Solving Complex Acoustic Problems Using High-Performance Computations

Gregory Bunting, Clark R. Dohrmann, Scott T. Miller, and Timothy F. Walsh

Introduction

Sound waves propagating in fluids (air, water, etc.) are a ubiquitous part of our everyday lives, from communication through speech to learning in classrooms to communicating underwater. The propagation of acoustic waves in these environments is well-understood and documented in the comprehensive history given in the book by Allan Pierce (2019). For example, one may wish to know the acoustic pressure field in a large expanse of water under the surface of the ocean (Duda et al., 2019), the sound field at every location and time in a large concert hall (Hochgraf, 2019), or the structural response of aerospace structures to high-intensity acoustic fields that are experienced in-flight (e.g., the Orion capsule in **Figure 1**, *left*). Unfortunately, when the geometry, boundary conditions, and/or given spatial distributions of material properties of the fluid are complex, the governing wave equations do not typically lend themselves to an analytic solution. The same holds true for wave equations in other areas of physics such as electromagnetism and optics. In these scenarios, numerical solution of the wave equations can be a powerful tool for computing the acoustic quantities of interest because otherwise there is no other means of obtaining this information.

Figure 1. *The Orion (**left**) is the new NASA spacecraft for astronauts to revisit the moon by 2024. Ground-based testing of the capsule can be modeled via the finite-element method (FEM). A FEM discretization of the acoustic domain surrounding the Orion (**right**) illustrates the domain discretization method. See text for discussion.*

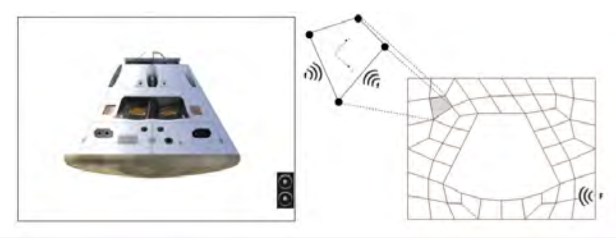

Computational acoustics (CA) has emerged as a subdiscipline of acoustics, concerned with combining mathematical modeling and numerical solution algorithms to approximate acoustic fields with computer-based models and simulation. Using CA, acoustic propagation is mathematically modeled via the wave equation, a continuous partial differential equation that admits wave solutions. The numerical methods of CA are focused on taking the continuous equations from calculus and turning them into discrete linear algebraic calculations, which are amenable to solution on digital computers. In the case of a concert hall or underwater domain with complex geometries that are not amenable to an analytic solution, CA would enable an acoustics engineer to compute a numerical solution to the wave equation to help the engineering design process. Some of the more popular of these methods are finite difference, finite volume, spectral element, boundary element, and finite-element methods (FEMs). Although each of the numerical strategies for solution of the acoustics equations has its own niche applications and advantages/disadvantages, in this article, we focus on the FEM and its application on modern high-performance computing platforms.

Example: Solving the Helmholtz Equation

As an illustration, one can consider the continuous and discrete versions of the acoustic Helmholtz equation for steady-state wave propagation in fluids. In the continuous form, when body loads are neglected, one has

$$\Delta p + k^2 p = 0 \tag{1}$$

where $p = p(x, y, z)$ is the acoustic steady-state pressure as a function of position and $k = \omega/c$ is the wave number. Applying one's favorite numerical method to solve Helmholtz's

©2020 Acoustical Society of America. All rights reserved.
https://doi.org/10.1121/AT.2020.16.2.22

equation (along with the associated boundary conditions) numerically yields a discrete set of n linear equations

$$A\,p = F \tag{2}$$

where A is an $n \times n$ matrix that contains a discrete representation of the continuous Helmholtz equation, p is a vector of unknown discrete acoustic pressures at the nodes of the discretized model, and F is an $n \times 1$ vector containing information about boundary conditions and energy/load sources for the acoustics problem at hand. For completeness, we note that in the case of boundary element methods, one would start with the *Helmholtz integral equation* instead of the differential form given in Eq. 1. The details are omitted here for brevity.

High-Performance Computing for Acoustics

In high-performance computing (HPC) for acoustics, a challenge is to solve equations in the form of Eq. 2 when the number of unknowns (n) becomes too large for a single computer. Modern HPC platforms and the corresponding software for domain decomposition and parallel communication are revolutionizing the numerical solution of acoustic wave equations. This is enabling the solution of practical engineering problems in acoustics such as airborne acoustic propagation (Hart et al., 2016), sonar applications, and aeroacoustic noise mitigation that were not possible using previous generations of computers. The enabling HPC technology allows one to resolve acoustic wave propagation in ever-increasing domain sizes (or, equivalently, ever-increasing frequency ranges, e.g., megahertz) of interest for the wave propagation. The number of discrete equations to be solved increases with the frequency range and domain size. Eventually, the growing number of degrees of freedom and corresponding matrix storage requirements preclude the solution of the problem on one's laptop or desktop as memory resources in the computer are exceeded.

