



Greedy Dissection Method for Shared Parallelism in Incomplete Factorization within INMOST Platform

Kirill Terekhov¹

¹Marchuk Institute of Numerical Mathematics of the Russian Academy of Sciences



September 28, 2021



INMOST

INMOST (www.inmost.org, www.inmost.ru) is a short for:

Integrated

Numerical

Modeling and

Object-oriented

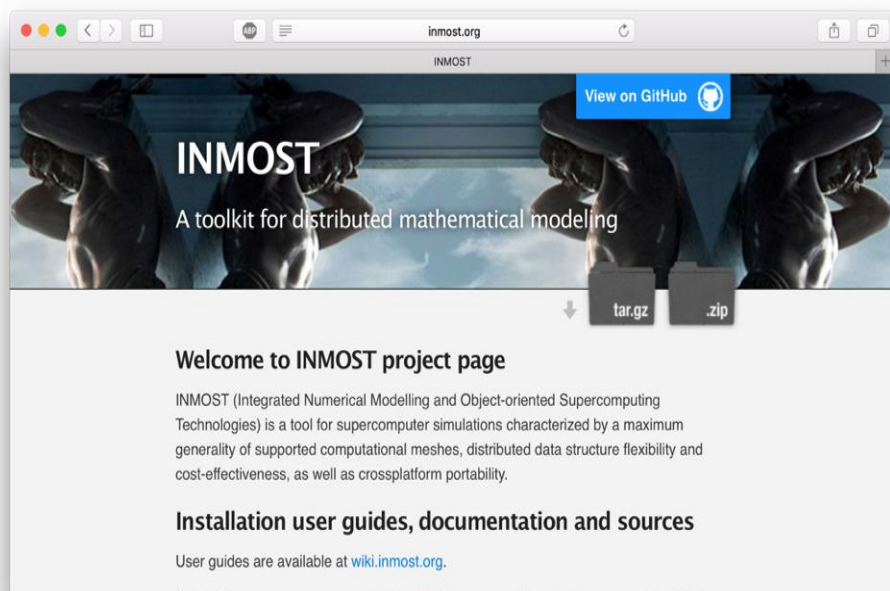
Supercomputing

Technologies

- Distributed meshes
- Distributed linear system assembly
- Parallel linear solver
- Automatic differentiation
- Nonlinear system assembly
- Coupling of unknowns and models


First version during 2012 internship at ExxonMobil

Contributors: Igor Konshin, Kirill Nikitin, Alexander Danilov, Ivan Kapyrin, Yuri Vassilevski, Alexei Chernyshenko (INM RAS, IBRAE RAS), Igor Kaporin (CMC RAS) Dmitri Bagaev, Andrei Burachkovski (MSU), Ruslan Yanbarisov, Alexei Logkiy, Sergei Petrov, Ivan Butakov (MIPT), German Kopytov (BFU), Timur Garipov, Pavel Tomin, Christine Mayer (Stanford), Ahmad Abushaikha, Longlong Li (HBKU), et al






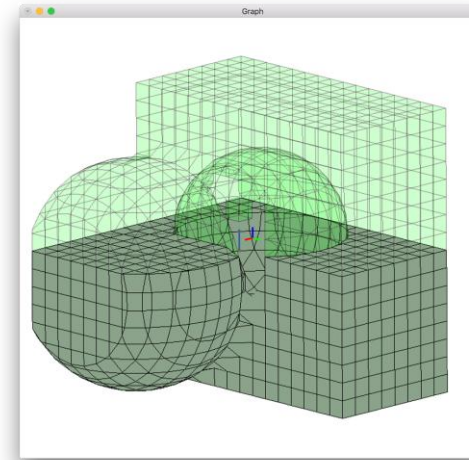
Grids




Human body





Geological grids with faults and pinch-outs,
support for commercial formats of oil & gas
simulators



Complex modifications




Fracture opening






Element collapse



Dynamic grids




OctreeCutcell example in INMOST-
Graphics repository



AdaptiveMesh example
for general grid
adaptation



Parmetis_AdaptiveReport






Models


Navier-Stokes

$$\rho \frac{\partial \mathbf{u}}{\partial t} + \text{div}(\rho \mathbf{u} \mathbf{u}^T - \mu \nabla \mathbf{u} + p \mathbf{I}) = 0$$

$$\text{div}(\mathbf{u}) = 0$$




Blood coagulation



Freesurface flows

$$\frac{\partial \varphi}{\partial t} + \text{div}(\varphi \mathbf{u}) = 0$$

$$|\nabla \varphi| = 1$$




Multiphase filtration

$$\frac{\partial \rho_w \theta S_w}{\partial t} - \nabla \cdot (\lambda_w \mathbb{K}(\nabla p - \rho_w g \nabla z)) = q_w$$

$$\frac{\partial \rho_o \theta S_o}{\partial t} - \nabla \cdot (\lambda_o \mathbb{K}(\nabla p - \nabla P_{C_o} - \rho_w g \nabla z)) = q_o$$

