

High Performance Simulation for Resonant-Mass Gravitational Radiation Antennas

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

University of Amsterdam, Department of Mathematics and Computer Science, Kruislaan 403,
1098 SJ Amsterdam, tel +31 20 525 7463, fax +31 20 525 7490,
<http://www.wins.uva.nl/research/pscs>, email: {janr, dick, peterslo}@wins.uva.nl,

Abstract. In this paper the design and validation of a high performance simulation is discussed that is of critical value to the feasibility study of the GRAIL project, the aim of which is to build a gravitational radiation antenna. Two relatively simple simulation models of this antenna are shown to be too restrictive for our purposes, necessitating the development of a simulation program that utilizes an explicit finite element kernel. The computational complexity of this simulation kernel requires the power that is offered by high performance computing methodology. Therefore it is tailored for execution on parallel systems. Since it is developed from scratch, we can circumvent notorious parallel programming pitfalls that usually are present in code migration. The simulation program is validated for its physical correctness as well as its performance gain. Performance results are presented for two distributed memory parallel systems: A Parsytec PowerXplorer (32 PowerPC's) and Parsytec CC (40 PowerPC+'s).

Keywords: High Performance Simulation, Development from Scratch, Explicit Parallel Finite Element, Gravitational Radiation Antenna's

Introduction

General Background

Although the existence of gravitational radiation (GR), being predicted by Einstein's theory of General Relativity, is unquestioned, the detection is a long-standing problem in experimental physics. The aim of the GRAIL² project 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 to validate the existence of GR that is emitted by astrophysical sources. These sources include supernovae, stellar collapses to black hole states and coalescence of binary neutron star systems.

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

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

According to general relativity, an oscillation of the quadrupole moment is the simplest mode of vibration generating gravitational waves. The counter statement for absorption

* presenting author

² GRAIL stands for *Gravitational Radiation Antenna in Leiden*. The participating Dutch (academic) institutes are RUL, UT, TUE, UvA, NIKHEF/FOM, SRON and ESA/ESTEC

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 GR antenna which essentially is a large resonant mass, that is forced in oscillation by impinging gravitational radiation.

In GRAIL, the antenna will consist of a spherical resonant mass of a copper 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 a way that the external vibrations at its resonant frequency (≈ 700 Hz) are attenuated by at least a factor of 10^{16} . Furthermore the sphere will be cooled to a temperature in the range of 10-20 millikelvin, in order to cancel out thermal noise as well as possible. 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 (like source type and direction of incidence) about the radiation which impinged on the sphere. The gravitational waves have extremely weak interaction with matter, the typical strains induced on a copper sphere of 3 m will be of the order of 10^{-20} m, which explains why the detection is such a difficult task and has not been realized yet.

Since the estimated cost (≈ 45 million Dutch guilders) of the project is quite high, a phase 0 study has been funded by NWO³, to study the technological feasibility of GRAIL. A variety of critical aspects concerning the design of the antenna was to be addressed in this study. Our task consisted of building detailed models of the antenna, which can give insight into its vibrational properties. A schematic design of the GRAIL antenna is shown in Fig. 1.




Fig. 1. A drawing of the GRAIL sphere in a heat-isolating vacuum chamber. On top the concrete arms to which it is suspended are visible.

³ Dutch Organization for Scientific Research

Critical Design Issues for GRAIL

The following subjects were to be investigated by means of simulation.

1. How does the suspension affect the eigen-modes and eigen-frequency for the fundamental quadrupole modes of a sphere?
2. How will material inhomogeneities affect these eigen-modes and eigen-frequencies?
3. Given a perfect sphere, how will it deform (into a droplet shape) 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 energy transducers on the antenna on the frequency spectrum of the system?
6. How will seismic noise, entering the sphere via the suspension, be visible at the antenna's surface? That is, how will the transducers "see it"?
7. Can energy, that has been deposited by cosmic rays or particles like muons, induce system eigen-modes that are not distinguishable from those induced by impinging gravitational waves?
8. How will the finger-print of typical gravitational wave sources (Like "chirps" and continuous waves) be seen by the sphere?

