A sparsified nested dissection algorithm

Léopold Cambier*, Eric Darve*, Chao Chen‡
Siva Rajamanickam†, Erik Boman†, Raymond Tuminaro†

*Stanford University, ‡ University of Texas at Austin, †Sandia National Laboratories

PDE's $\approx$ sparse $Ax = b$

$$-\nabla \cdot (a(x)\nabla u(x)) = f(x) \Rightarrow -U_{i-1} + 2U_i - U_{i+1} = f_i$$
$$\Rightarrow Ax = b$$

Nested Dissection & Linear Systems

Find sets of rows/cols L, R and S and order A such that

$$A = \begin{bmatrix} A_{LL} & & A_{LS} \\ & A_{RR} & A_{RS} \\ A_{SL} & A_{RS} & A_{SS} \end{bmatrix}$$

Then eliminate L and R. This creates fill-in on A_{SS}.

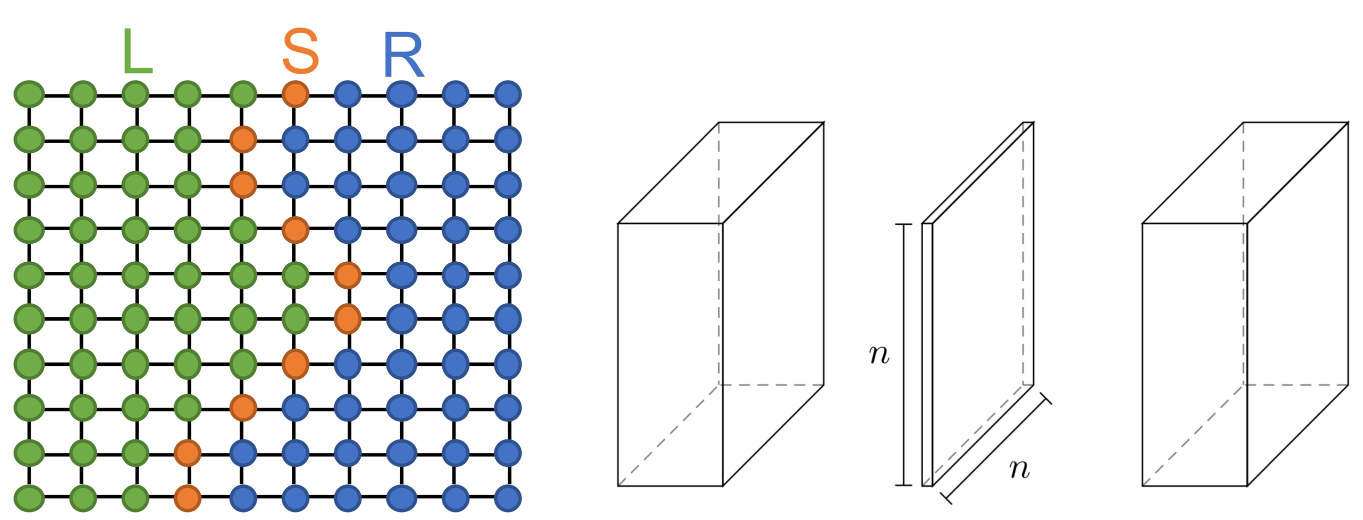

Same procedure is recursively applied on L and R. Issue? On "3D graphs" of size $N \approx n^3$, separators have size $N^{2/3} = n^2$. Hence, since A_{ss} is typically dense, the cost of the last elimination is $N^{2/3 \cdot 3} = N^2$. Too much.

Sparsification

Select p, a set of rows/cols at the interface between two eliminated interiors. Scale A_{pp} to I using Cholesky. Then consider all their neighbors n of p and approximate $A_{pn} \approx Q_c W_{cn} + Q_f W_{fn}$ with $\|W_{fn}\| \approx \varepsilon$. Then

$$\begin{bmatrix} Q^T & \\ & I \end{bmatrix} \begin{bmatrix} I & A_{pn} \\ A_{np} & A_{nn} \end{bmatrix} \begin{bmatrix} Q & \\ & I \end{bmatrix} = \begin{bmatrix} I & & \varepsilon \\ & I & W_{cn} \\ \varepsilon & W_{nc} & A_{nn} \end{bmatrix}$$

The size of p has been reduced to $|c|$, the ε-rank of A_{pn}.

