states is to use the primitive-variable matrix formulation to write

$$\frac{\Delta \mathbf{V}}{\Delta t} \approx \mathbf{A}_x(\mathbf{V}) \frac{\Delta \mathbf{V}}{\Delta x} \quad \text{and} \quad \mathbf{V}_{i\pm 1/2}^{n+1/2} = \mathbf{V}_{i\pm 1/2}^n + \frac{\Delta t}{2} \mathbf{A}_x(\mathbf{V}_{i\pm 1/2}^n) \mathbf{s}_{\text{lim}}, \quad (8.20)$$

where the matrix $\mathbf{A}_x(\mathbf{V})$ is a direct representation of the Euler equations (§B.2) and $\mathbf{s}_{\text{lim}} = \Delta \mathbf{V} / \Delta x$ by construction (e.g., Mignone & Tzeferacos 2010). Whether we use Equation 8.19 or 8.20 to perform the Hancock step, we take the time-advanced states at the cell edges as input to the Riemann solver (Figure 30). This so-called **MUSCL-Hancock** scheme is now second-order accurate in both space and time despite evaluating only one Riemann problem per timestep per interface. MUSCL-type schemes are the backbone of many hydro codes used in astrophysics (§8.7), as well as the ULULA solver (see Figure 31 and §C).

8.6 Multiple dimensions

Taking a Godunov scheme to more than one dimension adds significant complexity. For example, we need to take transverse velocity fluxes into account. There is however, a simple trick that simplifies things enormously, called “dimensional splitting” (Strang 1968). Here, we apply our solver separately in each dimension, neglecting transverse terms in the Riemann problem (although they are included in $\mathcal{F}(\mathbf{U})$, for example the $u_i u_j$ term in the momentum flux).

Intuitively, such a scheme would seem to incur large (first-order) errors because the x -fluxes are based on the previous timestep, the y -fluxes on the previous timestep updated with x -fluxes, and so on. However, we can magically eliminate all first-order error terms by switching up the order of directions. For example, in a 2D scheme, we apply our algorithm in the x - y - x order, in 3D x - y - z - z - y - x , and so on. The resulting algorithms are called **directionally split schemes**. Conversely, **unsplit schemes** must perform a timestep in all dimensions at once. The multidimensional, unsplit version of the MUSCL-Hancock scheme is called CTU for “Corner Transport Upwind” (Colella 1990). Both split MUSCL-Hancock and CTU are among the most popular methods for astrophysical hydrodynamics codes.

8.7 Popular hydrodynamics codes in astrophysics

We end this chapter by listing some of the most commonly used codes in astrophysics. Many fields (e.g., engineering) have established their own set of codes, and it seems impossible to write a high-performance hydro code that is general enough for all applications. Even within astrophysics, most codes serve a specialized sub-field. The following list is highly incomplete and slews toward galaxy formation.

First, there are a number of grid-based Godunov codes that all implement some form of adaptive refinement (which is beyond the scope of these notes). Most of these codes implement multiple reconstruction schemes, limiters, and Riemann solvers, as well as MHD. Popular examples include (sorted by age):

- **ZEUS / Athena / Athena++** (Stone & Norman 1992; Stone et al. 2008, 2020) are a series of highly advanced C/C++ frameworks for (M)HD simulations of all kinds.
- **ENZO** (Bryan et al. 1995, 2014) is a Fortran code with block-based AMR, meaning that blocks of cells are refined to the same level. ENZO is most commonly used for galaxy formation, feedback, and cluster simulations.
- **ART** (Kravtsov et al. 1997; Khokhlov 1998; Rudd et al. 2008) is a C-based AMR hydrodynamics+gravity code intended for galaxy formation, although an N -body version is also used for structure formation. The code uses oct-tree AMR, meaning that it refines cells adaptively into eight child cells.
- **FLASH** (Fryxell et al. 2000) is a flexible, multi-purpose, block-based AMR framework with numerous hydro+MHD solvers. FLASH is commonly used for explosions such as supernovae, but also for galaxy cluster simulations.
- **RAMSES** (Teyssier et al. 2006) is a tree-AMR code similar to ART in its layout but written in Fortran. It is one of the most common codes for galaxy simulations and contains numerous modules for feedback, radiative transfer, MHD, and so on.
- **PLUTO** (Mignone et al. 2007, 2012) is a C-based framework for (M)HD and even relativistic hydrodynamics. The code is oriented towards pure-hydro problems (such as accretion disks) since it does not solve for self-gravity.

In these notes, we have totally neglected an equally popular category of codes, namely, Lagrangian codes based on fluid elements (or “particles”). In **Smoothed Particle Hydrodynamics (SPH)**, the Euler equations are solved by following fluid particles that are influenced by their neighbors, e.g., via pressure. Popular codes of this type include:

- **Gadget (2/3/4)** (Springel et al. 2001; Springel 2005; Springel et al. 2021) is a series of extremely popular SPH codes that have been used in countless N -body and baryonic simulations of galaxy and structure formation, as well as other problems.
- **Gasoline** (Wadsley et al. 2017) is another C-based SPH code that is commonly used for galaxy formation simulations.

The main advantages of SPH methods are that they are automatically adaptive in resolution (since particles move to where the mass is), that they better conserve momentum and angular momentum, and that gas and collisionless particles (such as dark matter) can be treated on the same footing. However, SPH algorithms tend to struggle with steep density gradients and shocks (e.g., Agertz et al. 2007). To combine the advantages of a Lagrangian framework with shock-capturing Godunov schemes, the SPH and grid frameworks have more recently been merged in **moving-mesh codes**. These algorithms use Lagrangian particles to follow the matter in a simulation but draw a mesh between the particles and solve the Euler equations using Godunov-like schemes. Commonly used codes include:

- **Arepo** (Springel 2010; Weinberger et al. 2020) popularized the moving-mesh concept in astrophysics. It draws a mesh of non-cubic Voronoi cells between its Lagrangian particles. The mesh adapts as the particles move through space.
- **GIZMO** (Hopkins 2015) is another moving-mesh code based on Gadget3, with a flexible suite of hybrid Lagrangian-Eulerian hydro solvers.

This list is highly incomplete! Numerous other hydro codes are used in astrophysics, particularly for specialized applications such as 1D stellar evolution (MESA, Paxton et al. 2011), rotating disks (DISCO, Duffell 2016), nuclear reactions (CASTRO, Almgren et al. 2010) or relativistic problems (e.g., TESS, Duffell & MacFadyen 2011).

9 Fluid instabilities

In §4, we linearized the Euler equations and perturbed them around a uniform background state at rest. We found that a wave-like density and pressure perturbation travels as a sound wave, with no exponentially growing or decreasing modes. In this chapter, we use similar techniques to demonstrate the existence of a number of fluid instabilities that arise from gravitational collapse and at the boundary between two fluids. Table 3 gives an overview of the physical assumptions made in all of these cases.

9.1 The Jeans instability

As soon as a fluid is perturbed, a non-uniform gravitational potential arises due to the mutual attraction of gas particles. When considering sound waves, we neglected this **self-gravity** of the fluid because we implicitly assumed that pressure forces were much larger than gravitational accelerations. This situation is intuitive on Earth, where we experience a significant, static gravitational field but generally negligible self-gravity of gases. For example, the air in a room does not spontaneously collapse to a point at the center.

In astrophysics, however, we consider gases with much larger overall masses and much lower pressures. A cloud of gas is faced with a basic competition between pressure supporting the cloud and gravity contracting it. While the gravitational potential of the cloud grows with its mass, the pressure corresponds to a fixed energy per unit volume. Thus, we expect there to be a mass scale where gravity “wins” and the cloud collapses. This instability is known as the **Jeans instability**. The derivation follows in the same steps as that for sound waves (§4.1), but including self-gravity (see also CC §10.2). As previously, we start from the Eulerian continuity equation,

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (9.1)$$

and the velocity equation,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} - \nabla \Phi, \quad (9.2)$$

but we also add the Poisson equation for self-gravity, $\nabla^2 \Phi = 4\pi G \rho$. We assume that the system is adiabatic so that we can omit the internal energy equation. Once again, we express the perturbation as a small change to an unperturbed, stationary fluid state,

$$\rho = \rho_0 + \rho_1, \quad P = P_0 + P_1, \quad \mathbf{u} = \mathbf{u}_1, \quad \Phi = \Phi_0 + \Phi_1. \quad (9.3)$$

We linearize the equations, meaning that we keep only first-order terms in the perturbed quantities. This procedure gives the same equations as for sound waves but with a gravitational term,

$$\begin{aligned} \frac{\partial \rho_1}{\partial t} + \rho_0 \nabla \cdot \mathbf{u}_1 &= 0 \\ \frac{\partial \mathbf{u}_1}{\partial t} + \frac{\nabla P_1}{\rho_0} + \nabla \Phi_1 &= 0 \\ \nabla^2 \Phi_1 &= 4\pi G \rho_1, \end{aligned} \quad (9.4)$$

where we have used Equation 4.6 to show that $\nabla P_1 / \rho \approx \nabla P_1 / \rho_0$ to first order. We expand the perturbation in pressure and replace the derivative with the sound speed since $c_s^2 = \partial P / \partial \rho$,

$$P(\rho_0 + \rho_1) = P(\rho_0) + \left(\frac{\partial P}{\partial \rho} \right)_0 \rho_1 + \mathcal{O}(\rho_1^2) \approx P_0 + c_s^2 \rho_1. \quad (9.5)$$

Instability	Density ρ	Velocity u	Gravity Φ	Assumptions	§
Sound waves	$\rho_0 + \rho_1$	$u_0 = 0$	0	adiabatic	4
Jeans	$\rho_0 + \rho_1$	$u_0 = 0$	$\nabla^2 \Phi_1 \propto \rho_1$	adiabatic	9.1
Surface waves	$\rho_2 \ll \rho_1$	$u_1 = u_2 = 0$	gz	steady, $\boldsymbol{\omega} = 0$	9.3
Rayleigh-Taylor	$\rho_2 > \rho_1$	$u_1 = u_2 = 0$	gz	steady, $\boldsymbol{\omega} = 0$	9.4
Kelvin-Helmholtz	no constraint	$u_1 \neq u_2$	0 or gz	steady, $\boldsymbol{\omega} = 0$	9.5

Table 3: Assumptions made in deriving various fluid instabilities (or stability in the case of waves). We list the classic scenarios, but the derivation can be made more complicated by relaxing assumptions in some cases (e.g., a Rayleigh-Taylor instability with relative velocity). The meaning of the ρ and u variables differs from the first two cases (where we consider a single, perturbed fluid) to the latter three (where we consider a two-fluid interface).

