SAND2016-5317C

Exceptional service in the national interest

Sandia National Laboratories

$$\frac{\partial m_a}{\partial t} = -\nabla \cdot \left(\rho_l X_a^l \boldsymbol{q}_l + \rho_g X_a^g \boldsymbol{q}_g + \boldsymbol{J}_a^l + \boldsymbol{J}_a^g\right) + q_a^G,$$
$$\frac{\partial m_w}{\partial t} = -\nabla \cdot \left(\rho_l X_w^l \boldsymbol{q}_l + \rho_g X_w^g \boldsymbol{q}_g + \boldsymbol{J}_w^l + \boldsymbol{J}_w^g\right) + q_w^G,$$
$$\frac{\partial e}{\partial t} = -\nabla \cdot \left(\rho_l H_l \boldsymbol{q}_l + \rho_g H_g \boldsymbol{q}_g - \kappa_{\text{eff}} \nabla T\right) + q_e^G,$$

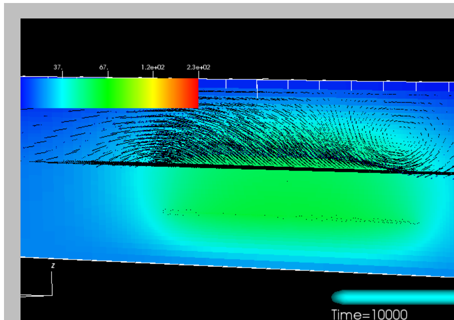

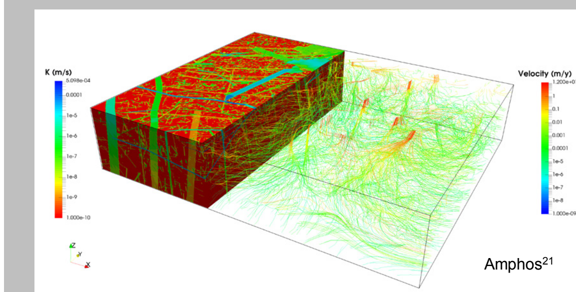

PFLOTRAN: Practical Application of High Performance Computing to Subsurface Simulation

Glenn Hammond

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

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. SAND NO. 2011-XXXXP

PFLOTRAN

Sandia National Laboratories

- Petascale reactive multiphase flow and transport code
- Open source license (GNU LGPL 2.0)
- Object-oriented Fortran 2003/2008
  - Pointers to procedures
  - Classes (extendable derived types with member procedures)
- Founded upon PETSc parallel framework
  - Parallel communication through MPI
  - Parallel I/O through binary HDF5
  - Unstructured domain decomposition through METIS/ParMETIS (Cmake)
- Demonstrated performance
  - Maximum # processor cores: 262,144 (Jaguar supercomputer)
  - Maximum problem size 3.34 billion degrees of freedom
  - Scales to over 10K cores

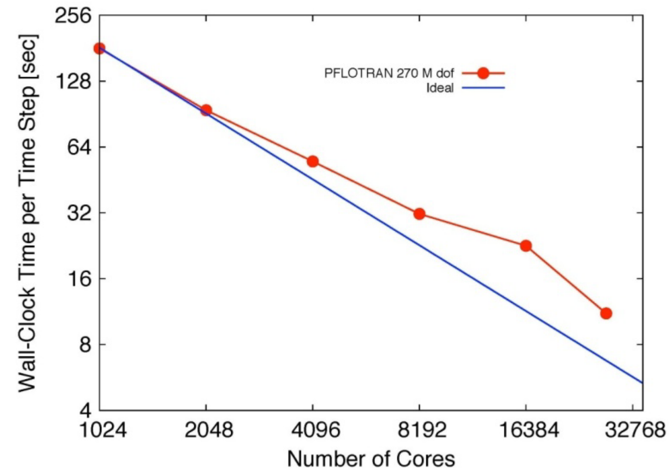

Application of PFLOTRAN

- Nuclear waste disposal
  - Waste Isolation Pilot Plant (WIPP) in Carlsbad, NM
  - DOE Used Fuel Disposition Program
  - SKB Forsmark Spent Fuel Nuclear Waste Repository (Sweden, Amphos[21])
- Climate: coupled overland/groundwater flow; CLM
  - Next Generation Ecosystem Experiments (NGEE) Arctic
  - DOE Earth System Modeling (ESM) Program
- Biogeochemical transport modeling
  - U(VI) fate and transport at Hanford 300 Area
  - Hyporheic zone biogeochemical cycling
    - Columbia River, WA, USA
    - East River, CO, USA
- CO_2 sequestration
- Enhanced geothermal energy
- Radioisotope tracers
- Colloid-facilitated transport