We can identify in principle two types of questions in this enumeration. Firstly, we have the design questions that have to be answered in order to tune the read-out of the system, like the effects of the suspension-hole and suspension-rod or material inhomogeneities. Secondly, we have the questions that also are important after the system has been built. How can we discriminate if the vibration that we observe is the result of a gravitational wave or something else? Clearly, a simulation that can help us in answering the questions above 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.

This paper is structured as follows. Section 1 introduces two models that can be used to answer a few basic questions but are too restrictive for any of the above questions. For this purpose we have to resort to the development of an explicit finite element (FE) simulation tool, that is tailored for our purposes. The accuracy that we aim for requires a very high resolution FE model, which is the reason for optimizing the FE solver, by exploiting the models' parallelism, which is the topic of section 2. Since the tool is developed from scratch, we can circumvent several pitfalls that are known to hamper the migration of existing simulation codes [7][2][6]. In section 3 we validate the tool's correctness and show that indeed we have a significant benefit from exploiting parallelism. Section 4 is reserved for a brief summary of the foregoing, followed by a discussion of the experimental results. We conclude with the observation that HPC technology enabled us to realize a simulation tool that runs efficiently, allowing us to perform accurate simulations within the restricted time-frame (1 year) of the GRAIL phase 0 study.

1 Simulation Methods

1.1 Analytical Solution

The equation of motion for elastic objects, also known as the *Navier* equation (Eq.(2)) can be solved analytically for a freely vibrating spherical object. It describes the variation in time of the displacement field (\mathbf{u}) in the object. The material of the object is

parameterized by the two *Lame* parameters (λ, μ) and the material density (ρ).

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

The general solution 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. For both types of vibration a frequency equation can be derived, which can be solved numerically [4]. The rotation free vibrations are known as *spheroidal eigen-modes*. A general spheroidal eigen-mode Ψ^s , is given by Eq.(3), which is the superposition of all particular rotation free solutions of Eq.(2) with traction free boundary conditions, note that we use polar coordinates.

$$\Psi^s = \sum_{n=0}^{\infty} \sum_{m=-n}^n [a_n(r) \hat{e}_r + b_n(r) R \nabla] Y_{nm}(\theta, \phi) \quad (3)$$

with R the sphere radius, $a_n(r)$ and $b_n(r)$ are dimension-less radial eigen-functions determined by the boundary conditions and Y_{nm} are spherical harmonics. For each value of n the frequency equation contains infinitely many solutions. The lowest frequency for $n = 2$, is the one that is most important for the GRAIL detector. A general quadrupole mode of the sphere can be described as the superposition of the five modes with $n = 2$ ($m = -2 \dots 2$). Since (according to GR) these modes interact with gravitational radiation, they are the primary subject of our investigations. The applicability of this analytical model fails as soon as we perturb the sphere in any way, although it provides us with a useful gauge for a simulation program. Let's take a look at an elegant numerical model which is less restrictive.

1.2 An Elegant Numerical Method for Eigen-mode Analysis

If we write down the Lagrangian (Eq.(4)) for a freely vibrating elastic object (following [8])⁴

$$L = \int_V \left(\frac{1}{2} \rho \omega^2 u_i u_i - \frac{1}{2} c_{ijkl} u_{i,j} u_{k,l} \right) dV \quad (4)$$

with 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

$$\delta L = \int_V (\rho \omega^2 u_i + c_{ijkl} u_{k,lj}) \delta u_i dV - \int_S (\nu_j c_{ijkl} u_{k,l}) \delta u_i dS = 0 \quad (5)$$

Eq.(5) expresses exactly the elastic wave equation (Eq.(2)), with traction free boundary conditions, that we solved in the previous section. If we expand the displacement vector in a set of basis functions according to Eq.(6),

$$\mathbf{u} = \sum_{lmn} a_{lmn} \Phi_{lmn} \quad (6)$$

with $\Phi_{lmn} = x^l y^m z^n$ and truncate the expansion to a certain order of accuracy, Eq.(5) can be rewritten in terms of a generalized eigen value problem (Eq.(7)).

