 CMES, vol.58, no.2, pp.131-158, 2010

Comparison of the Fast Multipole Method with Hierarchical Matrices for the Helmholtz-BEM

D. Brunner[1], M. Junge[1], P. Rapp[1], M. Bebendorf[2] and L. Gaul[1]

Abstract: The simulation of the hydroacoustic sound radiation of ship-like structures has an ever-growing importance due to legal regulations. Using the boundary element method, the overall dimension of the problem is reduced and only integrals over surfaces have to be considered. Additionally, the Sommerfeld radiation condition is automatically satisfied by proper choice of the fundamental solution. However, the resulting matrices are fully populated and the set-up time and memory consumption scale quadratically with respect to the degrees of freedom. Different fast boundary element methods have been introduced for the Helmholtz equation, resulting in a quasilinear complexity. Two of these methods are considered in this paper, namely the fast multipole method and hierarchical matrices. The first one applies a series expansion of the fundamental solution, whereas the second one is of pure algebraic nature and represents partitions of the original system matrix by low-rank approximations in outer-product form. The two methods are compared for a structure, which is partly immersed in water. The memory consumption, the set-up time and the time required for a matrix-vector product are investigated. Different frequency regimes are considered. Since the diagonal multipole expansion is known to be unstable in the low-frequency regime, two types of expansions are necessary for a wideband analysis.

Keywords: Acoustics, Boundary Element Method, Fast Multipole Method, Hierarchical Matrices, Adaptive Cross Approximation

1 Introduction

Nowadays, the sound quality has an ever-growing influence on the overall impression of a product. Depending on the size and complexity of the product, it may be cumbersome to perform experimental investigations. Therefore, the application of simulation tools is mandatory. Within this paper, the focus is on computing

[1] Institute of Applied and Experimental Mechanics, University of Stuttgart, Germany.
[2] Institute for Numerical Simulation, University of Bonn, Germany

the hydro-acoustic sound radiation of surface ships. Starting point is the time-harmonic Helmholtz equation (Kinsler, Frey, Coppens, and Sanders, 2000). For complex structures, discretization methods are necessary to investigate the sound radiation of vibrating bodies. Besides the finite element method (Zienkiewicz and Taylor, 2000; Bettess, 1992), the boundary element method (BEM) is well suited for exterior acoustic problems with unbounded domains (Bonnet, 1995; Wu, 2000; Gaul, Kögl, and Wagner, 2003). As main advantage, the Sommerfeld radiation condition is automatically satisfied at infinity (Nédélec, 2001; Steinbach, 2008). For simulating the hydro-acoustic sound radiation of surface ships, the pressure-free water surface has to be taken into account. This is favorably done using a half-space fundamental solution (Seybert and Soenarko, 1988; Seybert and Wu, 1989; Sladek, Tanaka, and Sladek, 2001). However, a big drawback of the classical BEM is the fully populated matrices resulting in a quadratical expense for setting-up and storing them.

In the last decades, several fast boundary element methods have been published, including the fast multipole method, hierarchical matrices and the panel clustering (Aimi, Diligenti, Lunardini, and Salvadori, 2003). One of the most popular technique is the fast multipole method (FMM). It is based on a series expansion of the fundamental solution and was first published by Rokhlin (1985) for the Laplacian. Later, a diagonal version for the Helmholtz operator followed (Rokhlin, 1993). In the early papers of Coifman, Rokhlin, and Wandzura (1993); Epton and Dembart (1995); Rahola (1996); Gyure and Stalzer (1998); Amini and Profit (1999), the complexity and implementation of the method is discussed. The development stage of the FMM at this time is summarized in the overview papers of Nishimura (2002) and Darve (2000). The latter one proposes efficient strategies for the memory management and interpolation schemes. Another introduction is given by Gumerov, Duraiswami, and Borovikov (2003). Their textbook presents an up-to-date discussion of the FMM including different expansions and convergence analyses (Gumerov and Duraiswami, 2004). The application to fluid-structure coupling is reported by Fischer (2004). An extension to half-space problems is discussed by Brunner (2009).

As pointed out by Greengard, Huang, Rokhlin, and Wandzura (1998), the diagonal form may become unstable in the low-frequency regime. They therefore propose a low frequency version, which uses a splitting in an evanescent and a propagating plane wave. A different approach is taken by Darve and Harvé (2004), who propose a stable plane wave expansion which is applicable over the whole frequency regime. The procedure is similar to the one published by Hu and Chew (2001), except that they use a different integration path. A combined wideband scheme is discussed by Cheng, Crutchfield, Gimbutas, Greengard, Ethridge, Huang, Rokhlin, Yarvin,

and Zhao (2006), which switches between different representations for low and high frequencies. Gumerov and Duraiswami (2007) propose a similar procedure based on a rotation – coaxial translation – back rotation scheme. As an advantage, no interpolation and filtering is needed any more. Furthermore, they use a recursion scheme to efficiently compute the coaxial translation coefficients (Gumerov and Duraiswami, 2003).

Besides acoustical problems, the FMM is applied to Poisson-type equations (He and ad S.-P. Lim, 2008), corrosion problems (Aoki, Amaya, Urago, and Nakayama, 2004) and for the analysis of the mechanical properties of composites (Wang and Yao, 2005). A review concerning the application of the FMM to electromagnetics is given by Chew, Song, Cui, Velamparambil, Hastriter, and Hu (2004).

A different kind of compression for the BE matrices is obtained by using hierarchical matrices ($\mathscr{H}$-matrices). They approximate subblocks of an appropriate partitioning of a given matrix using low-rank approximations. The general framework of $\mathscr{H}$-matrices goes back to Hackbusch (1999), who introduces $\mathscr{H}$-matrices and the basic algebraic operations. He shows that matrix-vector and matrix-matrix operations can be performed with almost linear complexity for a certain class of matrices. Early research is also reported within the papers of Tyrtyshnikov (1998) and Bebendorf, Rjasanow, and Tyrtyshnikov (2000). An algorithm called adaptive cross approximation (ACA) is proposed by Bebendorf (2000) and Bebendorf and Rjasanow (2003) to find the low-rank approximation within the $\mathscr{H}$-matrix while only using some of the entries of the BE matrix. An advantage of the ACA over the FMM is the easy parallelization of the algorithm, see Bebendorf and Kriemann (2005). Additionally, it is much simpler to implement ACA. An extensive introduction to $\mathscr{H}$-matrices is given by Bebendorf (2008) and Rjasanow and Steinbach (2007). As an advantage over the FMM, not only the near-field matrix but also an approximation of the global matrix can be used for preconditioning. A hierarchical LU decomposition-based preconditioner is discussed by Bebendorf (2005).

A comparison of the FMM and ACA is done for the Laplace equation by Buchau, Rucker, Rain, Rischmüller, Kurz, and Rjasanow (2003). A non-linear magnetostatic problem with BE-FE coupling is solved. The set-up of the FMM near-field matrix turns out to be faster than the initialization time of the $\mathscr{H}$-matrix. However, $\mathscr{H}$-matrices show their strength with the very fast computation of a matrix-vector product. A similar comparison is done by Forster, Schrefl, Dittrich, Scholz, and Fidler (2003), who compare a treecode algorithm with the $\mathscr{H}$-matrices for a simple model problem. They conclude, $\mathscr{H}$-matrices need more time during the set-up phase, but are much faster when computing matrix-vector products. A comparison of the ACA with the FMM for the Helmholtz equation is done by Brancati, Aliabadi, and Benedetti (2009), who use a partially pivoted ACA approach in com-

bination with the collocation method. The ACA turned out to be faster than the FMM. However, the used ACA and FMM versions were implemented by different authors and hence use different integration routines. Additionally, they do not apply a stable method like the Burton-Miller approach.

Within this paper, the FMM is compared with $\mathscr{H}$-matrices for both the low-frequency and the high-frequency regime using the Burton-Miller approach. The first scenario typically occurs, if the fluid-structure coupled system for a partly immersed ship is investigated. Usually, the finite element method is applied for the structure using shell elements. A fine discretization of the structure is needed to resolve the local vibration phenomena. Since the feedback of the pressure onto the structure cannot be neglected for fluids with a high density, a fully coupled system needs to be solved (Brunner, 2009). In case of a conforming coupling scheme, the boundary element discretization of the fluid typically has more than ten elements per wavelength. For this scenario, a low-frequency FMM version is chosen. In the second case, a discretization with ten elements per wavelength is assumed. As will be shown, in this high-frequency case, the diagonal form of the FMM is applicable. Throughout this paper, only the fluid part is investigated using fast boundary element methods. For a discussion of the fluid-structure coupled problem, see Brunner, Junge, Wilken, Cabos, and Gaul (2009).

The paper is organized as follows: In section two, the investigated acoustical half-space problem is presented. The third section gives an introduction to the BEM which is applied for the acoustical problem. In the subsequent fourth section, different fast BE methods are presented. The FMM is discussed in a high- and low-frequency version. Thereafter, the concept of $\mathscr{H}$-matrices is introduced. In the fifth section, numerical results are presented for a cylindrical test structure. Finally a conclusion is given.

2 The Acoustical Problem

In this section, the investigated acoustical problem is pointed out. Throughout this paper, all variables are assumed to have the time-harmonic behavior $\mathrm{e}^{-\mathrm{i}\omega t}$, where $\omega = 2\pi f$ is the circular frequency. Within this work, a vibrating body is assumed to be immersed in an exterior acoustic half-space Ω_a as visualized in Fig. 1.

