

Fast direct solvers for elliptic PDEs

Per-Gunnar Martinsson


Dept. of Mathematics & Oden Institute for Computational Sciences and Engineering

University of Texas at Austin

Students, postdocs, collaborators: Daniel Appelö, Chao Chen, Ke Chen, Yijun Dong, Adrianna Gillman, Abinand Gopal, James Levitt, Michael O'Neil, Heather Wilber, Bowei Wu, Anna Yesypenko.

Slides: http://users.oden.utexas.edu/~pgm/main_talks.html

Research support by:



Problem addressed: The talk concerns numerical methods for boundary value problems of the form

$$(BVP) \quad \begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma, \end{cases}$$

where Ω is a domain (2D or 3D) with boundary Γ , and where A is a linear elliptic differential operator; possibly with variable coefficients.

Examples of problems we are interested in:

- The equations of linear elasticity.
- Stokes' equation.
- Helmholtz' equation (at least at low and intermediate frequencies).
- Time-harmonic Maxwell (at least at low and intermediate frequencies).

Archetypical example: Poisson equation with Dirichlet boundary data:

$$\begin{cases} -\Delta u(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

Problem addressed: The talk concerns numerical methods for boundary value problems of the form

$$(BVP) \quad \begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma, \end{cases}$$

where Ω is a domain (2D or 3D) with boundary Γ , and where A is a linear elliptic differential operator; possibly with variable coefficients.

Examples of problems we are interested in:

- The equations of linear elasticity.
- Stokes' equation.
- Helmholtz' equation (at least at low and intermediate frequencies).
- Time-harmonic Maxwell (at least at low and intermediate frequencies).

Archetypical example: Poisson equation with Dirichlet boundary data:

$$\begin{cases} -\Delta u(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

Standard numerical recipe for (BVP): (1) Discretize via FD/FEM. (2) Iterative solver.

Focal point of this talk: The solution operator for (BVP). → Direct solvers.

Problem addressed: The talk concerns numerical methods for boundary value problems of the form

$$(BVP) \quad \begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma, \end{cases}$$

where Ω is a domain (2D or 3D) with boundary Γ , and where A is a linear elliptic differential operator; possibly with variable coefficients.

Linear solution operators: As a warmup, let us consider the Poisson equation

$$-\Delta u(\mathbf{x}) = g(\mathbf{x}) \quad \mathbf{x} \in \mathbb{R}^2$$

(with suitable decay conditions at infinity to ensure uniqueness). The solution is given by

$$(SLN) \quad u(\mathbf{x}) = \int_{\mathbb{R}^2} \phi(\mathbf{x} - \mathbf{y}) g(\mathbf{y}) d\mathbf{y}, \quad \mathbf{x} \in \mathbb{R}^2.$$

where the “fundamental solution” of the Laplace operator $-\Delta$ on \mathbb{R}^2 is defined by

$$\phi(\mathbf{x}) = -\frac{1}{2\pi} \log |\mathbf{x}|.$$

In principle very simple. Numerically non-trivial, however: The operator is *global*, so discretizing it leads to a *dense* matrix. (There is also the singular kernel to worry about!)

Problem addressed: The talk concerns numerical methods for boundary value problems of the form

$$(BVP) \quad \begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma, \end{cases}$$

where Ω is a domain (2D or 3D) with boundary Γ , and where A is a linear elliptic differential operator; possibly with variable coefficients.

Linear solution operators: A general solution operator for (BVP) takes the form

$$(SLN) \quad u(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) d\mathbf{y} + \int_{\Gamma} F(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) dS(\mathbf{y}), \quad \mathbf{x} \in \Omega,$$

where G and F are two kernel functions that depend on A , B , and Ω .

Good: The operators in (SLN) are friendly and nice.

Bounded, smoothing, often fairly stable, etc.

Bad: The kernels G and F in (SLN) are generally *unknown*.

(Other than in trivial cases — constant coefficients and very simple domains.)

Bad: The operators in (SLN) are *global*.


Dense matrices upon discretization. $O(N^2)$ cost? $O(N^3)$ cost?

Recall that we are interested in solving the PDE
$$\begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases} \quad (\text{BVP})$$

Explicit solution formula:
$$u(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) d\mathbf{y} + \int_{\Gamma} F(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) dS(\mathbf{y}), \quad \mathbf{x} \in \Omega. \quad (\text{SLN})$$

Recurring idea: Upon discretization, (SLN) leads to a matrix with *off-diagonal blocks of low numerical rank*.

This property can be exploited to attain linear or close to linear complexity for operations such as matrix-vector multiply, matrix-matrix multiply, LU factorization, matrix inversion, forming of Schur complements, etc.



All gray blocks have low rank.


Strong connections to Calderón-Zygmund theory for singular integral operators.

References: Fast Multipole Method (Greengard, Rokhlin); Panel Clustering (Hackbusch); \mathcal{H} - and \mathcal{H}^2 -matrices (Hackbusch et al); Hierarchically Block Separable matrices; Hierarchically Semi Separable matrices (Xia et al); HODLR matrices (Darve et al); BLR matrices (Buttari, Amestoy, Mary, ...); ...

Recall that we are interested in solving the PDE
$$\begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases} \quad (\text{BVP})$$

Explicit solution formula:
$$u(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) d\mathbf{y} + \int_{\Gamma} F(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) dS(\mathbf{y}), \quad \mathbf{x} \in \Omega. \quad (\text{SLN})$$

In real life, tessellation patterns of rank structured matrices tend to be more complex ...



Recall that we are interested in solving the PDE
$$\begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases} \quad (\text{BVP})$$

Explicit solution formula:
$$u(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) d\mathbf{y} + \int_{\Gamma} F(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) dS(\mathbf{y}), \quad \mathbf{x} \in \Omega. \quad (\text{SLN})$$

The talk will describe recent work on algorithms that numerically construct an approximation to (SLN).

When using these algorithms, the process of solving (BVP) splits into two stages:

1. **“Factorization” or “build” stage:** Build a representation of the inverse operator.
2. **“Solve” stage:** Apply the computed inverse to given data f or (and) g .

Typical characteristics of methods of this type:

- Memory usage tends to be high.
- Stage 1 tends to be slower than an iterative solve, *when convergence is fast*.
- Stage 2 is almost always VERY fast.

Fast Direct Solvers (FDS) are most competitive when:

- Getting iterative methods to converge rapidly is hard.
- When the cost of Stage 1 can be amortized over many solves.

→ scattering problems, time-stepping, optimization, ...

History:

1980s: Rokhlin and Greengard develop the Fast Multipole Method.

1991: Beylkin, Coifman, Rokhlin: Fast algorithms exist for most solution operators.

1996: Michielssen, Boag, Chew: Fast direct solvers for 3D scattering problems in certain geometries.

1998 onwards: Hackbusch, Bebendorf, Börm, Grasedyck, Khoromskij, Sauter, Tyrtyshnikov, ... develop \mathcal{H} and \mathcal{H}^2 -frameworks that provide explicit recipes for operator algebra in $O(n \log^r n)$ operations for r moderate.

Outline of talk:

- Introduction: Problem formulation & solution operators. *[Done!]*
- **Curse of dimensionality.**
- Interaction ranks — why are they small? How small are they?
- (Versions of fast direct solvers — “strong” versus “weak” etc.)
- High order discretizations and fast direct solvers.
- [New!] Randomized compression of rank structured matrices.

Curse of dimensionality

Algorithms involving rank-structured matrices scale *very* poorly with dimension.

For instance, for the classical Fast Multipole Method, key quantities scale as:

Dimension	Typical ranks	Number of “neighbors”	Length of “interaction list”
1	2	2	3
2	10–50	8	27
3	50–500	26	189
d	$(\log(1/\varepsilon))^{d-1}$	$3^d - 1$	$6^d - 3^d$

Curse of dimensionality

Algorithms involving rank-structured matrices scale *very* poorly with dimension.

For instance, for the classical Fast Multipole Method, key quantities scale as:

Dimension	Typical ranks	Number of “neighbors”	Length of “interaction list”
1	2	2	3
2	10–50	8	27
3	50–500	26	189
d	$(\log(1/\varepsilon))^{d-1}$	$3^d - 1$	$6^d - 3^d$

For fast direct solvers, the scaling with dimension is equally problematic:

Dimension	Current state of affairs
1	Extremely fast. Linear scaling is easy to attain.
2	Quite fast in practice, but simple methods do not scale linearly.
3	Slow. Basic methods cannot go much beyond $N \approx 10^7$.

Curse of dimensionality

Algorithms involving rank-structured matrices scale *very* poorly with dimension.

For instance, for the classical Fast Multipole Method, key quantities scale as:

Dimension	Typical ranks	Number of “neighbors”	Length of “interaction list”
1	2	2	3
2	10–50	8	27
3	50–500	26	189
d	$(\log(1/\varepsilon))^{d-1}$	$3^d - 1$	$6^d - 3^d$

For fast direct solvers, the scaling with dimension is equally problematic:


Effective Dimension	Current state of affairs
1	Extremely fast. Linear scaling is easy to attain.
2	Quite fast in practice, but simple methods do not scale linearly.
3	Slow. Basic methods cannot go much beyond $N \approx 10^7$.

Fortunately, we can often reduce the *effective* dimensionality.

For many 3D problems, the dense problems we need to invert “live” on 2D domains!

Curse of dimensionality: Dimension reduction via an integral equation

Recall that many boundary value problems can advantageously be recast as *boundary integral equations*. Consider, e.g., (sound-soft) acoustic scattering from a finite body:



$$(3) \quad \begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0 & \mathbf{x} \in \mathbb{R}^3 \setminus \bar{\Omega} \\ u(\mathbf{x}) = v(\mathbf{x}) & \mathbf{x} \in \partial\Omega \\ \lim_{|\mathbf{x}| \rightarrow \infty} |\mathbf{x}| (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0. \end{cases}$$

The BVP (3) has an alternative mathematical formulation in the BIE

$$(4) \quad -\pi i \sigma(\mathbf{x}) + \int_{\partial\Omega} \left(\left(\partial_{\mathbf{n}(\mathbf{y})} + i\kappa \right) \frac{e^{i\kappa|\mathbf{x}-\mathbf{y}|}}{|\mathbf{x}-\mathbf{y}|} \right) \sigma(\mathbf{y}) dS(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega.$$

The integral equation (4) has several advantages over the PDE (3), including:

- The domain of computation $\partial\Omega$ is finite.
- The domain of computation $\partial\Omega$ is 2D, while $\mathbb{R}^3 \setminus \bar{\Omega}$ is 3D.
- Equation (4) is inherently well-conditioned (as a “2nd kind Fredholm equation”).

A serious drawback of integral equations is that they lead to *dense coefficient matrices*.

Since we are interested in constructing inverses anyway, this is unproblematic for us!

Curse of dimensionality: Dimension reduction via sparse direct solvers

Let us next consider what happens if we directly discretize the PDE (using, say, finite elements or finite differences) to obtain a linear system

$$\mathbf{A}\mathbf{u} = \mathbf{b}$$

involving a *sparse* coefficient matrix \mathbf{A} .

Key idea: Do a sparse LU factorization based on a “nested dissection” ordering of the grid as an outer solver. Then use rank structured matrix algebra to deal with the dense matrices that arise.

Curse of dimensionality: Dimension reduction via sparse direct solvers

A 2D model problem: Let $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$. We seek to solve

$$(5) \quad \begin{cases} -\Delta u(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We introduce an $n \times n$ grid on Ω with nodes $\{\mathbf{x}_j\}_{j=1}^N$ where $N = n^2$, see Figure A. Letting $\mathbf{u} = [\mathbf{u}(j)]_{j=1}^N$ denote a vector of approximate solution values, $\mathbf{u}(j) \approx u(\mathbf{x}_j)$, and using the standard five-point stencil to discretize $-\Delta$, we end up with a sparse linear system

$$\mathbf{A}\mathbf{u} = \mathbf{b},$$

where $[\mathbf{A}\mathbf{u}](k) = \frac{1}{h^2}(4\mathbf{u}(k) - \mathbf{u}(k_s) - \mathbf{u}(k_e) - \mathbf{u}(k_n) - \mathbf{u}(k_w))$, see Figure B.




Figure A: The grid






Figure B: The 5-point stencil

Divide-and-conquer: Split the nodes in three groups as shown so that there are no connections between nodes in Ω_1 and Ω_2 . Then \mathbf{A} has zero blocks as shown:



$$\mathbf{A} = \begin{array}{|c|c|c|} \hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

Divide-and-conquer: Split the nodes in three groups as shown so that there are no connections between nodes in Ω_1 and Ω_2 . Then \mathbf{A} has zero blocks as shown:




$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}$$

Now suppose that we can somehow construct \mathbf{A}_{11}^{-1} and \mathbf{A}_{22}^{-1} . Then

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{A}_{31}\mathbf{A}_{11}^{-1} & \mathbf{A}_{32}\mathbf{A}_{22}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{S}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{A}_{11}^{-1}\mathbf{A}_{13} \\ \mathbf{0} & \mathbf{I} & \mathbf{A}_{22}^{-1}\mathbf{A}_{23} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}$$

where $\mathbf{S}_{33} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$ is a *Schur complement*.

Divide-and-conquer: Split the nodes in three groups as shown so that there are no connections between nodes in Ω_1 and Ω_2 . Then \mathbf{A} has zero blocks as shown:



$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}$$

Now suppose that we can somehow construct \mathbf{A}_{11}^{-1} and \mathbf{A}_{22}^{-1} . Then

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{A}_{31}\mathbf{A}_{11}^{-1} & \mathbf{A}_{32}\mathbf{A}_{22}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{S}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{A}_{11}^{-1}\mathbf{A}_{13} \\ \mathbf{0} & \mathbf{I} & \mathbf{A}_{22}^{-1}\mathbf{A}_{23} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}$$


where $\mathbf{S}_{33} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$ is a *Schur complement*.

In other words, in order to invert \mathbf{A} , we need to execute three steps:

- Invert \mathbf{A}_{11} to form \mathbf{A}_{11}^{-1} . size $\sim N/2 \times N/2$
- Invert \mathbf{A}_{22} to form \mathbf{A}_{22}^{-1} . size $\sim N/2 \times N/2$
- Invert $\mathbf{S}_{33} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$. size $\sim \sqrt{N} \times \sqrt{N}$

Notice the obvious recursion!

