Scaling an Open Source CFD Code by Leveraging Parallel Open Source Solvers within the Trilinos Framework – Preliminary Results

SAND2011-3432C

Aytekin Gel,

ALPEMI Consulting, LLC USA

Jonathan Hu*,

Sandia National Laboratories, USA

*Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Presentation Outline

- Objectives and Motivation
- Background & History
- Brief Introduction to Trilinos
- Current Implementation
- Summary & Conclusions
- Future Work
- Acknowledgments

ALPEMI CONSULTING

Objectives

- Achieve sustainable scalability for an open source CFD (MFIX) on modern HPC platforms, which are increasingly becoming heterogeneous many-core architectures.
- Off-load the burden to keep up with well tuned linear equation solver development and maintenance.
- Develop capability to interface an open source CFD code with available open source packages through standard APIs and reuse the features offered with these packages.

Motivation

- High-fidelity gasifier simulations are needed for several reasons:
  - Better understanding of the governing physics & scientific discovery.
  - Scalability study for new gasifier technology development form lab scale to commercial scale.
  - Optimization for robustness in performance.
  - Trouble shooting operational problems.
- Strong need for shorter cycle in time-to-market for gasifier technology development and deployment.
- Ability to conduct sufficiently accurate and fast gasifier simulations is important for the carbon constrained world.
- However, gasifier simulations are inherently compute intensive especially at the commercial scale.

ALPEMI CONSULTING

Challenge: How can we design commercial scale gasifiers for optimized operation?

Use validated computer models for answering scale up questions

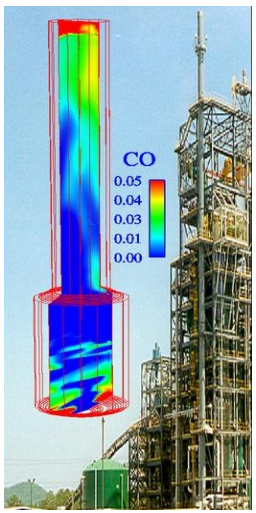

MFIX simulation of pilot scale 13 MW transport gasifier at Wilsonville, AL. Validation of the computer model with prototype system C. Guenther et al (2003)

<u>Parametric Study</u>

- Length/Diameter
- Coal feed rate
- Solids circulation rate
- Recycled syngas
- Coal jet penetration

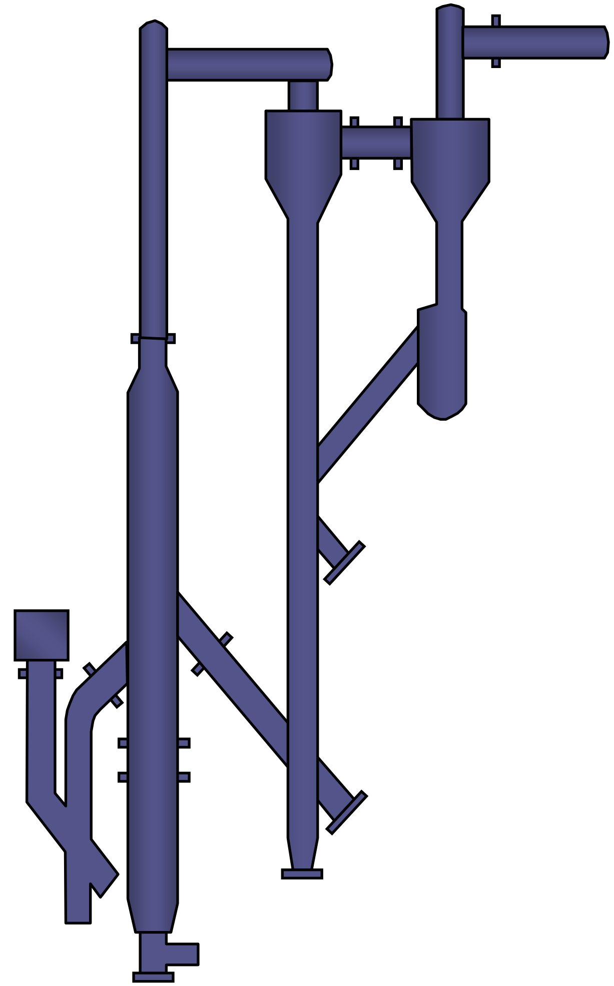

Commercial Gasifier

Source: Syamlal et al. (2009)

MFIX, Open Source Multiphase Flow Code

Mass conservation for phase m (m=g for gas and s for solids)

$$\frac{\partial}{\partial t}(\varepsilon_m \rho_m) + \nabla \cdot (\varepsilon_m \rho_m \vec{v}_m) = \sum_{l=1}^{N_m} R_{ml}$$

