THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Optimised 25-point finite difference schemes for the three-dimensional wave equation

Citation for published version:

Hamilton, B & Bilbao, S 2016, Optimised 25-point finite difference schemes for the three-dimensional wave equation. in *Proceedings of the 22nd International Congress on Acoustics*. International Congress on Acoustics, Buenos Aires, Argentina, 5/09/16.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the 22nd International Congress on Acoustics

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Numerical Techniques (others): Paper ICA2016-561**Optimised 25-point finite difference schemes for the three-dimensional wave equation****Brian Hamilton^(a), Stefan Bilbao^(b)**^(a)Acoustics & Audio Group, University of Edinburgh, UK, brian.hamilton@ed.ac.uk^(b)Acoustics & Audio Group, University of Edinburgh, UK, stefan.bilbao@ed.ac.uk**Abstract:**

Wave-based methods are increasingly viewed as necessary alternatives to geometric methods for room acoustics simulations, as they naturally capture wave phenomena like diffraction and interference. For methods that simulate the three-dimensional wave equation—and thus solve for the entire acoustic field in an enclosed space—computational costs can be high, so efficient algorithms are critical. In terms of computational complexity, finite difference schemes are possibly the simplest such algorithms, but they are known to suffer from numerical dispersion. High-order and optimised schemes can offer improved numerical dispersion, and thus, computationally efficient numerical solutions. In this paper, we consider two families of explicit finite difference schemes for the second-order wave equation in three spatial dimensions, using 25-point stencils on the Cartesian grid. We review known special cases that lead to high-order accuracy in space (and possibly in time), and we present new schemes with optimised stencil coefficients. These schemes provide accurate wave simulation using substantially less memory than the conventional scheme. Simulations are presented to demonstrate the performance of the optimised schemes.

Keywords: finite difference time domain, room acoustics, computational acoustics, numerical dispersion

Optimised 25-point finite difference schemes for the three-dimensional wave equation

1 Introduction

Wave-based numerical methods are important tools for room acoustics modelling and prediction. As opposed to geometric (ray-based) methods [1], wave-based methods offer a complete description of room acoustics, at least in theory [2]. Among the various types of wave-based methods available for room acoustics applications (e.g., [3–6]), finite difference methods that approximate the wave equation over regular grids, or *finite difference time domain* (FDTD) methods [7–9], are popular choices. These methods are relatively simple to implement and they are amenable to parallel computing architectures, such as graphics processing units (GPUs) [10–12]. In comparison to wave-based methods employing non-local spectral operators (e.g., [4, 6]), finite difference methods—operating locally in space—are more adaptable to impedance boundary truncation over irregular domains, possibly through extensions to finite volume techniques [13, 14].

Computational costs are of primary concern when using finite difference methods for room acoustics simulations in three spatial dimensions (3-D). Beyond the minimum grid spacing requirements dictated by sampling considerations, a particular finite difference scheme may require expensive grid refinements (memory costs scale cubically with the inverse of the grid spacing) in order to reduce erroneous numerical dispersion to a tolerable level. The standard (simplest) Cartesian scheme in 3-D, which uses a seven-point spatial stencil, is known to have poor dispersion properties, but various schemes offering improved numerical dispersion have been proposed in the room acoustics literature (e.g., [15–21]) and elsewhere (e.g., [22–26]). These schemes require more operations per point-wise update than the simplest scheme, but they can use coarser grids, which means smaller memory footprints and, generally, fewer operations in total. A conventional approach to designing such schemes is to make use of “leggy” stencils—i.e., stencils reaching only along axial directions, leading to high orders of accuracy in space [23, 25]. Stencils that also reach out along non-axial directions—tending to be compact in space—can lead to high orders of isotropy (enabling the use of frequency warping techniques [27]), high orders of accuracy in both space and time through modified equation methods [24, 25, 28], or optimised dispersion error [29].

In this paper we consider two families of explicit leapfrog schemes employing 25-point stencils, of the leggy type and compact in space. We review known special cases that lead to high-order accuracy in space and/or time [23, 25, 28]. Subsequently, we present new special cases with numerical dispersion optimised for a one-percent error threshold, leading to significant savings in memory and operation costs for accurate wave simulation.

