

Interfacing MFiX with PETSc and HYPRE Linear Solver libraries

PI: Gautham Krishnamoorthy (UND)
Lauren Clarke (Grad Student, UND)
KayLee Smith (Grad Student, UND)

Co-PI: Jeremy Thornock (U.Utah)
Surya Yamujala(Grad Student, U.Utah)
April 9th, 2019

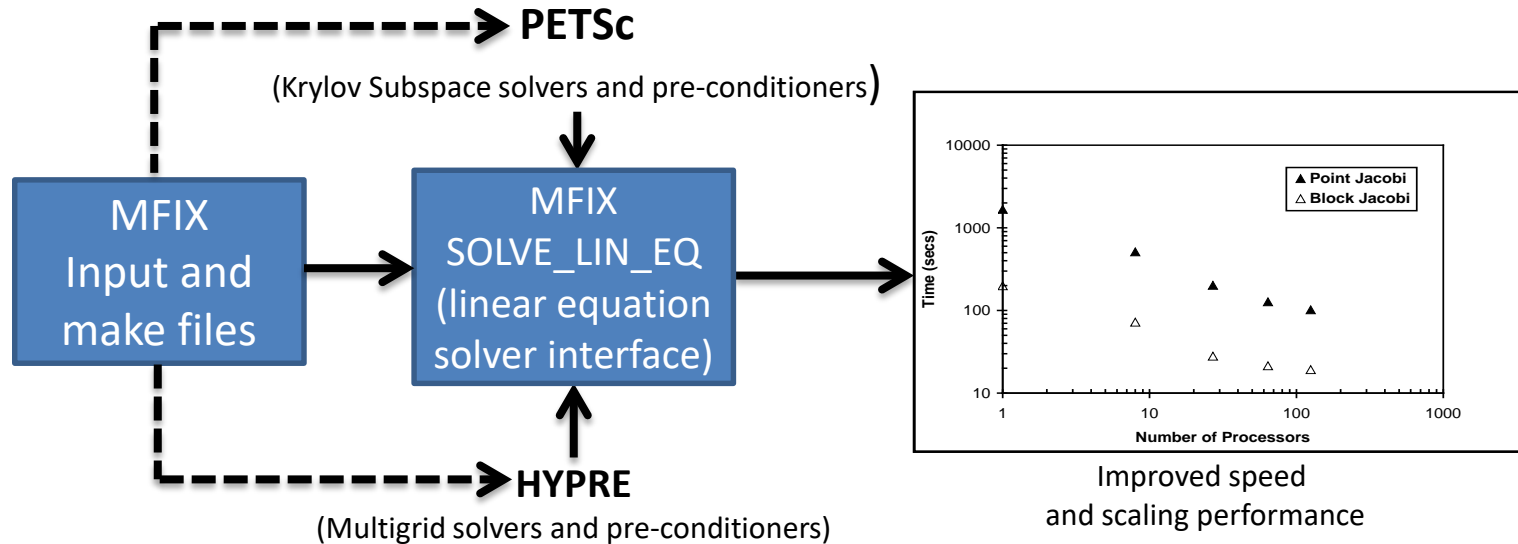


Acknowledgement: This research is being funded by the University Coal Research Program which is administered by DOE-NETL (Award #: DE-FE0026191)







Objective/Vision

- Build a robust, well-abstracted, interface to the PETSc, HYPRE linear solver libraries from MFiX



- Code verification against established MFiX solutions and code to code comparisons



Project Successes

- Lauren Elizabeth Clarke (Combined BS/MS Student, UND)
 - Mickey Leland Energy Fellow '17
 - 2 peer-reviewed papers (1 published, 1 under review)
 - Currently pursuing PhD at MIT
- KayLee Smith (Combined BS/MS Student, UND)
 - Mickey Leland Energy Fellow '18
 - Placed 1st (Engineering) in UND's Graduate Student Expo
- Surya Yamujala (MS Student, UUtah)
 - Quality and Reliability Assurance Engineer at IMFlash
- The MFiX-PETSc, MFiX-HYPRE interfaces have turned out to be faster and more robust in several problems including:
 - Domains involving local cell refinements and abrupt cell jumps
 - Complex geometries and large cell count (multi-tube BFB, multi-spout CLC reactors)
 - Large relaxation-factors
 - Large property variations (laminar sCO₂ flows)
 - Multi-fluid (different particle sizes) systems
- Increased collaboration between UND ↔ UUtah



Our First Step

(Identification of optimum solvers and pre-conditioners)

PETSc relative solve times for solution to the inhomogeneous Helmholtz Equation (3D) (Septadiagonal matrix, uniprocessor)

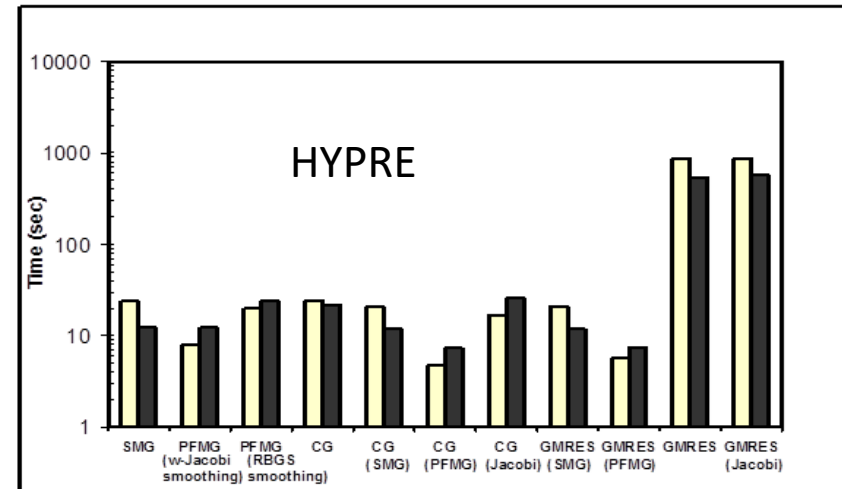
$$\nabla^2 A + k^2 A = -f$$

Stand alone solver timing studies


Degrees of Freedom	CG	GMRES	<u>BiCGSTAB</u>
150K	1.56	11.11	2.16
600K	23.45	700.00	35.56

