


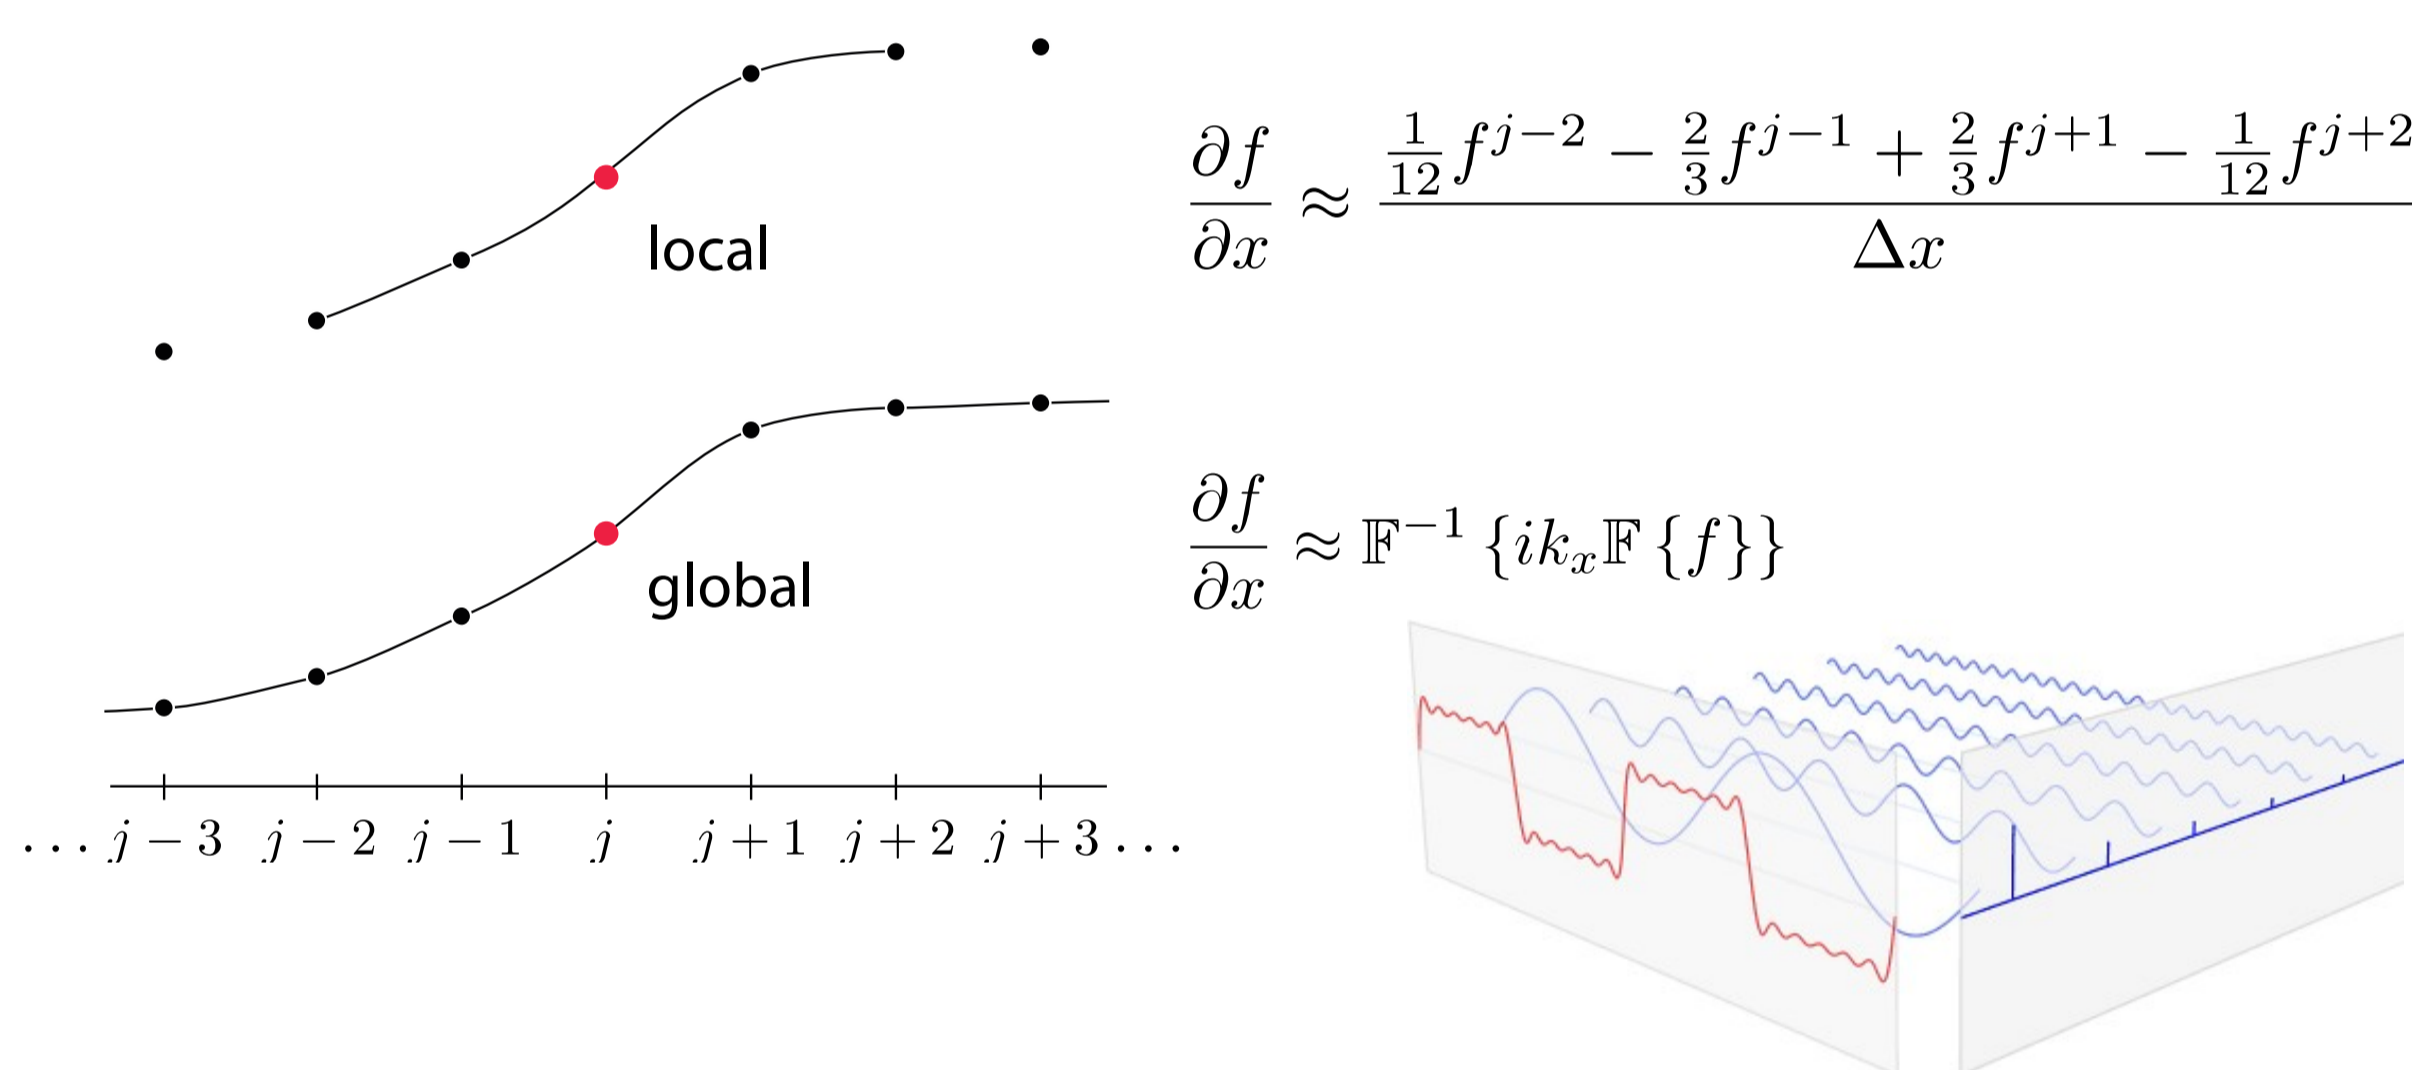
## Overview

The simulation of elastic wave propagation has many applications in ultrasonics, including the classification of bone diseases and non-destructive testing. In biomedical ultrasound in particular, elastic wave models have been used to investigate the propagation of ultrasound in the skull and brain, and to optimize the delivery of therapeutic ultrasound through the thoracic cage. Currently, there is an accurate elastic wave model written in MATLAB as part of the open-source k-Wave toolbox. However, the computational requirements of the model did not allow it to be deployed for realistic simulation cases. Therefore, we decided to create a native implementation in CUDA to accelerate the simulation.



## Proposed Method

The implementation of the native CUDA application follows the numerical model based on the explicit solution of coupled PDEs using the Fourier pseudospectral method. This uses the Fourier collocation spectral method to compute spatial derivatives, and a leapfrog finite-difference scheme to integrate forwards in time.

