UvA-DARE (Digital Academic Repository)

Simulation of gravitational wave detectors

de Ronde, J.F.; van Albada, G.D.; Sloot, P.M.A.

DOI

[10.1063/1.168591](https://doi.org/10.1063/1.168591)

Publication date

1997

Published in

Computers in Physics

[Link to publication](#)

Citation for published version (APA):

de Ronde, J. F., van Albada, G. D., & Sloot, P. M. A. (1997). Simulation of gravitational wave detectors. *Computers in Physics*, 11(5), 484-497. <https://doi.org/10.1063/1.168591>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Simulation of gravitational wave detectors

J. F. de Ronde,^{a)} G. D. van Albada, and P. M. A. Sloot

*Faculty of Mathematics, Computer Science, Physics, and Astronomy, University of Amsterdam,
1098 SJ Amsterdam, The Netherlands*

(Received 17 January 1997; accepted 11 June 1997)

A simulation program that provides insight into the vibrational properties of resonant mass gravitational radiation antennas is developed from scratch. The requirements that are set necessitate the use of an explicit finite element kernel. Since the computational complexity of this kernel requires significant computing power, it is tailored for execution on parallel computer systems. After validating the physical correctness of the program as well as the performance on distributed memory architectures, we present a number of “sample” simulation experiments to illustrate the simulation capabilities of the program. The development path of the code, consisting of problem definition, mathematical modeling, choosing an appropriate solution method, parallelization, physical validation, and performance validation, is argued to be typical for the design process of large-scale complex simulation codes. © 1997 American Institute of Physics. [S0894-1866(97)01805-1]

INTRODUCTION

General background

Although the existence of gravitational radiation, predicted by Einstein’s theory of general relativity,¹ is unquestioned, its detection is a long-standing problem in experimental physics. The aim of the GRAIL project² is to realize a spherical resonant mass detector with a sensitivity that is a few orders of magnitude higher than the present generation of detectors, thus offering the possibility of validating the existence of gravitational radiation that is emitted by astrophysical sources.³ These sources include supernovae, stellar collapses to black hole states, and coalescence of binary neutron star systems.⁴

The quadrupole moment of a mass distribution $\rho(\mathbf{x})$ is given by

$$D_{ij} = \int_V dV \rho(\mathbf{x}) \left(x_i x_j - \frac{1}{3} \delta_{ij} \mathbf{x}^2 \right). \quad (1)$$

According to general relativity (GR), an oscillation of this quadrupole moment is the simplest mode of vibration that can generate gravitational waves. The counterstatement for absorption of gravitational waves is that in the simplest case it also takes place via the excitation of the quadrupole modes of vibration of a massive object. This principle is the main argument for constructing a gravitational radiation antenna, which essentially is a large resonant mass that is forced into oscillation by impinging gravitational radiation.

In GRAIL, the antenna will consist of a spherical resonant mass of a Cu alloy, possibly CuAl, with a mass of about 100,000 kg and a diameter of 3 m. It will be suspended in vacuum inside a large cryostat in such a way that the external vibrations at its resonant frequency

(≈ 700 Hz) are attenuated by at least a factor of 10^{16} . The sphere will be cooled to a temperature in the range of 10–20 mK in order to cancel out thermal noise. Transducers will be attached to the surface of the sphere to detect the vibrations induced by gravitational radiation. An analysis of the electric signals from these transducers must give information (such as source type and direction of incidence) about the gravitational radiation that interacts with the sphere. The gravitational waves have extremely weak interaction with matter; the typical deformations induced on a Cu sphere of 3 m will be of the order of 10^{-20} m, which explains why detection is such a difficult task and why it has not yet been realized. Several arguments can be given for using a spherically shaped antenna^{5–8} instead of, for example, the cylinder or bar antennas in the pioneering days of gravitational radiation detection⁹ and present-day experimental setups like the NAUTILUS detector at the Frascati National Laboratory in Rome.¹⁰

Since the estimated cost (≈ 45 million Dutch guilders) of the project is quite high, a pilot study has been funded by NWO¹¹ to investigate the technological feasibility of GRAIL. In this project several critical aspects concerning the design of the antenna had to be addressed. The task of our simulation group was to construct detailed models of the antenna that provide insight into its vibrational properties. A schematic design of the GRAIL antenna concept is shown in Fig. 1.

The development and design of complex systems can be significantly enhanced by means of simulation. A recent example in which simulation was used in the analysis of the behavior of resonant mass antennas can be found in Ref. 8. In this work modal analysis was used to determine the frequency spectrum of a truncated icosahedral antenna. The accuracy of this simulation was fairly limited due to the low model resolution. The main result was that one can

^{a)}E-mail: janr@wins.uva.nl

Figure 1. The GRAIL antenna in a heat isolating vacuum chamber.

conclude that the behavior of a truncated icosahedral antenna is very similar to that of a sphere.

In this article we will describe the development of a high-resolution simulation system for resonant mass gravitational radiation antennas that can be applied to a number of important design issues.

Design issues

The following questions need to be addressed by means of extensive simulation experiments.

- (1) How does the suspension of the huge resonant mass affect the eigenmodes and eigenfrequency of the fundamental quadrupole modes of a sphere?
- (2) How will material inhomogeneities affect these eigenmodes and eigenfrequencies?
- (3) Given a perfect sphere, how will it deform under its own weight?
- (4) Does the coupling between the suspension rod and the sphere induce additional modes that will cause interference with the modes we are looking for?
- (5) What is the effect of mounting transducers on the antenna on the frequency spectrum of the system?
- (6) How will seismic noise, entering the sphere through the framework in which it is suspended, influence the transducer signals?
- (7) How will the fingerprint of typical gravitational wave sources be seen by the antenna?
- (8) Can energy that has been deposited by cosmic rays or particles like muons induce system eigenmodes that are not distinguishable from those induced by gravitational waves?

We can identify two types of questions in this list. First, the design questions that need to be answered in order to tune the readout of the system, like the effects of the suspension hole, suspension rod, or material inhomogeneities on the frequency spectrum of the antenna. Second, we

have those questions that pertain to the various sources of vibration: How can we discriminate whether the vibration that we observe is the result of a gravitational wave or of something else? Clearly a simulation that can help us in answering these questions can be of great value to the design and development phase of the GRAIL project. Eventually, when the system is operational, a (well-defined) simulation system can even be used to analyze the transducer output by means of reverse engineering of the transducer signal towards the original source. Incorporation of simulation in the experimentation cycle is sometimes referred to as “living simulations” or “simulation in the loop.”