Divide-and-conquer: Split the nodes in three groups as shown so that there are no connections between nodes in Ω_1 and Ω_2 . Then \mathbf{A} has zero blocks as shown:



$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}$$

Now suppose that we can somehow construct \mathbf{A}_{11}^{-1} and \mathbf{A}_{22}^{-1} . Then

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{A}_{31}\mathbf{A}_{11}^{-1} & \mathbf{A}_{32}\mathbf{A}_{22}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{S}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{A}_{11}^{-1}\mathbf{A}_{13} \\ \mathbf{0} & \mathbf{I} & \mathbf{A}_{22}^{-1}\mathbf{A}_{23} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}$$

where $\mathbf{S}_{33} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$ is a *Schur complement*.

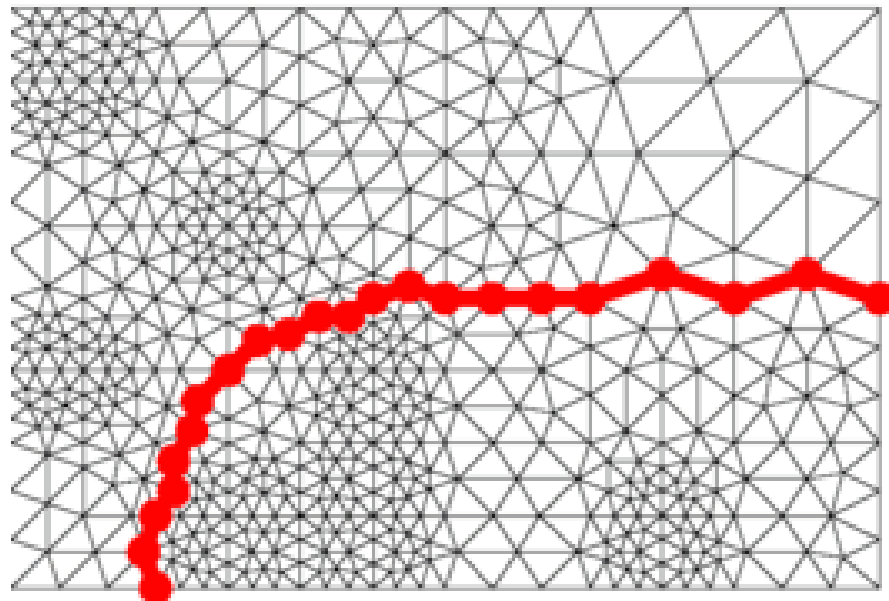
In other words, in order to invert \mathbf{A} , we need to execute three steps:

- Invert \mathbf{A}_{11} to form \mathbf{A}_{11}^{-1} . size $\sim N/2 \times N/2$
- Invert \mathbf{A}_{22} to form \mathbf{A}_{22}^{-1} . size $\sim N/2 \times N/2$
- Invert $\mathbf{S}_{33} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$. size $\sim \sqrt{N} \times \sqrt{N}$

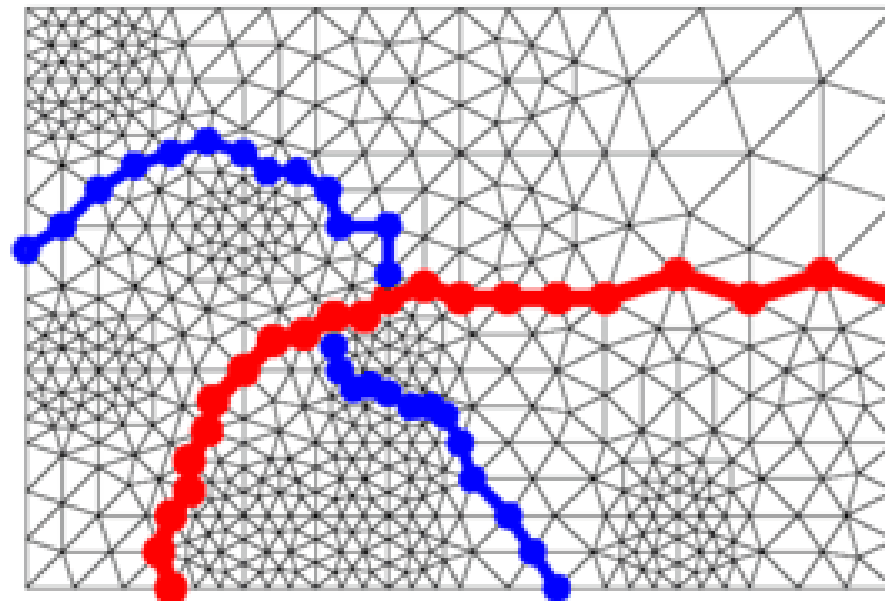
Notice the obvious recursion!

Sparse direct solvers with nested dissection ordering

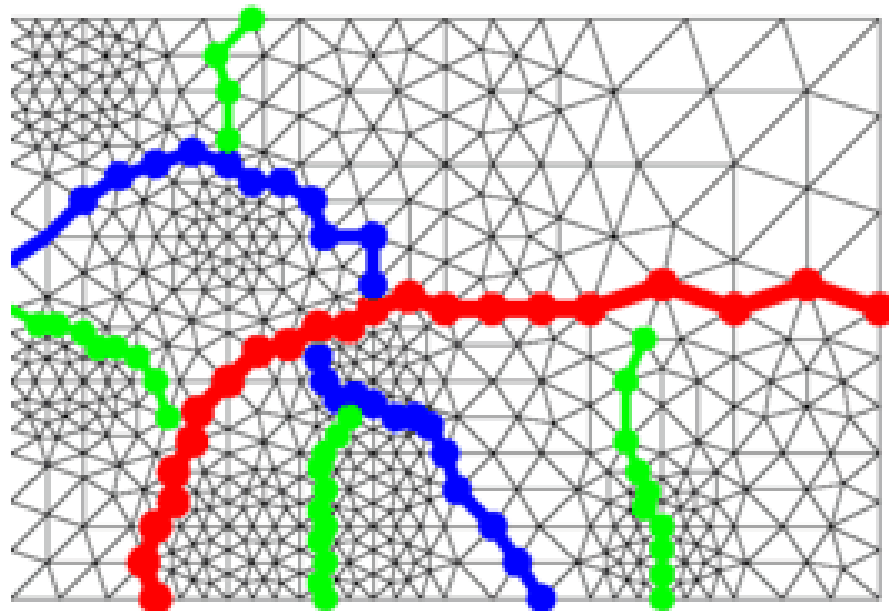
Typically, nested dissection orderings are more complicated:



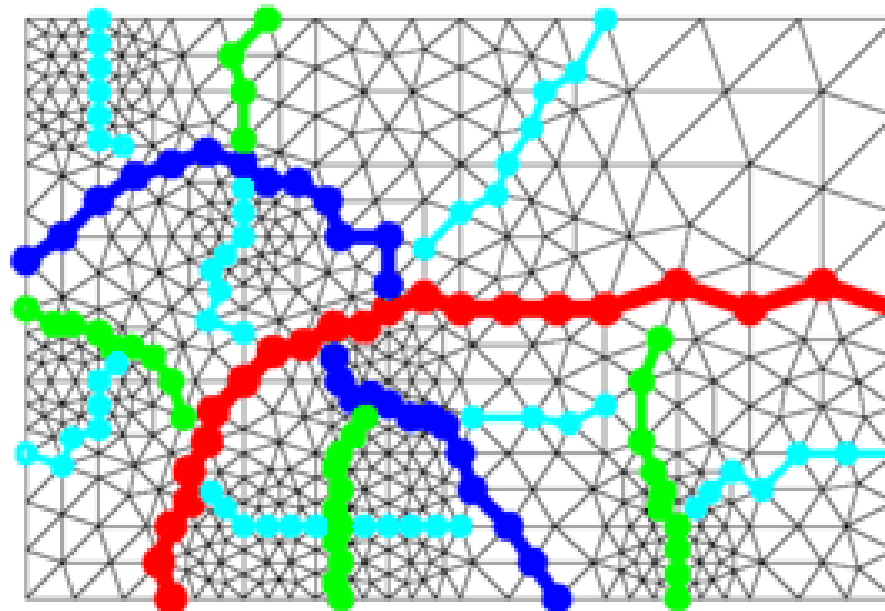
(i) *Top level separator*



(ii) *Two levels of separators*



(iii) *Three levels of separators*



(iv) *Four levels of separators*

Image credit: Jianlin Xia, "Robust and Efficient Multifrontal Solver for Large Discretized PDEs", 2012

Observe that while the computational domain is **2D** in this example, the rank structured matrices all live on the colored **1D** domains.

Curse of dimensionality: Dimension reduction

Key point: *When faced with a BVP in 3D, you can in most circumstances build direct solvers that rely only on dense operators associated with 2D domains.*

1. Constant coefficient problems: Reformulate as integral equation on boundary.
2. Variable coefficient problems: Use a sparse direct solver as an “outer” solver.

Outline of talk:

- Introduction: Problem formulation & solution operators. *[Done!]*
- Curse of dimensionality. *[Done!]*
- **Interaction ranks — why are they small? How small are they?**
- (Versions of fast direct solvers — “strong” versus “weak” etc.)
- High order discretizations and fast direct solvers.
- [New!] Randomized compression of rank structured matrices.

Recall that we are interested in solving the PDE
$$\begin{cases} Au(\mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \Omega, \\ Bu(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases} \quad (\text{BVP})$$

Explicit solution formula:
$$u(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) d\mathbf{y} + \int_{\Gamma} F(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) dS(\mathbf{y}), \quad \mathbf{x} \in \Omega. \quad (\text{SLN})$$

Question: Why do the dense matrices resulting upon discretization of (SLN) typically have *off-diagonal blocks of low numerical rank*?

(One) Answer: It is a consequence of the *smoothing effect* of elliptic differential equations; it can be interpreted as a *loss of information*.

This effect has many well known physical consequences:

- Rapid convergence of *multipole expansions* when the region of sources is far away from the observation point.
- The *St Venant principle* in mechanics.
- The inaccuracy of imaging at sub-wavelength scales.
- The intractability of solving the heat equation backwards.

Caveat: High-frequency problems present difficulties — no loss of information for length-scales $> \lambda$. Extreme accuracy of optics, high-frequency imaging, *etc.*

Interaction ranks: Boundary integral equations

Let us consider two simple boundary integral equations on a boundary Γ :

The first is a reformulation of a Dirichlet problem involving the Laplace equation:

$$\alpha\sigma(\mathbf{x}) + \int_{\Gamma} (d(\mathbf{x}, \mathbf{y}) + s(\mathbf{x}, \mathbf{y})) \sigma(\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma.$$

The second is a reformulation of a Dirichlet problem involving the Helmholtz equation:

$$\beta\sigma(\mathbf{x}) + \int_{\Gamma} (d_{\kappa}(\mathbf{x}, \mathbf{y}) + i\kappa s_{\kappa}(\mathbf{x}, \mathbf{y})) \sigma(\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma.$$

The kernels are derived from the corresponding fundamental solutions:

$$s(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x} - \mathbf{y}),$$

$$d(\mathbf{x}, \mathbf{y}) = \partial_{\mathbf{n}(\mathbf{y})} \phi(\mathbf{x} - \mathbf{y}),$$

$$s_{\kappa}(\mathbf{x}, \mathbf{y}) = \phi_{\kappa}(\mathbf{x} - \mathbf{y}),$$

$$d_{\kappa}(\mathbf{x}, \mathbf{y}) = \partial_{\mathbf{n}(\mathbf{y})} \phi_{\kappa}(\mathbf{x} - \mathbf{y}),$$

where, as before,

$$\phi(\mathbf{x}) = -\frac{1}{2\pi} \log |\mathbf{x}|,$$
$$\phi_{\kappa}(\mathbf{x}) = \frac{i}{4} H_0^{(1)}(\kappa |\mathbf{x}|).$$

Interaction ranks: Boundary integral equations

Let us consider two simple boundary integral equations on a boundary Γ :


The first is a reformulation of a Dirichlet problem involving the Laplace equation:

$$\alpha\sigma(\mathbf{x}) + \int_{\Gamma} (d(\mathbf{x}, \mathbf{y}) + s(\mathbf{x}, \mathbf{y})) \sigma(\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma.$$

The second is a reformulation of a Dirichlet problem involving the Helmholtz equation:

$$\beta\sigma(\mathbf{x}) + \int_{\Gamma} (d_{\kappa}(\mathbf{x}, \mathbf{y}) + i\kappa s_{\kappa}(\mathbf{x}, \mathbf{y})) \sigma(\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma.$$


Let \mathbf{A} denote the matrix resulting from discretization of either BIE.



On the next slide, we show the singular values of the off-diagonal block \mathbf{A}_{23} .

Interaction ranks: Boundary integral equations


The ranks of an off-diagonal block of \mathbf{A} :



This is all as expected. Somewhat accessible by analysis.

Interaction ranks: Boundary integral equations

The ranks of an off-diagonal block of \mathbf{A} :




This is all as expected. Somewhat accessible by analysis.

Now the fun part! We set $\mathbf{B} = \mathbf{A}^{-1}$, and plot the svds of the off-diagonal block \mathbf{B}_{23} .


Interaction ranks: Boundary integral equations

The ranks of an off-diagonal block of \mathbf{A} :



This is all as expected. Somewhat accessible by analysis.

Now the fun part! We set $\mathbf{B} = \mathbf{A}^{-1}$, and plot the svds of the off-diagonal block \mathbf{B}_{23} .




Remarkable similarity!

(Observe ill-conditioning due to close resonances for the Helmholtz BIE.)

Interaction ranks: Stiffness matrix from finite difference discretization


Recall our example of Laplace's equation discretized using the 5-point stencil.



$$\mathbf{A} = \begin{array}{|c|c|c|} \hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

We build the Schur complement $\mathbf{S} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$.

Then split the Schur complement into four parts:




$$\mathbf{S} = \begin{array}{|c|c|} \hline \mathbf{S}_{\alpha\alpha} & \mathbf{S}_{\alpha\beta} \\ \hline \mathbf{S}_{\beta\alpha} & \mathbf{S}_{\beta\beta} \\ \hline \end{array}$$

We explore the svds of $\mathbf{S}_{\alpha\beta}$ — encoding interactions between I_α and I_β .

Interaction ranks: Stiffness matrix from finite difference discretization


Recall our example of Laplace's equation discretized using the 5-point stencil.




