$\mathcal{DH}^2$-matrix compression for Helmholtz problems

Steffen Börm

Christian-Albrechts-Universität zu Kiel

GAMM ANLA 2016

Overview

Helmholtz equation

Goal: Solve the Helmholtz equation

$$\begin{aligned} \Delta u(x) + \kappa^2 u(x) &= 0 && \text{for all } x \in \Omega, \\ u(x) &= f(x) && \text{for all } x \in \partial\Omega. \end{aligned}$$

Challenge: Indefinite operator, pollution effect.

Approach: Boundary integral formulation

$$u(x) = \int_{\partial\Omega} g(x,y) \frac{\partial u}{\partial n}(y)\,dy - \int_{\partial\Omega} \frac{\partial g}{\partial n(y)}(x,y) u(y)\,dy \quad \text{for all } x \in \Omega$$

with fundamental solution

$$g(x,y) = \frac{\exp(\iota\kappa\|x-y\|)}{4\pi\|x-y\|}.$$

Challenge: Standard discretization leads to dense matrix.

Fast summation

Standard approach: Approximate g by degenerate kernel,

$$g(x, y) \approx \sum_{\nu=1}^{k} v_\nu(x) \sum_{\mu=1}^{k} s_{\nu\mu} \overline{w_\mu(y)}$$

Fast summation

Standard approach: Approximate g by degenerate kernel,

$$\int_{\partial\Omega} g(x,y)u(y)\,dy \approx \sum_{\nu=1}^{k} v_\nu(x) \int_{\partial\Omega} \sum_{\mu=1}^{k} s_{\nu\mu} \overline{w_\mu(y)} u(y)\,dy.$$

Fast summation

Standard approach: Approximate g by degenerate kernel,

$$\int_{\partial\Omega} g(x,y)u(y)\,dy \approx \sum_{\nu=1}^{k} v_\nu(x) \int_{\partial\Omega} \sum_{\mu=1}^{k} s_{\nu\mu}\overline{w_\mu(y)}u(y)\,dy.$$

Challenge: Helmholtz kernel oscillates rapidly if κ large.

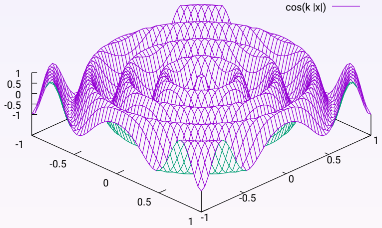

$$g(x,y) = \frac{\exp(\iota\kappa\|x-y\|)}{\|x-y\|}.$$

Consequence: Standard approximation schemes require $k \gg 1$.

Experiment: $\mathcal{H}^2$-matrix compression

Helmholtz matrix approximated by quasi-optimal $\mathcal{H}^2$-matrix.

n	κ	k	k/κ	M/n
2048	8	33	4.1	12.5
4608	12	47	3.9	15.4
8192	16	61	3.8	18.3
18432	24	95	4.0	22.2
32768	32	132	4.1	25.4
73728	48	263	5.5	30.1
131072	64	333	5.2	36.0

Observation: Maximal rank proportional to wave number κ, storage looks like $\mathcal{O}(n \log^2 \kappa)$.

Overview

Directional approximation

Idea: Divide spherical wave by plane wave in direction c, $\|c\| = 1$:

$$g(x,y) = \exp(\iota\kappa\langle x-y, c\rangle) \underbrace{\frac{\exp(\iota\kappa(\|x-y\| - \langle x-y, c\rangle))}{4\pi\|x-y\|}}_{=g_c(x,y)}.$$

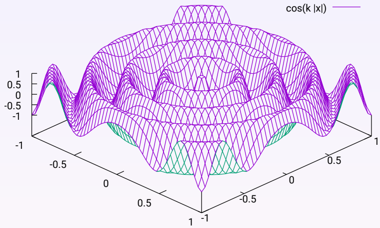

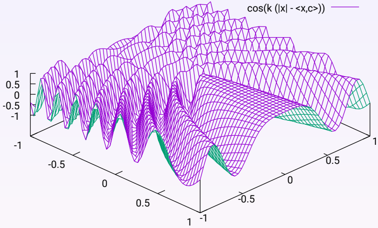

✎ Brandt (1991), Engquist/Ying (2007), Messner/Schanz/Darve (2012), Bebendorf/Kuske/Venn (2015)

