

Biomedical
Ultrasound
Group



k-Wave short course – Part 5

Accuracy and convergence

Bradley Treeby and Ben Cox

Biomedical Ultrasound Group (BUG)
Department of Medical Physics and Biomedical Engineering
University College London

Simulation accuracy

- For **linear** simulations in a **homogeneous, lossless, fluid** medium, the numerical method used in k-Wave is **exact** and **unconditionally stable**
- What about more complex scenarios?


Simulation accuracy

There are several factors to consider:





1. Perfectly matched layer
2. Numerical dispersion in heterogeneous media
3. Accuracy of reflection and transmission coefficients
4. Source staircasing
5. Medium staircasing
6. Acoustic absorption
7. Acoustic nonlinearity

1. Perfectly matched layer

- The effectiveness of the PML depends on:
 - PML size
 - PML absorption
 - Frequency
 - Angle of incidence
- Can assess the PML by measuring the **reflected** and **transmitted** waves




Normal incidence




Normal incidence

- At 3 PPW, signal is reduced by -100 dB ($1e-5$)
- At 4 PPW, signal is reduced by -120 dB ($1e-6$)





Time step

- Almost no dependence on time step for $\text{CFL} < 0.5$
- For $\text{CFL} > 1$, PML is ineffective



PML parameters

- The efficacy depends on the PML parameters, controlled by '`PMLSize`' and '`PMLAlpha`'
- Defaults are 20 grid points in 1D/2D and 10 grid points in 3D, with a `PMLAlpha` value of 2





PML parameters

- The efficacy depends on the PML parameters, controlled by 'PMLSize' and 'PMLAlpha'
- Defaults are 20 grid points in 1D/2D and 10 grid points in 3D, with a PMLAlpha value of 2
- Default parameters are typically ok
- Increasing the PML thickness will improve performance
- Significantly increasing PML absorption will not!