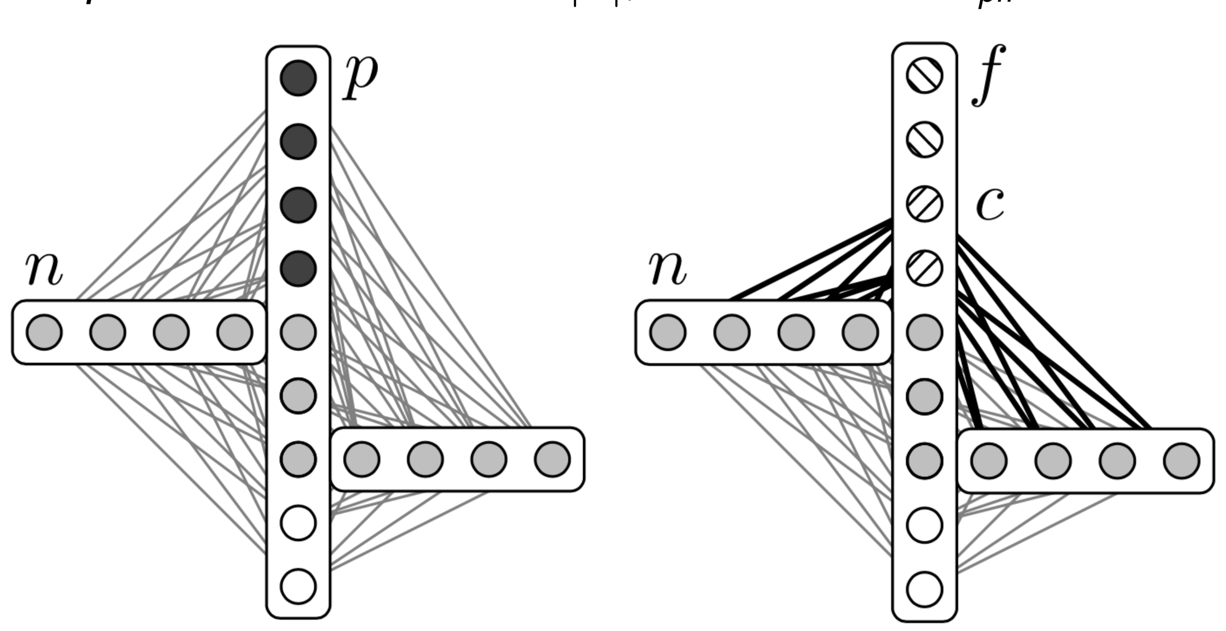

After the sparsification, all clusters & edges are smaller, but the matrix connectivity is unchanged. In particular, no fill-in is introduced.

