

MPARD: A High-Frequency Wave-Based Acoustic Solver for Very Large Compute Clusters

Nicolas Morales^{a,*}, Vivek Chavda^a, Ravish Mehra^a, Dinesh Manocha^a

^a*UNC Chapel Hill*
<http://gamma.cs.unc.edu/MPARD>

Abstract

We present a parallel time-domain wave solver designed for large and high frequency acoustic domains. Our approach is based on a novel scalable method for dividing acoustic field computations specifically for large-scale distributed memory clusters using parallel Adaptive Rectangular Decomposition (ARD).

In order to efficiently compute the acoustic field for large or high frequency domains, we need to take full advantage of the compute resources of large clusters. This is done with new algorithmic contributions, including a hypergraph partitioning scheme to reduce the communication cost between the cores on the cluster, a novel domain decomposition scheme that reduces the amount of numerical dispersion error introduced by the load balancing algorithm, and a revamped pipeline for parallel ARD computation that increases memory efficiency and reduces redundant computations.

Our resulting parallel algorithm makes it possible to compute the sound pressure field for high frequencies in large environments that are thousands of cubic meters in volume. We highlight the performance of our system on large clusters with 16000 cores on homogeneous indoor and outdoor benchmarks up to 10 kHz. To the best of our knowledge, this is the first time-domain parallel acoustic wave solver that can handle such large domains and frequencies.

Keywords: large-scale, wave-based methods, massively parallel

*Corresponding author
Email address: mmorales@cs.unc.edu (Nicolas Morales)

1. Introduction

Modeling and simulating acoustic wave propagation is one of the leading problems in scientific computing today [12]. Challenges in this area vary from real time constraints in video games and virtual reality systems to highly accurate offline techniques used in scientific computing and engineering. The simulation environment also varies; acoustic propagation problems can vary from small room acoustics problems to large, complicated outdoor scenes.

The sound we hear is the results of small changes in air pressure traveling as a wave. The propagation of these pressure waves is governed by the linear, second order partial differential acoustic wave equation:

$$\frac{\partial^2}{\partial t^2}p(\vec{x}, t) - c^2\nabla^2p(\vec{x}, t) = f(\vec{x}, t), \quad (1)$$

where \vec{x} is a 3D position, t is time, $p(\vec{x}, t)$ is the pressure at point \vec{x} and time t , f is a forcing term at point \vec{x} and time t , and c is the speed of sound. In this paper, we assume that the environment is homogeneous and the speed of sound is constant in the media.

There are two classes of methods known for solving the acoustic wave equation. The first class encompasses geometric approaches to solving the equation. These methods include such approaches as ray tracing techniques, image source methods, and beam tracing [2, 15, 16]. Geometric approaches are viable for real time applications, but cannot easily represent some wave phenomena such as diffraction or scattering.

The second class of methods, the wave-based techniques, directly solve the wave equation. These techniques provide accurate solutions to the acoustic wave equation, suitable for scientific and engineering applications, but at a large computational cost. These methods include finite difference time domain (FDTD) [6, 33], finite element method (FEM) [42], and the boundary element method (BEM) [11, 8]. Most current implementations of these methods are limited to small domains (less than 10 000 m³) or low frequencies (e.g., less than

2 kHz). There has also been some work in hybrid methods that combine wave-based and geometric techniques [46].

The main challenge is that the computational cost of computing the sound propagation in the environment scales with the 4th power of frequency and linearly with the volume of the scene, while memory use scales with the 3rd power of frequency and linearly with the scene volume. Since the human aural range scales from 20 Hz to 20 kHz, large scenes such as a cathedral (10 000 m³ to 15 000 m³) would require tens of Exaflops of computation and tens of terabytes of memory to compute using prior wave-based solvers. This makes high frequency acoustic wave simulation one of the more challenging problems in scientific computation [14].

Recently, there has been a lot of emphasis on reducing the computational cost of wave-based methods. These techniques, called *low dispersion techniques* use coarser computational meshes or decompositions in order to evaluate the wave equation. One example of a low dispersion algorithm is the Adaptive Rectangular Decomposition method (ARD)[29, 25], a solver for the 3D acoustic wave equation in homogeneous media. ARD is a domain decomposition method that subdivides the computational domain into rectangular regions. One of the main advantages of ARD is the greatly reduced computational and memory requirements over more traditional methods like FDTD [26].

Despite these computational advantages, the ARD method still requires a great deal of computing power and memory in order to evaluate the wave equation for the full range of human hearing on large architectural or outdoor scenes. In order to deal with these requirements, a parallel distributed version of ARD was developed [26]. However, this approach was limited to smaller scenes and compute clusters and still had high computational requirements. The approach did not take into account the cost of communication between cores and suffered from numerical instability.

Main results

We present MPARD, a Massively Parallel Adaptive Rectangular Decomposition method capable of computing 3D sound propagation through large architectural and outdoor scenes for pressure field computations at large frequencies (at least 10 kHz) and is designed to utilize tens of thousands of CPU cores.

Scaling to this number of CPU cores with such a large domain introduces many challenges, including communication cost, numerical stability, and reducing redundant computations. MPARD addresses all of these in a novel method that:

- Uses a hypergraph partitioning approach for load balance and communication reduction among multiple nodes
- Introduces a modified domain decomposition algorithm to improve the numerical accuracy of our wave-based simulator
- Uses a multi-stage program pipeline for scene preprocessing and runtime computation

MPARD is capable of computing the acoustic propagation in large indoor scenes ($20\,000\text{ m}^3$) up to at least 10 kHz at around 80 s per time step on 1024 cores. The algorithm has been tested to and scales up to at least 16000 cores, where each time step of a 5 kHz scene takes around 0.2 s. We have analyzed many aspects of our parallel algorithm including the scalability over tens of thousands of cores, the time spent in different computation stages, communication overhead, and numerical errors. To the best of our knowledge, this is one of the first wave-based solver that can handle such large domains and high frequencies.

2. Previous work

There has been considerable research in the area of wave-based acoustic solvers, including the field of parallel solvers, domain decomposition approaches,

and low-dispersion acoustic solvers. In this section, we provide a brief overview of some of these approaches.

2.1. Parallel wave-based solvers

Parallel wave-based solvers are used in a multitude of scientific domains, including the studying of seismic, electromagnetic, and acoustic waves. A large category of these solvers are parallel FDTD solvers either for large clusters [17, 43, 47, 48] or for GPUs [31, 36, 40, 41, 45, 35, 19]. A category of parallel methods are also based on finite-element schemes [13, 5].

Additionally, there are several parallel methods developed for specific applications of the wave equation. PetClaw [1] is a scalable distributed solver for time-dependent non-linear wave propagation. Other methods include distributed finite difference frequency-domain solvers for visco-acoustic wave propagation [27], discontinuous Galerkin solvers for heterogeneous electromagnetic and aeroacoustic wave propagation [4], scalable simulation of elastic wave propagation in heterogeneous media [3], etc.

2.2. Domain decomposition

MPARD, like ARD, is a domain decomposition method. Domain decomposition approaches subdivide the computational domain into smaller domains that can be solved locally. The solver uses these local solutions to compute a global solution of the scientific problem. Many of these approaches are designed for coarse grain parallelization where each subdomain or a set of subdomains is computed locally on a core or a node on a large cluster.

Domain decomposition approaches tend to fall into two categories: *overlapping subdomain* and *non-overlapping subdomain* methods [10].

Existing domain decomposition approaches in computational acoustics include parallel multigrid solvers for 3D underwater acoustics [32], pseudospectral acoustic solvers [19, 49], and the non-overlapping subdomain ARD approach [29]. Additional non-overlapping approaches include Trefftz methods that can be decomposed for parallelization [28] and line source decomposition methods for ESM (equivalent source method) techniques [18].

2.3. Low dispersion acoustic solvers

The goal of low dispersion methods is to reduce computation cost by using coarser computational meshes. A wide variety of approaches exist including an interpolated wideband scheme [22], extended pseudospectral methods for 3D domains [20], and a variety of low dispersion BEM approaches [34]. Waveguide approaches similarly reduce computation costs by reducing the computational grid size [44, 37, 38].

