Transient finite element/equivalent sources using direct coupling and treating the acoustic coupling matrix as sparse

John B. Fahnline and Micah R. Shepherd

ARTICLES YOU MAY BE INTERESTED IN

Solving transient acoustic boundary value problems with equivalent sources using a lumped parameter approach
The Journal of the Acoustical Society of America **140**, 4115 (2016); https://doi.org/10.1121/1.4968789

Efficient, wide-band rigid-body and elastic scattering computations using transient equivalent sources
The Journal of the Acoustical Society of America **146**, 2080 (2019); https://doi.org/10.1121/1.5125424

Machine learning in acoustics: Theory and applications
The Journal of the Acoustical Society of America **146**, 3590 (2019); https://doi.org/10.1121/1.5133944

A review of finite-element methods for time-harmonic acoustics
The Journal of the Acoustical Society of America **119**, 1315 (2006); https://doi.org/10.1121/1.2164987

The full-field equations for acoustic radiation and scattering
The Journal of the Acoustical Society of America **105**, 2574 (1999); https://doi.org/10.1121/1.426873

Improved accuracy for radiation damping in coupled finite element/equivalent source computations
The Journal of the Acoustical Society of America **150**, 2375 (2021); https://doi.org/10.1121/10.0006414

Transient finite element/equivalent sources using direct coupling and treating the acoustic coupling matrix as sparse

John B. Fahnline[a)] and Micah R. Shepherd

Applied Research Laboratory, The Pennsylvania State University, P.O. Box 30, State College, Pennsylvania 16804, USA

Transient structural-acoustic problems can be solved using time stepping procedures with the structure and fluid modeled using finite elements and equivalent sources, respectively. Limitations on the time step size for stable solutions have led to the current popularity of iterative coupling to enforce the boundary conditions at the fluid-structure interface, which also helps to alleviate difficulties caused by the fully populated acoustic coupling matrix. The research presented here examines a monolithic approach using a stabilized equivalent source formulation where the acoustic coupling matrix is either fully diagonal or treated as sparse. In theory, the matrix should be sparse because it relates nodal velocities to nodal acoustic pressure forces during a single time step, and the pressure waves can only travel a distance equal to the sound speed multiplied by the time step. The numerical results demonstrate that for the chosen example problems accurate results are obtained for either diagonal coupling matrices or with a large percentage of the terms set to zero. It is also demonstrated that the formulation adapts well to parallel processing environments and that the times associated with the equivalent source computations are proportional to the number of processors. © 2017 Acoustical Society of America. [http://dx.doi.org/10.1121/1.4998591]

I. INTRODUCTION

Transient coupled fluid-structure problems occur in a number of different disciplines and the available literature on the subject is quite large. When the fluid is considered to be infinite, boundary element methods are commonly used to model the fluid since only the boundary surface of the fluid has to be meshed, rather than the whole solution volume. Also, boundary element methods naturally satisfy the farfield radiation boundary conditions, whereas finite element methods require some technique, such as infinite elements or perfectly matched layers, to prevent spurious reflections at the outer mesh boundary. A good discussion of the origins of infinite elements and an assessment of their accuracy is given by Astley.[1] The paper by Zampolli *et al.*[2] gives a good overview of perfectly matched layers in the context of structural-acoustic scattering problems. For very large time-harmonic problems, these methods are advantageous primarily because they produce sparse matrices.

In contrast, boundary element and equivalent source computations for time-harmonic problems produce complex-valued, fully populated matrices. When coupled with finite element (FE) solutions for the structure, the combined equation system becomes very time consuming to solve, restricting the analyses to modal formulations and relatively low frequencies. Transient BE or ES solutions have not been considered because historically they have suffered from long-time instabilities. A number of researchers have shown that these instabilities are related to the well-known nonuniqueness/nonexistence problems in frequency domain boundary

element calculations, and that the CHIEF (Ref. 3) and Burton and Miller[4] formulations can help to eliminate the instabilities. A detailed discussion of nonuniqueness difficulties for transient equivalent source (ES) computations is given by Fahnline.[5]

Even using a stabilized boundary element formulation, limitations may still exist on the range of time step sizes that will produce stable and accurate results, as discussed by Soares.[6] It has been widely reported in the literature that the solutions are stable within a range of values for $c\Delta t/L$, where c is the sound speed in the fluid, Δt is the time step size, and L is the nominal element size of the boundary surface mesh. For large values, the solution remains stable, but does not produce accurate results. Only a few papers discuss the source of the instabilities when the small time step size is small. Yu *et al.*[7] and Frangi[8] suggest that the instabilities are related to the use of noncausal interpolation functions that instantaneously propagate nodal displacements within an element. These difficulties are exacerbated when performing coupled structural-acoustic analyses because structural analyses also impose restrictions on the time step size for accurate solutions. Since the wave speeds can differ considerably in the fluid and structure, the two sets of restrictions often conflict with each other. This has led to the current popularity of staggered solution formulations over more traditional direct coupling solutions. In a staggered approach, the structural and acoustic problems are solved independently, with iterative procedures used to derive the coupled solution. This helps to alleviate stability issues because appropriate time steps can be used for each solution domain. However, staggered formulations tend to be more complicated than direct coupling approaches.

a)Electronic mail: jbf103@arl.psu.edu

Aside from stability issues, the other reason often listed for not using a direct coupling approach is the resulting densely populated acoustic coupling matrices, as discussed by Soares[6] and Felippa et al.[9] This occurs because a matrix solution is generally required to compute the acoustic pressure field. In contrast to frequency domain computations, the coefficient matrix in a time domain boundary element solution using a marching-on-in-time (MOT) solution is sparse because the active variables only exist for the current time step. In both frequency and time domain formulations, the inverse of the coefficient matrix generally becomes fully populated. After some manipulations, the acoustic analysis can be used to compute an acoustic coupling matrix relating input nodal displacements to output nodal pressure forces. Physical arguments dictate that activating a single nodal degree-of-freedom will not produce a pressure field over the entire boundary surface within the time span of a single time step since sound waves can only travel a distance $c\Delta t$. Thus, in theory the acoustic coupling matrix should actually be sparse, and it becomes fully populated due to small numerical errors in the coefficient matrix and the process of solving the matrix equation. A literature review has not revealed any papers where this fact has been exploited by other researchers in the past, although it is difficult to make a definitive statement due to the vast number of papers on the subject.

Here, a MOT scheme will be derived for structural-acoustic problems using direct coupling and treating the acoustic coupling matrix as sparse. The formulation most closely resembles the coupled finite element/boundary element (FE/BE) frequency-domain solutions derived in the classic paper by Wilton,[10] where the acoustic variables are condensed out of the equations of motion. Ultimately, this leads to an efficient formulation that adapts well to parallel processing environments. The primary interest here is in metal structures in water, for which the formulation is well suited with reasonable choices for the mesh and time step sizes. Aside from treating the acoustic coupling matrix as sparse, the current formulation is novel because it uses a stabilized ES solution for the acoustic field, a detailed discussion of which can be found in the paper by Fahnline.[5]

II. NUMERICAL FORMULATION

The main challenge in solving transient coupled structural-acoustic problems is computing the pressure field due to the structural vibrations. Here, the ES formulation discussed by Fahnline[5] will be used to calculate the pressure field. The ES formulation uses average pressures and elemental volume velocities as variables, and these quantities must be converted to forces and velocities for the nodal degrees-of-freedom. Given a vector of nodal velocities at time t, the resulting volume velocity distribution can be computed by taking the dot product of the nodal velocity with the surface normal and integrating over each of the surface elements as

$$u_\nu(t) = \int\!\!\int_S \boldsymbol{v}(\boldsymbol{x},t) \cdot \boldsymbol{n} \, \mathrm{d}S_\nu(\boldsymbol{x}), \tag{1}$$

where $\nu = 1, 2,\ldots,NA$ and NA is the number of acoustic elements on the boundary surface. For linear elements, the volume velocity for an element is the average of the normal component of its nodal velocities multiplied by its area. The translation from nodal velocity to elemental volume velocity is noncausal because any nonzero nodal velocity is instantaneously translated into a nonzero volume velocity for the element as a whole.

A similar procedure can be used to translate an elemental pressure distribution on the exterior boundary surface in contact with the fluid to nodal forces. Knowing the pressure averaged over the surface of an element, the force vector for the nodal degrees-of-freedom can be computed by assigning the direction using the element normal, and dividing the resulting force amongst the nodes as

$$\boldsymbol{f}_{\mathrm{nodal}}(t) = \frac{1}{N} S_\nu p_\nu(t)\boldsymbol{n}_\nu, \tag{2}$$

where N is the number of nodes for element ν. The translation from elemental pressures to nodal forces is also noncausal because any nonzero pressure over an element is instantaneously translated to nodal forces for all the associated nodes.

Taking all the nodes and elements into account, Eqs. (1) and (2) can be written in matrix form as

$$\boldsymbol{u}(t) = \boldsymbol{V}\boldsymbol{v}(t) \tag{3}$$

and

$$\boldsymbol{f}_{\mathrm{acs}}(t) = \boldsymbol{F}\boldsymbol{p}(t), \tag{4}$$

respectively, where both $\boldsymbol{V}$ and $\boldsymbol{F}$ are sparsely populated. The matrix $\boldsymbol{V}$ has dimensions $NA \times NDOF$ and $\boldsymbol{F}$ has dimensions $NDOF \times NA$, with $NDOF$ being the number of degrees-of-freedom in the FE model.

A. Equivalent source formulation

For a specified nodal velocity vector, the nodal acoustic pressure forces can be computed once a matrix relating $\boldsymbol{p}(t)$ to $\boldsymbol{u}(t)$ is supplied by the equivalent source analysis. Here, only acoustic radiation problems will be considered. To derive the equivalent source solution, the boundary surface is input by the user as a collection of linear triangular and quadrilateral elements, and a discrete source is placed at the center of each element to represent the acoustic field. As a precursor to the time-stepping analysis, a triangular pulse of unit amplitude and length $2\Delta t$ is input to each source individually and the average pressure and volume velocity is computed over each element as a function of time, where Δt is the time step size. Because of the finite length of time for the triangular pulses, the resulting average pressures and volume velocities are only nonzero for a finite number of time steps. The data are stored as a three-dimensional matrix with indices time step, source number, and receiver number. In the equations, matrix notation is used with a single index for time step, and the rows and columns of the matrices correspond to receiver and source number, respectively. Also, a subscript is used to indicate the source type. Thus, $\boldsymbol{P}_s(n)$ and

John B. Fahnline and Micah R. Shepherd

$U_s(n)$ are the matrices of average pressure and volume velocities over the boundary surface elements at time step n due to triangular pulses of length $2\Delta t$ starting at $t=0$ from simple sources (subscripts d and t indicate dipole and tripole sources, respectively). The Appendix gives equations to define the matrix terms.

