High-frequency asymptotic compression of dense BEM matrices for general geometries without ray tracing

Daan Huybrechs · Peter Opsomer

Received: date / Accepted: date

Abstract Wave propagation and scattering problems in acoustics are often solved with boundary element methods. They lead to a discretization matrix that is typically dense and large: its size and condition number grow with increasing frequency. Yet, high frequency scattering problems are intrinsically local in nature, which is well represented by highly localized rays bouncing around. Asymptotic methods can be used to reduce the size of the linear system, even making it frequency independent, by explicitly extracting the oscillatory properties from the solution using ray tracing or analogous techniques. However, ray tracing becomes expensive or even intractable in the presence of (multiple) scattering obstacles with complicated geometries. In this paper, we start from the same discretization that constructs the fully resolved large and dense matrix, and achieve asymptotic compression by explicitly localizing the Green's function instead. This results in a large but sparse matrix, with a faster associated matrix-vector product and, as numerical experiments indicate, a much improved condition number. Though an appropriate localisation of the Green's function also depends on asymptotic information unavailable for general geometries, we can construct it adaptively in a frequency sweep from small to large frequencies in a way which automatically takes into account a general incident wave. We show that the approach is robust with respect to non-convex, multiple and even near-trapping domains, though the compression rate is clearly lower in the latter case. Furthermore, in spite of its asymptotic nature, the method is robust with respect to low-order discretizations such as piecewise constants, linears or cubics, commonly used in applications. On the other hand, we do not decrease the total number of degrees of freedom compared to a conventional classical discretization. The combination of the sparsifying modification of the Green's function with other accelerating schemes, such as the fast multipole method, appears possible in principle and is a future research topic.

Daan Huybrechs
Department of Computer Science, KU Leuven, Belgium
E-mail: daan.huybrechs@cs.kuleuven.be

Peter Opsomer
Department of Computer Science, KU Leuven, Belgium
E-mail: peter.opsomer@cs.kuleuven.be

Keywords Boundary Element Method · Oscillatory integration · High-frequency scattering · Compression · Condition number · Smooth window functions

Mathematics Subject Classification (2000) 65N38 · 65F50 · 45A05 · 45M05 · 65R20

1 Introduction

Numerical simulations in acoustics are often based on a Boundary Integral equation reformulation of the Helmholtz equation, see for example [12, 40]. An incoming wave that is scattered by an obstacle with boundary Γ results in a scattered field $u^s(\mathbf{x})$ that can be represented by the so-called *single layer potential*,

$$u^s(\mathbf{x}) = \int_{\Gamma} K(\mathbf{x}, \mathbf{y}) v(\mathbf{y}) ds(\mathbf{y}). \quad (1)$$

Here, $K(\mathbf{x}, \mathbf{y})$ is the Green's function of the Helmholtz equation with wavenumber k and $v(\mathbf{y})$ is the unknown *density function* defined on Γ . It is a physical quantity, namely the normal derivative of the total wave. Sound-soft and sound-hard obstacles give rise to Dirichlet or Neumann boundary conditions respectively, with zero pressure or zero normal velocity on Γ . Integral representations of the scattered field such as the one above, coupled with a boundary condition, give rise to an integral equation. In this paper, we present our results for the following integral equation of the first kind with a Dirichlet boundary condition:

$$\int_{\Gamma} K(\mathbf{x}, \mathbf{y}) v(\mathbf{y}) ds(\mathbf{y}) = -u^{\text{inc}}(\mathbf{x}), \quad \forall \mathbf{x} \in \Gamma. \quad (2)$$

Other representations of the scattered field and integral equations exist, such as the combined double and single layer potential [12, §3.6] and the Brakhage-Werner integral equation [7]. These also depend on the boundary condition and on whether we consider the interior or exterior problem, but the principle of our method applies similarly as long as the integral operator asymptotically localizes.

The function $u^{\text{inc}}(\mathbf{x})$ in the right hand side of (2) represents an incoming wave, and the integral equation expresses that the total field $u^{\text{tot}}(\mathbf{x}) = u^s(\mathbf{x}) + u^{\text{inc}}(\mathbf{x})$ vanishes on the boundary of the domain. The discretization of $v(\mathbf{y})$ in (2) results in a boundary element method.

As we shall consider high wavenumbers, this density function $v(y)$ is highly oscillatory, thus requiring many degrees of freedom. In order to resolve all the oscillations of the problem, we follow a standard rule of thumb and choose a number of degrees of freedom N that is proportional to the wavenumber k in 2D, and proportional to k^2 in 3D. A collocation or Galerkin discretization of (2) results in a linear system

$$Ac = b, \quad (3)$$

where b represents the boundary condition. For both kinds of discretizations, the matrix A is dense, large and increasingly ill-conditioned. The relationship between the condition number and the wavenumber is complicated, and depends on properties of the obstacle and of the integral operators, see [9, §5] and [32]. It should

be noted that integral equations of the second kind are more popular in applications because they are typically better conditioned for increasing N . However, the condition number also deteriorates when increasing the wavenumber.

A fast matrix-vector product with A in the high-frequency regime is achieved by the high-frequency Fast Multipole Method [35, 22, 11, 41]. Alternatively, so-called hybrid numerical-asymptotic methods aim to significantly reduce the size of the linear system by incorporating information about the solution from asymptotic analysis (see the review [9] and references therein). In particular, phase-extraction methods use information about the phase g of the solution $v(\mathbf{y})$ in order to discretize only the remaining non-oscillatory part f in the following factorization:

$$v(\mathbf{y}) = f(\mathbf{y}, k)e^{ikg(\mathbf{y})}. \quad (4)$$

Phase extraction methods are simplest for convex obstacles, and require ray-tracing or similar techniques for more complicated domains or multiple scattering configurations [8, 19, 15, 2, 24, 10]. These methods typically lead to matrices that are small, sometimes with frequency-independent size, but that are still dense. The use of specialized quadrature methods for highly oscillatory integrals in [27] led to a matrix that is both small and sparse. It is possible to compute asymptotic information (more specifically, stationary points) by computing where derivatives or gradients of the phase vanish, see [17]. This is not unlike ray tracing, where ray reflection points correspond to zeros of a derivative.

The goal of this paper is to reduce the computational cost in a BEM context by exploring the use of asymptotic properties of the integral operator, but without performing phase extraction or ray tracing techniques. In doing so, we aim for great generality and flexibility in the shape of the domain. We use no a priori asymptotic information about the solution, nor do we numerically estimate the phase of $v(\mathbf{y})$. For complicated domains, the solution could be a sum of many functions of the form (4), or it could have more complicated asymptotic structure, and that limits applicability of purely asymptotic methods.

The methodology of our approach is as follows. The linear system (3) represents the action of an oscillatory integral operator (discretized in A) on an oscillatory function (discretized in c). Hence, the product of a row of A with the column vector c discretizes a highly oscillatory integral. Without ray-tracing methods and phase estimation, the phase of the oscillations of this integral is unknown. Yet, each highly oscillatory integral of this kind is essentially local in nature, and one can localize oscillatory integrals explicitly using smooth cut-off functions. Conceptually, in our setting this corresponds to modifying the Green's function, or more generally the kernel of the integral equation at hand. That is, we replace $K(\mathbf{x}, \mathbf{y})$ by

$$\tilde{K}(\mathbf{x}, \mathbf{y}) = K(\mathbf{x}, \mathbf{y})w(\mathbf{x}, \mathbf{y}),$$

where $w(\mathbf{x}, \mathbf{y})$ is a smooth window function, or a sum of smooth window functions. This results in a modified discretization matrix $\tilde{A}$. The integral operator with kernel $\tilde{K}$ agrees with the original operator when $k \rightarrow \infty$ as long as $w(\mathbf{x}, \mathbf{y})$ equals 1 in regions that contribute to the oscillatory integrals, and smoothly decays to zero outside these regions. Using basis functions with compact support in the discretization, the corresponding entries of $\tilde{A}$ are exactly equal to those of A . For regions where $w(\mathbf{x}, \mathbf{y})$ vanishes, the entries of $\tilde{A}$ are exactly zero. In the intermediate regions, the entries of $\tilde{A}$ are perturbations of those of A .

The idea and implementation of this approach are fairly straightforward, and can be based on an existing BEM implementation. We have performed our computations in Matlab for 2D problems. Our implementation is publicly available on GitHub [34] and a copy of the code can always be obtained from the authors. The main challenge is estimating suitable locations for the window functions. In principle, this still requires asymptotic analysis of the scattering problem, and this remains complicated for complicated geometries. However, the task of locating windows is significantly simpler than the task of determining (multiple) phases of the solution. Indeed, some contributions to the oscillatory integrals intrinsically correspond to reflections of rays in the scattering configuration.

Exploiting a simple geometric visibility criterion results in some compression for any obstacle, or indeed any combination of obstacles in a multiple scattering configuration, regardless of the geometric shape. But the compression rate can be much improved still. We obtain the best results by numerically measuring the contributing regions of the integrals at a moderate value of k , and extrapolating the location of the windows to a higher value of k . Their location is refined with every significant increase of k , resulting in higher compression rates for high-frequency problems in an overall fairly efficient frequency sweeping solver. This does not quite achieve the effectiveness of a true asymptotic scheme in problems where those apply, but our scheme is applicable to a much wider range of problems, and the results always improve with increasing wavenumber k .

Sparsifying A results in lower CPU times and memory requirements, especially in a frequency sweep. Another important observation is that the condition number of $\tilde{A}$ seems much lower than that of A itself. The lower condition number has a positive effect on the number of iterations required in an iterative solver for (3). Indeed, the potentially high condition numbers of BEM matrices have previously led to the study of several preconditioners, such as tensor-based preconditioners [29] and adaptive cross approximation LU-based preconditioners [5].

Furthermore, we show that the method is robust in several ways. It is robust with respect to the geometry of the obstacle: some amount of compression can still be achieved even for near-trapping domains, which are a worst case for ray-tracing methods. It is also robust with respect to the order of the BEM discretization. In spite of the fact that the asymptotic analysis of oscillatory integrals is mostly based on C^∞ integrands, we show that asymptotic compression can be achieved even when $v(\mathbf{y})$ is discretized using piecewise linear or cubic functions. In this case, the oscillatory integrals that we localize have integrands that are merely continuous. Exploiting this robustness, we can approximate the matrix entries of $\tilde{A}$ by a simple rescaling of the entries of A that are computed in an existing code.

On the other hand, unlike ray tracing methods, our method is not frequency independent. Also, similar to asymptotic methods, many parts of our calculations depend on the nature and direction of the incident wave. However, the sparsifying modification of the Green's function is compatible with the Fast Multipole Method, since all we do is modify the Green's function, and a more efficient implementation in 3D exploiting that connection is a possible future research project. Other acceleration schemes compute convolutions via the easy Fourier transform of the Green function: the sparse cardinal sine decomposition relies on a nonuniform Fast Fourier Transform and spherical Gaussian quadrature in [1], while the quasi-periodic Green's function is split into Gaussians in the spatial domain and an exponentially decaying Fourier series in [16].

The structure of the paper is as follows. We review asymptotic properties of oscillatory integrals in § 2 and one can find the asymptotic compression scheme in § 3. The localisation of the smooth functions using adaptive computation of correlations and adaptive recompression for increasing frequency is described in § 4. Finally, we end the paper with a series of numerical experiments on a range of scattering configurations in § 5 and with some concluding remarks in § 7.

2 Oscillatory integrals

2.1 Asymptotic analysis of oscillatory integrals

Oscillatory integrals of the form

$$\int_a^b f(x) e^{ikg(x)} dx \quad (5)$$

are a classical topic in asymptotic analysis [6, 39]. Note that in our setting, the total phase of the integrand $g(x)$ is the sum of the phase of the Green's function, which is known explicitly in general, and the phase (or multiple phases) of a density function v of the form (4). The main observation is that the integral can be seen as a sum of contributions, originating in a small number of critical points p . They include points where the integration domain is not analytic (the endpoints of integration a and b in the integral above), any singularities of the integrand, points where contour lines of the phase are tangent to the boundary of the integration domain and so-called stationary points. The latter are points p where the derivative of the phase function $g(x)$ vanishes, i.e. $g'(p) = 0$. Intuitively, the explanation is that near a stationary point the integrand is locally non-oscillatory, but elsewhere oscillations cancel each other out. Similarly, the cancellation effect is reduced or absent near the endpoints of the interval and near singularities.

