

Finite difference methods for wave motion

Hans Petter Langtangen^{1,2}

¹Center for Biomedical Computing, Simula Research Laboratory

²Department of Informatics, University of Oslo

Dec 12, 2013

This is still a **preliminary version**.

Contents

Simulation of waves on a string	5
1.1 Discretizing the domain	6
1.2 The discrete solution	6
1.3 Fulfilling the equation at the mesh points	6
1.4 Replacing derivatives by finite differences	7
1.5 Formulating a recursive algorithm	8
1.6 Sketch of an implementation	9
Verification	10
2.1 A slightly generalized model problem	11
2.2 Using an analytical solution of physical significance	11
2.3 Manufactured solution	12
2.4 Constructing an exact solution of the discrete equations	13
Implementation	14
3.1 Making a solver function	15
3.2 Verification: exact quadratic solution	16
3.3 Visualization: animating the solution	16
3.4 Running a case	19
3.5 The benefits of scaling	20
Vectorization	21
4.1 Operations on slices of arrays	21
4.2 Finite difference schemes expressed as slices	23
4.3 Verification	24
4.4 Efficiency measurements	25

5 Exercises	
6 Generalization: reflecting boundaries	
6.1 Neumann boundary condition	
6.2 Discretization of derivatives at the boundary	
6.3 Implementation of Neumann conditions	
6.4 Index set notation	
6.5 Alternative implementation via ghost cells	
7 Generalization: variable wave velocity	
7.1 The model PDE with a variable coefficient	
7.2 Discretizing the variable coefficient	
7.3 Computing the coefficient between mesh points	
7.4 How a variable coefficient affects the stability	
7.5 Neumann condition and a variable coefficient	
7.6 Implementation of variable coefficients	
7.7 A more general model PDE with variable coefficients	
7.8 Generalization: damping	
8 Building a general 1D wave equation solver	
8.1 User action function as a class	
8.2 Pulse propagation in two media	
9 Exercises	
10 Analysis of the difference equations	
10.1 Properties of the solution of the wave equation	
10.2 More precise definition of Fourier representations	
10.3 Stability	
10.4 Numerical dispersion relation	
10.5 Extending the analysis to 2D and 3D	
11 Finite difference methods for 2D and 3D wave equations	
11.1 Multi-dimensional wave equations	
11.2 Mesh	
11.3 Discretization	
12 Implementation	
12.1 Scalar computations	
12.2 Vectorized computations	
12.3 Verification	
13 Migrating loops to Cython	
13.1 Declaring variables and annotating the code	
13.2 Visual inspection of the C translation	
13.3 Building the extension module	
13.4 Calling the Cython function from Python	

4 Migrating loops to Fortran	71
14.1 The Fortran subroutine	71
14.2 Building the Fortran module with f2py	72
14.3 How to avoid array copying	74
5 Migrating loops to C via Cython	75
15.1 Translating index pairs to single indices	75
15.2 The complete C code	76
15.3 The Cython interface file	77
15.4 Building the extension module	77
6 Migrating loops to C via f2py	78
16.1 Migrating loops to C++ via f2py	79
7 Using classes to implement a simulator	80
8 Exercises	80
9 Applications of wave equations	82
19.1 Waves on a string	82
19.2 Waves on a membrane	85
19.3 Elastic waves in a rod	85
19.4 The acoustic model for seismic waves	86
19.5 Sound waves in liquids and gases	87
19.6 Spherical waves	89
19.7 The linear shallow water equations	90
19.8 Waves in blood vessels	92
19.9 Electromagnetic waves	94
0 Exercises	94

List of Exercises, Problems, and Projects

Exercise	1	Simulate a standing wave
Exercise	2	Add storage of solution in a user action function ...
Exercise	3	Use a class for the user action function
Exercise	4	Compare several Courant numbers in one movie
Project	5	Calculus with 1D mesh functions
Exercise	6	Find the analytical solution to a damped wave ...
Problem	7	Explore symmetry boundary conditions
Exercise	8	Send pulse waves through a layered medium
Exercise	9	Compare discretizations of a Neumann condition ...
Exercise	10	Check that a solution fulfills the discrete ...
Project	11	Calculus with 2D/3D mesh functions
Exercise	12	Implement Neumann conditions in 2D
Exercise	13	Test the efficiency of compiled loops in 3D
Exercise	14	Simulate waves on a non-homogeneous string
Exercise	15	Simulate damped waves on a string
Exercise	16	Simulate elastic waves in a rod
Exercise	17	Simulate spherical waves
Exercise	18	Explain why numerical noise occurs
Exercise	19	Investigate harmonic averaging in a 1D model
Problem	20	Implement open boundary conditions
Problem	21	Earthquake-generated tsunami over a subsea ...
Problem	22	Earthquake-generated tsunami over a 3D hill
Problem	23	Investigate Matplotlib for visualization
Problem	24	Investigate visualization packages
Problem	25	Implement loops in compiled languages
Exercise	26	Simulate seismic waves in 2D
Project	27	Model 3D acoustic waves in a room
Project	28	Solve a 1D transport equation
Problem	29	General analytical solution of a 1D damped ...
Problem	30	General analytical solution of a 2D damped ...

A very wide range of physical processes lead to wave motion, where signals are propagated through a medium in space and time, normally with little or no permanent movement of the medium itself. The shape of the signals may undergo changes as they travel through matter, but usually not so much that the signals cannot be recognized at some later point in space and time. Many types of wave motion can be described by the equation $u_{tt} = \nabla \cdot (c^2 \nabla u) + f$, which we will solve in the forthcoming text by finite difference methods.

Simulation of waves on a string

We begin our study of wave equations by simulating one-dimensional waves on a string, say on a guitar or violin string. Let the string in the deformed state coincide with the interval $[0, L]$ on the x axis, and let $u(x, t)$ be the displacement at time t in the y direction of a point initially at x . The displacement function is governed by the mathematical model

$$\frac{\partial^2 u}{\partial t^2} = c^2 \frac{\partial^2 u}{\partial x^2}, \quad x \in (0, L), \quad t \in (0, T) \quad (1)$$

$$u(x, 0) = I(x), \quad x \in [0, L] \quad (2)$$

$$\frac{\partial}{\partial t} u(x, 0) = 0, \quad x \in [0, L] \quad (3)$$

$$u(0, t) = 0, \quad t \in (0, T) \quad (4)$$

$$u(L, t) = 0, \quad t \in (0, T) \quad (5)$$

The constant c and the function $I(x)$ must be prescribed.

Equation (1) is known as the one-dimensional *wave equation*. Since this PDE contains a second-order derivative in time, we need *two initial conditions*, here (2) specifying the initial shape of the string, $I(x)$, and (3) reflecting that the initial velocity of the string is zero. In addition, PDEs need *boundary conditions*, here (4) and (5), specifying that the string is fixed at the ends, i.e., that the displacement u is zero.

The solution $u(x, t)$ varies in space and time and describes waves that are moving with velocity c to the left and right.

Sometimes we will use a more compact notation for the partial derivatives to save space:

$$u_t = \frac{\partial u}{\partial t}, \quad u_{tt} = \frac{\partial^2 u}{\partial t^2}, \quad (6)$$

and similar expressions for derivatives with respect to other variables. Then the wave equation can be written compactly as $u_{tt} = c^2 u_{xx}$.

The PDE problem (1)-(5) will now be discretized in space and time by a finite difference method.

1.1 Discretizing the domain

The temporal domain $[0, T]$ is represented by a finite number of mesh points

$$0 = t_0 < t_1 < t_2 < \cdots < t_{N_t-1} < t_{N_t} = T.$$

Similarly, the spatial domain $[0, L]$ is replaced by a set of mesh points

$$0 = x_0 < x_1 < x_2 < \cdots < x_{N_x-1} < x_{N_x} = L.$$

One may view the mesh as two-dimensional in the x, t plane, consisting of points (x_i, t_n) , with $i = 0, \dots, N_x$ and $n = 0, \dots, N_t$.

Uniform meshes. For uniformly distributed mesh points we can in addition use the constant mesh spacings Δt and Δx . We have that

$$x_i = i\Delta x, \quad i = 0, \dots, N_x, \quad t_n = n\Delta t, \quad n = 0, \dots, N_t.$$

We also have that $\Delta x = x_i - x_{i-1}$, $i = 1, \dots, N_x$, and $\Delta t = t_n - t_{n-1}$, $n = 1, \dots, N_t$. Figure 1 displays a mesh in the x, t plane with $N_t = 5$, $N_x = 5$ and constant mesh spacings.

1.2 The discrete solution

The solution $u(x, t)$ is sought at the mesh points. We introduce the mesh function u_i^n , which approximates the exact solution at the mesh point (x_i, t_n) for $i = 0, \dots, N_x$ and $n = 0, \dots, N_t$. Using the finite difference method, we develop algebraic equations for computing the mesh function. The equations are shown in Figure 1 illustrate neighboring mesh points where values of u_i^n are connected through an algebraic equation. In this particular case, $u_2^1, u_1^2, u_2^2, u_3^2$, and u_2^3 are connected in an algebraic equation associated with the center point (2, 2). The term *stencil* is often used about the algebraic equation at a mesh point, and the geometry of a typical stencil is illustrated in Figure 1. One also often refers to the algebraic equations as *discrete equations*, (*finite*) *difference equations*, or *finite difference scheme*.

1.3 Fulfilling the equation at the mesh points

For a numerical solution by the finite difference method, we relax the requirement that (1) holds at all points in the space-time domain $(0, L) \times (0, T]$ and require that the PDE is fulfilled at the *interior* mesh points:

$$\frac{\partial^2}{\partial t^2} u(x_i, t_n) = c^2 \frac{\partial^2}{\partial x^2} u(x_i, t_n),$$

for $i = 1, \dots, N_x - 1$ and $n = 1, \dots, N_t - 1$. For $n = 0$ we have the initial conditions $u = I(x)$ and $u_t = 0$, and at the boundaries $i = 0, N_x$ we have the boundary condition $u = 0$.




Figure 1: Mesh in space and time for a 1D wave equation.

4 Replacing derivatives by finite differences

The second-order derivatives can be replaced by central differences. The most widely used difference approximation of the second-order derivative is

$$\frac{\partial^2}{\partial t^2} u(x_i, t_n) \approx \frac{u_i^{n+1} - 2u_i^n + u_i^{n-1}}{\Delta t^2}.$$

It is convenient to introduce the finite difference operator notation

$$[D_t D_t u]_i^n = \frac{u_i^{n+1} - 2u_i^n + u_i^{n-1}}{\Delta t^2}.$$

A similar approximation of the second-order derivative in the x direction reads

$$\frac{\partial^2}{\partial x^2} u(x_i, t_n) \approx \frac{u_{i+1}^n - 2u_i^n + u_{i-1}^n}{\Delta x^2} = [D_x D_x u]_i^n.$$

Algebraic version of the PDE. We can now replace the derivatives in (10) and get

$$\frac{u_i^{n+1} - 2u_i^n + u_i^{n-1}}{\Delta t^2} = c^2 \frac{u_{i+1}^n - 2u_i^n + u_{i-1}^n}{\Delta x^2}, \quad (11)$$

or written more compactly using the operator notation:

$$[D_t D_t u = c^2 D_x D_x]_i^n. \quad (12)$$

Algebraic version of the initial conditions. We also need to replace the derivative in the initial condition (3) by a finite difference approximation centered difference of the type

$$\frac{\partial}{\partial t} u(x_i, t_n) \approx \frac{u_i^1 - u_i^{-1}}{2\Delta t} = [D_{2t} u]_i^0,$$

seems appropriate. In operator notation the initial condition is written

$$[D_{2t} u]_i^n = 0, \quad n = 0.$$

Writing out this equation and ordering the terms give

$$u_i^{n-1} = u_i^{n+1}, \quad i = 0, \dots, N_x, \quad n = 0.$$

The other initial condition can be computed by

$$u_i^0 = I(x_i), \quad i = 0, \dots, N_x.$$

1.5 Formulating a recursive algorithm

We assume that u_i^n and u_i^{n-1} are already computed for $i = 0, \dots, N_x$. The unknown quantity in (11) is therefore u_i^{n+1} , which we can solve for:

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + C^2 (u_{i+1}^n - 2u_i^n + u_{i-1}^n),$$

where we have introduced the parameter

$$C = c \frac{\Delta t}{\Delta x},$$

known as the (dimensionless) *Courant number*. We see that the discrete version of the PDE features only one parameter, C , which is therefore the key parameter that governs the quality of the numerical solution. Both the primary parameter c and the numerical parameters Δx and Δt are lumped together into C .

Given that u_i^{n-1} and u_i^n are computed for $i = 0, \dots, N_x$, we find u_i^{n+1} at the next time level by applying the formula (14) for $i = 1, \dots, N_x - 1$. Figure 2 illustrates the points that are used to compute u_2^3 . For the boundary points $i = 0$ and $i = N_x$, we apply the boundary conditions $u_i^{n+1} = 0$.

A problem with (14) arises when $n = 0$ since the formula for u_i^1 involves u_i^{-1} which is an undefined quantity outside the time mesh (and the time domain). However, we can use the initial condition (13) in combination with (14) for $n = 0$ to arrive at a special formula for u_i^1 :

$$u_i^1 = u_i^0 - \frac{1}{2} C^2 (u_{i+1}^0 - 2u_i^0 + u_{i-1}^0).$$

Figure 2 illustrates how (16) connects four instead of five points: $u_2^1, u_1^0, u_3^0, u_2^0$.

We can now summarize the computational algorithm:




Figure 2: Modified stencil for the first time step.

1. Compute $u_i^0 = I(x_i)$ for $i = 0, \dots, N_x$
2. Compute u_i^1 by (16) and set $u_i^1 = 0$ for the boundary points $i = 0$ and $i = N_x$, for $n = 1, 2, \dots, N - 1$,
3. For each time level $n = 1, 2, \dots, N_t - 1$
 - (a) apply (14) to find u_i^{n+1} for $i = 1, \dots, N_x - 1$
 - (b) set $u_i^{n+1} = 0$ for the boundary points $i = 0, i = N_x$.

The algorithm essentially consists of moving a finite difference stencil through all the mesh points, which is illustrated by an animation in a [web page](#)¹ or a [movie file](#)².

6 Sketch of an implementation

In a Python implementation of this algorithm, we use the array elements $u[i]$ to store u_i^{n+1} , $u_1[i]$ to store u_i^n , and $u_2[i]$ to store u_i^{n-1} . Our naming convention is to use u for the unknown new spatial field to be computed, u_1 as the solution at one time step back in time, u_2 as the solution two time steps back in time and so forth.

The algorithm only needs to access the three most recent time levels, so we need only three arrays for u_i^{n+1} , u_i^n , and u_i^{n-1} , $i = 0, \dots, N_x$. Storing all the

¹http://tinyurl.com/k3sdbuv/pub/mov-wave/wave1D.PDE.Dirichlet.stencil_gpl/index.html

²http://tinyurl.com/k3sdbuv/pub/mov-wave/wave1D.PDE.Dirichlet.stencil_gpl/movie.ogg

solutions in a two-dimensional array of size $(N_x + 1) \times (N_t + 1)$ would be in this simple one-dimensional PDE problem, but is normally out of the question in three-dimensional (3D) and large two-dimensional (2D) problems. We therefore in all our programs for solving PDEs have the unknown in memory as few time levels as possible.

The following Python snippet realizes the steps in the computational algorithm.

```
# Given mesh points as arrays x and t (x[i], t[n])
dx = x[1] - x[0]
dt = t[1] - t[0]
C = c*dt/dx          # Courant number
Nt = len(t)-1
C2 = C**2           # Help variable in the scheme

# Set initial condition u(x,0) = I(x)
for i in range(0, Nx+1):
    u_1[i] = I(x[i])

# Apply special formula for first step, incorporating du/dt=0
for i in range(1, Nx):
    u[i] = u_1[i] - 0.5*C**2*(u_1[i+1] - 2*u_1[i] + u_1[i-1])
u[0] = 0; u[Nx] = 0 # Enforce boundary conditions

# Switch variables before next step
u_2[:, u_1[:]] = u_1, u

for n in range(1, Nt):
    # Update all inner mesh points at time t[n+1]
    for i in range(1, Nx):
        u[i] = 2*u_1[i] - u_2[i] - \
            C**2*(u_1[i+1] - 2*u_1[i] + u_1[i-1])

    # Insert boundary conditions
    u[0] = 0; u[Nx] = 0

    # Switch variables before next step
    u_2[:, u_1[:]] = u_1, u
```

2 Verification

Before implementing the algorithm, it is convenient to add a source term to the PDE (1) since it gives us more freedom in finding test problems for verification. In particular, the source term allows us to use *manufactured solutions* for testing, where we simply choose some function as solution, fit the corresponding source term, and define boundary and initial conditions consistent with the chosen solution. Such solutions will seldom fulfill the initial condition (2) so we need to generalize this condition to $u_t = V(x)$.

.1 A slightly generalized model problem

We now address the following extended initial-boundary value problem for one-dimensional wave phenomena:

$$u_{tt} = c^2 u_{xx} + f(x, t), \quad x \in (0, L), \quad t \in (0, T) \quad (17)$$

$$u(x, 0) = I(x), \quad x \in [0, L] \quad (18)$$

$$u_t(x, 0) = V(x), \quad x \in [0, L] \quad (19)$$

$$u(0, t) = 0, \quad t > 0 \quad (20)$$

$$u(L, t) = 0, \quad t > 0 \quad (21)$$

Sampling the PDE at (x_i, t_n) and using the same finite difference approximations as above, yields

$$[D_t D_t u = c^2 D_x D_x + f]_i^n. \quad (22)$$

Writing this out and solving for the unknown u_i^{n+1} results in

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + C^2(u_{i+1}^n - 2u_i^n + u_{i-1}^n) + \Delta t^2 f_i^n. \quad (23)$$

The equation for the first time step must be rederived. The discretization of the initial condition $u_t = V(x)$ at $t = 0$ becomes

$$[D_{2t} u = V]_i^0 \Rightarrow u_i^{-1} = u_i^1 - 2\Delta t V_i,$$

which, when inserted in (23) for $n = 0$, gives the special formula

$$u_i^1 = u_i^0 - \Delta t V_i + \frac{1}{2} C^2 (u_{i+1}^0 - 2u_i^0 + u_{i-1}^0) + \frac{1}{2} \Delta t^2 f_i^0. \quad (24)$$

.2 Using an analytical solution of physical significance

Many wave problems feature sinusoidal oscillations in time and space. For example, the original PDE problem (1)-(5) allows a solution

$$u_e(x, y, t) = A \sin\left(\frac{\pi}{L}x\right) \cos\left(\frac{\pi}{L}ct\right). \quad (25)$$

This u_e fulfills the PDE with $f = 0$, boundary conditions $u_e(0, t) = u_e(L, t) = 0$, as well as initial conditions $I(x) = A \sin\left(\frac{\pi}{L}x\right)$ and $V = 0$.

It is common to use such exact solutions of physical interest to verify implementations. However, the numerical solution u_i^n will only be an approximation to $u_e(x_i, t_n)$. We do not have knowledge of the precise size of the error in this approximation, and therefore we can never know if discrepancies between the computed u_i^n and $u_e(x_i, t_n)$ are caused by mathematical approximations or programming errors. In particular, if a plot of the computed solution u_i^n and the exact one (25) looks similar, many are tempted to claim that the implementation works, but there can still be serious programming errors although color plots look nice.

The only way to use exact physical solutions like (25) for serious and thorough verification is to run a series of finer and finer meshes, measure the integer error in each mesh, and from this information estimate the convergence rates. If these rates are very close to 2, we have strong evidence that the implementation works.

2.3 Manufactured solution

One problem with the exact solution (25) is that it requires a simple function ($V = 0, f = 0$) of the implemented problem (17)-(21). An advantage of a manufactured solution is that we can test all terms in the PDE and the boundary conditions, and initial conditions to be compatible with the chosen solution. Given that our boundary conditions in the implementation are $u(0, t) = u(L, t) = 0$, we must choose a solution that fulfills these conditions. One example is

$$u_e(x, t) = x(L - x) \sin t.$$

Inserted in the PDE $u_{tt} = c^2 u_{xx} + f$ we get

$$-x(L - x) \sin t = -2 \sin t + f \Rightarrow f = (2 - x(L - x)) \sin t.$$

The initial conditions become

$$u(x, 0) = I(x) = 0,$$

$$u_t(x, 0) = V(x) = (2 - x(L - x)) \cos t.$$

To verify the code, we run a series of refined meshes and compare convergence rates. In more detail, we keep $\Delta t/\Delta x$ constant for each experiment, implying that C is also constant throughout the experiments. A convergence parameter $h = \Delta t$ is introduced. For a given C (and c), L and T , we choose an initial time cell size h_0 and run experiments with decreasing $h_i = 2^{-i} h_0$, $i = 1, 2, \dots, m$. Halving the cell size in each experiment is necessary, but common. For each experiment we must record a scalar measure of the error. As will be shown later, it is expected that such error measures are proportional to h^2 . A standard choice of error measure is the ℓ^2 or ℓ^∞ norm of the error mesh function e_i^n :

$$\|e_i^n\|_{\ell^2} = \left(\Delta t \Delta x \sum_{n=0}^{N_t} \sum_{i=0}^{N_x} (e_i^n)^2 \right)^{\frac{1}{2}}, \quad e_i^n = u_e(x_i, t_n) - u_i^n,$$

$$\|e_i^n\|_{\ell^\infty} = \max_{i,n} |e_i^n|.$$

In Python, one can compute $\sum_i (e_i^{n+1})^2$ at each time step and accumulate the value in some sum variable, say `e2_sum`. At the final time step one

`sqrt(dt*dx*e2_sum)`. For the ℓ^∞ norm one must compare the maximum error at a time level (`e.max()`) with the global maximum over the time domain: `_max = max(e_max, e.max())`.

An alternative error measure is to use a spatial norm at one time step only, e.g., the end time T :

$$\|e_i^n\|_{\ell^2} = \left(\Delta x \sum_{i=0}^{N_x} (e_i^n)^2 \right)^{\frac{1}{2}}, \quad e_i^n = u_e(x_i, t_n) - u_i^n, \quad (28)$$

$$\|e_i^n\|_{\ell^\infty} = \max_{0 \leq i \leq N_x} |e_i^n|. \quad (29)$$

Let E_i be the error measure in experiment (mesh) number i and let h_i be the corresponding discretization parameter (h). We expect an error model $E_i = Ch_i^r$, where $r = 0$. To estimate r , we can compare two consecutive experiments and compute

$$r_i = \frac{\ln E_{i+1}/E_i}{\ln h_{i+1}/h_i}, \quad i = 0, \dots, m-1.$$

We should observe that r_i approaches 2 as i increases.

The next section describes a method of manufactured solutions where do not need to compute error measures and check that they converge as expected as the mesh is refined.

4 Constructing an exact solution of the discrete equations

For verification purposes we shall use a solution that is quadratic in space and linear in time. More specifically, our choice of the manufactured solution is

$$u_e(x, t) = x(L-x)\left(1 + \frac{1}{2}t\right), \quad (30)$$

which by insertion in the PDE leads to $f(x, t) = 2(1+t)c^2$. This u_e fulfills the boundary conditions and is compatible with $I(x) = x(L-x)$ and $V(x) = x(L-x)$.

A key feature of the chosen u_e is that it is also *an exact solution of the discrete equations*. To realize this very important result, we first establish the results

$$[D_t D_t t^2]^n = \frac{t_{n+1}^2 - 2t_n^2 + t_{n-1}^2}{\Delta t^2} = (n+1)^2 - n^2 + (n-1)^2 = 2, \quad (31)$$

$$[D_t D_t t]^n = \frac{t_{n+1} - 2t_n + t_{n-1}}{\Delta t^2} = \frac{((n+1) - n + (n-1))\Delta t}{\Delta t^2} = 0. \quad (32)$$

Hence,

$$[D_t D_t u_e]^n = x_i(L-x_i)[D_t D_t \left(1 + \frac{1}{2}t\right)]^n = x_i(L-x_i)\frac{1}{2}[D_t D_t t]^n = 0,$$

and

$$\begin{aligned} [D_x D_x u_e]^n &= (1 + \frac{1}{2}t_n)[D_x D_x (xL - x^2)]_i = (1 + \frac{1}{2}t_n)[LD_x D_x x - D_x^2 x] \\ &= -2(1 + \frac{1}{2}t_n). \end{aligned}$$

Now, $f_i^n = 2(1 + \frac{1}{2}t_n)c^2$ and we get

$$[D_t D_t u_e - c^2 D_x D_x u_e - f]^n = 0 - c^2(-1)2(1 + \frac{1}{2}t_n) + 2(1 + \frac{1}{2}t_n)c^2 = 0.$$

Moreover, $u_e(x_i, 0) = I(x_i)$, $\partial u_e / \partial t = V(x_i)$ at $t = 0$, and $u_e(x_{N_x}, 0) = 0$. Also the modified scheme for the first time step is full order accurate.

Therefore, the exact solution $u_e(x, t) = x(L-x)(1+t/2)$ of the PDE is also an exact solution of the discrete problem. We can use this result to verify that the computed u_i^n values from an implementation equals $u_e(x_i, t_n)$ to machine precision, *regardless of the mesh spacings Δx and Δt !* Nevertheless, there might be stability restrictions on Δx and Δt , so the test can only be used for a mesh that is compatible with the stability criterion (which in the case is $C \leq 1$, to be derived later).

Notice.

A product of quadratic or linear expressions in the various independent variables, as shown above, will often fulfill both the continuous and discrete PDE problem and can therefore be very useful solutions for verification implementations. However, for 1D wave equations of the type $u_{tt} = c^2 u_{xx}$, we shall see that there is always another much more powerful way of generating exact solutions (just set $C = 1$).

3 Implementation

This section presents the complete computational algorithm, its implementation in Python code, animation of the solution, and verification of the implementation.

A real implementation of the basic computational algorithm from Sect. 1.6 can be encapsulated in a function, taking all the input data for the problem as arguments. The physical input data consists of c , $I(x)$, $V(x)$, L , and T . The numerical input is the mesh parameters Δt and Δx . The possibility is to specify N_x and the Courant number $C = c\Delta t/\Delta x$. The more convenient to prescribe instead of Δt when performing numerical investigations because the numerical accuracy depends directly on C .

The solution at all spatial points at a new time level is stored in an array \mathbf{u} (of length $N_x + 1$). We need to decide what to do with this solution: visualize the curve, analyze the values, or write the array to file for later use. The decision what to do is left to the user in a supplied function.

```
def user_action(u, x, t, n):
```

here u is the solution at the spatial points x at time $t[n]$.

1.1 Making a solver function

first attempt at a solver function is listed below.

```
from numpy import *

def solver(I, V, f, c, L, Nx, C, T, user_action=None):
    """Solve  $u_{tt}=c^2u_{xx} + f$  on  $(0,L)x(0,T)$ ."""
    x = linspace(0, L, Nx+1) # Mesh points in space
    dx = x[1] - x[0]
    dt = C*dx/c
    Nt = int(round(T/dt))
    t = linspace(0, Nt*dt, Nt+1) # Mesh points in time
    C2 = C**2 # Help variable in the scheme
    if f is None or f == 0:
        f = lambda x, t: 0
    if V is None or V == 0:
        V = lambda x: 0

    u = zeros(Nx+1) # Solution array at new time level
    u_1 = zeros(Nx+1) # Solution at 1 time level back
    u_2 = zeros(Nx+1) # Solution at 2 time levels back

    import time; t0 = time.clock() # for measuring CPU time

    # Load initial condition into u_1
    for i in range(0, Nx+1):
        u_1[i] = I(x[i])

    if user_action is not None:
        user_action(u_1, x, t, 0)

    # Special formula for first time step
    n = 0
    for i in range(1, Nx):
        u[i] = u_1[i] + dt*V(x[i]) + \
            0.5*C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1]) + \
            0.5*dt**2*f(x[i], t[n])
    u[0] = 0; u[Nx] = 0

    if user_action is not None:
        user_action(u, x, t, 1)

    # Switch variables before next step
    u_2[:,], u_1[:,] = u_1, u

    for n in range(1, Nt):
        # Update all inner points at time t[n+1]
        for i in range(1, Nx):
            u[i] = -u_2[i] + 2*u_1[i] + \
                C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1]) + \
                dt**2*f(x[i], t[n])

        # Insert boundary conditions
```

```
u[0] = 0; u[Nx] = 0
if user_action is not None:
    if user_action(u, x, t, n+1):
        break

# Switch variables before next step
u_2[:,], u_1[:,] = u_1, u

cpu_time = t0 - time.clock()
return u, x, t, cpu_time
```

3.2 Verification: exact quadratic solution

We use the test problem derived in Section 2.1 for verification. Here is a realization of this verification as a nose test:

```
import nose.tools as nt

def test_quadratic():
    """Check that  $u(x,t)=x(L-x)(1+t/2)$  is exactly reproduced."""
    def exact_solution(x, t):
        return x*(L-x)*(1 + 0.5*t)

    def I(x):
        return exact_solution(x, 0)

    def V(x):
        return 0.5*exact_solution(x, 0)

    def f(x, t):
        return 2*(1 + 0.5*t)*c**2

    L = 2.5
    c = 1.5
    Nx = 3 # Very coarse mesh
    C = 0.75
    T = 18

    u, x, t, cpu = solver(I, V, f, c, L, Nx, C, T)
    u_e = exact_solution(x, t[-1])
    diff = abs(u - u_e).max()
    nt.assert_almost_equal(diff, 0, places=14)
```

3.3 Visualization: animating the solution

Now that we have verified the implementation it is time to do a real computation where we also display the evolution of the waves on the screen.

