Randomized algorithms for very large scale linear algebra

Gunnar Martinsson

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

University of Texas at Austin

Collaborators: Robert van de Geijn, Abinand Gopal, Francisco Igual, Yuji Nakatsukasa, Gregorio Quintana-Ortí, Vladimir Rokhlin, Joel Tropp, Heather Wilber.

Current and recent students/postdocs: Joar Bagge, Yunhui Cai, Chao Chen, Ke Chen, Simon Dirckx, Yijun Dong, Michael Han, Nathan Heavner, Joseph Kump, James Levitt, Kate Pearce, Anna Yesypenko.

Research support by:

Outline

- **High accuracy methods** *Main topic.*
  - Low rank approximation: The randomized SVD. *Key success story.*
  - Probabilistic performance analysis and connections to random matrix theory.
  - Streaming and single-pass methods.
  - Structured random maps ("fast Johnson-Lindenstrauss transforms").
  - Linear solvers: Sketch-to-solve vs. Sketch-to-precondition.

- **Sampling based methods for huge problems** *Very brief!*
  - Monte Carlo style methods $\rightarrow$ less reliable, less accurate, less robust.
  - Enable the solution of stupendously large problems that would otherwise be intractable.
  - Resolved some long-standing open theoretical questions in linear algebra.

- **Randomized methods for differential and integral operators** *Time permitting . . .*
  - Approximation of global operators of mathematical physics (solution operators, DtNs, . . .).
  - Tools for matrices that are not of global low rank, but have structure that can be exploited.
  - Conceptually related to Fast Multipole Methods, Fast Direct Solvers, Calderón-Zygmund theory . . .

Low rank approximation:

Let **A** be a given $m \times n$ matrix, and let k be an integer such that $1 \leq k \ll \min(m, n)$.

We seek to compute approximate factors **E** and **F** such that

$$\underset{m \times n}{\mathbf{A}} \approx \underset{m \times k}{\mathbf{E}} \; \underset{k \times n}{\mathbf{F}^*}.$$

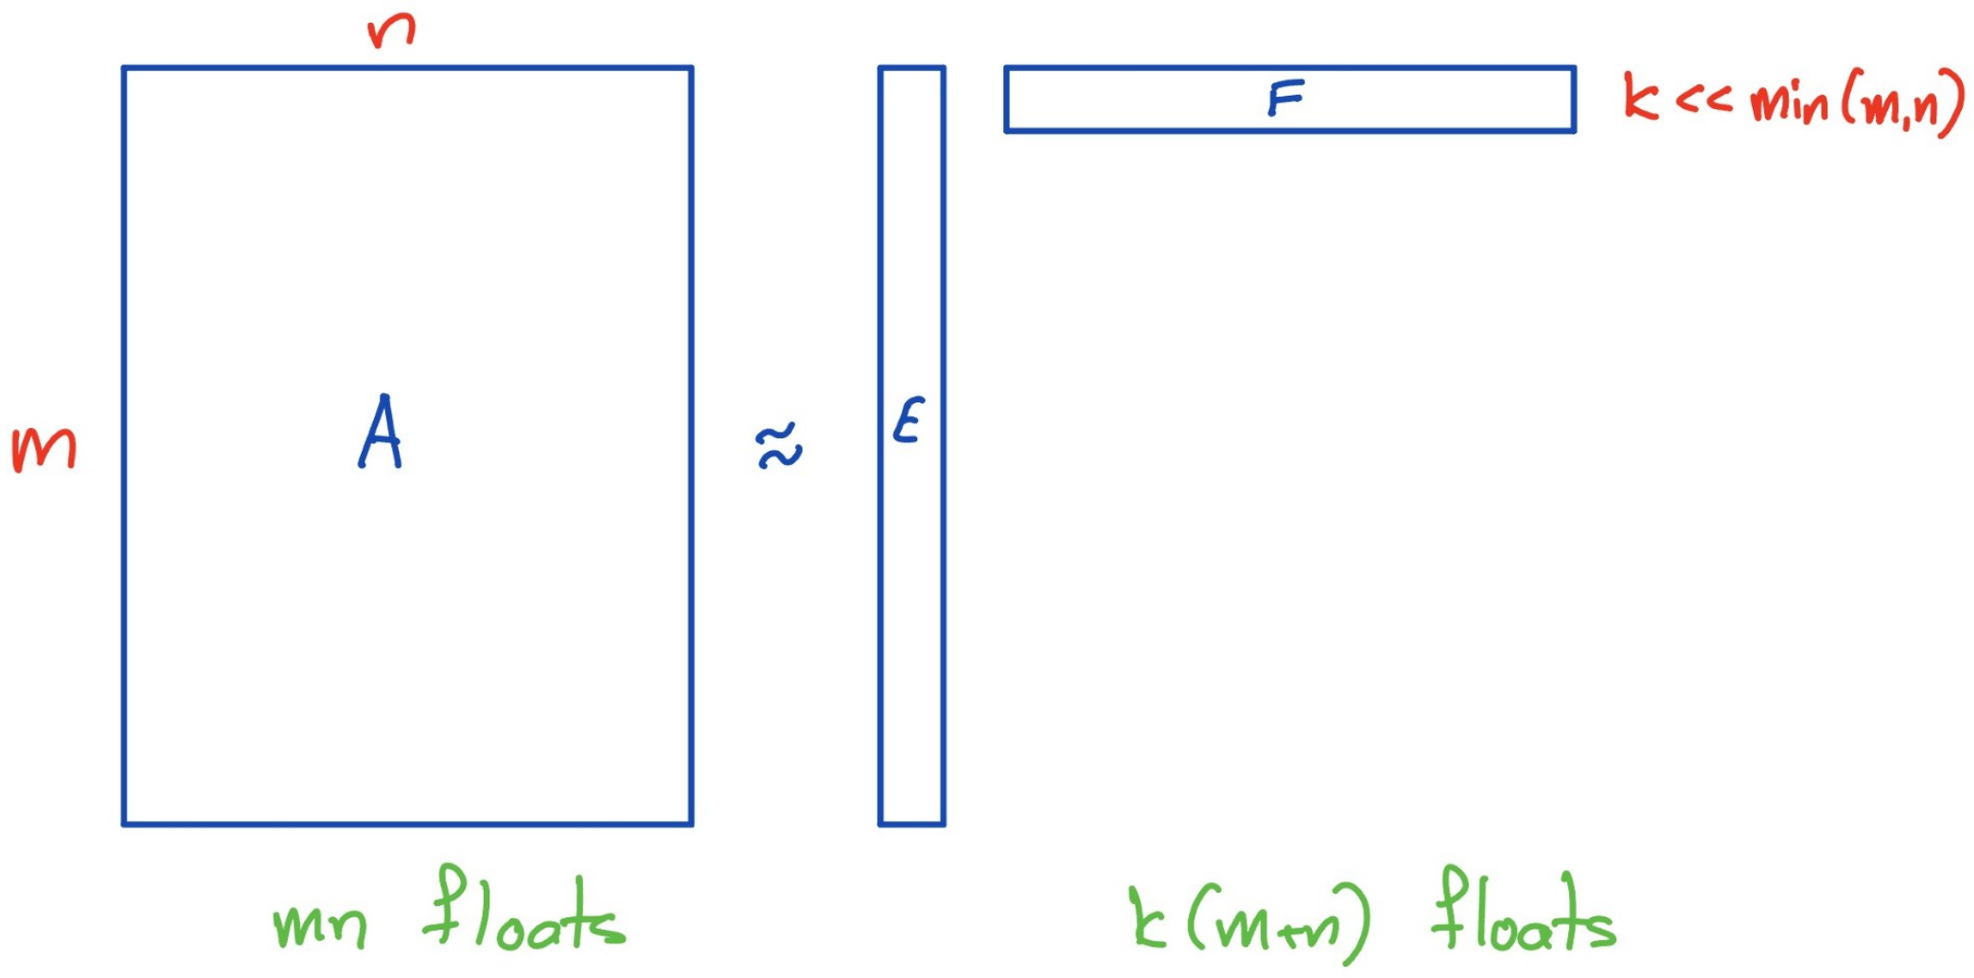

Low rank approximation:

Let **A** be a given $m \times n$ matrix, and let k be an integer such that $1 \leq k \ll \min(m, n)$.
We seek to compute approximate factors **E** and **F** such that

$$\underset{m \times n}{\mathbf{A}} \quad \approx \quad \underset{m \times k}{\mathbf{E}} \quad \underset{k \times n}{\mathbf{F}^*}.$$

Why?

- Fitting a hyperplane to a given set of points. Or fitting a multivariate normal distribution to measurements ("principal component analysis").
- Model reduction in scientific computing.
- Spectral algorithms in data analysis.
- "Fast" algorithms of various types: Fast Multipole Methods, generalizations of the Fast Fourier Transform, fast direct solvers, etc.
- Many, many, many more.

Note: We seek only to control the residual error $\|\mathbf{A} - \mathbf{E}\mathbf{F}^*\|$.

Low rank approximation:

Excellent algorithms for numerical low rank approximation already exist:

- **Compute the full SVD and truncate**
  - Eckart-Young theorem tells us this is optimal.
  - Algorithms are stable, accurate, and reliable. Routinely yield double precision accuracy.
  - Algorithms must be iterative, but converge *extremely* fast in practice.
  - Cost scales poorly with matrix size, $O(n^3)$ for $n \times n$ matrix. Hard to parallelize.
- **Krylov methods** For instance, $\text{ran}(\mathbf{Q}) = \text{ran}\big(\big[\mathbf{b},\, \mathbf{Ab},\, \mathbf{A}^2\mathbf{b},\, \ldots,\, \mathbf{A}^{k-1}\mathbf{b}\big]\big)$.
  - Particularly efficient when fast matrix-vector multiplication is available. (E.g. $\mathbf{A}$ is sparse.)
  - If the metric is number of matvecs, then Krylov methods are often "optimal".
  - Well understood – precise theory is available in many (but far from all!) cases.
- **Gram-Schmidt (with "column pivoting")**
  - Very simple!
  - If column pivoting is used, it is robust and typically works well. (Some intricacies arise.)
  - Can be stopped after k steps if required tolerance is met. $\rightarrow$ Complexity $O(mnk)$.
  - Communication intensive, so somewhat slow in practice. Does not preserve sparsity.
  - Output is quite far from optimal when singular values decay slowly.

Randomized algorithms for low rank approximation:

Problem: Given an $m \times n$ matrix $\mathbf{A}$, and a target rank k, where $k \ll \min(m, n)$, we seek to compute an approximate partial singular value decomposition:

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

with $\mathbf{U}$ and $\mathbf{V}$ having orthonormal columns, and $\mathbf{D}$ diagonal.

Randomized algorithms for low rank approximation:

Problem: Given an $m \times n$ matrix $\mathbf{A}$, and a target rank k, where $k \ll \min(m, n)$, we seek to compute an approximate partial singular value decomposition:

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

with $\mathbf{U}$ and $\mathbf{V}$ having orthonormal columns, and $\mathbf{D}$ diagonal.

Solution:

1. Draw an $n \times k$ Gaussian random matrix $\mathbf{G}$. `G = randn(n,k)`
2. Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AG}$. `Y = A * G`
3. Form an $m \times k$ orthonormal matrix $\mathbf{Q}$ s. t. $\text{col}(\mathbf{Y}) = \text{col}(\mathbf{Q})$. `[Q, ~] = qr(Y)`
4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^*\mathbf{A}$. `B = Q' * A`
5. Compute the SVD of $\mathbf{B}$ (small!): $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$. `[Uhat, Sigma, V] = svd(B,'econ')`
6. Form the matrix $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$. `U = Q * Uhat`

Why does it work? When $\mathbf{A}$ has exact rank k, the algorithm succeeds with probability 1.

Randomized algorithms for low rank approximation:

Problem: Given an $m \times n$ matrix $\mathbf{A}$, and a target rank k, where $k \ll \min(m, n)$, we seek to compute an approximate partial singular value decomposition:

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

with $\mathbf{U}$ and $\mathbf{V}$ having orthonormal columns, and $\mathbf{D}$ diagonal.

Solution:

1. Draw an $n \times k$ Gaussian random matrix $\mathbf{G}$. `G = randn(n,k)`
2. Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AG}$. `Y = A * G`
3. Form an $m \times k$ orthonormal matrix $\mathbf{Q}$ s. t. $\text{col}(\mathbf{Y}) = \text{col}(\mathbf{Q})$. `[Q, ~] = qr(Y)`
4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^*\mathbf{A}$. `B = Q' * A`
5. Compute the SVD of $\mathbf{B}$ (small!): $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$. `[Uhat, Sigma, V] = svd(B,'econ')`
6. Form the matrix $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$. `U = Q * Uhat`

Why does it work? When $\mathbf{A}$ has exact rank k, the algorithm succeeds with probability 1. In the general case, it works very well when the singular values decay rapidly. If they decay slowly, accuracy deteriorates.

Randomized algorithms for low rank approximation:

Problem: Given an $m \times n$ matrix $\mathbf{A}$, and a target rank k, where $k \ll \min(m, n)$, we seek to compute an approximate partial singular value decomposition:

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

with $\mathbf{U}$ and $\mathbf{V}$ having orthonormal columns, and $\mathbf{D}$ diagonal.

Solution:

1. Draw an $n \times k$ Gaussian random matrix $\mathbf{G}$. `G = randn(n,k)`
2. Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AG}$. `Y = A * G`
3. Form an $m \times k$ orthonormal matrix $\mathbf{Q}$ s. t. $\text{col}(\mathbf{Y}) = \text{col}(\mathbf{Q})$. `[Q, ~] = qr(Y)`
4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^*\mathbf{A}$. `B = Q' * A`
5. Compute the SVD of $\mathbf{B}$ (small!): $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$. `[Uhat, Sigma, V] = svd(B,'econ')`
6. Form the matrix $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$. `U = Q * Uhat`

Power iteration: When the singular values of $\mathbf{A}$ decay slowly, precision can be improved by replacing the formula $\mathbf{Y} = \mathbf{AG}$ on line 2 by $\mathbf{Y} = \mathbf{A}(\mathbf{A}^*\mathbf{G})$, or $\mathbf{Y} = \mathbf{A}(\mathbf{A}^*(\mathbf{AG}))$, or ...

Randomized low rank approximation:

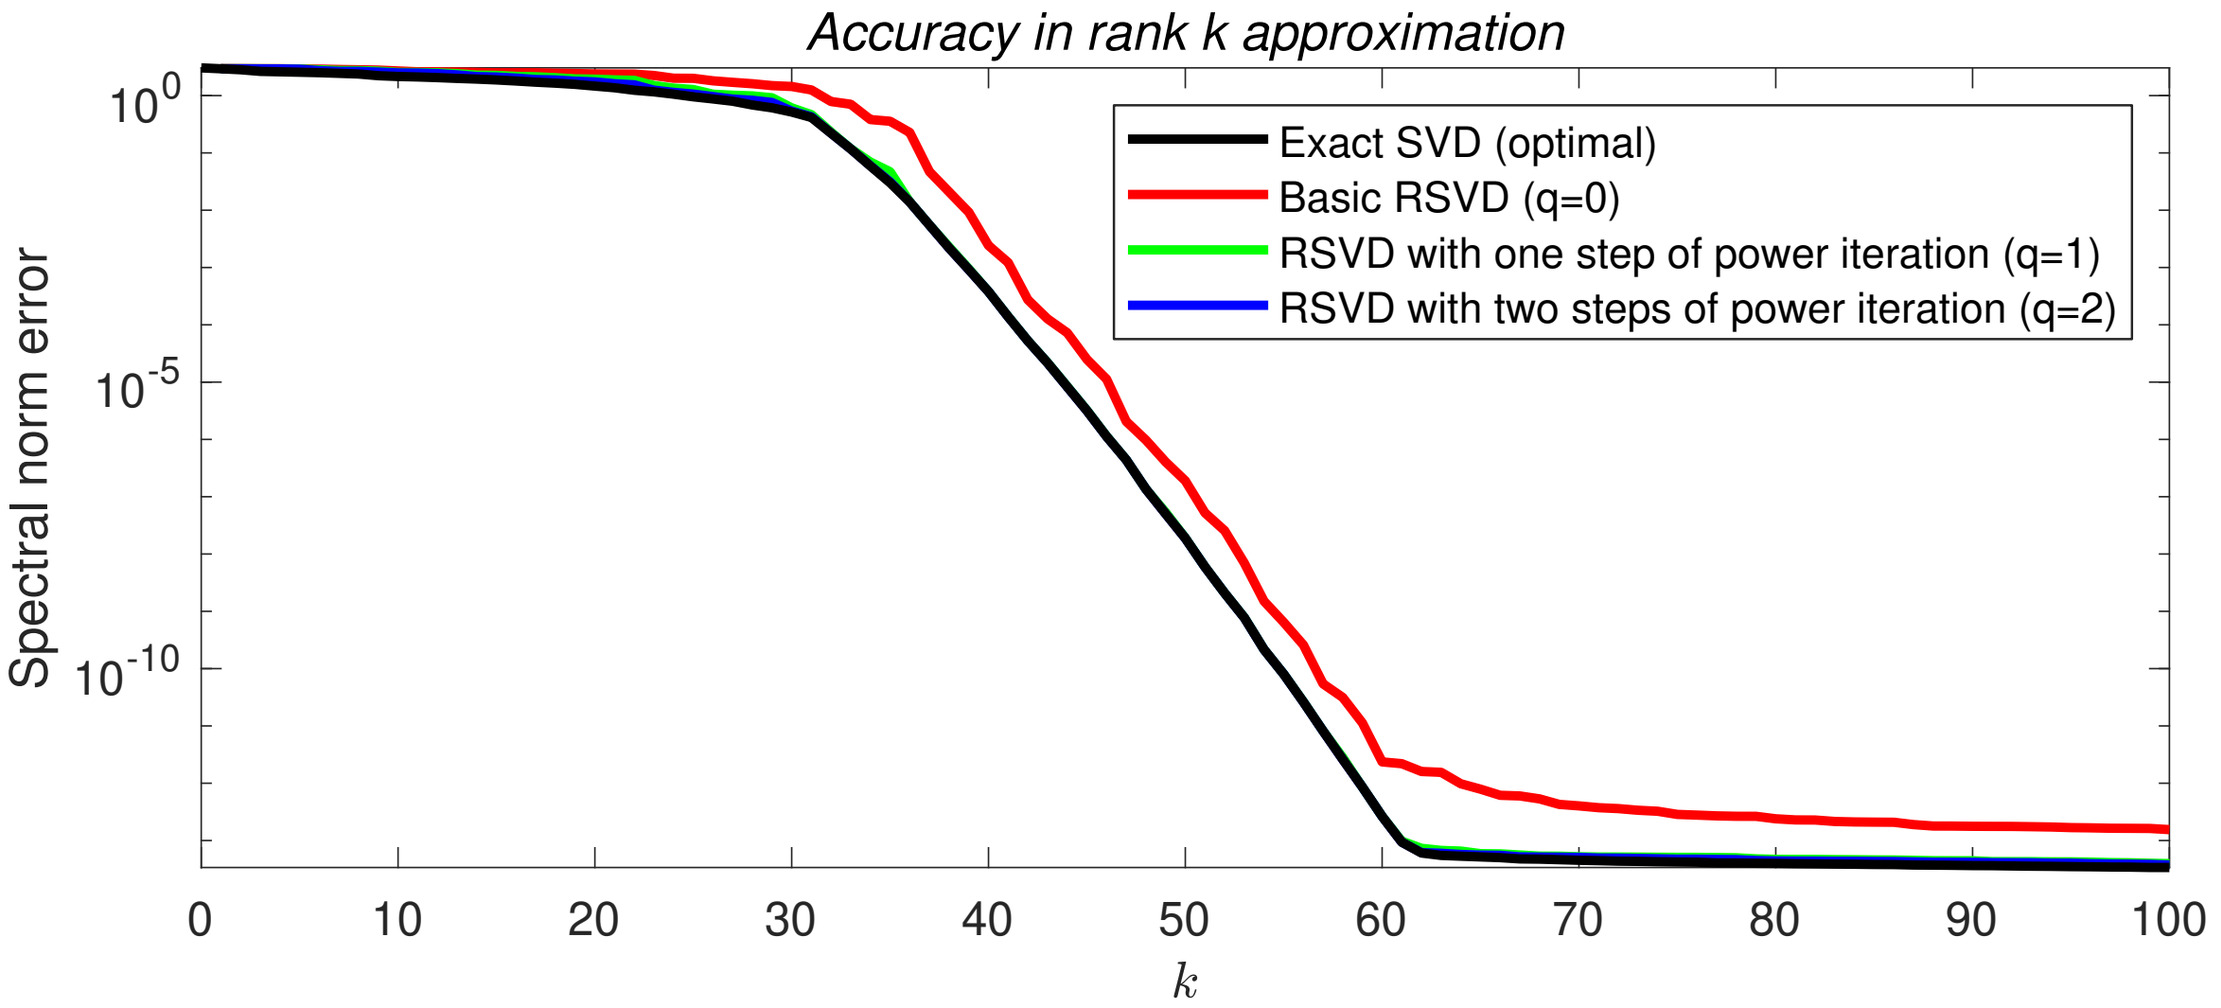

The plot shows the errors from the randomized range finder. To be precise, we plot

$$e_k = \|\mathbf{A} - \mathbf{P}_k\mathbf{A}\|,$$

where $\mathbf{P}_k$ is the orthogonal projection onto the first k columns of

$$\mathbf{Y} = (\mathbf{A}\mathbf{A}^*)^q\mathbf{A}\mathbf{G},$$

and where $\mathbf{G}$ is a Gaussian random matrix.

The matrix $\mathbf{A}$ is an approximation to a scattering operator for a Helmholtz problem.