Modern HPC platforms, based on either distributed central processing units (CPUs) and/or graphics processing units (GPUs), are built to optimize the use of memory resources on the largest problems in computational physics. In the case of acoustics, as the frequency range and/or domain size increases and the required in-core memory [aka random-access memory (RAM) for storing bits of information] resources correspondingly increase, an acoustics researcher can, in principle, simply employ larger numbers of CPUs and/or GPUs on computing clusters to enable the numerical solution. Modern HPC clusters deployed by the US Department of Defense (DOD) and Department of Energy (DOE) laboratories have access to tens of thousands of CPUs/GPUs, each with substantial in-core memory resources. The problem then becomes how to tailor the numerical method of interest so that it can be applied in these novel, distributed memory and architectural computing environments.

Because a wide range of acoustics applications encounter large-domain sizes and/or high frequencies of interest, the ability to numerically solve the acoustics equations in a scalable way is of broad interest across the field of acoustics engineering. Large-domain sizes present themselves in underwater acoustics (Duda et al., 2019), waves in atmospheric propagation scenarios (Hart et al., 2016), aeroacoustics for airborne structures, architectural acoustics in large concert halls, and large-scale acoustic chambers for testing aerospace structures (Schultz et al., 2015), to name a few. In these applications, the large size of area where the acoustic solution is desired translates to large matrices for the numerical solution and hence the need for HPC. Equivalently, applications with high frequencies present precisely the same computational challenges as large-domain sizes because in both cases the large number of wavelengths to be resolved requires more and more discrete elements and/or nodes to resolve the wave propagation. Ultrasound applications (Suslick, 2019) are an example where, due to the high frequencies and hence small wavelengths, numerical methods require large numbers of degrees of freedom for the solution. HPC has the potential to enable the solution of these and other acoustics problems across a variety of engineering disciplines.

In many acoustics applications, the FEM is an attractive numerical strategy. Some advantages of the method include

- The ability to construct unstructured, body-fitted meshes that capture curved interfaces between complex fluid/structural domains;
- Sparse systems (i.e., matrices wherein most entries are zeros) of algebraic equations that, when combined with a FEM of an elastic structure, render a coupled system of equations that is still sparse;
- The ability to solve either linear and/or nonlinear acoustic wave equations; and
- The ability to easily handle spatially varying material properties (e.g., capturing the speed of sound and

density that vary with vertical position in underwater or atmospheric acoustics).

By contrast, finite-difference approaches employ a structured grid that cannot easily capture curved interfaces. Boundary element methods present a dense linear system of equations, which makes coupling with finite-element-based structural models challenging because the latter present a sparse system of equations. In fairness, these alternative methods also have their own advantages over FEMs in certain applications. However, a common theme is the emergence of HPC resources and the benefits that are being presented to any numerical approach for solving acoustic problems.

Finite-Element Method

The FEM has been widely used as a tool for solving the acoustic wave equation. One of the earliest references is from Gladwell (1965), quickly followed by several follow-on efforts (Craggs, 1971). Additional references involving the coupling of an acoustic fluid with a structure followed in the late 1960s and early 1970s (Zienkiewicz and Newton, 1969; Craggs, 1972). More recent surveys on FEMs for acoustics and structural acoustics provide comprehensive technical reviews on the application of FEMs for solving acoustics problems (Atalla et al., 2017).

Finite-element technology solves partial differential equations (PDEs) by turning them into linear algebra. The FEM discretizes the physical domain of a problem into a finite number of elements. This discretization process is illustrated in **Figure 1**, *right,* for the Orion space capsule. In this case, the goal is to understand the structural response of the Orion capsule to high-intensity acoustic excitation as would be experienced in flight. The continuous physical domain (box) containing the Orion space capsule in **Figure 1**, *left*, is subdivided into a collection of elements in **Figure 1**, *right*. The solution is approximated by a polynomial with unknown coefficients, defined locally on each element, and is substituted into a suitable integral representation of the PDE. The result of this approximation methodology is a linear system of equations to be solved for the polynomial coefficients. Each of these unknowns is referred to as a degree of freedom.