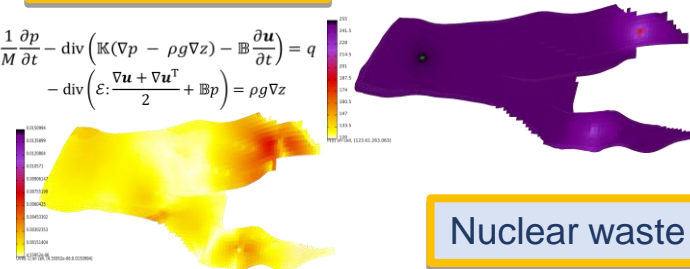
$$\frac{\partial \rho_g \theta (R S_o + S_g)}{\partial t} - \nabla \cdot (\lambda_g \mathbb{K}(\nabla p - \nabla P_{C_g} - \rho_g g \nabla z)) = q_g$$




Poromechanics

$$\frac{1}{M} \frac{\partial p}{\partial t} - \text{div}(\mathbb{K}(\nabla p - \rho g \nabla z) - \mathbb{B} \frac{\partial \mathbf{u}}{\partial t}) = q$$

$$-\text{div}\left(\varepsilon: \frac{\nabla \mathbf{u} + \nabla \mathbf{u}^T}{2} + \mathbb{B} p\right) = \rho g \nabla z$$




Nuclear waste




Mechanics


$$-\text{div}(\boldsymbol{\sigma}) = 0, \quad \mathbb{C}: \boldsymbol{\sigma} = \frac{\nabla \mathbf{u} + \nabla \mathbf{u}^T}{2}$$



Contact mechanics



Fracturing





Jacobian Structure in Multiphysics Problem

Models connect through

functions: mobility, density, porosity, properties,...

fluxes: Biot term, incompressibility conditions, capillary pressure, ...

right hand side: reactions, ...

