Interpolated Compressed Inverse Preconditioning: Fast and Accurate Simulation of Close-to-Touching Discs in Stokes Flow

Daniel Fortunato, Mariana Martínez Aguilar, Dhairya Malhotra

Abstract

We consider the flow of dense suspensions of rigid bodies in a Stokesian fluid. Such flows are difficult to compute numerically due to the presence of close-to-touching interactions, which may require a large number of unknowns to resolve sharply peaked surface forces, a large number of GMRES iterations to solve the discretized PDE, and an extremely small time step. A common way of dealing with these difficulties is to introduce a repulsion force between particles to prevent them from getting too close. However, this additional repulsion force is non-physical and may fundamentally alter the results of a simulation.

For suspensions of identical discs in 2D, we present a fast and accurate boundary integral method that mitigates these challenges without introducing artificial forces. Through precomputation, compression and interpolation of the close-to-touching part of the interaction operator, our method—termed *interpolated compressed inverse preconditioning*—efficiently handles close-to-touching interactions down to distances of 10^{-10} with only a coarse discretization of the boundary. Additionally, we present a preconditioner that significantly reduces the number of GMRES iterations required to solve the Stokes mobility problem at each time step by effectively reusing the Krylov subspace from previous time steps. Coupled with high-order, adaptive time-stepping using spectral deferred correction, we are able to take larger time steps, mitigating the temporal stiffness resulting from close-to-touching interactions.

For a graphical description of this work, see: https://danfortunato.com/talks/ICIP-poster.pdf.

1 Stokes mobility problem

We consider N_Ω rigid discs $\Omega = \{\Omega_1, \cdots, \Omega_{N_\Omega}\}$ embedded in a Stokesian fluid. The fluid velocity in the exterior of Ω is governed by the Stokes equations,

$$-\Delta \boldsymbol{u} + \nabla p = 0 \qquad \text{in } \mathbb{R}^2 \setminus \Omega, \tag{1}$$

$$\nabla \cdot \boldsymbol{u} = 0 \qquad \text{in } \mathbb{R}^2 \setminus \Omega, \tag{2}$$

where $\boldsymbol{u}$ is the fluid velocity and p is the fluid pressure. Equations (1) and (2) denote the momentum balance and incompressibility constraints, respectively. In addition, we also assume that the fluid velocity at infinity decays to zero,

$$\boldsymbol{u}(\boldsymbol{x}) \to 0 \qquad \text{as } |\boldsymbol{x}| \to \infty.$$

Each disc has a net force $\boldsymbol{F}_k$ and a net torque T_k acting about a point $\boldsymbol{x}_k^c$. The discs undergo rigid body motion with the velocity $\boldsymbol{V}$ given by,

$$\boldsymbol{V}(x) = \boldsymbol{v}_k + \omega_k(\boldsymbol{x} - \boldsymbol{x}_k^c)^\perp \qquad \text{for all } \boldsymbol{x} \in \Omega_k,$$

where $\boldsymbol{v}_k$ is the translational velocity and ω_k is the angular velocity of Ω_k about the point $\boldsymbol{x}_k^c$. A slip velocity boundary condition $\boldsymbol{u}_s$ between the rigid bodies and the fluid is prescribed. Therefore, the fluid velocity on the boundary $\partial\Omega$ is given by,

$$\boldsymbol{u} = \boldsymbol{V} + \boldsymbol{u}_s \qquad \text{on } \partial\Omega.$$

In the mobility problem, we are given $\boldsymbol{u}_s$, $\boldsymbol{F}_k$, and T_k about $\boldsymbol{x}_k^c$ for each Ω_k. The rigid body motion $\boldsymbol{V}$ (i.e., $\boldsymbol{v}_k$ and ω_k for each Ω_k) is not known and must be determined.

