

Improved Finite Difference Schemes for a 3-D Viscothermal Wave Equation on a GPU

Brian Hamilton

Acoustics and Audio Group, University of Edinburgh, Edinburgh, UK.

brian.hamilton@ed.ac.uk

Stefan Bilbao

Acoustics and Audio Group, University of Edinburgh, Edinburgh, UK.

Craig J. Webb

Oxford e-Research Centre, University of Oxford, Oxford, UK.

Summary

Viscothermal effects in air lead to a damping of high frequencies over time. Such effects cannot be neglected in large-scale room acoustics simulations for the full audible bandwidth. In this study, full-bandwidth room acoustics is modelled using a variant of the three-dimensional wave equation including viscothermal losses in air following from a simplification of the Navier-Stokes equations suitable for room acoustics applications. The model equation is numerically solved using time domain finite difference methods. A three-step parameterised finite difference scheme is proposed to model T_{60} decay times as a function of frequency more accurately than two-step schemes. Timing results from parallelised implementations on a graphics processing unit (GPU) device are presented.

PACS no. 43.55.+p

1. Introduction

Finite difference (FD) methods have been used for wave-based acoustics for some time [1, 2]. In the context of room acoustics, they are often called “finite difference time domain” (FDTD) methods [3], after the electromagnetics literature [4], although preexisting second-order (single-variable) wave equation formulations [5, 6, 7] are preferred to “FDTD” staggered formulations for simplicity and computational costs [8, 9].

It is well-known that the standard 3-D wave equation, or the equivalent first-order hyperbolic system, is only a simplified model to describe sound propagation in air [10, 11, 12]. The inclusion of sound attenuation effects can offer more convincing artificial reverberation from physical models of rooms [13, 14]. The classic model for sound attenuation is due to Stokes [15, 16], which lead to the further inclusion of thermal, bulk viscosity, and relaxation effects [10, 17, 11, 18, 12].

The purpose of this paper will be to improve upon the two-step FD scheme presented in [13] for a viscothermal 3-D wave equation, which is suitable for large-scale room acoustics simulations on graphics processing units (GPUs) [19, 9]. Related work includes

multistep FD schemes for nonlinear viscous wave equations in 1-D [20, 21], FD modelling of viscothermal and relaxation effects for Burgers’ equation in 1-D and 2-D [22], and 2-D FDTD modelling of relaxation effects [23] and sound attenuation in the upper atmosphere [24].

This paper is organised as follows. The simplified model for sound propagation in air with viscothermal effects is presented in Section 2, along with physical constants that are suitable for indoor sound propagation. Finite difference schemes and stability conditions are presented in Section 3. The numerical dispersion and dissipation of these schemes are analysed in Section 4. Simplified tests are conducted in Section 5 in order to verify the numerical T_{60} decay times predicted by frequency domain analysis. Timing results for various GPU implementations are presented in Section 6.

2. Models of Viscothermal Wave Propagation

A very general linear model of acoustic wave propagation follows from the linearization of the Navier-Stokes equations, including an assumption of Fourier heat conduction [10, 25]. It may be written as a system of three coupled partial differential equations in p , $\mathbf{v}$

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.

and τ , representing, respectively, pressure deviation from atmospheric pressure, the irrotational part of the velocity vector field $\mathbf{v}$, and temperature deviation τ , which are all functions of time t and coordinates $\mathbf{x} \in \mathbb{R}^3$, as:

$$\partial_t (p - d\tau) = -\frac{\rho c^2}{\gamma} \nabla \cdot \mathbf{v} \quad (1a)$$

$$\beta (\partial_t - cl_h \Delta) \tau = \frac{\gamma - 1}{\gamma} \partial_t p \quad (1b)$$

$$\rho (\partial_t - cl_v \Delta) \mathbf{v} = -\nabla p \quad (1c)$$

Here, the operator ∂_t represents partial differentiation with respect to time t , and ∇ , $\nabla \cdot$ and Δ are the three dimensional gradient, divergence and Laplacian operators, respectively. The various constants which appear here are: density ρ , wave speed c , ratio of specific heats γ , pressure increase per unit temperature increase at constant density β , and the two characteristic lengths l_h and l_v , defined as

$$l_h = \frac{\kappa}{\rho c C_p} \quad l_v = \frac{4\mu/3 + \mu_B}{\rho c} \quad (2)$$

where here, κ is the coefficient of thermal conductivity, C_p is heat coefficient at constant pressure, for a constant unit mass, and μ and μ_B are viscosity and bulk viscosity, respectively. Not included here is a separate uncoupled equation in the rotational part of the velocity vector field.

Because $\mathbf{v}$ is irrotational, it may be written as $\mathbf{v} = -\nabla \Psi$, for a scalar velocity potential Ψ , and the system (1) may be rewritten as

$$\partial_t (p - \beta \tau) = \frac{\rho c^2}{\gamma} \Delta \Psi \quad (3a)$$

$$\beta (\partial_t - cl_h \Delta) \tau = \frac{\gamma - 1}{\gamma} \partial_t p \quad (3b)$$

$$\rho (\partial_t - cl_v \Delta) \Psi = p \quad (3c)$$

In the interest of examining the frequency domain behaviour of this model, as well as, ultimately, a simplified model and resulting numerical schemes, consider the characteristic equation for system (3) resulting from the assumption of wave-like solutions of the form $e^{st+\boldsymbol{\xi} \cdot \mathbf{x}}$, where here, $s = \sigma + j\omega$ is a complex frequency domain variable, and $\boldsymbol{\xi} \in \mathbb{R}^3$ is a vector spatial frequency. It is given by

$$s^3 + c(\gamma l_h + l_v) |\boldsymbol{\xi}|^2 s^2 + c^2 (1 + \gamma l_v l_h |\boldsymbol{\xi}|^2) |\boldsymbol{\xi}|^2 s + c^3 l_h |\boldsymbol{\xi}|^4 = 0 \quad (4)$$

