Snapshots of modern mathematics
from Oberwolfach

№ 6/2018

Fast Solvers for Highly Oscillatory Problems

Alex Barnett

Waves of diverse types surround us. Sound, light and other waves, such as microwaves, are crucial for speech, mobile phones, and other communication technologies. Elastic waves propagating through the Earth bounce through the Earth's crust and enable us to "see" thousands of kilometres in depth. These propagating waves are highly oscillatory in time and space, and may scatter off obstacles or get "trapped" in cavities. Simulating these phenomena on computers is extremely important. However, the achievable speeds for accurate numerical modelling are low even on large modern computers. Our snapshot will introduce the reader to recent progress in designing algorithms that allow for much more rapid solutions.

1 Time harmonic waves in one and more dimensions

We are all familiar with the waves that spread out in growing circles when a raindrop hits a puddle, or a stone is thrown into a pond. This is an example of a *wave equation* in two dimensions. If x and y are Cartesian coordinates in the horizontal plane, then the height $U(x, y, t)$ of the water surface varies in space

(x, y) and time t. In fact the function U obeys a *partial differential equation* (PDE) [1], which relates its space and time derivatives.

1.1 Time harmonic waves in one dimension

Let us start here with a simpler case: waves in one dimension (1D). You can easily observe these by plucking a long elastic string, such as a washing line, and watching the waves bounce back and forth along it. Starting from Newton's second law, we can derive the displacement function $U(x, t)$, where x is the coordinate along the string, and show that it obeys the PDE

$$\frac{\partial^2 U}{\partial x^2} - \frac{1}{c(x)^2} \frac{\partial^2 U}{\partial t^2} = 0 \,, \qquad \text{(1D wave equation)} \qquad (1)$$

where $c(x)$ is the local speed of waves at the point x, which may vary with position (to understand how the speed of the wave might change in such a fashion, imagine that the rope is heavier in some places than others, thereby slowing down the waves in those locations). This 1D PDE is a good model for the majority of musical instruments [5], including strings, guitars, wind, brass, and pianos[2]. In many scenarios, one cares about a single wave that

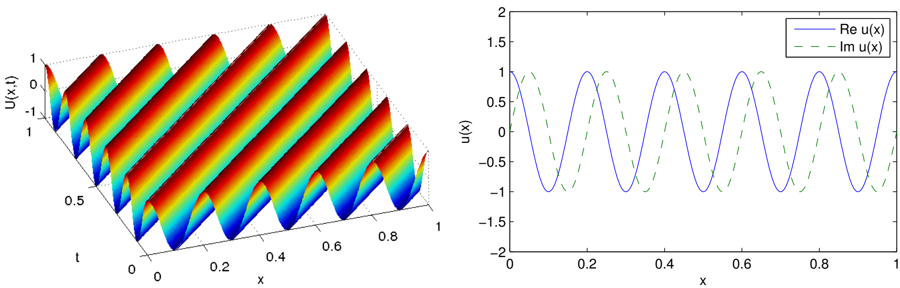

Figure 1: Simple traveling time harmonic wave in 1D with frequency $\omega/2\pi = 5$, and constant wave speed $c(x) = 1$, hence constant wavelength $\lambda = 2\pi/k = 0.2$. The left shows $U(x, t) = \text{Re}[e^{ikx}e^{-i\omega t}]$ as a 3D height plot in space and time. Note that the crests all move to the right with unit speed. The right graphs the real and imaginary parts of the complex-valued spatial function $u(x) = e^{ikx}$.

[1] There are many good books on PDEs; a basic one is [4] and a more mathematical one [7].

[2] For the fascinating reason that the spacing between resonant frequencies of *percussion* instruments is irregular, most of these instruments require instead either the 2D version of the wave equation (1), or wave equations for bending beams or plates (which are 4th order).

oscillates at a given *frequency*: this wave is called "time harmonic," which means that "everything vibrates with the same sinusoidal function of time". Imagine continuously vibrating the washing line, in which case its shape would settle into a steadily repeating pattern after a short time. A general sinusoidal (oscillating) function of time with frequency f (that is, repetition period $1/f$) can be written as $a\cos(\omega t) + b\sin(\omega t)$ for appropriate constants a and b, where $\omega = 2\pi f$. Including the changes, or displacements, in space we can also write $U(x,t) = u_1(x)\cos(\omega t) + u_2(x)\sin(\omega t)$. Mathematicians find it simpler to rewrite this expression using complex numbers as

$$U(x,t) = \mathrm{Re}[u(x)e^{-i\omega t}]\,, \qquad \textit{(definition of time harmonic solution)} \quad (2)$$

where you can check that $u(x) = u_1(x) - iu_2(x)$.