The structure of this paper is as follows. In Section 2 we review basics of finite difference schemes for the wave equation, numerical dispersion, and the simplest Cartesian scheme. In Section 3 we review two families of 25-point schemes and known special cases. The main contributions are found in Section 4, where we present 25-point schemes optimised a one-percent dispersion error tolerance. Numerical experiments comparing the performance of these schemes are provided in Section 5.

2 Background

2.1 The 3-D wave equation

The wave equation in three spatial dimensions serves as a basic model of sound propagation in rooms. This equation reads:

$$\partial_t^2 u = c^2 \Delta u \quad (1)$$

Here, $u = u(t, \mathbf{x})$ is a scalar field that could represent an acoustic velocity potential or sound pressure field [2], $t \geq 0$ and $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$ represent time and spatial position in 3-D, respectively, and c is the speed of sound in air (e.g., 340 m/s). The notation ∂_t denotes a partial derivative with respect to time, $\Delta = \partial_x^2 + \partial_y^2 + \partial_z^2$ is the 3-D Laplacian, and ∂_w denotes a partial derivative with respect to the spatial coordinate $w \in \{x, y, z\}$. Initial conditions are $u(0, \mathbf{x})$ and $\partial_t u(t, \mathbf{x})|_{t=0}$.

It is well-known that the wave equation describes lossless and non-dispersive wave propagation. Considering a plane-wave solution $u = \exp(j(\omega t + \mathbf{k} \cdot \mathbf{x}))$, where $j = \sqrt{-1}$, $\omega \in \mathbb{R}$ is the temporal angular frequency in rad/s and $\mathbf{k} = (k_x, k_y, k_z) \in \mathbb{R}^3$ is the wave vector with $k = |\mathbf{k}|$ the wavenumber in rad/m, the dispersion relation and phase velocity for the wave equation are, respectively:

$$\omega = \pm c|\mathbf{k}| \quad \Rightarrow \quad v_p := \omega/|\mathbf{k}| = \pm c \quad (2)$$

In other words, plane waves travel with the speed c regardless of their frequency.

2.2 Shift and difference operators

In order to discretise this equation on a Cartesian grid, let $u_i^n \approx u(nT, iX)$ represent a discrete sound field at time $t = nT$, $n \in \mathbb{Z}, n \geq 0$ and position iX , $i = (i_x, i_y, i_z) \in \mathbb{Z}^3$, where X is the Cartesian grid spacing and T is the time-step ($F_s = 1/T$ is the sample rate in Hz). Introduce the shift operators acting on u_i^n :

$$e_{t+} u_i^n := u_i^{n+1}, \quad e_{t-} u_i^n := u_i^{n-1}, \quad e_{w+} u_i^n := u_{i+\hat{e}_w}^n, \quad e_{w-} u_i^n := u_{i-\hat{e}_w}^n \quad (3)$$

where $w \in \{x, y, z\}$ and $\hat{e}_w$ are the standard unit vectors in $\mathbb{R}^3$. We also define the difference operators:

$$\delta_t^2 := e_{t+} - 2 + e_{t-}, \quad \delta_w^2 := e_{w+} - 2 + e_{w-} \quad \Rightarrow \quad \delta_t^2 = T^2 \partial_t^2 + O(T^4), \quad \delta_w^2 = X^2 \partial_w^2 + O(X^4) \quad (4)$$

2.3 Finite difference schemes

The finite difference schemes to follow in this study can be written in the form:

$$\delta_t^2 u_i^n = \lambda^2 (\delta_x^2 + \delta_y^2 + \delta_z^2) u_i^n + \lambda^2 P(\delta_x^2, \delta_y^2, \delta_z^2) u_i^n \quad (5)$$

where $\lambda = cT/X$ is the Courant number and $P = P(\delta_x^2, \delta_y^2, \delta_z^2)$ is a trinomial in $\delta_x^2, \delta_y^2, \delta_z^2$ that will be specified for each scheme. This general scheme is consistent with the 3-D wave equation provided that $P = O(X^4)$ (i.e., P is without zeroth- or first-order terms), which also leads to $O(T^2)$ and $O(X^2)$ accuracy in time and space for the scheme (at least). These schemes are explicit and can be operated as a two-step recursion, requiring two states to be stored in memory. The simplest scheme [30] is a special case with $P = 0$ and uses a seven-point stencil; see Fig. 1(a).

