DOI: 10.61782/fa.2023.0147

MEMORY EFFICIENT GPU-ACCELERATED 3D FDTD ACOUSTIC SIMULATIONS IN TURBULENT FIELDS

Hazael Montanaro[1]* **Dorothea Lincke**[1] **Timothy Van Renterghem**[2]
Kurt Heutschi[1] **Reto Pieren**[1]
[1] Empa, Swiss Federal Laboratories for Materials Science and Technology, Dubendorf, Switzerland
[2] Ghent University, WAVES Research Group, Gent-Zwijnaarde, Belgium

ABSTRACT

Sound attenuation calculations via engineering models for long-range propagation may exhibit significant variations. While some explicitly account for turbulence, others only consider it implicitly through factors such as screening and the ground effect. This makes an intercomparison challenging. Physics-based simulations capturing the appropriate turbulent spectral statistics could serve as a useful benchmark reference case in lieu of controlled experimental data under different meteorological conditions and topographies. To overcome memory limitations inherent in broadband simulations over large propagation distances, we extended our acoustic GPU-accelerated 3D Finite-Difference Time-Domain solver *Empapu* supporting arbitrary geometries and locally-reacting boundary conditions with a dynamic moving frame and turbulence as a position-dependent refractive index. The dynamic moving frame tracks the main wavefront in the direction of interest via index shifting and the turbulent field is calculated on-the-fly via an eddy (quasiwavelet) method reproducing the Kolmogorov spectrum superposed with a vertical sound speed profile in a memory-efficient manner. We test our approach by simulating propagation distances of up to 500 m with a 1.5 cm resolution in different turbulent configurations. This work will serve as the basis to assess engineering models in Switzerland such as CNOSSOS-EU, ISO 9613-2, and sonX.

Keywords: *FDTD, engineering model, turbulence, moving frame, HPC*

1. INTRODUCTION

Sound propagation calculations are vital for understanding and predicting acoustic wave behavior in outdoor environments, especially in assessing environmental noise impact. Engineering models like CNOSSOS-EU [1], ISO 9613-2 [2], and sonX [3], used for outdoor long-range sound attenuation calculations, have limitations and varying accuracy, particularly concerning meteorological influences such as atmospheric turbulence and sound speed gradients. They struggle with uneven terrains and multiple reflections, posing challenges for model validation. Due to the lack of controlled experimental data covering diverse meteorological conditions and topographies, physics-based simulations that capture turbulent spectral statistics can serve as valuable benchmarks for assessing and improving these models. However, high-resolution simulations over large distances demand considerable computational power and memory resources.

We addressed these challenges by extending our GPU-accelerated 3D Finite Difference Time Domain (FDTD) acoustic solver, *Empapu* [4], with a dynamic moving frame and position-dependent refractive index to model random turbulent fields. The moving frame tracks the main wavefront using index shifting and computes the turbulent field on-the-fly with an eddy method [5], reproducing the Kolmogorov spectrum in a memory-efficient manner. By managing memory constraints, offering a controlled environment for turbulence studies, and maintaining accuracy, our method provides a viable wave-based tool for validating and refining long-range propa-

---

**Corresponding author*: *hazael.montanaro@empa.ch.*
***Copyright**: ©2023 Hazael Montanaro et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.*

gation engineering models across various meteorological conditions and topographies.

2. BACKGROUND

2.1 Sound Attenuation Engineering Models

Engineering models such as CNOSSOS-EU, ISO 9613-2, and sonX are widely used to estimate sound attenuation from outdoor noise sources like roads and railways, and are crucial for assessing environmental noise impacts and designing effective mitigation strategies. Sound attenuation is typically calculated as the sum of individual propagation effects, including geometrical spreading, atmospheric absorption, ground effect, and possible shielding by obstacles. However, when dealing with reflections at surfaces of limited size, propagation over uneven terrain, or an inhomogeneous atmosphere, engineering models make various simplifying assumptions, leading to significant uncertainty. This uncertainty makes it challenging to validate the accuracy and applicability of these models in different outdoor environments, highlighting the need for a precise, physics-based tool to serve as a benchmark for comparisons and validation.