Randomized low rank approximation:

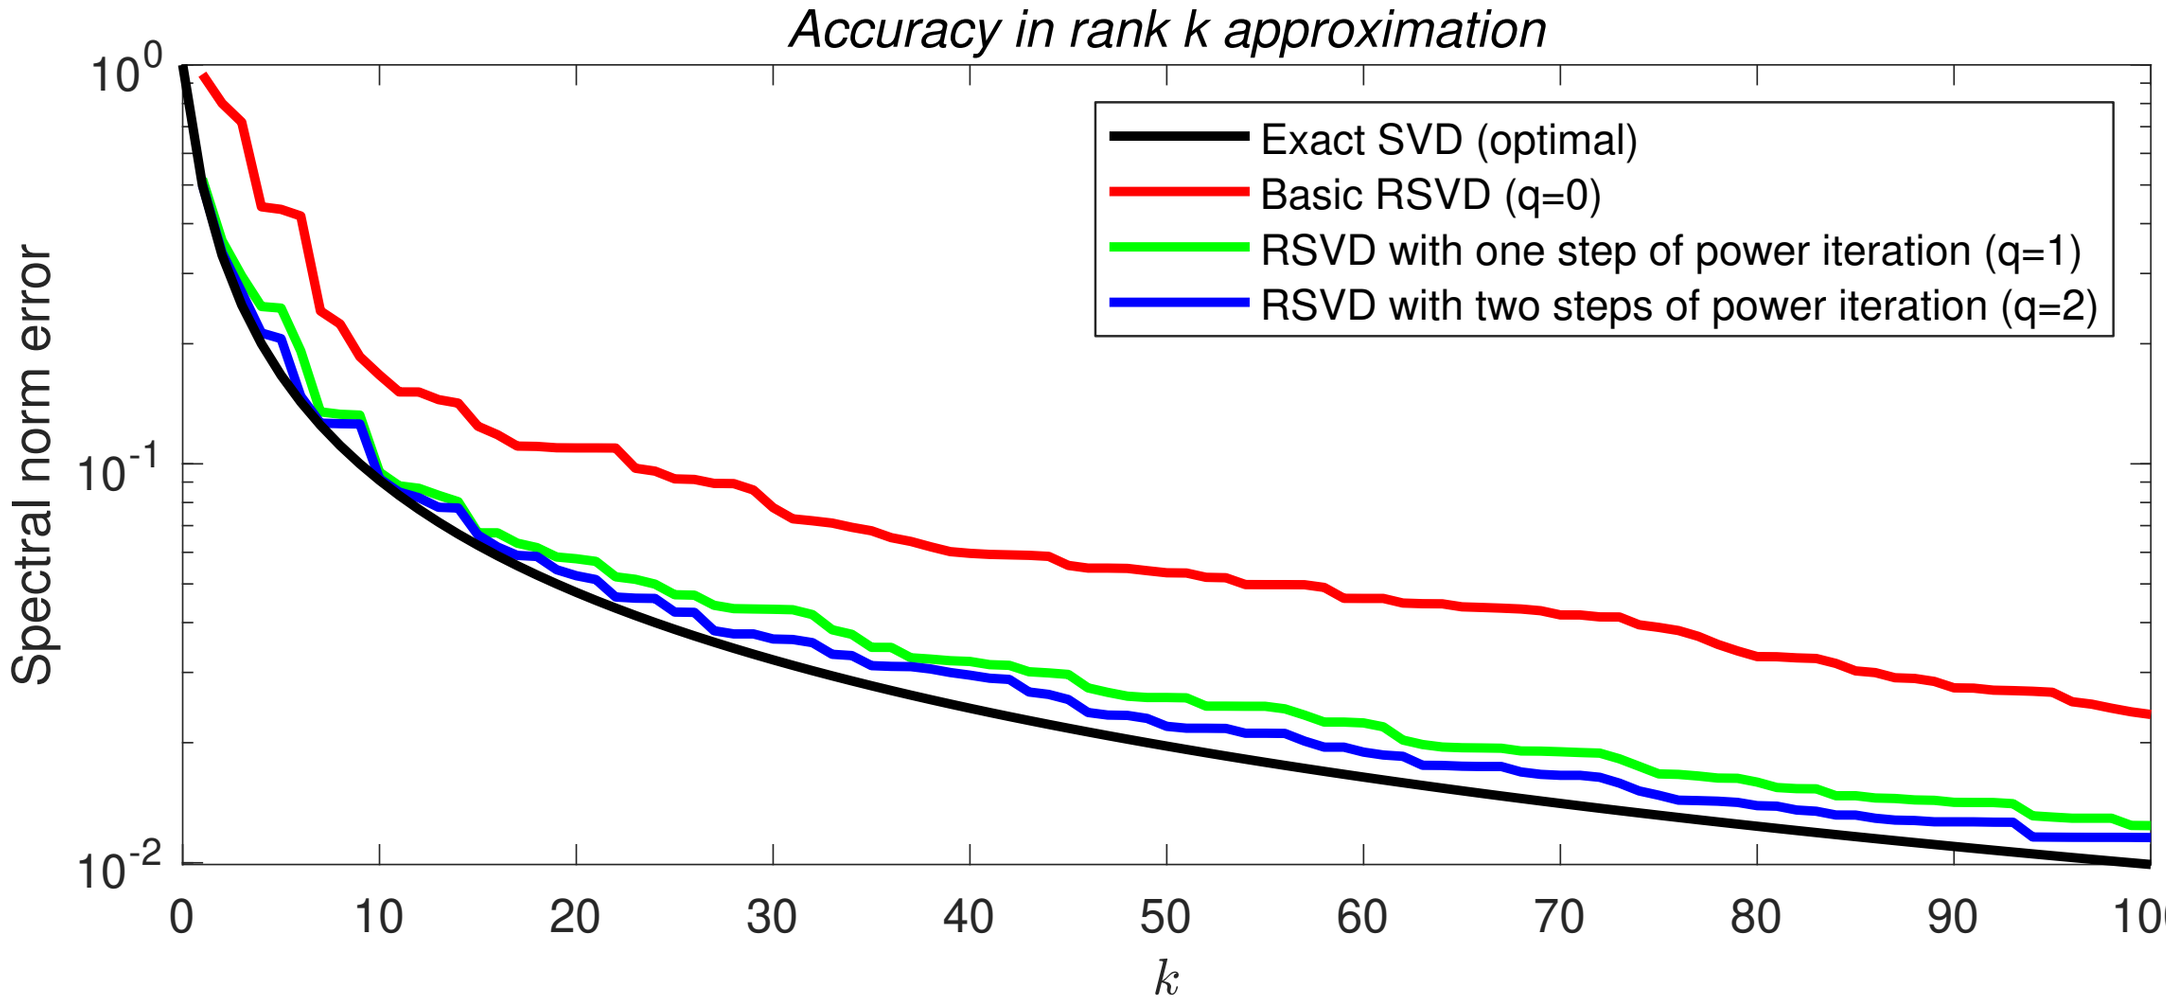

The plot shows the errors from the randomized range finder. To be precise, we plot

$$e_k = \|\mathbf{A} - \mathbf{P}_k\mathbf{A}\|,$$

where $\mathbf{P}_k$ is the orthogonal projection onto the first k columns of

$$\mathbf{Y} = (\mathbf{A}\mathbf{A}^*)^q\mathbf{A}\mathbf{G},$$

and where $\mathbf{G}$ is a Gaussian random matrix.

The matrix $\mathbf{A}$ now has singular values that decay slowly.

Randomized low rank approximation: The same plot, but showing 100 instantiations.

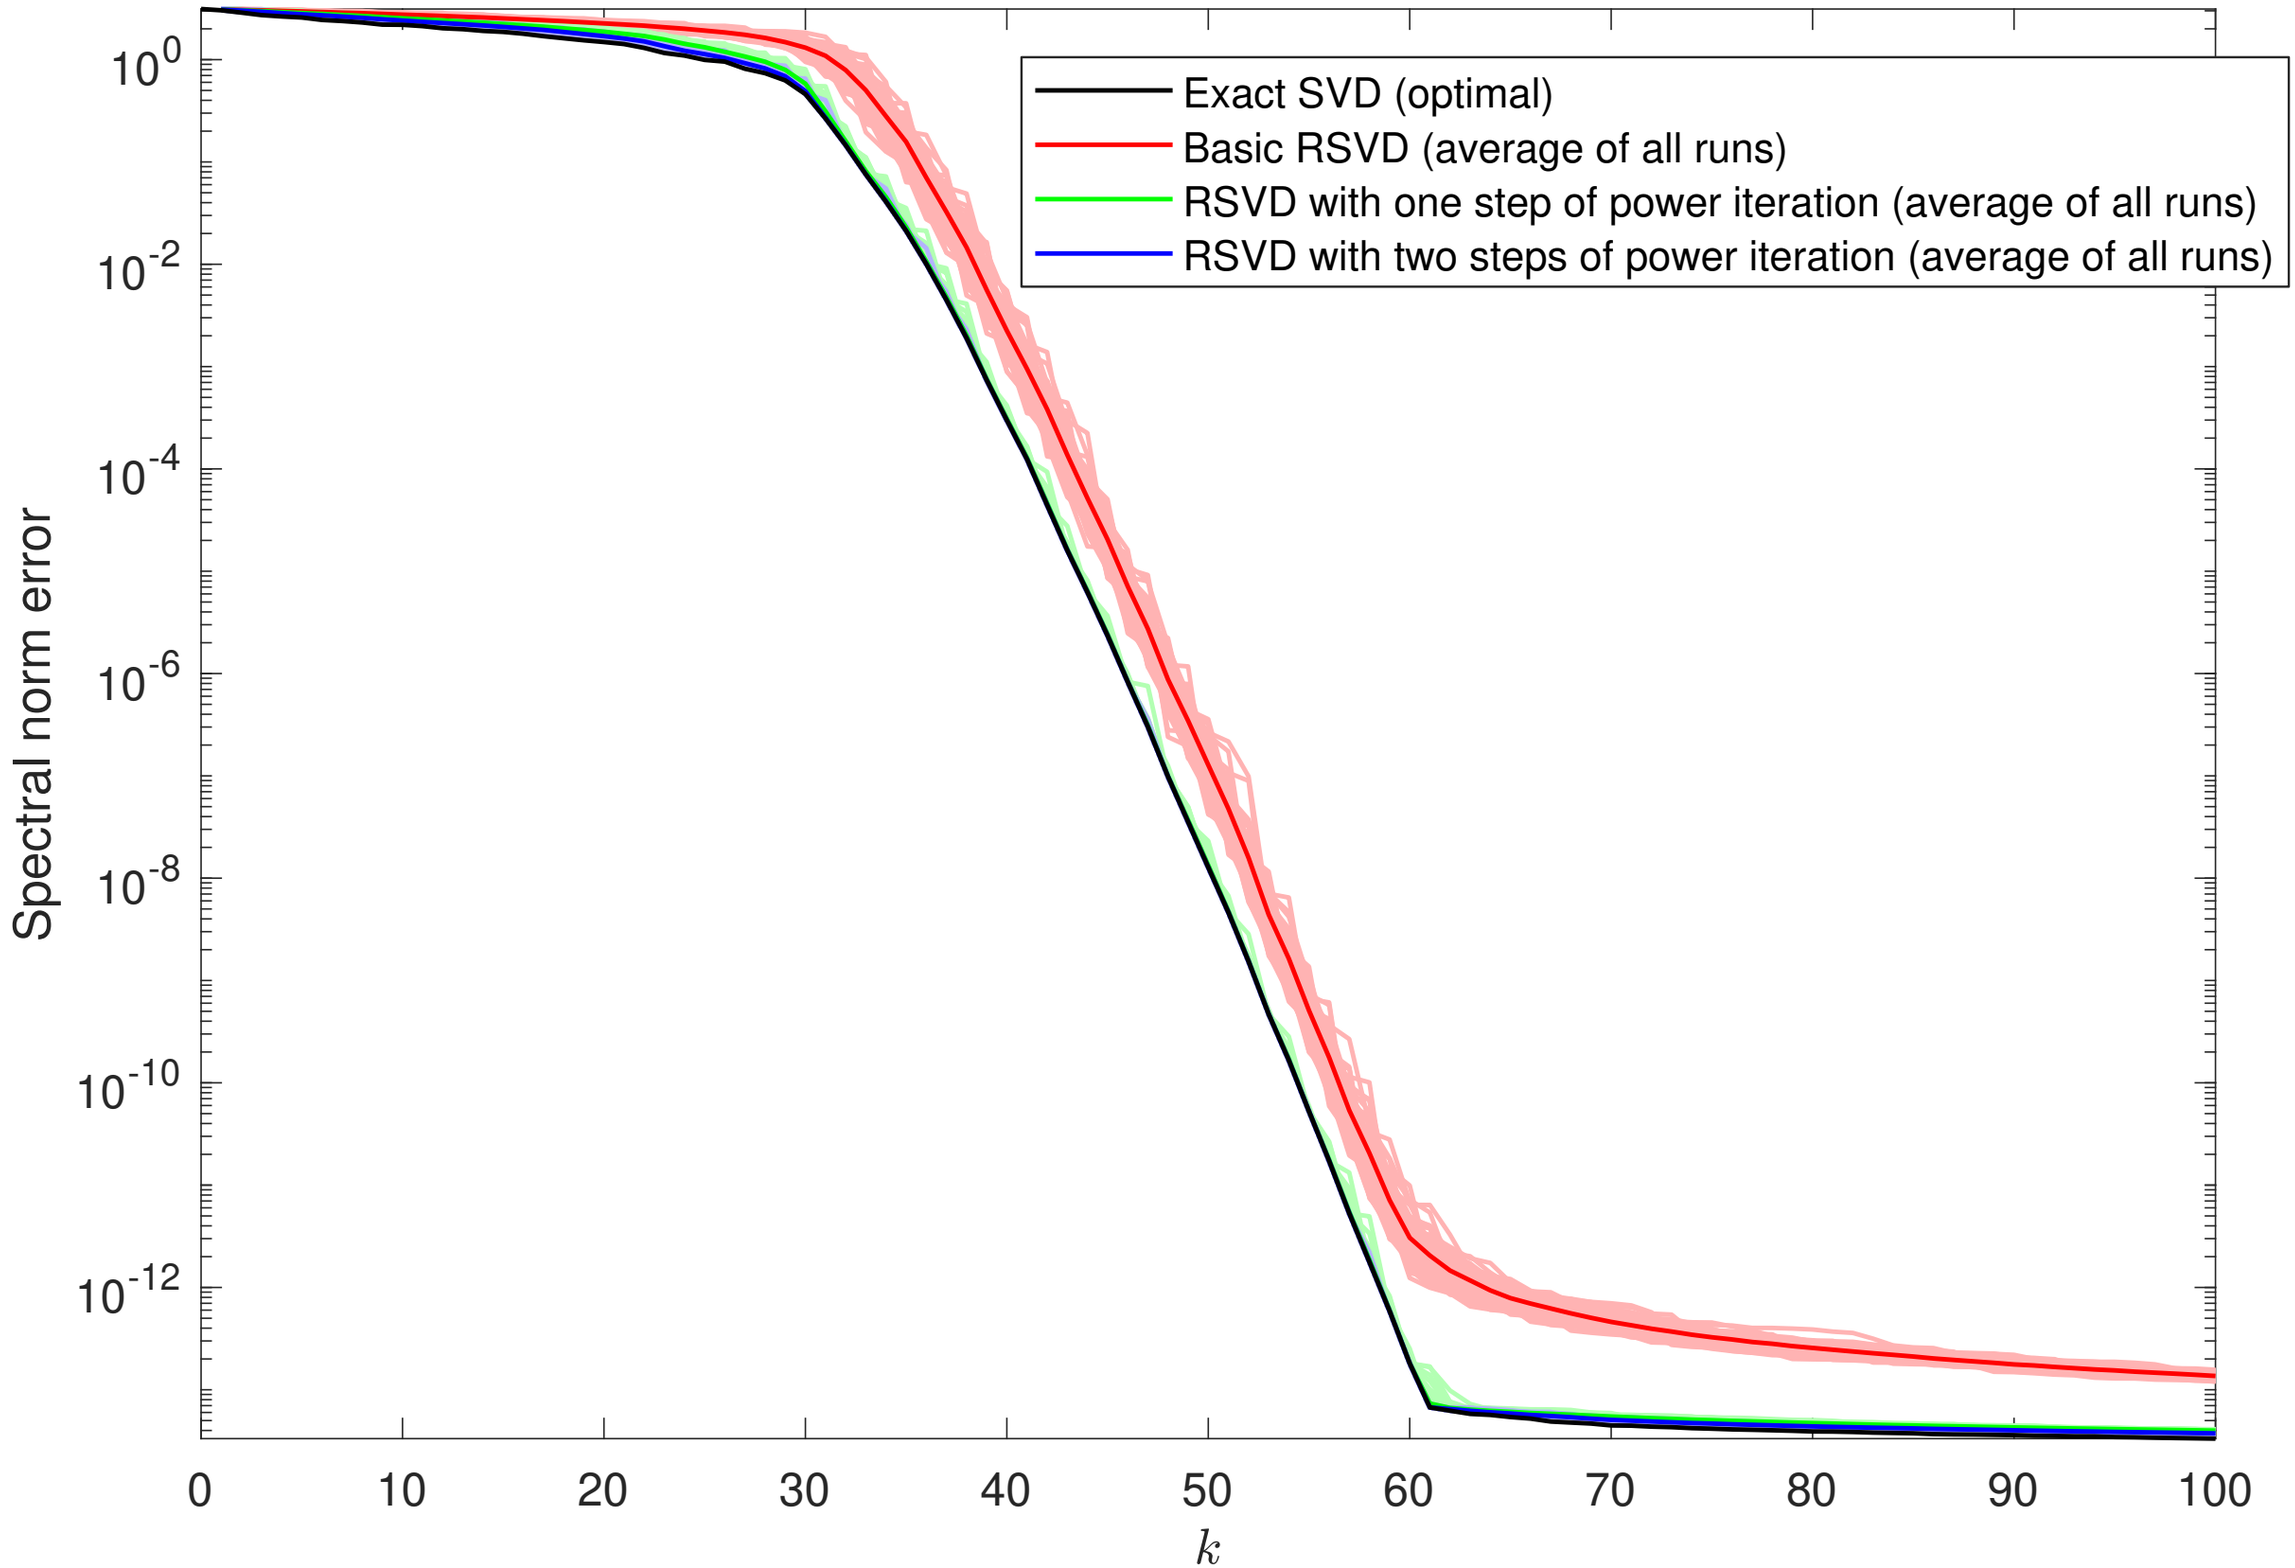

The darker lines show the mean errors across the 100 experiments.

Randomized low rank approximation: The same plot, but showing 100 instantiations.