2.1. A Simplified Model

The system (3), which is third order in time, possesses three independent solutions (a fact which is easily deduced from the characteristic equation, or complex dispersion relation (4)); only two of these correspond

to acoustic wave propagation. In the interest of simplification, consider the following approximation to (3b):

$$\beta \partial_t \tau = \frac{\gamma - 1}{\gamma} (\partial_t + cl_h \Delta) p \quad (5)$$

Under this approximation, and upon further neglecting terms of second order in l_h and l_v , the system (3) may be reduced to a single second order equation (a viscothermal wave equation) in the velocity potential Ψ as

$$\partial_t^2 \Psi - c\eta \Delta \partial_t \Psi - c^2 \Delta \Psi = 0 \quad (6)$$

where

$$\eta = l_v + (\gamma - 1) l_h \quad (7)$$

and with characteristic equation

$$s^2 + c\eta |\boldsymbol{\xi}|^2 s + c^2 |\boldsymbol{\xi}|^2 = 0 \quad (8)$$

The solutions to the above characteristic equation, defining the phase velocity and frequency-dependent loss characteristics for the system, are virtually identical to the two principal solutions of the characteristic equation (4) for the complete system when $\omega \ll c/\eta$, which is satisfied for the audible range of frequencies and physically reasonable choices of the constants c and η . For these frequencies of interest, plane wave solutions are of the form $e^{-\alpha t} e^{j(\omega t - \boldsymbol{\xi} \cdot \mathbf{x})}$, where α is the *temporal* absorption coefficient in Nepers/s

$$\alpha = \frac{\omega^2 \eta}{2c} = \frac{\omega^2}{2\rho c^2} \left(\frac{4}{3} \mu + \mu_B + (\gamma - 1) \frac{\kappa}{C_p} \right) \quad (9)$$

and the *spatial* absorption coefficient is

$$\alpha'_{cl} = \alpha/c \quad (10)$$

which expresses the attenuation over the distance ct in Nepers/m. The parameter α'_{cl} is also known as the *modified classical absorption coefficient* [11] (classical when $\mu_B = 0$ [12]). For the T_{60} decay time in seconds, a simple expression for wave-like solutions as a function of frequency is given by

$$T_{60} = \frac{3 \ln(10)}{\alpha} = \frac{3c \ln(10)}{2\eta \pi^2 f^2} \quad (11)$$

2.2. Numerical values for η

One can find specific values in standard texts [10, 11, 17] for most of the parameters featured here, except for the bulk viscosity μ_B , which seems to be unknown for air [26]. However, it has been estimated to be on the order of the viscosity μ [26], and in [11, 12] the estimate $\mu_B \approx 0.6\mu$ is given. This estimate for μ_B would result in a value for η on the order of 10^{-7} m for air at standard conditions. In any case, the “frequency-squared” absorption is lacking additional relaxation effects due to nitrogen and oxygen [11, 18]. A more




Figure 1: Absorption coefficients (from [18], and in agreement with [11]) at 20°C and standard pressure (one atm). Relaxation frequencies for oxygen and nitrogen (f_{r,O_2} and f_{r,N_2}) are marked by vertical dashed lines.

general absorption coefficient is often written as [11, 18, 22]

$$\alpha_{tot} = \alpha'_{cl} + \alpha_{N_2} + \alpha_{O_2} \quad (12)$$

for plane wave solutions $e^{-\alpha_{tot}ct}e^{j(\omega t - \xi \cdot \mathbf{x})}$, where α_{N_2} and α_{O_2} represent absorption due to relaxation processes in oxygen or nitrogen respectively. These processes are dependent on the ambient pressure, temperature, and relative humidity of the air [18, 11, 12]. Examples at standard conditions with various relative humidities are shown in Fig. 1. Below their respective *relaxation frequencies*, f_{r,N_2} and f_{r,O_2} , these absorption coefficients (α_{r,N_2} and α_{r,O_2}) behave as “frequency-squared” absorptions, and above they behave as constants; this can be seen in Fig. 1. It can also be seen that $\alpha_{tot} \approx \alpha_{O_2}$ in the frequency range 1 kHz-20 kHz for a relative humidity in the range of 30%-50%. This range of relative humidities is typical of indoor conditions.

It makes sense to use a first-order approximation to α_{O_2} and attribute this value to an additional increment to the bulk viscosity, as suggested in [11]. This can be accomplished from the expressions provided in [18, 11, 12]. Values of η , due to oxygen relaxation, are provided in Fig. 2, which should give relatively accurate absorption for 1 kHz-20 kHz frequencies when the relative humidity is between 30% and 80%. Indeed, this does not properly model frequencies near and below f_{r,N_2} , but the absorption at those frequencies should be negligible in comparison to wall absorption effects in practical room acoustics situations.

3. Finite difference schemes

In this section, finite difference schemes for the initial value problem (6) will be developed and stability conditions will be derived.