Momentum conservation

$$\frac{\partial}{\partial t}(\varepsilon_m \rho_m \vec{v}_m) + \nabla \cdot (\varepsilon_m \rho_m \vec{v}_m \vec{v}_m) = \nabla \cdot \overline{\overline{S}}_m + \varepsilon_m \rho_m \vec{g} + \sum_n \vec{I}_{mn}$$

Granular energy conservation (m ≠ g)

$$\frac{3}{2}\varepsilon_m \rho_m \left(\frac{\partial \Theta_m}{\partial t} + \vec{v}_m \cdot \nabla \Theta_m \right) = \nabla \cdot \vec{q}_{\Theta_m} + \overline{\overline{S}}_m : \nabla \vec{v}_m - \varepsilon_m \rho_m J_m + \Pi_{\Theta_m}$$

Energy conservation

$$\varepsilon_m \rho_m C_{pm} \left(\frac{\partial T_m}{\partial t} + \vec{v}_m \cdot \nabla T_m \right) = -\nabla \cdot \vec{q}_m + \sum_n \gamma_{mn} (T_n - T_m) - \Delta H_{rm}$$

Species mass conservation

$$\frac{\partial}{\partial t}(\varepsilon_m \rho_m X_{ml}) + \nabla \cdot (\varepsilon_m \rho_m X_{ml} \vec{v}_m) = R_{ml}$$

http://mfix.netl.doe.gov

R&D100
Award 2007

Mid-Atlantic Region
Tech-Transfer
Award 2006

Sources:

- Syamlal et al. "MFIX Documentation, Theory Guide," DOE/METC-94/1004, NTIS/DE94000087 (1993)
- Benyahia et al. "Summary of MFIX Equations 2005-4", From URL http://www.mfix.org/documentation/MfixEquations2005-4-3.pdf, July 2007.