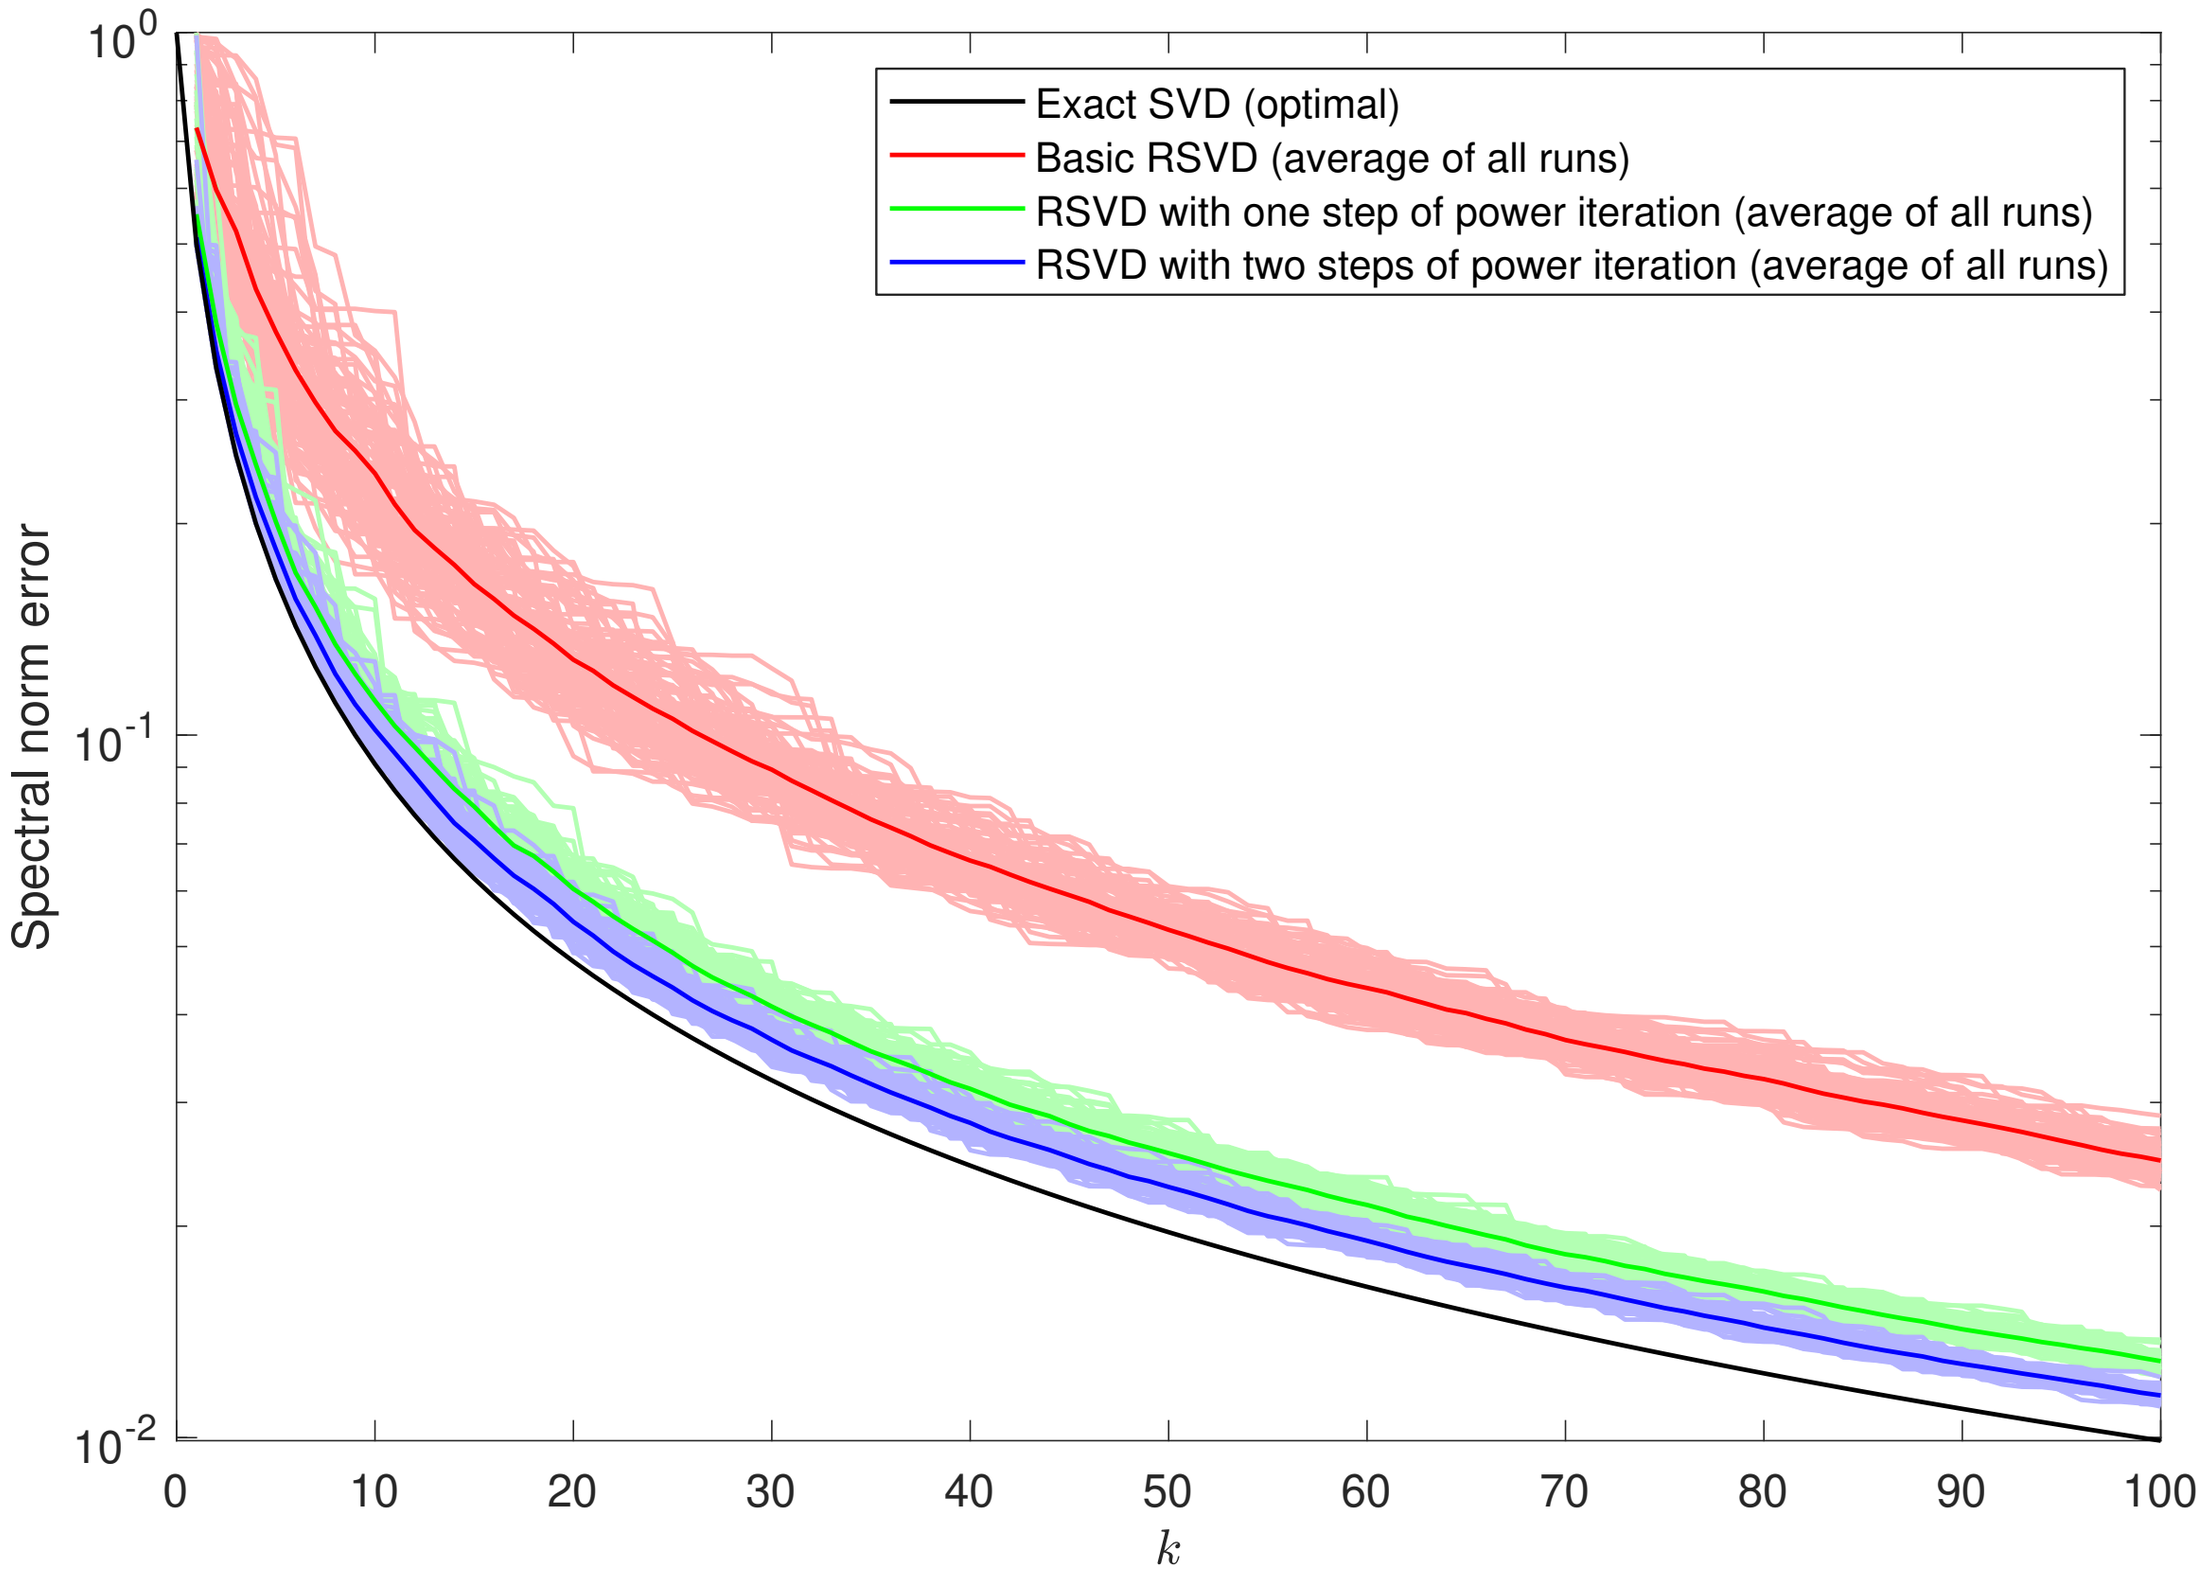

The darker lines show the mean errors across the 100 experiments.

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Error analysis: Observe first that steps (3)–(6) are all exact (up to floating point arithmetic). The error is entirely determined by the draw of the matrix $\mathbf{G}$, since

$$\mathbf{A}_{\text{approx}} = \mathbf{UDV}^* = \mathbf{Q}\hat{\mathbf{U}}\mathbf{DV}^* = \mathbf{QQ}^*\mathbf{A} = \mathbf{YY}^\dagger\mathbf{A}.$$

So:

$$\|\mathbf{A} - \mathbf{A}_{\text{approx}}\| = \|\mathbf{A} - \mathbf{YY}^\dagger\mathbf{A}\| = \|\mathbf{A} - (\mathbf{AG})(\mathbf{AG})^\dagger\mathbf{A}\|.$$

In other words, the question is how well the columns of $\mathbf{Y}$ span the column space of $\mathbf{A}$.

We recall from the Eckart-Young theorem that the optimal basis for col($\mathbf{A}$) is given by the dominant k left singular vectors of $\mathbf{A}$. Let $\mathbf{A} = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^*$ be the *exact* SVD of $\mathbf{A}$. Then

$$\mathbf{Y} = \mathbf{AG} = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^*\mathbf{G} = \{\text{Set } \tilde{\mathbf{G}} := \tilde{\mathbf{V}}^*\mathbf{G}\} = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{G}},$$

where $\tilde{\mathbf{G}}$ also has a Gaussian distribution. So the j'th column of $\mathbf{Y}$ is

$$\mathbf{y}(:,j) = \sum_{i=1}^{n} \tilde{\sigma}_i \, \tilde{g}_{i,j} \, \tilde{\mathbf{u}}_i.$$

Input: An $m \times n$ matrix **A**, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors **U**, **D**, and **V** in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix **G**.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix **Q** s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Recall that the j'th column of **Y** satisfies

$$\mathbf{y}(:,j) = \sum_{i=1}^{n} \tilde{\sigma}_i \, \tilde{g}_{i,j} \, \tilde{\mathbf{u}}_i.$$

where $\tilde{\sigma}_j$ is the j'th exact singular value, where $\tilde{\mathbf{u}}_j$ is the associated left singular vector, and where $\tilde{g}_{i,j}$ are i.i.d. $\in \mathcal{N}(0,1)$.

- When **A** has *exact* rank k, the algorithm succeeds (i.e. $\mathbf{A} = \mathbf{QQ}^*\mathbf{A}$) with probability 1.
- When the singular values *decay rapidly*, we expect close to optimal results.
- When the singular values *decay slowly*, the basis will likely not be good.

  This explains why powering helps in cases where the singular values decay slowly, since this boosts the separation between large/small svds. For instance, replacing $\mathbf{Y} = \mathbf{A\Omega}$ by $\mathbf{Y} = \mathbf{AA}^*\mathbf{A\Omega}$ results in $\mathbf{y}_j = \sum_{i=1}^{n} \left(\sigma_i^3 \psi_{i,j}\right) \mathbf{u}_i$.

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

The probabilistic behavior of the error has been studied extensively, and reasonably tight theoretical bounds have been established.

A small sample of references addressing this and related problems:

- Frieze, Kannan & Vempala, 2004.
- Martinsson, Rokhlin & Tygert,YALEU/DCS/RR-1361, 2006.
- Liberty, Woolfe, Martinsson, Rokhlin & Tygert, PNAS, 2008.
- Halko, Martinsson & Tropp, *SIAM Review*, 2011.
- Witten & Candès, *Algorithmica*, 2015.
- Gu, *SISC*, 2015. Analysis of randomized subspace iteration.
- Musco & Musco, *NIPS*, 2015. Analysis of block Krylov methods.
- Saibaba, *SIMAX*, 2019. Accuracy of singular vectors.
- Martinsson & Tropp, *Acta Numerica*, 2020.

Input: An $m \times n$ matrix **A**, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors **U**, **D**, and **V** in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix **G**.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.
(3) Compute an ON matrix **Q** s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

To illustrate, suppose that we seek to match the optimal error in a rank-k approximation.

We apply the basic randomized procedure involving $\ell := k + p$ samples.

1. Draw an $n \times \ell$ Gaussian random matrix **G**. `G = randn(n,l)`
2. Form the $m \times \ell$ sample matrix $\mathbf{Y} = \mathbf{AG}$. `Y = A * G`
3. Form an $m \times \ell$ orthonormal matrix **Q** such that $\text{ran}(\mathbf{Q}) = \text{ran}(\mathbf{Y})$. `[Q, ~] = qr(Y)`

The resulting error then satisfies

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2 \mathbf{\Psi}_{(n-k),\ell} \mathbf{\Psi}_{k,\ell}^\dagger\|^2$$

where "$\mathbf{\Psi}_{i,j}$" denotes an $i \times j$ matrix drawn from a Gaussian distribution, and where $\mathbf{D}_2 = \text{diag}(\sigma_{k+1}, \sigma_{k+2}, \ldots, \sigma_{\min(m,n)})$ holds the "tail" singular values of **A**.

Note: The term $\|\mathbf{D}_2\|^2$ is the minimal error, according to Eckart-Young.

The suboptimality term depends on $\sigma_{\min}(\mathbf{\Psi}_{k,\ell})$, where $\mathbf{\Psi}_{k,\ell}$ is a $k \times \ell$ Gaussian matrix.

- With no oversampling, $k = \ell$, we get terrible results!
- With lots of oversampling, $k \ll \ell$, things look good. ("Marchenko–Pastur" case)

What about the case where ℓ is only slightly larger than k?

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

To illustrate, suppose that we seek to match the optimal error in a rank-k approximation.

We apply the basic randomized procedure involving $\ell := k + p$ samples.

1. Draw an $n \times \ell$ Gaussian random matrix $\mathbf{G}$. `G = randn(n,l)`
2. Form the $m \times \ell$ sample matrix $\mathbf{Y} = \mathbf{AG}$. `Y = A * G`
3. Form an $m \times \ell$ orthonormal matrix $\mathbf{Q}$ such that $\text{ran}(\mathbf{Q}) = \text{ran}(\mathbf{Y})$. `[Q, ~] = qr(Y)`

The resulting error then satisfies

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2 \boldsymbol{\Psi}_{(n-k),\ell} \boldsymbol{\Psi}_{k,\ell}^\dagger\|^2$$

where "$\boldsymbol{\Psi}_{i,j}$" denotes an $i \times j$ matrix drawn from a Gaussian distribution, and where $\mathbf{D}_2 = \text{diag}(\sigma_{k+1}, \sigma_{k+2}, \ldots, \sigma_{\min(m,n)})$ holds the "tail" singular values of $\mathbf{A}$.

Note: The term $\|\mathbf{D}_2\|^2$ is the minimal error, according to Eckart-Young.

Proposition: (Chen & Dongarra 2005) Let $\boldsymbol{\Psi}$ be a Gaussian matrix of size $k \times k + p$ where $p \geq 2$. Then $\left(\mathbb{E}\left[\|\boldsymbol{\Psi}^\dagger\|_{\text{F}}^2\right]\right)^{1/2} \leq \sqrt{\frac{k}{p-1}}$ and $\mathbb{E}\left[\|\boldsymbol{\Psi}^\dagger\|\right] \leq \frac{e\sqrt{k+p}}{p}$.

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).

Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.

(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.

(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.

(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.

(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.

(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.

(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices.

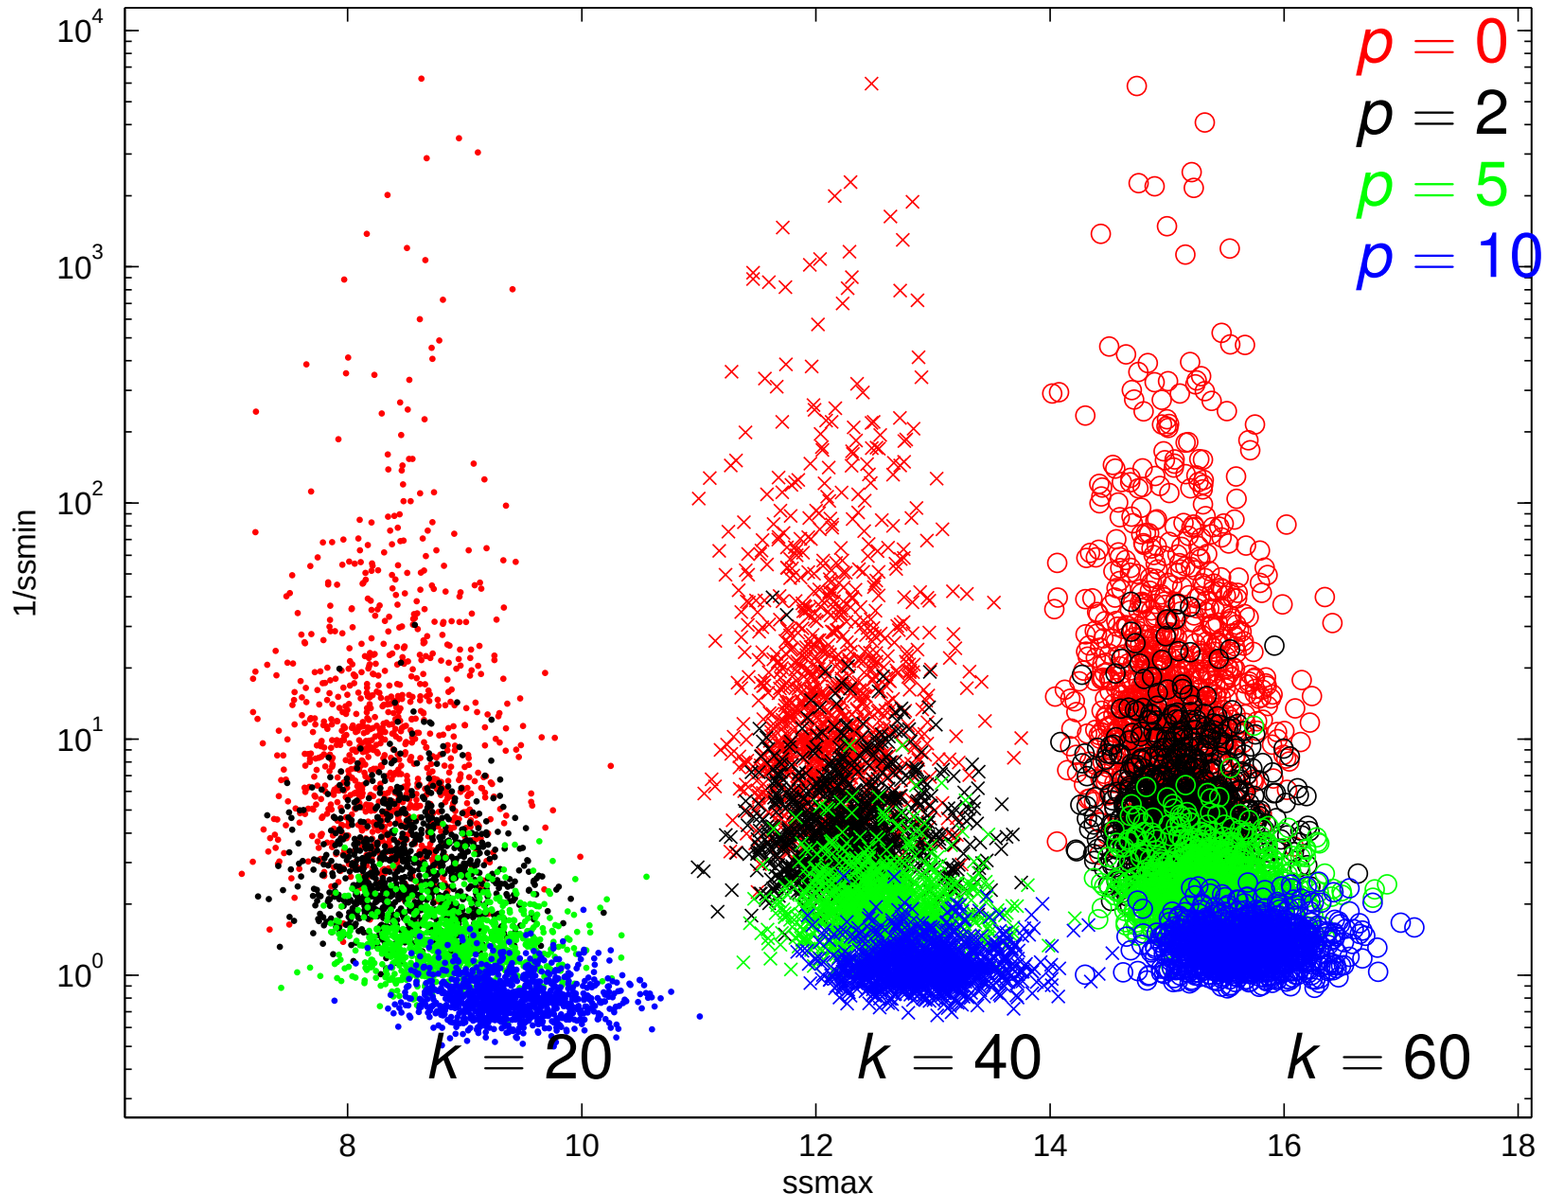

$1/\sigma_{\min}$ is plotted against $\sigma_{\max}$.

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Using the bound $\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2\mathbf{\Psi}_{(n-k),\ell}\mathbf{\Psi}_{k,\ell}^{\dagger}\|^2$, along with the bounds on $\sigma_{\min}(\Psi_{k,\ell})$, we can bound the expectation of the error:

Theorem: (Halko, Martinsson, Tropp 2011) Let $\mathbf{A}$ be an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$. Let k be a target rank, and let p be an over-sampling parameter such that $p \geq 2$ and $k + p \leq \min(m, n)$. Let $\mathbf{G}$ be a Gaussian random matrix of size $n \times (k+p)$ and set $\mathbf{Q} = \texttt{orth}(\mathbf{AG})$. Then the average error satisfies

$$\mathbb{E}\big[\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|_{\text{Fro}}\big] \leq \left(1 + \frac{k}{p-1}\right)^{1/2} \left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2\right)^{1/2},$$

$$\mathbb{E}\big[\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|\big] \leq \left(1 + \sqrt{\frac{k}{p-1}}\right) \sigma_{k+1} + \frac{e\sqrt{k+p}}{p} \left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2\right)^{1/2}.$$

There are also bounds on the likelihood of a large deviation from the expectation. (It turns out to decay super-exponentially fast as p increases!)

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{G}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Key points:

- High practical speed — interacts with $\mathbf{A}$ only through matrix-matrix multiplication.
- Communication efficient: Out-of-core, GPU, distributed memory, ...

Input: An $m \times n$ matrix **A**, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors **U**, **D**, and **V** in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix **G**.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix **Q** s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Key points:

- High practical speed — interacts with **A** only through matrix-matrix multiplication.
- Communication efficient: Out-of-core, GPU, distributed memory, ...
- Consider the problem of computing the dominant k eigenvectors/eigenvalues of a dense matrix of size $m \times n$. Reduction in complexity from $O(mnk)$ to $O(mn)$.

The key is to use a *Fast Johnson-Lindenstrauss transform*.

- Randomized trigonometric transforms (FFT, Hadamard, etc). Cost is $O(mn \log(k))$.
- Chains of Given's rotations ("Kac's random walk"). Cost is $O(mn \log(k))$.
- "Sparse sign matrix". Place r random entries in each row of Ω. (Say $r = 2$ or $r = 4$.) Cost is now $O(mn)$!

Practical acceleration is achieved at ordinary matrix sizes.

(Some additional modifications are required – we sketch from both sides, etc.)

Input: An $m \times n$ matrix **A**, a target rank k, and an over-sampling parameter p (say $p = 5$).

Output: Rank-$(k+p)$ factors **U**, **D**, and **V** in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.

(1) Draw an $n \times (k+p)$ random matrix **G**.

(2) Form th

(3) Compu

(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.

$\mathbf{\hat{U}DV}^*$.

