



Accelerating Two-Dimensional k-Wave Ultrasound Simulations Through Pruned FFT: A Treatment Planning Optimisation

Ondrej Olsak^(✉) , David Bayer , and Jiri Jaros 

Faculty of Information Technology, Brno University of Technology,
Brno, Czech Republic
iolsak@fit.vut.cz

Abstract. Wave propagation simulations are foundational tools across scientific and medical applications, yet their computational demands become significant for high-resolution simulations, particularly in medical applications where precise representation of different tissue geometries is crucial. This paper presents a novel approach to accelerate 2D wave propagation simulations in the k-Wave toolbox. Our method focuses on optimising Fourier transform computations through spectrum pruning. The Acoustic Field Propagator along with a bisection pruning algorithm to estimate the position of the spectral coefficients is used. Through these optimisations, our approach achieves significant performance gains, demonstrating speedups of up to 1.8x for large simulation domains. Experimental evaluation on medical ultrasound simulations demonstrates that the proposed method achieves focal point errors below 1% with minimal focus position shifts, while skipping up to 90% of spectral coefficients in large domains. This results in a significant simulation time reduction by half over the large simulation domains. Although the proposed method primarily focuses on accelerating k-Wave toolbox wave propagation simulation, it could be generally applied to wave propagation problems.

Keywords: Pruned Fast Fourier Transform · Ultrasound Simulation · Wave propagation simulation · k-Wave · Pseudo-spectral methods

1 Introduction

Wave propagation simulations play a crucial role in various fields of science and engineering. However, the time required for these simulations to complete can vary considerably, from a few seconds to days or even longer, depending on the complexity of the model and the level of resolution required. This paper focuses on accelerating 2D wave propagation simulations implemented by the k-Wave toolbox [19], which utilises *k*-space pseudo-spectral methods based on the Fourier Transform. The goal is to accelerate these simulations by optimising the computation of the Fourier Transform, which represents a significant part

of the simulation. This optimisation is particularly crucial for simulations over high-resolution domains, where computational demands can be significant.

The most common approach to performing the Fourier transform is the Fast Fourier Transform (FFT) algorithm [2]. Since some applications only require a specific subset of frequency components from the FFT algorithm's output, computing the complete set of spectral coefficients may be unnecessary. In such cases, it may be beneficial to reduce the computational cost by implementing either a Sparse Fourier Transform (sparse FFT) [13] or Pruned Fourier Transform (pruned FFT) [7] algorithm. The sparse FFT is suitable for signals with few non-zero/significant coefficients compared to the size of the input signal, where their position in the spectrum is unknown. Sparse FFT algorithms are usually specially designed using domain-specific knowledge. Conversely, pruned FFT algorithms are optimised to compute spectral coefficients within known patterns that occur in the spectral domain. This leads to bypassing unnecessary computations in the standard Fast Fourier Transform [14].

For wave propagation simulations, spectral coefficients typically cluster near low frequencies [10]. This characteristic makes the pruned FFT particularly suitable as a replacement for the standard FFT currently employed in the k-Wave toolbox. By implementing this change, we aim to accelerate wave propagation simulations through reduced computational time in the Fourier transform phase and subsequent operations in the spectral domain.

Given that k-Wave is designed for time domain acoustic and ultrasound simulations in complex and tissue-realistic media, the evaluation of the proposed approach will focus on simulations used for non-invasive treatment pre-planning, such as the application of focused high-intensity ultrasound. This application exemplifies situations where precise representation of the media geometry can affect the accuracy of the simulation result [12, 21]. Our evaluation will examine the method's impact on three critical aspects across various domain resolutions: simulation accuracy, focal point positioning and computational efficiency.

2 Mathematical Background

To compute the wave propagation, the k-Wave toolbox employs the pseudo-spectral method using Fourier basis functions. This technique involves representing the solution of the differential equation as a sum of specific basis functions. Unlike finite-difference time domain methods, which rely on local computations at neighbouring points, spectral methods use information from the entire domain, leading to higher accuracy [5]. The k-Wave toolbox runs simulations based on the following governing equations [19]:

$$\begin{aligned}
 \frac{\partial u}{\partial t} &= -\frac{1}{\rho_0} \nabla p \\
 \frac{\partial \rho}{\partial t} &= -(2\rho + \rho_0) \nabla \cdot u - u \cdot \nabla \rho_0 \\
 p &= c_0^2 (\rho + d \cdot \nabla \rho_0 + \frac{B}{2A} \frac{\rho^2}{\rho_0} - L_\rho)
 \end{aligned} \tag{1}$$