The result is a velocity equation in terms of only density and gravity,

$$\frac{\partial \mathbf{u}_1}{\partial t} + c_s^2 \frac{\nabla \rho_1}{\rho_0} + \nabla \Phi_1 = 0. \quad (9.6)$$

By taking the time derivative of the linearized continuity equation and substituting the velocity equation above we get our wave equation,

$$\begin{aligned} \frac{\partial^2 \rho_1}{\partial t^2} + \rho_0 \nabla \cdot \left(-c_s^2 \frac{\nabla \rho_1}{\rho_0} - \nabla \Phi_1 \right) &= 0 \\ \implies \frac{\partial^2 \rho_1}{\partial t^2} &= c_s^2 \nabla^2 \rho_1 + 4\pi G \rho_0 \rho_1. \end{aligned} \quad (9.7)$$

We insert the usual sinusoidal wave perturbation, which we can imagine as a series of slightly over- and underdense peaks and troughs like in a planar sound wave. Mathematically, we have

$$\rho_1 = A e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}, \quad (9.8)$$

where the wave vector $\mathbf{k} = k \hat{\mathbf{k}} = 2\pi/\lambda \hat{\mathbf{k}}$ and the angular frequency $\omega = 2\pi f$ (where f is the frequency in Hz). This results in the dispersion relation

$$-\omega^2 = -k^2 c_s^2 + 4\pi G \rho_0, \quad (9.9)$$

which we can cast in a more intuitive form by defining the Jeans wavenumber,

$$\omega^2 = c_s^2 (k^2 - k_J^2) \quad k_J \equiv \frac{\sqrt{4\pi G \rho_0}}{c_s}. \quad (9.10)$$

We see that ω becomes imaginary if $k < k_J$, in which case we have

$$\omega = c_s \sqrt{-1} \sqrt{k_J^2 - k^2} \implies e^{-i\omega t} = e^{c_s \sqrt{k_J^2 - k^2} t}. \quad (9.11)$$

This solution corresponds to a runaway collapse that happens faster for smaller k . More intuitively, this means that perturbations are unstable if their wavelength is greater than the **Jeans length**,

$$\boxed{\lambda_J \equiv \frac{2\pi}{k_J} = \sqrt{\frac{\pi c_s^2}{G \rho_0}}} \quad (9.12)$$

While the Jeans length technically refers to a plane wave in one direction, it is usually applied to spherical gas clouds. In this scenario, one Jeans length will roughly enclose a sphere of diameter λ_J , which allows us to define the **Jeans mass** within that sphere,

$$M_J \equiv \frac{4\pi\rho_0}{3} \left(\frac{\lambda_J}{2}\right)^3 = \frac{\pi^{5/2}c_s^3}{6G^{3/2}}\rho_0^{-1/2} = \frac{\pi^{5/2}}{6} \left(\frac{\gamma k_B T}{\mu m_p G}\right)^{3/2} \rho_0^{-1/2} \quad (9.13)$$

where we have inserted the ideal gas equation of state in the final step. The key dependences are those on density and temperature, given that all other numbers are constants. The Jeans mass is sometimes quoted with slightly different numerical pre-factors due to different definitions.

Unfortunately, our entire derivation suffers from a major issue: in assuming a uniform background potential with $\nabla\Phi_0 = 0$, we also assumed $\nabla^2\Phi_0 = 0$ and thus $\rho_0 = 0$. In other words, a uniform background can exist only in vacuum! Physically speaking, we have ignored all the mass around the perturbation. This omission is called the **Jeans swindle** (Binney & Tremaine 2008), but a more thorough calculation shows that the naive result is fortuitously close to the right answer (Falco et al. 2013).

How long does the collapse of a Jeans-unstable perturbation take? An estimate is given by the **free-fall time** for a sphere of uniform density,

$$t_{\text{ff}} = \sqrt{\frac{3\pi}{32G\rho_0}}. \quad (9.14)$$

Implicitly, we are assuming that the pressure becomes negligible once the perturbation has become Jeans-unstable. It makes sense that denser spheres collapse more quickly in this case. However, the simple, spherical collapse envisioned by the free-fall time cannot actually occur in the presence of angular momentum, magnetic fields, or non-thermal pressure sources.

9.2 Perturbations of a general two-fluid interface

Figure 32 shows a general two-fluid interface in pressure equilibrium. The densities ρ_1 and ρ_2 may differ, but we assume that the fluids are incompressible, i.e., that $\partial\rho_1/\partial t = \partial\rho_2/\partial t = 0$. Similarly, we allow the x -velocities u_1 and u_2 to differ, but both fluids share the same pressure, P . In the language of §8.2, we have a contact discontinuity which may or may not be moving differentially. This situation arises when we have two fluids moving past each other, or two static fluids on top of each other with a gravitational acceleration g pointing downwards. Our strategy is the same as for sound waves: we consider a small, linear perturbation to the fluid interface, parameterize it as a real and complex wave, and check whether the solution has any real and/or complex parts. The full derivation is laid out in §B.3; here, we give only a schematic outline.

We constrain the Euler equations by the condition that an initially vorticity-free, steady flow remains irrotational according to Kelvin’s circulation theorem (§3). Thus, we can write all velocities as gradients of a “velocity potential” Ψ , which itself splits into a background state and a perturbation (for each fluid). The continuity and energy equations give no information due to the incompressibility and pressure equilibrium, but the velocity equation translates into a relation between the velocity potential, pressure, density, and gravitational potential. However, when we parameterize the perturbation in the usual way,

$$\xi(x, t) = Ae^{i(kx - \omega t)}, \quad (9.15)$$




Figure 32: Setup of a general two-fluid interface with an initial boundary along the x -direction and a gravitational acceleration g pointing in the negative z -direction. The pairs of arrows highlight that the fluid velocities are along the x -direction but may point in opposite directions or be zero. We imagine that the interface is perturbed with a small wave $z = \xi(x, t)$. The goal is to find out whether this perturbation grows, moves as a wave, or vanishes.

we incur three unknowns: the amplitude A (how far the interface shifts in the z -direction) and the amplitude of the corresponding perturbation in the lower and upper velocity fields. The two additional pieces of information needed are provided by the Laplacian condition that $\nabla^2 \Psi = 0$ (again due the vorticity-free state of both fluids). Combining the resulting equations with the velocity equation and pressure equilibrium, we have three equations that relate the perturbation amplitudes to the known fluid quantities ρ and u . Combining these equations, we obtain the general **dispersion relation for perturbations of a two-fluid interface**,

$$\boxed{\frac{\omega}{k} = \frac{\rho_1 u_1 + \rho_2 u_2}{\rho_1 + \rho_2} \pm \left(\frac{g(\rho_1 - \rho_2)}{k(\rho_1 + \rho_2)} - \frac{\rho_1 \rho_2 (u_1 - u_2)^2}{(\rho_1 + \rho_2)^2} \right)^{1/2}} \quad (9.16)$$

We conclude that the nature of the instabilities depends on a balance between the input parameters and gravity: if the term inside the root becomes negative, there will be an instability. We also note that the only dependence on the wavenumber arises through the g/k term in the root. This makes sense because the system is otherwise self-similar, meaning that there is no preferred spatial scale that could give rise to a “special treatment” of perturbations of any particular wavenumber. Gravity breaks the self-similarity because we have specified our self-similar unit system by supplying the three input quantities ρ , u , and P .

9.3 Surface gravity waves

We now apply Equation 9.16 to some specific cases. For example, imagine a stable scenario where the two fluids are at rest and the lower fluid is heavier than the upper fluid ($\rho_2 < \rho_1$), e.g., water and air. The dispersion relation then reduces to

$$\frac{\omega}{k} = \pm \sqrt{\frac{g(\rho_1 - \rho_2)}{k(\rho_1 + \rho_2)}}. \quad (9.17)$$

Without gravity, $\omega/k = 0$, meaning that there are no perturbations. With gravity, we get only real values on the RHS, indicating that the perturbations propagate as waves. The wave properties simplify when $\rho_1 \gg \rho_2$, as is the case for water and air,

$$c_{\text{wave}} = \frac{\omega}{k} \approx \sqrt{\frac{g}{k}} = \sqrt{\frac{g\lambda}{2\pi}}, \quad (9.18)$$

where λ is the wavelength. This equation does not apply to so-called shallow-water waves like those breaking on ocean beaches, but we can imagine ripples in a pond. For a wavelength of 10 cm, we have $c_{\text{wave}} \approx 40$ cm/s, which sounds about right!




Figure 33: Experimental manifestation of the Rayleigh-Taylor instability. The darker liquid is heated at the bottom, lowering its density below that of the fluid on top. The mushroom-shaped plumes are characteristic of the Rayleigh-Taylor instability. The heating rate is higher in the right panel than in the left. Figure from van Dyke (1982).

9.4 The Rayleigh-Taylor instability

What happens if we invert the densities of the fluids so that the heavier fluid is on top, $\rho_2 > \rho_1$? In that case, Equation 9.17 is imaginary regardless of k , meaning that all wavelengths are unstable! Any perturbation at the interface leads to a plume of lower-density gas rising up or higher-density gas falling down (Figure 33), and a subsequent exchange where the fluid tries to arrange the higher density at the bottom. This is known as the Rayleigh-Taylor (RT) instability. It occurs in numerous astrophysical situations, for example in Type Ia supernovae: as nuclear burning consumes a white dwarf starting from the center, the hot, incinerated material wants to rise above the cold, dense material surrounding it.

9.5 The Kelvin-Helmholtz instability

Another classic two-fluid instability is the Kelvin-Helmholtz (KH) instability, which arises in the case of two fluids moving relative to each other ($u_1 \neq u_2$). Let us first consider the case of no gravity, $g = 0$. Since the second term in the root is negative, the dispersion relation becomes

$$\frac{\omega}{k} = \frac{\rho_1 u_1 + \rho_2 u_2}{\rho_1 + \rho_2} \pm i \sqrt{\frac{\rho_1 \rho_2 (u_1 - u_2)^2}{(\rho_1 + \rho_2)^2}} \quad (9.19)$$

Given that the input densities are positive and the velocity difference in the root is squared, all wavelengths are unstable! The instability is present even when $\rho_1 = \rho_2$, in which case

$$\frac{\omega}{k} = \frac{u_1 + u_2}{2} \pm i \frac{u_1 - u_2}{2}. \quad (9.20)$$

Physically, this relation means that the perturbation is traveling at the average speed of the two velocities and that the instability grows exponentially with a rate proportional to the difference of the speeds. For a commonly used setup with $\rho_1 = 1$, $\rho_2 = 2$, $u_1 = -1/2$, and $u_2 = +1/2$, we find

$$\frac{\omega}{k} = \frac{1}{6} \pm i \frac{\sqrt{2}}{3}. \quad (9.21)$$

The fact that any wavelength grows means that we “get out what we put in:” if we add perturbations of a certain size initially, those will grow into non-linear structures of roughly the same size. If we put in no perturbations, the instability will be seeded by random, numerical noise whose scale depends on the parameters of our simulation (Robertson et al. 2010). Similarly, in nature, we



Figure 34: Experimental confirmation of the Kelvin-Helmholtz instability. In a rectangular tube, a layer of water lies on top of colored brine. The tube is suddenly tilted, setting the fluids into motion. Figure from van Dyke (1982).

expect KH instabilities of any size to be possible, with the actual size depending on the physics of the situation (e.g., viscosity). If we add gravity to the situation, we find that the instability develops only if the term in the root in Equation 9.16 is still negative, which happens if

$$k > \frac{g(\rho_1^2 - \rho_2^2)}{\rho_1 \rho_2 (u_1 - u_2)^2}. \quad (9.22)$$

Thus, if $\rho_1 > \rho_2$, gravity stabilizes the interface to large-wavelength (small k) perturbations, but we can always find a wavelength below which perturbations grow. This situation is experimentally demonstrated in Figure 34, where a lighter fluid slides down on top of a heavier fluid and small-scale KH instabilities develop.

In astrophysics, the significance of the KH instability is that it is basically impossible to move to fluids past each other without mixing. For example, a jet of gas shooting out of an Active Galactic Nucleus (AGN) will be KH-unstable as it travels through the surrounding medium. The KH instability is implemented as a test case in the ULULA code (§C).

10 Magnetohydrodynamics

Besides the conditions that underlie the fluid picture, we have made another major assumption throughout our exploration of hydrodynamics: that fluids are electrically neutral. This assumption seems intuitive because the fluids we encounter on Earth tend to be neutral. If an atom is ionized, for example by a UV photon, it will recombine with an electron so quickly that there is no effect on the averaged fluid quantities. In astrophysics, this is not always the case because we often deal with diffuse gas and strong sources of radiation and heat. In fact, one could argue that the default state of cosmic gas is to be ionized, given that virtually all of the primordial gas in intergalactic space has been ionized since $z \approx 6$ (an epoch known as reionization). Other important plasmas include those in stars, in galaxy clusters, and in large parts of the ISM. In these systems, we cannot neglect the forces of electromagnetism (EM).

