



Corner singularities for elliptic problems: Integral equations, graded meshes, quadrature, and compressed inverse preconditioning[☆]

Johan Helsing^{*}, Rikard Ojala

Numerical Analysis, Centre for Mathematical Sciences, Lund University, P.O. Box 118, SE-221 00 Lund, Sweden

ARTICLE INFO

Article history:

Received 11 January 2008

Received in revised form 11 June 2008

Accepted 23 June 2008

Available online 1 July 2008

Keywords:

Corner singularity
Multiphase material
Triple-junction
Integral equation
Mesh grading
Conductivity

ABSTRACT

We take a fairly comprehensive approach to the problem of solving elliptic partial differential equations numerically using integral equation methods on domains where the boundary has a large number of corners and branching points. Use of non-standard integral equations, graded meshes, interpolatory quadrature, and compressed inverse preconditioning are techniques that are explored, developed, mixed, and tested on some familiar problems in materials science. The recursive compressed inverse preconditioning, the major novelty of the paper, turns out to be particularly powerful and, when it applies, eliminates the need for mesh grading completely. In an electrostatic example for a multiphase granular material with about two thousand corners and triple junctions and a conductivity ratio between phases up to a million we compute a common functional of the solution with an estimated relative error of 10^{-12} . In another example, five times as large but with a conductivity ratio of only a hundred, we achieve an estimated relative error of 10^{-14} .

© 2008 Elsevier Inc. All rights reserved.

1. Introduction

The last decades have seen huge progress on solving elliptic partial differential equations numerically using Fredholm second kind integral equation methods. Problems with smooth and well-separated boundaries and simple boundary conditions for Laplace's and similar equations in two dimensions are well understood. To solve them with fast iterative methods has become standard tasks. Recent research on high performance computing within this field is often focused on three dimensions [23,29] or on fast direct solvers [22].

In applications, particularly in materials science, geometries of interest, such as aggregates of grains and fractured specimens, almost never have smooth and well-separated boundaries. Branching points are frequent and second kind integral equations which otherwise are of Fredholm type lose important properties. Smooth kernels develop fixed (near) singularities. Layer densities, which are to be solved for, exhibit complicated asymptotic behavior. This, too, is well-understood theoretically [10]. Generalized Gaussian quadrature [3], graded meshes [1], tailor-made basis functions [7,8,14,26], Sinc quadrature [21], conformal mapping [5], and the method of images [2] are some techniques that have been suggested in order to avoid uneconomical mesh refinement. But when it comes to constructing efficient and versatile numerical schemes that simultaneously resolve kernels and densities close to boundary singularities of general types there still seems to be room for improvement, even in two dimensions [27]. Furthermore, there are reasons to reconsider the choices of the integral equations themselves. If one can find one second kind integral equation for a PDE on a given geometry, one can often find

[☆] This work was supported by the Swedish Research Science Council under contract 621-2007-6234.

^{*} Corresponding author.

E-mail addresses: helsing@maths.lth.se, Johan.Helsing@na.lu.se (J. Helsing), rikardo@maths.lth.se (R. Ojala).

many [12,17,19,24]. Which choice is the best depends on, for example, values of material parameters, needs for accuracy, regularity of boundary values, and what aspect of the solution is of interest. Sometimes, switching integral equation formulation could have greater impact on computational efficiency than advanced numerical techniques.

This paper treats two familiar boundary value problems on non-smooth infinite domains: one electrostatic and one elastostatic. The electrostatic problem is treated in detail, while the elastostatic problem is included merely to illustrate the broader applicability of our techniques. We settle for Nyström discretization, composite quadrature, and iterative solution and show how choices of non-standard integral equations in combination with special-purpose interpolatory quadrature [16], in some situations, leads to large improvement in convergence and achievable accuracy. Further speedup is achieved by graded meshes. Here the classic use of grading exponents [1] is replaced by a numerically controlled subdivision procedure. We also introduce a compressed inverse preconditioner. In short, this is a local change of variables that restores the Fredholm property of the integral equation across boundary singularities and makes the new layer density piecewise smooth. This, in turn, reduces the need for resolution and improves the spectral properties of the system matrix. Iterative solvers converge faster. Lastly, we present a recursive construction of the compressed inverse preconditioner. It can be viewed as a fast direct local solver and enables the treatment of geometries of unprecedented complexity. A pleasant side-effect is that the performance of otherwise less efficient integral equations can be substantially increased as the recursive compressed inverse preconditioner gets into action. One is more free to choose a formulation which suits a given problem from a modeling point of view.

A number of different techniques are introduced, combined, and tested in different environments. To help the reader navigate through the paper, we give a crude summary of concepts already at this early stage:

- Singular boundary points lead to problems with computational economy. The problems are particularly severe when high accuracy is sought. Mesh refinement makes spectra of system matrices grow. A finer discretization, thus, induces more iterations and eventually also less accurate results.
- Meshes close to singular boundary points could be refined according to various strategies. A *simply graded mesh* is refined using binary subdivision while an *aggressively graded mesh* is refined using a more advanced scheme, aiming at fewer subdivisions to reach a given resolution.
- *Recursive compressed inverse preconditioning* is a multilevel preconditioning technique based on a coarse grid and a hierarchy of grids on simply graded meshes. It aims at a well-conditioned equation at the top-level. It is extremely powerful when it applies and is the major novelty of the paper.
- *Direct compressed inverse preconditioning* is a two-level preconditioning technique. It has a larger setup cost than recursive compressed inverse preconditioning, but is easier to implement.
- *Plain recursive compression* is a multilevel compression technique. It could be seen as a use of special basis functions whose construction is purely numerical and free of asymptotic analysis.
- *Regularization* is a simple trick to enhance the performance of a quadrature rule in the discretization of integral operators which in a sense are (nearly) singular. It works best for continuous layer densities.
- *Special-purpose interpolatory quadrature* [16] is a more generally applicable and efficient alternative to regularization. It works also for layer densities with jump discontinuities. It has great impact on aggressively graded meshes and can be used to shorten recursion lengths in recursive compression.

When presenting our two boundary value problems we use symbols in accordance with previous work. Some variables have multiple meaning. In electrostatics λ , ρ and σ denote a contrast ratio, a single layer charge density, and conductivity. In elastostatics they denote a symmetric Williams exponent, a weight function, and stress. To simplify the transition between real and complex notation we shall make no distinction between points or vectors in the real plane \mathbb{R}^2 and points in the complex plane \mathbb{C} . We refer to the plane as D .

2. An electrostatic inclusion problem

Let an inclusion made out of N_{gr} grains with conductivities $\sigma_k, k = 1, \dots, N_{\text{gr}}$, be embedded in D while the remainder of D has conductivity σ_0 . The local conductivity $\sigma(z)$ is then a piecewise constant function on D . Let the boundary of all grains be denoted Γ and be given orientation. In general, for $N_{\text{gr}} > 1$, Γ will have branching points in the form of triple-junctions. Let $n_z = n(z)$ be the outward unit normal of Γ at z . Let the electric potential $U(z)$ at a point z far away from the origin obey

$$\lim_{z \rightarrow \infty} U(z) = \Re\{\bar{e}z\}, \quad z \in D, \tag{1}$$

where e is a unit vector which can be interpreted as the (negative) average electric field and \bar{e} is its complex conjugate. Assume now that we want to solve the electrostatic equation in the plane for the purpose of computing the quantity q given by

$$q = \int_{\Gamma} U(z) a(z) \Im\{\bar{e} dz\}, \tag{2}$$

where

$$a(z) = \sigma_+(z) - \sigma_-(z), \quad z \in \Gamma \tag{3}$$

and where $\sigma_+(z)$ is the conductivity on the positive side of Γ and $\sigma_-(z)$ is the conductivity on the negative side of Γ . The electrostatic equation means Laplace's equation in $D \setminus \Gamma$ with potential and normal current being continuous across Γ . The quantity q of (2) is important in homogenization problems and can, for a doubly periodic geometry, be interpreted as the contribution to the average electric current in the e -direction caused by the inclusion.

2.1. Three integral equations

Standard practice for electrostatic problems is to represent $U(z)$ as a sum of a single layer potential and a driving term [11]

$$U(z) = \Re\{\bar{e}z\} + \frac{1}{2\pi} \int_{\Gamma} \rho(\tau) \log|\tau - z| d|\tau|, \quad z \in D. \quad (4)$$

Here $\rho(z)$ is an unknown real valued layer density which can be solved from the integral equation

$$\rho(z) + \frac{\lambda(z)}{\pi} \int_{\Gamma} \rho(\tau) \Im\left\{\frac{n_z \bar{n}_\tau d\tau}{\tau - z}\right\} = 2\lambda(z) \Re\{\bar{e}n_z\}, \quad z \in \Gamma, \quad (5)$$

where

$$\lambda(z) = a(z)/b(z), \quad \text{and} \quad b(z) = \sigma_+(z) + \sigma_-(z). \quad (6)$$

Once $\rho(z)$ is available, q can be computed as

$$q = \sigma_0 \int_{\Gamma} \rho(z) \Re\{\bar{e}z\} d|z|. \quad (7)$$

An alternative integral equation can be derived from Green's second identity and expressed in terms of the potential $U(z)$ itself on Γ as

$$U(z) - \frac{1}{b(z)\pi} \int_{\Gamma} U(\tau) a(\tau) \Im\left\{\frac{d\tau}{\tau - z}\right\} = \frac{2\sigma_0}{b(z)} \Re\{\bar{e}z\}, \quad z \in \Gamma, \quad (8)$$

see [12] and references therein. We observe that after solving for $U(z)$, $z \in \Gamma$, Eq. (8) can also be used as direct evaluation formula for $U(z)$, $z \in D \setminus \Gamma$. One simply sets $b(z) = 2\sigma(z)$ and moves the integral to the right hand side.

Yet another integral equation can be derived for a density $\mu(z)$ using partial integration in the single layer potential equation (5):

$$\mu(z) + \frac{\lambda(z)}{\pi} \int_{\Gamma} \mu(\tau) \Im\left\{\frac{d\tau}{\tau - z}\right\} = 2\lambda(z) \Im\{\bar{e}z\}, \quad z \in \Gamma. \quad (9)$$

Once $\mu(z)$ is available, q can be computed as

$$q = -\sigma_0 \int_{\Gamma} \mu(z) \Re\{\bar{e} dz\}. \quad (10)$$

The integral equations (5), (8) and (9) are similar and related. They are all of Fredholm's second kind with unique solutions when $N_{\text{gr}} = 1$, Γ is smooth, and $|\lambda| < 1$. The operator on the left-hand side of (5) is then the adjoint of the operator on the left-hand side of (8). The operators on the left-hand sides of (8) and (9) transform into each other when the sign of λ is changed. Despite of their similarities, the integral equations and their solutions also exhibit differences which, in turn, have consequences for numerics.

First, for $N_{\text{gr}} = 1$ and $\lambda = 1$ the equations (5) and (8) admits non-trivial homogeneous solutions while (9) does not. For $\lambda = -1$ the opposite holds. This has consequences for the numerical solutions of discretized versions of (5), (8) and (9) when the contrast ratio of the material is high. Values of λ close to plus or minus one yield nearly singular system matrices and possibly very inaccurate solutions. Clearly, Eq. (9) is preferable whenever $\sigma_0 < \sigma_1$.

Second, the regularity of $\rho(z)$, $U(z)$ and $\mu(z)$ are not all the same and this has consequences when it comes to solving problems on domains where Γ has corners and triple-junctions. The three equations (5), (8) and (9) are still solvable when $|\lambda| < 1$, but none of them are of Fredholm's second kind. With the use of variable separation one can investigate the asymptotic behavior of $\rho(z)$, $U(z)$ and $\mu(z)$ close to a corner. One can show that the density $\rho(z)$ diverges and asymptotically behaves as s^{v-1} , where s is the distance to the corner, and $0.5 < v < 1$ is a real number. For the corners in, for example, an equilateral triangle we have for $\lambda > 0$ that v is the smallest positive non-zero solution to

$$\lambda \sin\left(\frac{2\pi v}{3}\right) = \sin(\pi v). \quad (11)$$

The densities $U(z)$ and $\mu(z)$, on the other hand, behave better. They are both continuous across corners when $N_{\text{gr}} = 1$. The density $U(z)$ is a cusp. It has a leading behavior which can be described by a constant plus a term proportional to s^v . When $N_{\text{gr}} > 1$, $U(z)$ is still continuous across corners and triple-junctions while $\mu(z)$ has jump discontinuities. The density $\mu(z)$ is a primitive function of $\rho(z)$.