Our approach is primarily based on the ARD method [29] which partitions the computational domain into rectangular regions. It utilizes the property that the wave equation has a closed form solution in a homogeneous rectangular domain. Therefore, the only numerical error originates from the interfaces between these rectangular regions, allowing a much coarser grid size.

3. ARD Background

In this section, we provide a brief overview of the ARD method and a description of the parallel ARD pipeline. More details about these methods are available in [29, 25, 26].

Table 1 summarizes the problem size limitations of previous ARD algorithms, including parallel ARD. Previous algorithms have been limited to problem sizes under 3 kHz on large architectural scenes. The MPARD algorithm, on the other hand, is capable of operating on frequencies an order of magnitude greater than previous ARD methods. However, all algorithms are based on the same fundamental principles.

3.1. Adaptive Rectangular Decomposition

Adaptive Rectangular Decomposition (ARD) is a sequential wave-based time domain method for solving the 3D acoustic wave equation (Equation 1). ARD is limited to homogeneous environments, where the speed of sound c does not vary.

Solver	No. Cores	Maximum Frequency on Cathedral
ARD [29]	1	1000 Hz to 2000 Hz
GPUARD [25]	480 CUDA cores	1650 Hz
Parallel ARD [26]	1024	3000 Hz
MPARD	16384	10 000 Hz

Table 1: Comparison between different ARD techniques include MPARD, the method we introduce. We use the Cathedral scene as a benchmark, which is $19\,177\text{ m}^3$ in volume. We show the highest frequency used to obtain results using these techniques. In the case of GPUARD memory becomes a limiting factor, while in ARD computation time is. Parallel ARD has difficulty scaling past 1000 cores because of the limitations listed in Section 3.2.3. In this paper, we introduce the MPARD technique, which can scale to much higher frequencies (10kHz) and a much higher number of cores.

ARD is a domain decomposition approach. It takes advantage of the fact that the analytic solution of a 3D rectangular domain in a homogeneous media is known [23]. The ARD solver computes this analytic solution inside the rectangular regions and patches the results across the boundaries between these regions using an FDTD stencil.

After the sound pressure in each rectangular subdomain is computed, sound propagation across the interfaces between subdomains must be calculated. ARD uses a (6,2) FDTD stencil in order to patch together subdomains.

The pressure field computation assumes a perfectly reflective boundary condition for the rectangular regions and the interface stencils. However, this is an unrealistic assumption for real-world scenes, which have a variety of different absorbing materials and walls (the free space boundary condition can be modeled as a fully absorptive wall).

ARD implements the Perfectly Matched Layer (PML) scheme, a modified form the wave equation that absorbs propagating waves [30]. PML subdomains, using this modified wave equation, are generated from wall regions and the boundary of the scene and use the same interface scheme to transfer pressure waves between subdomains. The boundary conditions of the ARD method are not frequency dependent, and ARD must be executed for each frequency band.

Frequency-dependent boundary conditions are an open problem for the ARD method.

3.2. Parallel ARD

The parallel ARD method is an adaptation of the ARD method for distributed compute clusters and shared memory machines [26]. MPARD is based on the same framework as parallel ARD, so this section will go over relevant implementation details.

3.2.1. Parallel algorithm

The parallel ARD pipeline is similar to the serial implementation discussed in Section 3.1. Because the subdomains in the ARD decomposition are non-overlapping, the local computation for each rectangular subdomain can be run independent of any other subdomain. Therefore, each core can update a set of subdomains independently of any other core. These subdomains are referred to as *local subdomains* for a particular core and are fully in memory for that core. Subdomains not present on a core are referred to as *remote subdomains*; only metadata is stored for these rectangular regions. This metadata generally includes which rank (core id) owns the subdomain and what interfaces the subdomain is spanned by. The core that is responsible for the local update of a rectangular region is referred to as the *owner* of that subdomain.

However, interfaces between subdomains still need to be evaluated. Each interface can span between two and six subdomains (the interface is three grid cells deep on either side of the boundary between subdomains), but only one core needs to carry out the interface computation. This core is referred to as the *owner* of the interface. The owner of the interface is generally also the owner of one of the rectangular regions that the interface spans. Before interface evaluation is performed, the owner of the interface needs all the pressure terms from the subdomains spanned by the interface. Not all of these subdomains are local; some may be resident on other cores. Therefore, the core evaluating the interface needs to retrieve the pressure data from other cores. Similarly, once

the interface is resolved, the results need to be sent back to the cores owning the subdomains spanned by the interface. Finally global update for each rectangular region is computed with the results from the interface handling. This means that the final global pressure computation requires two data transfers per remote interface.

3.2.2. Load balancing

Due to the greedy nature of the rectangular decomposition algorithm, some rectangles can be quite large. Additionally, because of the stairstepping artifacts of curved geometric surfaces, the decomposition algorithm can generate some very small subdomains that are a single grid cell in each dimension. Because the evaluation cost of the local update of a subdomain is linearly related to the volume of that subdomain, the discrepancy in subdomain sizes can cause load imbalance problems where all the cores wait for one core to finish computing the local update. In fact, without modifying subdomain sizes, increasing the number of cores does not provide any speedup at all [26].

Parallel ARD implements a load balancing scheme that modifies the sizes of the rectangles to be less than a certain threshold volume (for example the total scene volume divided by the number of cores). Splitting is implemented by dividing any rectangle greater than a certain volume into two smaller rectangles. The splitting plane is chosen so that one of the rectangles is just below the threshold volume. This approach can minimize the interface area added under some circumstances, but does not take into account the shape of the subdomains. This can cause numerical stability issues, as discussed in Section 4.2.

Finally rectangular subdomains are assigned to cores using a simple bin-packing approach where the total volume assigned to each core is roughly equal. The bin-packing approach only considers the total volume assigned to each core and does not take into account communication cost of the core assignment.

3.2.3. *Limitations of parallel ARD*

While parallel ARD has been shown to work well for smaller clusters up to 1000 cores, it has many limitations when scaling to very large clusters with tens of thousands of cores:

- Parallel ARD does not take into account the cost of communication between cores
- Parallel ARD suffers from numerical instability when a large number of cores is used
- Parallel ARD initializes interface information in the simulation itself, which can take hours on larger scenes

Large indoor and outdoor scenes at high frequencies on a large number of cores suffer from these issues. On some of our test scenes, hundreds of gigabytes of interface data are created. This means a significant increase in communication cost that is dependent on the assignment of subdomains to cores. Additionally, the Parallel ARD splitting scheme creates thin (less than 3 grid cells thick) and badly formed partitions when splitting the larger partitions. This occurs when the splitting threshold is lower than the smallest surface area of a given partition. Finally, the initialization time ARD requires to compute hundreds of gigabytes of interface data can take hours. On larger clusters where computation time is a valuable resource, this is time wasted.

4. **MPARD**

In order to compute the acoustic field over time for large or high-frequency scenes, MPARD introduces a new parallel algorithm for the ARD method designed to run on tens of thousands of cores. The algorithm includes a modified subdomain assignment scheme that takes into account the communication costs between cores, a modified domain decomposition and splitting stage that increases numerical stability, and new preprocessing stages for computing metadata for interfaces and PML subdomains.

4.1. Communication efficiency

Each interface in the scene covers two or more subdomains (for example, if the subdomains on either side are especially thin, the 6th order stencil may cover more than two). When subdomains on an interface are owned by different cores, the pressure terms required by the interface and the forcing terms generated by the interface need to be communicated.

MPARD uses asynchronous communication calls to avoid blocking while sending messages. When a core completes an operation (either a subdomain update or an interface update), it can send off a message to the cores that require the computed results. The sending core does not have to wait for the message to be received and can continue working on the next computation. When receiving data, a core can place the message on an internal queue until it needs the data for that operation. As a particular core finishes working on a subdomain, it sends off an asynchronous communication message to cores that require the pressure field of that subdomain for interface computation. As the message is being sent, that particular core begins to process the next subdomain, sending off a message upon completion. At the same time the receiving core is working on its own subdomains. The incoming message is stored on an internal queue for use when needed.