It is now easy to substitute (2) into (1) to obtain the differential equation satisfied by this complex function u, which reads

$$u''(x) + k(x)^2 u(x) = 0 \qquad \textit{(1D Helmholtz equation)} \quad (3)$$

where the known function $k(x) = \omega/c(x)$ is called the *wavenumber*.

Another key property of a time harmonic wave is its *wavelength* $\lambda = 2\pi/k$. The wavelength can be viewed as the "repetition distance", the distance covered by one full oscillation. Larger k (shorter λ) means more oscillations per given length. Remember that we are immersed in a bath of waves. For instance, sounds are transmitted by sound waves of wavelengths λ between about 15 mm and 15 m, while everything we see is light (electromagnetic waves) of wavelengths between 4×10^{-7} m and 7×10^{-7} m.

Notice that $u(x)$ is a function of only one variable. Therefore it is easier to find the solution to (3) than the solution to (1) which depends on two variables. Figure 1 shows an example U and u. Also notice that $u(x) = 0$ is a (very boring, also known as *trivial*) solution to (3).

In more practical scenarios, one adds a "source term" $g(x,t)$ to (3) which specifies the strength of vibrational driving at each point in space, or models waves sent in from far away so that they scatter or reflect. This gives

$$u'' + k(x)^2 u = g(x) \qquad \textit{(1D Helmholtz equation with source)} \quad (4)$$

Finally, one usually cares only about a bounded region of space, such as an interval $\Omega = (0, L)$ (the washing line tied between two points); on its endpoints one needs to impose "boundary conditions" which enforce that waves are only radiating away from the region. Another type of boundary condition – common for washing lines – is that u is pinned down to zero at some point; this is called a Dirichlet condition.

1.2 Time harmonic waves in more dimension

Many more wave phenomena occur in 2D (surface waves) or in 3D (acoustic, electromagnetic, and elastic waves). Waves that travel in more than one dimension are much harder to simulate on computers than waves that travel in 1D, essentially because of all the different directions waves can travel. The generalisation of (3) to more than 1D takes the form

$$\Delta u + k(\mathbf{x})^2 u = g(\mathbf{x}) \qquad \textit{(Helmholtz equation with source)} \qquad (5)$$

where $\mathbf{x} = (x, y)$ in 2D, or (x, y, z) in 3D. Here, the operator Δ is defined as $\Delta = \partial^2/\partial x^2 + \partial^2/\partial y^2$ in 2D, or $\partial^2/\partial x^2 + \partial^2/\partial y^2 + \partial^2/\partial z^2$ in 3D, and is called the *Laplacian operator*. To create a mathematically well-posed problem, we must also impose boundary or radiation conditions on a curve or surface enclosing Ω, the region of interest. This guarantees that the solutions to (5) are "good" solutions.

Figure 2 shows example "scattering and source" problems in 2D.

