European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS 2004
P. Neittaanmäki, T. Rossi, K. Majava, and O. Pironneau (eds.)
J. Périaux and M. Křížek (assoc. eds.)
Jyväskylä, 24–28 July 2004

PARALLELIZED UWVF METHOD FOR 3D HELMHOLTZ PROBLEMS

Tomi Huttunen*, Peter Monk†, and Jari P. Kaipio*

*Department of Applied Physics
P.O. Box 1627, FIN-70211 University of Kuopio, Finland
e-mails: tomi.huttunen@uku.fi, jari.kaipio@uku.fi,
web page: http://physics.uku.fi/ huttunen

† Department Mathematical Sciences
University of Delaware
Dewark, DE 19716, USA
e-mail: monk@math.udel.edu

Key words: Helmholtz equation, ultra weak variational formulation, perfectly matched layer, parallel computing

Abstract. *We investigate the parallelized ultra weak variational formulation (UWVF) method for large-scale 3D Helmholtz problems. The unbounded Helmholtz problem is truncated using the perfectly matched layers (PML). We propose a method to partition the problem in a balanced way and examine the scalability of the parallel UWVF method. The method is evaluated with numerical simulations that are performed on a PC cluster.*

1 INTRODUCTION

The use of oscillatory basis functions has opened new possibilities for solving large-scale wave problems. In articular, during the last decade, a variety of such techniques have been proposed. An extension of the finite element method, called the partition of unity finite element method (PUM) [1], has been successfully used with plane wave basis functions to solve Helmholtz problems at high wave numbers [16, 19]. Plane waves are also used in the discontinuous enrichment method [8], discontinuous Galerkin method [9], least-squares method [18], micro-local discretization [6], boundary element method [20] and in the ultra weak variational formulation (UWVF) [3, 4]. This study is dedicated to the parallel UWVF solver for 3D Helmholtz problems.

The UWVF in a general form was first introduced in [7] and further analyzed for Helmholtz and Maxwell problems in [3, 4, 5]. Numerical simulations of the UWVF method for Helmholtz problems in inhomogeneous media have been performed in [11, 13] and an extension of the method for Navier problems is proposed in [12]. Although the UWVF has given a promise for solving wave problems with reduced computational burden, it can suffer from numerical instability and a special method for choosing the basis is needed.

As shown by 2D examples in [11], the instability of the UWVF method can be reduced by using an element-wise varying number of plane wave basis functions. In addition, the truncation of unbounded problems with the perfectly matched layer (PML) [2] can essentially improve the accuracy of UWVF approximations [13]. However, the use of variable number of basis functions in the parallel UWVF method can lead to severe imbalance in memory requirements between processors if the UWVF problem is partitioned into subproblems with equal number of elements. Furthermore, the use of PML can further increase the imbalance since a reduced number of basis functions is needed in the PML.

In this study, we show how the memory requirement of partitioned UWVF problems can be estimated based on the knowledge on the element size and material properties. A relatively simple estimate can be obtained for the memory requirement for elements in the non-PML region. We also sketch how similar method could be extended into elements in the PML.

2 ULTRA WEAK VARIATIONAL FORMULATION

In this section, the UWVF for the inhomogeneous Helmholtz problem is outlined. The truncation of an unbounded problem is done using the perfectly matched layer (PML) [2]. Only a short review of the method is given here. A more thorough presentation of the coupled UWVF-PML method can be found from [13].

Let us define a domain Ω having the boundary Γ and outward unit normal n. We assume that domain can be divided into two parts $\Omega = \Omega_1 \cup \Omega_2$ so that Ω_1 is the actual computational domain and Ω_2 is the surrounding PML region. The Helmholtz problem

in an inhomogeneous medium occupying Ω can be formulated as follows

$$\nabla \cdot \left(\frac{1}{\rho} A \nabla p\right) + \frac{\kappa^2 \eta^2}{\rho} p = 0 \quad \text{in } \Omega, \tag{1}$$

$$\left(\frac{1}{\rho} n \cdot (A \nabla p) - i\varsigma p\right) = Q\left(-\frac{1}{\rho} n \cdot (A \nabla p) - i\varsigma p\right) + g \quad \text{on } \Gamma, \tag{2}$$

where $\kappa = \kappa(r) \in \mathbb{C}$ is the spatially varying wave number that satisfies $\Re(\kappa) > 0$ and $\Im(\kappa) \geq 0$. On the boundary the parameter $Q \in \mathbb{C}$ with $|Q| \leq 1$. The density $\rho = \rho(r)$ and parameter ς are real and positive. Finally, the complex valued source function on the exterior boundary Γ is denoted by g.