⁴ the summation convention is used where appropriate

$$\omega^2 E \mathbf{a} = \Gamma \mathbf{a} \quad (7)$$

The matrices E and Γ can be computed by integrating Eq.(5) properly. Solving for ω provides us with estimates of the eigen frequencies of the system in question.

VALIDATION

Table 1 shows several eigen-mode frequencies that are found using the numerical method as well as the corresponding values that can be found analytically, for a copper sphere. N denotes the truncation order of the basis function expansion, and $R \times R$ is the matrix size of the corresponding eigen value problem. We have used a LAPACK [1] routine to solve the system.

| 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 |

Table 1. The analytical and "numerical" frequencies in Hz for a few modes of a copper sphere.

SPHEROIDAL PERTURBATIONS

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

$$\frac{x^2}{d_x^2} + \frac{y^2}{d_y^2} + \frac{z^2}{d_z^2} = 1 \quad (8)$$

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. 2 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, again for a copper sphere. We can observe three lines that intersect at $d_z = 1.5$. Since the spherical symmetry is broken in the z -direction the five-fold degeneracy is only partially removed. It splits into two doublets and one singlet. Additional perturbations in x and y directions would break the symmetry completely leaving us with 5 different frequencies.




Fig. 2. The dependence of the quadrupole frequencies under variation of d_z .




Fig. 3. The dependence of the frequencies of the quadrupole modes versus the radius of the suspension hole. The five modes split into two doublets and a singlet.

1.3 A Finite Element Model

The numerical method also has been applied for a sphere with a suspension hole. Again, it can be observed that the five frequencies are split up, see Fig. 3. However, it leaves us with an uncomfortable feeling, since the sharp edges of the suspension hole possibly disrupt the numerical process, therefore we don't know if the splitting is an artifact or physically correct. Furthermore, we are stuck with a variety of questions that can simply not be answered by the analytical and numerical approaches, above. Therefore, we resort to the well known Finite Element (FE) method.

We discretize our antenna model in first order tetrahedral elements (each element has four nodal points). As an example, Fig. 4 shows the hull of a sphere with hole and suspension bar, consisting of 39684 tetrahedral elements. Now we can choose to approach the problem by means of modal analysis, that is, we have to solve the following eigenvalue problem.

$$[K - \omega^2 M]d = 0, \quad (9)$$

with K the system's stiffness matrix, M the mass matrix, d an eigen-mode and ω again an eigen frequency of the system. Although it is theoretically possible to investigate a linear system, such as ours, on the basis of modal analysis, the practical realization is a different story. First of all the model resolutions that we aim for are very high, necessitating the solution of huge eigen value 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 over 100 eigen modes at least is necessary to characterize a general low frequency solution of the system. Furthermore, it is difficult to characterize the effects that are due to coupling to external forces that are a-periodic, like seismic noise entering via the suspension and stochastic sources. Finally, the simulation model is likely to be extended such that it can incorporate non-linear effects, which can not be realized in a modal analysis. Therefore, we choose for explicit time integration of the equation of motion (Eq.(10)). Although it has some drawbacks, compared to modal-analysis, since 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.