Using the Stokes single- and double-layer potentials to represent the fluid velocity $\boldsymbol{u}$ in terms of an unknown surface density $\boldsymbol{\sigma}$, a boundary integral equation (BIE) for the Stokes mobility problem can be formulated as given in [1]:

$$\mathcal{K}\boldsymbol{\sigma} = g \qquad \text{on } \partial\Omega, \tag{3}$$

where $\mathcal{K}$ is a second-kind boundary integral operator and g encodes the given slip velocity, net force, and net torque.

2 Close-to-touching interactions

Consider the model problem of two discs separated by a distance d, with each disc discretized into a set of high-order Gauss–Legendre panels. The two disc problem serves as an effective pairwise preconditioner in a simulation with many close-touching discs. When the distance d between two discs gets small, the solution σ to the BIE in (3) becomes highly peaked. This requires an extremely fine discretization of the boundary in the close-to-touching region. We label the close-to-touching region as Γ_2 and the remaining boundary as $\Gamma_1 = \partial\Omega \setminus \Gamma_2$. Then, (3) can be discretized as a block linear system,

$$\begin{pmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{pmatrix} \begin{pmatrix} \sigma_1 \\ \sigma_2 \end{pmatrix} = \begin{pmatrix} g_1 \\ g_2 \end{pmatrix}, \tag{4}$$

where g_1 and σ_1 are the boundary conditions and unknowns on Γ_1, g_2 and σ_2 are the boundary conditions and unknowns on Γ_2, and K_{ij} represents a sub-block of the discretized operator $\mathcal{K}$ that computes interactions from sources on Γ_j to targets on Γ_i. Right preconditioning (4) with the block diagonal preconditioner $\left(\begin{smallmatrix} I & 0 \\ 0 & K_{22}^{-1} \end{smallmatrix}\right)$ yields the system

$$\begin{pmatrix} K_{11} & K_{12}K_{22}^{-1} \\ K_{21} & I \end{pmatrix} \begin{pmatrix} \sigma_1 \\ \overline{\sigma}_2 \end{pmatrix} = \begin{pmatrix} g_1 \\ g_2 \end{pmatrix}, \tag{5}$$

where $\overline{\sigma}_2 = K_{22}\sigma_2$ is a new unknown on Γ_2. While (5) may require fewer GMRES iterations to solve than (4), it still requires an excessively fine discretization in the close-to-touching region Γ_2. Additionally, computing K_{22}^{-1} on the fly can be expensive, especially for problems with moving boundaries. However, one may show that $\overline{\sigma}_2$ can be discretized on a coarse mesh and that the off-diagonal block $K_{12}K_{22}^{-1}$ is low rank.

2.1 Compressing close-to-touching interactions

Since Γ_1 and Γ_2 are disjoint, the discretized boundary integral operators K_{12} and K_{21} are low rank, with the numerical rank independent of the distance d. Hence, the column space of K_{21} is comprised of smooth functions that can be discretized using piecewise polynomials on a coarse mesh. From (5), we have $\overline{\sigma}_2 = g_2 - K_{21}\sigma_1$; therefore, $\overline{\sigma}_2$ is smooth whenever g_2 is smooth and it can be discretized on a coarse mesh. Since K_{12} is low rank, so is $K_{12}K_{22}^{-1}$ (with numerical rank independent of d). There are several ways of constructing a compressed representation for $K_{12}K_{22}^{-1}$. To retain the boundary integral structure and allow for acceleration by the fast multipole method (FMM), we use a representation of the form $K_{12}R$ where R is a low-rank operator such that

$K_{12}R \approx K_{12}K_{22}^{-1}$, up to a given numerical tolerance. We now describe the numerical construction of R.

Consider two different panelizations of Γ_2: a fine mesh where the panels on each disc are refined dyadically towards the closest point between the discs, and a coarse mesh with a small number of uniformly sized panels on each disc. We denote quantities on the coarse mesh with a superscript "c"; all other quantities are assumed to live on the fine mesh. Define the prolongation operator P that interpolates data from the coarse mesh to the fine mesh, and diagonal matrices W_f and W_c containing the weights for smooth integration on the fine and coarse meshes, respectively. Then, $W_c^{-1}P^TW_f$ computes an L^2 projection from the fine mesh to the coarse mesh.

