An oversampled collocation approach of the Wave Based Method for Helmholtz problems

Daan Huybrechs and Anda-Elena Olteanu

KU Leuven, Department of Computer Science, Celestijnenlaan 200A, BE-3001 Leuven, Belgium. Corresponding author: daan.huybrechs@cs.kuleuven.be

Abstract

The Wave Based Method (WBM) is a Trefftz method for the simulation of wave problems in vibroacoustics. Like other Trefftz methods, it employs a non-standard discretisation basis consisting of solutions of the partial differential equation (PDE) at hand. We analyse the convergence and numerical stability of the Wave Based Method for Helmholtz problems using tools from approximation theory. We show that the set of discretisation functions more closely resembles a frame, a redundant set of functions, than a basis. The redundancy of a frame typically leads to ill-conditioning, which indeed is common in Trefftz methods. Recent theoretical results on frames for function approximation suggest that the associated ill-conditioned system matrix can be successfully regularised, with error bounds available, when using a discrete least squares approach. While the original Wave Based Method is based on a weighted residual formulation, in this paper we pursue an oversampled collocation approach instead. We show that, for smooth scattering obstacles in two dimensions and without partitioning of the domain, the results closely follow the theory of frames. We identify cases where the method achieves very high accuracy whilst providing a solution with small norm coefficients, in spite of ill-conditioning. Moreover, the accurate results are reliably maintained even in parameter regimes associated with extremely high ill-conditioning.

Keywords: Wave Based Method, frames, collocation, Helmholtz
2010 MSC: 65N12, 65D15

1. Introduction

Numerical simulation methods for wave scattering and propagation problems lead to a wealth of mathematical and computational challenges. One of the primary concerns is the number of degrees of freedom that are required in order
5 to represent a wave field, in particular for problems involving high or moderately high frequencies. The class of Trefftz methods aims to reduce this number, by discretising the governing partial differential equation (PDE) using solutions of that same PDE. We refer the reader to [1] and references therein for a review

of Trefftz methods, such as the Method of Fundamental Solutions [2], the Wave
 10 Based Method [3, 4] and the Ultra-Weak Variational Formulation scheme [5].
 Since the ansatz of the solution satisfies the PDE by construction, it remains
 to enforce boundary and interface conditions. This leads to a linear system
 of equations that is typically dense, albeit with much smaller dimensions than
 a corresponding discretisation using finite elements. Unfortunately, in quite a
 15 few cases, the system exhibits a large or very large condition number [1]. This
 seems a cause of concern, and perhaps even prevents wider adoption of these
 methods. Yet, extensive experiments with and literature on Trefftz methods
 seem to indicate that, in spite of ill-conditioning of the linear system, high
 accuracy solutions are often found and high-order convergence is observed.

20 The analysis of Trefftz methods has seen great progress in recent years.
 A major focus of the mathematical literature on this topic lies with the best
 approximation to the solution in the approximation space, the space that is
 spanned by the chosen basis functions [6, 7]. Given the unconventional nature of
 the space, this is a non-trivial problem. The computation of the approximation
 25 itself, in particular the way it is obtained from a potentially ill-conditioned
 system, has received much less attention. Barnett and Betcke have studied the
 Method of Fundamental Solutions, and proposed an implementation strategy
 that is numerically stable by a judicious choice of the charge points in the
 method [8, 9]. An often quoted test is the Picard condition [10, §1.2.3]. This
 30 is a test that can provide confidence in a solution a-posteriori, but does not
 provide an a-priori guarantee of success.

Some recent methods in approximation theory have similar characteristics as
 Trefftz methods, similar in the sense that highly accurate solutions are found by
 solving ill-conditioned systems. These methods are based on using a so-called
 35 *frame*, rather than a basis. We review frames in §3 of this paper. Compared
 to a basis a frame is redundant, hence there are multiple representations of
 any given function in the frame. This leads to ill-conditioned matrices for the
 approximation problem. Yet, generic error bounds can be shown [11, 12]. More-
 over, convergence to high accuracy, up to machine precision even, can provably
 40 be guaranteed in spite of the ill-conditioning, under certain additional condi-
 tions. The main condition is *oversampling*: rather than solving square systems,
 accuracy and robustness is significantly improved by considering solutions in a
 least squares sense. Best results are achieved in particular with a discrete least
 squares approximation, based on a larger number M of function evaluations
 45 than degrees of freedom N (i.e., $M > N$, leading to a rectangular system). This
 approach is shown to be successful precisely when the approximation space is
 obtained from a frame. Hence the importance of the concept of frames in this
 context: if the redundant set is not a frame, functions can be found for which the
 approximation scheme may not converge using finite precision computations.

50 The analogue of a discrete least-squares approximation for Trefftz methods is
 an oversampled collocation approach. The goal of this paper is to describe such
 a collocation approach for the Wave Based Method, for Helmholtz problems in
 two dimensions. We fully address the problem of ill-conditioning of the linear
 systems and, through a sequence of experiments, illustrate how it relates to the

55 error bounds from the literature on frames. The experiments are inspired by
the prior analysis of the Method of Fundamental Solutions in [8] and, following
Barnett and Betcke, we interpret the requirement of having a convex domain in
WBM in terms of external singularities of the solution. We show much improved
accuracy of the oversampled collocation approach, compared to the weighted
60 residual formulation that is commonly employed. In addition, the approach is
more efficient since no calculation of integrals is required in the setup of the
system matrix.

On the other hand, our results are restricted to convex and certain non-
convex domains with a smooth boundary in two dimensions. We also restricted
65 ourselves to a single domain version of WBM, i.e., we did not consider a par-
titioning of the domain. Domain partitioning is comparable to h -refinement
and is common in WBM and other Trefftz methods. Whether the beneficial
properties of an oversampled collocation approach remain in the various other
settings in which WBM has been applied, in particular in the presence of corner
70 singularities, is not certain and is a topic of further research.

We also like to stress that, though we include no mathematical analysis in
this paper, the set of functions used in WBM does *not* satisfy the mathematical
conditions of a frame. The means that, mathematically, an accurate numerical
solution is not guaranteed for all possible boundary conditions. Correspond-
75 ingly, we show examples where full accuracy is not obtained.

Ongoing research efforts seem to indicate that this is the case for several
other Trefftz methods as well, even when it is known that the best approximation
from the approximation space converges. Yet, the connection to the literature
on frames does support the following observation. It is not sufficient that a good
80 approximation to the solution exists in the approximation space. However, it *is*
sufficient that a good approximation exists that has expansion coefficients with
(sufficiently) small discrete norm. Such a solution can be found numerically
using a linear algebra solver that produces solutions with small residual as well
as small norm, such as solvers based on SVD or QR decompositions (see §3.2). In
85 that case, if a good approximation with moderately small coefficients exists then,
with sufficient oversampling and using a sufficiently large number of degrees of
freedom, a numerical solution will be found with equal or with better accuracy.
Perhaps surprisingly, this is true regardless of the ill-conditioning of the linear
system. We believe this observation may underlie the success of Trefftz methods
90 in a variety of practical applications. In this paper we support the statements
above with numerical experiments and interpretations for WBM.