Best stand alone solver with pre-conditioning options in brackets


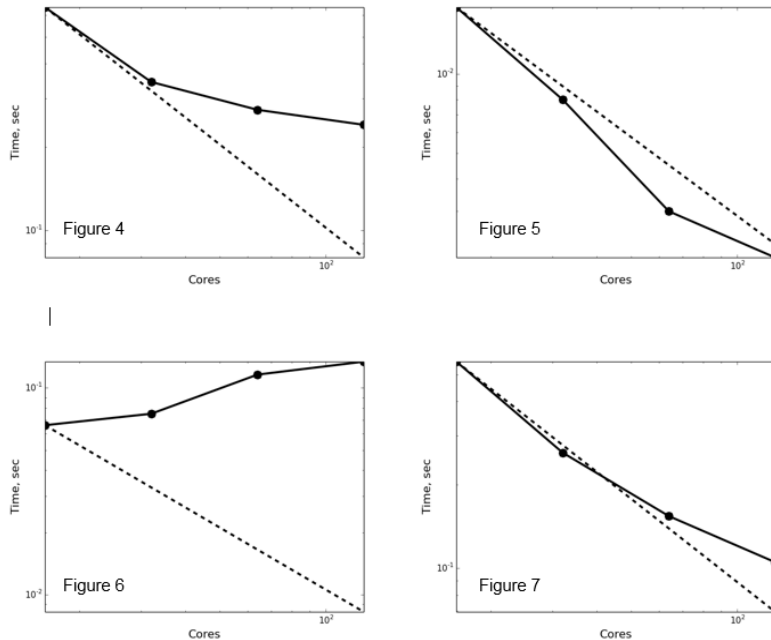
Degrees of Freedom	CG (Point Jacobi)	CG (Block Jacobi)	CG (ILU)	CG (SOR)
150K	1.29	1.06	1.06	1.00
600K	25.24	19.31	18.01	17.87
1.2M	57.64	42.94	41.76	40.00



A Road Block!

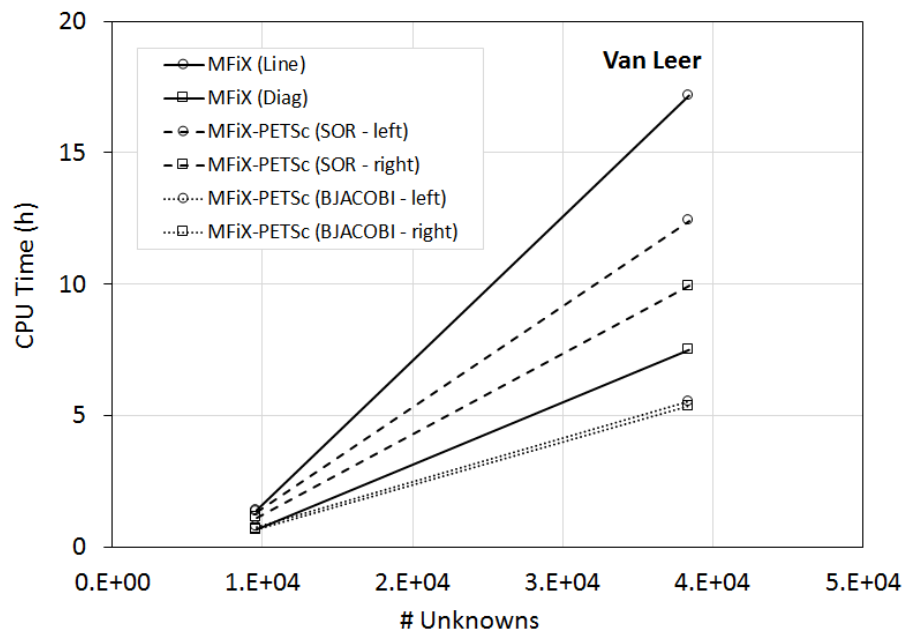
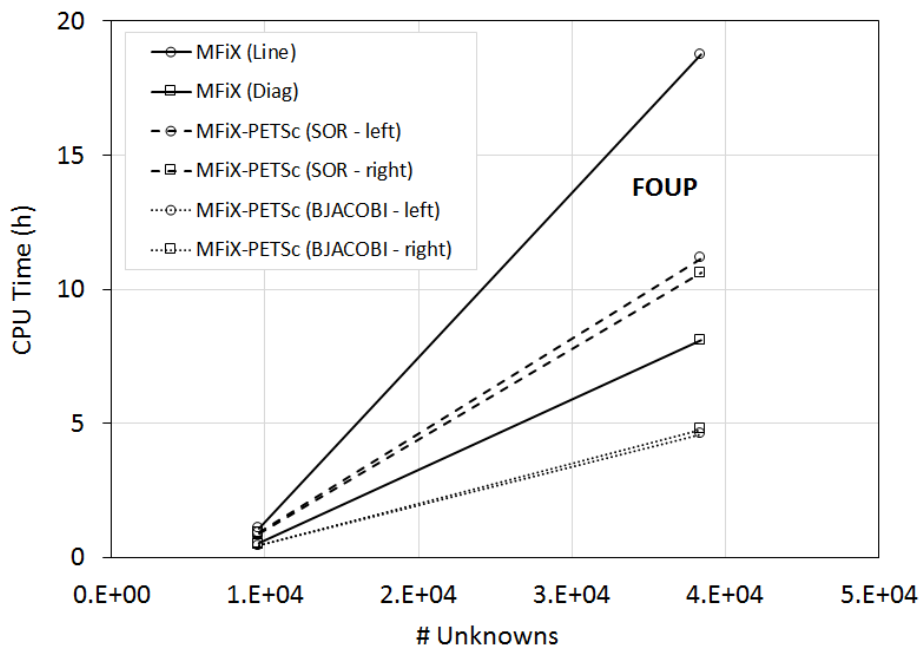
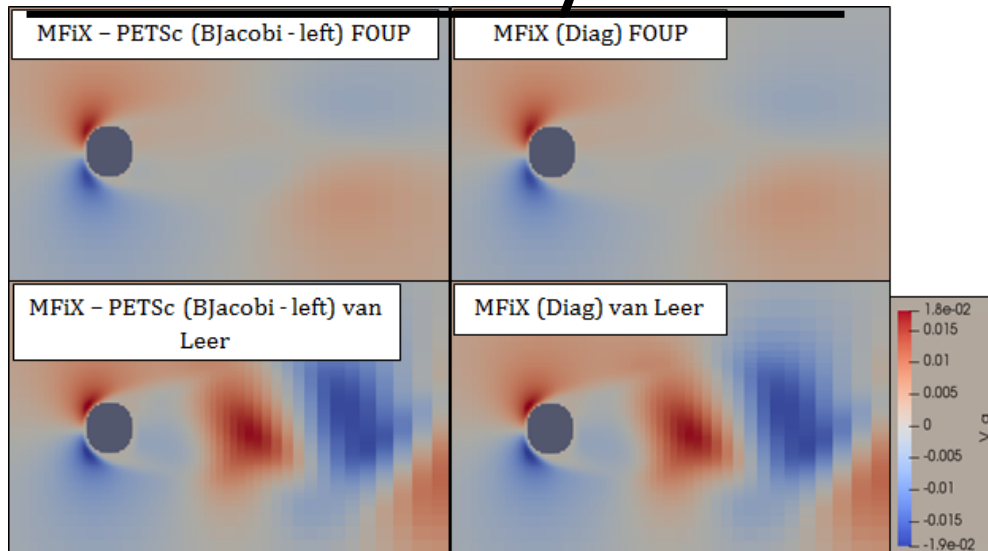