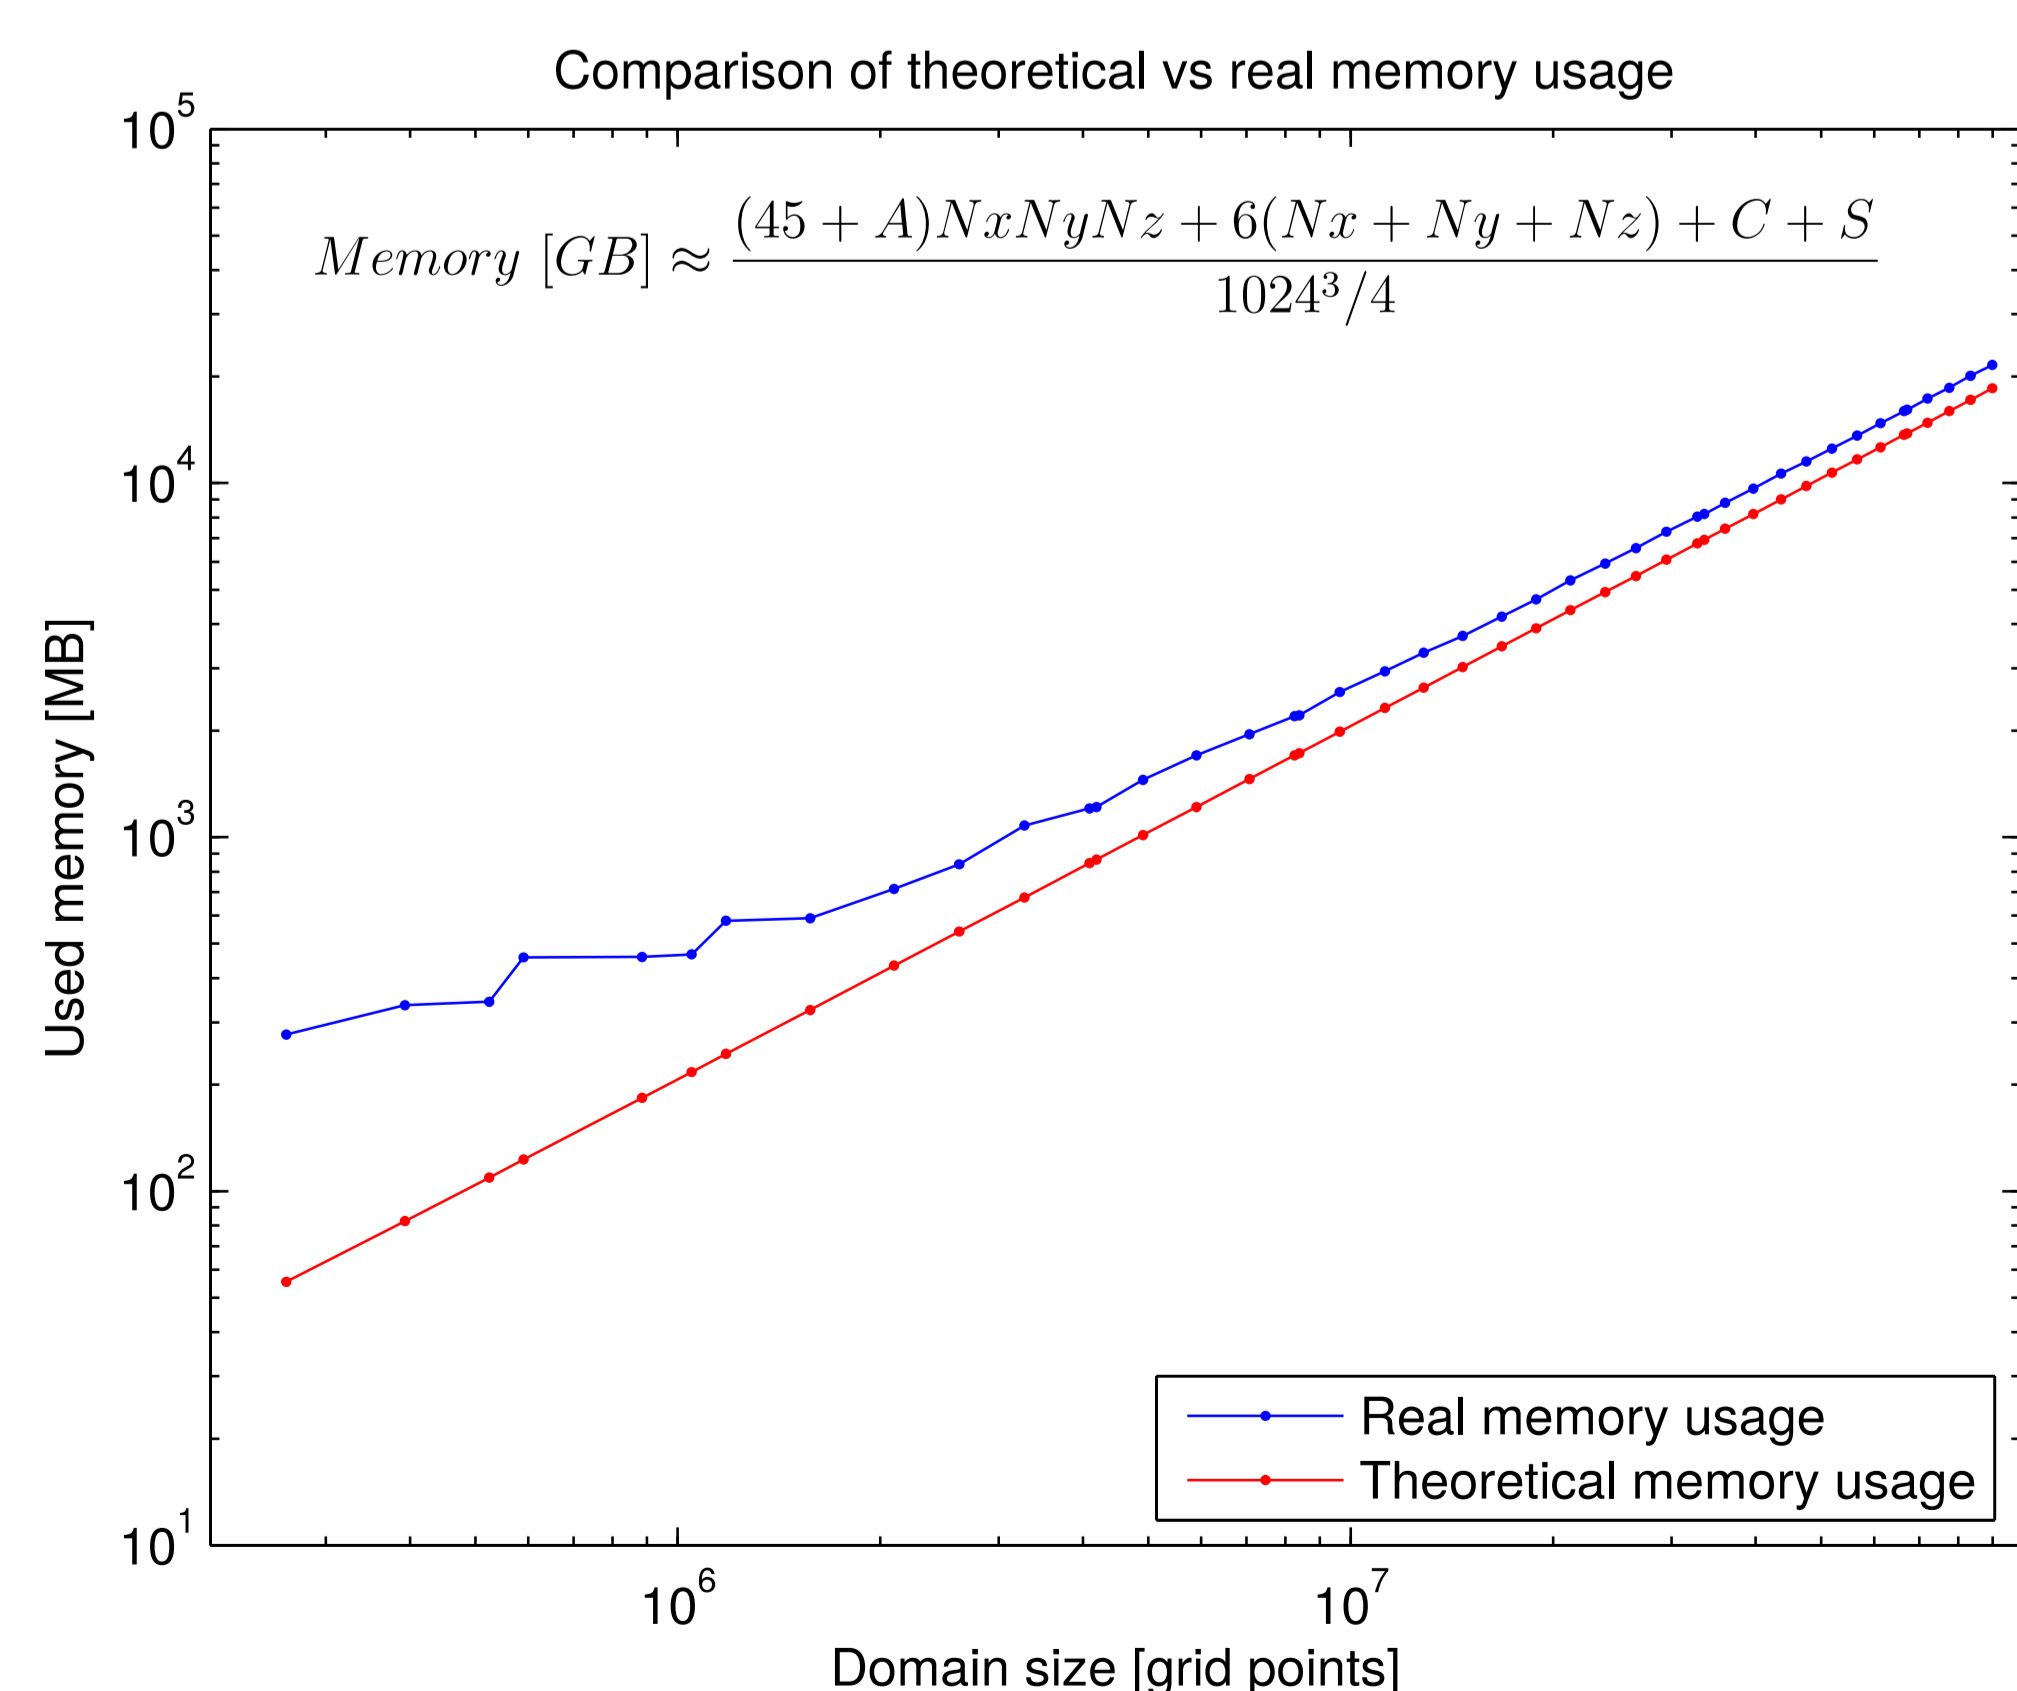


An example of such a calculation to compute the derivative of the 3D stress field in x dimension

$$\partial_x \sigma_{xx} = \mathcal{F}_x^{-1} \left\{ ik_x e^{+ik_x \Delta x/2} \mathcal{F}_x \{ \sigma_{xx} \} \right\}$$