The remainder of this article is structured as follows. Section I begins with a description of two solution methods that can be used to answer a few basic questions but that are too restrictive for our overall goal. For this purpose we resort to development of a simulation that is based on the explicit finite-element (FE) method. The simulation accuracy required necessitates high-resolution FE models. The performance of the FE solver is therefore enhanced by exploiting the models’ parallelism (Sec. II). Since the simulation is developed from scratch, we can circumvent several pitfalls that are known to hamper the migration of existing simulation codes to parallel computers.^{12–14} In Sec. III the physical correctness of our simulation is validated by a comparison with an analytical solution as well as by experiments that were carried out with a small prototype antenna. It is further shown that we benefit greatly by exploiting computational concurrency. In Sec. IV we present a few typical simulation experiments that can be carried out with the simulation program. Some concluding remarks regarding the computational science approach taken in our study are given in Sec. V.

I. SIMULATION METHODS

A. Analytical solution

The equation of motion for elastic objects, also known as the Navier equation [Eq. (2)] describes the variation in time of the displacement field ($\mathbf{u}$) in homogeneous isotropic elastic objects. The material of such objects is parameterized by the two Lamé parameters (λ, μ) and the material density (ρ).

$$\mu \Delta \mathbf{u} + (\lambda + \mu) \nabla \nabla \cdot \mathbf{u} = \rho \ddot{\mathbf{u}}. \quad (2)$$

Two parameters that are frequently used to describe the elastic properties of a material are the Poisson ratio (ν) and Young’s modulus (E). Both parameters can be expressed in terms of the Lamé parameters as follows:

$$\nu \equiv \frac{\lambda}{2(\lambda + \mu)}, \quad (3)$$

$$E \equiv \frac{\mu(3\lambda + 2\mu)}{\lambda + \mu}. \quad (4)$$

The general solution of Eq. (2) consists of two different contributions that correspond to divergence-free ($\nabla \cdot \mathbf{u} = 0$) and rotation-free ($\nabla \times \mathbf{u} = 0$) waves that propagate through the material with unique velocities.

Figure 2. The prolate stage of the $l=0$ mode.

Figure 3. The oblate stage of the $l=0$ mode.

For both types of vibration a frequency equation can be derived for a freely vibrating (traction-free) elastic sphere and can be solved numerically.¹⁵ The rotation-free vibrations are known as *spheroidal eigenmodes*. A general spheroidal eigenmode, Ψ^s , is given by Eq. (5), which is the superposition of all particular rotation-free solutions of Eq. (2) for a spherical object with traction-free boundary conditions; note that we use polar coordinates.

$$\Psi^s = \sum_{n=0}^{\infty} \sum_{m=-n}^n [a_n(r)\hat{e}_r + b_n(r)R\nabla] Y_{nm}(\theta, \phi), \quad (5)$$

with R the sphere radius and $\hat{e}_r$ the unit vector in the radial direction; $a_n(r)$ and $b_n(r)$ are dimensionless radial eigenfunctions determined by the boundary conditions, and Y_{nm} are spherical harmonics. For each value of n the frequency equation contains infinitely many solutions. A general quadrupole mode of the sphere can be described as the superposition of the five modes with $n=2$ ($m=-2\dots2$). Since (according to GR) only these modes interact with gravitational radiation, they are the primary subject of our investigations. Figures 2 and 3 show the shape of one of the quadrupole modes (denoted as the $l=0$ mode¹⁶) differing by half a period. This eigenmode oscillates between a prolate and an oblate shape.

The applicability of this analytical model fails as soon as we perturb the sphere in any way. However, it provides us with a useful gauge for a simulation program. First, we take a look at an elegant numerical model that is less restrictive than the analytical model described above.

B. Numerical method

If we write down the Lagrangian [Eq. (6)] for a freely vibrating elastic object (following Ref. 17) (the summation convention is used where appropriate)

$$L = \int_V \left(\frac{1}{2} \rho \omega^2 u_i u_i - \frac{1}{2} c_{ijkl} u_{i,j} u_{k,l} \right) dV, \quad (6)$$

with u_i the component of the displacement field $\mathbf{u}$ that is parallel to the i axis, $u_{i,j}$ the derivative of u_i in the j direction, c_{ijkl} the elastic tensor (dependent on λ and μ), and ω the eigen- (angular) frequency, and apply the principle of least action, setting $\delta L = 0$, we get

$$\begin{aligned} \delta L &= \int_V (\rho \omega^2 u_i + c_{ijkl} u_{k,l,j}) \delta u_i dV \\ &\quad - \int_S (v_j c_{ijkl} u_{k,l}) \delta u_i dS \\ &= 0, \end{aligned} \quad (7)$$

where v_j denotes the component along the j axis of the outward-pointing normal vector on the surface of the object. Equation (7) exactly expresses the elastic wave equation (2), with traction-free boundary conditions, of Sec. I A. If we expand the displacement vector in a set of basis functions according to

$$\mathbf{u} = \sum_{lmn} \mathbf{a}_{lmn} x^l y^m z^n, \quad (8)$$

and truncate the expansion to a certain polynomial order, Eq. (7) can be rewritten in terms of a generalized eigenvalue problem (9).

Table I. The analytical and “numerical” frequencies in Hz for a few spheroidal modes of a Cu sphere. The dashes denote that no frequencies in the same range were found for specific N . This is due to the fact that the polynomial expansion in question is not able to represent the corresponding eigenmode properly.

n	Analytical solution	$N=5$ $R=168$	$N=6$ $R=252$	$N=7$ $R=360$	$N=8$ $R=495$	$N=9$ $R=660$	$N=10$ $R=858$
2	654.0828	654.090	654.090	654.085	654.085	654.085	654.085
1	891.7504	893.488	891.796	891.796	891.788	891.788	891.788
3	975.4565	988.374	975.550	975.550	975.464	975.464	975.464
4	1251.997	...	1281.98	1252.35	1252.35	1252.01	1252.01
2	1263.632	1270.78	1270.78	1263.78	1263.78	1263.68	1263.68
0	1395.698	1396.06	1396.06	1396.06	1396.06	1396.06	1396.06
5	1510.716	...	...	...	1511.6	1511.65	1510.83
3	1661.032	...	1681.20	1681.20	1661.6	1661.63	1661.04
6	1760.268	...	...	...	...	1762.42	1762.42
1	1793.638	...	1843.70	1843.70	1795.91	1795.81	1793.67

$$\omega^2 E \mathbf{a} = \Gamma \mathbf{a}. \quad (9)$$

The matrices E and Γ can be computed by integrating Eq. (7) properly. Solving for ω provides us with estimates of the eigenfrequencies of the system in question.^{16,17}

1. Validation

Table I shows several eigenfrequencies found with the numerical method as well as the corresponding analytical values for a Cu sphere (material parameters: $\rho = 8933 \text{ kg/m}^3$, $E = 129.810^9 \text{ Pa}$, and $\nu = 0.343$). N denotes the truncation order of the basis function expansion, and $R \times R$ is the matrix size of the corresponding eigenvalue problem. We have used a LAPACK¹⁸ routine called DSYGV, which can be applied to symmetric generalized eigenvalue problems to solve the system.