A full discussion of the process of integrating the pressure and acoustic particle velocity over the surface elements was given in the previous paper by Fahnline,[5] and the discussion here will be limited to an illustration of the functional forms. In general, the integrations are difficult to perform accurately due to the singularities in pressure and velocity fields at the source locations. These difficulties can be alleviated by transforming to cylindrical coordinates with the origin located at the projection of the source in the plane of the element. In most cases, the radial integral can be performed analytically, leaving an integration over the azimuthal angle in the cylindrical coordinate system to be performed numerically using Gauss quadrature. The integration algorithm is limited to planar elements, and nonplanar quadrilateral elements are treated as being made up on an interior planar quadrilateral element using the midpoints of the element edges and four triangles using two midpoints and one of the corner nodes. The integration routines have been carefully validated, both with comparisons to analytical formulas for simple geometries and to very high order Gaussian integration over quadrilateral elements.

To illustrate the functions, a few terms of the matrices $P_s(t)$, $U_s(t)$, $P_t(t)$, and $U_t(t)$ will be shown. Figure 1 shows $P_s(t)$ and $U_s(t)$ for a simple source at the center of element 1 over several nearby receiver elements, the locations for which are illustrated in the inset diagram. The surface normal

for element 1 is assumed to point upwards, in the direction opposite element 4. The acoustic medium for the analysis is water with sound speed 1500 m/s and density 1000 kg/m^3. Also, the element size is 0.019 m, and the time step size is taken to be 1×10^{-6} s, so that $c\Delta t/L = 0.079$, which is typical for the subsequent computations. The curves in the figure are shown with resolution $t/\Delta t = 0.1$ to fully define the functional variations. In an actual computation, the curves would be sampled at integer values of $t/\Delta t$. To help check the integrations, it is often useful to calculate the times when the pulses should begin and end. For example, the element edges should begin to contribute to the average pressure over element 1 when $t/\Delta t = (L/2)/c\Delta t = 6.35$, and they should end when $t/\Delta t = (\sqrt{2}L/2)/c\Delta t + 2\Delta t/t = 10.98$, as confirmed in the figure. The differences in the times for the waves to travel across the elements and the differences in the magnitudes of the average pressures and volume velocities help to explain why the interior acoustic field causes numerical difficulties. Because the spherical wavefront first reaches the center of receiver element 4 and the source is aligned perpendicular to the element surface, it passes much more quickly has larger peak magnitude than elements 1 and 2.

The large magnitude for the acoustic quantities in the direction opposite the surface normal results in a large interior acoustic field, and, ultimately, instabilities. The tripole source formulation addresses this difficulty by creating a cardioid source, where the acoustic quantities go to zero in the direction opposite the surface normal as the distance from the source increases. Figure 2 shows $P_t(t)$ and $U_t(t)$ for a tripole source at the center of element 1 over several nearby receiver elements. Comparing to Fig. 1, the main difference is the reduction in the magnitude in the average pressure and
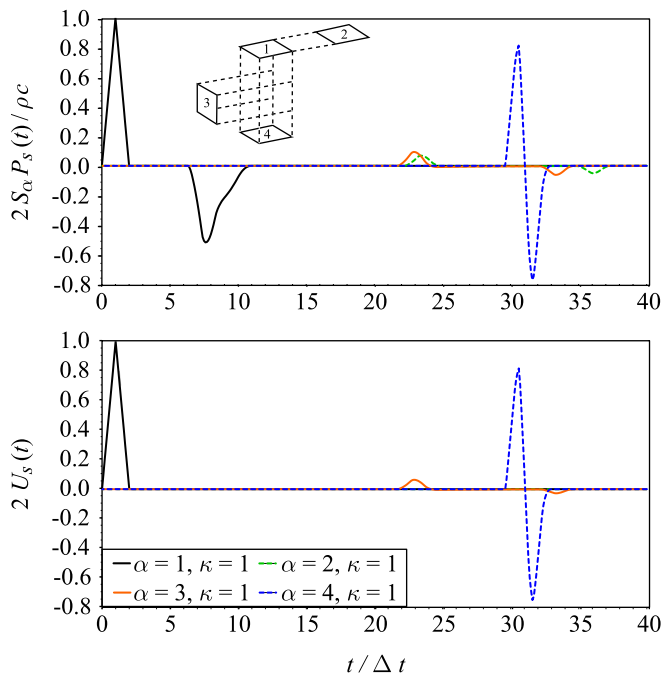

FIG. 1. (Color online) The normalized matrix coefficients for a simple source at the centroid of element 1 over elements 1–4, the locations for which are illustrated in the inset diagram. The definitions of $P_s(t)$ and $U_s(t)$ are given in the Appendix.
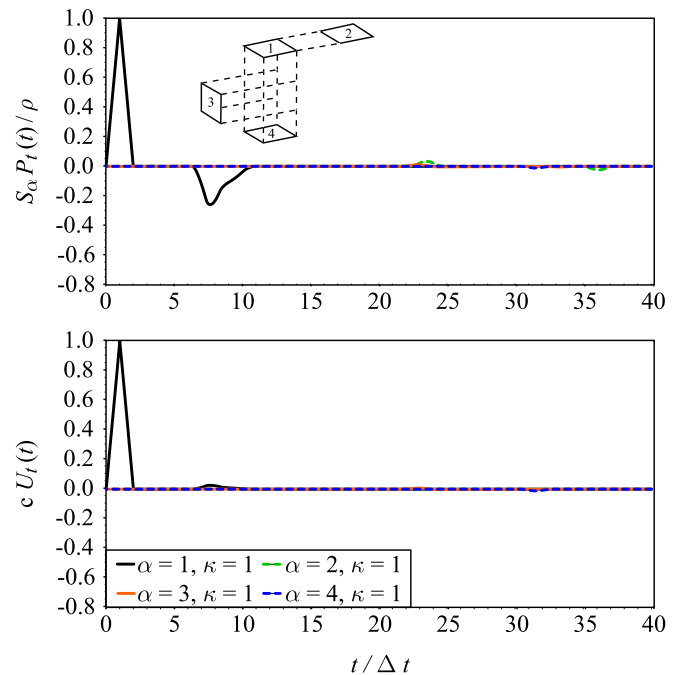
FIG. 2. (Color online) The normalized matrix coefficients for a tripole source at the centroid of element 1 over elements 1–4, the locations for which are illustrated in the inset diagram. The definitions of $P_t(t)$ and $U_t(t)$ are given in the Appendix.

volume velocity for receiver element 4. With that as a brief synopsis of the equivalent source formulation, the coupled FE/ES computations will now be addressed. For the purposes of this paper, all that is required from the equivalent source analysis are the matrices relating source amplitudes to the average pressures and volume velocities over the surface elements and it is assumed that they are computed accurately.

B. Simple source formulation

The equivalent source formulation is derived by assuming a linear interpolation of the sources amplitudes as a function of time, where the time step size remains constant. Since linear interpolation can also be interpreted as a summation of triangular pulses, the volume velocity distribution over the boundary surface for time step n can be written for the simple source formulation in the form

$$u(n) = \sum_{\mu=1}^{n} U_s(\beta)q(\mu), \tag{5}$$

where the number in parentheses indicates the time step, $q(n)$ is the vector of source amplitudes, and $\beta = n + 1 - \mu$. Also, $U_s(\beta)$ refers to the discrete matrix after it has been sampled at times $t = n\Delta t$. The letter β was chosen because it counts backwards as μ increases. Assuming $q(\mu)$, $\mu = 1, ..., n-1$ are known and $q(n)$ is unknown, the known quantities can be moved to the right-hand-side of the equation to yield

$$U_s(1)q(n) = u(n) - \sum_{\mu=1}^{n-1} U_s(\beta)q(\mu). \tag{6}$$

The summation on the right-hand-side accounts for sound radiated in the past that creates a volume velocity over the boundary surface elements at the current time step. Solving for $q(n)$ in terms of $u(n)$ gives

$$q(n) = [U_s(1)]^{-1} \left[u(n) - \sum_{\mu=1}^{n-1} U_s(\beta)q(\mu) \right]. \tag{7}$$

The matrix $U_s(1)$ represents the volume velocities created by the sources during the current time step.

Using the solution for the source amplitudes in Eq. (7), the pressures over the boundary surface can be computed as

$$p(n) = \sum_{\mu=1}^{n} P_s(\beta)q(\mu) = P_s(1)q(n) + \sum_{\mu=1}^{n-1} P_s(\beta)q(\mu)$$

$$= P_s(1)[U_s(1)]^{-1} \left[u(n) - \sum_{\mu=1}^{n-1} U_s(\beta)q(\mu) \right]$$

$$+ \sum_{\mu=1}^{n-1} P_s(\beta)q(\mu). \tag{8}$$

Substituting for $p(n)$ and $u(n)$ from Equations (3) and (4), the relationship between the nodal velocities and acoustic pressure forces can be written as

$$f_{\text{acs}}(n) = FP_s(1)[U_s(1)]^{-1} \left[Vv(n) - \sum_{\mu=1}^{n-1} U_s(\beta)q(\mu) \right]$$

$$+ F \sum_{\mu=1}^{n-1} P_s(\beta)q(\mu). \tag{9}$$

Taking the acoustic coupling matrix to be A, it can be identified as $A = -FP_s(1)[U_s(1)]^{-1}V$, where the minus sign is added so that A is a positive quantity when it is added into the equations of motion. The acoustic forces then become

$$f_{\text{acs}}(n) = -Av(n) + cs_s, \tag{10}$$

where cs stands for "convolution summations" and is given for the simple source formulation as

$$cs_s = -FP_s(1)[U_s(1)]^{-1} \sum_{\mu=1}^{n-1} U_s(\beta)q(\mu)$$

$$+ F \sum_{\mu=1}^{n-1} P_s(\beta)q(\mu). \tag{11}$$

In the numerical computations, the matrix becomes fully populated due to small errors in evaluating the integrals and finite precision arithmetic in factoring, solving, and multiplying the matrices. However, many of these terms should actually be identically zero since the distance the acoustic waves can travel during a single time step is limited by the time step size. Rather than processing the matrix as fully populated, the small terms in the matrix $P_s(1)[U_s(1)]^{-1}$ can be zeroed out, so that A becomes sparse.