Getting Around Moore's Law

Gordon Moore (1965) observed that the speed of a computer processor doubles about every two years. This became known as Moore's law and served as a target for computer chip manufacturers for several decades. During the reign of Moore's law from 1975 to 2012, larger computational problems could be solved by waiting for a faster processor to be produced. However, Moore's law could not reign forever because the physical limits of the microelectronics prevented such perpetual growth. Rather than wait for a faster processor, it became necessary to use many processors working together to solve larger problems; thus parallel computing is born.

The Advent of Parallel Computing

The early work in parallel computing for acoustics started in the 1990s and consisted of using many CPUs to solve a given problem. In the case of a finite-element solution of the wave equation, the approach was a divide-and-conquer strategy (aka domain decomposition), where the individual finite elements were evenly distributed across the CPUs and the solution of the global set of algebraic equations would be accomplished by many CPUs working in parallel. With the continued demise of Moore's law, manufacturers are now producing GPU-based computing platforms. A single GPU can have thousands of processor cores compared with tens of CPU cores per computational node. Seymour Cray, the father of supercomputing, once remarked, "If you were plowing a field, which would you rather use: two strong oxen or 1024 chickens?" (Cray, 2020). This antiquated quote reflects the opinion that it would be more advantageous to have one fast processor rather than restructuring work to be accomplished in a massively parallel fashion. Modern HPC is finding a way to harness the power of the chickens when the oxen are unavailable. Heterogeneous computing environments, where CPUs, GPUs, and possibly other processors coexist on a single piece of hardware are the future of scientific computing.

Key points for the parallelism of work are synchronization and independence. Tasks to be executed in parallel need to be independent from each other so that it does not matter which one gets completed first. Synchronization points in an algorithm provide a waiting point where all parallel tasks can meet up and exchange any information that might be needed for future work. Current models for heterogeneous computing rely on CPUs to organize and divide tasks to be parallel processed on the GPU. Hardware configurations dictate that multiple CPUs must be able to execute tasks simultaneously on a single GPU.

High-Performance Computing for Department of Energy/Department of Defense Applications

The emphasis on using HPC as a pillar of science and engineering for national security purposes can be traced back to 1992 when the United States passed a moratorium on nuclear testing. A consequence of this was the establishment of the Stockpile Stewardship Program (SSP) that was given the task of certifying the safety and reliability of the nuclear weapons stockpile without nuclear explosives testing. The Advanced Simulation and Computing Program (ASC) is a vital element of the SSP, creating the modeling and simulation capabilities necessary to combine theory and past experiments to create future engineering designs and assess aging components in the stockpile. Sierra Mechanics is one software in the ASC toolset and was developed at Sandia National Laboratories (see sandia.gov). Within the Sierra Mechanics suite, the Sierra-SD (structural dynamics) module includes capabilities for massively parallel acoustics and structural acoustics capabilities in both the time and frequency domains as well as eigenvalue capabilities for mode shape and frequency calculations (Bhardwaj et al., 2002; Bunting, 2019). In **Applications** we show examples of the use of these capabilities to solve acoustics problems on some of the world's largest supercomputers.

High-Performance Computing

Taking advantage of modern HPC platforms for acoustics requires an understanding of the architectures themselves to develop optimal software strategies to maximize performance. The evolution of computing platforms in the past couple of decades can be illustrated by comparing the top platforms in the late 1990s versus those of today. In 1997, the Advanced Strategic Computing Initiative (ASCI) Red machine of the DOE came online with over 9,000 processors and 1 terabyte of total memory, becoming the first supercomputer in the world to achieve the speed of 1 teraflop (i.e., 10^{12} floating point operations per second). We can compare that with the more recent DOE Summit (Oak Ridge National Laboratory, 2020) machine, which has over 200,000 CPU cores and 27,000 GPUs. Summit has a peak performance of 200 petaflops (or 200,000 teraflops), currently making it the fastest supercomputer in the world (Top500, 2019). Cloud computing, such as those offered by Google and Amazon, has a prohibitively slower communication time between CPU cores and thus are not optimal for scientific computing. Given this rapidly evolving hardware, navigating the last 20 years of supercomputing changes has led to some fundamental changes in the way the software is structured.

Parallel Scalability

An important concept in HPC is the notion of scalability, which encompasses both the *size* of the problem that can be solved and *how fast* the problem can be solved. The former is typically referred to as *weak scaling*, and the latter as *strong scaling*. In cases where the goal is to solve the problem very fast, the intent is to use HPC to achieve *strong scaling*. In other cases, the goal may be to solve very large problems (i.e., many degrees of freedom), in which case one wants to achieve *weak scaling* on the HPC platform.