Directional interpolation

Approach: Apply interpolation to smoothed function g_c.

$$\tilde{g}_{\tau\sigma,c}(x,y) := \sum_{\nu,\mu=1}^{k} \mathcal{L}_{\tau,\nu}(x) g_c(\xi_{\tau,\nu}, \xi_{\sigma,\mu}) \mathcal{L}_{\sigma,\mu}(y).$$

Directional approximation in domain $\tau \times \sigma$ given by

$$\begin{aligned} g(x,y) &= \exp(\iota\kappa\langle x-y, c\rangle) g_c(x,y) \approx \exp(\iota\kappa\langle x-y, c\rangle)\tilde{g}_{\tau\sigma,c}(x,y) \\ &= \sum_{\nu,\mu=1}^{k} \underbrace{\exp(\iota\kappa\langle x, c\rangle)\mathcal{L}_{\tau,\nu}(x)}_{=:\mathcal{L}_{\tau c,\nu}(x)} g_c(\xi_{\tau,\nu}, \xi_{\sigma\mu}) \underbrace{\exp(-\iota\kappa\langle y, c\rangle)\mathcal{L}_{\sigma,\mu}(y)}_{=:\overline{\mathcal{L}_{\sigma c,\mu}(y)}} \\ &= \sum_{\nu,\mu=1}^{k} \mathcal{L}_{\tau c,\nu}(x) g_c(\xi_{\tau,\nu}, \xi_{\sigma,\mu}) \overline{\mathcal{L}_{\sigma c,\mu}(y)} =: \tilde{g}_{\tau\sigma}(x,y). \end{aligned}$$

Matrix factorization

Goal: Approximate Galerkin matrix $G \in \mathbb{C}^{\mathcal{I}\times\mathcal{I}}$ given by

$$g_{ij} = \int_{\partial\Omega} \varphi_i(x) \int_{\partial\Omega} g(x,y)\varphi_j(y)\,dy\,dx.$$

Idea: If $g|_{\tau\times\sigma} \approx \tilde{g}_{\tau\sigma}$, choose

$$\widehat{\tau} := \{i \in \mathcal{I} \;:\; \operatorname{supp}\varphi_i \subseteq \tau\}, \qquad \widehat{\sigma} := \{j \in \mathcal{I} \;:\; \operatorname{supp}\varphi_j \subseteq \sigma\}$$

and approximate $G|_{\hat{\tau}\times\hat{\sigma}}$ by discretizing $\tilde{g}_{\tau\sigma}$ instead of g.

Result: Factorization $G|_{\hat{\tau}\times\hat{\sigma}} \approx V_{\tau c} S_{\tau\sigma} V_{\sigma c}^*$ with

$$v_{\tau c,i\nu} = \int_{\partial\Omega} \varphi_i(x)\mathcal{L}_{\tau c,\nu}(x)\,dx, \qquad s_{\tau\sigma,\nu\mu} = g_c(\xi_{\tau,\nu}, \xi_{\sigma,\mu}).$$

Directions

Kernel function $g_c(\cdot, y)$ for $y \in \sigma$
is smooth in a cone with axis c.

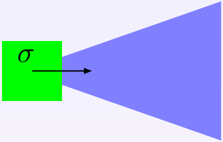

Directions

Kernel function $g_c(\cdot, y)$ for $y \in \sigma$ is smooth in a cone with axis c.

Interactions with clusters τ inside the cone can be approximated.

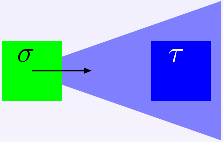

Directions

Kernel function $g_c(\cdot, y)$ for $y \in \sigma$ is smooth in a cone with axis c.

Interactions with clusters τ inside the cone can be approximated.

Multiple cones have to be used to cover the entire domain.

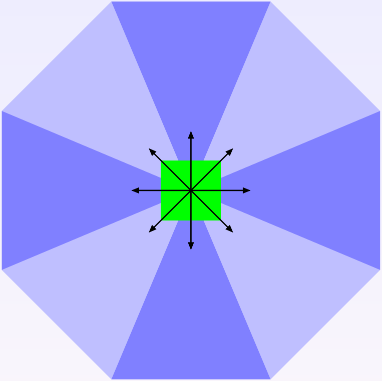

Directions

Kernel function $g_c(\cdot, y)$ for $y \in \sigma$ is smooth in a cone with axis c.