In conclusion, we can say that the integral equation (9) for $\mu(z)$ seems to be the most promising choice whenever $\sigma_0 < \sigma_1$ and $N_{gr} = 1$, and our objective is to compute q of (2) to high accuracy on a domain that contains corners. When $N_{gr} > 1$ the picture is more complex. But when $\sigma_0 < \sigma_k, k = 1, \dots, N_{gr}$, Eq. (9) still should give the most well-conditioned linear system.

3. An elastostatic crack problem

Let Γ be an oriented crack, a finite cut, in an infinite linearly elastic plane with stress

$$\lim_{z \rightarrow \infty} (\sigma_{xx}(z), \sigma_{yy}(z), \sigma_{xy}(z)) = (\sigma_{xx}^\infty, \sigma_{yy}^\infty, \sigma_{xy}^\infty) \tag{12}$$

applied at infinity. The crack starts at γ_s , ends at γ_e , is free of traction on both its faces, and allowed to respond to external loading with any finite crack opening displacement. The problem of determining the stress and displacement fields in this plane is a linear elastostatic problem. The stress state at infinity may be defined in terms of two constants α and β via $\sigma_{xx}^\infty + \sigma_{yy}^\infty = 2\alpha$ and $\sigma_{yy}^\infty - \sigma_{xx}^\infty + 2i\sigma_{xy}^\infty = 2\beta$.

The complex valued stress intensity factor $K = K_I + iK_{II}$ and the normalized stress intensity factor F at the crack tips γ_s and γ_e can be defined as [8,15]

$$K(\gamma_s) = \lim_{\epsilon \rightarrow 0^+} \sqrt{2\pi\epsilon} [\sigma_{y'y'}(\gamma_s - i\epsilon n(\gamma_s)) + i\sigma_{x'y'}(\gamma_s - i\epsilon n(\gamma_s))], \tag{13}$$

$$K(\gamma_e) = \lim_{\epsilon \rightarrow 0^+} \sqrt{2\pi\epsilon} [\sigma_{y'y'}(\gamma_e + i\epsilon n(\gamma_e)) + i\sigma_{x'y'}(\gamma_e + i\epsilon n(\gamma_e))], \tag{14}$$

$$F(\gamma_j) = \frac{K(\gamma_j)}{\sigma^\infty \sqrt{\pi a}}, \quad j = s, e, \tag{15}$$

where ϵ is a real number, x' and y' refer to local coordinate systems aligned with the crack at the crack tips, σ^∞ is related to the applied load, and a is a geometry dependent constant. Stress intensity factors can be computed once the elastostatic problem is solved. They are of interest in the study of fracture processes [9].

3.1. Two integral equations

Several second kind singular integral equations can be derived for the elastostatic crack problem using potential representations due to Muskhelishvili [17,24]. We pick two equations and only state results. For this, introduce the Cauchy singular operator M_1 whose action on a function $f(z)$ is given by

$$M_1 f(z) = \frac{1}{\pi i} \int_\Gamma \frac{f(\tau) d\tau}{\tau - z}, \quad z \in \Gamma \tag{16}$$

and the operators M_2 and M_3 , compact on smooth Γ , given by

$$M_2 f(z) = \frac{1}{\pi} \int_\Gamma f(\tau) \Im \left\{ \frac{d\tau}{\tau - z} \right\} + \frac{1}{2\pi i} \int_\Gamma \overline{f(\tau)} d \left[\frac{\tau - z}{\overline{\tau} - \overline{z}} \right], \quad z \in \Gamma \tag{17}$$

and

$$M_3 f(z) = \frac{1}{2\pi i} \left[\int_\Gamma \frac{f(\tau) d\tau}{\tau - z} + \frac{\bar{n}_z}{n_z} \int_\Gamma \frac{f(\tau) d\tau}{\overline{\tau} - \overline{z}} + \int_\Gamma \frac{\overline{f(\tau)} d\overline{\tau}}{\overline{\tau} - \overline{z}} + \frac{\bar{n}_z}{n_z} \int_\Gamma \frac{(\tau - z) \overline{f(\tau)} d\overline{\tau}}{(\overline{\tau} - \overline{z})^2} \right], \quad z \in \Gamma. \tag{18}$$

Let $\rho(z)$ be a weight given by

$$\rho(z) = ((z - \gamma_s)(z - \gamma_e))^{-\frac{1}{2}}, \tag{19}$$

whose value for $z \in \Gamma$ is defined as the limit from the right of the branch given by a branch cut along Γ and $z\rho(z) = 1$ at infinity [17]. The two Muskhelishvili integral equations then are

$$(I - M_3 \rho M_1 \rho^{-1}) \hat{\Omega}(z) = -\alpha + \beta \frac{\bar{n}_z}{n_z}, \quad z \in \Gamma, \tag{20}$$

$$(I - M_2 \rho^{-1} M_1 \rho) \hat{\omega}(z) = -\alpha z - \beta \bar{z}, \quad z \in \Gamma. \tag{21}$$

The stress or displacement field in the plane can be computed once $\hat{\Omega}(z)$ or $\hat{\omega}(z)$ is solved for on Γ . The normalized stress intensity factors follow from

$$F(\gamma_s) = \frac{i\sqrt{2}}{\sigma^\infty \sqrt{a}} \overline{M_1 \rho^{-1} \hat{\Omega}(\gamma_s)} \lim_{z \rightarrow \gamma_s} \overline{\rho(z)} \sqrt{\delta s(z)}, \quad z \in \Gamma, \tag{22}$$

$$F(\gamma_e) = -\frac{i\sqrt{2}}{\sigma^\infty \sqrt{a}} \overline{M_1 \rho^{-1} \hat{\Omega}(\gamma_e)} \lim_{z \rightarrow \gamma_e} \overline{\rho(z)} \sqrt{\delta s(z)}, \quad z \in \Gamma, \tag{23}$$

$$F(\gamma_j) = -\frac{n(\gamma_j) \overline{M_1 \rho \hat{\omega}(\gamma_j)}}{\sigma^\infty \sqrt{2a}} \lim_{z \rightarrow \gamma_j} (\overline{\rho(z)} \sqrt{\delta s(z)})^{-1}, \quad z \in \Gamma, \quad j = s, e, \tag{24}$$

where $\delta s(z)$ is arc length measured from the closest crack tip.

The integral equations (20) and (21) are both of Fredholm's second kind when Γ is smooth. Let us focus on their differences. One can show

$$\rho M_1 \rho^{-1} \hat{\Omega}(z) = \frac{d}{dz} (\rho^{-1} M_1 \rho \hat{\omega}(z)). \quad (25)$$

Assume that Γ contains a corner, called kink, with the asymptotic shape of a right angle and with vertex located at a point γ_k . Then (20) and (21) are not of Fredholm's second kind, but they are still solvable. Using variable separation one can show that close to γ_k the density $\hat{\Omega}(z)$ asymptotically behaves as $(z - \gamma_k)^{\lambda-1}$, where $0.5 < \lambda < 1$ is the solution with smallest positive non-zero real part to

$$\lambda = \sin\left(\frac{\lambda 3\pi}{2}\right). \quad (26)$$

The density $\hat{\omega}(z)$ is better behaved, as can be seen in (25). Its leading behavior can be described by a constant plus a term proportional to $(z - \gamma_k)^\lambda$.

4. Regularization and special-purpose interpolatory quadrature

The double-layer type kernels in (5), (8) and (9) and the related kernels of M_2 and M_3 in (20) and (21) are smooth on smooth Γ . Analytical limits exist for $\tau \rightarrow z$ and to use these limits in the discretization process is standard practice. Since also the layer densities are smooth, the Nyström method with composite Gauss–Legendre/Jacobi quadrature is very efficient. When Γ has corners, the kernels are not smooth even though analytical limits exist for $\tau \rightarrow z$ as long as z is not in a vertex. Neither are the layer densities smooth and Gaussian quadrature will not be accurate. The problem of simultaneously resolving the density and the kernel across a corner is difficult. One can use intense panel refinement or product integration with the densities represented in terms of tailor-made basis functions [7,15,17]. While tailor-made basis functions can yield accurate results with few discretization points, their construction relies on analytical methods which are hard to apply general situations. Therefore, we shall focus on panel refinement and try to make it as efficient as possible. The idea is to refine the mesh only as much as is required to resolve the density in terms of polynomials. The resolution of the kernel itself is taken care of by other means. Let us in the remainder of this section assume that the density is resolved.

A simple way to let the resolution of a double-layer kernel across a corner be controlled by the regularity of the layer density alone is by subtracting and adding the layer density, possibly multiplied with some smooth function, in the numerator of the kernel and then use a mix of analytic and numerical integration. This old technique, called *regularization* in this paper, is common for the computation of the Cauchy principal value of singular operators such as M_1 of (16). But it can be applied to all integral operators whose action on some function is known analytically. Some examples now follow:

The action of M_1 on the functions $\rho(z)$ and $1/\rho(z)$ of (19) on an open arc Γ can be evaluated analytically. For complex $f(z)$ belonging to a certain weighted L^2 space and by subtracting and adding $f(z)\rho(\tau)$ or $f(z)/\rho(\tau)$ in the numerator of the kernel one can show the regularizations [17]

$$M_1 \rho f(z) = \frac{1}{\pi i} \int_{\Gamma} \frac{(f(\tau) - f(z))\rho(\tau) d\tau}{\tau - z}, \quad z \in \Gamma, \quad (27)$$

$$M_1 \rho^{-1} f(z) = \frac{1}{\pi i} \int_{\Gamma} \frac{(f(\tau) - f(z)) d\tau}{\rho(\tau)(\tau - z)} + \left(z - \frac{\gamma_s + \gamma_e}{2}\right) f(z), \quad z \in \Gamma, \quad (28)$$

where the right-hand sides can be efficiently discretized with composite Gauss–Legendre/Jacobi quadrature. The action of the integral operator in (8) on the function $U(z) = 1$ can be evaluated analytically by rewriting it as a sum of the imaginary parts of $N_{gr} + 1$ different contour integrals of Cauchy type and with constant densities – one for each grain and one for the exterior domain – and use of the Cauchy integral formula. Subtraction and addition of $U(z)a(\tau)$ in the numerator of the kernel gives the regularization

$$\frac{1}{\pi} \int_{\Gamma} U(\tau) a(\tau) \Im \left\{ \frac{d\tau}{\tau - z} \right\} = (b(z) - 2\sigma_0) U(z) + \frac{1}{\pi} \int_{\Gamma} (U(\tau) - U(z)) a(\tau) \Im \left\{ \frac{d\tau}{\tau - z} \right\}, \quad z \in \Gamma. \quad (29)$$

The smoother $U(\tau)$ is across the corner, the better Gaussian quadrature should work. The limit of the kernel is not needed since $U(\tau) - U(z)$ vanishes for $\tau \rightarrow z$. It is also possible to regularize $M_2 \rho^{-1}$ by expressing it in terms of $M_1 \rho^{-1}$. However, even if the density is resolved across the corner, the kernel of the regularized $M_2 \rho^{-1}$ operator will contain components that vary a lot in their phase angles and Gaussian quadrature will not work well.

When regularization is not efficient, such as in the situation just described or when the layer density has a jump discontinuity across a corner, one can resort to more efficient *special-purpose interpolatory quadrature*. We refer the reader to Ref. [16], which gives a detailed account of how this recent technique can be applied to double-layer and Cauchy singular kernels.