The equations in this chapter use the so-called cgs-Gaussian unit system, which differs from the usual SI unit system of Volt, Coulomb, Ampère, and so on. In SI, the Maxwell and MHD equations are littered with factors of ϵ_0 and μ_0 , the permittivity and permeability of free space. These constants are essentially an artifact of the way the unit system is defined. Namely, μ_0 is caused by the definition of the Ampère, which is defined as the current that, when running through two parallel wires separated by one meter, causes a force per current of 2×10^{-7} N/A². This definition introduces a geometric factor of 4π and a conversion factor of 10^{-7} in SI units. Similarly, ϵ_0 arises from the definition of charge. If we re-define the Coulomb force between two charges q_1 and q_2 at a separation r from SI to Gaussian units, we get

$$F_{\text{SI}} = \frac{1}{4\pi\epsilon_0} \frac{q_1 q_2}{r^2} \quad \rightarrow \quad F_{\text{Gaussian}} = \frac{q_1 q_2}{r^2}. \quad (10.1)$$

The trick is that we have replaced the Coulomb with the “statcoulomb,” which has cgs units of $1 \text{ statC} \equiv 1 \text{ g}^{1/2} \text{ cm}^{3/2} \text{ s}^{-1}$. The other Gaussian units are defined similarly (see Tables 1 and 2). The major advantages of the Gaussian system are that the electric and magnetic fields now have the same units ($1 \text{ Gauss} = 1 \text{ G} \equiv 1 \text{ g}^{1/2} \text{ cm}^{-1/2} \text{ s}^{-1}$), that the only dimensional constant that appears in the equations is c , and that $4\pi\epsilon_0 \rightarrow 1$ and $\mu_0 \rightarrow c/4\pi$ (or rather, that those constants are meaningless altogether).

10.1 The interaction between fluids and electromagnetism

On a particle level, the interaction between matter and EM fields is relatively simple: the particles feel the **Lorentz force**,

$$\mathbf{F} = q_{\text{ptl}} \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right), \quad (10.2)$$

where \mathbf{E} is the electric field, \mathbf{B} the magnetic field, \mathbf{v} the velocity of an individual particle (§1.4), and q_{ptl} its EM charge (e.g., $+q_e$ for protons and $-q_e$ for electrons, Table 2). The perhaps unfamiliar factor of c arises due to the Gaussian unit system. The result is that charged particles are accelerated in the direction of electric fields (the Coulomb interaction). If magnetic fields dominate, the particles “gyrate” around field lines and can thus not move across them. Along the magnetic field lines, there is no Lorentz force and particles can move freely.

While the EM fields influence particle trajectories, the particles in a fluid determine the \mathbf{E} and \mathbf{B} fields via two mechanisms: charge imbalances source electric fields, and if particles move, the corresponding currents source magnetic fields. The exact relations, as well as the interaction

between \mathbf{E} and \mathbf{B} fields, are described by the **Maxwell equations**:

$$\begin{aligned} \nabla \cdot \mathbf{E} &= 4\pi q \\ \nabla \cdot \mathbf{B} &= 0 \\ \nabla \times \mathbf{E} &= -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t} \\ \nabla \times \mathbf{B} &= \frac{4\pi}{c} \mathbf{J} + \frac{1}{c} \frac{\partial \mathbf{E}}{\partial t} \end{aligned}$$

The factors of c and 4π once again arise from the Gaussian unit system. The first equation (Gauss' law) tells us that electric field lines begin and end on charges. In the context of fluids, q is the **net charge density of the fluid** as a function of space and time (which we will define more precisely in the next section). The second equation tells us that magnetic fields have zero divergence, or that there are no magnetic monopoles (charges). The third equation (Faraday's law of induction) tells us that changing \mathbf{B} fields generate rotational \mathbf{E} fields. Finally, Ampère's law tells us that rotational \mathbf{B} fields are caused by both changing \mathbf{E} fields and the net current density \mathbf{J} .

In the following, we consider a simple fluid treatment of plasmas that allows us to connect the Maxwell and Euler equations via the Lorentz force, charge, and current. Before embarking on this mission, however, we should briefly question whether a fluid treatment is still appropriate. In §1.3, we demanded that the mean free path should be small compared to the size of our fluid elements. In plasmas, the mean free path becomes hard to determine because particles feel numerous long-range EM interactions in addition to localized collisions. As a substitute, one can compute the distance over which particle trajectories are changed significantly by EM interactions. This distance can be dangerously large in astrophysical systems such as galaxy clusters. While this argument might make us question the validity of a fluid treatment, astrophysical plasmas are typically magnetized, and the magnetic field lines help to confine particles. In particular, the gyro-radius of the particles' motion around field lines is

$$r_{\text{gyro}} = \frac{m_{\text{ptl}} v_{\perp}}{q_{\text{ptl}} |\mathbf{B}|/c}, \quad (10.4)$$

where v_{\perp} is the velocity component perpendicular to the magnetic field. Assuming thermal velocities, where $v_{\perp}^2 = v^2/3 = k_{\text{B}}T/m_{\text{ptl}}$, the gyro radius of protons is

$$r_{\text{gyro}} = \frac{\sqrt{m_{\text{p}} k_{\text{B}} T}}{q_{\text{e}} B/c} \approx 3 \times 10^{-10} \text{ pc} \left(\frac{T}{10^6 \text{ K}} \right)^{1/2} \left(\frac{B}{\mu\text{G}} \right)^{-1}, \quad (10.5)$$

where q_{e} is the electron charge. This radius is tiny, implying that even weak magnetic fields are very good at confining charged particles or directing them along field lines. On the other hand, fields can be highly tangled, leading to somewhat random motion.

We conclude that a fluid approach to plasmas will be generally valid in astrophysical settings. In the following, we outline a complete derivation of ideal MHD but ignore details of the particle nature of plasmas. See Goedbloed & Poedts (2004) for a much more comprehensive treatment.

10.2 The two-fluid approach

To get a sense of the impact of charged particles in fluids, we imagine a plasma with two species of positively and negatively charged particles. In practice, those will be mostly protons and electrons,

but we will keep it general for now. The first condition we impose is conservation of number density, which we write as a continuity equation in n ,

$$\begin{aligned}\frac{\partial n^+}{\partial t} + \nabla \cdot (n^+ \mathbf{u}^+) &= 0 \\ \frac{\partial n^-}{\partial t} + \nabla \cdot (n^- \mathbf{u}^-) &= 0.\end{aligned}\tag{10.6}$$

We relate these expressions to the Euler equations by defining the plasma-averaged density,

$$\rho = m^+ n^+ + m^- n^-, \tag{10.7}$$

and average fluid velocity,

$$\mathbf{u} = \frac{m^+ n^+ \mathbf{u}^+ + m^- n^- \mathbf{u}^-}{m^+ n^+ + m^- n^-}.\tag{10.8}$$

By adding the two component equations 10.6 and multiplying by mass, we obtain the usual continuity equation for density. However, we also obtain a continuity equation for the electric charge,

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathbf{J} = 0, \tag{10.9}$$

where we have defined the net charge density (in units of statC/cm³),

$$q \equiv n^+ q^+ + n^- q^-, \tag{10.10}$$

and the electric current density (in units of statA/cm²),

$$\mathbf{J} \equiv n^+ q^+ \mathbf{u}^+ + n^- q^- \mathbf{u}^-. \tag{10.11}$$

If opposite charges ($q^+ = -q^-$) travel at the same velocity, there is no net current. We are now ready to connect these pieces to the Euler equations. We already saw that the continuity equation remains unchanged. Similarly, there is no reason why the internal energy equation should be changed. The velocity equation, however, will change; we derived it in §2.3 by considering the forces acting on a Lagrangian fluid element, and found that they were pressure gradients and gravity. We use the same logic to add the Lorentz force,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} - \nabla \Phi + \frac{1}{\rho} \left(q \mathbf{E} + \frac{\mathbf{J}}{c} \times \mathbf{B} \right).\tag{10.12}$$

In principle, this expression closes our system of equations because the Maxwell equations tell us how the EM fields depend on the particle quantities. However, our current description of MHD is much more complicated than the Euler equations: we are dealing with a total of seven differential equations that we would need to solve separately for the positive and negative species (since q and \mathbf{J} depend on their relative densities and velocities).

10.3 Eliminating electric fields and currents

The key insight in simplifying the Euler+EM equations is that electric fields play virtually no role in plasmas. Ironically, the reason is that electric fields are extremely efficient at smoothing out the very charge densities that create them: as soon as $q \neq 0$, a strong \mathbf{E} field accelerates negative and positive particles to even out the net charge. The corresponding currents are governed by Ohm's law, which tells us how an electric field translates into a current. The law is normally written

$\mathbf{J} = \sigma_e \mathbf{E}$, where σ_e is the conductivity of a material. In the presence of motion and magnetic fields, we need to add the Lorentz force, and **Ohm's law** becomes

$$\mathbf{J} = \sigma_e \left(\mathbf{E} + \frac{\mathbf{u}}{c} \times \mathbf{B} \right). \quad (10.13)$$

Since we imagine that currents are effective at negating net charges, we need to consider the possibility that the moving, charged particles could “overshoot” and develop charge-current oscillations, similar to the pressure-density overshoot in a sound wave. Such plasma oscillations do indeed exist, but it turns out that charges in the plasma are shielded at scales larger than the **Debye length** (e.g., CC §13.3),

$$\lambda_D = \sqrt{\frac{k_B T}{4\pi n q_e^2}} \approx 7 \times 10^3 \text{ cm} \left(\frac{T}{10^6 \text{ K}} \right)^{1/2} \left(\frac{n}{1/\text{cm}^3} \right)^{-1/2}. \quad (10.14)$$

Thermal motions smooth out the plasma oscillations on a scale of λ_D , which is small by astrophysical standards. Moreover, the plasma is effectively neutral on scales larger than λ_D , which means that the assumption of charge neutrality is very well justified in astrophysical scenarios. Given this physical insight, we now try to eliminate \mathbf{E} from the equations using a simple scaling argument. We guess the relative size of the terms in the induction equation,

$$\nabla \times \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t} \quad \Longrightarrow \quad \frac{|\mathbf{E}|}{l_0} \approx \frac{1}{c} \frac{|\mathbf{B}|}{t_0} \quad \Longrightarrow \quad \frac{|\mathbf{E}|}{|\mathbf{B}|} \approx \frac{l_0}{t_0 c} = \frac{u_0}{c}, \quad (10.15)$$

where l_0 and t_0 are characteristic length and timescales of our problem (§5.1). This argument seems rather hand-wavy, but there are no factors in the equation except for the derivatives. We use this insight to estimate the relative importance of the current and \mathbf{E} -field terms in Ampère's law,

$$\frac{\left| \frac{1}{c} \frac{\partial \mathbf{E}}{\partial t} \right|}{|\nabla \times \mathbf{B}|} \approx \frac{l_0}{c t_0} \frac{|\mathbf{E}|}{|\mathbf{B}|} \approx \left(\frac{u_0}{c} \right)^2 \ll 1. \quad (10.16)$$

Unless the characteristic velocities in our system are relativistic, the electric field term is strongly suppressed compared to the magnetic term. We thus cancel the electric term from Ampère's law,

$$\nabla \times \mathbf{B} = \frac{4\pi}{c} \mathbf{J}. \quad (10.17)$$

At this point, we suspect that we can also ignore the electric part of the Lorentz force term in our Euler+EM velocity equation (10.12). Indeed, we find that

$$\frac{q|\mathbf{E}|}{|\mathbf{J}|/c|\mathbf{B}|} \approx \frac{|\mathbf{E}|/l_0}{|\mathbf{B}|/l_0} \frac{|\mathbf{E}|}{|\mathbf{B}|} \approx \left(\frac{|\mathbf{E}|}{|\mathbf{B}|} \right)^2 \approx \left(\frac{u_0}{c} \right)^2 \ll 1, \quad (10.18)$$

where we have used Gauss' law to relate q to \mathbf{E} and Equation 10.17 to relate \mathbf{J} to \mathbf{B} . Since the \mathbf{E} term is once again strongly suppressed, we will omit it from our velocity equation. Moreover, we use Equation 10.17 to replace \mathbf{J} and get an expression in \mathbf{B} only,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} - \nabla \Phi + \frac{1}{4\pi\rho} (\nabla \times \mathbf{B}) \times \mathbf{B}. \quad (10.19)$$

Our logic has one strange consequence: we are allowing infinitely efficient currents to instantaneously wipe out net charges, yet we do allow non-zero charges. How can a current flow if the fluid cannot be net-charged? The answer lies in Equation 10.11: a current arises when electrons and ions have different velocities, even if their charge densities cancel out.