Visualization via SciTools. The following `viz` function defines a user callback function for plotting the solution at each time level:


```

def viz(I, V, f, c, L, Nx, C, T, umin, umax, animate=True):
    """Run solver and visualize u at each time level."""
    import scitools.std as plt
    import time, glob, os

    def plot_u(u, x, t, n):
        """user action function for solver."""
        plt.plot(x, u, 'r-',
                 xlabel='x', ylabel='u',
                 axis=[0, L, umin, umax],
                 title='t=%f' % t[n], show=True)
        # Let the initial condition stay on the screen for 2
        # seconds, else insert a pause of 0.2 s between each plot
        time.sleep(2) if t[n] == 0 else time.sleep(0.2)
        plt.savefig('frame_%04d.png' % n) # for movie making

    # Clean up old movie frames
    for filename in glob.glob('frame_*.png'):
        os.remove(filename)

    user_action = plot_u if animate else None
    u, x, t, cpu = solver(I, V, f, c, L, Nx, C, T, user_action)

    # Make movie files
    fps = 4 # Frames per second
    plt.movie('frame_*.png', encoder='html', fps=fps,
              output_file='movie.html')
    codec2ext = dict(flv='flv', libx64='mp4', libvpx='webm',
                    libtheora='ogg')
    filespec = 'frame_%04d.png'
    movie_program = 'avconv' # or 'ffmpeg'
    for codec in codec2ext:
        ext = codec2ext[codec]
        cmd = '%(movie_program)s -r %(fps)d -i %(filespec)s \' \
              -vcodec %(codec)s movie.%(ext)s' % vars()
        os.system(cmd)

```

A function inside another function, like `plot_u` in the above code segment, has access to *and remembers* all the local variables in the surrounding code inside the `viz` function (!). This is known in computer science as a *closure* and is very convenient to program with. For example, the `plt` and `time` modules defined outside `plot_u` are accessible for `plot_u` when the function is called (as `ser_action`) in the `solver` function. Some may think, however, that a class instead of a closure is a cleaner and easier-to-understand implementation of the `ser_action` function, see Section 8.

Making movie files. Several hardcopies of the animation are made from the `frame_*.png` files. The first movie, made by the SciTools function `plt.movie`, creates a `movie.html` file with a movie player for displaying the `frame_*.png` files. This movie player can be generated from the command line too

```

terminal> scitools movie encoder=html output_file=movie.html \
          fps=4 frame_*.png

```

We also use the `avconv` (or `ffmpeg`) programs to make movie files in formats: Flash, MP4, Webm, and Ogg. A typical `avconv` (or `ffmpeg`) call for creating a movie file in Ogg format with 4 frames per second built collection of plot files with names generated by `frame_%04d.png`, look

```

Terminal> avconv -r 4 -i frame_%04d.png -vcodec libtheora movie.og

```

The different formats require different video encoders (`-vcodec`) to be used: Flash applies `flv`, WebM applies `libvpx`, and MP4 applies `libx64`. Players `vlc`, `mplayer`, `gxine`, and `totem` can be used to play these movie files.

Note that padding the frame counter with zeros in the `frame_*.png` as specified by the `%04d` format, is essential so that the wildcard `frame_*.png` expands to the correct set of files.

Skipping frames for animation speed. Sometimes the time step Δt and T is large, leading to an inconveniently large number of plot files and animation on the screen. The solution to such a problem is to decide on a number of frames in the animation, `num_frames`, and plot the solution every `every` frame. The total number of time levels (i.e., maximum number of frames) is the length of `t`, `t.size`, and if we want `num_frames` need to plot every `t.size/num_frames` frame:

```

every = int(t.size/float(num_frames))
if n % every == 0 or n == t.size-1:
    st.plot(x, u, 'r-', ...)

```

The initial condition ($n=0$) is natural to include, and as `n % every == 0` very seldom be true for the very final frame, we also ensure that `n == t.size-1` and hence the final frame is included.

A simple choice of numbers may illustrate the formulas: say we have 801 frames in total (`t.size`) and we allow only 60 frames to be plotted. Then we need to plot every `801/60` frame, which with integer division yields 13 as `every`. Using the mod function, `n % every`, this operation is zero every time `n` is divided by 13 without a remainder. That is, the `if` test is true when `n` is 0, 13, 26, 39, ..., 780, 801. The associated code is included in the `plot_u` function in the file `wave1D_u0_sv.py`³.

Visualization via Matplotlib. The previous code based on the `plot_u` function from `scitools.std` can be run with Matplotlib as the visualization backend but if one desires to program directly with Matplotlib, quite different code is needed. Matplotlib's interactive mode must be turned on:

³http://tinyurl.com/jvzzcfn/wave/wave1D_u0_sv.py

```
import matplotlib.pyplot as plt
plt.ion() # interactive mode on
```

The most commonly used animation technique with Matplotlib is to update the data in the plot at each time level:

```
# Make a first plot
lines = plt.plot(t, u)
# call plt.axis, plt.xlabel, plt.ylabel, etc. as desired

# At later time levels
lines[0].set_ydata(u)
plt.legend('t=%g' % t[n])
plt.draw() # make updated plot
plt.savefig(...)
```

An alternative is to rebuild the plot at every time level:

```
plt.clf() # delete any previous curve(s)
plt.axis([...])
plt.plot(t, u)
# plt.xlabel, plt.legend and other decorations
plt.draw()
plt.savefig(...)
```

Many prefer to work with figure and axis objects as in MATLAB:

```
fig = plt.figure()
...
fig.clf()
ax = fig.gca()
ax.axis(...)
ax.plot(t, u)
# ax.set_xlabel, ax.legend and other decorations
plt.draw()
fig.savefig(...)
```

4 Running a case

The first demo of our 1D wave equation solver concerns vibrations of a string that is initially deformed to a triangular shape, like when picking a guitar string:

$$I(x) = \begin{cases} ax/x_0, & x < x_0, \\ a(L-x)/(L-x_0), & \text{otherwise} \end{cases} \quad (33)$$

We choose $L = 75$ cm, $x_0 = 0.8L$, $a = 5$ mm, $N_x = 50$, and a time frequency $\nu = 440$ Hz. The relation between the wave speed c and ν is $c = \nu\lambda$, where λ is the wavelength, taken as $2L$ because the longest wave on the string forms half a wavelength. There is no external force, so $f = 0$, and the string is at rest initially so that $V = 0$. A function setting these physical parameters and calling `iz` for this case goes as follows:

```
def guitar(C):
    """Triangular wave (pulled guitar string)."""
    L = 0.75
    x0 = 0.8*L
    a = 0.005
    freq = 440
    wavelength = 2*L
    c = freq*wavelength
    omega = 2*pi*freq
    num_periods = 1
    T = 2*pi/omega*num_periods
    Nx = 50

    def I(x):
        return a*x/x0 if x < x0 else a/(L-x0)*(L-x)

    umin = -1.2*a; umax = -umin
    cpu = viz(I, 0, 0, c, L, Nx, C, T, umin, umax, animate=True)
```

The associated program has the name `wave1D_u0_s.py`⁴. Run the program and watch the [movie of the vibrating string](#)⁵.

3.5 The benefits of scaling

The previous example demonstrated that quite some work is needed with listing relevant physical parameters for a case. By *scaling* the mathematical problem we can often reduce the need to estimate physical parameters locally. A scaling consists of introducing new independent and dependent variables with the aim that the absolute value of these vary between 0 and 1:

$$\bar{x} = \frac{x}{L}, \quad \bar{t} = \frac{c}{L}t, \quad \bar{u} = \frac{u}{a}.$$

Replacing old by new variables in the PDE, using $f = 0$, and dropping t results in the *scaled equation* $u_{\bar{t}\bar{t}} = u_{\bar{x}\bar{x}}$. This equation has no physical parameters (!).

If we have a program implemented for the physical wave equation in dimensions, we can obtain the dimensionless, scaled version by setting $a = 1$, $L = 1$. The initial condition corresponds to (183), but with setting $a = 1$, $L = 1$, $x_0 \in [0, 1]$. This means that we only need to decide on the x_0 value as a part of unity, because the scaled problem corresponds to setting all other parameters to unity! In the code we can just set `a=c=L=1`, `x0=0.8`, and there is no need to calculate with wavelengths and frequencies to estimate c .

The only non-trivial parameter to estimate in the scaled problem is the end time of the simulation, or more precisely, how it relates to periods in the solutions in time, since we often want to express the end time as a number of periods. Suppose u behaves as $\sin(\omega t)$ in time in variable dimension. The corresponding period is $P = 2\pi/\omega$. The frequency ω is related to the wavelength λ of the waves through the relations $\omega = kc$ and $k = 2\pi/\lambda$.

⁴http://tinyurl.com/jvzzcfn/wave/wave1D_u0_s.py

⁵http://tinyurl.com/k3sdbuv/pub/mov-wave/guitar_CO.8/index.html

$= 2\pi c/\lambda$ and $P = \lambda/c$. It remains to estimate λ . With $u(x, t) = F(x) \sin \omega t$ we find from $u_{tt} = c^2 u_{xx}$ that $c^2 F'' + \omega^2 F = 0$, and the boundary conditions demand $F(0) = F(L) = 0$. The solution is $F(x) = \sin(x\pi/L)$, which has wavelength $\lambda = 2\pi/(\pi/L) = 2L$. One period is therefore given by $P = 2L/c$. The dimensionless period is $\bar{P} = Pc/L = 2$.

Vectorization

The computational algorithm for solving the wave equation visits one mesh point at a time and evaluates a formula for the new value u_i^{n+1} at that point. Technically, this is implemented by a loop over array elements in a program. Such loops may run slowly in Python (and similar interpreted languages such as Perl and MATLAB). One technique for speeding up loops is to perform operations on entire arrays instead of working with one element at a time. This is referred to as *vectorization*, *vector computing*, or *array computing*. Operations on whole arrays are possible if the computations involving each element is independent of each other and therefore can, at least in principle, be performed simultaneously. Vectorization not only speeds up the code on serial computers, but it also makes it easy to exploit parallel computing.

1 Operations on slices of arrays

Efficient computing with `numpy` arrays demands that we avoid loops and compute with entire arrays at once (or at least large portions of them). Consider this calculation of differences $d_i = u_{i+1} - u_i$:

```
i = u.size
for i in range(0, n-1):
    d[i] = u[i+1] - u[i]
```

All the differences here are independent of each other. The computation of `d` can therefore alternatively be done by subtracting the array $(u_0, u_1, \dots, u_{n-1})$ from the array where the elements are shifted one index upwards: (u_1, u_2, \dots, u_n) , see Figure 3. The former subset of the array can be expressed by `u[0:n-1]`, `u[0:-1]`, or just `u[:-1]`, meaning from index 0 up to, but not including, the last element (-1). The latter subset is obtained by `u[1:n]` or `u[1:]`, meaning from index 1 and the rest of the array. The computation of `d` can now be done without an explicit Python loop:

```
d = u[1:] - u[:-1]
```

or with explicit limits if desired:

```
d = u[1:n] - u[0:n-1]
```

Indices with a colon, going from an index to (but not including) another index are called *slices*. With `numpy` arrays, the computations are still done by loops,

but in efficient, compiled, highly optimized code in C or Fortran. Such operations can also easily be distributed among many processors on supercomputers. We say that the *scalar code* above, working on an element (i, t) at a time, has been replaced by an equivalent *vectorized code*. The process of vectorizing code is called *vectorization*.




Figure 3: Illustration of subtracting two slices of two arrays.

Test the understanding.

Newcomers to vectorization are encouraged to choose a small array with five elements, and simulate with pen and paper both the loop version and the vectorized version.

Finite difference schemes basically contains differences between array elements with shifted indices. Consider the updating formula

```
for i in range(1, n-1):
    u2[i] = u[i-1] - 2*u[i] + u[i+1]
```

The vectorization consists of replacing the loop by arithmetics on slices of length `n-2`:

```
u2 = u[:-2] - 2*u[1:-1] + u[2:]
u2 = u[0:n-2] - 2*u[1:n-1] + u[2:n] # alternative
```

Note that `u2` here gets length `n-2`. If `u2` is already an array of length `n` we want to use the formula to update all the "inner" elements of `u2`, as when solving a 1D wave equation, we can write

```
u2[1:-1] = u[:-2] - 2*u[1:-1] + u[2:]
u2[1:n-1] = u[0:n-2] - 2*u[1:n-1] + u[2:n] # alternative
```

Pen and paper calculations with a small array will demonstrate what is going on. The expressions on the right-hand side are done in the following order involving temporary arrays with intermediate results, since we can only work with two arrays at a time in arithmetic expressions:

```
temp1 = 2*u[1:-1]
temp2 = u[0:-2] - temp1
temp3 = temp2 + u[2:]
u2[1:-1] = temp3
```

We can extend the previous example to a formula with an additional term computed by calling a function:

```
def f(x):
    return x**2 + 1

for i in range(1, n-1):
    u2[i] = u[i-1] - 2*u[i] + u[i+1] + f(x[i])
```

Assuming `u2`, `u`, and `x` all have length `n`, the vectorized version becomes

```
u2[1:-1] = u[:-2] - 2*u[1:-1] + u[2:] + f(x[1:-1])
```

2 Finite difference schemes expressed as slices

We now have the necessary tools to vectorize the algorithm for the wave equation. There are three loops: one for the initial condition, one for the first time step, and finally the loop that is repeated for all subsequent time levels. Since only the latter is repeated a potentially large number of times, we limit the efforts of vectorizing the code to this loop:

```
for i in range(1, Nx):
    u[i] = 2*u_1[i] - u_2[i] + \
        C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1])
```

The vectorized version becomes

```
u[1:-1] = - u_2[1:-1] + 2*u_1[1:-1] + \
    C2*(u_1[:-2] - 2*u_1[1:-1] + u_1[2:])
```

```
r
```

```
u[1:Nx] = 2*u_1[1:Nx] - u_2[1:Nx] + \
    C2*(u_1[0:Nx-1] - 2*u_1[1:Nx] + u_1[2:Nx+1])
```

The program `wave1D_u0_sv.py`⁶ contains a new version of the function `solver` where both the scalar and the vectorized loops are included (the argument `version` is set to `scalar` or `vectorized`, respectively).

⁶http://tinyurl.com/jvzzcfn/wave/wave1D_u0_sv.py

4.3 Verification

We may reuse the quadratic solution $u_e(x, t) = x(L-x)(1 + \frac{1}{2}t)$ to verify the vectorized code. A nose test can now test both the scalar and the vector version. Moreover, we may use a `user_action` function that compares the computed and exact solution at each time level and performs a test:

```
def test_quadratic():
    """
    Check the scalar and vectorized versions work for
    a quadratic u(x,t)=x(L-x)(1+t/2) that is exactly reproduced.
    """
    # The following function must work for x as array or scalar
    exact_solution = lambda x, t: x*(L - x)*(1 + 0.5*t)
    I = lambda x: exact_solution(x, 0)
    V = lambda x: 0.5*exact_solution(x, 0)
    # f is a scalar (zeros_like(x) works for scalar x too)
    f = lambda x, t: zeros_like(x) + 2*c**2*(1 + 0.5*t)

    L = 2.5
    c = 1.5
    Nx = 3 # Very coarse mesh
    C = 1
    T = 18 # Long time integration

    def assert_no_error(u, x, t, n):
        u_e = exact_solution(x, t[n])
        diff = abs(u - u_e).max()
        nt.assert_almost_equal(diff, 0, places=13)

    solver(I, V, f, c, L, Nx, C, T,
           user_action=assert_no_error, version='scalar')
    solver(I, V, f, c, L, Nx, C, T,
           user_action=assert_no_error, version='vectorized')
```

Lambda functions.

The code segment above demonstrates how to achieve very compact code with the use of lambda functions for the various input parameters. In essence,

```
f = lambda x, t: L*(x-t)**2
```

is equivalent to

```
def f(x, t):
    return L*(x-t)**2
```

Note that lambda functions can just contain a single expression and return statements.

One advantage with lambda functions is that they can be used directly in calls:

```
solver(I=lambda x: sin(pi*x/L), V=0, f=0, ...)
```

4 Efficiency measurements

Running the `wave1D_u0_sv.py` code with the previous string vibration example for $N_x = 50, 100, 200, 400, 800$, and measuring the CPU time (see the `un_efficiency_experiments` function), shows that the vectorized code runs substantially faster: the scalar code uses approximately a factor $N_x/5$ more me!

Exercises

Exercise 1: Simulate a standing wave

The purpose of this exercise is to simulate standing waves on $[0, L]$ and illustrate the error in the simulation. Standing waves arise from an initial condition

$$u(x, 0) = A \sin\left(\frac{\pi}{L}mx\right),$$

where m is an integer and A is a freely chosen amplitude. The corresponding exact solution can be computed and reads

$$u_e(x, t) = A \sin\left(\frac{\pi}{L}mx\right) \cos\left(\frac{\pi}{L}mct\right).$$

) Explain that for a function $\sin kx \cos \omega t$ the wave length in space is $\lambda = 2\pi/k$ and the period in time is $P = 2\pi/\omega$. Use these expressions to find the wave length in space and period in time of u_e above.

) Import the `solver` function `wave1D_u0_s.py` into a new file where the `viz` function is reimplemented such that it plots either the numerical *and* the exact solution, *or* the error.

) Make animations where you illustrate how the error $e_i^n = u_e(x_i, t_n) - u_i^n$ develops and increases in time. Also make animations of u and u_e simultaneously.

Hint 1. Quite long time simulations are needed in order to display significant discrepancies between the numerical and exact solution.

Hint 2. A possible set of parameters is $L = 12$, $m = 9$, $c = 2$, $A = 1$, $N_x = 80$, $\tau = 0.8$. The error mesh function e^n can be simulated for 10 periods, while 10-30 periods are needed to show significant differences between the curves for numerical and exact solution.

Filename: `wave_standing.py`.

Remarks. The important parameters for numerical quality are C and k where $C = c\Delta t/\Delta x$ is the Courant number and k is defined above (proportional to how many mesh points we have per wave length in space, see Section 10.4 for explanation).

Exercise 2: Add storage of solution in a user action function

Extend the `plot_u` function in the file `wave1D_u0_s.py` to also store the solution `u` in a list. To this end, declare `all_u` as an empty list in the `viz` function outside `plot_u`, and perform an append operation inside the `plot_u` function. Note that a function, like `plot_u`, inside another function, like `viz`, receives all local variables in `viz` function, including `all_u`, even when `plot_u` is called (as `user_action`) in the `solver` function. Test both `all_u.append(u)` and `all_u.append(u.copy())`. Why does one of these constructions fail to store the solution correctly? Let the `viz` function return the `all_u` list converted to a two-dimensional numpy array. Filename: `wave1D_u0_s_store.py`.

Exercise 3: Use a class for the user action function

Redo Exercise 2 using a class for the user action function. That is, define a class `Action` where the `all_u` list is an attribute, and implement the user action function as a method (the special method `__call__` is a natural choice). This version avoids that the user action function depends on parameters defined outside the function (such as `all_u` in Exercise 2). Filename: `wave1D_u0_s_class.py`.

Exercise 4: Compare several Courant numbers in one figure

The goal of this exercise is to make movies where several curves, corresponding to different Courant numbers, are visualized. Import the `solver` function from the `wave1D_u0_s` movie in a new file `wave_compare.py`. Reimplement the `plot_u` function such that it can take a list of `C` values as argument and create a movie with solutions corresponding to the given `C` values. The `plot_u` function must be changed to store the solution in an array (see Exercise 2 or 3 for details). The `solver` must be computed for each value of the Courant number, and one must run through each time step and plot all the spatial solution curves in one figure and store it in a file.

The challenge in such a visualization is to ensure that the curves in the movie corresponds to the same time point. The easiest remedy is to keep the time step and space resolution constant and change the wave velocity c to change the Courant number. Filename: `wave_numerics_comparison.py`.

Project 5: Calculus with 1D mesh functions

This project explores integration and differentiation of mesh functions using both scalar and vectorized implementations. We are given a mesh function u on a spatial one-dimensional mesh $x_i = i\Delta x$, $i = 0, \dots, N_x$, over the interval $[0, L]$.

) Define the discrete derivative of f_i by using centered differences at internal mesh points and one-sided differences at the end points. Implement a scalar version of the computation in a Python function and supply a nose test for the near case $f(x) = 4x - 2.5$ where the discrete derivative should be exact.

) Vectorize the implementation of the discrete derivative. Extend the nose test to check the validity of the implementation.

) To compute the discrete integral F_i of f_i , we assume that the mesh function f_i varies linearly between the mesh points. Let $f(x)$ be such a linear interpolant of f_i . We then have

$$F_i = \int_{x_0}^{x_i} f(x) dx.$$

The exact integral of a piecewise linear function $f(x)$ is given by the Trapezoidal rule. Show that if F_i is already computed, we can find F_{i+1} from

$$F_{i+1} = F_i + \frac{1}{2}(f_i + f_{i+1})\Delta x.$$

Write a function for a scalar implementation of the discrete integral as a mesh function. That is, the function should return F_i for $i = 0, \dots, N_x$. For a nose test one can use the fact that the above defined discrete integral of a linear function (say $f(x) = 4x - 2.5$) is exact.

) Vectorize the implementation of the discrete integral. Extend the nose test to check the validity of the implementation.

hint. Interpret the recursive formula for F_{i+1} as a sum. Make an array with each element of the sum and use the "cumsum" (`numpy.cumsum`) operation to compute the accumulative sum: `numpy.cumsum([1,3,5])` is `[1,4,9]`.

) Create a class `MeshCalculus` that can integrate and differentiate mesh functions. The class can just define some methods that call the previously implemented Python functions. Here is an example on the usage:

```
import numpy as np
calc = MeshCalculus(vectorized=True)
x = np.linspace(0, 1, 11)          # mesh
f = np.exp(x)                      # mesh function
df = calc.differentiate(f, x)      # discrete derivative
F = calc.integrate(f, x)           # discrete anti-derivative
```

Filename: `mesh_calculus_1D.py`.

6 Generalization: reflecting boundaries

The boundary condition $u = 0$ makes u change sign at the boundary; the condition $u_x = 0$ perfectly reflects the wave, see a [web page](#)⁷ or [file](#)⁸ for demonstration. Our next task is to explain how to implement boundary condition $u_x = 0$, which is more complicated to express numerically and also to implement than a given value of u . We shall present two recipes for implementing $u_x = 0$ in a finite difference scheme, one based on the modified stencil at the boundary, and another one based on extending the mesh with ghost cells and ghost points.

6.1 Neumann boundary condition

When a wave hits a boundary and is to be reflected back, one appropriate boundary condition

$$\frac{\partial u}{\partial n} \equiv \mathbf{n} \cdot \nabla u = 0.$$

The derivative $\partial/\partial n$ is in the outward normal direction from a general boundary. For a 1D domain $[0, L]$, we have that

$$\left. \frac{\partial}{\partial n} \right|_{x=L} = \frac{\partial}{\partial x}, \quad \left. \frac{\partial}{\partial n} \right|_{x=0} = -\frac{\partial}{\partial x}.$$

Boundary condition terminology.

Boundary conditions that specify the value of $\partial u/\partial n$, or shorter u_x , are known as [Neumann](#)^a conditions, while [Dirichlet conditions](#)^b refer to specifications of u . When the values are zero ($\partial u/\partial n = 0$ or $u = 0$) we speak about *homogeneous* Neumann or Dirichlet conditions.

^ahttp://en.wikipedia.org/wiki/Neumann_boundary_condition

^bhttp://en.wikipedia.org/wiki/Dirichlet_conditions

6.2 Discretization of derivatives at the boundary

How can we incorporate the condition (34) in the finite difference scheme? Since we have used central differences in all the other approximations to derivatives in the scheme, it is tempting to implement (34) at $x = 0$ and $t = t_n$ using a central difference

$$\frac{u_{-1}^n - u_1^n}{2\Delta x} = 0.$$

The problem is that u_{-1}^n is not a u value that is being computed since the point $x = -\Delta x$ is outside the mesh. However, if we combine (35) with the scheme for i

⁷http://tinyurl.com/k3sdbuv/pub/mov-wave/demo_BC-gaussian/index.html

⁸http://tinyurl.com/k3sdbuv/pub/mov-wave/demo_BC-gaussian/movie.flv

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + C^2 (u_{i+1}^n - 2u_i^n + u_{i-1}^n), \quad (36)$$

we can eliminate the fictitious value u_{-1}^n . We see that $u_{-1}^n = u_1^n$ from (35), which can be used in (36) to arrive at a modified scheme for the boundary point u_0^{n+1} :

$$u_0^{n+1} = -u_0^{n-1} + 2u_0^n + 2C^2 (u_1^n - u_0^n), \quad i = 0. \quad (37)$$

Figure 4 visualizes this equation for computing u_0^3 in terms of u_0^2 , u_0^1 , and u_1^2 .




Figure 4: Modified stencil at a boundary with a Neumann condition.

Similarly, (34) applied at $x = L$ is discretized by a central difference

$$\frac{u_{N_x+1}^n - u_{N_x-1}^n}{2\Delta x} = 0. \quad (38)$$

Combined with the scheme for $i = N_x$ we get a modified scheme for the boundary value $u_{N_x}^{n+1}$:

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + 2C^2 (u_{i-1}^n - u_i^n), \quad i = N_x. \quad (39)$$

The modification of the scheme at the boundary is also required for the special formula for the first time step. How the stencil moves through the mesh and is modified at the boundary can be illustrated by an animation in a [web page](#)⁹ or a [movie file](#)¹⁰.

⁹http://tinyurl.com/k3sdbuv/pub/mov-wave/wave1D.PDE.Neumann_stencil_gpl/index.html

¹⁰http://tinyurl.com/k3sdbuv/pub/mov-wave/wave1D.PDE.Neumann_stencil_gpl/movie.ogg

6.3 Implementation of Neumann conditions

The implementation of the special formulas for the boundary points can be done by using the general formula for the interior points also at the boundary but replacing u_{i-1}^n by u_{i+1}^n when computing u_i^{n+1} for $i = 0$ and u_{i+1}^n by u_{i-1}^n for $i = N_x$. This is achieved by just replacing the index $i - 1$ by $i + 1$ for $i = 0$ and $i + 1$ by $i - 1$ for $i = N_x$. In a program, we introduce variables to hold the offset indices: `im1` for $i-1$ and `ip1` for $i+1$. It is now just a matter of defining `im1` and `ip1` properly for the interior points and the boundary points. The coding for the latter reads

```
i = 0
ip1 = i+1
im1 = ip1 # i-1 -> i+1
u[i] = u_1[i] + C2*(u_1[im1] - 2*u_1[i] + u_1[ip1])

i = Nx
im1 = i-1
ip1 = im1 # i+1 -> i-1
u[i] = u_1[i] + C2*(u_1[im1] - 2*u_1[i] + u_1[ip1])
```

We can in fact create one loop over both the interior and boundary points and use only one updating formula:

```
for i in range(0, Nx+1):
    ip1 = i+1 if i < Nx else i-1
    im1 = i-1 if i > 0 else i+1
    u[i] = u_1[i] + C2*(u_1[im1] - 2*u_1[i] + u_1[ip1])
```

The program `wave1D_dn0.py`¹¹ contains a complete implementation of the 1D wave equation with boundary conditions $u_x = 0$ at $x = 0$ and $x = L$.

6.4 Index set notation

We shall introduce a special notation for index sets, consisting of writing $i \in \mathcal{I}_x$, instead of $i = 0, \dots, N_x$. Obviously, \mathcal{I}_x must be the set $\mathcal{I}_x = \{0, \dots, N_x\}$, but it is often advantageous to have a symbol for this set rather than specify all its elements. This saves writing and makes specification of algorithm implementation of computer code easier.

The first index in the set will be denoted \mathcal{I}_x^0 and the last \mathcal{I}_x^{-1} . Sometimes we need to count from the second element in the set, and the notation \mathcal{I}_x^+ is used. Correspondingly, \mathcal{I}_x^- means $\{0, \dots, N_x - 1\}$. All the indices corresponding to interior grid points are $\mathcal{I}_x^i = \{1, \dots, N_x - 1\}$. For the time domain it is natural to explicitly use 0 as the first index, so we will usually write $n = 0$ rather than $n = \mathcal{I}_t^0$. We also avoid notation like $x_{\mathcal{I}_x^{-1}}$ and will instead use $x_{\mathcal{I}_x^{-1}}$.