Historical Perspective: Reactor simulations with MFIX over the years

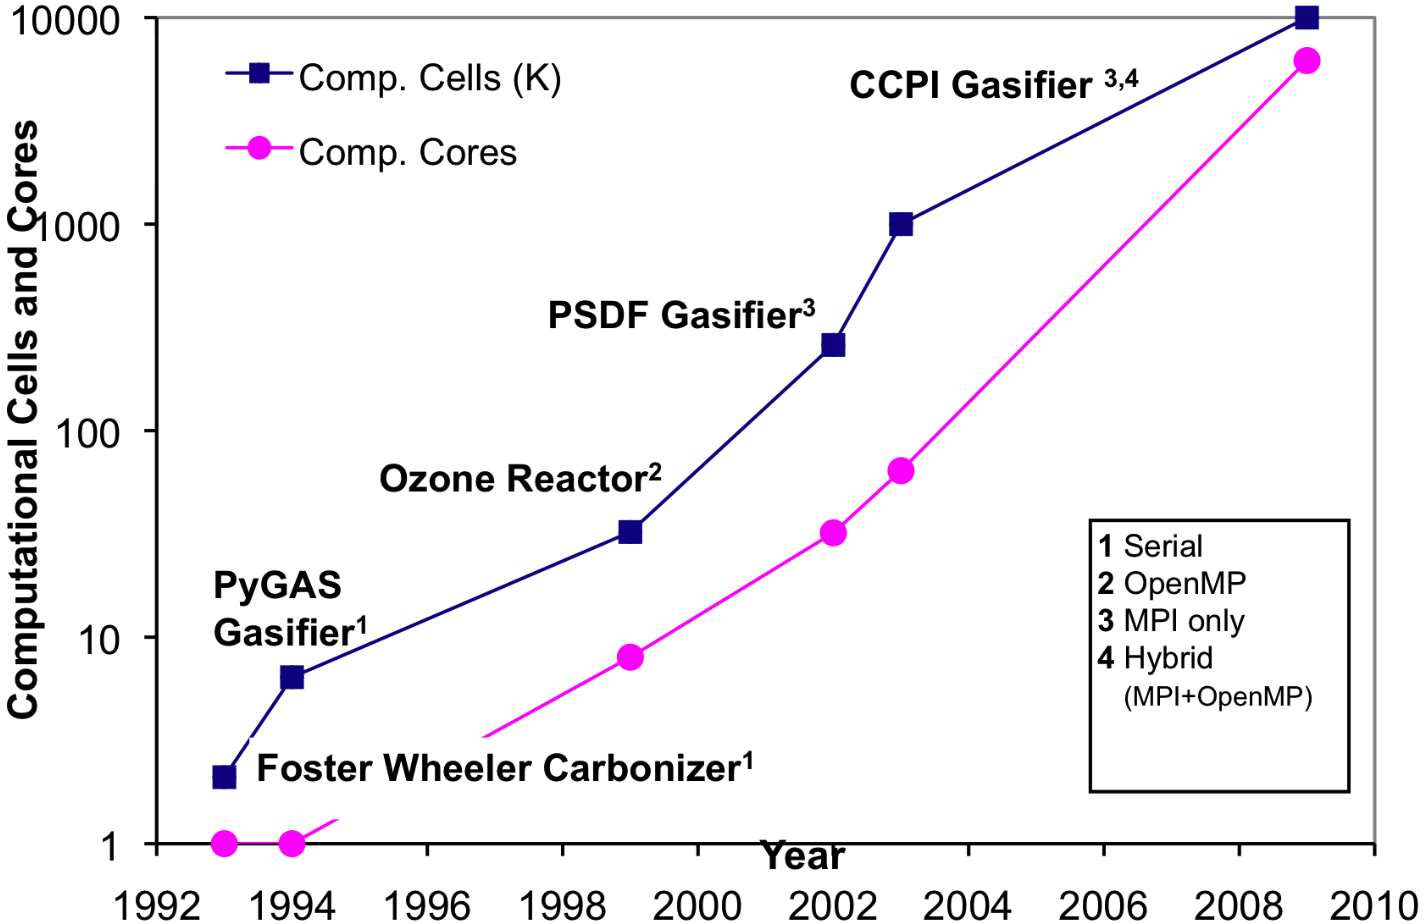

Source: Gel et al. (2010)

ALPEMI CONSULTING

Performance Analysis of MFIX:

- Typical commercial scale gasifier run requires 165K hrs per 1 second of simulation for a 10M cell resolution case running on 6192 cores. We need at least 15-20 second simulation per design configuration to extract useful information
- So reducing the this wall clock time is critical.
- Performance analysis of MFIX using tools like CrayPAT and TAU show our bottlenecks in linear equation solvers. Solution => replace our native BiCGSTAB linear equation solver with scalable solvers such as those from Trilinos framework.

ALPEMI CONSULTING

The Trilinos Project

- Open source project led by Sandia National Labs (USA)
  - Distributed development (labs, academia)
  - LGPL (transitioning to simplified BSD)
- Collection of high performance parallel numerical libraries
  - Predominantly C++ (some C and Fortran)
  - Developed by domain experts
- Capabilities
  - Fundamental linear algebra, interfaces
  - Discretizations, meshing, time integration, load balancing
  - Solvers: linear, nonlinear, eigen, stochastic, direct
  - Preconditioners: algebraic multigrid, incomplete factorizations, segregated
  - Optimization, automatic differentiation

The Trilinos Project

- Targets multicore CPU, mixed CPU/GPU architectures
- Support for MPI-only, MPI+threading
- Some applications that use Trilinos
  - Climate (Community Climate System Model)
    - HOMME (atmospheric component @ ORNL)
    - POP (Parallel Ocean Program @ LANL)
    - Glimmer (Ice sheet modeling @ U. Bristol)
    - Epetra, NOX, Loca, Belos, ML, Ifpack
  - Denovo (light water reactor sim. @ ORNL)
    - Anasazi
  - deal.II, LifeV, PETSc
    - Epetra, AztecOO, Belos, ML, Ifpack, …
- More information:
  - http://trilinos.sandia.gov