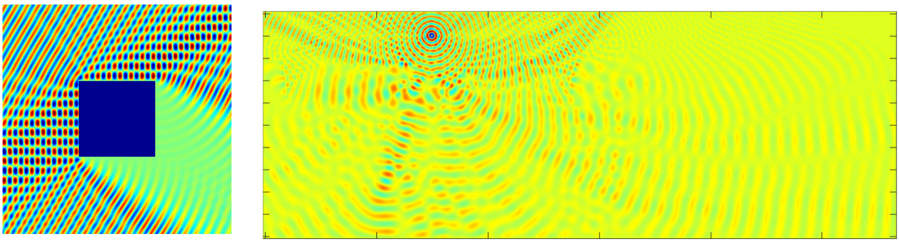

Figure 2: Highly oscillatory problems in 2D. Left: numerical computation of scattering of a "plane wave" (coming from about the 10 o'clock direction) from the unit square with Dirichlet boundary condition ($k = 50$ is constant in space), with error $\epsilon = 10^{-10}$, using the method of [1]. Note that the waves do not cast a hard shadow; they "diffract" around the corners. Right: numerical finite-difference computation of solution where g is a point source in a model for seismic wave propagation ($k(\mathbf{x})$ varies in space, causing bending and reflection of waves), from [9]. In both cases, Re$[u]$ is plotted using a colour scale where red is positive, blue is negative, and yellow or green is around zero.

2 The highly oscillatory case, and real world applications

When might we care about solving the above Helmholtz equations?

Imagine that you are a sound engineer who has been given the 3D geometry of a (small) concert hall (our region Ω) of typical dimension L =15 m, and you are asked to predict how sound emitted by the performers will be heard by each audience member (this will involve various "reflecting boundary conditions" due to the materials of the walls reflecting the sound waves). In this case, since the air density is almost the same everywhere, we can safely assume that $c(x)$, and hence $k(x)$, is constant. However, we note that the ratio of the shortest wavelength we can hear ($\lambda \approx 15$ mm) to the hall size is $L/\lambda \approx 10^3$, a big number. This regime where $L/\lambda \gg 1$ is called *highly oscillatory*. We will see below why solving such a problem accurately is very challenging, even on a big computer.

We have just seen an example of highly oscillatory waves in architectural acoustics. What others are there?[3]

Geology is studied, or oil searched for, using seismic (3D, elastic) waves emitted by earthquakes or by special heavy vibrator trucks.[4] In these cases, the wavelength $k(\mathbf{x})$ varies in unknown ways, and the goal is to reconstruct $k(\mathbf{x})$ given only a large number of detection events of reflected waves reaching the Earth's surface. This approach is called an *inverse problem*, and it is even harder than solving the Helmholtz or elastic equation itself (the *forward problem*). Figures. 2 and 3 include simulated seismic–wave solutions.

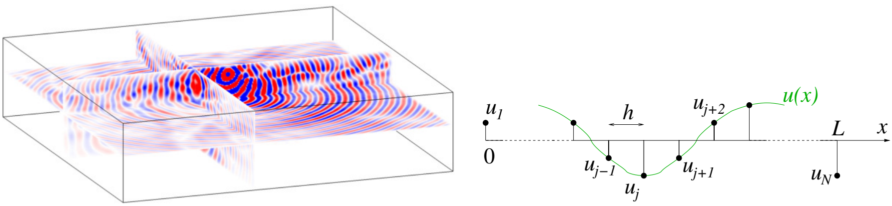

Figure 3: Left: Seismic wave solution u in a 3D domain (acoustic approximation), where g is a point source with frequency f of 6 cycles/second. The domain is 20 km $\times$ 20 km $\times$ 4.7 km. Red is positive and blue negative. Right: setup for 1D finite difference approximation in $[0, L]$. Each number u_j represents u at a gridpoint.

Another case is the following. Given the shape of an aircraft, engineers want to know the directional pattern that electromagnetic waves emitted from radars (electromagnetic waves in the 0.01 m to 1 m range) will scatter or reflect from it. Another goal is to find out how to design the aircraft shape to minimise wave-reflection or sound pollution.

[3] In some of these examples the Helmholtz equation must be replaced by the related Maxwell or elastodynamic equations.

[4] Really!

Other cases are the following. Whales communicate using underwater sound waves that propagate hundreds of kilometres through a depth-dependent $k(\mathbf{x})$. The human body can be safely imaged by ultrasound reflection (another inverse problem), yet to get the best pictures one needs accurate models of wave propagation in the various tissues. Light pulses are guided and switched on and off at high speeds in microscopic devices that enable the "internet backbone". One day, they might enable ultra-fast optical computing. The design of more efficient thin-film solar cells for renewable energy requires modelling light waves in complex geometries (that glass used there has a different wavelength k from air). Finally, at the microscopic scale, quantum mechanics uses wave functions to describe matter. These waves obey the Schrödinger equation, a multi-dimensional complex-valued version of (1) but with a single time derivative.[5]

This range of applications shows the importance of efficient numerical methods for solving highly oscillatory wave problems.