Another important thing to note is that in general ARD's communication cost of evaluating all the interfaces is proportional to the surface area of the rectangular region while the cost of evaluating each subdomain is proportional to the volume of the subdomain. This implies a rough $O(n^2)$ running time for interface evaluation (where n is the length of one side of the scene) while local update has an $O(n^3)$ running time. This means that as the size of the scene or the simulation frequency increases, the computation cost of local update dominates over the communication cost of transferring pressure and forcing values for interface evaluation. On the other hand the finer grid size of larger complex scenes can introduce many more interfaces which causes an increase in the cost of communication.

In order to evaluate these interfaces, communication is only required when

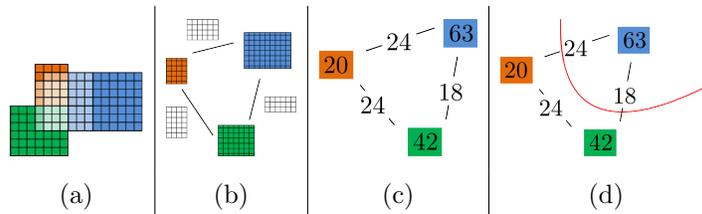


Figure 1: The relationship between computational elements of MPARD (rectangular subdomains and interfaces) and the hypergraph structure. (a) shows an example scene of three subdomains with three interfaces (shown in grey). The organization of the resulting hypergraph is shown in (b), while (c) shows the hypergraph with the node weights determined by the sizes of the rectangles and the hyperedge weights determined by the size of the interfaces. Subfigure (d) shows an example partitioning of the simple graph, taking into account both the computation cost of each node and the cost of each interface.

resolving the interface between two or more subdomains that lie on different cores. As a result, minimizing the number of these interactions can greatly reduce the total amount of communicated data.

This can be performed by ensuring that neighboring subdomains that are part of the same interface are located on the same core. Due to the complexity of the scene and the interactions between different rectangular regions, this is not feasible. However, we can use a heuristic algorithm to minimize the number of interfaces that span across multiple cores.

This problem can be reworded as a hypergraph partitioning problem. Our hypergraph can be represented as the pair $H = (X, E)$ where X , the nodes of the hypergraph are the rectangular regions of our decomposition, while the hyperedges E represent the interfaces between the rectangular regions. Hyperedges are used rather than regular edges because an interface can actually cover more than two subdomains (see Section 3.2.1).

The goal of the hypergraph partitioning algorithm is to divide the hypergraph into k regions such that the cost function of the hyperedges spanning the regions is minimized [9]. In ARD, the partitioning algorithm can be run to divide our computational elements (interfaces and rectangles) into k regions, where k is the number of processors used in the simulation. As a result, the

interface cost between cores is minimized.

Additionally, because the hypergraph partitioning algorithm attempts to generate k regions of equal cost, the heuristic serves as a way of load balancing the assignment of work to cores. The cost of evaluating a rectangular region is linearly related to the volume of the region. Therefore, we can input the volume of each rectangular subdomain as the weight parameter for a node in the hypergraph.

We implemented a hypergraph partitioning scheme through the PaToH (Partitioning Tool for Hypergraphs) library in order to minimize core-to-core communication and equally distribute computation load across all cores [9]. We pass the rectangular regions into the PaToH partitioner as vertices with weights equal to the respective subdomain volumes. The hyperedges are defined by the interfaces connecting separate subdomains. By partitioning the resulting hypergraph, the resulting core assignments will reduce overall computation cost, since neighboring vertices will tend to be assigned to the same core and they will not require communication. This also allows for a balanced load distribution by assigning each core roughly equal volume, which is linearly related to computation time. We use the following connectivity metric to determine the cost of communication across boundaries:

$$C(\Pi) = \sum_{n \in N_E} w_n (\lambda_n - 1), \quad (2)$$

where $C(\Pi)$ is the cost of a particular partitioning Π , N_E is the set of cut hyperedges, w_n is the weight of hyperedge n , and λ_n is the connectivity of the hyperedge. This metric is particularly useful because it prioritizes hyperedges that connect many vertices and is such directly proportional to the computation cost.

4.2. Load balancing and numerical stability

The splitting algorithm for load balancing introduced in parallel ARD minimizes the number of extra interfaces created by splitting in some cases. The

larger of the resulting subdomains is exactly below the maximum volume threshold of the splitting algorithm.

However, on a cluster with a higher number of cores and where the volume threshold can be relatively small, this splitting algorithm can introduce a series of very small and thin rectangles. Small and thin rectangles can create numerical instability during interface resolution. These errors are caused by the interaction between multiple overlapping interfaces.

In this particular case, it is more advantageous to have wider rectangular subdomains that may introduce more interface area rather than degenerate rectangles which can result in numerical issues.

We introduce a new splitting scheme for load balancing that is particularly useful for reducing numerical instability at interfaces for higher numbers of cores. Our new approach favors well-formed subdomains that are more cuboidal in shape rather than long and thin rectangles.

Each rectangular region that has a volume greater than some volume threshold Q is subdivided by the new algorithm. Q is determined by the equation

$$Q = \frac{V}{pf}, \quad (3)$$

where V is the total air volume of the scene in spatial discretization units, p is the number of cores the solver is to be run on, and f is the balance factor. It is usually the case that $f = 1$, although this can be changed to a higher value if smaller rectangles are desired for the hypergraph partitioning heuristic. However, in general, smaller values of Q increase the overall interface error of the solver since it increases the total interface area.

4.3. Interface and PML computation

An important step of the ARD and parallel ARD methods is the initialization of interfaces and PML regions. This generally involves determining subdomain adjacency and which grid cells are in an interface. This computation is linear with respect to the number of grid cells in the scene – that is, it scales linearly with the volume of the scene and with the 4th power of frequency. At higher

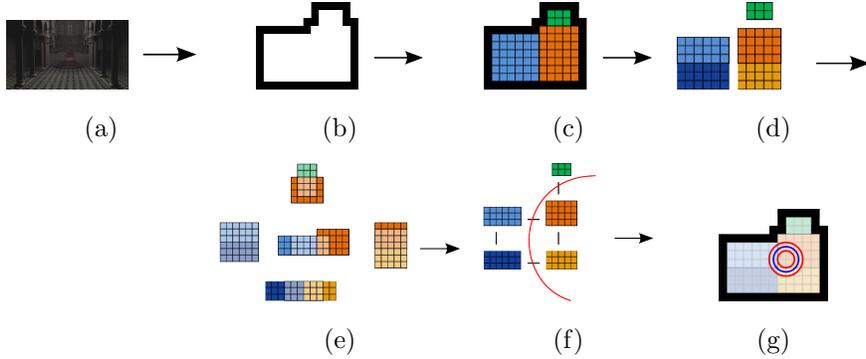


Figure 2: The MPARD pipeline. The input geometry (a) is voxelized in the first step (b). The rectangular decomposition step divides the domain into multiple non-overlapping subdomains (c). The splitting step then splits these subdomains when they are greater than the volume threshold Q (d). These partitions are then processed by the interface initialization stage which computes interface metadata (e). The final preprocessing stage allocates subdomains to nodes using the hypergraph partitioning (f). Finally, the simulation is run (g).

frequencies, the interface and PML setup can consume several hours of time before any actual simulation steps are run.

In order to reduce the running time of the simulation, MPARD introduces a new preprocessing step in which the interfaces can be initialized offline. This preprocessing step occurs after any splitting and load balancing, after the decomposition for the scene is final.

Additionally, this extra preprocessing step allows for further memory optimization in MPARD. With the kind of global metadata computed in the preprocessing, each core only needs to load the exact interfaces and PML regions it needs for its computations.