Trilinos ML/AztecOO Weak Scaling Example

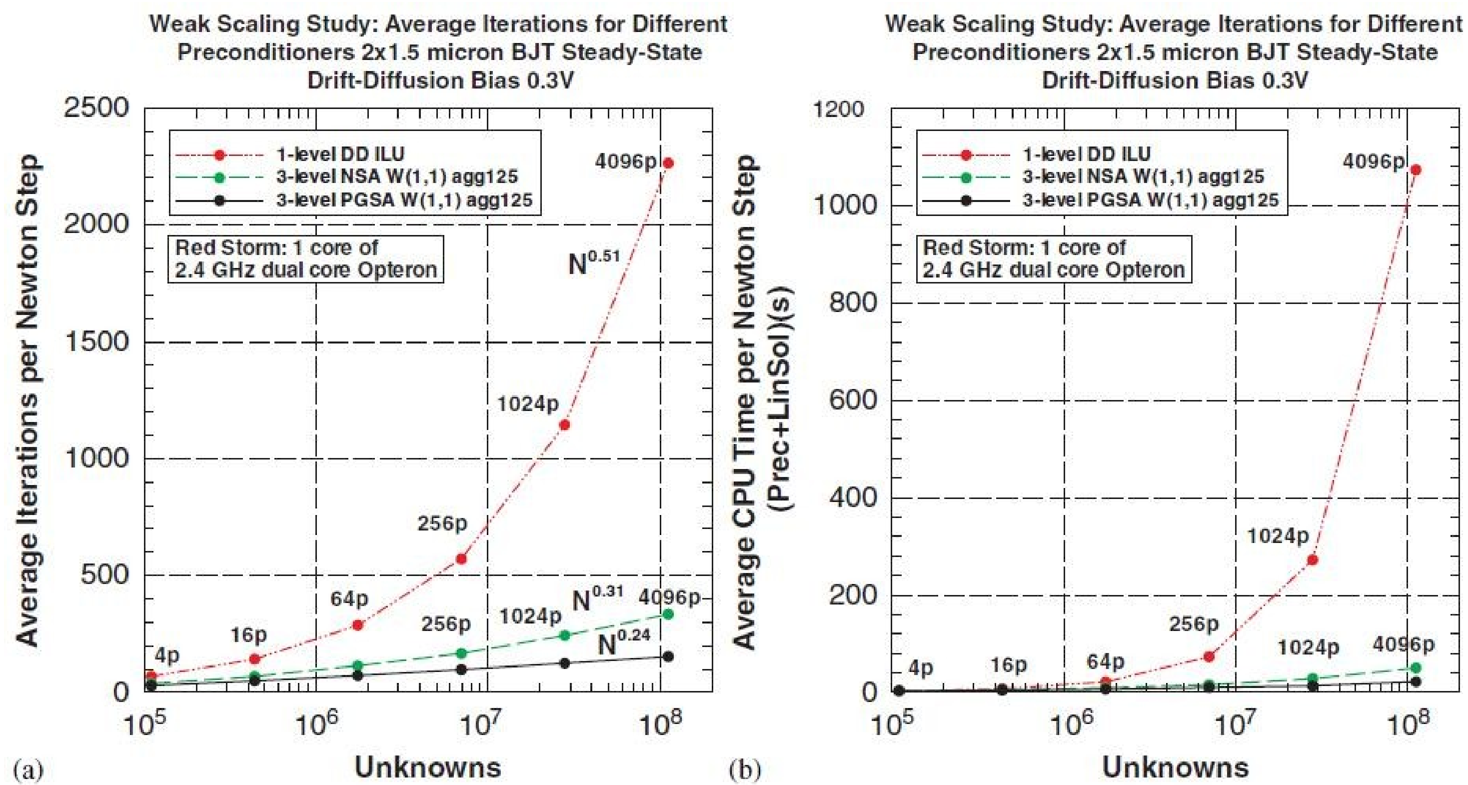

Figure 3. (a) Comparison of average number of iterations per Newton step vs problem size for the one-level and multilevel preconditioners for the 2D BJT and (b) average time per Newton step for the preconditioners.

From: P. Lin, et al., *Performance of a Petrov-Galerkin algebraic multilevel preconditioner for finite element modeling of the semiconductor device drift-diffusion equations*, Int. J. Numer. Meth. Engng, 84(4), p. 448-469, Sandia Technical Report 2009-3682 J

Brief Information on Trilinos:

Trilinos Package Summary

	Objective	Package(s)
Discretizations	Meshing & Spatial Discretizations	Intrepid, Phalanx, Shards, Pamgen, Sundance, Mesquite, STK, Moertel
	Time Integration	Rythmos
Optimization	Optimization (SAND)	MOOCHO, Aristos
Methods	Automatic Differentiation	Sacado
Core	Linear algebra objects	Epetra, Jpetra, Tpetra
	Abstract interfaces	Thyra, Stratimikos, RTOp, Piro
	Load Balancing	Zoltan, Isorropia
	"Skins"	PyTrilinos, WebTrilinos, Star-P, ForTrilinos, CTrilinos
	C++ utilities, (some) I/O	Teuchos, EpetraExt, Kokkos, Triutils
	GUIs	Optika
Preconditioners	Multigrid methods	ML
	Domain decomposition methods	CLAPS, IFPACK
	ILU-type methods	AztecOO, IFPACK, TIFPACK
	Block preconditioners	Teko, Meros
Solvers	Iterative (Krylov) linear solvers	AztecOO, Belos, Komplex
	Direct sparse linear solvers	Amesos
	Direct dense linear solvers	Epetra, Teuchos, Pliris
	Nonlinear system solvers	NOX, LOCA
	Iterative eigenvalue solvers	Anasazi
	Stochastic PDEs	Stokhos

9

Packages used in current implementation

Additional information on Trilinos can be found at http://trilinos.sandia.gov/

ALPEMI CONSULTING

Implementation

- A linear equation solver manager (LEQ_manager) is created
- Generate data in MFIX & passed through to the LEQ_manager in order to create an Epetra_CrsMatrix. Same is done for the right-hand side vector.
- Options are set for both ML (algebraic multigrid preconditioner) and AztecOO (Krylov iterative solver).
- The linear equation problem is solved & solution vector is passed back to Fortran side in MFIX.