In this study, the matrix A and the parameter η are either $A = I$ and $\eta = 1$ in the non-PML region Ω_1 or

$$A = \operatorname{diag}\left(\frac{d_y d_z}{d_x}, \frac{d_x d_z}{d_y}, \frac{d_x d_y}{d_z}\right) \quad \text{and} \quad \eta^2 = d_x d_y d_z \tag{3}$$

in the PML domain Ω_2. The parameters d_x, d_y and d_z arise from the complex stretching of the spatial variables so that

$$\frac{\partial x'}{\partial x} = d_x \tag{4}$$

where

$$x' = \begin{cases} x + \dfrac{i}{\kappa} \int_{x_0}^{x} \sigma_0 (|x| - x_0)^n dx, & |x| \geq x_0, \\ x, & |x| < x_0. \end{cases} \tag{5}$$

Similar expressions are used also for y and z. We also denote $r = (x, y, z)$ and $r' = (x', y', z')$.

For the UWVF we partition the domain Ω using a tetrahedral finite element mesh with elements K_k, $k = 1, \cdots, N$. We assume that the mesh is chosen so that the parameters ρ and κ are constants in each element and we denote $\rho_k \equiv \rho|_{K_k}$ and $\kappa_k \equiv \kappa|_{K_k}$. The problem (1)-(2) can now be decomposed into subproblems for each element K_k, $k = 1, \cdots, N$

$$\nabla \cdot (A_k \nabla p_k) + \kappa_k^2 \eta_k^2 p_k = 0 \qquad \text{in } K_k, \tag{6}$$

$$\frac{1}{\rho_k} n_k \cdot (A_k \nabla p_k) - i\varsigma p_k = -\frac{1}{\rho_j} n_j \cdot (A_j \nabla p_j) - i\varsigma p_j \qquad \text{on } \Sigma_{k,j}, \tag{7}$$

$$\frac{1}{\rho_k} n_k \cdot (A_k \nabla p_k) + i\varsigma p_k = -\frac{1}{\rho_j} n_j \cdot (A_j \nabla p_j) + i\varsigma p_j \qquad \text{on } \Sigma_{k,j}, \tag{8}$$

$$\left(\frac{1}{\rho_k} n_k \cdot (A_k \nabla p_k) - i\varsigma p_k\right) = Q\left(-\frac{1}{\rho_k} n_k \cdot (A_k \nabla p_k) - i\varsigma p_k\right) + g \qquad \text{on } \Gamma_k, \tag{9}$$

where $p_k = p|_{K_k}$. The coupling parameter for (7) and (8) on the interface $\Sigma_{k,j}$ and exterior boundary Γ_k is

$$\varsigma = \frac{1}{2}\left(\frac{\Re(\kappa_k)}{\rho_k} + \frac{\Re(\kappa_j)}{\rho_j}\right) \quad \text{and} \quad \varsigma = \frac{\Re(\kappa)}{\rho}, \tag{10}$$

respectively.

For the UWVF we define a new function χ_k on the skeleton of the mesh

$$\chi_k = \Big(\big(-\frac{1}{\rho_k} n_k \cdot (A_k \nabla) - i\varsigma\big) p_k\Big)\Big|_{\partial K_k}, \quad 1 \le k \le N. \tag{11}$$

It is shown in [3, 4] that χ_k satisfies the variational formulation.

$$\begin{aligned} &\sum_{k=1}^{N} \int_{\partial K_k} \frac{1}{\varsigma} \chi_k \overline{\left(-\frac{1}{\rho_k}\frac{\partial}{\partial n_k} - i\varsigma\right) v_k} - \sum_{k=1}^{N}\sum_{j=1}^{N} \int_{\Sigma_{k,j}} \frac{1}{\varsigma} \chi_j \overline{\left(\frac{1}{\rho_k}\frac{\partial}{\partial n_k} - i\varsigma\right) v_k} \\ - &\sum_{k=1}^{N} \int_{\Gamma_k} \frac{Q}{\varsigma} \chi_k \overline{\left(\frac{1}{\rho_k}\frac{\partial}{\partial n_k} - i\varsigma\right) v_k} = \sum_{k=1}^{N} \int_{\Gamma_k} \frac{1}{\varsigma} g \overline{\left(\frac{1}{\rho_k}\frac{\partial}{\partial n_k} - i\varsigma\right) v_k}, \end{aligned} \tag{12}$$

for all piecewise smooth functions v_k satisfying $\nabla \cdot (A_k \nabla \overline{v}_k) + \kappa_k^2 \eta^2 \overline{v}_k = 0$ in K_k. Terms of the double summation in the second term of the left hand side are assumed to exist only if two elements share a common face $\Sigma_{k,j}$. Equation (12) is called *the ultra weak variational formulation* of the Helmholtz problem.

