Simulation of Sound Wave Focusing by Variable Density Medium

S.-R. Koh, I. Zhukov, J.-H. Kim, W. Frings

published in

NIC Symposium 2022
M. Müller, Ch. Peter, A. Trautmann (Editors)

Forschungszentrum Jülich GmbH,
John von Neumann Institute for Computing (NIC),
Schriften des Forschungszentrums Jülich, NIC Series, Vol. 51,
ISBN 978-3-95806-646-5, pp. 385.
http://hdl.handle.net/2128/31840

Simulation of Sound Wave Focusing by Variable Density Medium

Seong-Ryong Koh[1], Ilya Zhukov[1], Jong-Hyun Kim[2], and Wolfgang Frings[1]

[1] Jülich Supercomputing Centre, Forschungszentrum Jülich, 52425 Jülich, Germany
E-mail: {s.koh, i.zhukov, w.frings}@fz-juelich.de

[2] Korea University Guro Medical Center, Gurodong-ro 148, 08308 Seoul, Korea

High-performance computing with GPU acceleration is performed to determine the acoustic field controlled by a diffraction grating approach. The numerical analysis considers synthesised structures which manipulate the wave propagation and the acoustic intensity in the targeted region. A simple convex lens showed that the wave beam is localised and focused down to a small spot. The acoustic intensity in the area is increased by 15dB compared to the result by a 5-slits configuration. The programming language exploits OpenACC directives to achieve a large amount of performance gain by using GPU on the JUWELS Booster. The computational speed of production run is doubled with GPU acceleration compared to the performance of a pure message-passing interface parallelisation.

1 Introduction

Flow-induced sound occurs in many engineering applications, yet it is one of the important energy and environment problems. Powerful high-frequency waves beyond the limit of human audibility are used in many ultrasonic imaging devices for medicine and industry. In the last decade, acoustofluidics, i.e., the fusion of micro- to nano-scale acoustics and fluid mechanics, is one of the fast-developing research topics in science and technology. The acoustic energy in this field is of ultrasonic nature such that the waves are controlled to interact with bio-tissues and micro-particles in fluidic structures. The challenges of acoustofluidics are closely related to physics, engineering, and biomedical applications. The precise wave control is essential to optimise the fluidic-acoustic interaction in real devices. Therefore, new concepts are essential to progress the noise-control techniques and to manufacture the industrial products in a form of new artificial materials.

The design of new acoustic materials requires abilities to manipulate and control sound waves such that the physical behaviours have been impossible with conventional materials. A review of metamaterials researches[1] introduced the designs and properties of acoustic materials with special parameters e.g., negative refractive index, which decreases the wave phase in the medium. The advances of material technologies require the identification of technical challenges related to the future engineering designs such as wave-control devices. Furthermore, the development of new wave-control structures influences expanding researches on new materials that control vibrations, waves and the motion of solid materials with a highly efficient noise reduction.

In the present study the physical problem is configured to develop wave control devices by using variable density media. The diffraction grating is designed to obtain wave focusing and acoustic energy transformation. As improving the mechanical performance the acoustic meta-material could bring a silent driving to vehicles as conventional sound

insulation materials have been suffering from heavy weight and cost constraints. The combination of transformation acoustics theory and highly anisotropic medium enables precise control over the deformation of sound fields, e.g., to hide or cloak objects from incident acoustic energy.

2 Motivation

Towards the further step to develop real-world applications the physical problem tackles the acoustic waves in a channel flow, e.g., frequently visited in the biomedical fluids[2], and employs synthesised structures to control the sound power in the targeted region.

For large-scale engineering applications, the local-reacting properties of fluid-contacting structures, e.g., metal or porous medium, allow their acoustic behaviour to be characterised by a single complex quantity called acoustic impedance. This quantity is greatly affected by high sound pressure levels and grazing flows. That is, the pressure fluctuations accompany the turbulent flow forces exerted on the surface such that the sound generation occurs as the radiation condition is satisfied at a given surface impedance. The numerical analysis of full scale acoustic fields, which frequently concern high Reynolds number flows, increases computational cost due to the large disparity between the flow and the acoustic scale.