Key points:

- High pra on.
- Commu
- Conside of a dense n

The key

- Rando
- Chains
- "Spars

  Cost is

Practica

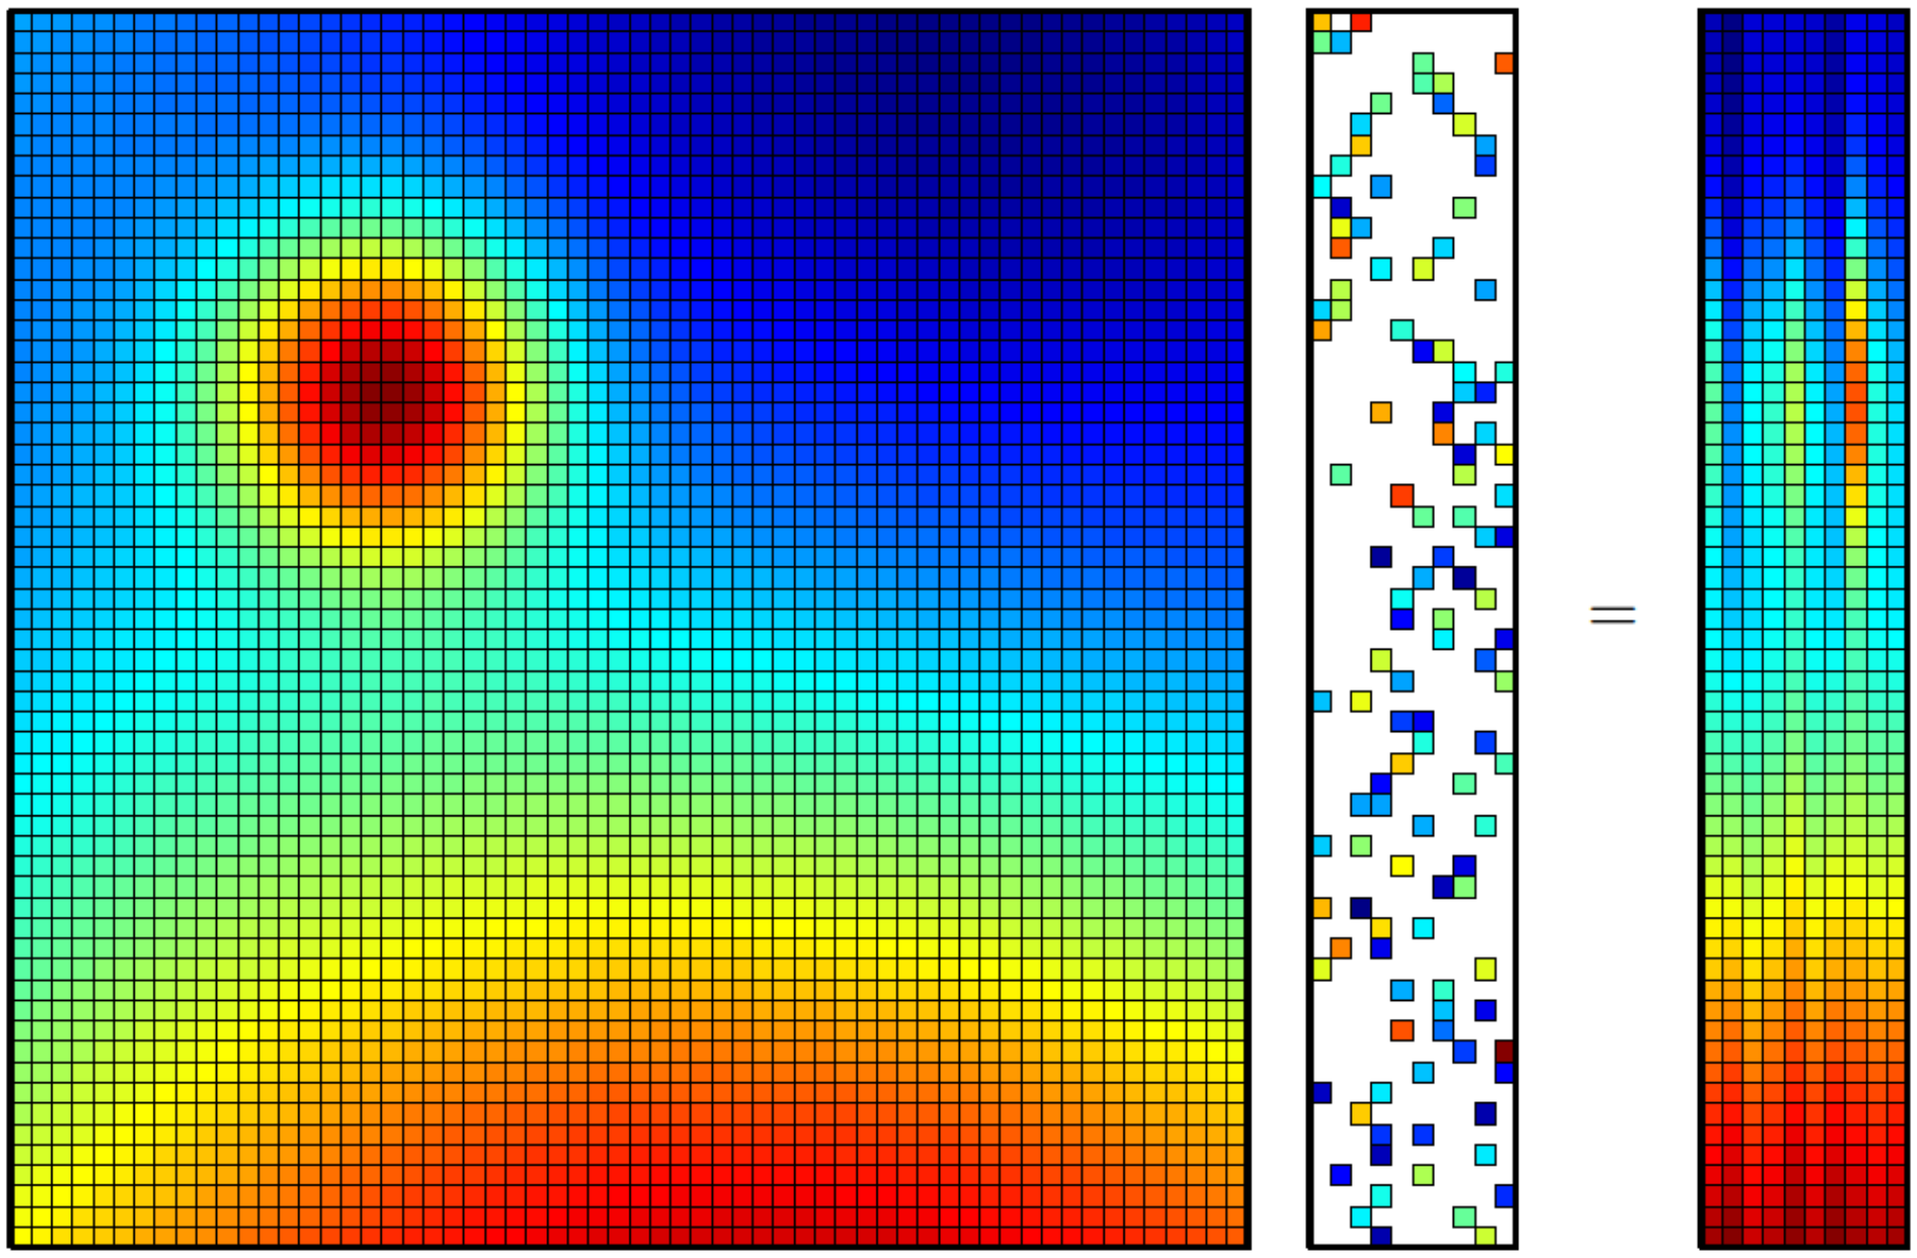

(Some additional modifications are required – we sketch from both sides, etc.)

Input: An $m \times n$ matrix **A**, a target rank k, and an over-sampling parameter p (say $p = 5$).	
Output: Rank-$(k+p)$ factors **U**, **D**, and **V** in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix **G**.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AG}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix **Q** s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Key points:

- High practical speed — interacts with **A** only through matrix-matrix multiplication.
- Communication efficient: Out-of-core, GPU, distributed memory, ...
- Consider the problem of computing the dominant k eigenvectors/eigenvalues of a dense matrix of size $m \times n$. Reduction in complexity from $O(mnk)$ to $O(mn)$.

  The key is to use a *Fast Johnson-Lindenstrauss transform*.
  - Randomized trigonometric transforms (FFT, Hadamard, etc). Cost is $O(mn \log(k))$.
  - Chains of Given's rotations ("Kac's random walk"). Cost is $O(mn \log(k))$.
  - "Sparse sign matrix". Place r random entries in each row of Ω. (Say $r = 2$ or $r = 4$.)
    Cost is now $O(mn)$!
- Single pass algorithms have been developed for *streaming environments.*

  The idea is that you are allowed to observe each matrix element only once.

  You cannot store the matrix. *Not possible with deterministic methods!*

Random mixing for solving linear systems & least squares problems

Fast random transforms open the door to other linear algebraic problems:

Computing *full* rank-revealing factorizations: Let $\mathbf{A}$ be an $n \times n$ matrix. To build a *full* rank-revealing factorization of $\mathbf{A}$, draw $\mathbf{\Omega}$ from a Haar distribution, and form

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

Then perform a QR factorization of $\mathbf{Y}$, so that

$$\mathbf{Y} = \mathbf{QR}.$$

One can prove that then a "rank-revealing" factorization is obtained through

$$\mathbf{A} = \mathbf{QR\Omega}^*.$$

The point here is communication efficiency. *(Demmel, Dumitriu, Holtz 2007)*

Setting $\mathbf{Y} = \mathbf{AA}^*\mathbf{A\Omega}$ is even better! *(Heavner, Chen, Gopal, Martinsson 2023)*

Solving linear systems without pivoting: Consider

$$\mathbf{Ax} = \mathbf{b}.$$

Randomized preconditioning results in a system that can be solved *without pivoting*

$$\left(\mathbf{\Psi}^*\mathbf{A\Omega}\right)\left(\mathbf{\Omega}^*\mathbf{x}\right) = \mathbf{\Psi}^*\mathbf{b}.$$

Again, the point is communication efficiency. *(Parker 1995)*

Random mixing for solving linear systems & least squares problems

Fast random transforms open the door to other linear algebraic problems:

Computing *full* rank-revealing factorizations: Let $\mathbf{A}$ be an $n \times n$ matrix. To build a *full* rank-revealing factorization of $\mathbf{A}$, draw $\mathbf{\Omega}$ from a Haar distribution, and form

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

Then perform a QR factorization of $\mathbf{Y}$, so that

$$\mathbf{Y} = \mathbf{QR}.$$

One can prove that then a "rank-revealing" factorization is obtained through

$$\mathbf{A} = \mathbf{QR\Omega}^*.$$

The point here is communication efficiency. *(Demmel, Dumitriu, Holtz 2007)*

Setting $\mathbf{Y} = \mathbf{AA}^*\mathbf{A\Omega}$ is even better! *(Heavner, Chen, Gopal, Martinsson 2023)*

General idea:

Randomized mixing puts you in a "generic position".

Standard Gram-Schmidt *without pivoting* becomes a reliable way to build an orthonormal basis.

Random mixing for solving linear systems & least squares problems

Suppose $\mathbf{A} \in \mathbb{R}^{m \times n}$ for $m \gg n$, and that you seek to solve $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|$.

Draw a random embedding $\mathbf{\Omega} \in \mathbb{R}^{m \times d}$ and construct a smaller *sketched* system.

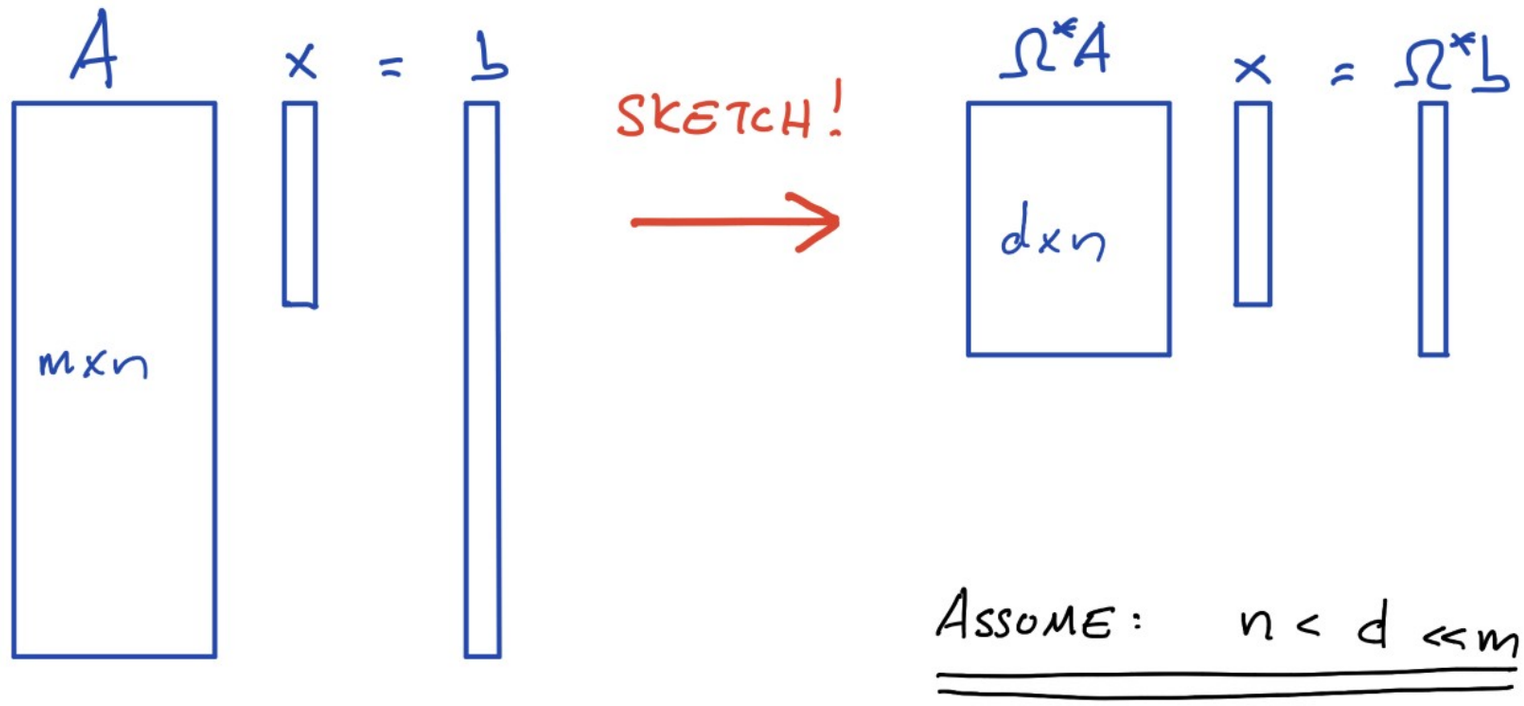

A bold approach — "sketch-to-solve":

Find the vector $\mathbf{x}$ that solves the sketched system.

A safe approach — "sketch-to-precondition":

Build a *preconditioner* $\mathbf{M} \in \mathbb{R}^{n \times n}$ by factorizing $\mathbf{\Omega}^*\mathbf{A}$ so that $\mathbf{\Omega}^*\mathbf{A} = \mathbf{QM}$.

Iterate on the preconditioned linear system $(\mathbf{AM}^{-1})(\mathbf{Mx}) = \mathbf{b}$.

Provable reduction in computational complexity: From $O(mn^2)$ to $O(mn + n^3)$.

Ideal use of randomization. Guaranteed accuracy if you evaluate the residual.

Rokhlin/Tygert (2008), Avron/Maymounkov/Toledo (2010), many more

Random mixing for solving linear systems & least squares problems

Suppose $\mathbf{A} \in \mathbb{R}^{m \times n}$ for $m \gg n$, and that you seek to solve $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|$.

Draw a random embedding $\mathbf{\Omega} \in \mathbb{R}^{m \times d}$ and construct a smaller *sketched* system.

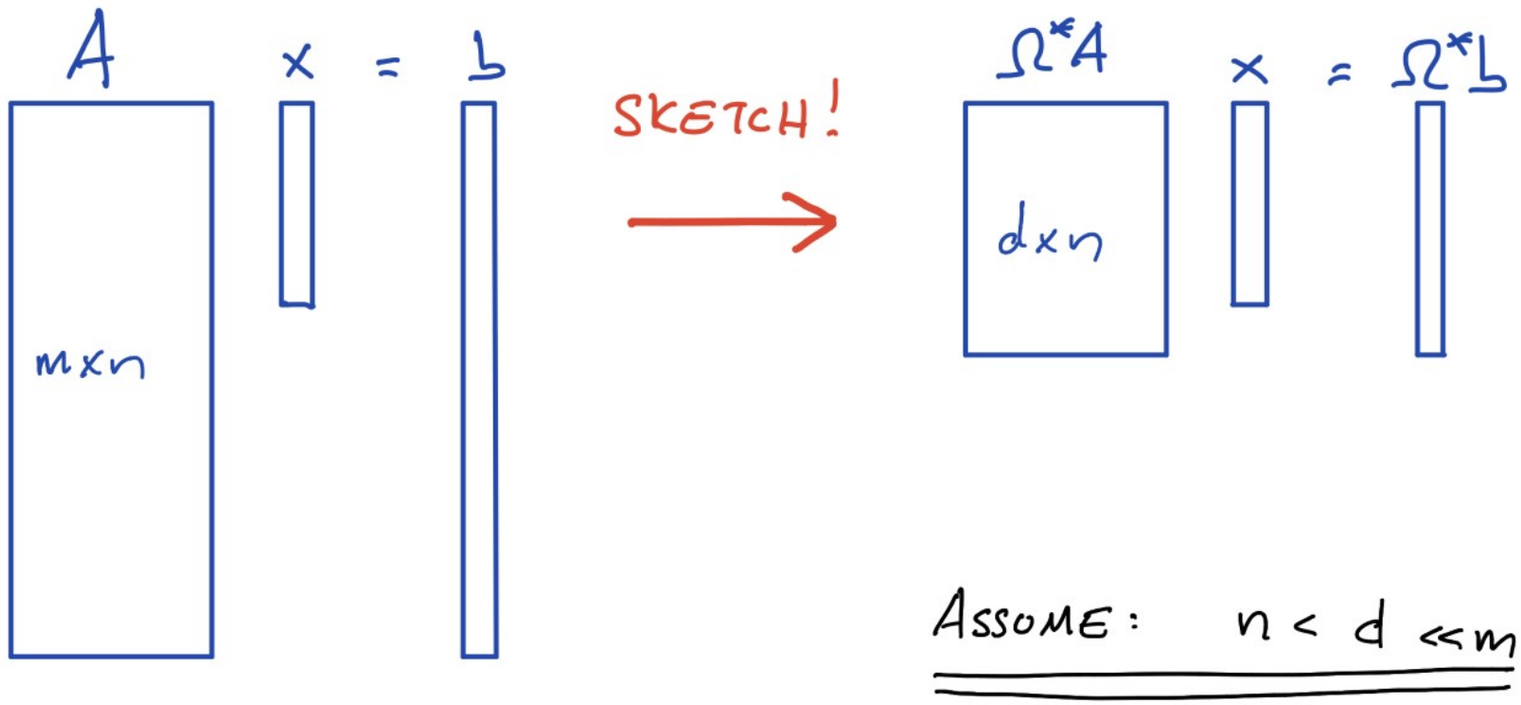

Question: Can the sketch be *down-sampling*? Simply pick d equations randomly?

Random mixing for solving linear systems & least squares problems

Suppose $\mathbf{A} \in \mathbb{R}^{m \times n}$ for $m \gg n$, and that you seek to solve $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|$.

Draw a random embedding $\mathbf{\Omega} \in \mathbb{R}^{m \times d}$ and construct a smaller *sketched* system.

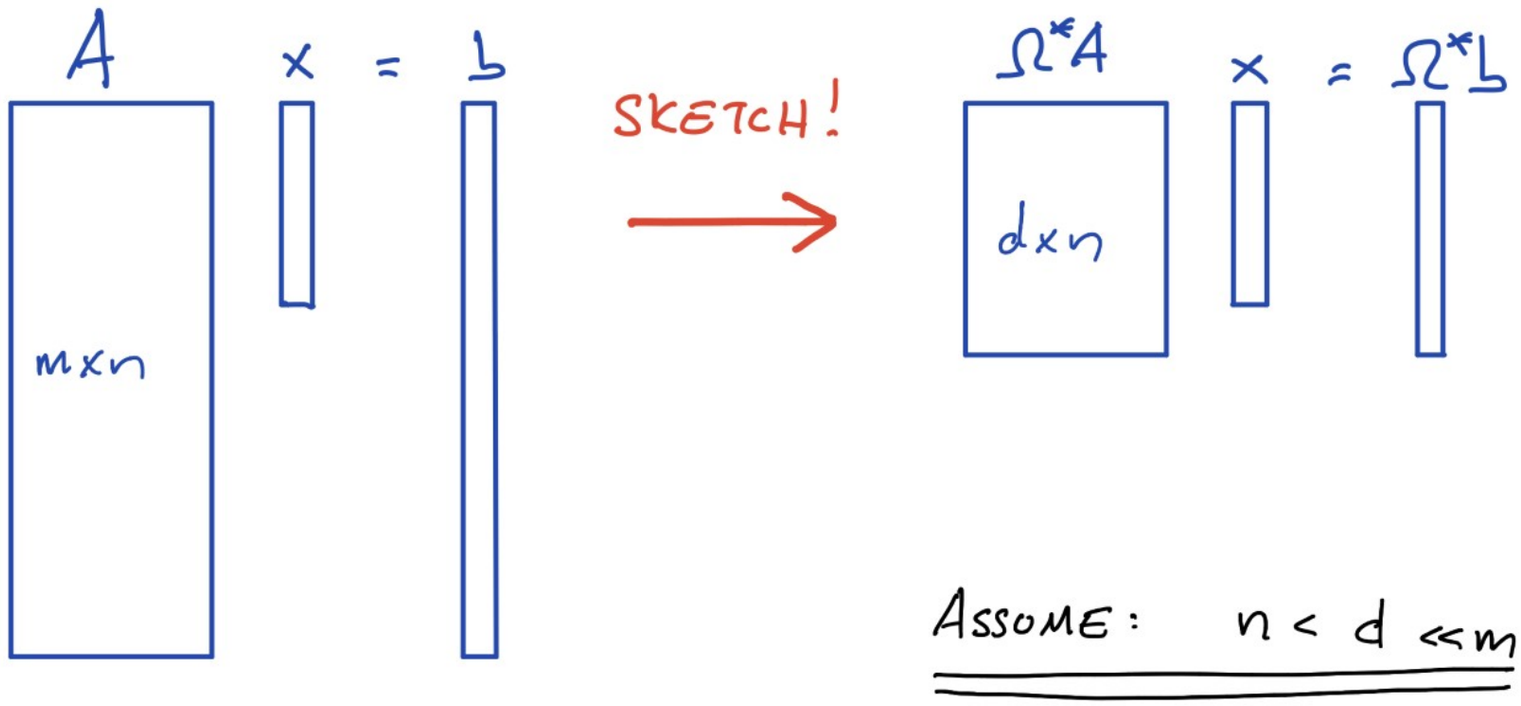

Question: Can the sketch be *down-sampling*? Simply pick d equations randomly?

No, in general. Can fail catastrophically.

Yes, if you first randomly mix the equations.