Now that an equation for the acoustic pressure forces has been derived, the full structural-acoustic problem can be addressed. The formulation will follow that used in NASTRAN for direct transient response analyses, as discussed by Blakely,[11] except with extra terms added to represent the acoustic pressure forces. This form has been found to be more stable than that discussed by Cook et al.[12] The equations of motion are

$$Kd(n) + Cv(n) + Ma(n) = f_{\text{ext}}(n) + f_{\text{acs}}(n), \tag{12}$$

where K, C, and M are the finite element stiffness, damping, and mass matrices, respectively, d and a are the nodal displacement and acceleration vectors, respectively, and f_{ext} is the applied external force vector. Substituting for $f_{\text{acs}}(n)$ from Eq. (10) gives

$$Kd(n) + (A + C)v(n) + Ma(n) = f_{\text{ext}}(n) + cs_s. \tag{13}$$

This can be converted to a time marching equation by substituting for the velocity and acceleration using the difference formulas

$$v(n-1) = \frac{1}{2\Delta t} \left[d(n) - d(n-2) \right] \tag{14}$$

and

John B. Fahnline and Micah R. Shepherd

$$a(n-1) = \frac{1}{\Delta t^2}\left[d(n) - d(n-1) + d(n-2)\right], \qquad (15)$$

respectively, and placing all the known quantities on the right-hand-side of the resulting equation. Averaging over three consecutive time steps, the time-stepping equation then becomes

$$
\begin{aligned}
K_{\mathrm{eff}}d(n) = & \frac{1}{3}\sum_{\mu=0}^{2}f_{\mathrm{ext}}(n-\mu) + cs_s \\
& + \frac{M}{\Delta t^2}[2d(n-1) - d(n-2)] \\
& + \frac{(A+C)}{2\Delta t}d(n-2) \\
& - \frac{K}{3}[d(n-1) - d(n-2)],
\end{aligned}
\qquad (16)
$$

where the effective stiffness matrix K_{eff} is given as

$$K_{\mathrm{eff}} = \frac{M}{\Delta t^2} + \frac{(A+C)}{2\Delta t} + \frac{K}{3}. \qquad (17)$$

Everything on the right-hand-side of Eq. (16) is known for time step n since the summations embedded within the cs_s term for the sources only extend to time step $n-1$. After solving the equation system for the displacement vector at time step n, the velocity and acceleration vectors can be computed using Eqs. (14) and (15), respectively. Once the velocity vector for the current time step is known, the source amplitude vector can be determined as

$$q(n) = [U_s(1)]^{-1}\left[Vv(n) - \sum_{\mu=1}^{n-1}U_s(\beta)q(\mu)\right]. \qquad (18)$$

The source amplitudes are the only data that need to be stored for each time step. In general, the formulation requires two matrix solutions at each time step; one to compute the nodal displacement vector in Eq. (16) and one to compute the source amplitudes in Eq. (18). Both solutions are fast because the coefficient matrices do not change as long as the time step remains constant and thus they only have to be factored once. The convolution summations in Eqs. (16) and (18) can be computed in parallel since they are all independent, and since each source and element only interact over a finite number of time steps due to the finite duration for the triangular pulses, the computation time per time step quickly reaches a static value.

C. Tripole source formulation

The average pressure over the boundary surface for the tripole source formulation is given as

$$p(n) = \sum_{\mu=1}^{n}\left[\frac{1}{c}P_s(\beta) + P_d(\beta)\right]\dot{f}(\mu) = \sum_{\mu=1}^{n}P_t(\beta)\dot{f}(\mu), \qquad (19)$$

where the subscript d indicates dipole sources. The amplitude for the simple source has been replaced with its time derivative and a factor of $1/c$. In the frequency domain, this translates to a coupling factor of $-i\omega/c$ with an assumed $e^{-i\omega t}$ time dependence.

The notation for the volume velocity of a dipole source is slightly different because it depends not only on source amplitude, but also on the time derivative of the source amplitude. The terms dependent on the source amplitude and its time derivative are given subscripts 0 and 1, respectively. Thus, $U_{d,0}(n)$ and $U_{d,1}(n)$ are the components of the volume velocities over the boundary surface elements at time step n due to triangular pulses of length $2\Delta t$ starting at $t=0$ for the dipole source amplitude and its time derivative, respectively. The dipole source amplitude for the current time step is then eliminated using second order backwards differences as $f(n) \approx (2\Delta t/3)\dot{f}(n) + \frac{4}{3}f(n-1) - \frac{1}{3}f(n-2)$, so that the volume velocity of a tripole source can be written as

$$
\begin{aligned}
u(n) = & \sum_{\mu=1}^{n}\left[\frac{1}{c}U_s(\beta) + U_{d,1}(\beta)\right]\dot{f}(\mu) + \sum_{\mu=1}^{n}U_{d,0}(\beta)f(\mu) \\
= & \sum_{\mu=1}^{n}U_t(\beta)\dot{f}(\mu) + \sum_{\mu=1}^{n}U_{d,0}(\beta) \\
& \times \left[\frac{4}{3}f(\mu-1) - \frac{1}{3}f(\mu-2)\right],
\end{aligned}
\qquad (20)
$$

where

$$U_t(n) = \frac{1}{c}U_s(n) + \frac{2\Delta t}{3}U_{d,0}(n) + U_{d,1}(n), \qquad (21)$$

and the primary variable of interest is $\dot{f}$. Solving Eq. (20) for $\dot{f}(n)$, substituting the result into Eq. (19), and taking $A = -FP_t(1)[U_t(1)]^{-1}V$, the acoustic force vector becomes

$$f_{\mathrm{acs}}(n) = -Av(n) + cs_t, \qquad (22)$$

where

$$
\begin{aligned}
cs_t = & -FP_t(1)[U_t(1)]^{-1}\left\{\sum_{\mu=1}^{n-1}U_t(\beta)\dot{f}(\mu)\right. \\
& + \sum_{\mu=1}^{n}U_{d,0}(\beta)\left[\frac{4}{3}f(\mu-1) - \frac{1}{3}f(\mu-2)\right]\right\} \\
& + F\sum_{\mu=1}^{n-1}P_t(\beta)\dot{f}(\mu).
\end{aligned}
\qquad (23)
$$

The time-stepping equations then are the same as given previously in Eq. (16) with a different version of A and cs_t substituted for cs_s. As for the simple source formulation, the equations of motion are first solved for the displacement vector, which is used to compute the velocity and acceleration vectors and the solution for the source amplitudes at the current time step.

J. Acoust. Soc. Am. **142** (2), August 2017

John B. Fahnline and Micah R. Shepherd   1015

D. Computational strategies

Aside from stability issues, which have plagued this type of computation in the past, the two main difficulties with the coupled FE/ES formulation are the computation and storage of the acoustic matrices and the subsequent computation of the convolution summations on the right-hand side of the equations. Similar conclusions were drawn by Walker and Yeung[13] and Ergin *et al.*[14] although the computational resources available have grown enormously since their papers were published. Both difficulties can be addressed in a cluster computing environment by distributing the data and performing the computations on multiple nodes and processors. The data can be stored as full matrices for a finite number of time steps, along with an integer matrix indicating the first nonzero time step for each source-receiver combination. This format works well in a distributed processing environment because all the data for any source-receiver combination will reside on a single processor, thus allowing the convolution summations for the acoustic field on the right-hand-side of Eq. (16) to be performed independently. Also, all the processors can be used to compute the matrices before entering the time-stepping phase of the computations. Along with the matrix coefficients, the source amplitudes associated with each source-receiver combination must be stored locally on the processors. It is thus beneficial to maximize the number of receiver locations stored with each source to reduce the number of copies of the source amplitudes that must be stored. This is especially important when the number of time steps grows large. Once the convolution summations have been computed, it is a simple matter to retrieve the data from each of the processors.

A number of different packages, including PARDISO,[15] MUMPS,[16] and SUPERLU,[17] can be used to solve sparse matrix equations, while SCALAPACK (Ref. 18) is appropriate for dense matrices. PARDISO is embedded within the Intel MKL libraries, while MUMPS and SUPERLU require separate installations. Since the data for $q(t)$, P, and U are more easily stored in SCALAPACK's block cyclic data distribution, the acoustic matrices are treated as dense and the $P(1)[U(1)]^{-1}$ computation is performed using SCALAPACK. The situation is different for the source amplitude solution at each time step because the number of input right-hand-sides is generally much smaller. In this case, the sparse matrix solutions generally outperform SCALAPACK, especially as the sparsity of $U(1)$ increases. For the sake of robustness and because the time required to solve the matrix system for a single time step is typically small, the solution times for the first time step are tabulated for each of the solvers, with the fastest one used for the subsequent time steps.

The coupled FE/ES equation system is much larger in size, and would require too much time to solve and too much storage to be treated as being densely populated. Timing tests for the various sparse solvers were then used to choose the best one for the current application. In most problems of interest, the coefficient matrix is nearly symmetric (MUMPS consistently lists the matrix's "structural symmetry" as 99%) and diagonally dominant. Ultimately MUMPS was chosen because it was more easily installed than SUPERLU and it

gave better performance than PARDISO. Counterintuitively, using more processors can actually increase solution times in MUMPS, and the fastest solutions were generally achieved using a single multi-processor node. Timing tests in the paper by Amestoy *et al.*[19] show similar behavior. For the problems considered here, all the processors on a single computational node are used for the MUMPS matrix solution.

III. EXAMPLE PROBLEMS WITH DIAGONAL COUPLING MATRICES

Two example problems will be solved in this section to illustrate the calculations and to demonstrate the accuracy of the formulation, including a flat plate and a thin sphere. The example problems were chosen because they could reasonably be modeled without including any structural damping. The computations are validated in two ways; first by either comparisons to experimental measurements or convergence tests in the time domain and second to transfer function computations and measurements in the frequency domain. In general, it is difficult to transform transient numerical computations of this type to the frequency domain because a large number of time steps must be computed to obtain adequate frequency resolution. Even when it is possible to perform a large number of time steps in a reasonable amount of time, the solution must remain stable or the transformed data will be dominated by the unstable component of the time series.