2. Spheroidal perturbations

An ellipsoid is described by its three semi-axes d_x , d_y , and d_z :

Figure 4. The fundamental quadrupole frequencies under variation of d_z . The spherical symmetry is broken in the $\hat{z}$ direction, resulting in a threefold splitting of the fivefold degenerate quadrupole frequency: two doublets and one singlet. The doublets are denoted by 1c_and_1s and 2c_and_2s and the singlet by 0.

$$\frac{x^2}{d_x^2} + \frac{y^2}{d_y^2} + \frac{z^2}{d_z^2} = 1. \quad (10)$$

For a sphere these axes have equal length. When we make one semi-axis longer and keep the other two constant, the object is called a prolate spheroid. If we make one semi-axis shorter than the other two, we have an oblate spheroid. In Fig. 4 the dependence of the lowest quadrupole frequency on d_z is depicted [varied in the range (0.1, 5.0) m], while the other two axes are kept at 1.5 m for a Cu sphere. We can observe three lines that intersect at $d_z = 1.5$. Since the spherical symmetry is broken in the z direction, the fivefold degeneracy is only partially removed. It splits into two doublets and one singlet. Additional perturbations in the x and y directions would break the symmetry completely, leaving us with five slightly different frequencies.

C. A finite-element model

The numerical method of Sec. I B also was applied to a sphere with a suspension hole. Again it can be observed that the five frequencies are split up (see Fig. 5). It leaves us with an uncomfortable feeling, however, since the sharp edges of the suspension hole possibly disrupt the numerical

Figure 5. The fundamental quadrupole frequencies vs the radius of the suspension hole. The spherical symmetry is broken in the $\hat{z}$ direction, resulting in a threefold splitting of the fivefold degenerate quadrupole frequency: two doublets and one singlet. The doublets are denoted by 1c_and_1s and 2c_and_2s and the singlet by 0.

Figure 6. The hull of a finite-element antenna model with the suspension rod discretized into 39,684 tetrahedral elements.

process. Therefore, we do not know whether the frequency splitting is an artifact or is physically correct. Furthermore, we are stuck with a variety of questions that cannot be answered by the analytical and numerical approaches above. Next we will consider the finite-element method as a third alternative to answer our questions. We discretize our antenna model in first-order tetrahedral elements (each element has four nodal points). As an example, Fig. 6 shows the hull of such a finite-element model, consisting of 39,684 tetrahedral elements. One approach to the problem is by means of modal analysis, that is, by studying the following eigenvalue problem:

$$K\mathbf{d} = \omega^2 M\mathbf{d}, \quad (11)$$

with K the system's stiffness matrix, M the mass matrix, $\mathbf{d}$ an eigenmode, and ω an eigenfrequency of the system. Although it is theoretically possible to investigate a linear system such as ours on the basis of modal analysis, its practical realization is a different story. First of all, the model resolutions that we aim for are very high, necessitating solution of huge eigenvalue problems. These may be solvable using state-of-the-art numerical solvers. If we succeed in doing that, we have to cope with a storage problem, since for each nodal point a superposition of at least over 100 eigenmodes is necessary to characterize a general low-frequency solution of the system. In addition, it is difficult to model the effects that are due to coupling to external forces that are aperiodic, like seismic noise entering via the suspension or stochastic sources.

On the other hand we can choose for explicit time integration of the equation of motion (12).

$$M\mathbf{a} + C\mathbf{v} + K\mathbf{d} + \mathbf{F} = 0. \quad (12)$$

The matrices M , C , and K are constructed by assembling the local mass, the viscosity, and the stiffness matrices of all individual finite elements in which the antenna model is divided; $\mathbf{a}$, $\mathbf{v}$, and $\mathbf{d}$ are vectors that describe, respectively, the acceleration, velocity, and displacement fields, and $\mathbf{F}$ is an external force vector. In order to integrate this system explicitly in time we adopt the following Newmark scheme:¹⁹

$$\begin{aligned} \mathbf{d}_{n+1} &= \mathbf{d}_n + \mathbf{v}_n \Delta t + \mathbf{a}_n \frac{(\Delta t)^2}{2}, \\ \mathbf{a}_{n+1} &= M^{-1} K(\mathbf{d}_{n+1}), \\ \mathbf{v}_{n+1} &= \mathbf{v}_n + (\mathbf{a}_n + \mathbf{a}_{n+1}) \frac{\Delta t}{2} \end{aligned} \quad (13)$$

with $\mathbf{a}_n$, $\mathbf{v}_n$, and $\mathbf{d}_n$ vectors that describe, respectively, the acceleration, the velocity, and the displacement fields in the FE system at time n . Viscous damping (C) and external forces ($\mathbf{F}$) can easily be added to this algorithm.^{19–21} The masses are “lumped” at the nodal points, resulting in a diagonal mass matrix. Therefore M^{-1} can be calculated quickly. For numerical stability Δt must be less than the Courant value, since the time integration scheme is conditionally stable. The choice of the algorithm is, on the one hand, motivated by the fact that it displays good numerical properties (low numerical dispersion and damping). This was found through a number of exploratory experiments on a one-dimensional finite-element model (data not shown). On the other hand, it is very well suited for exploiting computational concurrency, as we will see below. Although this explicit approach has the drawback that we have to solve our system in the time domain and afterwards perform a spectral analysis, it is much more flexible for our purposes, and is not hampered by the restrictions that apply to modal analysis. Furthermore, the simulation model can be extended such that it can incorporate nonlinear effects, such as *violin modes* in the suspension rod, that cannot be realized by modal analysis.

It was expected that the required simulation accuracy necessitates the use of a very-high-resolution FE model that is computationally and memory intensive. The simulation code was therefore designed and implemented for execution on distributed-memory computers.