East River, CO

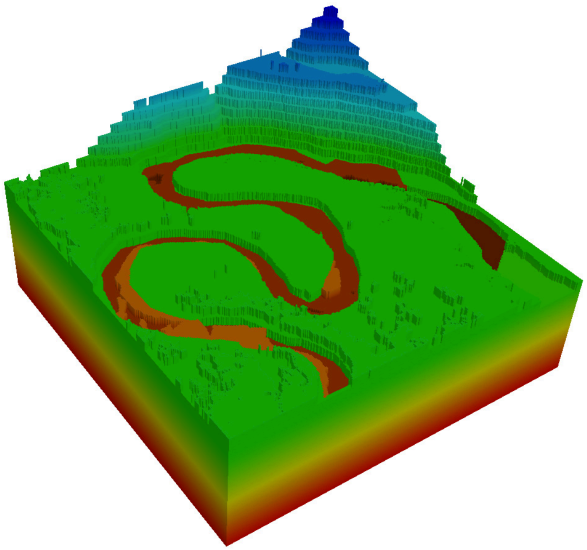

Dipankar Dwivedi, LBNL, 2016

Numerical Methods

Sandia National Laboratories

- Spatial discretization
  - Finite volume (2-point flux default)
  - Structured and unstructured grids
- Time discretization: backward Euler
- Nonlinear solver
  - Newton-Raphson
  - Line search/damping with custom convergence criteria
- Linear solver: direct (LU) or iterative (BiCGStab)
- Multi-physics coupling
  - Flow and transport/reaction: sequential
  - Transport and reaction: global implicit
  - Geomechanics and flow/transport: sequential
  - Geophysics and flow/transport: sequential

Deep Borehole Waste Disposal

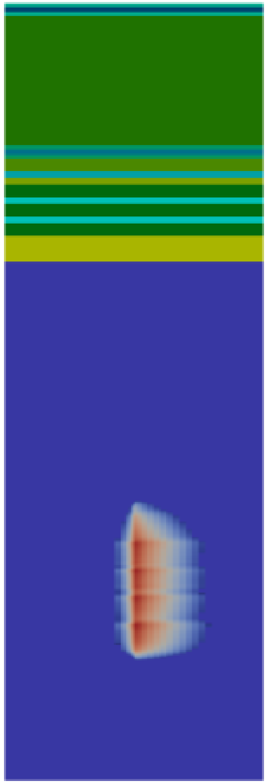

Emily Stein, SNL, 2015

PFLOTRAN Computing Capability

- High-Performance Computing (HPC)
  - Increasingly mechanistic process models
  - Highly-refined 3D discretizations
  - Massive probabilistic runs
- Open Source Collaboration
  - Leverages a diverse scientific community
  - Sharing among subject matter experts and stakeholders from labs/universities
- Modern Fortran (2003/2008)
  - Domain scientists remain engaged
  - Modular framework for customization
- Leverages Existing Capabilities
  - Meshing, visualization, HPC solvers, etc.
  - Configuration management, testing, and QA