2.4 Numerical dispersion relation

In order to analyse the numerical dispersion of a given scheme, we consider the plane-wave ansatz: $u_i^n = \exp(j(\omega nT + \mathbf{k} \cdot \mathbf{i}X))$, where $\omega \in [-\pi/T, \pi/T]$ and $\mathbf{k} \in [-\pi/X, \pi/X]^3$. The operators δ_t^2 and δ_w^2 act on the plane-wave ansatz as the following multipliers (Fourier symbols):

$$\delta_t^2 \Rightarrow -4\hat{s}_t, \quad \delta_w^2 \Rightarrow -4\hat{s}_w \quad (6)$$

where $\hat{s}_t = \sin^2(\omega T/2)$, and $\hat{s}_w = \sin^2((\mathbf{k}X) \cdot \hat{\mathbf{e}}_w/2)$ for $w \in \{x, y, z\}$. Scheme (5) then transforms to:

$$\hat{s}_t = \lambda^2(\hat{s}_x + \hat{s}_y + \hat{s}_z) - (\lambda^2/4)\hat{P}(-4\hat{s}_x, -4\hat{s}_y, -4\hat{s}_z) \quad (7)$$

where $\hat{P} = \hat{P}(-4\hat{s}_x, -4\hat{s}_y, -4\hat{s}_z)$ is a frequency-domain analogue of P . Eq. (7) is essentially the numerical dispersion relation of the scheme, but we can rewrite it in terms of normalised frequencies $\omega_T := \omega T$ and normalised vector-wavenumbers $\mathbf{k}_X := \mathbf{k}X$, leading to:

$$\omega_T(\mathbf{k}_X) = 2 \arcsin(\pm \lambda \sqrt{F}), \quad F := \hat{s}_x + \hat{s}_y + \hat{s}_z - \hat{P}/4 \quad (8)$$

2.5 von Neumann stability conditions

A numerical stability condition can be obtained using von Neumann analysis [30]. This leads the following conditions for numerical stability:

$$0 \leq \lambda^2 F \leq 1 \quad (9)$$

For example, for $\hat{P} = 0$ we obtain the well-known stability condition $\lambda \leq \sqrt{1/3} \approx 0.577$ for the simplest scheme [30]. For a fixed grid spacing, the second condition above places a restriction on the time-step: $T \leq \lambda_{\max} X/c$, where $\lambda_{\max}$ is the largest λ for which (9) is satisfied. It is often desirable to choose $T = \lambda_{\max} X/c$ in order to minimise the number of time-steps to compute for a given duration of time, provided that temporal sampling considerations are also satisfied.

2.6 Numerical phase velocity

Of particular interest in this study is the relative numerical phase velocity of the scheme, which we define as:

$$\hat{v}_p(\mathbf{k}_X) := \frac{\omega_T(\mathbf{k}_X)}{\lambda |\mathbf{k}_X|} \quad (10)$$

Considering only $\omega_T \geq 0$, $\hat{v}_p(\mathbf{k}_X)$ is ideally unity. In general for a discrete scheme, $\hat{v}_p(\mathbf{k}_X) \neq 1$ but $\hat{v}_p(\mathbf{k}_X) \rightarrow 1$ as $|\mathbf{k}_X| \rightarrow 0$ by consistency (and stability). Fig. 1(b) shows the relative phase velocity for the simplest scheme, as a function of wavenumber along three directions: axial ($k_x = |\mathbf{k}|, k_y = k_z = 0$), side-diagonal ($k_x = k_y, k_z = 0$), and space-diagonal ($k_x = k_y = k_z$); these tend to be the extreme-case directions and by symmetry we only need to consider $\mathbf{k} \in [0, \pi]^3$.

In order to compare the various schemes to follow, it helps to identify where specific dispersion error thresholds are met in the range of discrete wavenumbers. For example, the phase velocity in the simplest scheme deviates from unity by at most 1% up to $k_X \approx 0.190\pi$, and by at most 2% up to $k_X \approx 0.268\pi$, where $k_X := |\mathbf{k}_X|$. To express this more formally, let us define an L_∞ -error

Figure 1: (a) 7-point stencil in simplest scheme (right: view from above); (b) relative phase velocity for simplest scheme with $\lambda = \lambda_{\max}$, along axial, side-diagonal and space-diagonal directions (close-up on right).