The water within the half-space domain is modeled using the Helmholtz equation. On the Neumann boundary Γ_N, the acoustic flux $q := \frac{\partial p}{\partial n}$ is assumed to be prescribed. For reverberant surfaces, the flux is directly linked to the normal velocity v_n by the Euler's equation $q = \rho \mathrm{i}\omega v_n$, where ρ is the density of the fluid. The water surface is assumed to be pressure-free. The corresponding boundary value problem

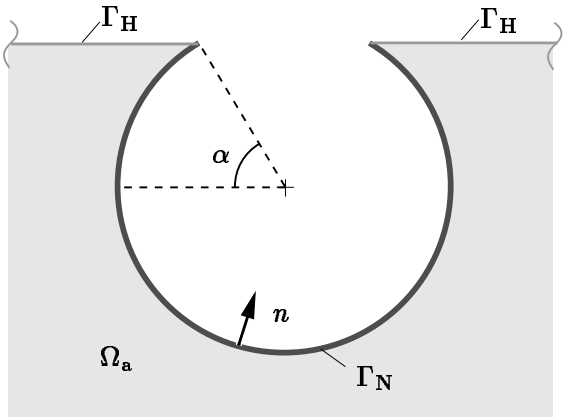

Figure 1: Half-space boundary value problem, where a body is partly immersed in a semi-infinite fluid domain Ω_a.

reads

$$\begin{aligned} \Delta p(x) + \kappa^2 p(x) &= 0 \qquad \text{for } x \in \Omega_a, \\ \frac{\partial p(x)}{\partial n_x} &= \bar{q}(x) \quad \text{for } x \in \Gamma_N, \\ p &= 0 \qquad \text{for } x \in \Gamma_H, \end{aligned} \tag{1}$$

where Δ denotes the Laplacian, $\kappa = \frac{\omega}{c}$ is the circular wavenumber with the speed of sound c of the fluid. Additionally the Sommerfeld radiation condition

$$\left| \frac{\partial p}{\partial r} - \mathrm{i}\,\kappa\, p \right| \leq \frac{C}{|r|^2} \quad \text{for } r = |x| \to \infty \tag{2}$$

has to be satisfied for this exterior problem.

3 Boundary Element Formulations

Acoustics is one of the areas where the boundary element method is most powerful. The main idea of the BEM is to transform the problem from a domain Ω to its boundary $\Gamma = \partial\Omega$. Hence, only a two-dimensional surface has to be discretized instead of a three-dimensional domain. The key idea is to use Green's second identity in combination with the screening property of the Dirac distribution. This way, the pressure p at an arbitrary point x within the exterior acoustic half-space domain Ω_a is given by the integral representation

$$p(x) = \int_{\Gamma_N} P(x,y)\, q(y)\, \mathrm{d}s_y - \int_{\Gamma_N} \frac{\partial P(x,y)}{\partial n_y}\, p(y)\, \mathrm{d}s_y, \qquad x \in \Omega_a. \tag{3}$$

The fundamental solution $P(x,y)$ depends on the type of the problem. For the half-space problem Eq. 1,

$$P(x,y) = \frac{\mathrm{e}^{\mathrm{i}\kappa r}}{4\pi r} - \frac{\mathrm{e}^{\mathrm{i}\kappa r'}}{4\pi r'} \tag{4}$$

is chosen with $r = |x-y|$ and $r' = |x'-y|$. Here, the mirror image of x with respect to the half-space plane is denoted by x'. For exterior acoustic problems, this choice ensures that the Sommerfeld radiation condition as stated in Eq. 2 is automatically fulfilled at infinity.

As pointed out before, this so-called representation formula is only valid for points x within the acoustic domain Ω_a. Shifting the point x onto the boundary Γ_N by a limit process yields the so-called boundary integral equation (BIE)

$$\frac{1}{2}\,p(x) = \underbrace{\int_{\Gamma_\mathrm{N}} P(x,y)\,q(y)\,\mathrm{d}s_y}_{=:(Vq)(x)} - \underbrace{\int_{\Gamma_\mathrm{N}} \frac{\partial P(x,y)}{\partial n_y}\,p(y)\,\mathrm{d}s_y}_{=:(Kp)(x)}, \quad x \in \Gamma_\mathrm{N}, \tag{5}$$

where $(Vp)(x)$ denotes the single layer potential and $(Kp)(x)$ is the double layer potential. In the same way, the hypersingular boundary integral equation (HBIE) is derived. Here, an additional derivative with respect to n_x has to be taken into account, yielding

$$\frac{1}{2}\,q(x) = \underbrace{\int_{\Gamma_\mathrm{N}} \frac{\partial P(x,y)}{\partial n_x}\,q(y)\,\mathrm{d}s_y}_{=:(K'q)(x)} - \underbrace{\frac{\partial}{\partial n_x}\int_{\Gamma_\mathrm{N}} \frac{\partial P(x,y)}{\partial n_y}\,p(y)\,\mathrm{d}s_y}_{=:-(Dp)(x)}, \quad x \in \Gamma_\mathrm{N}, \tag{6}$$

where $(K'p)(x)$ denotes the adjoint double layer potential and $(Dp)(x)$ is the hypersingular operator.

However, neither Eq. 5 nor Eq. 6 has a unique solution for all frequencies in case of an exterior problem. They fail at the resonance frequencies of the corresponding interior Dirichlet or Neumann problem. To overcome this drawback, the Burton-Miller approach is applied (Burton and Miller, 1971) which uses a linear combination of Eq. 5 and Eq. 6 with a coupling factor α. The Galerkin method is applied throughout this paper. A triangulation of the boundary is introduced. The resulting system of linear equations reads

$$\underbrace{\left(\frac{1}{2}\mathbf{M} + \mathbf{K} - \alpha\mathbf{D}\right)}_{\mathbf{K}_\mathrm{BE}}\mathbf{p} = \underbrace{\left(\mathbf{V} - \frac{1}{2}\alpha\mathbf{M}' + \alpha\mathbf{K}'\right)}_{\mathbf{C}_\mathrm{BE}}\bar{\mathbf{q}}, \tag{7}$$

where the bold matrices correspond to the Galerkin discretizations of the operators introduced above and $\mathbf{M}$, $\mathbf{M}'$ denote mass matrices, respectively. For inwards pointing normals, the factor $\alpha = -\frac{i}{\kappa}$ is chosen. The pressure is approximated using piecewise linear shape functions and the flux is approximated using piecewise constant shape functions. The weighting is done with piecewise linear functions. The numerical integration of the resulting double surface integrals is done according to the rules proposed by Sauter (1992). A regularization based on Stokes theorem is applied for the hypersingular operator.

As a disadvantage of classical boundary element methods, the matrices $\mathbf{K}_{\mathrm{BE}}$ and $\mathbf{C}_{\mathrm{BE}}$ are fully populated and setting up and storing the entries have a numerical complexity of order $\mathcal{O}(N^2)$ for a problem with N unknowns. If additionally a direct solver is applied for Eq. 7 even a cubical complexity results. For this work, an iterative solver is chosen. In what follows, a GMRES (Saad, 2003) is applied, where the numerical complexity is predominantly influenced by the complexity of the matrix-vector product. Hence, for classical boundary element methods, a complexity of order $\mathcal{O}(N^2)$ is achieved. This can be improved to $\mathcal{O}(N\log^k N)$ with small k by applying fast boundary element methods, which offer a matrix-vector product at almost linear expense.

4 Fast Boundary Element Methods

All fast boundary element methods try to reduce the numerical complexity in order to obtain an almost linear expense with respect to the number of unknowns. In this paper two of them are applied and therefore discussed in more detail.

4.1 Fast Multipole Method

One of the most commonly applied fast boundary element methods is the fast multipole method. It is based on a series expansion of the fundamental solution. Depending on the frequency regime of interest, different expansions are known which produce accurate results. In what follows, a high-frequency and a low-frequency version are considered.

4.1.1 High-Frequency FMM

Usually, a discretization of six to ten elements per wavelength is chosen for simulating the sound radiation of vibrating bodies. If the ratio of elements per wavelength is not significantly larger, a diagonal form of the multipole expansion is applicable. As is shown later, the diagonal expansion becomes unstable for a much larger ratio.

For the implementation of the fast multipole method, a hierarchical cluster tree is used, see Fig. 2. The clustering starts with a box containing all the elements. The

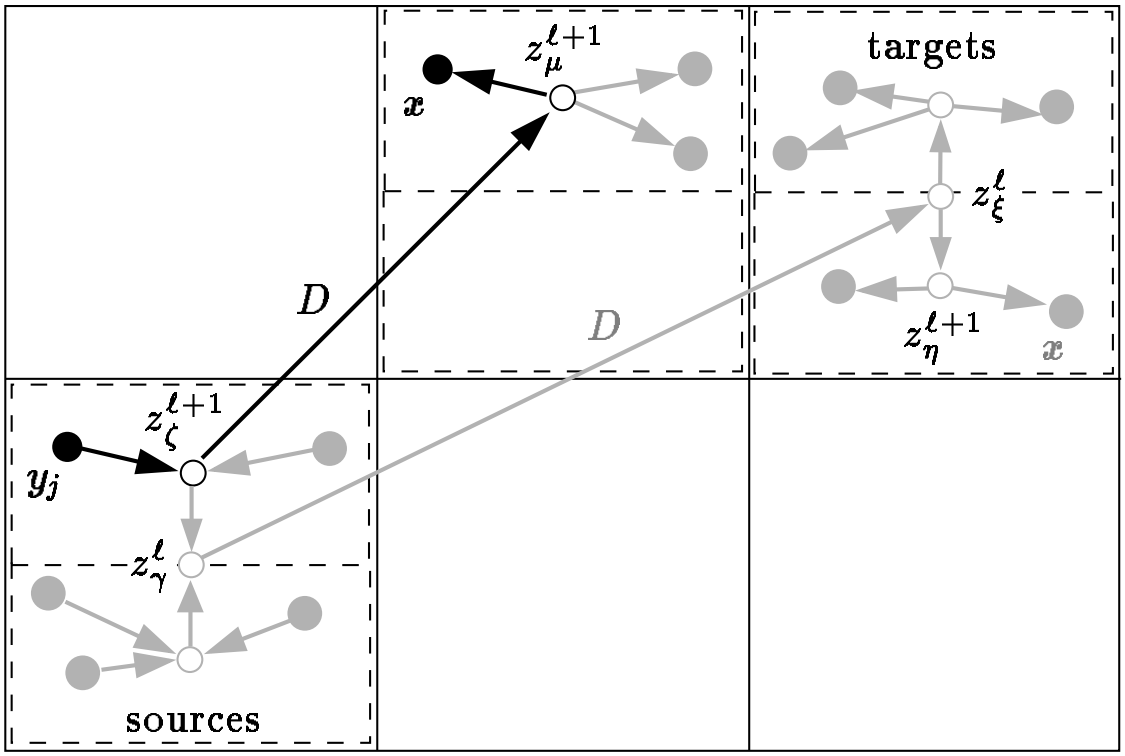

Figure 2: Translation of the information in a multilevel cluster tree. The clusters on level ℓ are plotted with solid lines whereas clusters on level $\ell+1$ are visualized with dashed lines.