The discrete UWVF is obtained by approximating the function χ_k using complex conjugated plane wave basis functions

$$\chi_k^a = \sum_{\ell=1}^{N_k} \chi_{k,\ell} \left(-\frac{1}{\rho_k} n_k \cdot (A_k \nabla) - i\varsigma\right) \varphi_{k,\ell}, \tag{13}$$

where

$$\varphi_{k,\ell} = \begin{cases} \exp(i\overline{\kappa}_k d_{k,\ell} \cdot r') & \text{in } K_k \\ 0 & \text{elsewhere,} \end{cases}$$

In addition, following the Galerkin method developed in [4], we choose $v_k = \varphi_{k,\ell}$. The angularly equidistributed directions $d_{k,\ell}$ are obtained by minimizing the maximum distance between N_k points on the unit sphere [10]. In this study, the maximum allowed number of directions per element is limited to 130.

By substituting the plane wave approximation to the equation (12), the discrete UWVF problem can be written in the form of the matrix equation

$$(I - D^{-1}C)X = D^{-1}b. \tag{14}$$

The matrix D is a block diagonal matrix arising from the first term of (12). It consists of blocks D_k, $1 \le k \le N$ so that the block D_k corresponds to the element K_k. The sparse block matrix C couples the solution in a single element to adjacent elements or includes terms arising from the boundary conditions. The data vector b corresponds to the right hand side of (12).

3 Parallel UWVF

In this section a parallelization strategy for the UWVF method is developed. First, we discuss on the partitioning of the problem for a distributed memory computers. Second, a load balancing method for the problem is proposed. The aim is to use compute an estimate for the number of basis functions in each element based on the knowledge on the material properties and the size of the element. This estimate can be used to partition the mesh so that the memory requirement for each processor is balanced.

3.1 Partitioning of the mesh

The idea behind the efficient parallelized UWVF code is that the computational mesh is first partitioned into n parts where n is the number of processors. Then each processor assembles the parts of the matrices D and C along with the right hand side b for the corresponding part of the mesh. The blocks D_k for the matrix D can be constructed using geometric and material property information of the element K_k only. However, for the blocks in the coupling matrix C, information from the elements adjacent to K_k are needed which necessitates communication between processors. Furthermore, for solving the matrix equation (14) with an iterative solver, the matrix-vector multiplication $(I - D^{-1}C)X$ is computed on each iteration which again calls for communication between subprocesses. The Bi-CGStab is used to solve (14). The parallel implementation of the Bi-CGStab methods are discussed in more detail in [21].

On one hand, it is known that the bottleneck in parallel iterations is often the amount of communication between processors. On the other hand, it is obvious that the amount of communication can be reduced with the number of faces of tetrahedra shared by different parts of the mesh. The method used in this study is based on the multilevel partitioning of the mesh [15] using the single processor version of METIS software package [17].

3.2 Load balancing

Let the element K_k be occupied by N_k basis functions. Furthermore, in the case of a tetrahedral mesh, the element has four faces shared with adjacent elements or at least one of the faces is on the exterior boundary. The number of entries in the matrices D and C of (14) for a single element is then

$$M_k = N_k^2 + \sum_{\ell=1}^{4} N_k N_{j_\ell} \tag{15}$$

where N_{j_ℓ} is the number of basis functions in the element K_j sharing the face $\Sigma_{k,j}$ or $N_j = N_k$ if the face is on exterior boundary. The storage needed for the right hand side and the vector to be solved in (14), together with vectors needed in the iterative solver, can be taken into account with an additional storage of mN_k coefficients. The value of m depends on the type of the iterative solver being used.

As discussed in [11] the method used to choose N_k in 2D examples relied on limiting the condition number of matrix blocks D_k below a predetermined limit. The number of basis functions N_k and the building of D_k was then obtained as a sequential procedure. More precisely, after an initial guess for N_k, the matrix block D_k was computed and the condition number of D_k was evaluated. This procedure was continued until the highest N_k which provided the condition number below a predetermined limit was found.