Strong scaling is demonstrated when doubling the amount of processing power available for a given problem cuts the solution time in half. Weak scaling represents the ability to solve very large problems with many degrees of freedom. In the case of a finite-element solution, this would imply that the model has many nodes and/or elements, which is commonly the case in acoustics applications when the domain size and/or frequency range of interest becomes large. Weak scaling is demonstrated when one simultaneously increases *both* the problem size and the processing power available and is able to solve the same problem in the same amount of time. More generally, it can be stated as the ability to solve an n times larger problem using n times more compute processors in nearly constant CPU time.

Linear Solvers

The HPC hardware described in **High-Performance Computing** is only useful for computational acoustics if one also has the software to solve Eq. 2 for large dimension n. A simple example of a linear system is given by the two equations $2x + y = 4$ and $x + 3y = 7$, where x and y are the unknowns. In this case, the solution $x = 1$ and $y = 2$ can be obtained by hand. When the dimension n of matrix A from Eq. 2 is of moderate size (less than several million), sparse direct solvers (Davis, 2006) can be used effectively to solve for the unknowns. For larger problems in computational acoustics when n exceeds several million, however, the computational resources required by a direct solver quickly become prohibitive. In the case of the Orion capsule example described in **Applications** with over 2 billion unknowns, a direct solution would take order of months on a hypothetical CPU processor with a peak performance of 1 teraflop and enough memory for the computations.

The predominant methods for solving linear systems with the large dimension n beyond the reach of direct solvers are preconditioned iterative approaches (Smith et al., 1996; Dohrmann et al., 2010). These methods solve the linear system in an iterative fashion rather than using a direct factorization. A common preconditioned iterative approach is based on the concept of a divide-and-conquer strategy where the physical domain is divided into disjoint partitions and each partition is handled by a separate CPU (Smith et al., 1996; Dohrmann and Widlund, 2010). Other preconditioners based on multigrid are popular in other applications, wherein each partition is further subdivided into another level of partitions.

Applications

Here, we present examples of applications where HPC has been utilized to solve acoustics problems of the form in Eq. 2, with a large dimension n that would not be possible with smaller scale computing platforms. Solutions of problems with over 2.2 billion degrees of freedom are presented.

Underwater Acoustics for Ship Shock

One application of underwater acoustics is ship-shock testing. Vessels in the Navy fleet must undergo ship-shock tests before they are certified for service. These tests involve setting off large underwater explosives near the vessel of interest, typically around 75% of the expected failure load. The purpose is not to sink the vessel but to find out what breaks when the ship is exposed to nearby explosions (e.g., electronics, chairs, pipes). Such at-sea tests are extremely expensive and take a vessel out of the fleet for many months or even years. The more a ship is damaged in such a test, the longer it takes to return to the fleet. Computational modeling of a ship-shock event is one strategy to design components that will survive a ship-shock test. In the far field, the large pressures generated by explosives can be modeled as acoustic pressure waves impinging on the ship. These underwater acoustics applications exhibit large simulation domains and frequency ranges of interest that result in many wavelengths in the domain. As such, they lend themselves well to solution via HPC (Moyer et al., 2016).

Figure 2 shows slices of the acoustic pressure field reflected from a stiffened cylinder in a transient-coupled structural-acoustics simulation. Acoustic loading is due to an underwater explosion away from the submerged, air-filled structure. Rings on the surface of the cylinder indicate the mechanical response. The near-field fluid domain consists of an ellipsoidal region composed of tetrahedral elements with an acoustic material formulation. The far-field, semi-infinite domain is approximated by infinite elements, which are not shown. **Figure 2**, *inset*, shows a typical gauge time history predicted by the analysis.

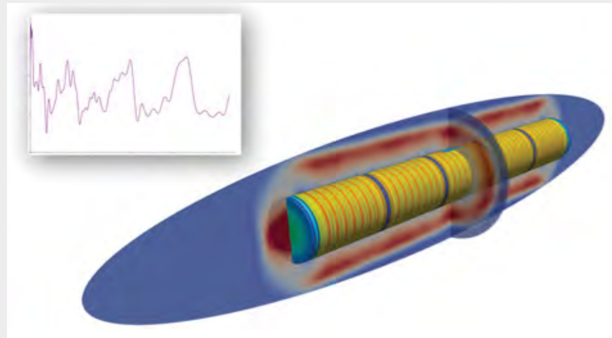

Figure 2. *The simulated response of a stiffened cylinder subject to underwater explosion is a demonstration of a typical Navy use of the FEM. The incident pressure wave excites the structure as it is reflected off the surface.* ***Inset:*** *a typical gauge time history. These results can be used to design structures to withstand explosive detonations in the surrounding medium (water).*

Simulation of Ground-Based Acoustic Tests

Qualification tests of aerospace structures and flight vehicles require that the structures be subjected to acoustic loads that are representative of the environments that will be experienced in-flight. One way to achieve this, of course, is to conduct a full-scale flight test on the structure. The associated accelerometer and/or pressure sensor measurements can then be used to assess the acoustic environment, and the resulting structural response.