10.4 The equations of ideal MHD

What remains is to find an expression for the evolution of the magnetic field itself. We cannot indiscriminately cross out the electric terms in the Maxwell equations unless another, additive term is much larger. Instead, we start from the simplified Ampère’s law and substitute \mathbf{B} terms for all \mathbf{E} and \mathbf{J} terms. Specifically, we take the curl of Equation 10.17 and substitute Ohm’s law for the current,

$$\nabla \times (\nabla \times \mathbf{B}) = \frac{4\pi}{c} \nabla \times \mathbf{J} = \frac{4\pi}{c} \sigma_e \left[\nabla \times \mathbf{E} + \frac{1}{c} \nabla \times (\mathbf{u} \times \mathbf{B}) \right]. \quad (10.20)$$

In Ohm’s law, we cannot neglect the electric field because $|\mathbf{E}|/(u/c)|\mathbf{B}| \sim 1$. We use the curl-of-curl identity (Equation A.20),

$$\nabla \times (\nabla \times \mathbf{B}) = \nabla(\nabla \cdot \mathbf{B}) - \nabla^2 \mathbf{B}, \quad (10.21)$$

and $\nabla \cdot \mathbf{B} = 0$, as well as Faraday’s law to replace $\nabla \times \mathbf{E}$,

$$\nabla \times (\nabla \times \mathbf{B}) = -\nabla^2 \mathbf{B} = \frac{4\pi\sigma_e}{c^2} \left[-\frac{\partial \mathbf{B}}{\partial t} + \nabla \times (\mathbf{u} \times \mathbf{B}) \right]. \quad (10.22)$$

We rearrange the terms into the **magnetic field evolution equation**,

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) = \frac{c^2}{4\pi\sigma_e} \nabla^2 \mathbf{B}. \quad (10.23)$$

This equation looks somewhat similar the Euler equations in that changes in \mathbf{B} are caused by a spatial derivative of a flux-like term, meaning motion with the fluid. The term on the RHS looks like a diffusion term comparably to the heat equation (5.16). Indeed, this “resistive” term has the effect of diffusing the magnetic field because a finite conductivity σ_e reduces the efficiency of currents to wipe out charge differences. The main approximation of **ideal MHD** is to assume **infinite conductivity**, or zero resistivity. Thus, we cancel the RHS term from Equation 10.23. This simplification is similar to ignoring viscosity and heat conduction in the Euler equations.

We have succeeded in writing equations that contain only our usual fluid quantities and \mathbf{B} , namely Equations 10.19 and 10.23 without its RHS. One fly in the ointment is the $1/4\pi$ factor in Equation 10.19. We absorb this factor into our base units for the magnetic field, i.e., we set our cgs magnetic units to $\sqrt{4\pi}$ G. More generally, in self-similar code units (§5.1), we have

$$B_0 \equiv \sqrt{\frac{4\pi m_0}{l_0 t_0^2}}. \quad (10.24)$$

With this final caveat out of the way, we collect the equations of ideal MHD:

Eulerian ideal MHD equations

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} + \frac{1}{\rho} (\nabla \times \mathbf{B}) \times \mathbf{B} - \nabla \Phi$$

$$\frac{\partial \varepsilon}{\partial t} + \mathbf{u} \cdot \nabla \varepsilon = -\frac{P}{\rho} \nabla \cdot \mathbf{u} + \frac{\Gamma}{\rho} - \frac{\Lambda}{\rho}$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) = 0$$

Note that the divergence-free condition of the magnetic field, $\nabla \cdot \mathbf{B} = 0$, must be enforced separately from the rest of the fluid system because the other equations do not automatically imply it! Maintaining a divergence-free field is one of the main challenges when solving the MHD equations numerically.

We can also express the MHD equations in conservation-law form, i.e., writing the spatial derivatives of the fluid quantities as divergences of conserved fluxes. For this purpose, we need to update our definition of the total energy to include the magnetic energy per unit volume,

$$E \equiv \rho \left(\frac{\mathbf{u}^2}{2} + \varepsilon + \frac{\mathbf{B}^2}{2} + \Phi \right) \quad (10.25)$$

The conservation laws follow after somewhat tedious vector math (see §4.3 in Goedbloed & Poedts 2004):

Ideal MHD equations in conservation-law form

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) &= 0 \\ \frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot \left(\rho \mathbf{u} \otimes \mathbf{u} + \mathbf{I} \left[P + \frac{\mathbf{B}^2}{2} \right] - \mathbf{B} \otimes \mathbf{B} \right) &= -\rho \nabla \Phi \\ \frac{\partial E}{\partial t} + \nabla \cdot \left(\left[E + P + \frac{\mathbf{B}^2}{2} \right] \mathbf{u} - \mathbf{u} \cdot \mathbf{B} \otimes \mathbf{B} \right) &= \rho \frac{\partial \Phi}{\partial t} + \Gamma - \Lambda \\ \frac{\partial \mathbf{B}}{\partial t} + \nabla \cdot (\mathbf{u} \otimes \mathbf{B} - \mathbf{B} \otimes \mathbf{u}) &= 0 \end{aligned}$$

10.5 Basic MHD dynamics: flux freezing and dynamos

Given the numerous curls and cross-products in the ideal MHD equations, it is not easy to develop an intuitive sense for how magnetic fields evolve. The most important takeaway is that magnetic fields and velocities are coupled: there are \mathbf{B} -terms in the equation for \mathbf{u} and vice versa. In the non-magnetic Euler equations, there was no mechanism to convert kinetic and thermal energy into one another (although viscosity provides such a mechanism in practice). Now, we can convert momentum into magnetic fields, and magnetic fields can accelerate the fluid.

In the momentum and energy equations, the thermal and magnetic terms “compete” with each other. Thus, we can get a first sense of the dynamics by comparing the relative size of those terms. This comparison is straightforward because we did not need to introduce a new electric scale unit into the problem. Equation 10.24 shows that the dimensions of the magnetic energy are

$$B_0^2 \propto \rho_0 u_0^2 = P_0 = E_0. \quad (10.26)$$

In other words, the square of the magnetic field has units of energy density or pressure! This allows us to directly compare the importance of the thermal and magnetic energy components with a dimensionless quantity called **plasma beta**,

$$\beta \equiv \frac{2P}{B^2} \approx 3.5 \left(\frac{T}{1000 \text{ K}} \right) \left(\frac{n}{\text{cm}^{-3}} \right) \left(\frac{B}{\mu\text{G}} \right)^{-2}. \quad (10.27)$$

System	n (cm ⁻³)	T (K)	B (G)	β
Magnetosphere of the Earth	10 ⁴	10 ⁴	0.3	4×10^{-6}
Coronal loop in the Sun	10 ¹⁰	10 ⁶	300	0.0004
Solar wind	10	10 ⁵	6×10^{-5}	1
Molecular cloud	100	20	5×10^{-6}	≈ 0.1
Warm HI in interstellar medium	1	10 ⁴	10 ⁻⁵	≈ 0.1
Galaxy cluster	10 ⁻³	10 ⁷	$10^{-6} - 10^{-5}$	≈ 1

Table 4: Rough estimates of typical properties of astrophysical plasmas. The values of β vary dramatically between systems. In the ISM, there is tentative evidence that magnetic fields are at least in equipartition with thermal pressure, though the uncertainties are large. In galaxy clusters, magnetic fields are thought to be important but perhaps not dominant. Some of the values for density, temperature, and magnetic field were taken from Goedbloed & Poedts (2004) and Draine (2011).

A large β means that a plasma is dominated by thermal energy, a small β that magnetic fields dominate the dynamics. The numbers in Equation 10.27 highlight that even small, μG magnetic fields can have a significant effect on the dynamics of typical astrophysical plasmas (see Table 4 for some examples).

Another fundamental insight into MHD dynamics is provided by the concept of “flux freezing.” Magnetic flux is defined as the amount of field lines piercing a 2D contour (not the closed surface of a fluid element!) that is moving with the fluid,

$$\Phi_B \equiv \oint \mathbf{B} \cdot d\mathbf{C}. \quad (10.28)$$

It can be shown that the only change in Φ_B is due to the diffusion term in Equation 10.23 (CC §13.3). Given that we assume infinite conductivity in ideal MHD, the **magnetic flux is frozen into a moving fluid** (Alfvén 1942). In other words, field lines and fluid move in unison, and field lines are made up by the same, moving particles. For example, if a magnetic plasma tube is compressed to a smaller radius, the magnetic flux density will increase to make up for the smaller area.

Flux freezing leads to a simple prediction for the scaling of the magnetic field with density. Imagine a spherical gas cloud with radius R_0 and fixed mass M that is initially threaded by a uniform magnetic field B_0 . The magnetic flux will be $\Phi_B \propto R_0^2 B_0$. For example, we can imagine the flux through the midplane of the spherical cloud. If the cloud collapses, e.g., due to gravity, we have $M \propto \rho R^3 = \text{const}$ and $\Phi_B = \text{const}$, and thus $B \propto \rho^{2/3}$ (Mestel 1965). This scaling has been confirmed numerically. For example, MHD simulations of the collapse of gas in the cosmic web show that the magnetic field in filaments roughly scales as $\rho^{2/3}$ (e.g., Marinacci et al. 2018).

This finding brings up a more fundamental question: where did the initial magnetic field come from in the first place? If we set $B = 0$ in the ideal MHD equations, there is no mechanism to create a magnetic field. The question of the origin of the “primordial seed fields” in the Universe is still subject to active research. It is thought that non-ideal processes such as the “Biermann battery” create extremely small seed fields, $B \approx 10^{-12}$ Gauss or smaller (Biermann 1950). The primordial seed fields are then amplified by so-called **magnetic dynamos**, an umbrella term for processes that convert kinetic to magnetic energy. The theory behind dynamo processes is complicated, but they are an almost ubiquitous phenomenon wherever magnetic fields are to be found: in stars such as the Sun, in the interiors of planets such as the Earth, and in galaxies. The magnetic field in galaxies is boosted by both small-scale turbulent motion and large-scale rotation, and is thought

to reach near-equipartition with thermal energy after a few Gyr (e.g. Wang & Abel 2009; Hanasz et al. 2009; Pakmor & Springel 2013).

10.6 MHD waves

Given that we have added new equations to the Euler system, we expect new characteristic wave speeds to emerge and the existing ones to change (§B.2). Indeed, the landscape of MHD waves is significantly richer than the sound waves we have encountered so far. Most notably, the propagation of waves now depends on whether they are parallel or perpendicular to the magnetic field direction. Since the full linearization procedure is extremely tedious (CC §13.4), we only quote the most important results.

First, the presence of magnetic fields enables an entirely new type of magneto-inertial waves called **Alfvén waves** (Alfvén 1942). The magnetic field acts as a tension that restores oscillations in density. This wave propagates transverse to the magnetic field direction, since there is no force along the field lines (unlike in sound waves, which are longitudinal). Following the pattern that $c_{\text{wave}}^2 = \text{restoring force/inertia}$ (§4.2), and given that B^2 is the magnetic equivalent of pressure, it makes sense that these waves propagate at the **Alfvén speed**,

$$c_a = \frac{|B|}{\sqrt{\rho}} \quad (10.29)$$

Second, the usual sound speed c_s is altered depending on the orientation relative to the magnetic field. It splits into the **fast and slow magneto-acoustic speeds**,

$$c_{\text{fs,ss}} = \sqrt{\frac{1}{2\rho} \left[\gamma P + B^2 \pm \sqrt{(\gamma P - B^2)^2 + 4\gamma P B_{\perp}^2} \right]}, \quad (10.30)$$

where B_{\perp} is the field component perpendicular to the wave vector. This expression looks complicated, but we can understand it based on a few limiting cases:

- when propagating along the field, we have $B_{\perp} \rightarrow 0$ and $c_{\text{fs}} \rightarrow c_s$ and $c_{\text{ss}} \rightarrow c_a$
- when propagating perpendicular to the field, we have $B_{\perp} \rightarrow B$ and thus $c_{\text{fs}} \rightarrow c_s + c_a$ and $c_{\text{ss}} \rightarrow 0$
- when $B^2 \ll P$, $c_{\text{fs}} \rightarrow c_s$ as expected
- when $B^2 \gg P$, $c_{\text{fs}} \rightarrow c_a$.