We remark that when the mesh is refined close to a singular point γ_k , there will generally be a loss of precision in the relative placement of the discretization points due to finite precision arithmetic. This effect is more noticeable for the special-purpose interpolatory quadrature than for regularization. The former technique uses relative distances between discret-

ization points more heavily than the latter. Therefore, when special-purpose interpolatory quadrature [16] is used in our numerical examples, we create local coordinate systems by translating a region of the mesh around each γ_k so that γ_k is at the origin. Relative distances within these regions are then calculated using local coordinates.

5. Graded meshes

Assume that we want to perform piecewise $(m - 1)$ th degree polynomial interpolation $P_f(t)$ of a function $f(t)$ on the interval $t \in [0, 1]$ and that $f(t)$ behaves as t^ν , $\nu > 0$, close to $t = 0$. Assume, further, that we divide $[0, 1]$ into $s + 1$ subintervals, specified by s breakpoints t_j and that we place m interpolation points on each subinterval with the same relative spacing. How should the subintervals be placed as to minimize the interpolation error $\|f(t) - P_f(t)\|_\infty$? This approximation problem is treated in Chapter 4.2.5 of Atkinson [1]. One idea is to place the breakpoints according to

$$t_j = \left(\frac{j}{s+1}\right)^g, \quad j = 1, 2, \dots, s, \tag{30}$$

where g is a grading exponent which depends on m and ν . In the context of solving integral equations with fixed singularities at points γ_k , an otherwise uniform mesh which is refined according to (30) on each side of a γ_k is called a graded mesh. Theoretical estimates for the grading exponent, which should restore the convergence rate (of the solution itself – not of the iterative solver) to that of a smooth kernel, are given in Chapter 8 of Ref. [1].

In numerical experiments we found it difficult to achieve good results with the grading (30). Better results were obtained with the following straight forward procedure: Choose $\nu = 0.6$ irrespective of the true value of ν , which could be difficult to obtain in general situations anyhow. Choose $m = 16$ and place the nodes according to Gauss–Legendre quadrature. Take $f(t) = t^{0.6}$ and a value of s and place the breakpoints so that $\|f(t) - P_f(t)\|_\infty$ is minimized via a purely numerical procedure which computes the interpolation at 1000 equispaced points on each subinterval. An example: the breakpoints for $s = 4$ are $t_1 \approx 9.49 \times 10^{-7}$, $t_2 \approx 9.61 \times 10^{-5}$, $t_3 \approx 3.48 \times 10^{-3}$, $t_4 \approx 7.13 \times 10^{-2}$. We call an otherwise uniform mesh which is refined according to this procedure on each side of a γ_k an *aggressively graded mesh*.

A less economical, but much simpler, way to make a graded mesh is to use binary subdivision of the interval $[0, 1]$ and place the breakpoints according to

$$t_j = 2^{j-s-1}, \quad j = 1, 2, \dots, s. \tag{31}$$

We call an otherwise uniform mesh which is refined according to (31) on each side of a γ_k a *simply graded mesh*. An advantage with the simply graded mesh is that, if at least moderately high order quadrature is used, the refined quadrature panels on opposite sides of a γ_k often can be considered as lying far away from each other [16]. Special-purpose interpolatory quadrature may only be needed on the two refined panels closest to γ_k .

6. Direct compressed inverse preconditioning

Several authors, for example Martinsson and Rokhlin, have presented compressed factorizations of inverses to entire $N \times N$ system matrices arising from Nyström discretizations of integral equations which can be used as fast direct solvers under certain conditions [22]. A representation of the inverse is constructed at a cost proportional to N . The inverse is applicable to a right-hand side at a cost proportional to N . The word ‘compressed’ refers to data structures in this procedure. The leading idea is to exploit, in an hierarchical way, that the system matrix has low-rank off-diagonal blocks [22].

We shall construct compressed inverses to $N_s \times N_s$ submatrices of system matrices. These inverses act as right preconditioners and are chosen so that the regularity of the layer density is improved. So, in addition to improving the spectral properties of the system matrix, the inverse preconditioner also enables a polynomial representation of the layer density with much fewer terms. ‘Compressed’ here refers to this latter property. For the electrostatic problem we, too, shall take advantage of hierarchies of low-rank blocks and make the cost of construction of our compressed inverse proportional to N_s .

6.1. Splitting

Assume that we have a second kind integral equation

$$(I + K)\mu(z) = f(z), \quad z \in \Gamma, \tag{32}$$

where $f(z)$ is a piecewise smooth right-hand side, $\mu(z)$ is an unknown layer density, and K is a integral operator which chiefly is compact but contains a finite number N_{fs} of fixed singularities, or similar complications, at boundary location γ_k , $k = 1, \dots, N_{fs}$. Let Γ_k denote a small part of the boundary surrounding γ_k and let the union of all Γ_k be Γ_s . Let Γ_0 denote the remaining part of the boundary so that $\Gamma = \Gamma_0 \cup \Gamma_s$. Let Υ_k be a projection operator that is zero everywhere on Γ except for on Γ_k , where it is identity. Now split K into a non-compact part K^* and a remaining compact part K° so that

$$K = K^* + K^\circ, \quad K^* = \sum_{k=1}^{N_{fs}} \Upsilon_k K \Upsilon_k. \tag{33}$$

In many applications of interest $(I + K^\star)^{-1}$ is a bounded operator. The preconditioned equation

$$(I + (I + K^\star)^{-1}K^\circ)\mu(z) = (I + K^\star)^{-1}f(z) \quad (34)$$

is then a Fredholm second kind integral equation, see Chapter 8 of [1]. This is standard, but we shall not use it. We shall instead use right preconditioning

$$(I + K^\circ(I + K^\star)^{-1})\tilde{\mu}(z) = f(z), \quad (35)$$

where $\tilde{\mu}(z) = (I + K^\star)\mu(z)$ is a new unknown. An important difference between the densities $\mu(z)$ of (34) and $\tilde{\mu}(z)$ of (35) is that the latter is more likely to be a piecewise smooth quantity than the former. Smoothness of the solution can be used for compression of the inverse $(I + K^\star)^{-1}$.

6.2. Discretization on two meshes

Assume that the boundary is parameterized as $z(t)$, where t is a parameter. We shall solve (35), and for comparison also (32), using Nyström discretization based on m -point composite quadrature in t . For this, assume that we have a coarse mesh of quadrature panels on Γ generating a grid of discretization points and a partition $\Gamma = \Gamma_S \cup \Gamma_0$ such that the right-hand side $f(z(t))$ is resolved on Γ and such that the kernel of K° is resolved in both variables on Γ . This means that the action of K° on a function, smooth on Γ_0 and non-smooth on Γ_S , can be accurately computed using composite Gaussian quadrature on Γ_0 and suitable composite product integration on Γ_S and the result is a function that is resolved on Γ . Now only the inverse $(I + K^\star)^{-1}$ needs to be resolved on a mesh that is further refined, but we shall temporarily also further refine K° . We call a mesh that coincides with the coarse mesh on Γ_0 but is simply or aggressively graded on Γ_S a *partially refined* mesh. The degree of refinement is determined by that $\mu(z(t))$ should be resolved by Nyström discretization.

Let $\mathbf{I}_{\text{pr}}, \mathbf{K}_{\text{pr}}, \mathbf{K}_{\text{pr}}^\star$ and $\mathbf{K}_{\text{pr}}^\circ$ be matrices corresponding to a discretization of the operators I, K, K^\star and K° relying on composite Gaussian quadrature, where subscript ‘pr’ indicates the partially refined mesh. Let \mathbf{f}_{pr} be a column vector corresponding to a discretization of $f(z)$. Then (32) and (35) are discretized as

$$(\mathbf{I}_{\text{pr}} + \mathbf{K}_{\text{pr}})\boldsymbol{\mu}_{\text{pr}} = \mathbf{f}_{\text{pr}}, \quad (36)$$

$$(\mathbf{I}_{\text{pr}} + \mathbf{K}_{\text{pr}}^\circ(\mathbf{I}_{\text{pr}} + \mathbf{K}_{\text{pr}}^\star)^{-1})\tilde{\boldsymbol{\mu}}_{\text{pr}} = \mathbf{f}_{\text{pr}}. \quad (37)$$

We shall also need the corresponding discretizations \mathbf{I}_{coa} and $\mathbf{K}_{\text{coa}}^\circ$ on the coarse mesh, and the diagonal matrices \mathbf{W}_{pr} and \mathbf{W}_{coa} containing the Gaussian quadrature weights used in the two discretizations.

6.3. Prolongation and restriction

Let \mathbf{P} be a discretization of a prolongation operator that performs piecewise $(m - 1)$ th degree polynomial interpolation in t from the coarse mesh to the partially refined mesh. This means that if \mathbf{a} is a column vector corresponding to the discretization of a piecewise polynomial $a(t)$ of degree $m - 1$ on the coarse mesh, then \mathbf{Pa} is a column vector of discretized values of $a(t)$ on the partially refined mesh – a prolonged vector. The trivial relation $(\mathbf{Pa})^T = \mathbf{a}^T\mathbf{P}^T$, where superscript ‘T’ denotes transposition, indicates how row vectors are prolonged.

Let \mathbf{Q} be a discretization of a restriction operator that performs piecewise $(m - 1)$ th degree polynomial interpolation in the other direction, that is, from the partially refined mesh to the coarse mesh. The sparse block-matrices \mathbf{P} and \mathbf{Q} coincide with the identity matrix on Γ_0 and

$$\mathbf{QP} = \mathbf{I}_{\text{coa}}. \quad (38)$$

The assumptions that $f(z)$ is resolved and the kernel of K° is resolved in both its variables on the coarse mesh can be expressed as

$$\mathbf{f}_{\text{pr}} = \mathbf{P}\mathbf{f}_{\text{coa}}, \quad (39)$$

$$\mathbf{K}_{\text{pr}}^\circ\mathbf{W}_{\text{pr}}^{-1} = \mathbf{P}\mathbf{K}_{\text{coa}}^\circ\mathbf{W}_{\text{coa}}^{-1}\mathbf{P}^T. \quad (40)$$

6.4. Compression

Use of (39) and (40) in (37) gives

$$\tilde{\boldsymbol{\mu}}_{\text{pr}} + \mathbf{P}\mathbf{K}_{\text{coa}}^\circ\mathbf{W}_{\text{coa}}^{-1}\mathbf{P}^T\mathbf{W}_{\text{pr}}(\mathbf{I}_{\text{pr}} + \mathbf{K}_{\text{pr}}^\star)^{-1}\tilde{\boldsymbol{\mu}}_{\text{pr}} = \mathbf{P}\mathbf{f}_{\text{coa}}. \quad (41)$$

Multiplication of (41) with \mathbf{PQ} from the left and (38) imply

$$\tilde{\boldsymbol{\mu}}_{\text{pr}} = \mathbf{PQ}\tilde{\boldsymbol{\mu}}_{\text{pr}}. \quad (42)$$

We define

$$\tilde{\mu}_{\text{coa}} = \mathbf{Q}\tilde{\mu}_{\text{prf}}. \tag{43}$$

Multiplication of (41) with \mathbf{Q} from the left and use of (38), (42) and (43) transforms (41) into an equation on the coarse mesh

$$(\mathbf{I}_{\text{coa}} + \mathbf{K}_{\text{coa}}^{\circ}\mathbf{R})\tilde{\mu}_{\text{coa}} = \mathbf{f}_{\text{coa}}, \tag{44}$$

where the matrix

$$\mathbf{R} = \mathbf{W}_{\text{coa}}^{-1}\mathbf{P}^T\mathbf{W}_{\text{prf}}(\mathbf{I}_{\text{prf}} + \mathbf{K}_{\text{prf}}^{\star})^{-1}\mathbf{P} \tag{45}$$

can be viewed as a compressed right inverse preconditioner, that can be computed and stored once and for all. The matrix \mathbf{R} is the identity matrix modified with blocks corresponding to interaction from Γ_k to $\Gamma_k, k = 1, \dots, N_{\text{fs}}$, on the coarse mesh.