However, flight tests tend to be very expensive, and due to instrumentation and telemetry limitations, only limited accelerometer data are typically available from such tests. As a result, ground-based acoustic testing is a common alternative wherein the structure is subjected to representative acoustic fields in an acoustic test chamber. Typically, high-powered speakers and other acoustic sources are used to generate the acoustic fields. The advantages of ground-based testing are that the cost is typically only a small fraction of that of a flight test and, perhaps more importantly, significantly more data can be gathered with ground-based acquisition systems.

The advantages of ground-based acoustic testing come with a challenge, however, in that one needs to engineer specific acoustic fields that emulate what would be seen in-flight. Given the large-domain sizes and high-frequency ranges of interest, these problems typically have many wavelengths in the domain, thus making HPC with finite elements an attractive solution strategy.

Ground-based testing goes hand in hand with computational acoustics modeling. **Figure 3A** shows a typical experimental setup of a ground-based acoustics test at Sandia National Labs, and **Figure 3B** shows a corresponding finite-element model of a representative test (Schultz et al., 2015). The ground-based test can tell us with high confidence the mechanical response from a certain loading environment at specific sensor locations. However, numerical simulations are used to test large numbers of different loading environments and provide the response at all points of a model, something impossible to do with experiments alone. These models typically reach sizes of hundreds of millions of finite elements.

Example: Orion Capsule in Ground-Based Acoustic Test

As an example of ground-based acoustic testing, we present a numerical model of a reverberation chamber test on the Orion capsule. Reverberation (reverb) chambers are rooms designed to produce a diffuse sound field around an object of interest, which is a common condition in flight environments. A diffuse sound field is an acoustic environment where the acoustic energy density is the same at all locations. By understanding structural response to a diffuse field, including absorption coefficients and transmission loss, the structural behavior during launch and reentry can be characterized.

The *sound absorbability* is determined by the change in reverb time of the test object. Acoustic excitation can be accomplished with a variety of different source arrangements. Reverb room tests are very important for ground testing flight objects that will be excited to uniformly high random pressure loads while in use.

To demonstrate a simulation of a reverb chamber test, we present a numerical simulation of a three-quarter scale version of the Orion capsule (crew module; National

Figure 3. ***A:*** *an engineer setting up a ground-based acoustic test of a weapon system. Instrumentation is being put into place in preparation for acoustic excitation to evaluate the structural response to high-amplitude acoustic fields.* ***B:*** *results of the computational acoustics simulation corresponding to the physical test in* ***A****. The simulation results provide pressure, acceleration, and stress values in the weapons system at every point in time in the simulation. These numerical results are used to evaluate the weapon response in the simulated acoustics environment. The finite-element mesh is represented by the grid, and the colors represent the magnitude of the acoustic pressure field at an instant in time.*