In the parallelized approach, the previous scheme that is used to choose N_k has two drawbacks. First, to use parallel processing efficiently, the mesh must be partitioned to each processor before starting of the dynamic procedure for selecting N_k. The partitioning of the mesh so that each part has approximately equal number of elements may lead to severe imbalance between loads at different processors. Namely, if the element size and material properties vary within the mesh, the number of basis functions that give an equal condition number of D_k varies significantly between elements. Second, in 2D examples the initial guess for N_k was chosen equal for each element. However, in 3D , due to the increased number of elements and number of basis function per an element, this approach becomes intolerably time-consuming. Both these drawbacks call for a method to estimate N_k prior to partitioning based on the element size and material properties in the element.

The numerical experiments in [4, 11] show that the condition number of D_k depends on the element size h_k and wave number κ_k. Let the element size h_k be defined as the average distance of element vertices $r_{k,\ell},\ \ell = 1, ..., 4$ from the centroid c_k of K_k

$$h_k = \frac{1}{4} \sum_{\ell=1}^{4} |c_k - r_{k,\ell}|. \tag{16}$$

The numerical examples of Section 4 show almost linear correlation between the number of basis function per element N_k and the dimensionless quantity $\kappa_k h_k$ (see Fig. 3). Therefore, the initial guess can be made, for example, in the form

$$N_k = A(\kappa_k h_k) + B, \tag{17}$$

where feasible parameters A and B can be estimated from the experimental cases.

The METIS allows weighed partitioning of the mesh. For load balancing, the approximately linear relation between $\kappa_k h_k$ and N_k provides means of estimating the number of matrix entries M_k corresponding to each element K_k. Furthermore, the use of M_k and mN_k to weight K_k in METIS gives a relatively balanced memory storage for all processors.

After the mesh is partitioned, the dynamic procedure introduced in [11] can be used to choose the stable number of basis functions N_k and to build the blocks D_k. A significantly reduced number of trials is needed for finding N_k due to the improved initial guess (17). However, the estimate (17) can be misleading if the medium in K_k is strongly absorbing since the effect of imaginary part of κ_k is not taken into account.

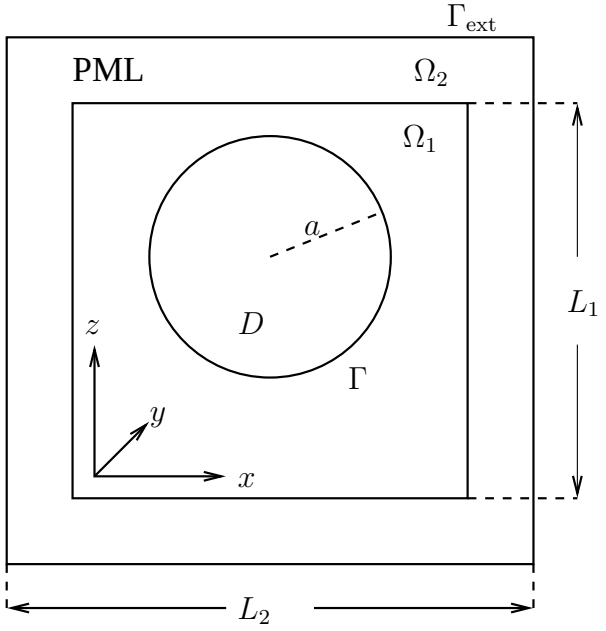

Figure 1: The spherical sound-hard scatterer D with radius $a = 1$ is located off-centered at point (0,0,0.2) inside a cubic computational domain ($L_1 = 3$). In addition, the computational domain is enclosed with a larger cube ($L_2 = 4$) and the region between the two cubes constitutes the PML. The whole computational domain, including the PML, is denoted by $\Omega = \Omega_1 \cup \Omega_2$ and the exterior boundary by Γ_{ext}. The boundary of the scatterer is denoted by Γ.

4 NUMERICAL SIMULATIONS

The numerical simulations are performed for the scattered wave field p_s from a sound-hard sphere with radius $a = 1$, see Fig. 1. The incident plane wave p_i propagates in the direction of the positive z axis and the boundary condition of the surface of the scatterer Γ is

$$\frac{\partial p}{\partial n} = \frac{\partial p_i}{\partial n} + \frac{\partial p_s}{\partial n} = 0,$$

where p_s is the scattered field.