Mathematically, the result of the analysis is an asymptotic expansion of the form

$$\int_a^b f(x) e^{ikg(x)} dx \sim \sum_{m=1}^M e^{ikg(p_m)} \sum_{j=1}^{\infty} A_{m,j}[f, g] k^{-r_{m,j}}, \quad k \rightarrow \infty.$$

Here, the points p_m are all the critical points of the integral, and the exponents $r_{m,j}$ form a strictly increasing sequence for each m . The coefficients $A_{m,j}$ depend linearly on f and non-linearly on g . Explicit expressions can be derived, showing that $A_{m,j}$ depends on the function values of f , g and their derivatives at the critical point p_m . The larger the index j , the higher the order of the derivative involved. Note that the expansion above assumes that f and g are infinitely differentiable. The exponents $r_{m,j}$ depend on the nature of the contributing point, and they are typically integer or rational numbers. We refer the reader to [13] for a detailed introduction to asymptotics of oscillatory integrals, examples of the computation of such expansions and numerical methods to evaluate such integrals efficiently for large frequency k .

An example of such a purely asymptotic method is given by computing asymptotic expansions of the solutions via asymptotic expansions of oscillatory integrals involving for example a complex exponential function for 3D problems. On the

Fig. 1: Illustration of the real part of an oscillatory integrand $e^{ik(y-1/2)^2}$ (dotted blue) with $k = 100$ and with a stationary point at $y = 1/2$. The stationary point contribution can be computed using a smooth window function (solid green) that covers it. The integral of the windowed function (dashed red) in the left panel asymptotically agrees with the integral of the oscillatory function. In contrast, the integral of the windowed function in the right panel is superalgebraically small in k .

other hand, the Green's function of the Helmholtz equation in 2D problems is given in terms of a Hankel function, such that expansions of similar integrals involving an oscillatory Hankel function can be used to deduce the asymptotic expansions of the solutions of the corresponding integral equations [23]. Yet, we are not concerned with the computation of asymptotic expansions in this paper. Instead, we solely rely on the fact that they exist and exploit the localization of oscillatory integrals by more robust means.

2.2 Explicit localization of oscillatory integrals

It is well-known, and common in applications, to single out the contributions of oscillatory integrals using smooth cut-off functions (see, for example, [8]). This is illustrated in Figure 1, which shows the real part of the oscillatory function $e^{ik(y-1/2)^2}$ with $k = 100$. This is a simple case of the general oscillatory integral (5), with a stationary point at $p = 1/2$. The oscillations rapidly cancel out for high k , so that we can neglect all but a small region around the stationary point. We can evaluate this contribution by multiplying the integrand with a C^∞ window centered around $1/2$ and then integrating the nonzero part.

Of course, the remaining integral is still highly oscillatory. One can shrink the support of the window function for increasing k , at a rate that depends on the nature of the contributing point. Loosely speaking, one wants to maintain a fixed number of oscillations within the support of the window, such that the integral can be evaluated at a fixed cost for increasing k . The size of the support of the window function around a stationary point is typically scaled as $\mathcal{O}(k^{-1/2+\epsilon})$, whereas for an endpoint contribution it can be smaller: $\mathcal{O}(k^{-1+\epsilon})$. The small and positive ϵ is needed to prove the small resulting error bounds. The size becomes larger for higher-order stationary points, where higher order derivatives of g vanish as well. The more derivatives of g vanish, the less oscillatory the integrand and the larger

the contribution. A stationary point p has order r if

$$g'(p) = g''(p) = \dots = g^{(r)}(p) = 0$$

but $g^{(r+1)}(p) \neq 0$. The associated contribution to the integral has size $\mathcal{O}(k^{-1/(r+1)})$ and the size of the window scales accordingly. For endpoints and singularities of the integrand, we have $r = 0$ and for non-degenerate stationary points, $r = 1$.

One typically uses C^∞ window functions in order to avoid introducing discontinuities in the derivatives of the integrand, as these would show up in the asymptotics of the modified integral. An extreme example of a non-smooth window is a block window on $[\lambda, \rho]$, which is discontinuous. This is equivalent to restricting an integral on $[a, b]$ to an integral on $[\lambda, \rho]$. Hence, one expects $\mathcal{O}(k^{-1})$ spurious contributions from the endpoints λ and ρ . It is important to note that the windows need not be symmetric.

3 Asymptotic compression of BEM matrices

We incorporate smooth window functions into the Green's function of an integral equation in order to isolate the contributions of the oscillatory integral operator. The most important consideration is where to place the windows, but first we introduce the specifics of the scheme.

3.1 Asymptotic compression method

We consider a domain whose boundary Γ is parametrised by κ mapping the parameter τ , which belongs to the domain $\mathcal{K}$, to Γ , such that the single layer potential becomes

$$\int_{\Gamma} K(\mathbf{x}, \mathbf{y}) v(\mathbf{y}) ds(\mathbf{y}) = \int_{\mathcal{K}} K(\mathbf{x}, \kappa(\tau)) \|\nabla \kappa(\tau)\| v(\tau) d\tau.$$

In a slight abuse of notation, we identify $v(\mathbf{y})$ with $v(\tau)$. In the following, we shall also use a Green's function in the parameter domain as defined by

$$K_{\kappa}(t, \tau) = K(\kappa(t), \kappa(\tau)) \|\nabla \kappa(\tau)\|.$$

For the discretization, we use a basis $\{\varphi_j\}_{j=1}^N$, defined in the parameter domain $\mathcal{K}$, of compactly supported functions φ_j with support S_j . We intend to use piecewise polynomials. Hence,

$$v(\tau) \approx v_N(\tau) = \sum_{j=1}^N c_j \varphi_j(\tau). \quad (6)$$

This setting can easily be modified to one where Γ itself is approximated by elements, as is common in BEM.

With collocation at the points $x_i = \kappa(t_i)$, $i = 1, \dots, N$, integral equation (2) leads to a discretization matrix A with entries

$$A_{i,j} = \int_{S_j} K_{\kappa}(t_i, \tau) \varphi_j(\tau) d\tau. \quad (7)$$

We continue with the collocation discretization – Galerkin is very similar. Recall that a row of A times a column vector c represents an oscillatory integration. Indeed, if $Ac = b$, then

$$b_i = \sum_{j=1}^N A_{i,j} c_j = \sum_{j=1}^N c_j \int_{S_j} K_\kappa(t_i, \tau) \varphi_j(\tau) d\tau = \int_\Gamma K_\kappa(t_i, \tau) \left(\sum_{j=1}^N c_j \varphi_j(\tau) \right) d\tau. \quad (8)$$

These integrals only have local contributions near certain critical points, which still depend on the collocation point t_i . Assuming we know these, we make this explicit using a smooth bivariate windowing function w ,

$$\tilde{K}_\kappa(t, \tau) = K_\kappa(t, \tau) w(t, \tau).$$

We introduce some more notation to specify w further. We divide the domain of w in three disjoint sets, $\mathcal{K} \times \mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2 \cup \mathcal{K}_3$, such that

$$w(t, \tau) = \begin{cases} 1, & \text{if } (t, \tau) \in \mathcal{K}_1, \\ w(t, \tau) \in [0, 1], & \text{if } (t, \tau) \in \mathcal{K}_2, \\ 0, & \text{if } (t, \tau) \in \mathcal{K}_3, \end{cases}$$

while $w(t, \tau)$ is C^∞ both in t and in τ .

Denote the compressed matrix corresponding to $\tilde{K}_\kappa$ by $\tilde{A}$, with entries

$$\tilde{A}_{i,j} = \int_{S_j} \tilde{K}_\kappa(t_i, \tau) \varphi_j(\tau) d\tau = \int_{S_j} K_\kappa(t_i, \tau) w(t_i, \tau) \varphi_j(\tau) d\tau. \quad (9)$$

It is clear that $\tilde{A}_{i,j} = A_{i,j}$ whenever $t_i \times S_j \in \mathcal{K}_1$, while $\tilde{A}_{i,j} = 0$ if $t_i \times S_j \in \mathcal{K}_3$. Recall that oscillatory integrals have only local contributions, so the set $\mathcal{K}_3$ on which the window vanishes will be large. This results in many zeros in $\tilde{A}$, hence *asymptotic compression*.

The matrix vector product $\tilde{A}c$ corresponds to the following oscillatory integral:

$$\sum_j \tilde{A}_{i,j} c_j = \int_{\mathcal{K}} K_\kappa(t_i, \tau) w(t_i, \tau) \sum_j c_j \varphi_j(\tau) d\tau = \int_{\mathcal{K}} K_\kappa(t_i, \tau) w(t_i, \tau) v_N(\tau) d\tau. \quad (10)$$

The two integrals for b_i , $\tilde{A}c$ above and Ac in (8), do not agree in general. However, the assumption is that when v_N is an oscillatory function, and since the Green's function is oscillatory as well, the two integrals do agree asymptotically for well chosen support sets $\mathcal{K}_1$ and $\mathcal{K}_3$. Their difference is an asymptotic error which is spectrally small in the wavenumber k even though we consider large perturbations of the matrix A , which may have a large condition number.

We proceed by replacing A by $\tilde{A}$ in (3). This leads to a new linear system of equations,

$$\tilde{A}\tilde{c} = b,$$

with a sparse matrix $\tilde{A}$. For collocation, we simply have $b_i = -u^{\text{inc}}(t_i)$. In this case, the weight function $w(t, \tau)$ need only be defined for the discrete points t_i .

3.2 Explicit expression of a smooth windowing function

Determining suitable locations for the window functions is a crucial problem, due to the fact that the phase of the solution is not known. We discuss this separately in the next section. Here, we describe the mathematical expression of the function.

The window function is a function of two variables, t and τ , where the integration is performed with respect to τ . Therefore, we consider the location of the window to be a function of t . In general, $w(t, \tau)$ is a sum of elementary window functions with locations depending on t :

$$w(t, \tau) = \sum_l \chi(\tau, \lambda_l(t), l_l(t), r_l(t), \rho_l(t)).$$

They are supported on $[\lambda, \rho]$ and equal to 1 on the smaller interval $[l, r]$, as can be seen in Figure 1. When two subsequent windows l and $l+1$ in w overlap or come too close, we join them into $\chi(\tau, \lambda_l, l_l, r_{l+1}, \rho_{l+1})$, see Figure 3 further on. The specific choice of χ can be important from a numerical point of view, but we expect very similar results as long as all derivatives have a reasonable bound. The extreme case would be a block window ($\lambda = l$ and $r = \rho$), introducing the aforementioned spurious endpoint contributions. We choose χ as the elementary window also used in [8], and shown in green in Figure 1:

$$\chi(\tau, \lambda, l, r, \rho) = \begin{cases} 0, & \tau \leq \lambda \\ \exp\left(\frac{2e^{\frac{l-\lambda}{\tau-l}-1}}{\frac{\tau-l}{\lambda-l}-1}\right), & \tau \in (\lambda, l) \\ 1, & \tau \in [l, r] \\ \exp\left(\frac{2e^{\frac{r-\rho}{\tau-r}-1}}{\frac{\tau-r}{\rho-r}-1}\right), & \tau \in (r, \rho) \\ 0, & \tau \geq \rho. \end{cases} \quad (11)$$

The width of the window will be determined automatically further on, even taking into account possible asymmetry. One expects the window to optimally scale as mentioned in the previous section on the asymptotics of oscillatory integrals, i.e., as $\mathcal{O}(k^{-1/(r+1)})$ around a stationary point of order r . Single reflections of rays bouncing off an obstacle are associated with stationary points of order 1 [23]. Degenerate stationary points do also arise frequently in scattering problems, for example at shadow boundaries where an incoming ray is tangential to Γ . Grazing rays, which give rise to creeping rays travelling along the boundary of the obstacle, are associated with stationary points of order 2 in asymptotic analysis [3, 33]. Note that these points do not depend on the wavenumber k , but they do depend on (the angle or structure of) the incoming wave. Hence, so do the contributing points and so does our computational scheme.