Interactions with clusters τ inside the cone can be approximated.

Multiple cones have to be used to cover the entire domain.

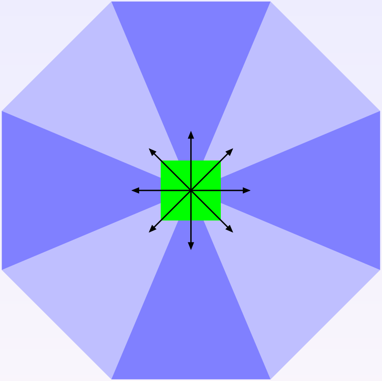

Approach: Choose fixed set of directions for each τ.
$\mathcal{O}(1 + \kappa^2 \operatorname{diam}^2(\tau))$ directions are sufficient.

Nested basis

Problem: Storing $V_{\tau c}$ for all τ, c requires $\sim n^2$ units of memory.

Solution: Organize domains τ in a cluster tree such that

$$\operatorname{sons}(\tau) \neq \emptyset \qquad \Longrightarrow \qquad \hat{\tau} = \bigcup_{\tau' \in \operatorname{sons}(\tau)} \hat{\tau}',$$

and approximate Lagrange polynomials by weighted interpolation

$$\mathcal{L}_{\tau c,\nu} \approx \sum_{\nu'=1}^{k} e_{\tau' \tau c,\nu' \nu} \, \mathcal{L}_{\tau' c',\nu'}.$$

Result: Nested representation $V_{\tau c}|_{\tau'} \approx V_{\tau' c'} E_{\tau' \tau c}$ if $\tau' \in \operatorname{sons}(\tau)$.

✎ Messner/Schanz/Darve (2012), B./Melenk (2015)

$\mathcal{DH}^2$-matrix

Approach:

- Fix set of directions $\mathcal{D}_\tau$ for all clusters τ.
- Choose admissible blocks $(\tau, \sigma, c) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.
- Compute leaf matrices $V_{\tau c} \in \mathbb{C}^{\hat{\tau}\times k}$ for all leaves τ and all $c \in \mathcal{D}_\tau$.
- Find transfer matrices $E_{\tau'\tau c} \in \mathbb{C}^{k\times k}$ for all $\tau' \in \operatorname{sons}(\tau)$, $c \in \mathcal{D}_\tau$.
- Create coupling matrices $S_{\tau\sigma} \in \mathbb{C}^{k\times k}$ for all $(\tau, \sigma, c) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.

Results:

- Storage $\mathcal{O}(kn + k^2\kappa^2 \log n)$.
- Matrix-vector multiplication in $\mathcal{O}(kn + k^2\kappa^2 \log n)$ operations.
- High-frequency case: $\kappa^2 \sim n$, complexity $\mathcal{O}(nk^2 \log n)$.

✎ B. (2015)

Overview

1. Introduction
2. Directional approximation
3. Algebraic compression
4. Experiments

Approximation of general matrices

Given: Matrix $G \in \mathbb{C}^{\mathcal{I}\times\mathcal{I}}$ and cluster/block structure.

Goal: Find $\mathcal{DH}^2$-matrix approximation of G,
i.e., directional cluster basis $(V_{\tau c})$ and coupling matrices $(S_{\tau\sigma})$.

Orthogonal basis: If we ensure $V_{\tau c}^* V_{\tau c} = I$, the best coupling matrices are given by orthogonal projection

$$G|_{\hat\tau\times\hat\sigma} \approx V_{\tau c} \underbrace{V_{\tau c}^* G|_{\hat\tau\times\hat\sigma} V_{\sigma c}}_{=:S_{\tau\sigma}} V_{\sigma c}^*,$$

Row and column projections can be analyzed separately,

$$\begin{aligned}\|G|_{\hat\tau\times\hat\sigma} - V_{\tau c} S_{\tau\sigma} V_{\sigma c}^*\|_2^2 \leq \|G|_{\hat\tau\times\hat\sigma} - V_{\tau c} V_{\tau c}^* G|_{\hat\tau\times\hat\sigma}\|_2^2 \\ + \|G|_{\hat\tau\times\hat\sigma} - G|_{\hat\tau\times\hat\sigma} V_{\sigma c} V_{\sigma c}^*\|_2^2.\end{aligned}$$

Low-rank structure

Question: What has to be approximated by $V_{\tau c}$?