The structure of the paper is as follows. In §2 we recall a formulation of the
Wave Based Method. We review some relevant concepts of the theory of frames
in §3. We describe in particular the generic error bounds on which our interpre-
95 tations are based. We formulate an oversampled collocation implementation of
the Wave Based Method in §4. We present a sequence of numerical experiments
and their results in §5. The results are interpreted using the above-mentioned
error bounds in §6. Finally, we make some concluding remarks in §7.

Figure 1: Circumscribing box B enclosing the domain Ω .

2. The Wave Based Method

We consider the homogeneous Helmholtz equation in the interior of a smooth and bounded domain $\Omega \subset \mathbb{R}^2$,

$$\Delta u(x, y) + k^2 u(x, y) = 0, \quad (x, y) \in \Omega, \quad (1)$$

with a given real constant wavenumber k , $[\text{m}^{-1}]$, and a corresponding boundary condition, e.g. of the commonly used Neumann type

$$\frac{\partial u}{\partial n} = w(s), \quad s \in \partial\Omega. \quad (2)$$

100 The meaning of the unknown function u depends on the wave phenomenon modeled at hand, and typically in the acoustic field, where the WBM has been developed, it denotes the steady state acoustic pressure.

Since the WBM belongs to the family of Trefftz methods, the solution of equation (1) on the domain of interest Ω is sought as an expansion of a complete
105 set of wave functions, which satisfy exactly the governing equation [3, 4].

In the case of WBM for Helmholtz problems in two dimensions, the functions are denoted by $\{\Phi_m^{(\pm)}, \Phi_n^{(\pm)}\}_{m,n=0}^\infty$. They are defined naturally on a rectangular domain $B \subset \mathbb{R}^2$ that circumscribes the domain of interest Ω . We assume without loss of generality that the box can be chosen to be $B = [0, L_x] \times [0, L_y]$, see Fig. 1. Thus, the solution $u(x, y)$ is approximated by $\tilde{u}(x, y)$ given by

$$u \approx \tilde{u} = \sum_{m=0}^{N_m} \left(\alpha_m^{(1)} \Phi_m^{(+)} + \alpha_m^{(2)} \Phi_m^{(-)} \right) + \sum_{n=0}^{N_n} \left(\alpha_n^{(3)} \Phi_n^{(+)} + \alpha_n^{(4)} \Phi_n^{(-)} \right), \quad (3)$$

where N_m and N_n represent the number of functions used in the x and y directions, respectively. The wave functions themselves are given explicitly as

$$\begin{aligned} \Phi_m^{(\pm)} &= \cos\left(k_x^{(m)} x\right) e^{-ik_y^{(m)} y}, \\ \Phi_n^{(\pm)} &= e^{-ik_x^{(n)} x} \cos\left(k_y^{(n)} y\right), \end{aligned} \quad (4)$$

with the wavenumber components

$$\begin{aligned} \left(k_x^{(m)}, k_y^{(m)}\right) &= \left(\frac{m\pi}{L_x}, \pm\sqrt{k^2 - \left(\frac{m\pi}{L_x}\right)^2}\right), \quad m = 0, \dots, N_m, \\ \left(k_x^{(n)}, k_y^{(n)}\right) &= \left(\pm\sqrt{k^2 - \left(\frac{n\pi}{L_y}\right)^2}, \frac{n\pi}{L_y}\right), \quad n = 0, \dots, N_n, \end{aligned} \quad (5)$$

such that $\left(k_x^{(m)}\right)^2 + \left(k_y^{(m)}\right)^2 = \left(k_x^{(n)}\right)^2 + \left(k_y^{(n)}\right)^2 = k^2$ holds and (1) is satisfied by the expansion functions. These are standing waves in one direction, and either propagating or evanescent waves in the other direction, depending on whether the corresponding wavenumber component is real or purely imaginary. Note that one component of the wavenumber always corresponds precisely to the frequencies $n\pi/L$ of a cosine series along an interval of length L . Evanescent waves are added once $n\pi/L$ exceeds the wavenumber k .

It should be noted that a complex value of the wavenumber components $k_y^{(m)}$ or $k_x^{(v)}$ of the wave functions $\Phi_m^{(+)}$ and $\Phi_n^{(+)}$ may lead to amplitudes significantly larger or significantly smaller than 1. For numerical considerations, these functions are typically scaled [3, §2.6.1]: the functions in (4) are multiplied by a constant, such that they are equal to 1 along one edge and exponentially decaying inwards.¹

Regarding truncation, a rule of thumb is to choose a user defined truncation parameter T such that the wavenumber components are linearly scaled with the physical wavenumber k by the factor T , leading to

$$N_m = \left\lceil \frac{kTL_x}{\pi} \right\rceil \quad \text{and} \quad N_n = \left\lceil \frac{kTL_y}{\pi} \right\rceil. \quad (6)$$

Using an analogy to Fourier series, one can think of increasing T as corresponding to increasing the resolution or the maximal frequency.

The problem then translates to finding the unknown coefficients of the expansion, enumerated as a single vector $\alpha = [\alpha_m^{(1)} \alpha_m^{(2)} \alpha_n^{(3)} \alpha_n^{(4)}]$, $m = 0, \dots, N_m$, $n = 0, \dots, N_n$. We enumerate the basis functions accordingly. A solution is found by imposing the boundary condition in a weighted residual formulation:

$$\int_{\partial\Omega} \Phi \left(\frac{\partial u}{\partial n} - w \right) ds = 0. \quad (7)$$

This leads to a linear system of equations $A\alpha = b$ of size $N = 2(N_m + 1) +$

¹This is not the only possible choice of scaling. An alternative may be to scale the evanescent waves such that they achieve the maximal absolute value 1 on the domain Ω , rather than on the bounding box B . The difference is minor when the box is close to the domain, which is usually the case.

$2(N_n + 1)$, where

$$(a_{i,j}) = \int_{\partial\Omega} \Phi_i \frac{\partial \Phi_j}{\partial n} ds, \quad (b_i) = \int_{\partial\Omega} \Phi_i w ds, \quad i, j = 1, \dots, N. \quad (8)$$

The resulting matrix A is symmetric, fully populated, complex valued and increasingly ill-conditioned as more wave functions are added to the approximation. Moreover, since the system is set up through evaluation of integrals, care needs to be taken to ensure sufficient accuracy in computations.

125 Regarding convergence, a sufficient condition stated in [3] for the approximation $\tilde{u}$ to converge to the exact solution u as $N \rightarrow \infty$ is that the domain Ω is convex. When this is not the case, several alternative approaches are available. A standard practice is to decompose the given domain into several convex subdomains, or, for more complex geometries, one can employ a multi-level
130 approach [13], or hybrid variants of the WBM [14].