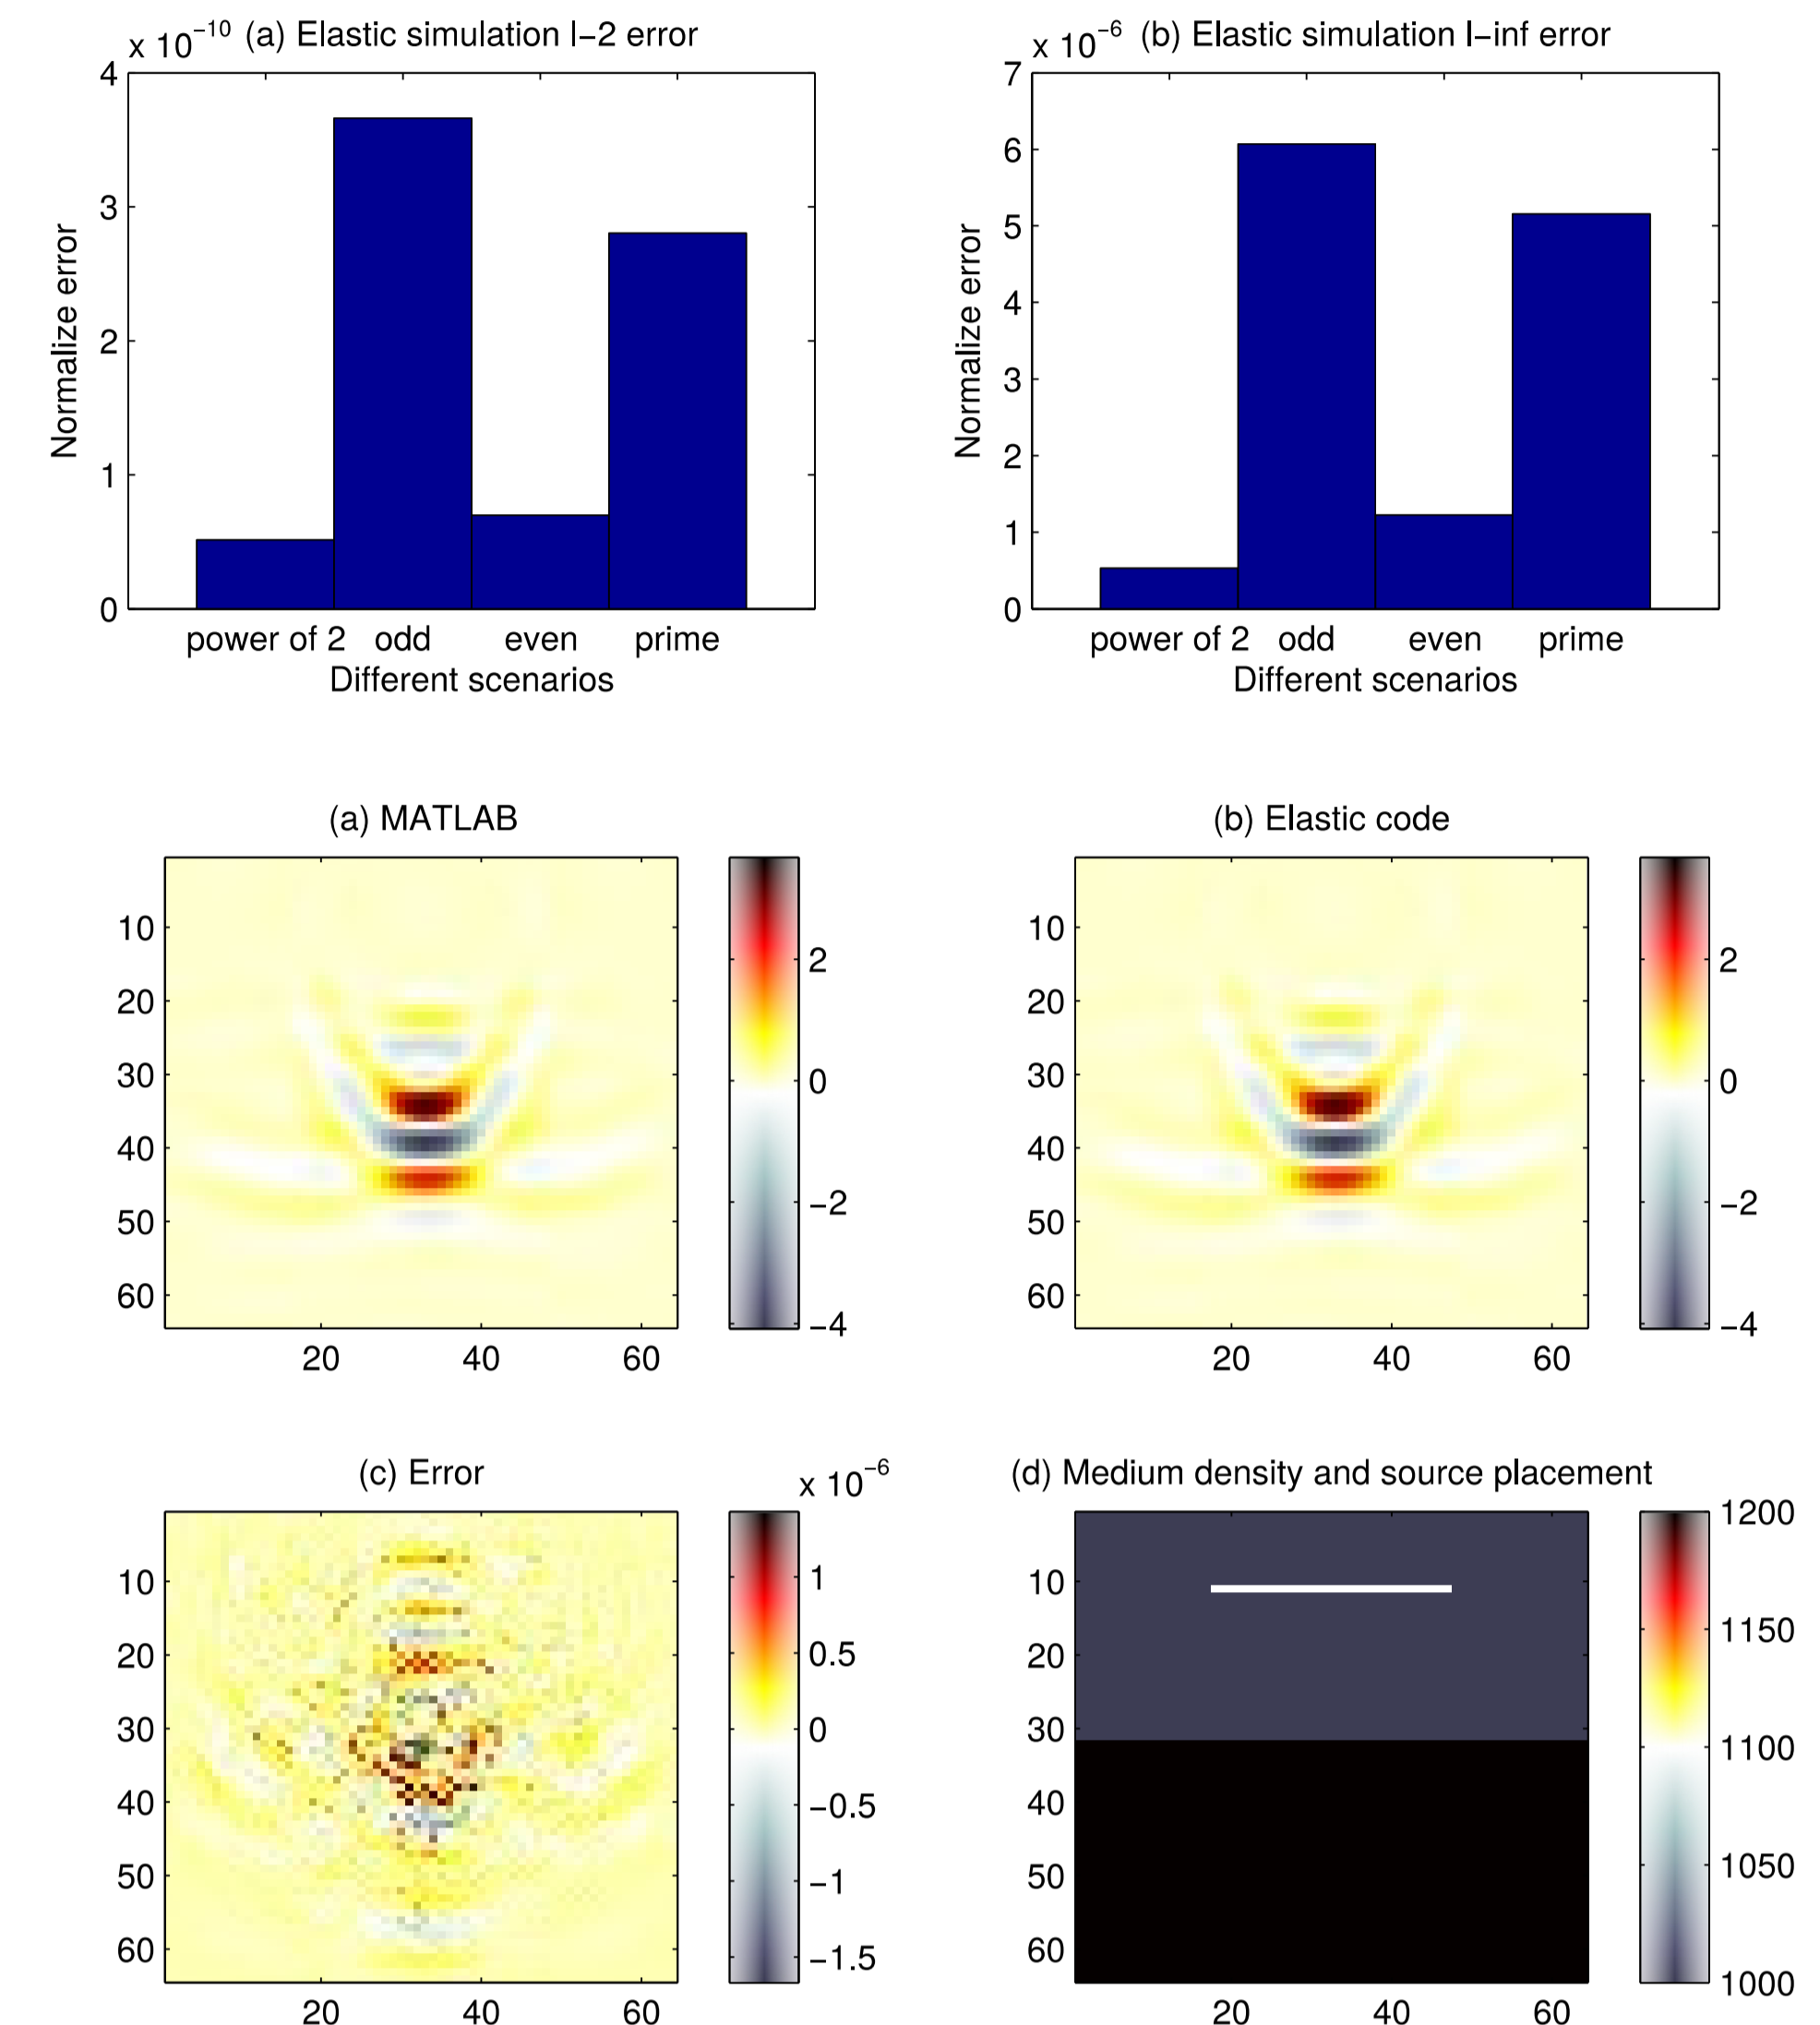
## Spatial Requirements of CUDA Implementation