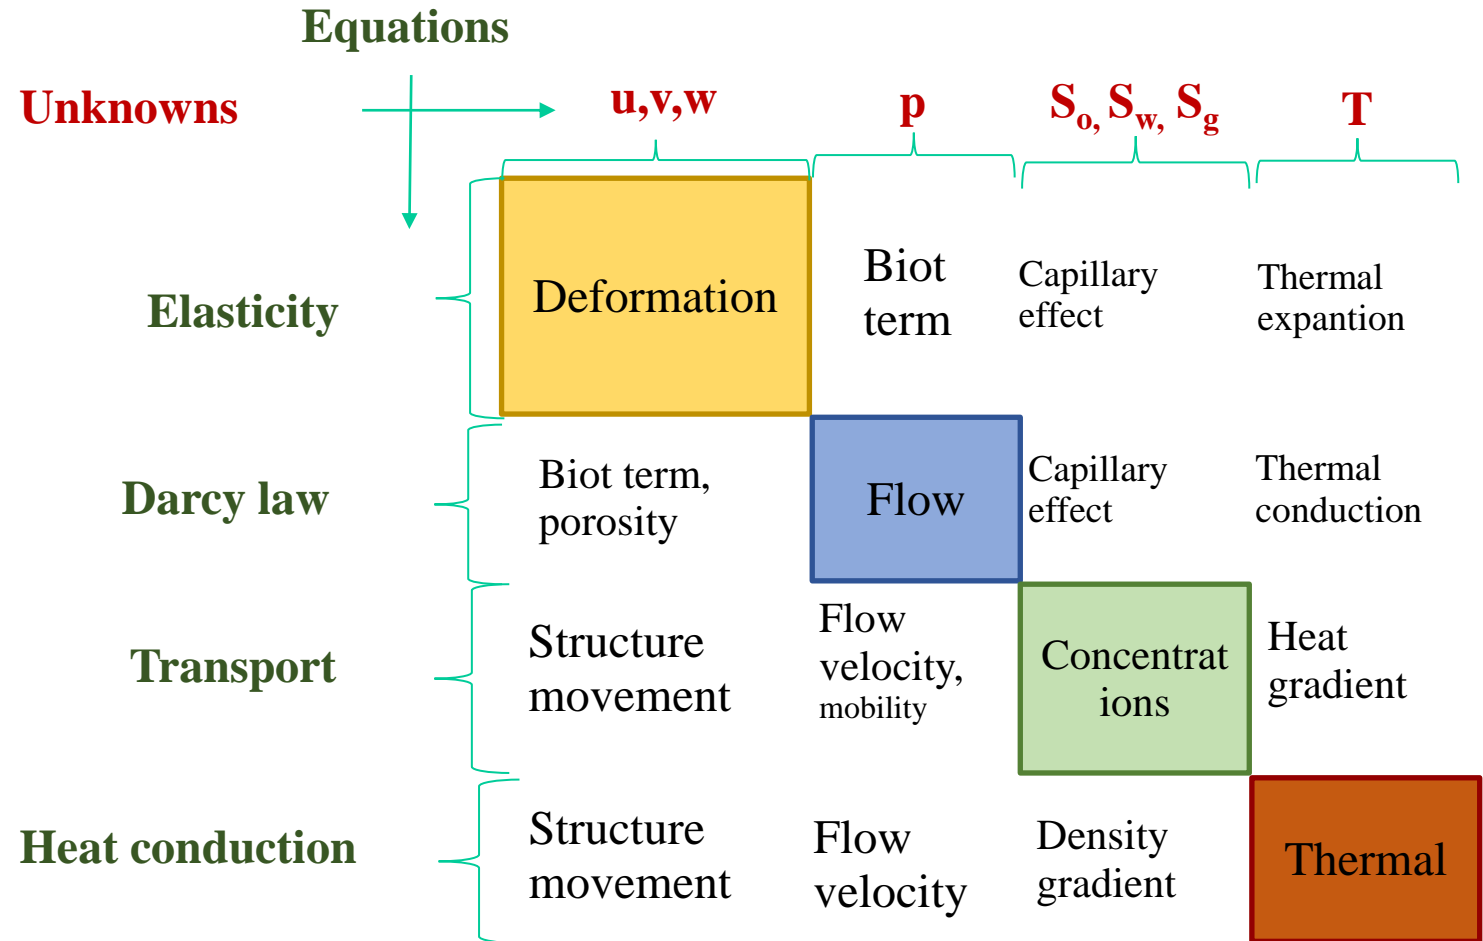
Ability to **control:**

models, unknowns, functions, coupling conditions, residual and Jacobian assembly.

Fully implicit solution:

saddle system, inf-sup stability issue and **complex linear systems**

Model decomposition for **reservoir simulator**:





Large systems

- Navier-Stokes system:

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \operatorname{div}(\rho \mathbf{u} \mathbf{u}^T - \mu \nabla \mathbf{u} + p \mathbb{I}) = -\frac{\mu}{K_f} \mathbf{u},$$

$$\operatorname{div}(\rho \mathbf{u}) = 0,$$

- Prothrombin (II):

$$\frac{\partial P}{\partial t} + \operatorname{div}(P \mathbf{u} - D \nabla P) = -(k_1 \phi_c + k_2 B_a + k_3 T + k_4 T^2 + k_5 T^3) P,$$

- Thrombin (IIa):

$$\frac{\partial T}{\partial t} + \operatorname{div}(T \mathbf{u} - D \nabla T) = (k_1 \phi_c + k_2 B_a + k_3 T + k_4 T^2 + k_5 T^3) P - k_6 A T,$$

- Clot formation factors (IXa, Xa):

$$\frac{\partial B_a}{\partial t} + \operatorname{div}(B_a \mathbf{u} - D \nabla B_a) = (k_7 \phi_c + k_8 T)(B_0 - B_a) - k_9 A B_a,$$

- Antithrombin (ATIII):


$$\frac{\partial A}{\partial t} + \operatorname{div}(A \mathbf{u} - D \nabla A) = -k_6 A - k_9 A B_a,$$

- Fibrinogen (I):

$$\frac{\partial F_g}{\partial t} + \operatorname{div}(F_g \mathbf{u} - D \nabla F_g) = -\frac{k_{10} T F_g}{K_{10} + F_g},$$

- Fibrin (Ia):

$$\frac{\partial F}{\partial t} + \operatorname{div}(F \mathbf{u} - D \nabla F) = \frac{k_{10} T F_g}{K_{10} + F_g} - k_{11} F,$$



- Fibrin-polymer:

$$\frac{\partial F_p}{\partial t} = k_{11} F,$$

- Inactivated platelets:

$$\frac{\partial \phi_f}{\partial t} + \operatorname{div}(k(\phi_c, \phi_f)(\phi_f \mathbf{u} - D_p \nabla \phi_f)) = (k_{12} T - k_{13} \phi_c) \phi_f,$$

- Activated platelets:

$$\frac{\partial \phi_c}{\partial t} + \operatorname{div}(k(\phi_c, \phi_f)(\phi_c \mathbf{u} - D_p \nabla \phi_c)) = -(k_{12} T - k_{13} \phi_c) \phi_f,$$

- Platelets mobility:

$$k(\phi_c, \phi_f) = \tanh\left(\pi \left(1 - \frac{\phi_c + \phi_f}{\phi_{\max}}\right)\right),$$

- Permeability:

$$\frac{1}{K_f} = \frac{16}{\alpha^2} \phi_p^{\frac{3}{2}} (1 + 56 \phi_p) \frac{\phi_{\max} + \phi_c}{\phi_{\max} - \phi_c}, \quad \phi_p = \min\left(\frac{7}{10}, \frac{F_p}{7000}\right).$$

Bouchnita, A., Terekhov, K., Nony, P., Vassilevski, Y., & Volpert, V.: **A mathematical model to quantify the effects of platelet count, shear rate, and injury size on the initiation of blood coagulation under venous flow conditions.** *PloS one*, 15(7), e0235392, 2020