II. PARALLELIZATION

In good tradition, we have chosen to base our parallelization method on the decomposition of the FE mesh. The consequences of the decomposition method for the time integration algorithm as well as the calculation of important system parameters (like energy) had to be considered in the design. In our discussion we will only consider the pure finite-element system. In the full-fledged simulation a variety of phenomena can be taken into account, such as transducer mounting, suspension rod, and sources of vibration. At this point it suffices to say that the parallelization method chosen does not have major consequences for the full-fledged simulation, and therefore may be left out of this discussion. We focus on two important aspects of the parallelization process: parallel time integration (the kernel) and parallel calculation of the energy.

Figure 7. Partitioning of a quadrilateral mesh into four parts. The nodal points at the shared boundaries are replicated while elements are assigned to unique processes.

The code was implemented in C with PVM²² message passing primitives according to the single-program, multiple-data paradigm.²³ The problem domain (the finite-element mesh consisting of N tetrahedral elements) is partitioned into P subdomains. Each subdomain is assigned to a single process and executes the same program applied to its local data. The process with identifier 0 is used to read in the whole mesh; it supplies the other $P-1$ processes with their locally essential data, and is used for input/output (I/O), a *primus inter pares* or *master* among the parallel processes.

A. Energy

In order to check energy conservation, we have to monitor the energy in the system. The total energy E_{tot} is given by

$$E_{\text{tot}} = \frac{1}{2}(\mathbf{v}^T \mathbf{M} \mathbf{v} + \mathbf{d}^T \mathbf{K} \mathbf{d}), \quad (14)$$

with $\mathbf{v}$ the velocity vector in which the individual velocity vectors of all nodal points are concatenated, and $\mathbf{d}$ the displacement vector, also formed by concatenating the fields at all nodal points. How do we calculate this in parallel?

The FE mesh is decomposed along the surfaces of the tetrahedral elements that constitute the mesh. First, a dual graph of the mesh is constructed; this then can be partitioned using any graph-partitioning method.²⁴ Nodal points that lie on shared boundaries are replicated on each process with this boundary in its local domain, whereas elements belong to a unique process. This procedure is schematically depicted in Fig. 7 for a simple quadrilateral mesh.

The global mass and stiffness matrices are constructed by assembling the individual contributions from each element in a global matrix. If we denote the assembly of two-element stiffness matrices k_1 and k_2 by $k_1 \oplus k_2$, K is constructed (analogously for the mass matrix) as follows:

$$K = k_1 \oplus k_2 \oplus \dots \oplus k_N, \quad (15)$$

with N the number of tetrahedral elements. If we monitor the energy values during a simulation run (it would be nice if this could be done completely separate), that is, for each subdomain (i), the energy content (E_{tot}^i) is calculated, and once in a while we assemble the subdomain energy contributions and accumulate them on process 0. This can be done as follows:

$$E_{\text{tot}}^i = \frac{1}{2}(\mathbf{v}^{iT} \mathbf{M}_i^i \mathbf{v}^i + \mathbf{d}^{iT} \mathbf{K}_i \mathbf{d}^i), \quad (16)$$

$\mathbf{v}^i$ denotes the velocity vector for the nodes in subdomain i , $\mathbf{d}^i$ is the displacement field, K_i is the stiffness matrix of

subdomain i , which is simply the assembly of its local element stiffness matrices, and finally $\mathbf{M}_i^i$ is an adjusted version of the (diagonal) mass matrix of subdomain i . The “normal” mass matrix contains the mass for each nodal point at the appropriate entry. In this adjusted version the masses of the internal nodal points are untouched. However, for the shared nodal points this value is divided by its “redundancy” value, which denotes how many processes have a replica of it. For example, in Fig. 7 nodal point 5 has a redundancy value of 4. As a consequence, the corresponding entry in each local adjusted mass matrix, either being on process 0, 1, 2, or 3, has to be divided by 4. In this way we can assure that we are not counting any energy contributions twice.

B. Time evolution

If we have constructed the mass and stiffness matrices for each subdomain (or process) correctly, we can parallelize the Newmark scheme above as follows: Each process gets a unique process identifier (ID) equal to the subdomain identifier (i) above. Nodal points that are shared among two or more processes are considered to have one real owner (the owning process considers it as a node of type A), while the other processes only have the replica of this node (considering the node as type B). Assigning a node to a specific process is based on the process ID. The process with the lowest ID considers a shared nodal point as A type, while the other processes see it as B type. Consequently, process 0 can only have shared nodes of type A, while process $P-1$ (assuming we have P processes) can have only shared nodes of type B. The following construction parallelizes the Newmark scheme correctly:

$$\begin{aligned} \mathbf{d}_{n+1} &= \mathbf{d}_n + \mathbf{v}_n \Delta t + \mathbf{a}_n (\Delta t)^2 / 2, \\ \text{Send } \mathbf{d}_{n+1} &\text{ from A nodes to B nodes,} \\ \text{Receive } \mathbf{d}_{n+1} &\text{ for B nodes from A nodes,} \\ \mathbf{a}_{n+1} &= \mathbf{M}^{-1} \mathbf{K} (\mathbf{d}_{n+1}), \\ \text{Send } \mathbf{a}_{n+1} &\text{ from B nodes to A nodes,} \\ \text{Receive } \mathbf{a}_{n+1} &\text{ for A nodes from B nodes,} \\ \text{Add contributions to } \mathbf{a}_{n+1} &\text{ from B nodes,} \\ \mathbf{v}_{n+1} &= \mathbf{v}_n + (\mathbf{a}_n + \mathbf{a}_{n+1}) \Delta t / 2. \end{aligned}$$

The A–B nodal point scheme was adopted from Lonsdale *et al.*²⁵

III. VALIDATION

Here we show a selection of experimental results that were produced for validation of our parallel FE simulation. First, we investigate whether the model is physically correct. Next, we measure the performance benefits that we get from exploiting parallelism for different mesh resolutions. Meshes are decomposed by two different graph-partitioning methods: recursive spectral bisection (RSB) and a generalized version of orthogonal recursive bisection (ORB),²⁴ which we call *generalized ORB* (GORB). GORB operates in such a way that consecutive partitionings are not necessarily orthogonal and that partitioning cardinalities are not necessarily a power of 2 (as is the case in ORB). Figure 8 shows partitioning in the case of GORB in four parts, where the consecutive partitionings are not performed orthogonally but at an angle of 80°. GORB recursively bi-

ORB partitioning

Figure 8. The hulls of the four partitions obtained for a sphere consisting of 79,556 tetrahedral elements obtained using the GORB partitioning method. Each of the four hulls has about 3140 triangular faces. The first bisection is along the $\hat{x}$ direction. Next, the bisection surface is rotated 80° and the two subdomains are both bisected individually along this second direction.