Implementation:

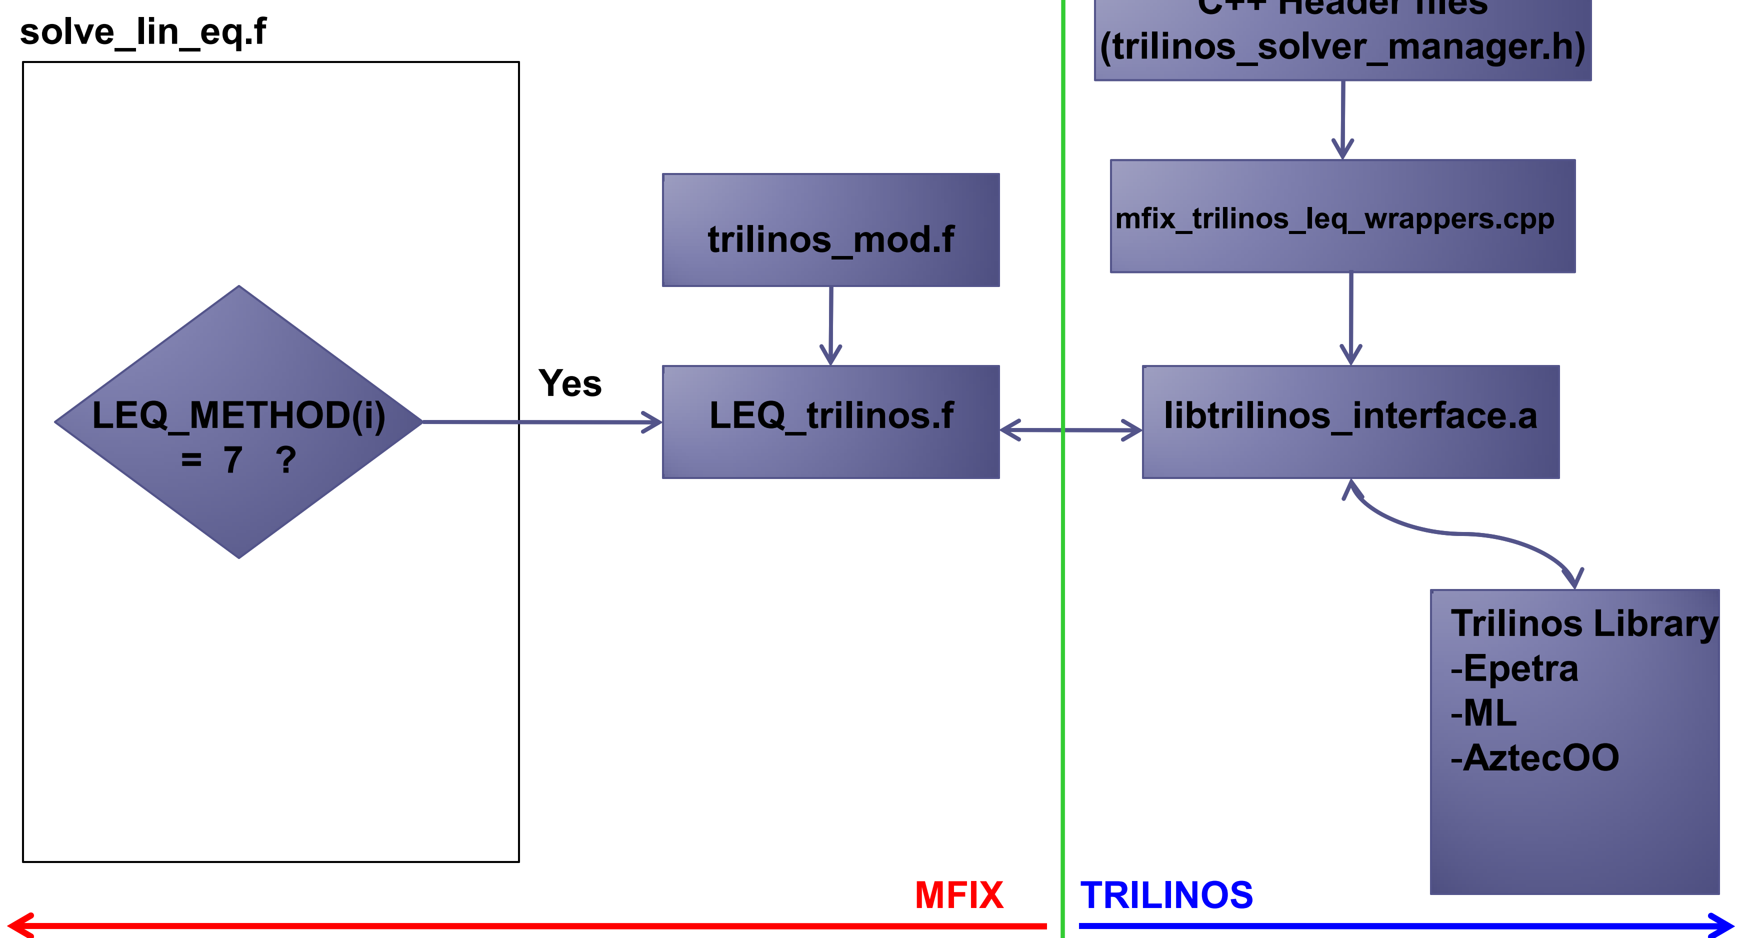

Future Improvement: Replace wrappers with ForTrilinos.

Alpemi Consulting

Parameter Input to Trilinos through “trilinos.dat”

Due to the string based input framework within Trilinos a new feature was added to MFIX to parse a separate input file*

trilinos.dat

```
#
### Input Parameters to be passed to Trilinos solvers from MFIX
#ML_Precond_Block
SA
ML output, 0, int
……
……
……
……
#Aztec_Solver_Block
AZ_solver, AZ_cg, string
AZ_output, AZ_last, int
AZ_max_iter, 250, int
AZ_tol, 1e-8, double
……
……
```

Preconditioner related settings for ML package

Linear Equation solver related settings for AztecOO

Note: Several open source Fortran string utilities were added to MFIX in order to parse the Trilinos data shown above

*** Future Improvement: merge trilinos.dat into mfix.dat.**

ALPEMI CONSULTING

Parameter List for ML Preconditioner:

+++ Printing ML parameter list "SA default values" on pid 0 :

```
ML debug mode = 0    [default]
ML output = 1    [unused]
XML input file = ml_ParameterList.xml    [default]
aggregation: damping factor = 1.333    [unused]
aggregation: type = Uncoupled-MIS    [unused]
coarse: max size = 128    [unused]
coarse: pre or post = post    [unused]
coarse: sweeps = 1    [unused]
coarse: type = Amesos-KLU    [unused]
default values = SA    [unused]
eigen-analysis: iterations = 10    [unused]