3.1. Grids and grid function

Let Ψ_i^n represent the discrete approximation (a grid function) to the solution of interest,




Figure 2: Values for η that can be attributed to first-order approximations of oxygen relaxation effects at 20°C and standard pressure (one atm), adapted from [18].

i.e. $\Psi_i^n \approx \Psi(nk, \mathbf{x}_i)$ for $n \in \mathbb{Z}^+$, $\mathbf{x}_i := ih, i = (i_x, i_y, i_z) \in \mathbb{Z}^3$. The grid spacing is h and the time-step is k .

3.2. Time difference operators

Shift operators in time and space may be defined as

$$s_{t\pm} \Psi_i^n := \Psi_i^{n\pm 1} \quad (13a)$$

$$s_{x\pm} \Psi_i^n := \Psi_{i\pm \hat{x}}^n \quad (13b)$$

where $\hat{x} = (1, 0, 0)$ and the operators $s_{y\pm}$ and $s_{z\pm}$ are similarly defined with $\hat{y} = (0, 1, 0)$ and $\hat{z} = (0, 0, 1)$ respectively. Averaging and finite difference operators can then be constructed as

$$\mu_{t-} := (1 + s_{t-})/2 = 1 + O(k) \quad (14a)$$

$$\delta_{t-} := (1 - s_{t-})/k = \partial_t + O(k) \quad (14b)$$

$$\delta_{tt} := (s_{t+} - 2 + s_{t-})/k^2 = \partial_t^2 + O(k^2) \quad (14c)$$

$$\delta_{xx} := (s_{x+} - 2 + s_{x-})/h^2 = \partial_x^2 + O(h^2) \quad (14d)$$

and similarly for δ_{yy} and δ_{zz} .

3.3. Two-step scheme

A simple approximation to (6) is the following explicit two-step scheme

$$\delta_{tt}\Psi_i^n - c^2\delta_\Delta\Psi_i^n - c\eta\delta_{t-}\delta_\Delta\Psi_i^n = 0 \quad (15)$$

where δ_Δ is a discrete Laplacian operator (stencil), of which there are many choices [27]; the simplest choice is the 7-point stencil $\delta_{\Delta,\text{SLF}} := \delta_{xx} + \delta_{yy} + \delta_{zz}$ [13], leading to a scheme that is sometimes known as “standard leapfrog” (SLF) when $\eta = 0$ [28]. The scheme (15) proceeds from the known or approximated values Ψ_i^0, Ψ_i^1 and the approximation is updated with the recursion:

$$\Psi_i^{n+1} = (2 + (\lambda^2 + 2\theta)\delta_\Delta^h)\Psi_i^n - (1 + 2\theta\delta_\Delta^h)\Psi_i^{n-1} \quad (16)$$

where $\delta_\Delta^h := h^2\delta_\Delta$, $\theta = \lambda\eta/(2h)$, and $\lambda = ck/h$. This scheme requires three states to be stored in memory rather than two for the lossless case (when $\eta = 0$). Choosing $\delta_\Delta = \delta_{\Delta,\text{SLF}}$, (16) can be further expanded to

$$\Psi_i^{n+1} = 2\Psi_i^n - \Psi_i^{n-1} + (\lambda^2 + 2\theta)Q_i^n - 2\theta Q_i^{n-1} \quad (17)$$

where Q_i^n is the sum defined as

$$Q_i^n = \Psi_{i+\hat{x}}^n + \Psi_{i-\hat{x}}^n + \Psi_{i+\hat{y}}^n + \Psi_{i-\hat{y}}^n + \Psi_{i+\hat{z}}^n + \Psi_{i-\hat{z}}^n - 6\Psi_i^n \quad (18)$$

Note that the variable Q_i^{n-1} does not need to be calculated at time-step n since it would have been calculated at time-step $n-1$. It is straightforward to work out that this redundancy can be exploited without the use of additional storage.

3.4. A parameterised three-step scheme

In the interest of providing a more accurate approximation to the viscothermal component, consider the following three-step scheme with the free parameter $\varepsilon \in \mathbb{R}$

$$\delta_{tt}\Psi_i^n - c^2\delta_\Delta\Psi_i^n - c\eta\varepsilon\delta_{t-}\delta_\Delta\Psi_i^n - c\eta(1-\varepsilon)\mu_{t-}\delta_{t-}\delta_\Delta\Psi_i^n = 0 \quad (19)$$

The two-step scheme (15) is recovered with $\varepsilon = 1$. The time-recursion for this three-step scheme is

$$\Psi_i^{n+1} = (2 + (\lambda^2 + \theta(1 + \varepsilon))\delta_\Delta^h)\Psi_i^n - (1 + 2\varepsilon\theta\delta_\Delta^h)\Psi_i^{n-1} + \theta(\varepsilon - 1)\delta_\Delta^h\Psi_i^{n-2} \quad (20)$$

This requires an extra state to be stored in memory for $\varepsilon \neq 1$, as well as one extra initial state to start the recursion. With the choice $\delta_\Delta = \delta_{\Delta,\text{SLF}}$, this is

$$\Psi_i^{n+1} = 2\Psi_i^n - \Psi_i^{n-1} + (\lambda^2 + \theta(1 + \varepsilon))Q_i^n - 2\varepsilon\theta Q_i^{n-1} + \theta(\varepsilon - 1)Q_i^{n-2} \quad (21)$$

Again, there is a redundancy here in the variables Q_i^{n-1} and Q_i^{n-2} , and only one stencil operation (Q_i^n) needs to be computed per time-step. This redundancy can also be exploited without the use of additional storage.