$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}$$

We build the Schur complement $\mathbf{S} = \mathbf{A}_{33} - \mathbf{A}_{31}\mathbf{A}_{11}^{-1}\mathbf{A}_{13} - \mathbf{A}_{32}\mathbf{A}_{22}^{-1}\mathbf{A}_{23}$.

Then split the Schur complement into four parts:




Singular values of $S_{\alpha\beta}$ for an 80×80 grid.




We explore the svds of $\mathbf{S}_{\alpha\beta}$ — encoding interactions between I_α and I_β .

Interaction ranks: Stiffness matrix from finite difference discretization

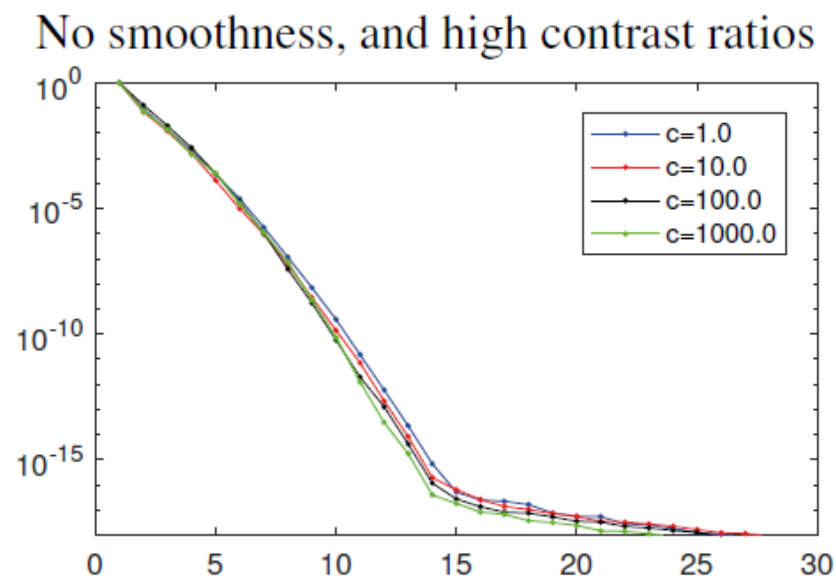
Let us try a few different PDEs, and different problem sizes:




(a)



(b)



(c)



(d)

Note: The rank decay property is remarkably stable!


Note: The decay continues to ϵ_{mach} — regardless of the discretization errors!

Interaction ranks: Stiffness matrix from finite difference discretization

Next, let us consider Helmholtz problems with increasing wave numbers.

Interaction ranks: Stiffness matrix from finite difference discretization

Next, let us consider Helmholtz problems with increasing wave numbers.




We recognize this pattern from the potential evaluation operator:

Fast decay once oscillations are resolved.

Interaction ranks: Stiffness matrix from finite difference discretization


Finally, let us consider the analogous 3D problem.



The geometry.

Interaction ranks: Stiffness matrix from finite difference discretization

Finally, let us consider the analogous 3D problem.



The singular values.

Outline of talk:

- Introduction: Problem formulation & solution operators. *[Done!]*
- Curse of dimensionality. *[Done!]*
- Interaction ranks — why are they small? How small are they? *[Done!]*
- (Versions of fast direct solvers — “strong” versus “weak” etc.)
- High order discretizations and fast direct solvers.
- [New!] Randomized compression of rank structured matrices.

Outline of talk:

- Introduction: Problem formulation & solution operators. *[Done!]*
- Curse of dimensionality. *[Done!]*
- Interaction ranks — why are they small? How small are they? *[Done!]*
- (Versions of fast direct solvers — “strong” versus “weak” etc.) *[Skipped.]*
- **High order discretizations and fast direct solvers.**
- [New!] Randomized compression of rank structured matrices.

Fast direct solvers and high order methods

Claim: Direct solvers are ideal for combining with *high order discretization*.

- Direct solvers use a lot of memory per degree of freedom.
→ *You want to maximize the oomph per DOF.*
- Direct solvers are particularly well suited for medium frequency wave problems.
→ *Need high accuracy due to ill-conditioned physics.*
- High order methods sometimes lead to more ill-conditioned systems.
→ *Can be hard to get iterative solvers to converge.*

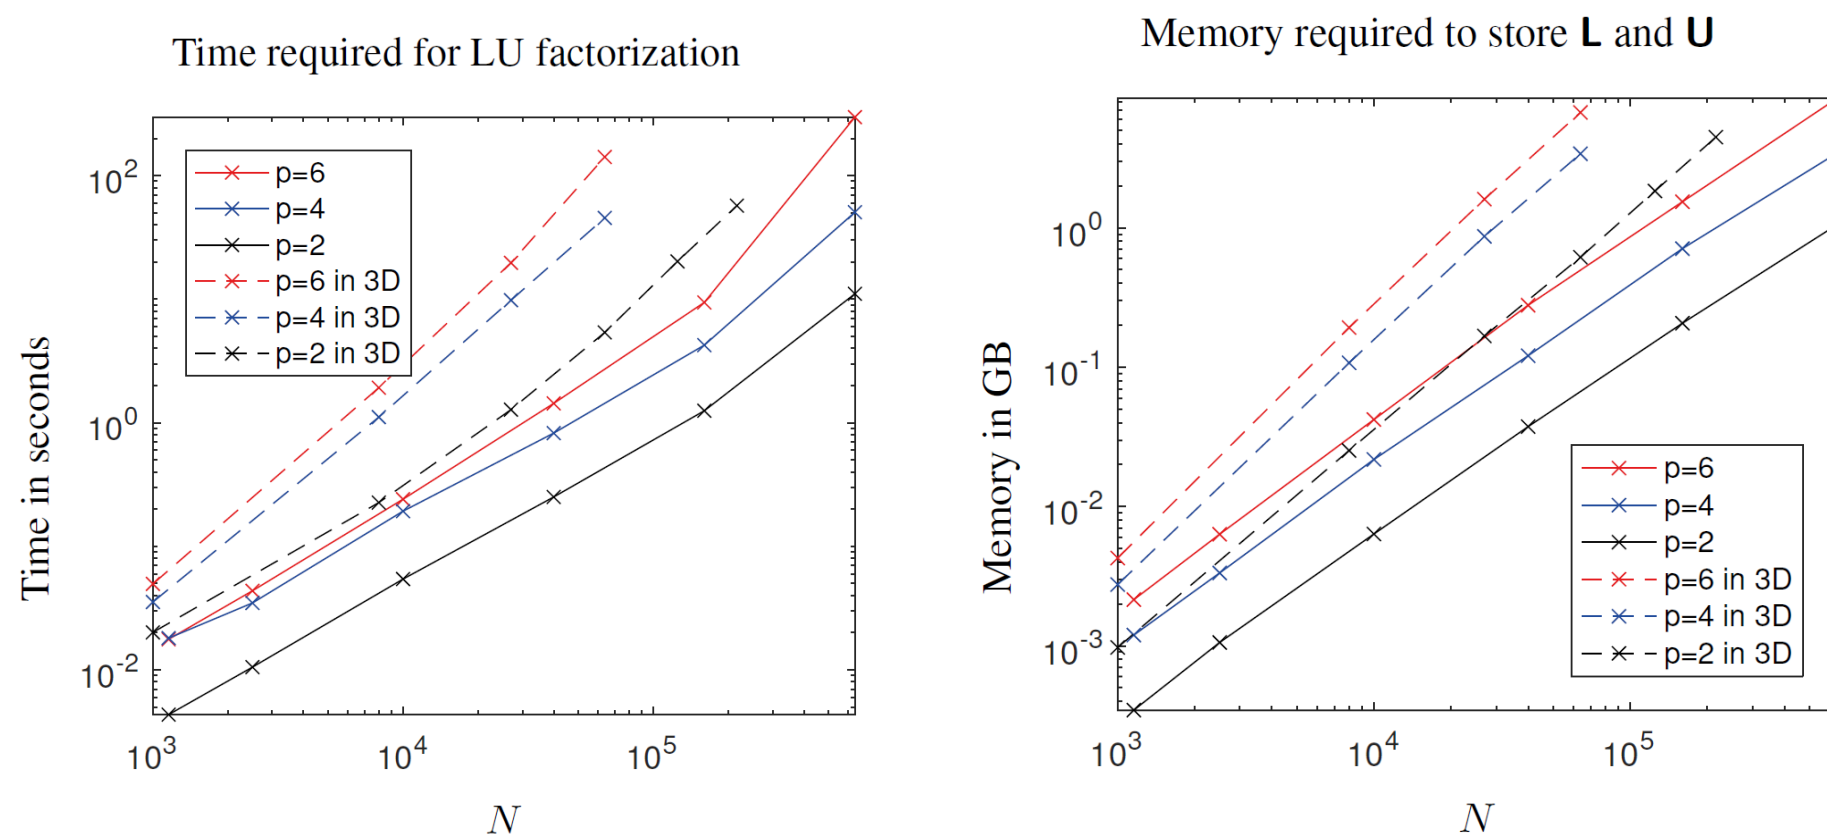
Problem: If you combine “nested dissection” with traditional discretization techniques (FD, FEM, etc), then the performance *plummets* as the order is increased.

Fast direct solvers and high order methods

Claim: Direct solvers are ideal for combining with *high order discretization*.

- Direct solvers use a lot of memory per degree of freedom.
→ *You want to maximize the oomph per DOF.*
- Direct solvers are particularly well suited for medium frequency wave problems.
→ *Need high accuracy due to ill-conditioned physics.*
- High order methods sometimes lead to more ill-conditioned systems.
→ *Can be hard to get iterative solvers to converge.*

Problem: If you combine “nested dissection” with traditional discretization techniques (FD, FEM, etc), then the performance *plummets* as the order is increased.



Fast direct solvers and high order methods

Claim: Direct solvers are ideal for combining with *high order discretization*.

- Direct solvers use a lot of memory per degree of freedom.
→ *You want to maximize the oomph per DOF.*
- Direct solvers are particularly well suited for medium frequency wave problems.
→ *Need high accuracy due to ill-conditioned physics.*
- High order methods sometimes lead to more ill-conditioned systems.
→ *Can be hard to get iterative solvers to converge.*

Problem: If you combine “nested dissection” with traditional discretization techniques (FD, FEM, etc), then the performance *plummets* as the order is increased.

Solution: Pick your discretization scheme carefully!


- When discretizing the PDE, use methods that play well with “static condensation”.
You want a clean separation between “interior” and “edge” degrees of freedom.
 - Multidomain spectral collocation methods (“HPS”, “ultraSEM”, etc.).
 - Discontinuous Galerkin. (Or so I speculate, at any rate.)
- When integral equation formulations are used, pick quadratures that have as localized “corrections” as possible. (Very technical point here!)

Fast direct solvers and high order methods: Multidomain spectral collocation

As a numerical illustration, let us consider the “Hierarchical Poincaré-Steklov (HPS)” method. We set $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$, and consider the problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We discretize using spectral collocation on a composite grid on Ω (Chebyshev nodes):




On patch boundaries, we enforce continuity of the potential and the normal derivative.

Fast direct solvers and high order methods: Multidomain spectral collocation

As a numerical illustration, let us consider the “Hierarchical Poincaré-Steklov (HPS)” method. We set $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$, and consider the problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We discretize using spectral collocation on a composite grid on Ω (Chebyshev nodes):



On patch boundaries, we enforce continuity of the potential and the normal derivative.

Fast direct solvers and high order methods: Multidomain spectral collocation

As a numerical illustration, let us consider the “Hierarchical Poincaré-Steklov (HPS)” method. We set $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$, and consider the problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We pick f as the restriction of a wave from a point source, $\mathbf{x} \mapsto Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

We then know the exact solution, $u_{\text{exact}}(\mathbf{x}) = Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

Fast direct solvers and high order methods: Multidomain spectral collocation


As a numerical illustration, let us consider the “Hierarchical Poincaré-Steklov (HPS)” method. We set $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$, and consider the problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We pick f as the restriction of a wave from a point source, $\mathbf{x} \mapsto Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

We then know the exact solution, $u_{\text{exact}}(\mathbf{x}) = Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

Approximate solution. ntot=1681 pts-per-wave=12.00



Fast direct solvers and high order methods: Multidomain spectral collocation

As a numerical illustration, let us consider the “Hierarchical Poincaré-Steklov (HPS)” method. We set $\Omega = [0, 1]^2$ and $\Gamma = \partial\Omega$, and consider the problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 u(\mathbf{x}) = 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{cases}$$

We pick f as the restriction of a wave from a point source, $\mathbf{x} \mapsto Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

We then know the exact solution, $u_{\text{exact}}(\mathbf{x}) = Y_0(\kappa|\mathbf{x} - \hat{\mathbf{x}}|)$.

The spectral computation on a leaf involves 21×21 points.

κ is chosen so that there are 12 points per wave-length.

p	N	N_{wave}	t_{build} (sec)	t_{solve} (sec)	E_{pot}	E_{grad}	M (MB)	M/N (reals/DOF)
21	6561	6.7	0.23	0.0011	2.56528e-10	1.01490e-08	4.4	87.1
21	25921	13.3	0.92	0.0044	5.24706e-10	4.44184e-08	18.8	95.2
21	103041	26.7	4.68	0.0173	9.49460e-10	1.56699e-07	80.8	102.7
21	410881	53.3	22.29	0.0727	1.21769e-09	3.99051e-07	344.9	110.0
21	1640961	106.7	99.20	0.2965	1.90502e-09	1.24859e-06	1467.2	117.2
21	6558721	213.3	551.32	20.9551	2.84554e-09	3.74616e-06	6218.7	124.3

Error is measured in sup-norm: $e = \max_{\mathbf{x} \in \Omega} |u(\mathbf{x}) - u_{\text{exact}}(\mathbf{x})|$.

Note: The times refer to a simple Matlab implementation executed on a \$1k laptop.

Note: Fixed number of points per wave-length. Almost no “pollution error”!


Fast direct solvers and high order methods: Multidomain spectral collocation

By incorporating rank-structured matrix algebra for the Schur complements, we can access larger problem sizes, and get linear scaling in important cases.

Problem	N	T_{build}	T_{solve}	MB
Laplace	1.7e6	91.68	0.34	1611.19
	6.9e6	371.15	1.803	6557.27
	2.8e7	1661.97	6.97	26503.29
	1.1e8	6894.31	30.67	106731.61
Helmholtz I	1.7e6	62.07	0.202	1611.41
	6.9e6	363.19	1.755	6557.12
	2.8e7	1677.92	6.92	26503.41
	1.1e8	7584.65	31.85	106738.85
Helmholtz II	1.7e6	93.96	0.29	1827.72
	6.9e6	525.92	2.13	7151.60
	2.8e7	2033.91	8.59	27985.41
Helmholtz III	1.7e6	105.58	0.44	1712.11
	6.9e6	510.37	2.085	7157.47
	2.8e7	2714.86	10.63	29632.89

(About six accurate digits in solution.)