The Python code associated with index sets applies the following convention

¹¹http://tinyurl.com/jvzccfn/wave/wave1D_dn0.py

Notation	Python
\mathcal{I}_x	<code>Ix</code>
\mathcal{I}_x^0	<code>Ix[0]</code>
\mathcal{I}_x^{-1}	<code>Ix[-1]</code>
\mathcal{I}_x^-	<code>Ix[: -1]</code>
\mathcal{I}_x^+	<code>Ix[1:]</code>
\mathcal{I}_x^i	<code>Ix[1: -1]</code>

an important feature of the index set notation is that it keeps our formulas and code independent of how we count mesh points. For example, the notation $i \in \mathcal{I}_x$ or $i = \mathcal{I}_x^0$ remains the same whether \mathcal{I}_x is defined as above or as starting at 1, i.e., $\mathcal{I}_x = \{1, \dots, Q\}$. Similarly, we can in the code define `Ix=range(Nx+1)` or `Ix=range(1,Q)`, and expressions like `Ix[0]` and `Ix[1:-1]` remain correct. The application where the index set notation is convenient is conversion of code from a language where arrays have base index 0 (e.g., Python and C) to languages where the base index is 1 (e.g., MATLAB and Fortran). Another important application is implementation of Neumann conditions via ghost points (see next section).

For the current problem setting in the x, t plane, we work with the index sets

$$\mathcal{I}_x = \{0, \dots, N_x\}, \quad \mathcal{I}_t = \{0, \dots, N_t\}, \quad (40)$$

defined in Python as

```
Ix = range(0, Nx+1)
It = range(0, Nt+1)
```

A finite difference scheme can with the index set notation be specified as

$$\begin{aligned} u_i^{n+1} &= -u_i^{n-1} + 2u_i^n + C^2 (u_{i+1}^n - 2u_i^n + u_{i-1}^n), \quad i \in \mathcal{I}_x^i, \quad n \in \mathcal{I}_t^i, \\ u_i &= 0, \quad i = \mathcal{I}_x^0, \quad n \in \mathcal{I}_t^i, \\ u_i &= 0, \quad i = \mathcal{I}_x^{-1}, \quad n \in \mathcal{I}_t^i, \end{aligned}$$

and implemented by code like

```
for n in It[1:-1]:
    for i in Ix[1:-1]:
        u[i] = -u_2[i] + 2*u_1[i] + \
            C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1])
    i = Ix[0]; u[i] = 0
    i = Ix[-1]; u[i] = 0
```

Notice.

The program `wave1D_dn.py`^a applies the index set notation and solves the 1D wave equation $u_{tt} = c^2 u_{xx} + f(x, t)$ with quite general boundary and initial conditions:

- $x = 0$: $u = U_0(t)$ or $u_x = 0$

- $x = L$: $u = U_L(t)$ or $u_x = 0$
- $t = 0$: $u = I(x)$
- $t = 0$: $u_t = I(x)$

The program combines Dirichlet and Neumann conditions, scalar and vectorized implementation of schemes, and the index notation into one file of code. A lot of test examples are also included in the program:

- A rectangular plug profile as initial condition (easy to use as an example as the rectangle should jump one cell per time step $C = 1$, without any numerical errors).
- A Gaussian function as initial condition.
- A triangular profile as initial condition, which resembles the typical initial shape of a guitar string.
- A sinusoidal variation of u at $x = 0$ and either $u = 0$ or $u_x = 0$ at $x = L$.
- An exact analytical solution $u(x, t) = \cos(m\pi t/L) \sin(\frac{1}{2}m\pi x/L)$, can be used for convergence rate tests.

^ahttp://tinyurl.com/jvzzcfn/wave/wave1D_dn.py

6.5 Alternative implementation via ghost cells

Idea. Instead of modifying the scheme at the boundary, we can introduce points outside the domain such that the fictitious values u_{-1}^n and $u_{N_x+1}^n$ defined in the mesh. Adding the intervals $[-\Delta x, 0]$ and $[L, L + \Delta x]$, often referred to as *ghost cells*, to the mesh gives us all the needed mesh points, corresponding to $i = -1, 0, \dots, N_x, N_x + 1$. The extra points $i = -1$ and $i = N_x + 1$ are called *ghost points*, and values at these points, u_{-1}^n and $u_{N_x+1}^n$, are called *ghost values*.

The important idea is to ensure that we always have

$$u_{-1}^n = u_1^n \quad \text{and} \quad u_{N_x+1}^n = u_{N_x-1}^n,$$

because then the application of the standard scheme at a boundary point $i = N_x$ will be correct and guarantee that the solution is compatible with the boundary condition $u_x = 0$.

Implementation. The `u` array now needs extra elements corresponding to the ghost cells and points. Two new point values are needed:


```
u = zeros(Nx+3)
```

The arrays `u_1` and `u_2` must be defined accordingly.

Unfortunately, a major indexing problem arises with ghost cells. The reason is that Python indices *must* start at 0 and `u[-1]` will always mean the last element in `u`. This fact gives, apparently, a mismatch between the mathematical indices $i = -1, 0, \dots, N_x + 1$ and the Python indices running over `u`: $0, \dots, Nx+2$. The remedy is to change the mathematical notation of the scheme, as in

$$u_i^{n+1} = \dots, \quad i = 1, \dots, N_x + 1,$$

meaning that the ghost points correspond to $i = 0$ and $i = N_x + 1$. A better solution is to use the ideas of Section 6.4: we hide the specific index value in an index set and operate with inner and boundary points using the index set notation.

To this end, we define `u` with proper length and `Ix` to be the corresponding indices for the real physical points:

```
u = zeros(Nx+3)
Ix = range(1, u.shape[0]-1)
```

That is, the boundary points have indices `Ix[0]` and `Ix[-1]` (as before). We first update the solution at all physical mesh points (i.e., interior points in the mesh extended with ghost cells):

```
for i in Ix:
    u[i] = - u_2[i] + 2*u_1[i] + \
           C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1])
```

It remains to update the ghost points. For a boundary condition $u_x = 0$, the ghost value must equal to the value at the associated inner mesh point. Computer code makes this statement precise:

```
i = Ix[0]          # x=0 boundary
u[i-1] = u[i+1]
i = Ix[-1]        # x=L boundary
u[i+1] = u[i-1]
```

The physical solution to be plotted is now in `u[1:-1]`, so this slice is the quantity to be returned from a solver function. A complete implementation appears in the program `wave1D_dn0_ghost.py`¹².

Warning.

We have to be careful with how the spatial and temporal mesh points are stored. Say we let `x` be the physical mesh points,

¹²http://tinyurl.com/jvzzcfn/wave/wave1D/wave1D_dn0_ghost.py

```
x = linspace(0, L, Nx+1)
```

“Standard coding” of the initial condition,

```
for i in Ix:
    u_1[i] = I(x[i])
```

becomes wrong, since `u_1` and `x` have different lengths and the index `i` corresponds to two different mesh points. In fact, `x[i]` corresponds to `u[1+i]`. A correct implementation is

```
for i in Ix:
    u_1[i] = I(x[i-Ix[0]])
```

Similarly, a source term usually coded as `f(x[i], t[n])` is incorrect and should be defined to be the physical points.

An alternative remedy is to let `x` also cover the ghost points such that `u[i]` is the value at `x[i]`. This is the recommended approach.

The ghost cell is only added to the boundary where we have a Neumann condition. Suppose we have a Dirichlet condition at $x = L$ and a homogeneous Neumann condition at $x = 0$. The relevant implementation then becomes

```
u = zeros(Nx+2)
Ix = range(0, u.shape[0]-1)
...
for i in Ix[:-1]:
    u[i] = - u_2[i] + 2*u_1[i] + \
           C2*(u_1[i-1] - 2*u_1[i] + u_1[i+1]) + \
           dt2*f(x[i], t[n])
i = Ix[-1]
u[i] = U_0          # set Dirichlet value
i = Ix[0]
u[i+1] = u[i-1]    # update ghost value
```

The physical solution to be plotted is now in `u[1:]`.

7 Generalization: variable wave velocity

Our next generalization of the 1D wave equation (1) or (17) is to allow for a variable wave velocity c : $c = c(x)$, usually motivated by wave motion in a medium composed of different physical media with different properties for propagating waves and hence different wave velocities c . Figure

7.1 The model PDE with a variable coefficient

Instead of working with the squared quantity $c^2(x)$ we shall for convenience introduce $q(x) = c^2(x)$. A 1D wave equation with variable wave velocity often takes the form




figure 5: Left: wave entering another medium; right: transmitted and reflected wave .

$$\frac{\partial^2 u}{\partial t^2} = \frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right) + f(x, t). \quad (41)$$

his equation sampled at a mesh point (x_i, t_n) reads

$$\frac{\partial^2}{\partial t^2} u(x_i, t_n) = \frac{\partial}{\partial x} \left(q(x_i) \frac{\partial}{\partial x} u(x_i, t_n) \right) + f(x_i, t_n),$$

here the only new term is

$$\frac{\partial}{\partial x} \left(q(x_i) \frac{\partial}{\partial x} u(x_i, t_n) \right) = \left[\frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right) \right]_i^n.$$

2 Discretizing the variable coefficient

he principal idea is to first discretize the outer derivative. Define

$$\phi = q(x) \frac{\partial u}{\partial x},$$

and use a centered derivative around $x = x_i$ for the derivative of ϕ :

$$\left[\frac{\partial \phi}{\partial x} \right]_i^n \approx \frac{\phi_{i+\frac{1}{2}} - \phi_{i-\frac{1}{2}}}{\Delta x} = [D_x \phi]_i^n.$$

hen discretize

$$\phi_{i+\frac{1}{2}} = q_{i+\frac{1}{2}} \left[\frac{\partial u}{\partial x} \right]_{i+\frac{1}{2}}^n \approx q_{i+\frac{1}{2}} \frac{u_{i+1}^n - u_i^n}{\Delta x} = [q D_x u]_{i+\frac{1}{2}}^n.$$

imilarly,

$$\phi_{i-\frac{1}{2}} = q_{i-\frac{1}{2}} \left[\frac{\partial u}{\partial x} \right]_{i-\frac{1}{2}}^n \approx q_{i-\frac{1}{2}} \frac{u_i^n - u_{i-1}^n}{\Delta x} = [q D_x u]_{i-\frac{1}{2}}^n.$$

These intermediate results are now combined to

$$\left[\frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right) \right]_i^n \approx \frac{1}{\Delta x^2} \left(q_{i+\frac{1}{2}} (u_{i+1}^n - u_i^n) - q_{i-\frac{1}{2}} (u_i^n - u_{i-1}^n) \right).$$

With operator notation we can write the discretization as

$$\left[\frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right) \right]_i^n \approx [D_x q D_x u]_i^n.$$

Remark.

Many are tempted to use the chain rule on the term $\frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right)$, but is not a good idea when discretizing such a term.

7.3 Computing the coefficient between mesh points

If q is a known function of x , we can easily evaluate $q_{i+\frac{1}{2}}$ simply as $q(x_{i+\frac{1}{2}}) = q(x_i + \frac{1}{2}\Delta x)$. However, in many cases c , and hence q , is only known as a discrete function, often at the mesh points x_i . Evaluating q between two mesh points x_i and x_{i+1} can then be done by averaging in three ways:

$$q_{i+\frac{1}{2}} \approx \frac{1}{2} (q_i + q_{i+1}) = [\bar{q}^x]_i, \quad (\text{arithmetic mean})$$

$$q_{i+\frac{1}{2}} \approx 2 \left(\frac{1}{q_i} + \frac{1}{q_{i+1}} \right)^{-1}, \quad (\text{harmonic mean})$$

$$q_{i+\frac{1}{2}} \approx (q_i q_{i+1})^{1/2}, \quad (\text{geometric mean})$$

The arithmetic mean in (44) is by far the most commonly used technique.

With the operator notation from (44) we can specify the discretized complete variable-coefficient wave equation in a compact way:

$$[D_t D_t u = D_x \bar{q}^x D_x u + f]_i^n.$$

From this notation we immediately see what kind of differences that each approximation with. The notation \bar{q}^x also specifies that the variable coefficient is approximated by an arithmetic mean, the definition being $[\bar{q}^x]_{i+\frac{1}{2}} = (q_i + q_{i+1})/2$. With the notation $[D_x q D_x u]_i^n$, we specify that q is evaluated directly at the mesh points: $q(x_{i-\frac{1}{2}})$ and $q(x_{i+\frac{1}{2}})$.

Before any implementation, it remains to solve (47) with respect to

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + \left(\frac{\Delta x}{\Delta t} \right)^2 \left(\frac{1}{2} (q_i + q_{i+1}) (u_{i+1}^n - u_i^n) - \frac{1}{2} (q_i + q_{i-1}) (u_i^n - u_{i-1}^n) \right) - \Delta t^2 f_i^n.$$

7.4 How a variable coefficient affects the stability

The stability criterion derived in Section 10.3 reads $\Delta t \leq \Delta x/c$. If $c = c(x)$, the criterion will depend on the spatial location. We must therefore choose a Δt that is small enough such that no mesh cell has $\Delta x/c(x) > \Delta t$. That is, we must use the largest c value in the criterion:

$$\Delta t \leq \beta \frac{\Delta x}{\max_{x \in [0, L]} c(x)}. \quad (49)$$

The parameter β is included as a safety factor: in some problems with a significantly varying c it turns out that one must choose $\beta < 1$ to have stable solutions ($\beta = 0.9$ may act as an all-round value).

7.5 Neumann condition and a variable coefficient

Consider a Neumann condition $\partial u/\partial x = 0$ at $x = L = N_x \Delta x$, discretized as

$$\frac{u_{i+1}^n - u_{i-1}^n}{2\Delta x} = 0 \quad u_{i+1}^n = u_{i-1}^n,$$

for $i = N_x$. Using the scheme (48) at the end point $i = N_x$ with $u_{i+1}^n = u_{i-1}^n$ results in

$$u_i^{n+1} = -u_i^{n-1} + 2u_i^n + \left(\frac{\Delta x}{\Delta t}\right)^2 \left(q_{i+\frac{1}{2}}(u_{i-1}^n - u_i^n) - q_{i-\frac{1}{2}}(u_i^n - u_{i-1}^n)\right) + \Delta t^2 f_i^n \quad (50)$$

$$= -u_i^{n-1} + 2u_i^n + \left(\frac{\Delta x}{\Delta t}\right)^2 (q_{i+\frac{1}{2}} + q_{i-\frac{1}{2}})(u_{i-1}^n - u_i^n) + \Delta t^2 f_i^n \quad (51)$$

$$\approx -u_i^{n-1} + 2u_i^n + \left(\frac{\Delta x}{\Delta t}\right)^2 2q_i(u_{i-1}^n - u_i^n) + \Delta t^2 f_i^n. \quad (52)$$

Here we used the approximation

$$\begin{aligned} q_{i+\frac{1}{2}} + q_{i-\frac{1}{2}} &= q_i + \left(\frac{dq}{dx}\right)_i \Delta x + \left(\frac{d^2q}{dx^2}\right)_i \Delta x^2 + \dots + \\ & \quad q_i - \left(\frac{dq}{dx}\right)_i \Delta x + \left(\frac{d^2q}{dx^2}\right)_i \Delta x^2 + \dots \\ &= 2q_i + 2\left(\frac{d^2q}{dx^2}\right)_i \Delta x^2 + \mathcal{O}(\Delta x^4) \\ &\approx 2q_i. \end{aligned} \quad (53)$$

An alternative derivation may apply the arithmetic mean of q in (48), leading to the term

$$(q_i + \frac{1}{2}(q_{i+1} + q_{i-1}))(u_{i-1}^n - u_i^n).$$

Since $\frac{1}{2}(q_{i+1} + q_{i-1}) = q_i + \mathcal{O}(\Delta x^2)$, we end up with $2q_i(u_{i-1}^n - u_i^n)$ for as we did above.

A common technique in implementations of $\partial u/\partial x = 0$ boundary conditions is to assume $dq/dx = 0$ as well. This implies $q_{i+1} = q_{i-1}$ and $q_{i+1/2} = q_i = N_x$. The implications for the scheme are

$$\begin{aligned} u_i^{n+1} &= -u_i^{n-1} + 2u_i^n + \\ & \quad \left(\frac{\Delta x}{\Delta t}\right)^2 \left(q_{i+\frac{1}{2}}(u_{i-1}^n - u_i^n) - q_{i-\frac{1}{2}}(u_i^n - u_{i-1}^n)\right) + \\ & \quad \Delta t^2 f_i^n \\ &= -u_i^{n-1} + 2u_i^n + \left(\frac{\Delta x}{\Delta t}\right)^2 2q_{i-\frac{1}{2}}(u_{i-1}^n - u_i^n) + \Delta t^2 f_i^n. \end{aligned}$$

7.6 Implementation of variable coefficients

The implementation of the scheme with a variable wave velocity may be done if the coefficient c is available as an array `c[i]` at the spatial mesh points. The following code is a straightforward implementation of the scheme (48):

```
for i in range(1, Nx):
    u[i] = -u_2[i] + 2*u_1[i] + \
           C2*(0.5*(q[i] + q[i+1])*(u_1[i+1] - u_1[i]) - \
              0.5*(q[i] + q[i-1])*(u_1[i] - u_1[i-1])) + \
           dt2*f(x[i], t[n])
```

The coefficient `C2` is now defined as `(dt/dx)**2` and *not* as the squared number since the wave velocity is variable and appears inside the parenthesis.

With Neumann conditions $u_x = 0$ at the boundary, we need to modify this scheme with the discrete version of the boundary condition, as shown in Section 7.5. Nevertheless, it would be convenient to reuse the formula for the interior points and just modify the indices `ip1=i+1` and `im1=i-1` as was done in Section 6.3. Assuming $dq/dx = 0$ at the boundaries, we can implement the scheme at the boundary with the following code.

```
i = 0
ip1 = i+1
im1 = ip1
u[i] = -u_2[i] + 2*u_1[i] + \
       C2*(0.5*(q[i] + q[ip1])*(u_1[ip1] - u_1[i]) - \
          0.5*(q[i] + q[im1])*(u_1[i] - u_1[im1])) + \
       dt2*f(x[i], t[n])
```

With ghost cells we can just reuse the formula for the interior points at the boundary, provided that the ghost values of both u and q are correctly updated to ensure $u_x = 0$ and $q_x = 0$.

A vectorized version of the scheme with a variable coefficient at internal points in the mesh becomes

```
u[1:-1] = - u_2[1:-1] + 2*u_1[1:-1] + \
          C2*(0.5*(q[1:-1] + q[2:])*u_1[2:] - u_1[1:-1]) - \
          0.5*(q[1:-1] + q[1:-2])*u_1[1:-1] - u_1[1:-2]) + \
          dt2*f(x[1:-1], t[n])
```

7 A more general model PDE with variable coefficients

Sometimes a wave PDE has a variable coefficient also in front of the time-derivative term:

$$\varrho(x) \frac{\partial^2 u}{\partial t^2} = \frac{\partial}{\partial x} \left(q(x) \frac{\partial u}{\partial x} \right) + f(x, t). \quad (56)$$

The natural scheme is

$$[\varrho D_t D_t u = D_x \bar{q}^x D_x u + f]_i^n. \quad (57)$$

We realize that the ϱ coefficient poses no particular difficulty because the only value ϱ_i^n enters the formula above (when written out). There is hence no need for any averaging of ϱ . Often, ϱ will be moved to the right-hand side, also without any difficulty:

$$[D_t D_t u = \varrho^{-1} D_x \bar{q}^x D_x u + f]_i^n. \quad (58)$$

8 Generalization: damping

Waves die out by two mechanisms. In 2D and 3D the energy of the wave spreads out in space, and energy conservation then requires the amplitude to decrease. This effect is not present in 1D. Damping is another cause of amplitude reduction. For example, the vibrations of a string die out because of damping due to air resistance and non-elastic effects in the string.

The simplest way of including damping is to add a first-order derivative to the equation (in the same way as friction forces enter a vibrating mechanical system):

$$\frac{\partial^2 u}{\partial t^2} + b \frac{\partial u}{\partial t} = c^2 \frac{\partial^2 u}{\partial x^2} + f(x, t), \quad (59)$$

where $b \geq 0$ is a prescribed damping coefficient.

A typical discretization of (59) in terms of centered differences reads

$$[D_t D_t u + b D_{2t} u = c^2 D_x D_x u + f]_i^n. \quad (60)$$

Writing out the equation and solving for the unknown u_i^{n+1} gives the scheme

$$u_i^{n+1} = (1 + \frac{1}{2} b \Delta t)^{-1} \left((\frac{1}{2} b \Delta t - 1) u_i^{n-1} + 2u_i^n + C^2 (u_{i+1}^n - 2u_i^n + u_{i-1}^n) \right) +$$

for $i \in \mathcal{I}_x^i$ and $n \geq 1$. New equations must be derived for u_i^1 , and for boundary points in case of Neumann conditions.

The damping is very small in many wave phenomena and then only for very long time simulations. This makes the standard wave equation damping relevant for a lot of applications.

8 Building a general 1D wave equation solver

The program `wave1D_dn_vc.py`¹³ is a fairly general code for 1D wave propagation problems that targets the following initial-boundary value problem

$$\begin{aligned} u_t &= (c^2(x) u_x)_x + f(x, t), & x \in (0, L), \quad t \in (0, T] \\ u(x, 0) &= I(x), & x \in [0, L] \\ u_t(x, 0) &= V(t), & x \in [0, L] \\ u(0, t) &= U_0(t) \text{ or } u_x(0, t) = 0, & t \in (0, T] \\ u(L, t) &= U_L(t) \text{ or } u_x(L, t) = 0, & t \in (0, T] \end{aligned}$$

The `solver` function is a natural extension of the simplest `solver1D` in the initial `wave1D_u0_s.py` program, extended with Neumann boundary conditions ($u_x = 0$), a possibly time-varying boundary condition on u ($U_L(t)$), and a variable wave velocity. The different code segments where these extensions are shown and commented upon in the preceding sections.

The vectorization is only applied inside the time loop, not for the boundary condition or the first time steps, since this initial work is negligible for long time simulations in 1D problems.

The following sections explain various more advanced programming techniques applied in the general 1D wave equation solver.

8.1 User action function as a class

A useful feature in the `wave1D_dn_vc.py` program is the specification of a `user_action` function as a class. Although the `plot_u` function in the previous `wave1D*.py` programs remembers the local variable `viz` function, it is a cleaner solution to store the needed variables together in the function, which is exactly what a class offers.

A class for flexible plotting, cleaning up files, and making a movie from the function `viz` and `plot_u` did can be coded as follows:

¹³http://tinyurl.com/jvzzcfn/wave/wave1D_dn_vc.py

```

class PlotSolution:
    """
    Class for the user_action function in solver.
    Visualizes the solution only.
    """
    def __init__(self,
                 casename='tmp', # Prefix in filenames
                 umin=-1, umax=1, # Fixed range of y axis
                 pause_between_frames=None, # Movie speed
                 backend='matplotlib', # or 'gnuplot'
                 screen_movie=True, # Show movie on screen?
                 title='', # Extra message in title
                 every_frame=1): # Show every_frame frame
        self.casename = casename
        self.yaxis = [umin, umax]
        self.pause = pause_between_frames
        module = 'scitools.easyviz.' + backend + '_'
        exec('import %s as plt' % module)
        self.plt = plt
        self.screen_movie = screen_movie
        self.title = title
        self.every_frame = every_frame

        # Clean up old movie frames
        for filename in glob('frame_*.png'):
            os.remove(filename)

    def __call__(self, u, x, t, n):
        if n % self.every_frame != 0:
            return
        title = 't=%f' % t[n]
        if self.title:
            title = self.title + ' ' + title
        self.plt.plot(x, u, 'r-',
                    xlabel='x', ylabel='u',
                    axis=[x[0], x[-1]],
                    self.yaxis[0], self.yaxis[1]],
                    title=title,
                    show=self.screen_movie)

        # pause
        if t[n] == 0:
            time.sleep(2) # let initial condition stay 2 s
        else:
            if self.pause is None:
                pause = 0.2 if u.size < 100 else 0
            time.sleep(pause)

        self.plt.savefig('%s_frame_%04d.png' % (self.casename, n))

```

Understanding this class requires quite some familiarity with Python in general and class programming in particular.

The constructor shows how we can flexibly import the plotting engine as typically) `scitools.easyviz.gnuplot_` or `scitools.easyviz.matplotlib_` (note the trailing underscore). With the `screen_movie` parameter we can suppress displaying each movie frame on the screen. Alternatively, for slow movies associated with fine meshes, one can set `every_frame` to, e.g., 10, causing every 10 frames to be shown.

The `__call__` method makes `PlotSolution` instances behave like `fit` so we can just pass an instance, say `p`, as the `user_action` argument to the `solver` function, and any call to `user_action` will be a call to `p.__call__`.

8.2 Pulse propagation in two media

The function `pulse` in `wave1D_dn_vc.py` demonstrates wave motion in heterogeneous media where c varies. One can specify an interval where the wave velocity is decreased by a factor `slowness_factor` (or increased by making this less than one). Four types of initial conditions are available: a rectangular pulse (`plug`), a Gaussian function (`gaussian`), a "cosine hat" consisting of a period of the cosine function (`cosinehat`), and half a period of a cosine function (`half-cosinehat`). These peak-shaped initial conditions can be placed in the middle (`loc='center'`) or at the left end (`loc='left'`) of the domain. The `pulse` function is a flexible tool for playing around with various wave velocities and location of a medium with a different wave velocity:

```

def pulse(C=1, Nx=200, animate=True, version='vectorized', T=2,
         loc='center', pulse_tp='gaussian', slowness_factor=2,
         medium=[0.7, 0.9], every_frame=1, sigma=0.05):
    """
    Various peaked-shaped initial conditions on [0,1].
    Wave velocity is decreased by the slowness_factor inside
    the medium. The loc parameter can be 'center' or 'left',
    depending on where the initial pulse is to be located.
    The sigma parameter governs the width of the pulse.
    """
    # Use scaled parameters: L=1 for domain length, c_0=1
    # for wave velocity outside the domain.
    L = 1.0
    c_0 = 1.0
    if loc == 'center':
        xc = L/2
    elif loc == 'left':
        xc = 0

    if pulse_tp in ('gaussian', 'Gaussian'):
        def I(x):
            return exp(-0.5*((x-xc)/sigma)**2)
    elif pulse_tp == 'plug':
        def I(x):
            return 0 if abs(x-xc) > sigma else 1
    elif pulse_tp == 'cosinehat':
        def I(x):
            # One period of a cosine
            w = 2
            a = w*sigma
            return 0.5*(1 + cos(pi*(x-xc)/a)) \
                if xc - a <= x <= xc + a else 0
    elif pulse_tp == 'half-cosinehat':
        def I(x):
            # Half a period of a cosine
            w = 4
            a = w*sigma

```

```

    return cos(pi*(x-xc)/a) \
        if xc - 0.5*a <= x <= xc + 0.5*a else 0
else:
    raise ValueError('Wrong pulse_tp="%s"' % pulse_tp)

def c(x):
    return c_0/slowness_factor \
        if medium[0] <= x <= medium[1] else c_0

umin=-0.5; umax=1.5*I(xc)
casename = '%s_Nx%s_sf%s' % \
    (pulse_tp, Nx, slowness_factor)
action = PlotMediumAndSolution(
    medium, casename=casename, umin=umin, umax=umax,
    every_frame=every_frame, screen_movie=animate)

solver(I=I, V=None, f=None, c=c, U_0=None, U_L=None,
    L=L, Nx=Nx, C=C, T=T,
    user_action=action, version=version,
    dt_safety_factor=1)

```

The `PlotMediumAndSolution` class used here is a subclass of `PlotSolution` here the medium with reduced c value, as specified by the `medium` interval, is visualized in the plots.

The reader is encouraged to play around with the `pulse` function:

```

>>> import wave1D_dn_vc as w
>>> w.pulse(loc='left', pulse_tp='cosinehat', Nx=50, every_frame=10)

```

to easily kill the graphics by Ctrl-C and restart a new simulation it might be easier to run the above two statements from the command line with

```
terminal> python -c 'import wave1D_dn_vc as w; w.pulse(...)'
```

Exercises

Exercise 6: Find the analytical solution to a damped wave equation

Consider the wave equation with damping (59). The goal is to find an exact solution to a wave problem with damping. A starting point is the standing wave solution from Exercise 1. It becomes necessary to include a damping term e^{-ct} and also have both a sine and cosine component in time:

$$u_e(x, t) = e^{-\beta t} \sin kx (A \cos \omega t + B \sin \omega t).$$

Find k from the boundary conditions $u(0, t) = u(L, t) = 0$. Then use the PDE to find constraints on β , ω , A , and B . Set up a complete initial-boundary value problem and its solution. Filename: `damped_waves.pdf`.