for the relative phase velocity:

$$\epsilon(k'_X) := \max_{k_X \in [0, k'_X]} |1 - \hat{v}_p| \quad (11)$$

where $0 < k'_X < \pi$, and let us also define the following two wavenumber quantities:

$$k_{1\%,X} := \max_{k'_X} \text{ such that } \epsilon \leq 0.01, \quad k_{2\%,X} := \max_{k'_X} \text{ such that } \epsilon \leq 0.02 \quad (12)$$

Then for the simplest scheme we have $k_{1\%,X} \approx 0.190\pi$ and $k_{2\%,X} \approx 0.268\pi$. For convenience, we also call $k_{1\%,X}$ the “1%-wavenumber”, and $k_{2\%,X}$ the “2%-wavenumber”. For the grid resolution, this means that there must be at least $2\pi/(k_{1\%,X}) \approx 10.52$ points per minimum wavelength of interest (PPW) if dispersion errors are to be kept below one-percent. For a two-percent error, the grid can be slightly coarser with $2\pi/(k_{2\%,X}) \approx 7.47$ PPW. The 25-point schemes to follow can meet these error tolerances with fewer PPW, leading to substantial savings in memory and total operations.

3 High-order 25-point schemes

In this section we review two 25-point schemes that are able achieve a high order of accuracy in space, and also possibly time [23,25,28,31], thereby offering improved numerical dispersion over the simplest scheme.

3.1 A 25-point leggy scheme

A standard approach to improving multidimensional finite difference schemes is to use high-order spatial differencing along each coordinate direction. In 3-D, this leads to $(6M+1)$ -point “leggy” stencils that reach out M points along each axial direction [23,25]. For $M=4$ we have a 25-point leggy scheme (see Fig. 2(a)), and it is able to achieve an eighth-order accuracy in space. This leggy scheme is of the form (5) with:

$$P = \sum_{m=2}^4 b_m (\delta_x^{2m} + \delta_y^{2m} + \delta_z^{2m}) \quad (13)$$

where $\delta_w^{2m} := (\delta_w^2)^m$ for $w \in \{x, y, z\}$. For $O(X^8)$ spatial accuracy, the coefficients $b_2, \dots, b_4$ are [31, 32]:

$$b_2 = -1/12, \quad b_3 = 1/90, \quad b_4 = -1/560 \quad (14)$$

The temporal accuracy of this scheme remains $O(T^2)$, so we will call this scheme “(2,8)-accurate” (the simplest scheme is (2,2)-accurate). It is important to remark that while the spatial accuracy is $O(X^8)$, in the limit of $X \rightarrow 0$ with λ fixed the temporal error dominates and consequently, $\hat{v}_p(\mathbf{k}_X) \rightarrow 1$ with $O(T^2) = O(X^2)$. Strictly speaking then, this scheme is only second-order accurate for the wave equation.

Using the plane-wave ansatz, we obtain a dispersion relation of the form (8) and we have for F :

$$F = \sum_{m=1}^4 \gamma_m (\hat{s}_x^m + \hat{s}_y^m + \hat{s}_z^m), \quad \gamma_m = (-4)^{m-1} b_m \quad (15)$$

with $b_1 = 1$. Since $\gamma_m \geq 0$, the conditions (9) are satisfied for $\lambda \leq \sqrt{105/512} \approx 0.452$. As is well-known, it can help to choose λ significantly lower than $\lambda_{\max}$ in such schemes in order to balance temporal errors with spatial errors, improving the dispersion error [23,25]. This will be treated as an optimisation and considered in Section 4.1.

The relative phase velocity of this scheme is shown in Fig. 2(b) for $\lambda = \lambda_{\max}$ (a naive, but safe choice). For this scheme we have $k_{1\%,X} \approx 0.340\pi$ and $k_{2\%,X} \approx 0.475\pi$. In other words, this scheme needs approximately 5.89 PPW and 4.21 PPW for at most one- and two-percent errors, respectively.

Figure 2: (a) 25-point leggy stencil (right: view from above); (b) relative phase velocity for (2,8)-scheme with $\lambda = \lambda_{\max}$, along axial, side-diagonal and space-diagonal directions (close-up on right).