3. Frames and numerical function approximation

The accuracy of a numerical simulation scheme for PDE's is greatly influenced by the extent to which the solution can be approximated by a linear combination of the chosen basis functions. In particular, the numerical solution
135 can not have greater accuracy than the best approximation to the solution in the approximation space, which is the span of the basis functions. Since Trefftz methods, and WBM in particular, employ non-standard basis functions (different from piecewise polynomials, say, which are well understood), we introduce a concept from approximation theory, *frames*. This will allow to investigate and
140 explain the convergence behaviour of WBM.

3.1. Definition and main properties

A frame $\Phi := \{\phi_k\}_{k=1}^{\infty}$ is a set of functions that is complete in a function space, such as the space $L^2(S)$ of square integrable functions on a domain S .² In addition, it should satisfy the so-called *frame property*:

$$A\|f\|_{L^2(S)}^2 \leq \sum_{k=1}^{\infty} \langle f, \phi_k \rangle^2 \leq B\|f\|_{L^2(S)}^2, \quad \forall f \in L^2(S). \quad (9)$$

This condition seems technical and not immediately intuitive: the main consequence of the frame property is that each function f in the space can be written as a linear combination of the frame elements ϕ_k , with coefficients that
145 are bounded in norm. The latter restriction, which we return to several times in this paper, is a crucial difference between a frame and any other redundant

²We denote the domain by S in this section and by Ω in the previous section, because they may refer to different domains. Indeed, the approximation problem in WBM is associated with the boundary $\partial\Omega$, hence in this context one has $S = \partial\Omega$.

Figure 2: If $S \subset B$ is a subdomain of the box B , then any basis for $L^2(B)$ is a frame for $L^2(S)$ [11]. This is true regardless of the geometric complexity of S . A popular choice on the box B is a tensor-product Fourier series. The approximation of a given function on S by a Fourier series on a larger bounding box is called the *Fourier extension problem*. While the construction of a basis on S is difficult for complex shapes, often leading to meshes and low-order piecewise polynomials, the construction of a frame is nearly trivial.

complete set of functions. We refer to [15] for a full mathematical treatment of frames, and to [11] for a recent review of the use of frames for numerical function approximation. In what follows, we focus on the main conclusions in [11].

Of importance for the current paper is the fact that frames may be overcomplete or redundant. Indeed, condition (9) does not even prevent linear dependencies between the functions ϕ_k . A simple frame with redundancy is obtained by adding two bases for the same space. Or, a basis can be augmented with a small number of additional functions. In both cases, a single basis is already complete and any addition to it is redundant. A third example is illustrated in Fig. 2: a frame on a domain can always be obtained from a basis on a larger domain. This setting is mathematically not the same as the construction of the basis functions in WBM, but it is certainly related.

Frames are much more flexible than bases, as the example of Fig. 2 shows where the domain S can have any shape. Unfortunately, the computation of the approximation to a function with N frame elements requires solving a linear system

$$Gc = B,$$

where the solution vector c yields the expansion coefficients for an approximation of the form $f \approx \sum_{k=1}^N c_k \phi_k(x)$. That is, unlike for an orthogonal basis, the coefficients c_k are not simply given by the inner products $\langle f, \phi_k \rangle$ (even though those inner products do appear in the frame property (9)). Moreover, if the frame is redundant, the matrix G is highly ill-conditioned or even singular. For these reasons the approximation in a frame is typically more computationally expensive than the approximation in a basis, though for the case of Fourier extensions fast algorithms have been devised [16, 17].

3.2. Computing a frame approximation

Two ways to compute a frame approximation are investigated in [11]. The first is the best approximation to f by projecting onto the span of the finite set $\Phi_N := \{\phi_k\}_{k=1}^N$. This leads to the linear system $G^p x = B^p$ where

$$G_{i,j}^p = \langle \phi_j, \phi_i \rangle_{L^2(S)} \quad \text{and} \quad B_i^p = \langle f, \phi_i \rangle_{L^2(S)}. \quad (10)$$

The square matrix $G^p \in \mathbb{C}^{N \times N}$ is called the *Gram matrix* in [11]. Both the elements of G^p and of the right hand side B^p require the numerical computation of integrals over the domain S , hence we refer to it as a continuous projection method.

The second approach is to compute the best approximation in a discrete, rather than continuous, least squares sense. To that end, one chooses a set $X_M := \{x_m\}_{m=1}^M$ of M points in S . This leads to the linear system $G^d x = B^d$,

$$G_{i,j}^d = \phi_j(x_i) \quad \text{and} \quad B_i^d = f(x_i). \quad (11)$$

Compared to the continuous projection, the discrete approach is both simpler and cheaper to implement, as it does not require the computation of integrals. However, the matrix $G^d \in \mathbb{C}^{M \times N}$ is now rectangular. We say that f is *oversampled* if $M > N$.

Both matrices G^p and G^d are ill-conditioned. When a large number of points M is used, the properties of the matrix G^p are similar to those of the normal equations $(G^d)^* G^d$ of the discrete least squares problem. This is common in least squares methods [18]. As a result, the condition number of G^p is almost the square of the condition number of G^d .

The systems $G^p c = B^p$ and $G^d c = B^d$ are usually solved with a direct solver. The method investigated in [11] is based on a truncated singular value decomposition (SVD) of the matrix. That is, the systems are regularised by discarding (truncating) all singular values below a threshold ϵ . The error bounds further on are derived for this truncated SVD solution. Numerical experiments show that, typically, slightly higher accuracy is obtained by using a rank-revealing pivoted QR-decomposition.³ For this reason, in our implementation, we have used a QR-based solver.

3.3. Error bounds

Can one expect highly accurate approximations, in spite of the potentially extreme ill-conditioning of G ? The surprising answer given in [11] is that, yes, in many cases you can, even when G is nearly singular. An intuitive explanation

³We like to stress here that the choice of the direct solver matters. For example, the solution $x = A \setminus B$ using the backslash-operator in MATLAB does not necessarily satisfy the bounds of this paper. The bounds are derived for a solution obtained using a truncated SVD, which yields a solution vector with minimal norm $\|x\|$. In fact, we found that for several examples in this paper, MATLAB's backslash solution was unstable due to growth of $\|x\|$ after an initial regime of convergence. In contrast, both the SVD and QR-based solution remain stable for larger systems and hence should be preferred for this class of ill-conditioned matrices.

is that, because of the redundancy of frames, there are many ways in which f can be expressed as a linear combination of the frame elements with comparable accuracy. Thus, it does not matter which solution of $Gx = B$ is found. Instead, the approximation error is governed by the residual $Gx - B$. Achieving a small residual is a considerably weaker condition than finding the exact solution of an ill-conditioned system of equations. Indeed, the ill-conditioning implies that the matrix G has a large numerical kernel, yet any element of the kernel (satisfying $Gx \approx 0$) has no impact on the residual.