Where u is the acoustic particle velocity, p is the acoustic pressure, ρ is the acoustic density, ρ_0 is the ambient (equilibrium) density, c_0 is the isentropic sound speed, d is the acoustic particle displacement, B/A is the nonlinearity parameter characterizing finite-amplitude effects, and L_ρ is a loss operator that accounts for acoustic absorption and dispersion in the medium. Equation (1) can be written in a discrete form using the k -space pseudo-spectral method [16]. The following equations are part of the spatial gradient calculations based on the Fourier collocation spectral method:

$$\frac{\partial}{\partial \xi} p^n = \mathcal{F}^{-1} \{ i k_\xi \kappa e^{i k_\xi \Delta \xi / 2} \mathcal{F} \{ p^n \} \}, \quad (2)$$

$$u_\xi^{n+\frac{1}{2}} = u_\xi^{n-\frac{1}{2}} - \frac{\Delta t}{\rho_0} \frac{\partial}{\partial \xi} p^n + \Delta t S_{F\xi}^n, \quad (3)$$

$$\frac{\partial}{\partial \xi} u_\xi^{n+\frac{1}{2}} = \mathcal{F}^{-1} \{ i k_\xi \kappa e^{-i k_\xi \Delta \xi / 2} \mathcal{F} \{ u_\xi^{n+\frac{1}{2}} \} \}, \quad (4)$$

$$\rho_\xi^{n+1} = \rho_\xi^n - \Delta t \rho_0 \frac{\partial}{\partial \xi} u_\xi^{n+\frac{1}{2}} + \Delta t S_{M\xi}^{n+\frac{1}{2}} \quad (5)$$

For the Cartesian direction $\xi = x, y$ in R^2 , \mathcal{F} and \mathcal{F}^{-1} denote the forward and inverse spatial Fourier transform, i is the imaginary unit, Δt is the time step, k_ξ represents the wave numbers in the ξ direction, and κ is the k -space operator defined as $\kappa = \text{sinc}(c_{ref} k \Delta t / 2)$, where c_{ref} is a scalar reference sound speed. Equations (2) and (4) are spatial gradient calculations based on the Fourier collocation spectral method. Equations (5) and (3) represent update steps utilising a k -space corrected first-order accurate forward difference.

The Fast Fourier Transform algorithm is used to convert signals from the spatial domain to the spectral domain. Each simulation step of the 2D wave propagation simulation involves 11 FFTs. This computation consumes approximately 60% of the total simulation time, making it a significant part of the overall simulation [6].

3 Transducer Position

Since most simulations utilise narrow bandwidth sources, the spectral coefficients cluster around low frequencies, leading to sparsity in the spectral domain, particularly in high-resolution simulation domains [3]. Despite this apparent sparsity, sparse FFT approaches prove unsuitable as they discard small but crucial coefficients that contribute to simulation accuracy, while also introducing additional computational complexity through filtering processes. The spectrum is also not sparse enough to benefit from the usage of sparse FFT. In contrast, pruned FFT offers a more efficient solution by precisely computing the specified region within the spectral domain, preserving all coefficients regardless of their magnitude and eliminating the need for additional computational steps. This makes pruned FFT

a more reliable choice for processing simulation data compared to sparse FFT methods [10].

One key factor that significantly impacts the position of computed spectral coefficients is the direction in which the wave propagates through the media [4]. Thus, the position of the coefficients in the spectrum (among other factors) depends on the orientation of the transducer that determines the direction of wave propagation. Figure 1 shows the wave propagation in water with different positions of an arc transducer, demonstrating different directions of ultrasound wave propagation. The direction of wave propagation clearly impacts the position of the coefficients in the spectrum (the zero-frequency components are shifted to the centre of the domain).




Fig. 1. Transducer position’s effect on wave propagation and spectral domain: (a) pressure field distribution for Y-axis propagation in space domain, (b) frequency domain representation of Y-axis propagation showing spectral coefficient distribution, (c) pressure field distribution for X-axis propagation in space domain, and (d) frequency domain representation of X-axis propagation showing spectral coefficient distribution.

For the application of the pruned FFT in wave propagation simulation, the most suitable transducer positions are those aligned with one of the axes. This alignment enables reduction of the area computed by the pruned FFT algorithm [11]. If the transducer is not aligned with either axis, alignment can be achieved by rotating the domain around its centre. To obtain a suitable shape for the simulation domain, which is typically rectangular or square, the rotated domain can be filled with surrounding media. At the end of the simulation, if needed, the domain with the wave propagation result can be rotated back to its original position and cropped to its original size. This operation requires the ability to fill the domain with surrounding media and the presence of an absorbing layer, such as Perfect Match Layer (PML) [19] adjacent to the rotated domain. This layer absorbs the propagated wave around the simulation domain and prevents possible reflections that might affect the result of the original simulation. In our proof-of-concept implementation of the two-dimensional pruned FFT algorithm, the computation reduction is made only in the second dimension (columns - X). This means that over the first dimension (rows - Y), the full FFT is computed over each row. In the second dimension, only a given number of columns is computed. For our implementation, the position of the transducer aligned with

the X-axis is the most suitable variant. However, the transducer can also be aligned with the Y-axis. In this case, the suitable reduced dimension would be Y (rows). This scenario is in actual implementation solved by rotating the domain by 90° to align the transducer with the X-axis.