Angle of incidence

- PML performance is similar across angles



PML position

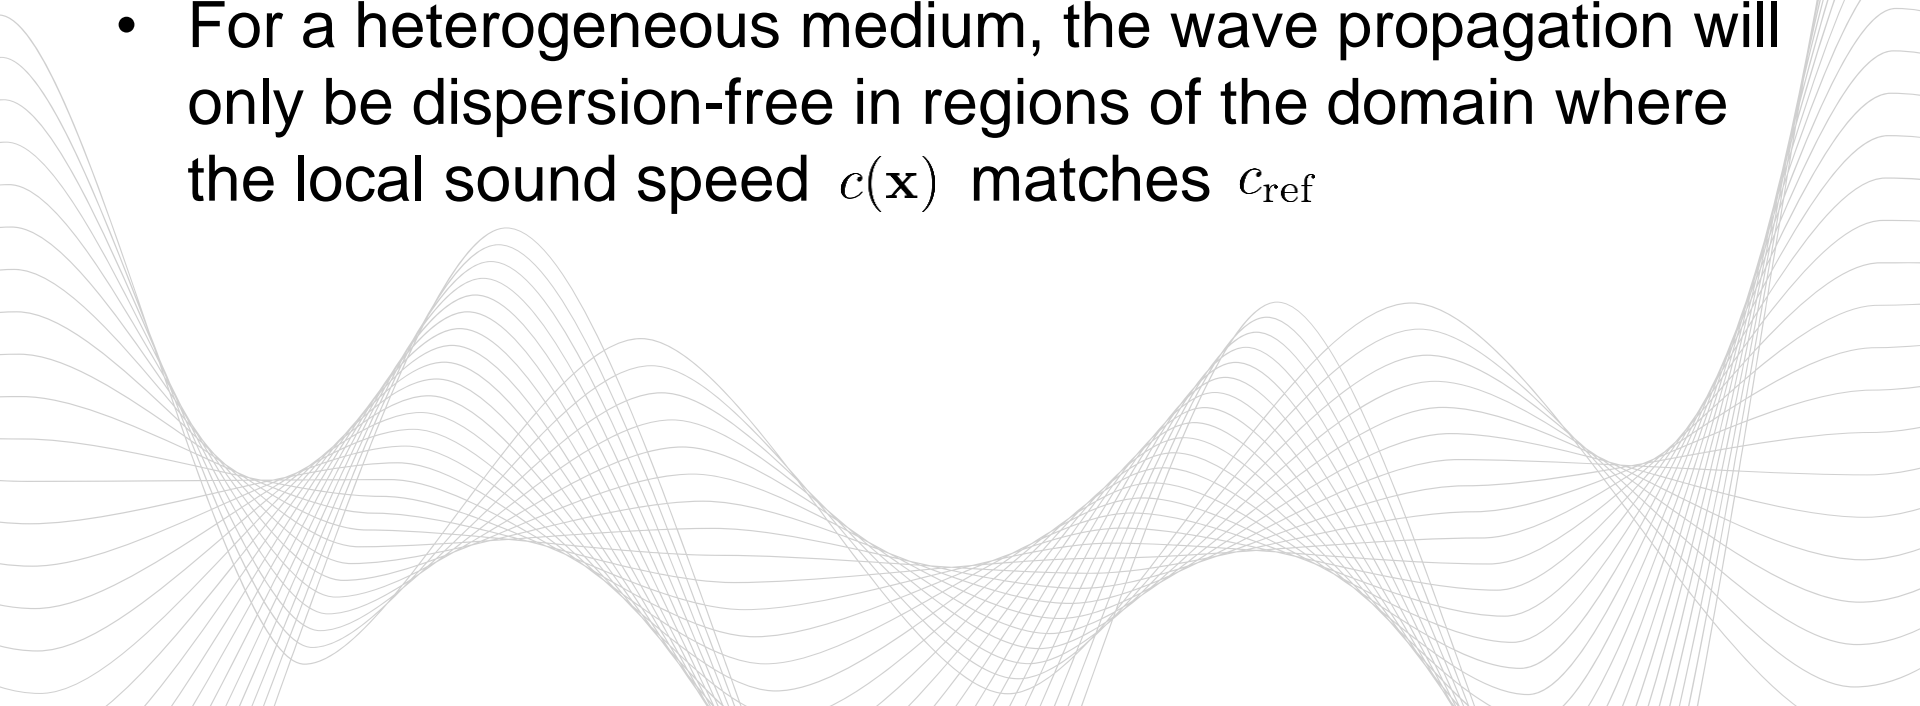
- By default, the PML is *inside* the domain specified by the user



- Care must be taken to position sources and sensors *outside* the PML
- The PML can also be placed outside by setting `'PMLInside'` to `false`


2. Numerical dispersion

- Recall the k-space operator $\kappa = \text{sinc}(c_{\text{ref}}k\Delta t/2)$ is used to correct for numerical dispersion in the finite difference time step
- This is applied in the Fourier domain, which means only a single value for the sound speed can be used
- For a heterogeneous medium, the wave propagation will only be dispersion-free in regions of the domain where the local sound speed $c(\mathbf{x})$ matches c_{ref}



Reference sound speed

- Phase error in the propagation of a plane wave after 50 wavelengths for $c_0 = 1500$ m/s



Treeby et al, JASA, 2012

Choice of reference sound speed

- If $c_{\text{ref}} \geq \max(c(\mathbf{x}))$, then the simulation is unconditionally stable, but the phase error can accumulate in regions of the domain where $c_{\text{ref}} \gg c(\mathbf{x})$
- By default, k-Wave sets $c_{\text{ref}} = \max(c(\mathbf{x}))$
- If $c_{\text{ref}} < \max(c(\mathbf{x}))$, the simulation is not unconditionally stable, but the phase error will never be worse than the leapfrog PSTD method ($\kappa \rightarrow 1$ as $c_{\text{ref}} \rightarrow 0$)
- For a linear and lossless simulation and $c_{\text{ref}} < \max(c(\mathbf{x}))$, the stability criterion is given by



$$\Delta t \leq \frac{2}{c_{\text{ref}} k_{\text{max}}} \sin^{-1} \left(\frac{c_{\text{ref}}}{\max(c(\mathbf{x}))} \right)$$