3.5. Stability conditions

Stability conditions for (20) can be derived through the use of the Schur-Cohn recursion [29]. Consider a polynomial in $z \in \mathbb{C}$ of degree n :

$$A_n(z) = a_{0,n} + a_{1,n}z + \dots + a_{n,n}z^n \quad (22)$$

and the associated reciprocal polynomial:

$$A_n^*(z) = a_{n,n} + a_{n-1,n}z + \dots + a_{0,n}z^n \quad (23)$$

and define

$$A_{n-1}(z) = z^{-1}(A_n^*(0)A_n(z) - A_n(0)A_n^*(z)) \quad (24)$$

According to Schur-Cohn theory [29], the roots of some polynomial $A_N(z)$ will be $|z| \leq 1$ as long as

$$\left| \frac{a_{0,n}}{a_{n,n}} \right| \leq 1, \quad n = N, \dots, 1 \quad (25)$$

In order to employ this recursion, the spatial Fourier transform and the Z-transform are applied to (20),

$$z^3 - (2 - (\lambda^2 + \theta(\varepsilon + 1))\Lambda)z^2 + (1 - 2\varepsilon\theta\Lambda)z + \theta(\varepsilon - 1)\Lambda = 0 \quad (26)$$

where $\Lambda = \Lambda(\boldsymbol{\xi})$ is the Fourier symbol of the operator $-\frac{1}{4}\delta_\Delta^h$ and $\boldsymbol{\xi} \in \mathbb{R}^3$ are the spatial frequencies. For now it will be assumed that $\Lambda \geq 0$ and $\Lambda \in \mathbb{R}$, which is a necessary condition for stability. Note that (26) has three roots. Two roots must appear in a complex conjugate pair and these pertain to wave-like solutions. The third root describes a parasitic mode that is heavily damped and will only appear for a vanishingly small period of time, so its existence is negligible.

Applying the Schur-Cohn recursion, the following stability condition on the time-step can be obtained from (26) for the special case of $\varepsilon \geq 1/2$

$$k \leq k_{\max} := \sqrt{h^2/(c^2\Lambda_{\max}) + (\varepsilon\eta/c)^2} - \varepsilon\eta/c \quad (27)$$

where $\Lambda_{\max} := \max_{\boldsymbol{\xi}} \Lambda$. This condition can be equivalently expressed in terms of the spatial step

$$h \geq h_{\min} := \sqrt{\Lambda_{\max}(c^2k^2 + 2\varepsilon\eta ck)} \quad (28)$$

which also requires that $\varepsilon \geq 1/2$. Stability conditions can be derived for $\varepsilon < 1/2$, but they become more complicated, and as will be seen, $\varepsilon \geq 1$ happens to be a useful range for the three-step scheme.

Condition (27) can be employed in order to choose k if one initially sets the spatial step h , whereas condition (28) can be used if one initially sets the time-step k according to audio sampling considerations, i.e.

$k = 1/F_s$ where F_s is a sample rate of interest, such as $F_s = 44.1$ kHz. Henceforth, it will be assumed that the latter approach is employed, with $k = 1/F_s$ for some F_s . It will also be assumed that $h = h_{\min}$, since this maximises the temporal bandwidth that can be simulated [27] and provides optimal numerical dispersion in certain schemes [28].

With the simplest choice $\delta_\Delta = \delta_{\Delta, \text{SLF}}$, the function Λ can be written as

$$\Lambda = \Lambda_{\text{SLF}} := S_x + S_y + S_z \quad (29)$$

where $S_x := \sin^2(\xi \cdot \hat{x}h/2)$ and similarly for S_y, S_z . In this case, $\Lambda \geq 0$ and $\Lambda_{\max} = 3$, and the usual Courant limit for the lossless scheme $\lambda \leq \sqrt{1/3}$ is obtained when $\eta = 0$. Furthermore, when $\varepsilon = 1$ the stability limit for the two-step viscothermal scheme [13] is obtained from (28).

4. Numerical dispersion and loss

In this section, conventional numerical dispersion analysis will be expanded to include the effects of loss. For this purpose, it will help to define a normalised spatial frequency $\xi_h := \xi h$ as well as a normalised angular frequency $\hat{\omega}_k := \hat{\omega}k \in \mathbb{R}$, and a normalised absorption coefficient $\alpha_k := \alpha k \in \mathbb{R}$. These are related by $\ln(z) = -\alpha_k + j\hat{\omega}_k$, where $z = z(\xi_h)$ is one of two wave-like solutions to (26).

4.1. Numerical phase velocity

For frequencies $f \ll c/\eta$, which is a range that comprises the full audible bandwidth, the dispersion relation (8) predicts a phase velocity that is virtually constant (non-dispersive) and an attenuation that scales with the frequency squared [16]. As is well known, finite difference schemes numerically introduce dispersion, and thus, a variable numerical phase velocity.