Note: There are situations where randomized sampling is an essential tool – primarily when the whole matrix cannot be assembled. Examples include kernel ridge regression and certain gigantic linear systems arising in electronic structure calculations. "Down-sample and solve" is unavoidable in such cases.

It can also be very helpful when $\mathbf{A}$ has structure, as in tensor approximation.

Outline

- **High accuracy methods** *Main topic.*
  - Low rank approximation: The randomized SVD. *Key success story.*
  - Probabilistic performance analysis and connections to random matrix theory.
  - Streaming and single-pass methods.
  - Structured random maps ("fast Johnson-Lindenstrauss transforms").
  - Linear solvers: Sketch-to-solve vs. Sketch-to-precondition.

- **Sampling based methods for huge problems**
  - Monte Carlo style methods $\rightarrow$ less reliable, less accurate, less robust.
  - Enable the solution of stupendously large problems that would otherwise be intractable.
  - Resolved some long-standing open theoretical questions in linear algebra.

- **Randomized methods for differential and integral operators**
  - Approximation of global operators of mathematical physics (solution operators, DtNs, ...).
  - Tools for matrices that are not of global low rank, but have structure that can be exploited.
  - Conceptually related to Fast Multipole Methods, Fast Direct Solvers, Calderón-Zygmund theory ...

Matrix approximation by sampling

Suppose that $\mathbf{A} = \sum_{t=1}^{T} \mathbf{A}_t$ where each $\mathbf{A}_t$ is “simple” in some sense.

Matrix approximation by sampling

Suppose that $\mathbf{A} = \sum_{t=1}^{T} \mathbf{A}_t$ where each $\mathbf{A}_t$ is "simple" in some sense.

Example: Sparse matrix written as a sum over its nonzero entries

$$\underbrace{\begin{bmatrix} 5 & -2 & 0 \\ 0 & 0 & -3 \\ 1 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}} = \underbrace{\begin{bmatrix} 5 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}_1} + \underbrace{\begin{bmatrix} 0 & -2 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}_2} + \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -3 \\ 0 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}_3} + \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}_4}$$

Example: Each $\mathbf{A}_i$ could be a column of the matrix

$$\underbrace{\begin{bmatrix} 5 & -2 & 7 \\ 1 & 3 & -3 \\ 1 & -1 & 1 \end{bmatrix}}_{=\mathbf{A}} = \underbrace{\begin{bmatrix} 5 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}}_{=\mathbf{A}_1} + \underbrace{\begin{bmatrix} 0 & -2 & 0 \\ 0 & 3 & 0 \\ 0 & -1 & 0 \end{bmatrix}}_{=\mathbf{A}_2} + \underbrace{\begin{bmatrix} 0 & 0 & 7 \\ 0 & 0 & -3 \\ 0 & 0 & 1 \end{bmatrix}}_{=\mathbf{A}_3}.$$

Example: Matrix-matrix multiplication broken up as a sum of rank-1 matrices:

$$\mathbf{A} = \mathbf{BC} = \sum_{t} \mathbf{B}(:,t)\mathbf{C}(t,:).$$

Matrix approximation by sampling

Suppose that $\mathbf{A} = \sum_{t=1}^{T} \mathbf{A}_t$ where each $\mathbf{A}_t$ is "simple" in some sense.

Let $\{p_t\}_{t=1}^{T}$ be a probability distribution on the index vector $\{1, 2, \ldots, T\}$.

Draw an index $t \in \{1, 2, \ldots, T\}$ according to the probability distribution given, and set

$$\mathbf{X} = \frac{1}{p_t}\mathbf{A}_t.$$

Then from the definition of the expectation, we have

$$\mathbb{E}[\mathbf{X}] = \sum_{t=1}^{T} p_t \times \frac{1}{p_t}\mathbf{A}_t = \sum_{t=1}^{T} \mathbf{A}_t = \mathbf{A},$$

so $\mathbf{X}$ is an *unbiased estimate* of $\mathbf{A}$.

Clearly, a single draw is not a good approximation — unrepresentative, *large variance*.

Instead, draw several samples and average:

$$\bar{\mathbf{X}} = \frac{1}{k}\sum_{t=1}^{k} \mathbf{X}_t,$$

where $\mathbf{X}_t$ are independent samples from the same distribution.

As k grows, the variance will decrease, as usual. Various Bernstein inequalities apply.

Matrix approximation by sampling

As an illustration of the theory, we cite a matrix-Bernstein result from J. Tropp (2015):

Theorem: Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. Construct a probability distribution for $\mathbf{X} \in \mathbb{R}^{m \times n}$ that satisfies

$$\mathbb{E}[\mathbf{X}] = \mathbf{A} \quad \text{and} \quad \|\mathbf{X}\| \leq R.$$

Define the per-sample second-moment: $v(\mathbf{X}) := \max\{\|\mathbb{E}[\mathbf{X}\mathbf{X}^*]\|, \|\mathbb{E}[\mathbf{X}^*\mathbf{X}]\|\}$.

Form the matrix sampling estimator: $\bar{\mathbf{X}}_k = \frac{1}{k}\sum_{t=1}^{k} \mathbf{X}_i$ where $\mathbf{X}_t \sim \mathbf{X}$ are iid.

Then $\mathbb{E}\|\bar{\mathbf{X}}_k - \mathbf{A}\| \leq \sqrt{\frac{2v(\mathbf{X})\log(m+n)}{k}} + \frac{2R\log(m+n)}{3k}$.

Furthermore, for all $s \geq 0$: $\mathbb{P}[\|\bar{\mathbf{X}}_k - \mathbf{A}\| \geq s] \leq (m+n)\exp\left(\frac{-ks^2/2}{v(\mathbf{X}) + 2Rs/3}\right)$.

Suppose that we want $\mathbb{E}\|\mathbf{A} - \bar{\mathbf{X}}\| \leq 2\epsilon$. The theorem says to pick

$$k \geq \max\left\{\frac{2v(\mathbf{X})\log(m+n)}{\epsilon^2}, \frac{2R\log(m+n)}{3\epsilon}\right\}$$

In other words, the number k of samples should be proportional to both $v(\mathbf{X})$ and to the upper bound R.

The scaling $k \sim \frac{1}{\epsilon^2}$ is discouraging, and unavoidable (since error $\epsilon \sim 1/\sqrt{k}$).

Matrix approximation by sampling – key points:

Sampling based methods provide a path forwards for problems where traditional techniques are not viable. Examples of applications:

- Kernel matrices in data analysis. Matrices are dense, and you cannot afford to form the entire matrix. "Kernel ridge regression".
- Kronecker product matrices that arise in tensor approximations.
- Vast linear systems arising in quantum chemistry.

In all these examples, there is reason to expect the data to be amenable to sampling.

Theory is well understood. Sampling based techniques provides a solution to certain problems that had proven intractable to traditional deterministic methods. Resolution to problem of finding subsets of columns/rows with maximal *spanning volume*.

When they are viable, techniques based on randomized embeddings tend to be preferable; they yield higher accuracy, and less variability in the outcome. ***Except:***

Matrix approximation by sampling – key points:

Sampling based methods provide a path forwards for problems where traditional techniques are not viable. Examples of applications:

- Kernel matrices in data analysis. Matrices are dense, and you cannot afford to form the entire matrix. "Kernel ridge regression".
- Kronecker product matrices that arise in tensor approximations.
- Vast linear systems arising in quantum chemistry.

In all these examples, there is reason to expect the data to be amenable to sampling.

Theory is well understood. Sampling based techniques provides a solution to certain problems that had proven intractable to traditional deterministic methods. Resolution to problem of finding subsets of columns/rows with maximal *spanning volume*.

When they are viable, techniques based on randomized embeddings tend to be preferable; they yield higher accuracy, and less variability in the outcome. ***Except:***

Sampling based methods excel at computing a *Cholesky factorization* of a sym. pos. def. matrix. (Use diagonal elements to form sampling probabilities.) When combined with rejection sampling, sublinear complexity and high accuracy become possible [Epperly, Tropp, Webber, SIMAX, 2025]. Extension to general matrices (through the Gram matrix) is also highly effective [Dong, Chen, Martinsson, and Pearce, SIMAX, 2025].

Outline

- **High accuracy methods** *Main topic.*
  - Low rank approximation: The randomized SVD. *Key success story.*
  - Probabilistic performance analysis and connections to random matrix theory.
  - Streaming and single-pass methods.
  - Structured random maps ("fast Johnson-Lindenstrauss transforms").
  - Linear solvers: Sketch-to-solve vs. Sketch-to-precondition.

- **Sampling based methods for huge problems**
  - Monte Carlo style methods $\rightarrow$ less reliable, less accurate, less robust.
  - Enable the solution of stupendously large problems that would otherwise be intractable.
  - Resolved some long-standing open theoretical questions in linear algebra.

- **Randomized methods for differential and integral operators**
  - Approximation of global operators of mathematical physics (solution operators, DtNs, ...).
  - Tools for matrices that are not of global low rank, but have structure that can be exploited.
  - Conceptually related to Fast Multipole Methods, Fast Direct Solvers, Calderón-Zygmund theory ...

Global operators in scientific computing

A key challenge in scientific computing is that many of the operators we seek to model computationally are *global* in the sense that every point in the domain talks to every other point. In consequence, they need to be represented as *dense* matrices.

Examples:

- Solution operators of linear elliptic PDEs.
- Boundary-to-boundary operators (e.g. Dirichlet-to-Neumann).
- Time-evolution operators of parabolic PDEs.
- Scattering matrices for many wave propagation problems.
- Schur complements in sparse direct solvers.

Algorithms:

- Methods for *applying* global operators rapidly are well established in specialized cases (FFT, Fast Multipole Methods etc). Methods exist for more general operators, but this remains a topic of research (e.g. "$\mathcal{H}$-matrices").
- Techniques for *inverting, factorizing, exponentiating, ...* such operators exist, with progress ongoing — "$\mathcal{H}$-matrices", "Fast Direct Solvers", "inverse FMM" ...
- Today: Randomized method for building data sparse representations.

Global operators in scientific computing

A key challenge in scientific computing is that many of the operators we seek to model computationally are *global* in the sense that every point in the domain talks to every other point. In consequence, they need to be represented as *dense* matrices.

Examples:

- Solution operators of linear elliptic PDEs.
- Boundary-to-boundary operators (e.g. Dirichlet-to-Neumann).
- Time-evolution operators of parabolic PDEs.
- Scattering matrices for many wave propagation problems.
- Schur complements in sparse direct solvers.

Vision:

Create a framework where we can explicitly form data sparse representations of global linear operators, use them to model complex physical systems, and then *directly* solve the resulting system of equations.

Example: Solution operators to linear elliptic boundary value problems

Consider a simple two-point boundary value problem

$$\text{(BVP)} \qquad \begin{cases} -u''(x) = g(x), & x \in \Omega = (0,1), \\ u(x) = 0, & x \in \Gamma = \partial\Omega = \{0,1\}. \end{cases}$$

We can write down the solution operator S explicitly as

$$\text{(SLN)} \qquad u(x) = [Sg](x) = \int_0^1 G(x,y)\, g(y)\, dy,$$

where the *Green's function* G takes the form

$$G(x,y) = \begin{cases} (1-x)y, & \text{when } x \geq y \quad \text{(on or below the diagonal)}, \\ x(1-y), & \text{when } x \leq y \quad \text{(on or above the diagonal)}. \end{cases}$$

Notes:

- The operator S is "benign" – compact, smoothing, ...
  (Cf. with the differential operator $d^2/(dx)^2$.)
- The operator S is *global* $\rightarrow$ *dense* matrix upon discretization.
- The kernel of S is (locally) *separable* $\rightarrow$ (locally) *low rank* matrices on discretization.
  Note the lack of regularity as you cross the diagonal, however.

Example: Solution operators to linear elliptic boundary value problems

Consider a simple elliptic PDE

$$\text{(PDE)} \qquad [-\Delta u](x) = g(x), \qquad x \in \Omega = \mathbb{R}^2,$$

combined with appropriate decay conditions at infinity. Again, we can write down the solution operator S explicitly:

$$\text{(SLN)} \qquad u(x) = [Sg](x) = \int_{\mathbb{R}^2} \phi(x-y)\, g(y)\, dy$$

where the *fundamental solution* ϕ takes the form $\phi(x-y) = \dfrac{1}{2\pi} \log |x-y|$.

Notes:

- The operator S is again more "benign" than the differential operator $-\Delta$.
- The operator S is *global* $\rightarrow$ *dense* matrix upon discretization.
- Is the kernel separable?

Example: Solution operators to linear elliptic boundary value problems

Consider a simple elliptic PDE

$$\text{(PDE)} \qquad [-\Delta u](x) = g(x), \qquad x \in \Omega = \mathbb{R}^2,$$

combined with appropriate decay conditions at infinity. Again, we can write down the solution operator S explicitly:

$$\text{(SLN)} \qquad u(x) = [Sg](x) = \int_{\mathbb{R}^2} \phi(x-y)\, g(y)\, dy$$

where the *fundamental solution* ϕ takes the form $\phi(x-y) = \dfrac{1}{2\pi} \log|x-y|$.

Approximate separability: Using complex notation, $x, y \in \mathbb{C}$, we have

$$\log(x-y) = \log(x) + \log(1-y/x) = \log(x) - \sum_{p=1}^{\infty} \frac{1}{p} \frac{y^p}{x^p}.$$

For a given precision ε, truncate at $P \sim \log(1/\varepsilon) \to$ (local) separable expansion of rank P.

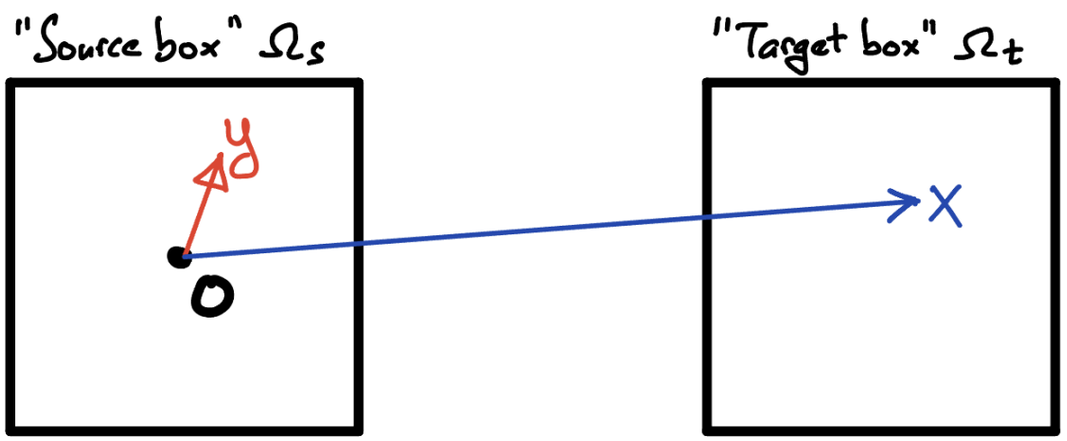

If $y \in \Omega_s$ and $x \in \Omega_t$, and $(\sqrt{2}/3)^P \leq \varepsilon$, then rank of interaction is $\leq P$ to precision ε.

Example: Solution operators to linear elliptic boundary value problems

Now consider a general linear boundary value problem of the form

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

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:

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

Example: Solution operators to linear elliptic boundary value problems

Now consider a general linear boundary value problem of the form

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

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

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

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?

Example: Solution operators to linear elliptic boundary value problems

Now consider a general linear boundary value problem of the form

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

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

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.

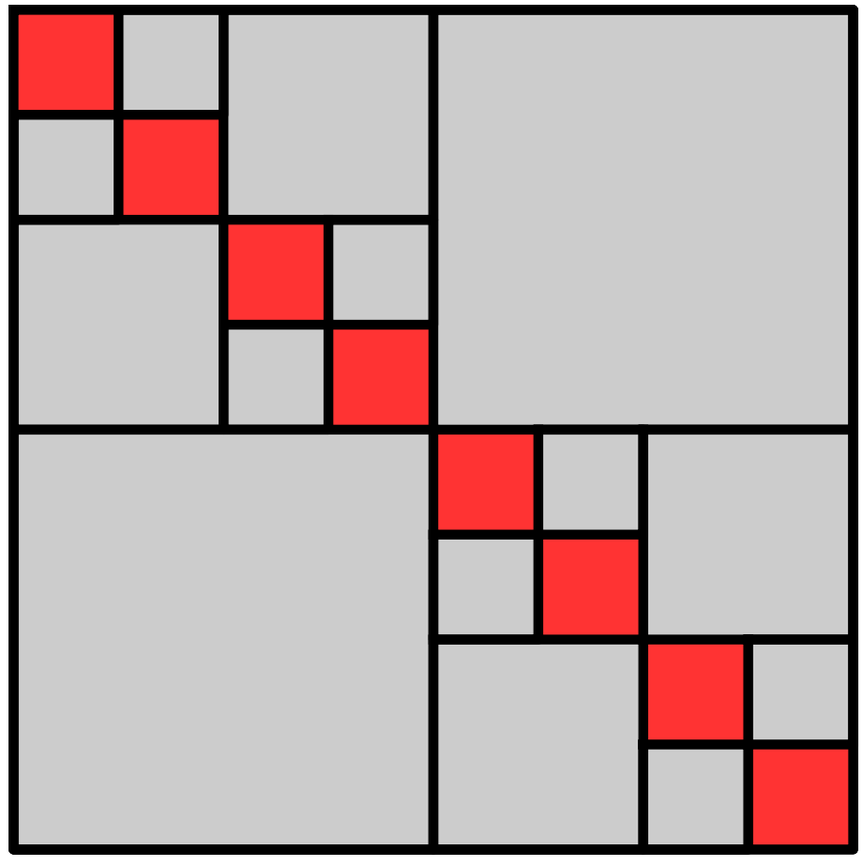

All gray blocks have low rank.

Example: Solution operators to linear elliptic boundary value problems

Now consider a general linear boundary value problem of the form

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

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

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

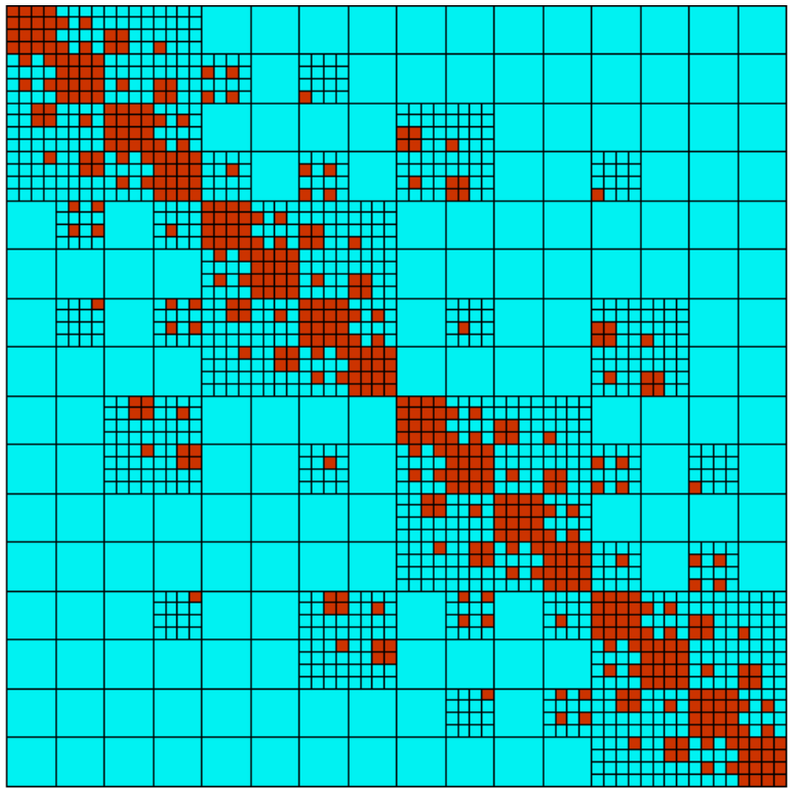

Image credit: Ambikasaran & Darve, arxiv.org #1407.1572

Example: Solution operators to linear elliptic boundary value problems

Now consider a general linear boundary value problem of the form

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

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

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, . . .); . . .

Example: Boundary 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:

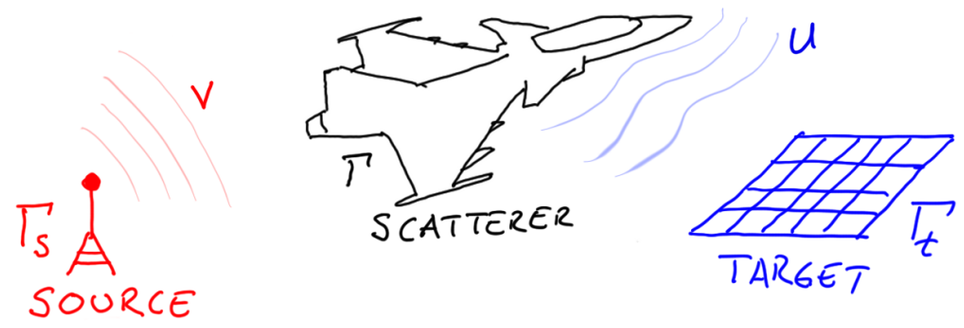

$$
(6) \qquad \begin{cases} -\Delta u(\boldsymbol{x}) - \kappa^2 u(\boldsymbol{x}) = 0 & \boldsymbol{x} \in \mathbb{R}^3 \backslash \overline{\Omega} \\ u(\boldsymbol{x}) = v(\boldsymbol{x}) & \boldsymbol{x} \in \partial\Omega \\ \lim\limits_{|\boldsymbol{x}|\to\infty} |\boldsymbol{x}| \left(\partial_{|\boldsymbol{x}|} u(\boldsymbol{x}) - i\kappa\, u(\boldsymbol{x}) \right) = 0. \end{cases}
$$

The BVP (6) has an alternative mathematical formulation in the BIE

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

The integral equation (7) has several advantages over the PDE (6), including:

- The domain of computation $\partial\Omega$ is finite.
- The domain of computation $\partial\Omega$ is 2D, while $\mathbb{R}^3 \backslash \overline{\Omega}$ is 3D.
- Equation (7) is inherently well-conditioned (as a "2^{nd} kind Fredholm equation").

The integral operator (7) is global, and a matrix resulting from discretizing it is dense.

But both this matrix and its inverse are rank structured $\rightarrow O(N)$ solvers possible.

Example: Poincaré-Steklov operators (Dirichlet-to-Neumann, etc)

Consider a well posed linear boundary value problem

$$\text{(BVP)} \qquad \begin{cases} Au(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = f(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma. \end{cases}$$

For $\boldsymbol{x} \in \Gamma$, let $h(\boldsymbol{x})$ denote the normal derivative of the solution u at $\boldsymbol{x}$. Then the map

$$T : f \mapsto h$$

is known as the *Dirichlet-to-Neumann (DtN)* map.