3.2 A 25-point compact scheme

Next we consider a 25-point stencil that is compact in space (with respect to the ℓ_1 -norm) rather than leggy; see Fig. 3(a). Using this stencil we can arrive at a fourth-order accuracy in both space and time [25,28]. This 25-point compact scheme can be written in the form (5) with:

$$P = a (\delta_x^2 \delta_y^2 + \delta_x^2 \delta_z^2 + \delta_y^2 \delta_z^2) + b (\delta_x^4 + \delta_y^4 + \delta_z^4) \quad (16)$$

and a, b are free parameters. Fourth-order accuracy in both space and time can be achieved, through modified equation methods, with the choices $a = \lambda^2/6$ and $b = (\lambda^2 - 1)/12$ [28].

For the dispersion relation of this scheme, F is:

$$F = \hat{s}_x + \hat{s}_y + \hat{s}_z - 4a (\hat{s}_x \hat{s}_y + \hat{s}_x \hat{s}_z + \hat{s}_y \hat{s}_z) - 4b (\hat{s}_x^2 + \hat{s}_y^2 + \hat{s}_z^2) \quad (17)$$

It has been shown that the stability condition in this (4,4)-accurate scheme is $\lambda \leq \sqrt{1/3} \approx$

0.577 [24]. It is optimal to choose $\lambda = \lambda_{\max}$ in this scheme since temporal and spatial errors are balanced. The relative phase velocity is then shown in Fig. 3(b). For this scheme we have $k_{1\%,X} \approx 0.424\pi$ and $k_{2\%,X} \approx 0.507\pi$. In other words, this scheme needs approximately 4.72 PPW and 3.94 PPW for at most one- and two-percent errors, respectively.

Figure 3: (a) 25-point compact stencil (right: view from above); (b) relative phase velocity for (4,4)-scheme with $\lambda = \lambda_{\max}$, along axial, side-diagonal and space-diagonal directions (close-up on right).

4 Optimised 25-point schemes

So far we have reviewed known schemes [23,25,28,31]. At this point we introduce the main contributions of this paper, which are three variations on the 25-point schemes incorporating further optimisations of numerical dispersion.

4.1 25-point leggy scheme with optimised Courant number

First, we consider the 25-point leggy scheme with high-order spatial accuracy ($b_2, \dots, b_4$ given by (14)), but now with the Courant number chosen away from the stability limit, to better balance temporal errors with spatial errors [23,25]. In particular, we focus on maximising the 1%-wavenumber, $k_{1\%,X}$. Carrying out a brute-force search of $\lambda \in [0, \lambda_{\max}]$, at $\lambda \approx 0.240$ we find a local maximum of $k_{1\%,X} \approx 0.655\pi$ (3.05 PPW for one-percent error); this is approximately 40 times more memory efficient than the simplest scheme for a one-percent dispersion error tolerance. The dispersion error for this scheme is shown in Fig. 4(a),

Figure 4: Relative phase velocity along axial, side-diagonal and space-diagonal directions for optimised 25-point schemes.

4.2 25-point schemes with optimised stencil weights and Courant numbers

Next we consider the two 25-point schemes as expressed by (13) and (16), but leaving b_2, b_3, b_4 and a, b (and λ) as free parameters to optimise, neglecting the constraints for high-order accuracy. We focus on maximising $k_{1\%,X}$, but in this case, we rescale this wavenumber quantity in order to take into account costs associated to the choice of λ , which affects the sample rate; i.e., we consider not only memory costs, but total operational costs (in space and time). This leads to a search for the free parameters b_2, b_3, b_4, λ or a, b, λ such that the rescaled 1%-wavenumber $\lambda^{1/4} k_{1\%,X}$ is maximised under the constraints (9).

This optimisation problem is generally non-linear due to the stability conditions (9). At this point in time, it is unknown if a solution to this optimisation problem exists or is unique, so we resort to standard optimisation tools in order to look for “optimised” schemes (though not necessarily “optimal”). Similarly to the approach in previous work [33], we use Matlab’s `fminsearch` (Nelder-Mead simplex algorithm) supplemented with non-linear constraints [34], starting from the parameters given in Section 3. Figs. 4(b) and 4(c) show the relative phase velocities of the resulting schemes. The optimised leggy scheme in Fig. 4(b) has $\lambda = 0.3979, b_4 = 0.0044125, b_3 = 0.0445444, b_2 = -0.019683$, and $k_{1\%,X} \approx 0.633\pi$ (3.16 PPW for one-percent error). The optimised compact scheme in Fig. 4(c) has $\lambda = 0.608, a = 0.10375, b = -0.0785$, and $k_{1\%,X} \approx 0.626\pi$ (3.19 PPW for one-percent error), leading to nearly 125 times fewer point-wise updates and nearly 40 times fewer total floating-point operations for a one-percent dispersion error tolerance. These wavenumber quantities are summarised in Table 1.

