Journal of Computational Physics 256 (2014) 214–219

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Short note

A simplified technique for the efficient and highly accurate discretization of boundary integral equations in 2D on domains with corners

A. Gillman [a], S. Hao [b], P.G. Martinsson [b,*]

[a] Department of Mathematics, Dartmouth College, Hanover, NH 03755-3551, United States
[b] Department of Applied Mathematics, Univ. of Colorado Boulder, Boulder, CO 80309-0526, United States

ARTICLE INFO

Article history:
Received 29 April 2013
Received in revised form 19 August 2013
Accepted 23 August 2013
Available online 11 September 2013

Keywords:
Boundary integral equation
Nyström discretization
High-order discretization
Corner singularity
Weakly singular integral equation
Quadrature
Grid refinement
Direct solver
Acoustic scattering

1. Introduction

This note comments on some recently developed techniques for computing an approximate solution to a Boundary Integral Equation (BIE) like

$$\alpha q(\boldsymbol{x}) + \int_{\Gamma} K(\boldsymbol{x}, \boldsymbol{y}) q(\boldsymbol{y})\, ds(\boldsymbol{y}) = f(\boldsymbol{x}), \quad \boldsymbol{x} \in \Gamma, \tag{1}$$

where Γ is a piecewise smooth contour in the plane, and where K is one of the standard kernels of potential theory such as, e.g., the single or double layer kernels associated with the Laplace or Helmholtz equations. A challenge in solving (1) is that its integrand exhibits complicated singular behavior near the corner points of Γ. A classical technique for dealing with this difficulty has been to expand the unknown q near the corner using specialized basis functions that incorporate analytical knowledge about the singularity [2]. Recently, however, a remarkable observation has been made [1,10,11,9] that there exist general purpose techniques that do not require any analytical à priori knowledge other than that the integrand of (1) be absolutely integrable.

* Corresponding author.
E-mail addresses: adrianna.gillman@dartmouth.edu (A. Gillman), sijia.hao@colorado.edu (S. Hao), martinss@colorado.edu (P.G. Martinsson).

0021-9991/$ – see front matter
http://dx.doi.org/10.1016/j.jcp.2013.08.049

In a nutshell, the idea of [1,10,11,9] is to use a standard Nyström discretization of (1) designed for a smooth contour. The discretization should use a panel based (i.e. non-global) quadrature rule such as, e.g., a composite Gaussian rule. Then simply refine the computational mesh near any corner. For any given computational tolerance ε (setting $\varepsilon = 10^{-10}$ or smaller is often entirely manageable), continue the refinement until the contribution from any panels directly touching a corner is bounded by ε (this is possible since the integrand in (1) is absolutely integrable). Then simply omit the two panels nearest to the corner from the discretization. Observe that on any remaining panel, the function q is smooth enough to be accurately represented by the interpolant implied by the chosen quadrature rule.

The apparent drawback of a simplistic refinement process like the one described is that it can dramatically increase the number of degrees of freedom required in the Nyström discretization. A key insight of [1,10,11,9] is that the "superfluous" degrees of freedom added by the refinement can be eliminated from the linear system via a *strictly local* process. Moreover, this local process can be executed in time that scales linearly with the number of degrees of freedom added. The end result is a linear system discretizing (1) that has about as many degrees of freedom as one would have needed had the corner not been present in the first place. (For the case of regular polygonal domains, the compression can even be performed in sublinear time [9].)

The task of "squeezing out" the degrees of freedom added by the local refinement near the corner is in [1,10,11,9] executed via purpose-built local compression techniques that can be somewhat challenging to implement. The purpose of this note is to demonstrate that this compression step can be executed via the general purpose direct solvers described in [5,7,12,14].