Large problems

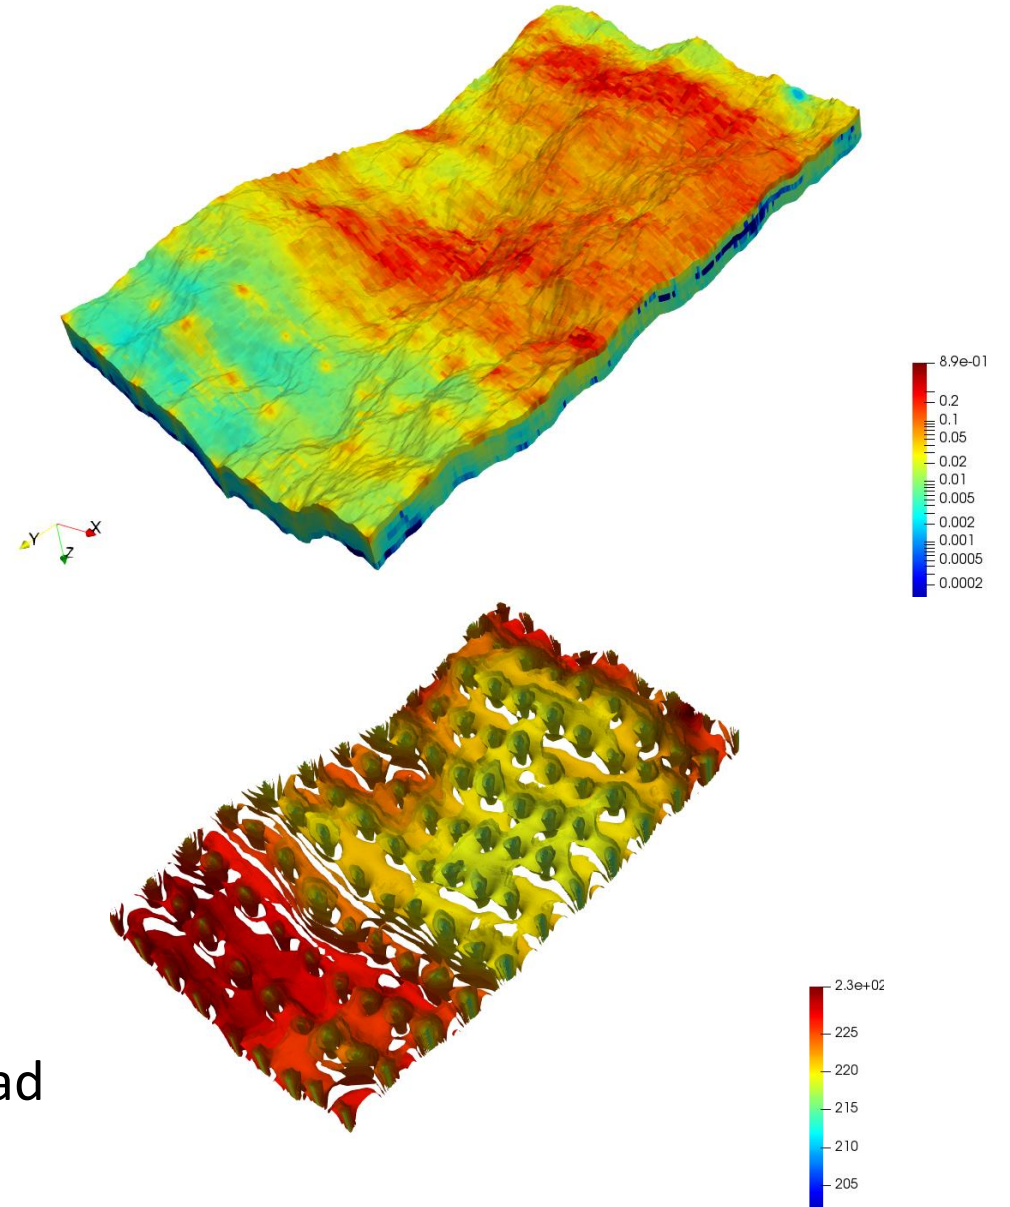
- Suitable for large problem solutions:
 - Black oil problem (previous talk)
 - 3x unknowns per cell
 - **100M** and **200M** cells (320 cores, INM RAS cluster):

Case	T_{mat}	T_{prec}	T_{iter}	T_{sol}	T_{upd}	N_n	N_l
SPE10_100M	14	18.5	55.4	78.6	0.2	402	3.5
SPE10_200M	29.6	34.7	64.1	107.5	0.38	428	3.96

- **Memory per core (significant issue!):**

Case	M_{grid}	M_{mat}	M_{prec}	M_{tot}
SPE10_100M	856.7	165.6	558.4	1943.6
SPE10_200M	1624	346.5	1054.3	4365.6

- Scaled up to **1B of cells** on 9600 Cray cores by Ahmad Abushaika and Longlong Li at HBKU, Qatar.