3 Numerical approximate solutions

3.1 Numerical approximate solutions: finite difference discretization

The PDEs presented involve mostly continuous functions. This means that, to describe them *exactly*, one would need to know the values of these functions at an infinite number of points. Of course, computers can handle only a finite, limited, number of real numbers.

An art in numerical analysis is to *approximate* the function u only to some desired accuracy ϵ, using a reasonably small number N of discrete unknowns. This procedure is called "discretization." This often involves relying on the fact that u is a "smooth" function, which means a continuous function without any sharp edges. This allows us to create a more efficient algorithm, thus faster computer solution time. Sometimes it is even possible to *prove* rigorously that the error in the final result is no larger than some maximum value ϵ.

We illustrate this procedure with a "finite difference" discretization of (4) (see [8] for more detail). Let the values u_j be defined by $u_j := u(x_j)$ at points x_j, for $j = 1, \ldots, N$, on a regular grid of spacing $h = L/(N-1)$, as in the right panel of Figure 3. The simplest way to approximate the second derivative in (4) is by

$$u''(x) \approx h^{-2}(u_{j-1} - 2u_j + u_{j+1}). \qquad \textit{(3-point stencil formula)} \qquad (6)$$

Enforcing (4) at each grid point and using (6), we obtain the linear system $A\,\mathbf{u} = \mathbf{g}$, where A is an $N \times N$ matrix with diagonal entries of the form

[5] You can find more on the Schrödinger equation in Snapshot 14/2015 *Quantum diffusion* by Antti Knowles.

$k(x_j) + 2h^{-2}$, entries of the form $-h^{-2}$ adjacent to the diagonal, and zero elsewhere. Here, $\mathbf{g}$ is a vector with entries $g(x_j)$, and $\mathbf{u}$ the unknown vector with entries u_j.

Since A has most entries zero it is called *sparse*. There are direct solution methods for this "tridiagonal" sparse structure that require only $\mathcal{O}(N)$ arithmetic operations.[6] This is *much* faster than the $\mathcal{O}(N^3)$ operations usually needed for the standard "Gaussian elimination".

3.2 Numerical approximate solutions: error estimation

The natural question at this point is: how accurate is this scheme? For simplicity, consider a source-free region (that is, $g = 0$), where $k(x)$ is constant. Then, let's take $u(x) = e^{ikx}$.[7] We know $(e^{ikx})'' = -k^2 e^{ikx}$, but (6) gives instead, using the Taylor series for cosine,

$$\begin{aligned}\frac{e^{i(k-h)x} - 2e^{ikx} + e^{i(k+h)x}}{h^2} &= \frac{2(\cos kh - 1)}{h^2} e^{ikx} \\ &= -\, k^2 e^{ikx} \left(1 - \frac{(kh)^2}{12} + R(kh),\right) \end{aligned} \tag{7}$$

where $R(k\,h)$ is a polynomial that contains terms of the form $(kh)^{2n}$ for $n = 2, 3, 4, \ldots$.

The term $-(kh)^2/12$ is thus the first term in the relative error of this discretization. Therefore, it is clear that $kh \ll 1$ is needed for high accuracy and to be able to discard the correction terms contained in $R(kh)$. In other words, there must be several grid points per wavelength, which means $N \gg L/\lambda$. How can we see this? We have argued that $kh \ll 1$ for high precision approximation methods. Then, using the expressions $h = L/(N-1)$ and $\lambda = 2\,\pi/k$ we see that $kh \ll 1$ implies $N \gg 2\,\pi\, L/\lambda$.

A rigorous error analysis is quite tricky, but we just mention that it would show that this scheme must have an even larger number of grid points per wavelength in order to maintain the same accuracy as k grows (a phenomenon known as the the *pollution effect*).

3.3 Numerical approximate solutions: extensions and generalisations

Researchers have invented much better ways to numerically solve the 1D Helmholtz equation, but the point is that the method described above can be

---

[6] The notation $\mathcal{O}(N)$ means that there are positive constants C and N_0 such that, for all $N > N_0$, the number of operations never exceeds CN.