4.4. MPARD Pipeline

MPARD introduces new preprocessing stages and a modified simulation stage. The pipeline overview can be found in Figure 2. The scene input (Figure 2(a)) is first voxelized (Figure 2(b)). Next, the rectangular decomposition fills the available air space with rectangular subdomains using a greedy algorithm (Figure 2(c)). Next, our new splitting algorithm that avoids degenerate

subdomains splits rectangular regions that are larger than the volume threshold (Figure 2(d)). After interface regions are calculated (Figure 2(e)), we assign subdomains to cores using hypergraph partitioning (f). Finally, our solver reads the preprocessing data and runs the simulation for a set number of time steps.

4.4.1. Voxelization

The voxelization stage takes in a triangle mesh representing the environment in which we want to compute the sound propagation. Because MPARD targets large and high frequency scenes that may consume a large amount of memory, we use a CPU method for voxelization. We implement an *accurate* and *minimal* method (meaning voxels should fully cover the geometry but not more than necessary) as introduced by Huang et al. [21].

The spatial discretization for the voxelization is determined by the minimum simulated wavelength and the required number of spatial samples per wavelength, which is typically between 2 and 4 [29]. Therefore, the voxelization only needs to be run once per desired maximum frequency.

4.4.2. Decomposition

The decomposition stage then reads the voxel field and determines the location of the different cuboidal subdomains. The process is a greedy approach, attempting to expand each rectangular subdomain into as large a volume as possible under the constraints of the wall voxels.

At very high frequencies, such as 10 kHz, this process can take several days to complete but only needs to run once for a given voxel input.

4.4.3. Core allocation and subdomain splitting

The next stage of the preprocessing is the core allocation and subdomain splitting stage. In addition to a decomposition computed in the previous stage, this step also requires the number of cores the solver will run on. This stage uses the input values to compute a hypergraph partitioning for the decomposition in addition to splitting any rectangular regions that have volumes greater than the volume threshold Q .

The core allocation stage then determines the assignment of subdomains to cores by using the hypergraph partitioning assignment or alternatively a simple bin-packing algorithm. This load balancing step ensures that each core has a roughly equal amount of work to complete during the acoustic simulation.

The allocation and splitting stage must be run for each decomposition for each desired core configuration.

4.4.4. Interface and PML preprocessing

The final stage of preprocessing computes interfaces and creates PML regions from wall voxels. This stage takes as input a modified decomposition from the core allocation and splitting stage in addition to a refinement parameter r that can be used to subdivide voxels in the final acoustic simulation. This allows us to run at r times the frequency the decomposition was run at. However, this is at the expense of some accuracy where high frequency geometric features of the scene that may affect sound propagation cannot be accurately represented.

One additional caveat of the interface and PML preprocessing file is file read performance in the simulator. The interface file can be several GBs in size, and thousands of CPU cores reading the file can cause a bottleneck. As a solution, we use the file striping feature of the Lustre file system [39] to increase file read performance over all cores.

The interface and PML initialization stage only needs to be run once for each core configuration.

5. Results and analysis

Our method was tested on two computing clusters: the large-scale Blue Waters supercomputer [7] at the University of Illinois and the UNC KillDevil cluster. Blue Waters is one of the world’s leading compute clusters, with 362240 XE Bulldozer CPU cores and 1.382 PB of memory. The KillDevil cluster has 9600 CPU cores and 41 TB of memory.

Our primary experiments were performed on the Sibenik Cathedral scene and the Village scene (Figure 2). Both scenes provide a challenge for the un-

Scene Name	Volume	Frequency	Number of Triangles
Cathedral	19 177 m ³	5 kHz, 10 kHz	55415
Village	362 987 m ³	1.5 kHz	358

Table 2: Dimensions and complexity of the scenes used in our experiments. The input triangle mesh is voxelized according to the simulation frequency.

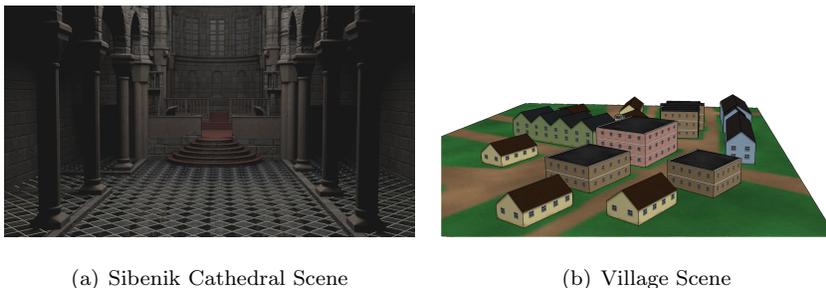
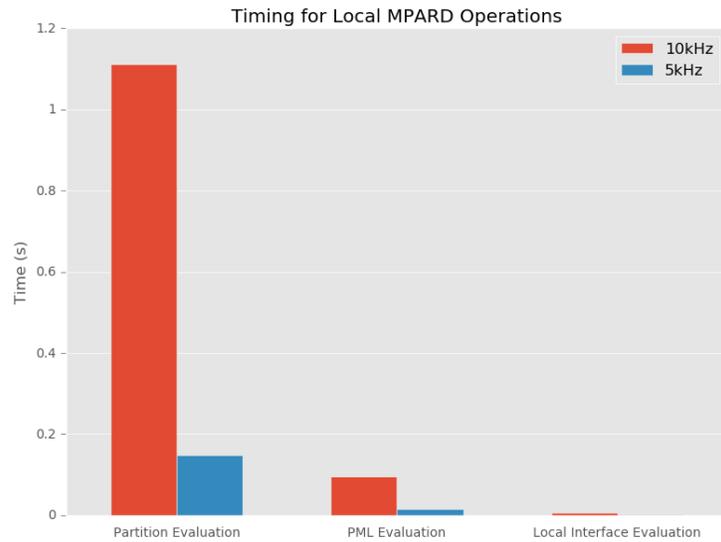


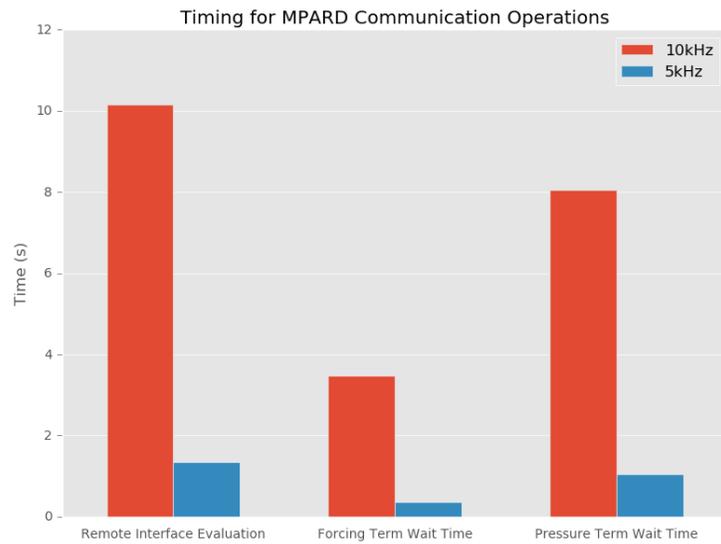
Figure 3: The scenes used in our experiments.

derlying ARD solver. Cathedral has many curved surfaces, creating very small rectangular regions in the rectangular decomposition. Additionally, the large areas in the center of the cathedral creates very large rectangular subdomains. Furthermore, the size of the scene is around 20 000 m³, making communicating the sound propagation of the scene at high frequencies with wave-based methods very challenging. For example, a 10 kHz voxelization of the cathedral has almost 4 billion voxels. On the other hand, Village is mostly a large open area with a few buildings (Figure 3). Village is also much larger than cathedral. The size of the scene (362 000 m³) presents many computational challenges, particularly with computing interfaces. The scene contains over 100GB of interface data compared to 40GB in the 10 kHz Cathedral scene. On the other hand, the cathedral scene contains many more partitions (176k compared to 31k), making communication more of a challenge.