A *relative phase velocity* for the scheme may be written as

$$v_{\text{rel}}(\xi_h) = \frac{\hat{\omega}_k}{\lambda|\xi_h|}, \quad \hat{\omega}_k = \Im(\ln(z(\xi_h))) \quad (30)$$

where z is one of two roots of (26) that pertain to wave-like solutions. v_{rel} is also known as the *dispersion coefficient* and ideally it should be unity for $f \ll c/\eta$. For audio and acoustics applications it is desirable to view the relative phase velocity as a function of temporal frequencies, but, even in the lossless FD scheme, a numerical analogue of the dispersion relation $\omega = c|\xi|$ does not hold. As in [30], the relative phase velocity along a fixed direction of plane wave propagation can be reassigned to temporal frequencies through the function that defines the numerical dispersion relation ($\hat{\omega}_k$ as a function of ξ_h). This reassigned relative phase velocity is plotted in Fig. 3a along the axial and diagonal directions with respect to the cubic grid (these are the directions where extreme cases are generally

found [28]). As can be seen, even with a relatively large choice of η , and for various choices of ε , the relative phase velocity is virtually identical to the lossless case for $c = 340$ m/s and $F_s = 44.1$ kHz.

It is important to point out that the axial directions will not propagate plane waves with frequencies beyond a certain cutoff frequency, which is approximately $0.2F_s$ [28] for the simplest scheme ($\delta_\Delta = \delta_{\Delta, \text{SLF}}$). To mitigate the effect of this cutoff frequency, as well as improve the relative phase velocity, one could choose a wider stencil than the standard one (at the cost of more computation), such as the 27-point stencil from the “interpolated wideband” (IWB) scheme [28]. This 27-point stencil can be defined as

$$\delta_{\Delta, \text{IWB}}^h = \delta_{xx}^h + \delta_{yy}^h + \delta_{zz}^h + \delta_{xx}^h \delta_{yy}^h \delta_{zz}^h + \delta_{xx}^h \delta_{yy}^h + \delta_{xx}^h \delta_{zz}^h + \delta_{yy}^h \delta_{zz}^h \quad (31)$$

where $\delta_{xx}^h := h^2 \delta_{xx}$ and similarly for δ_{yy}^h and δ_{zz}^h . The Fourier symbol of $-\frac{1}{4}\delta_{\Delta, \text{IWB}}^h$ is then

$$\Lambda_{\text{IWB}} = S_x + S_y + S_z + S_x S_y S_z - (S_x S_y + S_x S_z + S_y S_z) \quad (32)$$

and it is straightforward to show that $0 \leq \Lambda_{\text{IWB}} \leq 1$. By simply choosing $\delta_\Delta^h = \delta_{\Delta, \text{IWB}}^h$ the lossless IWB scheme [28] can be generalised to the viscothermal formulation (19), giving a three-step scheme with virtually no directional cutoff frequency aside from the Nyquist.¹ The relative phase velocity for this scheme, with various η and ε , is seen in Fig. 3b. Also, note that the presence of small viscothermal losses does not significantly affect the numerical phase velocity, as in the case of $\delta_\Delta = \delta_{\Delta, \text{SLF}}$.

4.2. Numerical dissipation

The numerical dissipation may be analysed by a *relative absorption coefficient*

$$\alpha_{\text{rel}}(\xi_h) = \frac{\hat{\alpha}_k}{\alpha_k}, \quad \hat{\alpha}_k = -\Re(\ln(z(\xi_h))) \quad (33)$$

The relative absorption coefficient can also be reassigned back to temporal frequencies via $\hat{\omega}_k$; this is displayed in Fig. 4a for various choices of ε with

¹ The minimum cutoff frequency may be used as a measure to compare different schemes, leading to the choice of the IWB scheme over other 27-point variants in some studies [31]. However, focussing on the cutoff frequency ignores the presence of phase velocity errors, which can be as much as 40% for the IWB scheme, as seen in Fig. 3b. The numerical phase (or group) velocity is, arguably, more critical in determining the validity of a simulated bandwidth. An error criterion in the range of 0.5%-2% may be necessary in order to ensure that reflection errors from modelled wall impedances are kept reasonably low [28, 32]. Whether or not such errors, or other artefacts caused by numerical dispersion, are *perceptible* is another matter completely. Currently, research on the perceptibility of dispersion error is limited [33, 34]. More will be said about this in Section 7.




Figure 3: Relative phase velocity for standard and IWB schemes with $\eta \in \{0, 1e-5\}$, $\varepsilon \in \{1, 1.3\}$, $c = 340$ m/s and $F_s = 44.1$ kHz

$\delta_\Delta = \delta_{\Delta,SLF}$. Notice that α_{rel} is correct at DC ($f = 0$), which follows from the consistency of the scheme. Otherwise, α_{rel} varies with frequency.

Unlike the numerical phase velocity, the numerical dissipation (as a function of temporal frequencies) appears to be isotropic, i.e. independent of the direction of plane wave propagation, since the axial and diagonal curves overlap for the same ε , aside from the effect of the aforementioned cutoff frequency near $0.2F_s$. This isotropy is also seen with various 27-point stencils, such as the case $\delta_\Delta^h = \delta_{\Delta,IWB}^h$, as displayed in Fig. 4b. As in the previous case, the curves pertaining to axial and diagonal directions (as well as other directions, left out for brevity) overlap for the same ε . Furthermore, these curves (Fig. 4b) are virtually identical to those along diagonal direction of the standard scheme (blue lines, Fig. 4a). As such, it should suffice to investigate only the diagonal direction of the standard scheme for the remainder of this study.

4.3. Normalising for memory costs

According to Figs. 4a and 4b, the three-step scheme with a choice of $\varepsilon \in [1.2, 1.3]$ results in an improved approximation to the viscothermal component, in comparison to the two-step scheme. However, the three-step scheme requires one extra state to be stored in memory. Thus, the increase in memory over the two-step scheme is a factor of $4/3$. It is always possible to oversample the grid of the two-step scheme, increasing its memory cost, in order to improve the approximation up to a given f (by consistency), so it makes sense to compare the two schemes for an equal memory cost.