Table 1: Properties of various schemes, including the figure where phase velocity is displayed, order of accuracy (OoA) in time and space, Courant number λ , the 1%-wavenumber ($k_{1\%,X}$), the number of PPW required for a max one-percent error in phase velocity ($2\pi/k_{1\%,X}$), and the rescaled 1%-wavenumber, $\lambda^{1/4} k_{1\%,X}$.

stencil size	stencil type	from Section	see Figure	OoA	λ (approx.)	$k_{1\%,X}$	$2\pi/k_{1\%,X}$	$\lambda^{1/4} k_{1\%,X}$
7-point	simplest	2	1(b)	(2, 2)	0.577	0.190π	10.52 PPW	0.166π
25-point	leggy	3.1	2(b)	(2, 8)	0.452	0.340π	5.89 PPW	0.279π
25-point	compact	3.2	3(b)	(4, 4)	0.577	0.424π	4.72 PPW	0.369π
25-point	leggy	4.1	4(a)	(2, 8)	0.240	0.655π	3.05 PPW	0.458π
25-point	leggy	4.2	4(b)	(2, 2)	0.397	0.633π	3.16 PPW	0.502π
25-point	compact	4.2	4(c)	(2, 2)	0.608	0.626π	3.19 PPW	0.553π

5 Numerical example

In this section we present a numerical example to compare the performance of the 25-point schemes under equivalent spatiotemporal grid densities $((X^3T)^{-1})$. For this example we consider a free-space initial value problem where $u(0, \mathbf{x})$ is a Ricker wavelet, i.e., the negative Laplacian of a spatial Gaussian, with variance $\sigma^2 = 0.05 \text{ m}^2$, and $\partial_t u(0, \mathbf{x})|_{t=0} = 0$. The approximate solutions are computed with $X = \lambda^{-1/4} \tilde{X}$ where $\tilde{X} = 0.15 \text{ m}$ and $c = 340 \text{ m/s}$ are fixed across schemes, and λ is as listed in Table 1. This leads to each scheme having a spatiotemporal density of $(X^3T)^{-1} = c\tilde{X}^{-4}$, and for a given volume and duration each 25-point scheme uses the same number of total floating point instructions when fused-multiply add (FMA) operations are available, such as on GPU devices [12]. The size of the domain and duration is chosen such that boundary truncation does not play a role. An ‘exact’ solution is computed using the fourth-order scheme with the spatial step sufficiently small for dispersion error to be virtually negligible in this problem; this is used to compute a residual at three outputs locations at a distance of $\sqrt{3} \text{ m}$ from the origin, along axial, side-diagonal and diagonal directions. Spline interpolation is used on the outputs

Figure 5: Outputs of simulations for various schemes read at a fixed distance along axial (blue), side-diagonal (green) and diagonal directions (red). (a) simplest scheme; (b) 25-point leggy (2,8)-accurate scheme with $\lambda = \lambda_{\max}$; (c) 25-point compact (4,4)-accurate scheme with $\lambda = \lambda_{\max}$; (d) 25-point leggy (2,8)-accurate scheme with λ optimised; (e) 25-point leggy scheme with stencil weights and λ optimised; (f) 25-point compact scheme with stencil weights and λ optimised.

(in space and time) to calculate residuals. The exact and numerical outputs are displayed in Fig. 5. The energy in each residual is calculated relative to the energy in the exact solution and the worst-case (out of the three directions) for each scheme appears in Table 2 (denoted by RE_{ℓ_2}). An analogous relative error is also calculated using the ℓ_∞ -norm and displayed in Table 2 (denoted by RE_{ℓ_∞}). It can be seen that the schemes with higher $\lambda^{1/4}k_{1\%}$ generally demonstrate the smallest errors, and the compact optimised 25-point scheme performs best in this example.