The main goal of the application is to make simulation of elastic wave propagation deployable on realistic cases. These cases require vast amount of memory. However, modern GPU's have only limited memory capacity at their disposal. Therefore, a memory consumption mode was created to determine what is, and what is not within the capabilities of certain GPU.



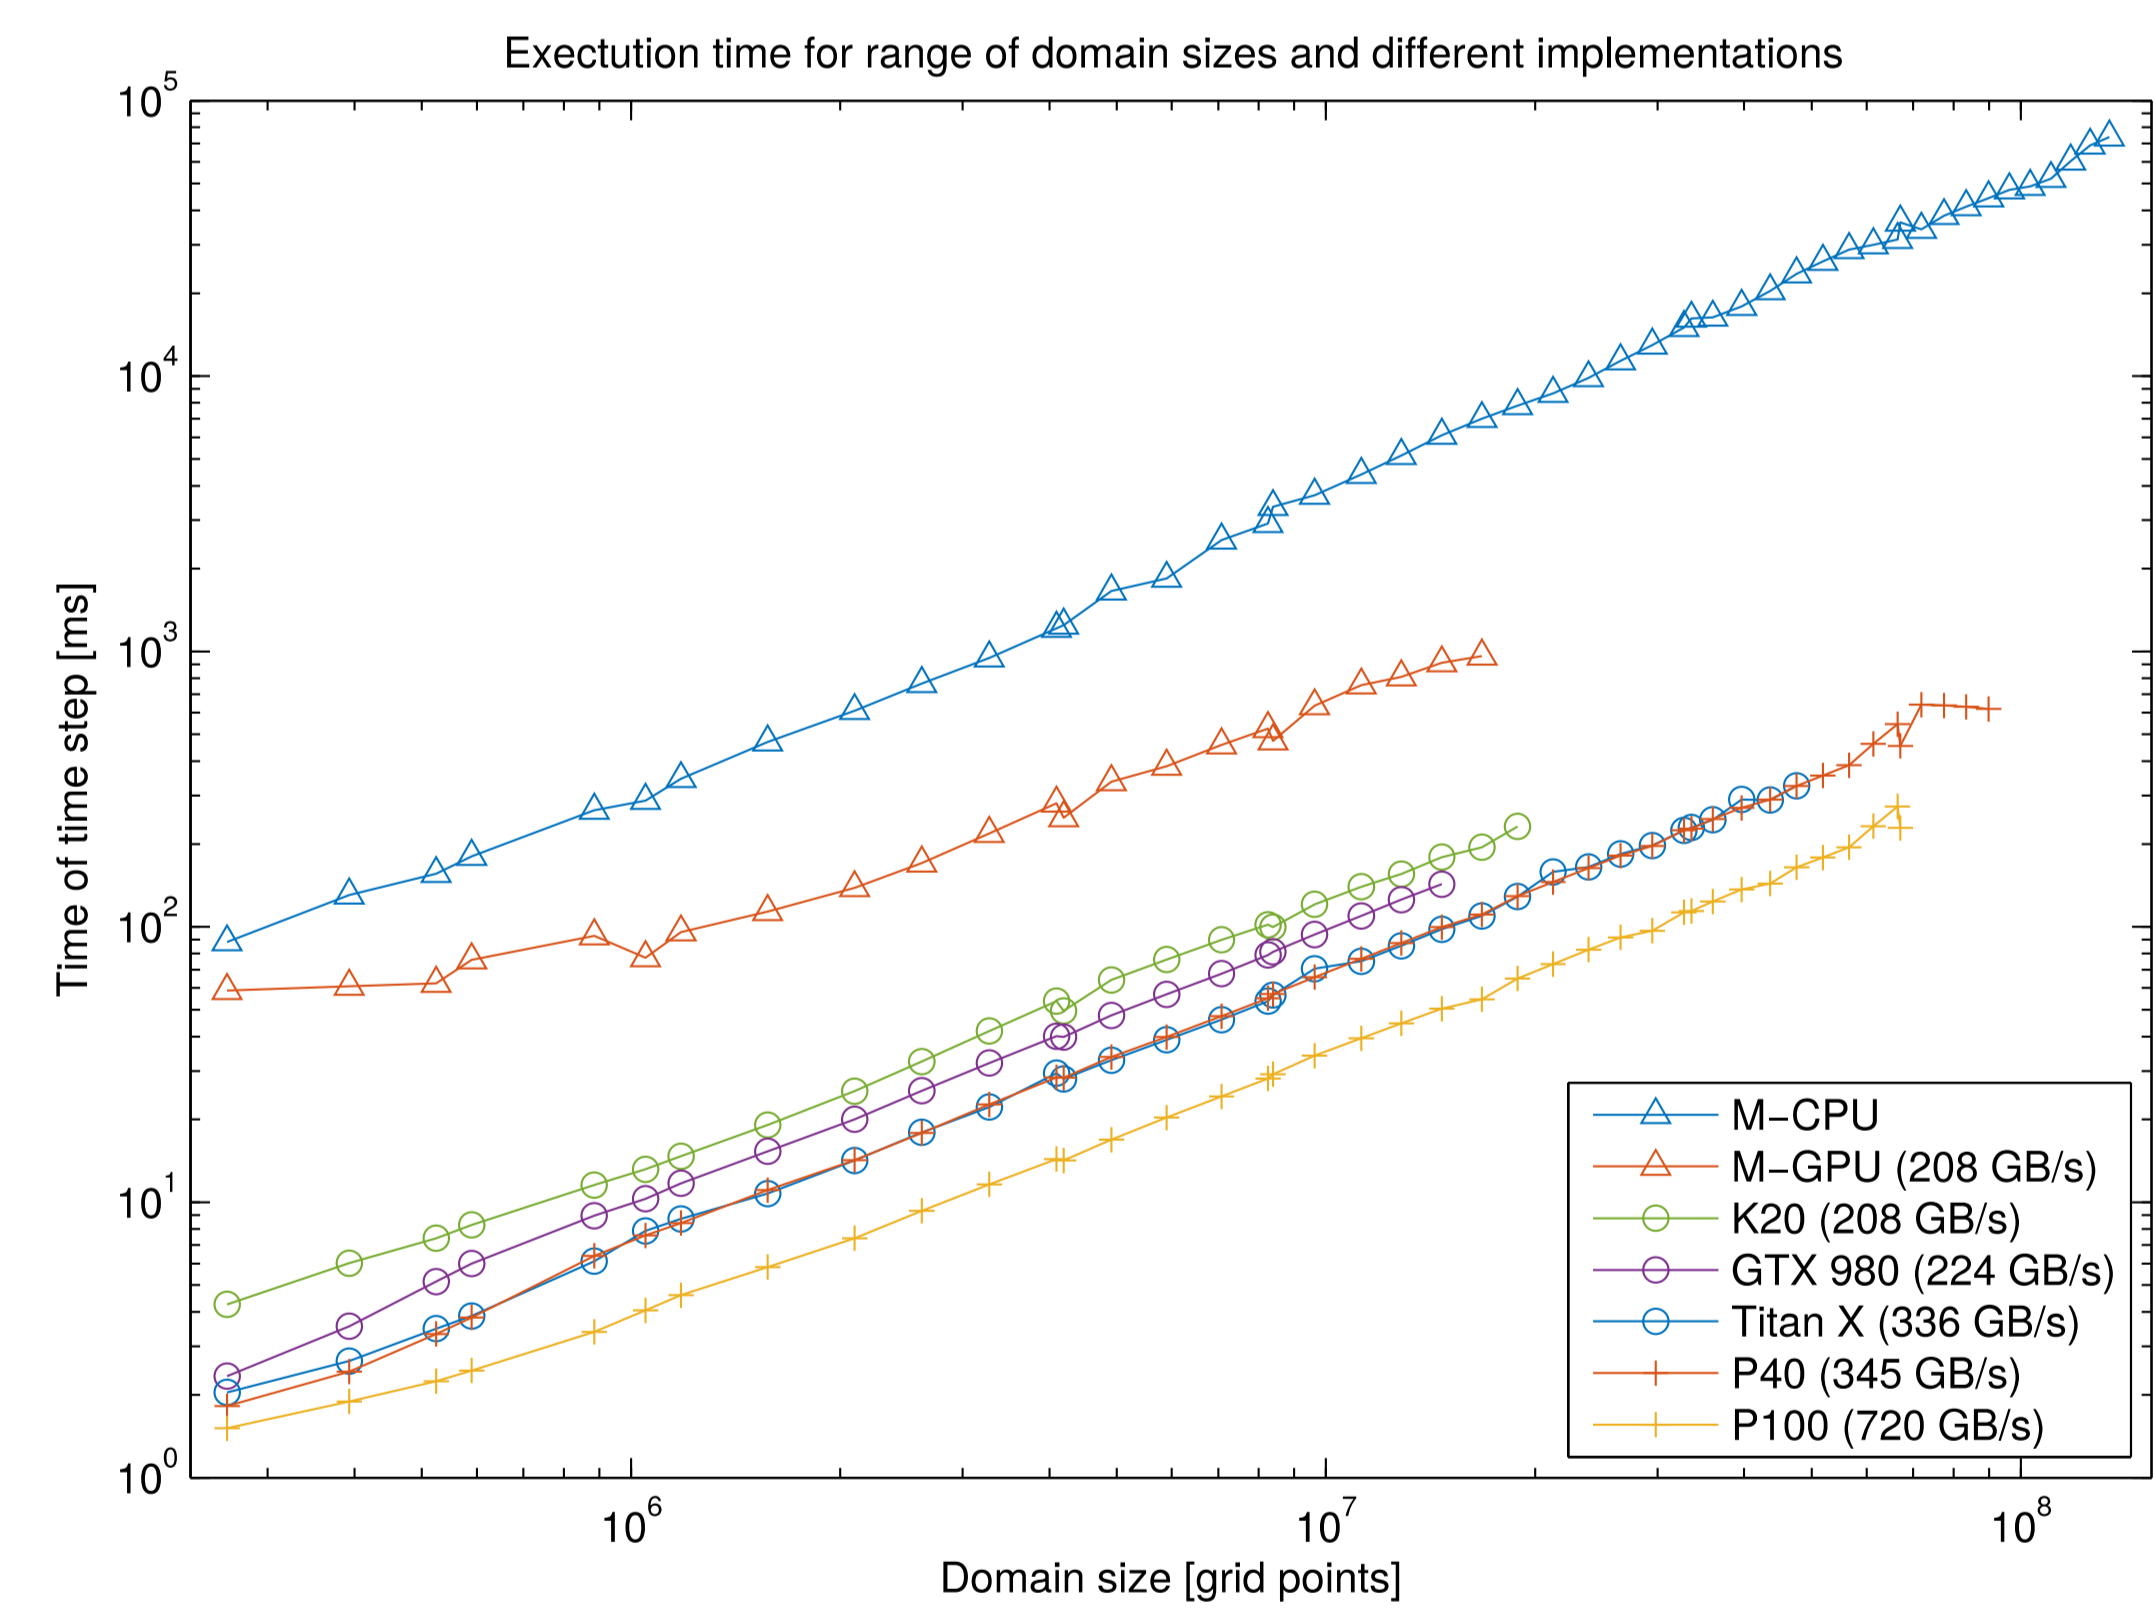
## Accuracy of Implementation

The accuracy of the implementation is determined by the accuracy of Fast Fourier Transform (FFT). The accuracy of FFT is strongly dependent on the number of elements of domain.



## Performance Investigation

The simulation time was investigated on a few different clusters in Europe with different hardware configurations. Application was benchmarked on graphic cards ranging from desktop models to GPU's specifically designed for HPC.



## Conclusions

By implementing simulation on graphic card, we were able to reduce iteration time by factor 158.5 at the best. It is worth noting that our implementation preforms about 5.8 times better than native MATLAB GPU implementation on the same hardware. The application is intended to be coupled with existing implementation of fluid model in the future, to be able to solve complex real-life scenarios. This process is not trivial because it introduces new challenges just as load balancing and many more.