Even if G is mathematically non-singular, and thus the system has a unique exact solution, finding this particular solution numerically is not feasible. The *regularized* solution found by truncating the SVD may in fact differ substantially from the exact solution, at least when comparing the coefficients. Yet, in most cases one does not a-priori care which set of coefficients is found, as long as the approximation $\sum_{k=1}^N c_k \phi_k(x)$ is accurate, *and* as long as the coefficients c_k do not grow too large. One typically prefers small coefficients c_k , since large coefficients with alternating signs may yield large cancellation errors in subsequent computations.

These observations lead to two criteria: we are looking for a solution x with small residual $Gx - B$, and with small norm $\|x\|$.

A first generic error bound formulated in [11] for the approximation of a function f via the truncated SVD of the projection matrix G^g is:

$$\|f - P_N^{p,\epsilon} f\|_{L^2(S)} \leq \|f - \sum_{k=1}^N z_k \phi_k\|_{L^2(S)} + \sqrt{\epsilon} \|z\|, \quad \forall z \in \mathbb{C}^N. \quad (12)$$

Here, the notation $\|f - P_N^{p,\epsilon} f\|_{L^2(S)}$ in the left hand side represents the approximation error, i.e., the difference between f and its approximation obtained via regularised SVD with truncation parameter ϵ , measured in the L^2 norm on S . The inequality (12) holds for any vector z , and hence shows that the solution will be at least as good as any approximation to f in the span of the truncated frame Φ_N , as long as the size of the coefficients $\|z\|$ remains small. One finds a best approximation, subject to having small coefficients. The maximal achievable accuracy is on the order of $\sqrt{\epsilon}$.

A second generic error bound applies to the oversampled least squares problem (11). It is slightly more involved and introduces more notation:

$$\begin{aligned} \|f - P_{M,N}^{d,\epsilon} f\|_{L^2(S)} &\leq \|f - \sum_{k=1}^N z_k \phi_k\|_{L^2(S)} + \kappa_{M,N}^\epsilon \|G^d z - B^d\| \\ &\quad + \epsilon \lambda_{M,N}^\epsilon \|z\|, \quad \forall z \in \mathbb{C}^N. \end{aligned} \quad (13)$$

As above, the quantity in the left hand side is the approximation error, measured on S . The right hand side also features the residual of the linear system corresponding to vector z . In addition, there are two constants, $\kappa_{M,N}^\epsilon$ and $\lambda_{M,N}^\epsilon$. These are difficult to analyse generically. However, in the case where Φ is a frame, these constants are bounded for sufficiently large M (with M large compared to N). This corresponds to sufficiently high oversampling.

Crucially, the maximal achievable error of the oversampled discrete scheme is on the order of ϵ in (13), rather than $\sqrt{\epsilon}$ in (12). This is the main reason for the observation that, when manipulating expansions that are possibly redundant, it is worthwhile to introduce oversampling into the discretisation of the problem.

Results are also included in [11] on the size of the solution vector, which we do not repeat here. Briefly, before the onset of convergence, the norm of the solution vector may in fact be as large as $1/\epsilon$ or $1/\sqrt{\epsilon}$. However, once N is sufficiently large such that the approximation starts to converge, the norm of the expansion coefficients decreases to a value that can be related to $\|f\|_{L^2(S)}$, the continuous norm of f itself on the domain S . Thus, after potential initial divergence, the regularised computation leads to a numerical approximation that is well-behaved for postprocessing purposes, i.e., it does not suffer from large cancellation errors. In view of the ill-conditioned linear system, this is a very desirable outcome.

4. A collocation approach of WBM

In the restricted setting of this paper, i.e., in the absence of a multilevel approach or hybrid couplings, the Wave Based Method can be seen as an approximation problem: one approximates a function on the boundary $\partial\Omega$ of a domain Ω , using a set of functions defined on a larger bounding box. Indeed, compare the linear system of the weighted residual formulation (8) to the Gram matrix in (10). If the expansion succeeds in matching the given Neumann data then, by construction, the expansion also satisfies the Helmholtz problem and a solution to the boundary value problem (1)–(2) is found.

Two drawbacks of the WBM are that (i) it relies on accurate integration techniques and (ii) it seems restricted to convex domains. The collocation formulation of the WBM starts by addressing the former one. The idea is straightforward, by extending the comparison to the frame approximation problem and considering a discrete least squares variant. To that end, the boundary condition is enforced in a set of M points distributed along the boundary, referred to as collocation points $\{s_k\}_{k=1}^M$. This leads to a linear system of equations $\tilde{A}\alpha = \tilde{b}$, where $\tilde{A}$ is an $M \times N$ matrix and $\tilde{b}$ is a vector, with entries given by

$$\tilde{a}_{i,j} = \phi_j(s_i) \quad \text{and} \quad \tilde{b}_i = v(s_i), \quad i = 1, \dots, M, \quad j = 1, \dots, N. \quad (14)$$

These are the expressions for a Dirichlet problem. The analogous formulation for Neumann boundary conditions is similar, but using the normal derivatives of the functions ϕ_j instead.

Bearing in mind the similarities to the discrete least squares approximation problem (11), it is important here to oversample. That is, we use more collocation points than there are degrees of freedom by a factor γ , $M = \gamma N$ with $\gamma > 1$. The theory of frames does not generically stipulate what γ should be, nor that such linear oversampling is even sufficient.⁴ The value of γ is problem

⁴By linear oversampling we mean the linear relation $M = \gamma N$ between M and N . In-

dependent. In this paper, unless stated otherwise, we shall use a value of $\gamma = 2$ and we observe numerically that this appears sufficient for our experiments.

260 A first immediate benefit from this approach is that the collocation matrix $\tilde{A}$ is less severely affected by ill-conditioning than the WBM matrix A . This follows from similar considerations discussed for frame approximations. A second benefit is computational: the cost of assembling $\tilde{A}$ is significantly lower, since each element of the matrix consists of one function evaluation, whereas in the
265 weighted residual formulation each element requires the numerical approximation of an integral over the boundary of the domain.

While our formulation of the oversampled collocation formulation is based on the connection to frames, we like to point out that other authors have made similar implementations for Trefftz methods. For example, the method of fun-
270 damental solutions as implemented in [8] is based on oversampled collocation.

5. Results

We apply both the weighted residual and the oversampled collocation formulations of WBM for a number of two-dimensional Helmholtz problems on smooth domains. We use either Dirichlet or Neumann boundary conditions. We focus
275 on two properties: convergence of the scheme as a function of the number of degrees of freedom, and size $\|x\|$ of the coefficient vector in the discrete l^2 -norm.

To that end, we devise a number of scattering configurations for which we can show a priori whether or not a solution vector with bounded norm exists. Inspired by the analysis in [8], the problems involve two kinds of singularities:
280 singularities induced by the boundary data, and singularities induced by the (non-convex) shape of the domain. For the time being, we do not consider singularities due to corners or other singularities of the domain shape itself.

All experiments were performed in MATLAB. The linear systems for the weighted residual formulation are solved with a custom rank-revealing QR-
285 decomposition with column pivoting, with threshold manually set to 2×10^{-13} , a value that was experimentally determined. The collocation systems were solved simply using MATLAB's backslash operator, which runs a comparable algorithm for this class of matrices. The integrals (8) involved with the classical WBM system are approximated using the trapezoidal rule with sufficiently
290 many points, since all integrands are smooth and periodic [19].