Note that, once (44) is solved, a discretized approximation to $\mu(z)$ on the coarse mesh can be recovered via

$$\mu_{\text{coa}} = \mathbf{S}\tilde{\mu}_{\text{coa}}, \tag{46}$$

$$\mathbf{S} = \mathbf{Q}(\mathbf{I}_{\text{prf}} + \mathbf{K}_{\text{prf}}^{\star})^{-1}\mathbf{P}. \tag{47}$$

Often, in applications, one is not interested in $\mu(z)$ itself. Rather, $\mu(z)$ is to be used in some integral. The matrices \mathbf{Q} and \mathbf{S} are then not needed. But $\mathbf{R}\tilde{\mu}_{\text{coa}}$ comes in handy as a discretized approximation of $\mu(z)$ multiplied with quadrature weight corrections appropriate for integration with piecewise smooth functions. For example, with N nodes t_k and weights w_k associated with integration in t on the coarse mesh, q of (10) can be obtained as

$$q = -\sigma_0 \sum_{k=1}^N (\mathbf{R}\tilde{\mu}_{\text{coa}})_k \Re\{\bar{e}z'(t_k)w_k\}. \tag{48}$$

A crude check if Γ_S is large enough so that K° is resolved is to examine the test vector \mathbf{t} given by

$$\mathbf{t} = \mathbf{K}_{\text{coa}}^{\circ}\mathbf{e}, \tag{49}$$

where \mathbf{e} is a vector of ones. If the coefficients $a_k, k = 0, \dots, m - 1$, of \mathbf{t} in an expansion in Legendre polynomials on each panel decay fast enough, then K° is resolved. We shall monitor the ratio $|a_{m-1}/a_0|$ on each panel and use the largest of these numbers as a measure of the resolution of K° .

7. Recursive compressed inverse preconditioning

The block-diagonal matrix \mathbf{R} of (45) could be expensive to compute as it stands. The computation of $\mathbf{K}_{\text{prf}}^{\star}$ may involve a fair amount of special-purpose interpolatory quadrature on an aggressively graded mesh and the cost of solving the system $(\mathbf{I}_{\text{prf}} + \mathbf{K}_{\text{prf}}^{\star})^{-1}\mathbf{P}$ grows as the number of subintervals on each Γ_k cubed. We shall now show how the blocks of \mathbf{R} may be computed in a much cheaper way, at a cost that grows only linearly in the number of subintervals. In fact, the cost will be so low so that the aggressively graded meshes are not needed. We can use simply graded meshes, which require less special-purpose interpolatory quadrature and no grading exponent whatsoever. For simplicity we first assume that there is only one fixed singularity, some corner at γ_1 , and that $\Gamma_S = \Gamma_1$ consists of two coarse panels – one on each side of γ_1 . See Fig. 1, upper left image, for an illustration with a corner in shape of a wedge. We use $m = 16$, that is, 16-point composite quadrature.

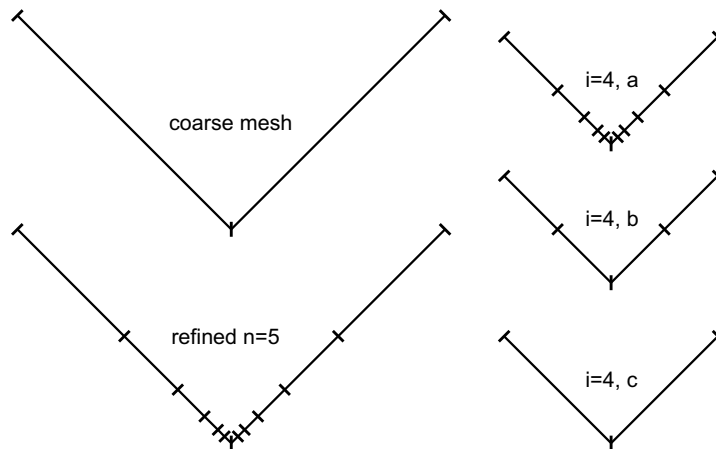


Fig. 1. Placement of panels on Γ_1 of a corner in the shape of a wedge. Upper left: Panels of the coarse mesh on Γ_1 . Lower left: Panels of a n -ply refined mesh on $\Gamma_1, n = 5$. Right: Panels on Γ_{ii} upon which the grids $\mathcal{G}_{ia}, \mathcal{G}_{ib}$, and \mathcal{G}_{ic} are constructed for $i = 4$ and $n = 5$. Note that \mathcal{G}_{ia} and \mathcal{G}_{ib} coincide for $i = 1$.

7.1. Grids and submatrices

Families of panels and grids on subsets of Γ_1 are now introduced. See the lower left and the right images of Fig. 1. Let Γ_1 be simply graded n times so that there are $2(n + 1)$ panels on Γ_1 . Let Γ_{i_1} denote the part of Γ_1 covered by the $2(i + 1)$ panels closest to γ_1 and let \mathcal{G}_{ia} denote a grid of $32(i + 1)$ quadrature points placed on these panels. Place four panels of equal length in parameter on Γ_{i_1} and let \mathcal{G}_{ib} denote a grid of 64 points placed on these panels. Place two panels of equal length in parameter on Γ_{i_1} and let \mathcal{G}_{ic} denote a grid of 32 points placed on these panels. Let \mathbf{P}_{iac} be the prolongation operator from \mathcal{G}_{ic} to \mathcal{G}_{ia} and let $\mathbf{P}_{iab}, \mathbf{P}_{ibc}$ be defined in a similar fashion. Actually, the index i in \mathbf{P}_{ibc} is not necessary. This operator always acts from a 32-point grid to a 64-point grid. Note also that the upper left and the lower right 16×16 blocks of \mathbf{P}_{iab} are identity matrices, $\mathbf{P}_{(i-1)ac}$ is a submatrix of \mathbf{P}_{iab} , and that $\mathbf{P}_{iac} = \mathbf{P}_{iab}\mathbf{P}_{bc}$.

We need to extract submatrices and the following convention will be used: Let \mathbf{A} be a $N_i \times N_i$ matrix corresponding to discretization of some operator on $\mathcal{G}_{ia}, \mathcal{G}_{ib}$, or \mathcal{G}_{ic} . With the split

$$\mathbf{A} = \mathbf{A}^\star + \mathbf{A}^\circ, \tag{50}$$

we mean that $\mathbf{A}_{jk}^\star = \mathbf{A}_{jk}, j, k = 17, \dots, N_i - 16$, and that the rest of the entries of \mathbf{A}^\star are zero. The symbol ‘ \star ’, thus, extracts elements in the center. The symbol ‘ \circ ’ extracts elements in a frame of width 16 around the center. The split (50) can be seen as a discrete analogy of (33) for a single corner. Now let \mathbf{K}_{ia} and \mathbf{K}_{ib} be discretizations of K^\star on \mathcal{G}_{ia} and \mathcal{G}_{ib} . Let $\mathbb{F}\{\cdot\}$ denote an operator which creates a frame of width 16 of zeros around its argument. Then, for example,

$$\mathbf{K}_{ia}^\star = \mathbb{F}\{\mathbf{K}_{(i-1)a}\}, \quad i = 2, \dots, n. \tag{51}$$

Let \mathbf{W}_{ia} and \mathbf{W}_{ib} be diagonal matrices with Gaussian weights from the discretizations on \mathcal{G}_{ia} and \mathcal{G}_{ib} as entries. We observe that

$$\mathbf{W}_{ib}^{-1} = 2^{(n-i)}\mathbf{W}_{nb}^{-1}. \tag{52}$$

7.2. The recursion

Define the 32×32 matrices

$$\mathbf{R}_i = 2^{(n-i)}\mathbf{P}_{iac}^\top \mathbf{W}_{ia} (\mathbf{I}_{ia} + \mathbf{K}_{ia})^{-1} \mathbf{P}_{iac}, \quad i = 1, \dots, n \tag{53}$$

and let \mathbf{W}_{nc} be a diagonal matrix of Gaussian weights for \mathcal{G}_{nc} . Then $\mathbf{W}_{nc}^{-1}\mathbf{R}_n$ is the block of \mathbf{R} of (45) corresponding to interaction from Γ_1 to Γ_1 . We now seek a relation between \mathbf{R}_i and \mathbf{R}_{i-1}

Eq. (53) can be rewritten using splitting of \mathbf{K}_{ia}

$$\mathbf{R}_i = 2^{(n-i)}\mathbf{P}_{iab}^\top \mathbf{P}_{iab}^\top ((\mathbf{W}_{ia} (\mathbf{I}_{ia} + \mathbf{K}_{ia}^\star)^{-1})^{-1} + \mathbf{K}_{ia}^\circ \mathbf{W}_{ia}^{-1})^{-1} \mathbf{P}_{iab} \mathbf{P}_{bc}. \tag{54}$$

The spectral radius of the matrix

$$\mathbf{K}_{ia}^\circ (\mathbf{I}_{ia} + \mathbf{K}_{ia}^\star)^{-1} \tag{55}$$

could in general depend on K , the corner opening angle, the choice of grading, and details in the split of \mathbf{K}_{ia} . In all our examples, however, this spectral radius is well below unity. Using an approximation similar to (40),

$$\mathbf{K}_{ia}^\circ \mathbf{W}_{ia}^{-1} = \mathbf{P}_{iab} \mathbf{K}_{ib}^\circ \mathbf{W}_{ib}^{-1} \mathbf{P}_{iab}^\top, \tag{56}$$

which holds if the operator corresponding to \mathbf{K}_{ib}° is resolved on \mathcal{G}_{ib} , and a Neumann series argument we can write (54) as

$$\mathbf{R}_i = 2^{(n-i)}\mathbf{P}_{bc}^\top ((\mathbf{P}_{iab}^\top \mathbf{W}_{ia} (\mathbf{I}_{ia} + \mathbf{K}_{ia}^\star)^{-1} \mathbf{P}_{iab})^{-1} + \mathbf{K}_{ib}^\circ \mathbf{W}_{ib}^{-1})^{-1} \mathbf{P}_{bc} \tag{57}$$

and use (51)–(53) and the structure of \mathbf{P}_{iab} to get the recursion

$$\mathbf{R}_i = \mathbf{P}_{bc}^\top (2\mathbb{F}\{\mathbf{R}_{i-1}^{-1}\} + (\mathbf{I}_{nb}^\circ + \mathbf{K}_{ib}^\circ) \mathbf{W}_{nb}^{-1})^{-1} \mathbf{P}_{bc}, \quad i = 1, \dots, n, \tag{58}$$

where

$$2\mathbb{F}\{\mathbf{R}_0^{-1}\} = (\mathbf{I}_{1b} + \mathbf{K}_{1b}^\star) (\mathbf{W}_{nb}^{-1})^\star \tag{59}$$

is used to start the recursion. We point out that each recursion step in (58) corresponds to one binary subdivision of a simply graded mesh.

The recursion (58) and (59) holds for a corner where the ‘arms’ may have curvature. For a corner in the shape of a wedge, that is, when the ‘arms’ are straight lines, scale invariance gives $\mathbf{K}_{ib}^\circ = \mathbf{K}_{nb}^\circ$ and $\mathbf{K}_{ib}^\star = \mathbf{K}_{nb}^\star$, and (58), (59) simplify to the fixed-point iteration

$$\mathbf{R}_i = \mathbf{P}_{bc}^\top (2\mathbb{F}\{\mathbf{R}_{i-1}^{-1}\} + (\mathbf{I}_{nb}^\circ + \mathbf{K}_{nb}^\circ) \mathbf{W}_{nb}^{-1})^{-1} \mathbf{P}_{bc}, \quad i = 1, \dots, \tag{60}$$

$$2\mathbb{F}\{\mathbf{R}_0^{-1}\} = (\mathbf{I}_{nb} + \mathbf{K}_{nb}^\star) (\mathbf{W}_{nb}^{-1})^\star. \tag{61}$$

A minor speedup could be achieved by using the Schur–Banachiewicz inverse formula [18] for (60).