Let us consider a multibody scattering problem involving multiple cavities:





Acoustic scattering on the exterior domain. Each bowl is about 5λ .

A hybrid direct/iterative solver is used (a highly accurate scattering matrix is computed for each body). On an office desktop, we achieved an accuracy of 10^{-5} , in about 6h (essentially all the time is spent in applying the inter-body interactions via the Fast Multipole Method). Accuracy 10^{-7} took 27h. [2015, CAMWA, Hao/M./Young]

Fast direct solvers and high order methods:

Boundary integral equations

Consider sound-soft scattering from a multi-body scatterer of size 4 wave-lengths:



The global scattering matrix is computed using the hierarchical direct solver described.
(The ellipsoids are not rotationally symmetric.)

[2015, BIT, with Bremer/Gillman/Martinsson.]

Fast direct solvers and high order methods:

Boundary integral equations

The local truncation error is set to 10^{-3} .

Grid dimensions	N	T	E	Ratio	Predicted
$2 \times 2 \times 2$	12 288	$1.02 \times 10^{+1}$	3.37×10^{-04}	-	-
$3 \times 3 \times 3$	41 472	$3.43 \times 10^{+1}$	4.81×10^{-04}	3.4	6.2
$4 \times 4 \times 4$	98 304	$7.92 \times 10^{+1}$	1.57×10^{-04}	2.3	3.7
$6 \times 6 \times 6$	331 776	$2.96 \times 10^{+2}$	7.03×10^{-04}	3.7	6.2
$8 \times 8 \times 8$	786 432	$6.70 \times 10^{+2}$	4.70×10^{-04}	2.3	3.7
$10 \times 10 \times 10$	1 536 000	$2.46 \times 10^{+3}$	3.53×10^{-04}	3.7	2.7

Increasing the accuracy is possible, but comes at a cost.




Now the local truncation error is set to 10^{-6} .

Grid dimensions	N	T	E	Ratio	Predicted
$2 \times 2 \times 2$	49 152	$1.61 \times 10^{+2}$	1.22×10^{-07}	-	-
$3 \times 3 \times 3$	165 888	$6.87 \times 10^{+2}$	4.92×10^{-07}	4.3	6.2
$4 \times 4 \times 4$	393 216	$1.68 \times 10^{+3}$	5.31×10^{-07}	2.4	3.6
$6 \times 6 \times 6$	1 327 104	$6.66 \times 10^{+3}$	4.60×10^{-06}	4.0	6.2
$8 \times 8 \times 8$	3 145 728	$1.59 \times 10^{+4}$	2.30×10^{-07}	2.4	3.6

[2015, BIT, with Bremer/Gillman/Martinsson.]

Fast direct solvers and high order methods:

Boundary integral equations



The domain is roughly $2 \times 2 \times 0.7$ wave-lengths in size.

$N_{\text{triangles}}$	N	T	E
32	1 664	$7.16 \times 10^{+00}$	3.51×10^{-02}
128	6 656	$6.29 \times 10^{+01}$	4.41×10^{-03}
512	26 624	$2.81 \times 10^{+02}$	4.08×10^{-05}
2 048	106 496	$2.60 \times 10^{+03}$	7.80×10^{-07}
8 192	425 984	$1.47 \times 10^{+04}$	3.25×10^{-08}

(Note: Laplace problems are much faster.)


[2015, BIT, with Bremer/Gillman/Martinsson.]

A surface Γ with corners and edges.

The grid has been refined to attain high accuracy.

Computing scattering matrices for the corners is conceptually easy (but laborious). The direct solver eliminates “extra” DOFs.

Compressing the edges takes effort!



N_{tris}	N	E	T	$N_{\text{out}} \times N_{\text{in}}$
192	21 504	2.60×10^{-08}	$6.11 \times 10^{+02}$	617×712
432	48 384	2.13×10^{-09}	$1.65 \times 10^{+03}$	620×694
768	86 016	3.13×10^{-10}	$3.58 \times 10^{+03}$	612×685

Results from a Helmholtz problem (acoustic scattering) on the domain exterior to the “edgy” cube.

The domain is about 3.5 wave-lengths in diameter.

[2015, BIT, with Bremer/Gillman/Martinsson.]

Fast direct solvers and high order methods:

“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

$$-\Delta u - \kappa^2 (1-b)u = -\kappa^2 b v$$

support(b)

Fast direct solvers and high order methods:


“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”



Introduce an artificial box Ω such that $\text{support}(b) \subseteq \Omega$.

$$-\Delta u - \kappa^2 u = 0 \text{ on } \Omega^c$$

Fast direct solvers and high order methods:

“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

$$-\Delta u - \kappa^2 (1 - b) u = -\kappa^2 b v$$

$$-\Delta u - \kappa^2 u = 0 \text{ on } \Omega^c$$

Introduce an artificial box Ω such that $\text{support}(b) \subseteq \Omega$.

On Ω :

- Variable coefficient PDE.

On Ω^c :

- Constant coefficient PDE.

Fast direct solvers and high order methods:

“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

$$-\Delta u - \kappa^2 (1 - b) u = -\kappa^2 b v$$

$$-\Delta u - \kappa^2 u = 0 \text{ on } \Omega^c$$

Introduce an artificial box Ω such that $\text{support}(b) \subseteq \Omega$.

On Ω :

- Variable coefficient PDE.
- Use HPS.

On Ω^c :

- Constant coefficient PDE.
- Use BIE.

Fast direct solvers and high order methods:

“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

$$-\Delta u - \kappa^2 (1 - b) u = -\kappa^2 b v$$

support(b)

$$-\Delta u - \kappa^2 u = 0 \text{ on } \Omega^c$$

Introduce an artificial box Ω such that $\text{support}(b) \subseteq \Omega$.

On Ω :

- Variable coefficient PDE.
- Use HPS.
- Build DtN for $\partial\Omega$.

On Ω^c :

- Constant coefficient PDE.
- Use BIE.
- Build DtN for $\partial\Omega^c$.

Fast direct solvers and high order methods:

“FEM-BEM coupling”

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support**, where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

$$-\Delta u - \kappa^2 (1 - b) u = -\kappa^2 b v$$

support(b)

$$-\Delta u - \kappa^2 u = 0 \text{ on } \Omega^c$$

Introduce an artificial box Ω such that $\text{support}(b) \subseteq \Omega$.

On Ω :

- Variable coefficient PDE.
- Use HPS.
- Build DtN for $\partial\Omega$.

On Ω^c :

- Constant coefficient PDE.
- Use BIE.
- Build DtN for $\partial\Omega^c$.


• Merge using fast operator algebra!

Fast direct solvers and high order methods:

“FEM-BEM coupling”

$$\begin{cases} -\Delta u_{\text{out}}(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u_{\text{out}}(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) u_{\text{in}}(\mathbf{x}) \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u_{\text{out}}(\mathbf{x}) - i\kappa u_{\text{out}}(\mathbf{x})) = 0 \end{cases}$$

The scattering potential b




Fast direct solvers and high order methods:

“FEM-BEM coupling”

$$\begin{cases} -\Delta u_{\text{out}}(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u_{\text{out}}(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) u_{\text{in}}(\mathbf{x}) \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u_{\text{out}}(\mathbf{x}) - i\kappa u_{\text{out}}(\mathbf{x})) = 0 \end{cases}$$

The outgoing field u_{out} (resulting from an incoming plane wave $u_{\text{in}}(x) = \cos(\kappa x_1)$)




$N = 231\,361$ $T_{\text{build}} = 7.2 \text{ sec}$ $T_{\text{solve}} = 0.06 \text{ sec}$ $E \approx 10^{-7}$ (estimated)

Fast direct solvers and high order methods:

“FEM-BEM coupling”

$$\begin{cases} -\Delta u_{\text{out}}(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u_{\text{out}}(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) u_{\text{in}}(\mathbf{x}) \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u_{\text{out}}(\mathbf{x}) - i\kappa u_{\text{out}}(\mathbf{x})) = 0 \end{cases}$$

The outgoing field u_{out} (resulting from an incoming plane wave $u_{\text{in}}(x) = \cos(\kappa x_1)$)




$N = 231\,361$ $T_{\text{build}} = 7.2 \text{ sec}$ $T_{\text{solve}} = 0.06 \text{ sec}$ $E \approx 10^{-7}$ (estimated)

Fast direct solvers and high order methods:

“FEM-BEM coupling”

$$\begin{cases} -\Delta u_{\text{out}}(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u_{\text{out}}(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) u_{\text{in}}(\mathbf{x}) \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u_{\text{out}}(\mathbf{x}) - i\kappa u_{\text{out}}(\mathbf{x})) = 0 \end{cases}$$

The scattering potential b — now a photonic crystal with a wave guide.




$N = 231\,361$ $T_{\text{build}} = 7.2 \text{ sec}$ $T_{\text{solve}} = 0.06 \text{ sec}$ $E \approx 10^{-6}$ (estimated)

Fast direct solvers and high order methods:

“FEM-BEM coupling”

$$\begin{cases} -\Delta u_{\text{out}}(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u_{\text{out}}(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) u_{\text{in}}(\mathbf{x}) \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u_{\text{out}}(\mathbf{x}) - i\kappa u_{\text{out}}(\mathbf{x})) = 0 \end{cases}$$

The total field $u = u_{\text{in}} + u_{\text{out}}$ (resulting from an incoming plane wave $u_{\text{in}}(x) = \cos(\kappa x_1)$).



Fast direct solvers and high order methods:

rank deficiencies in HPS


Recall that at the top level, we need to invert a dense matrix that is defined on the nodes of the interface high-lighted in red and blue below. This matrix holds restrictions of the Dirichlet-to-Neumann (DtN) operators for the two blocks. We have claimed that this matrix is rank-structured. *But what are the ranks?*

Fast direct solvers and high order methods:

rank deficiencies in HPS

Recall that at the top level, we need to invert a dense matrix that is defined on the nodes of the interface high-lighted in red and blue below. This matrix holds restrictions of the Dirichlet-to-Neumann (DtN) operators for the two blocks. We have claimed that this matrix is rank-structured. *But what are the ranks?*

Let \mathbf{T} denote the restriction of the DtN matrix mapping Dirichlet data on Γ_1 to Neumann data on Γ_2 for a $1\,089 \times 1\,089$ grid. Then \mathbf{T} is of size 512×512 .




Fast direct solvers and high order methods:


rank deficiencies in HPS

Recall that at the top level, we need to invert a dense matrix that is defined on the nodes of the interface high-lighted in red and blue below. This matrix holds restrictions of the Dirichlet-to-Neumann (DtN) operators for the two blocks. We have claimed that this matrix is rank-structured. *But what are the ranks?*

Let \mathbf{T} denote the restriction of the DtN matrix mapping Dirichlet data on Γ_1 to Neumann data on Γ_2 for a $1\,089 \times 1\,089$ grid. Then \mathbf{T} is of size 512×512 .



Singular values of \mathbf{T} .




Fast direct solvers and high order methods:

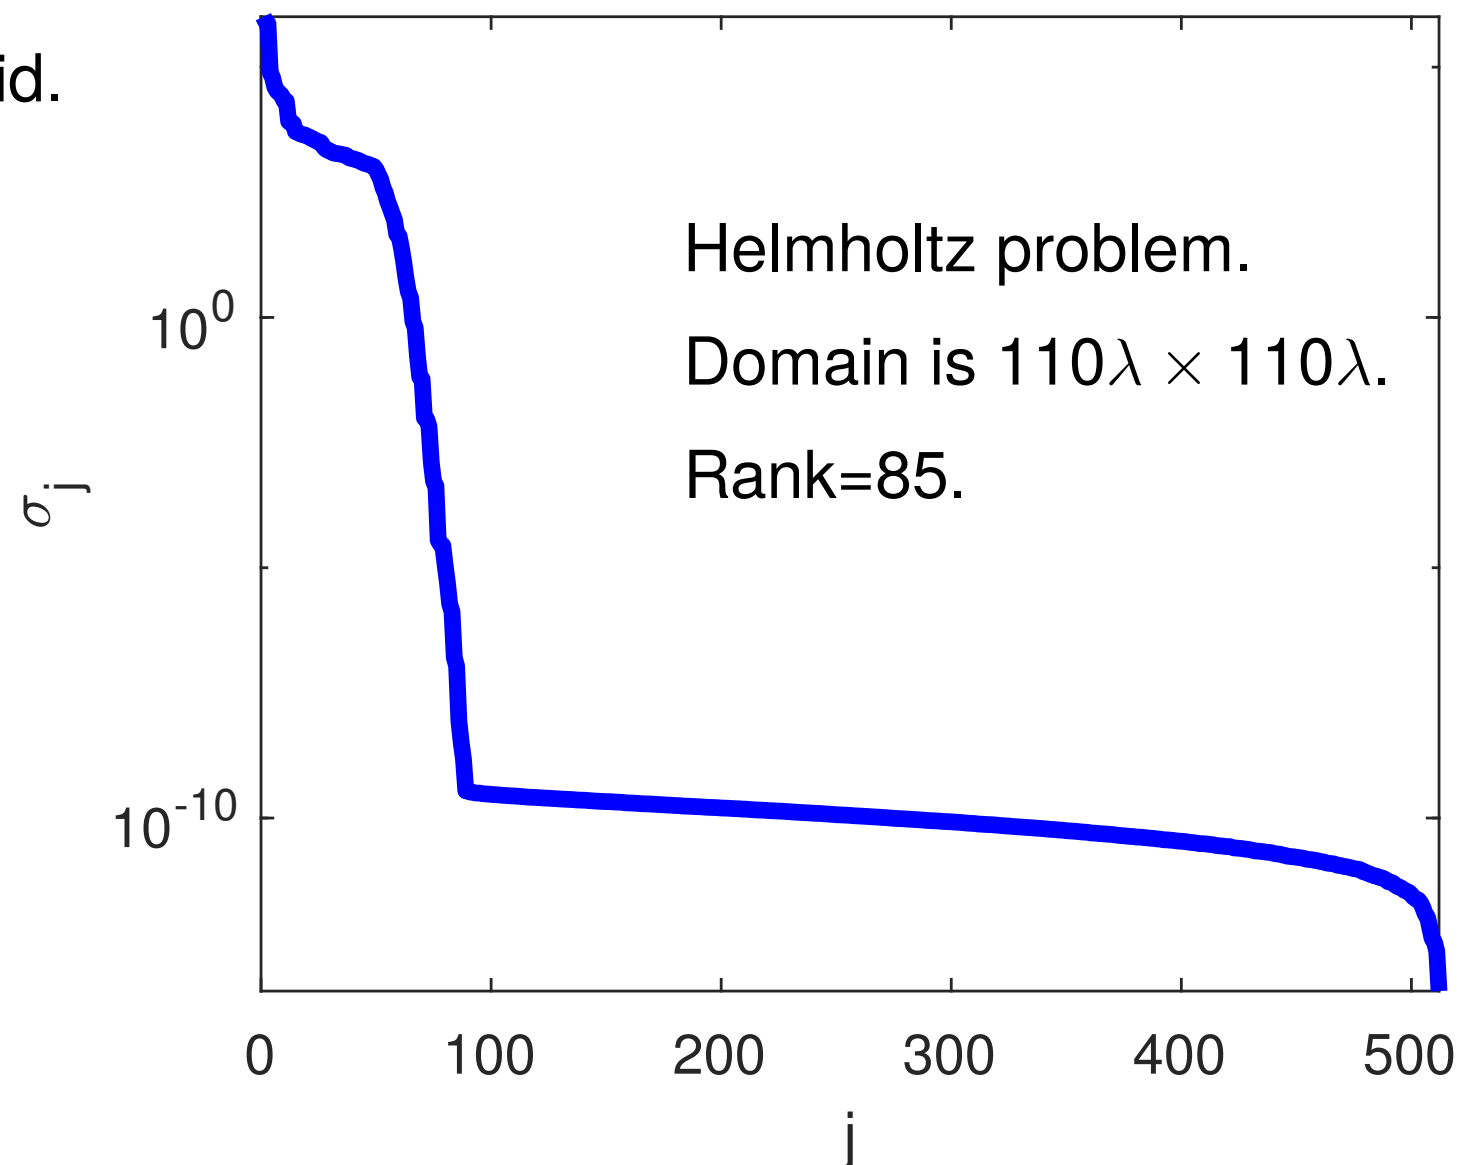
rank deficiencies in HPS

Recall that at the top level, we need to invert a dense matrix that is defined on the nodes of the interface high-lighted in red and blue below. This matrix holds restrictions of the Dirichlet-to-Neumann (DtN) operators for the two blocks. We have claimed that this matrix is rank-structured. *But what are the ranks?*

Let \mathbf{T} denote the restriction of the DtN matrix mapping Dirichlet data on Γ_1 to Neumann data on Γ_2 for a $1\,089 \times 1\,089$ grid. Then \mathbf{T} is of size 512×512 .



Singular values of \mathbf{T} .



Fast direct solvers and high order methods:


Lippmann-Schwinger

Consider the free space acoustic scattering problem

$$(6) \quad \begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support** in a domain Ω , where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”



In the figure, $v = u_{in}$.

Consider the free space acoustic scattering problem

$$(6) \quad \begin{cases} -\Delta u(\mathbf{x}) - \kappa^2 (1 - b(\mathbf{x})) u(\mathbf{x}) = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^2 \\ \lim_{|\mathbf{x}| \rightarrow \infty} \sqrt{|\mathbf{x}|} (\partial_{|\mathbf{x}|} u(\mathbf{x}) - i\kappa u(\mathbf{x})) = 0, \end{cases}$$

where

- b is a smooth scattering potential with **compact support** in a domain Ω , where
- v is a given “incoming potential” and where
- u is the sought “outgoing potential.”

Let us now use an integral equation formulation. It is well known that (6) has an alternative formulation in the *Lippmann-Schwinger integral equation*



$$(7) \quad \sigma(\mathbf{x}) + \kappa^2 b(\mathbf{x}) \int_{\Omega} \phi_{\kappa}(\mathbf{x} - \mathbf{y}) \sigma(\mathbf{y}) d\mathbf{y} = -\kappa^2 b(\mathbf{x}) v(\mathbf{x}), \quad \mathbf{x} \in \Omega,$$

where ϕ_{κ} is the free space fundamental solution of the Helmholtz equation.

We discretize (7) using the trapezoidal rule on a uniform grid, with Duan-Rokhlin quadrature corrections of order 10.


Fast direct solvers and high order methods:

Lippmann-Schwinger




The scattering potential

$\kappa = 50$



$\kappa = 201$



$\kappa = 804$

Fast direct solvers and high order methods:

Lippmann-Schwinger

Fixed 10 points per wavelength.

Direct solver is run at accuracy 10^{-3} and used as a preconditioner.

Weak admissibility is used.

N	κ	T_{build}	T_{inv}	T_{gmres}	mem	iter	res
6400	50.27	0.23	0.24	0.20	0.04	4	6.97e-11
25600	100.53	0.65	0.99	0.62	0.21	5	6.16e-12
102400	201.06	2.26	4.36	2.49	1.01	6	1.04e-12
409600	402.12	14.91	20.06	9.78	4.67	6	3.23e-11
1638400	804.25	99.01	91.37	56.13	21.16	9	8.12e-12
6553600	1608.50	430.60	398.88	330.91	94.63	13	3.93e-11
26214400	3216.99	3102.09	2024.16	2698.53	418.37	22	3.30e-11

The largest experiment is over 500λ in diameter: Less than 3h total run time.


Hardware: Workstation with dual Intel Xeon Gold 6254 (18 cores at 3.1GHz base frequency).

Joint with Abi Gopal, arxiv #2007.12718

Fast direct solvers and high order methods:

rough surface scattering

Consider acoustic scattering from an infinite half-plane with 8×8 bumps:



Each bump is 0.5λ tall and 2λ wide. Total domain is $44\lambda \times 44\lambda$.

The problem is formulated as a BIE, and is discretized using a “Zeta-corrected” quadrature rule. A direct solver with weak admissibility and $O(N^{1.5})$ scaling is used.

26M points total, $T_{\text{build}} = 1200\text{s}$, $T_{\text{solve}} = 184\text{s}$

About three digits of accuracy in the computed solution. (Tentative!)

Work in progress - with Abi Gopal and Bowei Wu.

Active research area!

HSS-accelerated multifrontal solvers: J. Xia, M. de Hoop, X. Li, ...

BLR-accelerated multifrontal solvers: P. Amestoy, C. Ashcraft, A. Buttari, J-Y. l'Excellent, T. Mary, ...

Linear complexity recursive skeletonization: K. Ho, L. Ying.

Linear complexity inverse FMM / strong recursive skeletonization: S. Ambikasaran, E. Darve; V. Minden, K. Ho, A. Damle, L. Ying; M. O'Neil, D. Sushnikova, M. Rachh, L. Greengard; ...

High frequency problems: Y. Liu, H. Guo, E. Michielssen; B. Engquist, L. Ying; S. Li, Y. Liu, P. Ghysels, L. Klaus; S. Börm, C. Börst; B. Bonev, J. Hesthaven; ...

HPC and heterogenous computing environments: D. Keyes, H. Ltaief, G. Turkiyyah; G. Biros, C. Chen; ...

High order spectral element methods: A. Townsend, D. Fortunato;

Apologies for omissions!

Outline of talk:

- Introduction: Problem formulation & solution operators. *[Done!]*
- Curse of dimensionality. *[Done!]*
- Interaction ranks — why are they small? How small are they? *[Done!]*
- (Versions of fast direct solvers — “strong” versus “weak” etc.) *[Skipped.]*
- High order discretizations and fast direct solvers. *[Done!]*
- **[New!] Randomized compression of rank structured matrices.**

Approximation of rank structured matrices

Environment: We are given a rank structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ (to be precise, \mathbf{A} is HBS/HSS of rank k). We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Question: Can you construct thin matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ such that \mathbf{A} can be completely reconstructed in $O(N)$ work from the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ where $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Approximation of rank structured matrices

Environment: We are given a rank structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ (to be precise, \mathbf{A} is HBS/HSS of rank k). We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Question: Can you construct thin matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ such that \mathbf{A} can be completely reconstructed in $O(N)$ work from the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ where $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Applications:

- Integral operators from classical physics. If you have a legacy method for the matrix-vector multiple (e.g. the Fast Multipole Method), then we could enable a range of operations – LU factorization, matrix inversion, etc.
- Multiplication of operators. Useful for forming Dirichlet-to-Neumann operators, for combining solvers of multi-physics problems, etc.
- Compression of Schur complements that arise in the LU or Cholesky factorization of sparse matrices. This lets us overcome key bottlenecks (e.g. LU factorization of a “finite element” matrix is accelerated from $O(N^2)$ to close to linear complexity.)

Approximation of rank structured matrices

Environment: We are given a rank structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ (to be precise, \mathbf{A} is HBS/HSS of rank k). We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Question: Can you construct thin matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ such that \mathbf{A} can be completely reconstructed in $O(N)$ work from the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ where $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Observation: In the simpler case where \mathbf{A} is of *globally low rank*, there are randomized algorithms that resolve the problem.

For instance, if you draw $\mathbf{\Omega}$ and $\mathbf{\Psi}$ from a Gaussian distribution, then (to high probability)

$$\mathbf{A} \approx (\mathbf{A}\mathbf{\Omega}) (\mathbf{\Psi}^* \mathbf{A}\mathbf{\Omega})^\dagger (\mathbf{\Psi}^* \mathbf{A}) = \mathbf{Y} (\mathbf{\Psi}^* \mathbf{Y})^\dagger \mathbf{Z}^*.$$

“Generalized Nyström method” of Nakatsukasa. $O(k)$ columns in $\mathbf{\Omega}$ and $\mathbf{\Psi}$ sufficient.

Alternatively, draw $\mathbf{\Omega} \in \mathbb{R}^{n \times (k+10)}$ from a Gaussian distribution and set $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$. Then let $\mathbf{\Psi} \in \mathbb{R}^{n \times (k+10)}$ be a matrix whose columns form an ON basis for $\text{col}(\mathbf{Y})$. Then

$$\mathbf{A} \approx \mathbf{\Psi} (\mathbf{\Psi}^* \mathbf{A}) = \mathbf{\Psi} \mathbf{Z}^*.$$

(This is the “randomized SVD” in slight disguise.)

Approximation of rank structured matrices

Environment: We are given a rank structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ (to be precise, \mathbf{A} is HBS/HSS of rank k). We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Question: Can you construct thin matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ such that \mathbf{A} can be completely reconstructed in $O(N)$ work from the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ where $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Existing algorithms: Algorithms that “almost” resolve the task:

Case 1: Suppose that in addition to matvec, we can also evaluate individual entries of \mathbf{A} . Then an HBS (a.k.a. HSS) representation can be computed in $O(N)$ operations.

Very computationally efficient in practice — requires only $O(k)$ matvecs.

- P.G. Martinsson, SIMAX, **32**(4), 2011.
- Later improvements by Jianlin Xia, Sherry Li, etc.

Case 2: If all we have is the matvec, then we can still compute a rank-structured representation of \mathbf{A} using so called “peeling” algorithms. The price we have to pay is that we now need $O(k \times \log N)$ matvecs involving \mathbf{A} and \mathbf{A}^* .

The method is still fast in many situations, and does save messy coding work. For instance, without this black-box method, implementing the matrix-matrix multiplication, or changing the partition tree, are quite hard to implement efficiently.

- L. Lin, J. Lu, L. Ying, *Fast construction of hierarchical matrix representation from matrix-vector multiplication*, JCP 2011.
- P.G. Martinsson, SISC, **38**(4), pp. A1959-A1986, 2016.

Approximation of rank structured matrices

Environment: We are given a rank structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ (to be precise, \mathbf{A} is HBS/HSS of rank k). We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Question: Can you construct thin matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ such that \mathbf{A} can be completely reconstructed in $O(N)$ work from the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ where $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

New!

J. Levitt & P.G. Martinsson, arxiv arXiv:2205.02990, May 8, 2022.

A method that is:

- True “black box” – no matrix entry evaluation required.
- True $O(N)$ complexity – no log factors.

Approximation of rank-structured matrices – mini review of RSVD

Brief review of a randomized technique for compressing a matrix of globally low rank.

Objective: Given an $m \times n$ matrix \mathbf{A} , find an approximate rank- k partial SVD:

$$\begin{array}{ccccccc} \mathbf{A} & \approx & \mathbf{U} & \mathbf{D} & \mathbf{V}^* & & \\ m \times n & & m \times k & k \times k & k \times n & & \end{array}$$

where \mathbf{U} and \mathbf{V} are orthonormal, and \mathbf{D} is diagonal. (We assume $k \ll \min(m, n)$.)

Fix an over-sampling parameter p (say $p = 10$) and set $r = k + p$.

1. Draw an $n \times r$ Gaussian random matrix Ω .

`Omega = randn(n,r)`

2. Form the $m \times r$ sample matrix $\mathbf{Y} = \mathbf{A}\Omega$.

`Y = A * Omega`

3. Form an $m \times r$ orthonormal matrix \mathbf{Q} such that $\text{ran}(\mathbf{Q}) = \text{ran}(\mathbf{Y})$.

`[Q, ~] = qr(Y)`

4. Form the $r \times n$ matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.

`B = Q' * A`

5. Form the full SVD of the small matrix \mathbf{B} : $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.

`[Uhat, Sigma, V] = svd(B, 'econ')`

6. Form the matrix $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

`U = Q * Uhat(:, 1:k)`


7. Truncate \mathbf{D} and \mathbf{V} .

`D = D(1:k, 1:k), V = V(:, 1:k)`

Key point: If \mathbf{A} has ε -rank k , and $\Omega \in \mathbb{R}^{n \times (k+10)}$ is drawn from a Gaussian distribution, then the columns of $\mathbf{Y} = \mathbf{A}\Omega$ form a “good” basis for the column space of \mathbf{A} .

Approximation of rank-structured matrices – A binary tree structure

An example binary tree structure for a matrix of size 400×400 . The levels of the tree represent successively refined partitions of the index vector $[1, \dots, 400]$.



Level 0: $l_1 = [1, 2, \dots, 400]$

Level 1:

$l_2 = [1, \dots, 200]$, $l_3 = [201, \dots, 400]$

Level 2:


$l_4 = [1, \dots, 100]$, $l_5 = [101, \dots, 200]$,
 $l_6 = [201, \dots, 300]$, $l_7 = [301, \dots, 400]$

Let m denote the leaf node size.