Since the wave functions satisfy exactly the Helmholtz equation inside the domain, errors are only introduced by not exactly satisfying the boundary conditions. Though these errors propagate towards the interior of the domain, we measure the relative error on the boundary data in a set of points $\{(x_i, y_i)\}_{i=1}^{n_p}$ on $\partial\Omega$, with n_p larger than the number of collocation points used to compute

deed there are known cases where M should grow like N^2 . Needless to say, the amount of oversampling negatively affects the computational cost of the scheme.

the approximation. For example, for Dirichlet data, this results in the metric

$$\varepsilon = \frac{\|\tilde{u}(x_i, y_i) - w(x_i, y_i)\|_{C^{n_p}}}{\|w(x_i, y_i)\|_{C^{n_p}}}, \quad (15)$$

where w denotes the function imposed on the boundary.

In several experiments, we choose boundary conditions that correspond to a known solution of the Helmholtz equation that is bounded in the interior of Ω . Hence, by construction, we know the exact solution of the interior boundary value problem and, moreover, we know the extension of this solution outside the domain Ω .

In this section we describe the numerical experiments and the results. We elaborate on their interpretation further on in §6.

5.1. Convex obstacles

We consider convex obstacles first and start by examining the effect of the size of the bounding box. This experiment is inspired by the analogy to Fourier extension frames, illustrated in Fig. 2 which should be compared with Fig. 1. For Fourier extension approximations, faster convergence rates are seen for larger bounding boxes [20, 21].

The domain is the unit disk centred at the midpoint of a square bounding box. We choose a wavenumber $k = 0.924$ and Neumann boundary conditions consisting of the normal derivative of a plane wave propagating at an angle of 0.3 radians. Thus, the exact solution on the disk is this plane wave. Results are shown in Fig. 3 as a function of N , the total number of degrees of freedom in the discretisation, and for different sizes of the edge length of the bounding square. The results show the accuracy, the condition number of the linear system and the l^2 -norm of the coefficient vector.

One observes in Fig. 3(a) that the collocation approach achieves higher accuracy than the weighted residual formulation, on the order of 10^{-15} compared to 10^{-8} . Higher convergence rates are seen for larger box sizes. Panel (b) shows very large condition numbers, and confirms that $\kappa(A) \approx \kappa(\tilde{A})^2$. Finally, panel (c) shows that the computed solution vectors have a norm that remains bounded as N increases, for both variations of WBM.

Next, we consider the same circular domain centred at $(1.5, 1.5)$, and enclosed in a $[0, 3] \times [0, 3]$ box of fixed size. This time we employ Dirichlet boundary conditions arising from a function satisfying the Helmholtz equation in the plane, yet with a singularity at a point (x_s, y_s) . In particular, we use the Hankel function of the first kind and order zero, $w(x, y) = H_0^{(1)}(k\|(x, y) - (x_s, y_s)\|)$. We examine convergence and norm of the solution when varying the location of the singularity, which corresponds to a point source. The ordinate $y_s = 1.5$ is kept constant, while $x_s \in [-1, -0.01, 0.01, 0.2, 0.4]$.

The results are shown in Fig. 4. Here, too, the collocation approach converges to higher accuracy as shown in panel (a). However, both implementations appear to have a maximal accuracy that depends on the location of the singularity. When the singularity is furthest away from the box ($x_s = -1$), the convergence curves are comparable to those in Fig. 3(a). As the source approaches

Figure 3: The solution on the unit disk with wavenumber $k = 0.924$, embedded in a square bounding box with edge lengths varying between 2 and 3.5. Solid lines correspond to weighted residual WBM, dashed line to the collocation approach.

the box, both outside ($x_s = -0.01$) and inside ($x_s = 0.01$) the box, a smaller convergence rate is seen, yet similar accuracy is ultimately achieved. However, in panel (b), a small bump is seen in the corresponding curves for small T : the solution vector $\|x\|$ initially grows, before settling down for increasing T to a small value comparable to the first case of a far away singularity. The bump is larger for the collocation approach.

The impact of the singularity worsens as it approaches the disk even more. For $x_s = 0.2$ and $x_s = 0.4$, the convergence curves in Fig. 4 level off at larger and larger values. It remains the case that collocation reaches higher accuracy than the weighted residual method. However, this comes at a price, seen in panel (b): the norms of the solution vectors settle down at significantly larger levels.

As the singularity approaches the domain, both methods become increasingly inaccurate and increasingly unstable. The best accuracy achieved for both formulations of the WBM is restricted by the growth of the norm of the coefficients. In this case, convergence stops when $\frac{\varepsilon}{\alpha} \approx 10^{-13}$ for collocation, and $\frac{\varepsilon}{\alpha} \approx 10^{-8}$ for the weighted residual approach, where α is the limiting value

Figure 4: WBM results using the weighted residual (line) and collocation (dashed line) formulations on a disk with a point source and $k = 0.924$; color varies with the abscissa of the source point: -1 (red), -0.01 (yellow), 0.01 (purple), 0.2 (green), 0.4 (light blue).

of $\|x\|$ for large T . These regimes corresponds precisely to an equal balance of
 350 the terms appearing in the right hand sides of (12) and (13), respectively.

5.2. Non-convex obstacles

In the previous experiment, the solution of the boundary value problem has a singularity in the exterior of the domain Ω , determined by the boundary condition. It is known that the solution of a Helmholtz problem on a non-convex
 355 domain also develops singularities in its extension to the exterior. However, in this case, the locations of the singularity are determined by the geometry of the domain, in particular its boundary, and not by the boundary condition data. We refer the reader to [8] and the references therein for an extensive description of known results. We adopt some examples from the same reference to illustrate
 360 the impact of these singularities in the Wave Based Method.

We show in Fig. 5 the domain referred to as ‘crescent’ in [8]. The parameterisation of its boundary is given by

$$f_1(t) = z_0 + e^{it} - \frac{a}{e^{it} + b}, \quad t \in [0, 2\pi], \quad (16)$$

where $z_0 = 1.5 + 1.5i$ and a and b are two parameters. Here, $f_1(t) = x(t) + iy(t)$ is a complex-valued function and we identify its real and imaginary parts with coordinates in the Euclidean plane $\mathbb{R}^2$.

The singularities associated with a non-convex domain that has an analytic
 365 parameterisation (analytic in the sense of complex analysis) are those of the so-called Schwartz function. For a given domain Ω with boundary $f(t)$, where f is an analytic function of t and $|f'(t)| \neq 0$, the Schwartz function is defined as $S(z) = \bar{f}(f^{-1}(z))$. It is a function of z in the complex plane, and may have singularities such as branch points and poles. We omit their computation, but
 370 highlight for the crescent domain in Fig. 5(a)–(c) the location of a pole in the exterior of Ω , and of two branch points of square root type in the interior of

Figure 5: The relevant singularities for a non-convex domain are those of the associated Schwartz function. Shown here is the crescent domain of [8] and its singularities: a pole in the exterior, denoted by '+', and two branch point singularities in the interior denoted by 'x'. The parameterisation of the boundary is given by (16) and contains two parameters a and b . In all cases shown, the crescent domain is enclosed by a square with edge length 3.