Large problems

- Suitable for large problem solutions:
 - Poroelasticity problem
 - 4x unknowns per cell
 - **1.2M** cells (INM RAS cluster, Lomonosov):

Machine	N_{proc}	T_{tot}	T_{asm}	T_{prec}	T_{iter}	T_{upd}
INM RAS cluster	100	15079.4	1119.8	7245.2	4463	479.7
	200	8791.2	582.9	3926.2	2800.9	252.4
	400	4637	300.3	1965.6	1374.2	127
Lomonosov supercomputer	700	3536	234.1	1071.1	1112.42	70.5

- **Target:** full-field geomechanical well stability test.





Flagship INMOST Linear Solver

- **Preconditioned BiCGStab(l)** method¹.
- **Preconditioner MPI-parallelization** using **Additive Schwarz Method**.
- **Preconditioner OpenMP-parallelization** using **Bordered Block-Diagonal Form**^{2,3}. **(New!)**
- **Multi-level preconditioner** with **deferred pivoting** based on the second-order **Crout-ILU**^{4,5}.
- **Condition estimation** of the inverse factors determines the **coarse system** and tunes dropping tolerances^{6,7}.
- **Scaling** and **reordering** of the local system before **factorization**^{8,9,10}.





References

- 1) Sleijpen, G.L.G., Diederik R. F.: *BiCGstab (l) for linear equations involving unsymmetric matrices with complex spectrum*. Electronic Transactions on Numerical Analysis 1.11 (1993): 2000. **(Krylov method)**
- 2) Grigori, L., Boman, E. G., Donfack, S., Davis, T. A.: *Hypergraph-based unsymmetric nested dissection ordering for sparse LU factorization*. SIAM Journal on Scientific Computing, 32.6 (2010): 3426-3446. **(Bordered block-diagonal form)**
- 3) Duff, I. S., Scott, J. A.: *Stabilized bordered block diagonal forms for parallel sparse solvers*. Parallel Computing, 31.3-4 (2005): 275-289. **(Bordered block-diagonal form)**
- 4) Li N., Saad Y., Chow E.: *Crout versions of ILU for general sparse matrices*. SIAM Journal on Scientific Computing 25.2 (2003): 716-728. **(Crout-ILU)**
- 5) Kaporin, I.E.: *High quality preconditioning of a general symmetric positive definite matrix based on its UTU+ UTR+ RTU-decomposition*. Numerical linear algebra with applications 5.6 (1998): 483-509. **(Second-order ILU)**
- 6) Bollhöfer, M.: *A robust ILU with pivoting based on monitoring the growth of the inverse factors*. Linear Algebra and its Applications 338.1-3 (2001): 201-218. **(Tuning dropping tolerances)**
- 7) Bollhöfer, M., Saad Y.: *Multilevel preconditioners constructed from inverse-based ILUs*. SIAM Journal on Scientific Computing 27.5 (2006): 1627-1650. **(Computing coarse system)**
- 8) Cuthill, E., McKee J.: *Reducing the bandwidth of sparse symmetric matrices*. Proceedings of the 1969 24th national conference. 1969. **(Reordering for bandwidth reduction)**
- 9) Olschowska, M., Arnold N.: *A new pivoting strategy for Gaussian elimination*. Linear Algebra and its Applications 240 (1996): 131-151. **(Maximizing diagonal product)**
- 10) Kaporin, I.E.: *Scaling, reordering, and diagonal pivoting in ILU preconditionings*. Russian Journal of Numerical Analysis and Mathematical Modelling 22.4 (2007): 341-375. **(Rescaling for condition reduction)**



Additive Schwarz Method

- For **MPI-parallelization**.
- Global matrix is composed of local **blocks**.
- Extend blocks to localize the **connection**.
- **Restricted** version.




Distributed system

- - Local partition outlier
- - Remote partition outlier
- - Local partition
- - Remote partitions
- ▨ - Extended rows



Bordered Block-Diagonal Form

- For **OpenMP-parallelization**.
- The problem is to find **minimum separator** for the **K**-way partitioning. (**NP-hard**)
- All algorithms are approximate:
 - Greedy dissection (**this work**).
 - Usual graph packages: Metis, Patch, Mondriaan. (**todo**)
- Work with each block in **parallel**.
- **Differed pivoting** for the **separator**.





Bordered Block-Diagonal Form

- Greedy algorithm:
 - G – graph of A , $H = G^T G$ – graph product.
 - $|G_i|$ predicts **block size** growth.
 - $|H_i|$ predicts **separator** growth.
 - Using **priority queue** pick a node with minimal growth in both **block** and **separator**.
 - Remove the node from G and H graphs and update **priority queue**.
 - Sequential** ☹️