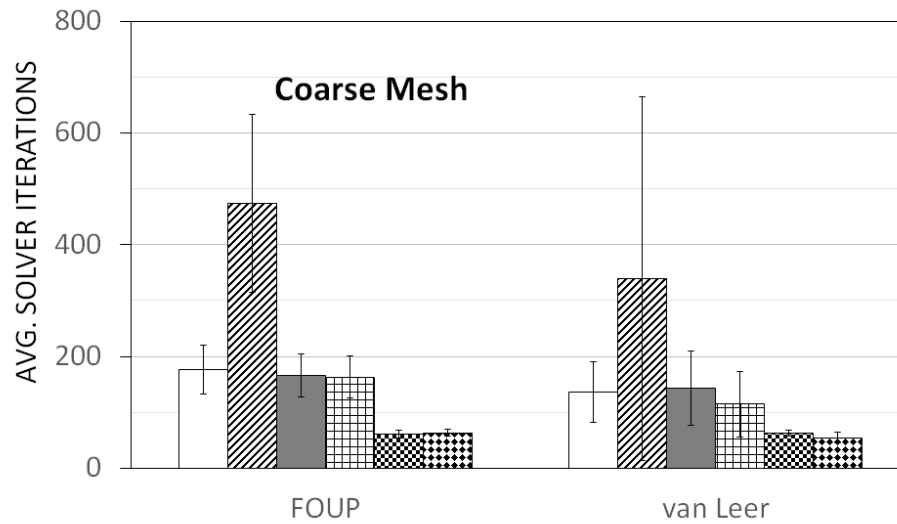
Asymmetry and the use of Conjugate Gradients



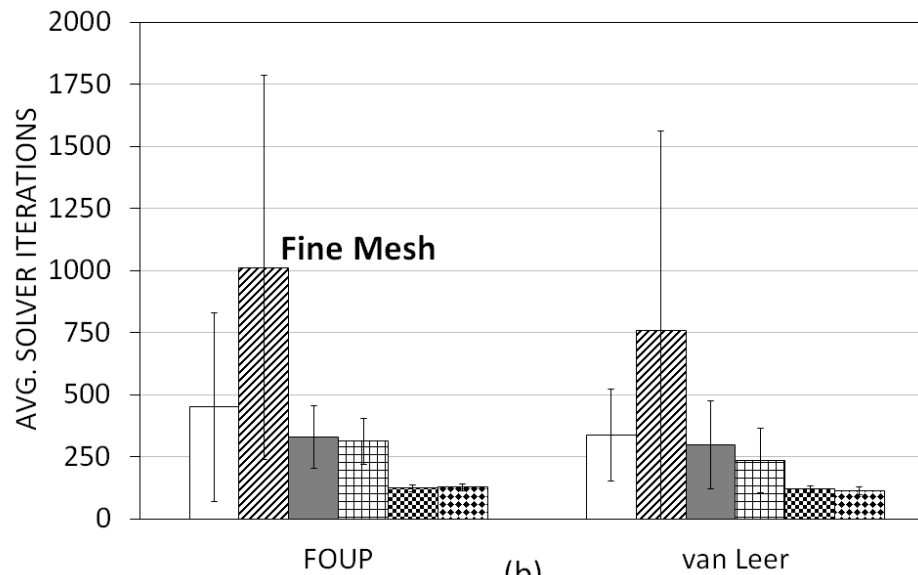
Figures 4 - 7: Plots of the (2) Total overall scaling (3) Matrix and vector object construction scaling (4) Solver object scaling and (5) BiCGStab scaling with an SMG preconditioner.