Data Assimilation

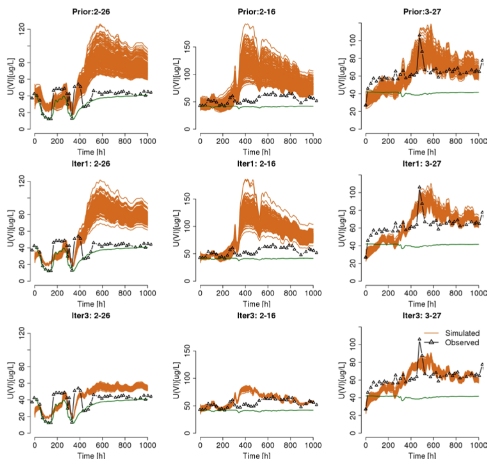

Xingyuan Chen, PNNL, 2011

PFLOTRAN Support Infrastructure

- **Mercurial:** distributed source control management tool
- **Bitbucket:** online PFLOTRAN repository
  - hg clone https://bitbucket.org/pflotran/pflotran-dev
  - Source tree
  - Commit logs
  - Wiki
    - Installation instructions
    - Quick guide
    - FAQ (entries motivated by questions on mailing list)
  - Pull requests
  - Issue tracker
- **Buildbot:** automated building and testing (regression and unit)
- **Google Groups:** pflotran-users and pflotran-dev mailing lists
- **Google Analytics:** tracks behavior on Bitbucket

