Development of a Terrestrial Dynamical Core for E3SM (TDycore)

Nathan Collier
Oak Ridge National Laboratory
https://github.com/TDycores-Project/TDycore

June 2019

Tale of Two Talks

Common Theme: Considerations in choosing a discretization method

I. The Effect of a Higher Continuous Basis on Solver Performance

Victor Calo (Curtin), David Pardo (Ikerbasque), Lisandro Dalcin (KAUST), Maciej Paszynski (AGH)

II. Selection of a Numerical Method for a Terrestrial Dynamical Core

Jed Brown (Colorado), Gautam Bisht (PNNL), Matthew Knepley (Buffalo), Jennifer Fredrick (SNL), Glenn Hammond (SNL), Satish Karra (LANL)

Higher Continuous Basis?

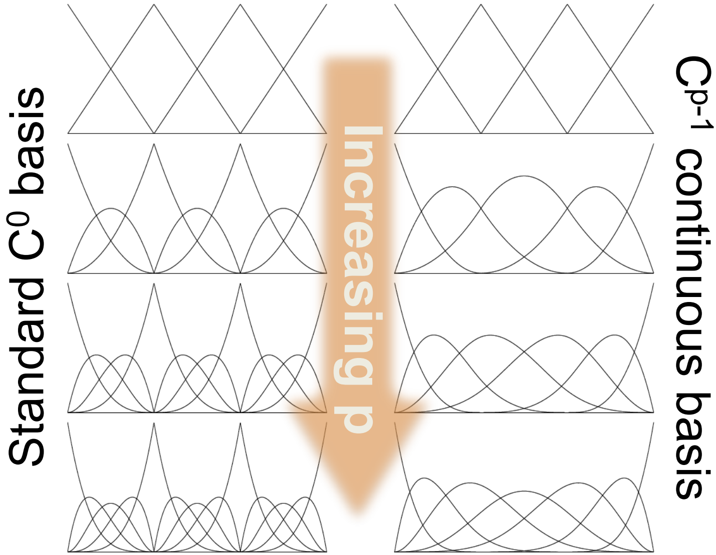

Poisson problem on unit cube

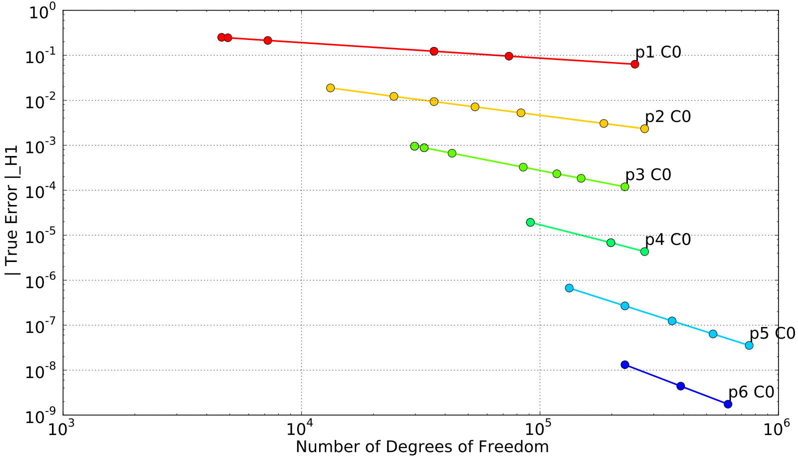

Poisson problem on unit cube

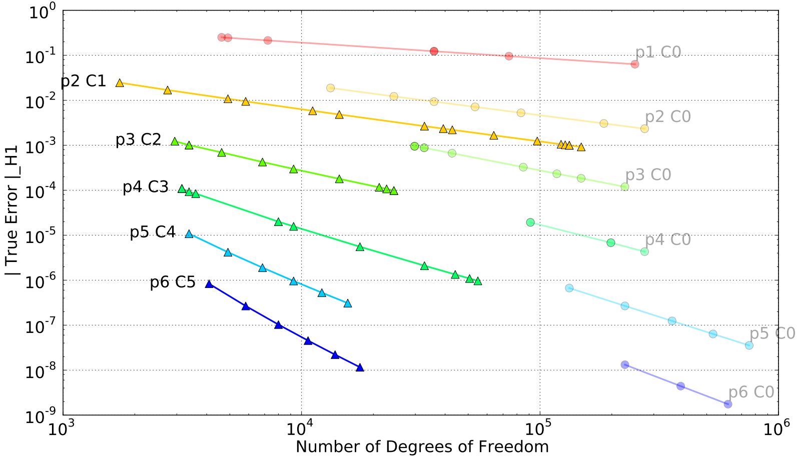

Are higher continuous spaces an efficient way to $p-$refine?