Either way, the wave speeds present in the solution are u , $u \pm c_{\text{ss}}$, $u \pm c_{\text{fs}}$, and $u \pm c_a$. They are always ordered in magnitude as

$$0 \leq c_{\text{ss}}^2 \leq c_a^2 \leq c_{\text{fs}}^2, \quad (10.31)$$

meaning that $|u| + c_{\text{fs}}$ is the fastest possible speed. In regions where $\beta \gg 1$, the Alfvén speed dominates, which means that the timestep in simulations of strongly magnetized plasmas will be set by the ratio of magnetic field to density rather than by the temperature.

10.7 Observing magnetic fields via the rotation measure

What is our knowledge of astrophysical magnetic fields based on? We cannot observe magnetic energy directly, but the **Faraday effect** comes to the rescue: the polarization of light is rotated by magnetic fields, giving rise to the **rotation measure**,

$$\text{RM} \propto \int_0^L n_e B_{\parallel} dl, \quad (10.32)$$

where n_e is the electron density and B_{\parallel} is the magnetic field component along the line of sight. By itself, observing the RM would not be terribly useful without knowing the intrinsic polarization of the source. Luckily, the actual change in polarization is proportional to $\lambda^2 \text{RM}$, meaning that we can observe the RM at two different wavelengths and thus eliminate the unknown intrinsic polarization.

11 Turbulence

Throughout this course, we have mostly considered orderly flows, in the sense that they could be encapsulated in some sort of mathematical description. When we arrived at fluid instabilities in §9, this picture started to break down: as soon as the instabilities begin to grow non-linearly, our first-order mathematical description is worthless. For example, the Kelvin-Helmholtz instability will grow into an arbitrarily complex network of whirls, given enough time.

In reality, most terrestrial and astrophysical flows are, at least to some extent, chaotic in nature. We call such flows “turbulent.” Prominent examples include the Earth’s atmosphere, where turbulence makes it impossible to predict the weather more than a week into the future, or the interstellar medium in galaxies, where supernova blastwaves keep stirring turbulent gas motions. Fortunately, all hope is not lost because, under certain conditions, turbulent systems can still be described statistically. In this chapter, we consider the very basics of this viewpoint. See, e.g., Lequeux 2005 §13 for a more thorough, yet accessible, treatment.

11.1 Self-similarity and the turbulent cascade

To explain what turbulence is, we can use the Kelvin-Helmholtz instability as an example: two parts of a fluid that flow past each other get entangled and create eddies (Figure 34). But what if the edges of those eddies are themselves Kelvin-Helmholtz unstable? They will generate new, smaller eddies, and so on. This **turbulent cascade** will only stop at a size scale where the fluid is stabilized against forming new KH eddies (Figure 35). We saw that this can happen due to gravity for fluids of different densities. However, we imagine a scenario where gravity points in no particular direction and where the morphology of the flow changes quickly, so that gravity is not necessarily acting perpendicular to fluid interfaces. Instead, the mechanism that typically cuts off the turbulent cascade is **viscosity**, which converts kinetic energy into thermal energy and thus diffuses eddies. This picture was first proposed by Lewis Richardson, who encapsulated the idea in a rhyme: *Big whirls have little whirls that feed on their velocity, and little whirls have lesser whirls and so on to viscosity.*

To understand the viscous cutoff of the turbulent cascade, we need to briefly review the very basics of viscosity (see CC §11 and vdB §10 for more details). Viscosity arises due to the exchange of momentum between neighboring fluid elements. Without derivation, we state that its main consequence is often an additional term in the velocity equation (which technically makes it a “Navier-Stokes equation”),

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla P}{\rho} + \frac{\eta}{\rho} \nabla^2 \mathbf{u}. \quad (11.1)$$

The η term is one of a number of terms that arise in the full treatment of viscosity, but it is the one that is most relevant for our purposes. It reminds us of the heat equation, where the Laplacian derivative serves to smooth out the quantity in question because it increases it in troughs (regions with positive curvature) and vice versa. The fluid property η is called the “dynamical coefficient of viscosity” and has units of g/cm/s. To gain physical insight into its effects, we non-dimensionalize the equation (§5.1) and rearrange it slightly (see also §5.1),

$$\tilde{\rho} \left(\frac{\partial \tilde{\mathbf{u}}}{\partial \tilde{t}} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\mathbf{u}} \right) = -\tilde{\nabla} \tilde{P} + \frac{1}{\text{Re}} \tilde{\nabla}^2 \tilde{\mathbf{u}}, \quad (11.2)$$

where the **Reynolds number** is defined as

$$\boxed{\frac{1}{\text{Re}} = \frac{\eta}{\rho_0 l_0 u_0}} \quad (11.3)$$




Figure 35: Turbulence develops in an initially ordered flow from left to right. The Reynolds number is about 2300, which is high enough for a significant cascade to develop, but sufficiently low that the viscous scale becomes apparent towards the right of the image. Figure from van Dyke (1982).

This dimensionless number tells us how important viscosity is compared to pressure gradients. In the following, we will consider **incompressible turbulence**, where $\rho_0 = \rho = \text{const.}$ While ρ_0 and η reflect the properties of a given fluid, l_0 and u_0 correspond to scales that depend on the problem at hand.

11.2 The Kolmogorov-Obukhov law

In the context of turbulence, l_0 represents the scale of an eddy and u_0 the velocity with which fluid moves in eddies of that size. The Reynolds number evolves as we march down the turbulent cascade, but we need to relate l_0 and u_0 to understand how. From now on, we will use the notation $l = l_0$ for size scales and $u_l = u_0(l)$ for the velocity at that scale. We will denote the Reynolds number corresponding to a given scale as $\text{Re}(l) = lu_l\rho/\eta$.

To make progress, we make three further assumptions. First, we consider only **fully developed turbulence**, meaning a system in statistical equilibrium. We do not concern ourselves with how the turbulence started or developed. Second, we assert that the cascade is **self-similar**, which makes sense given that the Euler equations are self-similar (§5.1). The only scale we have introduced is Re , which cuts off the cascade at the **diffusion scale**, l_d , where the velocity is u_d and $\text{Re}(l_d) = l_d u_d \rho / \eta \approx 1$. At the opposite end of the size scale, we imagine a large scale L where kinetic energy is injected. The range of scales in between, where $l_d \ll l \ll L$, is called the **inertial range** and must be self-similar in the absence of any additional physical scales (Figure 36). Third, we assume that the dynamics at a scale l are dominated by that scale and nearby scales, i.e., that the dynamics is **local in Fourier space**.

If the cascade is in equilibrium, the rate of kinetic energy transfer between scales must be constant in time and the same across all scales. For example, imagine a fan circulating the air in a room. The fan’s energy cascades down to the viscous scale and is converted to thermal energy. Assuming that this thermal energy diffuses out of the room (e.g., by conduction keeping the temperature fixed), the system is in equilibrium. Thus, we assert that the energy transfer rate per unit mass, ξ , is constant across scales. This rate must be proportional to the kinetic energy per unit mass at each scale, $\varepsilon_{\text{kin}} \propto u_l^2$, divided by some timescale t_l . Dimensionally, the only possibility for this timescale is $t_l = l/u_l$, which corresponds to the “eddy turnover time.” Thus,

$$\xi \propto \frac{\varepsilon_{\text{kin}}}{t_l} \propto \frac{u_l^2}{t_l} \propto \frac{u_l^3}{l} \quad \implies \quad u_l \propto (\xi l)^{1/3}. \quad (11.4)$$

Regardless of whether this argument seems plausible, it gives the only dimensionally correct com-




Figure 36: Schematic of the Kolmogorov-Obukhov turbulent power spectrum. The problem is self-similar in the inertial range between L and l_d , and the energy per scale follows a power law.

bination of variables that forms the desired energy transfer rate. Richardson (1926) also derived this scaling observationally. We can now understand the extent of the self-similar cascade more quantitatively. With the relation between l and u_l , we can establish a direct link between l and $\text{Re}(l)$,

$$\text{Re}(l) = \frac{lu_l\rho}{\eta} = \frac{\xi^{1/3}l^{4/3}\rho}{\eta} \implies l = \left(\frac{\eta\text{Re}(l)}{\rho}\right)^{3/4} \xi^{-1/4} \propto \text{Re}(l)^{3/4}, \quad (11.5)$$

where the last proportionality highlights that η , ρ , and ξ are constants (independent of l). Since $\text{Re}(l_d) \approx 1$ by definition, we have

$$\frac{L}{l_d} \propto \left(\frac{\text{Re}(L)}{\text{Re}(l_d)}\right)^{3/4} \propto \text{Re}(L)^{3/4}. \quad (11.6)$$

This makes sense: the larger the Reynolds number at the driving scale, the more levels of l there are before we encounter the diffusion scale. To obtain the famous Kolmogorov cascade,¹⁴ we switch to a Fourier space picture by defining the wavenumber $k \equiv 2\pi/l$ and the energy $E(k)$ in an infinitesimal interval k to $k + dk$. Expressed in these quantities, the average kinetic energy at scale l must be the integral over the contributions from all smaller scales,

$$\langle u_l^2 \rangle = \int_k^\infty E(k) dk \propto (\xi l)^{2/3}. \quad (11.7)$$

The logic for including all smaller scales is similar to, say, measuring the variance of a density field: all smaller scales contribute to the variance on a given scale, but larger ones do not because they would add an overall density or, in our case, a bulk flow that we are not counting. Thus, the average velocity $\langle u_l^2 \rangle$ is not the same as the typical velocity at the scale; the former must be lower since u_l decreases with l . However, dimensionally, we can still assume $\langle u_l^2 \rangle \propto (\xi l)^{1/3}$ because

¹⁴While Andrey Kolmogorov's name is synonymous with this result, Alexander Obukhov appears to have made similarly important contributions in collaboration with Kolmogorov. In fact, it was Obukhov who first wrote down the famous $k^{-5/3}$ scaling. While his name has historically been overlooked in the west, he is commonly credited in the Russian literature; see, e.g., Landau & Lifshitz (1987) §33 and notes by Andrey Kravtsov (private communication).

the difference between the two velocities cannot be scale-dependent if the problem is self-similar. Converting to k units we get

$$\int_k^\infty E(k)dk \propto (\xi l)^{2/3} \propto \xi^{2/3} k^{-2/3}. \quad (11.8)$$

We differentiate this expression with respect to k to obtain the **Kolmogorov-Obukhov law**,

$$\boxed{E(k) \propto \xi^{2/3} k^{-5/3}} \quad (11.9)$$

Figure 36 shows a schematic overview of the turbulent power spectrum. The Kolmogorov-Obukhov scaling has been confirmed in numerous experiments and simulations. While this result is extremely elegant, it describes only incompressible, non-MHD, fully developed turbulent systems in statistical equilibrium. Modern simulations of compressible MHD turbulence do indeed find different power spectra.

Appendix

A Mathematical Background

We use Q and \mathbf{Q} to denote arbitrary scalar and vector fields, and \mathbf{v} and \mathbf{w} if there is more than one field. All vectors are assumed to be spatial 3-vectors, $\mathbf{Q} = (Q_x, Q_y, Q_z) = Q_x \hat{\mathbf{x}} + Q_y \hat{\mathbf{y}} + Q_z \hat{\mathbf{z}}$.