Flow Over Cylinder






(a)




(b)

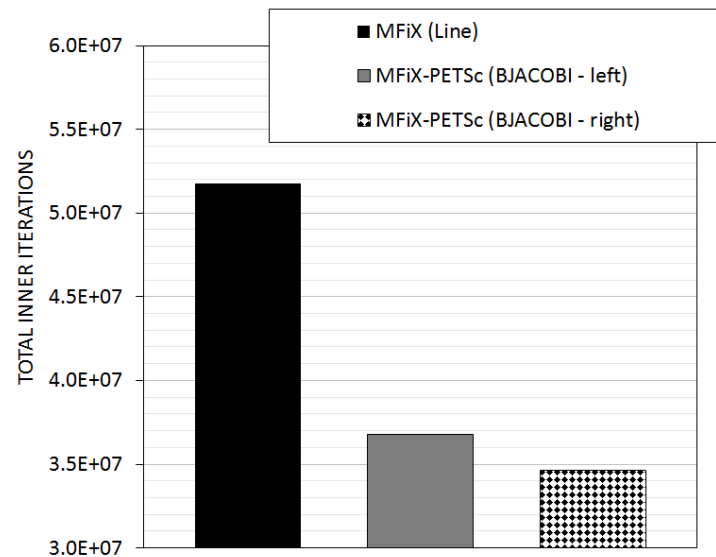
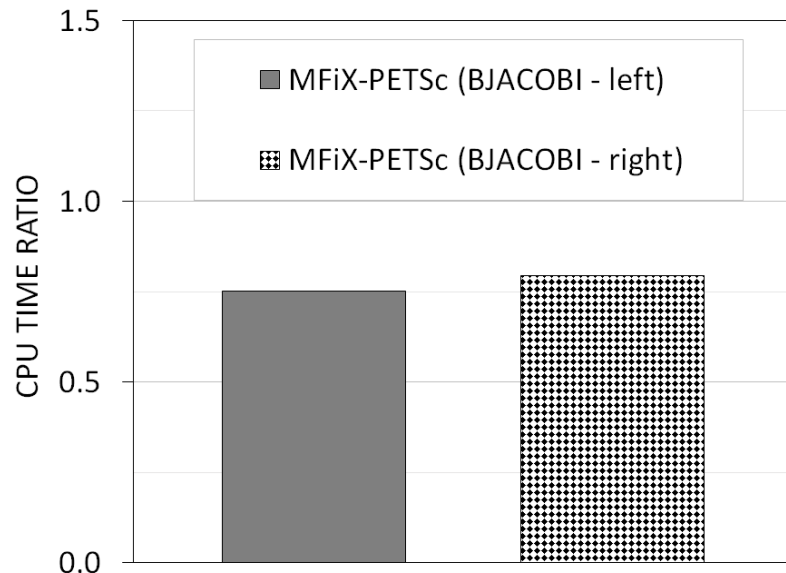





(a)




(b)




Polypropylene
(Non-Uniform Grid, 2D, H/D = 2)



MFiX (Line)




MFiX-PETSc
(BJACOBI - Left)




MFiX-PETSc
(BJACOBI - Right)


Glass
(Non-Uniform Grid, 2D, H/D = 2)



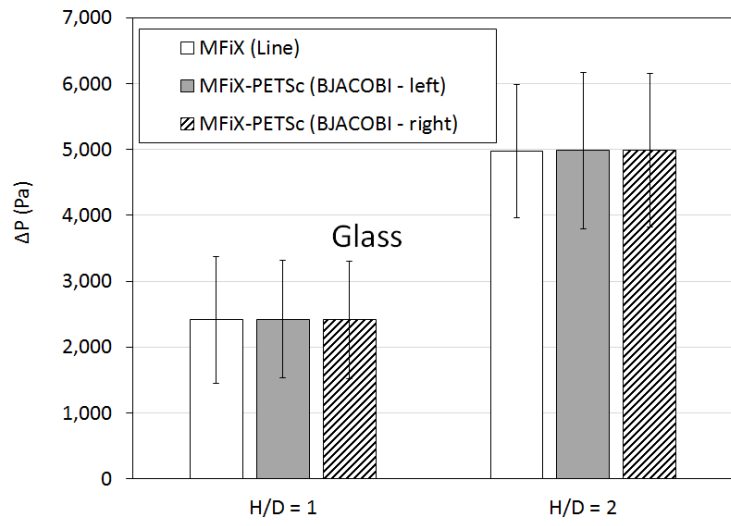
MFiX (Line)



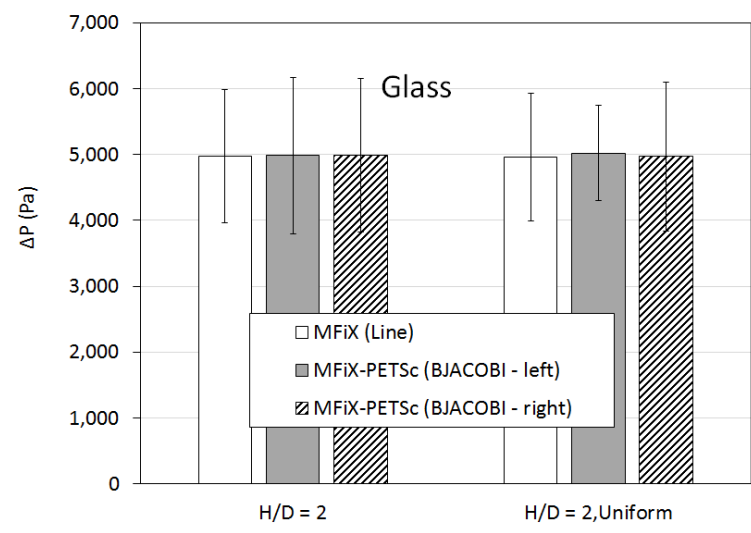
MFiX-PETSc
(BJACOBI - Left)