In the acoustofluidics the underpinning physics involves acoustics and fluid mechanics in micro-structures. The flow Reynolds number is usually very small, i.e., the viscous stresses easily dominate the flow motion. Furthermore, the applications are combinations of physics, electronics, medicine, and engineering. The numerical simulations are quite challenging to overcome the coupling problems which drastically increase the computational cost and the complexity of scientific interpretation. The computational analysis needs to resolve the temporal and the spatial scales which range from a micro unit to several meters in conventional mechanical designs. From the computational aspects it should be emphasised that the high-resolution numerical schemes and the high-efficiency high-performance computing (HPC) algorithms are essential to develop the noble wave control devices.

Therefore, this project aims for consolidation of the high-resolution computational fluid dynamics (CFD) workflow by exploiting graphics processing units (GPU) on a HPC system. The simulation of physical problems focuses on materials designed to control the propagation of acoustic waves generated in fluids. The GPU programming will include performance assessment and optimisation using OpenACC and CUDA-aware MPI libraries.

3 Sound Wave Control

A sound wave of wavelength λ is diffracted when the wavefront encounters an object or aperture. The sound intensity behind such obstructions is distributed by the shape of the aperture that the wave passes through. The diffraction characteristics are determined by interference between different portions of the wavefront. The resulting intensity distribution is called a diffraction pattern. When sound penetrates multiple apertures or slits with fixed spacing on a thin screen, the emerging wavefronts constructively interfere to produce a diffraction pattern with intensities peaked in certain directions as shown in Fig. 1. These

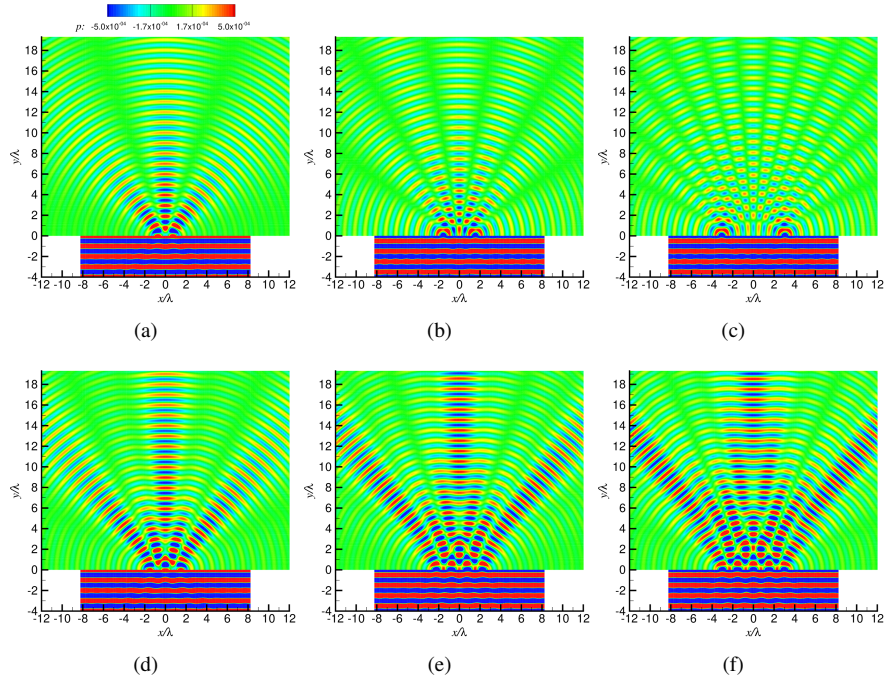

Figure 1. Instantaneous contours of acoustic pressure generated by various slits positioned at the coordinates $y = 0$, (a) 2 slits with $d_G = 1.5\lambda$, (b) 2 slits with $d_G = 3\lambda$, (c) 2 slits with $d_G = 6\lambda$, (d) 3 slits, (e) 4 slits, (f) 5 slits with $d_G = 1.5\lambda$.

directions are strongly dependent on both the slit spacing and wavelength of the incident wave. The direction of wave is determined by the basic grating equation for a wave with an incident angle α, i.e., $m\lambda = d_G(\sin\alpha + \sin\beta_m)$ where d_G is the spacing between slits, m is the order of diffraction, and β is the diffraction angle. Therefore, surfaces with well-defined slit locations are able to direct the acoustics of certain wavelengths into specific directions.