eigen-analysis: type = cg    [unused]
increasing or decreasing = increasing    [unused]
max levels = 3    [unused]
prec type = MGV    [unused]
print unused = 0    [unused]
read XML = 0    [default]
smoother: damping factor = 1    [unused]
smoother: pre or post = both    [unused]
smoother: sweeps = 1    [unused]
smoother: type = symmetric Gauss-Seidel    [unused]
----------- end of ML parameter list ------------
```

ALPEMI CONSULTING

Preliminary Results:

- Validation Case : 3d_square_cfb
- Grid Resolution : 30 x 200 x 30 [180K cells]

- Domain Decomposition : k-direction only
- Number of processors : 2, 4, 12, 24, 36
- Test Platform : mist@cs.uoregon.edu

Validation: Comparison of residuals & results

```
Time =  0.20000E-01  Dt =  0.20000E-03    CPU time left =  2081.758 s
 Nit      P0        P1       U0       V0       U1       V1     Max res
   1  5.3E-04  9.7E-02  1.5E-04  2.8E-03  3.6E-04  1.3E-02  P1
   2  1.0E-02  1.8E-03  5.6E-05  1.0E-03  1.8E-04  6.5E-03  P0
   3  2.4E-03  9.0E-04  2.7E-05  5.5E-04  9.0E-05  3.3E-03  V1
   4  1.4E-03  4.5E-04  1.3E-05  2.8E-04  4.5E-05  1.6E-03  V1
   5  7.8E-04  2.3E-04  6.7E-06  1.4E-04  2.3E-05  8.3E-04  V1
   6  4.5E-04  1.2E-04  3.4E-06  7.4E-05  1.1E-05  4.2E-04  P0
   7  2.6E-04  5.9E-05  1.7E-06  3.8E-05  5.7E-06  2.1E-04  P0
t=     0.0200 Dt=0.2000E-03 NIT=   7 Sm=11292.
  MbError%(0,MMAX): 0.9722E-07 -.9432E-12
```