Problem 7: Explore symmetry boundary conditions

Consider the simple "plug" wave where $\Omega = [-L, L]$ and

$$I(x) = \begin{cases} 1, & x \in [-\delta, \delta], \\ 0, & \text{otherwise} \end{cases}$$

for some number $0 < \delta < L$. The other initial condition is $u_t(x, 0) = 0$ as is no source term f . The boundary conditions can be set to $u = 0$. The problem to this is symmetric around $x = 0$. This means that we can solve the wave process in only the half of the domain $[0, L]$.

a) Argue why the symmetry boundary condition is $u_x = 0$ at $x = 0$.

Hint. Symmetry of a function about $x = x_0$ means that $f(x_0 + h) = f(x_0 - h)$.

b) Perform simulations of the complete wave problem from on $[-L, L]$ after, utilize the symmetry of the solution and run a simulation in half domain $[0, L]$, using a boundary condition at $x = 0$. Compare the two and make sure that they are the same.

c) Prove the symmetry property of the solution by setting up the complete initial-boundary value problem and showing that if $u(x, t)$ is a solution then also $u(-x, t)$ is a solution.

Filename: `wave1D_symmetric`.

Exercise 8: Send pulse waves through a layered medium

Use the `pulse` function in `wave1D_dn_vc.py` to investigate sending a pulse located with its peak at $x = 0$, through the medium to the right where another medium for $x \in [0.7, 0.9]$ where the wave velocity is decreased by a factor s_f . Report what happens with a Gaussian pulse, a "cosine hat" pulse, and a "plug" pulse for resolutions $N_x = 40, 80, 160$ and $s_f = 2, 4$. Use $C = 1$ in the medium outside $[0.7, 0.9]$. Simulate until $t = 1$. Filename: `pulse1D.py`.

Exercise 9: Compare discretizations of a Neumann condition

We have a 1D wave equation with variable wave velocity: $u_t = (q(x)u_x)_x$. The Neumann condition $u_x = 0$ at $x = 0, L$ can be discretized as shown in (52) as

The aim of this exercise is to examine the rate of the numerical error by using different ways of discretizing the Neumann condition. As test function $q = 1 + (x - L/2)^4$ can be used, with $f(x, t)$ adapted such that the solution has a simple form, say $u(x, t) = \cos(\pi x/L) \cos(\omega t)$ for some $\omega = \sqrt{q} \pi/L$.

) Perform numerical experiments and find the convergence rate of the error using the approximation and (55).

) Switch to $q(x) = \cos(\pi x/L)$, which is symmetric at $x = 0, L$, and check the convergence rate of the scheme (55). Now, $q_{i-1/2}$ is a 2nd-order approximation to q , $q_{i-1/2} = q_i + 0.25q_i''\Delta x^2 + \dots$, because $q_i' = 0$ for $i = N_x$ (a similar argument can be applied to the case $i = 0$).

) A third discretization can be based on a simple and convenient, but less accurate, one-sided difference: $u_i - u_{i-1} = 0$ at $i = N_x$ and $u_{i+1} - u_i = 0$ at $i = 0$. Derive the resulting scheme in detail and implement it. Run experiments to establish the rate of convergence.

) A fourth technique is to view the scheme as

$$[D_t D_t u]_i^n = \frac{1}{\Delta x} \left([q D_x u]_{i+\frac{1}{2}}^n - [q D_x u]_{i-\frac{1}{2}}^n \right) + [f]_i^n,$$

and place the boundary at $x_{i+\frac{1}{2}}$, $i = N_x$, instead of exactly at the physical boundary. With this idea, we can just set $[q D_x u]_{i+\frac{1}{2}}^n = 0$. Derive the complete scheme using this technique. The implementation of the boundary condition at $x = -\Delta x/2$ is $\mathcal{O}(\Delta x^2)$ accurate, but the interesting question is what impact the movement of the boundary has on the convergence rate (compute the errors as usual over the entire mesh).

10 Analysis of the difference equations

10.1 Properties of the solution of the wave equation

The wave equation

$$\frac{\partial^2 u}{\partial t^2} = c^2 \frac{\partial^2 u}{\partial x^2}$$

has solutions of the form

$$u(x, t) = g_R(x - ct) + g_L(x + ct), \quad (67)$$

for any functions g_R and g_L sufficiently smooth to be differentiated twice. The result follows from inserting (67) in the wave equation. A function of the form $g_R(x - ct)$ represents a signal moving to the right in time with constant velocity c .

This feature can be explained as follows. At time $t = 0$ the signal looks like $g_R(x)$. Introducing a moving x axis with coordinates $\xi = x - ct$, we see the function $g_R(\xi)$ is "at rest" in the ξ coordinate system, and the shape is always the same. Say the $g_R(\xi)$ function has a peak at $\xi = 0$. This peak is located at $x = ct$, which means that it moves with the velocity $dx/dt = c$ in the x coordinate system. Similarly, $g_L(x + ct)$ is a function initially with shape

$g_L(x)$ that moves in the negative x direction with constant velocity c (in $\xi = x + ct$, look at the point $\xi = 0$, $x = -ct$, which has velocity $dx/dt = -c$).

With the particular initial conditions

$$u(x, 0) = I(x), \quad \frac{\partial}{\partial t} u(x, 0) = 0,$$

we get, with u as in (67),

$$g_R(x) + g_L(x) = I(x), \quad -cg_R'(x) + cg_L'(x) = 0,$$

which have the solution $g_R = g_L = I/2$, and consequently

$$u(x, t) = \frac{1}{2}I(x - ct) + \frac{1}{2}I(x + ct).$$

The interpretation of (68) is that the initial shape of u is split into two each with the same shape as I but half of the initial amplitude. One traveling to the left and the other one to the right.

The solution has two important physical features: constant amplitude left and right wave, and constant velocity of these two waves. It turns out that the numerical solution will also preserve the constant amplitude, but this depends on the mesh parameters Δt and Δx .

The solution (68) will be influenced by boundary conditions when the $\frac{1}{2}I(x - ct)$ and $\frac{1}{2}I(x + ct)$ hit the boundaries and get, e.g., reflected by the domain. However, when $I(x)$ is nonzero only in a small part in the middle of the spatial domain $[0, L]$, which means that the boundaries are placed far from the initial disturbance of u , the solution (68) is very clearly observed in simulation.

A useful representation of solutions of wave equations is a linear combination of sine and/or cosine waves. Such a sum of waves is a solution if the governing PDE is linear and each sine or cosine wave fulfills the equation. To ease our calculations by hand we shall work with complex exponential functions instead of real-valued sine or cosine functions. The real part of complex exponential will typically be taken as the physical relevant quantity (whenever a physical relevant quantity is strictly needed). The idea now is to build $I(x)$ of complex wave components e^{ikx} :

$$I(x) \approx \sum_{k \in K} b_k e^{ikx}.$$

Here, k is the frequency of a component, K is some set of all the discrete frequencies needed to approximate $I(x)$ well, and b_k are constants that must be determined. We will very seldom need to compute the b_k coefficients: most of the time we look for and the understanding of the numerical methods we want to use come from investigating how the PDE and the scheme treat a single complex e^{ikx} wave.

Letting the number of k values in K tend to infinity makes the sum converge to $I(x)$, and this sum is known as a *Fourier series* representing $I(x)$. Looking at (68), we see that the solution $u(x, t)$, when $I(x)$ is represented

s in (69), is also built of basic complex exponential wave components of the form $e^{ik(x\pm ct)}$ according to

$$u(x, t) = \frac{1}{2} \sum_{k \in K} b_k e^{ik(x-ct)} + \frac{1}{2} \sum_{k \in K} b_k e^{ik(x+ct)}. \quad (70)$$

It is common to introduce the frequency in time $\omega = kc$ and assume that $u(x, t)$ is a sum of basic wave components written as $e^{ikx - \omega t}$. (Observe that inserting such a wave component in the governing PDE reveals that $\omega^2 = k^2 c^2$, or $\omega \pm kc$, reflecting the two solutions: one ($+kc$) traveling to the right and the other ($-kc$) traveling to the left.)

0.2 More precise definition of Fourier representations

The quick intuitive introduction above to representing a function by a sum of sine and cosine waves suffices as background for the forthcoming material on analyzing a single wave component. However, to understand all details of how different wave components sum up to the analytical and numerical solution, a more precise mathematical treatment is helpful and therefore summarized below.

It is well known that periodic functions can be represented by Fourier series. A generalization of the Fourier series idea to non-periodic functions defined on the real line is the *Fourier transform*:

$$I(x) = \int_{-\infty}^{\infty} A(k) e^{ikx} dk, \quad (71)$$

$$A(k) = \int_{-\infty}^{\infty} I(x) e^{-ikx} dx. \quad (72)$$

The function $A(k)$ reflects the weight of each wave component e^{ikx} in an infinite sum of such wave components. That is, $A(k)$ reflects the frequency content in the function $I(x)$. Fourier transforms are particularly fundamental for analyzing and understanding time-varying signals.

The solution of the linear 1D wave PDE can be expressed as

$$u(x, t) = \int_{-\infty}^{\infty} A(k) e^{i(kx - \omega(k)t)} dx.$$

In a finite difference method, we represent u by a mesh function u_q^n , where n counts temporal mesh points and q counts the spatial ones (the usual counter for spatial points, i , is here already used as imaginary unit). Similarly, $I(x)$ is approximated by the mesh function I_q , $q = 0, \dots, N_x$. On a mesh, it does not make sense to work with wave components e^{ikx} for very large k , because the shortest possible sine or cosine wave that can be represented on a mesh with spacing Δx is the wave with wavelength $2\Delta x$ (the sine/cosine signal bumps up and down between each mesh point). The corresponding k value is $= 2\pi/(2\Delta x) = \pi/\Delta x$, known as the *Nyquist frequency*. Within the range of

relevant frequencies $(0, \pi/\Delta x]$ one defines the *discrete Fourier transform*¹ $N_x + 1$ discrete frequencies:

$$I_q = \frac{1}{N_x + 1} \sum_{k=0}^{N_x} A_k e^{i2\pi k j / (N_x + 1)}, \quad i = 0, \dots, N_x,$$

$$A_k = \sum_{q=0}^{N_x} I_q e^{-i2\pi k q / (N_x + 1)}, \quad k = 0, \dots, N_x + 1.$$

The A_k values is the discrete Fourier transform of the I_q values, and they are the inverse discrete Fourier transform of the A_k values.

The discrete Fourier transform is efficiently computed by the *Fast Fourier Transform* algorithm. For a real function $I(x)$ the relevant Python code for computing and plotting the discrete Fourier transform appears in the code below.

```
import numpy as np
from numpy import sin

def I(x):
    return sin(2*pi*x) + 0.5*sin(4*pi*x) + 0.1*sin(6*pi*x)

# Mesh
L = 10; Nx = 100
x = np.linspace(0, L, Nx+1)
dx = L/float(Nx)

# Discrete Fourier transform
A = np.fft.rfft(I(x))
A_amplitude = np.abs(A)

# Compute the corresponding frequencies
freqs = np.linspace(0, pi/dx, A_amplitude.size)

import matplotlib.pyplot as plt
plt.plot(freqs, A_amplitude)
plt.show()
```

10.3 Stability

The scheme

$$[D_t D_t u = c^2 D_x D_x u]_q^n$$

for the wave equation $u_t = c^2 u_{xx}$ allows basic wave components

$$u_q^n = e^{i(kx_q - \tilde{\omega}t_n)}$$

¹⁴http://en.wikipedia.org/wiki/Discrete_Fourier_transform

s solution, but it turns out that the frequency in time, $\tilde{\omega}$, is not equal to the exact $\omega = kc$. The idea now is to study how the scheme treats an arbitrary wave component with a given k . We ask two key questions:

- How accurate is $\tilde{\omega}$ compared to ω ?
- Does the amplitude of such a wave component preserve its (unit) amplitude, as it should, or does it get amplified or damped in time (due to a complex $\tilde{\omega}$)?

The following analysis will answer these questions. Note the need for using q as counter for the mesh point in x direction since i is already used as the imaginary unit (in this analysis).

Preliminary results. A key result needed in the investigations is the finite difference approximation of a second-order derivative acting on a complex wave component:

$$[D_t D_t e^{i\omega t}]^n = -\frac{4}{\Delta t^2} \sin^2\left(\frac{\omega \Delta t}{2}\right) e^{i\omega n \Delta t}.$$

By just changing symbols ($\omega \rightarrow k$, $t \rightarrow x$, $n \rightarrow q$) it follows that

$$[D_x D_x e^{ikx}]_q = -\frac{4}{\Delta x^2} \sin^2\left(\frac{k \Delta x}{2}\right) e^{ikq \Delta x}.$$

Numerical wave propagation. Inserting a basic wave component $u_q^n = e^{i(kx_q - \tilde{\omega}t_n)}$ in (75) results in the need to evaluate two expressions:

$$\begin{aligned} [D_t D_t e^{ikx} e^{-i\tilde{\omega}t}]_q^n &= [D_t D_t e^{-i\tilde{\omega}t}]_q^n e^{ikq \Delta x} \\ &= -\frac{4}{\Delta t^2} \sin^2\left(\frac{\tilde{\omega} \Delta t}{2}\right) e^{-i\tilde{\omega} n \Delta t} e^{ikq \Delta x} \end{aligned} \quad (76)$$

$$\begin{aligned} [D_x D_x e^{ikx} e^{-i\tilde{\omega}t}]_q^n &= [D_x D_x e^{ikx}]_q^n e^{-i\tilde{\omega} n \Delta t} \\ &= -\frac{4}{\Delta x^2} \sin^2\left(\frac{k \Delta x}{2}\right) e^{ikq \Delta x} e^{-i\tilde{\omega} n \Delta t}. \end{aligned} \quad (77)$$

Then the complete scheme,

$$[D_t D_t e^{ikx} e^{-i\tilde{\omega}t}]_q^n = c^2 [D_x D_x e^{ikx} e^{-i\tilde{\omega}t}]_q^n$$

leads to the following equation for the unknown numerical frequency $\tilde{\omega}$ (after dividing by $-e^{ikx} e^{-i\tilde{\omega}t}$):

$$\frac{4}{\Delta t^2} \sin^2\left(\frac{\tilde{\omega} \Delta t}{2}\right) = c^2 \frac{4}{\Delta x^2} \sin^2\left(\frac{k \Delta x}{2}\right),$$

†

$$\sin^2\left(\frac{\tilde{\omega} \Delta t}{2}\right) = C^2 \sin^2\left(\frac{k \Delta x}{2}\right),$$

where

$$C = \frac{c \Delta t}{\Delta x}$$

is the Courant number. Taking the square root of (78) yields

$$\sin\left(\frac{\tilde{\omega} \Delta t}{2}\right) = C \sin\left(\frac{k \Delta x}{2}\right),$$

Since the exact ω is real it is reasonable to look for a real solution $\tilde{\omega}$. The right-hand side of (80) must then be in $[-1, 1]$ because the sine function on the left-hand side has values in $[-1, 1]$ for real $\tilde{\omega}$. The sine function on the right-hand side can attain the value 1 when

$$\frac{k \Delta x}{2} = m \frac{\pi}{2}, \quad m \in \mathbb{Z}.$$

With $m = 1$ we have $k \Delta x = \pi$, which means that the wavelength λ becomes $2 \Delta x$. This is the absolutely shortest wavelength that can be represented on the mesh: the wave jumps up and down between each mesh point. Values of $|m|$ are irrelevant since these correspond to k values whose wavelength is too short to be represented on a mesh with spacing Δx . For the shortest wave in the mesh, $\sin(k \Delta x / 2) = 1$, and we must require

$$C \leq 1.$$

Consider a right-hand side in (80) of magnitude larger than unity. The solution $\tilde{\omega}$ of (80) must then be a complex number $\tilde{\omega} = \tilde{\omega}_r + i \tilde{\omega}_i$ because the sine function is larger than unity for a complex argument. One can show that for any ω_i there will also be a corresponding solution with $-\omega_i$. The component with $\omega_i > 0$ gives an amplification factor $e^{\omega_i t}$ that grows exponentially. We cannot allow this and must therefore require $C \leq 1$ as a *stability criterion*.

Remark.

For smoother wave components with longer wave lengths per length Δx , the stability criterion (81) can in theory be relaxed. However, small round-off errors are always present in a numerical solution and these vary arbitrarily from mesh point to mesh point and can be viewed as unavoidable noise with wavelength $2 \Delta x$. As explained, $C > 1$ will for this very small noise lead to exponential growth of the shortest possible wave component in the mesh. This will therefore grow with time and destroy the whole solution.

0.4 Numerical dispersion relation

equation (80) can be solved with respect to $\tilde{\omega}$:

$$\tilde{\omega} = \frac{2}{\Delta t} \sin^{-1} \left(C \sin \left(\frac{k\Delta x}{2} \right) \right). \quad (82)$$

The relation between the numerical frequency $\tilde{\omega}$ and the other parameters k , c , Δx , and Δt is called a *numerical dispersion relation*. Correspondingly, $\omega = kc$ is the *analytical dispersion relation*.

The special case $C = 1$ deserves attention since then the right-hand side of (82) reduces to

$$\frac{2}{\Delta t} \frac{k\Delta x}{2} = \frac{1}{\Delta t} \frac{\omega\Delta x}{c} = \frac{\omega}{C} = \omega.$$

That is, $\tilde{\omega} = \omega$ and the numerical solution is exact at all mesh points regardless of x and Δt ! This implies that the numerical solution method is also an analytical solution method, at least for computing u at discrete points (the numerical method says nothing about the variation of u between the mesh points, and employing the common linear interpolation for extending the discrete solution gives a curve that deviates from the exact one).

For a closer examination of the error in the numerical dispersion relation when $C < 1$, we can study $\tilde{\omega} - \omega$, $\tilde{\omega}/\omega$, or the similar error measures in wave velocity: $\tilde{c} - c$ and \tilde{c}/c , where $c = \omega/k$ and $\tilde{c} = \tilde{\omega}/k$. It appears that the most convenient expression to work with is \tilde{c}/c :

$$\frac{\tilde{c}}{c} = \frac{1}{Cp} \sin^{-1} (C \sin p),$$

with $p = k\Delta x/2$ as a non-dimensional measure of the spatial frequency. In essence, p tells how many spatial mesh points we have per wave length in space of the wave component with frequency k (the wave length is $2\pi/k$). That is, p reflects how well the spatial variation of the wave component is resolved in the mesh. Wave components with wave length less than $2\Delta x$ ($2\pi/k < 2\Delta x$) are not visible in the mesh, so it does not make sense to have $p > \pi/2$.

We may introduce the function $r(C, p) = \tilde{c}/c$ for further investigation of numerical errors in the wave velocity:

$$r(C, p) = \frac{1}{Cp} \sin^{-1} (C \sin p), \quad C \in (0, 1], \quad p \in (0, \pi/2]. \quad (83)$$

This function is very well suited for plotting since it combines several parameters of the problem into a dependence on two non-dimensional numbers, C and p .

Defining

```
def r(C, p):
    return 2/(C*p)*asin(C*sin(p))
```




Figure 6: The fractional error in the wave velocity for different Courant numbers

we can plot $r(C, p)$ as a function of p for various values of C , see Figure 6. Note that the shortest waves have the most erroneous velocity, and that shorter waves move more slowly than they should.

With `sympy` we can also easily make a Taylor series expansion in the discretization parameter p :

```
>>> C, p = symbols('C p')
>>> rs = r(C, p).series(p, 0, 7)
>>> print rs
1 - p**2/6 + p**4/120 - p**6/5040 + C**2*p**2/6 -
C**2*p**4/12 + 13*C**2*p**6/720 + 3*C**4*p**4/40 -
C**4*p**6/16 + 5*C**6*p**6/112 + 0(p**7)
>>> # Factorize each term and drop the remainder 0(...) term
>>> rs_factored = [factor(term) for term in rs.lseries(p)]
>>> rs_factored = sum(rs_factored)
>>> print rs_factored
p**6*(C - 1)*(C + 1)*(225*C**4 - 90*C**2 + 1)/5040 +
p**4*(C - 1)*(C + 1)*(3*C - 1)*(3*C + 1)/120 +
p**2*(C - 1)*(C + 1)/6 + 1
```

We see that $C = 1$ makes all the terms in `rs_factored` vanish, except the constant term. Since we already know that the numerical solution is exact for $C = 1$, the remaining terms in the Taylor series expansion will also contain factors of $(C - 1)$ and cancel for $C = 1$.

From the `rs_factored` expression above we also see that the leading order terms in the error of this series expansion are

$$\frac{1}{6} \left(\frac{k\Delta x}{2} \right)^2 (C^2 - 1) = \frac{k^2}{24} (c^2\Delta t^2 - \Delta x^2), \quad (84)$$

pointing to an error $\mathcal{O}(\Delta t^2, \Delta x^2)$, which is compatible with the errors in the difference approximations ($D_t D_t$ and $D_x D_x$).

0.5 Extending the analysis to 2D and 3D

The typical analytical solution of a 2D wave equation

$$u_{tt} = c^2(u_{xx} + u_{yy}),$$

is a wave traveling in the direction of $\mathbf{k} = k_x \mathbf{i} + k_y \mathbf{j}$, where \mathbf{i} and \mathbf{j} are unit vectors in the x and y directions, respectively. Such a wave can be expressed by

$$u(x, y, t) = g(k_x x + k_y y - kct)$$

or some twice differentiable function g , or with $\omega = kc$, $k = |\mathbf{k}|$:

$$u(x, y, t) = g(k_x x + k_y y - \omega t).$$

We can in particular build a solution by adding complex Fourier components of the form

$$\exp(i(k_x x + k_y y - \omega t)).$$

A discrete 2D wave equation can be written as

$$[D_t D_t u = c^2(D_x D_x u + D_y D_y u)]_{q,r}^n. \quad (85)$$

This equation admits a Fourier component

$$u_{q,r}^n = \exp(i(k_x q \Delta x + k_y r \Delta y - \tilde{\omega} n \Delta t)), \quad (86)$$

as a solution. Letting the operators $D_t D_t$, $D_x D_x$, and $D_y D_y$ act on $u_{q,r}^n$ from (86) transforms (85) to

$$\frac{4}{\Delta t^2} \sin^2 \left(\frac{\tilde{\omega} \Delta t}{2} \right) = c^2 \frac{4}{\Delta x^2} \sin^2 \left(\frac{k_x \Delta x}{2} \right) + c^2 \frac{4}{\Delta y^2} \sin^2 \left(\frac{k_y \Delta y}{2} \right). \quad (87)$$

$$\sin^2 \left(\frac{\tilde{\omega} \Delta t}{2} \right) = C_x^2 \sin^2 p_x + C_y^2 \sin^2 p_y, \quad (88)$$

where we have eliminated the factor 4 and introduced the symbols

$$C_x = \frac{c^2 \Delta t^2}{\Delta x^2}, \quad C_y = \frac{c^2 \Delta t^2}{\Delta y^2}, \quad p_x = \frac{k_x \Delta x}{2}, \quad p_y = \frac{k_y \Delta y}{2}.$$

For a real-valued $\tilde{\omega}$ the right-hand side must be less than or equal to 1 in absolute value, requiring in general that

$$C_x^2 + C_y^2 \leq 1.$$

This gives the stability criterion, more commonly expressed directly as an inequality for the time step:

$$\Delta t \leq \frac{1}{c} \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} \right)^{-1/2}$$

A similar, straightforward analysis for the 3D case leads to

$$\Delta t \leq \frac{1}{c} \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right)^{-1/2}$$

In the case of a variable coefficient $c^2 = c^2(\mathbf{x})$, we must use the worst-case

$$\bar{c} = \sqrt{\max_{\mathbf{x} \in \Omega} c^2(\mathbf{x})}$$

in the stability criteria. Often, especially in the variable wave velocity case, it is wise to introduce a safety factor $\beta \in (0, 1]$ too:

$$\Delta t \leq \beta \frac{1}{\bar{c}} \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right)^{-1/2}$$

The exact numerical dispersion relations in 2D and 3D becomes, for a constant c ,

$$\begin{aligned} \tilde{\omega} &= \frac{2}{\Delta t} \sin^{-1} \left(\left(C_x^2 \sin^2 p_x + C_y^2 \sin^2 p_y \right)^{\frac{1}{2}} \right), \\ \tilde{\omega} &= \frac{2}{\Delta t} \sin^{-1} \left(\left(C_x^2 \sin^2 p_x + C_y^2 \sin^2 p_y + C_z^2 \sin^2 p_z \right)^{\frac{1}{2}} \right). \end{aligned}$$

We can visualize the numerical dispersion error in 2D much like we did in 1D. To this end, we need to reduce the number of parameters in $\tilde{\omega}$. The direction of the wave is parameterized by the polar angle θ , which means that

$$k_x = k \sin \theta, \quad k_y = k \cos \theta.$$

A simplification is to set $\Delta x = \Delta y = h$. Then $C_x = C_y = c\Delta t/h$, which we denote by C . Also,

$$p_x = \frac{1}{2} kh \cos \theta, \quad p_y = \frac{1}{2} kh \sin \theta.$$

The numerical frequency $\tilde{\omega}$ is now a function of three parameters:

- C reflecting the number cells a wave is displaced during a time step

- kh reflecting the number of cells per wave length in space
- θ expressing the direction of the wave

If we want to visualize the error in the numerical frequency. To avoid having Δt as a free parameter in $\tilde{\omega}$, we work with \tilde{c}/c , because the fraction $2/\Delta t$ is then rewritten as

$$\frac{2}{kc\Delta t} = \frac{2}{2kc\Delta th/h} = \frac{1}{Ckh}$$

and

$$\frac{\tilde{c}}{c} = \frac{1}{Ckh} \sin^{-1} \left(C \left(\sin^2\left(\frac{1}{2}kh \cos \theta\right) + \sin^2\left(\frac{1}{2}kh \sin \theta\right) \right)^{\frac{1}{2}} \right).$$

If we want to visualize this quantity as a function of kh and θ for some values of $C \leq 1$. It is instructive to make color contour plots of $1 - \tilde{c}/c$ in *polar coordinates* with θ as the angular coordinate and kh as the radial coordinate.

The stability criterion (89) becomes $C \leq C_{\max} = 1/\sqrt{2}$ in the present 2D case with the C defined above. Let us plot $1 - \tilde{c}/c$ in polar coordinates for $C_{\max}, 0.9C_{\max}, 0.5C_{\max}, 0.2C_{\max}$. The program below does the somewhat tricky work in Matplotlib, and the result appears in Figure 7. From the figure we can clearly see that the maximum C value gives the best results, and that waves whose propagation direction makes an angle of 45 degrees with an axis are the most accurate.

```
def dispersion_relation_2D(kh, theta, C):
    arg = C*sqrt(sin(0.5*kh*cos(theta))**2 +
                sin(0.5*kh*sin(theta))**2)
    c_frac = 2./(C*kh)*arcsin(arg)

    return c_frac

from numpy import exp, sin, cos, linspace, \
    pi, meshgrid, arcsin, sqrt
r = kh = linspace(0.001, pi, 101)
theta = linspace(0, 2*pi, 51)
r, theta = meshgrid(r, theta)

# Make 2x2 filled contour plots for 4 values of C
import matplotlib.pyplot as plt
C_max = 1/sqrt(2)
C = [[C_max, 0.9*C_max], [0.5*C_max, 0.2*C_max]]
fig, axes = plt.subplots(2, 2, subplot_kw=dict(polar=True))
for row in range(2):
    for column in range(2):
        error = 1 - dispersion_relation_2D(
            kh, theta, C[row][column])
        print error.min(), error.max()
        cax = axes[row][column].contourf(
            theta, r, error, 50, vmin=0, vmax=0.36)
        axes[row][column].set_xticks([])
        axes[row][column].set_yticks([])
```

```
# Add colorbar to the last plot
cbar = plt.colorbar(cax)
cbar.ax.set_ylabel('error in wave velocity')
plt.savefig('disp12D.png')
plt.savefig('disp12D.pdf')
plt.show()
```




Figure 7: Error in numerical dispersion in 2D.

11 Finite difference methods for 2D and 3D equations

A natural next step is to consider extensions of the methods for variants of the one-dimensional wave equation to two-dimensional (2D) and three-dimensional (3D) versions of the wave equation.