set of the corresponding indices is the so-called root cluster C_1^0. The clusters $C_\zeta^{\ell+1}$ on the next higher level $\ell+1$ are built by bisection of the corresponding box where the boundary elements τ_k are assigned to the clusters $C_\zeta^{\ell+1}$ with respect to their geometric center point. The geometric center point of a cluster C_γ^ℓ is denoted by z_γ^ℓ and the radius is r_γ^ℓ. All clusters are split along their dominant dimensions. The clusters $C_\zeta^{\ell+1}$ are called the sons of the father cluster C_γ^ℓ. The bisection process is repeated recursively and stops for clusters which have less than a prescribed number of elements. Within this work, this number of elements is set to twelve. Clusters which are not split any more are called leaf-clusters and are denoted by $\bar{C}_\gamma^\ell$.

The key idea is to split the distance vector between y_j and x into different parts which allows to write the fundamental solution in product form. The interaction is first considered for clusters on the same level which are represented by the indices μ and ζ in Fig. 2. The first vector $z_\zeta^{\ell+1} - y_j$ is local to the point y_j.The second vector D only depends on the clusters' centers. The third vector $x - z_\mu^{\ell+1}$ is local to the point x. Using this splitting, the series expansion of the free-space fundamental

solution reads (Rahola, 1996)

$$\frac{\mathrm{e}^{\mathrm{i}\kappa|x-y_j|}}{4\pi|x-y_j|} = \frac{\kappa\mathrm{i}}{(4\pi)^2}\int_{\mathbb{S}} \mathrm{e}^{\mathrm{i}\kappa(x-z_\mu^{\ell+1})\cdot\hat{s}} \underbrace{\sum_{l=0}^{L}(2l+1)\mathrm{i}^l\, h_l^{(1)}(\kappa|D|)\, P_l(\hat{s}\cdot\hat{D})}_{=:M_L(z_\zeta^{\ell+1},z_\mu^{\ell+1},\hat{s})} \mathrm{e}^{\mathrm{i}\kappa(z_\zeta^{\ell+1}-y_j)\cdot\hat{s}}\,\mathrm{d}\omega_s\,, \tag{8}$$

where $\hat{s} = (\sin\Theta\cos\Phi, \sin\Theta\sin\Phi, \cos\Theta)$ are the far-field directions on the unit sphere $\mathbb{S}$ and $\hat{D} := {}^{D}\!/_{|D|}$. The Hankel functions $h_l^{(1)}$ of first kind and order l as well as the Legendre polynomial P_l are defined as done by Rahola (1996). M_L is called translation operator with the expansion length L. Since the boxes of the cluster tree have the same dimensions on a given level, translations with an identical M_L occur several times within the multipole cycle. Hence, these translation operators only need to be stored once. According to Gyure and Stalzer (1998) L is estimated by the semi-empirical rule

$$L = \kappa d_{\max} + c_{\mathrm{e}}\log(\kappa d_{\max} + \pi)\,, \tag{9}$$

where $d_{\max}$ is the maximum cluster diameter on a given level ℓ and c_{e} denotes a constant which controls the accuracy of the expansion. The integration over the unit sphere is performed numerically (Rokhlin, 1993)

$$\int_{\mathbb{S}} g(\hat{s})\mathrm{d}\omega_s = \sum_{i=0}^{2L-1}\sum_{j=0}^{L-1} w_j \frac{\pi}{m}\, g(\hat{s}_{ij})\,, \tag{10}$$

where $\Theta_j = \arccos x_j$ and $\Phi_i = \frac{i\pi}{m}$. Here, x_j and w_j denote the abscissae and weights of the Gaussian quadrature.

Since the truncated expansion is only valid for well separated points, one has to distinguish between a far-field (FF), where the expansion is admissible and a near-field (NF), which has to be treated in the classical way. A cluster C_ζ^ℓ is in the near-field of cluster C_μ^ℓ, if the condition

$$|\mathrm{z}_\zeta^\ell - \mathrm{z}_\mu^\ell| \le c_{\mathrm{n}}\max\{r_\zeta^\ell, r_\mu^\ell\} \tag{11}$$

is fulfilled. The near-field constant c_{n} has to be chosen by the user.

In what follows, the evaluation of the matrix-vector product $\mathbf{w} = \mathbf{V}\bar{\mathbf{q}}$ is explained exemplary for the single-layer potential operator as introduced above. First, the multipole algorithm is described for a free-space problem. The extension to a half-space problem is discussed thereafter. For the FMM, the matrix-vector product

reads

$$w_m \approx \sum_{k\in \mathrm{NF}(m)} \tilde{\mathbf{V}}[m,k]\, q_k + \frac{\kappa \mathrm{i}}{(4\pi)^2} \sum_{\eta | m\in \tilde{C}_\eta^{\ell+1}} \int_{\Gamma_\eta} \nu_m(x) \int_{\mathbb{S}} \mathrm{e}^{(z_\eta^\ell - x)\cdot \hat{s}} N_\eta^\ell(\hat{s}) \mathrm{d}\omega_s \, \mathrm{d}s_x \,, \tag{12}$$

where $\tilde{\mathbf{V}}$ is the sparse near-field matrix of $\mathbf{V}$ and $\nu(x)$ are the weighting functions. The expression $\sum_{\eta|m\in\tilde{C}_\eta^{\ell+1}}$ implies that the sum runs over all leaf-clusters η, to which the node m belongs to. The so-called near-field signature $N_\eta^\ell(\hat{s})$ in Eq. 12 is computed as follows:

1. Compute the far-field signature $F(\hat{s})$ for every leaf cluster

$$F_\zeta^{\ell+1}(\hat{s}) = \sum_{k\in\tilde{C}_\zeta^{\ell+1}} \int_{\tau_k} \mathrm{e}^{\mathrm{i}\kappa(z_\zeta^{\ell+1}-y)\cdot\hat{s}} q_k \, \mathrm{d}s_y \,. \tag{13}$$

2. Translate the far-field signature to all interaction clusters using the translation operators M_L to obtain the near-field signatures $N(\hat{s})$

$$N_\mu^{\ell+1}(\hat{s}) = N_\mu^{\ell+1}(\hat{s}) + \sum_{C_\zeta^{\ell+1}\in \mathrm{IL}(C_\mu^{\ell+1})} M_L(z_\zeta^{\ell+1}, z_\mu^{\ell+1}, \hat{s})\, F_\zeta^{\ell+1}(\hat{s}) \,, \tag{14}$$

   The interaction list is formed by all clusters whose father clusters satisfy the near-field condition, but which are themselves not in each others near-fields.

3. Shift $F(\hat{s})$ to the parent cluster using

$$F_\gamma^\ell(\hat{s}) = \sum_{C_\zeta^{\ell+1}\in \mathrm{sons}(C_\gamma^\ell)} \mathrm{e}^{(z_\gamma^\ell - z_\zeta^{\ell+1})\cdot\hat{s}} F_\zeta^{\ell+1}(\hat{s}) \,. \tag{15}$$

4. Repeat the last two steps until the interaction list is empty.

5. Go the opposite direction and shift $N(\hat{s})$ to the child clusters using

$$N_\eta^{\ell+1}(\hat{s}) = N_\eta^{\ell+1}(\hat{s}) + \mathrm{e}^{(z_\eta^{\ell+1} - z_\xi^\ell)\cdot\hat{s}} N_\xi^\ell(\hat{s}) \tag{16}$$

   until a leaf clusters is reached. In this case, recover the solution using Eq. 12.

Please note: An interpolation or filtering strategy has to be applied in step 3 and 5 for shifting $F(\hat{s})$ and $N(\hat{s})$ if the expansion lengths on the two levels differ (Gyure and Stalzer, 1998).

In case of half-space problems, the half-space fundamental solution Eq. 4 has to be applied instead of the fundamental solution used in Eq. 8. The first term of Eq. 4 is represented by the expansion Eq. 8. Obviously, in the second term of Eq. 4 the point x is simply mirrored at the half-space plane. Hence, it turned out that a mirroring of the whole geometry with respect to the half-space plane results in an efficient implementation of the half-space FMM (Brunner, 2009). The cluster tree also needs to include the mirrored part now. Hence it contains twice as many clusters plus the new root cluster. As an advantage, the multipole expansion mentioned above is directly applicable. Step 1 only has to be done for the original clusters below the half-space plane. In step 2, the interactions from original clusters to mirrored clusters have to be included (only in this direction). Step 3 only has to be done for the original clusters. Step 5 has to be performed for the original and the mirrored clusters. After the solution has been recovered for a mirrored cluster, the result has to be subtracted from the corresponding non-mirrored entry. This strategy results in a good far-field compression for the mirrored clusters. A detailed description of the half-space FMM is found in the paper of Brunner, Of, Junge, Steinbach, and Gaul (2010).