The DtN is a powerful tool in many areas of scientific computing:

- It provides a compressed representation that "hides" interior dynamics in a subdomain from the rest of the model. A mathematically "ideal" *reduced model*.
- Essential tool for understanding *domain decomposition* methods.
- The basis of many methods for *inverse problems* where you seek to reconstruct variable coefficients in A by observing input-output pairs.
- Ideal for coupling *multi-physics* problems – e.g. fluid-solid interactions.
- Etc etc.

Example: Poincaré-Steklov operators (Dirichlet-to-Neumann, etc)

Let us consider a boundary value problem with Dirichlet data on a rectangular domain Ω partitioned into two subdomains $\Omega = \Omega_\alpha \cup \Omega_\beta$:

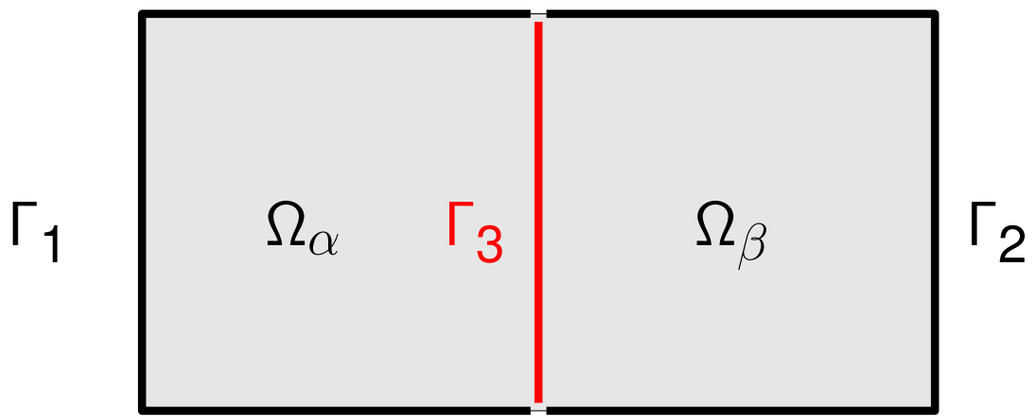

We partition the boundary of Ω as $\partial\Omega = \Gamma_1 \cup \Gamma_2$, and let Γ_3 denote the interior boundary. We know the Dirichlet data f_1 and f_2 on Γ_1 and Γ_2, but we do *not* know it on the "artificial" boundary Γ_3. However, if we know the DtN maps T^α and T^β for Ω_α and Ω_β, then we can decompose these as

$$\begin{bmatrix} T^\alpha_{11} & T^\alpha_{13} \\ T^\alpha_{31} & T^\alpha_{33} \end{bmatrix} \begin{bmatrix} f_1 \\ f_3 \end{bmatrix} = \begin{bmatrix} n_1 \\ n^\alpha_3 \end{bmatrix} \qquad \text{and} \qquad \begin{bmatrix} T^\beta_{22} & T^\beta_{23} \\ T^\beta_{32} & T^\beta_{33} \end{bmatrix} \begin{bmatrix} f_2 \\ f_3 \end{bmatrix} = \begin{bmatrix} n_2 \\ n^\beta_3 \end{bmatrix}$$

where n^α_3 and n^β_3 represent the boundary fluxes through Γ_3 from Ω_α and Ω_β, respectively. Since $n^\alpha_3 + n^\beta_3 = 0$, we can now form an equation for the unknown quantity f_3 as

$$(T^\alpha_{33} + T^\beta_{33}) f_3 = -T^\alpha_{31} f_1 - T^\beta_{32} f_2.$$

Example: Poincaré-Steklov operators (Dirichlet-to-Neumann, etc)

Illustration of how the DtN map can be used to weld together six subdomains in a domain decomposition problem.

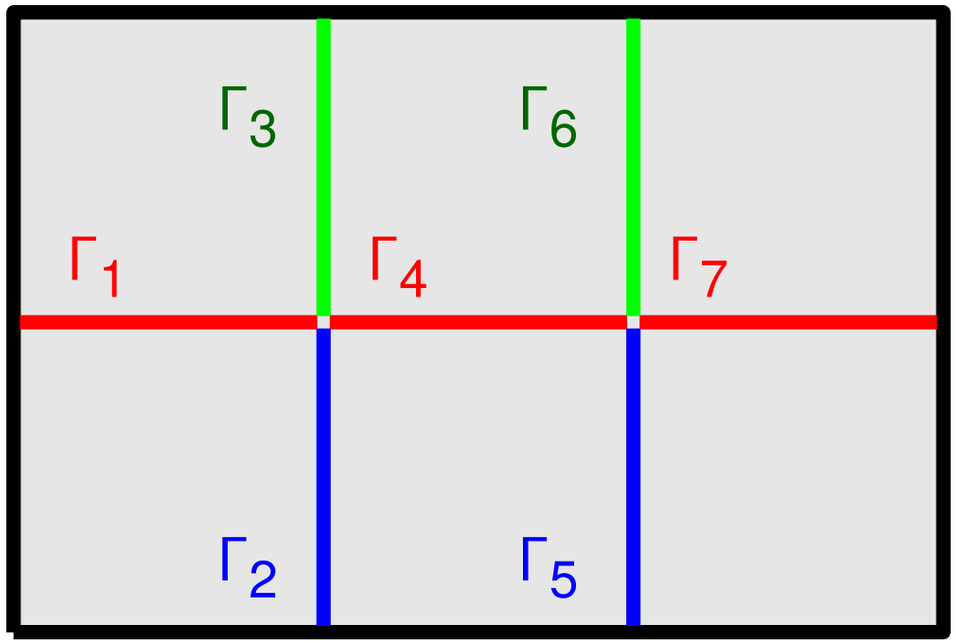

The domain.

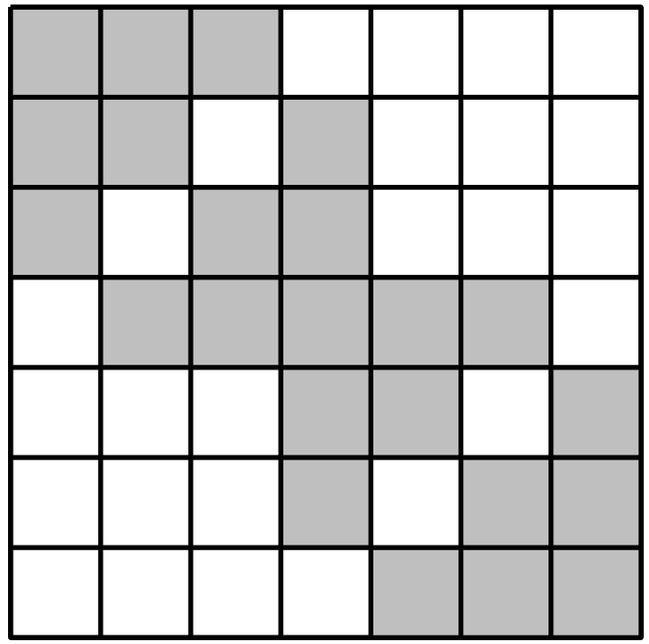

The sparsity pattern of the linear system that determines the Dirichlet data on the interior boundaries $\{\Gamma_i\}_{i=1}^{7}$.

Example: Poincaré-Steklov operators (Dirichlet-to-Neumann, etc)

Consider the free space acoustic scattering problem

$$\begin{cases} -\Delta u(\boldsymbol{x}) - \kappa^2 (1 - b(\boldsymbol{x}))\, u(\boldsymbol{x}) = -\kappa^2 b(\boldsymbol{x})\, v(\boldsymbol{x}), & \boldsymbol{x} \in \mathbb{R}^2 \\ \lim\limits_{|\boldsymbol{x}|\to\infty} \sqrt{|\boldsymbol{x}|}\big(\partial_{|\boldsymbol{x}|} u(\boldsymbol{x}) - i\kappa\, u(\boldsymbol{x})\big) = 0, \end{cases}$$

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."

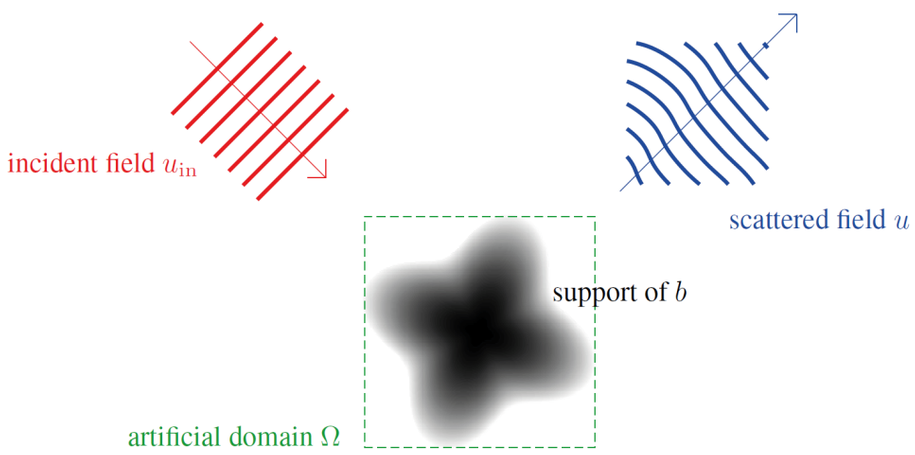

Introduce an artificial box Ω such that support(b) $\subseteq \Omega$.

On Ω:	On Ω^{c}:
• Variable coefficient PDE.	• Constant coefficient PDE.
• Discretize the PDE.	• Use BIE.
• Build DtN for $\partial\Omega$.	• Build DtN for $\partial\Omega^{\text{c}}$.

Glue the domains together using the DtNs.

(Actually, *impedance-to-impedance (ItI)* maps are better.)

[Gillman, Barnett, Martinsson, 2015]

Example: Scattering operators

A *scattering operator* is the linear operator that maps an incoming wave to an scattered field in acoustic or electromagnetic scattering problems.

Scattering operators are powerful tools for solving multibody scattering problems, as they break a problem into smaller parts:

- A local computation is used to build a scattering operator for each individual body. These computations are unconnected, so highly parallelizeable.
- Form a global system that uses the scattering matrices to describe how the bodies talk to each.

The benefit is that the global system you form this way is *far smaller* than a global system that fully resolves all individual scatters. (And often better conditioned too!)

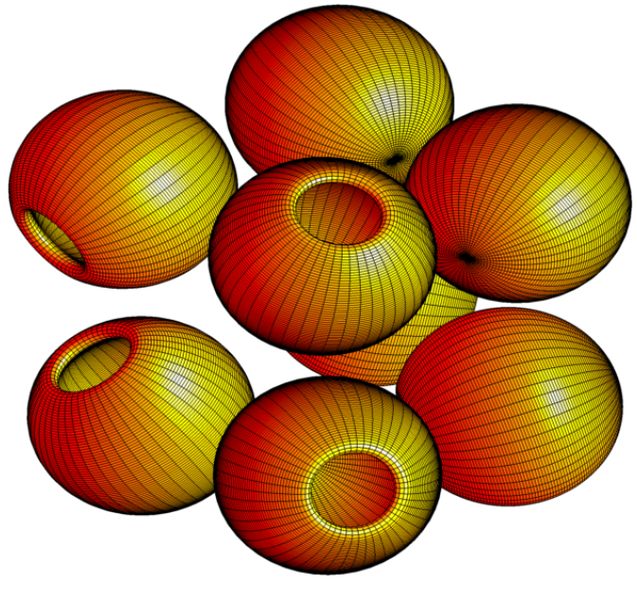

Example: Scattering operators

A *scattering operator* is the linear operator that maps an incoming wave to an scattered field in acoustic or electromagnetic scattering problems.

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

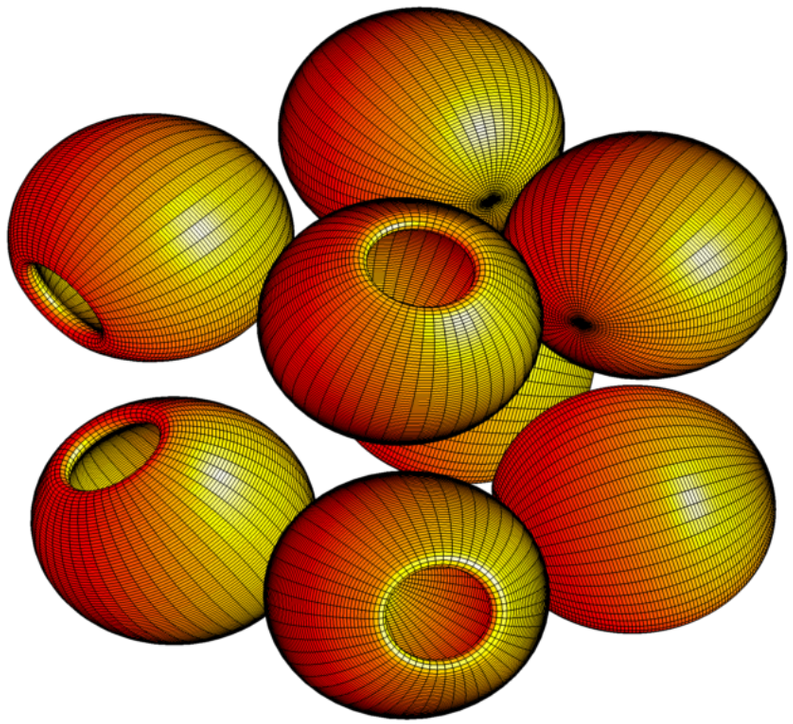

A hybrid direct/iterative solver is used (a highly accurate scattering matrix $\mathbf{S}_i$ is computed for body i) to form a global system

$$\tilde{\mathbf{q}}_i + \mathbf{S}_{ii}\left(\sum_{j\neq i}\mathbf{A}_{ij}\tilde{\mathbf{q}}_j\right) = \mathbf{S}_{ii}\tilde{\mathbf{v}}_i, \qquad i = 1,2,\ldots,J, \tag{8}$$

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]

Example: Schur complements Consider a finite difference discretization on a square resulting in a linear system $\mathbf{Au} = \mathbf{b}$. Let us partition the nodes into three sets as follows:

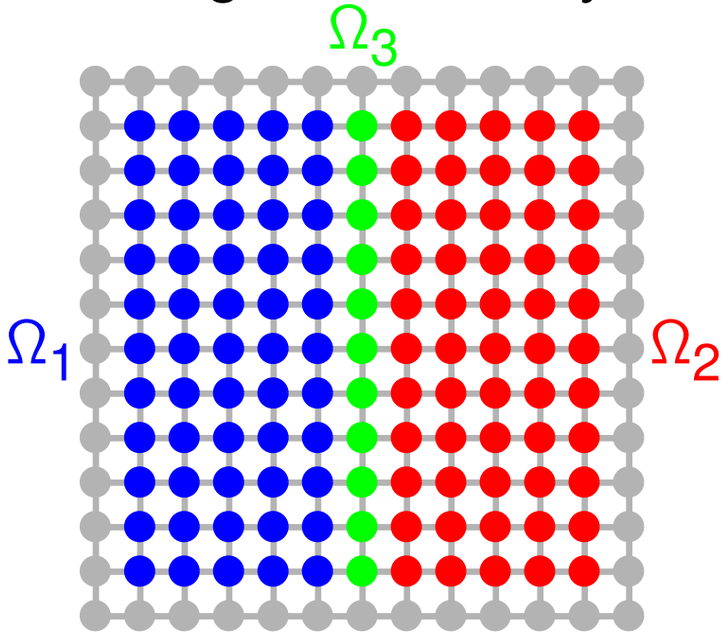

$$\mathbf{A} = \begin{array}{|c|c|c|} \hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

Example: Schur complements Consider a finite difference discretization on a square resulting in a linear system $\mathbf{Au} = \mathbf{b}$. Let us partition the nodes into three sets as follows:

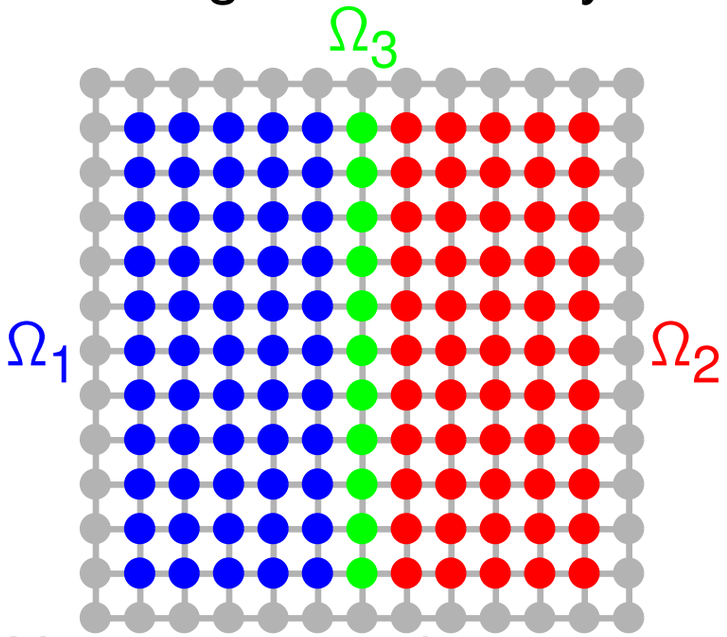

$$\mathbf{A} = \begin{array}{|c|c|c|}\hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

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

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

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.*

Example: Schur complements Consider a finite difference discretization on a square resulting in a linear system $\mathbf{Au} = \mathbf{b}$. Let us partition the nodes into three sets as follows:

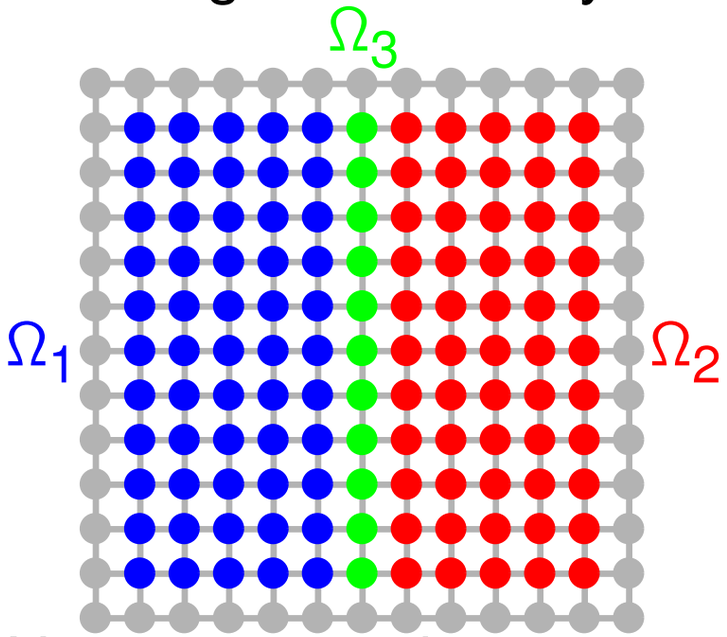

$$\mathbf{A} = \begin{array}{|c|c|c|} \hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

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

Notice the obvious recursion!

Example: Schur complements Consider a finite difference discretization on a square resulting in a linear system $\mathbf{Au} = \mathbf{b}$. Let us partition the nodes into three sets as follows:

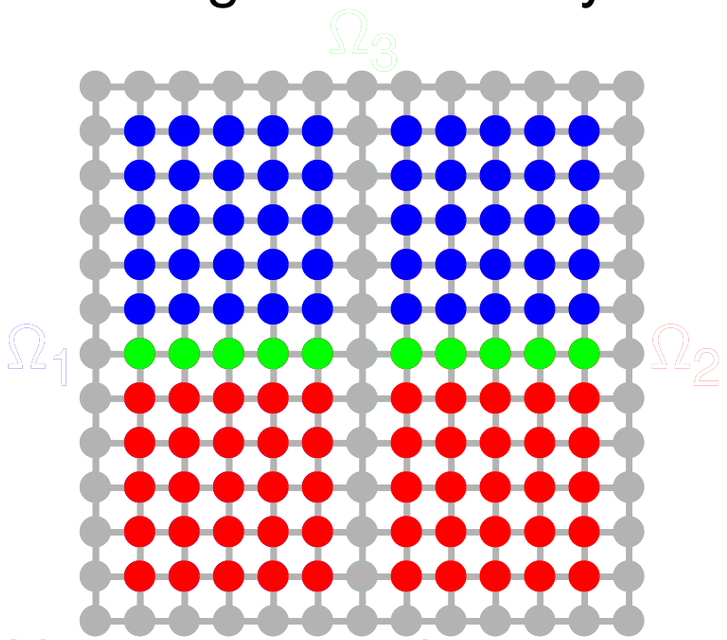

$$\mathbf{A} = \begin{array}{|c|c|c|} \hline \mathbf{A}_{11} & \mathbf{0} & \mathbf{A}_{13} \\ \hline \mathbf{0} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \\ \hline \end{array}$$

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

Notice the obvious recursion!