Antenna-Model




Fig. 4. The hull of a sphere with suspension-rod discretized into 39684 tetrahedral elements.

$$M\mathbf{a} + C\mathbf{v} + K\mathbf{d} + \mathbf{F} = 0 \quad (10)$$

The matrices M , C and K are constructed by assembling the local mass, viscosity and stiffness matrices of all individual finite elements in which the antenna model is divided and \mathbf{F} is an external force vector. In order to integrate this system explicitly in time we adopt the following Newmark [9] scheme (11).

$$\begin{aligned} \mathbf{d}_{n+1} &= \mathbf{d}_n + \mathbf{v}_n \Delta t + \mathbf{a}_n \frac{(\Delta t)^2}{2} \\ \mathbf{a}_{n+1} &= M^{-1} K(\mathbf{d}_{n+1}) \\ \mathbf{v}_{n+1} &= \mathbf{v}_n + (\mathbf{a}_n + \mathbf{a}_{n+1}) \frac{\Delta t}{2}, \end{aligned} \quad (11)$$

with \mathbf{a}_n , \mathbf{v}_n and \mathbf{d}_n vectors that respectively describe the acceleration, velocity and displacement fields in the FE system at time n . Viscous damping (C) and external forces (\mathbf{F}) can easily be added to this algorithm. The mass matrix is diagonal, since the masses are considered to be "lumped" at the nodal points. Therefore M^{-1} can be calculated fast. For numerical stability Δt has to be taken to be less than the Courant value, since the time integration scheme is conditionally stable. The choice of algorithm on the one hand is motivated by the fact that it displays good numerical properties (low numerical dispersion and damping). This was found from a number of exploratory experiments on a 1-dimensional finite element model (data not shown). On the other hand it is very well suited for incorporation of parallelism as we will see below.

It was expected that the required accuracy with which the eigen frequencies can be determined using this method necessitates the use of a very high resolution FE model, which

consequently, is quite computationally and memory intensive. Therefore the simulation code was designed and implemented for parallel execution.

2 Parallelizing the explicit FE method from Scratch

In good tradition, we have chosen to base our parallelization method on decomposition of the FE mesh. The consequences of the decomposition method for the time-stepping algorithm as well as for the calculation of system parameters like energy had to be considered in the design. In our discussion we will only consider the pure finite element system, consisting of an ideal sphere. In the full blown simulation a variety of phenomena can be taken into account, like transducer mounting, suspension rod and sources of vibration. At this point it suffices to say that the chosen parallelization method doesn't have major consequences for the hybrid (full fledged antenna) simulation, and therefore can be left out of this discussion. We focus on two important aspects of the parallelization process: parallel time-stepping (which essentially is the kernel of the program) and the parallel calculation of the system's energy. The code has been developed in C with PVM message passing primitives, utilizing the SPMD paradigm. The process with identifier 0 is used to read in the mesh, and is used for I/O, a sort of *primus inter pares* among the parallel processes.

2.1 Parallel Energy Calculations

In order to check if the system energy remains conserved if the antenna is not influenced by any external source we have to monitor the energy in the system. The total energy E_{tot} is given by:

$$E_{tot} = \frac{1}{2}(\mathbf{v}^T M \mathbf{v} + \mathbf{d}^T K \mathbf{d}) , \quad (12)$$

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. Firstly, a dual graph of the mesh is constructed, which then can be partitioned using any graph-partitioning method [5]. Nodal points that lie on shared boundaries are replicated on each process that has this boundary in its local domain, whereas elements belong to a unique process. This procedure is schematically depicted in Fig. 5 for a simple quadrilateral mesh.




Fig. 5. Partitioning of a quadrilateral mesh in four parts, nodal points at shared boundaries are replicated while elements are assigned to unique processes.

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 as follows (analogously for the mass matrix).

$$K = k_1 \oplus k_1 \oplus \dots \oplus k_N \quad (13)$$

When we monitor the energy values during a simulation run, it would be nice if this could be done completely separable, that is, for each sub-domain (i), the energy content (E_{tot}^i) is calculated, and once in a while we assemble the sub-domain energy contributions and accumulate them on process 0. This can be done as follows:

$$E_{tot}^i = \frac{1}{2}(\mathbf{v}^{iT} M_r^i \mathbf{v}^i + \mathbf{d}^{iT} K_i \mathbf{d}^i) \quad (14)$$

\mathbf{v}^i denotes the velocity vector for the nodes in sub-domain i , \mathbf{d}^i is the displacement field, K_i is the stiffness matrix of sub-domain i , which is simply the assembly of its local element stiffness matrices and finally M_r^i is an adjusted version of the (diagonal) mass matrix of sub-domain 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, denoting how many processes have a replica of it. For example in Fig. 5 nodal point 5 has a redundancy value of 4, therefore the entry for this point 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 double counting any energy contributions.