What effect does continuity have on the solver performance?

Are higher continuous spaces an efficient way to $p-$refine?

What effect does continuity have on the solver performance?

Spoiler Alert!

	C^0	C^{p-1}	C^{p-1}/C^0
Multifrontal direct solver	$\mathcal{O}(N^2 + Np^6)$	$\mathcal{O}(N^2 p^3)$	$\mathcal{O}(p^3)$
Iterative solvers*	$\mathcal{O}(Np^4)$	$\mathcal{O}(Np^6)$	$\mathcal{O}(p^2)$

*Estimates for Matrix-Vector products

Multi-frontal direct solver

Based on the concepts of the Schur complement and nested dissection.

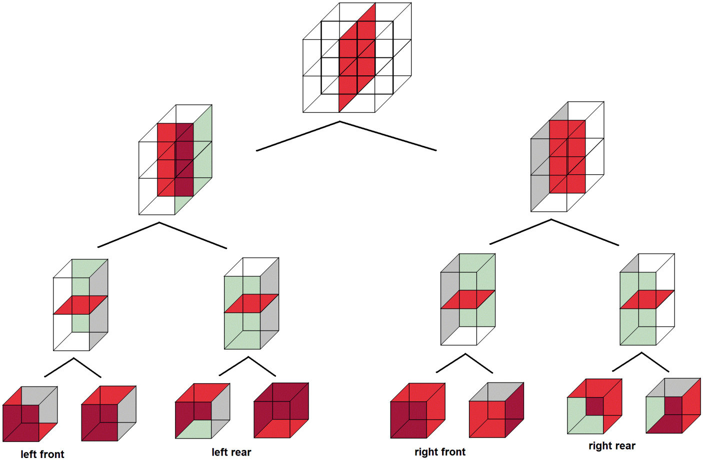

Key concept: size s of the separator

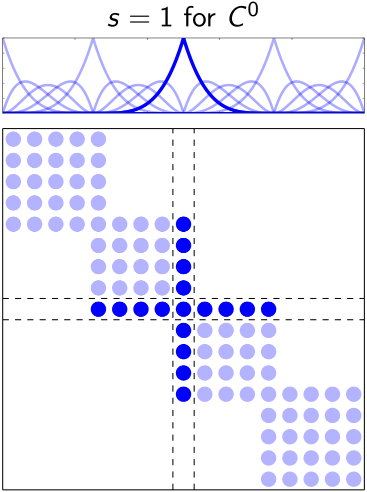

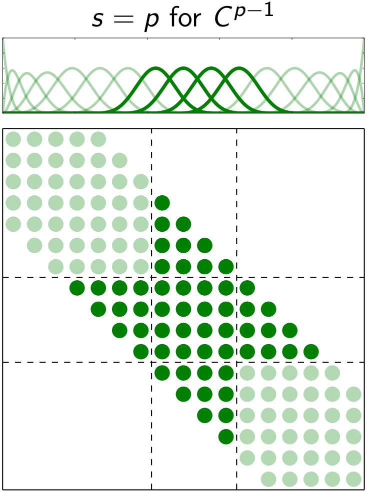

Estimates and Results ($d = 3$, $N = 30k$)

	C^0	C^{p-1}
Time	$\mathcal{O}(N^2 + Np^6)$	$\mathcal{O}(N^2 p^3)$
Memory	$\mathcal{O}(N^{4/3} + Np^3)$	$\mathcal{O}(N^{4/3} p^2)$

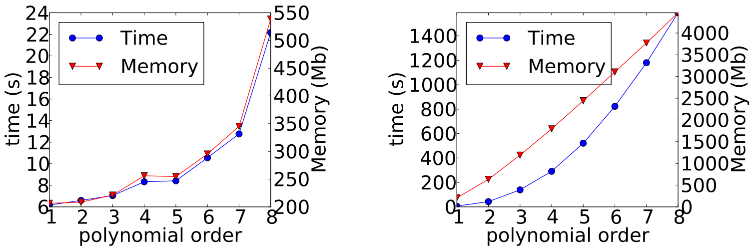

Solution time for C^0 vs C^{p-1} $(d = 3, N = 30k)$

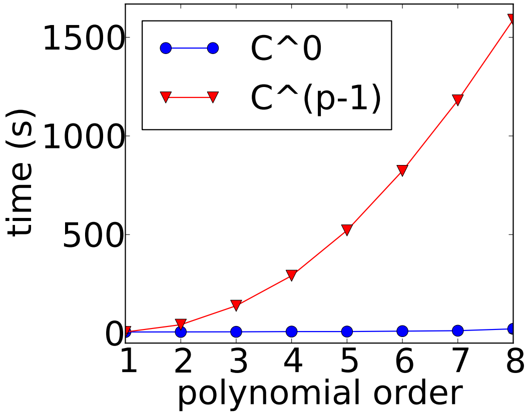

Iterative solvers

Much more complex to assess costs:

$$P\,(Ax - b) = 0$$

Need a model for:

- Matrix-vector multiplication
- Preconditioner (P) setup and application
- Convergence

Sample Linear Systems

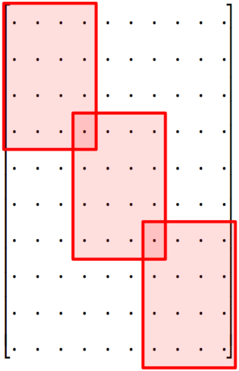

C^0 space

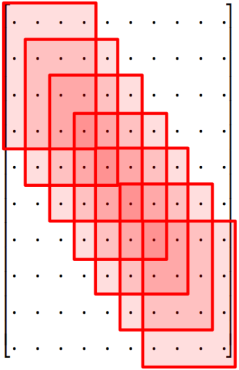

C^{p-1} space

Matrix-vector multiplication - C^0

The cost of a sparse matrix-vector multiply is proportional to the number of nonzero entries in the matrix.

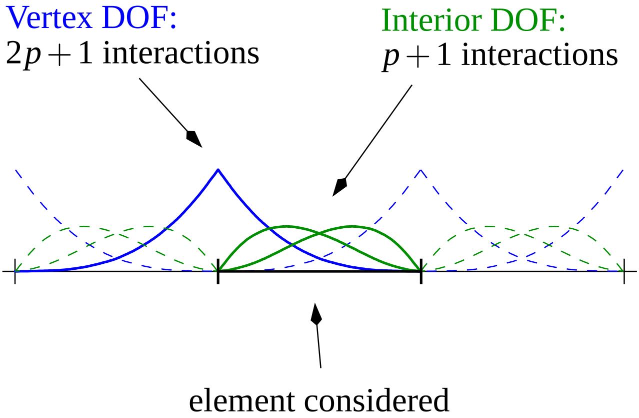

Matrix-vector multiplication - C^0

Dimension	Entity	Number of Entities	DOFs per Entity	Number of interactions
1D	vertex	1	1	$(2p+1)$
1D	interior	1	$(p-1)$	$(p+1)$
2D	vertex	1	1	$(2p+1)^2$
2D	edge	2	$(p-1)$	$(2p+1)(p+1)$
2D	interior	1	$(p-1)^2$	$(p+1)^2$
3D	vertex	1	1	$(2p+1)^3$
3D	edge	3	$(p-1)$	$(2p+1)^2(p+1)$
3D	face	3	$(p-1)^2$	$(2p+1)(p+1)^2$
3D	interior	1	$(p-1)^3$	$(p+1)^3$

Matrix-vector multiplication - C^0

$$
\begin{aligned}
\mathrm{nnz}^{C^0} &= \underbrace{(p-1)^3}_{\text{interior DOF}} \cdot (p+1)^3 \\
&+ \underbrace{3(p-1)^2}_{\text{face DOF}} \cdot (2p+1)(p+1)^2 \\
&+ \underbrace{3(p-1)}_{\text{edge DOF}} \cdot (2p+1)^2(p+1) \\
&+ \underbrace{1}_{\text{vertex DOF}} \cdot (2p+1)^3 \\
&= p^6 + 6p^5 + 12p^4 + 8p^3 \\
&= p^3(p+2)^3 = \mathcal{O}(p^6)
\end{aligned}
$$

Matrix-vector multiplication - C^{p-1}

The B-spline C^{p-1} basis is very regular, each DOF interacts with $2p+1$ others in 1D.

$$\mathrm{nnz}^{C^{p-1}} = p^3(2p+1)^3 = 8p^6 + 12p^5 + 6p^4 + p^3 = \mathcal{O}(8p^6)$$