As can be seen in Eq. 8, this expansion is efficient, since only one translation operator has to be applied for the conversion of a far-field signature to the corresponding near-field signature. For this reason, the expansion is called diagonal. As every signature consists of $2L^2$ coefficients (see Eq. 10), the overall complexity of the operations is of order $\mathscr{O}(L^2)$.

4.1.2 Low-Frequency FMM

In the low-frequency regime, the ratio of elements per wavelength is much greater than 10. Within this frequency regime, the diagonal form of the multipole expansion becomes unstable (Nishimura, 2002; Gumerov and Duraiswami, 2009) and is therefore not applicable. The series expansion described below follows the derivation given by Gumerov and Duraiswami (2004).

To simplify the notation, so-called singular spherical basis functions

$$\mathscr{S}_n^m(x) = h_n(\kappa r_x) Y_n^m(\theta_x, \varphi_x) \tag{17}$$

and regular spherical basis functions

$$\mathscr{R}_n^m(x) = j_n(\kappa r_x) Y_n^m(\theta_x, \varphi_x) \tag{18}$$

are introduced. As above, h_n denotes spherical Hankel functions, j_n are spherical Bessel functions and Y_n^m are spherical harmonics. If a point y is close to the expansion center and the evaluation point x is far away, the original Helmholtz fundamental solution can be written using the h-expansion or multipole expansion

$$\frac{\mathrm{e}^{\mathrm{i}\kappa|x-y|}}{4\pi|x-y|} \approx \mathrm{i}\kappa \sum_{n=0}^{N_t} \sum_{m=-n}^{n} \mathscr{M}_n^m(y)\, \mathscr{S}_n^m(x)\,, \tag{19}$$

where N_t is the expansion length to be specified later. The corresponding multipole coefficients are defined by

$$\mathscr{M}_n^m(y) = \mathscr{R}_n^{-m}(y). \tag{20}$$

In contrast to this, the j-expansion or local expansion is valid, if point x is close to the expansion center and the source point y is far away

$$\frac{\mathrm{e}^{\mathrm{i}\kappa|x-y|}}{4\pi|x-y|} \approx \mathrm{i}\kappa \sum_{n=0}^{N_t} \sum_{m=-n}^{n} \mathscr{L}_n^m(y)\, \mathscr{R}_n^m(x)\,, \tag{21}$$

where the local coefficients are denoted by $\mathscr{L}_n^m$. Using Eq. 19 or Eq. 21, the fundamental solution is split in a part which is only a function of x and a second one which is only a function of y. An important aspect for the FMM is the fast shifting of the expansion centers of the coefficients and the transformation of the multipole coefficients to the local coefficients:

To obtain an efficient multipole algorithm a rotation-coaxial translation-back rotation (RCR) scheme is applied, where three consecutive mappings are performed. This way, the overall operations for the FMM translations have a numerical complexity of order $\mathscr{O}(N_t^3)$. First a rotation around the expansion center has to be performed by applying rotation coefficients. Using the recursions proposed by Gumerov and Duraiswami (2003), the rotation coefficients themselves are computed for the cost of $\mathscr{O}(N_t^3)$. Formally the same transformation holds for the local coefficients $\mathscr{L}_n^m$. At the next step, coaxial translations have to be performed. Note, the translations are always in positive z-axis direction. The transformation between the $\mathscr{M}$ and $\mathscr{L}$ coefficients additionally converts the multipole coefficients to local coefficients. As before, setting up the coefficients and performing the transformation scales with order $\mathscr{O}(N_t^3)$. Hence, the overall RCR-scheme can be computed at cubical expense, too. A derivation of the translation and rotation coefficients is far beyond this paper. The interested reader is referred to the detailed paper of Gumerov and Duraiswami (2003). For the Burton-Miller formulation as stated in Eq. 7, also the directional derivative has to be computed. The corresponding recursion scheme is discussed by Gumerov and Duraiswami (2007).

As for the high frequency FMM, a near-field has to be set up by the classical BEM. The same near-field condition Eq. 11 is applied. However, for the RCR scheme cubical boxes which yield the same regular distances in every direction are advantageous. Therefore, an octree is applied instead of the bisectioning. Hence, only a minimal number of rotation and translation coefficients which scale with order $\mathscr{O}(N_t^3)$ have to be computed and stored. In contrast to this, the translation operator M_L for the HF FMM scales with order $\mathscr{O}(L^2)$. Choosing the bisection for the HF FMM has the advantage, that the far- and near-field conversion can be shifted to a coarser level and the overall conversions are reduced. This may speed up the computation.

The expansion length is estimated according to the recommendation of Gumerov and Duraiswami (2007)

$$N_t = \sqrt[4]{p_{\text{low}}^4 + p_{\text{high}}^4} \tag{22}$$

with

$$p_{\text{low}} = \frac{1}{\ln \gamma} \ln \frac{1}{\varepsilon(1-\gamma^{-1})^{3/2}} + 1 \quad , \quad p_{\text{high}} = \kappa a + \frac{1}{2}\left(3 \ln \frac{1}{\varepsilon}\right)^{2/3} \sqrt[3]{\kappa a}. \tag{23}$$

where $\gamma = 2$ is chosen and a denotes the radius of the smallest sphere surrounding a box on the corresponding level.

The overall multipole cycle is very similar compared to the one of the high frequency FMM. Instead of the far-field signatures $F(\hat{s})$, the multipole coefficients $\mathscr{M}_n^m$ have to be computed. The shifts to other levels and the translations to the interaction clusters are carried out by the RCR-scheme. This finally yields the local coefficients $\mathscr{L}_n^m$, which are used with Eq. 21 to recover the solution. For the implementation of the half-space problem, the same mirror technique as mentioned above is applied.

Obviously, the asymptotic complexity of the operations of the low frequency FMM is of order $\mathscr{O}(N_t^3)$ which is worse compared to $\mathscr{O}(L^2)$ for the high frequency FMM. Therefore, a combined scheme of the two approaches on the different levels of the FMM is proposed by Gumerov and Duraiswami (2007) and a similar one by Cheng, Crutchfield, Gimbutas, Greengard, Ethridge, Huang, Rokhlin, Yarvin, and Zhao (2006). In this paper, such a wideband scheme is not implemented, since either the very low- or high frequency regime is of interest.

4.2 Hierarchical Matrices

In contrast to the fast multipole method, $\mathscr{H}$-matrices do not require an expansion of the kernel. This procedure is based on low-rank approximations of matrix-

partitions. Only the integration routines are necessary which provide single entries of the system matrices.

4.2.1 Outer-Product Form and Low-Rank Approximation

The rank rk of a matrix $\mathbf{A}$ is defined as the number of linearly independent columns or rows of $\mathbf{A}$. For a matrix $\mathbf{A} \in \mathbb{C}^{m\times n}$ it holds

$$rk(\mathbf{A}) \leq \min\{m,n\}\,.$$

If the equals sign holds, matrix $\mathbf{A}$ is said to have full rank. In contrast to this, matrices $\mathbf{A} \in \mathbb{C}_k^{m\times n}$ consist of at most k linearly independent rows and columns and thus $rk(\mathbf{A}) \leq k$ holds. These matrices are called rank-k matrices. As alternative to the traditional entrywise representation, such a matrix can be stored in the outer-product form

$$\mathbf{A} = \sum_{i=1}^{k} \mathbf{u}_i \mathbf{v}_i^{\mathrm{H}} = \mathbf{U}\mathbf{V}^{\mathrm{H}}\,,$$

where the superscript H denotes the complex conjugate transpose. The storage requirement in outer-product form is $k(m+n)$ entries in contrast to mn entries in the traditional way. Hence, the outer-product form only makes sense, if k is small. A matrix is said to be of low rank, if the condition $k(m+n) < mn$ is satisfied. These matrices are favorably stored in outer-product form.

This also leads to an improvement of the matrix-vector multiplication $\mathbf{y} = \mathbf{A}\mathbf{x}$. In the traditional representation, a total of mn dominant arithmetic operations are required. In the outer-product form, the matrix-vector product is represented by

$$\mathbf{A}\mathbf{x} = \mathbf{U}(\mathbf{V}^{\mathrm{H}}\mathbf{x})\,. \tag{24}$$

It is efficiently evaluated in two steps

$$\mathbf{z} = \mathbf{V}^{\mathrm{H}}\mathbf{x} \qquad\qquad \mathbf{y} = \mathbf{U}\mathbf{z}\,.$$

The total number of essential operations is $k(m+n)$ for the outer-product form. For matrices of low rank, a matrix-vector multiplication is performed faster in outer-product form than in traditional entrywise representation.