RSB partitioning

Figure 9. The hulls of the four partitions obtained for the same sphere as in Fig. 8, but now using the RSB partitioning method. Although the edges in this partitioning appear more ragged than those obtained using ORB, they have only about 3010 faces. RSB optimizes the connectivity within partitions and minimizes the communication between them.

sects the geometrical domain in equal partitions (work load balancing), in alternating directions, not paying any attention to the mesh connectivity (communication). RSB applies a more complex, and consequently computationally more expensive, method to partition the mesh. In contrast with (G)ORB it also takes into account the mesh connectivity. It recursively bisects a finite-element mesh into equal numbers of elements (work load balancing), while (approximately) minimizing the number of edges between the resulting subdomains (communication load balancing). Figures 8 and 9, respectively, show the partitioning in the cases of GORB and RSB into four parts, with both bisections applied to a FE sphere of 79,556 elements. The difference in partitioning quality between GORB and RSB will have to be determined by explicitly measuring the kernel execution times because it cannot be found by simple visual inspection.

The performance measurements are carried out on two target platforms both installed at the IC³A in Amsterdam;²⁶ a Parsytec PowerXplorer [32 PowerPCs 601, with each 32 Mbytes of random access memory (RAM)] and a Parsytec CC (40 PowerPCs 604, each having 96 Mbytes of RAM). Our target material in the experiments below is the alloy CuAl (90–10) (with $\rho=7534.0 \text{ kg/m}^3$, $E=154.0 \times 10^9 \text{ Pa}$, and $\nu=0.3265$), and only considers spheres with radii of 1.5 m. All results with respect to the physical experiments given below are independent of the number of processors

used, which indicates the correctness of our parallel implementation.

A. Physical correctness

Let us consider simulations that are initiated with a pure quadrupole mode, that is, each nodal point in the system gets an initial velocity $\mathbf{v}^T = \dot{h}(x, -y, 0)$, with $\dot{h}$ the time derivative of h , a scale-free parameter denoting the deformation of space. Since the FE models are always slightly asymmetric, the fivefold degeneracy will be removed. Table II enumerates the average frequencies that are found around the analytical eigenfrequency f_0 ($f_0=780.25 \text{ Hz}$)

Table II. The average frequency $\langle f \rangle$ and the approximated error $\sigma(f)$ in % for the quadrupole principal mode for various model resolutions.

n	$\langle f \rangle$	$\sigma(f) \%$
84	$7.91e+02$	$8.72e-01$
124	$7.92e+02$	$1.70e+00$
698	$8.01e+02$	$4.87e-01$
1086	$7.97e+02$	$4.44e-01$
2724	$7.90e+02$	$3.23e-01$
2851	$7.96e+02$	$3.62e-01$
6114	$7.86e+02$	$2.25e-01$
9412	$7.85e+02$	$3.01e-01$
51501	$7.82e+02$	$2.85e-01$
79556	$7.82e+02$	$1.56e-01$

Figure 10. The time trajectory in the $\hat{x}$ direction of an arbitrary nodal point at the surface of a FE model of 79,556 elements initiated with a pure quadrupole velocity field.

and the deviation in % of the average value for 10 different FE models that consist of n tetrahedral elements. These results indicate that the dominating frequency for the FE models approaches the analytical frequency value for increasing model resolution (up to 0.16% from the analytical value for the highest resolution model).

Figure 10 shows the $\hat{x}$ displacement of an arbitrary nodal point at the sphere surface in the same simulation for the highest-resolution model. At first glance, it seems that the nodal point oscillates with a frequency of about 780 Hz, with a signal of ≈ 70 Hz superimposed. However, Fourier analysis (data not shown) has pointed out that this signal mainly results from superposition of the three oscillations that are most pronounced, that is, those oscillating at 781.8, 1493, and 1499 Hz, respectively. Possible beats that result from superposition of the principal quadrupole modes (around 780 Hz) are not visible in this picture.

We can also observe the system energy during simulation. It can be derived that the analytical energy (E_{quad}) for a sphere vibrating in a pure quadrupole mode can be expressed in terms of the radius (R) and the material density (ρ) as follows:

$$E_{\text{quad}} = \frac{4}{15} \pi \rho R^5 \dot{h}^2. \quad (17)$$

For the CuAl sphere with a radius of 1.5 m it follows that $E_{\text{quad}} = 47,929.2$ Nm. Within the finite-element simulations, it was found that the energy approaches this value very well for increasing model resolution (data not shown). For instance, the total energy for the highest-resolution model of 79,556 elements is approximately equal to 47,924.7, which approaches the analytical value up to 0.01%.

Within the GRAIL project, a prototype spherical antenna of the same CuAl alloy as that above, but with a radius of 0.075 m was fabricated. The following experiment was carried out. The prototype is excited by a radial impulse. Next, the time series of radial vibrations are recorded at several (arbitrarily chosen) places on the antenna surface. From these time series we have derived the power spectrum of the antenna vibrations by alternately applying a fast Fourier transform (FFT) and a nonlinear least-squares

Table III. Matches between frequencies computed using the explicit finite element code and experimental values for a 15 cm sphere that is excited by a purely radial excitation. In both sets of data, mode excitations may have been below the detection threshold, explaining frequencies occurring in a single set only. The analytical frequencies that are found for a pure sphere with the same material specifications are (within the same frequency range) 14830 (2,1), 15,708 (2,2), 21,243 (1,2), 22,915 (3,1), 23,407 (3,2), 30,025 (4,2) Hz, with, for instance, 14,830 (2,1) denoting a mode vibrating with 14,830 Hz, $n=2$ of toroidal type (1), and 15,708 (2,2) a mode vibrating with 15708 Hz, $n=2$ of spheroidal type (2).