2.2 3D GPU accelerated solver

Empapu is an in-house 3D GPU-accelerated FDTD software (see [6] for an overview of the numerical method) written in C++ and CUDA solving the acoustic wave equation

$$\nabla p = -\rho \frac{\partial \bar{v}}{\partial t} \tag{1}$$

$$\frac{\partial p}{\partial t} = -\rho c^2 \nabla \cdot \bar{v} \tag{2}$$

subject to various boundary conditions, where p represents sound pressure in Pa, $\bar{v}$ is particle velocity in m/s, ρ is media density in kg/m^3, and c is media speed of sound in m/s. The solver supports modeling of heterogeneous media, arbitrary geometries coupled with locally-reacting frequency-dependent boundary conditions, and perfectly matched layers (PML) on the outer boundary to simulate outdoor infinitely extended domains and terrains. The solver is further supported by a Python-based voxeling engine and pre- and post-processing tools. Further information about the solver and its implementation can be found in [4].

Simulating broadband sound propagation through large distances with high resolution requires significant computational power and memory resources. A modern commercial Graphical Processing Unit (GPU) with 24 GB of memory and an upper limiting frequency of 2 kHz limits a completely heterogeneous domain to approximately 6 000 m^3. As our goal is to simulate the wavefront in a 3D domain over complex terrains over distances up to 500 m, we have adapted our solver to efficiently handle long-range propagation simulations through a turbulent atmosphere.

3. SOLVER EXTENSIONS

3.1 Dynamic Moving Frame

We extended *Empapu* with a dynamic moving frame tracking the main wavefront in the direction of interest through index shifting. This approach enables us to simulate a smaller spatial domain without sacrificing accuracy, thus reducing memory consumption. The dynamic moving frame is implemented by storing domain data using sparse encoding. Terrain voxel indices are derived from a 1D or 2D surface representing a height profile for the full 3D simulated domain. Turbulence cannot be encoded sparsely as each voxel above the terrain has unique material properties. To address this, we discretize the turbulent sound speeds into 20 bins. Within the moving frame subdomain, we store a field of index values, where each value corresponds to a specific bin. As the frame shifts, new speeds of sound are computed based on a predefined distribution (see Sec. 3.2).

The moving computational domain shifts along the positive x-direction whenever the wavefront (moving at the speed c_{max}) propagates further than a grid step dx – regular interval shifts lead to accumulated errors over long distances. The moving frame tracks the main wavefront and discards data that falls outside the frame. This is only valid when back reflections are minimal and do not significantly affect the simulation results. When the frame moves, all fields within the frame are shifted accordingly, and the relevant data from the extended simulation domain are shifted in.

Special care is taken to model the boundary conditions. The X_- boundary layer domain (see Fig. 1) is updated irregularly. Depending on the set boundary condition, we either fix / compute the boundary layer nodes, or replace them with adjacent data. To minimize error back propagation due to this irregular behavior, we add a PML in the X_- direction, ensuring that the pressure within this region gets attenuated. We take additional care to make sure that the source pulse is sufficiently attenuated before entering the boundary / PML domain.

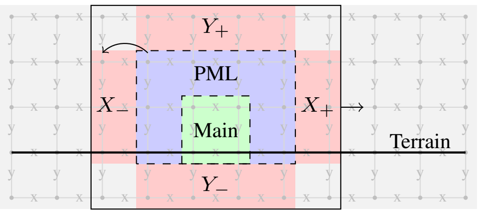

Figure 1: 2D representation of dynamic moving frame over staggered FDTD grid representing the different p, v_x, and v_y nodes. Inner domains represent relevant main, PML, and boundary layer subdomains. Arrows indicate direction of moving frame and index shifting.