4.2.2 Adaptive Cross Approximation

One possibility to find an outer-product form of a matrix is to apply a singular value decomposition (SVD). This is the best possible approximation. However, the expense of an SVD is of cubical order and therefore a SVD is too expensive. For

this reason, an algorithm called adaptive cross approximation (ACA) is applied as proposed by Bebendorf (2008). In the low frequency regime, one expects a quasi-linear complexity. The algorithm is described by the following recursive scheme: At the beginning, $\mathbf{R}^0 = \mathbf{A}$ is set. For a given pivot element with the row index i_k and column index j_k, the residual matrix after the next recursion step is

$$\mathbf{R}^{k+1} = \mathbf{R}^k - \mathbf{u}_k \mathbf{v}_k^{\mathrm{H}},$$

where the vectors $\mathbf{u}_k$ and $\mathbf{v}_k$ of the outer-product form are

$$\mathbf{u}_k = \mathbf{R}^k_{1:m,j_k} \qquad \text{and} \qquad \mathbf{v}_k^{\mathrm{H}} = \frac{1}{\mathbf{R}^k_{i_k,j_k}} \mathbf{R}^k_{i_k,1:n}.$$

The resulting approximation $\mathbf{A}_k = \sum_{\ell=1}^{k} \mathbf{u}_k \mathbf{v}_k^{\mathrm{H}}$ is a rank-k matrix. As stopping criterion, the following expression is used

$$||\mathbf{u}_{k+1}|| \, ||\mathbf{v}_{k+1}|| \leq \frac{\delta(1-\eta)}{1+\delta} ||\mathbf{A}_k||_{\mathrm{F}}. \tag{25}$$

The choice of δ and η is discussed below. As pointed out before, the best decomposition is given by the SVD. However, the SVD is too expensive to be computed for the original dense matrix. But if instead the SVD is computed using $\mathbf{A}_k$, the approximation can be further improved, see for instance Bebendorf (2005, 2008) for further details. In what follows, this is called recompression of a $\mathscr{H}$-matrix.

4.2.3 Partitioning of Matrices

So far, the approximation of a single matrix using ACA was discussed. Matrices arising from BE discretizations typically show a pronounced singularity, if a column degree of freedom is geometrically close to the row degree of freedom. An approximation of the total BE matrix is therefore not possible. Instead, ACA is only applied to matrix partitions, where the column and row degrees of freedom are far away from each other and thus do not touch the singularity. The diagonal partitions are stored in the classical entrywise representation. A visualization of the partitioning of a resulting $\mathscr{H}$-matrix is shown in Fig. 3. Blocks with an entrywise representation are shown in black color.

All other blocks are represented in outer-product form, where the number corresponds to the rank. Obviously, the better the row and column degrees of freedom are separated, the larger blocks can be approximated with a very low rank. There is only one question left: How is this partitioning set up? It is non-trivial for a three-dimensional problem to find separated variables x and y. For this purpose, a cluster tree T_I is built for the index set I of the row degrees of freedom and a corresponding cluster tree T_J for the index set J of the column, respectively. They are

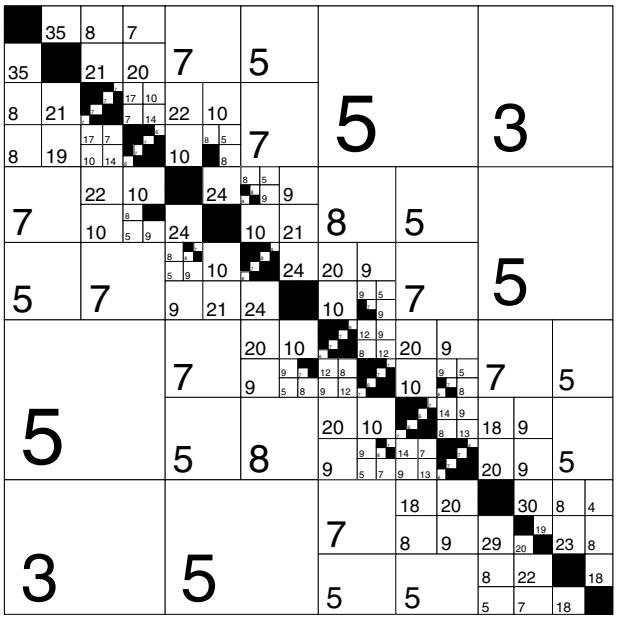

Figure 3: Example of a $\mathscr{H}$-matrix representation. Black partitions are stored in traditional entrywise representation. For all other partitions the outer-product form is applied, where the number corresponds to the rank k.

set up by a bisectioning process, which starts with the root cluster containing all the indices of the row or column degrees of freedom. All indices of elements/nodes on the first side of the hypersurface belong to the first son, whereas all other nodes are assigned to the second one. The division also defines the permutation to find the separated degrees of freedom, since all entries of the first son are given smaller indices compared to the second one. If this algorithm is called recursively, a cluster tree is obtained and a permutation of the sets I and J is found which reflects the distance to other entries within the actual geometry of the problem. The division stops, if a prescribed number of degrees of freedom b in a cluster is reached. Within this work, $b = 20$ is chosen.

To obtain the partitioning shown in Fig. 3, a block cluster tree $T_{I\times J}$ has to be set up from the two cluster trees T_I and T_J. The set-up of the block cluster tree consists of the following steps:

For a given pair of clusters $C^\ell_{I\gamma}$ and $C^\ell_{J\delta}$ on the level ℓ, the admissibility condition

$$2\min\{r^\ell_{I\gamma}, r^\ell_{J\delta}\} < \eta\,|z^\ell_{I\gamma} - z^\ell_{J\delta}| \tag{26}$$

is defined, where $r^\ell_{I\gamma}$ is the radius and $z^\ell_{I\gamma}$ the center of the cluster $C^\ell_{I\gamma}$ and η denotes the admissibility constant. If Eq. 26 is satisfied, the cluster pair is said to be admissible and their degrees of freedom are well separated. The corresponding block

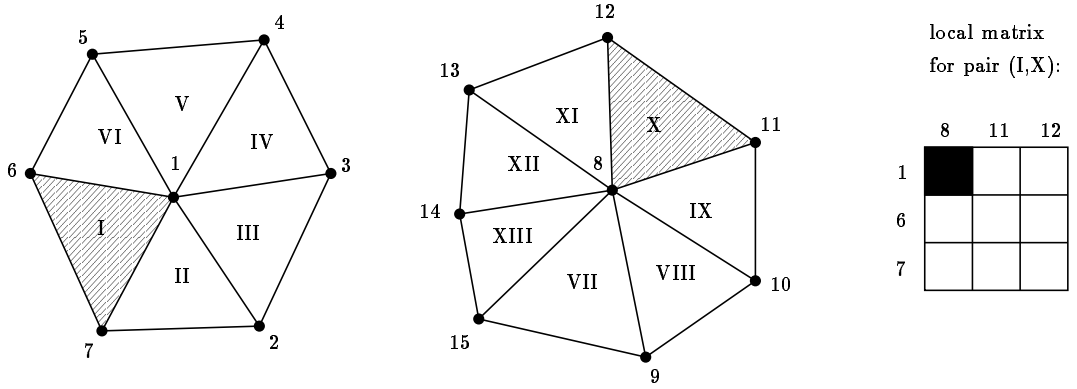

Figure 4: Integration of an element pair (I,X) and corresponding local matrix in case of $\mathbf{K}_{\mathrm{BE}}$ with piecewise linear test and shape functions.

of the matrix can be approximated in outer-product form and the algorithm stops. If condition Eq. 26 is not satisfied and neither $C^{\ell}_{I\gamma}$ nor $C^{\ell}_{J\delta}$ are leaf clusters, four new pairs are created of the 2×2 sons and the algorithm is called recursively. If instead one of the clusters is a leaf cluster, that is the number of degrees of freedom within this cluster is smaller than the threshold b, the corresponding block is small and can be represented exactly in the traditional entrywise manner. In this case, the algorithm stops as well. The overall procedure starts with the root clusters of the row and column degrees of freedom.

4.2.4 Modification of Integration Routines

As an advantage of ACA, only single entries of the matrices $\mathbf{K}_{\mathrm{BE}}$ and $\mathbf{C}_{\mathrm{BE}}$ are required for setting up the low-rank approximations. Basically, the integration routines of a classical approach can be reused. However, for piecewise linear test and shape functions this strategy is inefficient. This is visualized in Fig. 4. To obtain the entry where the row corresponds to node 1 and the column to node 8, one has to take into account all elements which share these nodes. An integration has to be performed for every element pair which yields nine entries. Only one entry is used and the remaining entries are rejected. To improve the performance, the integration routines are modified such that only the required entry is computed. This significantly reduces the integration time. Alternatively, the original routines may be used and the currently unnecessary entries are stored temporarily. Since always total rows or columns are needed for ACA, some of the entries are likely to be required later on. This can also help to accelerate the approximation process. However, within this paper such a scheme is not applied. Furthermore, only a single $\mathscr{H}$-matrix is set up for $\mathbf{K}_{\mathrm{BE}}$ which takes into account all operators. The same holds for $\mathbf{C}_{\mathrm{BE}}$. The implementation of the half-space problem is done by directly using the half-space

fundamental solution in the integration routines. In contrast to the FMM, no mirror technique is necessary. Except the modifications discussed above, the same kind of integration routines are applied for the LF FMM, the HF FMM and the ACA. The singular integrals are computed using the scheme according to Sauter (1992). Here, the inner surface integral is analytically solved and a numerical integration is applied for the outer surface integral surface integral (Fischer, 2004). For the non-singular integrals a simple numerical quadrature is applied. It is worth mentioning, an improvement of the integration routines is clearly in favor of the ACA, since more entries have to be computed compared to the near-field of the FMM.

4.2.5 Hierarchical Lower Upper (HLU) Decomposition

Besides the advantage that no series expansion of the fundamental solution is required, an approximation of the whole BE matrix is stored in an explicit way. This is important for setting up preconditioners, since also the far-field influence is taken into account. This is different in case of the FMM, where typically only the near-field is available for constructing preconditioners. For $\mathscr{H}$-matrices a hierarchical lower upper (HLU) factorization is available (Bebendorf, 2005) which has a quasilinear complexity. It is basically possible to apply an HLU as direct solver. However, typically it is more efficient to use an iterative solver with an HLU of lower precision for preconditiong. An HLU preconditioner is adaptive in the sense, that the accuracy is adjustable. This yields a high flexibility with such a black-box approach. In this paper no HLU is applied. However, for fluid-structure coupled problems an HLU may be advantageous (Brunner, 2009).