11.1 Multi-dimensional wave equations

The general wave equation in d space dimensions, with constant wave velocity c , can be written in the compact form

$$\frac{\partial^2 u}{\partial t^2} = c^2 \nabla^2 u \text{ for } \mathbf{x} \in \Omega \subset \mathbb{R}^d, t \in (0, T].$$

In a 2D problem ($d = 2$),

$$\nabla^2 u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2},$$

while in three space dimensions ($d = 3$),

$$\nabla^2 u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2}.$$

Many applications involve variable coefficients, and the general wave equation in d dimensions is in this case written as

$$\rho \frac{\partial^2 u}{\partial t^2} = \nabla \cdot (q \nabla u) + f \text{ for } \mathbf{x} \in \Omega \subset \mathbb{R}^d, t \in (0, T], \quad (97)$$

which in 2D becomes

$$\rho(x, y) \frac{\partial^2 u}{\partial t^2} = \frac{\partial}{\partial x} \left(q(x, y) \frac{\partial u}{\partial x} \right) + \frac{\partial}{\partial y} \left(q(x, y) \frac{\partial u}{\partial y} \right) + f(x, y, t). \quad (98)$$

to save some writing and space we may use the index notation, where subscript x , y , or z means differentiation with respect to that coordinate. For example,

$$\frac{\partial^2 u}{\partial t^2} = u_{tt},$$

$$\frac{\partial}{\partial y} \left(q(x, y) \frac{\partial u}{\partial y} \right) = (qu_y)_y.$$

The 3D versions of the two model PDEs, with and without variable coefficients, are now with the aid of the index notation for differentiation be stated as

$$u_{tt} = c^2(u_{xx} + u_{yy} + u_{zz}) + f, \quad (99)$$

$$\rho u_{tt} = (qu_x)_x + (qu_z)_z + f. \quad (100)$$

At each point of the boundary $\partial\Omega$ of Ω we need *one* boundary condition involving the unknown u . The boundary conditions are of three principal types:

1. u is prescribed ($u = 0$ or a known time variation for an incoming wave),
2. $\partial u / \partial n = \mathbf{n} \cdot \nabla u$ prescribed (zero for reflecting boundaries),
3. an open boundary condition (also called radiation condition) is specified to let waves travel undisturbed out of the domain, see Exercise ?? for details.

All the listed wave equations with *second-order* derivatives in time need *two* initial conditions:

1. $u = I$,
2. $u_t = V$.

11.2 Mesh

We introduce a mesh in time and in space. The mesh in time consists of points

$$t_0 = 0 < t_1 < \dots < t_{N_t},$$

often with a constant spacing $\Delta t = t_{n+1} - t_n$, $n \in \mathcal{I}_t^-$.

Finite difference methods are easy to implement on simple rectangle-shaped domains. More complicated shapes of the domain require substantially more advanced techniques and implementational efforts. On a rectangular-shaped domain mesh points are introduced separately in the various directions:

$$x_0 < x_1 < \dots < x_{N_x} \text{ in } x \text{ direction,}$$

$$y_0 < y_1 < \dots < y_{N_y} \text{ in } y \text{ direction,}$$

$$z_0 < z_1 < \dots < z_{N_z} \text{ in } z \text{ direction.}$$

We can write a general mesh point as (x_i, y_j, z_k, t_n) , with $i \in \mathcal{I}_x$, $j \in \mathcal{I}_y$, and $n \in \mathcal{I}_t$.

It is a very common choice to use constant mesh spacings: $\Delta x = x_i - x_{i-1}$, $\Delta y = y_j - y_{j-1}$, and $\Delta z = z_k - z_{k-1}$, $k \in \mathcal{I}_z^-$. With equal spacings one often introduces $h = \Delta x = \Delta y = \Delta z$.

The unknown u at mesh point (x_i, y_j, z_k, t_n) is denoted by $u_{i,j,k}^n$. In problems we just skip the z coordinate (by assuming no variation in the z direction: $\partial/\partial z = 0$) and write $u_{i,j}^n$.

11.3 Discretization

Two- and three-dimensional wave equations are easily discretized by using building blocks for discretization of 1D wave equations, because the multidimensional versions just contain terms of the same type that occurs in

Discretizing the PDEs. Equation (99) can be discretized as

$$[D_t D_t u = c^2(D_x D_x u + D_y D_y u + D_z D_z u) + f]_{i,j,k}^n.$$

A 2D version might be instructive to write out in detail:

$$[D_t D_t u = c^2(D_x D_x u + D_y D_y u) + f]_{i,j,k}^n,$$

which becomes

$$\frac{u_{i,j}^{n+1} - 2u_{i,j}^n + u_{i,j}^{n-1}}{\Delta t^2} = c^2 \frac{u_{i+1,j}^n - 2u_{i,j}^n + u_{i-1,j}^n}{\Delta x^2} + c^2 \frac{u_{i,j+1}^n - 2u_{i,j}^n + u_{i,j}^n}{\Delta y^2}$$

assuming as usual that all values at the time levels n and $n - 1$ are known, we can solve for the only unknown $u_{i,j}^{n+1}$. The result can be compactly written as

$$u_{i,j}^{n+1} = 2u_{i,j}^n + u_{i,j}^{n-1} + c^2 \Delta t^2 [D_x D_x u + D_y D_y u]_{i,j}^n. \quad (102)$$

As in the 1D case, we need to develop a special formula for $u_{i,j}^1$ where we combine the general scheme for $u_{i,j}^{n+1}$, when $n = 0$, with the discretization of the initial condition:

$$[D_{2t} u = V]_{i,j}^0 \Rightarrow u_{i,j}^{-1} = u_{i,j}^1 - 2\Delta t V_{i,j}.$$

The result becomes, in compact form,

$$u_{i,j}^{n+1} = u_{i,j}^n - 2\Delta V_{i,j} + \frac{1}{2} c^2 \Delta t^2 [D_x D_x u + D_y D_y u]_{i,j}^n. \quad (103)$$

The PDE (100) with variable coefficients is discretized term by term using the corresponding elements from the 1D case:

$$[\rho D_t D_t u = (D_x \bar{q}^x D_x u + D_y \bar{q}^y D_y u + D_z \bar{q}^z D_z u) + f]_{i,j,k}^n. \quad (104)$$

When written out and solved for the unknown $u_{i,j,k}^{n+1}$, one gets the scheme

$$\begin{aligned} u_{i,j,k}^{n+1} &= -u_{i,j,k}^{n-1} + 2u_{i,j,k}^n + \\ &= \frac{1}{\rho_{i,j,k}} \frac{1}{\Delta x^2} \left(\frac{1}{2} (q_{i,j,k} + q_{i+1,j,k}) (u_{i+1,j,k}^n - u_{i,j,k}^n) - \right. \\ &\quad \left. \frac{1}{2} (q_{i-1,j,k} + q_{i,j,k}) (u_{i,j,k}^n - u_{i-1,j,k}^n) \right) + \\ &= \frac{1}{\rho_{i,j,k}} \frac{1}{\Delta x^2} \left(\frac{1}{2} (q_{i,j,k} + q_{i,j+1,k}) (u_{i,j+1,k}^n - u_{i,j,k}^n) - \right. \\ &\quad \left. \frac{1}{2} (q_{i,j-1,k} + q_{i,j,k}) (u_{i,j,k}^n - u_{i,j-1,k}^n) \right) + \\ &= \frac{1}{\rho_{i,j,k}} \frac{1}{\Delta x^2} \left(\frac{1}{2} (q_{i,j,k} + q_{i,j,k+1}) (u_{i,j,k+1}^n - u_{i,j,k}^n) - \right. \\ &\quad \left. \frac{1}{2} (q_{i,j,k-1} + q_{i,j,k}) (u_{i,j,k}^n - u_{i,j,k-1}^n) \right) + \\ &+ \Delta t^2 f_{i,j,k}^n. \end{aligned}$$

Also here we need to develop a special formula for $u_{i,j,k}^1$ by combining the scheme for $n = 0$ with the discrete initial condition, which is just a matter of inserting $u_{i,j,k}^{-1} = u_{i,j,k}^1 - 2\Delta t V_{i,j,k}$ in the scheme and solving for $u_{i,j,k}^1$.

Handling boundary conditions where u is known. The schemes listed above are valid for the internal points in the mesh. After updating these, we need to visit all the mesh points at the boundaries and set the prescribed u value.

Discretizing the Neumann condition. The condition $\partial u / \partial n = 0$ implemented in 1D by discretizing it with a $D_{2x} u$ centered difference, and then eliminating the fictitious u point outside the mesh by using the general at the boundary point. Alternatively, one can introduce ghost cells and a ghost value to for use in the Neumann condition. Exactly the same is reused in multi dimensions.

Consider $\partial u / \partial n = 0$ at a boundary $y = 0$. The normal direction is $-y$ direction, so

$$\frac{\partial u}{\partial n} = -\frac{\partial u}{\partial y},$$

and we set

$$[-D_{2y} u = 0]_{i,0}^n \Rightarrow \frac{u_{i,1}^n - u_{i,-1}^n}{2\Delta y} = 0.$$

From this it follows that $u_{i,-1}^n = u_{i,1}^n$. The discretized PDE at the boundary point $(i, 0)$ reads

$$\frac{u_{i,0}^{n+1} - 2u_{i,0}^n + u_{i,0}^{n-1}}{\Delta t^2} = c^2 \frac{u_{i+1,0}^n - 2u_{i,0}^n + u_{i-1,0}^n}{\Delta x^2} + c^2 \frac{u_{i,1}^n - 2u_{i,0}^n + u_{i,-1}^n}{\Delta y^2}$$

We can then just insert $u_{i,-1}^n$ for $u_{i,1}^n$ in this equation and then solve for the boundary value $u_{i,0}^{n+1}$ as done in 1D.

From these calculations, we see a pattern: the general scheme at the boundary $j = 0$ too if we just replace $j - 1$ by $j + 1$. Such a pattern is particularly useful for implementations. The details follow from the exact 1D case in Section 6.3.

The alternative approach to eliminating fictitious values outside the domain is to have $u_{i,-1}^n$ available as a ghost value. The mesh is extended with one line (2D) or plane (3D) of ghost cells at a Neumann boundary. In the example it means that we need a line of ghost cells below the y axis. The ghost values must be updated according to $u_{i,-1}^{n+1} = u_{i,1}^{n+1}$.

12 Implementation

We shall now describe in detail various Python implementations for the standard 2D, linear wave equation with constant wave velocity and $u = 0$ at the boundary. The wave equation is to be solved in the space-time domain Ω where $\Omega = (0, L_x) \times (0, L_y)$ is a rectangular spatial domain. More precisely, the complete initial-boundary value problem is defined by

$$\begin{aligned} u_t &= c^2 (u_{xx} + u_{yy}) + f(x, y, t), & (x, y) \in \Omega, \quad t \in (0, T], \\ u(x, y, 0) &= I(x, y), & (x, y) \in \Omega, \\ u_t(x, y, 0) &= V(x, y), & (x, y) \in \Omega, \\ u &= 0, & (x, y) \in \partial\Omega, \quad t \in (0, T], \end{aligned}$$

here $\partial\Omega$ is the boundary of Ω , in this case the four sides of the rectangle $[0, L_x] \times [0, L_y]$: $x = 0$, $x = L_x$, $y = 0$, and $y = L_y$.

The PDE is discretized as

$$[D_t D_t u = c^2(D_x D_x u + D_y D_y u) + f]_{i,j}^n,$$

which leads to an explicit updating formula to be implemented in a program:

$$u^{n+1} = -u_{i,j}^{n-1} + 2u_{i,j}^n + C_x^2(u_{i+1,j}^n - 2u_{i,j}^n + u_{i-1,j}^n) + C_y^2(u_{i,j+1}^n - 2u_{i,j}^n + u_{i,j-1}^n) + \Delta t^2 f_{i,j}^n, \quad (109)$$

for all interior mesh points $i \in \mathcal{I}_x^i$ and $j \in \mathcal{I}_y^j$, and for $n \in \mathcal{I}_t^+$. The constants C_x and C_y are defined as

$$C_x = c \frac{\Delta t}{\Delta x}, \quad C_y = c \frac{\Delta t}{\Delta y}.$$

At the boundary we simply set $u_{i,j}^{n+1} = 0$ for $i = 0, j = 0, \dots, N_y$; $i = N_x, j = 0, \dots, N_y$; $j = 0, i = 0, \dots, N_x$; and $j = N_y, i = 0, \dots, N_x$. For the first step, $n = 0$, (110) is combined with the discretization of the initial condition $u_t = V$, $[D_{2t} u = V]_{i,j}^0$ to obtain a special formula for $u_{i,j}^1$ at the interior mesh points:

$$u^1 = u_{i,j}^0 + \Delta t V_{i,j} + \frac{1}{2} C_x^2 (u_{i+1,j}^0 - 2u_{i,j}^0 + u_{i-1,j}^0) + \frac{1}{2} C_y^2 (u_{i,j+1}^0 - 2u_{i,j}^0 + u_{i,j-1}^0) + \frac{1}{2} \Delta t^2 f_{i,j}^0, \quad (110)$$

The algorithm is very similar to the one in 1D:

1. Set initial condition $u_{i,j}^0 = I(x_i, y_j)$
2. Compute $u_{i,j}^1$ from (110)
3. Set $u_{i,j}^1 = 0$ for the boundaries $i = 0, N_x, j = 0, N_y$
4. For $n = 1, 2, \dots, N_t$:
 - (a) Find $u_{i,j}^{n+1}$ from (110) for all internal mesh points, $i \in \mathcal{I}_x^i, j \in \mathcal{I}_y^j$
 - (b) Set $u_{i,j}^{n+1} = 0$ for the boundaries $i = 0, N_x, j = 0, N_y$

2.1 Scalar computations

The `solver` function for a 2D case with constant wave velocity and $u = 0$ as boundary condition follows the setup from the similar function for the 1D case in `wave1D_u0_s.py`, but there are a few necessary extensions. The code is in the program `wave2D_u0.py`¹⁵.

¹⁵http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0.py

Domain and mesh. The spatial domain is now $[0, L_x] \times [0, L_y]$, specify the arguments `Lx` and `Ly`. Similarly, the number of mesh points in the directions, N_x and N_y , become the arguments `Nx` and `Ny`. In multi-dimensional problems it makes less sense to specify a Courant number as the wave is a vector and the mesh spacings may differ in the various spatial directions. We therefore give Δt explicitly. The signature of the `solver` function is:

```
def solver(I, V, f, c, Lx, Ly, Nx, Ny, dt, T,
          user_action=None, version='scalar'):
```

Key parameters used in the calculations are created as

```
x = linspace(0, Lx, Nx+1)      # mesh points in x dir
y = linspace(0, Ly, Ny+1)      # mesh points in y dir
dx = x[1] - x[0]
dy = y[1] - y[0]
Nt = int(round(T/float(dt)))
t = linspace(0, N*dt, N+1)     # mesh points in time
Cx2 = (c*dt/dx)**2; Cy2 = (c*dt/dy)**2 # help variables
dt2 = dt**2
```

Solution arrays. We store $u_{i,j}^{n+1}$, $u_{i,j}^n$, and $u_{i,j}^{n-1}$ in three two-dimensional arrays,

```
u = zeros((Nx+1,Ny+1)) # solution array
u_1 = zeros((Nx+1,Ny+1)) # solution at t-dt
u_2 = zeros((Nx+1,Ny+1)) # solution at t-2*dt
```

where $u_{i,j}^{n+1}$ corresponds to `u[i,j]`, $u_{i,j}^n$ to `u_1[i,j]`, and $u_{i,j}^{n-1}$ to `u_2`

Index sets. It is also convenient to introduce the index sets (cf. Section 1.6):

```
Ix = range(0, u.shape[0])
Iy = range(0, u.shape[1])
It = range(0, t.shape[0])
```

Computing the solution. Inserting the initial condition `I` in `u_1` and a callback to the user in terms of the `user_action` function is a straightforward generalization of the 1D code from Section 1.6:

```
for i in Ix:
    for j in Iy:
        u_1[i,j] = I(x[i], y[j])

if user_action is not None:
    user_action(u_1, x, xv, y, yv, t, 0)
```

he `user_action` function has additional arguments compared to the 1D case. The arguments `xv` and `yv` fact will be commented upon in Section 12.2.

The key finite difference formula (102) for updating the solution at a time `vel` is implemented in a separate function as

```
def advance_scalar(u, u_1, u_2, f, x, y, t, n, Cx2, Cy2, dt,
                 V=None, step1=False):
    Ix = range(0, u.shape[0]); Iy = range(0, u.shape[1])
    dt2 = dt**2
    if step1:
        Cx2 = 0.5*Cx2; Cy2 = 0.5*Cy2; dt2 = 0.5*dt2
        D1 = 1; D2 = 0
    else:
        D1 = 2; D2 = 1
    for i in Ix[1:-1]:
        for j in Iy[1:-1]:
            u_xx = u_1[i-1,j] - 2*u_1[i,j] + u_1[i+1,j]
            u_yy = u_1[i,j-1] - 2*u_1[i,j] + u_1[i,j+1]
            u[i,j] = D1*u_1[i,j] - D2*u_2[i,j] + \
                Cx2*u_xx + Cy2*u_yy + dt2*f(x[i], y[j], t[n])
            if step1:
                u[i,j] += dt*V(x[i], y[j])
    # Boundary condition u=0
    j = Iy[0]
    for i in Ix: u[i,j] = 0
    j = Iy[-1]
    for i in Ix: u[i,j] = 0
    i = Ix[0]
    for j in Iy: u[i,j] = 0
    i = Ix[-1]
    for j in Iy: u[i,j] = 0
    return u
```

The `step1` variable has been introduced to allow the formula to be reused for first step $u_{i,j}^1$:

```
u = advance_scalar(u, u_1, u_2, f, x, y, t,
                 n, Cx2, Cy2, dt, V, step1=True)
```

Now, we will make many alternative implementations of the `advance_scalar` function to speed up the code since most of the CPU time in simulations is spent in this function.

2.2 Vectorized computations

The scalar code above turns out to be extremely slow for large 2D meshes, and probably useless in 3D beyond debugging of small test cases. Vectorization is therefore a must for multi-dimensional finite difference computations in Python. For example, with a mesh consisting of 30×30 cells, vectorization brings down the CPU time by a factor of 70 (!).

In the vectorized case we must be able to evaluate user-given functions like (x, y) and $f(x, y, t)$, provided as Python functions $I(x, y)$ and $f(x, y, t)$, for the entire mesh in one array operation. Having the one-dimensional coordinate arrays `x` and `y` is not sufficient: these must be extended to vectorized versions,

```
from numpy import newaxis
xv = x[:,newaxis]
yv = y[newaxis,:]
# or
xv = x.reshape((x.size, 1))
yv = y.reshape((1, y.size))
```

This is a standard required technique when evaluating functions over a 2D array. For example, `sin(xv)*cos(yv)`, which then gives a result with shape (N_x+1, N_y+1) .

With the `xv` and `yv` arrays for vectorized computing, setting the boundary condition is just a matter of

```
u_1[:, :] = I(xv, yv)
```

One could also have written `u_1 = I(xv, yv)` and let `u_1` point to a new array, but vectorized operations often makes use of direct insertion in the origin through `u_1[:, :]` because sometimes not all of the array is to be filled with a function evaluation. This is the case with the computational scheme for

```
def advance_vectorized(u, u_1, u_2, f_a, Cx2, Cy2, dt,
                     V=None, step1=False):
    dt2 = dt**2
    if step1:
        Cx2 = 0.5*Cx2; Cy2 = 0.5*Cy2; dt2 = 0.5*dt2
        D1 = 1; D2 = 0
    else:
        D1 = 2; D2 = 1
    u_xx = u_1[:-2,1:-1] - 2*u_1[1:-1,1:-1] + u_1[2:,1:-1]
    u_yy = u_1[1:-1,:-2] - 2*u_1[1:-1,1:-1] + u_1[1:-1,2:]
    u[1:-1,1:-1] = D1*u_1[1:-1,1:-1] - D2*u_2[1:-1,1:-1] + \
        Cx2*u_xx + Cy2*u_yy + dt2*f_a[1:-1,1:-1]
    if step1:
        u[1:-1,1:-1] += dt*V[1:-1, 1:-1]
    # Boundary condition u=0
    j = 0
    u[:,j] = 0
    j = u.shape[1]-1
    u[:,j] = 0
    i = 0
    u[i,:] = 0
    i = u.shape[0]-1
    u[i,:] = 0
    return u
```

Array slices in 2D are more complicated to understand than those in 1D, but the logic from 1D applies to each dimension separately. For example, doing $u_{i,j}^n - u_{i-1,j}^n$ for $i \in \mathcal{I}_x^+$, we just keep `j` constant and make a slice in the first index: `u_1[1:,j] - u_1[:-1,j]`, exactly as in 1D. The `1:` slice selects all the indices $i = 1, 2, \dots, N_x$ (up to the last valid index), while `:-1` specifies the relevant indices for the second term: $0, 1, \dots, N_x - 1$ (up to, but not including, the last index).

In the above code segment, the situation is slightly more complicated, as each displaced slice in one direction is accompanied by a `1:-1` slice in the

irection. The reason is that we only work with the internal points for the index `i` and `j` is kept constant in a difference.

The boundary conditions along the four sides makes use of a slice consisting of all indices along a boundary:

```
i[:,0] = 0
i[:,Ny] = 0
i[0,:] = 0
i[Nx,:] = 0
```

The `f` function is in the above vectorized update of `u` first computed as an array over all mesh points:

```
f_a = f(xv, yv, t[n])
```

One could, alternatively, use the call `f(xv, yv, t[n])[1:-1,1:-1]` in the last term of the update statement, but other implementations in compiled languages benefit from having `f` available in an array rather than calling our Python function `f(x,y,t)` for every point.

Also in the `advance_vectorized` function we have introduced a boolean `step1` to reuse the formula for the first time step in the same way as we did with `advance_scalar`. We refer to the `solver` function in `wave2D_u0.py` for more details on how the overall algorithm is implemented.

The callback function now has the arguments `u, x, xv, y, yv, t, n`. The inclusion of `xv` and `yv` makes it easy to, e.g., compute an exact 2D solution in the callback function and compute errors, through an expression like `- exact_solution(xv, yv, t[n])`.

2.3 Verification

Testing a quadratic solution. The 1D solution from Section 2.4 can be generalized to multi-dimensions and provides a test case where the exact solution also fulfills the discrete equations such that we know (to machine precision) what numbers the solver function should produce. In 2D we use the following generalization of (30):

$$u_e(x, y, t) = x(L_x - x)y(L_y - y)\left(1 + \frac{1}{2}t\right). \quad (111)$$

This solution fulfills the PDE problem if $I(x, y) = u_e(x, y, 0)$, $V = \frac{1}{2}u_e(x, y, 0)$, and $f = 2c^2(1 + \frac{1}{2}t)(y(L_y - y) + x(L_x - x))$. To show that u_e also solves the discrete equations, we start with the general results $[D_t D_t 1]^n = 0$, $[D_t D_t t]^n = 0$, and $[D_t D_t t^2] = 2$, and use these to compute

$$D_x D_x u_e^n_{i,j} = [y(L_y - y)\left(1 + \frac{1}{2}t\right)D_x D_x x(L_x - x)]^n_{i,j} = y_j(L_y - y_j)\left(1 + \frac{1}{2}t_n\right)2.$$

A similar calculation must be carried out for the $[D_y D_y u_e^n]_{i,j}$ and $[D_t D_t u_e^n]_{i,j}$ terms. One must also show that the quadratic solution fits the special formula

for $u_{i,j}^1$. The details are left as Exercise 10. The `test_quadratic` function in the `wave2D_u0.py`¹⁶ program implements this verification as a nose test.

13 Migrating loops to Cython

Although vectorization can bring down the CPU time dramatically compared to scalar code, there is still some factor 5-10 to win in these types of applications. Implementing the finite difference scheme in compiled code, typically in C, or C++. This can quite easily be done by adding a little extra code to the program. Cython is an extension of Python that offers the easiest way to migrate our Python loops in the scalar code down to machine code and the efficient C.

Cython can be viewed as an extended Python language where variables are declared with types and where functions are marked to be implemented in C. Migrating Python code to Cython is done by copying the desired code snippets to functions (or classes) and placing them in one or more separate files with the extension `.pyx`.

13.1 Declaring variables and annotating the code

Our starting point is the plain `advance_scalar` function for a scalar implementation of the updating algorithm for new values $u_{i,j}^{n+1}$:

```
def advance_scalar(u, u_1, u_2, f, x, y, t, n, Cx2, Cy2, dt,
                  V=None, step1=False):
    Ix = range(0, u.shape[0]); Iy = range(0, u.shape[1])
    dt2 = dt**2
    if step1:
        Cx2 = 0.5*Cx2; Cy2 = 0.5*Cy2; dt2 = 0.5*dt2
        D1 = 1; D2 = 0
    else:
        D1 = 2; D2 = 1
    for i in Ix[1:-1]:
        for j in Iy[1:-1]:
            u_xx = u_1[i-1,j] - 2*u_1[i,j] + u_1[i+1,j]
            u_yy = u_1[i,j-1] - 2*u_1[i,j] + u_1[i,j+1]
            u[i,j] = D1*u_1[i,j] - D2*u_2[i,j] + \
                    Cx2*u_xx + Cy2*u_yy + dt2*f(x[i], y[j], t[n])
            if step1:
                u[i,j] += dt*V(x[i], y[j])
    # Boundary condition u=0
    j = Iy[0]
    for i in Ix: u[i,j] = 0
    j = Iy[-1]
    for i in Ix: u[i,j] = 0
    i = Ix[0]
    for j in Iy: u[i,j] = 0
    i = Ix[-1]
    for j in Iy: u[i,j] = 0
    return u
```

¹⁶http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0.py

We simply take a copy of this function and put it in a file `wave2D_u0_loop_cy.pyx`. The relevant Cython implementation arises from declaring variables with types and adding some important annotations to speed up array computing in Cython. Let us first list the complete code in the `.pyx` file:

```
import numpy as np
cimport numpy as np
cimport cython
ctypedef np.float64_t DT # data type

@cython.boundscheck(False) # turn off array bounds check
@cython.wraparound(False) # turn off negative indices (u[-1,-1])
cpdef advance(
    np.ndarray[DT, ndim=2, mode='c'] u,
    np.ndarray[DT, ndim=2, mode='c'] u_1,
    np.ndarray[DT, ndim=2, mode='c'] u_2,
    np.ndarray[DT, ndim=2, mode='c'] f,
    double Cx2, double Cy2, double dt2):

    cdef int Ix_start = 0
    cdef int Iy_start = 0
    cdef int Ix_end = u.shape[0]-1
    cdef int Iy_end = u.shape[1]-1
    cdef int i, j
    cdef double u_xx, u_yy

    for i in range(Ix_start+1, Ix_end):
        for j in range(Iy_start+1, Iy_end):
            u_xx = u_1[i-1,j] - 2*u_1[i,j] + u_1[i+1,j]
            u_yy = u_1[i,j-1] - 2*u_1[i,j] + u_1[i,j+1]
            u[i,j] = 2*u_1[i,j] - u_2[i,j] + \
                Cx2*u_xx + Cy2*u_yy + dt2*f[i,j]

    # Boundary condition u=0
    j = Iy_start
    for i in range(Ix_start, Ix_end+1): u[i,j] = 0
    j = Iy_end
    for i in range(Ix_start, Ix_end+1): u[i,j] = 0
    i = Ix_start
    for j in range(Iy_start, Iy_end+1): u[i,j] = 0
    i = Iy_end
    for j in range(Iy_start, Iy_end+1): u[i,j] = 0
    return u
```

This example may act as a recipe on how to transform array-intensive code with loops into Cython.

1. Variables are declared with types: for example, `double v` in the argument list instead of just `v`, and `cdef double v` for a variable `v` in the body of the function. A Python `float` object is declared as `double` for translation to C by Cython, while an `int` object is declared by `int`.
2. Arrays need a comprehensive type declaration involving
 - the type `np.ndarray`,
 - the data type of the elements, here 64-bit floats, abbreviated as `DT` through `ctypedef np.float64_t DT` (instead of `DT` we could use the

full name of the data type: `np.float64_t`, which is a Cython type),

- the dimensions of the array, here `ndim=2` and `ndim=1`,
- specification of contiguous memory for the array (`mode='c'`)

3. Functions declared with `cpdef` are translated to C but also accessible in Python.
4. In addition to the standard `numpy` import we also need a special import of `numpy`: `cimport numpy as np`, to appear *after* the `numpy` import.
5. By default, array indices are checked to be within their legal limits. To speed up the code one should turn off this feature for a specific function by placing `@cython.boundscheck(False)` above the function header.
6. Also by default, array indices can be negative (counting from the end). This feature has a performance penalty and is therefore here turned off by writing `@cython.wraparound(False)` right above the function header.
7. The use of index sets `Ix` and `Iy` in the scalar code cannot be fully translated to C. One reason is that constructions like `Ix` involve negative indices, and these are now turned off. Another reason is that Cython loops must take the form `for i in xrange` or `for i in range` for being translated into efficient C loops. We have therefore introduced `Ix_start` as `Ix[0]` and `Ix_end` as `Ix[-1]` to mark the start and end of the values of index `i`. Similar variables are introduced for the `j` index. A loop `for i in Ix` is written with these new variables as `for i in range(Ix_start, Ix_end+1)`.

Array declaration syntax in Cython.

We have used the syntax `np.ndarray[DT, ndim=2, mode='c']` to declare `numpy` arrays in Cython. There is a simpler, alternative syntax, employing `typed memory views`^a, where the declaration looks like `double [I]`. However, the full support for this functionality is not yet ready, and in this text we use the full array declaration syntax.