Discrete Fracture Networks

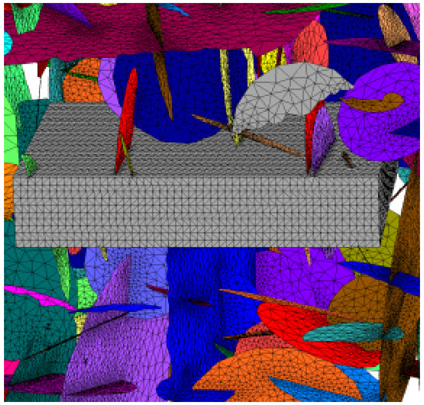

Satish Karra, LANL, 2015

PFLOTRAN Bitbucket Wiki

Sandia National Laboratories

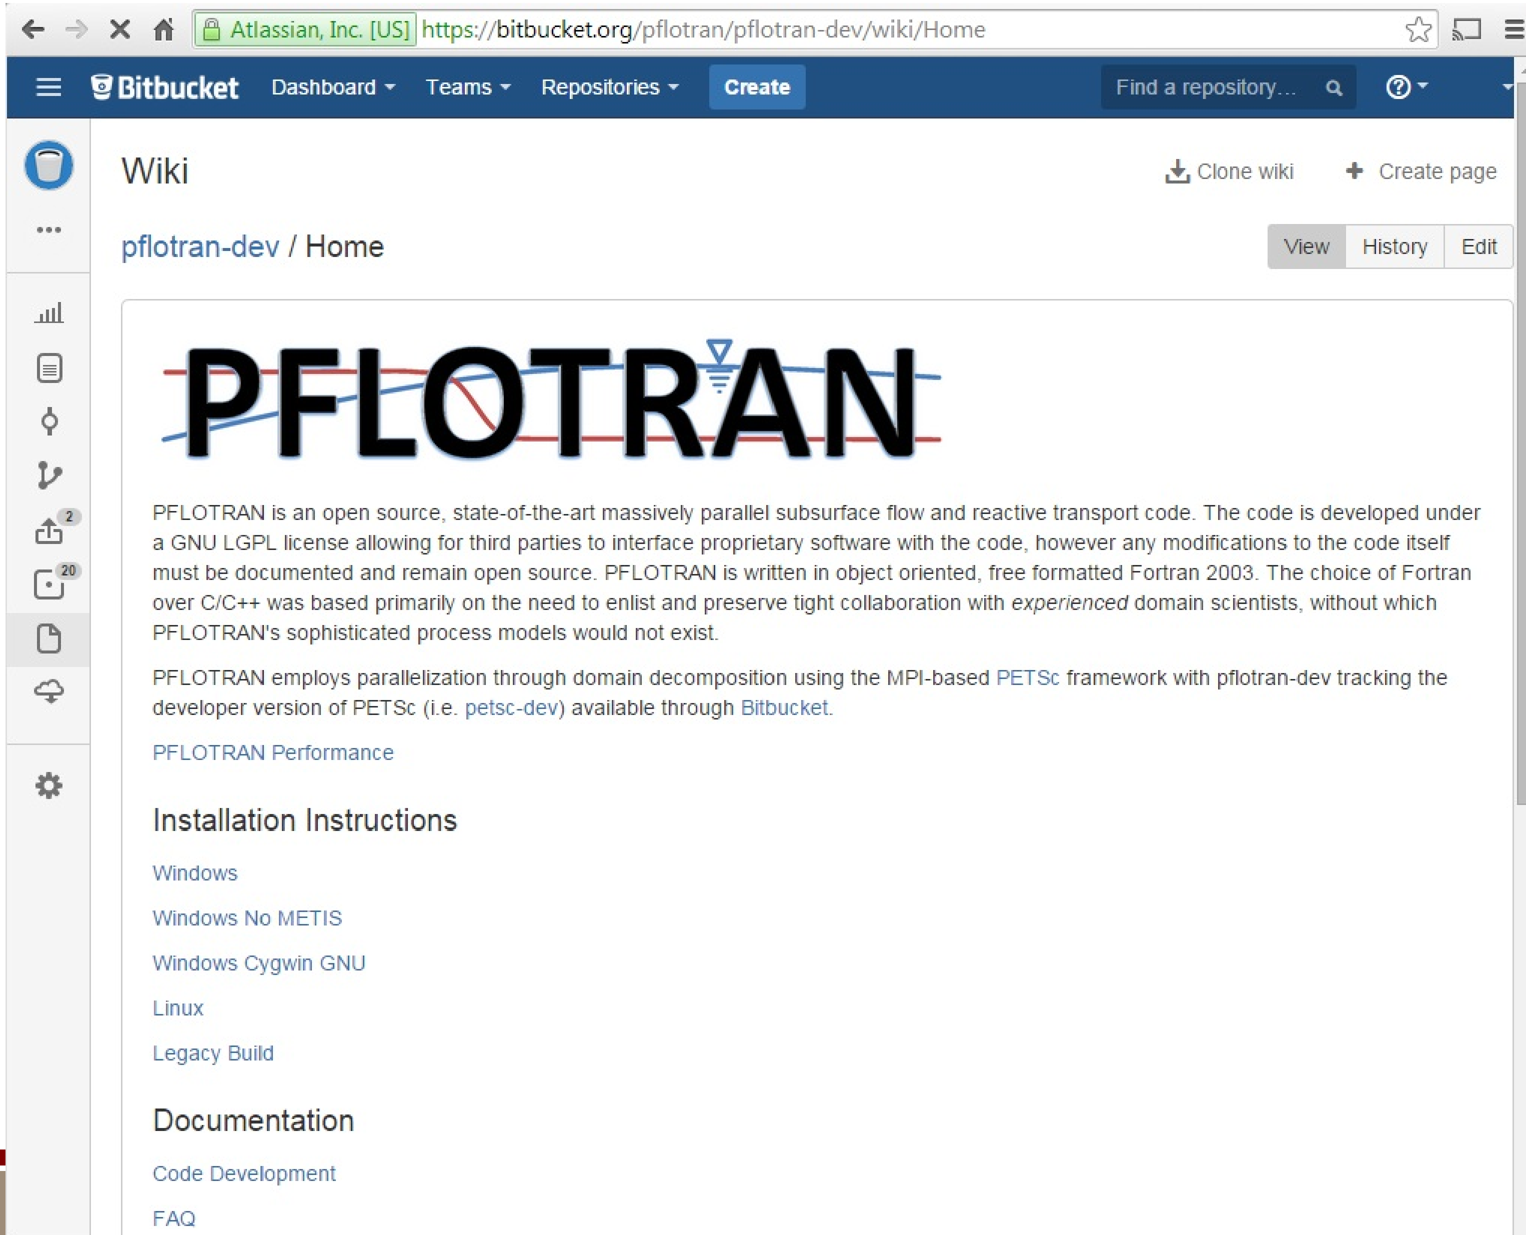
Atlassian, Inc. [US] https://bitbucket.org/pflotran/pflotran-dev/wiki/Home

Bitbucket Dashboard Teams Repositories Create Find a repository…

Wiki

Clone wiki Create page

pflotran-dev / Home

View History Edit

PFLOTRAN

PFLOTRAN is an open source, state-of-the-art massively parallel subsurface flow and reactive transport code. The code is developed under a GNU LGPL license allowing for third parties to interface proprietary software with the code, however any modifications to the code itself must be documented and remain open source. PFLOTRAN is written in object oriented, free formatted Fortran 2003. The choice of Fortran over C/C++ was based primarily on the need to enlist and preserve tight collaboration with *experienced* domain scientists, without which PFLOTRAN's sophisticated process models would not exist.