In what follows, for all simulations the $\mathscr{H}$-matrix library AHMED[1] is applied.

5 Numerical Results

The proposed fast BE methods are compared for the cylindrical test structure depicted in Fig. 5.

The structure has a diameter of 2 m and a draft of 1.866 m, which is equivalent to an angle of $\alpha = 60°$, see Fig. 1. The total length is 22 m. Five different discretizations using triangular elements are considered. The coarsest one denoted "original" has approximately 11k nodes. In every refinement step, the number of nodes is approximately doubled. Hence, the finest model "refined 4×" has approximately 170k nodes. To be able to compare the BEM results with an analytical solution, a sound field is synthezised using nine monopole sources which are situated along the centerline of the structure. Since this sound field has to be pressure-free on

[1] AHMED: "**A**nother software library on **H**ierarchical **M**atrices for **E**lliptic **D**ifferential equations", written by M. Bebendorf, http://www.bebendorf.ins.uni-bonn.de/AHMED.html

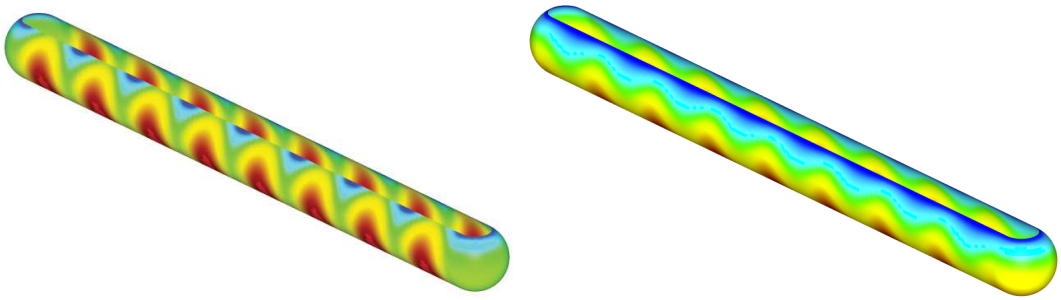

Figure 5: Cylindrical test structure: Analytical flux boundary conditions (left) and computed pressure field (right) at the frequency $f = 100\,\mathrm{Hz}$.

the water-surface, the monopole sources are also mirrored with respect to the water surface and their field subtracted from the non-mirrored one. The resulting flux- and pressure distributions are depicted in Fig. 5 left and right. In what follows, the analytical flux field is used as boundary condition. For the pressure field, the Dirichlet error

$$e_{\mathrm{D}} = \frac{||\mathbf{p}_{\mathrm{BEM}} - \mathbf{p}_{\mathrm{mono}}||}{||\mathbf{p}_{\mathrm{mono}}||} \tag{27}$$

is evaluated, where $\mathbf{p}_{\mathrm{BEM}}$ is the nodal pressure vector computed by the BEM and $||\cdot||$ denotes the L_2-norm. For all following investigations, a GMRES is applied as iterative solver. A simple diagonal scaling is used as preconditioner. All timing tests are run on an Intel Xeon 5160 3GHz CPU with 16GB RAM. C++ is used as programming language using the standard GNU compiler and the Math Kernel Library (MKL) of Intel.

5.1 Low-Frequency Stability

First the accuracy of the different methods is compared for the "original" model with 11k nodes and the frequency range $f = [5\,\mathrm{Hz}, 200\,\mathrm{Hz}]$. This discretization results in approximately 950 Elements per wavelength at $f = 15\,\mathrm{Hz}$. The near-field parameter of the FMM implementations is chosen as $c_{\mathrm{n}} = 3$. The corresponding Dirichlet errors e_{D} of the low-frequency (LF) FMM and high-frequency (HF) FMM are depicted in Fig. 6 (left). In case of HF FMM, the errors increase rapidly for small frequencies. Even for a very large expansion length parameter $c_{\mathrm{e}} = 6$, there is a significant discrepancy compared to the results obtained by the classical BEM, where the fully populated matrices are set up. This clearly shows the infeasibility of the diagonal expansion for this frequency regime. For the refined models, this effect is even stronger within the investigated frequency range, since the ratio of elements per wavelength is even higher in these cases. In contrast to this, the LF FMM yields

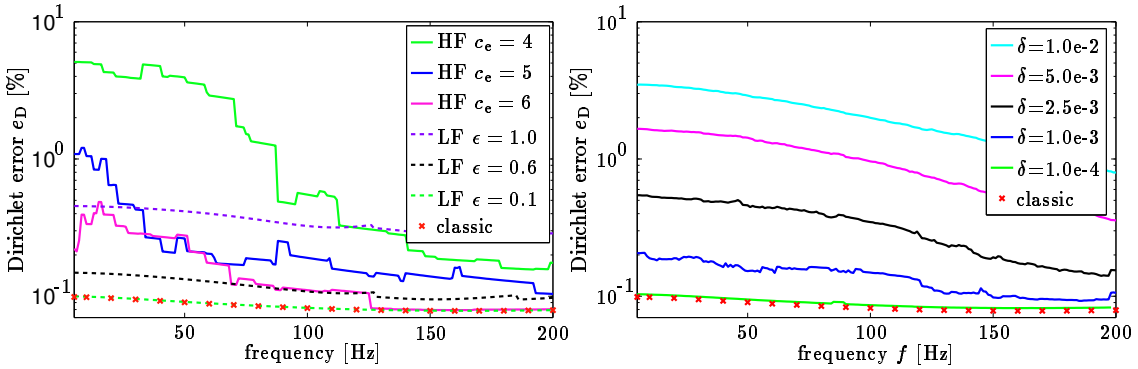

Figure 6: Cylidrical test structure: Dirichlet errors for the FMM (left) and when using $\mathscr{H}$-matrices (right).

Table 1: Scenario 1: Dirichlet errors in % for the LF FMM in case of a constant frequency $f = 15$ Hz.

		original	refined 1×	refined 2×	refined 3×	refined 4×
c_n	ε \ λ/h	944.9	1330.2	1889.8	2661.6	3779.5
3.0	1.0	0.383	0.925	2.192	2.745	5.312
3.0	**0.6**	**0.166**	**0.139**	**0.055**	**0.265**	**0.248**
3.0	0.1	0.104	0.057	0.028	0.030	0.001
3.0	0.03	0.098	0.052	0.025	0.014	0.031

accurate results and the errors are comparable with those of the classical BEM for $\varepsilon = 0.1$. For the $\mathscr{H}$-matrices, a parameter $\eta = 0.9$ is chosen. The corresponding Dirichlet errors e_D are depicted in Fig. 6 (right) for different accuracies δ. The level of the classical BEM is reached for $\delta = 1e-4$. Hence, the $\mathscr{H}$-matrices do not suffer from stability problems in the low-frequency range.

5.2 Simulation Time and Memory Consumption

For the comparison of the simulation time and memory consumption, two scenarios are chosen: In the first one, the frequency is kept fixed at 15 Hz for the five models. In this case, the LF FMM is applied. In the second scenario, the same models are used, but this time the frequency is chosen such that the ratio of wavelength to element size is $\lambda/h = 10$. Here, the HF FMM is applied.

5.2.1 Scenario 1: Constant Frequency 15 Hz

The Dirichlet errors for different ε are summarized in Tab. 1. For $\varepsilon = 0.03$ the quadratic convergence behavior of the BEM is observable except for the finest

Table 2: Scenario 1: Dirichlet errors in % for the $\mathscr{H}$-matrices in case of a constant frequency $f = 15\,\mathrm{Hz}$.

δ \ λ/h	original 944.9	refined 1× 1330.2	refined 2× 1889.8	refined 3× 2661.6	refined 4× 3779.5
$5.0 \cdot 10^{-3}$	1.402	3.126	5.341	8.367	13.00
$\mathbf{1.0 \cdot 10^{-3}}$	**0.121**	**0.218**	**0.503**	**0.091**	**1.598**
$5.0 \cdot 10^{-4}$	0.096	0.075	0.169	0.281	0.541
$1.0 \cdot 10^{-5}$	0.099	0.055	0.032	0.020	0.013

Table 3: Scenario 1: CPU time [in s] and memory consumption [in MB] for the LF FMM and $\mathscr{H}$-matrices in case of a constant frequency $f = 15\,\mathrm{Hz}$.

λ/h	original 944.9	refined 1× 1330.2	refined 2× 1889.8	refined 3× 2661.6	refined 4× 3779.5
LF FMM:					
set-up time	164	359	657	1616	2624
MV-prod. time	2.8	5.4	12.1	20.4	46.8
memory	47	92	166	440	772
$\mathscr{H}$-matrices:					
set-up time	1117	2759	6442	14552	32025
MV-prod. time	0.029	0.069	0.156	0.350	0.790
memory	86	199	458	1029	2304
crossover its.	349	453	484	644	639

model. The outlier may be caused by the regularization based on Stokes theorem which is only applied for the near-field (Fischer, 2004). The corresponding Dirichlet errors for the $\mathscr{H}$-matrices are shown in Tab. 2. Except for the finest model, the errors for an accuracy $\delta = 1.0 \cdot 10^{-5}$ are comparable with the ones of the LF FMM. This shows the correct implementation of the methods.