The Fourier series solution for p_s in the spherical coordinates (R, θ, ϕ) can be computed as [14]

$$p_s(R, \theta, \phi) = p_0 \sum_{n=0}^{\infty} i^n (2n+1) \frac{j'_n(\kappa a)}{h'_n(\kappa a)} P_n(\cos\theta) h_n(\kappa R),$$

where j'_n and h'_n are derivatives of the nth order spherical Bessel and Hankel functions. The Legendre polynomials of order n are denoted by P_n and p_0 is the complex valued amplitude of the plane wave. The series is truncated when the changes in the solution due to an additional Fourier mode is below 10^{-6}.

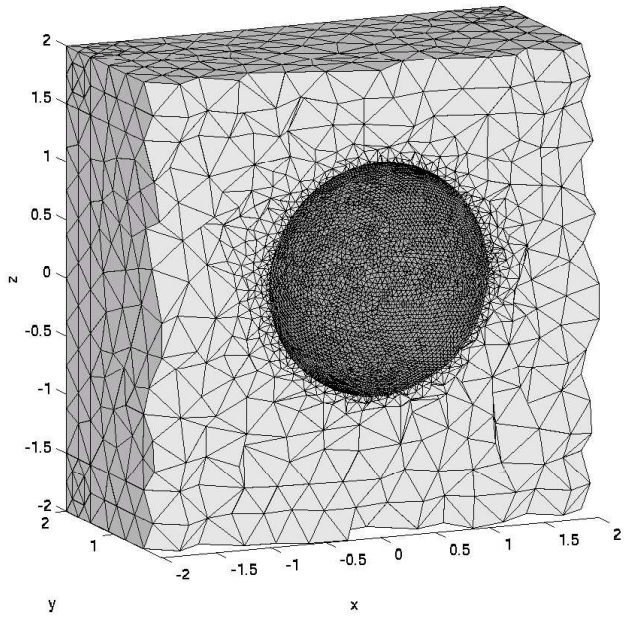

Figure 2: The mesh used in all simulations of this study. Relatively dense discretization is used near the scatterer. The PML consists about 2 layers of elements.

The UWVF approximation of the problem is obtained by setting $Q = 1$ and $g = \frac{1}{2}p_i$ on Γ. On the exterior boundary of the PML region Γ_{ext} we use $Q = 0$ and $g = 0$. In all simulations, the decay parameter in the PML is $\sigma_0 = 5$. The Bi-CGStab iterations are terminated when the relative residual is below 10^{-5}. The computer code used in the simulations is coded with Fortran90 and parallelized using MPI (Message Passing Interface). The computations are done with a Beowulf PC cluster consisting of 24 2.6 GHz Pentium 4 processors and having 48.0 GB total RAM. The processors are connected with a 1GB ethernet switch.

All simulations of this study are computed in the mesh of Fig. 2. It consists of 19264 vertices and 84515 tetrahedra. To reduce the error arising from the triangulation of the curved surface, we use relatively dense discretization near the scatterer. However, the element size increases rapidly with the distance from the sphere. Using the indicator (16) for the element size, for the largest element in the mesh we have $\max(h_k)$=0.2658.

In Fig. 3, we present the number of basis functions per element as a function of the dimensionless quantity $\kappa_k h_k$. The results are computed for three different value of the maximum condition number of the blocks D_k inside the non-PML domain. The results are computed using $\kappa a = 10\pi$. The figure shows the highest number of basis functions in each element which provides the condition number that is below the predetermined value.

A somewhat linear relationship between the condition number and $\kappa_n h_n$ can be observed. A linear least-squares fit to the data of Fig. 3 for the curve (17) is given in Table