2.2 Parallel Time Evolution

If we have constructed the mass and stiffness matrices for each sub-domain (or process) correctly, we can parallelize the Newmark scheme above as follows. Each process gets a unique process identifier (ID), equal to the sub-domain 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 processors 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 rest of the processes only may see it as B-type. Consequently, process 0 can only have shared nodes of type A, while on the other hand process $P-1$ (assuming we have P processes) can have only shared nodes of type-B. The Newmark scheme is parallelized as follows.

$$\mathbf{d}_{n+1} = \mathbf{d}_n + \mathbf{v}_n \Delta t + \mathbf{a}_n \frac{(\Delta t)^2}{2}$$

Send \mathbf{d}_{n+1} from A-nodes to B-nodes
 Receive \mathbf{d}_{n+1} for B-nodes from A-nodes
 $\mathbf{a}_{n+1} = M^{-1}K(\mathbf{d}_{n+1})$
 Send \mathbf{a}_{n+1} from B-nodes to A-nodes
 Receive \mathbf{a}_{n+1} for A-nodes from B-nodes
 Add contributions to \mathbf{a}_{n+1} from B-nodes
 $\mathbf{v}_{n+1} = \mathbf{v}_n + (\mathbf{a}_n + \mathbf{a}_{n+1}) \frac{\Delta t}{2}$

The A-B nodal point scheme has been adopted from Lonsdale *et al.* [3]. This construction parallelizes the time-stepping kernel correctly.

3 Experiments

In this section we show a selection of the experimental results that have been produced for the validation of our parallel FE simulation. First, we investigate if the model is physically correct. Next, we measure the performance benefits that we get from exploiting parallelism for different mesh resolutions. We consider two different graph partitioning methods for decomposing the meshes; recursive spectral bisection (RSB) [5] and a method that generalizes the concept of orthogonal recursive bisection (ORB) to partitioning cardinalities that are not necessarily a power of 2. The performance measurements are done on two target platforms: a Parsytec PowerXplorer (32 PowerPC's, with each 32 Mbyte of RAM) and a Parsytec CC (40 PowerPC+'s, each having 96 Mbyte of RAM). Our target material in the experiments below is the alloy CuAl (90-10), while we only consider spheres of radius 1.5 m.

3.1 Physical Correctness

Figure 6 displays the x -displacement versus time of an arbitrary nodal point (in this case number 1713) on the surface of a FE sphere consisting of 79556 elements, which vibrates in a pure quadrupole mode, that is, each nodal point in the system gets an initial velocity $\mathbf{v}^T = (x, -y, 0)$. The beats are produced by the superposition of the principal quadrupole frequencies that are slightly split up. Table 2 enumerates the average of frequencies that are found around the analytical eigen frequency f_0 ($f_0 = 780.25$ Hz) and the deviation in % of the average value quadrupole mode for 10 different FE models consisting of n tetrahedral elements. We can observe that the main frequency found for the highest resolution model approaches the analytical frequency up to 0.15%. Figure 7 shows us the evolution of the total energy for two FE models that are initiated with a quadrupole field. Clearly, the high resolution model has a lower fluctuation in energy during the simulation. Furthermore it approaches the analytical value better.





Fig. 6. The time trajectory in the x -direction of an arbitrary nodal point at the surface of a FE model of 79556 elements, initiated with a pure quadrupole velocity field.