We were able to run the MPARD solver on the cathedral scene up to 10 kHz. Figure 4 shows the average running time of each stage of our algorithm on this scene in comparison to the 5 kHz. The wait time for interface terms shows the



(a) Local Timings



(b) Communication Timings

Figure 4: The average running time of each stage of our solver on a 10 kHz scene compared to the 5 kHz scene. These results were obtained on 1024 cores. The communication timings show the various stages of interface evaluation and communication. Remote Interface evaluation is the compute cost of the interface, while the Forcing Term Wait time is the communication cost of transferring forcing terms from the interface to the spanned subdomains, and the Pressure Term Wait time is the communication cost of transferring pressure terms to the interface from the spanned subdomains.

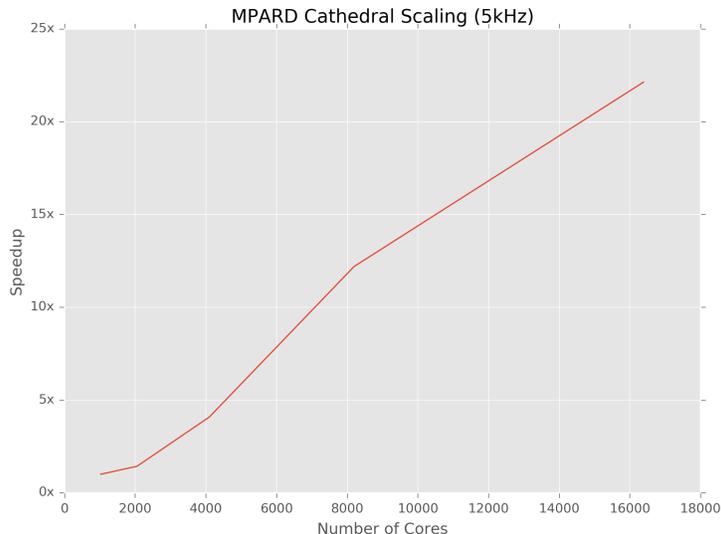


Figure 5: Scalability results from 1024 to 16384 cores on the 5 kHz cathedral scene. We obtain close-to-linear scaling in this result. The base speedup is 1024 on 1024 cores (assuming 1024 cores is 1024 times faster than one core) since the scene will not fit in memory on a lower number of cores.

necessity for optimizing communication at higher frequencies.

5.1. Scalability results

The primary scalability experiment was done on the cathedral scene at 5 kHz. The experiment was executed on the Blue Waters supercomputer up to 16384 cores. The main purpose of this experiment is to understand the performance of MPARD at very high number of cores. Figure 5 shows the performance of MPARD on the cathedral scene for 1024 cores all the way up to 16384 cores. With this kind of compute power, we are able to compute each time step on Cathedral in 0.193 75 s for a 5 kHz scene.

The Village scene also shows scalability up to 8192 cores. The dominating cost in the village scene is the computation time of the interfaces (Figure 7). We show sublinear scaling in this case, with compute times as fast as 0.6761 s per time step for a 1.5 kHz scene.

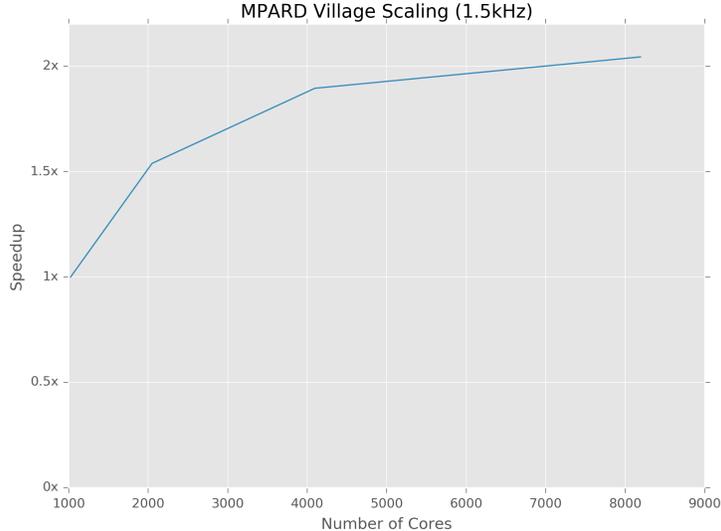


Figure 6: Scalability results from 1024 to 8192 cores on the 1.5 kHz village scene. We obtain sublinear scaling in this result. The base speedup is 1024 on 1024 cores since the scene will not fit in memory on a lower number of cores.

Figure 7 shows a plot of how the interface area generated by our splitting scheme increases as the number of cores used in the splitting algorithm increases. We show a comparison between the Cathedral and Village scenes.

5.2. Comparisons and benefits

In comparison to standard methods like FDTD, ARD does not require as fine of a computational grid. Traditional FDTD methods generally require a spatial discretization that is $1/10$ the minimum wavelength (although the low-dispersion methods listed in Section 2.3 aim to lower this requirement). In comparison ARD can use a much coarser grid size, around $1/2.6$ times the minimum wavelength [29, 25]. This means that ARD can inherently be 24 – 50 times more memory efficient than FDTD and up to 75 – 100 times faster [26].

5.2.1. Comparison with GPU ARD

MPARD provides advantages over previous GPU parallel ARD approaches. Large scenes at high frequencies require terabytes of memory that are easily

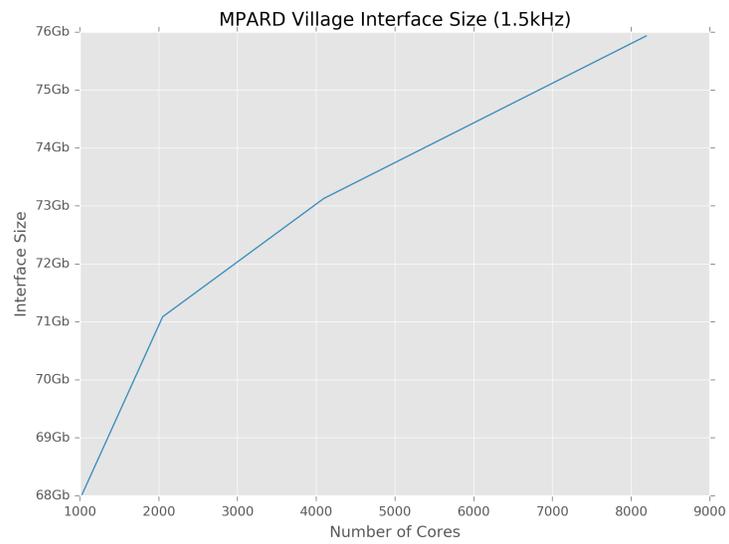
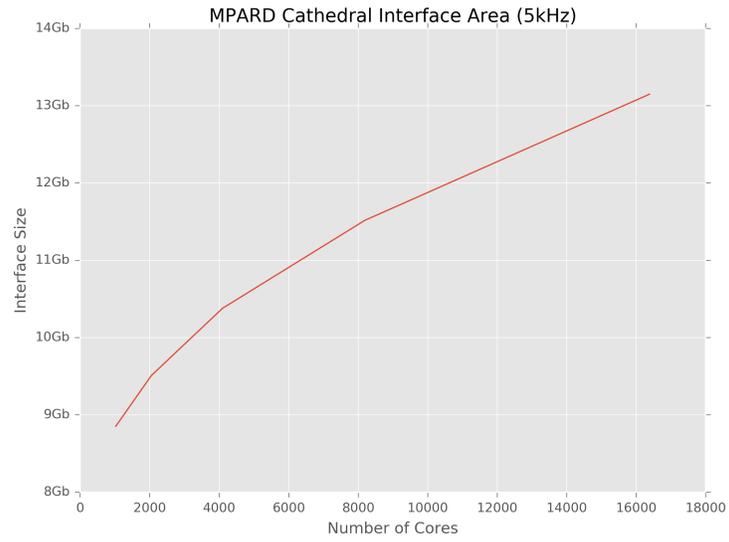


Figure 7: Interface area generated by our splitting scheme for the Village and Cathedral scenes. The interface area tapers off as the number of cores increases.