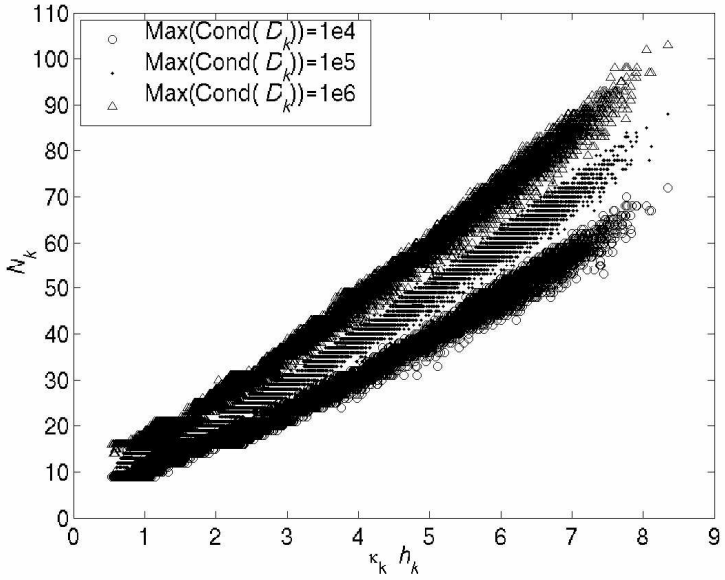

Figure 3: The number of basis functions per element N_k as function of product of the wave number κ_k and element size h_k for different condition number limits. The results are computed for $\kappa a = 10\pi$.

1. We also list the relative error of the UWVF approximation for each case.

Max(Cond(D_k))	A	B	Err %
10^4	7.27	3.17	2.94
10^5	8.74	5.12	1.21
10^6	10.76	5.34	0.96

Table 1: The parameters A and B of the equation (17) that are obtained a linear least-squares fit to the data of Fig. 3. In the last column we show the error of the UWVF approximation for each case.

We examined next the storage distribution between processors when estimates from Table 1 are used. A test problem ($\kappa a = 10\pi$) is computed using the linear initial guess for N_k from Table 1 in the case of Max(Cond(D_k))=10^5 for the non-PML elements. Since notably less basis functions is needed in the PML elements [13], a constant initial guess $N_k = 10$ is used in the PML elements. In Fig. 4 we show the normalized memory requirement for 11 of 12 processors used in the simulation. The processor 1 is not included since is has additional storage such as the whole computational mesh. The results of Fig. 4 indicate that an improved method for load balancing is needed for elements in the PML. Such a method will be outlined later in this study.

We also studied the scalability of the UWVF solver. In Fig. 5 the computation time is plotted against the number of processors in the case of $\kappa a = 4\pi$ and Max(Cond(D_k)) = 10^5. Again, the results are computed using the initial guess of the form (17) where

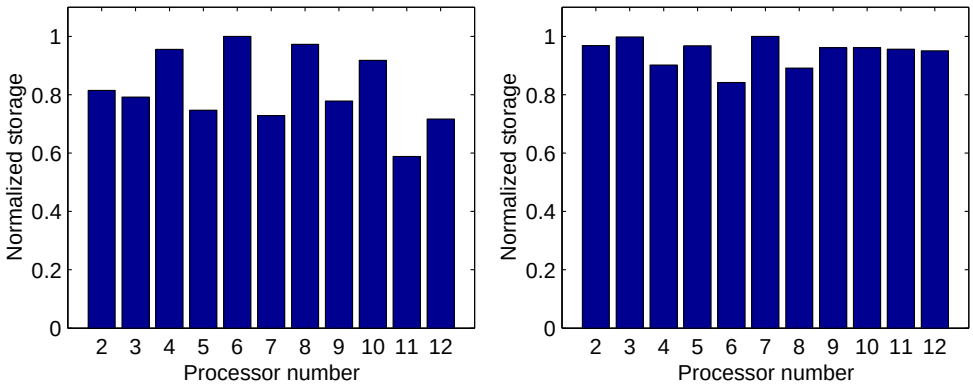

Figure 4: Normalized load in each processors when estimate for Max(Cond(D_k))=10^5 from Table 1 is used in the non-PML elements. The results are computed with $\kappa a = 10\pi$. Left: The PML has decay parameter $\sigma_0 = 5$ and the initial guess $N_k = 10$ is used in the PML elements. Right: The PML has decay parameter $\sigma_0 = 0$ (i.e. all elements are non-absorbing) and the initial guess for Max(Cond(D_k))=10^5 from Table 1 is used for all elements.

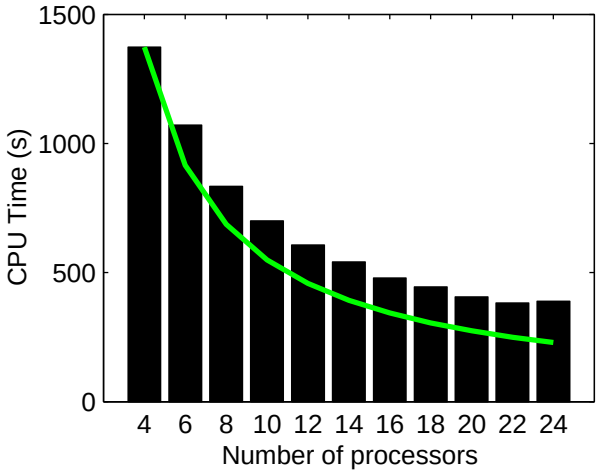

Figure 5: Bars show the scalability of the parallel UWVF solver for the scattering problem $\kappa a = 4\pi$. The solid line represents the ideal scalability, i.e. the CPU time is inversely proportional to the number of processors.