4 Estimation of the Coefficient Area

An integral part of the pruned FFT algorithm is localising the area containing coefficients crucial for computing the wave propagation simulation. Most of these coefficients represent low frequencies, thus they are present at the corners of the spectral domain. Since the symmetry of the real-to-complex FFT in the Y dimension is used, only the first m columns of the X domain will be computed. To estimate the value of m , the Acoustic Field Propagator (AFP) [17] together with a bisection pruning algorithm is used.

The AFP enables the calculation of the wave field at all spatial positions at a given time in a single step. The advantage of AFP is its computational speed compared to a full wave propagation simulation. However, this method can only be used in homogeneous media with a single-frequency transducer and cannot compute the reflections and absorption of the propagated wave. Thus, it cannot replace a full ultrasound wave propagation simulation. When complex tissue interfaces create significant reflections or scattering, the AFP may fail to capture high-frequency components generated at these boundaries. Despite these limitations, it provides an acceptable estimation of spectral coefficient positions. The AFP is executed with the same transducer and homogeneous media, with a sound speed equal to the minimum value present in the original simulation. The lower the sound speed, the higher the frequencies that may occur in the spectral domain. The resulting spectrum of the propagated wave is used to estimate the position of spectral coefficients that will be computed by the pruned FFT in the wave propagation simulation.

To determine the first m columns for the pruned FFT, a bisection pruned algorithm is employed. In this algorithm, the dimension X is considered as an interval $[0, N/2]$, to find an optimal cutoff point. First, the AFP spectrum is shifted so that low frequencies are centred. Then, bisection is applied symmetrically to both halves along the X-axis. In each iteration, a midpoint $m = (lower + upper)/2$ is computed. Coefficients below m are temporarily set to zero, preserving only the spectral information in the interval $[m, N/2]$, after which an inverse FFT is performed. Due to the symmetrical properties of the FFT, this cutoff has a corresponding effect on the right half of the full domain $[N/2, N]$, creating a mirror image of the preserved region. The resulting spatial domain is compared to the original. Based on a user-defined error threshold, the iteration repeat with upper half of the interval if the error is below the threshold or lower half interval if it is above. The algorithm terminates when the border position stabilizes.

Three error thresholds were considered for the bisection pruning algorithm: Mean Absolute Percentage Error (MAPE), Root Mean Squared Percentage Error

(RMSPE), and Normalised Percentage L_∞ Error. However, RMSPE proved unsuitable due to its quadratic nature, which led to irregular error changes during the threshold search and limited spectral coefficient reduction. Table 1 compares the number of coefficients skipped by the pruned FFT algorithm using MAPE and Normalised Percentage L_∞ Error in a simulation with an arc transducer in water aligned along the X-axis. Note that the L_∞ error reflects the spatial domain error after applying the pruned FFT, excluding transducer-related discrepancies.

Table 1. Comparison of error and skipped rows/columns under different levels of Mean Absolute Percentage Error (MAPE) and Normalised Percentage L_∞ error in homogeneous domain with edge size of 1024.

	MAPE					Norm. Perc. L_∞				
	10%	20%	30%	40%	50%	1%	2%	3%	4%	5%
Rows skip	1	1	614	800	866	670	810	856	886	906
Rows skip [%]	0.10	0.10	59.96	78.12	84.57	65.43	79.10	83.59	86.52	88.48
L_∞ error [%]	0.003	0.003	0.005	0.054	0.312	0.005	0.072	0.240	0.498	0.768

The MAPE appears to yield better accuracy in the simulation results by computing significantly more coefficients than the Normalised Percentage L_∞ Error. However, the MAPE was found inadequate as it fails to accommodate zero values that may be encountered in spatial analysis. In contrast, the L_∞ Error provides a more reliable measurement by focusing on the maximum difference, without being influenced by the distribution of smaller errors. Additionally, this error metric is more intuitive for potential users, as it represents the maximum error occurring at a single grid point in the entire domain, making it easier to adjust based on specific needs. Thus, the Normalised Percentage L_∞ Error is more suitable for area estimation and will be used in all subsequent experiments.