Pressure Eqns. solved with MFIX's native BiCGSTAB

LEQ_METHOD(1) = 2

```
t=     0.0200  Wrote       SPx: 1,
t=     0.0200  Wrote RES;
Time =  0.20000E-01  Dt =  0.20000E-03    CPU time left =  2628.584 s
 Nit      P0        P1       U0       V0       U1       V1     Max res
   1  5.3E-04  9.7E-02  1.5E-04  2.8E-03  3.6E-04  1.3E-02  P1
   2  5.0E-03  1.8E-03  5.5E-05  1.0E-03  1.8E-04  6.5E-03  V1
   3  2.3E-03  9.0E-04  2.7E-05  5.5E-04  9.0E-05  3.3E-03  V1
   4  1.3E-03  4.5E-04  1.3E-05  2.8E-04  4.5E-05  1.6E-03  V1
   5  7.6E-04  2.3E-04  6.7E-06  1.4E-04  2.3E-05  8.3E-04  V1
   6  4.4E-04  1.2E-04  3.4E-06  7.4E-05  1.1E-05  4.2E-04  P0
   7  2.5E-04  5.9E-05  1.7E-06  3.8E-05  5.7E-06  2.1E-04  P0
t=     0.0200 Dt=0.2000E-03 NIT=   7 Sm=11292.
  MbError%(0,MMAX): -.4494E-13 -.4261E-13
```

Pressure Eqns. solved with Trilinos (ML+AztecOO)

LEQ_METHOD(1) = 7

ALPEMI CONSULTING

Preliminary Timing Results for 4 core runs:

With Native BiCGSTAB solver LEQ_METHOD(1) = 2	With Trilinos (ML/BiCGSTAB) LEQ_METHOD(1) = 7
812 s.	1057 s.

Input Parameters:	Input Parameters:
LEQ_PC(1)= ‘DIAG’	smoother: type = symmetric Gauss-Seidel
LEQ_IT(1) = 100	AZ_max_iter= 100
LEQ_TOL(1) = 1.0e-4	AZ_tol = 1e-6

Two order of magnitude in convergence tolerance and symm. Gauss-Seidel smoother, which degrades in parallel execution appears longer time needed for Trilinos solvers. Comparison for timings during the bench-marking study need to made based on equivalent settings for apples to apples comparison.

ALPEMI CONSULTING

Summary & Conclusions

- A preliminary version of MFIX using Trilinos framework for solving linear equations has been developed using MFIX-Hypre interface as the baseline template.
- To expedite the development current version has some simplifications such cartesian grid and I or K direction decomposition only.
- Need to conduct extensive benchmarking tests to determine best set of options from Trilinos solvers for comparable problem types and sizes.

Future Work: Linear Equation Solvers

- Remove the assumptions used in initial version (i.e., enable for 3D decomposition, cylindrical grid, etc.).
- Employ pointers instead of memory copies from MFIX to create A matrix and B vector in Trilinos.
- Test various options in ML and AztecOO to determine a range of settings that gives the best performance for various MFIX application cases (stiff problems needed).
- Profile & Optimize (e.g. save EpetraCrsMatrix object)
- Create Trilinos based version of test and tutorial cases.
- Make the implementation ready for distribution to MFIX users through CVS.
- Explore PETSc library using the current interface, which enables PETSc access through Trilinos.

ALPEMI CONSULTING

Future Work: NonLinear Equation Solvers

- Initiated the discussion with Trilinos team on how to interface MFIX with NOX from Trilinos (nonlinear equation solver), and the necessary structural changes.
- NOX has been hooked up to several Fortran codes, and Andy was working towards NOX integration with an F90 code (Ice Sheet) in October.
- A general outline for MFIX implementation:
  - Need to write two functions in Fortran that can be called by NOX and each take a solution vector as input:

    (1) Function to calculate residuals,

    (2) Function to calculate Jacobian.

Acknowledgments

- Additional technical support was provided by Trilinos Development Team.

- This research used resources of the Performance Research Lab at the University of Oregon and the TAU project for the resources.

References:

Gel, A., Li, J., Guenther, C. and Syamlal, M. (2010) "High Resolution Numerical Simulations of Coal Gasifiers Using High Performance Computing" 2010 Workshop on Multiphase Flow Science, May 4-6, 2010, Pittsburgh, PA. URL: http://www.netl.doe.gov/publications/proceedings/10/mfsw/index.html

Syamlal, M., Guenther, C., Gel, A. & Pannala, S. (2009) "Advanced Coal Gasifier Designs Using Large-Scale Simulations", SciDAC 2009, June 14-18, San Diego URL: https://hpcrd.lbl.gov/scidac09/talks/Syamlal_AdvancedGasifierModeling_v3.pdf

Gel et al. (2007) "Comparison of frameworks for a next-generation multiphase flow solver, MFIX: a group decision-making exercise" Journal Concurrency and Computation: Practice & Experience - Component and Framework Technology in High-Performance and Scientific Computing archive, Vol. 19, Issue 5, doi>10.1002/cpe.v19:5

An Overview of Trilinos, Eleventh DOE ACTS Collection Workshop 2010, URL: http://acts.nersc.gov/events/Workshop2010/Agenda10_files/Trilinos.pdf

Appendix

- Native BiCGSTAB solver performance statistics.