3.3 Robustness for low-order discretizations and computation of matrix entries

The asymptotic analysis of integrals of the form (5) usually assumes a smooth integrand. Typically, later terms in the asymptotic expansion depend on higher order derivatives of the integrand. Yet, in our setting we are faced with the discretized integral (10) that is not particularly smooth: the smoothness is limited at

least by the smoothness of the basis functions φ_j . We can only expect low-order asymptotics because of the piecewise polynomial basis functions used in this article. In this sense, it is somewhat unexpected that the results further on show that asymptotic compression is also effective for piecewise linear basis functions, as smoothness of the integrand is important for asymptotic behaviour as $k \rightarrow \infty$.

This apparent robustness also allows a significant simplification, that makes the implementation much cheaper. Rather than integrating the window function exactly, we will explore approximating the compressed matrix element of $\tilde{A}$ simply by rescaling the uncompressed element of A :

$$\begin{aligned}\tilde{A}_{i,j} &= \int_{S_j} K_\kappa(t_i, \tau) w(t_i, \tau) \varphi_j(\tau) d\tau \\ &\approx w(t_i, \tau_j) \int_{S_j} K_\kappa(t_i, \tau) \varphi_j(\tau) d\tau \\ &= w(t_i, \tau_j) A_{i,j}.\end{aligned}\tag{12}$$

Here, τ_j is a point within the support of the basis function φ_j , for example the collocation point corresponding to φ_j . This is a rather crude approximation of the exact weighted integral. However, it has the important advantage that the element of $\tilde{A}$ is a simple rescaling of the corresponding element of A . Thus, existing code for the computation of entries of A can maximally be reused.

Note that this approach is still decidedly different from simply compressing A by putting entries to zero. Even with our simplified approach, the elements of A are weighted, and thus decay to zero with some regularity. Simply discarding some matrix entries would correspond to using block-window functions, leading to the spurious endpoint contributions mentioned in § 2.2, which § 5.6 indicates still works but with much larger errors.

4 Adaptive location of the window functions

We consider three ways of finding a suitable support of the window functions. The first is based on the simple geometric notion of visibility. This approach is inferior to the following ones, but it does illustrate that determining supports is much easier than performing ray tracing. Second, we numerically estimate suitable windows by solving a complete scattering problem at a moderate value of the wavenumber. These windows are kept constant and reused for larger frequencies. Finally, we compute windows adaptively in a frequency sweep, refining the windows for increasing frequency. The latter scheme is most efficient, and we can start at a smaller value of the wavenumber for the initial windows.

4.1 Visibility criterion

Asymptotically, a wave field can be seen as a collection of rays along straight paths (in a homogeneous medium) [4]. Points of reflections of these rays physically correspond to stationary points in the integral equation formulation. For scatterers with a moderately complicated geometry, it becomes unwieldy to trace all reflecting rays through the scene. However, in the integral equation formulation, the

oscillatory integral in equation (2) corresponding to a point $\mathbf{x} \in \Gamma$ on the boundary always asymptotically corresponds to a sum of contributions from a limited number of points.

If a point B on the boundary is not directly visible from a point A on the illuminated part of the boundary, then B can not be a stationary point of the oscillatory integral at A . Indeed, there can be no *direct* ray from B to A (or vice-versa). There can only be indirect connections via intermediate reflection points C elsewhere on the boundary, or via creeping rays along the surface of the obstacle. In the former case, B is a stationary point for the integral at C which is subsequently linked with A , but there is no direct connection from B to A via a stationary point which has to be taken into account. This means that for any given point A on the illuminated part of the boundary, we can asymptotically discard the density in all non-visible parts of the boundary from that point A , unless that part is close to A due to the contribution from the singularity.

More specifically, we choose a fixed parameter $T = l - \lambda = \rho - r$ to denote the length of the decaying part of the window. The window $w(t_i, \tau) \equiv 1$ when $\kappa(\tau)$ can see the point $\kappa(t_i)$ or when $t_i = \tau$ and else equals $\chi(\tau, l - T, l, l + T)$ where l is the closest parameter to τ such that $\kappa(l)$ can see $\kappa(t_i)$. Examples for two specific geometries are given in § 5.5. Note that this simple criterion does not require explicit information on the phase of the density $v(\mathbf{y})$. However, this reasoning does not immediately hold for points on the shadow side of the boundary, for which the situation is more complicated (see §4.4 further on).

In a collocation solver, we have to decide for each collocation point t_i where to place the window function $w(t_i, \tau)$. For objects with self-reflections, one could trace a ray from the source through the scattering scene to the collocation point, much like ray tracing in computer graphics. Each point of reflection is expected to be a stationary point for the next point of reflection, or ultimately for the collocation point. Hence, each of these points needs a window around it, and any two subsequent points have to be able to ‘see’ each other for a reflection to be possible. We do not consider any ray tracing in our visibility criterion, but simply exploit the latter visibility condition. Computing which points can see one another gives the regions where critical points can lie. The visibility problem is also common in computer graphics on the GPU, and the computation can be sped up using Kd-trees or Bounding Volume Hierarchies [38].

4.2 Adaptive asymptotic compression

Points that contribute to a general oscillatory integral can be detected automatically using a sliding window $\zeta(x - z)$ centered around a point z . If the window contains no contributing points, asymptotic theory predicts the windowed integral

$$F(z) = \int_a^b \zeta(x - z) f(x) e^{ikg(x)} dx \quad (13)$$

to be superalgebraically small in the frequency parameter k (for C^∞ integrands). As explained in § 2, the integrand would look like the red dashed line at the right of Figure 1 where the positive and negative parts cancel. Conversely, the function $F(z)$ can be expected to be large if the window around z does contain a

Fig. 2: Schematic depiction of the computation of the correlations $R_{n,q}$ given by (15). The matrix and the solution vector c are computed for a wavenumber k_1 . Each row n of A is multiplied by a sliding window centered around σ_q , and the inner product with the solution vector c results in a localised correlation value $R_{n,q}$. Recall that row n corresponds to collocation point t_n . For more implementation details, we refer the reader to the code available online [34].

contributing point. This corresponds to shifting the solid green window function to z near $1/2$ in Figure 1, where the integrand is locally less oscillatory.

In the context of integral equation methods, one can compute correlations between any two points on the boundary as integrals of the unmodified integrand times a moving window function. This can only be done a posteriori, since the integrand in (2) involves the density function which is unknown a priori.

Having computed v_N for a (moderate) value of the wavenumber k_1 , we can compute correlations comparable to (13) above. We do so for each collocation point $t_n^{(k_1)}$, $m = 1, \dots, N$, and use windows around center points $z = \sigma_q$, $q = 1, \dots, Q$ where we choose $Q = 1.5N$. We use more centers than collocation points by a factor of 1.5 in order to determine more accurately where the important contributions are located. The sliding window function we use is

$$\zeta(\tau - \sigma_q) = \chi(\tau - \sigma_q, -T, 0, 0, T).$$

Here, χ is the elementary window function introduced in § 3.2 with center σ_q , support width $2T$, and τ corresponds to the 1D parametrisation of the obstacle. We compute a correlation matrix R , with elements given by:

$$R_{n,q} = \int_{\mathcal{K}} \zeta(\tau - \sigma_q) K_{\mathcal{K}}(t_n, \tau) v_N(\tau) d\tau \quad (14)$$

$$\begin{aligned} &= \sum_{j=1}^N c_j \int_{\mathcal{K}} \zeta(\tau - \sigma_q) K_{\mathcal{K}}(t_n, \tau) \varphi_j(\tau) d\tau \\ &\approx \sum_{l=1}^N \zeta(t_l - \sigma_q) A_{n,l} c_l. \end{aligned} \quad (15)$$

Recall from (6) that $v_N(\tau)$ is our approximation to the density function. In the latter step, we have again approximated a windowed integral by a weighted sum of matrix elements of the dense discretization matrix A , which makes the computation of all correlations a cheap operation. No additional numerical integration is required beyond what was already computed for the regular BEM matrix A . This

computation is visualised at the top of Figure 2: as σ_q shifts, other parts of the matrix-vector product Ac will be windowed.

For the compression phase at a higher wavenumber $k_2 > k_1$, we identify the collocation point $t_i^{(k_2)}$ with the closest collocation point $t_n^{(k_1)}$ at the lower wavenumber k_1 at which the correlations were computed and we identify the interaction point $t_j^{(k_2)}$ with the closest window center $\sigma_q^{(k_1)}$. We retain all unmodified matrix elements corresponding to correlations higher than a fixed percentage ξ times the maximal correlation on each row.

It is possible to determine suitable threshold percentages ξ adaptive to the geometry by monitoring the relative compression error $\|\tilde{A}c - b\|/\|b\|$, or even only for a few representative elements of b . One can also compare correlations to ξ times the maximal $|R_{n,q}^{(k_1)}|$ for all n and q , but we found a row-wise threshold to be more robust with respect to changes in the maximum between rows. For example, the correlations are much smaller in the shadow region because the windows which compute them are small. Although this contribution is not exponentially small as the integrand is locally nonoscillatory, such a global threshold would compare them with a higher correlation from the Green singularity in the illuminated region.

The goal is to maintain a compression error of the same order as the discretization error. Without any assumption on the density $v_N(\tau)$, one might expect large errors from throwing away a part of the matrix. However, if the incident wave changes to give significant contributions elsewhere, the correlations $R_{n,q}$ will also increase there as we compute the density at a moderate frequency. We provide numerical results further on for various types of incident waves.

The procedure is illustrated for a specific example in Figure 3, where $n = 100$ and $i = 200$. More specifically, to compute the i -th row of $\tilde{A}_{i,j}^{(k_2)}$ in general, we first determine the collocation point used to compute the correlations at wavenumber k_1 that is closest to the one corresponding to this row

$$n = \arg \min_n |t_n^{(k_1)} - t_i^{(k_2)}|.$$

Next, we define a set of intervals I consisting of a region around the collocation point and points which are between σ_q 's where the correlation $R_{n,q}$ lies above the threshold percentage $\xi = 4\%$ times the maximal correlation on the row n (which is 0.126 in Figure 3, indeed at the τ_j where the matrix-vector product $\tilde{A}_{200,j}\tilde{c}_j$ is least oscillatory):

$$I = \left\{ t_i^{(k_2)} - T, t_i^{(k_2)} + T \right\} \cup \left\{ [\sigma_p, \sigma_{p+1-Q\lfloor l/Q \rfloor}] \mid \min(|R_{n,p}|, |R_{n,p+1-Q\lfloor l/Q \rfloor}|) \geq \xi \max_l |R_{n,l}| \right\}. \quad (16)$$

We always include the Green's function singularity t_i , even when the computed correlations are small there, since otherwise $\tilde{A}$ may not always have full rank. Windows are merged if they would come too close to each other after adding a decaying part T :

$$C = I \cup \left\{ [a, b] \mid a \in I \quad \& \quad b \in I \quad \& \quad |b - a| \leq 2T + \frac{5}{N_2} \right\}. \quad (17)$$

Fig. 3: Illustration of the algorithm for the computation of the 200-th row of the compressed matrix $\tilde{A}_{i,j}^{(k_2)}$ given by (18) at wavenumber $k_2 = 2^8$ for the near-convex obstacle shown in Figure 7. This row corresponds to an oscillatory integral over the domain and the figure shows the real part of the integrand. The correlations were computed at $k_1 = 2^7$ and are shown in blue. The sets I and C are indicated: I is where the correlations are higher than a threshold (here 5.03×10^{-3}), and C is like I but with nearby intervals merged in order to define a smooth window function. The final window is shown in red. No compression is applied within C . The matrix element is set to zero when further than $T = 0.02$ away from this set. Finally, matrix elements in between are weighted.

The additional $5/N_2$ ensures that gaps between windows correspond to at least five zero matrix entries in $\tilde{A} \in \mathbb{C}^{N_2 \times N_2}$. Of course, this constant is somewhat arbitrary.

Finally, we compute the entries of our compressed matrix (9) as the approximation (12): a window factor times the original BEM matrix entry

$$\tilde{A}_{i,j}^{(k_2)} = \zeta(|t_j^{(k_2)} - C|) A_{i,j}^{(k_2)} = \begin{cases} A_{i,j}^{(k_2)} & \text{if } t_j^{(k_2)} \in C \\ \exp\left(\frac{2e^{\frac{T}{|t_j^{(k_2)} - C| - T}}}{\frac{T - |t_j^{(k_2)} - C|}{T} - 1}\right) A_{i,j}^{(k_2)} & \text{if } |t_j^{(k_2)} - C| < T \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

Here, $|t_j^{(k_2)} - C|$ is the distance of $t_j^{(k_2)}$ to the set C , which is the distance to the closest point belonging within C . One can see in Figure 3 that $\zeta(|t_j^{(k_2)} - C|)$ equals 1 for τ_j on the set C , is a factor $\in [0, 1]$ in the decaying part of the window with width T and zero otherwise.