Let $L \approx \log(N/m)$ denote the depth of the tree.

Approximation of rank-structured matrices – HBS (a.k.a. HSS) structure

Consider the following tessellation of a matrix, where each block represents interactions between two leaf nodes of the tree.



Approximation of rank-structured matrices – HBS (a.k.a. HSS) structure

HBS requirements for the finest level: for every leaf node τ , there must exist basis matrices \mathbf{U}_τ and \mathbf{V}_τ such that for every leaf node $\tau' \neq \tau$, we have

$$\mathbf{A}_{\tau,\tau'} = \mathbf{U}_\tau \tilde{\mathbf{A}}_{\tau,\tau'} \mathbf{V}_{\tau'}^*.$$

$m \times m \quad m \times k \quad k \times k \quad k \times m$

$\mathbf{A}_{4,4}$	$\mathbf{A}_{4,5}$	$\mathbf{A}_{4,6}$	$\mathbf{A}_{4,7}$
$\mathbf{A}_{5,4}$	$\mathbf{A}_{5,5}$	$\mathbf{A}_{5,6}$	$\mathbf{A}_{5,7}$
$\mathbf{A}_{6,4}$	$\mathbf{A}_{6,5}$	$\mathbf{A}_{6,6}$	$\mathbf{A}_{6,7}$
$\mathbf{A}_{7,4}$	$\mathbf{A}_{7,5}$	$\mathbf{A}_{7,6}$	$\mathbf{A}_{7,7}$

Approximation of rank-structured matrices – HBS (a.k.a. HSS) structure

HBS requirements for the finest level: for every leaf node τ , there must exist basis matrices \mathbf{U}_τ and \mathbf{V}_τ such that for every leaf node $\tau' \neq \tau$, we have

$$\mathbf{A}_{\tau,\tau'} = \mathbf{U}_\tau \tilde{\mathbf{A}}_{\tau,\tau'} \mathbf{V}_{\tau'}^*.$$

$m \times m \quad m \times k \quad k \times k \quad k \times m$

$\mathbf{A}_{4,4}$	$\mathbf{A}_{4,5}$	$\mathbf{A}_{4,6}$	$\mathbf{A}_{4,7}$
$\mathbf{A}_{5,4}$	$\mathbf{A}_{5,5}$	$\mathbf{A}_{5,6}$	$\mathbf{A}_{5,7}$
$\mathbf{A}_{6,4}$	$\mathbf{A}_{6,5}$	$\mathbf{A}_{6,6}$	$\mathbf{A}_{6,7}$
$\mathbf{A}_{7,4}$	$\mathbf{A}_{7,5}$	$\mathbf{A}_{7,6}$	$\mathbf{A}_{7,7}$

The on-diagonal blocks are not assumed to be low-rank, and pose the main challenge for black-box compression.

Approximation of rank-structured matrices – Telescoping factorization

This leads to a factorization of \mathbf{A} .

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

The diagram illustrates the telescoping factorization of matrix \mathbf{A} . Matrix \mathbf{A} is a 4x4 grid with red and blue blocks. Matrix $\mathbf{U}^{(L)}$ is a 4x4 grid with blue blocks on the diagonal. Matrix $\tilde{\mathbf{A}}^{(L)}$ is a 4x4 grid with red and blue blocks. Matrix $(\mathbf{V}^{(L)})^*$ is a 4x4 grid with blue blocks on the diagonal. Matrix $\mathbf{D}^{(L)}$ is a 4x4 grid with red blocks on the diagonal.

Approximation of rank-structured matrices – Telescoping factorization

This leads to a factorization of \mathbf{A} .

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

The diagram shows the factorization of matrix \mathbf{A} into four components: $\mathbf{U}^{(L)}$, $\tilde{\mathbf{A}}^{(L)}$, $(\mathbf{V}^{(L)})^*$, and $\mathbf{D}^{(L)}$. Each matrix is represented as a 4x4 grid of colored blocks. \mathbf{A} has red blocks at (1,1), (2,2), (3,3), and (4,4), and blue blocks elsewhere. $\mathbf{U}^{(L)}$ has blue blocks on the diagonal. $\tilde{\mathbf{A}}^{(L)}$ has red blocks at (1,1), (1,2), (2,1), and (2,2), and blue blocks elsewhere. $(\mathbf{V}^{(L)})^*$ has blue blocks on the diagonal. $\mathbf{D}^{(L)}$ has red blocks on the diagonal.

$\tilde{\mathbf{A}}^{(L)}$ is also an HBS matrix, and it can be factorized similarly, leading to a *telescoping factorization*.

Approximation of rank-structured matrices – Telescoping factorization


This leads to a factorization of \mathbf{A} .

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

$\tilde{\mathbf{A}}^{(L)}$ is also an HBS matrix, and it can be factorized similarly, leading to a *telescoping factorization*.

For example, a factorization of an HBS matrix with a tree of depth $L = 3$ takes the form

$$\mathbf{A} = \mathbf{U}^{(3)} (\mathbf{U}^{(2)} (\mathbf{U}^{(1)} \mathbf{D}^{(0)} (\mathbf{V}^{(1)})^* + \mathbf{D}^{(1)}) (\mathbf{V}^{(2)})^* + \mathbf{D}^{(2)}) (\mathbf{V}^{(3)})^* + \mathbf{D}^{(3)}.$$




Approximation of rank-structured matrices – A naive approach

Consider the task of finding basis matrix \mathbf{U}_4 for node 4 using randomized sampling. We seek a sample of $\mathbf{A}(I_4, I_4^C)$, the HBS row block of node 4.

Approximation of rank-structured matrices – A naive approach

Consider the task of finding basis matrix \mathbf{U}_4 for node 4 using randomized sampling. We seek a sample of $\mathbf{A}(I_4, I_4^c)$, the HBS row block of node 4.


The naive approach is to sample with a random matrix $\mathbf{\Omega} \in \mathbb{R}^{N \times r}$, $r = k + 10$, that has a block of zeros in rows indexed by I_4 . Then $\mathbf{Y}(I_4, :)$ will contain a sample of $\mathbf{A}(I_4, I_4^c)$.



Approximation of rank-structured matrices – A naive approach

Consider the task of finding basis matrix \mathbf{U}_4 for node 4 using randomized sampling. We seek a sample of $\mathbf{A}(I_4, I_4^C)$, the HBS row block of node 4.


The naive approach is to sample with a random matrix $\mathbf{\Omega} \in \mathbb{R}^{N \times r}$, $r = k + 10$, that has a block of zeros in rows indexed by I_4 . Then $\mathbf{Y}(I_4, :)$ will contain a sample of $\mathbf{A}(I_4, I_4^C)$.



This scheme requires taking a separate set of r samples for each leaf node, for a total of $\sim rN/m$ samples. There is a lot of wasted information in \mathbf{Y} .

Approximation of rank-structured matrices – the “not quite black box” approach

Sample \mathbf{A} with a completely dense random matrix $\Omega \in \mathbb{R}^{N \times r}$.



Afterwards, subtract unwanted contributions back out of \mathbf{Y} .


Subtracting the contribution of $\mathbf{A}_{4,4}$ gives the desired sample of $\mathbf{A}(I_4, I_4^c)$,

$$\mathbf{Y}(I_4, :) - \mathbf{A}_{4,4} \Omega(I_4, :).$$

This scheme requires only r samples in total, but it also requires direct access to a small number of entries of \mathbf{A} .

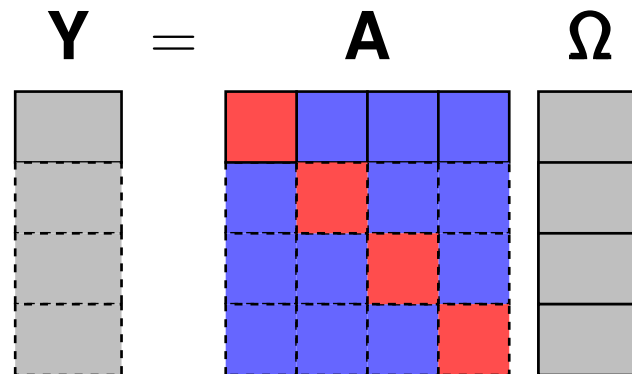
Approximation of rank-structured matrices – Finding \mathbf{U}

Sample \mathbf{A} with a completely dense random matrix $\Omega \in \mathbb{R}^{N \times (r+m)}$, where m is the leaf node size. (Think $m \approx 2k$.)



Approximation of rank-structured matrices – Finding \mathbf{U}

Sample \mathbf{A} with a completely dense random matrix $\Omega \in \mathbb{R}^{N \times (r+m)}$, where m is the leaf node size. (Think $m \approx 2k$.)

$$\mathbf{Y} = \mathbf{A} \Omega$$


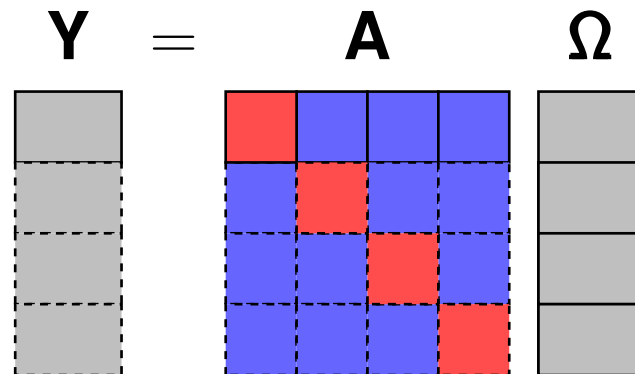
Since $\Omega(l_4, :)$ is of size $m \times (r+m)$, it has a nullspace of dimension at least r . Let

$$\mathbf{Q}_4 = \text{nullspace}(\Omega(l_4, :), r)$$

be an $(r+m) \times r$ orthonormal basis of the nullspace of $\Omega(l_4, :)$.

Approximation of rank-structured matrices – Finding U

Sample \mathbf{A} with a completely dense random matrix $\Omega \in \mathbb{R}^{N \times (r+m)}$, where m is the leaf node size. (Think $m \approx 2k$.)

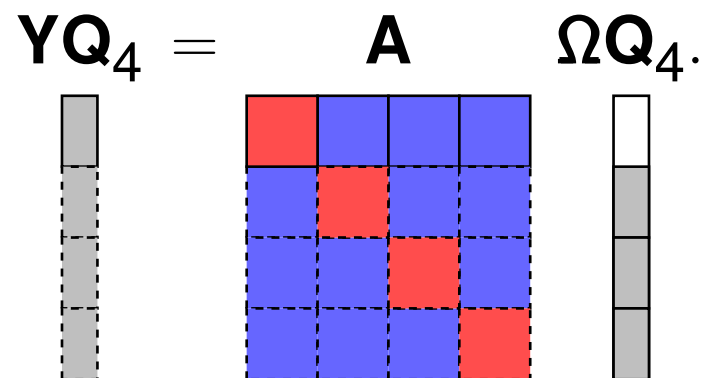
$$\mathbf{Y} = \mathbf{A} \Omega$$


Since $\Omega(l_4, :)$ is of size $m \times (r+m)$, it has a nullspace of dimension at least r . Let

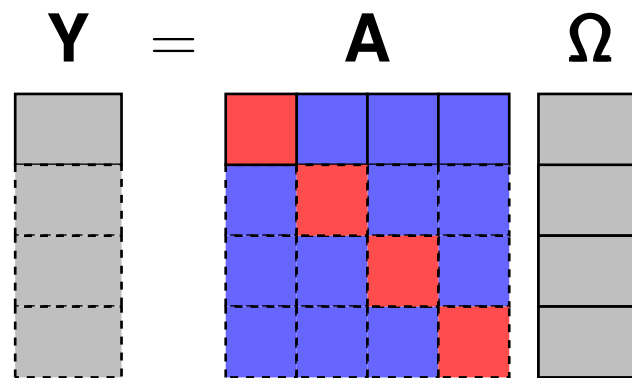
$$\mathbf{Q}_4 = \text{nullspace}(\Omega(l_4, :), r)$$

be an $(r+m) \times r$ orthonormal basis of the nullspace of $\Omega(l_4, :)$.

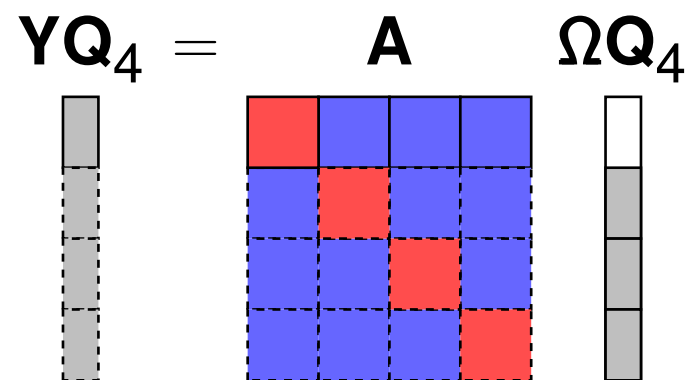
Then

$$\mathbf{YQ}_4 = \mathbf{A} \Omega \mathbf{Q}_4$$


Approximation of rank-structured matrices – Finding \mathbf{U}

$$\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$$


$$\mathbf{Q}_4 = \text{nullspace}(\mathbf{\Omega}(I_4, :), r)$$

$$\mathbf{YQ}_4 = \mathbf{A} \mathbf{\Omega Q}_4$$


We find the sample by multiplying $\mathbf{Y}(I_4, :)\mathbf{Q}_4$.

Orthonormalizing the sample gives basis matrix \mathbf{U}_4 ,

$$\mathbf{U}_4 = \text{qr}(\mathbf{Y}(I_4, :)\mathbf{Q}_4).$$

Approximation of rank-structured matrices – Finding \mathbf{U}

- For each leaf node τ , we compute

$$\mathbf{Q}_\tau = \text{nullspace}(\mathbf{\Omega}(I_\tau, :), r)$$

$$\mathbf{U}_\tau = \text{qr}(\mathbf{Y}(I_\tau, :)\mathbf{Q}_\tau).$$

- \mathbf{U}_τ only depends on $\mathbf{\Omega}(I_\tau, :)$ and $\mathbf{Y}(I_\tau, :)$.
- We only need $r + m$ samples to find \mathbf{U}_τ for every leaf node τ .
- $\mathbf{\Omega}\mathbf{Q}_\tau$ is a Gaussian random matrix, except for the block intentionally zeroed out.