Choice of reference sound speed

- The reference sound speed c_{ref} is specified by `medium.sound_speed_ref`
- Can be set to a scalar value (in m/s), or `'max'` (the default) , `'min'` , or `'mean'`
- To minimise dispersion, a good choice is the sound speed of the background medium

3. Heterogeneous media

- Consider the propagation of a pulse through a step change in material properties
- When the pulse crosses the interface, the pressure is influenced by the band-limited interpolant of the step





3. Heterogeneous media

- The accuracy of the reflection and transmission coefficients depends on:
 1. The size of the impedance change
 2. The number of points per wavelength (PPW)
 3. The size of the time step / CFL



Effect of points per wavelength

- Step change of 10% in sound speed and density





Effect of points per wavelength

- Step change of 100% in sound speed and density





Effect of time step / CFL

- There is also a weak dependence on the size of the time step / CFL





Effect of time step / CFL

- There is also a weak dependence on the size of the time step / CFL





Effect of impedance change

- More PPW are needed for larger impedance changes





Effect of impedance change

- For impedance changes typically seen in biological tissue (including bone), the reflection and transmission coefficients are encoded accurately





Sound speed vs density

- Sound speed change causes the transmission coefficient to be **underestimated**
- Density change causes the transmission coefficient to be **overestimated**





Extreme impedance changes – Gold

- Water \rightarrow Gold
- $(1500 \text{ m/s}, 1000 \text{ kg/m}^3) \rightarrow (3240 \text{ m/s}, 19300 \text{ kg/m}^3)$





Extreme impedance changes – Air

- Water \rightarrow Air
- $(1500 \text{ m/s}, 1000 \text{ kg/m}^3) \rightarrow (343 \text{ m/s}, 1.225 \text{ kg/m}^3)$



Extreme impedance changes – Air

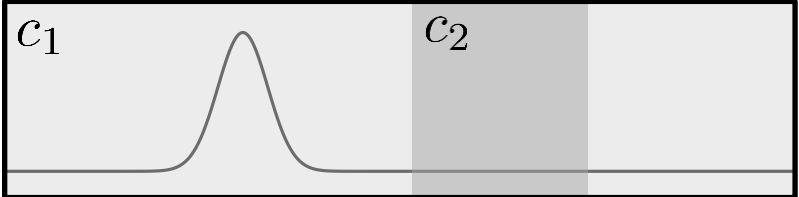
- Water \rightarrow Air
- $(1500 \text{ m/s}, 1000 \text{ kg/m}^3) \rightarrow (343 \text{ m/s}, 10 \times 1.225 \text{ kg/m}^3)$



- Trick: using a **lower density change** can reduce the error
- $(R = -0.999 \rightarrow -0.994 \text{ and } T = 0.000560 \rightarrow 0.00559)$


Maximum supported frequency

- A heterogeneous medium will have a different maximum supported frequency in different regions of the grid

$$f_{\max,1} = \frac{c_1}{2\Delta x} \quad \begin{array}{|c|c|c|} \hline c_1 & c_2 & \\ \hline \end{array} \quad f_{\max,2} = \frac{c_2}{2\Delta x}$$


The diagram shows a horizontal rectangular grid divided into three sections. The first section on the left is labeled c_1 and contains a smooth, bell-shaped wave pulse. The middle section is shaded gray and labeled c_2 . The third section on the right is also labeled c_2 . To the left of the grid is the formula $f_{\max,1} = \frac{c_1}{2\Delta x}$, and to the right is $f_{\max,2} = \frac{c_2}{2\Delta x}$.


- This can lead to very strange effects!




Comparison with finite differences

- These errors are not unique to the k-space pseudospectral method or k-Wave

Spectral Gradient




Finite Difference Gradient



- Errors for finite difference method are typically worse

4. Source staircasing


- Continuous surface must be represented on a discrete Cartesian grid
- Leads to errors in the **density** and **position** of source points



- Can be solved using **off-grid sources** (discussed earlier)



5. Medium staircasing

- Medium staircasing arises for the same reason as source staircasing
- Can be one of the largest sources of error for methods based on fixed Cartesian meshes
- Example: scattering of a 3-cycle tone-burst by a 2D cylinder created using `makeDisc`