8. Direct compressed inverse preconditioning for the crack problem

We choose (21). For brevity we include the weight ρ in the definition of M_1 and the weight ρ^{-1} in the definition of M_2 . Operator splitting in fixed singular and compact parts gives $M_2 = M_2^* + M_2^\circ$. We also split $M_1 = M_1^* + M_1^\circ$ in a similar fashion where M_1^* is larger than M_2^* . Let Γ_{S_1} be the part of Γ where M_1^* is non-zero and let Γ_{S_2} be the (smaller) part where M_2^* is non-zero. We demand that M_1° is resolved in the variable of integration on Γ and in the other variable on $\Gamma_0 \cup \Gamma_{S_2}$ in the sense of Section 6.2. Right preconditioning of (20) in the style of (35) results in

$$\left(I - (M_2^\circ M_1^\circ + M_2^\circ M_1^* + M_2^* M_1^\circ) (I - M_2^* M_1^*)^{-1} \right) \tilde{\omega}(z) = -\alpha z - \bar{\beta} \bar{z}. \tag{62}$$

While all three composed operators in the first closed parenthesis on the left-hand side are compact, only the two first produce smooth output. The operator $M_2^* M_1^\circ$ produces non-smooth output so $\tilde{\omega}(z)$ cannot be a piecewise smooth function. If we, however, proceed to split

$$\tilde{\omega}(z) = \tilde{\omega}_1(z) + M_2^* \tilde{\omega}_2(z), \tag{63}$$

where $\tilde{\omega}_2(z)$ is defined as zero except for on Γ_{S_2} , we can solve the system

$$\tilde{\omega}_1(z) - M_2^\circ (M_1^\circ + M_1^*) (I - M_2^* M_1^*)^{-1} (\tilde{\omega}_1 + M_2^* \tilde{\omega}_2)(z) = -\alpha z - \bar{\beta} \bar{z}, \quad z \in \Gamma, \tag{64}$$

$$\tilde{\omega}_2(z) - M_1^\circ (I - M_2^* M_1^*)^{-1} (\tilde{\omega}_1 + M_2^* \tilde{\omega}_2)(z) = 0, \quad z \in \Gamma_{S_2} \tag{65}$$

and get two smooth solutions $\tilde{\omega}_1(z)$ and $\tilde{\omega}_2(z)$.

Discretization of (64) and (65) in the style of (37) gives a system where four quantities per corner need to be compressed:

$$\begin{aligned} & \mathbf{W}(\mathbf{I} - \mathbf{M}_2^* \mathbf{M}_1^*)^{-1}, \quad \mathbf{W} \mathbf{M}_1^* (\mathbf{I} - \mathbf{M}_2^* \mathbf{M}_1^*)^{-1}, \\ & \mathbf{W}(\mathbf{I} - \mathbf{M}_2^* \mathbf{M}_1^*)^{-1} \mathbf{M}_2^*, \quad \text{and} \quad \mathbf{W} \mathbf{M}_1^* (\mathbf{I} - \mathbf{M}_2^* \mathbf{M}_1^*)^{-1} \mathbf{M}_2^* \end{aligned} \tag{66}$$

and where we have omitted the subscript ‘pr’. Direct compression as in (45) is straight-forward, while recursive compression is harder to implement.

9. Plain recursive compression

The compressed preconditioned equation (44) can, via (46), be written as a compressed un-preconditioned equation for the original variable μ_{coa}

$$(\mathbf{I}_{\text{coa}} + (\mathbf{S}^{-1} - \mathbf{I}_{\text{coa}}) + \mathbf{K}_{\text{coa}}^\circ \mathbf{R} \mathbf{S}^{-1}) \mu_{\text{coa}} = \mathbf{f}_{\text{coa}}. \tag{67}$$

Here the columns of the blocks of \mathbf{S} can be viewed as special basis functions, the entries of $\mathbf{R} \mathbf{S}^{-1}$ correspond to quadrature weight corrections, and $\mathbf{S}^{-1} - \mathbf{I}_{\text{coa}}$ corresponds to a discretization of K^* . While clumsier than (44) the compressed un-preconditioned equation (67) might have advantages in complicated settings where K is not a just compact operator but a sum of a compact operator and an operator whose spectrum one does not wish to alter.

In order to make (67) efficient we need a recursive scheme for \mathbf{S} . For this define, in analogy with (53),

$$\mathbf{S}_i = \mathbf{Q}_{ica} (\mathbf{I}_{ia} + \mathbf{K}_{ia})^{-1} \mathbf{P}_{iac}, \quad i = 1, \dots, n, \tag{68}$$

where \mathbf{S}_n is the block of \mathbf{S} of (47) corresponding to interaction from Γ_1 to Γ_1 . Use of splitting, (56), and a Neumann series argument gives

$$\mathbf{S}_i = \mathbf{Q}_{cb} \mathbf{Q}_{iba} (\mathbf{I}_{ia} + \mathbf{K}_{ia}^*)^{-1} \mathbf{P}_{iab} \left(\mathbf{P}_{iab}^\top \mathbf{W}_{ia} (\mathbf{I}_{ia} + \mathbf{K}_{ia}^*)^{-1} \mathbf{P}_{iab} \right)^{-1} \left(\left(\mathbf{P}_{iab}^\top \mathbf{W}_{ia} (\mathbf{I}_{ia} + \mathbf{K}_{ia}^*)^{-1} \mathbf{P}_{iab} \right)^{-1} + \mathbf{K}_{ib}^\circ \mathbf{W}_{ib}^{-1} \right)^{-1} \mathbf{P}_{bc}, \tag{69}$$

which can be written

$$\mathbf{S}_i = \mathbf{Q}_{cb} \left(2\mathbb{F}\{\mathbf{S}_{i-1}\} \mathbb{F}\{\mathbf{R}_{i-1}^{-1}\} + (\mathbf{W}_{nb}^{-1})^\circ \right) \left(2\mathbb{F}\{\mathbf{R}_{i-1}^{-1}\} + (\mathbf{I}_{nb}^\circ + \mathbf{K}_{ib}^\circ) \mathbf{W}_{nb}^{-1} \right)^{-1} \mathbf{P}_{bc}, \quad i = 1, \dots, n \tag{70}$$

and together with

$$\mathbb{F}\{\mathbf{S}_0\} = (\mathbf{I}_{1b} + \mathbf{K}_{1b}^*)^{-1} \tag{71}$$

be used in tandem with (58) and (59) as a recursion for \mathbf{S}_n . For a corner in the shape of a wedge, \mathbf{K}_{ib}° and \mathbf{K}_{1b}^* can be replaced with \mathbf{K}_{nb}° and \mathbf{K}_{nb}^* .

10. Small-scale numerical examples

We first demonstrate the performance of our techniques for a few small setups, each exemplifying some interesting geometric feature. The goal is to compute the quantity q of (2) or F of (15) as accurately and cheaply as possible using the GMRES iterative solver [25] with a low-threshold stagnation avoiding technique [16] for the main linear system. The code is implemented in MATLAB and executed on a SunBlade 100 workstation. System matrices are formed explicitly. Naturally, with this mix of built-in and coded MATLAB software and outdated hardware, timings will not be indicative of how fast things really can go. We still give timings to demonstrate the speed gains of the methods proposed. But bear in mind, that memory requirements in GMRES (Krylov subspace vectors plus the diagonal block matrix \mathbf{R}), rate of convergence, and achievable accuracy are the important properties that are carried over to more elaborate implementations and larger problems. As reference solutions for q and F we shall use estimated 'best values' from convergence studies on a large number of meshes.

10.1. The Triangle Problem

We solve an electrostatic problem. Fig. 2 depicts an equilateral triangle with side length two and $\sigma = 1000$ placed, away from the origin, in a surrounding medium with $\sigma = 1$ and with electric field $e = i$. The opening angles $\pi/3$ suggest quite strong self-interaction of the layer densities across the corners. According to the analysis of Section 2, the densities $U(z)$ of (8) and $\mu(z)$ of (9) are continuous while $\rho(z)$ of (5) has singularities and (9) should give the most well-conditioned system. We construct a coarse mesh with p panels of equal length on each triangle side. Experiments show that $p = 4$ is sufficient to resolve all densities away from the corners. We shall use partially refined meshes that are either simply graded, 'sig', or aggressively graded, 'agg'. The 'sig' meshes are constructed by repeated binary subdivision of the two coarse panels closest to each corner. Regularization 'reg' can be applied to all three equations (5), (8) and (9). Regularization for (8) is given by (29). The same formula holds for (9) since $a(\tau)$ and $b(z)$ are constant in this example, while (5) can be regularized using $\rho(\tau)n_z\bar{n}_\tau = \rho(\tau)n_z\bar{n}_\tau - \rho(z) + \rho(z)$.

Fig. 3 compares results obtained with the unpreconditioned formulation (36) for (5), (8) and (9) using 'sig' with and without 'reg'. Regularization has a major advantage over the common use of analytical limits in terms of rate of convergence of q for (8), as predicted, and also for (5), which is unexpected and may be explained by the high symmetry of the problem. Regularization also allows for higher achievable accuracy. The most accurate results are produced by (9). It outperforms (5) and (8) with a factor equal to the conductivity ratio. The left image of Fig. 4, where also special-purpose interpolatory quadrature 'spq' [16] is used, shows that 'agg' save a factor of about three in terms of discretization points compared to 'sig'. The standard equation (5) is particularly poor. For one thing, it needs about twice the number of GMRES iterations as do the other two equations, see the right image of Fig. 4.

Clearly, Eq. (9) is the most efficient in this case. With 35 binary subdivisions, that is, 3552 discretization points on the 'sig' mesh, and 39 GMRES iterations the estimated relative error in q is 5×10^{-15} . This corresponds to one of the '* Eq. (9) sig+reg' experiments in Fig. 3. The computing time is around 35 s. With 33 GMRES iterations and 1152 discretization points on the 'agg' mesh the estimated relative error in q is 2×10^{-15} . This corresponds to one of the ' Δ Eq. (9) agg+spq' experiments in Fig. 4. The computing time is around 5 s.




Fig. 2. The Triangle Problem.

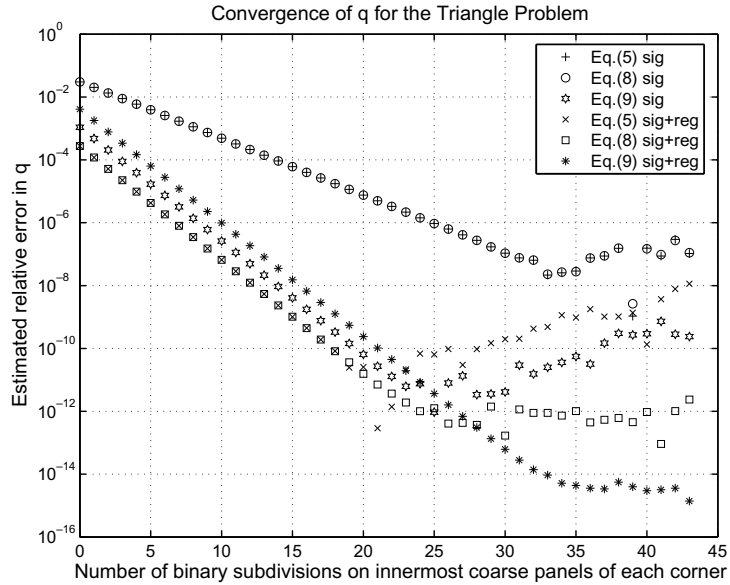


Fig. 3. Results for the Triangle Problem with a simply graded mesh. The reference solution is taken as $q_{ref} = 4.455790323135370$.