Matrix-vector multiplication

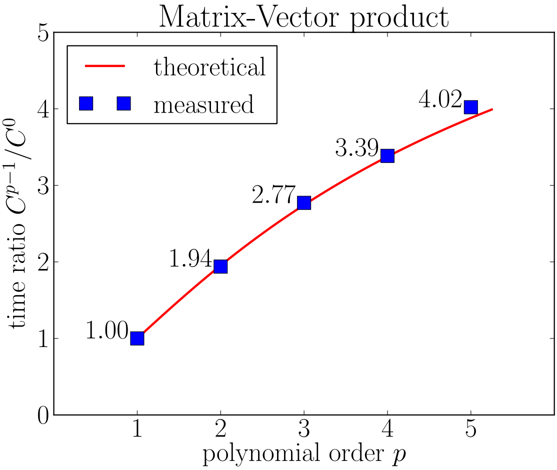

Matrix-vector multiplication

However, for C^0 spaces, we can use static condensation as in the multifrontal direct solver.

Entity	Number of Entities	DOFs per Entity	Number of interactions	Statically condensed
vertex	1	1	$(2p+1)^3$	$-8(p-1)^3$
edge	3	$(p-1)$	$(2p+1)^2(p+1)$	$-4(p-1)^3$
face	3	$(p-1)^2$	$(2p+1)(p+1)^2$	$-2(p-1)^3$