Ω , for different combinations of the parameters a and b . The branch points are located at $z_{b1,2} = z_0 - b \pm 2i\sqrt{a}$ and the pole at $z_p = z_0 - \frac{a}{b}$.

The solution of equation 1 with constant Dirichlet boundary condition, $w(x, y) = 1$, for $(x, y) \in \partial\Omega$, is approximated for each variation of the crescent domain. The results are shown in Fig. 6. The error is calculated using $n_p = 3000$ boundary points.

As the domain becomes more non-convex, which corresponds to the pole of the Schwartz function moving closer to the domain, the maximal accuracy of WBM decreases. Similar to the case of convex domains with nearby singularities, the collocation approach yields higher accuracy, at the cost of producing a solution with larger coefficient norm.

Figure 6: WBM results using the weighted residual (line) and collocation (dashed line) formulations with constant Dirichlet boundary data $w(x, y) = 1$ and $k = 0.924$, for the three variations of the crescent domain shown in Fig. 5: I (blue), II (red), III (yellow).

In the case of crescents I and II, it is possible to choose a bounding box

that encloses the domain but not the poles of the Schwartz function. In the
 385 next experiment the box is centred at $z_0 = 0.64 + 1.75i$, with sides of length
 $L_x = 1.4$ and $L_y = 3.5$. The results in Fig. 7 show that in both cases full
 expected accuracy is achieved (of order $\sqrt{\epsilon}$ and order ϵ , respectively, for weighted
 residual and collocation formulations). This is in spite of the non-convexity
 of the domain and in spite of the ill-conditioning of the linear systems. The
 390 coefficient vector has small norm for a wide range of the truncation parameter
 T . The convergence rate is lower compared to the results in Fig. 6, because the
 bounding box is tighter. This is in agreement with the first experiment of this
 section, in which we varied the size of the box.

Figure 7: WBM results using weighted residual (line) and collocation (dashed line) formulations with constant Dirichlet boundary data $w(x, y) = 1$ and $k = 0.924$, for crescents I (blue) and II (red), and with a bounding box chosen such that the singularities are located outside.

This experiment suggests that the convexity requirements in the Wave Based
 395 Method is not a fundamental restriction. Rather, at least in the current setting
 of smooth obstacles, convergence and stability are simultaneously hampered by
 the presence of singularities in the extension of the Helmholtz solution, exterior
 to Ω but within the bounding box. We elaborate on this point in §6.

We further illustrate this property with another example domain from [8],
 the ‘inverted ellipse’. The domain is illustrated in Fig. 8. Its boundary is
 parameterised by the complex-valued function

$$f_2(t) = z_0 + \frac{e^{it}}{1 + \tau e^{2it}}, \quad t \in [0, 2\pi],$$

where we choose $z_0 = 1 + 1.75i$ and τ is a parameter.

400 We consider two values $\tau = 0.25$ and $\tau = 0.35$ and choose the bounding box
 to be $[0, 2] \times [0, 3.5]$. Computation of the Schwarz function shows two branch
 point singularities at $z_{b1,2} = z_0 \pm \sqrt{\frac{1}{4\tau}}$. For the first ellipse ($\tau = 0.25$), the
 singularities are exactly on the edge of the box, whereas in the second case
 ($\tau = 0.35$) they are located in the interior of the box. Indeed, the second
 405 inverted ellipse deviates more from a convex domain than the first one.

The corresponding simulation results are shown in Fig. 9. At this stage, the pattern looks familiar. The collocation method achieves higher accuracy. Although the convergence rate is rather slow, and although the corresponding matrices for increasing T become larger and increasingly ill-conditioned, for the first ellipse the method ultimately achieves machine precision accuracy. This does not happen for the second ellipse, where the Schwarz singularity is inside the box, and accuracy is limited to about 10^{-7} . The limiting size of the coefficient norm is correspondingly larger. The results for weighted residual formulation are similar, but with maximal accuracy of $\sqrt{\epsilon}$ and 10^{-5} respectively and, as before, with a smaller limiting norm of the solution vector compared to the collocation approach.

In this final experiment we have increased the oversampling factor to $\gamma = 4$. This may be due to the fact that we chose equispaced points in the parameter domain $t \in [0, 2\pi]$, yet the corresponding points on the boundary $\partial\Omega$ are unevenly distributed. Still, recall that the unknown constants $\kappa_{M,N}^\epsilon$ and $\lambda_{M,N}^\epsilon$ in (13) are known to be bounded only in the limit of $M \rightarrow \infty$. As of yet, there is no theoretical guarantee that linear oversampling is sufficient, unless γ can be arbitrarily increased when convergence is not seen.

Figure 8: Two inverted ellipses and the location of the two branch points of the corresponding Schwartz functions: they are on the box boundary (left panel) or in its interior (right panel).

6. Discussion

We focus on the interpretation of the results in this paper, rather than on the mathematical analysis. Though the set of WBM basis functions does not constitute a frame for the relevant function spaces involved, which we do not show in this paper, the generic error bounds (12) and (13) do apply.

Figure 9: WBM results using weighted residual (line) and collocation (dashed line) formulations with constant boundary data and wavenumber $k = 0.924$ for the inverted ellipse.

6.1. Convergence to accuracy on the order of ϵ versus $\sqrt{\epsilon}$

Arguably, the matter of numerical stability is more important than the distinction between very high and even higher precision. Indeed, our choices of the truncation parameter are much larger than is customary in the literature on WBM, since extremely high accuracy is rarely requested in applications. Still, the bound (13) for the discrete least squares approximation shows maximal convergence to a value that is proportional to machine precision, and the results in Fig. 3 show that this maximal accuracy is obtained using computations in standard double floating point precision only. This is remarkable, in view of the extreme ill-conditioning of the matrices involved as illustrated in Fig. 3(b). Indeed, at some point the matrices are numerically indistinguishable from being singular.

The results for the weighted residual formulation show convergence to roughly $\sqrt{\epsilon}$, as predicted by (12). We reiterate that, in order to obtain these results, we resorted to a direct solver that minimises the norm $\|x\|$ of the solution vector, while also minimising the approximation error which corresponds to the residual. Achieving both objectives simultaneously is possible only due to the redundancy in the discretisation: both systems are numerically *underdetermined*, i.e., the number of singular values of A or $\tilde{A}$ larger than a threshold is actually smaller than the number of degrees of freedom, even when oversampling. This redundancy arises in WBM from restricting the bounding box to a smaller domain and it is crucial: it enables the freedom in a solver to minimise $\|x\|$, while maintaining a small residual. We refer to [11] for a detailed analysis of this effect.