The second method considered for area estimation involved computing the norm over each column of the spectrum and skipping the columns where the norm is below a given threshold. However, this approach has two main disadvantages compared to the bisection method. First, the resulting area of spectral coefficients may not be continuous. In contrast, when the estimated area is continuous, the pruned FFT result can be stored in a reduced-size matrix (with one dimension equal to the original size), which benefits element-wise operations and memory optimisation, especially for large domain simulations. Second, the preprocessing time is higher. The bisection pruning algorithm has a time complexity of $O(\log N)$, whereas the norm computation has a time complexity of $O(N)$ (with N representing half the number of columns, due to the spectral domain's symmetry). Moreover, when considering the computation of forward and backward FFTs to determine the threshold error for the AFP, the preprocessing time becomes significant. Figure 2 illustrates the experimental pipeline used to evaluate the proposed algorithm. The red dashed rectangle highlights the

preprocessing stage, consisting of the AFP execution followed by the bisection pruning.




Fig. 2. The pipeline with preprocessing operations to estimate area in the spectral domain and the comparison of the reference and modified simulation.

5 Implementation

All principles described in the previous section result in a proof-of-concept version of the two-dimensional k-Wave wave propagation simulation. The original CUDA implementation of k-Wave was modified to incorporate the pruned FFT algorithm instead of the FFT.

The pruned FFT was implemented using the cuFFT [9] and VkFFT [15] libraries. The computation is divided into two stages: first, a full real-to-complex FFT along the X-axis using cuFFT, and second, a Y-axis FFT using VkFFT. Although cuFFT could perform both, VkFFT offers better flexibility for complex memory layouts—beneficial for future 3D extensions. Because spectral elements with an X coordinate greater than m are zero, only first m Y-axis FFTs are executed. The inverse FFT reverses this process: starting with a Z-axis transformation, then a Y-axis FFT for the first m X slices with VkFFT, followed by zeroing coefficients in the remaining X slices to eliminate artifacts from the initial transform, and finally an inverse X-axis transform with cuFFT.

To evaluate performance and error, the standard FFTs in the k-Wave toolbox were replaced by our pruned FFT in the wave propagation simulation. This approach accelerates computations for acoustic pressure, velocity gradients, and the absorption term during each simulation step. Additionally, by confining non-zero coefficients to a single spectral area, we can optimize element-wise matrix multiplications (highlighted in Eqs. 2 and 4), potentially further reducing simulation times, while leaving CUDA kernels for real domain operations unchanged.

6 Evaluation of the Method over Real Data

In this section, we evaluate the modified version of the k-Wave wave propagation algorithm that utilizes the pruned FFT. To simplify the experimental setup and focus on performance and computational error, all experiments align the wave




Fig. 3. Ultrasound simulation setups and results: (a) human skull configuration, (b) human liver configuration, (c) skull acoustic pressure distribution, and (d) liver acoustic pressure distribution. Orange rectangle shows error measurement area; green arc indicates transducer position. (Color figure online)

propagation direction with the X-axis, eliminating the need for domain rotation and additional preprocessing.

These examples represent practical clinical scenarios where precise targeting of specific locations within the human body with focused ultrasound waves is essential. Such applications are particularly relevant in therapeutic procedures that rely on ultrasound focusing techniques, allowing us to evaluate the algorithm’s performance and precision in contexts closely mirroring real-world treatments.

In the search for the optimal focus position, multiple simulations may be required. To enhance computational efficiency, a two-stage approach is employed: initially, accelerated simulations with an acceptable error margin are performed to identify promising transducer positions. Once an approximate optimal position is determined, a full simulation is executed using that position to ensure accuracy and reliability. This hybrid strategy significantly reduces the overall computational time during the transducer position search while maintaining the necessary precision for medical applications [12].

Furthermore, high simulation resolution is crucial not only in medical applications but also in accurately representing diverse material geometries, especially for structures with large differences in material properties. A higher domain resolution helps prevent stair-casing artifacts and phase shifts [12,21]. Conse-

Table 2. Properties of the skull simulation across different domain resolutions.

	1x	2x	4x	8x	16x	32x
Nx	288	576	1152	2304	4608	9216
Ny	384	768	1536	3072	6144	12288
dx/dy [m]	9.375e-4	4.6875e-4	2.34375e-4	1.17187e-4	5.85937e-5	2.92969e-5
CLF	0.3	0.3	0.3	0.15	0.1	0.05
PPW (water)	5.36	10.73	21.46	42.92	85.85	171.69
Time steps	2798	5595	11189	44753	134258	537031

quently, high-resolution simulations tend to contain many more zero or negligible spectral coefficients, making them particularly suitable for the proposed spectrum pruning.

In all subsequent experiments, the domain properties—sound speed, density, and absorption coefficients—reflect those of real tissue. Figure 3a shows the transducer’s position relative to the skull, while Fig. 3b illustrates its position relative to the liver. An additive transducer operating at 300 kHz with an amplitude of 100 kPa was used [20]. The same figures indicate the area within which acoustic pressure error is evaluated. Although the pruned FFT removes high frequencies—resulting in significant error at the transducer itself—the primary concern in ultrasound applications is the error within the tissue. This approach for assessing the accuracy of nonlinear wave propagation in layered, absorbing fluid media follows the method used in [8]. All experiments were executed on an NVIDIA RTX A5000.