$$33p^4 - 12p^3 + 9p^2 - 6p + 3 = \mathcal{O}(33p^4)$$

Matrix-vector multiplication

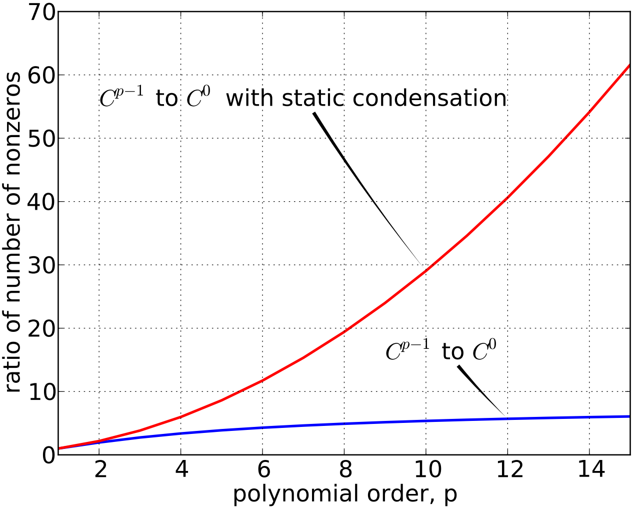

3D Poisson + CG + ILU

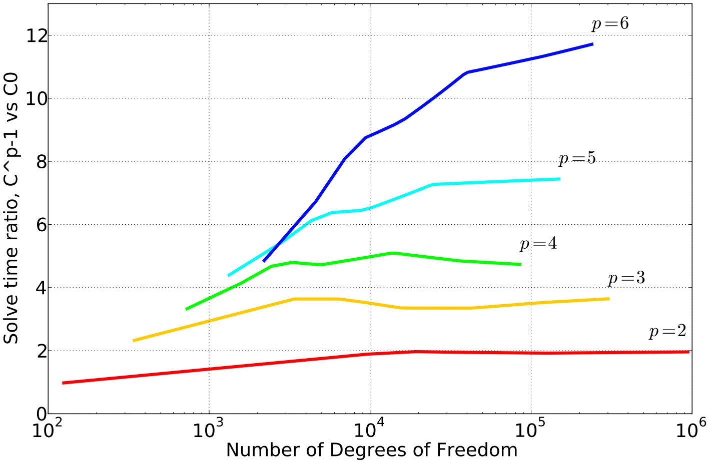

3D Poisson + CG + ILU + static condensation

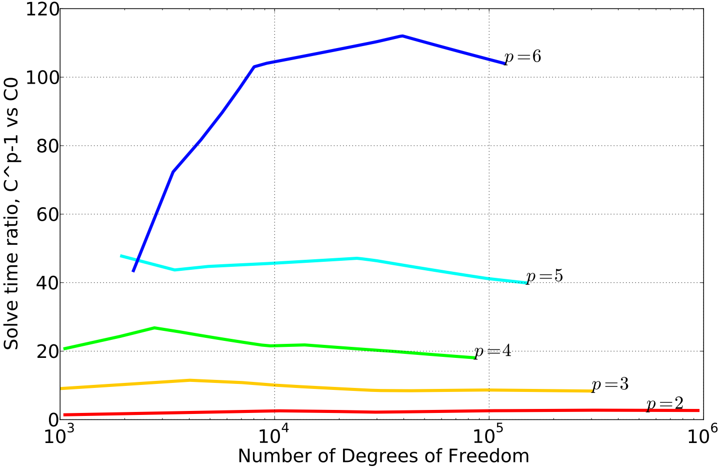

Related Publications

- N Collier, D Pardo, L Dalcin, M Paszynski, VM Calo, The cost of continuity: A study of the performance of isogeometric finite elements using direct solvers, Computer Methods in Applied Mechanics and Engineering 213, 353-361, 2012. 10.1016/j.cma.2011.11.002
- N Collier, L Dalcin, D Pardo, VM Calo, The cost of continuity: performance of iterative solvers on isogeometric finite elements, SIAM Journal on Scientific Computing 35 (2), A767-A784, 2013. 10.1137/120881038
- N Collier, L Dalcin, VM Calo, On the computational efficiency of isogeometric methods for smooth elliptic problems using direct solvers, International Journal for Numerical Methods in Engineering 100 (8), 620-632. 10.1002/nme.4769

Tale of Two Talks

Common Theme: Considerations in choosing a discretization method

I. The Effect of a Higher Continuous Basis on Solver Performance

Victor Calo (Curtin), David Pardo (Ikerbasque), Lisandro Dalcin (KAUST), Maciej Paszynski (AGH)

II. Selection of a Numerical Method for a Terrestrial Dynamical Core

Jed Brown (Colorado), Gautam Bisht (PNNL), Matthew Knepley (Buffalo), Jennifer Fredrick (SNL), Glenn Hammond (SNL), Satish Karra (LANL)

Energy Exascale Earth System Model (E3SM)

- The terrestrial water cycle is a key component of the Earth system model
- While conceptually key processes transport water laterally, the representation is 1D in current models
- Requirements: accurate velocities on distorted grids with uncertain and rough coefficients at global scale
- Naturally think of mixed finite elements

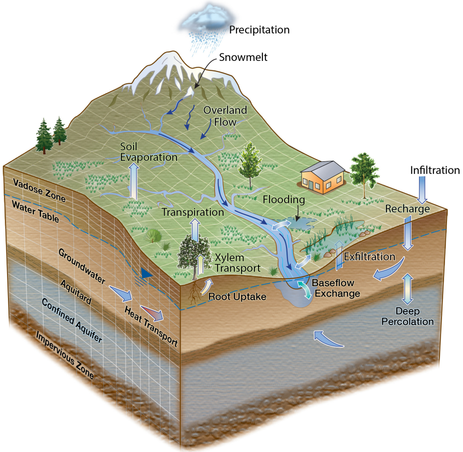

Simplified Problem Statement

Strong form Find $\mathbf{u}$ and p such that,

$$\begin{aligned}
\mathbf{u} &= -K\nabla p && \text{in } \Omega \\
\nabla \cdot \mathbf{u} &= f && \text{in } \Omega \\
p &= g && \text{on } \Gamma_D \\
\mathbf{u} \cdot \mathbf{n} &= 0 && \text{on } \Gamma_N
\end{aligned}$$

Simplified Problem Statement

Strong form Find $\mathbf{u}$ and p such that,

$$\begin{aligned} \mathbf{u} &= -K\nabla p && \text{in } \Omega \\ \nabla \cdot \mathbf{u} &= f && \text{in } \Omega \\ p &= g && \text{on } \Gamma_D \\ \mathbf{u} \cdot \mathbf{n} &= 0 && \text{on } \Gamma_N \end{aligned}$$

Candidate approaches:

- Mixed finite elements (BDM) + FieldSplit/BDDC/hybridization
- Wheeler-Yotov (WY) + AMG
- Arnold-Boffi-Falk (ABF) + FieldSplit/BDDC/hybridization
- Multipoint flux approximation (MFPA) + AMG

Simplified Problem Statement

Strong form Find $\mathbf{u}$ and p such that,

$$\begin{aligned} \mathbf{u} &= -K\nabla p && \text{in } \Omega \\ \nabla \cdot \mathbf{u} &= f && \text{in } \Omega \\ p &= g && \text{on } \Gamma_D \\ \mathbf{u} \cdot \mathbf{n} &= 0 && \text{on } \Gamma_N \end{aligned}$$

Weak form Find $\mathbf{u} \in \mathbf{V}$ and $p \in W$ such that,

$$\begin{aligned} \left(K^{-1}\mathbf{u}, \mathbf{v}\right) &= (p, \nabla \cdot \mathbf{v}) - \langle g, \mathbf{v} \cdot \mathbf{n} \rangle_{\Gamma_D}, && \mathbf{v} \in \mathbf{V} \\ (\nabla \cdot \mathbf{u}, w) &= (f, w), && w \in W \end{aligned}$$

where $\mathbf{V} = \{\mathbf{v} \in H^{\mathrm{div}}(\Omega) : \mathbf{v} \cdot \mathbf{n} = 0 \text{ on } \Gamma_N\}$, $W = L^2(\Omega)$