^a<http://docs.cython.org/src/userguide/memoryviews.html>

13.2 Visual inspection of the C translation

Cython can visually explain how successfully it can translate a code from Python to C. The command

```
Terminal> cython -a wave2D_u0_loop_cy.pyx
```

reduces an HTML file `wave2D_u0_loop_cy.html`, which can be loaded into a web browser to illustrate which lines of the code that have been translated to C. Figure 8 shows the illustrated code. Yellow lines indicate the lines that Cython did not manage to translate to efficient C code and that remain in Python. For the present code we see that Cython is able to translate all the loops with array computing to C, which is our primary goal.

```
Raw output: wave2D_u0_loop_cy.c
1: import numpy as np
2: cimport numpy as np
3: cimport cython
4: ctypedef np.float64_t DT # data type
5:
6: @cython.boundscheck(False) # turn off array bounds check
7: @cython.wraparound(False) # turn off negative indices (u[-1,-1])
8: cdef advance
9:
10: np.ndarray[DT, ndim=2, mode='c'] u
11: np.ndarray[DT, ndim=2, mode='c'] u_1
12: np.ndarray[DT, ndim=2, mode='c'] u_2
13: double Cx2, double Cy2, double dt2
14:
15: cdef int ix_start = 0
16: cdef int iy_start = 0
17: cdef int ix_end = u.shape[0]-1
18: cdef int iy_end = u.shape[1]-1
19: cdef int i, j
20: cdef double u_xx, u_yy
21:
22: for i in range(ix_start+1, ix_end):
23:     for j in range(iy_start+1, iy_end):
24:         u_xx = u_1[i-1, j] - 2*u_1[i, j] + u_1[i+1, j]
25:         u_yy = u_1[i, j-1] - 2*u_1[i, j] + u_1[i, j+1]
26:         u[i, j] = 2*u_1[i, j] - u_2[i, j] + \
27:             Cx2*u_xx + Cy2*u_yy + dt2*f(i, j)
28: # Boundary condition u=0
29: j = iy_start
30: for i in range(ix_start, ix_end+1): u[i, j] = 0
31: j = iy_end
32: for i in range(ix_start, ix_end+1): u[i, j] = 0
33: i = ix_start
34: for j in range(iy_start, iy_end+1): u[i, j] = 0
35: i = iy_end
36: for j in range(iy_start, iy_end+1): u[i, j] = 0
37: return u
```

Figure 8: Visual illustration of Cython’s ability to translate Python to C.

You can also inspect the generated C code directly, as it appears in the file `wave2D_u0_loop_cy.c`. Nevertheless, understanding this C code requires some familiarity with writing Python extension modules in C by hand. Deep down in the file we can see in detail how the compute-intensive statements are translated to some complex C code that is quite different from what we a human would write at least if a direct correspondence to the mathematics was in mind).

3.3 Building the extension module

Cython code must be translated to C, compiled, and linked to form what is known to the Python world as a *C extension module*. This is usually done by making a `setup.py` script, which is the standard way of building and installing Python software. For an extension module arising from Cython code, the following `setup.py` script is all we need to build and install the module:

```
from distutils.core import setup
from distutils.extension import Extension
from Cython.Distutils import build_ext

cymodule = 'wave2D_u0_loop_cy'
setup(
    name=cymodule,
    ext_modules=[Extension(cymodule, [cymodule + '.pyx'],)],
    cmdclass={'build_ext': build_ext},
)
```

We run the script by

```
Terminal> python setup.py build_ext --inplace
```

The `--inplace` option makes the extension module available in the directory as the file `wave2D_u0_loop_cy.so`. This file acts as a normal module that can be imported and inspected:

```
>>> import wave2D_u0_loop_cy
>>> dir(wave2D_u0_loop_cy)
['__builtins__', '__doc__', '__file__', '__name__',
 '__package__', '__test__', 'advance', 'np']
```

The important output from the `dir` function is our Cython function `advance` (the module also features the imported `numpy` module under the name `np` as many standard Python objects with double underscores in their name).

The `setup.py` file makes use of the `distutils` package in Python and Cython’s extension of this package. These tools know how Python was installed on the computer and will use compatible compiler(s) and options when linking other code in Cython, C, or C++. Quite some experience with building program systems is needed to do the build process manually, so using a script is strongly recommended.

Simplified build of a Cython module.

When there is no need to link the C code with special libraries, Cython offers a shortcut for generating and importing the extension module:

```
import pyximport; pyximport.install()
```

This makes the `setup.py` script redundant. However, in the `wave2D_u0_loop_cy` code we do not use `pyximport` and require an explicit build process and many other modules.

13.4 Calling the Cython function from Python

The `wave2D_u0_loop_cy` module contains our `advance` function, which we can call from the Python program for the wave equation:

```
import wave2D_u0_loop_cy
advance = wave2D_u0_loop_cy.advance
...
for n in It[1:-1]:
    # time loop
    f_a[:, :] = f(xv, yv, t[n]) # precompute, size as u
    u = advance(u, u_1, u_2, f_a, x, y, t, Cx2, Cy2, dt2)
```

Efficiency. For a mesh consisting of 120×120 cells, the scalar Python code require 1370 CPU time units, the vectorized version requires 5.5, while the Cython version requires only 1! For a smaller mesh with 60×60 cells Cython is about 1000 times faster than the scalar Python code, and the vectorized version about 6 times slower than the Cython version.

4 Migrating loops to Fortran

Instead of relying on Cython's (excellent) ability to translate Python to C, we can instead invoke a compiled language directly and write the loops ourselves. Let us start with Fortran 77, because this is a language with more convenient array handling than C (or plain C++). Or more precisely, we can write a program with the same multi-dimensional indices in the Fortran code as in the NumPy arrays in the Python code, while in C these arrays are one-dimensional and require us to reduce multi-dimensional indices to a single index.

4.1 The Fortran subroutine

We write a Fortran subroutine `advance` in a file `wave2D_u0_loop_f77.f`¹⁷ for implementing the updating formula (110) and setting the solution to zero at the boundaries:

```

subroutine advance(u, u_1, u_2, f, Cx2, Cy2, dt2, Nx, Ny)
  integer Nx, Ny
  real*8 u(0:Nx,0:Ny), u_1(0:Nx,0:Ny), u_2(0:Nx,0:Ny)
  real*8 f(0:Nx, 0:Ny), Cx2, Cy2, dt2
  integer i, j
  Cf2py intent(in, out) u

  C Scheme at interior points
  do j = 1, Ny-1
    do i = 1, Nx-1
      u(i,j) = 2*u_1(i,j) - u_2(i,j) +
&      Cx2*(u_1(i-1,j) - 2*u_1(i,j) + u_1(i+1,j)) +
&      Cy2*(u_1(i,j-1) - 2*u_1(i,j) + u_1(i,j+1)) +
&      dt2*f(i,j)
    end do
  end do

  C Boundary conditions
  j = 0
  do i = 0, Nx
    u(i,j) = 0
  end do
  j = Ny
  do i = 0, Nx
    u(i,j) = 0
  end do
  i = 0
  do j = 0, Ny
    u(i,j) = 0
  end do

```

¹⁷http://tinyurl.com/jvzcf/wave/wave2D_u0/wave2D_u0_loop_f77.f

```

end do
i = Nx
do j = 0, Ny
  u(i,j) = 0
end do
return
end

```

This code is plain Fortran 77, except for the special `Cf2py` comment line here specifies that `u` is both an input argument *and* an object to be returned from the `advance` routine. Or more precisely, Fortran is not able to return a value from a function, but we need a *wrapper code* in C for the Fortran subroutine to be able to call it from Python, and in this wrapper code one can return a value by calling Python code.

Remark.

It is not strictly necessary to return `u` to the calling Python code because the `advance` function will modify the elements of `u`, but the conventional way in Python is to get all output from a function as returned values. That is the right way of calling the above Fortran subroutine from Python is

```
u = advance(u, u_1, u_2, f, Cx2, Cy2, dt2)
```

The less encouraged style, which works and resembles the way the Fortran subroutine is called from Fortran, reads

```
advance(u, u_1, u_2, f, Cx2, Cy2, dt2)
```

14.2 Building the Fortran module with `f2py`

The nice feature of writing loops in Fortran is that the tool `f2py` can with a little work produce a C extension module such that we can call the Fortran version of `advance` from Python. The necessary commands to run are

```

Terminal> f2py -m wave2D_u0_loop_f77 -h wave2D_u0_loop_f77.pyf \
--overwrite-signature wave2D_u0_loop_f77.f
Terminal> f2py -c wave2D_u0_loop_f77.pyf --build-dir build_f77 \
-DF2PY_REPORT_ON_ARRAY_COPY=1 wave2D_u0_loop_f77.f

```

The first command asks `f2py` to interpret the Fortran code and make a Fortran specification of the extension module in the file `wave2D_u0_loop_f77.pyf`. The second command makes `f2py` generate all necessary wrapper code, compile the Fortran file and the wrapper code, and finally build the module. The build takes place in the specified subdirectory `build_f77` so that files can be in place if something goes wrong. The option `-DF2PY_REPORT_ON_ARRAY_COPY=`

f2py write a message for every array that is copied in the communication between Fortran and Python, which is very useful for avoiding unnecessary array copying (see below). The name of the module file is `wave2D_u0_loop_f77.so`, and this file can be imported and inspected as any other Python module:

```
>>> import wave2D_u0_loop_f77
>>> dir(wave2D_u0_loop_f77)
['__doc__', '__file__', '__name__', '__package__',
 '__version__', 'advance']
>>> print wave2D_u0_loop_f77.__doc__
This module 'wave2D_u0_loop_f77' is auto-generated with f2py....
Functions:
  u = advance(u,u_1,u_2,f,cx2,cy2,dt2,
             nx=(shape(u,0)-1),ny=(shape(u,1)-1))
```

Examine the doc strings!

Printing the doc strings of the module and its functions is extremely important after having created a module with `f2py`, because `f2py` makes Python interfaces to the Fortran functions that are different from how the functions are declared in the Fortran code (!). The rationale for this behavior is that `f2py` creates *Pythonic* interfaces such that Fortran routines can be called in the same way as one calls Python functions. Output data from Python functions is always returned to the calling code, but this is technically impossible in Fortran. Also, arrays in Python are passed to Python functions without their dimensions because that information is packed with the array data in the array objects, but this is not possible in Fortran. Therefore, `f2py` removes array dimensions from the argument list, and `f2py` makes it possible to return objects back to Python.

Let us follow the advice of examining the doc strings and take a close look at the documentation `f2py` has generated for our Fortran `advance` subroutine:

```
>>> print wave2D_u0_loop_f77.advance.__doc__
This module 'wave2D_u0_loop_f77' is auto-generated with f2py
Functions:
  u = advance(u,u_1,u_2,f,cx2,cy2,dt2,
             nx=(shape(u,0)-1),ny=(shape(u,1)-1))

advance - Function signature:
  u = advance(u,u_1,u_2,f,cx2,cy2,dt2,[nx,ny])
Required arguments:
  u : input rank-2 array('d') with bounds (nx + 1,ny + 1)
  u_1 : input rank-2 array('d') with bounds (nx + 1,ny + 1)
  u_2 : input rank-2 array('d') with bounds (nx + 1,ny + 1)
  f : input rank-2 array('d') with bounds (nx + 1,ny + 1)
  cx2 : input float
  cy2 : input float
  dt2 : input float
Optional arguments:
  nx := (shape(u,0)-1) input int
  ny := (shape(u,1)-1) input int
Return objects:
  u : rank-2 array('d') with bounds (nx + 1,ny + 1)
```

Here we see that the `nx` and `ny` parameters declared in Fortran are `C` arguments that can be omitted when calling `advance` from Python.

We strongly recommend to print out the documentation of *every* function to be called from Python and make sure the call syntax is explicitly listed in the documentation.

14.3 How to avoid array copying

Multi-dimensional arrays are stored as a stream of numbers in memory. For a two-dimensional array consisting of rows and columns there are two ways of creating such a stream: *row-major ordering*, which means that rows are stored consecutively in memory, or *column-major ordering*, which means that columns are stored one after each other. All programming languages inherit from C, including Python, apply the row-major ordering, but Fortran uses column-major storage. Thinking of a two-dimensional array in Python as a matrix, it means that Fortran works with the transposed matrix.

Fortunately, `f2py` creates extra code so that accessing `u(i,j)` in the subroutine corresponds to the element `u[i,j]` in the underlying `numpy` array (without the extra code, `u(i,j)` in Fortran would access `u[j,i]` in the `numpy` array). Technically, `f2py` takes a copy of our `numpy` array and reorders it before sending the array to Fortran. Such copying can be costly. For 2D simulations on a 60×60 grid the overhead of copying is a factor of 5 means that almost the whole performance gain of Fortran over vectorized code is lost!

To avoid having `f2py` to copy arrays with C storage to the corresponding Fortran storage, we declare the arrays with Fortran storage:

```
order = 'Fortran' if version == 'f77' else 'C'
u = zeros((Nx+1,Ny+1), order=order) # solution array
u_1 = zeros((Nx+1,Ny+1), order=order) # solution at t-dt
u_2 = zeros((Nx+1,Ny+1), order=order) # solution at t-2*dt
```

In the compile and build step of using `f2py`, it is recommended to use the extra option for making `f2py` report on array copying:

```
Terminal> f2py -c wave2D_u0_loop_f77.pyf --build-dir build_f77 \
           -DF2PY_REPORT_ON_ARRAY_COPY=1 wave2D_u0_loop_f77.f
```

It can sometimes be a challenge to track down which array that causes copying. There are two principal reasons for copying array data: either the array does not have Fortran storage or the element types do not match those declared in the Fortran code. The latter cause is usually effectively eliminated by using `real*8` data in the Fortran code and `float64` (the default `float` type in Python) in the arrays on the Python side. The former reason is more common, so it is recommended to check whether an array before a Fortran call has the right storage one can use by checking the result of `isfortran(a)`, which is `True` if the array `a` has Fortran storage.

Let us look at an example where we face problems with array storage. A typical problem in the `wave2D_u0.py` code is to set

```
f_a = f(xv, yv, t[n])
```

before the call to the Fortran `advance` routine. This computation creates a new array with C storage. An undesired copy of `f_a` will be produced when sending `f_a` to a Fortran routine. There are two remedies, either direct insertion of data into an array with Fortran storage,

```
f_a = zeros((Nx+1, Ny+1), order='Fortran')
...
f_a[:,:] = f(xv, yv, t[n])
```

or remaking the `f(xv, yv, t[n])` array,

```
f_a = asarray(f(xv, yv, t[n]), order='Fortran')
```

The former remedy is most efficient if the `asarray` operation is to be performed a large number of times.

Efficiency. The efficiency of this Fortran code is very similar to the Cython code. There is usually nothing more to gain, from a computational efficiency point of view, by implementing the *complete* Python program in Fortran or C. That will just be a lot more code for all administering work that is needed in scientific software, especially if we extend our sample program `wave2D_u0.py` to handle a real scientific problem. Then only a small portion will consist of loops with intensive array calculations. These can be migrated to Cython or Fortran as explained, while the rest of the programming can be more conveniently done in Python.

5 Migrating loops to C via Cython

The computationally intensive loops can alternatively be implemented in C code. Just as Fortran calls for care regarding the storage of two-dimensional arrays, working with two-dimensional arrays in C is a bit tricky. The reason is that `numpy` arrays are viewed as one-dimensional arrays when transferred to C, while C programmers will think of `u`, `u_1`, and `u_2` as two dimensional arrays and index them like `u[i][j]`. The C code must declare `u` as `double* u` and translate an index pair `[i][j]` to a corresponding single index when `u` is viewed as one-dimensional. This translation requires knowledge of how the numbers in `u` are stored in memory.

5.1 Translating index pairs to single indices

Two-dimensional `numpy` arrays with the default C storage are stored row by row. In general, multi-dimensional arrays with C storage are stored such that the last

index has the fastest variation, then the next last index, and so on, ending with the slowest variation in the first index. For a two-dimensional `u` created as `zeros((Nx+1, Ny+1))` in Python, the individual elements are stored in the following order:

```
u[0,0], u[0,1], u[0,2], ..., u[0,Ny], u[1,0], u[1,1], ...,
u[1,Ny], u[2,0], ..., u[Nx,0], u[Nx,1], ..., u[Nx, Ny]
```

Viewing `u` as one-dimensional, the index pair (i, j) translates to $i(N_y + j)$. So, where a C programmer would naturally write an index `u[i][j]`, the C programmer must read `u[i*(Ny+1) + j]`. This is tedious to write, so it can be helped by defining a C macro,

```
#define idx(i,j) (i)*(Ny+1) + j
```

so that we can write `u[idx(i,j)]`, which reads much better and is easier to debug.

Be careful with macro definitions.

Macros just perform simple text substitutions: `idx(hello,world)` expands to `(hello)*(Ny+1) + world`. The parenthesis in `(i)` are essential - using the natural mathematical formula $i(N_y + j)$ in the macro definition, `idx(i-1,j)` would expand to `i-1*(Ny+1) + j`, which is the wrong formula. Macros are handy, but requires careful use. In C++, lambda functions are safer and replace the need for macros.

15.2 The complete C code

The C version of our function `advance` can be coded as follows.

```
#define idx(i,j) (i)*(Ny+1) + j

void advance(double* u, double* u_1, double* u_2, double* f,
             double Cx2, double Cy2, double dt2,
             int Nx, int Ny)
{
    int i, j;
    /* Scheme at interior points */
    for (i=1; i<=Nx-1; i++) {
        for (j=1; j<=Ny-1; j++) {
            u[idx(i,j)] = 2*u_1[idx(i,j)] - u_2[idx(i,j)] +
                Cx2*(u_1[idx(i-1,j)] - 2*u_1[idx(i,j)] + u_1[idx(i+1,j)]) +
                Cy2*(u_1[idx(i,j-1)] - 2*u_1[idx(i,j)] + u_1[idx(i,j+1)]) +
                dt2*f[idx(i,j)];
        }
    }
    /* Boundary conditions */
    j = 0; for (i=0; i<=Nx; i++) u[idx(i,j)] = 0;
    j = Ny; for (i=0; i<=Nx; i++) u[idx(i,j)] = 0;
    i = 0; for (j=0; j<=Ny; j++) u[idx(i,j)] = 0;
    i = Nx; for (j=0; j<=Ny; j++) u[idx(i,j)] = 0;
}
```

5.3 The Cython interface file

If the code above appears in a file `wave2D_u0_loop_c.c`¹⁸. We need to compile this file together with C wrapper code such that `advance` can be called from Python. Cython can be used to generate appropriate wrapper code. The relevant Python code for interfacing C is placed in a file with extension `.pyx`. Here this file, called `wave2D_u0_loop_c_cy.pyx`¹⁹, looks like

```
import numpy as np
cimport numpy as np
cimport cython

cdef extern from "wave2D_u0_loop_c.h":
    void advance(double* u, double* u_1, double* u_2, double* f,
                double Cx2, double Cy2, double dt2,
                int Nx, int Ny)

@cython.boundscheck(False)
@cython.wraparound(False)
def advance_cwrap(
    np.ndarray[double, ndim=2, mode='c'] u,
    np.ndarray[double, ndim=2, mode='c'] u_1,
    np.ndarray[double, ndim=2, mode='c'] u_2,
    np.ndarray[double, ndim=2, mode='c'] f,
    double Cx2, double Cy2, double dt2):
    advance(&u[0,0], &u_1[0,0], &u_2[0,0], &f[0,0],
           Cx2, Cy2, dt2,
           u.shape[0]-1, u.shape[1]-1)
    return u
```

We first declare the C functions to be interfaced. These must also appear in a C header file, `wave2D_u0_loop_c.h`²⁰,

```
extern void advance(double* u, double* u_1, double* u_2, double* f,
                  double Cx2, double Cy2, double dt2,
                  int Nx, int Ny);
```

The next step is to write a Cython function with Python objects as arguments. The name `advance` is already used for the C function so the function to be called from Python is named `advance_cwrap`. The contents of this function is simply a call to the `advance` version in C. To this end, the right information from the Python objects must be passed on as arguments to `advance`. Arrays are sent with their C pointers to the first element, obtained in Cython as `&u[0,0]` (the `&` takes the address of a C variable). The `Nx` and `Ny` arguments in `advance` are easily obtained from the shape of the `numpy` array `u`. Finally, `u` must be returned such that we can set `u = advance(...)` in Python.

5.4 Building the extension module

It remains to build the extension module. An appropriate `setup.py` file is

¹⁸http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0_loop_c.c

¹⁹http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0_loop_c_cy.pyx

²⁰http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0_loop_c.h

```
from distutils.core import setup
from distutils.extension import Extension
from Cython.Distutils import build_ext

sources = ['wave2D_u0_loop_c.c', 'wave2D_u0_loop_c_cy.pyx']
module = 'wave2D_u0_loop_c_cy'
setup(
    name=module,
    ext_modules=[Extension(module, sources,
                           libraries=[], # C libs to link with
                           )],
    cmdclass={'build_ext': build_ext},
)
```

All we need to specify is the `.c` file(s) and the `.pyx` interface file. Cython automatically runs to generate the necessary wrapper code. Files are then compiled and linked to an extension module residing in the file `wave2D_u0_loop_c_cy.pyd`. Here is a session with running `setup.py` and examining the resulting module in Python

```
Terminal> python setup.py build_ext --inplace
Terminal> python
>>> import wave2D_u0_loop_c_cy as m
>>> dir(m)
['__builtins__', '__doc__', '__file__', '__name__', '__package__',
 '__test__', 'advance_cwrap', 'np']
```

The call to the C version of `advance` can go like this in Python:

```
import wave2D_u0_loop_c_cy
advance = wave2D_u0_loop_c_cy.advance_cwrap
...
f_a[:,:] = f(xv, yv, t[n])
u = advance(u, u_1, u_2, f_a, Cx2, Cy2, dt2)
```

Efficiency. In this example, the C and Fortran code runs at the same time and there are no significant differences in the efficiency of the wrapper code. The overhead implied by the wrapper code is negligible as long as we do not use very small meshes and consequently little numerical work in the `advance` function.

16 Migrating loops to C via f2py

An alternative to using Cython for interfacing C code is to apply `f2py`. The `f2py` code is the same, just the details of specifying how it is to be called from Python differ. The `f2py` tool requires the call specification to be a Fortran 90 subroutine defined in a `.pyf` file. This file was automatically generated when we used `f2py` on a Fortran subroutine. With a C function we need to write this module our-

we can use a trick and let `f2py` generate it for us. The trick consists in writing the signature of the C function with Fortran syntax and place it in a Fortran file, here `wave2D_u0_loop_c_f2py_signature.f`:

```

subroutine advance(u, u_1, u_2, f, Cx2, Cy2, dt2, Nx, Ny)
Cf2py intent(c) advance
integer Nx, Ny, N
real*8 u(0:Nx,0:Ny), u_1(0:Nx,0:Ny), u_2(0:Nx,0:Ny)
real*8 f(0:Nx, 0:Ny), Cx2, Cy2, dt2
Cf2py intent(in, out) u
Cf2py intent(c) u, u_1, u_2, f, Cx2, Cy2, dt2, Nx, Ny
return
end

```

Note that we need a special `f2py` instruction, through a `Cf2py` comment line, for telling that all the function arguments are C variables. We also need to specify that the function is actually in C: `intent(c) advance`.

Since `f2py` is just concerned with the function signature and not the complete contents of the function body, it can easily generate the Fortran 90 module specification based solely on the signature above:

```

terminal> f2py -m wave2D_u0_loop_c_f2py \
-h wave2D_u0_loop_c_f2py.pyf --overwrite-signature \
wave2D_u0_loop_c_f2py_signature.f

```

The compile and build step is as for the Fortran code, except that we list C files instead of Fortran files:

```

terminal> f2py -c wave2D_u0_loop_c_f2py.pyf \
--build-dir tmp_build_c \
-DF2PY_REPORT_ON_ARRAY_COPY=1 wave2D_u0_loop_c.c

```

As when interfacing Fortran code with `f2py`, we need to print out the doc string to see the exact call syntax from the Python side. This doc string is identical to the C and Fortran versions of `advance`.

6.1 Migrating loops to C++ via `f2py`

C++ is a much more versatile language than C or Fortran and has over the last two decades become very popular for numerical computing. Many will therefore prefer to migrate compute-intensive Python code to C++. This is, in principle, easy: just write the desired C++ code and use some tool for interfacing it from Python. A tool like `SWIG`²¹ can interpret the C++ code and generate interfaces for a wide range of languages, including Python, Perl, Ruby, and Java. However, `SWIG` is a comprehensive tool with a correspondingly steep learning

²¹<http://swig.org/>

curve. Alternative tools, such as `Boost Python`²², `SIP`²³, and `Shiboken` are similarly comprehensive. Simpler tools include `PyBindGen`²⁵,

A technically much easier way of interfacing C++ code is to use the possibility to use C++ classes directly from Python, but instead of directly interface to the C++ code. The C interface can be handled by `f2py` as in the example with pure C code. Such a solution means that classes in Python and C++ cannot be mixed and that only primitive data types like numbers, strings, and arrays can be transferred between Python and C++. Another solution is often a very good solution because it forces the C++ code to work with primitive data, which usually gives faster code than if fancy data structures with pointers are used. The arrays coming from Python, and looking like plain C/C++ arrays, can be efficiently wrapped in more user-friendly C++ array classes in the C++ code, if desired.

17 Using classes to implement a simulator

- Introduce classes `Mesh`, `Function`, `Problem`, `Solver`, `Visualizer`

18 Exercises

Exercise 10: Check that a solution fulfills the discrete model

Carry out all mathematical details to show that (111) is indeed a solution of the discrete model for a 2D wave equation with $u = 0$ on the boundary. You must check the boundary conditions, the initial conditions, the general wave equation at a time level and the special version of this equation for the fixed time level. Filename: `check_quadratic_solution.pdf`.

Project 11: Calculus with 2D/3D mesh functions

The goal of this project is to redo Project 5 with 2D and 3D mesh functions ($f_{i,j}$ and $f_{i,j,k}$).

Differentiation. The differentiation results in a discrete gradient field which in the 2D case can be represented by a three-dimensional array `df[d,i,j]` where `d` represents the direction of the derivative and `i` and `j` are mesh indices or counters in 2D (the 3D counterpart is `df[d,i,j,k]`).

²²http://www.boost.org/doc/libs/1_51_0/libs/python/doc/index.html

²³<http://riverbankcomputing.co.uk/software/sip/intro>

²⁴<http://qt-project.org/wiki/Category:LanguageBindings::PySide::Shiboken>

²⁵<http://code.google.com/p/pybindgen/>

Integration. The integral of a 2D mesh function $f_{i,j}$ is defined as

$$F_{i,j} = \int_{y_0}^{y_j} \int_{x_0}^{x_i} f(x,y) dx dy,$$

here $f(x,y)$ is a function that takes on the values of the discrete mesh function $f_{i,j}$ at the mesh points, but can also be evaluated in between the mesh points. The particular variation between mesh points can be taken as bilinear, but this is not important as we will use a product Trapezoidal rule to approximate the integral over a cell in the mesh and then we only need to evaluate $f(x,y)$ at the mesh points.

Suppose $F_{i,j}$ is computed. The calculation of $F_{i+1,j}$ is then

$$\begin{aligned} F_{i+1,j} &= F_{i,j} + \int_{x_i}^{x_{i+1}} \int_{y_0}^{y_j} f(x,y) dy dx \\ &\approx \Delta x \int_{y_0}^{y_j} f(x_{i+\frac{1}{2}}, y) dy \\ &\approx \Delta x \frac{1}{2} \left(\int_{y_0}^{y_j} f(x_i, y) dy + \int_{y_0}^{y_j} f(x_{i+1}, y) dy \right) \end{aligned}$$

The integrals in the y direction can be approximated by a Trapezoidal rule. A similar idea can be used to compute $F_{i,j+1}$. Thereafter, $F_{i+1,j+1}$ can be computed by adding the integral over the final corner cell to $F_{i+1,j} + F_{i,j+1} - F_{i,j}$. Carry out the details of these computations and extend the ideas to 3D. Filename: `esh_calculus_3D.py`.

Exercise 12: Implement Neumann conditions in 2D

Modify the `wave2D_u0.py`²⁶ program, which solves the 2D wave equation $u_{tt} = c^2(u_{xx} + u_{yy})$ with constant wave velocity c and $u = 0$ on the boundary, to have Neumann boundary conditions: $\partial u / \partial n = 0$. Include both scalar code (for debugging and reference) and vectorized code (for speed).

To test the code, use $u = 1.2$ as solution ($I(x,y) = 1.2$, $V = f = 0$, and c arbitrary), which should be exactly reproduced with any mesh as long as the stability criterion is satisfied. Another test is to use the plug-shaped pulse in the `ulse` function from Section 8 and the `wave1D_dn_vc.py`²⁷ program. This pulse is exactly propagated in 1D if $c\Delta t / \Delta x = 1$. Check that also the 2D program can propagate this pulse exactly in x direction ($c\Delta t / \Delta x = 1$, Δy arbitrary) and in y direction ($c\Delta t / \Delta y = 1$, Δx arbitrary). Filename: `wave2D_dn.py`.