For the measurements, a modal impact hammer was used to excite the structure into vibration. The resulting acceleration was monitored using accelerometers along with the measured input force from the hammer. For the problems discussed here, the hammer produces an input force profile resembling a squared half sine wave, as discussed by Akay and Latcha.[20] One of the main analysis goals is to demonstrate that fluid coupling is included correctly, and the measurements are made in water so that fluid coupling effects are relatively large. The measurements were performed in a reverberation water tank, with the test articles suspended using either fishing line or bungee cords. Simultaneous pressure measurements were also made during the measurements, but the test articles were submerged only slightly below the air-water interface, and reflections of the waves from the water surface quickly invalidate the time domain comparisons (the numerical computations consider the structure to be submerged in an infinite fluid). However, the reflections did not have a significant effect on the acceleration measurements. In the subsequent time domain comparisons between experimental measurements and numerical computations, the length and magnitude of the input force is extracted from the experimental measurements and used as input to the numerical computations. For the frequency domain comparisons, both the input force and output acceleration were first transformed to the frequency domain, and then acceleration-to-force transfer functions are computed for each frequency.

One of the main goals for the examples is to demonstrate that the acoustic coupling matrix can be treated as sparse. Generally, this requires nonzero terms of the $P(1)[U(1)]^{-1}$ matrix to be set to zero. However, if the time step is small

John B. Fahnline and Micah R. Shepherd

enough, waves cannot travel from the source locations at the element centers to any other element in a single time step and the matrices $P(1)$ and $U(1)$ become diagonal. In this case, there is no approximation involved because all of the off-diagonal terms are identically zero. Knowing the smallest element side length for the acoustic mesh L_{min}, the time step size required to make the matrices diagonal can be calculated as $\Delta t < L_{min}/2\,c$. In this limit, the acoustic coupling matrix adds no extra nonzero terms to the effective stiffness matrix and the times to solve the coupled equation system are the same for uncoupled and coupled analyses. Decreasing the time step size further unnecessarily increases the number of nonzero terms for the convolution summations because more time steps are required for triangular pulses of length $2\Delta t$ to travel across the elements. It is also possible for rows and columns of $P(1)$ and $U(1)$ to be diagonal for some of the surface elements, and nondiagonal for other elements. Since it involves no approximation, the time step size will be taken to be small enough to produce diagonal $P(1)$ and $U(1)$ matrices for the examples in this section. In general, the solutions are most stable with $c\Delta t/L = 1$, as discussed by Stutz and Ochmann,[22] and the time step sizes required to give diagonal $P(1)$ and $U(1)$ matrices are considerably smaller.

A. Bronze plate

The first example problem is a bronze plate in water. The plate has width 0.3048 m, length 0.762 m, and thickness 0.04445 m, as illustrated in Fig. 3. The direction for the drive force is assumed to be positive in the outward normal direction, as illustrated in the figure. In the numerical analysis, the material properties for the plate are taken as follows: Young's modulus 117 GPa, Poisson's ratio 0.3, and density 7468 kg/m³. The drive point is at one of the top corners and the response point is the middle point of the shorter edge, as illustrated by the accelerometer location in Fig. 3. The peak measured input force is 500 N and its duration is approximately 5.5×10^{-4} s. The hydrophone is located on the underside of the plate a distance 0.0508 m from the plate surface, 0.42863 m from the driven corner along the longer dimension of the plate, and 0.12383 m from the driven corner along the shorter dimension. No structural damping is included in the analysis, but acoustic radiation damping is included in the matrix A.

For the numerical computations, the minimum element size is 0.0148 m, so that Δt must be less than 4.9×10^{-6} s for $P(1)$ and $U(1)$ to be diagonal. Taking $\Delta t = 4 \times 10^{-6}$, and with the maximum element size equal to 0.019 m, then $c\Delta t/L = 0.32$. Figure 4 shows experimental measurements and numerical simulations of the transient acceleration of the
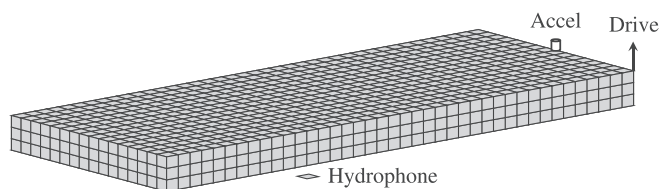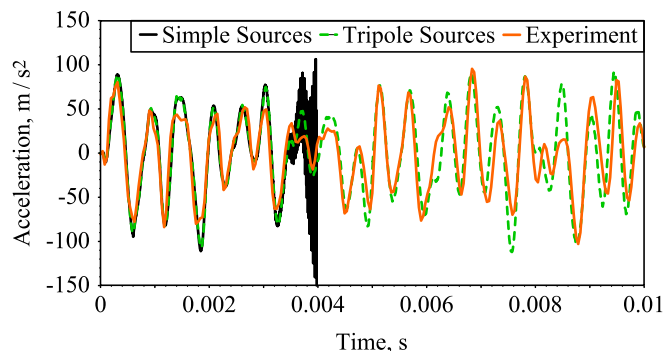

FIG. 4. (Color online) Acceleration as a function of time for the bronze plate.

plate in water for the drive and response points illustrated in Fig. 3. The discrepancies between the experimental and numerical results are likely due to the hammer hit point and accelerometer location not being exactly at the plate corner and edge.

In general, it is easier to compare results in the frequency domain because differences in the resonance frequencies and damping levels are readily apparent. In transforming the time domain results to the frequency domain, the frequency resolution is given as $\Delta f = 1/N\Delta t$, where N is the number of time steps. Because the solution only includes acoustic radiation damping, the lowest frequency resonances of the plate have very little damping, and Δf must be small to resolve the peaks. For this problem, 524 288 time steps of length 4×10^{-6} s were computed, giving a frequency resolution of approximately 0.5 Hz. Even after this many time steps, the acceleration had only decayed to approximately 6.9% of its peak value. The numerical results thus show some "windowing" effects due to cropping the data before it has completely decayed to zero. Recalling that the drive force is input as a half period of a squared sine wave, both the input force and nodal accelerations are transformed to the frequency domain before computing transfer functions, which mimics the techniques used to process the experimental data. Figure 5 shows accelerance comparisons for the bronze plate as a function of frequency, with the drive and response points as illustrated previously in Fig. 3. The frequency domain FE/ES calculations are performed using the formulation discussed in the book chapter by Fahnline.[21] The solution is written in terms of the *in vacuo* structural modes with residual vectors added

FIG. 3. Illustration of the plate mesh and the drive point, response point, and hydrophone locations.
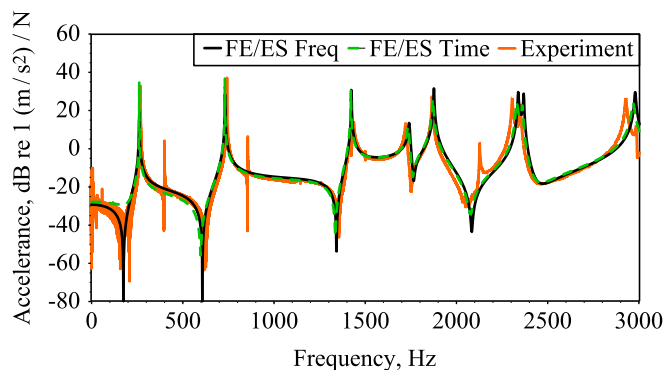

FIG. 5. (Color online) Acceleration as a function of frequency for the bronze plate.

in to represent the contributions from truncated modes. The resonance peaks at 401, 857, and 2144 Hz in the experimental measurements are anti-symmetric modes, and are not evident in the numerical computations due to the response point location on plate centerline.

As mentioned previously, reflections from the air-water interface in the experiments make it difficult to compare pressure measurements to numerical computations in the time domain. However, for this example the reflections do not change the frequency content substantially. Figure 6 shows comparisons of the measured and computed pressure at the hydrophone location illustrated previously in Fig. 3. The measurements show some evidence of electrical noise at 60 Hz harmonics. To test the stability limits, the time step size can be reduced; producing a stable solution for $c\Delta t/L = 0.158$ and an unstable solution for $c\Delta t/L = 0.079$.

B. Thin steel sphere

The previous problem is not a completely general test of the numerical formulation because the acoustic resonances within the boundary surface occur at relatively high frequencies. As a problem where nonuniqueness issues should be more prevalent, a thin steel sphere is considered next. The structural model of the sphere is constructed from QUAD4 shell elements in NASTRAN, where the radius and thickness of the sphere are 1 and 0.05 m, respectively. The material properties are taken as: Young's modulus 200 GPa, Poisson's ratio 0.3, and density 7860 kg/m³. The peak input force is taken to be 1 N and its duration is 0.002 s. Due to the assumption that the input force is half a period of a squared sine wave, the duration must be small enough for no nulls in the input force spectrum to occur in the frequency range of interest. With the duration equal to 0.002 s, the first null occurs at 2000 Hz. Unlike the previous example, physical hardware does not exist for the sphere, so convergence studies and comparisons between frequency and time domain calculations are used to validate the analysis.

Before choosing the mesh for the full analysis, several different meshes were used to test convergence. In general, coarsening a finite element mesh makes the model stiffer, leading to overestimates of the resonance frequencies. In the time domain, this causes the response curves for meshes with different resolutions to resemble each other with a change in the time scale, where overestimates of the resonance frequencies
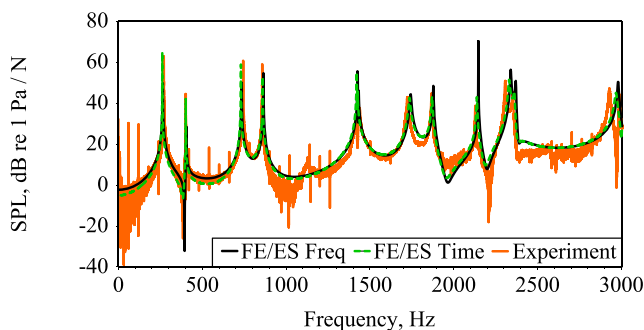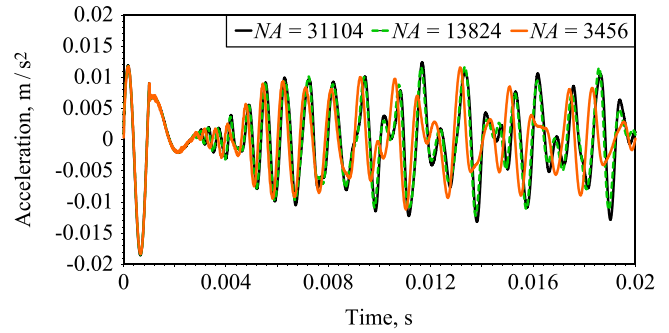

FIG. 7. (Color online) Acceleration as a function of time for the steel sphere using the tripole source formulation.

lead to shorter time scales. To illustrate, Fig. 7 shows comparisons of drive point acceleration for the sphere as a function of time for meshes with $NA = 3546$, 13 824, and 31 104. The smallest element size for the finest mesh is 0.0137 m, and taking the time step to be 4×10^{-6} s gives fully diagonal coupling matrices for all three meshes. The results in the figure show that the meshes with $NA = 13\,824$ and 31 104 give nearly the same solution. In general, refining the mesh incurs a considerable computational cost, so the mesh with $NA = 13\,824$ is used for the full analysis, and it is illustrated in Fig. 8. Ultimately, the process of testing mesh refinement is somewhat more efficient in the time domain than in the frequency domain because only a relatively small number of time steps is required.