Figure 4: Relative absorption coefficients for standard and IWB schemes with various ε , $c = 340$ m/s and $F_s = 44.1$ kHz




Figure 5: α_{rel} after normalising for memory costs. The two-step scheme with $F_s = 48.5$ kHz has the same memory usage as the three-step scheme with $F_s = 44.1$ kHz.

For an oversampling factor of X in the two-step scheme (λ fixed), i.e. choosing $F_s = 44.1X$ kHz, the memory storage increases by a factor of X^3 . An oversampling factor of $X = \sqrt[3]{4/3} \approx 1.101 \Rightarrow F_s = 48.5$ kHz would then normalise for memory costs. Relative absorption coefficients are plotted in Fig 5 with this normalisation taken into account. It can be seen that the three-step scheme offers a better approximation to the viscothermal component than the two-step scheme even at the same memory cost.

It is also worth plotting the numerical absorption on a logarithmic scale in order to see how it deviates from the classical frequency-power law (a straight line on the log-log scale) for a particular choice of η . This is shown in Fig. 7.

5. Numerical Experiments

In this section, T_{60} decay times obtained from simulations with the two-step and three-step schemes will be compared to the expression (11). It is expected that




Figure 6: T_{60} numerical tests for $c = 340$ m/s, $\eta = 1e-6$ m. White represents a level of -60dB and below.




Figure 7: Classical and numerical absorption for $c = 340$ m/s and $\eta = 1.5e-6$ m

the simulated T_{60} times will be close to the following expression obtained from a frequency-domain analysis:

$$\hat{T}_{60} = \frac{3 \ln(10)}{\hat{\alpha}_k / k} \quad (34)$$

A small box of air (approximately one cubic meter) was modelled with fully reflective boundaries and $\eta = 1e-6$ m. As there is no loss at the boundaries, the decay of the numerical solution is solely due to the viscothermal component. In order to excite the entire spectrum, uniform random vectors were used to initialise the time recursions. A typical response, obtained from a point within the box (the location is immaterial) is displayed in Fig. 8.

The spectrograms displayed in Fig. 6 were calculated using a Hann window of 512 samples with a 50 sample hop size. In these spectrograms, zero decibels represents the average magnitude of the spectrum of the output at the reference time $t = 0$ and white represents a magnitude equal to or less than -60 dB. Thus, the simulated T_{60} decay time is the demarcation between greys and pure white. In each spectrogram, this demarcation coincides with what is predicted by (34). As expected, the three-step schemes with $F_s = 44.1$ kHz (Figs. 6b and 6c) have simulated T_{60} times that are in better agreement with the model (11) than the two-step scheme with $F_s = 48.8$ kHz (normalised for computational cost).




Figure 8: Response from simulation with viscothermal losses ($\eta = 1e-6$ m) and lossless walls

6. GPU Implementation

The viscosity schemes considered here were implemented for the CUDA architecture and tested with an Nvidia Tesla K20 GPU device. The general threading approach was to issue enough threads to cover the entire 3-D data grid, as opposed to using a 2-D set of threads and iterating over the last dimension. A thread block size of $32 \times 2 \times 2$ was used. In terms of memory usage, the Kepler read-only data cache was used where possible, but the use of shared memory was not investigated. Constant memory was used to store global coefficients (stencil weights, θ , λ , ε).

The viscosity schemes were tested with 7-, 13-, 19-, and 27-point stencils operating on a cubic grid. These stencils are special cases of the 3-D “interpolated schemes” [35, 28, 27] (formulations for these stencils are provided in the Appendix). Two versions of the viscosity schemes were tested. The first version used a standard implementation whereby each stencil operation was explicitly calculated from its standard data grids. In the second version, the redundancy of the algorithm (mentioned in Section 3) was exploited in order to reduce the number of stencil operations, and thus, the amount of data that was read from the grids. The algorithm is memory bandwidth limited on such GPU devices [9] and thus the reduction in memory usage was exploited as opposed to the reduction in floating-point operations.

Tests were conducted for a $640 \times 480 \times 480$ sized grid in both single and double precision, and in each case,

Table I: Single precision timing results

stencil size	lossless	$\eta > 0$, two-step		$\eta > 0$, three-step	
	time (s)	CTI-A	CTI-B	CTI-A	CTI-B
7-point	42	1.17x	1.31x	1.62x	1.69x
13-point	47	1.55x	1.21x	2.38x	1.55x
19-point	51	1.69x	1.14x	3.33x	1.49x
27-point	75	1.83x	1.12x	3.41x	1.27x

Table II: Double precision timing results

stencil size	lossless	$\eta > 0$, two-step		$\eta > 0$, three-step	
	time (s)	CTI-A	CTI-B	CTI-A	CTI-B
7-point	68	1.31x	1.47x	1.90x	1.93x
13-point	78	1.54x	1.38x	2.54x	1.79x
19-point	79	2.27x	1.37x	3.63x	1.75x
27-point	89	2.73x	1.29x	4.62x	1.60x

2000 time-steps were computed. The timing results are presented in Table I for single precision and Table II for double precision. In these tables, the timings for the two- and three-step viscosity schemes implemented using the standard approach (version ‘A’) and the data-reduced approach (version ‘B’) are compared to the timings for the respective lossless schemes. The comparison is presented in terms of a “compute time increase” (CTI) (slow-down) of the viscosity timings over the respective lossless timings. It should be mentioned that frequency-independent impedance boundary conditions (not featured here, similar to [13, 19, 9]) were included in these GPU implementations.