We add that it is also possible to oversample the weighted residual formulation. Indeed, the number of functions of the form (4) is infinite and we have truncated the set using (6). The formulation that was used in this paper leads to square matrices, because we used an equal number of functions in the expansion (3) as in the weighting function in (7). However, we could use more functions for the weight, leading to a rectangular system as in the oversampled

collocation approach. A similar bound to (13) holds in this case [12], with maximal accuracy on the order of ϵ . This could be a viable approach for Trefftz methods based on a variational formulation, for which a collocation approach is not easily available.

6.2. Singularities of the solution

As mentioned in the introduction, a crucial question is the following: does a solution to the linear system exist, with small residual and with small norm? In the absence of singularities on the boundary of the domain itself, the answer to that question lies in the location of singularities in the extension of the Helmholtz solution from the domain S to the bounding box B (recall Fig. 1).

The set of functions in the Wave Based Method is the truncation of a complete set in the bounding box, complete in the sense that the solution to any boundary value problem with Neumann data on the four edges of the rectangle can be represented [3, §2.5.2]. The extension of the solution of the Helmholtz equation on a domain Ω to a larger domain is unique and analytic, when the domain is sufficiently smooth such that it admits a parameterisation that is analytic. Note that the expansion (3) satisfies the Helmholtz equation not only on Ω , but on the whole of B . Therefore, expansion (3) necessarily approximates not only the true solution on Ω , but also its extension to B . A sufficient condition for the existence of an approximate expansion with bounded norm coefficients is that this extension is bounded and free of singularities on B .⁵

In practice, the situation is slightly more involved. Though the extension of the Helmholtz problem is mathematically unique, there is some ‘wobble room’, for lack of a better terminology. A weaker condition than the above, but also sufficient, is that an extension exists whose Dirichlet or Neumann trace on Ω agrees to within an ϵ tolerance with the given boundary condition. Finally, it is sufficient for an approximation to such an extension to exist for finite N , corresponding to the chosen value of the truncation parameter T , and the extension may vary with N (or equivalently with T).

It seems that the *wobble room* increases quantitatively if the singularity is further away from the domain Ω . An important qualitative difference is whether the singularity of the Helmholtz extension lies within the box B or outside.

These considerations were illustrated in Fig. 4. We solve an interior problem, for which we know that the unique extension has a point source singularity outside Ω and we vary the location of that singularity. Convergence to high accuracy is shown in Fig. 4(a) when the point source is outside the box. The minimal achievable error deteriorates as the point source moves inside the box and closer to the domain Ω . Correspondingly, the size $\|x\|$ of the solution vector grows as shown in Fig. 4(b). The maximal accuracy of the approximation problem is negatively correlated with the distance to the nearest singularity.

⁵The set of plane waves actually satisfy the Helmholtz equation on the whole plane $\mathbb{R}^2$. However, we only have a guarantee that they can represent any solution to a Neumann problem on B , not on any larger domain. Note that the size of the box affects the distribution of the plane waves, via the edge lengths L_x and L_y in (5).

The fact that problems exist for which the solver does not converge to ϵ accuracy indicates numerically that the WBM set of functions is not a frame. Indeed, if it were, then convergence to maximal precision should be seen for sufficiently large N for any boundary condition. Mathematically, the existence of such approximations is precisely guaranteed by the frame property (9).

If nearby singularities arise in a practical computation, the solution is to enrich the approximation space, i.e., to add other solutions of the Helmholtz equation to the expansion. The goal should be to build an approximation space in which the exact solution can be approximated with small norm coefficients.

6.3. Non-convex domains

The results show that the Wave Based Method does not converge to high accuracy for all non-convex domains, but it does for some. Exterior extensions of solutions to the Helmholtz equation almost always develop singularities at certain points, determined by the geometry of the domain. If these points are outside the bounding box, then convergence is not at stake. If they are inside the box then, as for convex domains with exterior singularities, there is a maximal achievable accuracy that deteriorates with decreasing distance of the singularity. The underlying cause is the same as in the convex case: convergence is limited by the growth of the norm of the solution vector, i.e., of the size of the coefficients in expansion (3).

Whether exterior singularities present a problem in practice can perhaps be determined a-priori based on physical insights. Unfortunately, the computation of their location for non-convex obstacles is not straightforward, and for this reason it remains advisable to restrict WBM to convex obstacles. Still, if a singularity can be expected inside the box, a possible solution may be to enrich the approximation space. For the method of fundamental solutions, the solution proposed in [8] is to position the charge points in between the domain Ω and the singularities. It seems plausible to assume that adding such singular solutions to the approximation space of WBM would improve convergence.

Alternatively, the issues associated with non-convex domains and solutions with nearby singularities may be treated using h-refinement (subdividing the computational domain), rather than p-refinement (increasing the number of degrees of freedom on a single domain). Indeed this is an approach taken by several Trefftz methods [1]. The interplay between h- or hp-refinement on the one hand, and frames and redundancy on the other hand, remains to be investigated.

6.4. The frame property

We have not included a mathematical analysis of the way in which the frame property (9) is violated. In the case of a Dirichlet boundary condition with Dirichlet data in the function space $L^2(\partial\Omega)$, say, one should investigate whether the traces of the chosen plane waves along the boundary $\partial\Omega$ constitute a frame for that space. The relevant function space is associated with the boundary of the domain, not with the domain itself.

Apart from the singularities that we focused on in our experiments, a major reason for why these traces may not be a frame in several Trefftz methods is

the existence of evanescent waves. If the PDE allows solutions with shorter wavelengths in one direction, at the cost of exponential decay or increase in
545 the other direction, then the traces of the Trefftz functions have to be able to approximate these solutions with bounded coefficients. However, if all Trefftz functions have a wavelength on the order of the inverse of the wavenumber, which is the case when they are plane waves with the same wavenumber, then so do their traces. As a result, approximating functions with shorter wavelengths
550 necessarily leads to large coefficients. Hence, the frame property can not hold for *all* functions in the relevant space.

Interestingly, the Wave Based Method does include evanescent waves in its discretization whenever k_x or k_y become imaginary, recall (5). However, by construction these evanescent waves are aligned with the X and Y-axes, whereas
555 presumably they would have to be distributed along the boundary of the domain for the frame property to hold for $L^2(\partial\Omega)$. The analysis of the frame property for plane wave basis functions is a topic of future research.

We would like to emphasize that our comments merely represent the mathematical manifestations of a known problem. Indeed, several methods have
560 already been proposed based on adding evanescent waves to discretizations to improve accuracy, for example [22, 23]. We only remark that the inclusion of evanescent waves may improve accuracy and numerical stability of a scheme, even when such waves are seemingly not physically present in the system under consideration.

565 7. Conclusions

Ill-conditioning of the linear systems arising in Trefftz methods is often perceived as a cause of concern. We show in this paper that, for at least one Trefftz method, this need not be so: numerical convergence can be guaranteed for a range of problems, in spite of potentially extreme ill-conditioning. We hasten
570 to add that this is certainly not a general observation, and in particular it is limited only to approximation problems involving an approximation space that is, in some sense that can be made precise, redundant. For any other type of problems, ill-conditioning is, and always will be, a major potential source of loss of accuracy in numerical computations.