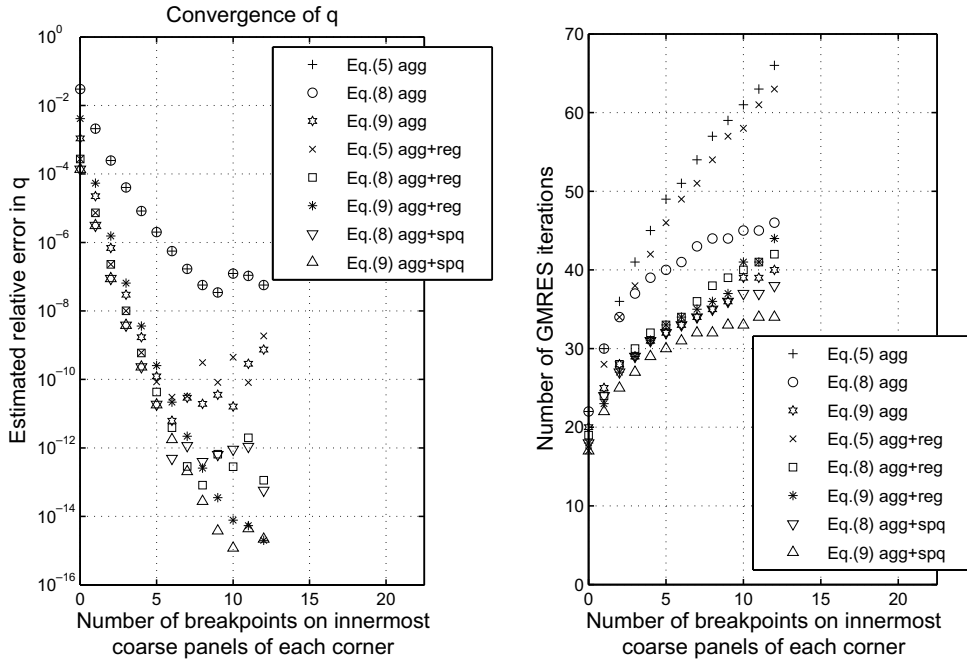


Fig. 4. Results for the Triangle Problem with an aggressively graded mesh. See Section 5 for the definition of breakpoints. The right image shows the number of GMRES iterations needed to reach a stopping criterion threshold of ϵ_{mach} in the relative residual.

We now apply direct compressed inverse preconditioning (44) and (45) to (9) in combination with ‘agg+spq’. We first need to determine Γ_S . If we let Γ_S include only the part of the boundary covered by the six coarse panels upon which grading is performed, the resolution of K° , see (49), is only 10^{-9} . The relative error in q is about 10^{-12} . Instead we let Γ_S cover twelve coarse panels, two on each side of each corner. Now the resolution of K° is 10^{-14} . With ten breakpoints on the coarse panels closest to each corner, only nine GMRES iterations, and 192 discretization points on the coarse mesh we solve (44) and (45) and get an estimated relative error in q of 4×10^{-15} . The memory requirements in GMRES are reduced by 90% compared to (36) with ‘sig+reg’ and by 65% compared to (36) with ‘agg+spq’. The computing time is down to around two and a half seconds.

Recursive compressed inverse preconditioning (44) and (60) can, slightly modified, also be applied to (9). Since Γ_S now is larger than assumed in Section 7, the definitions of Γ_{ki} and \mathcal{G}_{ia} , \mathcal{G}_{ib} , and \mathcal{G}_{ic} for each corner must be changed: Γ_{ki} is now the part of the boundary covered by the $2(i+2)$ panels closest to γ_k , each \mathcal{G}_{ia} is a grid of $32(i+2)$ points, each \mathcal{G}_{ib} is a grid of 96 points constructed by placing six panels on Γ_{ki} and they are not of equal size – two are twice as long as the remaining four, each \mathcal{G}_{ic} is a grid of 64 points constructed by placing four panels of equal size on Γ_{ki} . With 35 steps in (60) and 'spq' in (61) we produce results at least as good as to those obtained with (44) and (45). The computing time is down to less than 1 s.

We finally apply plain recursive compression (67), (60) and (70) to (9). The value of q , with 35 steps in (60) and (70) and 'spq' in (61) and (71), is almost identical to that produced by (44) and (60). The number of GMRES iterations is 17, which corresponds to the number required on a coarse mesh using only 'spq'. See Fig. 4, right image and zero breakpoints. One can say that plain recursive compression retains the conditioning and the number of unknowns of the original equation discretized on a coarse mesh, while giving the accuracy of a discretization on a partially refined mesh. The computing time is around one and a half second.

10.2. The Segmented Circle Problem

Fig. 5 depicts an inclusion with unit circular boundary consisting of three equisized segments with different conductivities, centered off-origin, in a surrounding medium. The opening angles of the corners are larger than in the Triangle Problem, but some boundary parts are curved and there is less symmetry. The density $U(z)$ is continuous and makes (8) a candidate for regularization 'reg'. The density $\mu(z)$ of (9) has jump discontinuities and should benefit from special-purpose interpolatory quadrature 'spq' [16].

We construct a coarse mesh with p panels of equal length on each straight boundary part and $2p$ panels of equal length on each curved boundary part. Experiments show that $p=2$ is sufficient to resolve all densities away from the corners. Fig. 6 compares results obtained with (36) on partially refined meshes that are simply graded, 'sig', with and without 'reg' and 'spq'. The expected positive impact of 'reg' for (8) and 'spq' for (9) in terms of convergence are clearly seen, while (5), with its singular density, is not improved at all. Fig. 7 shows that aggressive grading 'agg' saves a factor of about three in terms of memory requirements. A difference compared to the Triangle Problem occurs for (8): 'spq' is now an even better choice than 'reg' both in terms of convergence, see Fig. 6, and in terms of the number of GMRES iterations needed, see the right image of Fig. 7.

Again, Eq. (9) is the most efficient. It outperforms (5) and (8) in terms of achievable accuracy with a factor equal to the highest conductivity ratio. With 14 binary subdivisions, that is, 2976 discretization points on the 'sig' mesh, and 21 GMRES iterations the estimated error in q is 10^{-15} . This corresponds to one of the ' Δ Eq. (9) sig+spq' experiments in Fig. 6. The computing time is around 25 s. With 21 iterations and 1284 discretization points on the 'agg' mesh we get a value of q which coincides with the reference solution. This corresponds to one of the ' Δ Eq. (9) agg+spq' experiments in Fig. 7. The computing time is around 5 s.

We now apply direct compressed inverse preconditioning (44) and (45) to (9) in combination with 'agg+spq'. Γ_S includes the part of the boundary covered by the coarse panels upon which 'agg' is performed. With five breakpoints on the coarse panels closest to each triple-junction, only 14 GMRES iterations, and 288 discretization points on the coarse mesh we solve




Fig. 5. The Segmented Circle Problem.

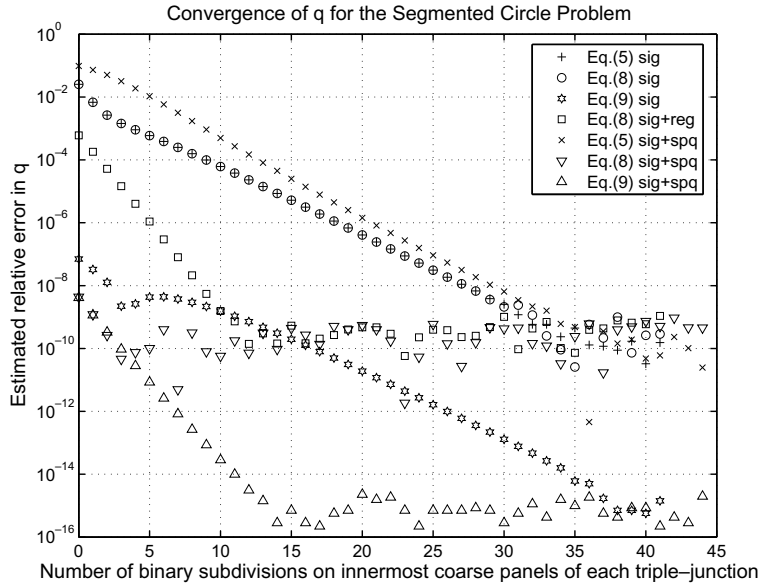


Fig. 6. Results for the Segmented Circle Problem with a simply graded mesh. The reference solution is taken as $q_{ref} = 6.25303717710058$.

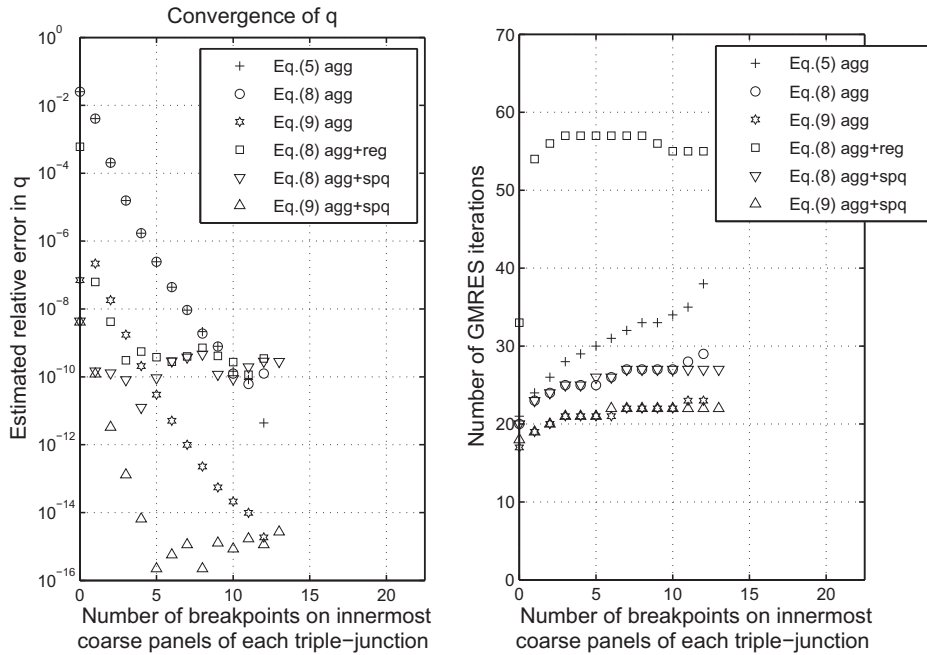


Fig. 7. Results for the Segmented Circle Problem with an aggressively graded mesh.

(44) and (45) and again get a value of q which coincides with the reference solution. The memory requirements in GMRES are reduced by 80% compared to (36) with 'sig+spq' and by 50% compared to (36) with 'agg+spq'. The computing time is down to around two and a half seconds. Recursive compressed inverse preconditioning (44) and (58) can, slightly modified since we have triple-junctions rather than corners, also be applied. With $n = 14$ steps in (58) and 'spq' in (59) we produce almost identical results to those obtained with (44) and (45). The computing time is down to less than 2 s.

We finally apply plain recursive compression (67), (58) and (70) to (9). The result, with $n = 14$ steps in (58) and (70) and 'spq' in (59) and (71), is almost identical to that produced by (44) and (58). The number of GMRES iterations is 18, which corresponds to the number of iterations required on a coarse mesh using only 'spq'. See Fig. 7, right image and zero breakpoints.

10.3. The two circle problem

It is not necessary that the points γ_k and the intervals Γ_k , in the derivation of (44), correspond to actual singularities in K . Compressed inverse preconditioning is also applicable when Γ_k corresponds to a region that in a more general sense is difficult, such as when the boundary falls back on itself or when disjoint boundary parts lie close to each other. Consider the Two Circle Problem of Fig. 8. We use (9) with ‘spq’, a coarse mesh with 32 panels, and eight binary subdivisions on the four coarse panels lying closest to the origin. We compare the performance of (36) to that of direct compressed inverse preconditioning (44) and (45) and with Γ_s being the eight coarse panels closest to the origin. Both methods produce q with an estimated relative error of about 10^{-14} . But the number of GMRES iterations needed for full convergence decrease from 161 to nine, the memory requirements in GMRES are reduced by 85%, and the execution time is reduced by 75% as we switch from (36)–(44). In this particular example, however, the hybrid method of images for circles developed by Cheng and Greengard [2] is even more efficient.

10.4. The V-problem

Fig. 9 depicts an elastic plane with a crack in the shape of a V turned upside down. The computation of the normalized stress intensity factor $F(\gamma_s)$ of (15) in this setup with $\sigma^\infty = \sigma_{yy}^\infty$ is a benchmark problem treated by several authors [4,7,8,28]. The problem is well-conditioned. The most accurate result seems to be $F(\gamma_s) = 0.5207675522 + 0.6411159455i$ from [7].