Table 3. The measurements in the human skull with different bisection threshold.

	1% Threshold					
	1x	2x	4x	8x	16x	32x
Skip [%]	22.92	43.40	63.02	77.78	87.41	93.51
L_∞ domain [%]	3.13	10.72	14.15	15.43	16.56	17.62
L_∞ focal point [%]	0.23	1.05	1.13	0.89	0.60	0.76
Focal point shift [mm]	0	0.469	0	0	0.059	0
Time original [s]	0.38	1.69	15.84	243.75	3636.13	54366.42
Time modified [s]	0.37	1.53	12.21	169.62	1916.76	29333.16
Step time original [ms]	0.136	0.296	1.386	5.325	26.48	98.97
Step time modified [ms]	0.132	0.268	1.069	3.706	13.96	53.40
AFP time [s]	0.23	0.74	2.62	9.10	32.34	42.34
Speedup	1.03	1.10	1.30	1.44	1.90	1.85
	2% Threshold					
Skip [%]	39.58	61.11	77.43	87.59	93.27	96.44
L_∞ domain [%]	5.63	13.82	17.92	20.23	21.67	22.69
L_∞ focal point [%]	0.67	1.21	1.30	1.14	1.19	0.93
Focal point shift [mm]	0	0.469	0	0	0.059	0
Time original [s]	0.38	1.68	15.87	243.67	3654.35	54374.44
Time modified [s]	0.35	1.43	11.93	166.06	1847.03	28782.14
Step time original [ms]	0.136	0.294	1.389	5.324	26.61	98.98
Step time modified [ms]	0.125	0.250	1.044	3.628	13.45	52.39
AFP time [s]	0.26	0.62	2.66	9.47	33.25	174.20
Speedup	1.09	1.17	1.33	1.47	1.98	1.89

To measure the impact of the optimisation on the differently sized simulations, the resolution of the original simulation domains was upsampled using *nearest neighbour* approximation to maintain the domain in its segmented form. This ensures that no artificial material properties are introduced during the upscaling process. To show the impact of the bisection threshold on the simulation result, the measurements were made for 1% and 2% Normalised Percentage L_∞ bisection threshold. The simulation properties for both the skull and liver setup are presented in Tables 2 and 4 respectively.

Table 4. Properties of the liver simulation across different domain resolutions.

	1x	2x	4x	8x	16x
Nx	480	960	1920	3840	7680
Ny	480	960	1920	3840	7680
dx/dy [m]	3.33333e-4	1.66667e-4	8.33333e-5	4.16667e-5	2.08333e-5
CLF	0.3	0.3	0.3	0.3	0.15
PPW (water)	15.09	30.18	60.36	120.72	241.44
Time steps	5524	11047	22094	44187	176746

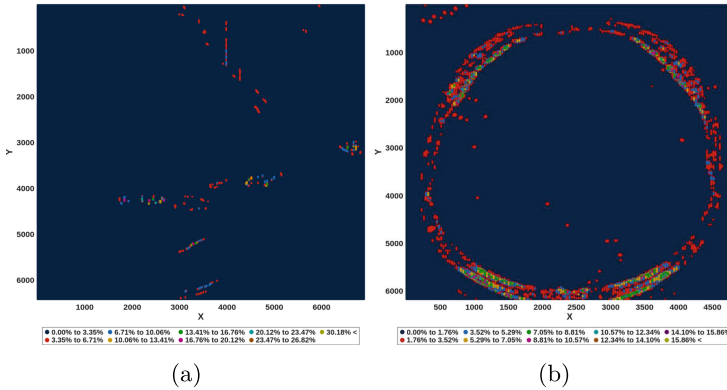
The original size of the simulation in skull was 288×384 with uniform grid spacing of 9.375×10^{-4} m. This simulation was upsampled up to 32 times while keeping the physical size the same as the original simulation. The original size of the liver simulation was 480×480 with uniform grid spacing of 3.333×10^{-4} m. This simulation was upsampled up to 16 times. Since absorption was present in all experiments, it was necessary to adjust the Courant-Friedrichs-Lewy (CFL) number to maintain simulation stability [18]. The CFL number affects the simulation time step, which can result in longer simulation times. The simulation time was chosen based on the time it takes the wave to travel from one corner of the grid to the geometrically opposite one.

When we examine the results in Table 3 (skull simulation) and Table 5 (liver simulation), we observe that the number of skipped spectral coefficients increases with the simulation domain resolution, reaching up to approximately 90%. The Fig. 4a and 4b show the error distribution of the normalized L_∞ error using a 1% bisection threshold and reveal that significant errors primarily occur at the tissue boundaries where sound speed and density change dramatically. Throughout most of the simulated media, errors remain manageable at just a few percentage points. As expected, increasing to a 2% bisection threshold leads to higher computation errors across all simulations, but also results in a more significant reduction in computed coefficients due to the greater loss of spectral information.