Medium staircasing example

- Cylinder contrast of $c_{\text{cyl}} = 2c_0$ and $\rho_{\text{cyl}} = 2\rho_0$
- Reference simulation shown at 50 PPW





Convergence

- Error converges as the number of PPW is increased



Source smoothing



- Smoothing the medium can reduce spurious reflections from staircased edges



- However, the impedance change is no longer sharp

Source smoothing



- Smoothing the medium can reduce spurious reflections from staircased edges



- However, the impedance change is no longer sharp

Source smoothing

- Cylinder contrast of $c_{\text{cyl}} = 2c_0$ and $\rho_{\text{cyl}} = 2\rho_0$
- Smoothed solution is not accurate at low PPW



6. Acoustic absorption

- The power law acoustic absorption term $L\rho$ is given by

$$L = \tau \frac{\partial}{\partial t} (-\nabla^2)^{\frac{y}{2}-1} + \eta (-\nabla^2)^{\frac{y+1}{2}-1}$$

- The fractional Laplacian terms can be easily calculated in the Fourier domain using

$$\mathcal{F} \left\{ (-\nabla^2)^a \rho \right\} = k^{2a} \mathcal{F} \{ \rho \}$$

- Which gives

$$L\rho^{n+1} = \tau \mathcal{F}^{-1} \left\{ k^{y-2} \mathcal{F} \left\{ \frac{\partial \rho^{n+1}}{\partial t} \right\} \right\} + \eta \mathcal{F}^{-1} \left\{ k^{y-1} \mathcal{F} \left\{ \rho^{n+1} \right\} \right\}$$

- This still leaves $\partial \rho / \partial t$ to be discretised

Acoustic absorption



- This can be solved using the substitution $\partial\rho/\partial t = -\rho_0\nabla\cdot\mathbf{u}$

$$\begin{aligned} L\rho^{n+1} = & -\tau \mathcal{F}^{-1} \left\{ k^{y-2} \mathcal{F} \left\{ \rho_0 \sum_{\xi} \frac{\partial}{\partial \xi} u_{\xi}^{n+\frac{1}{2}} \right\} \right\} \\ & + \eta \mathcal{F}^{-1} \left\{ k^{y-1} \mathcal{F} \left\{ \rho^{n+1} \right\} \right\} \end{aligned}$$

- However, the velocity and density being added are at different time points due to grid staggering
- This leads to numerical dispersion
- Note: the absorption term in `kspaceSecondOrder` is exact

Convergence

- Accuracy of absorption and dispersion values for $0.1 f^{1.9}$, $0.25 f^{1.5}$, and $0.5 f^{1.1}$



Restrictions

- The absorption is specified by `medium.alpha_coeff` and `medium.alpha_power`
- The implementation of the absorption term in the Fourier domain means `alpha_power` must be scalar (a single value used for the whole domain)
- Select a reference frequency and `alpha_power`, then recompute `alpha_coeff` values

$$\alpha_{0,\text{new}} = \alpha_0 f_{\text{ref}}^{y - y_{\text{ref}}}$$

7. Nonlinearity

- The **convective nonlinearity** in k-Wave is implemented using spatial derivatives, and the **material nonlinearity** doesn't have derivatives, so no finite difference error

$$\frac{\partial \rho}{\partial t} = -\rho_0 \nabla \cdot \mathbf{u} - 2\rho \nabla \cdot \mathbf{u}$$