Example: Schur complements

Typically, nested dissection orderings are more complicated:

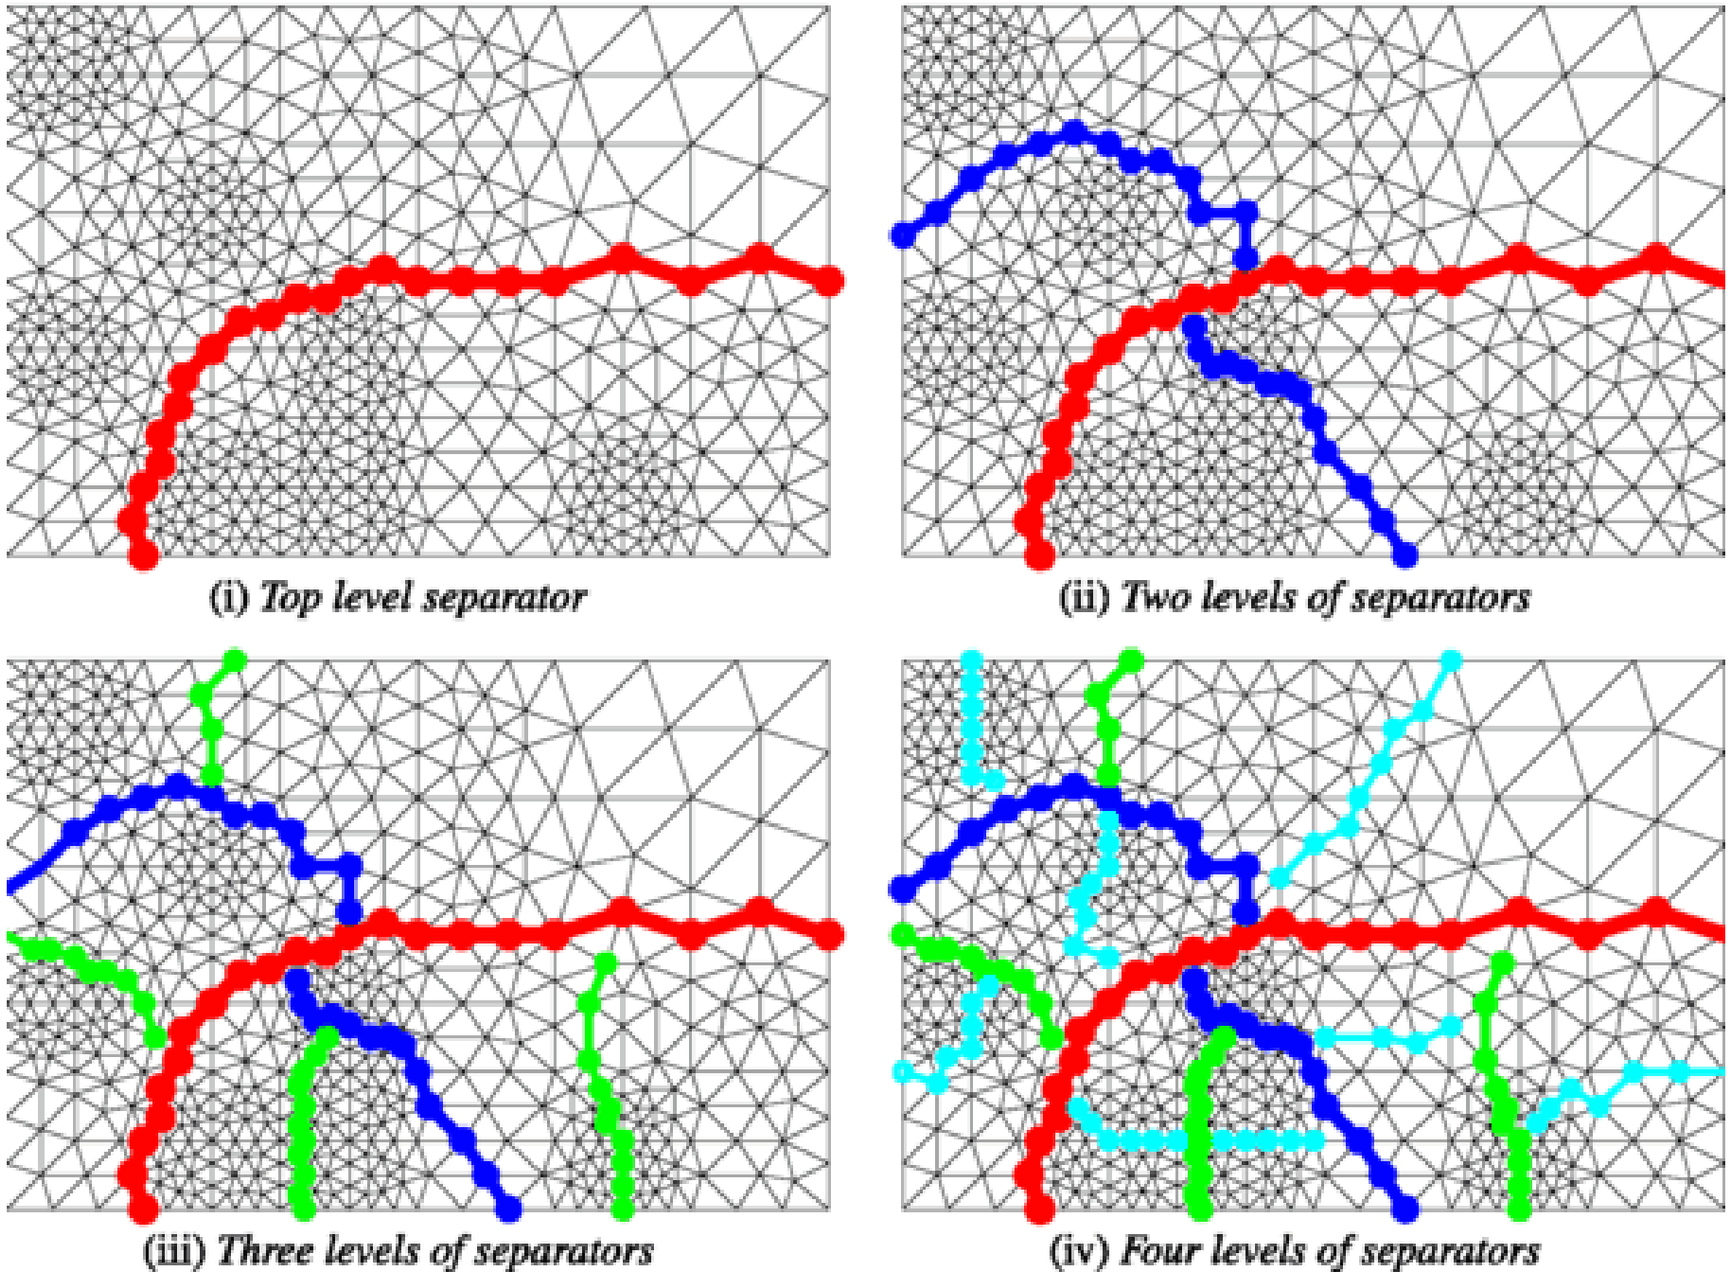

Image credit: Jianlin Xia, "Robust and Efficient Multifrontal Solver for Large Discretized PDEs", 2012

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

Example: Schur complements

It is well known that the dense factorization of the largest Schur complements is the dominant cost in sparse LU factorization.

For problems in 2D, the asymptotic flop count is $O(N^{1.5})$.

For problems in 3D, the asymptotic flop count is $O(N^2)$.

Example: Schur complements

It is well known that the dense factorization of the largest Schur complements is the dominant cost in sparse LU factorization.

For problems in 2D, the asymptotic flop count is $O(N^{1.5})$.

For problems in 3D, the asymptotic flop count is $O(N^2)$.

Assertion: These Schur complements very often behave like discretizated integral operators. (E.g. Dirichlet-to-Neumann.)

They are rank-structured, and are amenable to "fast" matrix algebra. Exploiting this, the complexity of sparse direct solvers for elliptic PDEs has in the past 10 – 20 years been decreased dramatically:

	Build stage		*Solve stage*	
2D	$N^{3/2}$	$\rightarrow N$	$N \log N$	$\rightarrow N$
3D	N^2	$\rightarrow N$	$N^{4/3}$	$\rightarrow N$

Example: Schur complements

It is well known that the dense factorization of the largest Schur complements is the dominant cost in sparse LU factorization.

For problems in 2D, the asymptotic flop count is $O(N^{1.5})$.

For problems in 3D, the asymptotic flop count is $O(N^2)$.

Nested dissection solvers with $O(N)$ complexity — *Le Borne, Grasedyck, & Kriemann (2007)*, *Martinsson (2009)*, *J. Xia, Chandrasekaran, Gu, & Li (2009)*, *Gillman & Martinsson (2011)*, *Schmitz & L. Ying (2012)*, *Darve & Ambikasaran (2013)*, *Ho & Ying (2015)*, *Amestoy, Ashcraft, et al (2015)*, *Oseledets & Suchnikova (2015)*, etc.
$O(N)$ direct solvers for integral equations were developed by *Martinsson & Rokhlin (2005)*, *Greengard, Gueyffier, Martinsson, & Rokhlin (2009)*, *Gillman, Young, & Martinsson (2012)*, *Ho & Greengard (2012)*, *Ho & Ying (2015)*. Work in 1990's *Y. Chen, P. Starr, V. Rokhlin, L. Greengard, E. Michielssen.* Related to work on $\mathcal{H}$ and $\mathcal{H}^2$ matrix methods (1998 and forwards) by *Börm, Bebendorf, Hackbusch, Khoromskij, Sauter, etc.*

Example: Schur complements

It is well known that the dense factorization of the largest Schur complements is the dominant cost in sparse LU factorization.

For problems in 2D, the asymptotic flop count is $O(N^{1.5})$.

For problems in 3D, the asymptotic flop count is $O(N^2)$.

Note: Complexity is not $O(N)$ if the nr. of "points-per-wavelength" is fixed as $N \to \infty$. This limits direct solvers to problems of size a couple hundreds of wave-lengths or so.

More complicated rank-structured formats — "butterfly matrices" — offer promise here, and initial results show great promise.

Interaction ranks: Why are they small?

We have claimed that a wide range of global operators that arise in scientific computing have rank structure. Specifically, the claim is that the numerical rank of interaction between two subdomains that are separated in space is low.

Why is this the case?

Interaction ranks: Why are they small?

We have claimed that a wide range of global operators that arise in scientific computing have rank structure. Specifically, the claim is that the numerical rank of interaction between two subdomains that are separated in space is low.

Why is this the case?

(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*.

Compression of a matrix discretizing a continuum operator

Question: How do you obtain a data sparse representation of a dense matrix discretizing a continuum operator?

Compression of a matrix discretizing a continuum operator

Question: How do you obtain a data sparse representation of a dense matrix discretizing a continuum operator?

- In some standard environments (convolution with a known fundamental solution, say), there exist analytic techniques that work very well. "Abramowitz & Stegun" or "proxy surface method". *Classical FMM environment*
- In some cases where the kernel is known explicitly, heuristic techniques such as "adaptive cross approximation" are often used. These are fast, but not 100% reliable. Ok for building pre-conditioners. *$\mathcal{H}$-matrix environment*

Compression of a matrix discretizing a continuum operator

Question: How do you obtain a data sparse representation of a dense matrix discretizing a continuum operator?

- In some standard environments (convolution with a known fundamental solution, say), there exist analytic techniques that work very well. "Abramowitz & Stegun" or "proxy surface method". *Classical FMM environment*
- In some cases where the kernel is known explicitly, heuristic techniques such as "adaptive cross approximation" are often used. These are fast, but not 100% reliable. Ok for building pre-conditioners. *$\mathcal{H}$-matrix environment*
- In more general cases, *randomized* algorithms are very competitive. These methods require that you have some means of applying the operator (e.g. a legacy PDE solver), so that you can observe input-output pairs.

Compression of a rank-structured matrix

Environment: We are given a rank-structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Objective: 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 $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Sample the column space of the matrix:

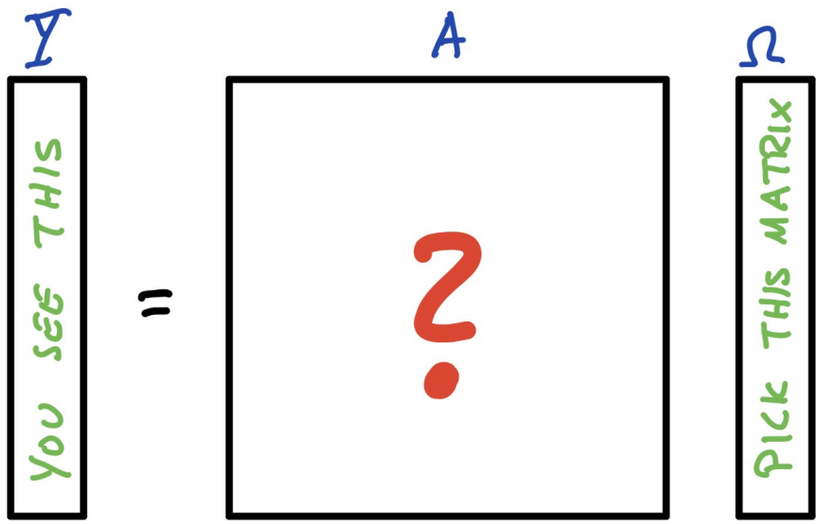

If $\mathbf{A} \neq \mathbf{A}^*$, then sample the row space too:

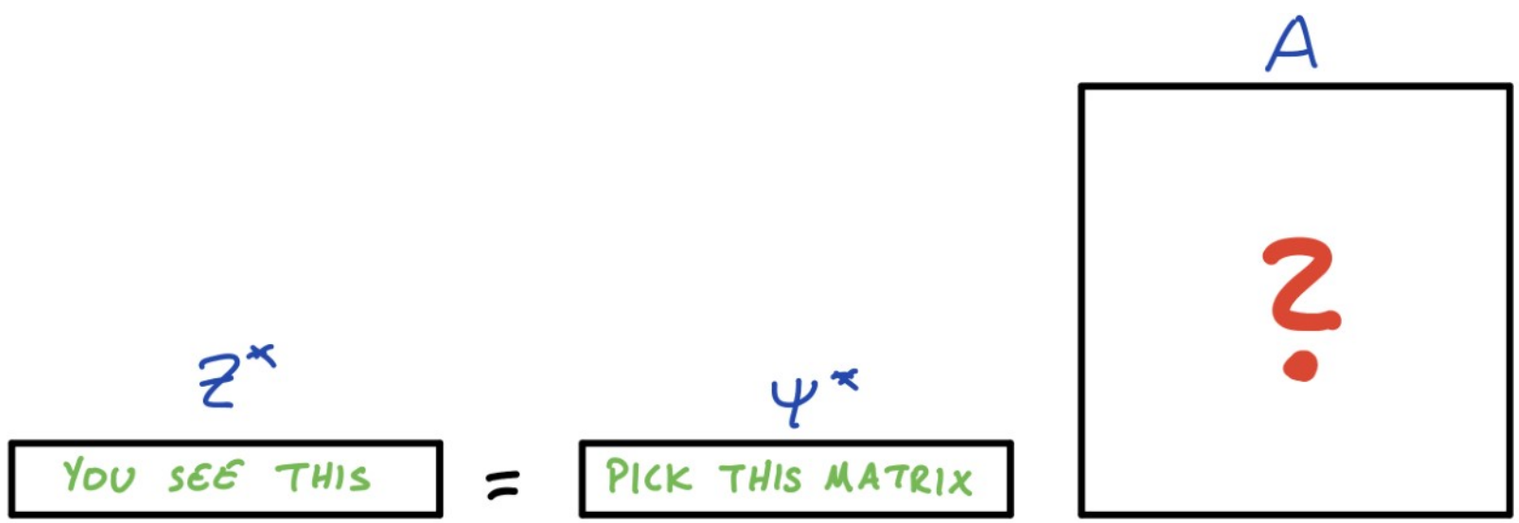

Compression of a rank-structured matrix

Environment: We are given a rank-structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Objective: 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 $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

The low rank case: In the particularly simple case where $\mathbf{A}$ has *global* rank k, we revert to the case we considered in the first part of the talk.

In the current framework, the randomized SVD takes the form:

- Set $s = k$ and draw a "test matrix" $\mathbf{\Omega} \in \mathbb{R}^{N \times s}$ from a Gaussian distribution.
- Form the "sample matrix" $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$.
- Build $\mathbf{\Psi}$ to hold an ON basis for ran($\mathbf{Y}$), e.g., $[\mathbf{\Psi}, \sim] = \mathtt{qr}(\mathbf{Y}, 0)$.
- Form $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$.

Then $\mathbf{A} = \mathbf{\Psi}\left(\mathbf{\Psi}^*\mathbf{A}\right) = \mathbf{\Psi}\mathbf{Z}^*$ with probability 1.

In the more typical case where $\mathbf{A}$ is only *approximately* of rank k, some *oversampling* is required to get a reliable scheme. (Say $s = k + 10$, or $s = 2k$, or some such.)

Rank structured case: Extract *all* the low-rank matrices, and *all* the dense blocks, from a very limited set of global "probes". How do you disentangle the mixed samples?

Compression of a rank-structured matrix

Environment: We are given a rank-structured matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. We assume that we can evaluate $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ and $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$ fast.

Objective: 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 $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$?

Why generalize from "global low rank" to "rank-structured":

- 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.
- Compute Dirichlet-to-Neumann (or Impedance-to-Impedance) operators explicitly whenever you have access to a fast PDE solvers.
- Compression of Schur complements that arise in the LU or Cholesky factorization of sparse matrices. This overcomes the key computational bottleneck, and for instance admits the acceleration of the LU factorization of a "finite element" matrix *from $O(N^2)$ to near linear complexity.*
- Multiplication of operators. $\rightarrow$ Multiphysics, multi-modal discretizations, etc.

Compression of a rank-structured matrix: How do you do it?

1. **With entry evaluation:** If you can evaluate individual entries of the target matrix, then you can simply subtract off the dense blocks from the sample matrix.
   - First algorithms of this nature to be developed (2008).
   - Leads to optimal complexity, and high practical speed.
   - Restrictive assumption! Excludes many important applications.

2. **With zero blocks in test matrix:** Create test matrices with strategically placed zero blocks that hit the dense blocks of the target matrix.
   - First complete "black box" algorithms (2011).
   - Every level must be processed independently $\rightarrow O(N \log N)$ complexity.
   - Pre-factors tend to be large. Can be reduced through graph coloring algorithms that design bespoke zero patterns for any given geometry.

3. **With algebraic structure in the test matrix:** It turns out to be possible to form test matrices that allow you to retro-actively go back and create zeros blocks where you like *after* extracting a universal sample.
   - High practical speed, and first method that is $O(N)$ and completely black box (2022).
   - Original method works best for 2D problems, or problems on surfaces in 3D.
   - Current work is aimed at finding different algebraic structures that lead to high efficiency for fully general 3D geometries.
   - It turns out to be possible to *invert* the matrix as you go!

Compression of a rank-structured matrix: How do you do it?

1. **With entry evaluation:** If you can evaluate individual entries of the target matrix, then you can simply subtract off the dense blocks from the sample matrix.

   *[P.G. Martinsson, SIMAX, **32**(4), 2011. Later improvements by Jianlin Xia, Xiaoye Sherry Li, . . .*]

2. **With zero blocks in test matrix:** Create test matrices with strategically placed zero blocks that hit the dense blocks of the target matrix.

   *[L. Lin, J. Lu, L. Ying, JCP, **230**(10), pp. 4071–4087, 2011; P.G. Martinsson, SISC, **38**(4), pp. A1959-A1986, 2016;*

   *J. Levitt & P.G. Martinsson, JCAM **451**(1), 2024.]*

3. **With algebraic structure in the test matrix:** It turns out to be possible to form test matrices that allow you to retro-actively go back and create zeros blocks where you like *after* extracting a universal sample.

   *[J. Levitt & P.G. Martinsson, SISC, **46**(3), 2024. K. Pearce, A. Yesypenko, J. Levitt, P.G. Martinsson, to appear in SIMAX*

   (arXiv:2501.05528). A. Yesypenko & P.G. Martinsson, to appear in J. Sci. Comp. (arxiv:2311.01451).]

4. **Plenty of related work recently!**

   Exploiting the continuum structure of the underlying operator: A. Townsend & N. Boullé (2022).

   Complexity analysis – what is the minimal number of matvecs required? Recent work by A. Townsend, D. Halikias, C. Musco & C. Musco, D. Persson, N. Boullé.

   Randomized compression of butterfly matrices: Path towards medium/high frequency problems. *[Y. Liu, X. Xing, H. Guo, E. Michielssen, P. Ghysels, X.S. Li. 2021], [Y. Li, H. Yang, E. Martin, K. Ho, and L. Ying, 2015].*

Compression of a rank-structured matrix: With entry evaluation

If you can evaluate individual entries of the target matrix, then you can simply subtract off the dense blocks from the sample matrix.

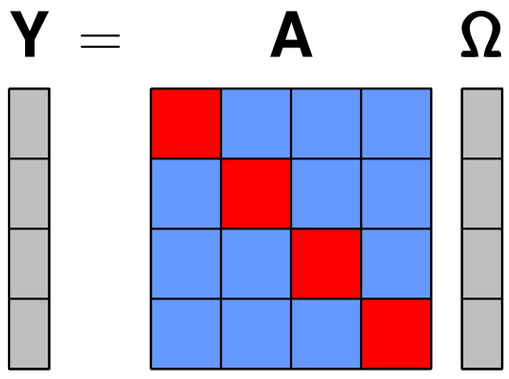

Explicitly form the diagonal blocks, and subtract their contributions:

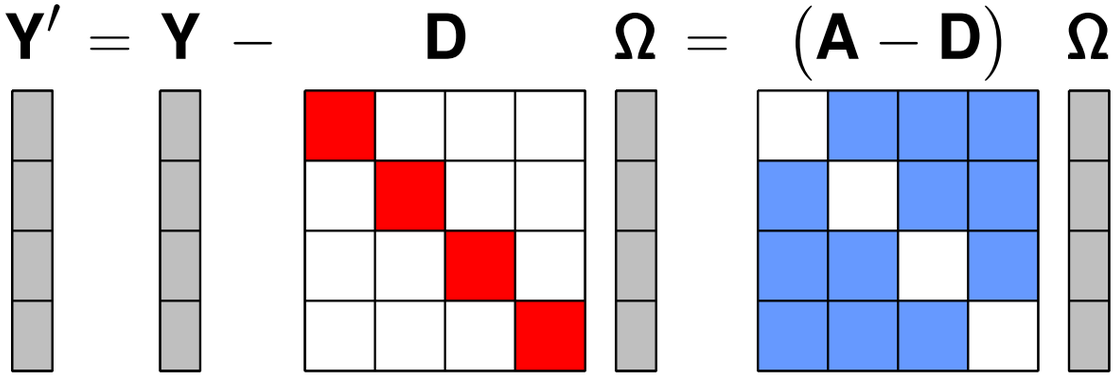

- First algorithms of this nature to be developed (2008).
- Leads to optimal complexity.
- High practical speed.
- Restrictive assumption! Excludes many important applications.

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

Compression of a rank-structured matrix: With zero blocks

The most straight-forward way to accomplish the task is to build test matrices with strategically placed zero blocks.