Fig. 4: Similar illustration as in Figure 3, but for a more complicated trapping domain (as shown in Figure 7). The set I consists of two separate parts, one of which covers the Green singularity. For the value $k = 2^7$ of the wavenumber used here, the intervals are joined in C . Not much compression is possible in this complicated case. However, the method is robust with respect to the complicated nature, and compression will still improve with increasing wavenumber.

In summary, we have devised a window that covers the regions with high correlations and the Green singularity, and smoothly decays to zero away from them, while taking care to merge windows that would (almost) overlap.

The example in Figure 3 is illustrative, but represents only a simple case. We show robustness of the correlation procedure in Figure 4. Here, the domain is a trapping domain (shown in Figure 7) with complicated high-frequency behaviour that is hard to analyze analytically. The set I consists of two disjoint intervals, which are merged in the set C . The effect of the complicated scattering behaviour is quantitative, indeed the compression is not very high in this case. Yet, by construction the method still applies with only small error and one still expects better compression at higher frequencies.

Note that the procedure leads to a window for each collocation point at the wavenumber k_1 where the correlations are computed. At larger frequencies, the discretization may be refined and there may be more collocation points. For each of those points, we simply use the window of the closest collocation point at k_1 . We can reuse the correlation matrix $R_{n,q}$ for larger wavenumbers k because the list of stationary points or contributing regions that $R_{n,q}$ provides has a structure which is essentially independent of k . That is because stationary points are essentially a geometrical feature, and hence their location is independent of the frequency.

The approach using correlations has an additional advantage. When determining the asymptotic expansion of the oscillatory integral as in [33], one computes the points where the gradient of the phase of the integrand is exactly zero and determines the contributions coming from them. This is implicitly done by the reflecting points in ray tracing. The correlations also find regions where the gradient of the phase is small (in absolute value) but does not cross zero. These regions

have no contribution as $k \rightarrow \infty$, but a finite k might have to be very large for this contribution to become negligible. Thus, including these regions is also advantageous, which the correlation technique does. This has also been remarked in [18, §5.1.1]. We also retain contributions from creeping rays if they are still significant at the frequency at which the correlations are computed.

We illustrate the method in § 5 with numerical results, but we already point out two issues. The first issue relates to the value of the ‘moderate’ wavenumber k_1 . The adaptive asymptotic compression method is adaptive with respect to the geometry and computes the correlation matrix R only at this lowest wavenumber. These correlations provide valuable high-frequency information only as long as there is ‘enough’ high-frequency behaviour present at that value of k . Solving the full boundary element problem at this value of k might thus be expensive as well. Indeed, this is the problem we are trying to speed up in the first place. The second issue is that thus far we are using fixed windows for each value of k , computed at the small value of k_1 . In doing so, we are not exploiting the fact that windows could be smaller at higher frequencies as described in § 3.2.

Both of these issues are addressed by adaptively recompressing the matrix as the frequency increases. This is especially appealing when one has to solve the same problem for a range of frequencies. As will be explained in § 4.3, we can take the smallest frequency to be smaller, and we can efficiently obtain smaller windows for larger values of the wavenumber k .

4.3 Adaptive recompression at higher frequencies

In the adaptive asymptotic compression technique thus far, we have been using fixed windows computed at a moderate wavenumber. This leads to a fixed compression ratio, independent of the frequency and a decreasing asymptotic error. However, it should be possible to use smaller windows at larger frequencies for a constant error, hence improving compression with increasing frequency. In the absence of prior asymptotic information, one way to achieve higher compression at higher frequencies in general is to recompute the correlations $R_{n,q}$ at each wavenumber.

Of course, computing the full correlation matrix is prohibitively expensive, since it requires the dense matrix A . However, if we have already computed a compressed matrix $\tilde{A}^{(k_j)}$ at a smaller wavenumber k_j , then we can omit a large part of the computation. Indeed, if the correlation between two points on the boundary is small at a certain frequency, indicating that they are not asymptotically linked, it is safe to assume that their correlation will only be smaller at higher frequencies. Thus, we have to compute correlations only for the non-zero entries of $\tilde{A}^{(k_j)}$. Assuming some of these will be small, because the integrals are more oscillatory, this leads to a sparser matrix $\tilde{A}^{(k_{j+1})}$ at a higher wavenumber k_{j+1} . We can repeat this reasoning, where at each higher frequency we benefit from the improved sparsity of the previous stage.

To be precise, we introduce some notation. Say we have a range of strictly increasing wavenumbers k_j , $j = 1, \dots, J$, and we use N_j degrees of freedom at each wavenumber k_j . In a full correlation technique, we would use $1.5N_j$ centers $\sigma_q^{(k_j)}$, $q = 1, \dots, Q_j$, for each collocation point $t_n^{k_j}$, $n = 1, \dots, N_j$. The corresponding correlation matrix $R_{n,q}^{(k_j)}$ determines window functions $w_n^{(k_j)}$ to compute the

compressed matrix $\tilde{A}^{(k_j)}$, and the solution of the discretized integral equation is $\tilde{c}^{(k_j)} = \tilde{A}^{(k_j)} \setminus b^{(k_j)}$.

In the adaptive asymptotic recompression scheme, we only compute the full correlation matrix $R^{(k_1)}$ at the smallest wavenumber k_1 . We reuse $R^{(k_1)}$ for the next few wavenumbers k_l up to and including $k_j = 2k_1$. For this larger wavenumber k_j , after having computed the solution vector $\tilde{c}^{(k_j)}$ and the associated density function $v^{(N_j)}$, we only compute some of the elements of $R^{(k_j)}$. We can restrict our attention to the support of the previously computed window functions $w_{m'}^{(j-1)}$:

$$R_{n,q}^{(k_j)} = \begin{cases} \int_{\mathcal{K}} \zeta(\tau - \sigma_q^{(k_j)}) K_{\mathcal{K}}(t_n^{(k_j)}, \tau) v^{(N_j)}(\tau) d\tau, & \text{if } \sigma_q^{(k_j)} \in \text{supp } w_{m'}^{(j-1)}, \\ 0, & \text{otherwise.} \end{cases}$$

where m' is such that $t_{m'}^{(k_{j-1})}$ is the closest collocation point to $t_n^{(k_j)}$. Subsequently, the element $R_{n,q}^{(k_j)}$ can be approximated by a weighted sum of elements of $\tilde{A}^{(k_j)}$ instead of $A^{(k_j)}$ as before in (15). Moreover, we only need to sum elements where $\tilde{A}$ is nonzero, which is given by the support of $w_{m'}^{(j-1)}$, and where ζ is nonzero. This is captured in the set D in

$$D = \left\{ j \mid T < |t_j^{(k_l)} - \sigma_q| \right\},$$

$$R_{n,q}^{(k_j)} = \begin{cases} \sum_{l \in D} \zeta(t_l - \sigma_q^{(k_j)}) \tilde{A}_{m,l}^{(k_j)} \tilde{c}_l^{(k_j)}, & \text{if } \sigma_q^{(k_j)} \in \text{supp } w_{m'}^{(j-1)}, \\ 0, & \text{otherwise.} \end{cases} \quad (19)$$

Finally, as before, new windows are computed based on thresholding the correlations using a percentage ξ . These are used for the compressed matrices $\tilde{A}^{(k_l)}$ at the next wavenumbers $k_l > k_j$ until the wavenumber has doubled again, so $k_l = 2k_j$ and so on. This loop can be visualised by inserting $R_{n,q}$ of Figure 2 into Figure 3 and by inserting $\tilde{A}^{(k_l)}$ of the latter as $A^{(k_1)}$ in the former for $k_l = 2k_j$.

The sweeping strategy is illustrated in Figure 5. At the lowest wavenumber k_1 , one computes the standard collocation BEM matrix $A_{i,j}^{(k_1)}$ (7), the right hand side $b_i = u^{\text{inc}}(t_i)$ and the solution $c = A^{-1}b$. The matrix and solution are then used to compute the correlations $R_{n,q}^{(k_1)}$ through (15) for some choice of the window centers σ_q : for example $\sigma_n = t_n$, the collocation points at k_1 . The windows are computed using the same algorithm as described in § 4.2.

The benefits of the adaptive recompression are two-fold. First, we can restrict the (possibly expensive) computation of correlations to subsets of the domains, namely to the support of the window functions identified at lower frequencies. Second, the support of the windows decreases with increasing frequency, leading to improved compression. The exponential spacing of the wavenumbers at which correlations are recomputed ensures that the overhead of adaptive recompression remains modest at high frequencies, since the numerical solutions themselves are typically computed at linearly spaced frequencies.

4.4 Collocation points in the shadow region(s) and Fermat's principle

In some cases, the location of stationary points has been extensively studied, most notably for single convex smooth scatterers in 2D and 3D [8, 14, 18, 17]. In this

As an example of the references given above, the computation can be done exactly for an incident plane wave in the x direction on a circle of radius 1 parametrised as $\kappa(\tau) = [\cos(2\pi\tau), \sin(2\pi\tau)]$ for $\tau \in [0, 1]$. Each collocation point has at least one critical point given by $\tau = t$, which corresponds to the singularity of the Green's function. For points in the illuminated region and the two points in the shadow boundary, this is the only one. There is exactly one other critical point, a stationary point, when the collocation point lies in the shadow region, and it arises as follows. One can assume that the phase of $v(\tau)$ on the illuminated side is equal to the phase of the incident wave [27, Rem. 4.1]. The phase of $v(\tau)$

Algorithm 1 Pseudocode for the adaptive recompression algorithm

Require: Threshold percentage ξ , window decay length T , wavenumbers k_l

$c^{(k_1)} = (A^{(k_1)})^{-1} b^{(k_1)}$ $\triangleright$ Standard BEM with collocation points $t_n^{(k_1)}$

$\zeta(\tau) = \begin{cases} \exp\left(\frac{2e^{\frac{\tau-T}{T}}}{\frac{\tau-T}{T}-1}\right), & \tau \in [0, T) \\ 0, & \tau \geq T. \end{cases}$ $\triangleright$ Window function

$\sigma_q = (q-1)/(1.5N_1)$ $\triangleright$ For parametrisation $\tau \in [0, 1]$, or use $\sigma_q = t_q^{(k_1)}$

$R_{n,q}^{(k_1)} = \sum_{j=1}^N \zeta(|t_j^{(k_1)} - \sigma_q|) A_{n,j}^{(k_1)} c_j^{(k_1)}$ $\triangleright$ Initial correlations

$k_{\text{corr}} = k_1$

for $l = 2, \dots, L$ **do**

for $i = 1, \dots, N_l$ **do**

$n = \arg \min_n |t_n^{(k_{\text{corr}})} - t_i^{(k_l)}|$

 (16): $I = \left\{ t_i^{(k_l)} - T, t_i^{(k_l)} + T \right\} \cup \left\{ [\sigma_p, \sigma_{p+1}] \mid \min(|R_{n,p}|, |R_{n,p+1}|) \geq \xi \max_m |R_{n,m}| \right\}$

 (17): $C = I \cup \left\{ [a, b] \mid a \in I \ \& \ b \in I \ \& \ |b-a| \leq 2T + \frac{5}{N_l} \right\}$

for all $t_j^{(k_l)} \in C$ **do**

$\tilde{A}_{i,j}^{(k_l)} = A_{i,j}^{(k_l)}$ $\triangleright$ Standard BEM

end for

for all $|t_j^{(k_l)} - C| < T$ **do**

$\tilde{A}_{i,j}^{(k_l)} = \zeta(|t_j^{(k_l)} - C|) A_{i,j}^{(k_l)}$ $\triangleright$ Weighted standard BEM

end for

$\tilde{c}^{(k_l)} = (\tilde{A}^{(k_l)})^{-1} b^{(k_l)}$ $\triangleright$ Backslash, GMRES, ...