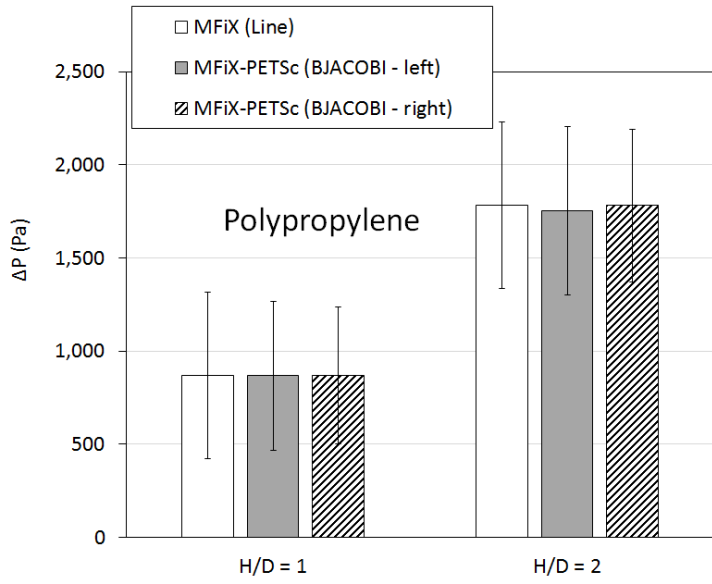
MFiX-PETSc
(BJACOBI - Right)




(a)



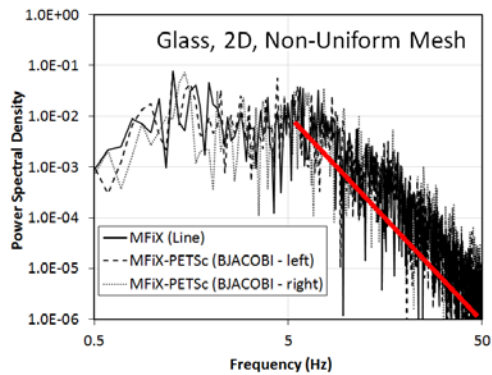
(b)



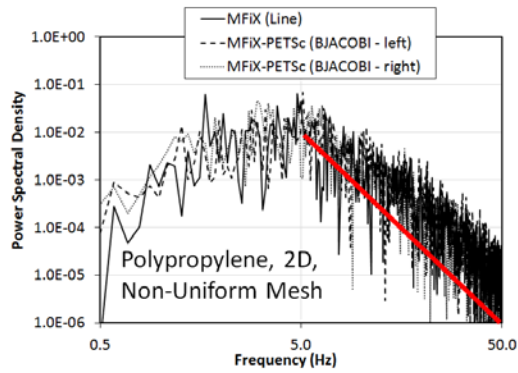
(c)



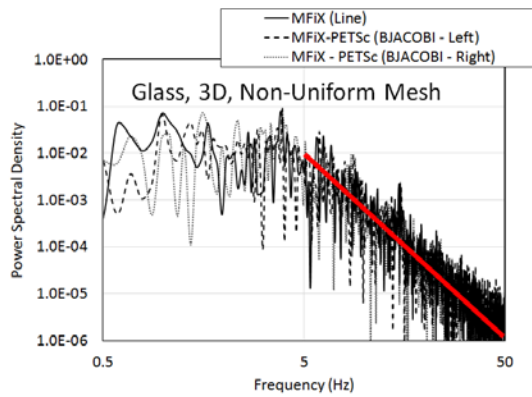
(d)



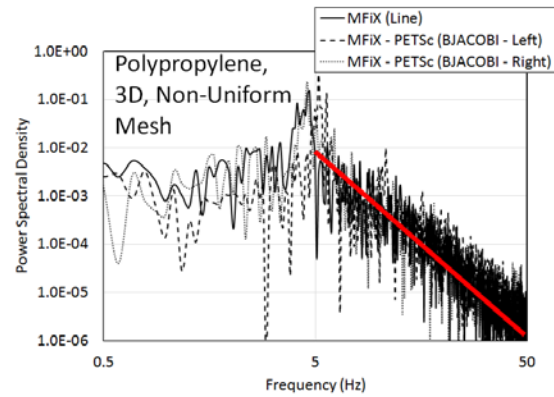
(a)



(b)




(c)




(d)