2. A linear algebraic observation

The compression technique that allows us to eliminate the superfluous degrees of freedom is based on the observation that certain off-diagonal blocks of the coefficient matrix resulting from the discretization of (1) have low numerical rank. Critically, the important ranks do not depend on how many degrees of freedom are used in the refinement near the corner. To illustrate how such rank-deficiencies can be exploited, consider in general the task of solving the linear system

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{q}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix}, \tag{2}$$

where $\mathbf{A}_{11}$ is of size $n_1 \times n_1$ and $\mathbf{A}_{22}$ is of size $n_2 \times n_2$. Now assume that $\mathbf{A}_{12}$ and $\mathbf{A}_{21}$ each are of rank k. Think of n_1 as a large number (e.g. the number of degrees of freedom used in the refinement of the corner, say $n_1 \sim 10^3$), and k as a small number (often in the range 20–50). Then $\mathbf{A}_{12}$ and $\mathbf{A}_{21}$ admit factorizations

$$\underset{n_1 \times n_2}{\mathbf{A}_{12}} = \underset{n_1 \times k}{\mathbf{U}_1} \; \underset{k \times n_2}{\mathbf{B}_{12}} \quad \text{and} \quad \underset{n_2 \times n_1}{\mathbf{A}_{21}} = \underset{n_2 \times k}{\mathbf{B}_{21}} \; \underset{k \times n_1}{\mathbf{V}_1^*}, \tag{3}$$

where $\mathbf{U}_1$ and $\mathbf{V}_1$ are well-conditioned matrices. We further assume that the data vector $\mathbf{f}_1$ belongs to the same k-dimensional space as the columns of $\mathbf{A}_{12}$ (if it does not, then the space can be extended as needed),

$$\mathbf{f}_1 = \mathbf{U}_1 \tilde{\mathbf{f}}_1. \tag{4}$$

When (3) and (4) hold, the linear system (2) with $n_1 + n_2$ unknowns is in a certain sense equivalent to the smaller system

$$\begin{bmatrix} \mathbf{D}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{q}}_1 \\ \mathbf{q}_2 \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{f}}_1 \\ \mathbf{f}_2 \end{bmatrix} \tag{5}$$

with only $k + n_2$ unknowns. In (5), $\mathbf{D}_{11}$ and $\tilde{\mathbf{q}}_1$ are defined by

$$\mathbf{D}_{11} = \left(\mathbf{V}_1^* \mathbf{A}_{11}^{-1} \mathbf{U}_1\right)^{-1} \quad \text{and} \quad \tilde{\mathbf{q}}_1 = \mathbf{V}_1^* \mathbf{q}_1. \tag{6}$$

When we say that (2) and (5) are "equivalent" we mean that the solution $\{\mathbf{q}_1, \mathbf{q}_2\}$ of the larger system (2) can be obtained from the solution $\{\tilde{\mathbf{q}}_1, \mathbf{q}_2\}$ of the smaller system (5) via the formula

$$\mathbf{q}_1 = \mathbf{A}_{11}^{-1} \mathbf{U}_1 \mathbf{D}_{11} \tilde{\mathbf{q}}_1. \tag{7}$$

To be precise, the equivalence holds when $\mathbf{A}_{11}$ and $\mathbf{V}_1^* \mathbf{A}_{11}^{-1} \mathbf{U}_1$ are both non-singular.

3. Matrix skeletons

For the low-rank factorizations (3), it is convenient to use a so-called *interpolative decomposition* (ID) [4] in which $\mathbf{B}_{12}$ is a $k \times n_2$ matrix consisting of k rows of $\mathbf{A}_{12}$ and $\mathbf{B}_{21}$ is an $n_2 \times k$ matrix consisting of k columns of $\mathbf{A}_{21}$. The matrices $\mathbf{U}_1$ and $\mathbf{V}_1$ each hold a $k \times k$ identity matrix as a submatrix, and have no entries whose magnitude exceeds 1.

The advantage of using an ID is that the matrices $\mathbf{A}_{12}$ and $\mathbf{A}_{21}$ *need never be formed*. Instead, a local computation determines the index vectors pointing out which columns and rows are needed, and then only those entries need to be computed