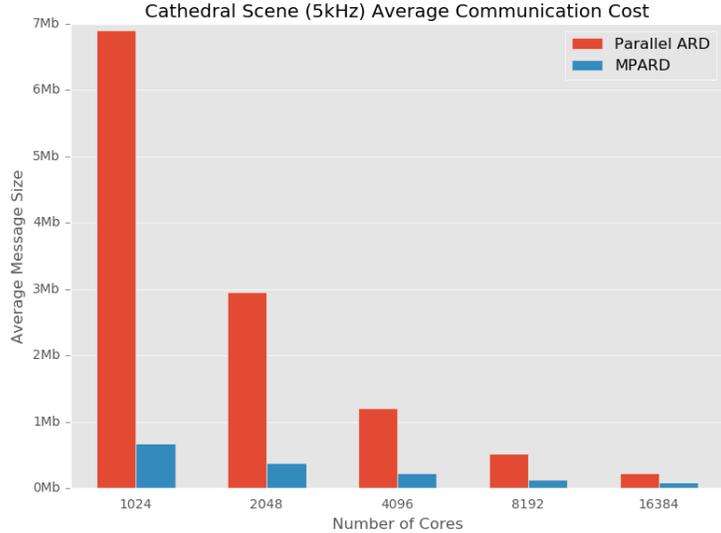


Figure 8: Communication cost comparison between MPARD and parallel ARD on a 5 kHz cathedral scene. In this case, the large number of interfaces in the cathedral scene means that the reduction in communication cost using MPARD is much larger.

available on large compute clusters but are not available on GPUs [25]. Secondly, MPARD scales over a much larger number of cores while GPU ARD is limited to a single machine and a shared memory architecture.

5.2.2. Comparison with Parallel ARD

Communication costs In order to examine the benefits of hypergraph partitioning, the total size of all messages averaged per core sent during a single simulation time step of both scenes was computed. The scenes had cores assigned through the parallel ARD bin-packing approach and the other had cores assigned through the MPARD hypergraph partitioning approach [26]. Figure 8 shows the results of this experiment on a 5 kHz scene where the hypergraph partitioning approach has a 10x reduction in communication costs per core for 1024 cores and a 3x reduction per core for 16384 cores. The village scene, half as large in memory also had a reduction in communication costs (Figure 9).

Subdomain splitting Figure 10 shows the difference in decomposition be-

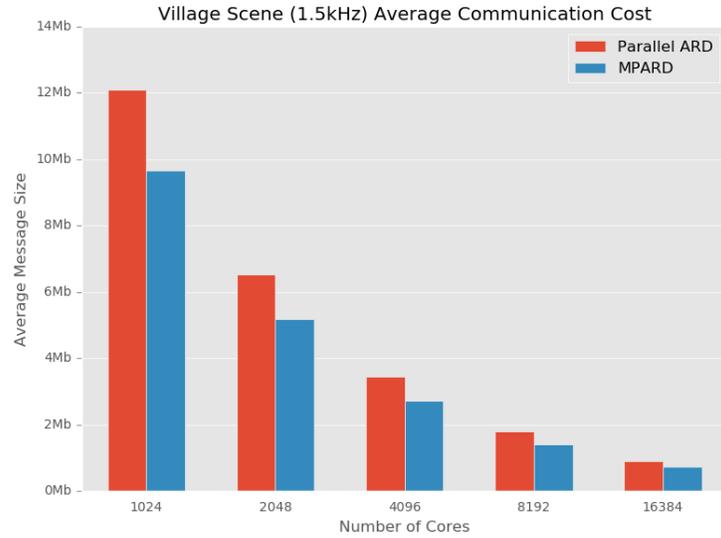


Figure 9: Communication cost comparison between MPARD and parallel ARD on a 1.5 kHz village scene. The village scene has fewer interfaces so there is less reduction in communication cost, but MPARD still is more efficient in this case.

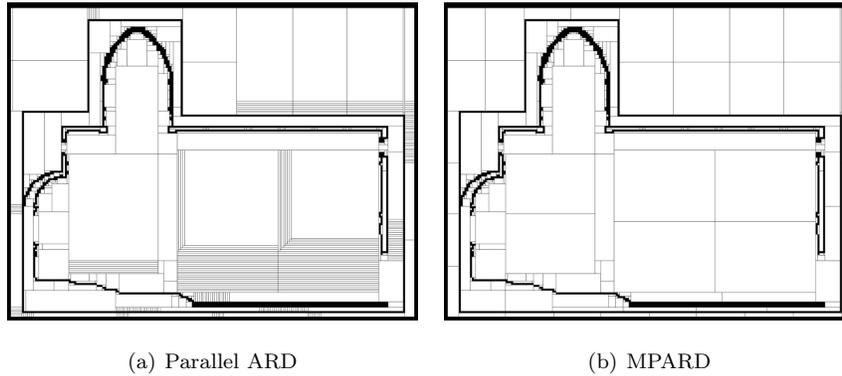


Figure 10: 2D slice comparison between parallel ARD and MPARD splitting methods. Notice the series of thin rectangles that are generated in parallel ARD [26] and can result in numerical stability problems because of the width of the partition. MPARD does not have similar problems and is more stable than parallel ARD. This result was computed on a decomposition for 1024 cores on a 500 Hz scene.

tween the old conservative approach and the new subdivision approach on the Cathedral scene. Both scene decompositions were run on 1024 cores to show the difference between the two approaches at a high number of cores. The parallel ARD approach creates a series of very thin rectangles along the edge of the interface while the new splitting scheme creates more cuboidal subdomains.

As discussed in Section 3.2.3, the thin rectangular decomposition can cause numerical instability. This numerical instability is not present in MPARD simulations. Even as the number of cores increase to 16000, the error of sampled impulse responses does not exceed 5% (see Figure 11). The error was computed using the following metric:

$$\frac{\sum_{i=1}^n \left(p'_i(\vec{x}) - p_i(\vec{x}) \right)^2}{\sum_{i=1}^n p_i^2}, \quad (4)$$

where p_i is the reference pressure computed in Equation 1 at listener position \vec{x} without partition splitting, p'_i is the pressure computed with partition splitting, and n is the total number of time steps.

A comparison in Figure 12 between a reference single-threaded impulse response and the impulse response calculated on a 16384 core run shows how various features of the impulse responses match. Figure 13 shows how the error of the full pressure field varies over 2000 time steps, until most of the sound dissipates. For the full-field error, we used the mean-squared error metric for each time-step:

$$\frac{1}{m} \sum_{j=1}^m \left(p'_j - p_j \right)^2 \quad (5)$$

where m is number of grid points in the volume, p_j is the reference pressure at grid index j , and p'_j is the computed pressure at grid index j .

This shows MPARD's stability and accuracy over the whole course of a full simulation. For these experiments, we followed a similar experimental setup as previous work that have validated the serial ARD technique [24]. This work focused on the village scene, a digital reproduction of a village in which the

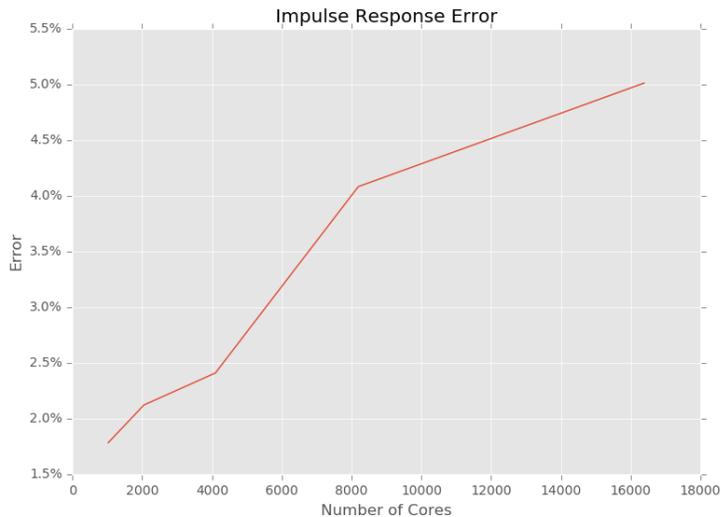


Figure 11: Error of an impulse response taken around 10m away from the source on the Village scene using the MPARD method for a single band impulse at 225 Hz. The impulse response is calculated for 1024 cores all the way to 16384 cores. Subdomain splitting increases as the number of cores increases, so the error also increases. Even for an extreme number of cores (16384), the error is limited to around 5%.

propagation of various sound sources was measured. The measured impulse responses at various listener positions were compared to impulse responses computed by the ARD method. The comparison was done by comparing average dB and comparing spectrograms with dynamic time warping to account for slight differences in the arrival times of different wave forms. We compared our results directly to the serial ARD implementation of this scene, and so did not have to use dynamic time warping techniques.