Algorithm 1 Greedy dissection.

```


1: function GREEDYDISSECTION( $A, parts$ )
2:   Let  $G = (V, E)$ , where                                ▷ Graph induced by matrix
3:    $V = \{1, \dots, n\}$ ,                                   ▷ Vertex set
4:    $G_i = \{\delta_{ij} | a_{ij} \neq 0\}$ .                       ▷ Edge set
5:   Let  $H = G^T G$ .                                         ▷ Graph product
6:   Let  $\omega_B = \{\omega_B^i = \|G_i\| \mid \forall i \in V\}$ .           ▷ Weight for block growth
7:   Let  $\omega_S = \{\omega_S^i = \|S_i\| \mid \forall i \in V\}$ .           ▷ Weight for separator growth
8:   Let  $B_i = \{\emptyset\}, \forall i \in \{1, \dots, parts\}$ .    ▷ Blocks
9:   Let  $S = \emptyset$ .                                       ▷ Separator
10:  Let  $R = C = V$ .                                         ▷ Candidate row and column sets
11:  Let  $P_i = Q_i = 0, \forall i \in \{1, \dots, n\}$ .             ▷ Rows and columns reordering matrices
12:   $r = c = q = 1$ .                                         ▷ Row and column index and current block
13:  while  $R \neq \emptyset$  do
14:     $i = \arg \min_{j \in R} (\omega_B^j + \omega_H^j)$ ,                ▷ Minimal growth of both the block and separator
15:     $R = R \setminus \{i\}$ ,                                  ▷ Remove from the candidate set
16:     $P_i = r$ ,
17:     $r = r + 1$ ,                                             ▷ Enumerate row
18:     $B_q = B_q \cup \{i\}$ ,                                  ▷ Augment current block
19:    if  $i \in S$  then  $S = S \setminus \{i\}$ .
20:    for  $m \in H_i \cap R$  do
21:       $\omega_m^S = \omega_m^S - 1$ .                                ▷ Update weights for separator growth
22:    end for
23:    for  $j \in G_i \cap C$  do
24:       $C = C \setminus \{j\}$ ,
25:       $Q_j = c$ ,
26:       $c = c + 1$ .                                         ▷ Enumerate column
27:      for  $k \in G_j^T$  do
28:         $\omega_k^B = \omega_k^B - 1$ .                                ▷ Update weights for block growth
29:        if  $k \notin S$  then
30:           $S = S \cup \{k\}$ ,
31:           $\omega_k^S = \omega_k^S - 1$ .                                ▷ Add row to separator
32:          for  $m \in H_k \cap R$  do
33:             $\omega_m^S = \omega_m^S - 1$ .                                ▷ Update weights for separator growth
34:          end for
35:        end if
36:      end for
37:    end for
38:    if  $\|B_q\| \geq \|R\|/parts$  then
39:       $R = R \setminus S$ ,                                     ▷ Remove current separator from candidate set
40:       $q = q + 1$ .                                         ▷ Consider next block
41:    end if
42:  end while
43:  for all  $i \in S$  do
44:     $P_i = r$ ,
45:     $r = r + 1$ .
46:  end for
47:   $Q = P$ .
48:  return  $[P, Q, B, S]$ 
49: end function

```




Singly-Bordered Block-Diagonal Form

First level




Schur complement



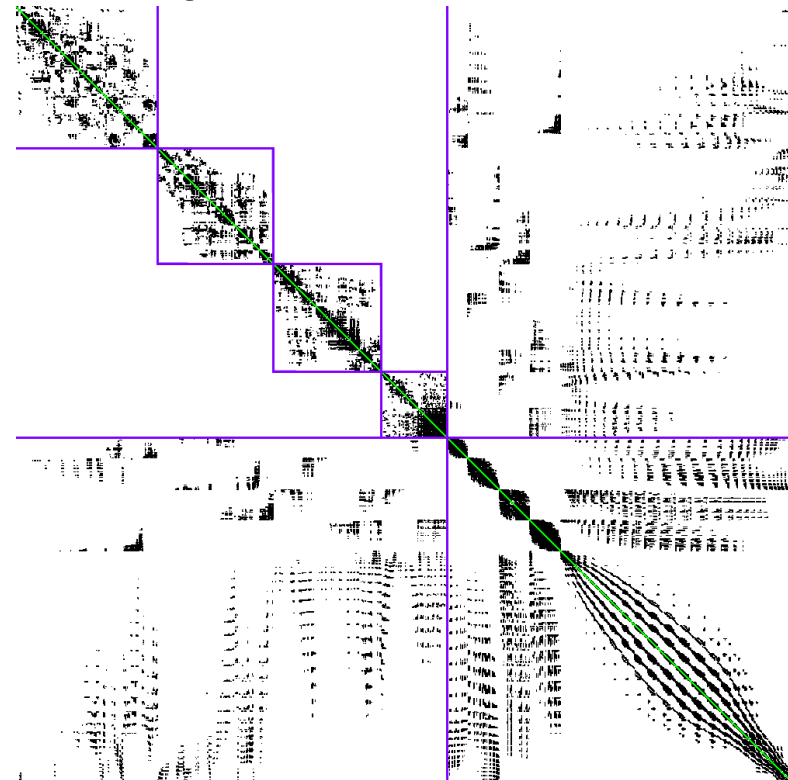


Doubly-Bordered Block-Diagonal Form

First level




Schur complement






SBBD vs DBBD forms

SBBD Schur (**no diagonal**)



DBBD Schur (**has diagonal**)







Ordering and Scaling

- Each successive level is **reordered** and **rescaled**:
 - **Reorder** to maximize diagonal product. (famous MC64)
 - partially **sequential** 😞
 - may apply **in parallel** to independent blocks, but **reduce** Schur quality.
 - **Rescale** into I-dominant matrix.
 - **Reorder** symmetrically with bandwidth or fill-in reduction algorithm in each block.