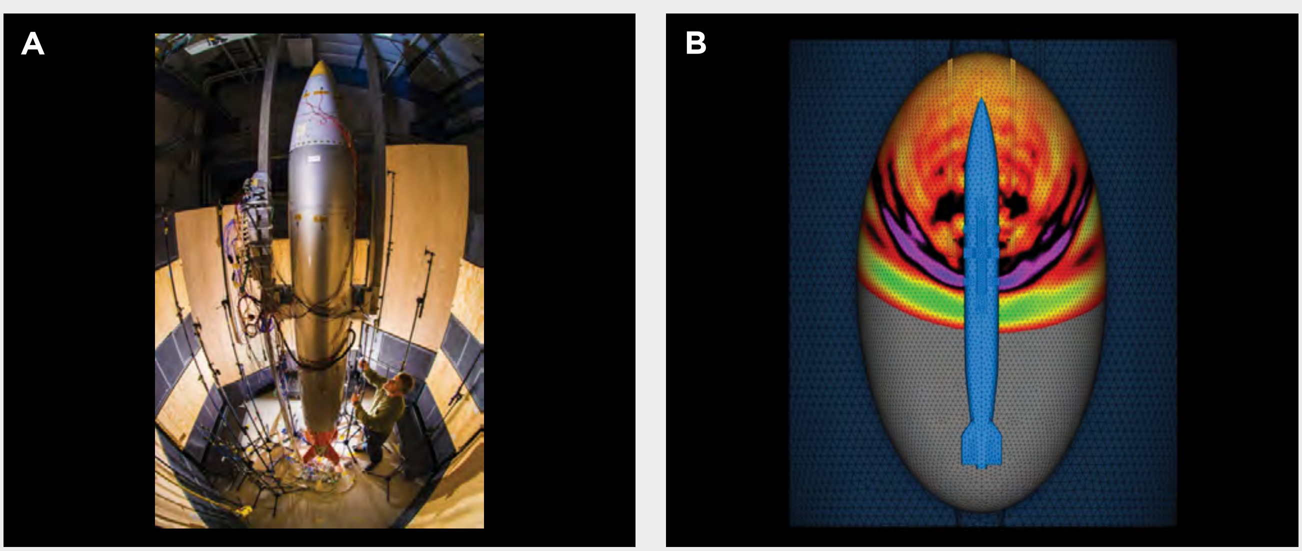

Aeronautics and Space Administration, 2019) in the middle of the vibroacoustic test facility (VATF) of Sandia National Labs (Schultz et al., 2015). We note that this is purely a numerical study, not an actual experimental test. The VATF is a rectangular box 6.58 m × 7.50 m × 9.17 m, making the volume ratio of capsule to room approximately 0.1. Acoustic excitation to the 140 dB level is provided by a 0.1-m^2 loudspeaker in the bottom corner of the room. It provides a sinusoidal acoustic velocity loading with an amplitude of 3.4 m/s and a frequency of 350 Hz.

The accuracy of a finite-element solution is dependent on the size of the elements used to obtain the solution. In acoustics, the element size used in the mesh will limit the frequencies resolved. For instance, computing a sound field using finite elements would require the mesh size

$$h = \frac{c_o}{\lambda f_{max}} \tag{3}$$

where h is the size of the finite element, c_0 is the speed of sound in air, f_{max} is the highest frequency requested, and λ is the number of elements per wavelength needed. Typically, for low-frequency excitation, we select λ = 10 for linear hexahedral elements. Fewer elements can be used in conjunction with high-order polynomial interpolation as the basis functions that are able to approximate a waveform with less error, but for simplicity, we do not cover those details here.

The acoustic domain for this problem is the air enclosed by the reverb chamber and surrounding the Orion capsule. Using the Sierra-SD software, a transient reverb simulation for over 2.2 billion unknowns was solved on the Serrano supercomputer using over 22,000 computing cores in under 8 hours! **Figure 4A** shows an early time instance of the developing acoustic pressure field on the Orion capsule. The sinusoidal excitation is clearly visible on the surface. **Figure 4B** illustrates the acoustic pressure field on the Orion and in a cutout of the chamber at the final time of 0.2 s. The acoustic pressure field has become visibly diffuse at this instant.

Conclusions

This article has presented a discussion of modern HPC hardware and software advances for modeling acoustic problems with large numbers of degrees of freedom, which arise in a wide range of applications. Example applications in ground-based acoustics testing and underwater acous-

Figure 4. ***A:*** *acoustic pressure field on NASA's Orion capsule suspended in Sandia's acoustic reverberation chamber at an instant early in time. The sound field is generated by an acoustic source in the nearest bottom corner of the room. The source generates acoustic waves that are visible and coherent at this early time instant (not yet diffuse). The red represents high acoustic pressure, and the blue is low acoustic pressure.* ***B:*** *diffuse acoustic pressure field on NASA's Orion capsule after 0.2 s of simulated time. The pressure field from* ***A*** *has evolved into a diffuse field, which is needed to model the statistical behavior and structural response to a random input, such as launch and reentry. The distribution of this diffuse field is the desired output of a reverberation chamber experiment.*