Exercise 13: Test the efficiency of compiled loops in 3D

Extend the `wave2D_u0.py` code and the Cython, Fortran, and C versions to 3D. Set up an efficiency experiment to determine the relative efficiency of pure scalar

²⁶http://tinyurl.com/jvzzcfn/wave/wave2D_u0/wave2D_u0.py

²⁷http://tinyurl.com/jvzzcfn/wave/wave1D_dn_vc.py

Python code, vectorized code, Cython-compiled loops, Fortran-compiled and C-compiled loops. Normalize the CPU time for each mesh by the version. Filename: `wave3D_u0.py`.

19 Applications of wave equations

This section presents a range of wave equation models for different physical phenomena. Although many wave motion problems in physics can be modeled by the standard linear wave equation, or a similar formulation with a system of second-order equations, there are some exceptions. Perhaps the most important examples are waves: these are modeled by the Laplace equation with time-dependent boundary conditions at the water surface (long water waves, however, can be approximated by a standard wave equation, see Section 19.7). Quantum mechanics provides another example where the waves are governed by the Schrödinger equation and not a standard wave equation. Many wave phenomena also take nonlinear effects into account when the wave amplitude is significant. Shock waves in the air is a primary example.

The derivations in the following are very brief. Those with a firm background in continuum mechanics will probably have enough information to fill in the details, while other readers will hopefully get some impression of the physical approximations involved when establishing wave equation models.

19.1 Waves on a string

Figure 9 shows a model we may use to derive the equation for waves on a string. The string is modeled as a set of discrete point masses (at mesh points) connected by elastic strings in between. The strings are at a high constant tension T . Let the mass at mesh point x_i be m_i . The displacement of this mass in the y direction is denoted by $u_i(t)$.

The motion of mass m_i is governed by Newton's second law of motion. The position of the mass at time t is $x_i \mathbf{i} + u_i(t) \mathbf{j}$, where \mathbf{i} and \mathbf{j} are unit vectors in the x and y direction, respectively. The acceleration is then $u_i''(t) \mathbf{j}$. Two forces are acting on the mass as indicated in Figure 9. The force \mathbf{T}^- acting on point x_{i-1} can be decomposed as

$$\mathbf{T}^- = -T \sin \phi \mathbf{i} - T \cos \phi \mathbf{j},$$

where ϕ is the angle between the force and the line $x = x_i$. Let $\Delta u_i = u_i - u_{i-1}$ and let $\Delta s_i = \sqrt{\Delta u_i^2 + (x_i - x_{i-1})^2}$ be the distance from mass m_{i-1} to mass m_i . It is seen that $\cos \phi = \Delta u_i / \Delta s_i$ and $\sin \phi = (x_i - x_{i-1}) / \Delta s_i$. To simplify we introduce a constant mesh spacing $\Delta x = x_i - x_{i-1}$. The force can be written

$$\mathbf{T}^- = -T \frac{\Delta x}{\Delta s_i} \mathbf{i} - T \frac{\Delta u_i}{\Delta s_i} \mathbf{j}.$$

The force \mathbf{T}^+ acting toward x_{i+1} can be calculated in a similar way:




figure 9: Discrete string model with point masses connected by elastic strings.

$$\mathbf{T}^+ = T \frac{\Delta x}{\Delta s_{i+1}} \mathbf{i} + T \frac{\Delta u_{i+1}}{\Delta s_{i+1}} \mathbf{j}.$$

Newton's second law becomes

$$m_i u_i''(t) \mathbf{j} = \mathbf{T}^+ + \mathbf{T}^-,$$

which gives the component equations

$$T \frac{\Delta x}{\Delta s_i} = T \frac{\Delta x}{\Delta s_{i+1}}, \quad (112)$$

$$m_i u_i''(t) = T \frac{\Delta u_{i+1}}{\Delta s_{i+1}} - T \frac{\Delta u_i}{\Delta s_i}. \quad (113)$$

A basic reasonable assumption for a string is small displacements and small displacement gradients $\Delta u_i / \Delta x$. For small $g = \Delta u_i / \Delta x$ we have

$$\Delta s_i = \sqrt{\Delta u_i^2 + \Delta x^2} = \Delta x \sqrt{1 + g^2} + \Delta x \left(1 + \frac{1}{2}g^2 + \mathcal{O}(g^4)\right) \approx \Delta x$$

Equation (112) is then simply the identity $T = T$, while (113) can be written as

$$m_i u_i''(t) = T \frac{\Delta u_{i+1}}{\Delta x} - T \frac{\Delta u_i}{\Delta x},$$

which upon division by Δx and introducing the density $\rho_i = m_i / \Delta x$ becomes

$$\rho_i u_i''(t) = T \frac{1}{\Delta x^2} (u_{i+1} - 2u_i + u_{i-1}).$$

We can now choose to approximate u_i'' by a finite difference in time and discretized wave equation,

$$\rho_i \frac{1}{\Delta t^2} (u_i^{n+1} - 2u_i^n - u_i^{n-1}) = T \frac{1}{\Delta x^2} (u_{i+1} - 2u_i + u_{i-1}).$$

On the other hand, we may go to the continuum limit $\Delta x \rightarrow 0$ and replace $u(x, t)$, ρ_i by $\rho(x)$, and recognize that the right-hand side of (114) approaches $\partial^2 u / \partial x^2$ as $\Delta x \rightarrow 0$. We end up with the continuous model for a wave string:

$$\rho \frac{\partial^2 u}{\partial t^2} = T \frac{\partial^2 u}{\partial x^2}.$$

Note that the density ρ may change along the string, while the tension T is constant. With variable wave velocity $c(x) = \sqrt{T/\rho(x)}$ we can write the wave equation in the more standard form

$$\frac{\partial^2 u}{\partial t^2} = c^2(x) \frac{\partial^2 u}{\partial x^2}.$$

Because of the way ρ enters the equations, the variable wave velocity c appears inside the derivatives as in many other versions of the wave equation. However, most strings of interest have constant ρ .

The end points of a string are fixed so that the displacement u is zero. The boundary conditions are therefore $u = 0$.

Damping. Air resistance and non-elastic effects in the string will contribute to reduce the amplitudes of the waves so that the motion dies out after a certain time. This damping effect can be modeled by a term bu_t on the left-hand side of the equation

$$\rho \frac{\partial^2 u}{\partial t^2} + b \frac{\partial u}{\partial t} = T \frac{\partial^2 u}{\partial x^2}.$$

The parameter b must normally be determined from physical experiments.

External forcing. It is easy to include an external force acting on the string. If we have a vertical force $\tilde{f}_i \mathbf{j}$ acting on mass m_i . This force affects the vertical component of Newton's law and gives rise to an extra term $\tilde{f}(x, t)$ on the right-hand side of (116). In the model (117) we would add a term $\mathbf{f}(x, t) = \tilde{f}(x, y)/\varrho(x)$.

Modeling the tension via springs. We assumed, in the derivation above, that the tension in the string, T , was constant. It is easy to check this assumption by modeling the string segments between the masses as standard springs, where the force (tension T) is proportional to the elongation of the spring segment. Let k be the spring constant, and set $T_i = k\Delta\ell$ for the tension in the spring segment between x_{i-1} and x_i , where $\Delta\ell$ is the elongation of this segment from the tension-free state. A basic feature of a string is that it has high tension in the equilibrium position $u = 0$. Let the string segment have an elongation $\Delta\ell_0$ at the equilibrium position. After deformation of the string, the elongation is $\ell = \Delta\ell_0 + \Delta s_i$: $T_i = k(\Delta\ell_0 + \Delta s_i) \approx k(\Delta\ell_0 + \Delta x)$. This shows that T_i is independent of i . Moreover, the extra approximate elongation Δx is very small compared to $\Delta\ell_0$, so we may well set $T_i = T = k\Delta\ell_0$. This means that the tension is completely dominated by the initial tension determined by the tuning of the string. The additional deformations of the spring during the vibrations do not introduce significant changes in the tension.

9.2 Waves on a membrane

9.3 Elastic waves in a rod

Consider an elastic rod subject to a hammer impact at the end. This experiment will give rise to an elastic deformation pulse that travels through the rod. A mathematical model for longitudinal waves along an elastic rod starts with the general equation for deformations and stresses in an elastic medium,

$$\varrho \mathbf{u}_{tt} = \nabla \cdot \boldsymbol{\sigma} + \varrho \mathbf{f}, \quad (119)$$

where ϱ is the density, \mathbf{u} the displacement field, $\boldsymbol{\sigma}$ the stress tensor, and \mathbf{f} body forces. The latter has normally no impact on elastic waves.

For stationary deformation of an elastic rod, one has that $\sigma_{xx} = E u_x$, with all other stress components being zero. Moreover, $\mathbf{u} = u(x)\mathbf{i}$. The parameter E is known as Young's modulus. Assuming that this simple stress and deformation field, which is exact in the stationary case, is a good approximation in the transient case with wave motion, (119) simplifies to

$$\varrho \frac{\partial^2 u}{\partial t^2} = \frac{\partial}{\partial x} \left(E \frac{\partial u}{\partial x} \right). \quad (120)$$

The associated boundary conditions are u or $\sigma_{xx} = E u_x$ known, typically $u = 0$ for a clamped end and $\sigma_{xx} = 0$ for a free end.

19.4 The acoustic model for seismic waves

Seismic waves are used to infer properties of subsurface geological structures. The physical model is a heterogeneous elastic medium where sound is produced by small elastic vibrations. The general mathematical model for deformation in an elastic medium is based on Newton's second law,

$$\varrho \mathbf{u}_{tt} = \nabla \cdot \boldsymbol{\sigma} + \varrho \mathbf{f},$$

and a constitutive law relating $\boldsymbol{\sigma}$ to \mathbf{u} , often Hooke's generalized law,

$$\boldsymbol{\sigma} = K \nabla \cdot \mathbf{u} \mathbf{I} + G(\nabla \mathbf{u} + (\nabla \mathbf{u})^T - \frac{2}{3} \nabla \cdot \mathbf{u} \mathbf{I}).$$

Here, \mathbf{u} is the displacement field, $\boldsymbol{\sigma}$ is the stress tensor, \mathbf{I} is the identity tensor, ϱ is the medium's density, \mathbf{f} are body forces (such as gravity), K is the bulk modulus and G is the shear modulus. All these quantities may vary in space, while \mathbf{u} and $\boldsymbol{\sigma}$ will also show significant variation in time during motion.

The acoustic approximation to elastic waves arises from a basic assumption that the second term in Hooke's law, representing the deformations that give rise to shear stresses, can be neglected. This assumption can be interpreted as approximating the geological medium by a fluid. Neglecting also the body forces, (121) becomes

$$\varrho \mathbf{u}_{tt} = \nabla(K \nabla \cdot \mathbf{u})$$

Introducing p as a pressure via

$$p = -K \nabla \cdot \mathbf{u},$$

and dividing (123) by ϱ , we get

$$\mathbf{u}_{tt} = -\frac{1}{\varrho} \nabla p.$$

Taking the divergence of this equation, using $\nabla \cdot \mathbf{u} = -p/K$ from (124) and the *acoustic approximation to elastic waves*:

$$p_{tt} = K \nabla \cdot \left(\frac{1}{\varrho} \nabla p \right).$$

This is a standard, linear wave equation with variable coefficients. It is common to add a source term $s(x, y, z, t)$ to model the generation of sound waves

$$p_{tt} = K \nabla \cdot \left(\frac{1}{\varrho} \nabla p \right) + s.$$

A common additional approximation of (127) is based on using the divergence rule on the right-hand side,

$$K\nabla \cdot \left(\frac{1}{\varrho} \nabla p \right) = \frac{K}{\varrho} \nabla^2 p + K\nabla \left(\frac{1}{\varrho} \right) \cdot \nabla p \approx \frac{K}{\varrho} \nabla^2 p,$$

Under the assumption that the relative spatial gradient $\nabla \varrho^{-1} = -\varrho^{-2} \nabla \varrho$ is small, this approximation results in the simplified equation

$$p_{tt} = \frac{K}{\varrho} \nabla^2 p + s. \quad (128)$$

The acoustic approximations to seismic waves are used for sound waves in the ground, and the Earth's surface is then a boundary where p equals the atmospheric pressure p_0 such that the boundary condition becomes $p = p_0$.

anisotropy. Quite often in geological materials, the effective wave velocity $= \sqrt{K/\varrho}$ is different in different spatial directions because geological layers are compacted such that the properties in the horizontal and vertical direction differ. With z as the vertical coordinate, we can introduce a vertical wave velocity c_z and a horizontal wave velocity c_h , and generalize (128) to

$$p_{tt} = c_z^2 p_{zz} + c_h^2 (p_{xx} + p_{yy}) + s. \quad (129)$$

9.5 Sound waves in liquids and gases

Sound waves arise from pressure and density variations in fluids. The starting point of modeling sound waves is the basic equations for a compressible fluid here we omit viscous (frictional) forces, body forces (gravity, for instance), and temperature effects:

$$\varrho_t + \nabla \cdot (\varrho \mathbf{u}) = 0, \quad (130)$$

$$\varrho \mathbf{u}_t + \varrho \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p, \quad (131)$$

$$\varrho = \varrho(p). \quad (132)$$

These equations are often referred to as the Euler equations for the motion of a fluid. The parameters involved are the density ϱ , the velocity \mathbf{u} , and the pressure p . Equation (131) reflects mass balance, (130) is Newton's second law for a fluid, with frictional and body forces omitted, and (132) is a constitutive law relating density to pressure by thermodynamics considerations. A typical model for (132) is the so-called **isentropic relation**²⁸, valid for adiabatic processes where there is no heat transfer:

$$\varrho = \varrho_0 \left(\frac{p}{p_0} \right)^{1/\gamma}. \quad (133)$$

²⁸http://en.wikipedia.org/wiki/Isentropic_process

Here, p_0 and ϱ_0 are reference values for p and ϱ when the fluid is at rest. γ is the ratio of specific heat at constant pressure and constant volume (for air).

The key approximation in a mathematical model for sound waves is to assume that these waves are small perturbations to the density, pressure, and velocity. We therefore write

$$\begin{aligned} p &= p_0 + \hat{p}, \\ \varrho &= \varrho_0 + \hat{\varrho}, \\ \mathbf{u} &= \hat{\mathbf{u}}, \end{aligned}$$

where we have decomposed the fields into a constant equilibrium value corresponding to $\mathbf{u} = 0$, and a small perturbation marked with a hat symbol. Inserting these decompositions in (130) and (131), neglecting all products of small perturbations and/or their derivatives, and dropping the hat symbols, one gets the following linearized PDE system for the small perturbations of density, pressure, and velocity:

$$\begin{aligned} \varrho_t + \varrho_0 \nabla \cdot \mathbf{u} &= 0, \\ \varrho_0 \mathbf{u}_t &= -\nabla p. \end{aligned}$$

Now we can eliminate ϱ_t by differentiating the relation $\varrho(p)$,

$$\varrho_t = \varrho_0 \frac{1}{\gamma} \left(\frac{p}{p_0} \right)^{1/\gamma-1} \frac{1}{p_0} p_t = \frac{\varrho_0}{\gamma p_0} \left(\frac{p}{p_0} \right)^{1/\gamma-1} p_t.$$

The product term $p^{1/\gamma-1} p_t$ can be linearized as $p_0^{1/\gamma-1} p_t$, resulting in

$$\varrho_t \approx \frac{\varrho_0}{\gamma p_0} p_t.$$

We then get

$$\begin{aligned} p_t + \gamma p_0 \nabla \cdot \mathbf{u} &= 0, \\ \mathbf{u}_t &= -\frac{1}{\varrho_0} \nabla p. \end{aligned}$$

Taking the divergence of (137) and differentiating (136) with respect to time gives the possibility to easily eliminate $\nabla \cdot \mathbf{u}_t$ and arrive at a standard wave equation for p :

$$p_{tt} = c^2 \nabla^2 p,$$

where $c = \sqrt{\gamma p_0 / \varrho_0}$ is the speed of sound in the fluid.

9.6 Spherical waves

spherically symmetric three-dimensional waves propagate in the radial direction only so that $u = u(r, t)$. The fully three-dimensional wave equation

$$\frac{\partial^2 u}{\partial t^2} = \nabla \cdot (c^2 \nabla u) + f$$

then reduces to the spherically symmetric wave equation

$$\frac{\partial^2 u}{\partial t^2} = \frac{1}{r^2} \frac{\partial}{\partial r} \left(c^2(r) r^2 \frac{\partial u}{\partial r} \right) + f(r, t), \quad r \in (0, R), \quad t > 0. \quad (139)$$

We can easily show that the function $v(r, t) = ru(r, t)$ fulfills a standard wave equation in Cartesian coordinates if c is constant. To this end, insert $u = v/r$ in

$$\frac{1}{r^2} \frac{\partial}{\partial r} \left(c^2(r) r^2 \frac{\partial u}{\partial r} \right)$$

to obtain

$$r \left(\frac{dc^2}{dr} \frac{\partial v}{\partial r} + c^2 \frac{\partial^2 v}{\partial r^2} \right) - \frac{dc^2}{dr} v.$$

The two terms in the parenthesis can be combined to

$$r \frac{\partial}{\partial r} \left(c^2 \frac{\partial v}{\partial r} \right),$$

which is recognized as the variable-coefficient Laplace operator in one Cartesian coordinate. The spherically symmetric wave equation in terms of $v(r, t)$ now becomes

$$\frac{\partial^2 v}{\partial t^2} = \frac{\partial}{\partial r} \left(c^2(r) \frac{\partial v}{\partial r} \right) - \frac{1}{r} \frac{dc^2}{dr} v + rf(r, t), \quad r \in (0, R), \quad t > 0. \quad (140)$$

In the case of constant wave velocity c , this equation reduces to the wave equation in a single Cartesian coordinate called r :

$$\frac{\partial^2 v}{\partial t^2} = c^2 \frac{\partial^2 v}{\partial r^2} + rf(r, t), \quad r \in (0, R), \quad t > 0. \quad (141)$$

That is, any program for solving the one-dimensional wave equation in a Cartesian coordinate system can be used to solve (141), provided the source term is multiplied by the coordinate, and that we divide the Cartesian mesh solution by r to get the spherically symmetric solution. Moreover, if $r = 0$ is included in the domain, spherical symmetry demands that $\partial u / \partial r = 0$ at $r = 0$, which means that

$$\frac{\partial u}{\partial r} = \frac{1}{r^2} \left(r \frac{\partial v}{\partial r} - v \right) = 0, \quad r = 0,$$

implying $v(0, t) = 0$ as a necessary condition. For practical applications, we exclude $r = 0$ from the domain and assume that some boundary condition is assigned at $r = \epsilon$, for some $\epsilon > 0$.

19.7 The linear shallow water equations

The next example considers water waves whose wavelengths are much larger than the depth and whose wave amplitudes are small. This class of waves is generated by catastrophic geophysical events, such as earthquakes at the bottom, landslides moving into water, or underwater slides (or a combination of earthquakes frequently release avalanches of masses). For example, a large earthquake will normally have an extension of many kilometers but lift the water only a few meters. The wave length will have a size dictated by the earthquake area, which is much larger than the water depth, and compared to the wavelength, an amplitude of a few meters is very small. The water is essentially a thin film, and mathematically we can average the problem in the vertical direction and approximate the 3D wave phenomenon by 2D PDEs. Instead of a 3D water domain in three space dimensions, we get a horizontal 2D domain with an unknown function for the surface elevation and the water depth as a coefficient in the PDEs.

Let $\eta(x, y, t)$ be the elevation of the water surface, $H(x, y)$ the water depth corresponding to a flat surface ($\eta = 0$), $u(x, y, t)$ and $v(x, y, t)$ the depth-averaged horizontal velocities of the water. Mass and momentum balance of the water volume give rise to the PDEs involving these quantities:

$$\begin{aligned} \eta_t &= -(Hu)_x - (Hv)_y \\ u_t &= -g\eta_x, \\ v_t &= -g\eta_y, \end{aligned}$$

where g is the acceleration of gravity. Equation (142) corresponds to mass balance while the other two are derived from momentum balance (Newton's second law).

The initial conditions associated with (142)-(144) are η , u , and v prescribed at $t = 0$. A common condition is to have some water elevation $\eta = I(x, y)$ and assume that the surface is at rest: $u = v = 0$. A submarine earthquake means a sufficiently rapid motion of the bottom and the water volume such that the bottom deformation is mirrored at the water surface as an elevation $I(x, y)$ and that $u = v = 0$.

Boundary conditions may be η prescribed for incoming, known water depth, zero normal velocity at reflecting boundaries (steep mountains, for incompressible water $un_x + vn_y = 0$, where (n_x, n_y) is the outward unit normal to the boundary), sophisticated boundary conditions are needed when waves run up at the shore and at open boundaries where we want the waves to leave the computational domain undisturbed.

Equations (142), (143), and (144) can be transformed to a standard, linear wave equation. First, multiply (143) and (144) by H , differentiate (143) with respect to x and (144) with respect to y . Second, differentiate (142) with respect to t and use that $(Hu)_{xt} = (Hu_t)_x$ and $(Hv)_{yt} = (Hv_t)_y$ when H is independent of t . Third, eliminate $(Hu_t)_x$ and $(Hv_t)_y$ with the aid of the other two differentiated equations. These manipulations results in a standard, linear wave equation for η :

$$\eta_{tt} = (gH\eta_x)_x + (gH\eta_y)_y = \nabla \cdot (gH\nabla\eta). \quad (145)$$

In the case we have an initial non-flat water surface at rest, the initial conditions become $\eta = I(x, y)$ and $\eta_t = 0$. The latter follows from (142) if $u = v = 0$, or simply from the fact that the vertical velocity of the surface is η_t , which is zero for a surface at rest.

The system (142)-(144) can be extended to handle a time-varying bottom topography, which is relevant for modeling long waves generated by underwater slides. In such cases the water depth function H is also a function of t , due to the moving slide, and one must add a time-derivative term H_t to the left-hand side of (142). A moving bottom is best described by introducing $z = H_0$ as the still-water level, $z = B(x, y, t)$ as the time- and space-varying bottom topography, so that $H = H_0 - B(x, y, t)$. In the elimination of u and v one may assume that the dependence of H on t can be neglected in the terms $(Hu)_{xt}$ and $(Hv)_{yt}$. We then end up with a source term in (145), because of the moving (accelerating) bottom:

$$\eta_{tt} = \nabla \cdot (gH\nabla\eta) + B_{tt}. \quad (146)$$

The reduction of (146) to 1D, for long waves in a straight channel, or for approximately plane waves in the ocean, is trivial by assuming no change in y direction ($\partial/\partial y = 0$):

$$\eta_{tt} = (gH\eta_x)_x + B_{tt}. \quad (147)$$

Wind drag on the surface. Surface waves are influenced by the drag of the wind, and if the wind velocity some meters above the surface is (U, V) , the wind drag gives contributions $C_V\sqrt{U^2 + V^2}U$ and $C_V\sqrt{U^2 + V^2}V$ to (143) and (144), respectively, on the right-hand sides.

Bottom drag. The waves will experience a drag from the bottom, often roughly modeled by a term similar to the wind drag: $C_B\sqrt{u^2 + v^2}u$ on the right-hand side of (143) and $C_B\sqrt{u^2 + v^2}v$ on the right-hand side of (144). Note that in this case the PDEs (143) and (144) become nonlinear and the elimination of u and v to arrive at a 2nd-order wave equation for η is not possible anymore.

Effect of the Earth's rotation. Long geophysical waves will often be affected by the rotation of the Earth because of the Coriolis force. This force gives rise to a term fv on the right-hand side of (143) and $-fu$ on the right-hand side of (144). Also in this case one cannot eliminate u and v to work with an equation for η . The Coriolis parameter is $f = 2\Omega \sin \phi$, where Ω is the angular velocity of the earth and ϕ is the latitude.

19.8 Waves in blood vessels

The flow of blood in our bodies is basically fluid flow in a network of vessels. Unlike rigid pipes, the walls in the blood vessels are elastic and will increase their diameter when the pressure rises. The elastic forces will then push back and accelerate the fluid. This interaction between the flow of blood and the deformation of the vessel wall results in waves traveling along our blood vessels.

A model for one-dimensional waves along blood vessels can be derived by averaging the fluid flow over the cross section of the blood vessels. Let r be the radial coordinate along the blood vessel and assume that all cross sections are circular, although with different radius $R(x, t)$. The main quantities to compute are the cross section area $A(x, t)$, the averaged pressure $P(x, t)$, and the total flux $Q(x, t)$. The area of this cross section is

$$A(x, t) = 2\pi \int_0^{R(x,t)} r dr,$$

Let $v_x(x, t)$ be the velocity of blood averaged over the cross section at x . The volume flux, being the total volume of blood passing a cross section in a time unit, becomes

$$Q(x, t) = A(x, t)v_x(x, t)$$

Mass balance and Newton's second law lead to the PDEs

$$\begin{aligned} \frac{\partial A}{\partial t} + \frac{\partial Q}{\partial x} &= 0, \\ \frac{\partial Q}{\partial t} + \frac{\gamma + 2}{\gamma + 1} \frac{\partial}{\partial x} \left(\frac{Q^2}{A} \right) + \frac{A}{\rho} \frac{\partial P}{\partial x} &= -2\pi(\gamma + 2) \frac{\mu}{\rho} \frac{Q}{A}, \end{aligned}$$

where γ is a parameter related to the velocity profile, ρ is the density of blood, and μ is the dynamic viscosity of blood.

We have three unknowns A , Q , and P , and two equations (150) and (151). A third equation is needed to relate the flow to the deformations of the vessel wall. A common form for this equation is

$$\frac{\partial P}{\partial t} + \frac{1}{C} \frac{\partial Q}{\partial x} = 0,$$

where C is the compliance of the wall, given by the constitutive relation

$$C = \frac{\partial A}{\partial P} + \frac{\partial A}{\partial t},$$

hich require a relationship between A and P . One common model is to view the vessel wall, locally, as a thin elastic tube subject to an internal pressure. This gives the relation

$$P = P_0 + \frac{\pi h E}{(1 - \nu^2) A_0} (\sqrt{A} - \sqrt{A_0}),$$

here P_0 and A_0 are corresponding reference values when the wall is not deformed, h is the thickness of the wall, and E and ν are Young's modulus and Poisson's ratio of the elastic material in the wall. The derivative becomes

$$C = \frac{\partial A}{\partial P} = \frac{2(1 - \nu^2) A_0}{\pi h E} \sqrt{A_0} + 2 \left(\frac{(1 - \nu^2) A_0}{\pi h E} \right)^2 (P - P_0). \quad (154)$$

Another (nonlinear) deformation model of the wall, which has a better fit with experiments, is

$$P = P_0 \exp(\beta(A/A_0 - 1)),$$

here β is some parameter to be estimated. This law leads to

$$C = \frac{\partial A}{\partial P} = \frac{A_0}{\beta P}. \quad (155)$$

Reduction to standard wave equation. It is not uncommon to neglect the viscous term on the right-hand side of (151) and also the quadratic term with Q^2 on the left-hand side. The reduced equations (151) and (152) form a first-order near wave equation system:

$$C \frac{\partial P}{\partial t} = -\frac{\partial Q}{\partial x}, \quad (156)$$

$$\frac{\partial Q}{\partial t} = -\frac{A}{\rho} \frac{\partial P}{\partial x}. \quad (157)$$

These can be combined into standard 1D wave equation PDE by differentiating the first equation with respect to t and the second with respect to x ,

$$\frac{\partial}{\partial t} \left(C C \frac{\partial P}{\partial t} \right) = \frac{\partial}{\partial x} \left(\frac{A}{\rho} \frac{\partial P}{\partial x} \right),$$

which can be approximated by

$$\frac{\partial^2 Q}{\partial t^2} = c^2 \frac{\partial^2 Q}{\partial x^2}, \quad c = \sqrt{\frac{A}{\rho C}}, \quad (158)$$

here the A and C in the expression for c are taken as constant reference values.

19.9 Electromagnetic waves

Light and radio waves are governed by standard wave equations arising from Maxwell's general equations. When there are no charges and no currents in a vacuum, Maxwell's equations take the form

$$\begin{aligned} \nabla \cdot \mathbf{E} &= 0, \\ \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t}, \\ \nabla \times \mathbf{B} &= \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}, \end{aligned}$$

where $\epsilon_0 = 8.854187817620 \cdot 10^{-12}$ (F/m) is the permittivity of free space known as the electric constant, and $\mu_0 = 1.2566370614 \cdot 10^{-6}$ (H/m) is the permeability of free space, also known as the magnetic constant. Taking the curl of the two last equations and using the identity

$$\nabla \times (\nabla \times \mathbf{E}) = \nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = -\nabla^2 \mathbf{E} \text{ when } \nabla \cdot \mathbf{E} = 0,$$

immediately gives the wave equation governing the electric and magnetic fields

$$\begin{aligned} \frac{\partial^2 \mathbf{E}}{\partial t^2} &= c^2 \frac{\partial^2 \mathbf{E}}{\partial x^2}, \\ \frac{\partial^2 \mathbf{B}}{\partial t^2} &= c^2 \frac{\partial^2 \mathbf{B}}{\partial x^2}, \end{aligned}$$

with $c = 1/\sqrt{\mu_0 \epsilon_0}$ as the velocity of light. Each component of \mathbf{E} and \mathbf{B} satisfies a wave equation and can hence be solved independently.