PFLOTRAN employs parallelization through domain decomposition using the MPI-based PETSc framework with pflotran-dev tracking the developer version of PETSc (i.e. petsc-dev) available through Bitbucket.

PFLOTRAN Performance

Installation Instructions

Windows

Windows No METIS

Windows Cygwin GNU

Linux

Legacy Build

Documentation

Code Development

FAQ

PFLOTRAN Bitbucket Commits

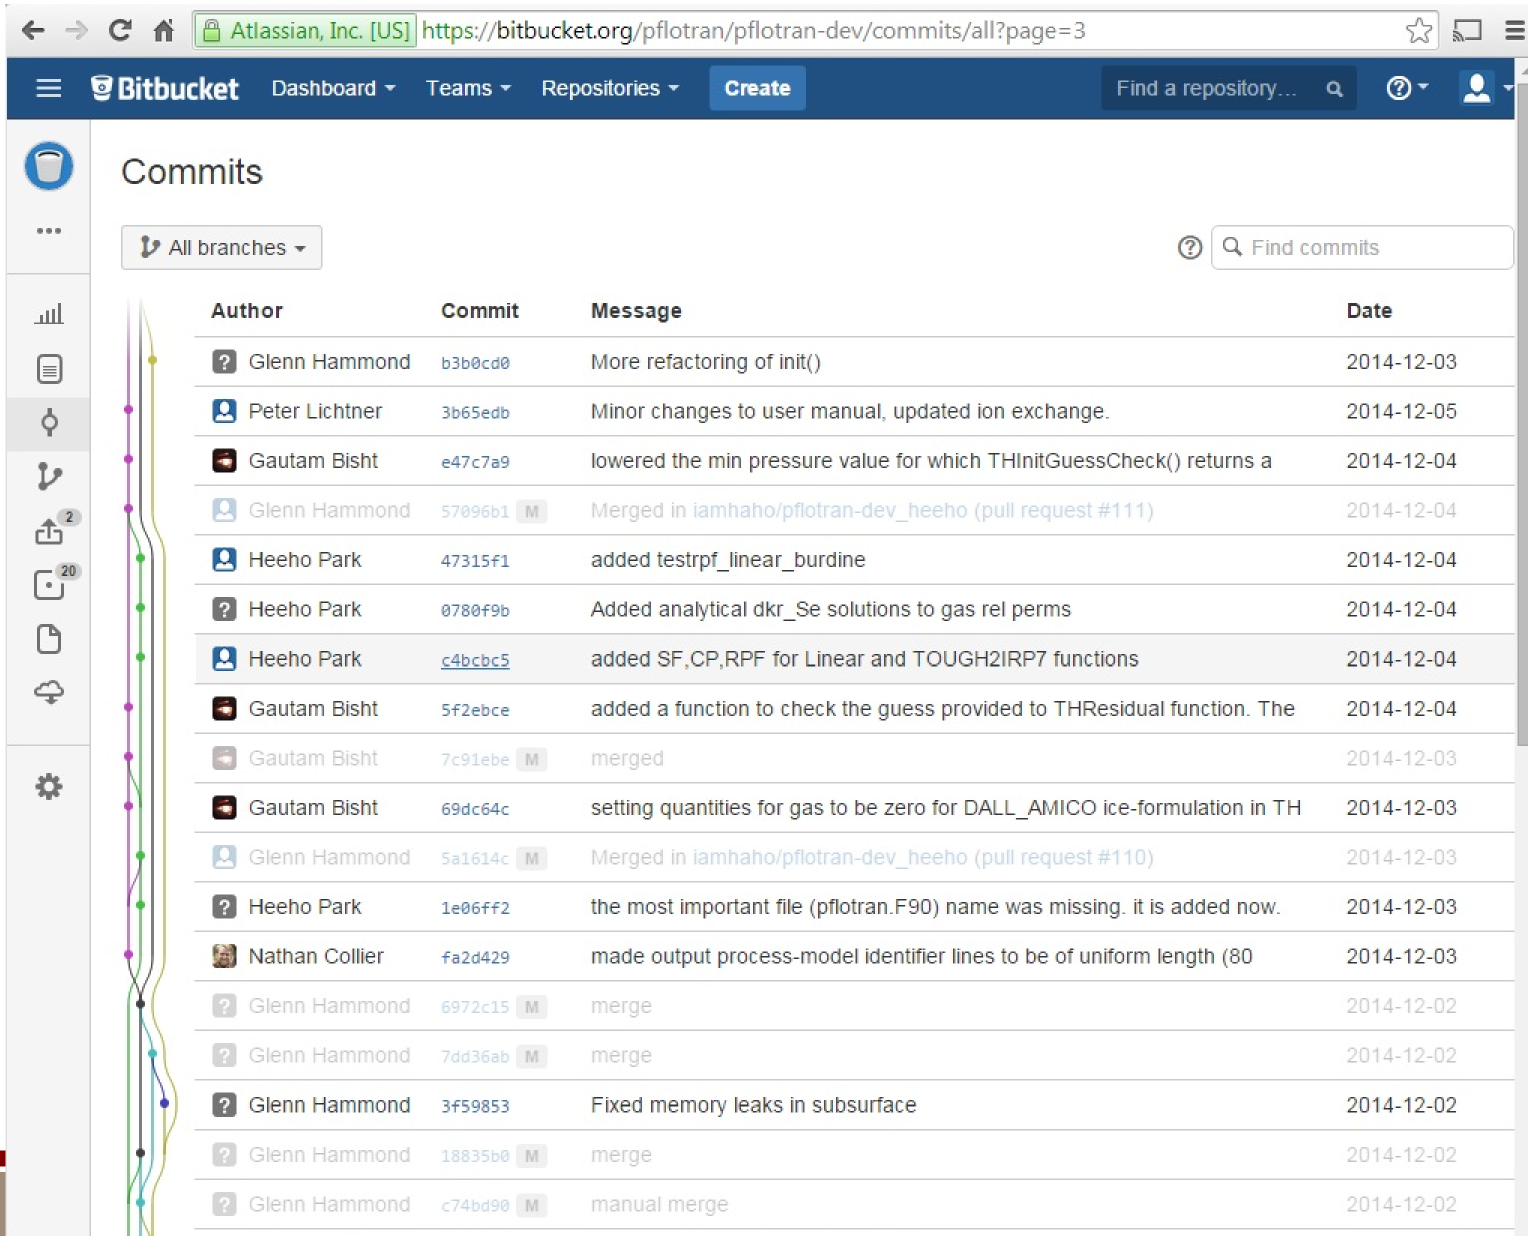

Author	Commit	Message	Date
Glenn Hammond	b3b0cd0	More refactoring of init()	2014-12-03
Peter Lichtner	3b65edb	Minor changes to user manual, updated ion exchange.	2014-12-05
Gautam Bisht	e47c7a9	lowered the min pressure value for which THInitGuessCheck() returns a	2014-12-04
Glenn Hammond	57096b1 M	Merged in iamhaho/pflotran-dev_heeho (pull request #111)	2014-12-04
Heeho Park	47315f1	added testrpf_linear_burdine	2014-12-04
Heeho Park	0780f9b	Added analytical dkr_Se solutions to gas rel perms	2014-12-04
Heeho Park	c4bcbc5	added SF,CP,RPF for Linear and TOUGH2IRP7 functions	2014-12-04
Gautam Bisht	5f2ebce	added a function to check the guess provided to THResidual function. The	2014-12-04
Gautam Bisht	7c91ebe M	merged	2014-12-03
Gautam Bisht	69dc64c	setting quantities for gas to be zero for DALL_AMICO ice-formulation in TH	2014-12-03
Glenn Hammond	5a1614c M	Merged in iamhaho/pflotran-dev_heeho (pull request #110)	2014-12-03
Heeho Park	1e06ff2	the most important file (pflotran.F90) name was missing. it is added now.	2014-12-03
Nathan Collier	fa2d429	made output process-model identifier lines to be of uniform length (80	2014-12-03
Glenn Hammond	6972c15 M	merge	2014-12-02
Glenn Hammond	7dd36ab M	merge	2014-12-02
Glenn Hammond	3f59853	Fixed memory leaks in subsurface	2014-12-02
Glenn Hammond	18835b0 M	merge	2014-12-02
Glenn Hammond	c74bd90 M	manual merge	2014-12-02