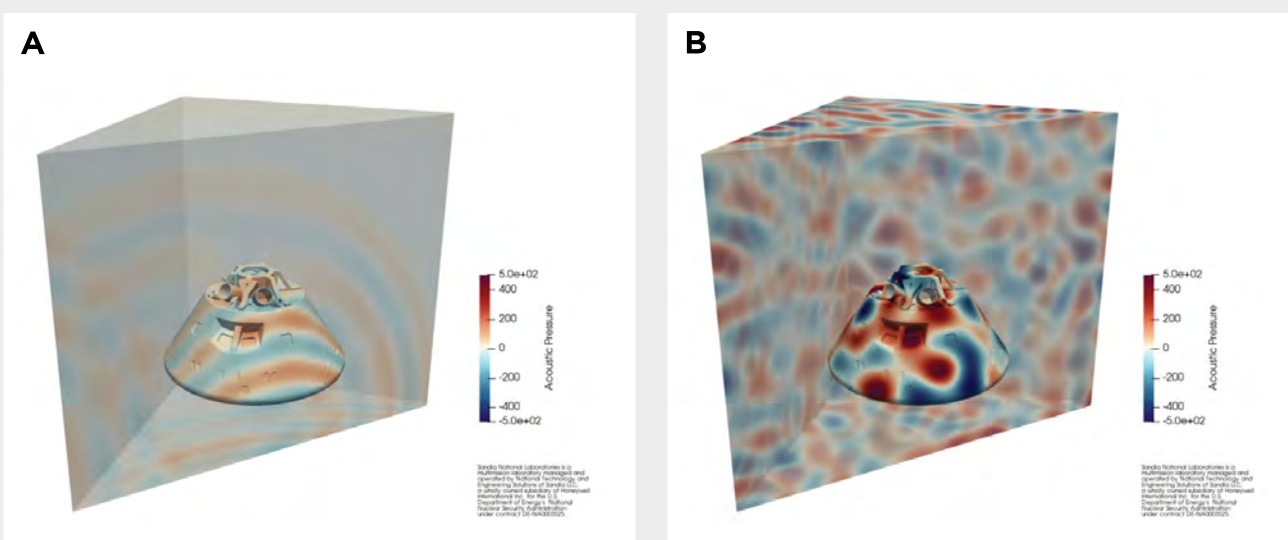

tics on models and corresponding linear systems with over 2 billion degrees of freedom were demonstrated, showing the potential for HPC to enable acoustics solutions that were not previously possible. As future software and hardware advances continue to evolve, one can expect HPC to continue to expand the range of acoustics problems that can be solved for realistic engineering applications.

Although the technology has made great strides in recent decades, there is still significant research that is ongoing and more that is required for HPC to continue to expand the boundaries of large-scale acoustic modeling. Some areas where emerging computational research is essential include

- The optimal use of GPU-based architectures with high-order finite elements;
- Continued development of GPU-aware multilevel domain decomposition and multigrid solvers for acoustics and structural acoustics problems;
- Advances in mesh creation for large-scale acoustics problems; and
- Advances in HPC for large-scale optimization (e.g., design and inverse problems) problems in structural acoustics, wherein the solution of the acoustics equations is inside of an optimization loop.

Acknowledgments

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the US Department of Energy National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the US Department of Energy or the United States Government. The capabilities reported herein are the work of many people including the authors, Nathan Crane, David Day, Sean Hardesty, Payton Lindsay, Lynn Munday, Kendall Pierson, Jerry Rouse, Scott Gampert, Ryan Schultz, Nicholas Reynolds, Michael Miraglia, and Jonathan Stergiou.

References

Atalla, N., and Sgard, F. (2017). *Finite Element and Boundary Methods in Structural Acoustics and Vibration*. CRC Press, Boca Raton, FL.

Bhardwaj, M., Pierson, K., Reese, G., Walsh T., Day, D., Alvin, K., Peery, J., Farhat, C., and Lesoinne, M. (2002). Salinas: A scalable software for high-performance structural and solid mechanics simulations. *In Proceedings of the 2002 ACM/IEEE Conference on Supercomputing*, Baltimore, MD, November 16–22, 2002.

Bunting, G. (2019). *Strong and Weak Scaling of the Sierra/SD Eigenvector Problem to a Billion Degrees of Freedom*. Technical Report SAND2019-1217, Sandia National Laboratories, Albuquerque, NM. Available at https://www.osti.gov/biblio/1494162.

Craggs, A. (1971). The transient response of a coupled plate-acoustic system using plate and acoustic finite elements. *Journal of Sound and Vibration* 15, 509-528.

Craggs, A. (1972). An acoustic finite element approach for studying boundary flexibility and sound transmission between irregular enclosures. *Journal of Sound and Vibration* 30, 331-339.

Cray, S. (2020) *Seymour Cray*. Wikipedia. Available at https://en.wikipedia.org/wiki/Seymour_Cray.

Davis, T. (2006). *Direct Methods for Sparse Linear Systems* (*Fundamentals of Algorithms 2*). Society for Industrial and Applied Mathematics, Philadelphia, PA.

Dohrmann, C., and Widlund, O. (2010). Hybrid domain decomposition algorithms for compressible and almost incompressible elasticity. *International Journal for Numerical Methods in Engineering* 82, 157-183.