20 Exercises

Exercise 14: Simulate waves on a non-homogeneous string

Simulate waves on a string that consists of two materials with different densities. The tension in the string is constant, but the density has a jump at the interface. Experiment with different sizes of the jump and produce animations that visualize the effect of the jump on the wave motion.

Hint. According to Section 19.1, the density enters the mathematical model through the term ρ in $\rho u_{tt} = T u_{xx}$, where T is the string tension. Modify, e.g., the `wave1D_u0` code to incorporate the tension and two density values. Make a mesh for `rho` with density values at each spatial mesh point. A value for the tension should be 150 N. Corresponding density values can be computed from the wave equation estimations in the `guitar` function in the `wave1D_u0_sv.py` file.

Filename: `wave1D_u0_sv_discont.py`.

Exercise 15: Simulate damped waves on a string

Formulate a mathematical model for damped waves on a string. Use data from Section 3.4, and tune the damping parameter so that the string is very close to the rest state after 15 s. Make a movie of the wave motion. Filename: `wave1D_u0_sv_damping.py`.

Exercise 16: Simulate elastic waves in a rod

A hammer hits the end of an elastic rod. The exercise is to simulate the resulting wave motion using the model (120) from Section 19.3. Let the rod have length L and let the boundary $x = L$ be stress free so that $\sigma_{xx} = 0$, implying that $u/\partial x = 0$. The left end $x = 0$ is subject to a strong stress pulse (the hammer), modeled as

$$\sigma_{xx}(t) = \begin{cases} S, & 0 < t \leq t_s, \\ 0, & t > t_s \end{cases}$$

The corresponding condition on u becomes $u_x = S/E$ for $t \leq t_s$ and zero afterwards (recall that $\sigma_{xx} = Eu_x$). This is a non-homogeneous Neumann condition, and you will need to approximate this condition and combine it with the scheme (the ideas and manipulations follow closely the handling of a non-zero initial condition $u_t = V$ in wave PDEs or the corresponding second-order ODEs or vibrations). Filename: `wave_rod.py`.

Exercise 17: Simulate spherical waves

Implement a model for spherically symmetric waves using the method described in Section 19.6. The boundary condition at $r = 0$ must be $\partial u/\partial r = 0$, while the condition at $r = R$ can either be $u = 0$ or a radiation condition as described in Problem 20. The $u = 0$ condition is sufficient if R is so large that the amplitude of the spherical wave has become insignificant. Make movie(s) of the case where the source term is located around $r = 0$ and sends out pulses

$$f(r, t) = \begin{cases} Q \exp(-\frac{r^2}{2\Delta r^2}) \sin \omega t, & \sin \omega t \geq 0 \\ 0, & \sin \omega t < 0 \end{cases}$$

where, Q and ω are constants to be chosen.

Hint. Use the program `wave1D_u0_sv.py` as a starting point. Let `solver` compute the v function and then set $u = v/r$. However, $u = v/r$ for $r = 0$ requires special treatment. One possibility is to compute `u[1:] = v[1:]/r[1:]` and then set `u[0]=u[1]`. The latter makes it evident that $\partial u/\partial r = 0$ in a plot.

Filename: `wave1D_spherical.py`.

Exercise 18: Explain why numerical noise occurs

The experiments performed in Exercise 8 shows considerable numeric noise in the form of non-physical waves, especially for $s_f = 4$ and the plot of the half a "cosinehat" pulse. The noise is much less visible for a Gaussian pulse. Run the case with the plug and half a "cosinehat" pulses for $C = 0.9, 0.25$, and $N_x = 40, 80, 160$. Use the numerical dispersion relation to explain the observations. Filename: `pulse1D_analysis.pdf`.

Exercise 19: Investigate harmonic averaging in a 1D problem

Harmonic means are often used if the wave velocity is non-smooth or discontinuous. Will harmonic averaging of the wave velocity give less numerical noise for the case $s_f = 4$ in Exercise 8? Filenames: `pulse1D_harmonic.py`, `pulse1D_harmonic.py`.

Problem 20: Implement open boundary conditions

To enable a wave to leave the computational domain and travel undisturbed through the boundary $x = L$, one can in a one-dimensional problem impose the following condition, called a *radiation condition* or *open boundary condition*

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = 0.$$

The parameter c is the wave velocity.

Show that (161) accepts a solution $u = g_R(x - ct)$, but not $u = g_L(x + ct)$. This means that (161) will allow any right-going wave $g_R(x - ct)$ to pass through the boundary.

A corresponding open boundary condition for a left-going wave through the boundary is

$$\frac{\partial u}{\partial t} - c \frac{\partial u}{\partial x} = 0.$$

The condition (161) can be discretized by centered differences at the end point $i = N_x$, corresponding to $x = x_R$:

$$[D_{2t}u + cD_{2x}u = 0]_{N_x}^n.$$

Eliminate the fictitious value $u_{N_x+1}^n$ by using the discrete equation at the end point. The equation for the first step, u_i^1 , is in principle affected, but then use the condition $u_{N_x} = 0$ since the wave has not yet reached the boundary.

A corresponding open boundary condition for a left-going wave through the boundary is

$$\frac{\partial u}{\partial t} - c \frac{\partial u}{\partial x} = 0.$$

plement a solver that incorporates the conditions (163) and (164). Start with the peak-shaped Gaussian function in the middle of the domain as $I(x)$ and demonstrate that waves travel undisturbed out of the domain at $x = L$ and $x = 0$. Make a nose test for checking that the surface is flat after a certain time.

Remark. The condition (161) works perfectly in 1D when c is known. In 2D and 3D, however, the condition reads $u_t + c_x u_x + c_y u_y = 0$, where c_x and c_y are the wave speeds in the x and y directions. Estimating these components (i.e., the direction of the wave) is often challenging. Other methods are normally used in 2D and 3D to let waves move out of a computational domain. Filename: `ave1D_open_BC.py`.

Problem 21: Earthquake-generated tsunami over a subsea hill

A subsea earthquake leads to an immediate lift of the water surface, see Figure 10. The lifted water surface splits into two tsunamis, one traveling to the right and one to the left, as depicted in Figure 11. Since tsunamis are normally very long waves, compared to the depth, with a small amplitude, compared to the wavelength, the wave equation model described in Section 19.7 is relevant:

$$\eta_{tt} = (gH(x)\eta_x)_x,$$

where g is the acceleration of gravity, and $H(x)$ is the still water depth.




Figure 10: Sketch of initial water surface due to a subsea earthquake.

To simulate the right-going tsunami, we can impose a symmetry boundary at $x = 0$: $\partial\eta/\partial x = 0$. We then simulate the wave motion in $[0, L]$. Unless the ocean ends at $x = L$, the waves should travel undisturbed through the boundary at $x = L$. A radiation condition as explained in Problem 20 can be used for this




Figure 11: An initial surface elevation is split into two waves.

purpose. Alternatively, one can just stop the simulations before the wave reaches the boundary at $x = L$. In that case it does not matter what kind of boundary condition we use at $x = L$. Imposing $\eta = 0$ and stopping the simulation when $|\eta_i^n| > \epsilon$, $i = N_x - 1$, is a possibility (ϵ is a small parameter).

The shape of the initial surface can be taken as a Gaussian function

$$I(x; I_0, I_a, I_m, I_s) = I_0 + I_a \exp\left(-\left(\frac{x - I_m}{I_s}\right)^2\right),$$

with $I_m = 0$ reflecting the location of the peak of $I(x)$ and I_s being a measure of the width of the function $I(x)$ (I_s is $\sqrt{2}$ times the standard deviation of the familiar normal distribution curve).

Now we extend the problem with a hill at the sea bottom, see Figure 12. The wave speed $c = \sqrt{gH(x)} = \sqrt{g(H_0 - B(x))}$ will then be reduced in the water above the hill.

One possible form of the hill is a Gaussian function,

$$B(x; B_0, B_a, B_m, B_s) = B_0 + B_a \exp\left(-\left(\frac{x - B_m}{B_s}\right)^2\right),$$

but many other shapes are also possible, e.g., a "cosine hat" where

$$B(x; B_0, B_a, B_m, B_s) = B_0 + B_a \cos\left(\pi \frac{x - B_m}{2B_s}\right),$$

when $x \in [B_m - B_s, B_m + B_s]$ while $B = B_0$ outside this interval.

Also an abrupt construction may be tried:

$$B(x; B_0, B_a, B_m, B_s) = B_0 + B_a,$$




Figure 12: Sketch of an earthquake-generated tsunami passing over a subsea hill.

for $x \in [B_m - B_s, B_m + B_s]$ while $B = B_0$ outside this interval.

The `wave1D_dn_vc.py`²⁹ program can be used as starting point for the implementation. Visualize both the bottom topography and the water surface elevation in the same plot. Allow for a flexible choice of bottom shape: (166), (167), (168), or $B(x) = B_0$ (flat).

The purpose of this problem is to explore the quality of the numerical solution for different shapes of the bottom obstruction. The "cosine hat" and the box-shaped hills have abrupt changes in the derivative of $H(x)$ and are more likely to generate numerical noise than the smooth Gaussian shape of the hill. Investigate if this is true. Filenames: `tsunami1D_hill.py`, `tsunami1D_hill.pdf`.

Problem 22: Earthquake-generated tsunami over a 3D hill

This problem extends Problem 21 to a three-dimensional wave phenomenon, governed by the 2D PDE (145). We assume that the earthquake arises from a fault along the line $x = 0$ in the xy -plane so that the initial lift of the surface can be taken as $I(x)$ in Problem 21. That is, a plane wave is propagating to the right, but will experience bending because of the bottom.

The bottom shape is now a function of x and y . An "elliptic" Gaussian function in two dimensions, with its peak at (B_{mx}, B_{my}) , generalizes (166):

$$B(x; B_0, B_a, B_{mx}, B_{my}, B_s, b) = B_0 + B_a \exp\left(-\left(\frac{x - B_{mx}}{B_s}\right)^2 - \left(\frac{y - B_{my}}{bB_s}\right)^2\right), \quad (169)$$

²⁹http://tinyurl.com/jvzzcfn/wave/wave1D/wave1D_dn_vc.py

where b is a scaling parameter: $b = 1$ gives a circular Gaussian function, circular contour lines, while $b \neq 1$ gives an elliptic shape with elliptic contour lines.

The "cosine hat" (167) can also be generalized to

$$B(x; B_0, B_a, B_{mx}, B_{my}, B_s) = B_0 + B_a \cos\left(\pi \frac{x - B_{mx}}{2B_s}\right) \cos\left(\pi \frac{y - B_{my}}{2B_s}\right)$$

when $0 \leq \sqrt{x^2 + y^2} \leq B_s$ and $B = B_0$ outside this circle.

A box-shaped obstacle means that

$$B(x; B_0, B_a, B_m, B_s, b) = B_0 + B_a$$

for x and y inside a rectangle

$$B_{mx} - B_s \leq x \leq B_{mx} + B_s, \quad B_{my} - bB_s \leq y \leq B_{my} + bB_s,$$

and $B = B_0$ outside this rectangle. The b parameter controls the rectangle shape of the cross section of the box.

Note that the initial condition and the listed bottom shapes are symmetric around the line $y = B_{my}$. We therefore expect the surface elevation to be symmetric with respect to this line. This means that we can halve the computational domain by working with $[0, L_x] \times [0, B_{my}]$. Along the right boundary, $y = B_{my}$, we must impose the symmetry condition $\partial\eta/\partial n = 0$. A symmetry condition ($-\eta_x = 0$) is also needed at the $x = 0$ boundary. The initial condition has a symmetry here. At the lower boundary $y = 0$ set a Neumann condition (which becomes $-\eta_y = 0$). The wave motion is simulated until the wave hits the reflecting boundaries where $\partial\eta/\partial n = 0$ (one can also set $\eta = 0$ - the particular condition does not matter as long as the simulation is stopped before the wave is influenced by the boundary condition).

Visualize the surface elevation. Investigate how different hill shapes and different sizes of the water gap above the hill, and different resolutions $\Delta y = h$ and Δt influence the numerical quality of the solution. Filenames: `tsunami2D_hill.py`, `tsunami2D_hill.pdf`.

Problem 23: Investigate Matplotlib for visualization

Play with native Matplotlib code for visualizing 2D solutions of the wave equation with variable wave velocity. See if there are effective ways to visualize the solution and the wave velocity. Filename: `tsunami2D_hill_mpl.py`.

Problem 24: Investigate visualization packages

Create some fancy 3D visualization of the water waves and the subsea topography. Problem 22. Try to make the hill transparent. Possible visualization to

- [Mayavi](#)³⁰
- [Paraview](#)³¹
- [OpenDX](#)³²

filename: `tsunami2D_hill_viz.py`.

Problem 25: Implement loops in compiled languages

Extend the program from Problem 22 such that the loops over mesh points, inside the time loop, are implemented in compiled languages. Consider implementations in Cython, Fortran via `f2py`, C via Cython, C via `f2py`, C/C++ via Instant, and C/C++ via `scipy.weave`. Perform efficiency experiments to investigate the relative performance of the various implementations. It is often advantageous to normalize CPU times by the fastest method on a given mesh. Filename: `tsunami2D_hill_compiled.py`.

Exercise 26: Simulate seismic waves in 2D

The goal of this exercise is to simulate seismic waves using the PDE model (29) in a 2D xz domain with geological layers. Introduce m horizontal layers of thickness h_i , $i = 0, \dots, m - 1$. Inside layer number i we have a vertical wave velocity $c_{z,i}$ and a horizontal wave velocity $c_{h,i}$. Make a program for simulating such 2D waves. Test it on a case with 3 layers where

$$c_{z,0} = c_{z,1} = c_{z,2}, \quad c_{h,0} = c_{h,2}, \quad c_{h,1} \ll c_{h,0}.$$

Let s be a localized point source at the middle of the Earth's surface (the upper boundary) and investigate how the resulting wave travels through the medium. The source can be a localized Gaussian peak that oscillates in time over some time interval. Place the boundaries far enough from the expanding wave so that the boundary conditions do not disturb the wave. Then the type of boundary condition does not matter, except that we physically need to have $p = p_0$, where p_0 is the atmospheric pressure, at the upper boundary. Filename: `seismic2D.py`.

Project 27: Model 3D acoustic waves in a room

The equation for sound waves in air is derived in Section 19.5 and reads

$$p_{tt} = c^2 \nabla^2 p,$$

where $p(x, y, z, t)$ is the pressure and c is the speed of sound, taken as 340 m/s. However, sound is absorbed in the air due to relaxation of molecules in the gas.

³⁰<http://code.enthought.com/projects/mayavi/>

³¹<http://www.paraview.org/>

³²<http://www.opendx.org/>

A model for simple relaxation, valid for gases consisting only of one molecule, is a term $c^2 \tau_s \nabla^2 p_t$ in the PDE, where τ_s is the relaxation time to generate sound from, e.g., a loudspeaker in the room, this sound source also be added to the governing equation.

The PDE with the mentioned type of damping and source then becomes

$$p_{tt} = c^2 \nabla^2 p + c^2 \tau_s \nabla^2 p_t + f,$$

where $f(x, y, z, t)$ is the source term.

The walls can absorb some sound. A possible model is to have a "wall" (thicker than the physical wall) outside the room where c is changed such that some of the wave energy is reflected and some is absorbed in the wall. The absorption of energy can be taken care of by adding a damping term bp_t to the equation:

$$p_{tt} + bp_t = c^2 \nabla^2 p + c^2 \tau_s \nabla^2 p_t + f.$$

Typically, $b = 0$ in the room and $b > 0$ in the wall. A discontinuity in c will give rise to reflections. It can be wise to use a constant c in the room to control reflections because of the discontinuity between c in the air and in the wall, while b is gradually increased as we go into the wall to avoid reflections because of rapid changes in b . At the outer boundary of the wall the conditions $p = 0$ or $\partial p / \partial n = 0$ can be imposed. The waves should anyway be approximately damped to $p = 0$ this far out in the wall layer.

There are two strategies for discretizing the $\nabla^2 p_t$ term: using a central difference between times $n + 1$ and $n - 1$ (if the equation is sampled at n) or use a one-sided difference based on levels n and $n - 1$. The latter has the advantage of not leading to any equation system, while the former is second order accurate as the scheme for the simple wave equation $p_{tt} = c^2 \nabla^2 p$. To make the equation system, go for the one-sided difference such that the overall scheme becomes explicit and only of first order in time.

Develop a 3D solver for the specified PDE and introduce a wall layer. Implement a solver with the method of manufactured solutions. Make some demonstrations where the wall reflects and absorbs the waves (reflection because of discontinuity in c and absorption because of growing b). Experiment with the impact of the τ_s parameter. Filename: `acoustics.py`.

Project 28: Solve a 1D transport equation

We shall study the wave equation

$$u_t + cu_x = 0, \quad x \in (0, L], \quad t \in (0, T],$$

with initial condition

$$u(x, 0) = I(x), \quad x \in [0, L],$$

and one periodic boundary condition

$$u(0, t) = u(L, t). \quad (176)$$

his boundary condition means that what goes out of the domain at $x = L$ comes in at $x = 0$. Roughly speaking, we need only one boundary condition because of the spatial derivative is of first order only.

Physical interpretation. The parameter c can be constant or variable, $c = c(x)$. The equation (174) arises in *transport* problems where a quantity u , which could be temperature or concentration of some contaminant, is transported with the velocity c of a fluid. In addition to the transport imposed by "travelling with the fluid", u may also be transported by diffusion (such as heat conduction or Fickian diffusion), but we have in the model $u_t + cu_x$ assumed that diffusion effects are negligible, which they often are.

A widely used numerical scheme for (174) applies a forward difference in time and a backward difference in space when $c > 0$:

$$[D_t^+ u + cD_x^- u = 0]_i^n. \quad (177)$$

or $c < 0$ we use a forward difference in space: $[cD_x^+ u]_i^n$.

We shall hereafter assume that $c = c(x) > 0$.

To compute (182) we need to integrate $1/c$ to obtain C and then compute the inverse of C .

The inverse function computation can be easily done if we first think discretely. Say we have some function $y = g(x)$ and seek its inverse. Plotting (x_i, y_i) , here $y_i = g(x_i)$ for some mesh points x_i , displays g as a function of x . The inverse function is simply x as a function of y , i.e., the curve with points (y_i, x_i) . We can therefore quickly compute points at the curve of the inverse function. One way of extending these points to a continuous function is to assume a linear variation (known as linear interpolation) between the points (which actually means to draw straight lines between the points, exactly as done by a plotting program).

The function `wrap2callable` in `scitools.std` can take a set of points and return a continuous function that corresponds to linear variation between the points. The computation of the inverse of a function g on $[0, L]$ can then be done by

```
def inverse(g, domain, resolution=101):
    x = linspace(domain[0], domain[L], resolution)
    y = g(x)
    from scitools.std import wrap2callable
    g_inverse = wrap2callable((y, x))
    return g_inverse
```

To compute $C(x)$ we need to integrate $1/c$, which can be done by a Trapezoidal rule. Suppose we have computed $C(x_i)$ and need to compute $C(x_{i+1})$. Using the Trapezoidal rule with m subintervals over the integration domain $[x_i, x_{i+1}]$ gives

$$C(x_{i+1}) = C(x_i) + \int_{x_i}^{x_{i+1}} \frac{dx}{c} \approx h \left(\frac{1}{2} \frac{1}{c(x_i)} + \frac{1}{2} \frac{1}{c(x_{i+1})} + \sum_{j=1}^{m-1} \frac{1}{c(x_i + \dots)}$$

where $h = (x_{i+1} - x_i)/m$ is the length of the subintervals used for the integral over $[x_i, x_{i+1}]$. We observe that (178) is a *difference equation* which we can solve by repeatedly applying (178) for $i = 0, 1, \dots, N_x - 1$ if a mesh x_0, x_1, \dots is prescribed. Note that $C(0) = 0$.

a) Show that under the assumption of $a = \text{const}$,

$$u(x, t) = I(x - ct)$$

fulfills the PDE as well as the initial and boundary condition (provided $I(L) = 0$).

b) Set up a computational algorithm and implement it in a function. Assume a is constant and positive.

c) Test implementation by using the remarkable property that the numerical solution is exact at the mesh points if $\Delta t = c^{-1} \Delta x$.

d) Make a movie comparing the numerical and exact solution for the PDE for two choices of initial conditions:

$$I(x) = \left[\sin \left(\pi \frac{x}{L} \right) \right]^{2n}$$

where n is an integer, typically $n = 5$, and

$$I(x) = \exp \left(- \frac{(x - L/2)^2}{2\sigma^2} \right).$$

Choose $\Delta t = c^{-1} \Delta x, 0.9c^{-1} \Delta x, 0.5c^{-1} \Delta x$.

e) The performance of the suggested numerical scheme can be investigated by analyzing the numerical dispersion relation. Analytically, we have the *Fourier component*

$$u(x, t) = e^{i(kx - \omega t)},$$

is a solution of the PDE if $\omega = kc$. This is the *analytical dispersion relation*. A complete solution of the PDE can be built by adding up such Fourier components with different amplitudes, where the initial condition I determines the amplitudes. The solution u is then represented by a Fourier series.

A similar discrete Fourier component at (x_p, t_n) is

$$u_p^q = e^{i(kp\Delta x - \tilde{\omega}n\Delta t)},$$

here in general $\tilde{\omega}$ is a function of k , Δt , and Δx , and differs from the exact $\omega = kc$.

Insert the discrete Fourier component in the numerical scheme and derive an expression for $\tilde{\omega}$, i.e., the discrete dispersion relation. Show in particular that the $\Delta t/(c\Delta x) = 1$, the discrete solution coincides with the exact solution at the mesh points, regardless of the mesh resolution (!). Show that if the stability condition

$$\frac{\Delta t}{c\Delta x} \leq 1,$$

the discrete Fourier component cannot grow (i.e., $\tilde{\omega}$ is real).

) Write a test for your implementation where you try to use information from the numerical dispersion relation.

) Set up a computational algorithm for the variable coefficient case and implement it in a function. Make a test that the function works for constant a .

) It can be shown that for an observer moving with velocity $c(x)$, u is constant. This can be used to derive an exact solution when a varies with x . Show first that

$$u(x, t) = f(C(x) - t), \quad (182)$$

here

$$C'(x) = \frac{1}{c(x)},$$

a solution of (174) for any differentiable function f .

olution. Let $\xi = C(x) - t$. We have that

$$u_t = f'(\xi)(-1),$$

while

$$u_x = f'(\xi)C'(x) = f'(\xi)\frac{1}{c(x)},$$

implying that $au_x = f'(\xi)$. Then we have $u_t + cu_x = -f'(\xi) + f'(\xi) = 0$.

) Use the initial condition to show that an exact solution is

$$u(x, t) = I(C^{-1}(C(x) - t)),$$

with C^{-1} being the inverse function of $C = \int c^1 dx$. Since $C(x)$ is an integral $\int (1/c) dx$, $C(x)$ is monotonically increasing and there exists hence an inverse function C^{-1} with values in $[0, L]$.

Solution. In general we have $u(x, t) = f(C(x) - t)$ and the solution in form with $f(\xi) = I(C^{-1}(\xi))$. Moreover, at $t = 0$ we have $I(C^{-1}(C(x)))$ which is the required initial condition.

j) Implement a function for computing $C(x_i)$ and one for computing C^{-1} at any x . Use these two functions for computing the exact solution $I(C^{-1}(C(x) - t))$. End up with a function `u_exact_variable_c(x, n, c, I)` that returns the value of $I(C^{-1}(C(x) - t_n))$.

k) Make movies showing a comparison of the numerical and exact solutions for the two initial conditions (180) and (28). Choose $\Delta t = \Delta x / \max_{0,L} c$ and the velocity of the medium as

$$1. \quad c(x) = 1 + \epsilon \sin(k\pi x/L), \quad \epsilon < 1,$$

$$2. \quad c(x) = 1 + I(x), \text{ where } I \text{ is given by (180) or (28).}$$

The PDE $u_t + cu_x = 0$ expresses that the initial condition $I(x)$ is transported with velocity $c(x)$.

Filename: `advect1D.py`.

Problem 29: General analytical solution of a 1D damped wave equation

We consider an initial-boundary value problem for the damped wave equation

$$\begin{aligned} u_{tt} + bu_t &= c^2 u_{xx}, & x \in (0, L), \quad t \in (0, T] \\ u(0, t) &= 0, \\ u(L, t) &= 0, \\ u(x, 0) &= I(x), \\ u_t(x, 0) &= V(x). \end{aligned}$$

Here, $b \geq 0$ and c are given constants. The aim is to derive a general analytical solution of this problem. Familiarity with the method of separation of variables for solving PDEs will be assumed.

a) Seek a solution on the form $u(x, t) = X(x)T(t)$. Insert this solution into the PDE and show that it leads to two differential equations for X and T :

$$T'' + bT' + \lambda T = 0, \quad c^2 X'' + \lambda X = 0,$$

with $X(0) = X(L) = 0$ as boundary conditions, and λ as a constant to be determined.

) Show that $X(x)$ is on the form

$$X_n(x) = C_n \sin kx, \quad k = \frac{n\pi}{L}, \quad n = 1, 2, \dots$$

here C_n is an arbitrary constant.

) Under the assumption that $(b/2)^2 < k^2$, show that $T(t)$ is on the form

$$T_n(t) = e^{-\frac{1}{2}bt}(a_n \cos \omega t + b_n \sin \omega t), \quad \omega = \sqrt{k^2 - \frac{1}{4}b^2}, \quad n = 1, 2, \dots$$

the complete solution is then

$$u(x, t) = \sum_{n=1}^{\infty} \sin kx e^{-\frac{1}{2}bt} (A_n \cos \omega t + B_n \sin \omega t),$$

here the constants A_n and B_n must be computed from the initial conditions.

) Derive a formula for A_n from $u(x, 0) = I(x)$ and developing $I(x)$ as a sine Fourier series on $[0, L]$.

) Derive a formula for B_n from $u_t(x, 0) = V(x)$ and developing $V(x)$ as a sine Fourier series on $[0, L]$.

) Calculate A_n and B_n from vibrations of a string where $V(x) = 0$ and

$$I(x) = \begin{cases} ax/x_0, & x < x_0, \\ a(L-x)/(L-x_0), & \text{otherwise} \end{cases} \quad (183)$$

) Implement the series for $u(x, t)$ in a function `u_series(x, t, tol=1E-10)`, here `tol` is a tolerance for truncating the series. Simply sum the terms until $|a_n|$ and $|b_n|$ both are less than `tol`.

) What will change in the derivation of the analytical solution if we have $u(0, t) = u_x(L, t) = 0$ as boundary conditions? And how will you solve the problem with $u(0, t) = 0$ and $u_x(L, t) = 0$?

Filename: `damped_wave1D.pdf`.

Problem 30: General analytical solution of a 2D damped wave equation

Carry out Problem 29 in the 2D case: $u_{tt} + bu_t = c^2(u_{xx} + u_{yy})$, where $(x, y) \in [0, L_x] \times (0, L_y)$. Assume a solution on the form $u(x, y, t) = X(x)Y(y)T(t)$.

Filename: `damped_wave2D.pdf`.

Index

- arithmetic mean, 36
- array slices, 21
- averaging
 - arithmetic, 36
 - geometric, 36
 - harmonic, 36
- C extension module, 69
- C/Python array storage, 74
- column-major ordering, 74
- Courant number, 50
- Cython, 66
- cython -a (Python-C translation in HTML), 68
- declaration of variables in Cython, 67
- Dirichlet conditions, 28
- discrete Fourier transform, 47
- `distutils`, 69
- Fortran array storage, 74
- Fortran subroutine, 71
- Fourier series, 47
- Fourier transform, 47
- geometric mean, 36
- harmonic average, 36
- homogeneous Dirichlet conditions, 28
- homogeneous Neumann conditions, 28
- index set notation, 30, 62
- lambda function (Python), 24
- mesh
 - finite differences, 5
 - mesh function, 6
- Neumann conditions, 28
- `nose` tests, 16
- open boundary condition, 96
- radiation condition, 96
- row-major ordering, 74
- scalar code, 21
- `setup.py`, 69
- slice, 21
- software testing
 - nose, 16
- stability criterion, 50
- stencil
 - 1D wave equation, 6
 - Neumann boundary, 29
- unit testing, 16
- vectorization, 21
- wave equation
 - 1D, 5
 - 1D, analytical properties,
 - 1D, exact numerical solution,
 - 1D, finite difference method,
 - 1D, implementation, 14
 - 1D, stability, 50
 - 2D, implementation, 60
- waves
 - on a string, 5
- wrapper code, 71