Approximation of rank-structured matrices – Compression overview

Recall the telescoping factorization $\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$.

Steps:

1. Find $\mathbf{U}^{(L)}, \mathbf{V}^{(L)}$.
2. Find $\mathbf{D}^{(L)}$.
3. Compress $\tilde{\mathbf{A}}^{(L)}$ recursively.

Compute randomized samples of \mathbf{A} and \mathbf{A}^ .*

- 1: Form Gaussian random matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ of size $N \times s$.
- 2: Multiply $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$.

Compress level by level from finest to coarsest.

- 3: **for** level $\ell = L, L - 1, \dots, 0$ **do**
- 4: **for** node τ in level ℓ **do**
- 5: **if** τ is a leaf node **then**
- 6: $\mathbf{\Omega}_\tau = \mathbf{\Omega}(I_\tau, :), \quad \mathbf{\Psi}_\tau = \mathbf{\Psi}(I_\tau, :)$
- 7: $\mathbf{Y}_\tau = \mathbf{Y}(I_\tau, :), \quad \mathbf{Z}_\tau = \mathbf{Z}(I_\tau, :)$
- 8: **else**
- 9: Let α and β denote the children of τ .
- 10: $\mathbf{\Omega}_\tau = \begin{bmatrix} \mathbf{V}_\alpha^* \mathbf{\Omega}_\alpha \\ \mathbf{V}_\beta^* \mathbf{\Omega}_\beta \end{bmatrix}, \quad \mathbf{\Psi}_\tau = \begin{bmatrix} \mathbf{U}_\alpha^* \mathbf{\Psi}_\alpha \\ \mathbf{U}_\beta^* \mathbf{\Psi}_\beta \end{bmatrix}$
- 11: $\mathbf{Y}_\tau = \begin{bmatrix} \mathbf{U}_\alpha^* (\mathbf{Y}_\alpha - \mathbf{D}_\alpha \mathbf{\Omega}_\alpha) \\ \mathbf{U}_\beta^* (\mathbf{Y}_\beta - \mathbf{D}_\beta \mathbf{\Omega}_\beta) \end{bmatrix}, \quad \mathbf{Z}_\tau = \begin{bmatrix} \mathbf{V}_\alpha^* (\mathbf{Z}_\alpha - \mathbf{D}_\alpha^* \mathbf{\Psi}_\alpha) \\ \mathbf{V}_\beta^* (\mathbf{Z}_\beta - \mathbf{D}_\beta^* \mathbf{\Psi}_\beta) \end{bmatrix}$
- 12: **if** level $\ell > 0$ **then**
- 13: $\mathbf{Q}_\tau = \text{nullspace}(\mathbf{\Omega}_\tau, r), \quad \mathbf{P}_\tau = \text{nullspace}(\mathbf{\Psi}_\tau, r)$
- 14: $\mathbf{U}_\tau = \text{qr}(\mathbf{Y}_\tau \mathbf{Q}_\tau, r), \quad \mathbf{V}_\tau = \text{qr}(\mathbf{Z}_\tau \mathbf{P}_\tau, r)$
- 15: $\mathbf{D}_\tau = (\mathbf{I} - \mathbf{U}_\tau \mathbf{U}_\tau^*) \mathbf{Y}_\tau \mathbf{\Omega}_\tau^\dagger + \mathbf{U}_\tau \mathbf{U}_\tau^* ((\mathbf{I} - \mathbf{V}_\tau \mathbf{V}_\tau^*) \mathbf{Z}_\tau \mathbf{\Psi}_\tau^\dagger)^*$
- 16: **else**
- 17: $\mathbf{D}_\tau = \mathbf{Y}_\tau \mathbf{\Omega}_\tau^\dagger$

Approximation of rank-structured matrices – Finding \mathbf{D}

From the telescoping factorization

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

we define $\tilde{\mathbf{A}}^{(L)}$ and $\mathbf{D}^{(L)}$ as follows.

Approximation of rank-structured matrices – Finding \mathbf{D}

From the telescoping factorization

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

we define $\tilde{\mathbf{A}}^{(L)}$ and $\mathbf{D}^{(L)}$ as follows.

$$\mathbf{A} = \mathbf{U}^{(L)} \overbrace{(\mathbf{U}^{(L)})^* \mathbf{A} \mathbf{V}^{(L)}}^{\tilde{\mathbf{A}}^{(L)}} (\mathbf{V}^{(L)})^* + \overbrace{\mathbf{A} - \mathbf{U}^{(L)} (\mathbf{U}^{(L)})^* \mathbf{A} \mathbf{V}^{(L)} (\mathbf{V}^{(L)})^*}_{\mathbf{D}^{(L)}}$$

Approximation of rank-structured matrices – Finding \mathbf{D}

From the telescoping factorization

$$\mathbf{A} = \mathbf{U}^{(L)} \tilde{\mathbf{A}}^{(L)} (\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)}$$

we define $\tilde{\mathbf{A}}^{(L)}$ and $\mathbf{D}^{(L)}$ as follows.

$$\mathbf{A} = \mathbf{U}^{(L)} \overbrace{(\mathbf{U}^{(L)})^* \mathbf{A} \mathbf{V}^{(L)} (\mathbf{V}^{(L)})^*}^{\tilde{\mathbf{A}}^{(L)}} + \overbrace{\mathbf{A} - \mathbf{U}^{(L)} (\mathbf{U}^{(L)})^* \mathbf{A} \mathbf{V}^{(L)} (\mathbf{V}^{(L)})^*}^{\mathbf{D}^{(L)}}$$

Block \mathbf{D}_τ of $\mathbf{D}^{(L)}$ is given by

$$\begin{aligned} \mathbf{D}_\tau &= \mathbf{A}_{\tau,\tau} - \mathbf{U}_\tau \mathbf{U}_\tau^* \mathbf{A}_{\tau,\tau} \mathbf{V}_\tau \mathbf{V}_\tau^* \\ &= \dots \\ &= (\mathbf{I} - \mathbf{U}_\tau \mathbf{U}_\tau^*) \mathbf{Y}_\tau \boldsymbol{\Omega}_\tau^\dagger + \mathbf{U}_\tau \mathbf{U}_\tau^* \left((\mathbf{I} - \mathbf{V}_\tau \mathbf{V}_\tau^*) \mathbf{Z}_\tau \boldsymbol{\Psi}_\tau^\dagger \right)^* \end{aligned}$$

Approximation of rank-structured matrices – Compressing $\tilde{\mathbf{A}}^{(L)}$

To compute randomized samples of $\tilde{\mathbf{A}}^{(L)}$, we multiply the telescoping factorization with Ω to obtain

$$\mathbf{Y} = \mathbf{A}\Omega = (\mathbf{U}^{(L)}\tilde{\mathbf{A}}^{(L)}(\mathbf{V}^{(L)})^* + \mathbf{D}^{(L)})\Omega,$$


and rearrange to obtain

$$\underbrace{(\mathbf{U}^{(L)})^*(\mathbf{Y} - \mathbf{D}^{(L)}\Omega)}_{\text{sample matrix}} = \tilde{\mathbf{A}}^{(L)} \underbrace{(\mathbf{V}^{(L)})^*\Omega}_{\text{test matrix}}.$$

Approximation of rank-structured matrices: Sparse LU

Let \mathbf{C} be the stiffness matrix for the standard five-point stencil finite difference approximation to the Poisson equation on a rectangular grid.

We partition the grid as shown and tessellate \mathbf{C} accordingly.



$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{0} & \mathbf{C}_{13} \\ \mathbf{0} & \mathbf{C}_{22} & \mathbf{C}_{23} \\ \mathbf{C}_{31} & \mathbf{C}_{32} & \mathbf{C}_{33} \end{bmatrix}$$


The matrix we seek to compress is the Schur complement

$$\mathbf{A} = \mathbf{C}_{33} - \mathbf{C}_{31}\mathbf{C}_{11}^{-1}\mathbf{C}_{31} - \mathbf{C}_{32}\mathbf{C}_{22}^{-1}\mathbf{C}_{23}.$$

Approximation of rank-structured matrices: Sparse LU

$r = 30, m = 60$


Time for matrix compression



Approximation of rank-structured matrices: Sparse LU


$r = 30, m = 60$

Time for matrix-vector multiplication



Approximation of rank-structured matrices: Sparse LU


$r = 30, m = 60$



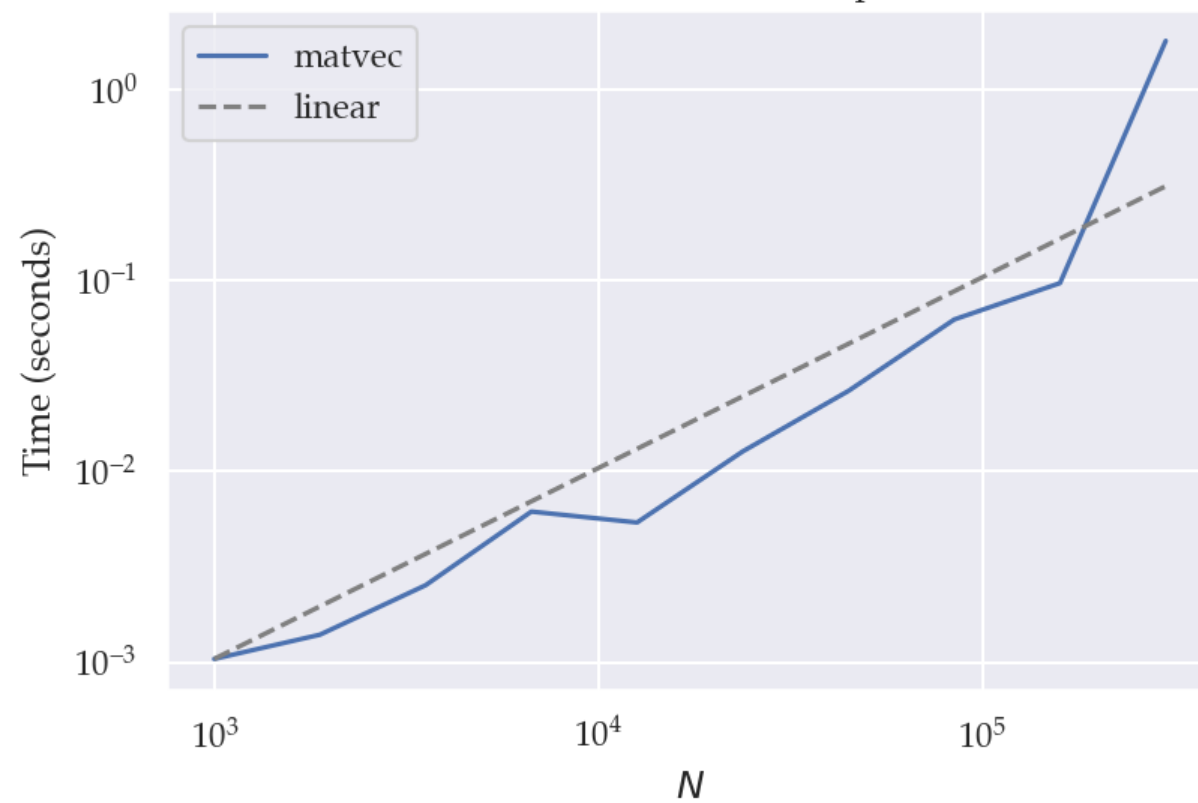
Approximation of rank-structured matrices: FMM

