Techniques for Efficient Fourier Transform Computation in Ultrasound Simulations

Ondrej Olsak
iolsak@fit.vutbr.cz
Faculty of Information Technology
Brno University of Technology
Brno, Czech Republic

Jiri Jaros
jarosjir@fit.vutbr.cz
Faculty of Information Technology
Brno University of Technology
Brno, Czech Republic

ABSTRACT

Noninvasive ultrasound surgeries represent a rapidly growing field in medical applications. Preoperative planning often relies on computationally expensive ultrasound simulations. This paper explores methods to accelerate these simulations by reducing the computation time of the Fourier transform, which is an integral part of the simulation in the k-Wave toolbox. Two experiments and their results will be presented. The first investigates substituting the standard Fast Fourier Transform (FFT) with a Sparse Fourier Transform (SFT). The second approach utilises filtering of the frequency spectrum, inspired by image compression algorithms. The aim of both experiments is to find a suitable method for accelerating the Fourier transform while utilising the sparsity of the spectrum in acoustic pressure. Our findings show that filtering offers significantly better results in terms of computation error, leading to a substantial reduction in overall simulation runtime.

CCS CONCEPTS

• **Computing methodologies → Massively parallel and high-performance simulations**.

KEYWORDS

Ultrasound wave propagation, k-Wave, Sparse Fourier Transform

ACM Reference Format:
Ondrej Olsak and Jiri Jaros. 2024. Techniques for Efficient Fourier Transform Computation in Ultrasound Simulations. In *The 33rd International Symposium on High-Performance Parallel and Distributed Computing (HPDC '24), June 3–7, 2024, Pisa, Italy.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3625549.3658825

1 INTRODUCTION

The Fourier transform is a fundamental mathematical tool used to convert signals from the time domain to the frequency domain. This transformation is crucial for applications like image compression, ultrasound imaging, or numerical calculation of derivation.

Widely used algorithm for computation of the Fourier transform is Fast Fourier transform algorithm (FFT) [1]. This algorithm reduces time complexity of the Fourier transform from $O(N^2)$ to $O(N \log N)$ where N is length of input signal, which is significant improvement compared to the original algorithm. While the Fast Fourier Transform revolutionised the speed of calculating Fourier transforms, it is not always necessary to compute all N coefficients. In some problems the spectrum contains only a few nonzero coefficients or significant coefficients (that we are interested in) and noise. An example of such a problem is GPS synchronisation [3].

This property of the signal is used by algorithm called the Sparse Fourier transform (SFT). This algorithm expect the signal to contain at most k significant coefficients, where $k \ll N$ for the signal of length N. This leads to lower time complexities compared to the FFT that may be beneficial for some applications working with signals with sparse spectral domain and large datasets or in real-time environment. There are multiple implementations of SFT algorithms that use various techniques to estimate Fourier coefficients from sparse signals. Some of these implementations are universal, aiming to identify the k most significant coefficients across general signals, while others exploit specific signal characteristics or operate within defined domains of knowledge.

This paper explores possibility to accelerate ultrasound wave propagation simulations in k-Wave toolbox [6] by involving some sort of the Sparse Fourier transform algorithm that utilises information about the simulation, such as transmitter frequency, media density, sound speed, etc. In these simulations, calculating the Fourier transform typically consumes about 60% of the time in each simulation step [4]. We specifically focus on applying the SFT to the spatial pressure distribution. During the preplanning of ultrasound surgeries, multiple simulations are executed to find the required position of the transmitter (focus). By involving the SFT algorithm, the overall time required to locate the position of the transmitter, and thus the cost paid for computation resources, should decrease. The following sections present two experiments and their results. The first investigates the usage of the classic SFT approach by imitating its functionality, computing the simulation step only with the k most significant coefficients. The second approach takes inspiration from image compression techniques, utilising filtering to reduce the number of computed Fourier coefficients. The results from these experiments will be used in ongoing research focused on this topic.

2 PHYSICAL PROBLEM DESCRIPTION

To compute the ultrasound wave propagation simulation, the k-Wave toolbox uses the pseudo-spectral method with the Fourier basis function. The idea of this method is to transform the solution of the differential equation into a sum of certain basis functions. In the spectral methods, the solution depends on the entire domain compared to the finite-difference time domain methods where gradient is computed base on the function value at the neighbour

HPDC '24, June 3–7, 2024, Pisa, Italy

ACM ISBN 979-8-4007-0413-0/24/06
https://doi.org/10.1145/3625549.3658825

points. This makes the spectral method more accurate than local methods [2].

The k-Wave toolbox is employed to run simulations based on the following governing equations [6].