A crucial aspect for focused ultrasound procedures is the accuracy of the focal point and its position. With a 1% bisection threshold, the focal point error remains mostly below 1%. In skull simulations, the focal point shift is minimal, either zero or limited to a single grid point relative to the original domain size.

Table 5. The measurements in the human liver with different bisection threshold.

	1% Threshold				
	1x	2x	4x	8x	16x
Skip [%]	39.17	64.17	76.04	84.64	91.54
L_∞ domain [%]	10.84	20.49	25.25	29.78	33.53
L_∞ focal point [%]	0.47	0.99	0.86	0.05	2.07
Focal point shift [mm]	0	0.236	0.755	0.750	0.750
Time original [s]	1.03	6.75	72.35	537.40	9549.67
Time modified [s]	0.95	5.56	41.38	339.19	5106.51
Step time original [ms]	0.186	0.425	2.14	7.80	34.52
Step time modified [ms]	0.172	0.350	1.22	4.93	18.46
AFP time [s]	0.48	1.43	4.77	15.92	55.81
Speedup	1.08	1.21	1.75	1.58	1.87
	2% Threshold				
Skip [%]	59.58	75.42	85.00	91.41	95.47
L_∞ domain [%]	18.14	25.24	29.31	33.03	35.12
L_∞ focal point [%]	2.60	0.33	1.48	2.02	2.37
Focal point shift [mm]	0	0.850	0.755	0.750	0.750
Time original [s]	1.03	6.77	72.06	539.87	9554.37
Time modified [s]	0.92	5.40	40.13	323.92	4984.51
Step time original [ms]	0.186	0.426	2.13	7.84	34.54
Step time modified [ms]	0.167	0.340	1.19	4.70	18.02
AFP time [s]	0.48	1.32	4.60	16.69	94.15
Speedup	1.12	1.25	1.80	1.67	1.92

**Fig. 4.** The distribution of the Normalised L_∞ error over grid points of the final acoustic pressure distribution in (a) 16 times upscaled liver and (b) 32 times upscaled skull, grouped into 10 intervals.

In contrast, liver simulations can experience focal point shifts of several grid points due to the heterogeneous nature of bone tissue. Notably, simulation times improved significantly, especially for larger domains. With a 1% bisection error, a speed-up of up to 1.8 times was observed, reducing simulation time from roughly 15 to 8 h—a substantial saving when multiple simulations are required. However, when including the AFP preprocessing time, the pruned FFT’s benefits diminish for small domain sizes due to overhead, while for large domains, AFP accounts for only a negligible portion of the total time saved.

Profiling

The profiling of the proof-of-concept pruned FFT implementation was performed on the same GPU used for the experiments, employing the identical input dataset utilized in the resolution evaluations described in Sect. 6.

Figure 5a compares the overall time spent on FFT computations in both directions between the original and modified implementations. The results indicate that acceleration was achieved for all simulation sizes, with larger simulations exhibiting greater speedup. The only exception is the largest simulation, where a decrease in speedup may be due to less efficient FFT algorithm selection. This trend is easily explained by the decreasing percentage of coefficients processed in the Y dimension. As noted in Sect. 2, FFTs account for approximately 60% of the total simulation time. According to Amdahl’s law, with 60% of the computation optimisable, the maximum theoretical speedup is 2.5 times. Comparing this limit to the actual speedup shown in Fig. 5a highlights the high efficiency of the optimization.

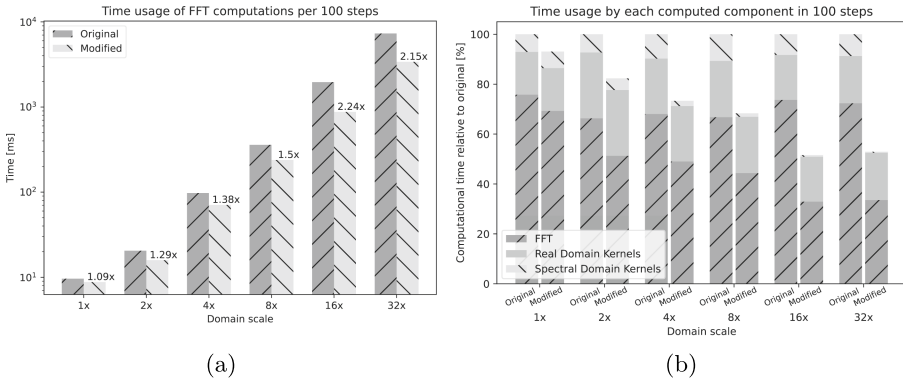


Fig. 5. Performance comparison between original and modified implementations: (a) FFT computation duration; (b) Time percentage breakdown of FFT and kernel computations.