$r = 50, m = 100$

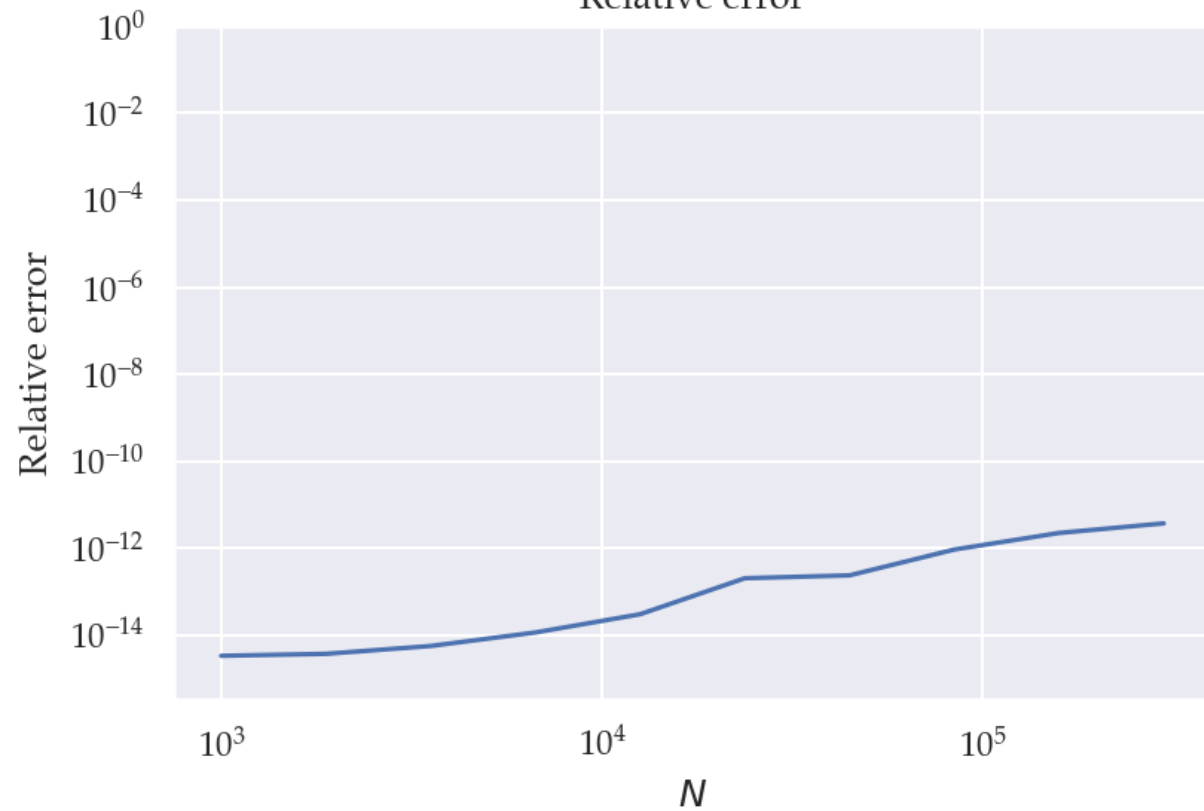
Time for matrix compression




Time for matrix-vector multiplication



Relative error

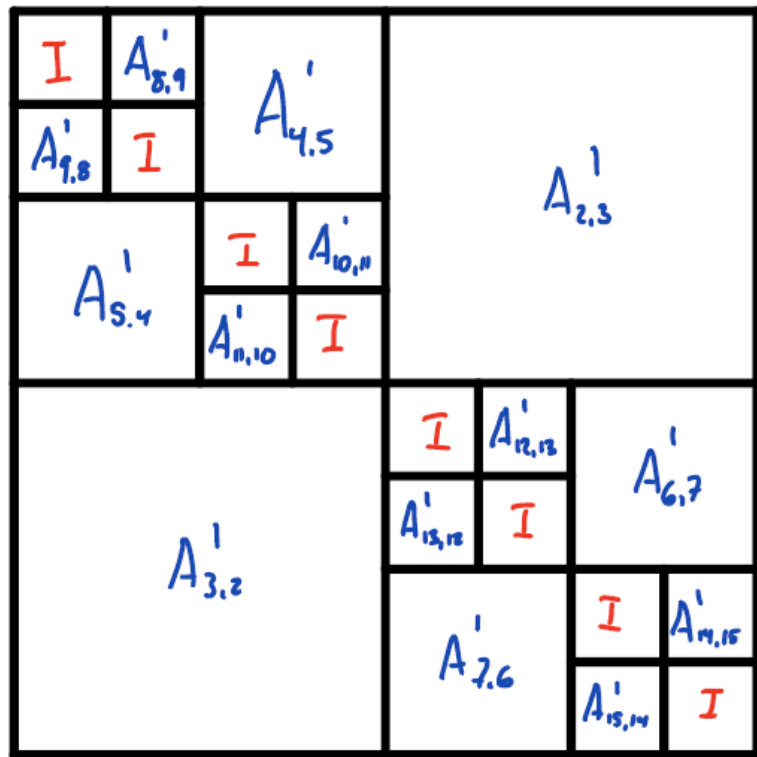


Memory requirement / N



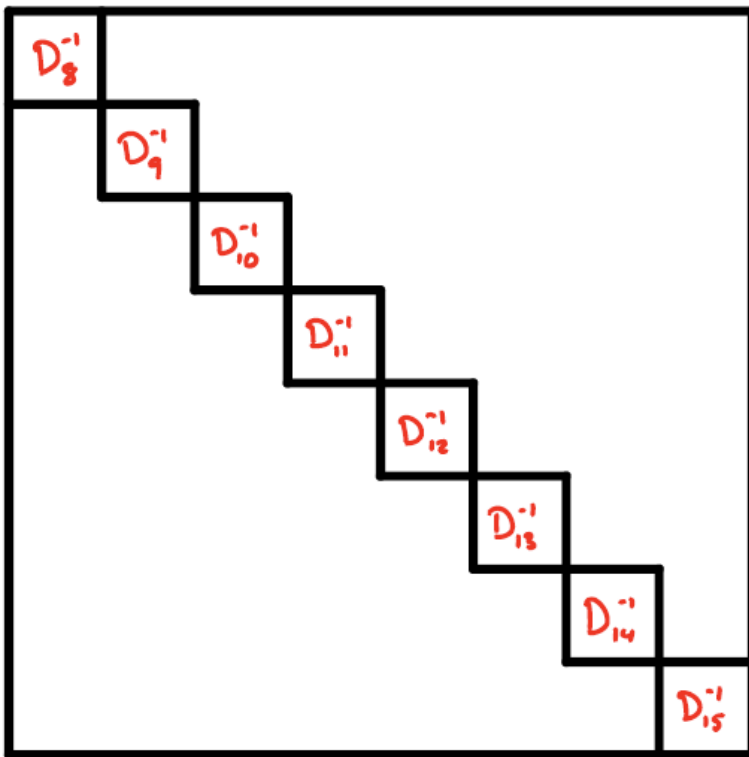
Approximation of rank-structured matrices: Key points

- Fully “black box”. Interacts with \mathbf{A} only via the matvec.
- True linear complexity. Requires only $O(k)$ samples from \mathbf{A} and \mathbf{A}^* .
Much faster in practice than existing black box algorithms.
(However, prefactor in # samples is slightly suboptimal – unlike Townsend/Halikias.)
- Ideal tool for acceleration of sparse direct solvers.




A_3

=



B_3




A

Key Ideas:

The solution operator of a linear elliptic PDE is “friendly.”

- Smoothing.
- Stable.





Key Ideas:

The solution operator of a linear elliptic PDE is “friendly.”

- Smoothing.
- Stable.

Long range interactions are low rank.

- Cf. St Venant principle, multipole expansions, etc.
- Smoothness is *not* necessary.
- Numerical compression is essential.
- Wave problems with small λ remain challenging.



Key Ideas:

The solution operator of a linear elliptic PDE is “friendly.”

- Smoothing.
- Stable.



Long range interactions are low rank.

- Cf. St Venant principle, multipole expansions, etc.
- Smoothness is *not* necessary.
- Numerical compression is essential.
- Wave problems with small λ remain challenging.

High-order discretizations + FDS.

- Maximize the work done by each DOF to save memory.
- Perfect for ill-conditioned problems with oscillatory solutions.
- Requires care in choosing discretization scheme.

New randomized methods for matrix algebra \rightarrow acceleration & simplification.



Where we are now:

Postdoc position(s) available!

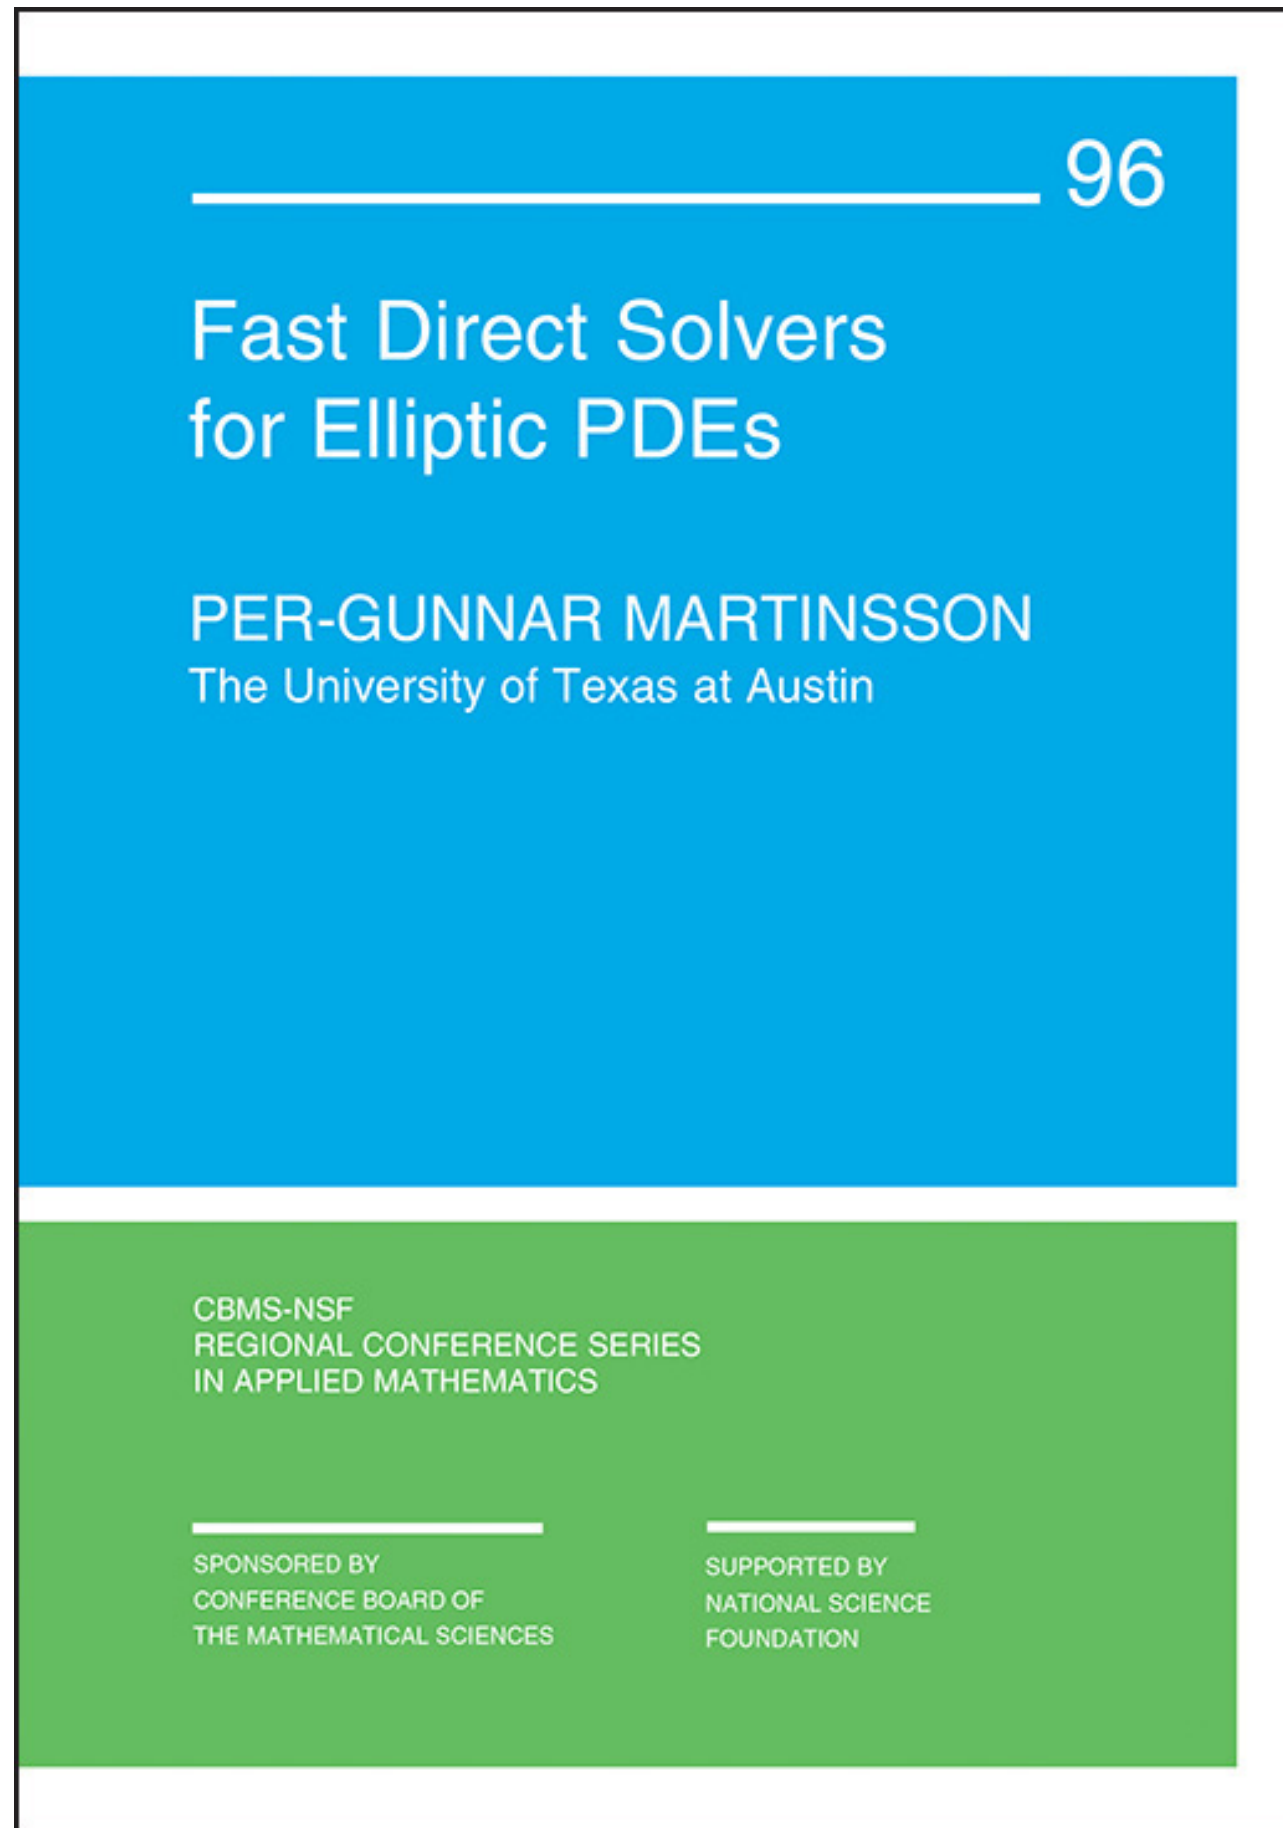
- We have developed direct solvers with $O(N)$ complexity for elliptic PDEs with non-oscillatory (or “mildly oscillatory”) solutions for most standard environments:
 - Sparse matrices from FEM/FD/composite spectral/... in both 2D and 3D.
 - Boundary integral equations in 2D and 3D. (Work in progress ...)
- Advantages of direct solvers:
 - Often instantaneous solves once a solution operator has been built.
 - Can eliminate problems with slow convergence of iterative solvers.
 - Communication efficient.
- Disadvantages of direct solvers:
 - Memory hogs. (But distributed memory is OK.)
 - The build stage is still slow for many 3D problems. (I am optimistic that we will fix this!)

Where to go next:

New powerful tool available → lots of opportunities!

- Explore happy couplings:
 - Direct solver + high order discretization. *(Helps with memory. Wave problems.)*
 - Direct solver + integral equation formulations. *(Need dense matrices anyway.)*
 - Direct solver + parallelization. *(Root of tree is cheap!)*
 - Direct solver + numerical coarse graining. *(Another talk...)*
- Parabolic and hyperbolic problems. Parallel-in-time methods?

More details in a 2019 monograph:



96

Fast Direct Solvers
for Elliptic PDEs

PER-GUNNAR MARTINSSON
The University of Texas at Austin

CBMS-NSF
REGIONAL CONFERENCE SERIES
IN APPLIED MATHEMATICS

SPONSORED BY
CONFERENCE BOARD OF
THE MATHEMATICAL SCIENCES

SUPPORTED BY
NATIONAL SCIENCE
FOUNDATION