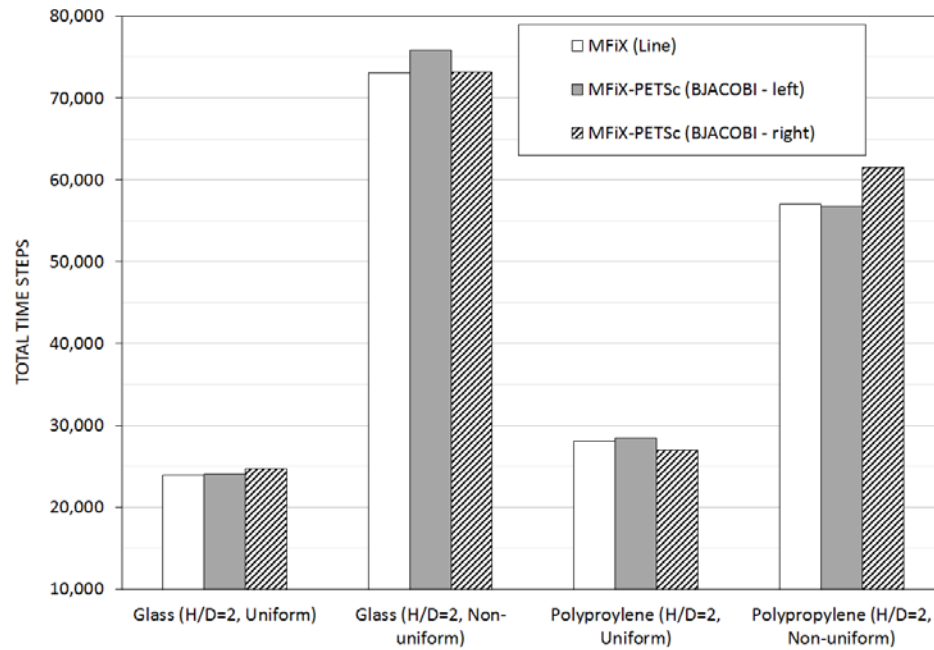
	Glass (H/D = 2)		PP (H/D = 2)	
Preconditioner	2D, Non-Uniform	3D, Non-Uniform	2D, Non-Uniform	3D, Non-Uniform
MFiX – line relaxation	1.3	3.8	4.8	4.7
MFiX-PETSc – BJACOBI (left)	1.3	3.9	5.0	5.2
MFiX PETSc – BJACOBI (right)	1.5	1.6	3.0	4.6




(a)




(b)




(a)




(b)




(a)



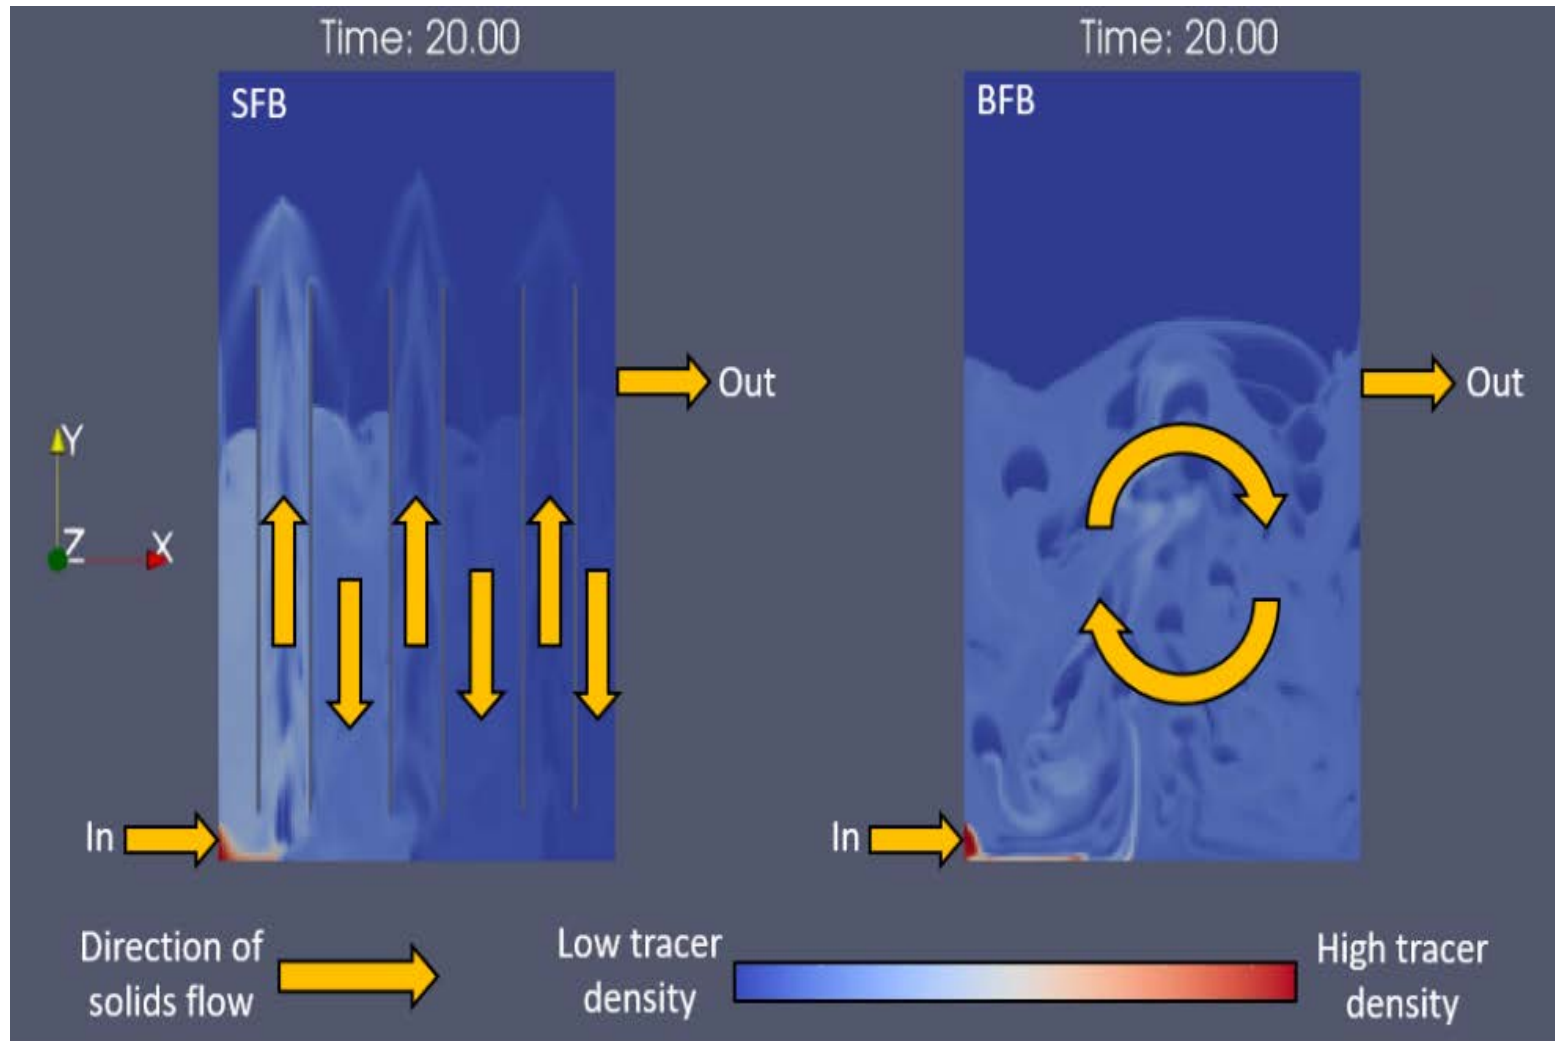
(b)