It can be seen that the data-reduced approach provides lower computation times than the standard approach for 13-, 19-, and 27-point stencils. The improvements obtained from the data-reduced approach become more substantial as the stencil size increases. In the case of the 27-point scheme in double precision, the data-reduced implementation is almost three times as efficient as the standard approach. However, in the special case of the 7-point stencil ($\delta_{\Delta,SLF}$) the standard approach runs slightly faster than the data-reduced approach. This may be due to caching effects for these memory bandwidth limited algorithms. Comparing the lossless case to the data-reduced form, the reason for the increase in computation times is due to the extra read and write requirements of the data-reduced approach over and above that of the lossless case, even though both compute only one stencil operation per time-step.

7. Conclusions and Future Directions

Finite difference schemes were presented for a viscothermal wave equation. Absorption parameters were provided which are suitable for audible frequencies in the context of room acoustics with standard indoor conditions (temperature, pressure, relative humidity)

in air. Stability conditions for a three-step parameterised FD scheme were derived for the initial-value problem. Numerical dispersion and numerical dissipation were analysed in the frequency domain and optimal parameters for the three-step scheme were selected, demonstrating improved accuracy in the approximation to the viscothermal component over two-step schemes. This improvement was confirmed by numerical experiments analysing the numerical T_{60} decay times as a function of frequency via spectrograms. Timing results from a GPU implementation on an Nvidia Tesla K20 GPU device were presented and it was shown that using a data-reduced approach, the computation times for the three-step schemes were reasonable in comparison to lossless counterparts.

It should be pointed out that an additional redundancy in the 13-point stencil operating over the cubic grid could be exploited by implementing the 13-point stencil on a face-centered cubic (FCC) grid [36]. This comes at the expense of a slightly more complex kernel operation and memory decomposition [37, 9].

It is important to consider the amount of numerical dispersion error in these FD schemes, which may limit the amount of “usable bandwidth”. The improvements offered by the three-step schemes are generally in high frequencies where dispersion error is more pronounced. Thus, it may not be critical to employ three-step schemes over two-step viscothermal counterparts when the interest is only in a lower portion of the simulated bandwidth. However, as the numerical dissipation seems to be independent of the stencil (discrete Laplacian) employed (aside from cutoff frequencies), these three-step schemes will become more useful for stencils which offer a wider bandwidth under which dispersion error is kept to a reasonable level [27].

Other directions of future study are to incorporate complex-impedance boundary conditions, taking into account boundary layer effects due to viscosity [10, 12], and to model complex geometries using a more general finite volume framework [38].

8. Appendix

As in [28], a parameterised discrete Laplacian (normalised by h^2) using up to 27 points may be written as:

$$\delta_{\Delta}^h = \delta_{xx}^h + \delta_{yy}^h + \delta_{zz}^h + a(\delta_{xx}^h \delta_{yy}^h + \delta_{xx}^h \delta_{zz}^h + \delta_{yy}^h \delta_{zz}^h) + b(\delta_{xx}^h \delta_{yy}^h \delta_{zz}^h)$$

The coefficients a, b and respective values for $\Lambda_{\max}$ for the various stencils used in the GPU tests are provided in Table III.

References

- [1] R. M. Alford, K. R. Kelly, and D. M. Boore, “Accuracy of finite-difference modeling of the acoustic wave equation,” *Geophysics*, vol. 39, no. 6, pp. 834–842, 1974.
- [2] M. A. Dablain, “The application of high-order differencing to the scalar wave equation,” *Geophysics*, vol. 51, no. 1, pp. 54–66, 1986.

Table III: Stencil parameters

stencil size	a	b	$\Lambda_{\max}$
7-point	0	0	3
13-point	1/4	0	1
19-point	1/6	0	4/3
27-point	1/6	1/48	4/3