if $k_l \geq 2k_{\text{corr}}$ **then** $\triangleright$ Recompute correlations

for all $n=1, \dots, N_l$ **do**

 (16): $I = \dots$, (17): $C = \dots$

$\sigma_q = (q-1)/(1.5N_l)$

for all $\sigma_q \in C$ **do** $\triangleright$ Points where $\tilde{A}_{n,j}^{(k_l)}$ is nonzero

$D = \left\{ j \mid T < |t_j^{(k_l)} - \sigma_q| \right\}$ $\triangleright$ Points in the moving window

$R_{n,q}^{(k_l)} = \sum_{j \in D} \zeta(|t_j^{(k_l)} - \sigma_q|) \tilde{A}_{n,j}^{(k_l)} \tilde{c}_j^{(k_l)}$ $\triangleright$ Smaller window at k_{l+1}

end for

end for

$k_{\text{corr}} = k_l$

end if

end for

is thus equal to $\cos(2\pi\tau)$ for $\mathbf{y} = \kappa(\tau)$ in the illuminated region, and the phase of $K_\kappa(t, \tau)$ is the Euclidean distance to the collocation point. Hence, there exists a slowly varying f such that the integrand

$$K_\kappa(t, \tau)v(\tau) = f(t, \tau, k) \exp[ikg(t, \tau)]$$

with

$$g = \cos 2\pi\tau + \sqrt{(\cos 2\pi t - \cos 2\pi\tau)^2 + (\sin 2\pi t - \sin 2\pi\tau)^2}$$

and therefore

$$\frac{\partial g}{\partial \tau} = \frac{-\sin 2\pi\tau}{2\pi} + \frac{\sin 2\pi\tau \cos 2\pi t - \cos 2\pi\tau \sin 2\pi t}{2\pi \sqrt{(\cos 2\pi t - \cos 2\pi\tau)^2 + (\sin 2\pi t - \sin 2\pi\tau)^2}}.$$

Fig. 6: A collocation point in the shadow region ($\mathbf{x}$ in the figure) always has a stationary point $\mathbf{p}$ on the boundary in the illuminated region, that lies on the line connecting $\mathbf{x}$ to the source $\mathbf{s}$. By Fermat's principle this point minimizes the sum $\|\mathbf{s} - \mathbf{y}\| + \|\mathbf{y} - \mathbf{x}\|$ for $\mathbf{y}$ on the boundary.

The derivative of the phase is zero at $\tau = 1/2 - t$, hence $1/2 - t$ is a stationary point $\mathbf{p}$ of the integral (1). It lies in the illuminated region on the horizontal line through the collocation point.

While the computation of the circle is specific, the conclusion appears to be generic: it holds for more general objects that for any collocation point in the shadow region, there is a stationary point in the illuminated region on the line connecting the collocation point with the source. In Figure 6, we illustrate the existence of a stationary point on this line according to Fermat's principle, which states that rays travel along paths that minimize the total travel time. Consider a point $\mathbf{x}$ in the shadow region, and a line from that point to the point source (or to infinity in case of a plane wave). Among all possible points on the boundary $\mathbf{y}$, the distance from the point source to $\mathbf{y}$ plus the distance from $\mathbf{y}$ to $\mathbf{x}$ is trivially minimized precisely by the point $\mathbf{p}$ shown, where the line crosses the illuminated region.

Since we have assumed a non-penetrable scatterer in this paper, no ray physically passes through the interior of the obstacle. The stationary point arises because the Green's function only takes into account the distance between two points, and not the relative location of those points to the obstacle. Intuitively, one can say that the density $v(\mathbf{y})$ is asymptotically very small in the shadow region, hence the integral in (2) needs a contribution from the illuminated region to satisfy the boundary condition. Our adaptive strategy for general geometries confirms this and automatically recovers the stationary point. In our (re)compression scheme, we do include the Green singularity from the shadow region into our windows since although the contribution is asymptotically very small, $\tilde{A}$ may otherwise not always have full rank.

Fig. 7: These are the obstacles used in the numerical experiments, with their respective names, threshold percentages ξ , relative number of points and incident waves. For the superposition of incident plane waves in different directions, the drawn sizes of the waves indicate their relative magnitude. Some points of relevance to the discussion are explicitly indicated in the figures and referenced in the text.

5 Numerical results

Throughout this section, we perform experiments in 2D using the incident waves on Dirichlet obstacles shown and named in Figure 7, where we also indicate the number of points per wavelength and the incident wave. These include a highly non-convex obstacle, multiple scattering configurations and a superposition of three incident plane waves with $2\pi/3$ angle differences. A smooth superposition of plane waves would be provided by a Herglotz-type incidence, and beam-type incidence would result from a density on a certain contour, but plane waves and point sources were selected. We will illustrate different aspects of our methods with these different scattering configurations.

We use piecewise linear basis functions in the reported experiments, since we explicitly aim to investigate asymptotic compression in low-order discretizations. The use of higher order basis functions leads to similar results in § 5.6. The basis functions are approximately evenly spread out over the arclength of the obstacles, as adding more collocation points in the shadow region would unfairly increase our compression rates and performance. More specifically, that would increase the number of columns in $\tilde{A}$ where we only include the Green singularity as the density is too low in the shadow to give large contributions to the integral, and the singularity only needs one small window.

Throughout this section, the parameter T denoting the decaying part of the window is chosen to be 0.02 for the computation of $R_{n,q}$ as well as $\tilde{A}$. The only exception is that we use $T = 0.02k_1/k$ for the Green singularity in $\tilde{A}$, where k_1 is the lowest wavenumber of the simulation. These values could be further optimised as mentioned in § 7.

The scheme is implemented in Matlab, where we compute the integrals in (7) for computing the full matrix and (1) for evaluation of the scattered field using the Sweldens quadrature scheme [37]. The BEM implementation that we started from is described in more detail in [28]. The integrals for $\tilde{A}_{m,n}$ and $R_{n,q}$ are replaced by their approximations (12) and (15).

We use two error criteria in the examples. The first is the relative error in the solution vector $\tilde{c} = \tilde{A} \setminus b$ using the compressed matrix $\tilde{A}$ versus the exact numerical solution $c = A \setminus b$. The second criterion is the residual error of the integral equation, which for equation (2) corresponds to the accuracy with which the field satisfies the boundary condition. We have measured this error at 100 points randomly distributed over the boundary of the scatterer and computed the relative 1-norm.

5.1 Sparsity pattern and compression rate

We start by illustrating the sparsity of the discretization matrix $\tilde{A}$ and the compression rates for increasing k for the obstacles shown in Figure 7. We have used the adaptive recompression scheme outlined in §4.3, starting from small values of the wavenumber and using the corresponding threshold and relative numbers of collocation points as indicated in Figure 7. Thresholds were chosen to maintain roughly the same accuracy for increasing k , approximately equal to the discretization error. Precise results are given in § 5.3 further on.

The sparsity pattern of $\tilde{A}$ is shown for the almost convex obstacle in the left panel of Figure 8. The diagonal clearly shows the contribution of the singularity of the Green's function on each row. There are two sidelobes, corresponding to collocation points in the shadow region having a stationary point in the illuminated region (see §4.4); however, these were determined automatically. The structure of this matrix resembles very much the structure of compressed matrices that appeared in literature for a strictly circular obstacle, for example in the asymptotic methods of [27] and [21]. One might expect that the whole region between the Green singularity on the diagonal and the stationary point in the sidelobe needs to be included, as creeping rays travel along this path. However, those rays are exponentially small in k and we obtain low errors when not including this region throughout this paper.

Fig. 8: Left: Structure of the compressed matrix $\tilde{A}$ using adaptive recompression on the almost convex obstacle at $k = 2^{11}$. Right: fraction of nonzero elements in $\tilde{A}$ using adaptive recompression for increasing wavenumbers. The compression rate improves with increasing k in all but one case and although still fairly modest, it is expected to keep improving for even higher wavenumbers, even for general obstacles with unknown high-frequency behaviour.

Yet, there are two important main differences with respect to those asymptotic methods. On the one hand, the matrix in this paper is larger, since its dimensions scale linearly with increasing k . The methods of [27] and [21] used phase extraction to reduce this complexity, and as a result are limited to convex obstacles. On the other hand, the method of this paper does not require any a priori explicit knowledge of the nature of the solution. Therein lies a major gain. Though phase extraction is perhaps straightforward for the illuminated part of a convex obstacle, it is much less so on the shadow side. In the transitional region in between, the correct asymptotic behaviour becomes rather challenging, even for convex obstacles [3, 33]. Our approach is also different from [26], where compression is obtained by discarding small elements in a transformed matrix via wavelet packets.

The right panel of Figure 8 shows the improvement in compression rate as a function of the increasing wavenumber for all obstacles and it paints an encouraging picture. It is evident that the compression is complexity reducing for all but one configuration, including even the configurations with complicated incident waves, a trapping domain and with multiple obstacles – though at the shown range of the wavenumber, the percentage of non-zero elements remains fairly high.

For the trapping domain and circle multiple scattering configuration, there seems to be no compression. This is due to the combination of a complicated multiple scattering configuration and a superposition of three incident plane waves with $2\pi/3$ angle differences. The nonconvex polygon does not suffer as much from this, but the ellipse does at low wavenumbers. As there are 99.83% nonzeros for the trapping domain and circle at $k = 512$ and 99.64% at $k = 1024$, visible compression is only expected to appear at much higher wavenumbers for this combination of a complicated geometry with a ‘difficult’ boundary condition.

However, in spite of these issues, the example does illustrate that asymptotic information can be exploited even in the presence of highly non-trivial scattering patterns. Though compression may not result in large computational savings for challenging cases, later examples will show that the gain in condition number typically improves rapidly with increasing frequency and this may prove to be more important.

As expected, the results are best for convex and almost convex obstacles. These are the curves with decidedly better compression rates, corresponding to ‘easier’ scattering problems: a circle, the almost convex problem and the non-convex polygon, and perhaps the ellipse and the two circles. For cases where asymptotics apply, asymptotic methods are likely superior, but our compression rate does seem to improve like $\mathcal{O}(k^{-1/2})$. This corresponds closely to the expected scaling of $\mathcal{O}(k^{-1/2})$ width of a window around a stationary point, as explained in § 2. Indeed, because we refine the discretization with increasing wavenumber, there are $\mathcal{O}(k)$ rows in the matrix. Each row has $\mathcal{O}(k)$ elements, but each window covers only $\mathcal{O}(k^{1/2})$ of them due to their decreasing support width. That leads to $\mathcal{O}(k^{3/2})$ nonzero elements in $\tilde{A}$ in total, compared to $\mathcal{O}(k^2)$ elements overall. Points in the transition region need larger $\mathcal{O}(k^{-1/3})$ windows, but this region itself shrinks like $\mathcal{O}(k^{-1/3})$ [27]. Although this reasoning is a crude computation, it supports a compression rate of $\mathcal{O}(k^{-1/2})$. This scaling seems to be matched nicely in the adaptive scheme. It is less clear what the compression rate in the other configurations is, as a function of the wavenumber. Yet, the downward trend is evident in all cases. One might speculate that a sufficient condition for an $\mathcal{O}(k^{-1/2})$ compression rate could be that the obstacle is star-shaped, which means that there is a center point from which a straight line can reach all points on the boundary Γ without intersecting Γ in between.

5.2 The correlation matrix and adaptive localization of stationary points

Next, we illustrate the correlation matrix (introduced in § 4.2) and its relation to the structure of a scattering problem. As described earlier in the paper, first we solve the full scattering problem for some value of the wavenumber. The solution is used to estimate the relevance of parts of the boundary for each collocation point using a moving window in the oscillatory integral. In these experiments, we have used the value $T = 0.02$ for the parameter that determines the width of the moving window.

An informative example of a correlation matrix is shown in the left panel of Figure 9, for the multiple scattering configuration with three ellipses (the ellipses are shown and numbered 1 – 2 – 3 in Figure 7). Here, rows correspond to collocation points and columns to the centers of the moving windows. The correlation matrix has a 3×3 blocked structure, with each block corresponding to the coupling between two of the three ellipses. Asymptotic analysis of multiple scattering configurations shows that ultimately, rays bounce around in periodic orbits that minimize the distance between different obstacles [15, 2]. These periodic orbits appear as the vertical patterns in the coupling submatrices in Figure 9.