In the 1930s the variable density was reported as a new method of producing acoustic spectra[3]. The method already employed a variable-density sound-film as a diffraction grating for optical waves and created an instantaneous separation of sound into its components. Starting from a pure tone the diffraction angle of which increases with frequency, the diffraction pattern of numerous frequencies was considered in a diffraction grating which could record the range of frequencies over seven octaves. In recent studies an analytical model for a surface wave diffraction[4] was formulated to predict the diffraction pattern generated by the phased diffraction grating. The direction of waves is guided by a unit element of the diffraction grating. The diffraction waves are generated by the entire grating composed of a linear array of unit elements. The width, thickness, and grating length of the unit grating are w, δ, and d. The incident plane wave with an angular frequency $\omega = 2\pi f$ and amplitude A_i is defined by $A_i e^{i\omega t}$. Assuming the transmission coefficient is C, transmission waves at the exit of the unit element can then be expressed as $A_i C e^{i(\omega t - kd + \Delta\phi)}$ with

a phase difference $-kd + \Delta\phi$ by adding superstrate. The additional superstrates manipulate the dispersion relation of the travelling waves on the surface. The guided waves are diffracted at the exit of a unit element. A broadband metagrating design[5] detailed the density of periodical micro-structures which drives the wave diffraction and the Bessel beam implementation. An iterative algorithm was used to control the wave form by recording the phase profiles. The theoretical phase distribution and the simulation results show a good agreement in development of focused beams.

4 Numerical Method

The flow and the acoustic field of variable-density diffraction grating are analysed by solving the conservation equations for compressible fluid flows. The equations describing the sound propagation are the acoustic perturbation equations (APE)[6]. Since a compressible flow problem is tackled in this study the APE-4 system is used, i.e., the formulation of which is derived from the viscous conservation equations. The non-linear terms containing entropy fluctuations occur as additional source terms in the right-hand side of the governing equations. The contribution of viscous dissipation is negligible in the momentum equation of the APE system. To accurately resolve the acoustic wave propagation, a sixth-order finite difference scheme with the summation by parts property[7,8] is used for the spatial discretisation and an alternating 5-6-stage low-dispersion and low-dissipation Runge-Kutta method for the temporal integration[9]. The high-order discretisation with the summation by parts property (SBP)[8] improves the numerical stability of continuous problems via SBP operators. The dispersion relation preserving scheme (DRP)[7] is reformulated with the SBP properties. The sixth order discretisation in space uses the 9-points numerical stencil with a high-order 11-points spatial filter to prevent the spurious wave generation[10].

The flow and the acoustic simulation easily demand numerical meshes of $\mathcal{O}(10^9)$ degree of freedom which is only manageable by HPC systems. Due to the high-order discretisation schemes a typical message-passing-interface (MPI) approach suffers from the serious increase of computational overhead in mesh decomposition. To obtain the large gain of performance acceleration with a minimum number of MPI domains, a GPU programming with OpenACC is adopted to enhance the parallel efficiencies and to accelerate the computational performance in the simulation of the flow and the acoustic field of an acoustofluidic problem.

5 GPGPU Programming

GPUs are not new in the HPC ecosystems and most HPC centres in the world are equipped with GPUs. Usage of GPUs allows researchers and data scientists to improve their computations in orders of magnitude. An initial version of the CFD/CAA solver supports MPI and OpenMP on HPC systems. The behaviour of code subroutines in runtime is evaluated on the JUWELS[11] with GPU accelerators, i.e., JUWELS Booster.

5.1 General Structure

The numerical solver has the typical structure of a scientific application, i.e. initialisation, computational loop and post-processing. A graphical representation of 100 iterations executed on 4 MPI ranks on JUWELS Booster is shown in Fig. 2. The long green area in

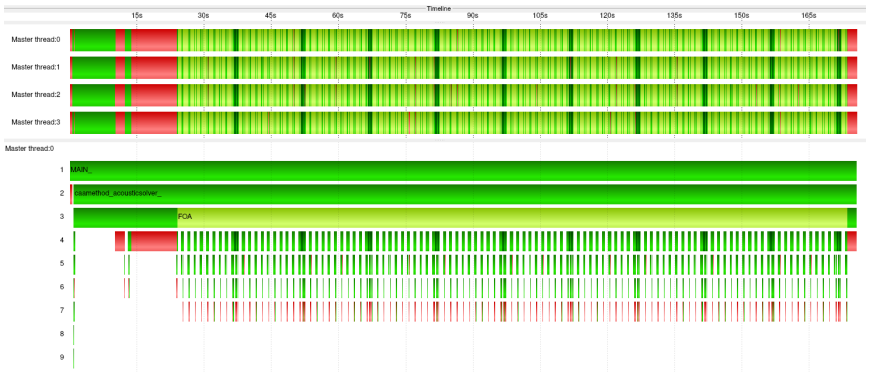

Figure 2. Execution timeline of the CFD/CAA solver, each row represents a single MPI process. A process timeline shows the different levels of function calls in a stacked bar chart at MPI rank 0 by MPI routines (red) and the user functions (green).