3.2 Turbulence Modeling: Eddy (Quasiwavelet) Method

We model the effects of turbulence, wind speed and temperature profiles on sound propagation as an effective sound speed. We employ the eddy (quasiwavelet) method [5] superposed with a vertical sound speed profile [7]. The former approach reproduces the Kolmogorov spectrum, capturing features of turbulent atmospheres such as energy cascades and intermittency. The latter captures variations with height above the ground due to factors such as the surface roughness and atmospheric stability. Note that this method serves as an initial model, and a more realistic representation would require a vectorial field.

The quasiwavelet method synthesizes wavelet-like basis functions that resemble eddies in a turbulent flow. We compute a 3D scalar field of speeds of sound $c(\bar{r})$ by superposing Gaussian basis functions Ψ (the quasi-wavelets) with random amplitudes $b_{\alpha n}$ (with zero mean), uniformly distributed positions $\bar{r}^{\alpha n}$, and fixed widths a_α:

$$c_{\text{turb}}(\bar{r}) = \sum_{\alpha=1}^{N} \sum_{n=1}^{N_\alpha} b_{\alpha n} \Psi \left(\frac{|\bar{r} - \bar{r}_{\alpha n}|}{a_\alpha} \right) \tag{3}$$

where α represents different turbulence scales, and n denotes the index of basis functions for each size class. We apply this generalized approach to model temperature-induced turbulence which we thentransform to effective sound speed [5]. We chose this spatial domain technique because it locally and compactly represents the prescribed power spectral density and requires few degrees of freedom across the entire simulation domain and can be efficiently implemented for real-time computations.

We incorporate the terrain-height dependent sound speed profile according to [7]. We compute the linear-logarithmic dependency of sound speed c on height z' above the ground as

$$c(\bar{r}) = c_0 + c_{\text{lin}} z' + c_{\text{log}} \ln \left(1 + \frac{z'}{z_0} \right) + c_{\text{turb}}(\bar{r}) \tag{4}$$

where c_0 is the mean speed of sound, z_0 is the aerodynamic roughness length, representing the ground's surface roughness, and the $c_{\text{lin / log}}$ parameters are deduced from measurements.

Our GPU-accelerated solver efficiently computes the binned speeds of sound within the moving computational domain. As the dynamic window shifts, we update the sound speed field by generating new turbulence data on-the-fly with the aforementioned parameters that is spatially and temporally consistent with the entire simulation domain.

4. SIMULATIONS

Simulations with varying turbulent configurations were performed for a total propagation distance of 500 m (simulation time 1.5 s), employing a dynamic moving frame of size $6 \times 6 \times 10\,\text{m}^3$ with a 1.5 cm resolution ($\sim$ 150 MVoxels). We employ a broadband Gaussian point source excitation over a rough terrain representing natural grassland (height variation 16 cm; flow resistivity $\sigma_0 = 200\,\text{kPa s/m}^2$; $z_0 = 0.03\,\text{m}$). 64 PML layers are used to minimize back reflections. 10 turbulent configurations are generated according to the nominal parameters found in [5] and [6] with an outer length scale of 10 m and a total of 5080 quasiwavelets over the full $550 \times 6 \times 10\,\text{m}^3$ extended simulation domain with $c_{\text{lin}} = 0.0015$ and $c_{\text{log}} = 0.15$ (see Fig. 2). Effective sound speed varied between 342 and 345 m/s. We report on differences between a non-turbulent atmosphere and 10 different turbulent configurations with the same parameters.

5. RESULTS

Initial validation was done by recording the third octave band spectra up to 2 kHz of a point source at 0.5 m height on a receiver at a distance of 500 m in the presence of soft ground (flow resistivity $\sigma_0 = 200\,\text{kPa s/m}^2$) and comparing it to the analytical solution of the ground-effect