Native BiCGSTAB solver performance:

Single timestep run with native BiCGSTAB solver and solver statistics:

```
....
 Time =    0.0000       Dt =   0.20000E-03    CPU time left =     0.000 s
  Nit      P0          P1          U0          V0          U1         V1       Max res
    1  1.2E-03   1.6E-02   2.3E-03   0.4        0.        0.1        V0
    2  0.3       3.3E-02   2.1E-03   0.2       1.6E-02   0.8        V1
    3  0.2       2.4E-02   1.4E-03   0.1       7.2E-03   0.2        V1
    4  0.1       1.7E-02   9.0E-04   4.6E-02   4.5E-03   8.1E-02   P0
    5  6.8E-02   1.1E-02   5.7E-04   2.0E-02   2.9E-03   3.0E-02   P0
    6  4.1E-02   7.3E-03   3.5E-04   8.6E-03   1.9E-03   1.1E-02   P0
    7  2.4E-02   4.6E-03   2.1E-04   3.7E-03   1.3E-03   4.4E-03   P0
    8  1.4E-02   2.9E-03   1.3E-04   1.6E-03   8.0E-04   1.7E-03   P0
    9  8.6E-03   1.8E-03   7.9E-05   7.1E-04   5.1E-04   8.7E-04   P0
   10  5.1E-03   1.1E-03   4.7E-05   3.1E-04   3.2E-04   5.8E-04   P0
   11  3.0E-03   6.9E-04   2.9E-05   1.4E-04   2.0E-04   3.6E-04   P0
   12  1.8E-03   4.3E-04   1.7E-05   6.4E-05   1.2E-04   2.2E-04   P0
   13  1.1E-03   2.6E-04   1.0E-05   3.0E-05   7.4E-05   1.3E-04   P0
   14  6.4E-04   1.6E-04   6.3E-06   1.4E-05   4.5E-05   7.5E-05   P0
 t=     0.0000 Dt=0.2000E-03 NIT= 14 Sm=11255.           CPU=      30. s
   MbError%(0,MMAX): 0.1473E-06 0.3979E-10
Elapsed CPU time =     0.312993E+02 sec
 t=     0.0002  Wrote      SPx: 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B,              Disk=
25.25 Mb
```

ALPEMI CONSULTING

Native BiCGSTAB solver performance:

Single timestep run with native BiCGSTAB solver and solver statistics:

```
....
Total number of non-linear iterations           14
 Average number per time-step           14
 Equation number-----Number of linear solves
            1 ---------         714
            2 ---------           4
            3 ---------          54
            4 ---------          56
            5 ---------          54
            6 ---------           0
            7 ---------           0
            8 ---------           0
            9 ---------           0
           10 ---------           0
 Equation number-----Avg. number of linear solves for NIT
            1 ---------          51
            2 ---------           0
            3 ---------           3
            4 ---------           4
            5 ---------           3
….
…….
```

Pressure solution requires approx. 51 linear equation solver iterations

Native BiCGSTAB solver performance:

Run (TSTOP = 0.05 s) with native BiCGSTAB solver & solver statistics:

```
....
Total number of non-linear iterations          1692
 Average number per time-step                 6
 Equation number-----Number of linear solves
            1 ---------        86292
            2 ---------         3354
            3 ---------         6766
            4 ---------         6768
            5 ---------         7119
            6 ---------            0
            7 ---------            0
            8 ---------            0
            9 ---------            0
           10 ---------            0
 Equation number-----Avg. number of linear solves for NIT
            1 ---------           51
            2 ---------            1
            3 ---------            3
            4 ---------            4
            5 ---------            4
              ….
              …….
```

Pressure solution requires approx. 51 linear equation solver iterations.

Trilinos (ML/BICGSTAB) solver performance:

Run with ML preconditioning and AztecOO BiCGSTAB :

```
....
Time =    0.0000      Dt =   0.20000E-03    CPU time left =     0.000 s
                *******************************************************
                ***** Problem: Epetra::CrsMatrix
                ***** Preconditioned BICGSTAB solution
                ***** ML (L=3, IFPACK_pre0/IFPACK_post0, ~/Amesos_KLU_2)
                ***** No scaling
                *******************************************************

                iter:    0          residual = 1.000000e+00
                iter:    1          residual = 2.511064e-02
                iter:    2          residual = 3.708502e-04
                iter:    3          residual = 5.017193e-06
                iter:    4          residual = 6.164090e-08


                Solution time: 0.539572 (sec.)
                total iterations: 4
  Nit      P0        P1        U0        V0        U1        V1     Max res
    1   1.2E-03   1.6E-02   2.3E-03   0.4        0.        0.1      V0
          ….
          …….
```

Pressure solution requires 4 linear equation solver iterations after preconditioning phase.

ALPEMI CONSULTING