6. Conclusion and future work

MPARD is a massively parallel approach showing scalability up to 16000 cores and is capable of calculating pressure fields for large scenes with frequencies up to 10 kHz. MPARD introduces several improvements over parallel ARD, reducing communication cost by assigning cores with a hypergraph partition-

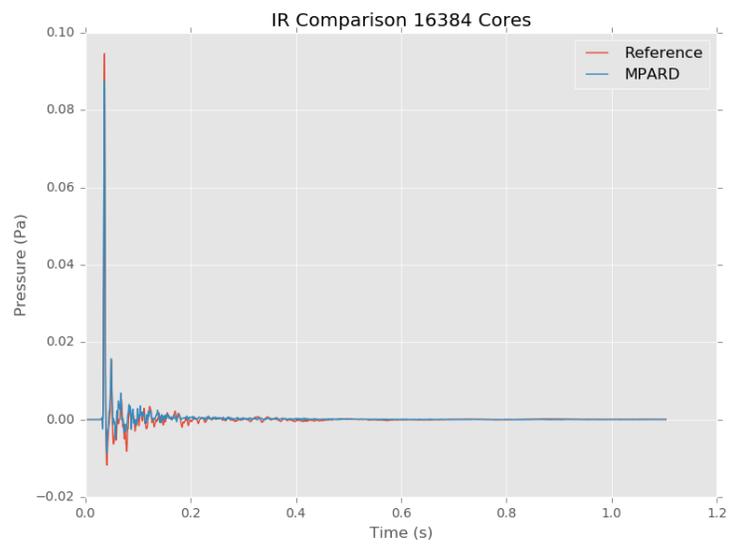


Figure 12: Comparison of impulse responses 10 m away from a 225 Hz source. The reference impulse response uses the serial implementation of ARD [24], while the simulated impulse response uses a subdomain partitioning for 16384 cores. We show that the impulse response matches closely with the serial implementation validated in Mehra et al. [24].

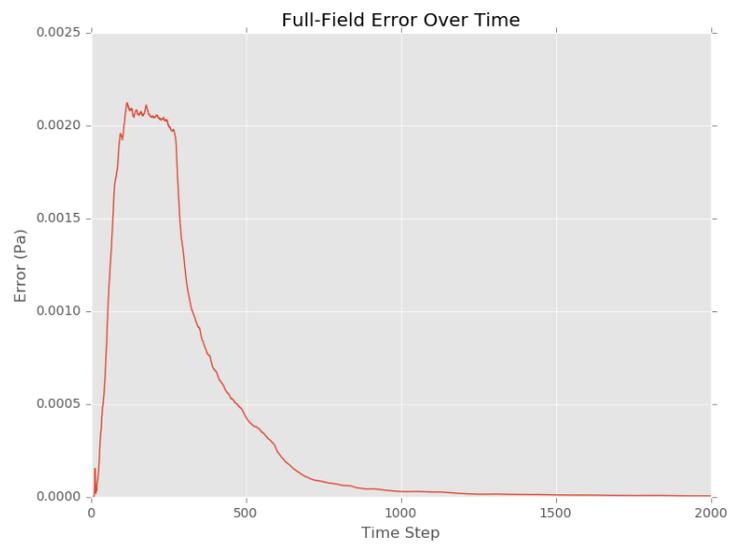


Figure 13: Error of the full field MPARD simulation over time on a 16384 core simulation excited by a single source at 225 Hz. The error is introduced by spurious reflections and dispersion, but does not accumulate as time increases. The error was computed using the mean-squared error for each time step.

ing scheme, providing better numerical stability for higher numbers of cores, and implementing an extended preprocessing pipeline that reduces the usage of valuable cluster resources for redundant calculations.

In the future, although we have shown the ability to compute very high frequency scenes, we would like to test the method on more scenes including even larger architectural and outdoor scenes at high frequencies. There has already been some work on hybrid sound propagation, using geometric methods for higher frequencies and ARD for lower frequencies [46], so we would like to expand on this for large outdoor scenes.

Scaling to large outdoor scenes such as cities or towns poses many technical challenges. Although we reduce communication costs between cores with a hypergraph partitioning technique, a hybrid shared/distributed memory approach may reduce communication costs even further. Additionally, we could take advantage of the computational power of GPUs for a fully heterogenous computing algorithm.

Acknowledgments

This research was supported in part by ARO Contracts W911NF-13-C-0037, W911NF-14-1-0437 and the National Science Foundation (NSF awards 1345913 and 1456299). This research is part of the Blue Waters sustained-petascale computing project and the Blue Waters Student Internship Program, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. We want to thank Allan Porterfield, Alok Meshram, Lakulish Antani, Anish Chandak, Joseph Digerness, Alban Bassuet, Keith Wilson, and Don Albert for useful discussions and models.

- [1] A. Alghamdi, A. Ahmadi, D. I. Ketcheson, M. G. Knepley, K. T. Mandli, and L. Dalcin. Petclaw: a scalable parallel nonlinear wave propagation solver for python. In *Proceedings of the 19th High Performance Computing*

- Symposia*, HPC '11, pages 96–103, San Diego, CA, USA, 2011. Society for Computer Simulation International.
- [2] J. B. Allen and D. A. Berkley. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America*, 65(4):943–950, April 1979.
- [3] H. Bao, J. Bielak, O. Ghattas, L. F. Kallivokas, D. R. O’Hallaron, J. R. Shewchuk, and J. Xu. Large-scale simulation of elastic wave propagation in heterogeneous media on parallel computers. *Computer methods in applied mechanics and engineering*, 152(1):85–102, 1998.
- [4] M. Bernacki, L. Fezoui, S. Lanteri, and S. Piperno. Parallel discontinuous galerkin unstructured mesh solvers for the calculation of three-dimensional wave propagation problems. *Applied mathematical modelling*, 30(8):744–763, 2006.
- [5] M. A. Bhandarkar and L. V. Kalé. A parallel framework for explicit fem. In *High Performance Computing HiPC 2000*, pages 385–394. Springer, 2000.
- [6] S. Bilbao. Modeling of complex geometries and boundary conditions in finite difference/finite volume time domain room acoustics simulation. *Audio, Speech, and Language Processing, IEEE Transactions on*, 21(7):1524–1533, 2013.
- [7] B. Bode, M. Butler, T. Dunning, W. Gropp, T. Hoe-fler, W.-m. Hwu, and W. Kramer. The blue waters super-system for super-science. In C. Jeffrey S . Vetter and H. 2013, editors, *Contemporary HPC Architectures*. Sitka Publications, 2012.
- [8] C. A. Brebbia. *Boundary Element Methods in Acoustics*. Springer, 1 edition, Oct. 1991.
- [9] Ü. Çatalyürek and C. Aykanat. Patoh (partitioning tool for hypergraphs). In *Encyclopedia of Parallel Computing*, pages 1479–1487. Springer, 2011.