Problem statement

Strong form Find $\mathbf{u}$ and p such that,

$$\begin{aligned} \mathbf{u} &= -K\nabla p && \text{in } \Omega \\ \nabla\cdot\mathbf{u} &= f && \text{in } \Omega \\ p &= g && \text{on } \Gamma_D \\ \mathbf{u}\cdot\mathbf{n} &= 0 && \text{on } \Gamma_N \end{aligned}$$

Weak form Find $\mathbf{u}\in\mathbf{V}$ and $p\in W$ such that,

$$\begin{aligned} \left(K^{-1}\mathbf{u},\mathbf{v}\right) &= (p,\nabla\cdot\mathbf{v}) - \langle g,\mathbf{v}\cdot\mathbf{n}\rangle_{\Gamma_D}, && \mathbf{v}\in\mathbf{V} \\ (\nabla\cdot\mathbf{u},w) &= (f,w), && w\in W \end{aligned}$$

where $\mathbf{V}=\{\mathbf{v}\in H^{\mathrm{div}}(\Omega) : \mathbf{v}\cdot\mathbf{n}=0 \text{ on } \Gamma_N\}$, $W=L^2(\Omega)$

Wheeler & Yotov 2006

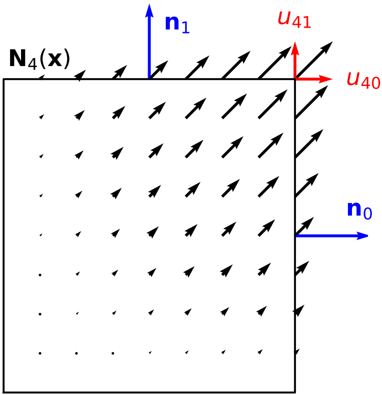

Ingredients:

- Brezzi–Douglas–Marini (BDM_1) velocity space
- Basis interpolatory at corners
  $\mathbf{N}_4(\mathbf{x}_4) \cdot \mathbf{n}_0 = u_{40}$
  $\mathbf{N}_4(\mathbf{x}_4) \cdot \mathbf{n}_1 = u_{41}$
- Vertex-based quadrature (under-integrated)
- Constant pressure space

Wheeler & Yotov 2006

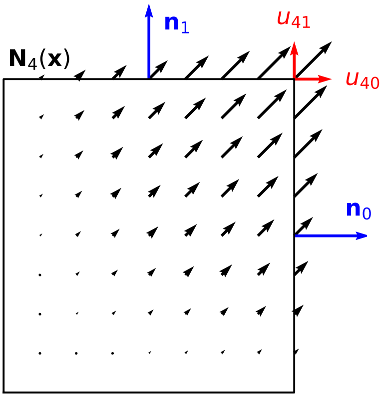

Ingredients:

- Brezzi–Douglas–Marini (BDM_1) velocity space
- Basis interpolatory at corners
  $\mathbf{N}_4(\mathbf{x}_4) \cdot \mathbf{n}_0 = u_{40}$
  $\mathbf{N}_4(\mathbf{x}_4) \cdot \mathbf{n}_1 = u_{41}$
- Vertex-based quadrature (under-integrated)
- Constant pressure space

This means that velocity DOFs only couple to each other at vertices.

Wheeler & Yotov Assembly

for vertex v in mesh **do**

- setup vertex local problem
  $$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} U \\ P \end{bmatrix} = \begin{bmatrix} G \\ F \end{bmatrix}$$
- **for** element e connected to v **do**
  - $A \leftarrow \left(K^{-1}\mathbf{u}_v, \mathbf{v}_v\right)_{\Omega_e}$
  - $B^T \leftarrow -\left(p_e, \nabla \cdot \mathbf{v}_v\right)_{\Omega_e}$
  - $G \leftarrow -\langle g, \mathbf{v}_v \cdot \mathbf{n}\rangle_{\Gamma_{D,e}}$
  - $F \leftarrow \left(f_e, w_e\right)_{\Omega_e}$
- **end for**
- Assemble Schur complement
  $(BA^{-1}B^T)P = F - BA^{-1}G$