[7] In fact u is a linear combination of this and its conjugate. The argument still applies.

easily generalised to 2D and 3D, giving the commonly used 5-point and 7-point stencils. The resulting linear systems are characterised by matrices that are sparse but not tridiagonal, therefore the linear systems are harder to solve.

We can say that in 2D, $N \gg (L/\lambda)^2$, and a direct solution takes a number of operations of the order $\mathcal{O}\left(N^{3/2}\right)$. In 3D, one has $N \gg (L/\lambda)^3$, and a direct solution takes $\mathcal{O}\left(N^2\right)$ operations. Returning to the concert hall problem, where $L/\lambda = 10^3$, we see that at least a billion unknowns would be needed, and 10^{18} arithmetic operations (which would take up to many months on a desktop computer!). Fortunately, mathematicians and engineers have developed improved solution methods that are more efficient. Incidentally, this problem size ($N \sim 10^9$) is about the largest that can be currently solved in 3D variable-$k(\mathbf{x})$ seismic applications.

4 Modern progress and open questions

We saw above that highly oscillatory problems can lead to massive linear systems when discretized. The other standard discretization approach is called the *finite element method*, which is useful when the geometry of the domain (concert hall) and/or the variations of the wavenumber $k(\mathbf{x})$ are complicated.

There are plenty of *iterative* methods to solve such systems (that is, these massive linear systems) that rely only on the ability to compute $A\,\mathbf{x}$ for a given vector $\mathbf{x}$. We briefly report on recent progress of the tools that allow us to tackle the problem of simulating useful PDEs.

Domain decomposition. In this procedure we find direct solutions in subregions (areas of the concert hall) where these can be easily obtained, and then use iterative methods to combine the solutions in these different regions (for instance [9]).

Sweeping preconditioners. Here one exploits the fact that, in many applications, waves do not reflect very strongly from the medium [3].

Higher-order finite difference. Another direction is to discretize the PDEs with a higher order of accuracy. In this procedure, the stencil formulas are bigger (that is, there are more terms), and the resulting linear systems become more difficult to handle. However, the accuracy is higher. The error decreases very fast with the increase of number of divisions, if one is working in the regime of applicability of “spectral methods” or “spectral collocation”.

Boundary integral equations. If the wavenumber $k(\mathbf{x})$ is piecewise constant[8], as in the concert hall example or in the case where light transitions between

[8] This means that it is constant in some region of space, such as air, and then it is constant in a neighboring region of space, such as glass.

air and glass, we can obtain analytical predictions of how the waves propagate through each region of constant wavenumber k (such as air-to-glass).
Let us consider a point-like source g and the 3D Helmholtz equation (5) for propagating waves generated by the source. It can be shown that the usual grid-like approach to solving the Helmholtz equation within the domain of interest (concert hall) can be mapped to a scenario where we employ grid-like information on the boundary of the domain (walls of the concert hall). The advantage in this approach is that fewer grid points suffice now, since the boundaries of a surface are at least "one dimension lower" than those of the volume they contain. In other words, the advantage is that it reduces the 3D system (the entire concert hall) to a 2D system (the walls of the concert hall). When the wavenumber k is large (or, equivalently, wavelengths are short), this can be a huge reduction.
Careful design leads to a *well-conditioned* (that is, a linear system whose solution is stable with respect to changes in the data) linear system for which iterative methods converge rapidly. This contrasts with the direct PDE discretizations described above, which are always *ill-conditioned* (a linear system whose solution is unstable with respect to changes in the data). However, the $N \times N$ linear system is now *dense* rather than sparse (that is, it can have non-zero entries everywhere), therefore computing $A\,\mathbf{x}$ from $\mathbf{x}$ would naively take $\mathcal{O}(N^2)$ work.
Amazingly, by clever hierarchical use of the fact that the interaction between distant clusters of points is well approximated by a *low rank* matrix, one may reduce this work to only $\mathcal{O}(N)$, or $\mathcal{O}(N \log N)$. This is called the *fast multipole method*, see [6, 2]. The extension of such ideas to fast direct and "butterfly" solvers for integral equations is an active area of research.
Ray optics approximation. When L/λ is huge, for example $L/\lambda > 100$, one can often get a decent solution using the ray optics approximation. This approximation allows us to view light not as an electromagnetic wave but simply as a set of rays that propagates in a straight line from the source. These rays might reflect against obstacles and refract when they propagate between two different media. This approximation is useful to explain why it is very easy (and not surprising) to predict that, in a concert hall, you will see the performers clearly (unless obstructed by an object, of course). Note that λ is much larger for sound waves than for light waves, which means that the ratio L/λ might not be very big for the typical sound wave emitted, say, by a cello. The ray optics approximation, therefore, does not apply and this explains why sound engineering for concert halls is required and represents an extremely important task.