Computed (Hz)	Measured (Hz)	Computed/measured
13,650	13,687	0.997
13,740	13,720	1.001
	13,797	
13,849	13,815	1.002
14,391	14,424	0.998
14,500	14,470	1.002
14,626	14,587	1.003
19,580		
19,689	19,690	1.000
19,797	19,812	0.999
21,045		
21,171	21,138	1.002
21,207	21,289	0.996
21,316	21,307	1.000
	21,343	
	21,371	
21,623	21,647	0.999
	21,667	
21,750	21,777	0.999
21,840	21,809	1.000
27,716	27,719	1.000
	27,739	
	27,802	
27,843	27,863	0.999
	27,919	
	27,954	
	27,972	
27,987	27,991	1.000

fitting procedure on the signal residue. Simultaneously, a simulation experiment was performed in which the same procedure was carried out. That is, a radial displacement is induced on an arbitrary surface point of a detailed finite-element model consisting of 191,310 tetrahedral elements. The model is realistic in the sense that it incorporates the suspension rod and the suspension hole. For several points at the model surface we record the time series, analogous to the real experiment, followed by the same procedure to produce a power spectrum. Table III lists a part of the frequency spectrum that is obtained for a radial impulse on the prototype antenna, accompanied by the frequencies found from the finite element simulation. The third column in Table III corresponds to the ratio of model frequency to experimental frequency. We can observe that the frequency spectrum found in the simulation model is quite close to experimental values. The residue of the experimental signal and the simulated signal turns out to be less than 0.1%.

Table IV. The kernel execution times in seconds on the PowerXplorer, for four different mesh sizes and two different decomposition methods (GORB denoted by suffix G and RSB denoted by suffix R) versus the partitioning cardinality.

	6114-G	6114-R	9412-G	9412-R	51501-G	51501-R	79556-G	79556-R
1	$3.46e-01$	$3.46e-01$	$5.25e-01$	...	...	...	...	...
4	$1.47e-01$	$1.45e-01$	$2.05e-01$	$2.02e-01$	$8.89e-01$	$8.85e-01$	$1.31e+00$	$1.30e+00$
6	$1.58e-01$	...	$2.02e-01$	...	$6.95e-01$	...	$1.01e+00$	...
8	$1.22e-01$	$1.13e-01$	$1.77e-01$	$1.72e-01$	$5.98e-01$	$6.00e-01$	$7.65e-01$	$7.84e-01$
9	$1.52e-01$	...	$2.17e-01$	...	$5.99e-01$	...	$7.80e-01$	...
12	$1.25e-01$	...	$1.84e-01$	...	$5.14e-01$	...	$6.55e-01$	...
16	$1.10e-01$	$1.04e-01$	$1.64e-01$	$1.63e-01$	$5.00e-01$	$4.53e-01$	$6.34e-01$	$6.41e-01$
18	$1.02e-01$	...	$1.57e-01$	...	$4.78e-01$	...	$6.71e-01$	...
24	$7.99e-02$	...	$1.31e-01$	...	$4.61e-01$	...	$6.93e-01$	...
27	$8.01e-02$	...	$9.75e-02$	...	$4.72e-01$	...	$6.59e-01$	...
32	$7.00e-02$	$8.58e-02$	$9.07e-02$	$1.21e-01$	$4.42e-01$	$4.59e-01$	$6.11e-01$	$6.36e-01$

B. Parallel performance

Next we consider the parallel performance of the simulation kernel. We apply GORB (partitioning cardinalities of 2, 4, 6, 8, 9, 12, 16, 18, 24, 27, and 32) and RSB (partitioning cardinalities of 4, 8, 16, and 32) to four different meshes consisting of 6114, 9412, 51,501, and 79,556 elements, respectively. Table IV displays the single-kernel execution times (one time step in the Newmark scheme) that were obtained for simulations on the PowerXplorer for a varying number of processors. Table V shows similar results obtained on the Parsytec CC. The execution times are the best case, that is, energy and external force calculations (modeling gravitational radiation sources) are left out.

In both Tables IV and V we can observe that the difference in execution speed between GORB and RSB is not significant. In fact GORB is often slightly better than RSB for partitioning cardinalities that are powers of 2. This can be ascribed to the fact that both partitioning methods create compact three-dimensional subdomains that do not contain any disconnected parts, and this leads to an average surface-volume ratio (or, equivalently, communication-calculation ratio) of these subdomains that is approximately the same.

The working memory of the CC processors is much larger than that of the PowerXplorer. Therefore on the CC

all problems can be executed on a single processor, in contrast to the PowerXplorer, where single-processor performance cannot be determined for high-resolution problems. The CC communication network and CPU are significantly faster than those of the PowerXplorer, which accounts for the fact that the kernel execution time for each problem is shorter on the CC.

For the PowerXplorer the kernel execution times are not significantly decreased for more than 16 partitions, whereas in case of the CC increasing partitioning cardinality consequently leads to faster execution times, at least for high-resolution models. Since we have considered the best performance, we can expect that the absolute performance of the simulation code in a more realistic simulation (taking other phenomena into account) will be degraded compared to the best case and, consequently, result in a relatively better “scalability” of the code, since the amount of work per subdomain increases whereas the communication pattern and size remain unchanged.

IV. “SAMPLE” SIMULATION EXPERIMENTS

Since the main theme of this article is the concept of a complex parallel simulation program, we will not present detailed quantitative results with respect to the feasibility of

Table V. The kernel execution times in seconds on the CC for four different mesh sizes and two different decomposition methods (GORB denoted by suffix G and RSB denoted by suffix R) versus the partitioning cardinality.

	6114-G	6114-R	9412-G	9412-R	51501-G	51501-R	79556-G	79556-R
1	$1.38e-01$	$1.38e-01$	$2.10e-01$	$2.10e-01$	$1.12e+00$	$1.12e+00$	$1.71e+00$	$1.71e+00$
2	$8.68e-02$	...	$1.27e-01$	...	$6.23e-01$	...	$9.20e-01$	...
4	$5.41e-02$	$5.34e-02$	$7.63e-02$	$7.48e-02$	$3.47e-01$	$3.39e-01$	$4.99e-01$	$4.93e-01$
6	$6.26e-02$	...	$7.61e-02$	...	$2.39e-01$	...	$3.47e-01$	...
8	$5.49e-02$	$5.40e-02$	$7.94e-02$	$7.84e-02$	$2.23e-01$	$2.26e-01$	$2.92e-01$	$2.93e-01$
9	$4.85e-02$	...	$7.14e-02$	...	$2.17e-01$	...	$2.72e-01$	...
12	$4.26e-02$	...	$5.93e-02$	...	$1.83e-01$	...	$2.72e-01$	...
16	$3.74e-02$	$3.99e-02$	$5.02e-02$	$4.88e-02$	$1.83e-01$	$1.87e-01$	$2.64e-01$	$2.37e-01$
18	$3.66e-02$	...	$4.84e-02$	...	$1.64e-01$	...	$2.30e-01$	...
24	$3.73e-02$	...	$4.36e-02$	...	$1.17e-01$	...	$1.80e-01$	...
27	$3.75e-02$	...	$4.36e-02$	...	$1.24e-01$	...	$1.79e-01$	...
32	$3.98e-02$	$4.40e-02$	$4.42e-02$	$4.66e-02$	$1.15e-01$	$1.19e-01$	$1.62e-01$	$1.66e-01$

Figure 11. The spherical antenna with one transducer system mounted.

building a gravitational wave detector. Nevertheless, we will illustrate the capabilities of the FE simulation program using a few “sample” experiments.