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

Equation (1) can be written in a discrete form using the k-space pseudo-spectral method [5]. This equation is part of the spatial gradient calculations based on the Fourier collocation spectral method.

$$
\frac{\partial}{\partial \xi} p^n = \mathcal{F}^{-1}\{ik_\xi \kappa e^{ik_\xi} \mathcal{F}\{p^n\}\} \tag{2}
$$

In Equation (2) for the Cartesian direction $\xi = x$ in R^1, $\xi = x, y$ in R^2, $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the forward and inverse spatial Fourier transform , i is the imaginary unit, k_ξ represents the wave numbers in the ξ direction, and κ is the k-space operator defined as $\kappa = sinc(c_{ref} k \Delta t/2)$, where c_{ref} is a scalar reference sound speed.

In the actual implementation the Fast Fourier transform algorithm is used to transform signal from spatial to spectral domain. Each simulation step of the ultrasound propagation simulation consists of 14 FFTs, that take significant part of the simulation [4].

3 EXPERIMENTS

All simulations presented in the following sections are computed using the k-Wave toolbox. To evaluate the quality of experimental simulations, a reference ultrasound wave propagation simulation was created. This simulation will be used to compute the error in acoustic pressure against the modified version of the simulation. The simulation itself has a size of 1024^2 grid points and consists of a parabolic ultrasound transmitter and a medium with four layers representing water, skin, skull, and brain. The reference acoustic pressure at the end of the simulation, along with its spectrum, can be seen in Figure 1. For a correct interpretation of the results, especially the images in the frequency domain, it is important to mention that all images of the acoustic pressure contain the state after the last simulation step has finished.

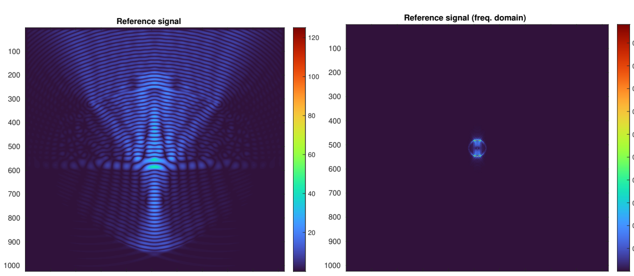

(a) Time domain **(b) Frequency domain**

Figure 1: The last step of the reference ultrasound wave propagation in heterogeneous media.

3.1 Computing with k most significant Fourier coefficients

The goal of this experiment is to simulate the behaviour of the Sparse Fourier transform algorithm, which computes only a given number of the most significant coefficients. This was achieved by using the actual implementation of the k-Wave, where immediately after the computation of the FFT over the acoustic pressure, a certain percentage of coefficients were zeroed out. The non-zero coefficients left in the spectrum represent the significant coefficients that would be found by the SFT algorithm.

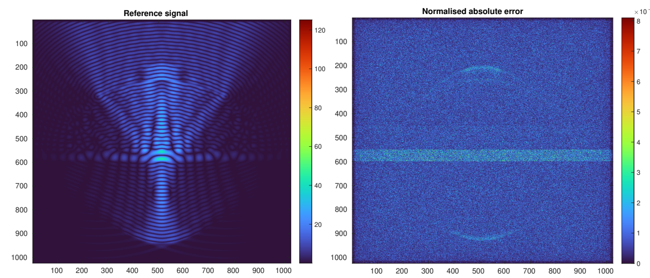

(a) Time domain **(b) Normalised absolute error**

Figure 2: The last simulation step of the ultrasound wave propagation in heterogeneous media effected by filtering out 30% of the lowest-amplitude coefficients.

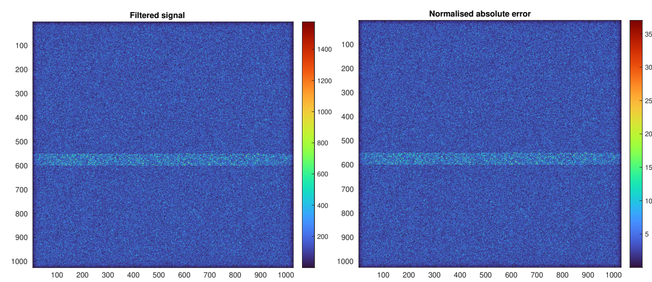

(a) Time domain **(b) Normalised absolute error**

Figure 3: The last simulation step of the ultrasound wave propagation in heterogeneous media effected by filtering out 70% of the lowest-amplitude coefficients.

The results of this experiment have shown that by filtering out 30% of lowest coefficients (see Figure 2), the normalised absolute error against the reference solution is on the order of 10^{-3}, thus the maximum error in focus is around 0.8%. On the other hand, in the experiment where 70% of all coefficients were zeroed out (see Figure 3), the final acoustic pressure distribution consisted only of noise.

In the case of using the SFT in ultrasound wave propagation, the number of significant coefficients that need to be found is relatively high compared to the size of the domain. This could lead to longer execution times due to the overhead caused by the nature of the SFT algorithm, rather than using the actual solution in the form of FFT. In this case, by reducing the number of significant coefficients in the spectral domain that need to be estimated by SFT, the computation error of the simulation increases rapidly.