Preprocessing methods

Initial System



Fill-in reduction (METIS)

Bandwidth reduction



Bandwidth reduction


- Complete **blood coagulation** problem.
- 13 equations per cell.

- **Navier-Stokes** problem only.
- 4 equations per cell.
- Fill-in reduction (METIS) is 4 times slower than bandwidth reduction (RCM) algorithm.
- Quite dense matrix graph, big separator.

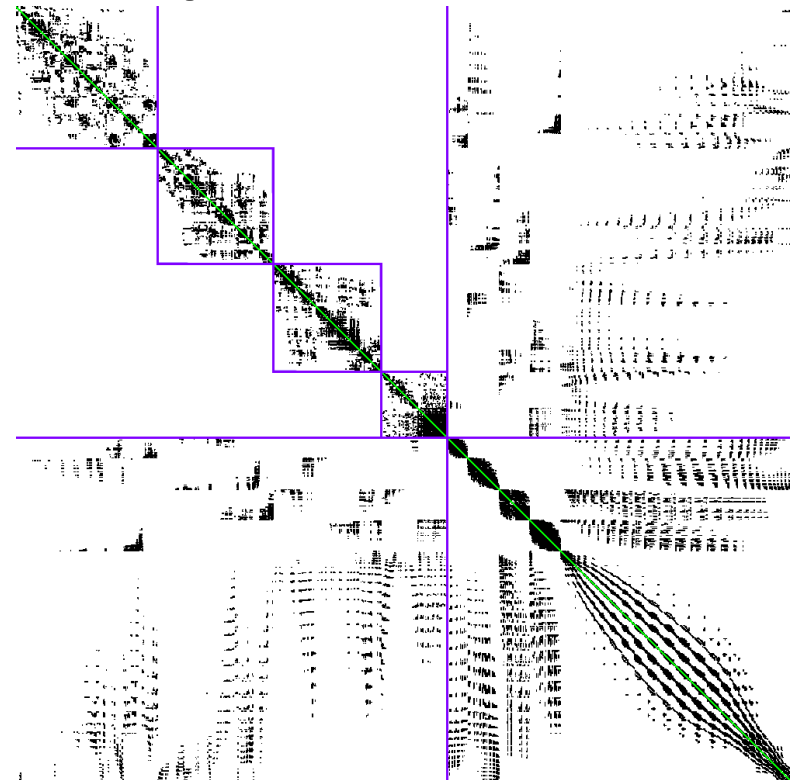


Doubly-Bordered Block-Diagonal Form (initial)

First level




Schur complement






Doubly-Bordered Block-Diagonal Form (RCM)

First level





Schur complement






Second-order Crout Incomplete LU

- Dual-threshold dropping:
 - \mathbf{T}^2 for factorization.
 - \mathbf{T} for iterations.
- Running condition estimation:
 - $\mathbf{K} = \max(|L^{-1}|, |U^{-1}|)$
 - $\mathbf{T}/\mathbf{K} = \text{const}$ tuning.
 - Limit growth of \mathbf{K} .




U-factor elimination




L-factor elimination

Dense row accumulator:



Transposed matrix traversal:






Schur Complement

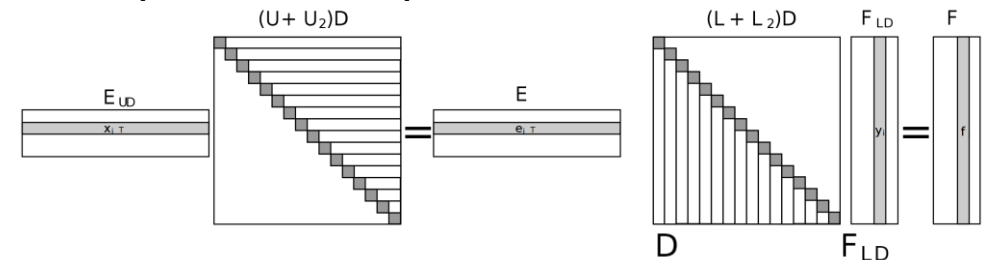
- Part that leads to growth of \mathbf{K} is accumulated in \mathbf{C} :
 - system reordering after factorization.
- Next level system is the Schur complement:
 - $\mathbf{S} = \mathbf{C} - \mathbf{E} (\mathbf{D}\mathbf{U})^{-1} \mathbf{D}(\mathbf{L}\mathbf{D})^{-1} \mathbf{F}$.
 - Requires forward and backward substitution with sparse right hand side.
 - Fill-in control is critical.