All simulations are carried out with full-fledged antenna models of high resolution (approximately 130,000 elements) on the two high-performance distributed-memory platforms discussed in Sec. III B. The antenna material resembles CuAl (90–10). Six resonant mass transducers are mounted on the antenna in somewhat of an arrangement that places the transducers at the pentagons on the lower half of a truncated icosahedron; this was also done in the TIGA project.²⁷ Each transducer consists of two masses. The first mass of $M_1 = 100$ kg is attached to the antenna by a spring that is described by the diagonal stiffness matrix $K_1 = \text{diag}(2.27 \times 10^9, 2.27 \times 10^{10}, 2.27 \times 10^{10})$ N/m. The second mass of $M_2 = 0.385$ kg is attached to the first mass by a spring with stiffness matrix $K_2 = \text{diag}(8.70 \times 10^6, 8.70 \times 10^7, 8.70 \times 10^7)$ N/m.

Figure 11 schematically depicts a transducer that is attached to the antenna. The principal resonance frequency f_r of the transducer system is equal to $f_r = 780$ Hz, which coincides with the principal quadrupole frequency of the ideal detector (see Sec. III A). The two masses are chosen such that the amplitude of vibration detected at the antenna’s surface will be amplified by two orders of magnitude. The addition of the multimode resonant transducers results in a significant splitting of the resonant frequencies of the system (data not shown). Although the transducer models that we have implemented are rather simplistic (point masses coupled with ideal springs), they have already given quite satisfactory results in several applications (see, for example, Sec. IV B).

A. Gravity

Earth’s gravity will induce internal strain in the antenna. The sphere that is suspended is expected to deform due to this strain. We would like to quantify the amount of deformation that results due to the presence of gravity. In order to do this we perform the following simulation. At time zero we “switch on” gravity ($g = 9.8 \text{ m/s}^2$), and let the antenna fall down. The suspension rod is fixed at the top to prevent the sphere from dramatically accelerating towards the Earth’s surface. The velocities in the sphere are quenched by a small damping factor in order to get rid of the kinetic energy in the system. When the system has come to rest, we take a look at the displacement and strain distribution within the sphere. Figure 12 displays a contour

Figure 12. Contour plot of the displacement in the $\hat{z}$ direction of the suspended CuAl sphere due to gravity ($g = 9.8 \text{ m/s}^2$). The displacement values of the various contours correspond to (1) -0.35 , (2) -1.96 , (3) -2.18 ; further contours do not differ significantly from the value of contour (3); all values are in 10^{-3} m . The minus signs indicate displacements towards the Earth’s surface.

plot of the displacement field in the $\hat{z}$ direction within the antenna, after this deformation process. Figure 13 zooms in on the suspension region, again showing a contour plot of the displacement in the $\hat{z}$ direction. Overall, the sphere moves about 2 mm downwards, which results from the fact that the suspension rod is elongated by this amount in the $\hat{z}$ direction. The displacement values within the sphere vary from about 2×10^{-3} to $3 \times 10^{-4} \text{ m}$ when we traverse the sphere from the outside towards the suspension rod. Another interesting view is provided by the vertical strain

Figure 13. The same contour plot as that displayed in Fig. 12, but zoomed in on the suspension region. The displacement values of the various contours correspond to (1) -1.60 , (2) -1.82 , (3) -2.16 , (4) -2.18 ; further contours do not differ significantly from the value of contour (4); all values are in 10^{-3} m. The minus signs indicate displacements towards the Earth's surface.

$(\partial u_z / \partial z)$ within the antenna, which is shown in Fig. 14. As expected, the suspension rod is subjected to the highest strain, as is the region where the sphere and rod make contact.

A shortcoming of our simulation program is that it relies on a linear finite element solver. As a consequence, we are not able to model nonlinear effects, such as so called *violin modes* in the suspension rod, or show how the internal strain of the sphere affects the eigenmodes. Incorporating these into our simulation program is a future challenge.

Figure 14. The vertical strain $(\partial u_z / \partial z)$ within the antenna. High strain values (dark areas) are found in the suspension rod, that is carrying the entire sphere, and near the contact region of the sphere and the rod. The strain in the sphere itself is much lower (light colored areas).

B. Seismic noise

The suspension system of the antenna must attenuate vibrations at the principal resonance frequency to a factor of 320 dB.²⁸ If we let seismic noise enter the system via the suspension, how does it appear at the transducer readout? The following experiment is carried out. The suspension points (connections to the outside world) of the suspension rod are driven by forced oscillation in the $\hat{z}$ direction. That is, the z coordinate of the displacement field $u_z(t)$ of each suspension point oscillates as follows:

Figure 15. Narrow-banded seismic noise around the principal quadrupole frequency.

$$u_z(t) = \sum_{n=1}^N [A_n \cos(2\pi f_n t) + B_n \sin(2\pi f_n t)]. \quad (18)$$

The amplitudes A_n and B_n and the frequency spectrum f_n model the seismic noise. In the following experiment the narrow-banded noise spectrum depicted in Fig. 15 was used. The $\hat{z}$ displacement versus time, which is driving the suspension points according to this spectrum, is displayed in Fig. 16.

Figures 17 (inner mass) and 18 (outer mass) show the radial displacement of the two transducer masses starting at $t=0$ in one of the six transducers. The maximum amplitude in the noise spectrum (Fig. 15) has the same order of magnitude as the maximum amplitude of the outer mass. Consequently, it is of paramount importance that all seismic vibrations in the target frequency range are attenuated below the expected amplitude of vibrations that are induced by gravitational radiation in order to assure that we have a decent signal-to-noise ratio.

Figure 16. The $\hat{z}$ displacement for the suspension point driven by the noise spectrum of Fig. 15.

Figure 17. Time trajectory of the inner transducer mass on one transducer. The system is forced into oscillation by the seismic noise spectrum of Fig. 15.