Assuming that the boundary data g_2 is smooth and therefore representable on the coarse mesh, we have

$$g_2 = P\, g_2^c, \tag{6}$$
$$\overline{\sigma}_2 = P\, \overline{\sigma}_2^c. \tag{7}$$

Since K_{12} and K_{21} are low rank, they can be approximated accurately by their coarse discretizations,

$$K_{12} = K_{12}^c W_c^{-1} P^T W_f, \tag{8}$$
$$K_{21} = P K_{21}^c, \tag{9}$$

Substituting (6)–(9) in (5), we obtain

$$\begin{pmatrix} K_{11} & K_{12}^c R \\ K_{21}^c & I \end{pmatrix} \begin{pmatrix} \sigma_1 \\ \overline{\sigma}_2^c \end{pmatrix} = \begin{pmatrix} g_1 \\ g_2^c \end{pmatrix} \tag{10}$$

where $R = W_c^{-1}P^TW_fK_{22}^{-1}P$. This definition of R is used in the RCIP (recursively compressed inverse preconditioning) method [2]. For an order-p fine mesh with $\mathcal{O}(\log d)$ levels of refinement, direct construction of R takes $\mathcal{O}(p^3(\log d)^3)$ operations; the RCIP method provides a faster algorithm to construct R without directly computing K_{22}^{-1}, taking $\mathcal{O}(p^3 \log d)$ operations. Our main result—termed ICIP (interpolated compressed inverse preconditioning)—instead constructs R through precomputation and interpolation, requiring only $\mathcal{O}(p^2)$ work.

2.2 Interpolated compressed inverse preconditioning (ICIP)

Constructing $R = R(d)$ each time for a different value of d can be expensive since it requires computing K_{22}^{-1}. Instead, we construct a polynomial interpolant for $R(d)$ over a range $d \in [d_{\min}, d_{\max}]$ (where $0 < d_{\min} < d_{\max}$). Then for any value of d in the interval, we construct $R(d)$ through entrywise interpolation. We use Chebyshev polynomials in $\log d$ as our interpolation basis, i.e.,

$$[R(d)]_{ij} \approx \sum_{k=0}^{q} [R_k]_{ij} T_k(\log d)$$

where $[R_k]_{ij}$ is the kth Chebyshev coefficient for the ijth entry of $R(d)$. For accurate interpolation of $R(d)$ over a large dynamic range of $10^{-10} < d < 10^{-1}$, only a moderate interpolation order of $q = 32$ is required. After an offline precomputation to generate $\{R_k\}_{k=0}^q$, constructing $R(d)$ at each time step costs $\mathcal{O}(p^2 q)$ operations.

3 Accelerating timestepping with subspace recycling

While the two-disc preconditioner is effective at lowering the number of GMRES iterations induced by close-to-touching interactions, a significant number of GMRES iterations may still be required at each time step for problems with many discs. To ameliorate this effect, we propose a preconditioner which effectively reuses the Krylov subspace from previous time steps.

After the kth iteration of GMRES, the Krylov matrix is given by $X = [b \; Ab \; \cdots \; A^{k-1}b]$. Let $QR = AX$ be the QR decomposition of AX. Then the matrix P given by

$$P = I - QQ^T + XR^{-1}Q^T$$

has the following properties:

$$PAx = x \text{ for all } x \in \text{span}(X),$$
$$Py = y \text{ for all } y \perp \text{span}(X).$$

Hence, P effectively reuses the given Krylov subspace X when used as a preconditioner in a Krylov method. In a high-order time-stepping scheme based on spectral deferred corrections, this preconditioner can drastically reduce the number GMRES iterations by recycling the Krylov subspace between time steps.

References

[1] C. Pozrikidis, *Boundary Integral and Singularity Methods for Linearized Viscous Flow*, Cambridge University Press, 1992.

[2] J. Helsing, *Solving integral equations on piecewise smooth boundaries using the RCIP method: a tutorial*, 2022.