For example, consider the two blocks in the upper right of the computed correlation matrix, on the first and second row blocks. These two blocks correspond to collocation points on ellipse 1 or ellipse 2, with an observation point varying on

the third obstacle. The center point $\mathbf{p}$ that is indicated in the left part of Figure 9 (and in Figure 7) corresponds to the leftmost point on the third ellipse, which is in the shadow of the first one. The correlations are high when concentrated near $\mathbf{p}$, and this is physically the only region where rays can reflect off the third ellipse onto the other two along the shortest path, for a plane wave incidence from the left. Hence, that is where any of the stationary points must lie. The compressed matrix $\tilde{A}$ that follows from this $R_{n,q}$ can be rather sparse away from $\mathbf{p}$ in these two subblocks.

Fig. 9: Correlation matrix $R_{n,q}$ at $k = 2^7$ for the multiple scattering configuration with three ellipses (left), and for the scattering configuration with the trapping obstacle (right). Each row corresponds to a collocation point τ_m on one of the obstacles. Each column corresponds to a center of a moving window σ_q . The value shown is the value of the localized oscillatory integral, i.e., of the oscillatory integral operator applied to a windowed part of the exact solution through the approximation (15).

The Green singularity is clearly visible in the correlation matrix along the diagonal. Note that the three diagonal blocks are the self-interaction of the ellipses. These display a pattern that relates to that shown in Figure 8 before. The first and second diagonal block exhibit large correlations in much wider anti-diagonal side-lobes above and below the main diagonal. These correspond to a stationary point in the illuminated region, for a collocation point in the shadow region of the obstacles, much like for the single scattering problem shown in Figure 8. The additional side-lobe for obstacle 2 may seem counterintuitive as it is in its lower shadow region, but it represents the reflections with obstacle 1 through c_j in (15). It has also been interpreted in § 5.5 as obstacle 1 acting as a source, creating a shadow region on obstacle 2. These side-lobes are less pronounced on the third obstacle, because it lies in the shadow of the first one.

The results are similar for the non-trapping single scattering obstacles. More specifically, a collocation point in the shadow region has a stationary point in the illuminated region, in agreement with § 4.4. The computed correlations as in Figure 9 seem quite good at following the stationary points that one can deduce from § 4.4 or reflection arguments.

For the trapping obstacle in the right part of Figure 9, $R_{n,q}$ is high for all $q \in [600, 1300]$ for most n . So in $\tilde{A}$, almost all of the illuminated region is included in the windows, due to the trapping of rays inside the cavity. As a result, there is very little compression. Of course, one could not hope for much better in view of the complicated high-frequency behaviour of this obstacle, which figures like these could help to understand. Also, one could gain insight by viewing this trapping domain as a mix of an exterior with an interior scattering problem. Additionally, there is a very high correlation near the points $\mathbf{q}$, $\mathbf{r}$ and $\mathbf{s}$ which are also indicated in Figure 7. An important side-effect of the complicated high-frequency behaviour is that some of the correlations near the diagonal, covering the Green function's singularity, were rather small. However, in the implementation we always include the Green singularity into the support of the window function, regardless of the value of the correlations.

Fig. 10: Total acoustic field with interference patterns for the nonconvex polygon at $k = 2^{11}$ (left) and for the two circles at $k = 2^7$ (right). Although the figure resolution is not high enough to accurately distinguish the wavelength, this does illustrate that interference appears, which may influence the correlations.

One may understand these oscillations of the correlations near the diagonal by viewing the total field obtained after solving the system. The simple example at the left of Figure 10 shows an interference pattern within the nonconvex part of the obstacle. Although the trapping obstacle exhibits significantly more complicated high-frequency scattering behaviour, the nonconvex polygon also has some small correlations near the diagonal, so we expect these to originate in interference. Similarly, there is interference between the two circles shown in the right of Figure 10. This is the case for the lower right part of the upper circle and the upper right part of the lower circle, as they can receive rays from the point source at $(1, 1)$ as well as from the other circle.

5.3 Compression error and discretization error

Table 1: Accuracy as a function of the wavenumber for the scattering configuration with two circular obstacles. The windows were computed using correlations at $k = 128$ and reused for larger wavenumbers. The errors indicate how well the scattered field matches the boundary condition, for the exact solution $c = A^{-1}b$ and the approximate solution $\tilde{c} = \tilde{A}^{-1}b$ computed after asymptotic compression.

Residual int. eq.	$k = 128$	$k = 256$	$k = 512$	$k = 1024$
c	6.78e-4	4.82e-4	2.55e-4	2.34e-4
$\tilde{c}$	7.10e-4	4.86e-4	2.56e-4	2.34e-4

The overall goal is to match the error because of asymptotic compression to the discretization error, as one can not expect to improve on that: the correlation thresholds were chosen accordingly. The accuracy results for the multiple scattering configuration with two disks are summarized in Table 1. Here, the discretization error is around 0.02% using the number of points per wavelength indicated in Figure 7 in our standard BEM implementation. Using a threshold percentage of $\xi = 0.3\%$, we can compute the scattered field using the compressed matrix $\tilde{A}$ without additional loss of accuracy. In this example, we have used a fixed window computed at $k = 128$ and reused that window for larger values of the wavenumber k .

Table 2: Results similar to those in Table 1, but for the trapping obstacle and using adaptive recompression. An advantage of recompression is that we can start at a smaller value of the wavenumber.

Residual int. eq.	$k = 16$	$k = 32$	$k = 64$	$k = 128$	$k = 256$	$k = 512$	$k = 1024$	$k = 2048$
c	1.81e-2	1.65e-2	1.32e-2	1.39e-2	1.16e-2	1.62e-2	1.59e-2	1.87e-2
$\tilde{c}$	1.81e-2	1.75e-2	1.42e-2	1.45e-2	1.19e-2	1.65e-2	1.61e-2	1.90e-2

The results when using adaptive recompression at higher frequencies are similar, since the recompressed correlation matrix $R_{n,q}$ at a wavenumber k_j behaves in the same way as when we had computed full correlations at k_j : we just recursively leave out negligible contributions. One example is shown in Table 2 for the trapping obstacle. This example shows that the errors of the method do not deteriorate when using adaptive recompression, even for an obstacle with complicated high-frequency behaviour. While the error of $\tilde{c}$ approached that of c in Table 1 for a fixed percentage of nonzeros, recompression in Table 2 has an approximately constant error for a decreasing percentage of nonzeros, see Figure 8. Both tables used linear basis functions, but the errors in the latter were higher as the trapping domain is more difficult.

Figure 11 shows the percentage of nonzero elements as a function of the threshold percentage ξ , for the almost convex obstacle. For low ξ , $\tilde{c}$ provides a good ap-

Fig. 11: Errors as a function of the threshold percentage ξ at wavenumber $k = 2^8$ for the almost convex obstacle. The errors rise with increasing ξ , but the percentage of nonzeros decreases from 99.93% at $\xi = 1.3 \times 10^{-5}$ to 17.5% at $\xi = 0.546$.

proximation of the solution vector, as our compressed matrix $\tilde{A}$ has a low forward compression error then. The errors on the scattered field also decrease proportionally to ξ until they reach the order of the error when using the full matrix A . This is shown in two ways: the error in satisfying the boundary conditions (the residual error of the integral equation), and the total interior field. The latter should be zero for this particular Dirichlet problem and it saturates at about 0.004% with and without compression, averaged over 100 points inside the obstacle chosen randomly for each ξ . We have explored some values of ξ and chosen the ones in Figure 7, but these could be further optimised as mentioned in § 7.

5.4 Computational efficiency and condition numbers

Now, we summarize performance aspects of the adaptive asymptotic recompression scheme that was introduced in § 4.3. First, we measure the efficiency of the frequency sweep with adaptive recompression, compared to one where the full discretization matrix is constructed. We have used Matlab2017a on a 64-bit machine with 66 GB memory and 32 Intel(R) Xeon(R) CPU E5-2650 v2 CPU's at 2.6 Ghz: the results are shown in Figure 12 for the nonconvex polygon. Although no other tasks were burdening this machine, there is an outlier at $k = 1904$. At each other wavenumber, it is faster to construct $\tilde{A}$ (estimated slope $k^{1.680}$ using the times for $k > 256$ excluding the outlier) than the full matrix A ($k^{1.927}$), though the factor is not very large in this example. However, a large gain is observed when one would solve the resulting linear system using an iterative solver as the time to solve the

standard system becomes about the time to compute A . The estimated slope for GMRES on $\tilde{A}$ is $k^{1.878}$, while the direct method takes about $\mathcal{O}(k^{2.5})$ time for both matrices. Moreover, the gains accumulate in a frequency sweep.

Fig. 12: Time (s) at each wavenumber for the different computations needed for the nonconvex polygon, where we can choose between gmres and the MATLAB backslash. We use linear basis functions and recompression in $\tilde{A}$. The wavenumber k increases linearly in steps of 16, while our correlations are efficiently recomputed for k at integer powers of two.

For this example, the total cost for standard BEM summed over the 128 linearly spaced (as in an FFT) wavenumbers is 1.20e5 seconds. This consists in constructing a full $9k \times 9k$ matrix and solving the associated linear system using a direct solver at wavenumber $k_1 = 16$, then $k_2 = 32$ and so on until $k_{128} = 2048$. In contrast, the total cost of computing all correlations, constructing all 128 compressed matrices and solving the resulting linear systems, is 4.04e4 seconds excluding the outlier, or one third of the cost of standard BEM. This gap is only expected to increase for larger wavenumbers as the compression rates will improve. The computations are independent of the results at other k for standard BEM, while in order to compute the cost of our algorithm at one single wavenumber, one should include those of computing the previous $R_{n,q}$ as well as $\tilde{A}$ at those exponentially spaced lower wavenumbers. Especially when one would use iterative solvers, solving the linear system will become the dominant cost, which our method decreases through a sparse matrix and low number of iterations.

Finally, Figure 13 shows that the compression leads to an improvement in condition number of $\tilde{A}$ with respect to the standard BEM matrix A , which then lowers the number of iterations of an iterative solver. For comparison purposes, we have added lines with respective slopes $k^{3/2}$ and $k^{1/2}$ for the condition numbers, and k and $k^{1/4}$ for the number of iterations. In general, however, the obstacle and/or the incident wave may be too intricate to allow any significant compression or such speedup of the iterative solution at the wavenumbers considered. We have used

Fig. 13: Estimated condition number (left) and number of GMRES iterations (right) for some obstacles when using adaptive recompression ($\tilde{A}$) and without compression (A) for wavenumbers at integer powers of two.

a non-restarted GMRES in Matlab with a tolerance of 10^{-5} for this experiment. The same qualitative behaviour is observed for all obstacles, using piecewise linear basis functions. The condition number improves by a factor of 1.1 to 3 compared to using fixed windows from correlations computed at $k = 2^7$, although the number of iterations varies less predictably about that range.

5.5 Visibility criterion

Next, we illustrate the possibilities and limitations of compression using just the visibility criterion outlined in § 4.1. We focus on scatterers with complicated asymptotic structure of the rays to illustrate its general applicability, although one can not gain much from asymptotics in such cases. We use the trapping domain as a first example. The domain and its parameterization are shown in Figure 14. The corresponding sparsity pattern of $\tilde{A}$ is shown in the right of Figure 14, where we have used the visibility criterion to perform compression only for collocation points in the illuminated region.

We have highlighted a collocation point corresponding to $\tau = 0.5$, and Figure 14 shows the part of the obstacle that is directly visible from this point. It includes much of the cavity, but nothing outside of it. We add a dashed blue region where the window can smoothly decay to zero. This window translates into the large dense square in the centre of the sparsity pattern of $\tilde{A}$: all points inside the cavity are asymptotically linked directly to each other. The dense block captures the complicated multiple scattering going on inside the cavity. Yet, this whole block is asymptotically independent of the solution outside the cavity and in the illuminated region, and this yields the zero elements outside the range $[0.365 - T, 0.665 + T]$ for collocation $\tau \in [0.3, 0.72]$, where the one-sided support of the decaying part $T = 0.1$.

Fig. 14: Left: Obstacle, parametrisation and region which is visible from $\tau = 1/2$ for the trapping obstacle. Right: Structure of the compressed matrix using the visibility criterion.