- [3] D. Botteldooren, "Finite-difference time-domain simulation of low-frequency room acoustic problems," *J. Acoustical Society of America*, vol. 98, pp. 3302–3308, 1995.
- [4] A. Taflov, "Application of the finite-difference time-domain method to sinusoidal steady-state electromagnetic-penetration problems," *IEEE Trans. Electromagnetic Compatibility*, vol. EMC-22, no. 3, pp. 191–202, 1980.
- [5] R. Courant, K. Friedrichs, and H. Lewy, "Über die partiellen differenzengleichungen der mathematischen physik," *Mathematische Annalen*, vol. 100, no. 1, pp. 32–74, 1928.
- [6] G. E. Forsythe and W. R. Wasow, *Finite-difference methods for partial differential equations*. New York: Wiley, 1960.
- [7] R. D. Richtmyer and K. W. Morton, *Difference methods for initial-value problems*, 2nd ed. Interscience Publishers, 1967, ch. 10: Sound Waves.
- [8] J. Botts and L. Savioja, "Integrating finite difference schemes for scalar and vector wave equations," in *Proc. IEEE ICASSP*, Vancouver, Canada, 2013, pp. 171–175.
- [9] C. J. Webb, "Parallel computation techniques for virtual acoustics and physical modelling synthesis," Ph.D. thesis, University of Edinburgh, 2014.
- [10] P. M. Morse and K. U. Ingard, *Theoretical acoustics*. Princeton University Press, 1968.
- [11] A. D. Pierce, *Acoustics: an introduction to its physical principles and applications*. McGraw-Hill New York, 1989.
- [12] D. T. Blackstock, *Fundamentals of Physical Acoustics*. John Wiley & Sons, 2000.
- [13] C. J. Webb and S. Bilbao, "Computing room acoustics with CUDA - 3D FDTD schemes with boundary losses and viscosity," in *Proc. IEEE ICASSP*, Prague, Czech Republic, 2011, pp. 317–320.
- [14] A. Southern, S. Siltanen, D. Murphy, and L. Savioja, "Room impulse response synthesis and validation using a hybrid acoustic model," *IEEE Trans. Audio, Speech, and Language Processing*, vol. 21, no. 9, Sep. 2013.
- [15] G. G. Stokes, "Sound attenuation due to viscosity," *Trans. Camb. Phil. Soc.*, vol. 8, pp. 75–102, 1845.
- [16] M. J. Buckingham, "Causality, Stokes' wave equation, and acoustic pulse propagation in a viscous fluid," *Physical Review E*, vol. 72, no. 2, p. 026610, 2005.
- [17] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of acoustics*, 4th ed. John Wiley & Sons, 2000.
- [18] "Acoustics – attenuation of sound during propagation outdoors. Part 1: Calculation of the absorption of sound by the atmosphere," International Organization for Standardization, Geneva, Switzerland, Standard ISO 9613-1: 1993.
- [19] C. J. Webb and A. Gray, "Large-scale virtual acoustics simulation at audio rates using three dimensional finite difference time domain and multiple GPUs," in *Proc. Int. Cong. Acoustics (ICA)*, Montréal, Canada, 2013.
- [20] Z. Zhang, "The multistep finite difference fractional steps method for a class of viscous wave equations," *Mathematical Methods in the Applied Sciences*, vol. 34, no. 4, pp. 442–454, 2011.
- [21] D. Deng, "Accuracy improvement of a multistep splitting method for nonlinear viscous wave equations," *International Journal of Computer Mathematics*, no. ahead-of-print, pp. 1–20, 2014.
- [22] M. Wochner, "Numerical simulation of multi-dimensional acoustic propagation in air including the effects of molecular relaxation," Ph.D. dissertation, The Pennsylvania State University, 2006.
- [23] X. Yuan, D. Borup, J. Wiskin, M. Berggren, and S. Johnsson, "Simulation of acoustic wave propagation in dispersive media with relaxation losses by using FDTD method with PML absorbing boundary condition," *IEEE Transactions on, Ultrasonics, Ferroelectrics and Frequency Control*, *IEEE Transactions on*, vol. 46, no. 1, pp. 14–23, 1999.
- [24] C. de Groot-Hedlin, "Finite-difference time-domain synthesis of infrasound propagation through an absorbing atmosphere," *J. Acoustical Society of America*, vol. 124, no. 3, pp. 1430–1441, 2008.
- [25] M. Bruneau, P. Herzog, J. Kergomard, and J. Polack, "General formulation of the dispersion equation in bounded visco-thermal fluid, and application to some simple geometries," *Wave motion*, vol. 11, no. 5, pp. 441–451, 1989.
- [26] M. S. Cramer, "Numerical estimates for the bulk viscosity of ideal gases," *Physics of Fluids*, vol. 24, no. 6, p. 066102, 2012.
- [27] 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.
- [28] 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.
- [29] J. Strikwerda, *Finite difference schemes and partial differential equations*. Philadelphia, PA: SIAM, 2004.
- [30] M. van Walstijn and K. Kowalczyk, "On the numerical solution of the 2D wave equation with compact FDTD schemes," in *Proc. Digital Audio Effects (DAFx)*, Espoo, Finland, 2008, pp. 205–212.
- [31] T. Ishii, T. Tsuchiya, and K. Okubo, "Three-dimensional sound field analysis using compact explicit-finite difference time domain method with graphics processing unit cluster system," *Japanese Journal of Applied Physics*, vol. 52, no. 7S, p. 07HC11, 2013.
- [32] N. Borrel-Jensen, "Real-time auralisation of the lower frequency sound field using numerical methods on the GPU," Master's thesis, RWTH Aachen University and University of Copenhagen, 2012.
- [33] A. Southern, D. T. Murphy, T. Lokki, and L. Savioja, "The perceptual effects of dispersion error on room acoustic model auralization," in *Proc. Forum Acusticum, Aalborg, Denmark*, 2011, pp. 1553–1558.
- [34] M. Cobos, J. Escolano, J. J. López, and B. Pueo, "Subjective effects of dispersion in the simulation of room acoustics using Digital Waveguide Mesh," in *Audio Engineering Society Convention 124*, 2008.
- [35] S. Bilbao, "Wave and scattering methods for the numerical integration of partial differential equations," Ph.D. thesis, Stanford University, 2001.
- [36] B. Hamilton and S. Bilbao, "On finite difference schemes for the 3-D wave equation using non-Cartesian grids," in *Proc. Sound and Music Computing (SMC) Conf.*, Stockholm, Sweden, 2013, pp. 592–599.
- [37] 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, Sep. 2013, pp. 336–343.
- [38] S. Bilbao, "Modeling of complex geometries and boundary conditions in finite difference/finite volume time domain room acoustics simulation," *IEEE Trans. Audio, Speech, and Language Processing*, vol. 21, no. 7, pp. 1524–1533, Jul. 2013.