■ **Parallel** 😊


Partially factorized DBBD matrix:



Computation of operators:



Schur computation:





Analogy to the Algebraic Multigrid

- **Coarse** system should contain the **largest error** of the **smoother**.
- Condition estimation reveals the **error** in the **smoother** and provides the ***coarse-fine splitting*** of the system.
- Ideal prolongation $P=(-EB^{-1}, I)$ and restriction $R=(-FB^{-1}, I)^T$.
 - (not satisfied by the present method).
- Schur complement corresponds to the **coarse** system.
- **Universal but much more computationally complex.**
 - (definitely not linear computational complexity)



Results

100x100x100 Poisson problem

Black-oil problem with 1M cells

threads	1	1	2	4	8	16
T_{tot}	42.7	47	29.2 (1.6x)	20.2 (2.3x)	13.4 (3.5x)	10.8 (4.3x)
iters	17	23	27	25	27	29
T_{iter}	6	7.8	4.77 (1.6x)	3.2 (2.4x)	2 (3.9x)	1.64 (4.8x)
levels	1	2	5	6	7	7
pivots	0	40564	63588	85689	109995	142988
T_{prec}	36.2	38.7	23.9 (1.6x)	16.5 (2.3x)	10.8 (3.6x)	8.6 (4.5x)
T_{ord}	0.5 (1.4%)	0.7 (2%)	2.6 (11%)	2.4 (15%)	2.3 (21%)	2.55 (29%)
T_{mpt}	0.3 (0.8%)	0.4 (1%)	0.4 (1.6%)	0.4 (2.4%)	0.41 (3.7%)	0.44 (5%)
T_{snd}	- (-%)	- (-%)	1.9 (8%)	1.9 (11%)	1.8 (16.6%)	2.04 (24%)
T_{rcm}	0.2 (0.6%)	0.4 (1%)	0.3 (1.3%)	0.17 (1%)	0.1 (0.9%)	0.06 (0.8%)
T_{sc}	0.9 (2.5%)	1 (2.7%)	0.5 (2.1%)	0.33 (2%)	0.2 (1.9%)	0.17 (1.9%)
T_{fact}	34.2 (95%)	12.7 (23%)	6.3 (26.4%)	4 (24%)	2.48 (23%)	1.32 (15%)
T_{schur}	- (-%)	23.7 (71%)	13.9 (58.2%)	9.2 (55%)	5.35 (49%)	4.04 (47%)

threads	1	1	2	4	8	16
T_{tot}	46.3	56.6	44.7 (1.3x)	37.2 (1.5x)	30.5 (1.9x)	29.4 (1.9x)
iters	19	30	30	32	29	35
T_{iter}	11.3	25	14.2 (1.8x)	10.7 (2.3x)	6.3 (4x)	6.1 (4.1x)
levels	1	5	6	8	9	10
pivots	0	167799	201296	278060	355152	458293
T_{prec}	33.7	30.2	29.1 (1x)	25.2 (1.2x)	22.9 (1.3x)	22 (1.4x)
T_{ord}	5.4 (16%)	5.5 (18%)	13.1 (45%)	13.3 (53%)	12.3 (54%)	11.5 (52%)
T_{mpt}	2.9 (9%)	2.8 (9%)	2.4 (8%)	2.4 (9%)	2.3 (10%)	2.4 (11%)
T_{snd}	- (-%)	- (-%)	7.9 (27%)	8 (32%)	7.8 (34%)	7.8 (36%)
T_{rcm}	2.5 (7%)	2.7 (9%)	2.8 (10%)	2.9 (11%)	2.3 (10%)	1.2 (5.5%)
T_{sc}	4.4 (13%)	4.6 (15%)	2.6 (9%)	1.7 (7%)	1 (4%)	0.8 (4%)
T_{fact}	20.3 (60%)	9.4 (31%)	5.2 (18%)	3 (12%)	3.1 (13%)	2.4 (11%)
T_{schur}	- (-%)	7.1 (23%)	5.2 (18%)	4.8 (19%)	4.3 (19%)	5.2 (24%)

Single-level algorithm



Conclusions and future directions

- Solved memory but not efficiency issues:
 - Sequential algorithm for DBBD form and Schur growth due to increasing deferred pivoting.
- Future directions:
 - Faster algorithm for **DBBD** form.
 - Parallel maximal product transversal.
 - **Block** elimination, **block** pivoting (*reduce deferred pivoting and Schur growth*).
 - **Schur** computation using ideal restriction and prolongation (*more accurate Schur*).
- We are working on a very flexible linear solvers framework S^3M to address multiphysics problems:
 - Konshin, I., Terekhov, K.: *Sparse System Solution Methods for Complex Problems*. In *International Conference on Parallel Computing Technologies*, (2021, September): 53-73. Springer, Cham.

Supported by RSF grant №21-71-20024

Thank you for your attention

Contacts

- KIRILL.TEREHOV@GMAIL.COM

Links

- WWW.INMOST.ORG
- WWW.INMOST.RU