Collocation points with τ approximately in $[0.2, 0.3]$ or $[0.72, 0.8]$ correspond to a locally convex part of the obstacle in the illuminated region. These points would only be influenced by themselves, as they cannot ‘see’ other points. This results in a small band $[\tau - 0.02 - T, \tau + 0.02 + T]$ around the diagonal in the discretization, above and below the central square at the right of Figure 14. As mentioned, we have not attempted to apply compression based on visibility in the shadow region ($\tau \notin [0.2, 0.8]$).

We compute $\tilde{A}$ at $k = 2^9$ using this structure, which results in a relative error on c of 12.1% and a worsening of the residual of the integral equation on 100 random points from 1.19% to 5.84%. This gives a qualitatively accurate solution, but the errors are quite large. The errors do respectively improve to 3.85% and 2.38% by not applying compression in the locally convex parts of the obstacle in the illuminated region. Or one can increase the frequency, as it seems that the trapping region still significantly influences those parts at $k = 2^9$ due to creeping rays.

It can be expected that the visibility criterium might be easier to validate for multiple scattering of convex obstacles, or obstacles with a finite number of reflections. Here, one has to be careful in what it means to be ‘visible’. We test the multiple scattering configuration with two circles shown in Figure 7 at $k = 256$ where we change the point source to a plane wave coming from the left to obtain the more pronounced behaviour described below. First, we compress parts of the rows that correspond to self-interactions in the region that is directly illuminated by the source. More precisely, if the x -coordinate is less than $-0.5 - T$ at $\kappa(t)$, then $\lambda = t - 2T, l = t - T, r = t + T$ and $\rho = t + 2T$ for $T = 0.15$ and there is no compression elsewhere. However, the 92% nonzeros worsen the relative residue of the integral equation from 1.0835% to 5.41% and give a relative error on c of 21.28%.

These high errors are explained by the correlation matrix shown in Figure 15. In the middle rows of the self-interaction blocks, one not only notices the Green singularity on the diagonal, but also a stationary point elsewhere. For example,

Fig. 15: Correlations for the two circles with an incident plane wave from the left at $k = 256$ (left) and sparsity pattern using visibility (right).

in the upper left block of Figure 15, there is a side lobe along the antidiagonal direction that ranges approximately from the row labelled 1500 to the one labelled 2500. These rows correspond to collocation points in the lower half of the lower obstacle. These points are in the shadow with respect to the upper obstacle, and therefore they have a stationary point in the illuminated part, i.e., the upper half of the lower circle. The concepts of illumination and shadow are relative to the other obstacle here, not to the incoming plane wave.

Applying compression solely for collocation points on parts of the obstacle that are illuminated both by the incoming wave and by the other obstacle is successful: the residual of the integral equation and the error on c become 1.0832% and 0.0816% respectively for 96% nonzeros. This indicates that each obstacle in a multiple scattering configuration should be viewed as a source on top of the incoming wave when applying the visibility criterion.

Yet, much more compression is possible: it is quite clear from the correlation matrix in Figure 15 that there is large potential for compression in the coupling blocks. Indeed, any point in the upper half of the upper obstacle can not ‘see’ the whole lower obstacle. Similarly, no point in the lower half of the lower obstacle can be a stationary point for a collocation point anywhere on the whole upper obstacle. In the upper right and lower left blocks of the correlation matrix, this yields large blue regions amenable to compression. The additional compression leads to the 85% nonzeros in the right of Figure 15, a residual of the integral equation of 1.0839% and an error on c of 0.0817%. The compression rate could improve with increasing frequency as the width of the decaying part T can decrease but has a lower bound independently of k . This bound is due to geometric considerations and is quite high, also because of this concept of generalised shadow.

The visibility approach could be made automatic and more robust, and whether two points can ‘see’ each other is independent of the incident wave. However, one can only gain a compression rate which does not improve with k and which deteriorates with more complex geometries or more obstacles, as each of them

acts as a source giving a shadow. The shadow region is also dependent on the incident wave and we applied no compression there; moreover, it is not clear how a visibility criterion would incorporate a trapping region of an obstacle for moderate wavenumbers. The visibility criterion would benefit from a better understanding of the stationary points of collocation points that are in the shadow (recall § 4.4).

5.6 Additional robustness results

To begin, compressing the BEM matrix also helps against ill-conditioning of the integral formulation of the first kind near resonances. As a simple example, the resonant frequencies for a circle of radius $1/2$ are $k = 2j_{n,m}$ (the m th zero of the n th unmodified Bessel function of the first kind). The full discretization matrix of this problem in our code has condition number 1020 for $n = 50$ and $m = 1$. The condition number is large but not infinite, due to the discretization error. Still, after adaptive compression with the point source from Figure 7, the condition number is reduced to 36.9 with 34% nonzeros, while for an incident plane wave from the left, the condition number is reduced to 37.3. The eigenfunction of the problem is inherent to the domain, not to the boundary condition, and it seems entirely eliminated by the compression.

Secondly, we illustrate a simplification of the compression scheme based on hard truncation: we simply discard all elements of A where they would be zero in $\tilde{A}$ in the adaptive scheme, but we make no modifications to the other entries of A . Indeed, this corresponds closely to using a block window function, rather than the smooth function (11). As mentioned in § 2, this should lead to spurious $\mathcal{O}(k^{-1})$ endpoint contributions for a general oscillatory integral. We consider the trapping obstacle at $k = 128$. Without compression, there is a residual of 1.09% on the integral equation; with smooth windows, this becomes 1.26% for 2.80% error on c ; with the non-smooth compression windows, only 3.33% and 12.1% respectively. The errors clearly grow, but not excessively so. The reason is, at least in this case, that the contributions of the Green's singularity and stationary points dominate. The spurious contributions caused by the discontinuous compression scheme are more than an order of magnitude smaller, such that one still obtains the correct qualitative behaviour of the solution.

Finally, the previous experiment is repeated with piecewise cubic basis functions $\varphi_j(\tau)$. The condition number improves from 1.28e3 to 295, while this was 1.27e3 to 116 for linear basis functions. The residual on the integral equation with cubic basis functions marginally changes from 1.16% to 1.33%, for an error on c of 4.66%. Although we mainly considered piecewise linear basis functions in this article, adaptive asymptotic (re)compression thus can be applied to more general basis functions.

6 Preliminary results in 3D: asymptotic compression on a sphere

The principles underlying the asymptotic compression scheme explained in § 3 are valid in higher dimensions as well. In this section, we perform a limited experiment to illustrate the validity of the asymptotic compression scheme in 3D. We illustrate the performance by applying it to scattering off a unit ball using the

openly available software package BEMPP [36]. Our 3D experiments can be found in the CPP code of Release 1.0 of [34].

6.1 Galerkin discretization of a spherical obstacle

Motivated by the reasonable results when using piecewise linear basis functions in 2D, in spite of the typical smoothness of integrands that is assumed in asymptotic analysis, we stretch the limits and choose piecewise constant basis functions in 3D. We have used a mesh of the unit sphere with 3219 piecewise constant triangular elements for $k = 8$ and approximately four times more elements for $k = 16$. We respect the scaling $N = \mathcal{O}(k^2)$ for the total number of degrees of freedom for larger wavenumbers.

The BEMPP software is based on Galerkin discretization, and hence we adapt our scheme to use Galerkin rather than collocation. This change only affects the choice of the window function: in a collocation scheme we only have to choose a suitable window for each collocation point t_i , whereas in a Galerkin scheme, t is a continuous variable in the window $w(t, \tau)$. In order to avoid this additional complexity, and since we are dealing with a non-smooth integrand anyway, we can simplify the computation of matrix entries considerably using the technique suggested in § 3.3 and in particular equation (12). That is, rather than recomputing the matrix entries using windowed integrands, we simply use the matrix elements as computed by BEMPP and we weigh them according to the window function:

$$\tilde{A}_{i,j} = \int_{S_i} \int_{S_j} K_\kappa(t, \tau) w(t, \tau) \varphi_j(\tau) \varphi_i(t) \|\kappa(t)\| d\tau dt \quad (20)$$

$$\approx w(t_i, \tau_j) \int_{S_i} \int_{S_j} K_\kappa(t, \tau) \varphi_j(\tau) \varphi_i(t) \|\kappa(t)\| d\tau dt = w(t_i, \tau_j) A_{i,j}. \quad (21)$$

Here, we choose the points t_i and τ_j to be the centres of S_i and S_j , the supports of the basis functions φ_i and φ_j . This very practical approximation is quite crude, and yet experiments indicate it is nearly as accurate for the overall solution error as when computing the windowed integrals: see § 6.3 further on.

6.2 Compression and convergence for small wavenumbers

We consider a plane wave incidence in the x -direction, analogously to § 4.4 for scattering by a disk. The asymptotic behaviour of the solution is rather similar. There are no stationary points of the integrand when $\kappa(t)$ lies in the illuminated region. There is a single stationary point in the illuminated region when t is in the shadow, and the stationary point has the same y and z coordinate as $\kappa(t)$.

Because of the lack of a natural ordering of points on a 2D boundary, we define a window function in space, rather than in the parameter domain as before. The window depends on a parameter T that determines its size, and is shaped like a ball around the point $\mathbf{x} - \mathbf{y}$:

$$w(\mathbf{x}, \mathbf{y}) = \chi \left(\sqrt{(x_x - x_y)^2 + (y_x - y_y)^2 + (z_x - z_y)^2}, 0, 0, 0.8T, T \right).$$

Fig. 16: 3D example at low wavenumbers: Condition number (left) and relative error on the solution vector c (right).

For the purposes of our experiment, a more convenient choice avoids defining two windows via

$$w(\mathbf{x}, \mathbf{y}) = \chi \left(\sqrt{(y_x - y_y)^2 + (z_x - z_y)^2}, 0, 0, 0.8T, T \right).$$

since that window represents a circular tube in the x -direction through the obstacle that automatically includes the stationary point in the illuminated region when t is in the shadow. That is because the y and z coordinates of t and its stationary point agree in analogy to § 4.4.

The left side of Figure 16 shows that the condition number improves markedly when T is small. Figure 16 also shows that the relative error decreases with increasing support size T , and the error is smaller for larger wavenumbers k (at least for small T). Yet, for small T the errors are rather large and this is due to the fact that the wavenumbers are too small to observe asymptotic error decay. Using larger meshes at higher wavenumbers gives rise to memory issues in our direct implementation, when initialising A or $\tilde{A}$ as a full $N \times N = \mathcal{O}(k^4)$ matrix. For the latter, it should be possible to use a compressed representation such as a sparse or $\mathcal{H}^2$ matrix. Thus, memory usage would reduce significantly and a hierarchical structure allows checking more efficiently whether $\mathbf{y}$ is in the support of $w(\mathbf{x}, \mathbf{y})$. This approach shows promise, but its practical implementation is quite involved. Instead, in the next section, we merely illustrate asymptotic accuracy of the scheme for larger wavenumbers using the fact that we know the exact solution.

6.3 Compression accuracy at larger wavenumbers

In order to avoid having to compute an expensive reference solution of the scattering problem at larger wavenumbers, we restrict ourselves to a limited experiment using the known exact solution. First, we compute a reference solution by projecting the analytical solution onto the discretization space of piecewise constant functions. Next, we compute a full row of the discretization matrix and a compressed one, and assess the corresponding error in the matrix-vector product.

An exact solution for the scattered field as an infinite expansion involving spherical harmonics is given in, e.g., [25, pg. 292-294] for Neumann boundary

Table 3: Average relative residual of the reference solution measured at 40 points.

k	8	16	32	64	128	256
Residual	0.0383	0.0356	0.0341	0.0338	0.0317	0.0276

Fig. 17: Relative error, as a function of the wavenumber, of the matrix vector product $Ac = b$ for a row i corresponding to φ_i supported in the illuminated region (left panel) and in the shadow region (right panel) of the sphere. ‘Full row’ means we do not use a windowed integral, ‘Kernel 0.6’ means we integrate the windowed kernel function with $T = 0.6$, and ‘Element 0.2’ means we use approximation (21) with $T = 0.2$.

conditions. The normal derivative of a truncation of this expansion yields a highly accurate approximation of the density function, which we subsequently project onto a basis of piecewise constants. We refine the mesh in each spatial dimension as the wavenumber increases using Gmsh [20]: more specifically, the mesh has approximately $50k^2$ elements. The accuracy of this reference solution is illustrated in Table 3. We apply the integral operator to the density function and evaluate at a point on the boundary. This value is compared to the boundary condition at 40 points placed along the sphere. The error is roughly constant for increasing wavenumbers, and it matches the discretization error of the integral operator.