Fig. 7. The energy evolution for two FE models that are initiated with a pure quadrupole field. Model resolutions are 9412 and 79556 respectively.

| n | $\langle f \rangle$ | $\sigma(f)$ in % |
|-------|---------------------|------------------|
| 84 | 7.910e+02 | 8.720e-01 |
| 124 | 7.922e+02 | 1.695e+00 |
| 698 | 8.011e+02 | 4.868e-01 |
| 1086 | 7.976e+02 | 4.437e-01 |
| 2724 | 7.899e+02 | 3.231e-01 |
| 2851 | 7.965e+02 | 3.623e-01 |
| 6114 | 7.862e+02 | 2.250e-01 |
| 9412 | 7.853e+02 | 3.012e-01 |
| 51501 | 7.821e+02 | 2.847e-01 |
| 79556 | 7.817e+02 | 1.550e-01 |

Table 2. The average frequency and the approximated error in % for the quadrupole principal mode for various model resolutions.

3.2 Parallel Performance

Next, consider parallel performance of the simulation kernel. We apply ORB (partitioning cardinalities of 2, 4, 6, 8, 9, 12, 16, 18, 24, 27, 32) and RSB (partitioning cardinalities of 4, 8, 16, 32) respectively to 4 different meshes of size 6114, 9412, 51501 and 79556. Figure 8 displays the single kernel execution times that were obtained for simulations on the PowerXplorer for varying number of processors. Figure 9 shows similar results obtained on the Parsytec CC. The execution times are best-case, that is, energy calculations and for example external force calculations (modelling gravitational radiation sources) are left out. Both figures are shown in log-log style, to clarify the regions where the performance gain is linear.

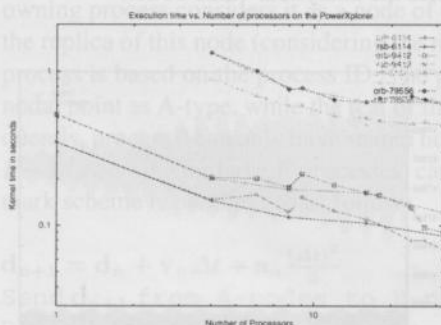


Fig. 8. The kernel execution time on the PowerXplorer for 4 different mesh sizes and two different decomposition methods (ORB and RSB) versus the partitioning cardinality.

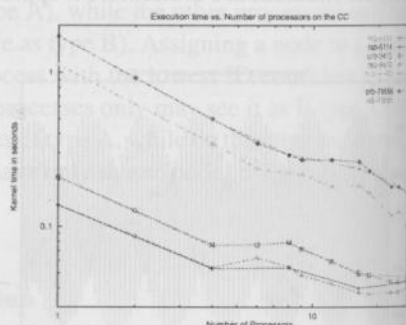


Fig. 9. The kernel execution time on the CC for 4 different mesh sizes and two different decomposition methods (ORB and RSB) versus the partitioning cardinality.

4 Discussion and Conclusions

4.1 Summary

We have constructed three consecutively more refined simulation models for the calculation of the vibrational properties of a gravitational radiation antenna. The first two models were shown to be too restrictive for our purposes, which forced us to build a third model based on the FE method. The design of the corresponding simulation tool was tailored for distributed memory computer systems, since we expected that the required simulation accuracy would necessitate significant computing power. In more detail we have discussed the consequences of domain decomposition as the parallelization strategy for the FE kernel.

4.2 Discussion

Table 2 indicates that the dominating frequency for FE simulations initiated with a pure quadrupole mode, approaches the analytical frequency value for increasing model resolution. Since the FE models always are slightly asymmetric some additional modes will be present of which the frequency has shifted from the analytical value. However, the deviation of these frequencies from the analytical value becomes smaller with the model resolution. Fig. 6 shows the x -displacement of an arbitrary nodal point at the sphere surface in the simulation for the highest resolution model. We can observe the beats, that result from the frequency splitting. Figure 7 shows us the evolution of the total energy for two FE models that are initiated with a quadrupole field. Clearly, the high resolution model has a lower fluctuation in energy during the simulation. Furthermore it approaches the analytical value better. The energy values that are found are completely independent of the number of processors used, which enforces our confidence in the correctness of our implementation.