end for

Wheeler & Yotov Assembly

for vertex v in mesh **do**

  setup vertex local problem $\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} U \\ P \end{bmatrix} = \begin{bmatrix} G \\ F \end{bmatrix}$

  for element e connected to v **do**

    $A \leftarrow \left(K^{-1}\mathbf{u}_v, \mathbf{v}_v\right)_{\Omega_e}$

    $B^T \leftarrow -\left(p_e, \nabla \cdot \mathbf{v}_v\right)_{\Omega_e}$

    $G \leftarrow -\langle g, \mathbf{v}_v \cdot \mathbf{n}\rangle_{\Gamma_{D,e}}$

    $F \leftarrow \left(f_e, w_e\right)_{\Omega_e}$

  end for

  Assemble Schur complement $(BA^{-1}B^T)P = F - BA^{-1}G$

end for

Global cell-centered pressure system which is SPD

Sample Wheeler-Yotov Stencils

$$K = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

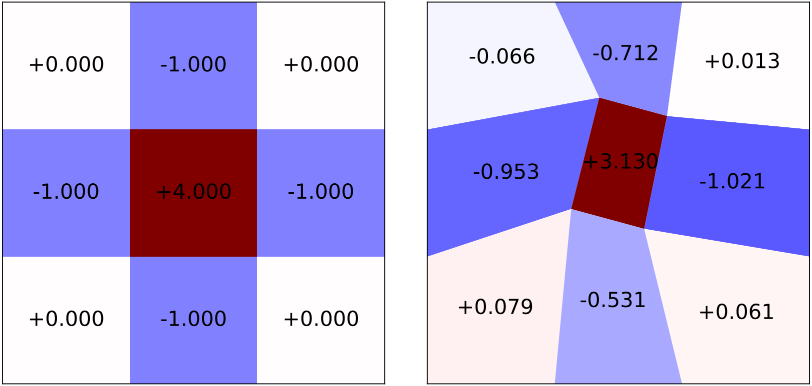

Sample Wheeler-Yotov Stencils

$$K = \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$$

+0.000	-1.000	+0.000
-3.000	+8.000	-3.000
+0.000	-1.000	+0.000

$$R_{10}\ K\ R_{10}^T$$

+0.133	-0.985	-0.209
-2.865	+7.850	-2.865
-0.209	-0.985	+0.133

$$R_{45}\ K\ R_{45}^T$$

+0.250	-1.500	-0.750
-1.500	+7.000	-1.500
-0.750	-1.500	+0.250

Sample Wheeler-Yotov Stencils

$$K = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

+0.000	-1.000	+0.000
-1.000	+4.000	-1.000
+0.000	-1.000	+0.000

$$K \cdot (10^{-3} \text{ if } x > 2/3)$$

+0.000	-1.000	+0.000
-1.000	+3.002	-0.002
+0.000	-1.000	+0.000

SPE10 Test Problem

We use the permeabilities from the SPE10 problem:

- $60 \times 220 \times 85 = 1{,}122{,}000$ cells
- Diagonal permeability $K_{xx} = K_{yy} \neq K_{zz}$
- We induce flow by Dirichlet conditions
- Solve on original permeabilities and also rotate around two axes

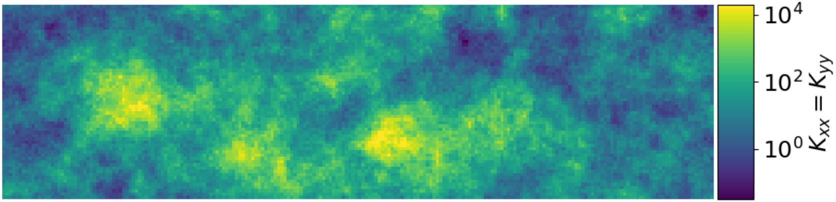

Sample slice of the permeability field

Solver Options

WY Options

```
-ksp_type cg
-pc_type hypre
```

BDM Options

```
-ksp_type gmres
-pc_type fieldsplit
-pc_fieldsplit_type schur
-pc_fieldsplit_schur_fact_type full
-pc_fieldsplit_schur_precondition selfp
-fieldsplit_0_ksp_type cg
-fieldsplit_0_pc_type jacobi
-fieldsplit_1_ksp_type cg
-fieldsplit_1_pc_type hypre
```