For engineering applications, a Dirichlet error of 1.5% is acceptable. Therefore, in order to compare the computation time and memory consumption, the parameter sets which are highlighted in bold style in the tables Tab. 1 and Tab. 2 are chosen. The simulation times and memory consumptions are summarized in Tab. 3. Concerning the simulation time, the set-up time and the time for a matrix-vector product with $\mathbf{K}_{\mathrm{BE}}$ (MV-prod. time) as required by the iterative solver are distinguished. The set-up time for the LF FMM consists of the precomputation of all translation and rotation coefficients as mentioned above. Also the time for setting up the near-field and computing the right hand side vector $\mathbf{b} = \mathbf{C}_{\mathrm{BE}}\,\bar{\mathbf{q}}$ is included.

In case of the $\mathscr{H}$-matrices, the set-up time is composed of the time for assembling and recompressing the matrices $\mathbf{K}_{\mathrm{BE}}$ and $\mathbf{C}_{\mathrm{BE}}$. The set-up time for the cluster trees is not included in both cases. Obviously, the set-up time of the LF FMM is faster by a factor of 7 to 12 compared with the $\mathscr{H}$-matrices. However, one should keep in mind that at this stage, the whole compression is already done for the $\mathscr{H}$-matrices. This becomes clear by comparing the time for a single matrix-vector product with $\mathbf{K}_{\mathrm{BE}}$. Here, $\mathscr{H}$-matrices are faster by a factor of 100 to 60 compared with the LF FMM. Thus, for the choice of the most efficient method, the number of required iterations of the GMRES which are equal to the number of required matrix-vector products is crucial. This crossover point is shown in the last line of Tab. 3. It is reached at 349 iterations for the smallest model and increases to approximately 640 iterations for the larger models. For a simulation of the sound radiation of vibrating structures where the flux $\bar{\mathbf{q}}$ is known, the number of necessary iterations are typically below this crossover point. This is especially the case for half-space problems which usually have a good condition number. The situation changes for fluid-structure coupled problems, where the convergence is often much worse (Brunner, 2009). Here, the use of $\mathscr{H}$-matrices may be advantageous. In Tab. 3 also the memory consumption is listed. For the LF FMM, all rotation and translation coefficients mentioned above, the local coefficients $\mathscr{L}_n^m$ and the near-field are included. A near-field matrix is only set up for $\mathbf{K}_{\mathrm{BE}}$. In case of the $\mathscr{H}$-matrices, the listed memory consumption corresponds to the approximation of $\mathbf{K}_{\mathrm{BE}}$. The required memory of $\mathbf{C}_{\mathrm{BE}}$ is greater by a factor of approximately 1.3, but this matrix is only needed once for computing the right hand side vector b. During the solution procedure where the Krylov vectors of the GMRES additionally have to be stored, only $\mathbf{K}_{\mathrm{BE}}$ has to be kept in the memory and is therefore used for the comparison. As seen in Tab. 3, the memory consumption of the $\mathscr{H}$-matrices is larger by a factor of 1.8 to 3 compared to the LF FMM.

5.2.2 Scenario 2: Ratio of 10 Elements per Wavelength

In the second case, the same models are used. However, this time the frequency is chosen such that the ratio of wavelength to element size is $\lambda/h = 10$. Here, the HF FMM is applied. As shown in Tab. 4 and Tab. 5, the results reach the discretization error of the BEM even for a small expansion length parameter $c_{\mathrm{e}} = 2$. A further increase of the FMM accuracy does not show any effect. Concerning the $\mathscr{H}$-matrices, a parameter $\delta = 1e-3$ yields accurate results as well. Again the errors of the HF FMM are comparable with those of the $\mathscr{H}$-matrices. As above, the parameters printed in bold style are applied for the following comparison.

The definition of the set-up time and memory consumption for the $\mathscr{H}$-matrices is the same as before. For the HF FMM, the set-up time consists of the computation

Table 4: Dirichlet errors in % for the HF FMM in case of a constant ratio of $\lambda/h = 10.0$ elements per wavelength.

c_n	$c_e \setminus f$	original 1411.7 Hz	refined 1× 1995.4 Hz	refined 2× 2823.0 Hz	refined 3× 3992.5 Hz	refined 4× 5655.0 Hz
3.0	1.0	1.670	1.732	1.603	1.600	1.707
3.0	**2.0**	**1.570**	**1.641**	**1.535**	**1.547**	**1.637**
3.0	3.0	1.568	1.640	1.534	1.547	1.637
3.0	4.0	1.568	1.640	1.534	1.547	1.637

Table 5: Dirichlet errors in % for the $\mathscr{H}$-matrices in case of a constant ratio of $\lambda/h = 10.0$ elements per wavelength.

$\delta \setminus f$	original 1411.7 Hz	refined 1× 1995.4 Hz	refined 2× 2823.0 Hz	refined 3× 3992.5 Hz	refined 4× 5655.0 Hz
$1.0 \cdot 10^{-2}$	1.771	1.903	1.999	2.408	3.302
$5.0 \cdot 10^{-3}$	1.658	1.710	1.666	1.880	2.198
$\mathbf{1.0 \cdot 10^{-3}}$	**1.573**	**1.647**	**1.542**	**1.567**	**1.667**
$5.0 \cdot 10^{-4}$	1.572	1.643	1.538	1.555	1.651

time for the translation operators M_L, the near-field and the right hand side vector $\mathbf{b}$. The memory consumption of the translation operators, near-field signatures and the near-field itself are considered for the HF FMM. The results are summarized in Tab. 6. The situation is similar as for the first scenario: The set-up time of the HF FMM is by a factor of 9 to 23 shorter compared to the $\mathscr{H}$-matrices. However, the matrix-vector product of the latter ones is faster by a factor of 164 to 66. Hence, the crossover points are situated between 288 and 647 iteration steps.

Concerning the memory consumption, $\mathscr{H}$-matrices require 1.75 to 3.87 times more memory compared to the HF FMM. The ratio increases for larger models. Thus, the crossover point and the memory consumption become more favorable for the HF FMM for the larger models. However, it is worth mentioning, the FMM implementation of the investigated half-space problem is very efficient when using the mirror technique (Brunner, 2009). This is because of the good far-field compression in the mirrored part. In case of the $\mathscr{H}$-matrices, the kernel within the integration routines becomes more complex and one does not gain an additional compression of the mirrored part. Hence, a comparison of the FMM and $\mathscr{H}$-matrices for a non-half-space problem probably yield an even better result for the latter ones.

Table 6: Scenario 2: CPU time [in s] and memory consumption [in MB] for the HF FMM and $\mathscr{H}$-matrices in case of a constant ratio of 10 elements per wavelength.

	original	refined 1×	refined 2×	refined 3×	refined 4×
f [Hz]	1411.7	1995.4	2823.0	3992.5	5655.0
HF FMM:					
set-up time	225	491	912	1945	3918
MV-prod. time	6.4	14.7	30.6	67.2	136.8
memory	78	164	345	773	1627
$\mathscr{H}$**-matrices:**					
set-up time	2051	5258	13352	34696	91136
MV-prod. time	0.039	0.103	0.275	0.737	2.050
memory	137	312	839	2258	6265
crossover its.	288	327	411	493	647

6 Conclusion

This paper compares the fast multipole method with hierarchical matrices for a half-space problem as it occurs when investigating the vibro-acoustic behavior of surface ships. Depending on the frequency regime, two different multipole expansions are applied. However a single implementation of hierarchical matrices is applied for all frequencies. Only for low frequencies a quasilinear numerical complexity can be expected. In a first scenario, the frequency is kept fix at 15 Hz and the model is successively refined. This yields much more than ten elements per wavelength for the largest model. This scenario is of interest when investigating fluid-structure coupled problems, where a conforming coupling scheme is applied between the structure and the fluid. In the second case, a ratio of ten elements per wavelength is applied. In both cases, the set-up time of the hierarchical matrices is longer compared to the fast multipole method. However, the computation time for a single matrix-vector product is significantly faster for the hierarchical matrices. Hence, the choice of the most efficient method depends on the case of application. The crossover point is between 288 and 647 matrix-vector products in case of the considered examples and used implementations. Concerning the memory consumption, the fast multipole method yields a better compression. The required memory is up to a factor 3.87 larger for the hierarchical matrices for the considered examples. It should also be kept in mind, the construction of efficient preconditioners may be significantly easier for hierarchical matrices, since also the far-field is included in the explicit representation. Additionally, the implementation and parallelization is much easier for the hierarchical matrices compared to the fast multiple

methods. The investigations in this paper are made for a half-space formulation, which is implemented in different ways for the fast multipole method and the hierarchical matrices. The results of a comparison for a non-half-space problem might look different.

Acknowledgement: This research was financially supported by the German Research Foundation (DFG) under the transfer project SFB404/T3. The authors acknowledge valuable contributions by Dr. Christian Cabos and Marc Wilken of Germanischer Lloyd AG, Hamburg and by Prof. Olaf Steinbach and Dr. Günther Of from the Technical University of Graz within this transfer project.

References

Aimi, A.; Diligenti, M.; Lunardini, F.; Salvadori, A. (2003): A new application of the panel clustering method for 3D SGBEM. *CMES: Computer Modeling & Engineering Sciences*, vol. 4, pp. 31–49.

Amini, S.; Profit, A. T. J. (1999): Analysis of a diagonal form of the fast multipole algorithm for scattering theory. *BIT*, vol. 39, no. 4, pp. 585–602.

Aoki, S.; Amaya, K.; Urago, M.; Nakayama, A. (2004): Fast multipole boundary element analysis of corrosion problems. *CMES: Computer Modeling & Engineering Sciences*, vol. 6, pp. 123–131.

Bebendorf, M. (2000): Approximation of boundary element matrices. *Numer. Math.*, vol. 86, pp. 565–589.

Bebendorf, M. (2005): Hierarchical LU decomposition-based preconditioners for BEM. *Computing*, vol. 74, pp. 225–247.