A.1 Vector operators

We begin with fundamental operations that involve two vectors. First, the dot product returns a scalar

$$\mathbf{v} \cdot \mathbf{w} = v_x w_x + v_y w_y + v_z w_z, \quad (\text{A.1})$$

which tells us to what extent the two vectors are aligned ($\mathbf{v} \cdot \mathbf{w} = 0$ if they are orthogonal, and the product of their lengths if they are perfectly aligned). The second vector operation is the cross product,

$$\mathbf{v} \times \mathbf{w} = \begin{vmatrix} \hat{\mathbf{x}} & \hat{\mathbf{y}} & \hat{\mathbf{z}} \\ v_x & v_y & v_z \\ w_x & w_y & w_z \end{vmatrix} = (v_y w_z - v_z w_y) \hat{\mathbf{x}} + (v_z w_x - v_x w_z) \hat{\mathbf{y}} + (v_x w_y - v_y w_x) \hat{\mathbf{z}}. \quad (\text{A.2})$$

Differential vector operators follow the same logic. The gradient operator is

$$\nabla \equiv \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z} \right) \quad (\text{A.3})$$

so that the gradient of a scalar field is

$$\nabla Q = \left(\frac{\partial Q}{\partial x}, \frac{\partial Q}{\partial y}, \frac{\partial Q}{\partial z} \right) = \frac{\partial Q}{\partial x} \hat{\mathbf{x}} + \frac{\partial Q}{\partial y} \hat{\mathbf{y}} + \frac{\partial Q}{\partial z} \hat{\mathbf{z}}. \quad (\text{A.4})$$

The gradient of a vector field, $\nabla \mathbf{Q}$, is a tensor that is easiest understood in index notation (§A.2). The divergence of a vector field is defined as

$$\nabla \cdot \mathbf{Q} = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z} \right) \cdot (Q_x, Q_y, Q_z) = \frac{\partial Q_x}{\partial x} + \frac{\partial Q_y}{\partial y} + \frac{\partial Q_z}{\partial z}. \quad (\text{A.5})$$

Similarly, the curl follows the rule for the cross product,

$$\nabla \times \mathbf{Q} = \begin{vmatrix} \hat{\mathbf{x}} & \hat{\mathbf{y}} & \hat{\mathbf{z}} \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ Q_x & Q_y & Q_z \end{vmatrix} = \left(\frac{\partial Q_z}{\partial y} - \frac{\partial Q_y}{\partial z} \right) \hat{\mathbf{x}} + \left(\frac{\partial Q_x}{\partial z} - \frac{\partial Q_z}{\partial x} \right) \hat{\mathbf{y}} + \left(\frac{\partial Q_y}{\partial x} - \frac{\partial Q_x}{\partial y} \right) \hat{\mathbf{z}}. \quad (\text{A.6})$$

Finally, the Laplacian is defined as

$$\nabla^2 Q = (\nabla \cdot \nabla) Q = \frac{\partial^2 Q}{\partial x^2} + \frac{\partial^2 Q}{\partial y^2} + \frac{\partial^2 Q}{\partial z^2}. \quad (\text{A.7})$$

A.2 Index notation and tensors

While most equations in these notes are written in vector notation, the frequently used Einstein summation convention is clearer in some cases. Here, we assume that any repeated index is summed

over in all three dimensions (x, y, z). Any remaining index indicates that the result has vector dimensions (rather than scalar). In this notation, the dot product is

$$\mathbf{v} \cdot \mathbf{w} = v_i w_i \equiv \sum_i v_i w_i = v_x w_x + v_y w_y + v_z w_z. \quad (\text{A.8})$$

The ∇ operator takes on the particularly simple form

$$\nabla = \frac{\partial}{\partial x_i}, \quad (\text{A.9})$$

which can then be applied as a gradient,

$$\nabla Q = \frac{\partial Q}{\partial x_i} = \left(\frac{\partial Q}{\partial x}, \frac{\partial Q}{\partial y}, \frac{\partial Q}{\partial z} \right), \quad (\text{A.10})$$

which is still a vector because there are no matching indices to sum over. The divergence, however, does introduce matching indices and is thus reduced to a scalar,

$$\nabla \cdot \mathbf{Q} = \frac{\partial v_i}{\partial x_i} = \frac{\partial Q_x}{\partial x} + \frac{\partial Q_y}{\partial y} + \frac{\partial Q_z}{\partial z}. \quad (\text{A.11})$$

The index notation becomes particularly useful when we are dealing with objects with more than one index, such as tensors like

$$\mathbf{vw} \equiv (\mathbf{v} \otimes \mathbf{w})_{ij} = v_i w_j. \quad (\text{A.12})$$

The gradient operator for tensors follows the same rule,

$$\nabla \mathbf{Q} = \frac{\partial Q_i}{\partial x_j} \quad (\text{A.13})$$

and so on for the divergence and curl operators. We sometimes wish to recast a gradient as a divergence, for example in the derivation of the conservation form of the momentum equation (§B.1). We can use the identity matrix to make a tensor out of a scalar field, where

$$\nabla \cdot (\mathbf{I}Q) = \frac{\partial}{\partial x_i} (\delta_{ij} Q)_i = \frac{\partial Q}{\partial x_j} = \nabla Q \quad (\text{A.14})$$

Here, the Kronecker δ_{ij} is unity if $i = j$ and zero otherwise.

A.3 Vector identities

In this section, we list some vector-related identities relevant to this course. One foundational relation is the **divergence theorem**, which says that we can think of divergence as the volume equivalent of a vector field \mathbf{Q} 's flux through the surfaces enclosing the volume,

$$\int_S \mathbf{Q} \cdot \hat{\mathbf{n}} dS \equiv \int_S \mathbf{Q} \cdot d\mathbf{S} = \int_V \nabla \cdot \mathbf{Q} dV. \quad (\text{A.15})$$

When working with divergence, we frequently encounter the **product rule for divergence**,

$$\nabla \cdot (Q\mathbf{v}) = Q\nabla \cdot \mathbf{v} + \mathbf{v} \cdot \nabla Q. \quad (\text{A.16})$$

The **curl of a gradient** is always zero,

$$\nabla \times (\nabla Q) = 0, \quad (\text{A.17})$$

as is the **divergence of curl**,

$$\nabla \cdot (\nabla \times \mathbf{Q}) = 0, \quad (\text{A.18})$$

but the **cross-product with curl** is not,

$$\mathbf{Q} \times (\nabla \times \mathbf{Q}) = \nabla \left(\frac{1}{2} Q^2 \right) - \mathbf{Q} \cdot \nabla \mathbf{Q} \quad (\text{A.19})$$

and neither is the curl of the curl,

$$\nabla \times (\nabla \times \mathbf{Q}) = \nabla(\nabla \cdot \mathbf{Q}) - \nabla^2 \mathbf{Q}. \quad (\text{A.20})$$

A.4 Eigenvalues and eigenvectors

If an $n \times n$ matrix \mathbf{A} is diagonalizable, there are n **eigenvalues** λ_i and eigenvectors \mathbf{V}_i such that

$$\mathbf{A}\mathbf{V}_i = \lambda_i \mathbf{V}_i. \quad (\text{A.21})$$

The eigenvalues can be found from the determinant, $|\mathbf{A} - \lambda \mathbf{I}| = 0$. For more details, see any textbook on linear algebra.

B Derivations

In this appendix, we go through some of the more tedious derivations that, while important, would interrupt the flow of the main text.

B.1 Conservation-law form of the Euler equations

The goal of this derivation is to convert the Eulerian equations for velocity (2.15) and internal energy (2.22) into conservation laws similar to the continuity equation (2.6). When we recast the equation for \mathbf{u} as an equation for $\rho\mathbf{u}$, we get

$$\frac{\partial(\rho\mathbf{u})}{\partial t} = \rho \frac{\partial\mathbf{u}}{\partial t} + \mathbf{u} \frac{\partial\rho}{\partial t} = -\rho\mathbf{u} \cdot \nabla\mathbf{u} - \mathbf{u}(\mathbf{u} \cdot \nabla\rho) - \mathbf{u}\rho\nabla \cdot \mathbf{u} - \nabla P - \rho\nabla\Phi, \quad (\text{B.1})$$

where we have substituted the expressions from the continuity and “momentum” equations 2.6 and 2.15 and multiplied through. The \mathbf{u} -terms in this equation seem difficult to interpret until we write them in index notation,

$$\rho\mathbf{u} \cdot \nabla\mathbf{u} + \mathbf{u}(\mathbf{u} \cdot \nabla\rho) + \mathbf{u}\rho\nabla \cdot \mathbf{u} = \rho u_i \frac{\partial u_j}{\partial x_i} + u_i u_j \frac{\partial \rho}{\partial x_i} + \rho u_j \frac{\partial u_i}{\partial x_i} = \frac{\partial}{\partial x_i} (\rho u_i u_j). \quad (\text{B.2})$$

The object on the right is a tensor, $\mathbf{u} \otimes \mathbf{u} = u_i u_j$, and the contraction of the derivative with the first index corresponds to the divergence. We also include the pressure term into the so-called momentum flux density tensor (for ideal fluids),

$$\Pi_{ij} \equiv \rho u_i u_j + P\delta_{ij} = \rho\mathbf{u} \otimes \mathbf{u} + \mathbf{I}P \quad (\text{B.3})$$

where \mathbf{I} is the identity matrix (§A.2). We can now rewrite Equation B.1 as a conservation law,

$$\boxed{\frac{\partial(\rho\mathbf{u})}{\partial t} + \nabla \cdot (\rho\mathbf{u} \otimes \mathbf{u} + \mathbf{I}P) = -\rho\nabla\Phi} \quad (\text{B.4})$$

To derive a total energy equation, we recall the definition $E = \rho(|\mathbf{u}|^2/2 + \varepsilon + \Phi)$ and write

$$\frac{DE}{Dt} = \frac{E}{\rho} \frac{D\rho}{Dt} + \rho \left(\mathbf{u} \cdot \frac{D\mathbf{u}}{Dt} + \frac{D\varepsilon}{Dt} + \frac{D\Phi}{Dt} \right) = \frac{\partial E}{\partial t} + \mathbf{u} \cdot \nabla E. \quad (\text{B.5})$$

We rearrange and substitute the RHS of the Lagrangian fluid equations to find

$$\begin{aligned} \frac{\partial E}{\partial t} &= -\mathbf{u} \cdot \nabla E - E\nabla \cdot \mathbf{u} - \mathbf{u} \cdot \nabla P - \rho\mathbf{u} \cdot \nabla\Phi + \rho \frac{D\Phi}{Dt} - P\nabla \cdot \mathbf{u} + \Gamma - \Lambda \\ &= -\mathbf{u} \cdot \nabla(E + P) - (E + P)\nabla \cdot \mathbf{u} + \rho \frac{\partial\Phi}{\partial t} + \Gamma - \Lambda \\ &= -\nabla \cdot [(E + P)\mathbf{u}] + \rho \frac{\partial\Phi}{\partial t} + \Gamma - \Lambda \end{aligned} \quad (\text{B.6})$$

and thus

$$\boxed{\frac{\partial E}{\partial t} + \nabla \cdot [(E + P)\mathbf{u}] = \rho \frac{\partial\Phi}{\partial t} + \Gamma - \Lambda} \quad (\text{B.7})$$

The source terms on the RHS include possible time changes in the gravitational potential, as well as heating and cooling.

B.2 The matrix form of the fluid equations and the eigenvalue perspective

In this section, we use linear algebra to understand a fundamental property of the Euler equations: the characteristic speeds with which different waves and perturbations travel. In §5.3, we saw that a quantity purely moving along with some velocity is described by the advection equation. Thus, we will decompose the full Euler equations into advection-like parts. We begin with the conservation-law form and recall Equation 5.10,

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathcal{F}(\mathbf{U}) = \mathbf{S}. \quad (\text{B.8})$$

For simplicity, we will set $\mathbf{S} = 0$ in this section, i.e., assume that there is no heating, cooling, or gravity. In one dimension, the conservation law form now reads

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathcal{F}(\mathbf{U})}{\partial x} = 0. \quad (\text{B.9})$$

We apply the chain rule to the derivative to write

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathcal{F}}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial x} = \frac{\partial \mathbf{U}}{\partial t} + \mathbf{A}(\mathbf{U}) \frac{\partial \mathbf{U}}{\partial x} = 0, \quad (\text{B.10})$$

where we have introduced the **Jacobian matrix of the flux vector**, $\mathbf{A} \equiv \partial \mathcal{F} / \partial \mathbf{U}$, or $A_{ij} = \partial \mathcal{F}_i / \partial U_j$. We note a curious consequence of this formulation: the eigenvalues of \mathbf{A} have a special meaning. If \mathbf{A} is diagonalizable, there are eigenvectors of the fluid state, $\mathbf{U}^{(k)}$, and eigenvalues λ_k ,¹⁵

$$\mathbf{A} \mathbf{U}^{(k)} = \lambda_k \mathbf{U}^{(k)} \quad \implies \quad \mathbf{A} \frac{\partial \mathbf{U}^{(k)}}{\partial x} = \lambda_k \frac{\partial \mathbf{U}^{(k)}}{\partial x} \quad \implies \quad \frac{\partial \mathbf{U}^{(k)}}{\partial t} + \lambda_k \frac{\partial \mathbf{U}^{(k)}}{\partial x} = 0. \quad (\text{B.11})$$

The last expression represents k advection equations with speeds λ_k ! We conclude that **the eigenvalues of the Jacobian of the fluxes are the wave speeds** that are present in a given physical setup. For the 1D advection equation, there is only the fluid speed u , indicating a wave moving to the left or right.