For the mesh with $NA = 13\,824$, the maximum element size is equal to 0.039 m, so that $c\Delta t/L = 0.154$. In transforming the results to the time domain, 131 072 time steps of size 4×10^{-6} s were computed, giving a frequency resolution of approximately 1.9 Hz. Figure 9 shows comparisons of the numerical computations for the acceleration-to-force transfer functions at the drive point, where the results for the frequency domain FE/ES analysis are shown for $NA = 13\,824$.
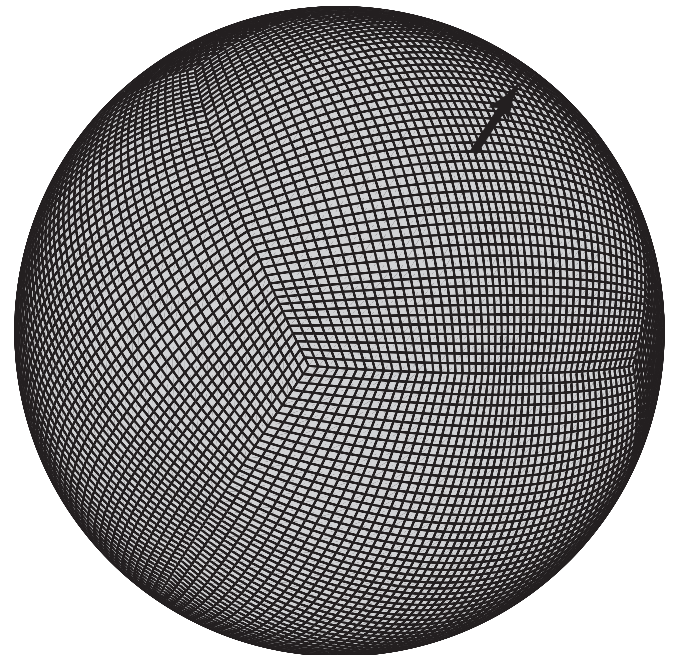

FIG. 8. Illustration of the finite and boundary element meshes for the sphere and the drive and response point.
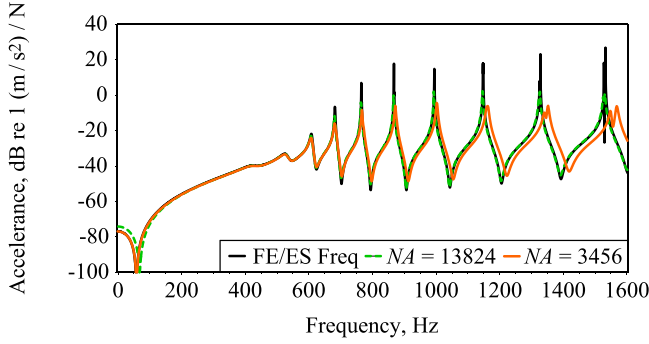

FIG. 6. (Color online) Normalized sound pressure level as a function of frequency for the bronze plate.

FIG. 9. (Color online) Acceleration as a function of frequency for the steel sphere.

FIG. 10. Illustration of the cylinder mesh and the drive point location.

The results confirm that acoustic radiation damping, which is very strong for the first few resonances, is included correctly in the analysis. For this particular problem, the tripole source formulation does not become unstable for small time step sizes (the smallest value tested was 1×10^{-6} s, corresponding to $c\Delta t/L = 0.015$). This may be a consequence of the large radiation damping for the lowest axisymmetric modes. The damping levels are higher for the transient solution, which is a well-known characteristic of time-stepping procedures.

IV. EXAMPLE WITH TERMS OF THE COUPLING MATRIX SET TO ZERO

It is not always possible or desirable to choose the time step small enough to make $P(1)$ and $U(1)$ fully diagonal, especially if stability is an issue for small step sizes or the boundary surface mesh contains a few small elements. In this section, the possibility of making the matrix A sparse by setting terms of $P_t(1)[U_t(1)]^{-1}$ to zero will be examined. It is difficult to precisely calculate which matrix terms should be nonzero due to the noncausal interpolation functions. Rather than trying to determine which terms should theoretically be zero based on the boundary surface mesh geometry, the time step size, and sound speed, a simple ratio of the matrix terms can be used instead, as defined as

$$\text{Ratio} = |A_{\mu\nu}|/\|A\|_{\max}. \qquad (24)$$

If this ratio is less than a specified value, x_{filter}, the matrix element is set zero. Varying x_{filter} gives a trade-off between accuracy and solution time, where the number of terms set to zero decreases as x_{filter} decreases. The time step size also influences solution accuracy, so time step variations will be considered first.

A thin circular steel cylinder, the mesh for which is shown in Fig. 10, will be used as an example to test the algorithm when terms of $P(1)$ and $U(1)$ are set to zero. The cylinder has outer diameter 0.169 m, length 0.427 m, and thickness 0.0064 m. The material properties for the cylinder are taken as Young's modulus 190 GPa, Poisson's ratio 0.3, and density 7830 kg/m³. The drive and response points are in the radial direction on the outside of the cylinder a distance 0.0427 m from the closest edge. The peak measured input force is 2.5 N and its duration is approximately 2.73×10^{-4} s.

The input force for the cylinder had to be much lower than for the bronze plate considered previously to keep the accelerometer from overloading since the drive and response points are in close proximity and the cylinder is not nearly as thick or heavy as the plate.

The cylinder is a good problem to illustrate the sparsity of A because, depending on the time step size, it is possible for a triangular pulse of length $2\Delta t$ to travel across the thickness in a single time step. The minimum and maximum element dimensions for the surface mesh are 0.0021 and 0.0037 m, respectively. The time step size must be less than 7.1×10^{-7} s for $P(1)$ and $U(1)$ to be fully diagonal, which is impractically small. A range of values for the time step size will be considered to illustrate how it affects the solution accuracy and computation time. The stability parameter $c\Delta t/L$ is equal to 2, 1.6, 0.8, and 0.4 for time step sizes 5×10^{-6}, 4×10^{-6}, 2×10^{-6}, and 1×10^{-6} s, respectively. In the following results, the parameter x_{filter} is set to 1×10^{-6}, which corresponds to a difference in magnitude of 120 dB. Table I lists the percentage of nonzero terms in K, A, $P_t(1)$, $P_t(1)[U_t(1)]^{-1}$, and K_{eff} for various time step sizes along with the maximum number of active time steps, $NATS_{\max}$, for any source-receiver combination. The total number of terms is $(NDOF)^2$ for K, A, and K_{eff} and $(NA)^2$ for $P_t(1)$ and $P_t(1)[U_t(1)]^{-1}$, where $NDOF = 209\,088$ and $NA = 35\,424$ for the cylinder. The number of active time steps grows as the time step size decreases because more time steps are required for a triangular pulse to travel across the boundary surface elements. The number of nonzero terms in A more than doubles in changing the time step from 4×10^{-6} s to 5×10^{-6} s because the shell thickness is 0.0064 m and sound waves travel across the thickness in 4.3×10^{-6} s. Also,

TABLE I. The sparseness of the matrices as a function of time step size with x_{filter} set to 1×10^{-6} s. The matrix $U_t(1)$ has the same sparsity pattern as $P_t(1)$.

Δt	$NATS_{\max}$	K	$P_t(1)$	$P_t(1)[U_t(1)]^{-1}$	A	K_{eff}
				% nonzero		
5×10^{-6} s	3	0.032	0.091	0.962	0.149	0.177
4×10^{-6} s	3	0.032	0.060	0.457	0.073	0.102
2×10^{-6} s	4	0.032	0.025	0.150	0.030	0.059
1×10^{-6} s	6	0.032	0.003	0.003	0.003	0.032

J. Acoust. Soc. Am. **142** (2), August 2017

John B. Fahnline and Micah R. Shepherd   1019

when the time step is equal to 1×10^{-6} s, the sparsity of $\boldsymbol{K}$ and $\boldsymbol{K}_{\text{eff}}$ are nearly the same, indicating that the acoustic analysis does not add many nonzero terms into $\boldsymbol{K}_{\text{eff}}$.

Changes in the number of active time steps and the sparsity of the matrices have large effects on the computation times. Table II summarizes the computation times for 1000 time steps using the various time step sizes, where the calculations were performed on a Linux cluster using four nodes, each with with 48 processors and 256 GB of RAM. The computation of the source amplitude vector and its time derivative require convolution summations, but these data are already computed and not included as part of the column labeled $\boldsymbol{q}, \dot{\boldsymbol{q}}$. Several trends are evident in the table. The times to compute the ES matrices and the convolution summations increase with the number of active time steps simply because there are more nonzero terms to calculate. The times to compute $\boldsymbol{P}(1)[\boldsymbol{U}(1)]^{-1}$ remain essentially constant because the calculations are performed using subroutines that assume the matrices are fully populated, without taking advantage of their sparsity. As discussed previously in the section on computational strategies, the coupled equation system is solved using only one computational node because inter-node communication can slow MUMPS considerably. The factorization of the coupled equation system only occurs during the first time step and occupies a small fraction of the time, especially as the number of time steps grows. The data in the table indicate that the solution times for the coupled equation system are a strong function of the number of nonzero terms in the effective stiffness matrix.

The times required to compute $\boldsymbol{q}$ and $\dot{\boldsymbol{q}}$ warrant some discussion. Reducing the time step size makes the matrix $\boldsymbol{U}(1)$ sparser, such that the times to compute $\boldsymbol{q}$ and $\dot{\boldsymbol{q}}$ should decrease along with the time step size. However, the MUMPS solver (version 5.0.0) actually becomes much slower for very sparse matrices. This issue is discussed on the MUMPS website as a FAQ topic. The website suggests that this difficulty has to do with retaining very small numbers in the solution. As a consequence, PARDISO far outperforms MUMPS for very sparse matrices. In the data listed in the table, the solver switches from MUMPS to PARDISO for the time step sizes 1×10^{-6} s and 2×10^{-6} s. Overall, the results show that the user can trade-off computation times between the convolution summations and the matrix solution for the coupled system somewhat by varying the time step size.