- [10] T. F. Chan and O. B. Widlund. Domain decomposition algorithms. In *Acta Numerica*, pages 61–143. 1994.
- [11] J.-T. Chen, Y.-T. Lee, and Y.-J. Lin. Analysis of multiple-sources radiation and scattering problems by using a null-field integral equation approach. *Applied Acoustics*, 71(8):690–700, 2010.
- [12] M. J. Crocker. *Handbook of Acoustics*. Wiley-IEEE, 1998.
- [13] R. Diekmann, U. Dralle, F. Neugebauer, and T. Römke. Padfem: a portable parallel fem-tool. In *High-Performance Computing and Networking*, pages 580–585. Springer, 1996.
- [14] B. Engquist and O. Runborg. Computational high frequency wave propagation. *Acta Numerica*, 12:181–266, 2003.
- [15] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, and J. West. A beam tracing approach to acoustic modeling for interactive virtual environments. In *Proc. of ACM SIGGRAPH*, pages 21–32, 1998.
- [16] T. Funkhouser, N. Tsingos, and J.-M. Jot. Survey of Methods for Modeling Sound Propagation in Interactive Virtual Environment Systems. *Presence and Teleoperation*, 2003.
- [17] C. Guiffaut and K. Mahdjoubi. A parallel fdtd algorithm using the mpi library. *Antennas and Propagation Magazine, IEEE*, 43(2):94–103, 2001.
- [18] M. Hornikx and J. Forssén. The 2.5-dimensional equivalent sources method for directly exposed and shielded urban canyons. *The Journal of the Acoustical Society of America*, 122(5):2532–2541, 2007.
- [19] M. Hornikx, T. Krijnen, and L. van Harten. openpstd: The open source pseudospectral time-domain method for acoustic propagation. *Computer Physics Communications*, 203:298–308, 2016.

- [20] M. Hornikx, R. Waxler, and J. Forssén. The extended fourier pseudospectral time-domain method for atmospheric sound propagation. *The Journal of the Acoustical Society of America*, 128(4):1632–1646, 2010.
- [21] J. Huang, R. Yagel, V. Filippov, and Y. Kurzion. An accurate method for voxelizing polygon meshes. In *Volume Visualization, 1998. IEEE Symposium on*, pages 119–126. IEEE, 1998.
- [22] K. Kowalczyk and M. van Walstijn. Room acoustics simulation using 3-d compact explicit ftd schemes. *Audio, Speech, and Language Processing, IEEE Transactions on*, 19(1):34–46, 2011.
- [23] H. Kuttruff. *Room acoustics*. CRC Press, 2009.
- [24] R. Mehra, N. Raghuvanshi, A. Chandak, D. G. Albert, D. K. Wilson, and D. Manocha. Acoustic pulse propagation in an urban environment using a three-dimensional numerical simulation. *The Journal of the Acoustical Society of America*, 135(6):3231–3242, 2014.
- [25] R. Mehra, N. Raghuvanshi, L. Savioja, M. C. Lin, and D. Manocha. An efficient gpu-based time domain solver for the acoustic wave equation. *Applied Acoustics*, 73(2):83 – 94, 2012.
- [26] N. Morales, R. Mehra, and D. Manocha. A parallel time-domain wave simulator based on rectangular decomposition for distributed memory architectures. *Applied Acoustics*, 97:104–114, 2015.
- [27] S. Operto, J. Virieux, P. Amestoy, J.-Y. L'Excellent, L. Giraud, and H. B. H. Ali. 3d finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study. *Geophysics*, 72(5):SM195–SM211, 2007.
- [28] B. Pluymers, B. Van Hal, D. Vandepitte, and W. Desmet. Trefftz-based methods for time-harmonic acoustics. *Archives of Computational Methods in Engineering*, 14(4):343–381, 2007.

- [29] N. Raghuvanshi, R. Narain, and M. C. Lin. Efficient and Accurate Sound Propagation Using Adaptive Rectangular Decomposition. *IEEE Transactions on Visualization and Computer Graphics*, 15(5):789–801, 2009.
- [30] Y. S. Rickard, N. K. Georgieva, and W.-P. Huang. Application and optimization of pml abc for the 3-d wave equation in the time domain. *Antennas and Propagation, IEEE Transactions on*, 51(2):286–295, 2003.
- [31] J. Saarela and L. Savioja. An open source finite-difference time-domain solver for room acoustics using graphics processing units. *Acta Acustica united with Acustica*, 2014.
- [32] F. Saied and M. J. Holst. Multigrid methods for computational acoustics on vector and parallel computers. *Urbana*, 51:61801, 1991.
- [33] S. Sakamoto, A. Ushiyama, and H. Nagatomo. Numerical analysis of sound propagation in rooms using the finite difference time domain method. *The Journal of the Acoustical Society of America*, 120(5):3008, 2006.
- [34] T. Sakuma, S. Sakamoto, and T. Otsuru. *Computational simulation in architectural and environmental acoustics*. Springer, 2014.
- [35] E. M. Salomons, W. J. Lohman, and H. Zhou. Simulation of sound waves using the lattice boltzmann method for fluid flow: Benchmark cases for outdoor sound propagation. *PloS one*, 11(1):e0147206, 2016.
- [36] L. Savioja. Real-time 3d finite-difference time-domain simulation of low- and mid-frequency room acoustics. In *13th Int. Conf on Digital Audio Effects*, volume 1, page 75, 2010.
- [37] L. Savioja, J. Backman, A. Järvinen, and T. Takala. Waveguide mesh method for low-frequency simulation of room acoustics. 1995.
- [38] L. Savioja and V. Valimäki. Reducing the dispersion error in the digital waveguide mesh using interpolation and frequency-warping techniques. *Speech and Audio Processing, IEEE Transactions on*, 8(2):184–194, 2000.

- [39] P. Schwan. Lustre: Building a file system for 1000-node clusters. In *Proceedings of the 2003 Linux Symposium*, volume 2003, 2003.
- [40] J. Sheaffer and B. Fazenda. Wavecloud: an open source room acoustics simulator using the finite difference time domain method. *Acta Acustica united with Acustica*, 2014.
- [41] P. Sypek, A. Dziekonski, and M. Mrozowski. How to render ftdt computations more effective using a graphics accelerator. *Magnetics, IEEE Transactions on*, 45(3):1324–1327, 2009.
- [42] L. L. Thompson. A review of finite-element methods for time-harmonic acoustics. *The Journal of the Acoustical Society of America*, 119(3):1315–1330, 2006.
- [43] A. Vaccari, A. Cala’Lesina, L. Cristoforetti, and R. Pontalti. Parallel implementation of a 3d subgridding ftdt algorithm for large simulations. *Progress In Electromagnetics Research*, 120:263–292, 2011.
- [44] S. Van Duyne and J. O. Smith. The 2-d digital waveguide mesh. In *Applications of Signal Processing to Audio and Acoustics, 1993. Final Program and Paper Summaries., 1993 IEEE Workshop on*, pages 177–180. IEEE, 1993.
- [45] C. J. Webb and S. Bilbao. Binaural simulations using audio rate ftdt schemes and cuda. In *Proc. of the 15th Int. Conference on Digital Audio Effects (DAFx-12), York, United Kingdom*, 2012.
- [46] H. Yeh, R. Mehra, Z. Ren, L. Antani, D. Manocha, and M. Lin. Wave-ray coupling for interactive sound propagation in large complex scenes. *ACM Transactions on Graphics (TOG)*, 32(6):165, 2013.
- [47] W. Yu, Y. Liu, T. Su, N.-T. Hunag, and R. Mittra. A robust parallel conformal finite-difference time-domain processing package using the mpi library. *Antennas and Propagation Magazine, IEEE*, 47(3):39–59, 2005.

- [48] W. Yu, X. Yang, Y. Liu, L.-C. Ma, T. Sul, N.-T. Huang, R. Mittra, R. Maaskant, Y. Lu, Q. Che, et al. A new direction in computational electromagnetics: Solving large problems using the parallel fdtd on the bluegene/l supercomputer providing teraflop-level performance. *Antennas and Propagation Magazine, IEEE*, 50(2):26–44, 2008.
- [49] Y. Q. Zeng, Q. H. Liu, and G. Zhao. Multidomain pseudospectral time-domain (pstd) method for acoustic waves in lossy media. *Journal of Computational Acoustics*, 12(03):277–299, 2004.