Duda, T., Bonnel, J., Coelho, E., and Heaney, K. (2019). Computational acoustics in oceanography: The research roles of sound field simulations. *Acoustics Today* 15(3), 28-37. https://doi.org/10.1121/AT.2019.15.3.28.

Gladwell, G. (1965). A finite element method for acoustics. In *Proceedings of the Fifth International Conference on Acoustics*, Liege, Belgium, Paper L33.

Hart, C., Reznicek, N., and Wilson, C. (2016). Comparison between physics-based, engineering, and statistical learning models for outdoor sound propagation. *The Journal of the Acoustical Society of America* 139, 2640-2655.

Hochgraf, K. (2019). The art of concert hall acoustics: Current trends and questions in research and design. *Acoustics Today* 15(1), 28-36.

Oak Ridge National Laboratory. (2020). *Summit*. Leadership Computing Facility, Oak Ridge National Laboratory, Oak Ridge, TN. Available at https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit.

Moore, G. (1965). Cramming more components onto integrated circuits. *Electronics* 38(8), 114-117.

Moyer, T., Stergiou, J., Reese, G., Luton, J., and Abboud, N. (2016). Navy enhanced sierra mechanics (NESM): Toolbox for predicting Navy shock and damage. *Computing in Science & Engineering* 18(6), 10-18.

National Aeronautics and Space Administration (2019). *Orion Capsule*. National Aeronautics and Space Administration, Washington, DC. Available at https://nasa3d.arc.nasa.gov/detail/orion-capsule.

Pierce, A. (2019). *Acoustics*, 3rd ed. Springer International Publishing, Cham, Switzerland.

Schultz, R., Ross, M., Stasiunas, E., and Walsh, T. (2015). Finite element simulation of a direct-field acoustic test of a flight system using acoustic source inversion. In *Proceedings of the 86th Shock and Vibration Symposium, Shock and Vibration Exchange*, Orlando, FL, October 5–8, 2015.

Smith, B., Bjørstad, P., and Gropp, W. (1996). *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, New York.

Suslick, K. (2019). The dawn of ultrasonics and the palace of science. *Acoustics Today* 15(4), 38-46. https://doi.org/10.1121/AT.2019.15.4.38.

Top500. (2020). Summit up and running at Oak Ridge, claims first exascale application. Available at https://www.top500.org.

Zienkiewicz, O., and Newton, R. (1969). Coupled vibrations of a structure submerged in a compressible fluid. In *Proceedings of International Symposium on Finite Element Techniques*, Stuttgart, Germany, May 1–15, 1969, pp. 359-379.

About the Authors

Gregory Bunting
gbuntin@sandia.gov

Sandia National Laboratories
PO Box 5800, MS 0845
Albuquerque, New Mexico 87185-0845, USA

Gregory Bunting earned a PhD from Purdue University (West Lafayette, IN) in 2016. Gregory currently works in the Solid Mechanics and Structural Dynamics Group, Sandia National Laboratories (Albuquerque, NM), which develops and maintains simulation and modeling capabilities for real-world problems. His research interests include real-time hybrid simulation, structural acoustics, perfectly matched layers, inverse problems, and higher order methods.

Clark R. Dohrmann
crdohrm@sandia.gov

Sandia National Laboratories
PO Box 5800, MS 0845
Albuquerque, New Mexico 87185-0845, USA

Clark R. Dohrmann is a staff member in the Computational Solid Mechanics and Structural Dynamics Department, Sandia National Laboratories (Albuquerque, NM), having received his degrees from Ohio State University (Columbus). His research interests include preconditioners for iterative solvers, finite-element methods, and structural dynamics. He currently resides in Tennessee with his wife and seven children.

Scott T. Miller stmille@sandia.gov

Sandia National Laboratories
PO Box 5800, MS 0845
Albuquerque, New Mexico 87185-0845, USA

Scott T. Miller is a computational mechanician with research interests in fluid-structure interaction, blast-induced traumatic brain injuries, and scientific and engineering software development. He completed his PhD in theoretical and applied mechanics, University of Illinois (Champaign-Urbana), previously worked in Applied Research Laboratory, Pennsylvania State University (University Park), and is currently a staff member at Sandia National Laboratories (Albuquerque, NM).

Timothy F. Walsh
tfwalsh@sandia.gov

Sandia National Laboratories
PO Box 5800, MS 0897
Albuquerque, New Mexico 87185-0897, USA

Timothy F. Walsh received a PhD in computational and applied mathematics from the University of Texas at Austin in 2001. He has been a staff member at Sandia National Laboratories (Albuquerque, NM) for 19 years, with research interests in computational acoustics, structural dynamics, inverse problems, and design optimization.