Solver Performance

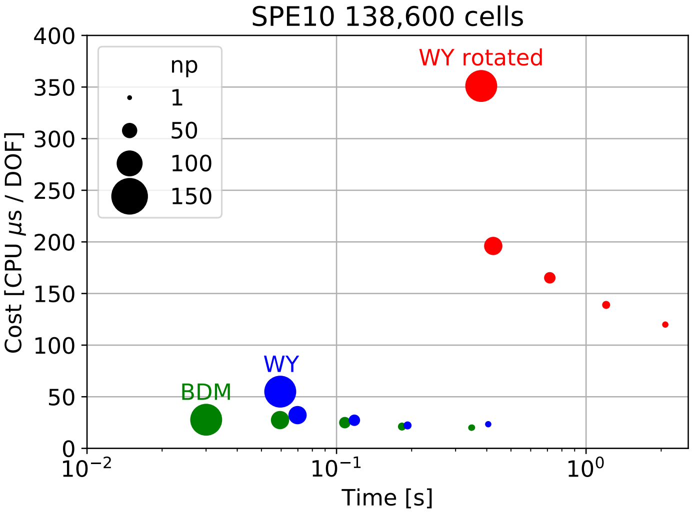

Solver Performance

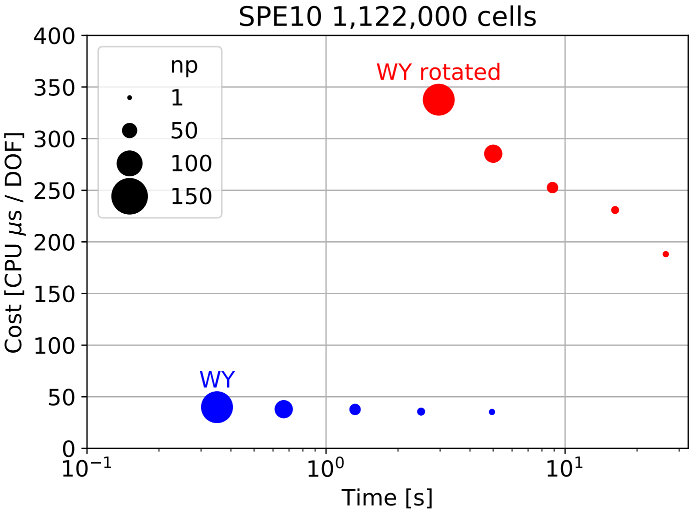

Assembly Performance (not optimized)

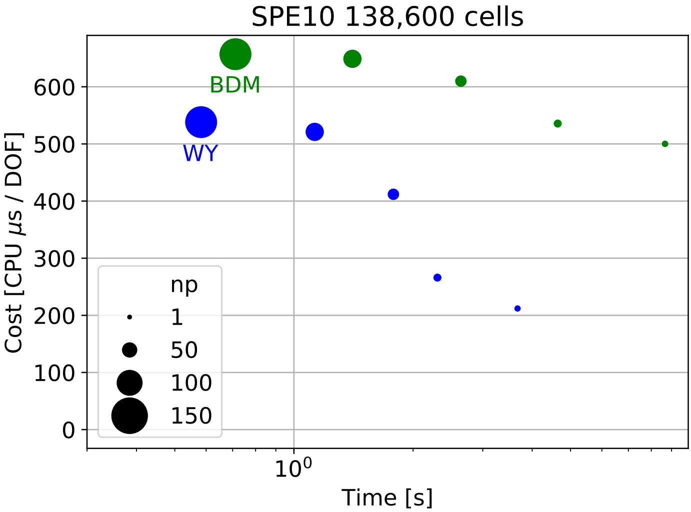

Concluding Remarks

TDycore:

- From limited results, WY approach is at least competitive although it appears to hit the strong scaling limit before BDM/fieldsplit
- BDM appears to use more memory than WY ($\approx$ 10 times)
- WY assembly is competitive although BDM lends itself to easier vectorization
- Experimentation is key: `-tdy_method {wy|bdm|...}`

Concluding Remarks

TDycore:

- From limited results, WY approach is at least competitive although it appears to hit the strong scaling limit before BDM/fieldsplit
- BDM appears to use more memory than WY ($\approx$ 10 times)
- WY assembly is competitive although BDM lends itself to easier vectorization
- Experimentation is key: `-tdy_method {wy|bdm|...}`

Talk/Meeting:

- All of the presented work uses PETSc (PetIGA + DMPlex/Section)
- Using DMPlex/Section opens doors for solver approaches
- Most of my exposure to solvers comes from using PETSc
- Originally exposed to PETSc $\approx$ 11 years ago at DOE ACTS workshops

Important Links

- PetIGA: https://bitbucket.org/dalcinl/petiga
- TDycore: https://github.com/TDycores-Project/TDycore