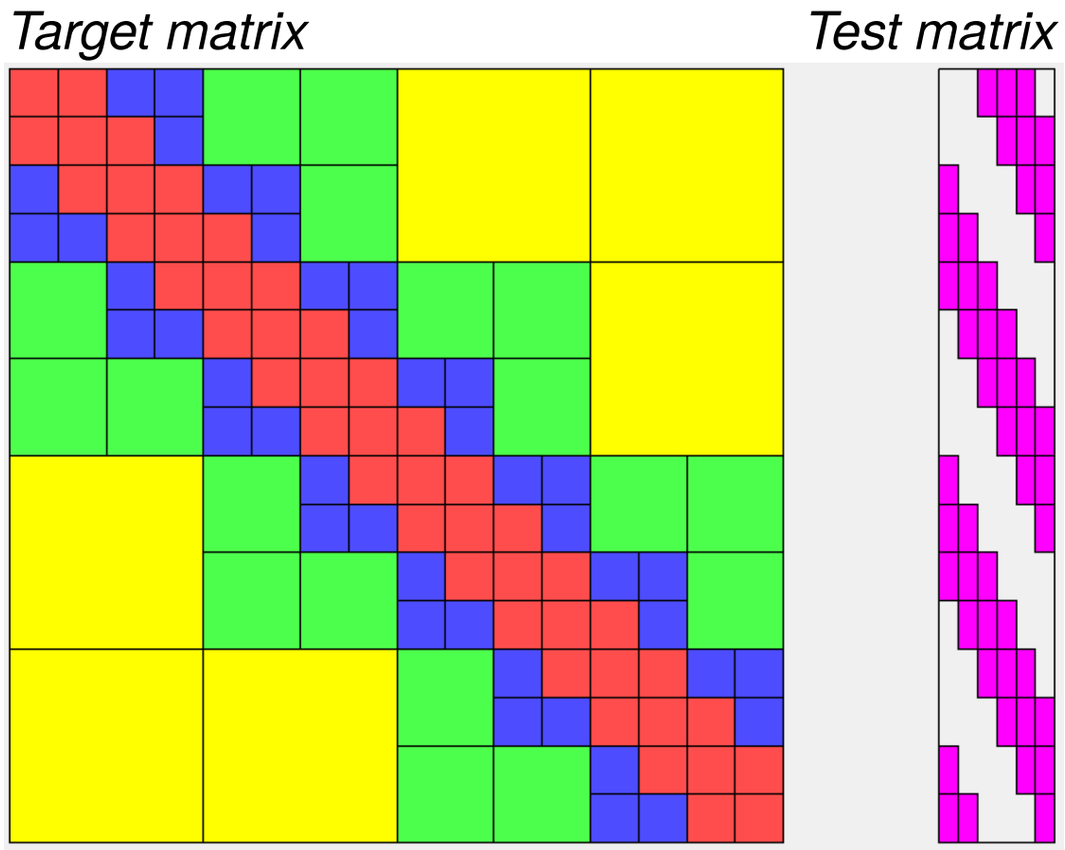

Low rank on level 2. *Low rank on level 3.* *Low rank on level 4.* *Dense blocks.* *Gaussian blocks.*

- First complete “black box” algorithms (2011).
- Every level must be processed independently $\rightarrow O(N \log N)$ complexity.
- Pre-factors tend to be large. Can be reduced through graph coloring algorithms that design bespoke zero patterns for any given geometry.

*L. Lin, J. Lu, L. Ying, JCP, **230**(10), pp. 4071–4087, 2011; P.G. Martinsson, SISC, **38**(4), pp. A1959-A1986, 2016; J. Levitt & P.G. Martinsson, JCAM **451**(1), 2024.*

Compression of a rank-structured matrix – A naive approach

Consider the task of finding a 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.

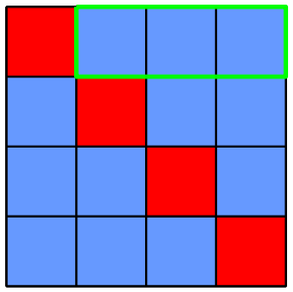

Compression of a rank-structured matrix – A naive approach

Consider the task of finding a 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.

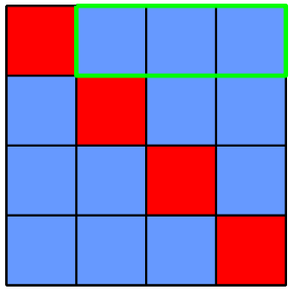

The naive approach is to sample with a random matrix $\mathbf{\Omega} \in \mathbb{R}^{N \times r}$, $r = k + 10$, that has a block of zeros in rows indexed by I_4. Then $\mathbf{Y}(I_4, :)$ will contain a sample of $\mathbf{A}(I_4, I_4^c)$.

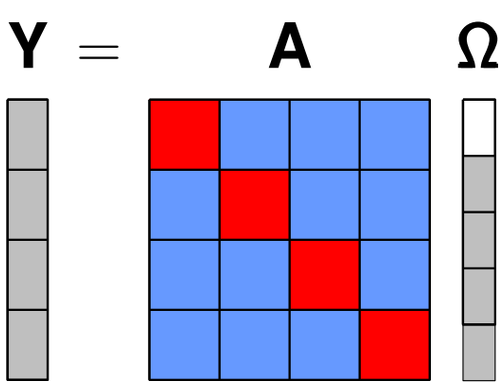

Compression of a rank-structured matrix – A naive approach

Consider the task of finding a 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.

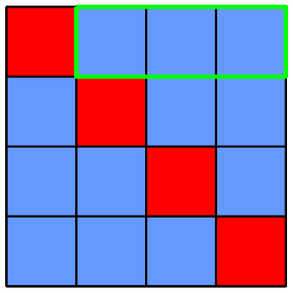

The naive approach is to sample with a random matrix $\mathbf{\Omega} \in \mathbb{R}^{N \times r}$, $r = k + 10$, that has a block of zeros in rows indexed by I_4. Then $\mathbf{Y}(I_4, :)$ will contain a sample of $\mathbf{A}(I_4, I_4^c)$.

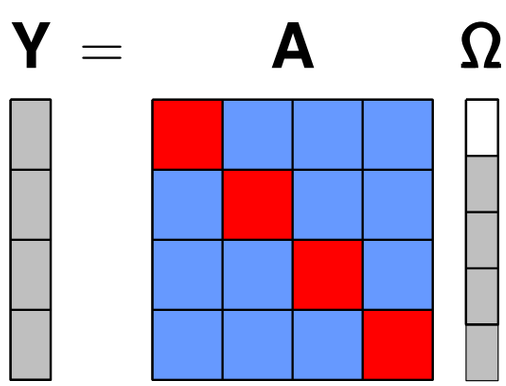

This scheme requires taking a separate set of r samples *for each leaf node*, for a total of $\sim rN/m$ samples. There is a lot of wasted information in $\mathbf{Y}$.

Compression of a rank-structured matrix – the "almost" black-box case

Sample $\mathbf{A}$ with *fixed* dense random matrices $\mathbf{\Omega} \in \mathbb{R}^{N \times r}$ and $\mathbf{\Psi} \in \mathbb{R}^{N \times r}$:

Assumption: You can do matvecs and *entry evaluation*. (More general soon.)

Compression of a rank-structured matrix – the "almost" black-box case

Sample $\mathbf{A}$ with *fixed* dense random matrices $\mathbf{\Omega} \in \mathbb{R}^{N\times r}$ and $\mathbf{\Psi} \in \mathbb{R}^{N\times r}$:

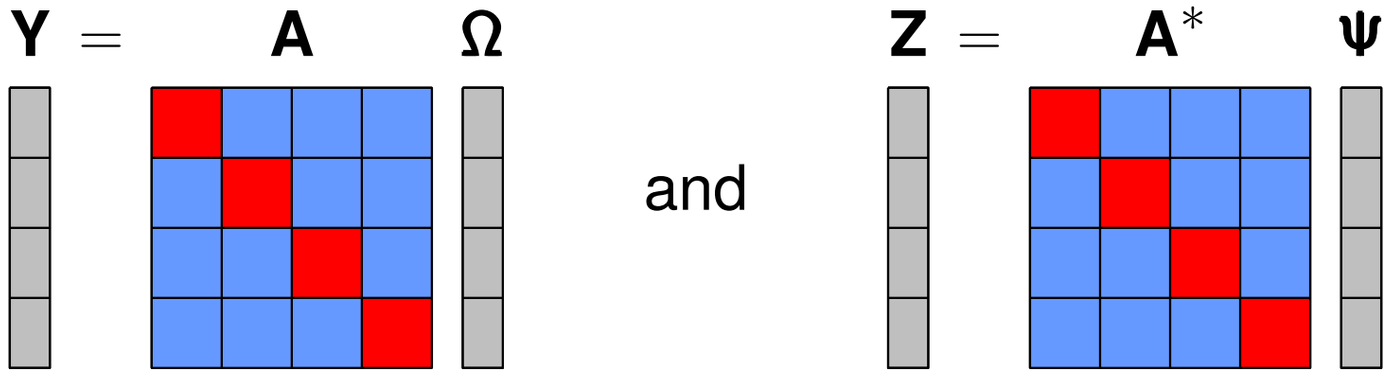

Assumption: You can do matvecs and *entry evaluation*. (More general soon.)

In this case, we can explicitly form the diagonal blocks, and subtract their contributions:

$$\mathbf{Y}' = \mathbf{Y} - \mathbf{D}\,\mathbf{\Omega} = (\mathbf{A} - \mathbf{D})\,\mathbf{\Omega}$$

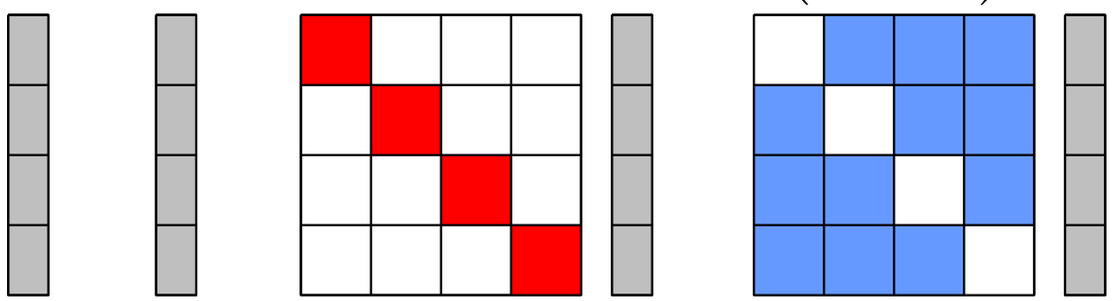

Processing $\mathbf{Z}$ analogously, we obtain basis matrices $\mathbf{Y}_j$ and $\mathbf{Z}_j$ for $j \in \{4, 5, 6, 7\}$ such that

$$\mathbf{A}_{i,j} \approx \mathbf{Y}_i\,\mathbf{B}_{i,j}\,\mathbf{Z}_j^*, \qquad i \neq j,$$

for *some* small matrices $\mathbf{B}_{i,j}$.

In a final step, use the ID to build the matrices $\mathbf{B}_{i,j}$ via entry evaluation.

Compression of a rank-structured matrix – fully black box

Sample $\mathbf{A}$ with a completely dense random matrix $\mathbf{\Omega} \in \mathbb{R}^{N \times (r+m)}$, where m is the leaf node size. (Think $m \approx 2k$ and $r = k + 10$.)

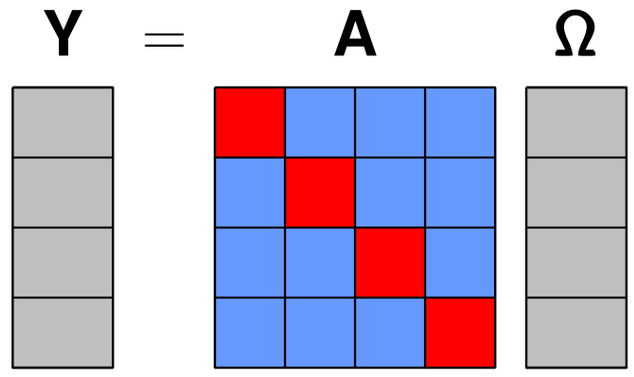

Let us consider the problem of finding a basis matrix $\mathbf{U}_4$ for the block $\mathbf{A}(I_4, I_4^c)$.

Compression of a rank-structured matrix – fully black box

Sample $\mathbf{A}$ with a completely dense random matrix $\mathbf{\Omega} \in \mathbb{R}^{N\times(r+m)}$, where m is the leaf node size. (Think $m \approx 2k$ and $r = k + 10$.)

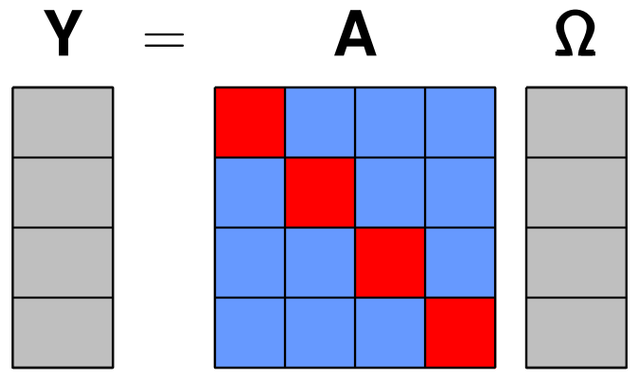

Let us consider the problem of finding a basis matrix $\mathbf{U}_4$ for the block $\mathbf{A}(I_4, I_4^c)$.

Since $\mathbf{\Omega}(I_4, :)$ is of size $m \times (r + m)$, it has a nullspace of dimension at least r. Let

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

be an $(r + m) \times r$ orthonormal basis of the nullspace of $\mathbf{\Omega}(I_4, :)$.

Compression of a rank-structured matrix – fully black box

Sample **A** with a completely dense random matrix $\mathbf{\Omega} \in \mathbb{R}^{N\times(r+m)}$, where m is the leaf node size. (Think $m \approx 2k$ and $r = k + 10$.)

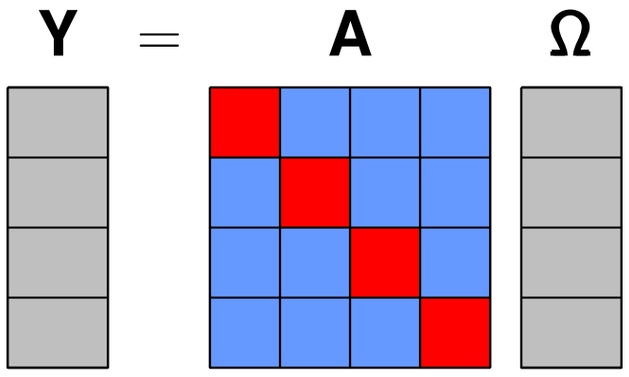

Let us consider the problem of finding a basis matrix $\mathbf{U}_4$ for the block $\mathbf{A}(I_4, I_4^c)$.
Since $\mathbf{\Omega}(I_4, :)$ is of size $m \times (r + m)$, it has a nullspace of dimension at least r. Let

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

be an $(r + m) \times r$ orthonormal basis of the nullspace of $\mathbf{\Omega}(I_4, :)$.Then

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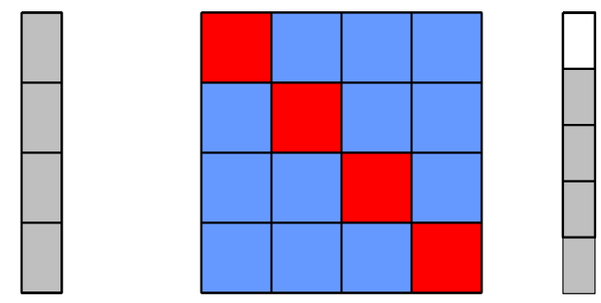 $\mathbf{\Omega Q}_4$.

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

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

Compression of a rank-structured matrix – fully black box

- For each leaf node τ, we compute

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

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

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

Compression of a rank-structured matrix – fully black box

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* $\mathbf{A}^*$*.*

Form Gaussian random random matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ of size $N \times s$.

Multiply $\mathbf{Y} = \mathbf{A\Omega}$ and $\mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}$.

Compress level by level from finest to coarsest.

for level $\ell = L, L-1, \ldots, 0$ **do**

  for node τ in level ℓ **do**

    if τ is a leaf node **then**

$$\mathbf{\Omega}_\tau = \mathbf{\Omega}(I_\tau, :), \qquad \mathbf{\Psi}_\tau = \mathbf{\Psi}(I_\tau, :)$$

$$\mathbf{Y}_\tau = \mathbf{Y}(I_\tau, :), \qquad \mathbf{Z}_\tau = \mathbf{Z}(I_\tau, :)$$

    else

Let α and β denote the children of τ.

$$\mathbf{\Omega}_\tau = \begin{bmatrix} \mathbf{V}_\alpha^* \mathbf{\Omega}_\alpha \\ \mathbf{V}_\beta^* \mathbf{\Omega}_\beta \end{bmatrix}, \qquad \mathbf{\Psi}_\tau = \begin{bmatrix} \mathbf{U}_\alpha^* \mathbf{\Psi}_\alpha \\ \mathbf{U}_\beta^* \mathbf{\Psi}_\beta \end{bmatrix}$$

$$\mathbf{Y}_\tau = \begin{bmatrix} \mathbf{U}_\alpha^* (\mathbf{Y}_\alpha - \mathbf{D}_\alpha \mathbf{\Omega}_\alpha) \\ \mathbf{U}_\beta^* (\mathbf{Y}_\beta - \mathbf{D}_\beta \mathbf{\Omega}_\beta) \end{bmatrix}, \quad \mathbf{Z}_\tau = \begin{bmatrix} \mathbf{V}_\alpha^* (\mathbf{Z}_\alpha - \mathbf{D}_\alpha^* \mathbf{\Psi}_\alpha) \\ \mathbf{V}_\beta^* (\mathbf{Z}_\beta - \mathbf{D}_\beta^* \mathbf{\Psi}_\beta) \end{bmatrix}$$

    if level $\ell > 0$ **then**

$$\mathbf{Q}_\tau = \texttt{nullspace}(\mathbf{\Omega}_\tau, r), \qquad \mathbf{P}_\tau = \texttt{nullspace}(\mathbf{\Psi}_\tau, r)$$

$$\mathbf{U}_\tau = \texttt{qr}(\mathbf{Y}_\tau \mathbf{Q}_\tau, r), \qquad \mathbf{V}_\tau = \texttt{qr}(\mathbf{Z}_\tau \mathbf{P}_\tau, r)$$

$$\mathbf{D}_\tau = (\mathbf{I} - \mathbf{U}_\tau \mathbf{U}_\tau^*)\mathbf{Y}_\tau \mathbf{\Omega}_\tau^\dagger + \mathbf{U}_\tau \mathbf{U}_\tau^* \left((\mathbf{I} - \mathbf{V}_\tau \mathbf{V}_\tau^*)\mathbf{Z}_\tau \mathbf{\Psi}_\tau^\dagger\right)^*$$

    else

$$\mathbf{D}_\tau = \mathbf{Y}_\tau \mathbf{\Omega}_\tau^\dagger$$

Compression of a rank-structured matrix – Finding 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.

Compression of a rank-structured matrix – Finding 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.

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

Compression of a rank-structured matrix – Finding 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.

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

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

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

Compression of a rank-structured matrix – Compressing $\tilde{\mathbf{A}}^{(L)}$

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

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

and rearrange to obtain

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

Compression of a rank structured matrix: RSRS

Suppose you have extracted samples

$$\mathbf{Y} = \mathbf{A}\mathbf{\Omega} \qquad \text{and} \qquad \mathbf{Z} = \mathbf{A}^*\mathbf{\Psi}.$$

We use the set $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$ to extract the information to compress the first block using "block nullification" and "block extraction".

Then do one step of strong recursive skeletonization to obtain a partial factorization

$$\mathbf{A} = \mathbf{L}\tilde{\mathbf{A}}\mathbf{R},$$

where $\mathbf{L}$ and $\mathbf{R}$ each consists of two block elimination steps, and $\tilde{\mathbf{A}}$ is a matrix where some blocks have been zeroed out, and some have been modified.

We next seek a sample of $\tilde{\mathbf{A}}$. To do this, we form

$$\tilde{\mathbf{Y}} := \mathbf{L}^{-1}\mathbf{Y} = \mathbf{L}^{-1}\mathbf{A}\mathbf{\Omega} = \mathbf{L}^{-1}\big(\mathbf{L}\tilde{\mathbf{A}}\mathbf{R}\big)\mathbf{\Omega} = \tilde{\mathbf{A}}\tilde{\mathbf{\Omega}},$$

where we defined

$$\tilde{\mathbf{\Omega}} := \mathbf{R}\mathbf{\Omega}.$$

Analogously, form $\{\tilde{\mathbf{Z}}, \tilde{\mathbf{\Psi}}\}$.

Proceed to the next box using the set $\{\tilde{\mathbf{Y}}, \tilde{\mathbf{\Omega}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{\Psi}}\}$ in place of $\{\mathbf{Y}, \mathbf{\Omega}, \mathbf{Z}, \mathbf{\Psi}\}$.

Key points

- Randomized algorithms for low rank approximation are highly efficient.
  - Interaction with target matrix only through matrix-matrix multiplication $\rightarrow$ very high practical speed.
  - Particularly efficient for GPUs, out-of-core computing, distributed memory, etc.
  - Structured random maps ("fast J-L transform"): Improves on $O(mnk)$ complexity.
  - Single pass algorithms have been developed for *streaming environments.*

    Not possible with deterministic methods!
  - Extension to the case of *tensors* is active area of research.

- Randomized algorithms for solving linear systems
  - Overdetermined least squares is particularly successful.
  - The talk only scratched the surface – randomized Kaczmarz, acceleration of Krylov, ...

- Randomized algorithms based on *sampling* make (some) huge problems tractable.
  - Success stories: kernel ridge regression, computational chemistry, tensor approximation, ...

- Randomized compression of global operators.
  - Black box randomized algorithms for compressing global operators have been established.
  - The combination of "fully black box" and "true linear complexity" was realized only recently.
  - Powerful tools in the construction of fast direct solvers for elliptic PDEs.
  - Path to solving multi-physics (and "multi-discretization") problems.
  - Going from the discrete to the continuum in randomized compression.

Surveys:

- P.G. Martinsson and M. O'Neil, "Fast Direct Solvers for Elliptic PDEs".
  In review, 2025. (Arxiv report 2511.07773)
  Review of direct algorithms for global operators in scientific computing.
- K. Pearce and P.G. Martinsson ...
- P.G. Martinsson and J. Tropp, "Randomized Numerical Linear Algebra: Foundations & Algorithms".
  Acta Numerica, 2020. (Arxiv report 2002.01387)
  Long survey summarizing major findings in the field in the past decade.
- P.G. Martinsson, "Randomized methods for matrix computations." The Mathematics of Data, IAS/Park City Mathematics Series, 25(4), pp. 187 - 231, 2018.
  Book chapter that is written to be accessible to a broad audience. Focused on practical aspects.
- N. Halko, P.G. Martinsson, J. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions." *SIAM Review*, 53(2), 2011, pp. 217-288.
  Survey that describes the randomized SVD and its variations.

Tutorials, summer schools, etc:

- 2020: 3 lecture mini course on randomized linear algebra, KTH, Stockholm. Videos available.
- 2016: Park City Math Institute (IAS): *The Mathematics of Data.*
- 2014: CBMS summer school at Dartmouth College. 10 lectures on YouTube.
- 2009: NIPS tutorial lecture, Vancouver, 2009. Online video available.

DOE report on randomized algorithms: `https://arxiv.org/abs/2104.11079` (2021)