5 Conclusions and considerations

I end this snapshot with a couple of open questions to think about. Solutions to these questions could revolutionize current understanding of high frequency wave problems.

What is the most efficient way to numerically represent oscillatory solutions in 2D or 3D when the wavenumber k is constant? And which one when the wavenumber $k(\mathbf{x})$ changes?

How can we best employ *distributed* computer architectures (that is, a network of computers which take on parts of a complex task in parallel to each other) to solve wave problems with huge amounts of data?

It is generally believed that, when a wave is trapped in a resonant (reflective) cavity, the complexity of the numerical problem is at least $\mathcal{O}(k^3)$ in 2D, or $\mathcal{O}(k^6)$ in 3D. Can these bounds be beaten for resonant cavities that operate in the highly oscillatory regime?

Image credits

Fig. 2, right panel PhD thesis of Leo Zepeda-Nuñez, MIT, 2015.

Fig. 3, left panel Until summer 2018: Seismic Laboratory for Imaging and Modelling (SLIM), UBC, https://www.slim.eos.ubc.ca/SLIM.Projects.ResearchWebInfo/Modelling/modelling.html, visited on January 7, 2017. From summer 2018: Seismic Laboratory for Imaging and Modelling (SLIM), GATech, https://www.slim.eas.gatech.edu/SLIM.Projects.ResearchWebInfo/Modelling/modelling.html.

All other figures Produced by the author.

References

[1] A. H. Barnett and T. Betcke, *An exponentially convergent non-polynomial finite element method for time-harmonic scattering from polygons*, SIAM Journal on Scientific Computing **32** (2010), no. 3, 1417–1441.

[2] H. Cheng, W. Y. Crutchfield, Z. Gimbutas, L. Greengard, F. Ethridge, J. Huang, V. Rokhlin, N. Yarvin, and J. Zhao, *A wideband fast multipole method for the helmholtz equation in three dimensions*, Journal of Computational Physics **216** (2006), 300–325.

[3] B. Engquist and L. Ying, *Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers*, Journal on Multiscale Modeling and Simulation **9** (2011), no. 2, 686–710.

[4] S. J. Farlow, *Partial differential equations for scientists and engineers*, Dover, 1993.

[5] N. H. Fletcher and T. Rossing, *The physics of musical instruments*, 2nd edition ed., Springer-Verlag, 1998.

[6] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, Journal of Computational Physics **73** (1987), 325–348.

[7] R. B. Guenther and J. W. Lee, *Partial differential equations of mathematical physics and integral equations*, Prentice Hall, 1988.

[8] R. J. LeVeque, *Finite difference methods for ordinary and partial differential equations*, SIAM, 2007.

[9] L. Zepeda-Nuñez and L. Demanet, *Nested domain decomposition with polarized traces for the 2D Helmholtz equation*, arxiv:1510.01831 (2016).

Alex Barnett *is an associate professor of mathematics at Dartmouth College, New Hampshire, and in the Numerical Algorithms Group at the Flatiron Institute, Simons Foundation, New York.*

Mathematical subjects
Numerics and Scientific Computing

Connections to other fields
Engineering and Technology, Physics

DOI
10.14760/SNAP-2018-006-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9 – 11
77709 Oberwolfach
Germany

Director
Gerhard Huisken

Mathematisches Forschungsinstitut Oberwolfach

Member of the Leibniz Association

IMAGINARY
open mathematics