(a)



(b)



Scalability

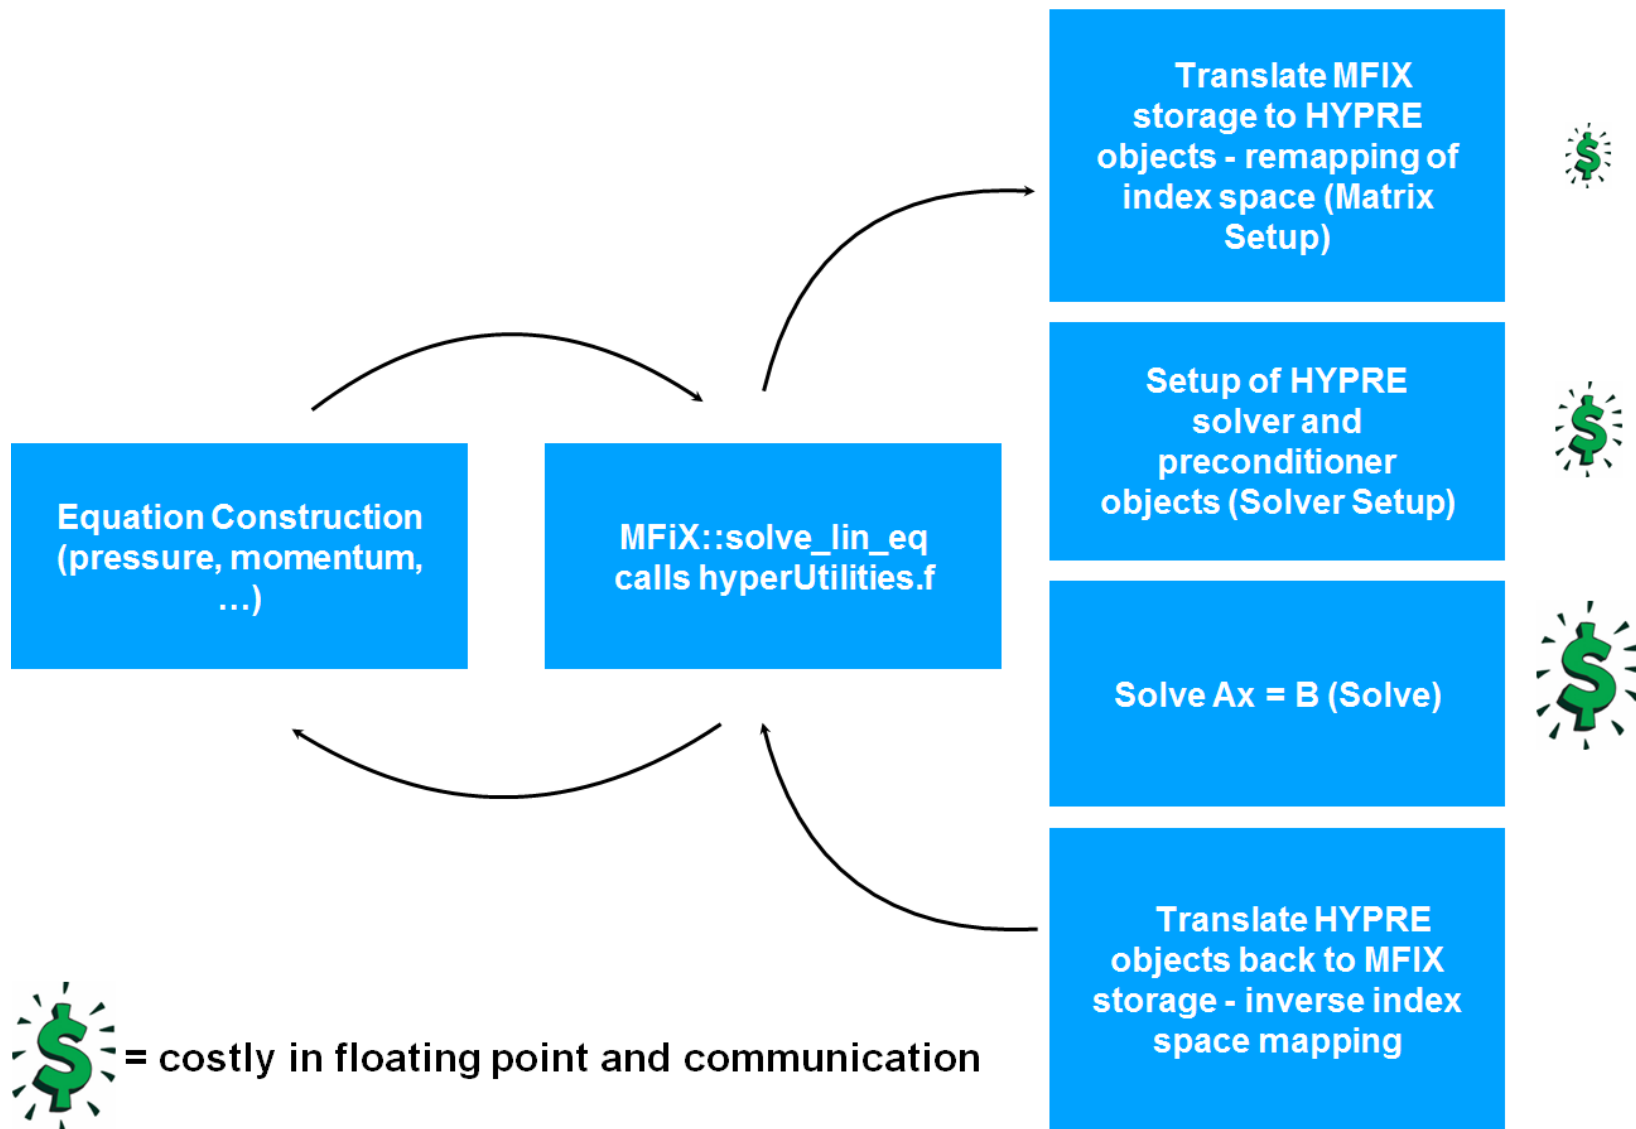
- Machine latency is a fixed cost - when computational work gets close or runs below this cost, parallel scalability (strong) breaks down
- Algorithmic decisions/design/methods contribute to scalability - some methods/algorithms (weak) scale better than others
- Multigrid solvers scales like $O(n)$ whereas CG scales like $O(n^{1.5})$ - $O(n^2)$
- Multigrid scalability has shown problem dependency