Aside from computation time, the time step size also influences solution accuracy. Figure 11 shows the drive point acceleration for the cylinder as a function of time for step sizes 1×10^{-6} and 4×10^{-6} s. The results in the figure show
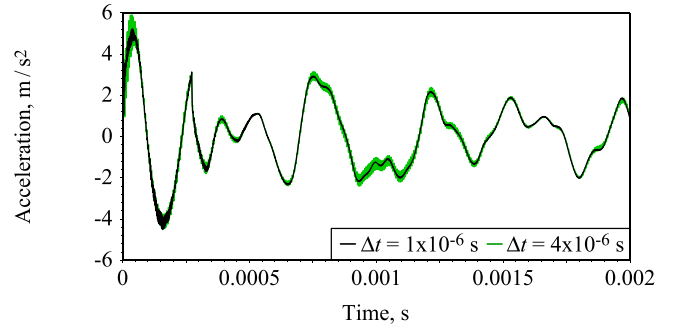

FIG. 11. (Color online) Acceleration for the thin cylinder as a function of time.

that refining the time step size reduces high frequency oscillations in the response. The best results are achieved with Δt set to 1×10^{-6} s, for which the matrices $\boldsymbol{P}(1)$ and $\boldsymbol{U}(1)$ are nearly diagonal. Unfortunately, for this time step size, the solution becomes unstable after approximately 0.011 s. The subsequent computations are thus performed with Δt set to 2×10^{-6} s. Ultimately, since the time required for the convolution summations scales well with number of processors, it is generally beneficial to reduce the time step size as small as possible to increase the solution accuracy and to reduce the number of nonzero terms in $\boldsymbol{P}(1)$ and $\boldsymbol{U}(1)$.

In varying the time step size, x_{filter} was set to 1×10^{-6}, which corresponds to a difference in magnitude of 120 dB. The trade-off between solution accuracy and computation time as a function of x_{filter} will now be considered with the time step fixed. Zeroing out more terms of the acoustic coupling matrix reduces the time required to solve the coupled equation, but may negatively impact the solution accuracy. The goal will be to determine the maximum acceptable value that can be used for x_{filter} without causing significant errors. Figure 12 shows the drive point acceleration for the cylinder as a function of time for various values of x_{filter} with Δt equal to 2×10^{-6} s. The figure shows that accurate results are obtained for this problem with $x_{\text{filter}} < 1 \times 10^{-3}$. The small shift in the data for $x_{\text{filter}} = 1 \times 10^{-3}$ slowly accumulates over time, so that a noticeable shift occurs on the scale of 100 000 time steps. Ultimately, this corresponds to a slight frequency shift in the resonance peaks. Table III lists the percentage of nonzero terms in $\boldsymbol{K}$, $\boldsymbol{A}$, $\boldsymbol{P}_t(1)$, $\boldsymbol{P}_t(1)[\boldsymbol{U}_t(1)]^{-1}$, and $\boldsymbol{K}_{\text{eff}}$ for the various values of x_{filter} along with the times to factor the coupled equation system and solve it for 1000

TABLE II. Computation times as a function of the time step size, where M-V stands for matrix-vector multiplication.

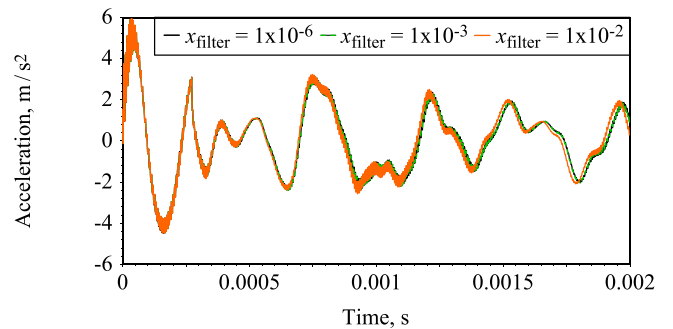| Δt | $\boldsymbol{P}, \boldsymbol{U}$ | $\boldsymbol{P}_t(1)[\boldsymbol{U}_t(1)]^{-1}$ | Coupled system | | | cs_t | $\boldsymbol{q}, \dot{\boldsymbol{q}}$ |
			Factor	Solve	M-V		
5×10^{-6} s	1029.5 s	246.6 s	104.6 s	219.9 s	225.9 s	412.8 s	11.8 s
4×10^{-6} s	1079.5 s	246.2 s	69.0 s	184.0 s	141.2 s	404.8 s	12.5 s
2×10^{-6} s	1297.0 s	246.2 s	32.9 s	123.6 s	97.9 s	384.3 s	19.4 s
1×10^{-6} s	1754.5 s	247.6 s	7.2 s	55.3 s	72.6 s	515.6 s	22.3 s

FIG. 12. (Color online) Acceleration for the thin cylinder as a function of time for various values for x_{filter}.

John B. Fahnline and Micah R. Shepherd

TABLE III. Computation times for the coupled matrix solutions and sparseness of the matrices as a function of time step size with Δt equal to 2×10^{-6} s. The matrix $U_t(1)$ has the same sparsity pattern as $P_t(1)$.

	Solution time		% nonzero				
x_{filter}	Factor	Solve	K	$P_t(1)$	$P_t(1)[U_t(1)]^{-1}$	A	K_{eff}
1×10^{-6}	32.9 s	123.6 s	0.032	0.025	0.150	0.030	0.059
1×10^{-5}	24.1 s	111.2 s	0.032	0.025	0.105	0.023	0.051
1×10^{-4}	19.7 s	91.8 s	0.032	0.025	0.062	0.016	0.045
1×10^{-3}	17.4 s	88.4 s	0.032	0.025	0.037	0.012	0.041
1×10^{-2}	11.9 s	67.0 s	0.032	0.025	0.014	0.007	0.036
1×10^{-1}	6.9 s	57.4 s	0.032	0.025	0.003	0.003	0.032

time steps with the time step size set to 2×10^{-6} s. The results again show the strong dependence of the matrix factorization and solution times on the number of nonzero terms in the effective stiffness matrix. Also, the differences in the matrix solution times for x_{filter} equal to 1×10^{-3} and 1×10^{-4} are fairly small, so that in practice is better to choose the more accurate solution. Clearly, $P_t(1)[U_t(1)]^{-1}$ should not be more sparse than $P_t(1)$ or $U_t(1)$, individually, so that x_{filter} should be taken smaller than 1×10^{-2}. To illustrate the number of active elements, Fig. 13 shows the elements with nonzero pressure when one acoustic element on the outer surface of the cylinder is vibrating for four different choices of x_{filter}. The illustration demonstrates that even with $x_{filter} = 1 \times 10^{-6}$, which retains many more terms, the pressure is very nearly zero over almost all of the boundary surface. In all the problems considered by the authors to date, choosing x_{filter} equal to 1×10^{-4} yields accurate results. However, the accuracy

can always be confirmed by reducing x_{filter} and checking convergence.

As a final step, the time domain computations will be compared to frequency domain computations and experimental measurements. With the time step chosen as 2×10^{-6} s and with x_{filter} set to 1×10^{-4}, Fig. 14 shows comparisons of the drive point acceleration as a function of time computed using simple sources, tripole sources, along with experimental measurements. For the full analysis 524 288 time steps were computed, giving a frequency resolution of approximately 1 Hz. Figure 15 shows comparisons of the experimental measurements and numerical computations for the acceleration-to-force transfer functions at the drive point as a function of frequency. The experimental measurements exhibit considerably higher damping than either the time or frequency domain FE/ES results, which is likely due to friction from contact between the cylinder edges and the bungee cords used to suspend it. As for the thin sphere, the time domain computations show some extra algorithmic damping in comparison to the frequency domain computations.

V. COMPUTATION TIME SCALING

The goal in this section is to test how well the computation times scale as the number of processors increases for the cylinder mesh discussed in Sec. IV. Table IV lists the significant computation times for 1000 time steps using the tripole source formulation with $x_{filter} = 1 \times 10^{-4}$ and a time step size of $\Delta t = 2 \times 10^{-6}$ s, where the calculations were performed on a Linux cluster with nodes that each have 36 processors and 128 GB of RAM. Other jobs were running while the
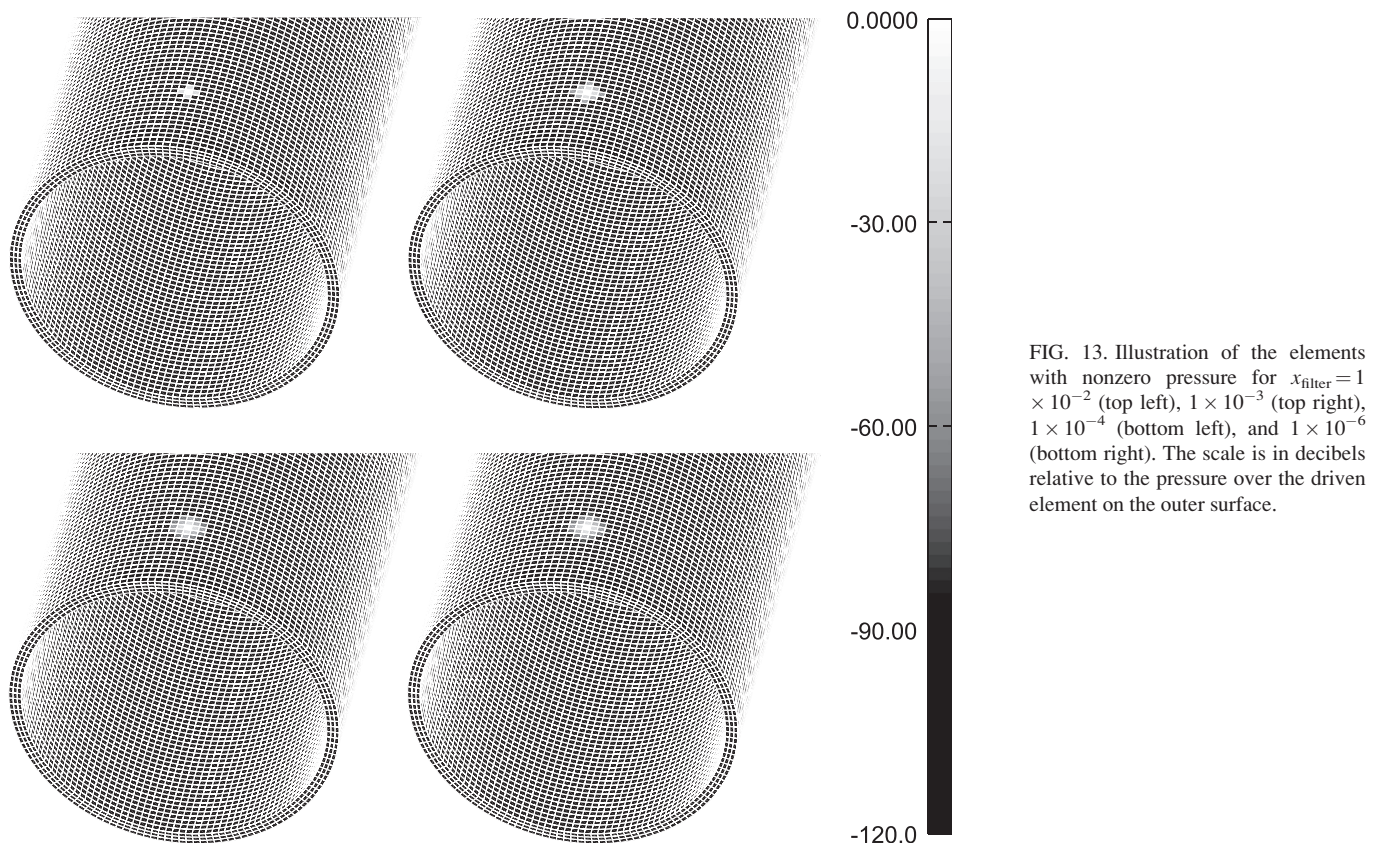