Table 2: Relative errors for each simulation output (worst-case among three directions) and grid spacings used in order to normalise spatiotemporal densities (with λ as in Table 1).

stencil size	stencil type	from Section	see Figure	X (approx. in m)	RE_{ℓ_2}	RE_{ℓ_∞}
7-point	simplest	2	5(a)	0.1721	-2.61 dB	-2.70 dB
25-point	leggy	3.1	5(b)	0.1829	-7.63 dB	-5.77 dB
25-point	compact	3.2	5(c)	0.1721	-7.56 dB	-7.12 dB
25-point	leggy	4.1	5(d)	0.2143	-8.01 dB	-8.90 dB
25-point	leggy	4.2	5(e)	0.1889	-11.04 dB	-10.57 dB
25-point	compact	4.2	5(f)	0.1699	-14.56 dB	-14.36 dB

6 Conclusions and Final Remarks

In this study we considered explicit finite difference schemes for the 3-D wave equation, using leggy and compact-in-space 25-point stencils. High-order (in space and/or time) special cases were analysed in terms of numerical dispersion. New 25-point schemes were presented with Courant numbers and stencil weights optimised for a one-percent dispersion error threshold. Whereas the conventional seven-point scheme requires more than 10 PPW for a max one-percent dispersion error, the optimised schemes achieve similar accuracy using as few as 3 PPW, leading to significant savings in memory and total operational costs. A numerical example was presented to compare the various 25-point schemes under equal computational costs, and the optimised compact scheme produced the most accurate results. The formulation of impedance boundary conditions for these schemes will be the subject of future work.