A coarse mesh with p panels on each branch of the crack is constructed. Experiments show that $p = 3$ is more than enough to resolve the densities in (20) and (21) away from the kink. Fig. 10 compares results for simply and aggressively graded meshes ‘sig’ and ‘agg’ with and without ‘spq’ for the M_2 and M_3 operators. The M_1 operator is discretized via (27) and (28). Note that without ‘spq’, Eq. (20), whose density is diverging, produces slightly better results for F than (21), whose density is continuous. With ‘spq’ the situation is reversed. There is a large positive impact in the convergence rate for (21), while (20) is not affected and therefore not shown. The ‘spq’ also bounds the number of GMRES iterations needed to meet the stopping criterion threshold for (21) as the mesh is refined, see Fig. 11. With 28 binary subdivisions in (21), that is, 992 points on the ‘sig’ mesh, ‘spq’, and 21 iterations, the estimated relative error in $F(\gamma_s)$ is 10^{-14} . The computing time is 16 s. A similar estimated error is achieved with 480 points on the ‘agg’ mesh, ‘spq’, and 22 iterations. The computing time is 6 s.

We now test inverse preconditioning (64) and (65) with direct compression of (66) on the ‘agg’ mesh. Γ_{s2} is the part of the boundary covered by the two coarse panels closest to the kink and Γ_{s1} is the part of the boundary covered by the four coarse panels closest to the kink. With 13 iterations and $96 + 32 = 128$ discretization points on the coarse mesh we get an estimated relative error in $F(\gamma_s)$ of 10^{-13} . The computing time is 11 s. The memory requirements in GMRES are unchanged compared to the use of (21) on the ‘sig’ mesh and twice as large compared to the use of (21) on the ‘agg’ mesh. So in this particular example there is not much to gain by using compression. Compression could, however, be profitable speed-wise as well as memory-wise for multiple right-hand sides and for large problems with relatively few kinks.




Fig. 8. Two conducting circles with radii $R = 0.499999$ separated by a distance of $\delta = 0.000002$ in a conducting plane. The panel placement on the coarse mesh is indicated. The reference solution is taken as $q_{\text{ref}} = 5.1613733230816$.




Fig. 9. The V-problem.

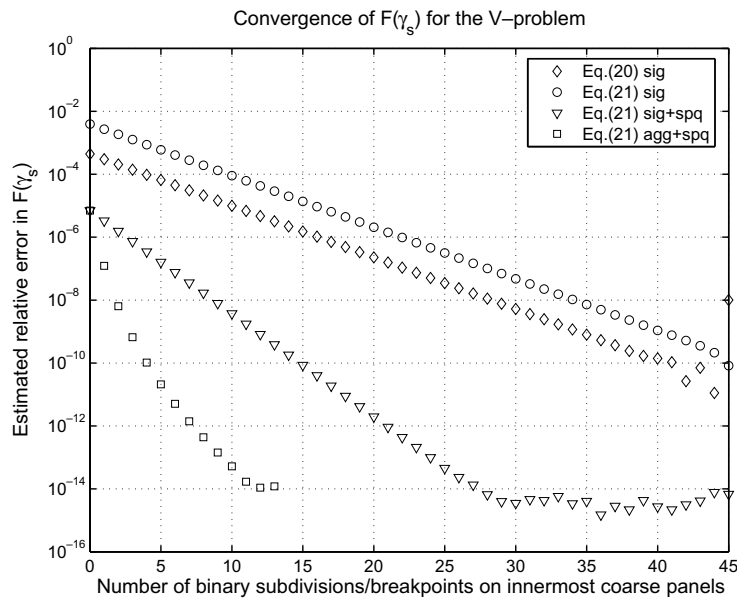


Fig. 10. Results for the V-problem. The reference solution is taken as $F(\gamma_s)_{\text{ref}} = 0.520767552218259 + 0.641115945496149i$.

11. Large-scale numerical examples

We conclude the experiments with some large-scale electrostatic computations for inclusions built up of perturbed hexagons. The overall goal is to compute q of (2) via (5), (8) and (9) as accurately and rapidly as possible. We use the recursive compressed inverse preconditioning (44) and (60) in all corners and triple-junctions. Special-purpose interpolatory quadrature [16] is used in (61) and local coordinate systems consist of one coarse panel from each edge emanating from a vertex. Linear systems are solved using GMRES as in Section 10, but the computer now is an Intel Core2 6400 at 2.13 GHz. The code is mainly written in MATLAB, with time-critical parts written in C, such as the fast multipole method [13] with precision $\epsilon = 10^{-13}$ used for matrix-vector multiplication.

A number N_{gr} of regular hexagonal grains are placed in a honeycomb-like pattern in the unit square. The structure is randomized slightly as to avoid symmetries that may allow for a simpler solution. This is done by moving the grain vertices a distance of one-tenth of the length of a grain side in a random direction. The largest structure used, with $N_{\text{gr}} = 5293$, is dis-

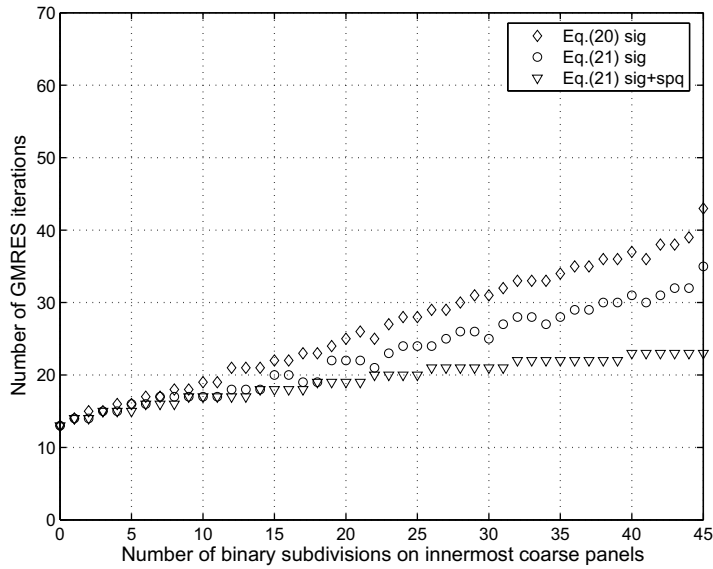


Fig. 11. GMRES iterations for the V-problem.

played in Fig. 12. The conductivity σ_k of an individual grain is given by 10^{c_k} , where c_k is a random variable. Two cases are studied, differing in the range of the conductivities assigned. In the first case, c_k is uniformly distributed in $[-1, 1]$. In the second case, c_k is multiplied by a factor of three, highlighting the impact of higher conductivity ratios, while keeping the structure the same as in the first case in all other respects. The conductivity of the background material is always given by $\sigma_0 = 1$ and the electric field is $e = i$ in both cases. We implement (5) and (9) as they stand. For (8) we introduce

$$U^*(z) = U(z)a(z), \quad z \in \Gamma \tag{72}$$

and solve for $U^*(z)$ in

$$U^*(z) - \frac{\lambda(z)}{\pi} \int_{\Gamma} U^*(\tau) \Im \left\{ \frac{d\tau}{\tau - z} \right\} = 2\sigma_0 \lambda(z) \Re \{ \bar{e}z \}, \quad z \in \Gamma. \tag{73}$$

The formula (2) then reads

$$q = \int_{\Gamma} U^*(z) \Im \{ \bar{e} dz \}. \tag{74}$$

The reason for choosing (73), which from a numerical viewpoint corresponds to a diagonal similarity transformation of (8), is that (8), despite having a continuous unknown, has some negative characteristics when combined with the recursion formula (60). These problems are alleviated as we switch to (73). To check for convergence in (60), we use the condition that convergence has occurred at step i if

$$\| \mathbf{R}_i - \mathbf{R}_{i-1} \|_F / \| \mathbf{R}_{i-1} \|_F < \epsilon_{\text{mach}}, \tag{75}$$

where subscript ‘F’ denotes Frobenius norm. We construct a coarse mesh on Γ by placing p panels of equal length in parameter as well as in arc length on the shortest grain edge in the structure. On the other grain edges panels are placed so that all the panels in the mesh have approximately the same length. That is, each edge in the structure consists of at least p panels. Each panel, in turn, has 16 Gauss–Legendre nodes with accompanying suitably scaled weights.

We first vary the minimum number of panels p on a structure with $N_{\text{gr}} = 390$ and $\sigma_k \in [10^{-1}, 10^1]$. We calculate q via the three formulations for $p = 2, 3, \dots, 20$. As a reference solution we take the arithmetic mean of the 12 values of q computed for $p = 4, 5, 6, 7$. Throughout this section we shall compute estimated absolute errors, rather than relative errors. The reason is that as N_{gr} grows, the grained inclusion will, in a macroscopic sense, be close to a homogenized inclusion with conductivity one. This, in turn, causes q to be small in magnitude. The estimated relative errors may oscillate as some structures will, by chance, have q close to zero. One could interpret the absolute error in q as a relative error with respect to the geometric mean of all conductivities involved. Fig. 13 shows that the proposed methods are stable for all three integral equation formulations under rather serious overresolution and that the geometry is resolved already at $p = 3$.

Having investigated resolution in p for a problem with a fixed number N_{gr} of grains we now turn to setups where we vary N_{gr} while holding p fixed. We use $p = 3$. A reference solution q_{ref} for each structure is computed as the arithmetic mean of the q obtained using all three integral equation formulations with $p = 4$. Figs. 14 and 15 show results for $\sigma_k \in [10^{-1}, 10^1]$. The error grows slowly with N_{gr} , settling at around 10^{-14} for $N_{\text{gr}} = 5293$. The number of iterations required to reach a relative




Fig. 12. A perturbed honeycomb structure consisting of 5293 grains. The conductivity varies between $\sigma = 0.1$ (light) and $\sigma = 10$ (dark). The background material has conductivity $\sigma_0 = 1$. There are 10,878 corners and triple-junctions in the structure.

residual threshold of 10^{-16} seems to converge to around 105. The values of q_{ref} , not shown, range from about -0.9 to 0.5 and slowly converge with N_{gr} to about 0.04 ± 0.04 .

Figs. 14 and 15 do not disclose any significant difference between the formulations (5), (9) and (73) for the present conductivity distribution, neither in terms of achievable accuracy nor in terms of iterations needed. However, the computations of q via (9) and via (73) are somewhat faster than the computation via (5). The formulation (5), with diverging density, requires between 40 and 70 steps in recursion (60) to reach the fixed point. The formulations (9) and (73), with bounded densities, only require between 20 and 40 steps. Furthermore, some extra complex arithmetic is needed in (5) compared to (9) and (73). As for memory requirements, considering the largest setup with $N_{gr} = 5293$, the number of discretization points is

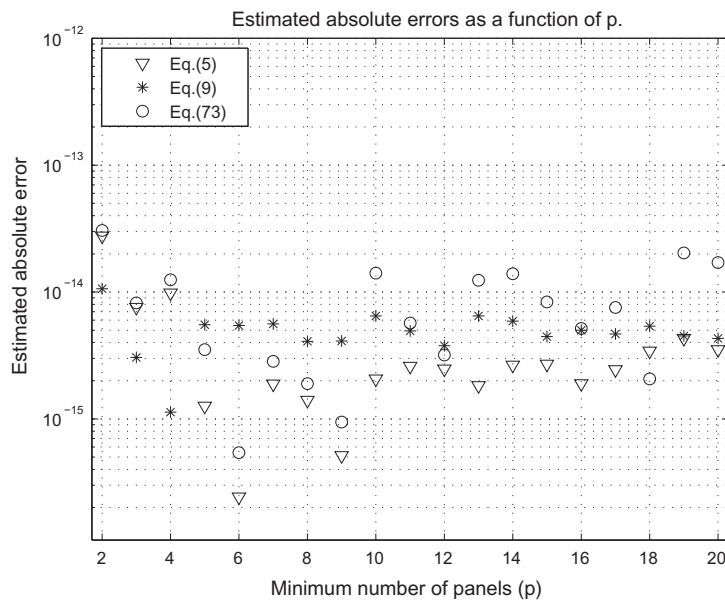


Fig. 13. Estimated absolute errors for q of (2) using the integral equation formulations (5), (9) and (73) with varying degree of overresolution. The inclusion consists of 390 grains. The reference value is taken as $q_{ref} = -0.0326509925268292$.