FIG. 13. Illustration of the elements with nonzero pressure for $x_{filter} = 1 \times 10^{-2}$ (top left), 1×10^{-3} (top right), 1×10^{-4} (bottom left), and 1×10^{-6} (bottom right). The scale is in decibels relative to the pressure over the driven element on the outer surface.
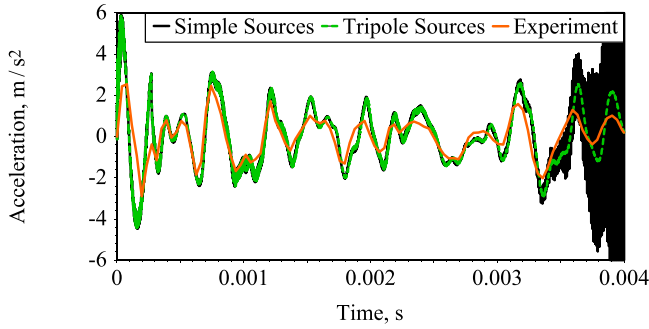
FIG. 14. (Color online) Acceleration as a function of time for the thin cylinder.

TABLE IV. Computation times as a function of the number of processors, NP, where M-V stands for matrix-vector multiplication.

			Coupled system				
NP	P, U	$P_t(1)[U_t(1)]^{-1}$	Factor	Solve	M-V	cs_t	$q, \dot{q}$
360	305.0 s	52.4 s	11.2 s	53.6 s	45.0 s	156.1 s	7.1 s
720	161.8 s	33.0 s	11.9 s	53.6 s	45.1 s	88.0 s	7.0 s
1080	109.1 s	28.8 s	11.3 s	53.2 s	45.4 s	59.4 s	6.9 s
1440	107.4 s	30.0 s	11.7 s	53.5 s	45.7 s	52.7 s	7.4 s
1800	84.7 s	51.2 s	11.5 s	53.1 s	46.0 s	46.8 s	7.4 s

computations were being performed, so that some small variation exists due to differences in the communication times. The time required to compute A is not listed because given a sparse version of $P(1)U(1)^{-1}$, the subsequent multiplications by F and V are fast. The results show that, for a problem of this size, even though the acoustic matrices are processed as if they are densely populated, the times required to compute $U(1)^{-1}$ and $P(1)U(1)^{-1}$ are reasonably short because the SCALAPACK routines scale well with number of processors and the 35 424 acoustic elements are far less than the 209 088 total number of degrees-of-freedom. In general, the expectation is that the computation times for $P(1)U(1)^{-1}$ will be reduced as the number of processors increases, but the times actually increase using 1440 and 1800 processors. This is likely due to either increases in the amount of inter-node communication or the processor grids and blocking factors for the SCALAPACK computations. The sparse matrix computations, including factoring and solving the coupled system and the matrix-vector multiplications, are performed on a single multiprocessor node and remain essentially constant. The other steps scale well with number of processors, such that the computation time decreases almost in direct proportion. The analysis steps in the last four columns are repeated for every time step, and thus their percentage of the total computation time will increase as the number of time steps increase. The important conclusion is that given enough processors, the time to compute the convolution summations can be made as small as the time to compute the solution of the coupled equation system.

For very large problems, the $P(1)U(1)^{-1}$ computations become much more time consuming due to the cubic dependence on matrix size for densely populated matrix solutions
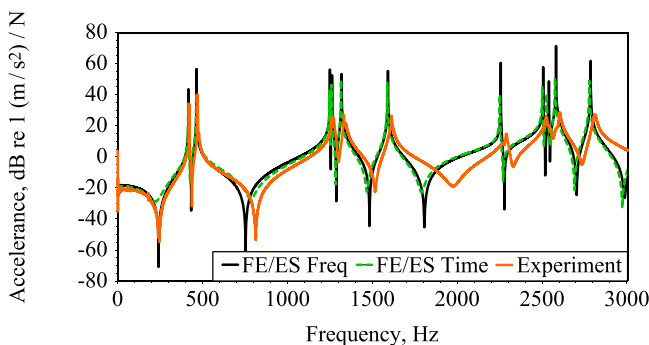

FIG. 15. (Color online) Acceleration as a function of frequency for the thin cylinder.

and multiplications. However, in a practical situation, the processor count will also have to increase considerably due to the memory requirements for storing the P and U matrices in RAM. To illustrate the memory requirements, the tripole source formulation stores four matrices of size $NATS \times (NA)^2 \times 4$ bytes, where 4 bytes represents the memory required to store a single matrix element in single precision and $NATS$ is the number of active time steps. The largest problem the authors have solved to date has $NA = 400\,000$ and $NATS = 8$, such that approximately 20 TB of RAM was required to store the matrices. The analysis was performed on a Linux cluster using 300 nodes, each with 36 processors and 128 GB RAM, giving 10 800 processors and 38.4 TB of RAM. For this problem, the matrix factorization and solutions required 70 and 14 min, respectively. Thus, large problems can be addressed using the formulation, but only with vast computing resources. For reference purposes, the analysis was run on the Air Force Research Laboratory cluster Thunder, and represented approximately a tenth of the cluster's total resources. Of course, this discussion becomes irrelevant when the time step can be chosen small enough to make $P(1)$ and $U(1)$ fully diagonal, in which case the computation of $P(1)U(1)^{-1}$ is trivial. When the time step size is small enough for some, but not all, of the rows and columns of $P(1)$ and $U(1)$ to be diagonal, it is possible to partition the matrices into block diagonal form. Matrix inversion and multiplication is then trivial for the diagonal part of the matrices, and the full calculation is performed for the rest, potentially reducing the computation times considerably.

VI. DISCUSSION OF STABILITY

In the previous analyses of the bronze plate and thin steel cylinder, instabilities occur for small values of $c\Delta t/L$. These boundary surfaces can be categorized as "thin shapes," where the thickness in one direction is much smaller than the others. It is likely that the instabilities are due to "thin-shape breakdown" difficulties, as discussed by Hargreaves and Cox[23] for transient boundary element analyses. Since the length-to-thickness ratio is smaller for the cylinder (0.015) than for the plate (0.058), numerical solutions for the cylinder are expected to be less stable than those for the plate, and, indeed, the cylinder becomes unstable for $c\Delta t/L \approx 0.4$ while the plate becomes unstable for $c\Delta t/L \approx 0.08$. Ultimately, these types of problems are more stable when analyzed using a "dipole source formulation," where the thickness direction is represented as infinitesimally small and it is assumed that plate and

John B. Fahnline and Micah R. Shepherd

shell elements deform only in bending. A discussion of equivalent source calculations using dipole sources alone is given in the paper by Fahnline,[5] and it can be coupled to finite element analyses to give a coupled FE/ES formulation using the same techniques discussed in Sec. II of this paper. It is also possible to combine mixed regular and thin bodies together in a single analysis using tripole and dipole sources to represent the regular and thin bodies, respectively. A detailed discussion of these formulations is not included here due to space limitations.

VII. CONCLUSIONS

In this paper, a monolithic time-marching scheme for structural-acoustic problems has been derived where the acoustic coupling matrix is either fully diagonal or treated as sparse. This basic idea originated with the realization that the coupling matrix should be sparse in theory because it relates nodal velocities to nodal acoustic pressure forces, and the pressure waves can only travel a distance $c\Delta t$ during a single time step. Computationally, the matrix becomes fully populated due to small errors in the coefficient matrix, the matrix solution, and the subsequent matrix multiplications. To make the matrix sparse, terms are set to zero that are smaller than a specified fraction of the largest term. The numerical results demonstrate that this is a reasonable approximation, and that for the chosen example problem accurate results are obtained with the fraction set to 1×10^{-4}, which is equivalent to neglecting terms 80 dB smaller than largest term. It is also demonstrated that the formulation adapts well to parallel processing environments and that the computation times associated with the equivalent source computations are proportional to the number of processors. Several numerical examples are solved to illustrate the computations and validate the formulation.

ACKNOWLEDGMENTS

APPENDIX

In this appendix, the matrix terms derived from the equivalent source analysis are explicitly defined. The actual method for performing the integration is discussed in the paper by Fahnline.[5] The source amplitudes are assumed as triangular pulses of length two time steps Δt starting at time $t = 0$ or, in mathematical terms,

$$T(t) = \left(1 - \frac{|\Delta t - t|}{\Delta t}\right)[H(t) - H(t - 2\Delta t)], \quad (A1)$$

where H is the Heaviside function. The subsequent equations are written for a source at point q, and the surface integrations are performed for the field point x located on an element with surface area S, which represents one term of a matrix as a function of time. Other variables include the outward unit normal direction n, the unit direction for the dipole source n_q, the vector distance from the source to the field point $R = x - q$, and its magnitude $R = |x - q|$.