The eigenvalue technique generalizes to any equation that looks like an advection equation; it does not necessarily have to refer to the conserved flux vector \mathcal{F} . For example, we can consider the slightly more complicated, non-conservative linearized fluid equations 4.4 and 4.6 that we used to derive sound waves,

$$\begin{aligned} \frac{\partial \rho_1}{\partial t} + \rho_0 \nabla \cdot \mathbf{u}_1 &= 0 \\ \frac{\partial \mathbf{u}_1}{\partial t} + \left(\frac{\partial P}{\partial \rho} \right)_0 \frac{\nabla \rho_1}{\rho_0} &= 0. \end{aligned}$$

We convert to 1D by writing $\nabla \rightarrow \partial / \partial x$ and $\mathbf{u}_1 \rightarrow u$, and we insert the ideal gas sound speed $c_s = \sqrt{\partial P / \partial \rho}$,

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \rho_0 \frac{\partial u}{\partial x} &= 0 \\ \frac{\partial u}{\partial t} + \frac{c_s^2}{\rho_0} \frac{\partial \rho}{\partial x} &= 0. \end{aligned} \quad (\text{B.12})$$

¹⁵We use superscript notation to not confuse the numbering of eigenvectors with their components, as in A_{ij} and so on. The eigenvalues are scalars, so there is no confusion.

We can write this as

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{A}(\mathbf{V}) \frac{\partial \mathbf{V}}{\partial x} = 0 \quad (\text{B.13})$$

where $\mathbf{V} \equiv (\rho, u)$ and

$$\mathbf{A}(\mathbf{V}) = \begin{pmatrix} 0 & \rho_0 \\ c_s^2/\rho_0 & 0 \end{pmatrix}.$$

We use the symbol \mathbf{V} rather than \mathbf{U} to indicate that the variables are not conserved in this case (since the second equation is in u rather than ρu). Nevertheless, the rest of the math still works out the same. To find the wave speeds in the problem, we calculate the eigenvalues using the usual determinant procedure,

$$|\mathbf{A} - \lambda \mathbf{I}| = 0 \quad \implies \quad \begin{vmatrix} 0 - \lambda & \rho_0 \\ c_s^2/\rho_0 & 0 - \lambda \end{vmatrix} = \lambda^2 - c_s^2 = 0 \quad (\text{B.14})$$

and thus (see, e.g., [Toro §2.1](#))

$$\lambda_1 = +c_s, \quad \lambda_2 = -c_s, \quad \mathbf{V}^{(1)} = (\rho_0, -c_s), \quad \mathbf{V}^{(2)} = (\rho_0, +c_s). \quad (\text{B.15})$$

This tells us that there are two possible waves in the problem, going to the left or right with speed c_s . That's exactly what we would have expected for our wave system! Note that there is no “wave” with speed u here because we set up the problem such that the overall velocity $u_0 = 0$.

We can find a similar formulation for the full Euler equations in 1D (and with $\mathbf{S} = 0$). This can, in principle, be done either in conservative variables, $\mathbf{U} \equiv (\rho, \rho u, E)$, or in primitive variables, $\mathbf{V} \equiv (\rho, u, P)$; the math works out easier in the latter case. Since we are using P instead of ε , we need to convert the internal energy equation to a pressure equation. We can do this as long as we assume the equation of state for an ideal gas. The derivation is a little lengthy; we find

$$\frac{\partial P}{\partial t} + \mathbf{u} \cdot \nabla P + \gamma P \nabla \cdot \mathbf{u} = 0. \quad (\text{B.16})$$

We use that $\gamma P = c_s^2 \rho$ for an ideal fluid to replace pressure with the sound speed. Adding the continuity and velocity equations from [§2.5](#), we can write down our primitive system in 1D,

$$\boxed{\begin{aligned} \frac{\partial \rho}{\partial t} + u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} &= 0 \\ \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + \frac{1}{\rho} \frac{\partial P}{\partial x} &= 0 \\ \frac{\partial P}{\partial t} + u \frac{\partial P}{\partial x} + c_s^2 \rho \frac{\partial u}{\partial x} &= 0 \end{aligned}}$$

This system is now in the desired form of Equation [B.13](#): a time derivative of each quantity plus a combination of spatial derivatives summing to zero. We note that the first spatial term in each equation is already suggestive of advection, namely, the velocity u times the spatial derivative of the quantity in question. Thus, the 3×3 matrix \mathbf{A} has u 's along its diagonal,

$$\mathbf{A}(\mathbf{V}) = \begin{pmatrix} u & \rho & 0 \\ 0 & u & 1/\rho \\ 0 & c_s^2/\rho & u \end{pmatrix}. \quad (\text{B.18})$$

We compute the eigenvalues as in Equation B.14,

$$|\mathbf{A} - \lambda \mathbf{I}| = \begin{vmatrix} u - \lambda & \rho & 0 \\ 0 & u - \lambda & 1/\rho \\ 0 & c_s^2/\rho & u - \lambda \end{vmatrix} = (u - \lambda)(u^2 - 2u\lambda + \lambda^2 - c_s^2) = 0. \quad (\text{B.19})$$

Although not immediately obvious, this equation has three eigenvalue solutions

$$\lambda_1 = u - c_s, \quad \lambda_2 = u, \quad \lambda_3 = u + c_s. \quad (\text{B.20})$$

As expected, we have three characteristic speeds now: the fluid velocity u as well as sound waves going to the left and right, which are added to u . We could continue by analyzing the left and right eigenvectors, but their meaning is not as easy to discern (Toro §3.1.2).

For many practical applications, e.g., the numerical hydro schemes described in §8.5, we wish to extend the matrix formalism to higher dimensions. There, things get a little confusing with the different velocity components. We split the equation into two operators,

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{A}_x(\mathbf{V}) \frac{\partial \mathbf{V}}{\partial x} + \mathbf{A}_y(\mathbf{V}) \frac{\partial \mathbf{V}}{\partial y} = 0 \quad (\text{B.21})$$

with separate matrices

$$\mathbf{A}_x(\mathbf{V}) = \begin{pmatrix} u_x & \rho & 0 & 0 \\ 0 & u_x & 0 & 1/\rho \\ 0 & 0 & u_x & 0 \\ 0 & c_s^2/\rho & 0 & u_x \end{pmatrix} \quad \mathbf{A}_y(\mathbf{V}) = \begin{pmatrix} u_y & 0 & \rho & 0 \\ 0 & u_y & 0 & 0 \\ 0 & 0 & u_y & 1/\rho \\ 0 & 0 & c_s^2/\rho & u_y \end{pmatrix}. \quad (\text{B.22})$$

For each of these matrices, we find additional “eigen-speeds” u , e.g.,

$$\lambda_1 = u_x - c_s, \quad \lambda_2 = \lambda_3 = u_x, \quad \lambda_4 = u_x + c_s \quad (\text{B.23})$$

for \mathbf{A}_x , and so on. The extra eigenvalues tell us that the fluid can travel in both the x and y directions. We proceed similarly to extend the formalism to 3D (e.g., Toro §3.2), or even to the ideal MHD system of equations (e.g., Mignone & Tzeferacos 2010).

B.3 Dispersion relation of perturbations at a two-fluid interface

Our goal is to derive a dispersion relation between the wavenumber k and frequency ω of a wave-like perturbation of a two-fluid interface (Figure 32). This derivation follows CC §10.1.2. The first key insight is that the flow is irrotational, meaning that its vorticity is zero. According to Kelvin’s circulation theorem, a steady, adiabatic flow conserves its vorticity, so that it stays irrotational if it starts out that way. To enforce this condition throughout our calculations, we express the velocities as gradients of an imaginary “velocity potential,” $\mathbf{u} \equiv -\nabla \Psi$. The continuity equation is zero for incompressible fluids, and the energy equation will not help us since the fluids are in pressure equilibrium, which means that their relative internal energy depends only on the densities. The velocity equation, however, takes on an interesting form,

$$\frac{\partial(-\nabla \Psi)}{\partial t} + \nabla \left(\frac{\mathbf{u}^2}{2} \right) = -\frac{\nabla P}{\rho} - \nabla \Phi, \quad (\text{B.24})$$

where we have written $\mathbf{u} \cdot \nabla \mathbf{u}$ as the gradient of $\mathbf{u}^2/2$. Note that throughout this derivation the vector \mathbf{u} includes the velocity perturbation, so $\mathbf{u}_1 \neq u_1$ and $\mathbf{u}_2 \neq u_2$; thus, we write \mathbf{u}^2 rather than u^2 for the total velocity. Since ρ is constant, we can pull it into the gradient of P so that all terms are gradients,

$$-\nabla \frac{\partial \Psi}{\partial t} + \nabla \left(\frac{\mathbf{u}^2}{2} \right) + \nabla \left(\frac{P}{\rho} \right) + \nabla \Phi = 0. \quad (\text{B.25})$$

We integrate this equation to remove the ∇ symbols, but an integration constant remains,

$$-\frac{\partial \Psi}{\partial t} + \frac{\mathbf{u}^2}{2} + \frac{P}{\rho} + \Phi = f(t). \quad (\text{B.26})$$

This expression is reminiscent of the Bernoulli constant, but we are allowing the velocity potential to change with time, at least to linear order. We will come back to this equation, but first we need to consider the functional form of Ψ . It is constrained by the incompressible condition, which, by the continuity equation, means that $\nabla \cdot \mathbf{u} = 0$ so that there are no inflows into and outflows out of fluid elements. This means that $\nabla^2 \Psi_1 = \nabla^2 \Psi_2 = 0$, so that the velocity potentials for the two fluids can only have linear terms,

$$\Psi_1 = -u_1 x + \psi_1(x, z, t) \quad \Psi_2 = -u_2 x + \psi_2(x, z, t). \quad (\text{B.27})$$

When differentiated, this gives us

$$\mathbf{u}_1 = -\nabla \Psi_1 = u_1 \hat{x} - \nabla \psi_1 \quad \mathbf{u}_2 = -\nabla \Psi_2 = u_2 \hat{x} - \nabla \psi_2, \quad (\text{B.28})$$

that is, the u_1 and u_2 velocities in the x -direction and no divergence as long as $\nabla^2 \psi_1 = \nabla^2 \psi_2 = 0$. We are now in a position to connect the velocity potential to the motion of the perturbation $\xi(x, t)$. Let's imagine a Lagrangian fluid element moving along the interface as part of the lower fluid. From the perspective of the fluid element, its Lagrangian derivative is just the up-down motion of the boundary, $-\partial \Psi_1 / \partial z$. We can also use the usual expression for the Eulerian-Lagrangian conversion, Equation 2.3,

$$\frac{D\xi}{Dt} = -\frac{\partial \psi_1}{\partial z} = \frac{\partial \xi}{\partial t} + \mathbf{u}_1 \cdot \nabla \xi \approx \frac{\partial \xi}{\partial t} + (u_1 \hat{x}) \cdot \nabla \xi = \frac{\partial \xi}{\partial t} + u_1 \frac{\partial \xi}{\partial x}. \quad (\text{B.29})$$

We have used that \mathbf{u}_1 is almost entirely in the x -direction since the perturbation is small; the equation above holds to first order. We get the same equation for the other fluid,

$$-\frac{\partial \psi_2}{\partial z} = \frac{\partial \xi}{\partial t} + u_2 \frac{\partial \xi}{\partial x}. \quad (\text{B.30})$$