the top four bars is an actual compute loop (marked as a focus of analysis (FOA) in the stack diagram at the bottom) and is the most significant part of CFD/CAA solver. FOA consists of a large loop (100 iterations in the depicted testcase), where after each iteration occurs short non-blocking communication within `dom_dom_upgrategl_` routine, and every tenth iteration has a 3d filtering step (short dark green stripes). Short communication based on MPI non-blocking communication allowed to overlap communication with computation in a very efficient way. At the end of the computation loop all results were aggregated with MPI collective operations, i.e, `MPI_Reduce` and `MPI_Allreduce`. As computation is a lion's share of FOA, it is possible and reasonable to parallelise it, e.g., by distributing computations over available CPU cores or GPUs. The CFD/CAA solver can benefit already from OpenMP parallelisation which was implemented in the past.

5.2 Parallelisation Decisions and Strategies

There are several alternatives for GPU parallelisation, i.e., CUDA, OpenACC, HIP[a] and various libraries. Considering all these options we were tempted by OpenACC due to multiple reasons, e.g., OpenACC is relatively easy to implement as it uses compiler directives that are very similar to OpenMP; it gives better portability and readability of the code in comparison to other methods; OpenACC code can be compiled to CPU and GPU versions easily and can co-exist with OpenMP directives. Note that there are several drawbacks of OpenACC, i.e., so far only NVHPC compilers provide full OpenACC support, whereas newer GCC compilers provide only basic support; the low-level tuning for a specific computer architecture can be difficult. To port the CFD/CAA solver to the GPUs the following general rule was used, i.e. avoid costly memory transfers between CPUs and GPUs as much as possible. To accomplish this rule we reused data on GPUs as long as possible and used a special MPI library that allows exchanging memory data from one GPU to another without an intermediate copy to CPU memory (such libraries are called CUDA-aware MPI). CUDA-aware MPI routines can be included into OpenACC applications.

[a]HIP is a C++ Heterogeneous-Compute Interface for Portability:
`https://github.com/ROCm-Developer-Tools/HIP`

5.3 Current Implementation and Results

To integrate OpenACC in the CFD/CAA solver in desired fashion the following approach was used: allocate and initialise all necessary data before FOA and copy results back to the CPU at the end of the computation loop. All communication across GPUs performs via CUDA-aware MPI. A small part of the computation, e.g. 3d filtering step was not ported to GPUs due to the necessity of intensive communication across domains and frequent data transfer from host to device and backwards. This defect will be addressed in a follow-up project.

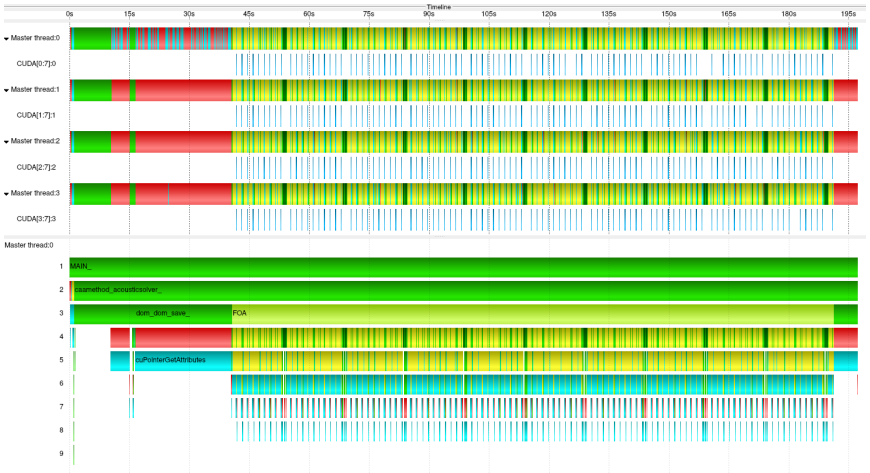

Figure 3. Execution timeline of the CFD/CAA solver (at the top): depicts a sequence of events in time, each row represents a single MPI process or CUDA stream. A process timeline (at the bottom) shows the different levels of function calls in a stacked bar chart for MPI rank 0. MPI routines are depicted in red, user functions are shown in all shades of green, OpenACC routines are yellow, CUDA kernels are dark blue, CUDA API routines are light blue.