Before sparsification | After sparsification

Matrix graph

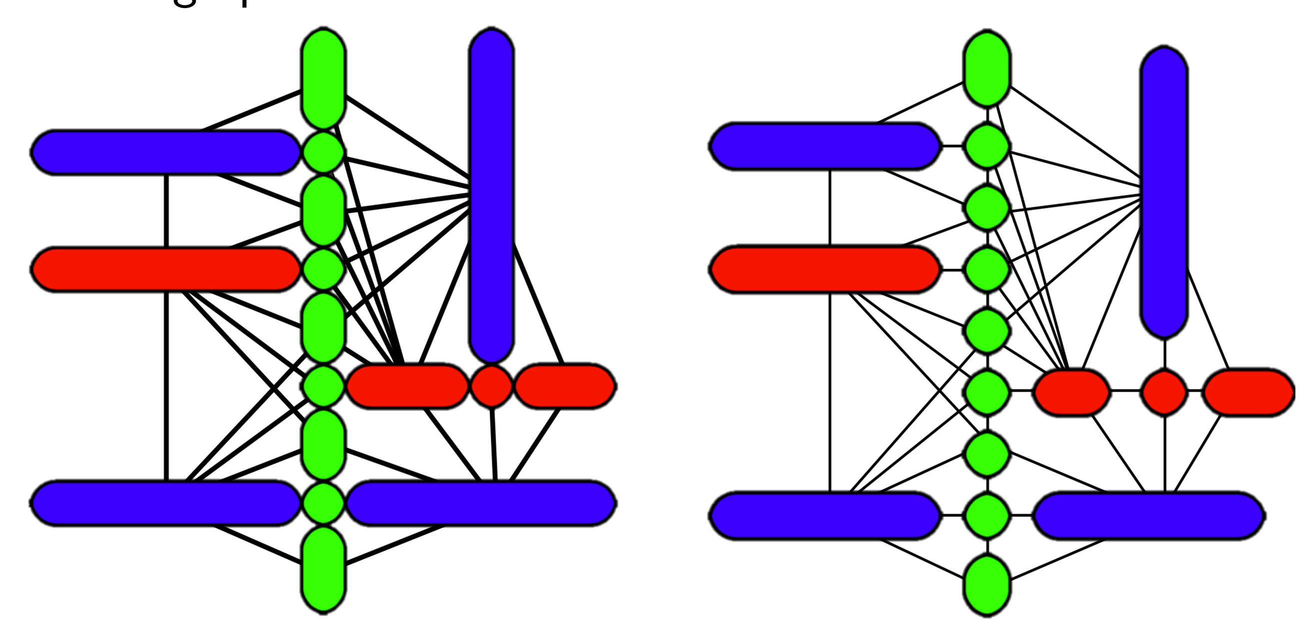

Trailing matrix

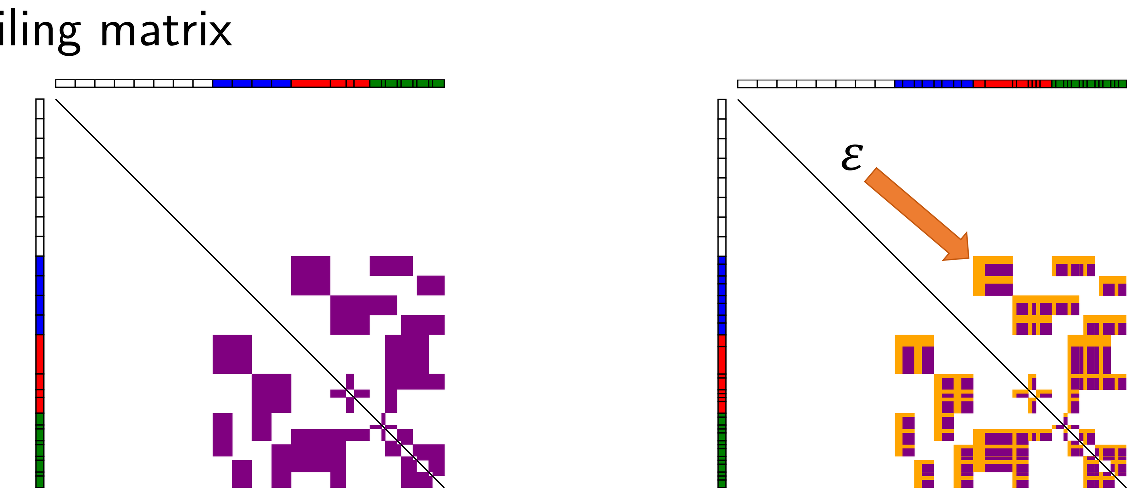

Elliptic PDEs (SPD)

Theorem: if A is SPD, for all $\varepsilon \geq 0$, the factorization never breaks down.

- For 2D graphs, top separator size becomes $\approx \mathcal{O}(1)$ (vs $\mathcal{O}\left(N^{1/2}\right)$)
- For 3D graphs, top separator size becomes $\approx \mathcal{O}\left(N^{1/3}\right)$ (vs $\mathcal{O}\left(N^{2/3}\right)$).

