

# Performance, Portability and Productivity for Room Acoustics Codes

Larisa Stoltzfus



THE UNIVERSITY of EDINBURGH

EPSRC Centre for Doctoral Training in Pervasive Parallelism

## Motivation

### Problems:

- Parallel programming is becoming more complex and less portable
- Re-writing and re-tuning code is a tedious task
- Computational scientists should not have to be experts in parallelisation methods
- Porting to new architectures often means research groups must maintain multiple code bases

### Potential Solution:

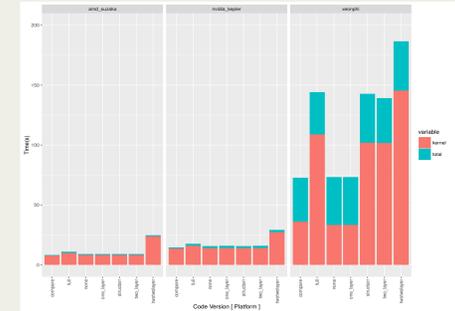
High level parallel abstraction layers could decouple parallelisation from simulation codes for more portable software without a compromise in performance

### Context:

Acoustics models written by NESS are ported separately to run on CPUs in MATLAB and on NVIDIA GPUs in C/CUDA

*Is there a productive, high-level way to write these codes that will maintain performance across multiple architectures?*

## Abstraction -vs- Performance



Graph of variation in performance of different data abstractions across multiple platforms

Removing memory layout management from programmer responsibility could control performance repercussions for data abstractions

## Approach

### Low-level:

Begin with comparison of different implementations of the simplest room acoustics benchmark run over multiple platforms, looking at performance, portability and programmability

Lay down groundwork for future comparisons

Using an abstracted low-level version, test out and compare potential optimisations

Ascertain biggest issues and best optimisations for simple benchmarks, then investigate more advanced room codes

### High-level:

Investigate feasibility of higher level frameworks - are there solutions available that could be tailored to stencil codes like NESS?

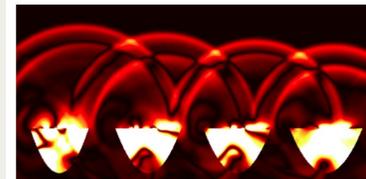
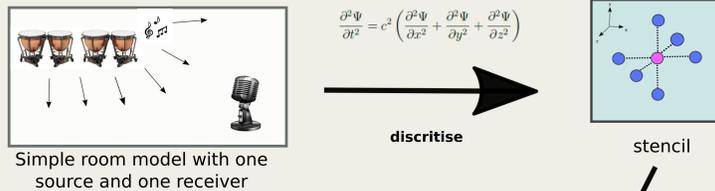
## Case Study: NESS

The NESS (*Next gEneration Sound Synthesis*) project aims to develop simulations for modelling musical instruments in different acoustic settings without the use of samples

The room acoustics models are discretised using a finite difference method across the time domain to simulate the behaviour of sound waves

This type of simulation uses neighbouring points to make updates (an algorithm known as a **stencil**) and is quite common in HPC to calculate new values across a grid over a number of timesteps

The 3D wave equation (below) is discretised down to stencil form and the shape of the room is simulated with a grid



Snapshot of four timpani drums as modelled in a 3D space<sup>1</sup>

## Advanced Rooms

Current work has only been done on the simplest room acoustics benchmarks

Real rooms have multiple sound sources, more complicated room structures, boundary conditions, viscosity, obstacles and multiple time steps

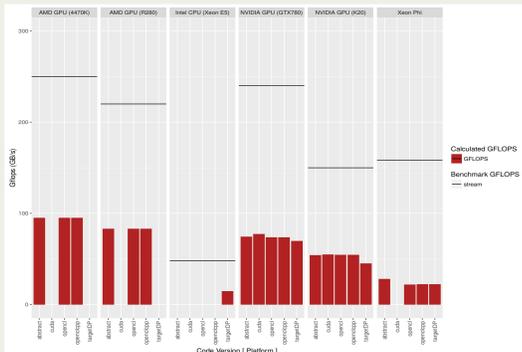
How do these additional features affect results?

Are these stencils just as or more difficult to abstract?



Real spaces are much more complex to model

## Low-level Results



Graph of different low-level implementations run across multiple platforms

- Room acoustics codes are memory bandwidth-bound
- Data throughput could be higher based on STREAM results

epcc

EPSRC

Engineering and Physical Sciences Research Council

[1] Craig J. Webb PhD Thesis: [http://www.ness-music.eu/wp-content/uploads/2014/07/CJWebb\\_thesis-1.pdf](http://www.ness-music.eu/wp-content/uploads/2014/07/CJWebb_thesis-1.pdf)

## Future Work

Along with room acoustics, many other processes can be modelled through stencils:

How well can these be abstracted using current higher level frameworks (ie. Kokkos, SYCL, parallel skeletons)? What are their limitations? Could these frameworks be augmented to better accommodate physical simulations?

How do room acoustics codes compare with other stencil applications - do they face similar bottlenecks?

What is the simplest way of abstracting stencils that would allow them to be run in a portable, performant and productive way?