We show that accurate solutions are found, if they exist in the span of the chosen basis functions, and if their expansion coefficients are not too large. In turn, this means that the limitations of a method can be uncovered by studying when this is not the case. With the experiments of this paper, we have focused on the presence of singularities. There may be other reasons why the exact
580 solution can not be well approximated with small coefficients, and a prominent difficulty arises in applications involving evanescent waves.

The beneficial aspects of oversampled collocation approaches have been advocated in literature before, e.g. [8]. Enriching approximation spaces to capture properties of a solution (resulting in smaller coefficients in the corresponding
585 approximate expansion) have also been explored, for example by adding corner singularities [24]. The main contribution of this paper is the increased level of

(mathematical) confidence with which one can make statements about convergence, accuracy and numerical stability of a Trefftz method in the presence of numerical ill-conditioned linear systems.

590 Through experiments, we could show applicability of WBM beyond convex domains, a limitation that is commonly accepted in literature, as long as the non-convex domain is not too non-convex. We showed improved accuracy of oversampled collocation. Still, the scope of this paper was limited and the contents led to several topics of ongoing and future research. In particular,
 595 we did not examine the influence of corner singularities in the solution. We did not analyse the oversampling factor γ (recall that $M = \gamma N$ with $\gamma > 1$). The generic error bounds we used do not indicate the convergence rate of the solution. We have focused solely on p-refinement. Finally, we have studied only a very restricted subset of the many available settings in which the Wave Based
 600 Method has been successfully applied.

Acknowledgements

The authors greatly appreciate discussions on the topic of this paper with Ben Adcock, Alex Barnett, Timo Betcke, Vincent Coppé, Elke Deckers, Wim Desmet, Andrew Gibbs, David Hewett, Ralf Hiptmair, Stijn Jonckheere, Roel
 605 Matthysen, Andrea Moiola and Marcus Webb. The first author was supported by FWO-Flanders projects G.0641.11 and G.A004.14, as well as by KU Leuven project C14/15/055.

- [1] R. Hiptmair, A. Moiola, I. Perugia, Building Bridges: Connections
 610 and Challenges in Modern Approaches to Numerical Partial Differential Equations, Vol. Lecture notes in computational science and engineering, Springer, 2016, Ch. A Survey of Trefftz Methods for the Helmholtz Equation, pp. 237–279.
- [2] G. Fairweather, A. Karageorghis, P. A. Martin, The method of fundamental solutions for scattering and radiation problems, Eng. Anal. Boundary Elem.
 615 27 (2003) 759–769.
- [3] W. Desmet, A wave based prediction technique for coupled vibroacoustic analysis, Ph.D. thesis, KU Leuven (1998).
- [4] E. Deckers, O. Atak, L. Coox, R. D’Amico, H. Devriendt, S. Jonckheere, K. Koo, B. Pluymers, D. Vandepitte, W. Desmet, The Wave Based Method:
 620 An overview of 15 years of research, Wave Motion 51 (2014) 550–565.
- [5] O. Cessenat, B. Després, Application of an Ultra Weak Variational Formulation of elliptic PDEs to the two-dimensional Helmholtz problem, SIAM J. Numer. Anal. 35 (1) (1998) pp. 255–299.
- [6] R. Hiptmair, A. Moiola, I. Perugia, C. Schwab, Approximation by harmonic
 625 polynomials in star-shaped domains and exponential convergence of Trefftz hp-dGFEM, Math. Model. Numer. Anal. (M2AN) 48 (2014) 727–752.

- [7] R. Hiptmair, A. Moiola, I. Perugia, Plane wave discontinuous Galerkin methods: exponential convergence of the hp-version, *Found. Comput. Math.* 16 (3) (2016) 637–675.
- 630 [8] A. H. Barnett, T. Betcke, Stability and convergence of the Method of Fundamental Solutions for Helmholtz problems on analytic domains, *J. Comput. Phys.* 227 (2008) 7003–7026.
- [9] A. H. Barnett, T. Betcke, An exponentially convergent nonpolynomial finite element method for time-harmonic scattering from polygons, *SIAM J. Sci. Comput.* 32 (3) (2010) 1417–1441.
- 635 [10] P. C. Hansen, Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion, SIAM, Philadelphia, 2005.
- [11] B. Adcock, D. Huybrechs, Frames and numerical approximation, Tech. Rep. TW-674, KU Leuven (December 2016).
- 640 [12] B. Adcock, D. Huybrechs, Frames and numerical approximation II: generalized sampling, Tech. rep. (2018).
- [13] B. V. Genechten, K. Vergote, D. Vandepitte, W. Desmet, A multi-level wave based numerical modelling framework for the steady-state dynamic analysis of bounded Helmholtz problems with multiple inclusions, *Comput. Methods Appl. Mech. Eng.* 199 (2010) 1881–1905.
- 645 [14] B. van Hal, Automation and performance optimization of the Wave Based Method for interior structural-acoustic problems, Ph.D. thesis, KU Leuven (2004).
- [15] O. Christensen, Frames and Bases: an Introductory course, Springer, Basel, 2008.
- 650 [16] R. Matthysen, D. Huybrechs, Fast algorithms for the computation of Fourier extensions of arbitrary length, *SIAM J. Sci. Comput.* 38 (2) (2016) A899–A922.
- [17] R. Matthysen, D. Huybrechs, Function approximation on arbitrary domains using Fourier frames, *SIAM J. Numer. Anal.*
- 655 [18] C. L. Lawson, R. J. Hanson, Solving least squares problems, Classics in Applied Mathematics, SIAM, Philadelphia, 1996.
- [19] L. N. Trefethen, J. A. C. Weideman, The exponentially convergent trapezoidal rule, *SIAM Review* 56 (2014) 385–458.
- 660 [20] D. Huybrechs, On the Fourier extension of non-periodic functions, *SIAM J. Numer. Anal.* 47 (6) (2010) 4326–4355.
- [21] B. Adcock, D. Huybrechs, On the resolution power of Fourier extensions for oscillatory functions, *J. Comput. Appl. Math.* 260 (2014) 312–336.

- 665 [22] T. Luostari, T. Huttunen, P. Monk, Improvements for the ultra weak variational formulation, *Internat. J. Numer. Methods Engrng* 94 (2013) 598–624.
- [23] R. Tezaur, L. Zhang, C. Farhat, A discontinuous enrichment method for capturing evanescent waves in multiscale fluid and fluidsolid problems, *Comput. Methods Appl. Methc. Engrg.* 197 (2008) 1680–1698.
- 670 [24] E. Deckers, B. Bergen, B. V. Genechten, D. Vandepitte, W. Desmet, An efficient Wave Based Method for 2D acoustic problems containing corner singularities, *Comput. Methods Appl. Mech. Engrg.* 241–244 (2014) 286–301.