$A = 8.74$ and $B = 5.12$ is used in the non-PML region while a constant initial guess $N_k = 10$ is used in the PML.

The performance of the UWVF solver as a function of wave number was investigated next. In Fig. 6 the accuracy of the UWVF approximation is shown as a function of the parameter κa. A slight increase in error can be observed as the frequency of the problem increases. For the largest wave number the error jumps. The jump is most likely due to limit in the highest allowed number of basis functions per element (=130) which is reached in many element at the highest wave number.

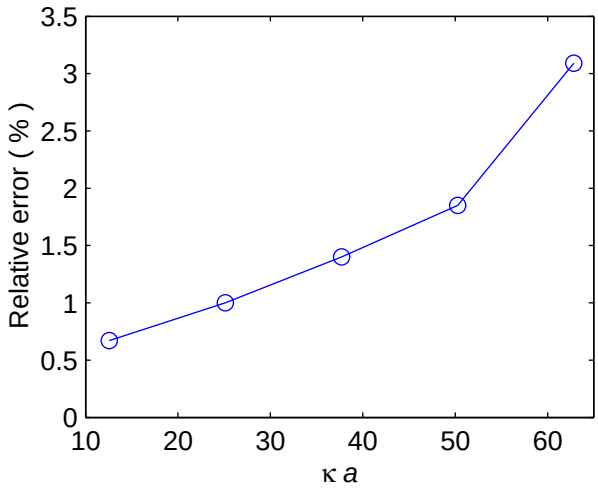

Figure 6: The accuracy of the UWVF approximation as the function of κa in the mesh of Fig. 2 and limiting the maximum condition number of D_k below 10^5.

Finally, we studied whether it could be possible to estimate the number of basis functions in the PML elements based on the elements size and material parameters. The numerical simulations of this study and the simulations presented in [13] indicate that the condition number of the block D_k depends also on the absorption in the element K_k. In Fig. 7 we present the number of basis function in each element in the PML as a function of $\kappa_k h_k$ and a decay parameter α_k. The decay amplitude for the element K_k is computed as

$$\alpha_k = \sqrt{\frac{1}{3}\sum_{j=1}^{3}[\kappa_k \Im(\xi_j')]^2}$$

at the centroid of the element. Here $\xi_j' = x', y', z'$ are the stretched spatial functions given by the equation (5).

5 CONCLUSIONS

We investigated the performance of the parallel 3D UWVF solver for Helmholtz problems at high wave numbers. The results indicate that it is possible to obtain a balanced partitioning of the problem by using the knowledge of the element size and material properties in each element. Although the partitioning method of this study was applied for problems in non-absorbing media, the final part of the study (Fig. 7) suggest that a similar approach can be extended also for absorbing media and elements in the PML. Furthermore, the results suggest that a stable number of basis functions for each element could be obtained directly using element and material properties without "the trial and error" method used in this study. This would notably speed up the computational procedure.

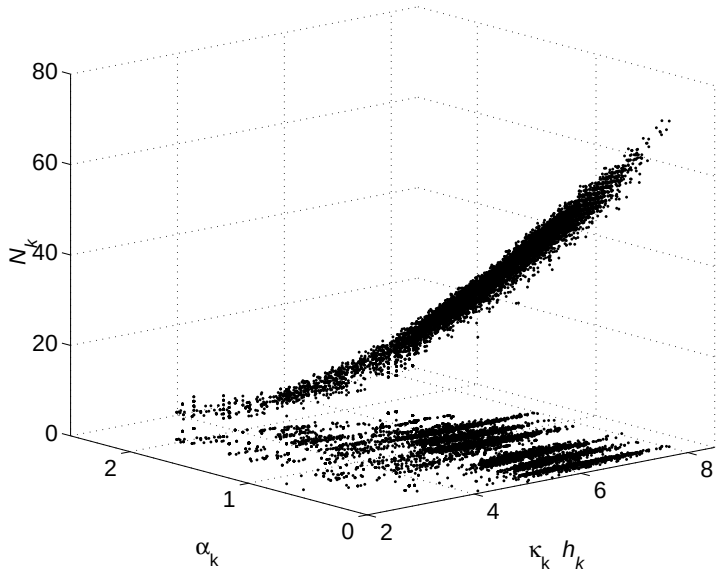

Figure 7: The number of basis functions per element N_k as a function of $\kappa_k h_k$ and decay amplitude α_k when the condition number $\mathrm{Max}(\mathrm{Cond}(D_k))$ is limited below 10^5 and $\kappa a = 10\pi$. Dots in the $(\kappa_k h_k, \alpha_k)$-plane are projections of the data points to that plane.