In Figs. 8 and 9 we can observe that the difference in execution speed between ORB and RSB is not significant. In fact ORB consequently is slightly better than RSB for partitioning cardinalities that are powers of two. Both methods create compact 3-dimensional sub domains not containing any disconnected parts, leading to an average surface-volume ratio of these sub-domains 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 with the PowerXplorer, where the single processor performance can not be determined for the high-resolution problems. The CC communication network as well as the CPU is significantly faster than that of the PowerXplorer, which accounts for the fact that the kernel execution time for each problem instance is shorter on the CC.

For the PowerXplorer the kernel execution times are not significantly decreased for more than 16 partitions, while in case of the CC increasing partitioning cardinality consequently leads to faster execution times, at least in case of the high resolution models. Note that we have considered best case performance, leaving for example energy calculations out. Therefore we can expect that the absolute performance of the simulation code in a more realistic simulation (taking other phenomena into account) will degrade compared to the best case and consequently, it will result into better "scalability" of the code, since the calculation - communication ratio of the application increases.

4.3 Conclusions

We have seen that application of high performance computing methodology to our simulation can lead to a significant performance gain. Large models can now be simulated efficiently on parallel systems that are built from commodity processors like the PowerPC. Typical simulation runs require up to 100,000 time steps (kernel iterations), necessitating execution times of the order of one to two weeks per simulation run on an Ultra-Sparc (left aside the fact that the problem requires a significant amount of memory, which for the high resolution models becomes a bottleneck on a single processor, as we could see for the PowerXplorer). On the two parallel machines in our study, the simulation time can be reduced to the order of a few hours, by fully exploiting the problem's parallelism.

The simulation program will be used for the feasibility study of calculating the effects of a variety of different phenomena, enumerated in the introductory section. In the next project phase it will be utilized in the analysis and validation of the signals that are picked up by the GRAIL antenna.

Acknowledgments

The work presented in this paper was made possible by funding of NWO in the GRAIL project fase 0 feasibility study.

References

1. E. Anderson et al. *LAPACK Users' Guide*, second edition, 1994.
2. B.D. Kandhai, P.M.A. Sloot, and J.P. Huot. Constrained migration of an atmospheric circulation model. In B.Hertzberger H.Liddel, A.Colbrook and P.M.A. Sloot, editors, *high performance computing and networking*, number 1067 in ISBN 3-540-61142-8, pages 269-275. Springer-Verlag, April 1996.
3. G. Lonsdale, J. Clinckemallie, S. Vlachoutsis, J. F. de Ronde, P. M. A. Sloot, N. Floros, and J. Reeve. Crashworthiness simulation migration to distributed memory, mimd machines. In *Conference on Supercomputing Applications in the Automotive Industries of the 26th ISATA*, September 1993.
4. A.E.H. Love. *A Treatise on the Mathematical Theory of Elasticity*. Dover Publications, 1944.
5. H. D. Simon. Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering*, 2(2/3):135-148, 1991.
6. P.M.A. Sloot. Modelling for parallel simulation: Possibilities and pitfalls, invited lecture. In *Eurosim '95, Simulation congress*, pages 29-44, Amsterdam, the Netherlands, 1995.
7. P.M.A. Sloot 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.
8. T. M. Bell W. M. Visscher, A. Migliori and R. A. Reinert. On the normal modes of free vibration of inhomogeneous and anisotropic elastic objects. *Journal of the Acoustic Society America*, 90(4):2154-2162, October 1991.
9. O.C. Zienkiewicz and R.L. Taylor. *The Finite Element Method*. McGraw-Hill Book company, 1994.