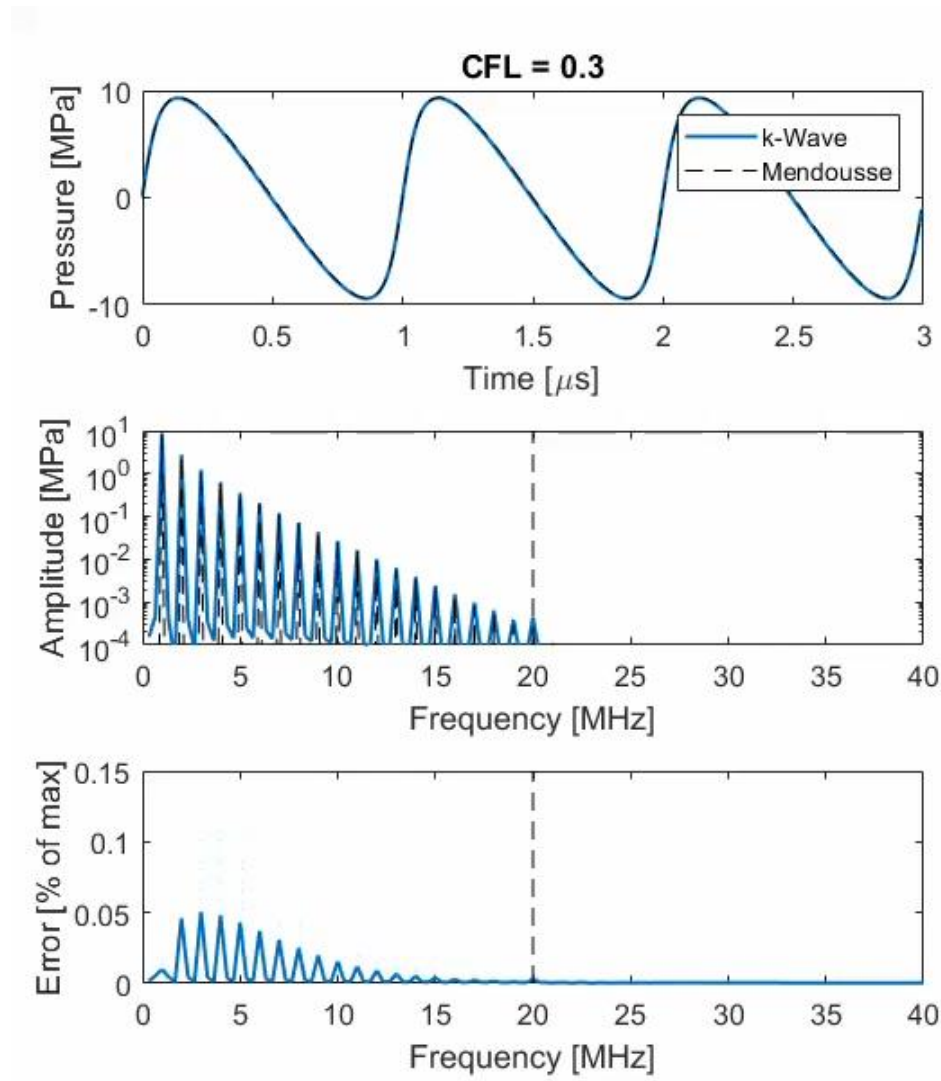
mass conservation

$$p = c_0^2 \left(\rho + \frac{B}{2A} \frac{\rho^2}{\rho} - L\rho \right)$$

pressure density


- However, the k-space correction is not exact when nonlinearity is included, so some error is introduced

Convergence with CFL



Convergence with CFL

- Reducing the time step increases the accuracy of the harmonic amplitudes




Spectral blocking

- For nonlinear simulations, the grid spacing must be chosen to give at least **2 PPW at the highest harmonic** at which there is significant energy

```
% compute dx and Nx based on x_size and f_max  
ppw = 3;  
dx   = c0_min / (ppw * f_max);  
Nx   = round(x_size / dx);
```


- Otherwise, energy that would be transferred to higher frequency harmonics is **aliased** to lower frequencies
- This is sometimes called **spectral blocking**, and causes a distortion of the wave and build up of error

Spectral blocking example



Spectral blocking example

- Increasing the shock parameter pushes energy into frequencies not supported by the grid
- Once this happens, there is a rapid increase in the error



- Not an intrinsic problem – just increase grid size


Real-world example of spectral blocking





Source properties

- 30 mm diameter
- 30 mm ROC
- 1 MHz CW source
- Source amplitudes from 0.1 to 1.2 MPa

Computational parameters





- Axisymmetric code
- Grid size: 3456 x 1372 (100 harmonics)
- Memory usage: 1 GB







Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**

100 kPa





200 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





300 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





400 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





500 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





600 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





700 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





800 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





900 kPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**





1 MPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**

1.1 MPa




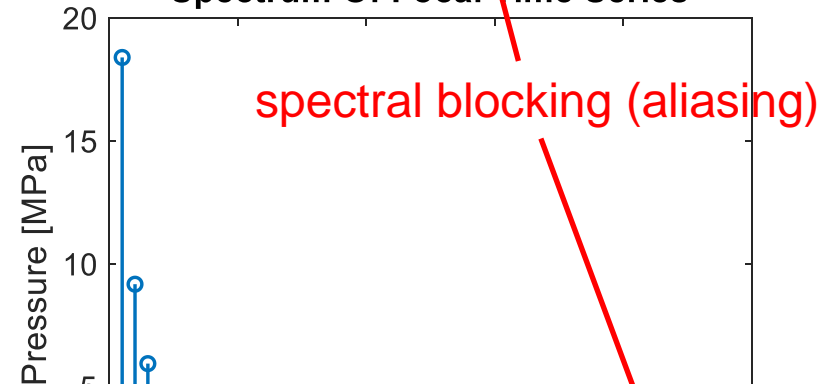
Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**

1.2 MPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**

spectral blocking (aliasing)

1.2 MPa

Axial Peak Pressure**Time Series At Spatial Peak P₊****Axial Peak Pressure At First 10 Harmonics****Spectrum Of Focal Time Series**

Spectral blocking can be eliminated by increasing the grid size
(always run a convergence test)

Summary

Factors affecting accuracy:

1. Perfectly matched layer:
2. Numerical dispersion in heterogeneous media
3. Accuracy of reflection and transmission coefficients
4. Source staircasing
5. Medium staircasing
6. Acoustic absorption
7. Acoustic nonlinearity

Δx	Δt
✓	
	✓
✓	~
~	
✓	
✓	✓
✓	✓

How can we trust simulation results aren't influenced by numerical parameters?



Perform a convergence test!




Challenge to modellers:

Repeat your simulation with reduced Δx and Δt
Do you get the same answer?

Problem sizes

- Many simulations involve very large domain sizes



- 20 cm $\approx 133 \lambda$ at 1 MHz, 1333 λ at 10 MHz
- In 3D, using 3 PPW, each matrix is 240 GB!

Target systems


- There are now several optimised versions of k-Wave written in C++
- These are replacements for `kspaceFirstOrder3D`
- Three target systems:



1. Desktop Computer



2. Graphics Processing Unit (GPU)



3. High Performance Computing (HPC) Cluster

Running the C++ codes

- MATLAB and C++ codes are linked by input and output files stored to disk in HDF5 format



- This means MATLAB is not required to run the simulation

Running the C++ codes

Two options to run the C++ code:

1. Use `kspaceFirstOrder3DC` (C++ / OpenMP) and `kspaceFirstOrder3DG` (C++ / CUDA) to call the code blindly (automatically save the input file, run the simulation, and load the outputs)
2. Save the input file using the '`SaveToDisk`', `Filename` option, then run the C++ code from the command line (windows) or terminal (linux)

More details given in the k-Wave manual

Assessing memory usage

- Memory can be calculated by

$$\text{memory usage [GB]} \approx \frac{(13 + A) N_x N_y N_z + (7 + B) \frac{N_x}{2} N_y N_z}{1024^3 / 4} \\ + \text{input} + \text{output}$$


- Here $A = [0, 8]$,
 - 1 if c_0 is heterogeneous
 - 4 if ρ_0 is heterogeneous
 - 1 if B/A is given and heterogeneous
 - 2 if absorbing and c_0 or α_0 are heterogeneous
- Here $B = [0, 2]$, 0 if the medium is lossless, or 2 if the medium is absorbing

Assessing memory usage

- Example grid sizes and total memory usage:
 - $128^3 \rightarrow 200 \text{ MB}$ Desktop / GPU
 - $256^3 \rightarrow 1.6 \text{ GB}$ Desktop / GPU
 - $512^3 \rightarrow 13 \text{ GB}$ Desktop / GPU
 - $1024^3 \rightarrow 100 \text{ GB}$ Server
 - $2048^3 \rightarrow 0.80 \text{ TB}$ Cluster
 - $4096^3 \rightarrow 6.4 \text{ TB}$ Cluster
- The number of time steps also increases as the grid size increases (e.g., $N_t \approx 5 * N_x$)