Figure 5b illustrates the relative time usage of the original and modified implementations, divided into three segments: FFTs, real domain kernels, and

spectral domain kernels, with the overall duration normalized to that of the original implementation. The significant reduction in computation time is observed in the FFT segment, which decreases in accordance with the domain scale. The time spent on spectral domain kernels is also greatly reduced; for the largest simulation, it drops from approximately 7% of the overall simulation time to less than 1%. As expected, no speedup was observed for the real domain kernels, since the optimizations were applied solely to the spectral part of the computation.

7 Discussion

Comparing our implementation with other acoustic pressure wave solvers [1], experiments at various resolutions show computation errors within acceptable cross-comparison ranges. For human head simulations similar to ours, relative L_∞ error ranges from 10% to 100% compared to k-Wave. Focal point shifts (0–2mm) and acoustic pressure errors (10^{-2} – 10^1) also align with benchmarks. While direct comparison is challenging, the results achieved in the experiments presented here are promising and provide the first insight into the accuracy of this method over high-resolution simulations.

A limitation of this approach that requires further investigation is its performance in heterogeneous simulations where the difference between the properties of two media is so high that removing high frequencies from the spectrum prevents the wave from propagating or reflecting correctly. This can lead to significant errors at the boundaries of such media. Another limitation lies in the type of transducer used. For a piston transducer, the overall simulation error should be lower since most of the spectral coefficients representing this transducer are included in the computed part of the spectrum. However, if the piston transducer is unaligned, high computation errors may occur. In contrast, a point transducer is not suitable for this optimization, as accurately representing it requires the inclusion of high frequencies. The worst-case scenario represents the signals with widely distributed spectral content (sharp impulses or broadband noise). Similarly, simulation involving point source rather than distributed transducers may retain significant energy in high frequencies, limiting coefficient reduction.

There are several avenues for further improvement. Implementing the pruned FFT in the first dimension would reduce arithmetic operations and memory accesses. Applying the pruned FFT across all dimensions of the two-dimensional domain would allow storing the resulting spectral coefficients in reduced matrices that correspond to the computed area, potentially enhancing memory performance during simulation steps. Additionally, reducing the spectrum in both dimensions, rather than only along the X-axis, could be considered. However, this would likely increase computation error due to the additional removal of coefficients, and the performance gains may not justify losing the current advantage of maintaining a single continuous area of coefficients, which simplifies subsequent operations.

Despite FFT-pruning being well-established and the impact of spatial resolution on image quality being well understood, integrating these techniques for

spectral methods in wave propagation simulations is novel. This approach bridges computational efficiency with high-resolution accuracy in complex wave models, opening new research opportunities in ultrasound simulation and acoustics.

8 Conclusion

This paper presented an approach to accelerate k-Wave's wave propagation simulation by replacing the standard FFT algorithm with a pruned FFT. The proposed method was demonstrated via a proof-of-concept implementation of the pruned FFT integrated into the k-Wave simulation framework. Experiments using simulation data from human skull and liver models showed significant improvements in computational time, particularly for high-resolution domains. Although transitions between media with significant differences in properties such as sound speed and density may lead to high computation errors at their boundaries, the overall impact on focal point accuracy was minimal, resulting in negligible focal shifts.

This approach shows promising potential for focused ultrasound procedures, especially in scenarios where multiple simulations are required to determine the optimal transducer position for targeting specific tissue areas. The improved computational efficiency makes it particularly suitable for treatment planning, where rapid iteration through various transducer configurations is necessary while maintaining acceptable accuracy. Notably, despite the challenges associated with ultrasound penetration through bone, our method achieved up to a 1.8x speedup in large simulation domains with a 1% bisection threshold, demonstrating its robustness in demanding scenarios.

Future research will focus on further improving the pruned FFT implementation and its integration into the k-Wave toolbox, including enhancements to the preprocessing phase, comprehensive evaluation and validation of the proposed approach, and potential extension to three-dimensional simulations. This work contributes to the ongoing effort to enhance the efficiency of spectral methods in wave propagation simulations, particularly for medical ultrasound applications. The promising results pave the way for more efficient high-resolution simulations, potentially enabling faster and more accurate treatment planning in clinical settings.

Acknowledgments. This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254). This project has received funding from the European Unions Horizon Europe research and innovation programme under grant agreement No 101071008. This work was supported by Brno University of Technology under project number FIT-S-23-8141.