REFERENCES

[1] I. Babuška and J.M. Melenk. The partition of unity method. *International Journal for Numerical Methods in Engineering*, 40:727–758, 1997.

[2] J. Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114:185–200, 1994.

[3] O. Cessenat. *Application d'une nouvelle formulation variationelle des equations d'ondes harmonigues, Problemes de Helmholtz 2D et de Maxwell 3D.* PhD thesis, Paris IX Dauphine, 1996.

[4] O. Cessenat and B. Després. Application of an ultra weak variational formulation of elliptic PDEs to the two-dimensional Helmholtz problem. *SIAM Journal of Numerical Analysis*, 35(1):255–299, 1998.

[5] O. Cessenat and B. Després. Using plane waves as base functions for solving time harmonic equations with the ultra weak variational formulation. *Journal of Computational Acoustics*, 11(2):227–238, 2003.

[6] A. de La Bourdonnaye. Convergence of the approximation of wave functions by oscillatory functions in the high frequency limit. *Comptes Rendus de l'Academie des Science. Serie I*, 318:765–768, 1994.

[7] B. Després. Sur une formulation variationnelle de type ultra-faible. *Comptes Rendus de l'Academie des Sciences - Series I*, 318:939–944, 1994.

[8] C. Farhat, I. Harari, and L. Franca. The discontinuous enrichment method. *Computer Methods in Applied Mechanics and Engineering*, 190:6455–6479, 2001.

[9] C. Farhat, I. Harari, and U. Hetmaniuk. A discontinuous Galerkin method with Lagrange multipliers for the solution of Helmholtz problems in the mid-frequency regime. *Computer Methods in Applied Mechanics and Engineering*, 192:1389–1419, 2003.

[10] R.H. Hardin, N.J.A. Sloane, and W. D. Smith. Spherical coverings, May 1997. Available on: http://www.research.att.com/ njas/coverings/index.html.

[11] T. Huttunen, P. Monk, and J.P. Kaipio. Computational aspects of the ultra-weak variational formulation. *Journal of Computational Physics*, 182:27–46, 2002.

[12] T. Huttunen, P. Monk, F. Collino, and J.P. Kaipio. The ultra weak variational formulation for elastic wave problems. *SIAM Journal on Scientific Computing*, 2003. Accepted.

[13] T. Huttunen, J.P. Kaipio, and P. Monk. The perfectly matched layer for the ultra weak variational formulation of the 3D Helmholtz problem. 2004. *International Journal for Numerical Methods in Engineering*, Accepted.

[14] F. Ihlenburg. *Finite Element Analysis of Acoustic Scattering*. Springer, 1998.

[15] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1998.

[16] O. Laghrouche, P. Bettess, and R.J. Astley. Modelling of short wave diffraction problems using approximating systems of plane waves. *International Journal for Numerical Methods in Engineering*, 54:1501–1533, 2002.

[17] G. Karypis. METIS - family of multilevel partitioning algorithms, June 2003. Available on: http://www-users.cs.umn.edu/ karypis/metis/.

[18] P. Monk and D. Wang. A least squares method for the Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering*, 175:121–136, 1999.

[19] P. Ortiz and E. Sanchez. An improved partition of the unity finite element model for diffraction problems. *International Journal for Numerical Methods in Engineering*, 50:2727–2740, 2001.

[20] E. Perrey-Debain, J. Trevelyan, and P. Bettess. Plane wave basis in integral equation for 3D scattering. In G.C. Cohen, E. Heikkola, P. Joly, and P. Neittaanmäki, editors, *Mathematical and Numerical Aspects of Wave Propagation - Waves 2003*, pages 292–297. Springer, 2003.

[21] L. T. Yang and R. P. Brent. The improved BiCGStab method for large and sparse unsymmetric linear systems on parallel distributed memory architectures. In *Proceedings of the Fifth International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP-02)*, pages 324–328, 2002.