A graphical representation of 100 iterations executed on 4 MPI ranks with A100 GPU each on JUWELS Booster is illustrated in Fig. 3. The total runtime with OpenACC implementation in Fig. 3 is longer than in the case in Fig. 2 due to the measurement overhead. The overhead is caused by instrumentation of the code by Score-P measurement infrastructure[b]. The non-instrumented code executions with a relatively heavy computational load on 2 nodes show that the runtime of FOA with a pure MPI approach on JUWELS Booster (96 MPI ranks) is 348 seconds and the one with the OpenACC also on JUWELS Booster (16 MPI ranks with 8 A100 GPUs) is 153 seconds, i.e., 2.3 times faster with GPU acceleration. With production runs where GPUs' load is more realistic, one can get even better speedup. The result shows a good agreement with the measurement in the GPU Hackathon 2021[12]. Note that the number of MPI ranks per GPU is considered an important factor to tune the code performance on the JUWELS Booster.

[b]Score-P measurement infrastructure: `https://www.score-p.org/`

5.4 Future Work

Although we achieved quite good results during the project, there are some aspects we want to improve or investigate, i.e. port as much as possible computations from CPUs to GPUs, investigate possibilities of OpenMP offloading, investigate GPU kernels for possible low-level optimisations.

6 Hardware and Software Configuration

All measurements were performed on the JUWELS Booster[c] which is a part of the high-performance cluster JUWELS. The JUWELS Booster consists of 936 compute nodes equipped with four NVIDIA A100 GPUs for each node. The GPUs are hosted by AMD EPYC Rome CPUs. The compute nodes are connected with HDR-200 InfiniBand in a DragonFly+ topology. The following software was used for the tests on JUWELS Booster: NVHPC/22.3, ParaStationMPI/5.5.0-1, CUDA/11.5 and HDF5/1.12.1. Furthermore, Score-P/7.1 was used for measurements and Vampir/10.0.0[d] for timeline visualisations.

7 Simulation Results

7.1 Simulation Setup

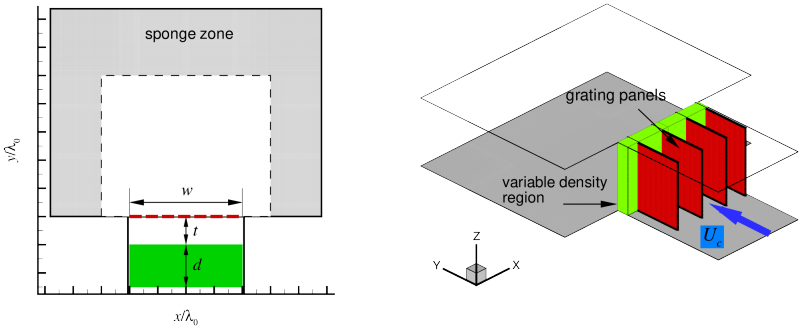

Figure 4. Schematic of the acoustic domain, w denotes the width of diffraction grating, t is the tolerance between the grating exit and the variable-density region, of which the thickness is d, the waves move in the positive y direction and the convection speed of a channel flow is U_c.

The computational study is performed to obtain the flow and the acoustic field of diffraction gratings with variable density distribution. The test cases included slits defined by three different widths d_G. The number of slits varies up to five with which the

[c] JUWELS Booster: `https://apps.fz-juelich.de/jsc/hps/juwels/booster-overview.html`

[d] Vampir trace visualizer: `https://vampir.eu/`

diffraction pattern is obtained in Fig. 1. Fig. 4 shows the schematic of an acoustic domain illustrated with a set of parameters. The grating panels are defined with the thickness d which varies from λ_0 to $4\lambda_0$ where λ_0 is the reference length of incident waves. The diffraction grating consists of eight elements distributed in the x direction. Each element is defined with a variable-density function which enables phase-shifts of propagating waves. The phase-shift difference is related with the distribution function and the distance d. The mixed forms of a quadratic and a linear function are used to define the distribution of variable density inside the grating elements.