Acknowledgements This work was supported by the European Research Council under grant StG-2011-279068-NESS, and by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] L. Savioja and U. P. Svensson, "Overview of geometrical room acoustic modeling techniques," *JASA*, vol. 138, no. 2, pp. 708–730, 2015.
- [2] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of acoustics*. John Wiley & Sons, 4th ed., 2000.
- [3] Y. W. Lam and J. Hargreaves, "Time domain modelling of room acoustics," in *Acoustics 2012*, (Nantes, France), 2012.
- [4] R. Mehra, N. Raghuvanshi, L. Savioja, M. C. Lin, and D. Manocha, "An efficient GPU-based time domain solver for the acoustic wave equation," *Applied Acoustics*, vol. 73, no. 2, pp. 83–94, 2012.
- [5] T. Sakuma, S. Sakamoto, and T. Otsuru, *Computational simulation in architectural and environmental acoustics*. Springer, 2014.
- [6] M. Hornikx, T. Krijnen, and L. van Harten, "openPSTD: The open source pseudospectral time-domain method for acoustic propagation," *Computer Physics Communications*, vol. 203, pp. 298–308, 2016.
- [7] O. Chiba, T. Kashiwa, H. Shimoda, S. Kagami, and I. Fukai, "Analysis of sound fields in three dimensional space by the time-dependent finite-difference method based on the leap frog algorithm," *J. Acoust. Soc. Jpn.(J)*, vol. 49, pp. 551–562, 1993.
- [8] D. Botteldooren, "Finite-difference time-domain simulation of low-frequency room acoustic problems," *JASA*, vol. 98, pp. 3302–3308, 1995.
- [9] L. Savioja, M. Karjalainen, and T. Takala, "DSP formulation of a finite difference method for room acoustics simulation," in *IEEE Nordic Signal Processing Symposium*, pp. 455–458, 1996.
- [10] J. Saarelma and L. Savioja, "An open source finite difference time-domain solver for room acoustics using graphics processing units," in *Proceedings of Forum Acusticum*, (Krakow, Poland), 2014.
- [11] C. Webb, *Parallel computation techniques for virtual acoustics and physical modelling synthesis*. PhD diss., University of Edinburgh, 2014.
- [12] B. Hamilton, C. J. Webb, A. Gray, and S. Bilbao, "Large stencil operations for GPU-based 3-D acoustics simulations," in *Proc. Digital Audio Effects (DAFx)*, (Trondheim, Norway), Dec. 2015.
- [13] S. Bilbao, B. Hamilton, J. Botts, and L. Savioja, "Finite volume time domain room acoustics simulation under general impedance boundary conditions," *IEEE/ACM Trans. Audio, Speech, and Language Processing*, vol. 24, no. 1, pp. 161–173, 2016.
- [14] S. Bilbao and B. Hamilton, "Finite volume modeling of viscothermal losses and frequency-dependent boundaries in room acoustics simulations," in *AES 60th International Conference*, (Leuven, Belgium), Feb. 2016.
- [15] L. Savioja, "Improving the three-dimensional digital waveguide mesh by interpolation," *Proc. Nordic Acoustical Meeting*, pp. 265–268, 1998.
- [16] S. Bilbao, *Wave and scattering methods for numerical simulation*, pp. 318–321. Wiley, 2004.
- [17] S. Sakamoto, "Phase-error analysis of high-order finite difference time domain scheme and its influence on calculation results of impulse response in closed sound field," *Acoustical science and technology*, vol. 28, no. 5, pp. 295–309, 2007.
- [18] K. Kowalczyk and M. van Walstijn, "Room acoustics simulation using 3-D compact explicit FDTD schemes," *IEEE Trans. Audio, Speech, and Language Processing*, vol. 19, no. 1, pp. 34–46, 2011.
- [19] B. Hamilton and C. J. Webb, "Room acoustics modelling using GPU-accelerated finite difference and finite volume methods on a face-centered cubic grid," in *Proc. Digital Audio Effects (DAFx)*, (Maynooth, Ireland), pp. 336–343, Sept. 2013.
- [20] B. Hamilton, S. Bilbao, and C. J. Webb, "Revisiting implicit finite difference schemes for 3-D room acoustics simulations on GPU," in *Proc. Digital Audio Effects (DAFx)*, (Erlangen, Germany), 2014.
- [21] J. van Mourik and D. Murphy, "Explicit higher-order FDTD schemes for 3D room acoustic simulation," *IEEE/ACM Trans. Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 2003–2011, 2014.
- [22] G. Fairweather and A. Mitchell, "A high accuracy alternating direction method for the wave equation," *IMA J. Applied Math.*, vol. 1, no. 4, pp. 309–316, 1965.
- [23] M. A. Dablain, "The application of high-order differencing to the scalar wave equation," *Geophysics*, vol. 51, no. 1, pp. 54–66, 1986.
- [24] J. Tuomela, "Fourth-order schemes for the wave equation, Maxwell equations, and linearized elastodynamic equations," *Numerical Methods for PDEs*, vol. 10, no. 1, pp. 33–63, 1994.
- [25] G. Cohen, *Higher-order numerical methods for transient wave equations*. Springer-Verlag, 2002.
- [26] X.-H. Cai, Y. Liu, Z.-M. Ren, J.-M. Wang, Z.-D. Chen, K.-Y. Chen, and C. Wang, "Three-dimensional acoustic wave equation modeling based on the optimal finite-difference scheme," *Applied Geophysics*, vol. 12, no. 3, pp. 409–420, 2015.
- [27] L. Savioja and V. Välimäki, "Interpolated rectangular 3-D digital waveguide mesh algorithms with frequency warping," *IEEE Trans. Speech and Audio Processing*, vol. 11, no. 6, pp. 783–790, 2003.
- [28] J. Tuomela, "On the construction of arbitrary order schemes for the many dimensional wave equation," *BIT Num. Math.*, vol. 36, no. 1, pp. 158–165, 1996.
- [29] S. Bilbao, "Optimized FDTD schemes for 3-D acoustic wave propagation," *IEEE Trans. Audio, Speech, and Language Processing*, vol. 20, no. 5, pp. 1658–1663, 2012.
- [30] G. E. Forsythe and W. R. Wasow, *Finite-difference methods for partial differential equations*, pp. 378–382. New York: Wiley, 1960.
- [31] L. Anné, P. Joly, and Q. H. Tran, "Construction and analysis of higher order finite difference schemes for the 1D wave equation," *Computational Geosciences*, vol. 4, no. 3, pp. 207–249, 2000.
- [32] J. Tuomela, "A note on high order schemes for the one dimensional wave equation," *BIT Num. Math.*, vol. 35, no. 3, pp. 394–405, 1995.
- [33] B. Hamilton and S. Bilbao, "Fourth-order and optimised finite difference schemes for the 2-D wave equation," in *Proc. Digital Audio Effects (DAFx)*, (Maynooth, Ireland), 2013.
- [34] J. D'Errico, "Bound constrained optimization using fminsearch." <http://mathworks.com/matlabcentral/fileexchange/8277>, 2012. Accessed: 2016-04-28.