We have computed a number of rows of the dense matrix A , as well as of the compressed matrix $\tilde{A}$ using (21). We can also evaluate the right hand side b of the linear system, using the boundary condition. The exact solution vector c of the linear system satisfies $Ac = B$. The projection of the known analytical solution onto the basis yields a vector $\bar{c}$. We define $\tilde{b} = A\bar{c}$, and note that Figure 17 shows that $\tilde{b} - b$ is small – in particular it is on the order of the discretization error.

Using the compressed matrix $\tilde{A}$, we have the approximation

$$\tilde{A}\bar{c} = \tilde{b}.$$

In Figure 17, we show the relative error of the compressed matrix-vector product element $\tilde{b}_i$ compared to the exact one b_i ,

$$\left| \frac{b_i - \tilde{b}_i}{b_i} \right|.$$

For reference, the figure also includes the relative error of $\bar{b}_i$ versus b_i ('Full row'). These values correspond to the results in Table 3 and they are indicative of the discretization error.

In this approach, we can cheaply consider larger wavenumbers. We have used a row corresponding to φ_i in the illuminated region in the left panel of Figure 17, and in the shadow region in the right panel. In the latter case, the major contribution to the integral originates in a stationary point in the illuminated region. The results show that the error decreases with increasing k in both cases. We have illustrated different values of T , the width of the window function. The error is initially larger for smaller T , but still converges to the level of the discretization error for increasing k . The convergence is slower for the case of the shadow region, but recall from the 2D experiment that a stationary point requires a larger window and we are comparing the same support in these results.

Finally, the figure also includes the results when integrating the windowed integrand (20), rather than using the approximation (21). It shows that no accuracy is lost when using the much cheaper approximation, as the results are very close to each other. In fact, the errors in Figure 17 for $T = 0.6$ agree up to at least seven digits, which is the case for $k = 8$.

7 Concluding remarks

Our results show that exploiting asymptotic behaviour in scattering problems is possible without a priori knowledge of the asymptotic behaviour of the solution. Compared to true asymptotic methods, our approach is costly as all oscillations are still resolved with a fine discretization. Yet, we demonstrated that asymptotic compression can be applied simply by rescaling standard BEM matrix elements and is robust and automatic in several ways for problems with highly complicated high-frequency asymptotics (3D Galerkin in § 6 and 2D collocation in § 5). The principles underlying the scheme appear to be valid for other integral equation formulations, boundary conditions, general objects, arbitrary incident waves, basis functions, higher dimensions, other physical problems, orders of the BEM basis functions and alternative discretisation schemes as well. Some modifications to our algorithm can be made to increase the efficiency: T could decrease with k and the influence of other parameters can be further investigated. Furthermore, the windowing procedure may be applied in the computation of oscillatory integrals that represent the scattered field, after the boundary integral equation has been solved.

There are several opportunities for further research. First, though we have avoided it in this paper, the present method may be combined with asymptotic methods. Any available asymptotic information can be used to reduce the cost of computing the solution. One may combine approximate asymptotic knowledge with the robustness of the present approach, for example if a small number of geometrically complicated parts of the boundary prevent a full asymptotic scheme. Figures of the correlations could help obtaining this asymptotic knowledge for intricate high-frequency behaviour. For example, Figure 9 indicates different contributions to the integral, whose phases are equal to the distance to physically relevant stationary point plus the phase(s) there. The extension of the principle

of Fermat may help too, as well as the localisation of regions where the gradient of the phase is small but does not cross zero.

Alternatively, the method can be combined with non-asymptotic methods. The Fast Multipole Method and $\mathcal{H}^2$ matrices yield a fast matrix vector product for discretizations of scattering problems, with optimal time and memory complexity. Since our method essentially amounts to modifying the Green's function, the methods may be compatible. The sparsification might not provide an important improvement as FMM already reduces memory usage, but the significant reduction in the condition number may make the linear system much easier to solve. Both may also be combined with local quadrature correction or quadrature by expansion [31], which are expected to improve the accuracy of the integrals we compute. The computation of the integrals may also be sped up using other methods for the evaluation of oscillatory integrals, such as the quantized-tensor train in combination with a cross-approximation scheme [30] or some of the methods discussed in [13].

Acknowledgements The authors would like to thank Samuel Groth, Stephen Langdon, Niels Billen, Philip Dutré, Karl Meerbergen, Laurent Jacques and Dave Hewett for interesting and helpful discussions on this paper and related topics. The authors were supported by FWO Flanders [projects G.0617.10, G.0641.11 and G.A004.14].

References

1. Alouges, F., Aussal, M.: The sparse cardinal sine decomposition and its application for fast numerical convolution. *Numer. Algor.* **70**, 427–448 (2015)
2. Anand, A., Boubendir, Y., Ecevit, F., Reitich, F.: Analysis of multiple scattering iterations for high-frequency scattering problems. II: The three-dimensional scalar case. *Numer. Math.* **114**, 373–427 (2010)
3. Asheim, A., Huybrechs, D.: Extraction of uniformly accurate phase functions across smooth shadow boundaries in high frequency scattering problems. *SIAM J. Appl. Math.* **74**(2), 454–476 (2014)
4. Babich, V.M., Buldyrev, V.S.: Short-wavelength diffraction theory. Springer-Verlag, Berlin (1991)
5. Bebendorf, M.: Hierarchical LU decomposition-based preconditioners for BEM. *Computing* **74**, 225–247 (2005)
6. Bleistein, N., Handelsman, R.A.: Asymptotic expansions of integrals. Dover Publications, Inc., 31 East 2nd Street, Mineola, N.Y. (1986)
7. Brakhage, H., Werner, P.: Über das Dirichletsche Außenraumproblem für die Helmholtzsche Schwingungsgleichung. *Arch. Math.* **16**, 325–329 (1965)
8. Bruno, O., Geuzaine, C., Monro, J.J., Reitich, F.: Prescribed error tolerances within fixed computational times for scattering problems of arbitrarily high frequency: the convex case. *Phil. Trans. R. Soc. Lond. A* **362**, 629–645 (2004)
9. Chandler-Wilde, S., Graham, I., Langdon, S., Spence, E.: Numerical-asymptotic boundary integral methods in high-frequency acoustic scattering. *Acta Numer.* pp. 89–305 (2012)
10. Chandler-Wilde, S.N., Hewett, D.P., Langdon, S., Twigger, A.: A High Frequency Boundary Element Method for Scattering by a Class of Nonconvex Obstacles. *Numerische Mathematik* **129**(4), 647–689 (2015)
11. Cheng, H., Crutchfield, W.Y., Gimbutas, Z., Greengard, L.F., Ethridge, J.F., Huang, J., Rokhlin, V., Yarvin, N., Zhao, J.: A wideband fast multipole method for the helmholtz equation in three dimensions. *J. Comput. Phys.* **216**(1), 300 – 325 (2006)
12. Colton, D.L., Kress, R.: Integral equation methods in scattering theory. John Wiley & Sons, New York (N.Y.) (1983)
13. Deaño, A., Huybrechs, D., Iserles, A.: Computing highly oscillatory integrals, vol. 155. SIAM, Philadelphia (2018)

14. Domínguez, V., Graham, I.G., Smyshlyaev, V.: A hybrid numerical-asymptotic boundary integral method for high-frequency acoustic scattering. *Numer. Math.* **106**, 471–510 (2007)
15. Ecevit, F., Reitich, F.: Analysis of multiple scattering iterations for high-frequency scattering problems. I: The two-dimensional case. *Numer. Math.* **114**, 271–354 (2009)
16. G. Beylkin, C.K., Monzón, L.: Fast algorithms for Helmholtz Green’s functions. *Proc. Roy. Soc., Ser. A* **464**, 3301–3326 (2008)
17. Ganesh, M., Hawkins, S.: A fully discrete Galerkin method for high frequency exterior acoustic scattering in three dimensions. *J. Comput. Phys.* **230**, 104–125 (2011)
18. Ganesh, M., Langdon, S., Sloan, I.: Efficient evaluation of highly oscillatory acoustic scattering surface integrals. *J. Comput. Appl. Math.* **204**, 363–374 (2007)
19. Geuzaine, C., Bruno, O., Reitich, F.: On the $O(1)$ solution of multiple-scattering problems. *IEEE Trans. Magn.* **41**(5), 1488–1491 (2005)
20. Geuzaine, C., Remacle, J.F.: Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering* **79**(11), 1309–1331 (2009)
21. Giladi, E.: Asymptotically derived boundary elements for the Helmholtz equation in high frequencies. *J. Comput. Appl. Math.* **198**, 52–74 (2007)
22. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**(2), 325–348 (1987)
23. Groth, S., Huybrechs, D., Opsomer, P.: High-order terms in the ray expansion for high frequency scattering by single and multiple obstacles (2017). In preparation
24. Groth, S.P., Hewett, D.P., Langdon, S.: Hybrid numerical-asymptotic approximation for high-frequency scattering by penetrable convex polygons. *IMA J. Appl. Math.* **80**, 324–353 (2015)
25. Harrington, R.F.: Time-harmonic electromagnetic fields. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331 (1961)
26. Huybrechs, D., Vandewalle, S.: A two-dimensional wavelet-packet transform for matrix compression of integral equations with highly oscillatory kernel. *J. Comput. Appl. Math.* **197**, 218–232 (2006)
27. Huybrechs, D., Vandewalle, S.: A sparse discretization for integral equation formulations of high frequency scattering problems. *SIAM J. Sci. Comput.* **29**(6), 2305–2328 (2007)
28. Huybrechs, D., Vandewalle, S.: An efficient implementation of boundary element methods for computationally expensive Green’s functions. *Eng. Anal. Boundary Elem.* **32**(8), 621–632 (2008)
29. Khoromskij, B.N.: Tensor-structured preconditioners and approximate inverse of elliptic operators in $\mathbb{R}^d$. *J. Constr. Approx.* **30**, 599–620 (2009)
30. Khoromskij, B.N., Veit, A.: Efficient computation of highly oscillatory integrals by using qtt tensor approximation. *Comput. Methods Appl. Math.* **16**(1), 145–159 (2016)
31. Klöckner, A., Barnett, A., Greengard, L., O’Neil, M.: Quadrature by expansion: A new method for the evaluation of layer potentials. *J. Comput. Phys.* **252**, 332–349 (2013)
32. Kress, R., Spassov, W.T.: On the condition number of boundary integral operators in acoustic and electromagnetic scattering. *Numer. Math.* **42**, 77–95 (1983)
33. Melrose, R.B., Taylor, M.E.: Near peak Scattering and the corrected Kirchhoff Approximation for a Convex Obstacle. *Advances in Mathematics* **55**, 242–315 (1985)
34. Opsomer, P.: <https://github.com/popsomer/bempp.git> (2016). URL <https://github.com/popsomer/bempp.git>. Release: Asymptotic compression version 3
35. Rokhlin, V.: Rapid Solution of Integral Equations of Classic Potential Theory. *J. Comput. Phys.* **60**, 187–207 (1985)
36. Śmigaj, W., Betcke, T., Arridge, S., Phillips, J., Schweiger, M.: Solving Boundary Integral Problems with BEM++. *ACM T. Math. Software* **41**(2), 6:1–6:40 (2015)
37. Sweldens, W., Piessens, R.: Quadrature formulae and asymptotic error expansions for wavelet approximations of smooth functions. *SIAM J. Numer. Anal.* **31**(4), 1240–1264 (1994)
38. Wald, I., Mark, W.R., Gunther, J., Boulos, S., Thiago, I., Hunt, W., Parker, S.G., Shirley, P.: State of the Art in Ray Tracing Animated Scenes. *Eurograph.* (2007)
39. Wong, R.S.: Asymptotic approximations of integrals. SIAM, 3600 Market Street, 6th floor, Philadelphia (Pa.) 19104-2688 USA (2001). Republication of 1944
40. Wu, T.: Boundary element acoustics. WIT Press (2000). Reprint 2005
41. Ying, L.: Fast directional computation of high frequency boundary integrals via local FFTs. *SIAM Multiscale Model. Simul.* **13**(1), 423–439 (2015)