7.2 Physical Solution

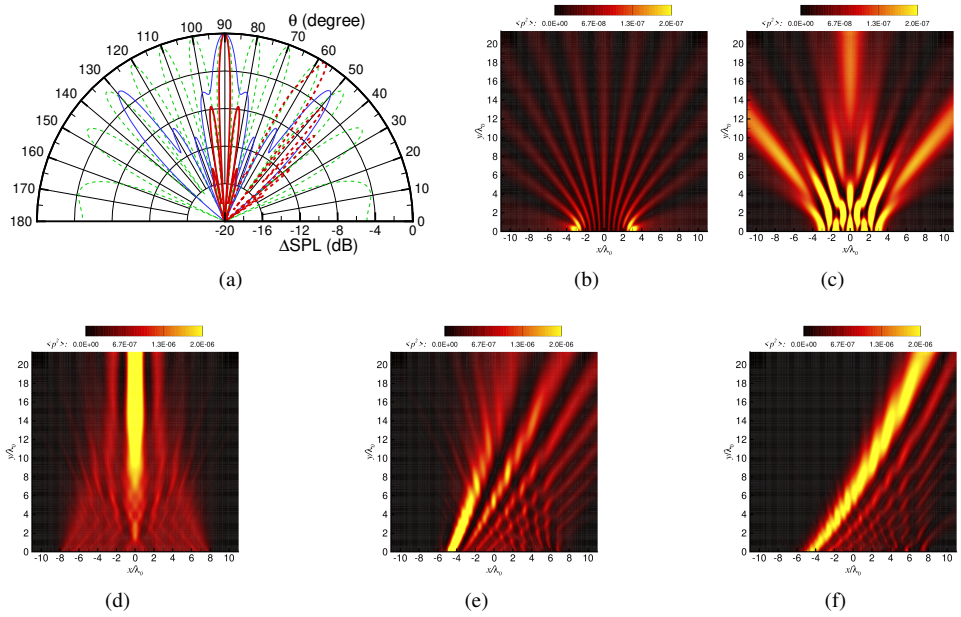

Figure 5. Contours of acoustic intensity obtained by guided waves, (a) directivity with the azimuthal angle θ, the intensity contours determined by (b) 2 slits with $d_G = 6\lambda$, (c) 5 slits with $d_G = 1.5\lambda$, (d) a density lens, (e) grating with a quadratic distribution of density, and (f) grating with a linear distribution of density.

In Fig. 5 the distributions of acoustic intensity are determined in the downstream of the diffraction gratings. The directivity in Fig. 5(a) is calculated with the pressure signals obtained by the virtual microphones at 20λ away from the grating exits. The origin is locates at the xyz-coordinates (0,0,0) for (b), (c), and (d) and (-5,0,0) for (e) and (f). The intensity maximum of each case is used for a normalised sound pressure level ΔSPL, i.e., 18 dB at the propagation angle $\theta = 90°$ for green dashes by (b), 26 dB for blue solid lines by (c), 41 dB for red solid lines by (d), and 38 dB at $\theta = 58°$ for red dashes by (f).

Fig. 6 illustrates the guided circular waves transformed to a focused beam (waves coloured by red and blue) using phased diffraction grating. Colours indicate contours of

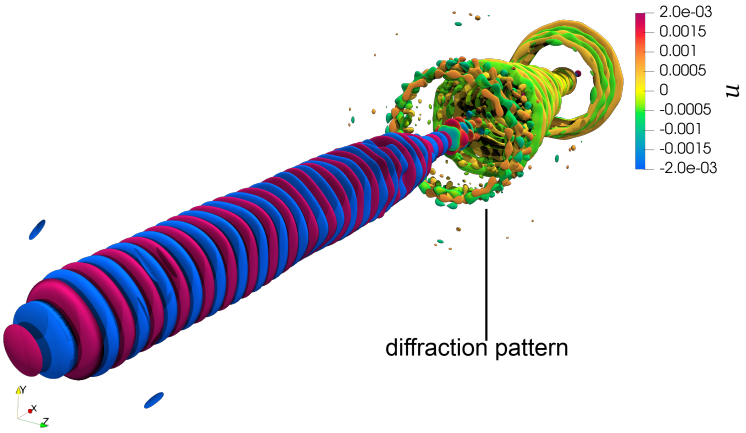

Figure 6. Contours of the acoustic pressure determined by phase diffraction grating. Colours indicate the perturbed velocity component in the x direction.

perturbed velocity component u. The full three-dimensional computational domain consists of ca. 1.77 billion meshes with the minimum cell spacing of $0.02\lambda_0$.