C. Chirps

A so-called gravitational radiation “chirp” will be emitted (expected theoretically) just before coalescence of a binary neutron-star system.²⁹ The effects of this phenomenon on the detector can be modeled by applying a time-varying “chirp force” $\mathbf{f}(t)$ on the antenna, which acts along a pure quadrupole field, according to

$$\mathbf{f}^T(t) \propto \ddot{h}(x, -y, 0), \quad (19)$$

with $\ddot{h}$ the second-order time derivative of the deformation of space, derived in a Newtonian approximation.

Let us consider a chirp that enters the antenna at time $t = -0.04$ s. Figure 19 models its time varying amplitude $\ddot{h}$. The normalization of the chirp signal has an order of magnitude that is representative for an event in the Virgo cluster. Since the antenna responds purely linearly, the shape of the response signal resulting from interaction with any other identically shaped incoming signal will be exactly the same.

Figure 18. Time trajectory of the outer transducer mass on one transducer. The system is forced into oscillation by the seismic noise spectrum of Fig. 15.

Figure 19. The time dependent amplitude $\ddot{h}$ of a chirp initiating at time $t = -0.04$ s. The chirp ends at $t = 0$.

Figure 20 shows the evolution of the energy in the antenna from the moment that the chirp enters until it has passed. The energy of the chirp gets deposited in the antenna. After the chirp has passed, the energy stabilizes (at $t = 0$ s).

The response to this chirp can again be measured. Figure 21 shows the time trajectory in the $\ddot{x}$ direction of one of the outer transducer masses. We observe the tiny ($\approx 10^{-19}$ m) displacements of the transducer masses.

V. CONCLUSIONS

In this article we have sketched the complete path along which a complex high-performance simulation was developed. We were able to identify the problem definition, modeling choices, the code design for parallel systems, validation, and experimentation. In order to realize such a system, it is mandatory that we have knowledge of a variety of disciplines, like computer science, numerical mathematics, and (computational) physics. It is our feeling that

Figure 20. The time evolution of the total energy deposited by the chirp of Fig. 19. After $t = 0$ the energy stabilizes. Note that we have used a log scale on the energy axis.

Figure 21. The $\hat{x}$ displacement of one of the outer transducer masses. The transducer already “notices” the chirp’s presence before $t = 0$.

the whole process that is described above is typical for the development of new simulation systems that utilize high-performance methodology and therefore can serve as an example of *computational science*, which inherently is an interdisciplinary research field. Mixing the best of several worlds allowed us to realize a full-fledged simulation system that meets our purposes.

In the future we will focus on embedding the simulation program within the experimentation cycle of GRAIL, referred to as “simulation in the loop” in the Introduction.

ACKNOWLEDGMENTS

The work presented in this article was made possible by funding from NWO of the GRAIL project pilot study. The contributions of J. de Rue to Sec. IB are gratefully acknowledged.

REFERENCES

1. C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (Freeman, San Francisco, 1973).
2. <http://qv3pluto.leidenuniv.nl/grail/grail.html>. GRAIL stands for “gravitational radiation antenna in Leiden.”
3. The GRAIL Project Team. GRAIL: A proposal for a gravitational radiation antenna in The Netherlands. Technical Report No. NIKHEF/95-005, Nationaal Instituut voor Kernfysica en Hoge-Energie Fysica, 1995.
4. K. S. Thorn, *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987).
5. N. Ashby and J. Dreitlein, *Phys. Rev. D* **12**, 336 (July 1975).
6. J. A. Lobo, *Phys. Rev. D* **52**, 591 (July 1995).
7. S. M. Merkowitz and W. W. Johnson, *Phys. Rev. D* **51**, 2546 (March 1995).
8. S. M. Merkowitz, Ph.D. thesis, Louisiana State University, 1995.
9. J. Weber, *Phys. Rev.* **117**, 307 (1960).
10. <http://www.roma1.infn.it/rog/rogmain.html>. Rome gravitational wave experiment, NAUTILUS, 1996.
11. Dutch organization of scientific research.

12. P. M. A. Slood and J. Reeve, Camas-tr-2.3.7 executive report on the Camas workbench; Technical report, University of Amsterdam and University of Southampton, October 1995.
13. B. D. Kandhai, P. M. A. Slood, and J. P. Huot, in *High-performance Computing and Networking*, edited by B. Hertzberger, H. Liddel, A. Colbrook, and P. M. A. Slood, No. 1067 in ISBN 3-540-61142-8 (Springer, Berlin, 1996), pp. 269–275.
14. P. M. A. Slood, Eurosim '95, Simulation Congress, Amsterdam, The Netherlands, 1995, pp. 29–44.
15. A. E. H. Love, *A Treatise on the Mathematical Theory of Elasticity* (Dover, New York, 1944).
16. J. de Rue, MS. thesis, University of Amsterdam, 1996.
17. T. M. Bell, W. M. Visscher, A. Migliori, and R. A. Reinert, J. Acoust. Soc. Am. **90**, 2154 (1991).
18. E. Andersen, Z. Bai, and C. Bischof, LAPACK User's Guide, 2nd ed. 1994.
19. O. C. Zienkiewicz and R. L. Taylor, *The Finite Element Method* (McGraw-Hill, New York, 1994).
20. T. J. R. Hughes, *The Finite Element Method, Linear and Dynamic Finite Element Analysis* (Prentice-Hall, Englewood Cliffs, NJ, 1987).
21. R. H. MacNeal, *Finite Elements: Their Design and Performance* (Dekker, New York, 1994).
22. A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, PVM 3 User's Guide and Reference Manual, Oak Ridge National Laboratory, May 1993.
23. G. Fox, M. Johnson, G. Lyzenga, S. Otto, J. Salmon, and D. Walker, *Solving Problems on Concurrent Processors* (Prentice-Hall, Englewood Cliffs, NJ, 1988), Vol. 1.
24. H. D. Simon, Partitioning of unstructured problems for parallel processing. Comput. Syst. Engr. **2**, 135 (1991).
25. G. Lonsdale, J. Clinckemaulle, S. Vlachoutsis, J. F. de Ronde, P. M. A. Slood, N. Floros, and J. Reeve, Conference on Supercomputing Applications in the Automotive Industries of the 26th ISATA, September 1993.
26. Interdisciplinary Center for Complex Computer Facilities. <http://www.wins.uva.nl/research/ic3a/>. Amsterdam, 1996.
27. W. W. Johnson and S. M. Merkwitz, Phys. Rev. Lett. **70**, 2367 (April 1993).
28. T. L. Aldcroft, Rev. Sci. Instrum. **63**, 3815 (1992).
29. L. Blanchet, B. R. Iyer, C. M. Clifford, and A. G. Wiseman, Class. Quantum Grav. **13**, 575 (1996).