A very ill-conditioned 2D problem, modeling ice-flows on Antarctica

N	spaND t_F (s.)	t_S (s.)	n_{CG}	size_{Top}	mem_F (10^9)
5 layers					
629 544 (16 km)	6	3	7	78	0.15
2 521 872 (8 km)	27	19	8	88	0.63
10 096 080 (4 km)	107	114	10	99	2.61
10 layers					
1 154 164 (16 km)	24	8	8	137	0.42
4 623 432 (8 km)	94	44	8	147	1.73
18 509 480 (4 km)	500	384	10	159	7.59

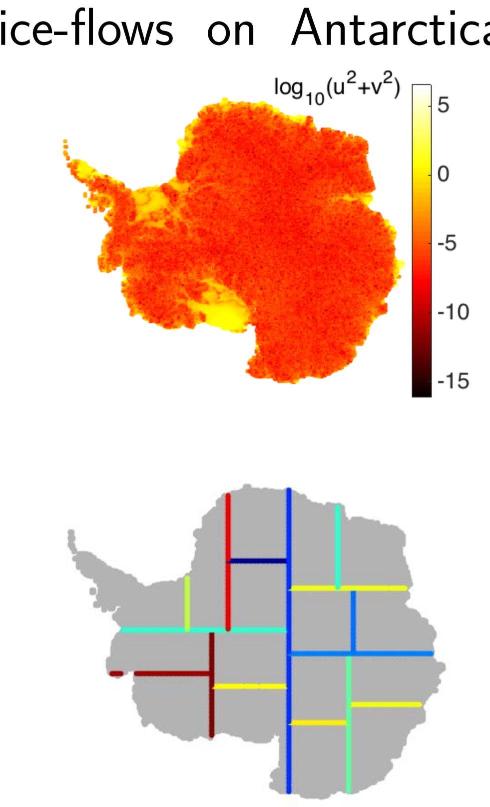

General matrices

For problems with $\approx$ symmetric sparsity patterns, the same algorithm can be applied with some changes:

- Scale pivot from A_{pp} to I using SVD or full-pivoting LU. Partial pivoted LU can amplify low-rank approximation errors excessively.
- Compress both lower and upper parts $Q_c \left[W_{cn} W_{nc}^\top\right] \approx \left[A_{pn}\ A_{np}^\top\right]$

Advection-diffusion (unsymmetric)

$$\text{3D:} \quad -a\nabla u + b \cdot \nabla u = f, u|_\Omega = 0$$

Boundary value problem.