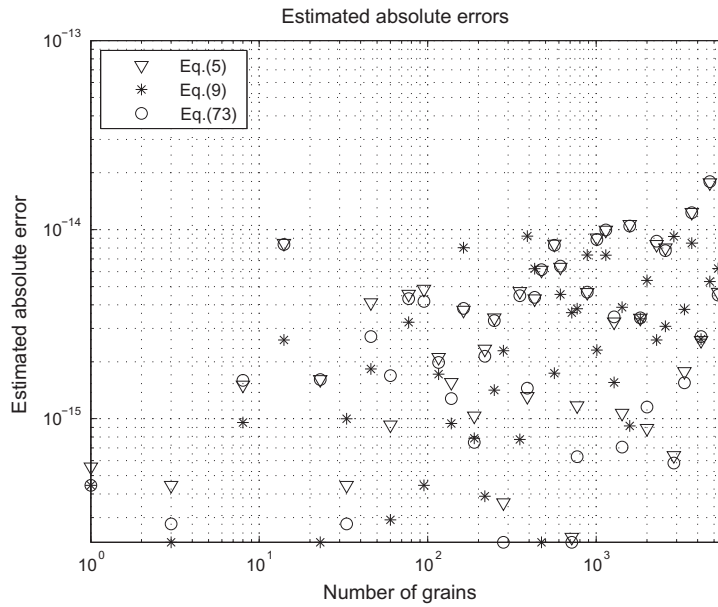


Fig. 14. Estimated absolute errors when computing q of (2) for honeycomb structures using the integral equations (5), (9) and (73). The grain conductivities vary between 0.1 and 10.

1141904. At 109 iterations this requires 1 Gb of memory for the storage of the Krylov subspace vectors. The 10,878 block-matrices in the recursive compressed inverse preconditioner, one for each corner and triple-junction, need also be stored and require about 200 Mb of RAM. A further 100 Mb is used by various vectors pertaining to discretization points, conductivity, et cetera, giving a total of about 1.3 Gb of RAM required for the computation. The time required to compute q is dominated by GMRES. For example, computing q on the structure depicted in Fig. 12 using (9), (10) takes 75 minutes total, of which 9% is due to the recursive compressed inverse preconditioning.

We now turn to the second case with $\sigma_k \in [10^{-3}, 10^3]$. Stronger inhomogeneity makes the problem considerably harder and more ill-conditioned. An error plot, similar to Fig. 14 but not shown, indicates that our methods still are equally stable for the three formulations (5), (9) and (73), although two additional digits of accuracy are lost. Fig. 16, analogous to Fig. 15, shows that the number of GMRES iterations keeps climbing without any tendency to have approached an upper limit at




Fig. 15. GMRES iterations for Fig. 14. The stopping criterion threshold in the relative residual is set at 10^{-16} .




Fig. 16. Same as in Fig. 15, but the grain conductivities vary between 10^{-3} and 10^3 .

$N_{\text{gr}} = 885$, which is the largest setup that can be handled due to memory constraints. This phenomenon is well known and sometimes called ‘critical slowing down’ in the applied physics literature, where numerical experiments on large structures often are done on discrete lattice networks [6,20]. Some long-range preconditioning would help, in addition to the local preconditioning supplied by our recursive compressed inverse preconditioner.

The efficiency of (5), (8) and (9) for the inclusion problem depends on the conductivity distribution, see Section 2. In the small-scale experiments of Section 10 we used distributions intended to favor (9) in order to demonstrate this point. In the present section we use conductivity distributions for the grains with a geometric mean of approximately σ_0 , hoping that neither of the formulations should be obviously favored. And indeed, no major differences in efficiency were discovered. This is remarkable, considering the different regularities of the involved densities and only possible thanks to the recursive compressed inverse preconditioner. Equipped with this tool, one is more free to choose a formulation which suits a given problem from a modeling point of view.

12. Conclusions

The literature on solving elliptic boundary value problems using integral equation methods is rich on ‘standard practice’ accumulated from successful experience with setups involving smooth and well-separated boundaries. This paper shows that when it comes to problems on domains with corners, standard practice on fundamental issues can be much improved upon. At least if high accuracy in combination with speed is of concern. We give a series of electrostatic and elastostatic examples in familiar two-dimensional settings. General actions, such as choosing an alternative integral equation formulation more fit for the problem at hand, using special-purpose interpolatory quadrature, and a mesh whose grading is determined in a novel manner, are shown to improve on the order of convergence and lead to large savings in storage and computing time.

The use of compressed inverse preconditioners, free from special basis functions and grading exponents, can improve performance even further and also revive integral equation formulations which normally have to be discarded due to excessive demands for refinement. To demonstrate the potential of this new technique, we present large-scale electrostatic examples on idealized, yet challenging, granular materials with more than 5000 grains. In applications, of course, other boundary conditions than ours may arise and the presence of body forces may cause complications. Furthermore, corner opening angles can be smaller than those of our perturbed hexagons. An automated approach to determine how many quadrature panels should be included in Γ_k , the part of the boundary surrounding a singularity γ_k on which the compressed preconditioner acts, based on rigorous proofs, is left for future work.

It would be nice to include more comparison with previous work. Unfortunately, papers on fast and highly accurate computations on domains with a large number of corners and triple-junctions are scarce. One example, however, is Ref. [26], on stress driven grain boundary diffusion in two-dimensional domains containing up to twelve grains without cracks. Here the most difficult numerical aspect involves solving an elastostatic equation repeatedly using a least squares finite element method incorporating special basis functions and it is extremely important that the singularities in the stress fields near grain boundary junctions are properly resolved.

References

- [1] K.E. Atkinson, *The numerical Solution of Integral Equations of the Second Kind*, Cambridge University Press, Cambridge, 1997.
- [2] H. Cheng, L. Greengard, On the numerical evaluation of electrostatic fields in dense random dispersions of cylinders, *J. Comput. Phys.* 136 (2) (1997) 629–639, doi:10.1006/jcph.1997.5787.
- [3] H. Cheng, V. Rokhlin, N. Yarvin, Nonlinear optimization, quadrature, and interpolation, *SIAM J. Optim.* 9 (4) (1999) 901–923, doi:10.1137/S1052623498349796.
- [4] Y.Z. Chen, Stress intensity factors for curved and kinked cracks in plane extension, *Theor. Appl. Fract. Mech.* 31 (3) (1999) 223–232, doi:10.1016/S0167-844(99)00016-6.
- [5] R.V. Craster, Y.V. Obnosov, Checkerboard composites with separated phases, *J. Math. Phys.* 42 (11) (2001) 5379–5388, doi:10.1063/1.1398336.
- [6] R.G. Edwards, J. Goodman, A.D. Sokal, Multigrid method for the random-resistor problem, *Phys. Rev. Lett.* 61 (12) (1988) 1333–1335, doi:10.1103/PhysRevLett.61.1333.
- [7] J. Englund, Large scale computations for cracks with corners, in: R. Gallego, M.H. Aliabadi (Eds.), *Advances in Boundary Element Techniques*, BeTeQ IV, Queen Mary University of London, 2003, pp. 71–76.
- [8] J. Englund, Fast accurate and stable algorithm for the stress field around a zig-zag-shaped crack, *Eng. Fract. Mech.* 70 (2) (2003) 355–364, doi:10.1016/S0013-7944(02)00040-1.
- [9] J. Englund, A higher order scheme for two-dimensional quasi-static crack growth simulations, *Comput. Methods Appl. Mech. Eng.* 196 (21–24) (2007) 2527–2538, doi:10.1016/j.cma.2007.01.007.
- [10] L.G. Fel, I.V. Kaganov, Relation between effective conductivity and susceptibility of two-component rhombic checkerboard, *J. Phys. A* 36 (19) (2003) 5349–5358, doi:10.1088/0305-4470/36/19/311.
- [11] L. Greengard, J.-Y. Lee, Electrostatics and heat conduction in high contrast composite materials, *J. Comput. Phys.* 211 (1) (2006) 64–76, doi:10.1016/j.jcp.2005.05.004.
- [12] L. Greengard, M. Moura, *On the Numerical Evaluation of Electrostatic Fields in Composite Materials*, Acta Numerica 1994, Cambridge University Press, Cambridge, 1994, pp. 379–410.
- [13] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348, doi:10.1016/0021-999(87)90140-9.
- [14] J. Helsing, A. Jonsson, On the computation of stress fields on polygonal domains with V-notches, *Int. J. Num. Meth. Eng.* 53 (2) (2002) 433–453, doi:10.1002/nme.291.
- [15] J. Helsing, A. Jonsson, A seventh order accurate and stable algorithm for the computation of stress inside cracked rectangular domains, *Int. J. Multiscale Comput. Eng.* 2 (1) (2004) 47–68, doi:10.1615/IntJMultCompEng.v2.i1.
- [16] J. Helsing, R. Ojala, On the evaluation of layer potentials close to their sources, *J. Comput. Phys.* 227 (5) (2008) 2899–2921, doi:10.1016/j.jcp.2007.11.024.
- [17] J. Helsing, G. Peters, Integral equation methods and numerical solutions of crack and inclusion problems in planar elastostatics, *SIAM J. Appl. Math.* 59 (3) (1999) 965–982, doi:10.1137/S0036139998332938.
- [18] H.V. Henderson, S.R. Searle, On deriving the inverse of a sum of matrices, *SIAM Rev.* 23 (1) (1981) 53–60.
- [19] P. Kolm, S.D. Jiang, V. Rokhlin, Quadruple and octuple layer potentials in two dimensions I: analytical apparatus, *Appl. Comput. Harmon. Anal.* 14 (1) (2003) 47–74, doi:10.1016/S1063-520(03)00004-6.
- [20] P. Kumar, V.V. Nukala, S. Šimunović, M.N. Guddati, An efficient algorithm for modelling progressive damage accumulation in disordered materials, *Int. J. Num. Meth. Eng.* 62 (14) (2005) 1982–2008, doi:10.1002/nme.1257.
- [21] R. Kress, I.H. Sloan, F. Stenger, A Sinc quadrature method for the double-layer integral equation in planar domains with corners, *J. Integr. Eqn. Appl.* 10 (3) (1998) 291–317.
- [22] P.G. Martinsson, V. Rokhlin, A fast direct solver for boundary integral equations in two dimensions, *J. Comput. Phys.* 205 (1) (2005) 1–23, doi:10.1016/j.jcp.2004.10.033.
- [23] A. Mayo, A. Greenbaum, Fourth order accurate evaluation of integrals in potential theory on exterior 3D regions, *J. Comput. Phys.* 220 (2) (2007) 900–914, doi:10.1016/j.jcp.2006.05.042.
- [24] N.I. Muskhelishvili, *Some Basic Problems of the Mathematical Theory of Elasticity*, P. Noordhoff Ltd., Groningen, 1953.
- [25] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comp.* 7 (3) (1986) 856–869, doi:10.1137/0907058.
- [26] J. A. Sethian, J. Wilkening, A numerical model of stress driven grain boundary diffusion, *J. Comput. Phys.* 193 (1) (2004) 275–305, doi:10.1016/j.jcp.2003.08.015.
- [27] E. Tuncer, Y.V. Serdyuk, S.M. Gubanski, Dielectric mixtures: electrical properties and modeling, *IEEE Trans. Dielect. Electr. Insul.* 9 (5) (2002) 809–828, doi:10.1109/TDEI.2002.1038664.
- [28] A.K. Yavuz, S.L. Phoenix, S.C. TerMaath, An accurate and fast analysis for strongly interacting multiple crack configurations including kinked (V) and branched (Y) cracks, *Int. J. Solids Struct.* 43 (22–23) (2006) 6727–6750, doi:10.1016/j.ijsolstr.2006.02.005.
- [29] L. Ying, G. Biros, D. Zorin, A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains, *J. Comput. Phys.* 219 (1) (2006) 247–275, doi:10.1016/j.jcp.2006.03.021.