Example of production simulation

- Nonlinear simulations performed using the [Visible Human Female Data Set](#) and the open-source [AustinWoman](https://tiny.cc/AustinWoman) segmentation (tiny.cc/AustinWoman)




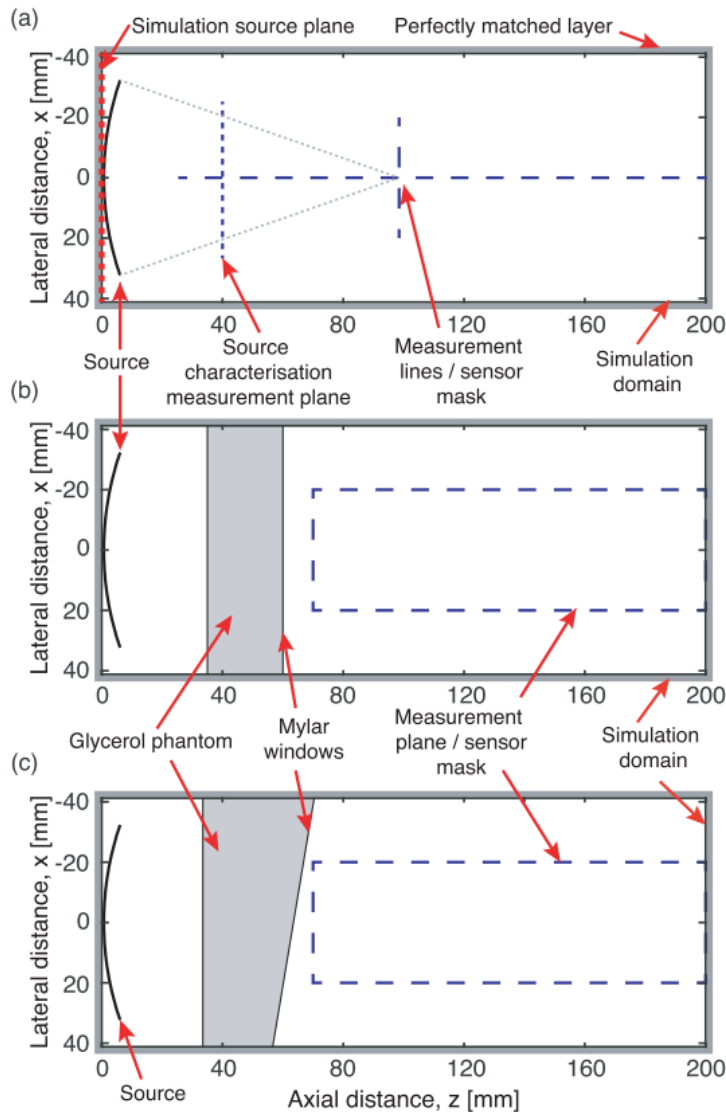


Inferior pole
of left kidney
(90 mm deep)