We now have all the pieces in place, namely, three equations (B.26, B.29, B.30) that contain only the input quantities (ρ_1 , u_1 , etc) and the three unknowns ψ_1 , ψ_2 , and ξ . As usual, we express the perturbation as a general sum of waves,

$$\xi = A e^{i(kx - \omega t)}, \quad (\text{B.31})$$

but it suffices to analyze a single mode because different frequencies do not mix in a system of linear equations. The velocity potential associated with the perturbation has to have the same x and t dependence but could also contain z -dependent terms. We assume a general expression,

$$\psi_1 = C_1 e^{i(kx - \omega t) + h_1(z)}, \quad (\text{B.32})$$

where $h_1(z)$ is some function. The solution once again lies in the incompressibility condition, which told us that

$$\nabla^2 \psi_1 = 0 \quad \implies \quad (ik)^2 \psi_1 + \frac{\partial^2 h_1}{\partial z^2} \psi_1 = 0 \quad \implies \quad h_1(z) = \pm kz. \quad (\text{B.33})$$

We set the signs so that the perturbation stays small: for ψ_1 , $z < 0$, so $h_1 = +kz$ and, similarly, $h_2 = -kz$. Putting it all together we have

$$\psi_1 = C_1 e^{i(kx - \omega t) + kz} \quad \psi_2 = C_2 e^{i(kx - \omega t) - kz}. \quad (\text{B.34})$$

We insert these expressions (as well as the expression for ξ) into Equations B.29 and B.30 to find two dispersion relations,

$$-kC_1 = iA(ku_1 - \omega) \quad kC_2 = iA(ku_2 - \omega). \quad (\text{B.35})$$

However, we have three unknowns for two equations because we have not related the amplitude of the velocity perturbations (C_1 and C_2) to the amplitude of the perturbation in the boundary position (A). The missing piece of information is contained in Equation B.26, which relates the densities and velocities to the shared pressure and gravity. We convert it into an equation for pressure,

$$P = \rho_1 \left(\frac{\partial \psi_1}{\partial t} - \frac{\mathbf{u}_1^2}{2} - \Phi + f_1(t) \right) = \rho_2 \left(\frac{\partial \psi_2}{\partial t} - \frac{\mathbf{u}_2^2}{2} - \Phi + f_2(t) \right), \quad (\text{B.36})$$

where we have used that $\partial \Psi_1 / \partial t = \partial \psi_1 / \partial t$ and the same for Ψ_2 . Two terms are not in a useful form yet: we need to determine \mathbf{u} in terms of the perturbation and fix $f(t)$. We get the latter by considering the fluids far from the boundary where the perturbation vanishes, $\mathbf{u}_1 \rightarrow u_1$ and $\mathbf{u}_2 \rightarrow u_2$. Thus, we find $f_1(t) = u_1^2/2$ and $f_2(t) = u_2^2/2$. We approximate \mathbf{u}_1^2 to linear order,

$$\mathbf{u}_1^2 = (u_1 \hat{x} - \nabla \psi_1)^2 \approx u_1^2 - 2u_1 \frac{\partial \psi_1}{\partial x} \quad (\text{B.37})$$

and similarly for \mathbf{u}_2^2 . Finally, we put in the particular gravitational potential $\Phi = g\xi$ since we are considering fluid elements as they are moving up and down with the perturbation ξ . Putting it all together, Equation B.36 becomes

$$\rho_1 \left(\frac{\partial \psi_1}{\partial t} + u_1 \frac{\partial \psi_1}{\partial x} + g\xi \right) = \rho_2 \left(\frac{\partial \psi_2}{\partial t} + u_2 \frac{\partial \psi_2}{\partial x} + g\xi \right). \quad (\text{B.38})$$

We note that the velocity-squared terms from Equation B.37 and $f(t)$ have canceled to give us a linear equation once again. We substitute the waveform expressions for ξ and ψ to get our third relation between the properties of the perturbation,

$$\rho_1 [-iC_1(\omega - u_1 k) - gA] = \rho_2 [-iC_2(\omega - u_2 k) - gA], \quad (\text{B.39})$$

which we combine with Equation B.35 to find our dispersion relation,

$$\boxed{\frac{\omega}{k} = \frac{\rho_1 u_1 + \rho_2 u_2}{\rho_1 + \rho_2} \pm \left(\frac{g}{k} \frac{\rho_1 - \rho_2}{\rho_1 + \rho_2} - \frac{\rho_1 \rho_2 (u_1 - u_2)^2}{(\rho_1 + \rho_2)^2} \right)^{1/2}} \quad (\text{B.40})$$

C ULULA: a lightweight hydro code in Python

Given their focus on numerical hydrodynamics, these notes are accompanied by a 2D Python code called ULULA.¹⁶ The code is available in a public BitBucket repository, bitbucket.org/bdiemer/ulula, and the documentation can be found at bdiemer.bitbucket.io/ulula.

The purpose of ULULA is not to create a hydro code for research problems but to provide a simple, quick platform for experimenting with hydro algorithms beyond extremely simplified problems such as the 1D advection equation. On the other hand, ULULA is not meant to compete with already existing Python codes such as PYRO (Zingale 2014), which offers a sizable suite of hydro solvers (including complex schemes such as unsplit CTU). In a nutshell, the main design goals of ULULA are:

- **Brevity:** full-scale hydrodynamics codes tend to be large pieces of software that are hard to understand by looking at the code. The actual hydro solver of ULULA fits into a few hundred lines of Python.
- **Easy to use:** pre-implemented test problems can be executed with a couple of lines of code, and ULULA comes with a detailed online documentation.
- **Easy to understand:** in addition to the documentation, critical parts of the code are generously commented with the hope of making them easy to modify and extend.
- **Speed:** while ULULA is not meant for high-performance scientific simulations, execution speed is a major consideration. ULULA does not use Cython or other compiler extensions to Python, but all array operations use numpy indexing. As a result, typical test problems such as a Kelvin-Helmholtz instability can be run using a midsize domain (e.g., 200^2 cells) in under a minute on a modern laptop.

The code is written in an object-oriented fashion, with two main classes for initial setups and for the simulation itself. The recommended hydro scheme in ULULA is a dimensionally split MUSCL-Hancock algorithm with linear reconstruction, a choice of slope limiters, and an HLLC Riemann solver (a more advanced version of the HLL solver described in §8). A number of initial setups are pre-implemented, with a focus on the kinds of tests described in these notes.

¹⁶The name is a play on words. ULULA solves Euler's equations, whose name contains the German word for owl (*Eule*). Ulula means owl in Latin, the language in which Euler wrote most of his papers.

References

- Agertz, O., Moore, B., Stadel, J., et al. 2007, *MNRAS*, 380, 963
- Alfvén, H. 1942, *Nature*, 150, 405
- Almgren, A. S., Beckner, V. E., Bell, J. B., et al. 2010, *ApJ*, 715, 1221
- Biermann, L. 1950, *Zeitschrift Naturforschung Teil A*, 5, 65
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics: Second Edition* (Princeton University Press)
- Bryan, G. L., Norman, M. L., Stone, J. M., Cen, R., & Ostriker, J. P. 1995, *Computer Physics Communications*, 89, 149
- Bryan, G. L., Norman, M. L., O’Shea, B. W., et al. 2014, *ApJS*, 211, 19
- Clarke, C., & Carswell, B. 2014, *Principles of Astrophysical Fluid Dynamics* (Cambridge University Press)
- Colella, P. 1990, *Journal of Computational Physics*, 87, 171
- Colella, P., & Glaz, H. M. 1985, *Journal of Computational Physics*, 59, 264
- Colella, P., & Woodward, P. R. 1984, *Journal of Computational Physics*, 54, 174
- Courant, R., Friedrichs, K., & Lewy, H. 1928, *Mathematische Annalen*, 100, 32
- Draine, B. T. 2011, *Physics of the Interstellar and Intergalactic Medium* (Princeton University Press)
- Duffell, P. C. 2016, *ApJS*, 226, 2
- Duffell, P. C., & MacFadyen, A. I. 2011, *ApJS*, 197, 15
- Falco, M., Hansen, S. H., Wojtak, R., & Mamon, G. A. 2013, *MNRAS*, 431, L6
- Fryxell, B., Olson, K., Ricker, P., et al. 2000, *ApJS*, 131, 273
- Godunov, S. K. 1959, *Matematicheskii Sbornik*, 47, 271, english translation in: *J. Comput. Phys.*, vol. 3, 1965, pp. 251–266
- Goedbloed, J. P. H., & Poedts, S. 2004, *Principles of Magnetohydrodynamics* (Cambridge University Press)
- Hanasz, M., Wóltański, D., & Kowalik, K. 2009, *ApJ*, 706, L155
- Harten, A., Lax, P. D., & van Leer, B. 1983, *SIAM Review*, 25, 35
- Hopkins, P. F. 2015, *MNRAS*, 450, 53
- Khokhlov, A. 1998, *Journal of Computational Physics*, 143, 519
- Kravtsov, A. V., Klypin, A. A., & Khokhlov, A. M. 1997, *ApJS*, 111, 73
- Landau, L. D., & Lifshitz, E. M. 1987, *Fluid Mechanics*, 2nd edn. (Elsevier)
- Lequeux, J. 2005, *The Interstellar Medium* (Springer), doi:10.1007/b137959
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, 480, 5113
- Mestel, L. 1965, *QJRAS*, 6, 265
- Mignone, A., Bodo, G., Massaglia, S., et al. 2007, *ApJS*, 170, 228
- Mignone, A., & Tzeferacos, P. 2010, *Journal of Computational Physics*, 229, 2117
- Mignone, A., Zanni, C., Tzeferacos, P., et al. 2012, *The Astrophysical Journal Supplement Series*, 198, 7
- Pakmor, R., & Springel, V. 2013, *MNRAS*, 432, 176
- Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJS*, 192, 3
- Richardson, L. F. 1926, *Proceedings of the Royal Society of London Series A*, 110, 709
- Robertson, B. E., Kravtsov, A. V., Gnedin, N. Y., Abel, T., & Rudd, D. H. 2010, *MNRAS*, 401, 2463
- Rudd, D. H., Zentner, A. R., & Kravtsov, A. V. 2008, *ApJ*, 672, 19
- Shu, F. H. 1992, *The physics of astrophysics. Volume II: Gas dynamics*. (University Science Books)
- Sod, G. A. 1978, *Journal of Computational Physics*, 27, 1
- Springel, V. 2005, *MNRAS*, 364, 1105
- . 2010, *MNRAS*, 401, 791

- Springel, V., Pakmor, R., Zier, O., & Reinecke, M. 2021, MNRAS, 506, 2871
- Springel, V., Yoshida, N., & White, S. D. M. 2001, New A, 6, 79
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137
- Stone, J. M., & Norman, M. L. 1992, ApJS, 80, 753
- Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, ApJS, 249, 4
- Strang, G. 1968, SIAM Journal on Numerical Analysis, 5, 506
- Teyssier, R., Fromang, S., & Dormy, E. 2006, Journal of Computational Physics, 218, 44
- Toro, E. 2009, Riemann solvers and numerical methods for fluid dynamics, 3rd edn. (Springer)
- Toro, E. F., Spruce, M., & Speares, W. 1994, Shock Waves, 4, 25
- van den Bosch, F. 2020, ASTR 595: Astrophysical Flow (Lecture notes)
- van Dyke, M. 1982, NASA STI/Recon Technical Report A, 82, 36549
- van Leer, B. 1974, Journal of Computational Physics, 14, 361
- . 1979, Journal of Computational Physics, 32, 101
- Wadsley, J. W., Keller, B. W., & Quinn, T. R. 2017, MNRAS, 471, 2357
- Wang, P., & Abel, T. 2009, ApJ, 696, 96
- Weinberger, R., Springel, V., & Pakmor, R. 2020, ApJS, 248, 32
- Zhang, J., Johnson, P. C., & Popel, A. S. 2007, Physical Biology, 4, 285
- Zingale, M. 2014, Astronomy and Computing, 6, 52
- . 2021, Introduction to Computational Astrophysical Hydrodynamics (The Open Astrophysics Bookshelf).
https://github.com/Open-Astrophysics-Bookshelf/numerical_exercises