Answer: All blocks $G|_{\hat{\tau}\times\hat{\sigma}}$ with $(\tau, \sigma, c) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.

Low-rank structure

Question: What has to be approximated by $V_{\tau c}$?

Partial answer: All blocks $G|_{\hat\tau\times\hat\sigma}$ with $(\tau,\sigma,c) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.

Important: Directional cluster basis is nested.
Anything that is not approximated in the son clusters cannot be approximated in the father.

Low-rank structure

Question: What has to be approximated by $V_{\tau c}$?

Partial answer: All blocks $G|_{\hat\tau\times\hat\sigma}$ with $(\tau,\sigma,c)\in\mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.

Important: Directional cluster basis is nested.
Anything that is not approximated in the son clusters cannot be approximated in the father.

Complete answer: All blocks $G|_{\hat\tau\times\hat\sigma^*}$ where

- $(\tau^*,\sigma^*,c^*)\in\mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$,
- τ is a descendant of τ^*, and
- c^* approximates c.

We collect these tuples in the set $\mathcal{R}_{\tau c}\subseteq\mathcal{L}^+_{\mathcal{I}\times\mathcal{I}}$.

Leaf matrices

Assume that τ is a leaf cluster and that $\mathcal{R}_{\tau c}$ is available for all $c \in \mathcal{D}_\tau$.

Goal: Find $V_{\tau c}$ with k columns such that

$$G|_{\hat{\tau} \times \hat{\sigma}^*} \approx V_{\tau c} V_{\tau c}^* G|_{\hat{\tau} \times \hat{\sigma}^*} \qquad \text{for all } (\tau^*, \sigma^*, c^*) \in \mathcal{R}_{\tau c}.$$

Leaf matrices

Assume that τ is a leaf cluster and that $\mathcal{R}_{\tau c}$ is available for all $c \in \mathcal{D}_\tau$.

Goal: Find $V_{\tau c}$ with k columns such that

$$G|_{\hat\tau\times\hat\sigma^*} \approx V_{\tau c} V_{\tau c}^* G|_{\hat\tau\times\hat\sigma^*} \qquad \text{for all } (\tau^*, \sigma^*, c^*) \in \mathcal{R}_{\tau c}.$$

Reformulation: Define the submatrix

$$G_{\tau c} := G|_{\hat\tau\times\mathcal{F}_{\tau c}}, \qquad \mathcal{F}_{\tau c} := \bigcup\{\hat\sigma^* \ :\ (\tau^*, \sigma^*, c^*) \in \mathcal{R}_{\tau c}\}$$

with farfield matrix $G_{\tau c}$ and consider only

$$G_{\tau c} \approx V_{\tau c} V_{\tau c}^* G_{\tau c}.$$

$\rightarrow$ Compute SVD of $G_{\tau c}$, remove small singular values.

Transfer matrices

Assume $\operatorname{sons}(\tau) = \{\tau_1, \tau_2\}$, $c_1 \in \mathcal{D}_{\tau_1}$, $c_2 \in \mathcal{D}_{\tau_2}$ approximate $c \in \mathcal{D}_\tau$.

Recursion: Ensure that $V_{\tau_1 c_1}$, $V_{\tau_2 c_2}$ have already been computed.

Reduced matrices given by

$$\widehat{G}_{\tau c} = \begin{pmatrix} V_{\tau_1 c_1}^* G|_{\hat{\tau}_1 \times \mathcal{F}_{\tau c}} \\ V_{\tau_2 c_2}^* G|_{\hat{\tau}_2 \times \mathcal{F}_{\tau c}} \end{pmatrix}, \qquad \widehat{V}_{\tau c} = \begin{pmatrix} E_{\tau_1 \tau c} \\ E_{\tau_2 \tau c} \end{pmatrix}$$