Example of production simulation




Experimental validation: Nonlinear




Martin et al, IEEE TUFFC, 2019

Experimental validation: Bone phantoms

- Bone phantoms with realistic cortical properties

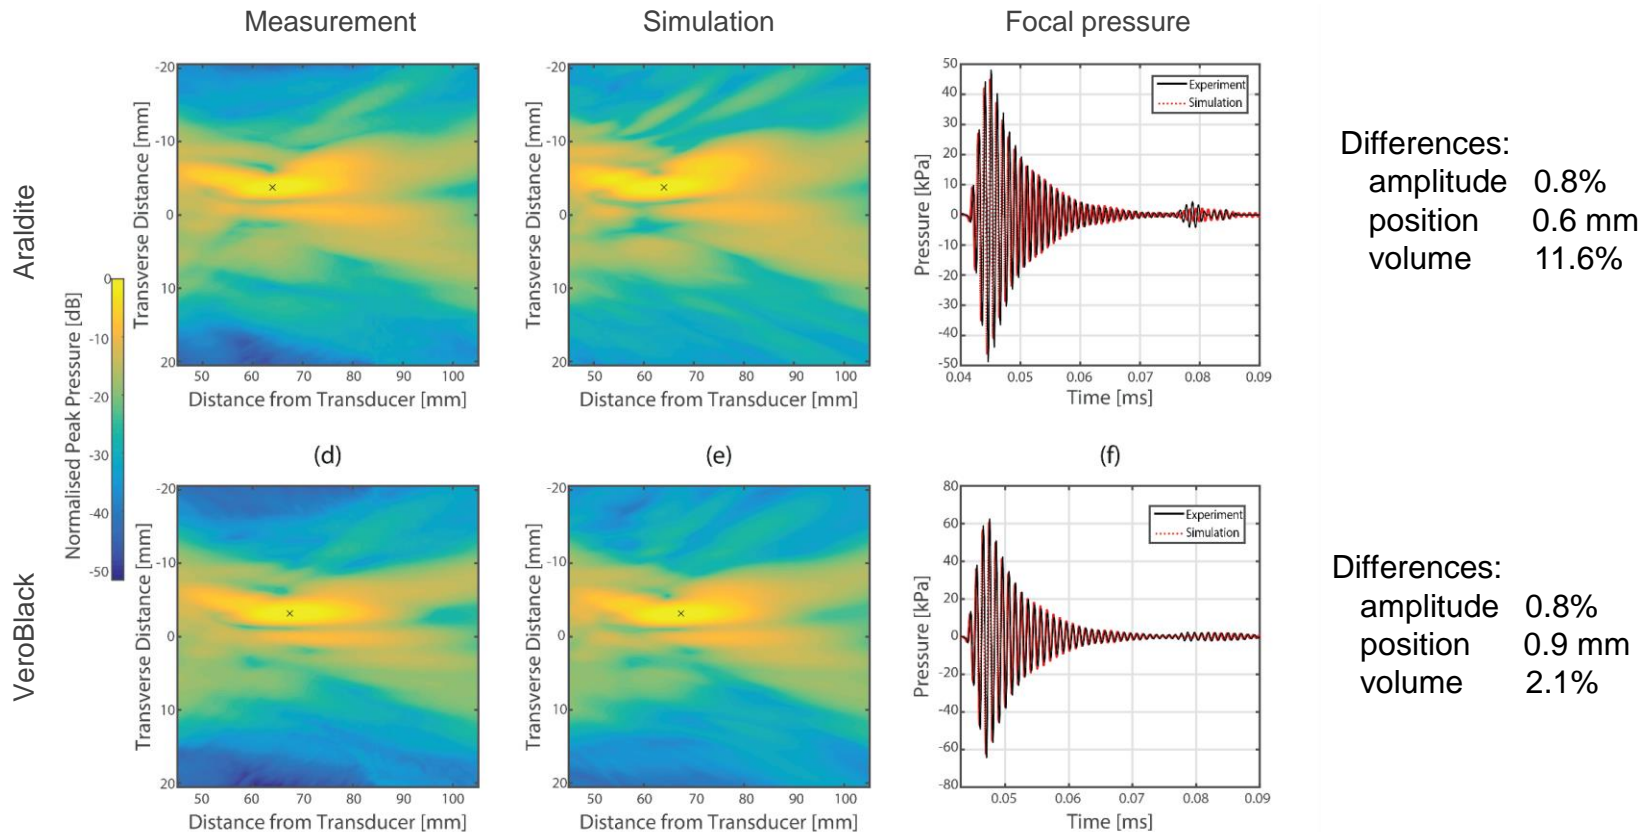


	Araldite	VeroBlack
compressional speed [m/s]	3008	2495
shear speed [m/s]	1562	1081
density [kg/m ³]	1637	1180
absorption [dB/cm at 1 MHz]	1.8	3.7



Experimental validation: Bone phantoms

- With known geometry and properties, measurement and model agree quantitatively



1 MHz, spherically focused transducer

Closing remarks

- k-Wave is a powerful tool and widely used, but not the best tool for every job!
- We are just about to start work on k-Wave II
 - Object oriented and extensible
 - Differentiable models and deep learning
 - Model coupling
 - MATLAB front-end will also be compiled as a Python package
- Codes and slides available from:
<http://www.k-wave.org/downloads/isna2022.zip>