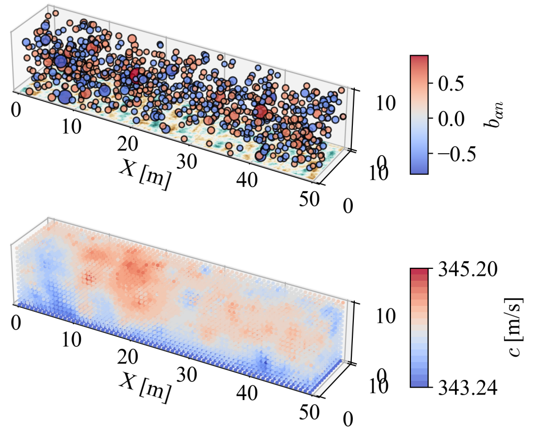

Figure 2: (Top) Quasiwavelets in space with radius a_α and color coded with amplitudes $b_{\alpha n}$ overlaid with rough terrain and (Bottom) corresponding synthesized discretized speed of sound distribution used to model turbulent field for the first 50 m.

[6] using the locally-reacting impedance model from [4]. Results were within 1 dB for a moving window of size $4 \times 2 \times 2\,\mathrm{m}^3$ at 1.5 cm resolution.

Simulations as described in Sec. 4 of 10 different configurations ran on an NVIDIA RTX 3090 Ti GPU. Each took 55 mins. Simulation times differed by $<1\%$ with respect to the homogeneous case, indicating that on-the-fly turbulent generation does not adversely affect simulation times. The impact of turbulence is illustrated in Fig. 3, where we see the combined impact of the rough terrain and turbulence. The impact of the terrain geometry dominates over the turbulence. Point recordings 500 m from the source differed from the homogeneous configuration by 0.5 dB in sound pressure level. Different turbulent configurations differed on average by 0.2 dB.

6. CONCLUSION AND FUTURE WORK

We have presented a novel approach for simulating sound attenuation in turbulent atmospheres by extending *Empapu*'s solver with a dynamic moving frame and an eddy (quasiwavelet) method for turbulence modeling. Our method demonstrates promising results in handling large-scale, high-resolution simulations and provides a valuable benchmark for assessing the validity and applicability of engineering models. Future work will focus on comparing simulation results with experimental data and engineering models, such as CNOSSOS-EU, ISO 9613-2, and sonX.

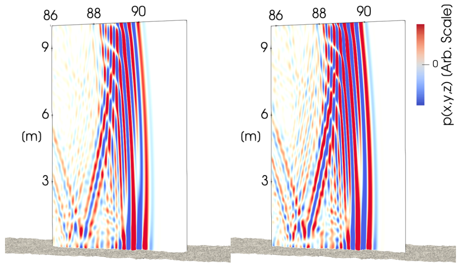

Figure 3: Pressure field distribution in compressed color scale overlaid with rough terrain for turbulent (left) and non-turbulent (right) atmosphere.

7. ACKNOWLEDGMENTS

This work was performed in the ATLAS project which is funded by the Swiss Federal Office for the Environment.

8. REFERENCES

[1] S. Kephalopoulos, M. Paviotti, and F. Anfosso-Lédée, "Common noise assessment methods in europe (CNOSSOS-EU)," 2012.

[2] ISO 9613-2:1996, "Acoustics – Attenuation of sound during propagation outdoors – Part 2: General method of calculation," 1996.

[3] J. Wunderli, R. Pieren, and K. Heutschi, "The Swiss shooting sound calculation model sonARMS," *Noise Control Engineering Journal*, vol. 60, no. 3, pp. 224–235, 2012.

[4] H. Montanaro, J. Knöpfel, K. Heutschi, and R. Pieren, "Parallelized FDTD simulations to improve engineering models of urban road traffic noise," in *Proc. 24th International Congress on Acoustics*, (Gyeongju, Korea), pp. ABS–0712 (8 pp.), 2022.

[5] V. E. Ostashev and D. K. Wilson, *Acoustics in moving inhomogeneous media*. CRC Press, 2015.

[6] K. Attenborough and T. Van Renterghem, *Predicting outdoor sound*. CRC Press, 2021.

[7] D. Heimann, M. Bakermans, J. Defrance, and D. Kühner, "Vertical sound speed profiles determined from meteorological measurements near the ground," *Acta Acustica united with Acustica*, vol. 93, no. 2, pp. 228–240, 2007.