Bebendorf, M. (2008): *Hierarchical Matrices, A Means to Efficienctly Solve Elliptic Boundary Value Problems*. Springer Berlin, Heidelberg.

Bebendorf, M.; Kriemann, R. (2005): Fast parallel solution of boundary integral equations and related problems. *Comp. Vis. Sci.*, vol. 8, pp. 121–135.

Bebendorf, M.; Rjasanow, S. (2003): Adaptive low-rank approximation of collocation matrices. *Computing*, vol. 70, pp. 1–24.

Bebendorf, M.; Rjasanow, S.; Tyrtyshnikov, E. E. (2000): Approximation using diagonal-plus-skeleton matrices. In *Mathematical aspects of boundary element methods (Palaiseau, 1998)*, volume 414 of *Chapman & Hall/CRC Res. Notes Math.*, pp. 45–52. Chapman & Hall/CRC, Boca Raton, FL.

Bettess, P. (1992): *Infinite Elements*. Penshaw Press, U.K.

Bonnet, M. (1995): *Boundary Integral Equation Methods for Solids and Fluids*. John Wiley & Sons LTD, Chichester.

Brancati, A.; Aliabadi, M. H.; Benedetti, I. (2009): Hierarchical adaptive cross approximation GMRES technique for solution of acoustic problems using the boundary element method. *CMES: Computer Modeling & Engineering Sciences*, vol. 43, pp. 149–172.

Brunner, D. (2009): *Fast Boundary Element Methods for Large-Scale Simulations of the Vibro-Acoustic Behavior of Ship-Like Structures*. PhD thesis, University of Stuttgart, 2009.

Brunner, D.; Junge, M.; Wilken, M.; Cabos, C.; Gaul, L. (2009): Vibro-acoustic simulations of ships by coupled fast BE-FE approaches. *Proceedings of IMAC XXVII: Conference & Exposition on Structural Dynamics, Orlando, CDROM*.

Brunner, D.; Of, G.; Junge, M.; Steinbach, O.; Gaul, L. (2010): A fast BE-FE coupling for partly immersed bodies. *Int. J. Numer. Meth. Eng.*, vol. 81, pp. 28–47.

Buchau, A.; Rucker, W.; Rain, O.; Rischmüller, V.; Kurz, S.; Rjasanow, S. (2003): Comparison between different approaches for fast and efficient 3-D BEM computations. *IEEE T. Magn.*, vol. 39, pp. 1107–1110.

Burton, A. J.; Miller, G. F. (1971): The application of integral equation methods for the numerical solution of boundary value problems. *Proc. Roy. Soc. London A*, vol. 232, pp. 201–210.

Cheng, H.; Crutchfield, W. Y.; Gimbutas, Z.; Greengard, L. F.; Ethridge, J. F.; Huang, J.; Rokhlin, V.; Yarvin, N.; Zhao, J. (2006): A wideband fast multipole method for the Helmholtz equation in three dimensions. *J. Comput. Phys.*, vol. 216, pp. 300–325.

Chew, W. C.; Song, J. M.; Cui, T. J.; Velamparambil, S.; Hastriter, M. L.; Hu, B. (2004): Review of large scale computing in electromagnetics with fast integral equation solvers. *CMES: Computer Modeling & Engineering Sciences*, vol. 5, pp. 361–372.

Coifman, R.; Rokhlin, V.; Wandzura, S. (1993): The fast multipole method for the wave equation: A pedestrian description. *IEEE Antenn. Progag. M.*, vol. 35, pp. 7–12.

Darve, E. (2000): The fast multipole method: numerical implementation. *J. Comput. Phys.*, vol. 160, pp. 195–240.

Darve, E.; Harvé, P. (2004): Efficient fast multipole method for low-frequency scattering. *J. Comput. Phys.*, vol. 197, pp. 341–363.

Epton, M. A.; Dembart, B. (1995): Multipole translation theory for the tree-dimensional Laplace and Helmholtz equation. *SIAM J. Sci. Comput.*, vol. 16, no. 4, pp. 865–897.

Fischer, M. (2004): *The Fast Multipole Boundary Element Method and its Application to Structure-Acoustic Field Interaction*. PhD thesis, University of Stuttgart, 2004.

Forster, H.; Schrefl, T.; Dittrich, R.; Scholz, W.; Fidler, J. (2003): Fast boundary methods for magnetostatic interactions in micromagnetics. *IEEE T. Magn.*, vol. 39, pp. 2513–2515.

Gaul, L.; Kögl, M.; Wagner, M. (2003): *Boundary Element Methods for Engineers and Scientists. An Introductory Course with Advanced Topics.* Springer Berlin.

Greengard, L.; Huang, J.; Rokhlin, V.; Wandzura, S. (1998): Accelerating fast multipole methods for the Helmholtz equation at low frequencies. *IEEE Comput. Sci. Eng.*, vol. 5, pp. 32–38.

Gumerov, N. A.; Duraiswami, R. (2003): Recursions for the computation of multipole translations and rotation coefficients for the 3-D Helmholtz equation. *SIAM J. Sci. Stat. Comp.*, vol. 25, pp. 1344–1381.

Gumerov, N. A.; Duraiswami, R. (2004): *Fast Multipole Methods for the Helmholtz Equation in Three Dimensions*. Elsevier, Oxford.

Gumerov, N. A.; Duraiswami, R. (2007): Fast multipole accelerated boundary element method for the 3D Helmholtz equation. Technical report, Institute for Advanced Computer Studies, University Maryland, published online: http://hdl.handle.net/1903/7504, 2007.

Gumerov, N. A.; Duraiswami, R. (2009): A broadband fast multipole accelerated boundary element method for the 3D Helmholtz equation. *J. Acoust. Soc. Am.*, vol. 125, pp. 191–205.

Gumerov, N. A.; Duraiswami, R.; Borovikov, E. A. (2003): Data structures, optimal choice of parameters, and complexity results for generalized multilevel fast multipole methods in d dimensions. Technical report, Institute for Advanced Computer Studies, University Maryland, published online: http://hdl.handle.net/1903/1270, 2003.

Gyure, M. F.; Stalzer, M. A. (1998): A prescription for the multilevel Helmholtz FMM. *IEEE Comput. Sci. Eng.*, vol. 5, pp. 39–47.

Hackbusch, W. (1999): A sparse matrix arithmetic based on H-matrices I. introduction to H-matrices. *Computing*, vol. 62, no. 2, pp. 89–108.

He, X.; ad S.-P. Lim, K.-M. L. (2008): Fast BEM solvers for 3D Possion-Type equations. *CMES: Computer Modeling & Engineering Sciences*, vol. 35, pp. 21–48.

Hu, B. H.; Chew, W. C. (2001): Fast inhomogeneous plane wave algorithm for scattering from objects above the multilayered medium. *IEEE T. Geosci. Remote*, vol. 39, pp. 1028–1038.

Kinsler, L. E.; Frey, A. R.; Coppens, A. B.; Sanders, J. V. (2000): *Fundamentals of Acoustics*. John Wiley & Sons, Inc., 4th edition.

Nédélec, J.-C. (2001): *Acoustic and Electromagnetic Equations: Integral Representations for Harmonic Problems*. Applied Mathematical Sciences 144. Springer New York.

Nishimura, N. (2002): Fast multipole accelerated boundary integral equation methods. *Appl. Mech. Rev.*, vol. 55, pp. 299–324.

Rahola, J. (1996): Diagonal forms of the translation operators in the fast multipole algorithm for scattering problems. *BIT*, vol. 36, no. 2, pp. 333–358.

Rjasanow, S.; Steinbach, O. (2007): *The Fast Solution of Boundary Integral Equations*. Springer New York.

Rokhlin, V. (1985): Rapid solution of integral equations of classical potential theory. *J. Comput. Phys.*, vol. 60, pp. 187–207.

Rokhlin, V. (1993): Diagonal forms of translation operators for the Helmholtz equation. *Appl. Comput. Harmon. A.*, vol. 1, pp. 82–93.

Saad, Y. (2003): *Iterative Methods for Sparse Linear Systems*. PWS.

Sauter, S. (1992): *Über die effiziente Verwendung des Galerkinverfahrens zur Lösung Fredholmscher Integralgleichungen*. PhD thesis, Christian-Albrechts-Universität zu Kiel, 1992.

Seybert, A.; Wu, T. (1989): Modified Helmholtz integral equation for bodies sitting on an infinite plane. *J. Acoust. Soc. Am.*, vol. 85, pp. 19–23.

Seybert, A. F.; Soenarko, B. (1988): Radiation and scattering of acoustic waves from bodies of arbitrary shape in a three-dimensional half space. *Transactions of the ASME*, vol. 110, pp. 112–117.

Sladek, V.; Tanaka, M.; Sladek, J. (2001): Revised Helmholtz integral equation for bodies sitting on an infinite plane. *Trans. of JSCES*, vol. 3, pp. 35–40.

Steinbach, O. (2008): *Numerical Approximation Methods for Elliptic Boundary Value Problems. Finite and Boundary Elements*. Springer New York.

Tyrtyshnikov, E. (1998): Mosaic-skeleton approximations. *Calcolo*, vol. 33, pp. 47–57.

Wang, H.; Yao, Z. (2005): A new fast multipole boundary element method for large scale analysis of mechanical properties in 3d particle-reinforced composites. *CMES: Computer Modeling & Engineering Sciences*, vol. 7, pp. 85–95.

Wu, T. W.(Ed): *Boundary Element Acoustics. Fundamentals and Computer Codes*. WIT Press Southampton.

Zienkiewicz, O. C.; Taylor, R. L. (2000): *The Finite Element Method, Volume 1: The Basis*. Butterworth Heinemann Oxford.