3.2 Efficient Fourier Transform Computation via Filtering

This experiment was inspired by the process usually used in image compression. When we look at the image in Figure 1b, it can be observed that most of the coefficients are located at the centre of the spectral domain, indicating that they represent low frequencies. The idea is to use a mask, as shown in Figure 4, where the yellow circle represents coefficients that will be computed (1) and the blue coefficients that will be zeroed out (0).

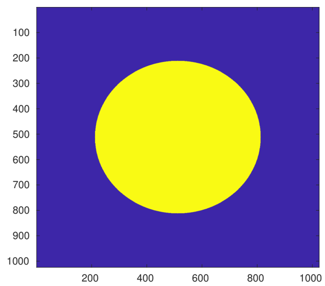

Figure 4: Filtration mask

The original simulation was modified so that after each computation of the full FFT over the acoustic pressure, the spectral domain is multiplied by the mentioned mask, and the result is used within the actual simulation step. This process is repeated for each simulation step.

The result of the simulation with this modification is shown in Figure 5. It can be observed that most of the computation error is located at the position of the transmitter and at the boundaries of the media. This is caused by removing high frequencies that represent sharp edges. Since the additive transmitter is used, even though we filter out high frequencies, the transmitter persists in the domain.

The number of coefficients that need to be computed is equal to the number of yellow grid points in the filtration mask. The radius of the circle is 256 grid points, which means the circle contains 205 861 grid points. The number of grid points in the domain is 1024^2. This means that approximately 80% of the coefficients can be zeroed out, and approximately 20% of the coefficients are computed. This gives us significantly better results than using the classic SFT approach described in the first experiment, as the maximum error in focus in the spatial domain is around 8.9%. Another advantage is removing overhead of the SFT algorithm by specifying coefficients (using the mask) that will be computed directly based on the transmitter and domain properties.

4 CONCLUSION

The experiments have shown that by computing only coefficients within a given mask that covers specific areas in the spectral domain, the accuracy of this simulation is significantly better than by computing the same number of coefficients selected based on the amplitude value. This may lead to a reduction in the time complexity of the FFT used in simulations. The following research will focus on solving challenges that come with this approach, such as transmitter error reduction and the selection of the spectral coefficients that will be computed and used during simulation.

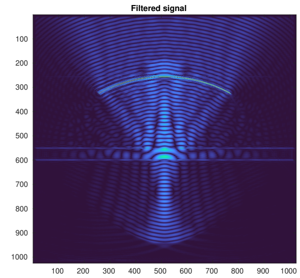

(a) Time domain

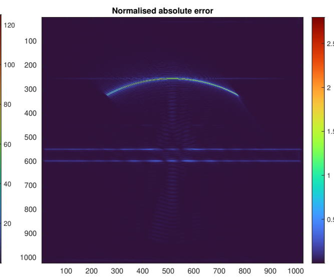

(b) Normalised absolute error

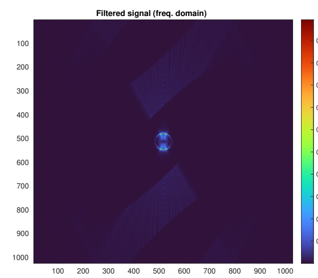

(c) Frequency domain

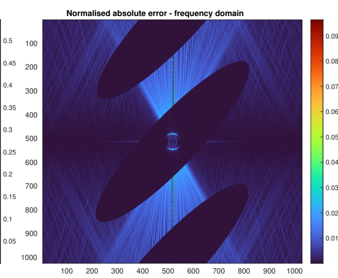

(d) Normalised absolute error in frequency domain

Figure 5: The last simulation step of the ultrasound wave propagation in heterogeneous media effected by filtering out coefficient using mask.

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] James W. Cooley and John W. Tukey. 1965. An Algorithm for the Machine Calculation of Complex Fourier Series. *Math. Comp.* 19, 90 (1965), 297–301. http://www.jstor.org/stable/2003354

[2] S. Gottlieb and D. Gottlieb. 2009. Spectral methods. *Scholarpedia* 4, 9 (2009), 7504. https://doi.org/10.4249/scholarpedia.7504 revision #91796.

[3] Haitham Hassanieh, Fadel Adib, Dina Katabi, and Piotr Indyk. 2012. Faster GPS via the Sparse Fourier Transform. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking* (Istanbul, Turkey) *(Mobicom '12)*. Association for Computing Machinery, New York, NY, USA, 353–364. https://doi.org/10.1145/2348543.2348587

[4] Jiri Jaros, Bradley Treeby, and Alistair Rendell. 2012. Use of multiple GPUs on shared memory multiprocessors for ultrasound propagation simulations. *Conferences in Research and Practice in Information Technology Series* 127, 43–52.

[5] Bradley Treeby, Ben Cox, and Jiri Jaros. 2016. 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

[6] E. Bradley Treeby, Jiri Jaros, P. Alistair Rendell, and T. Ben Cox. 2012. Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *Journal of the Acoustical Society of America* 131, 6 (2012), 4324–4336. https://doi.org/10.1121/1.4712021

Received 21 March 2024; revised 17 April 2024; accepted 12 April 2024