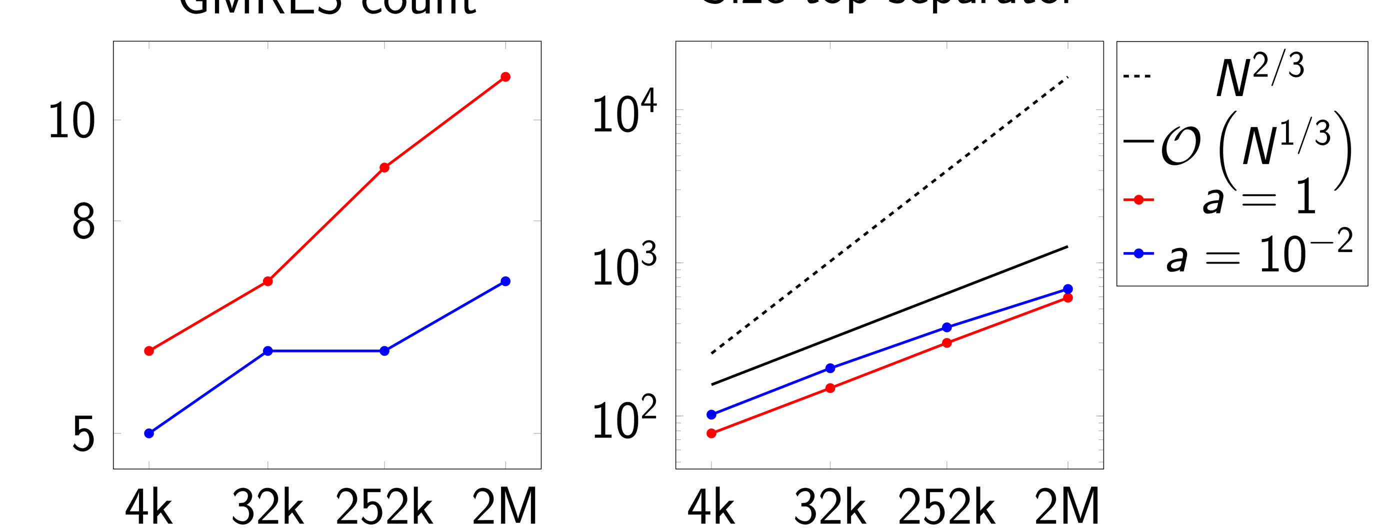

$$\text{2D+time:} \quad \frac{\partial u}{\partial t} = b \cdot \nabla u, \text{ periodic BCs}$$

Implicit Euler with time step dt, spatial discretization using DG with $dx \approx 1/N^{1/2}$.

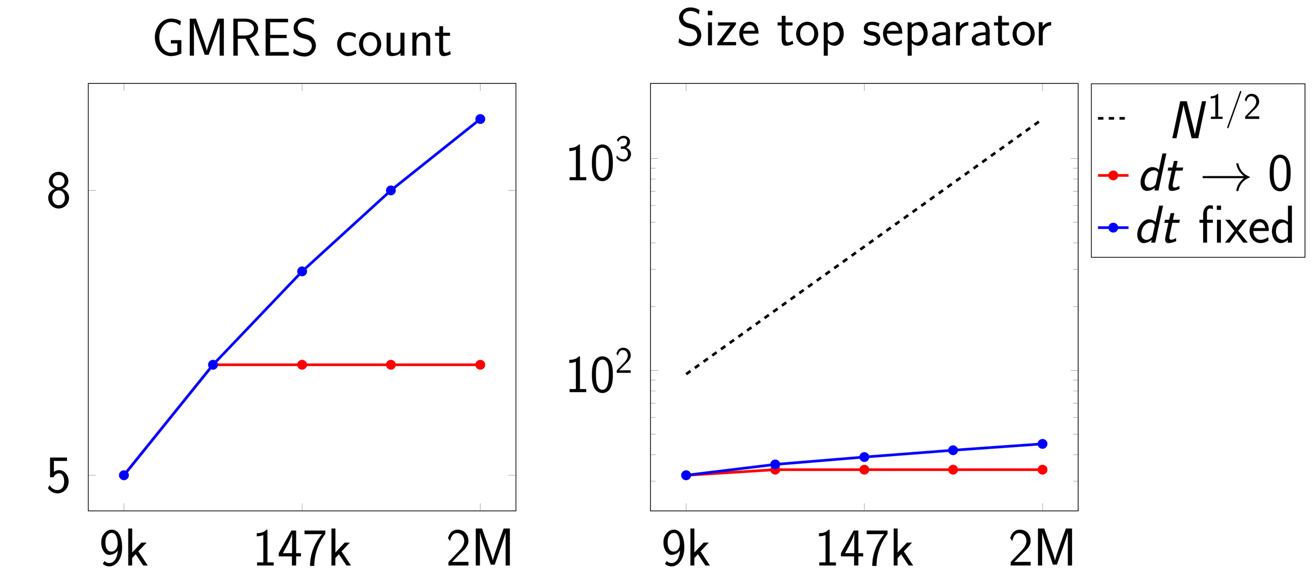

References & Acknowledgements

L. Cambier, C. Chen, E. G. Boman, S. Rajamanickam, R. S. Tuminaro, and E. Darve. An algebraic sparsified nested dissection algorithm using low-rank approximations. arXiv preprint arXiv:1901.02971, 2019.

Work funded by Sandia National Laboratories and Total