lead to

$$\begin{aligned} \|G_{\tau c} - V_{\tau c} V_{\tau c}^* G_{\tau c}\|_2^2 \leq\ & \|G|_{\hat{\tau}_1 \times \mathcal{F}_{\tau c}} - V_{\tau_1 c_1} V_{\tau_1 c_1}^* G|_{\hat{\tau}_1 \times \mathcal{F}_{\tau c}}\|_2^2 \\ & + \|G|_{\hat{\tau}_2 \times \mathcal{F}_{\tau c}} - V_{\tau_2 c_2} V_{\tau_2 c_2}^* G|_{\hat{\tau}_2 \times \mathcal{F}_{\tau c}}\|_2^2 \\ & + \|\widehat{G}_{\tau c} - \widehat{V}_{\tau c} \widehat{V}_{\tau c}^* \widehat{G}_{\tau c}\|_2^2. \end{aligned}$$

The first two terms correspond to the sons' influence, the last can be controlled via SVD and provides us with the transfer matrices.

$\mathcal{DH}^2$-matrix compression algorithm

Recursion following the cluster tree.

- Set up $\mathcal{R}_{\tau c}$ during the downward pass.
- Construct $V_{\tau c}$, $E_{\tau' \tau c}$, and $V_{\tau c}^* G_{\tau c}$ during the upward pass.

Compared to $\mathcal{H}^2$-matrices:

- Similar complexity estimates.
- Parallelization tricky, large clusters very time-consuming.
- Complexity $\mathcal{O}(n^2 k)$

$\mathcal{DH}^2$-matrix compression algorithm

Recursion following the cluster tree.

- Set up $\mathcal{R}_{\tau c}$ during the downward pass.
- Construct $V_{\tau c}$, $E_{\tau' \tau c}$, and $V_{\tau c}^* G_{\tau c}$ during the upward pass.

Compared to $\mathcal{H}^2$-matrices:

- Similar complexity estimates.
- Parallelization tricky, large clusters very time-consuming.
- Complexity $\mathcal{O}(n^2 k)$, but
- efficient recompression of $\mathcal{DH}^2$-matrices in $\mathcal{O}(nk^2 \log n)$.

Overview

Experiment: Single and double layer potential

Approximate the matrix G corresponding to SLP and DLP operators for the unit sphere.

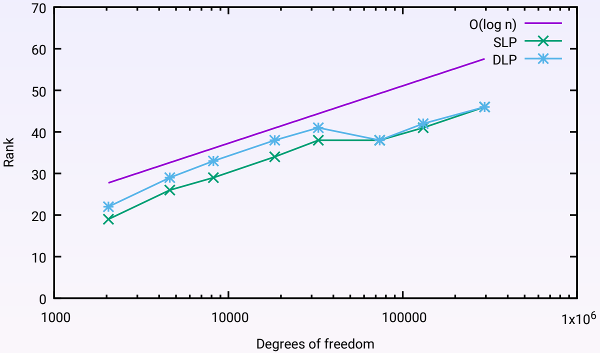

Result: Maximal rank grows like $\mathcal{O}(\log n)$.

Experiment: Preconditioner

Approximate the LR factorization of the Brakhage-Werner operator.

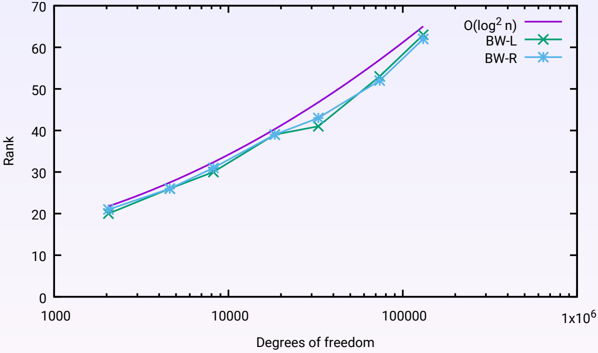

Result: Maximal rank grows like $\mathcal{O}(\log^2 n)$.

Conclusion

Challenge: Approximate high-frequency Helmholtz operators.

Approach: Directional approximation, $\mathcal{DH}^2$-matrix.
$\rightarrow$ Complexity $\mathcal{O}(nk^2 \log n)$.

Compression: Recursive algorithm, $\mathcal{O}(n^2 k)$ operations.

Work in progress: On-the-fly recompression in $\mathcal{O}(nk^2 \log n)$, $\mathcal{DH}^2$-matrix arithmetic operations, inclusion in H2Lib package

http://www.h2lib.org