The velocity potential of a simple source can be written as

$$\Phi(x, t) = -Q(\tau)/4\pi R, \quad (A2)$$

where Q is the simple source amplitude and $\tau = t - R/c$. The pressure field of a simple source can be computed from Eq. (A2) as

$$p(x, t) = -\rho \frac{\partial \Phi}{\partial t} = \frac{\rho}{4\pi R} \frac{\partial Q(\tau)}{\partial t} = -\frac{\rho c}{4\pi R} \frac{\partial Q(\tau)}{\partial R}, \quad (A3)$$

where ρ is the fluid density. Evaluating the average pressure over surface element α at time t due to a pulse $T(t)$ for source κ gives one element of the matrix $P_s(t)$ as

$$P_s(t) = -\frac{\rho c}{4\pi S_\alpha} \int\int_{S_\alpha} \left\{\frac{1}{R} \frac{\partial T(\tau)}{\partial R}\right\}_{R=R_\kappa} dS(x), \quad (A4)$$

where the subscript s indicates simple source, the row number is α, the column number is κ, and $R_\kappa = |x - q_\kappa|$ is the distance from source κ. The discrete matrices $P_s(n)$ are computed by sampling $P_s(t)$ at $t = n\Delta t$. The component of the acoustic particle velocity in the direction n can be computed from Eq. (A1) using Euler's equation as

$$v_n(x, t) = (\nabla R \cdot n) \frac{\partial}{\partial R}\left[-\frac{Q(\tau)}{4\pi R}\right] = -\frac{R \cdot n}{4\pi R} \frac{\partial}{\partial R}\left[\frac{Q(\tau)}{R}\right]. \quad (A5)$$

Evaluating the volume velocity over surface element α at time t due to a pulse $T(t)$ for source κ gives one element of the matrix $U_s(t)$ as

$$U_s(t) = -\frac{1}{4\pi} \int\int_{S_\alpha} \left\{\frac{R \cdot n}{R} \frac{\partial}{\partial R}\left[\frac{T(\tau)}{R}\right]\right\}_{R=R_\kappa} dS(x). \quad (A6)$$

The velocity potential of a dipole source can be written as

$$\Phi(x, t) = \frac{R \cdot n_q}{4\pi R} \frac{\partial}{\partial R}\left[\frac{F(\tau)}{R}\right], \quad (A7)$$

where F is the dipole source amplitude (which has units of force), and n_q is a unit vector indicating the source's direction. The pressure field of a dipole source can be computed from the velocity potential as

$$p(x, t) = -\rho \frac{R \cdot n_q}{4\pi R} \frac{\partial}{\partial R}\left[\frac{\dot{F}(\tau)}{R}\right], \quad (A8)$$

where $\dot{F}$ is the derivative of F with respect to time. One element of the matrix $P_d(t)$ for element α and source κ is given as

$$P_d(t) = -\frac{\rho}{4\pi S_\alpha} \int\int_{S_\alpha} \left\{\frac{R \cdot n_{q_\kappa}}{R} \frac{\partial}{\partial R}\left[\frac{T(\tau)}{R}\right]\right\}_{R=R_\kappa} dS(x), \quad (A9)$$

where n_{q_κ} is the direction for dipole source κ. The component of the acoustic particle velocity in the direction n can be computed as

$$v_n(\boldsymbol{x}, t) = \frac{\boldsymbol{n} \cdot \boldsymbol{n}_q}{4\pi R} \frac{\partial}{\partial R} \left[\frac{F(\tau)}{R} \right]$$

$$+ \frac{\boldsymbol{R} \cdot \boldsymbol{n}_q}{4\pi} \frac{\boldsymbol{R} \cdot \boldsymbol{n}}{R} \frac{\partial}{\partial R} \left\{ \frac{1}{R} \frac{\partial}{\partial R} \left[\frac{F(\tau)}{R} \right] \right\}$$

$$= \frac{\boldsymbol{n} \cdot \boldsymbol{n}_q}{4\pi R} \frac{\partial}{\partial R} \left[\frac{F(\tau)}{R} \right]$$

$$- \frac{\boldsymbol{R} \cdot \boldsymbol{n}_q}{4\pi} \frac{\boldsymbol{R} \cdot \boldsymbol{n}}{R} \frac{\partial}{\partial R} \left[\frac{F(\tau)}{R^3} + \frac{\dot{F}(\tau)}{cR^2} \right]. \qquad \text{(A10)}$$

The volume velocity of a dipole source depends on both F and $\dot{F}$, which are given subscripts 0 and 1, respectively. One element of the matrices $\boldsymbol{U}_{d,0}(t)$ and $\boldsymbol{U}_{d,1}(t)$ for element α and source κ is given as

$$U_{d,0}(t) = \frac{1}{4\pi} \iint_{S_\alpha} \left\{ \frac{\boldsymbol{n} \cdot \boldsymbol{n}_{q_\kappa}}{R} \frac{\partial}{\partial R} \left[\frac{T(\tau)}{R} \right] \right.$$

$$\left. - (\boldsymbol{R} \cdot \boldsymbol{n}_{q_\kappa}) \frac{\boldsymbol{R} \cdot \boldsymbol{n}}{R} \frac{\partial}{\partial R} \left[\frac{T(\tau)}{R^3} \right] \right\}_{R=R_\kappa} \mathrm{d}S(\boldsymbol{x})$$

$$\text{(A11)}$$

and

$$U_{d,1}(t) = -\frac{1}{4\pi c} \iint_{S_\alpha} \left\{ (\boldsymbol{R} \cdot \boldsymbol{n}_{q_\kappa}) \frac{\boldsymbol{R} \cdot \boldsymbol{n}}{R} \frac{\partial}{\partial R} \left[\frac{T(\tau)}{R^2} \right] \right\}_{R=R_\kappa} \mathrm{d}S(\boldsymbol{x}),$$

$$\text{(A12)}$$

respectively.

For the tripole source formulation, one element of the matrix $\boldsymbol{P}_t(t)$ is given as

$$P_t(t) = \frac{1}{c} P_s(t) + P_d(t), \qquad \text{(A13)}$$

and one element of the matrix $\boldsymbol{U}_t(t)$ is given as

$$U_t(t) = \frac{1}{c} U_s(t) + \frac{2\Delta t}{3} U_{d,0}(t) + U_{d,1}(t). \qquad \text{(A14)}$$

[1] R. J. Astley, "Infinite elements for wave problems: A review of current formulations and an assessment of accuracy," Int. J. Numer. Meth. Eng. 49(7), 951–976 (2000).

[2] M. Zampolli, A. Tesei, F. B. Jensen, N. Malm, and J. B. Blottman III, "A computationally efficient finite element model with perfectly matched layers applied to scattering from axially symmetric objects," J. Acoust. Soc. Am. 122(3), 1472–1485 (2007).

[3] H. A. Schenck, "Improved integral formulation for acoustic radiation problems," J. Acoust. Soc. Am. 44(1), 41–58 (1968).

[4] A. J. Burton and G. F. Miller, "The application of integral equation methods to the numerical solution of some exterior boundary-value problems," Proc. R. Soc. London, Ser. A: Math. Phys. Eng. Sci. 323, 201–210 (1971).

[5] J. B. Fahnline, "Solving transient acoustic boundary value problems with equivalent sources using a lumped parameter approach," J. Acoust. Soc. Am. 140(6), 4115–4129 (2016).

[6] D. Soares, "Acoustic modeling by BEM–FEM coupling procedures taking into account explicit and implicit multi-domain decomposition techniques," Int. J. Numer. Meth. Eng. 78(9), 1076–1093 (2009).

[7] G. Yu, W. J. Mansur, J. A. M. Carrer, and L. Gong, "Time weighting in time domain BEM," Eng. Anal. Bound. Elem. 22(3), 175–181 (1998).

[8] A. Frangi, "Causal shape functions in the time domain boundary element method," Comput. Mech. 25(6), 533–541 (2000).

[9] C. A. Felippa, K. C. Park, and C. Farhat, "Partitioned analysis of coupled mechanical systems," Comput. Methods Appl. Mech. Eng. 190(24), 3247–3270 (2001).

[10] D. T. Wilton, "Acoustic radiation and scattering from elastic structures," Int. J. Numer. Meth. Eng. 13(1), 123–138 (1978).

[11] K. Blakely, MSC/NASTRAN Basic Dynamic Analysis: User's Guide, Version 68 (MacNeal-Schwendler, 1993), pp. 154–155.

[12] R. D. Cook, D. S. Malkus, M. E. Plesha, and R. J. Witt, Concepts and Applications of Finite Element Analysis, 4th ed. (Wiley, New York, 2001), pp. 416–421.

[13] S. P. Walker and C. Y. Yeung, "Parallel computation of integral equation methods for three-dimensional transient wave propagation," Commun. Numer. Meth. Eng. 11(6), 515–524 (1995).

[14] A. A. Ergin, B. Shanker, and E. Michielssen, "Fast analysis of transient acoustic wave scattering from rigid bodies using the multilevel plane wave time domain algorithm," J. Acoust. Soc. Am. 107(3), 1168–1178 (2000).

[15] O. Schenk and K. Gärtner, "Solving unsymmetric sparse systems of linear equations with Pardiso," Future Gen. Comput. Syst. 20(3), 475–487 (2004).

[16] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster, "A fully asynchronous multifrontal solver using distributed dynamic scheduling," SIAM J. Matrix Anal. Appl. 23(1), 15–41 (2001).

[17] X. S. Li and J. W. Demmel, "Superlu_dist: A scalable distributed-memory sparse direct solver for unsymmetric linear systems," ACM Trans. Math. Softw. (TOMS) 29(2), 110–140 (2003).

[18] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, and A. Petitet, ScaLAPACK Users' Guide (SIAM, Philadelphia, PA, 1997), Vol. 4.

[19] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and X. S. Li, "Analysis and comparison of two general sparse solvers for distributed memory computers," ACM Trans. Math. Softw. (TOMS) 27(4), 388–421 (2001).

[20] A. Akay and M. Latcha, "Sound radiation from an impact-excited clamped circular plate in an infinite baffle," J. Acoust. Soc. Am. 74(2), 640–648 (1983).

[21] J. B. Fahnline, "Boundary-element analysis," in Engineering Vibroacoustic Analysis: Methods and Applications (Wiley Online Library, New York, 2016), pp. 179–229.

[22] M. Stutz and M. Ochmann, "Stability behavior and results of a transient boundary element method for exterior radiation problems," J. Acoust. Soc. Am. 123(5), 3530 (2008).

[23] J. A. Hargreaves and T. Cox, "A transient boundary element method for acoustic scattering from mixed regular and thin rigid bodies," Acta Acust. Acust. 95(4), 678–689 (2009).

John B. Fahnline and Micah R. Shepherd