Method	Storage	Flops
GE(banded)	n^5	n^7
Gauss-Seidel	n^3	$n^5 \log n$
Optimal SOR	n^3	$n^4 \log n$
CG	n^3	$n^{3.5} \log n$
Full MG	n^3	n^3

To solve the Poisson problem in 3D ($N = n^3$)



Source: Kab Seok Kang, GIST HPC Summer school 2015

Modified “inner solve” - MFIX coupled with HYPRE



HYPRE Scaling in MFiX


HYPRE Solver and Matrix setup costs must be overcome with enough work




Comparison with MFiX on a single core

Where is the HYPRE overhead?


Matrix Setup




Solver Setup



Actual Solve

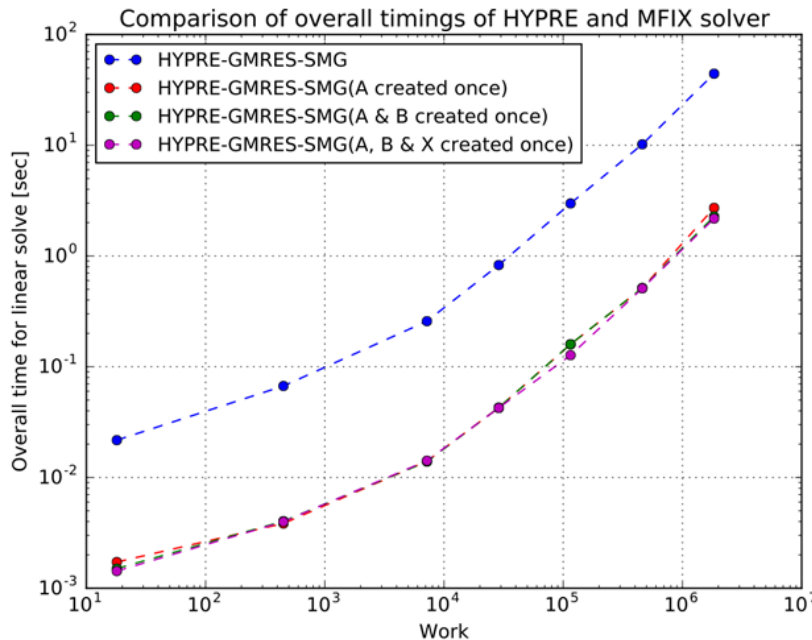


Actual Solve - strong and weak scaling






Setup Cost Mitigation

- Can setup costs be mitigated by:
 - reusing the solver/matrix objects every iteration?
 - over the lifetime of the simulation?



- Reuse of matrix objects shows same scaling behavior with a reduction in time
- Reuse of solver objects was met with memory leaks and other issues - still needs investigation
- Solver object problems could be due to underlying communication requirements in HYPRE for Struct interface
- For the Hype-Struct interface, no solver reinitialize option was available (as opposed to other interfaces, e.g. IJ-Interface)

Fluidized Bed Problem





- Initial apple-to-apples comparison (same solver tolerance, max iterations)
- MFiX native solver diverged while HYPRE converged
- PFMG preconditioner consistently offers better performance (well known)

Outer Solve Parameters Impact on Efficiency

Can HYPRE inner solver robustness be exploited?

Benchmark TFM03 Multiple Cores




Pressure Convergence with high outer relaxation factors




- HYPRE appears to support higher relaxation factors (it converges) for the outer solve
- Stalling or slow/noisy convergence seen with MFiX native solver with higher relaxation factors

Challenge Problem: Flow over a tube bank


Same Workload per Core



Increasing problem size -> -> ->




- Workload/Core $O(10^4)$
- Promising at lower core counts
- Should repeat with more workload/core



should be here

was here

Challenge Problem: Flow over a tube bank



- 5X more work/core
- ~3.5-2.5X speedup
- Needs more work, further study