References

1. Aubry, J.F., Bates, O., Boehm, E.A.: Benchmark problems for transcranial ultrasound simulation: intercomparison of compressional wave models. *J. Acoust. Soc. Am.* **152**(2), 1003–1019 (2022). <https://doi.org/10.1121/10.0013426>

2. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex fourier series. *Math. Comput.* **19**(90), 297–301 (1965). <http://www.jstor.org/stable/2003354>
3. Engholm, M., Stepinski, T.: Designing and evaluating transducers for narrow-band ultrasonic spectroscopy. *NDT & E Int.* **40**(1), 49–56 (2007). <https://doi.org/10.1016/j.ndteint.2006.07.006> <https://www.sciencedirect.com/science/article/pii/S0963869506000673>
4. Gircys, M., Ross, B.J.: Image evolution using 2D power spectra. *Complexity* **2019**(1), 7293193 (2019). <https://doi.org/10.1155/2019/7293193>
5. Gottlieb, S., Gottlieb, D.: Spectral methods. *Scholarpedia* **4**(9), 7504 (2009). <https://doi.org/10.4249/scholarpedia.7504>
6. Jaros, J., Treeby, B., Rendell, A.: Use of multiple GPUs on shared memory multi-processors for ultrasound propagation simulations, vol. 127, pp. 43–52 (2012)
7. Markel, J.: FFT pruning. *IEEE Trans. Audio Electroacoust.* **19**(4), 305–311 (1971). <https://doi.org/10.1109/TAU.1971.1162205>
8. Martin, E., Jaros, J., Treeby, B.E.: Experimental validation of k-wave: nonlinear wave propagation in layered, absorbing fluid media. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **67**(1), 81–91 (2020). <https://doi.org/10.1109/TUFFC.2019.2941795>
9. NVIDIA Corporation: cuFFT library. <https://developer.nvidia.com/cufft>. Accessed 25 Sept 2024
10. Olsak, O., Jaros, J.: Techniques for efficient fourier transform computation in ultrasound simulations. In: *Proceedings of the 33rd International Symposium on High-Performance Parallel and Distributed Computing. HPDC '24*, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3625549.3658825>
11. Olsak, O., Jaros, J.: Accelerating ultrasound wave propagation simulations using pruned FFT. In: *2024 IEEE International Conference on High Performance Computing and Communications (HPCC)*, pp. 168–173. <https://doi.org/10.1109/HPCC64274.2024.00032>
12. Robertson, J.L.B., Cox, B.T., Jaros, J., Treeby, B.E.: Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *J. Acoust. Soc. Am.* **141**(3), 1726–1738 (2017). <https://doi.org/10.1121/1.4976339>
13. Schumacher, J., Püschel, M.: High-performance sparse fast fourier transforms. In: *2014 IEEE Workshop on Signal Processing Systems (SiPS)*, pp. 1–6 (2014). <https://doi.org/10.1109/SiPS.2014.6986055>
14. Sorensen, H., Burrus, C.: Efficient computation of the DFT with only a subset of input or output points. *IEEE Trans. Sig. Process.* **41**(3), 1184–1200 (1993). <https://doi.org/10.1109/78.205723>
15. Tolmachev, D.: VkFFT—a performant, cross-platform and open-source GPU FFT library. *IEEE Access* **11**, 12039–12058 (2023). <https://doi.org/10.1109/ACCESS.2023.3242240>
16. Treeby, B., Cox, B., Jaros, J.: k-wave a matlab toolbox for the time domain simulation of acoustic wave fields user manual (2016). http://www.k-wave.org/manual/k-wave_user_manual_1.1.pdf
17. Treeby, B.E., Budisky, J., Wise, E.S., Jaros, J., Cox, B.T.: Rapid calculation of acoustic fields from arbitrary continuous-wave sources. *J. Acoust. Soc. Am.* **143**(1), 529–537 (2018). <https://doi.org/10.1121/1.5021245>

18. Treeby, B.E., Jaros, J., Rendell, A.P., Cox, B.T.: Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *J. Acoust. Soc. Am.* **131**(6), 4324–4336 (2012). <https://doi.org/10.1121/1.4712021>
19. Treeby, E.B., Jaros, J., Rendell, P.A., Cox, T.B.: Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *J. Acoust. Soc. Am.* **131**(6), 4324–4336 (2012). <https://doi.org/10.1121/1.4712021>
20. Ye, P.P., Brown, J.R., Pauly, K.B.: Frequency dependence of ultrasound neurostimulation in the mouse brain. *Ultrasound Med. Biol.* **42**(7), 1512–1530 (2016). <https://doi.org/10.1016/j.ultrasmedbio.2016.02.012>
21. Yoon, K., Lee, W., Croce, P., Cammalleri, A., Yoo, S.S.: Multi-resolution simulation of focused ultrasound propagation through ovine skull from a single-element transducer. *Phys. Med. Biol.* **63**(10), 105001 (2018). <https://doi.org/10.1088/1361-6560/aabe37>