8 Summary and Outlook

A computational aeroacoustics approach is used to analyse the acoustic field generated by a diffraction grating. Simulations showed that the synthesised structures are able to control the wave propagation and the acoustic intensity in the targeted region. A numerical solver exploits the OpenACC directives to achieve a significant performance gain by using GPGPU on the JUWELS Booster. A performance analysis was detailed by using Score-P and the visualisation via Vampir. Production runs on the JUWELS Booster show a 2.3-fold runtime improvement in FOA, which is consistent with results from the Helmholtz GPU Hackathon[12].

The outlook addresses the relationship of the material parameters with the acoustic variables which is essential to define an optimal structure. The correlation of the flow variables with the characteristic length scales of the optimum design will be deduced by quantifying the structure effects on the flow field and the acoustic impedance. This prospective study includes a deep learning technique that is available via the open-source reinforcement learning library *smarties*[13]. The library performance has been proved in studying various unsteady dynamics problems. In simulation of acoustic fields the agents are distributed evenly along a spherical surface perpendicular to the three-dimensional incident waves, with each agent obtaining state information (local acceleration of fluid particles) at a polar angle ϕ^a, computing the reward on the surface and feeding into the policy function to obtain actions in the next step. The details will be presented with providing an optimised phased diffraction grating for targeted acoustics.

Acknowledgements

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (`www.gauss-centre.eu`) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS[11] at Jülich Supercomputing Centre (JSC).

References

1. S. Cummer, J. Christensen, and A. Alù, *Controlling Sound with Acoustic Metamaterials*, Nature Reviews Materials, **1**, 16001, February 2016.
2. S. R. Koh, J.-H. Kim, and A. Lintermann, *Numerical Analysis of Oscillatory Flows in the Human Brain by a Lattice-Boltzmann Method*, in: Proceedings of 14th WCCM-ECCOMAS Congress 2020, 2021.
3. D. Brown, *Sound-Films as Diffraction Gratings for the Visual Fourier Analysis of Sound-Waves*, Nature, **140**, 1099, 1937.
4. Z. Tian and L. Yu, *Elastic Phased Diffraction Gratings for Manipulation of Ultrasonic Guided Waves in Solids*, Physical Review Applied, **11**, 024052, 2019.
5. Y. Wang, Y. Cheng, and X. Liu, *Modulation of Acoustic Waves by a Broadband Metagrating*, Science Reports, **9**, 7271, 2019.
6. R. Ewert and W. Schröder, *Acoustic Perturbation Equations Based on Flow Decomposition via Source Filtering*, Journal of Computational Physics, **188**, 365–398, 2003.
7. C. K. W. Tam and J. C. Webb, *Dispersion-Relation-Preserving Finite Difference Schemes for Computational Acoustics*, Journal of Computational Physics, **107**, no. 2, 262–281, 1993.
8. S. Johansson, *High Order Finite Difference Operators with the Summation by Parts Property Based on DRP Schemes*, Tech. Rep. 2004-035, Uppsala University, 2004.
9. F. Q. Hu, M. Y. Hussaini, and J. L. Manthey, *Low-Dissipation and Low-Dispersion Runge-Kutta Schemes for Computational Acoustics*, Journal of Computational Physics, **124**, 177–191, 1996.
10. C. Bogey and C. Bailly, *A Family of Low Dispersive and Low Dissipative Explicit Schemes for Flow and Noise Computations*, Journal of Computational Physics, **194**, no. 1, 194–214, 2004.
11. Jülich Supercomputing Centre, *JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Juelich Supercomputing Centre*, Journal of large-scale research facilities, **7**, no. A138, 2021.
12. OpenACC Organization, "Helmholtz GPU Hackathon 2021", 2021, March 15–24 (Digital Event), Jülich Supercomputing Centre.
13. G. Novati and P. Koumoutsakos, *Remember and Forget for Experience Replay*, in: Proceedings of the 36th International Conference on Machine Learning, 4851–4860, 2019.