PFLOTRAN Buildbot

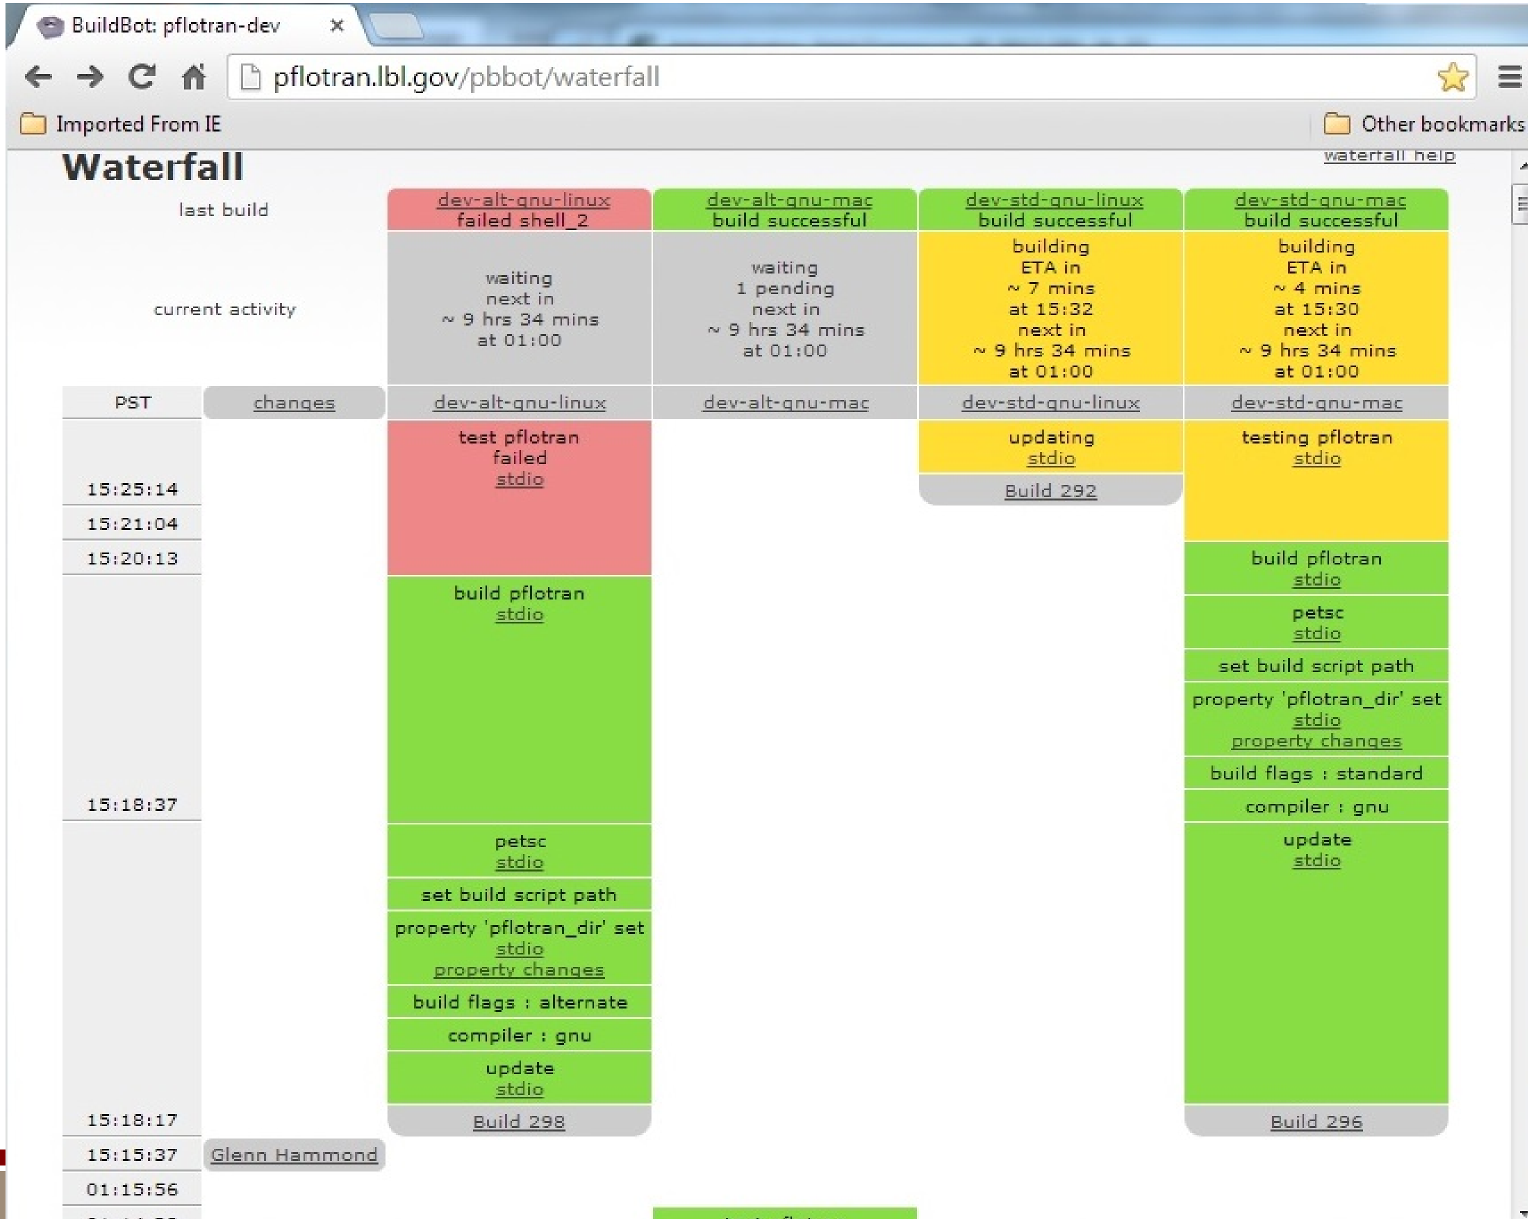

Google Analytics

Hits on PFLOTRAN Bitbucket site in 2015

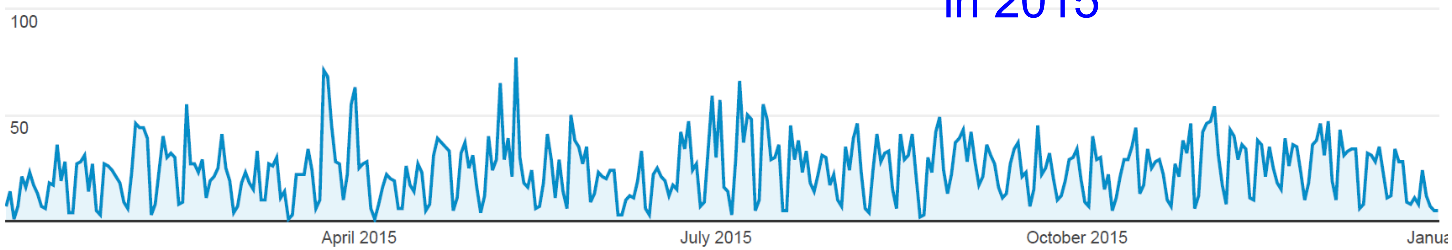

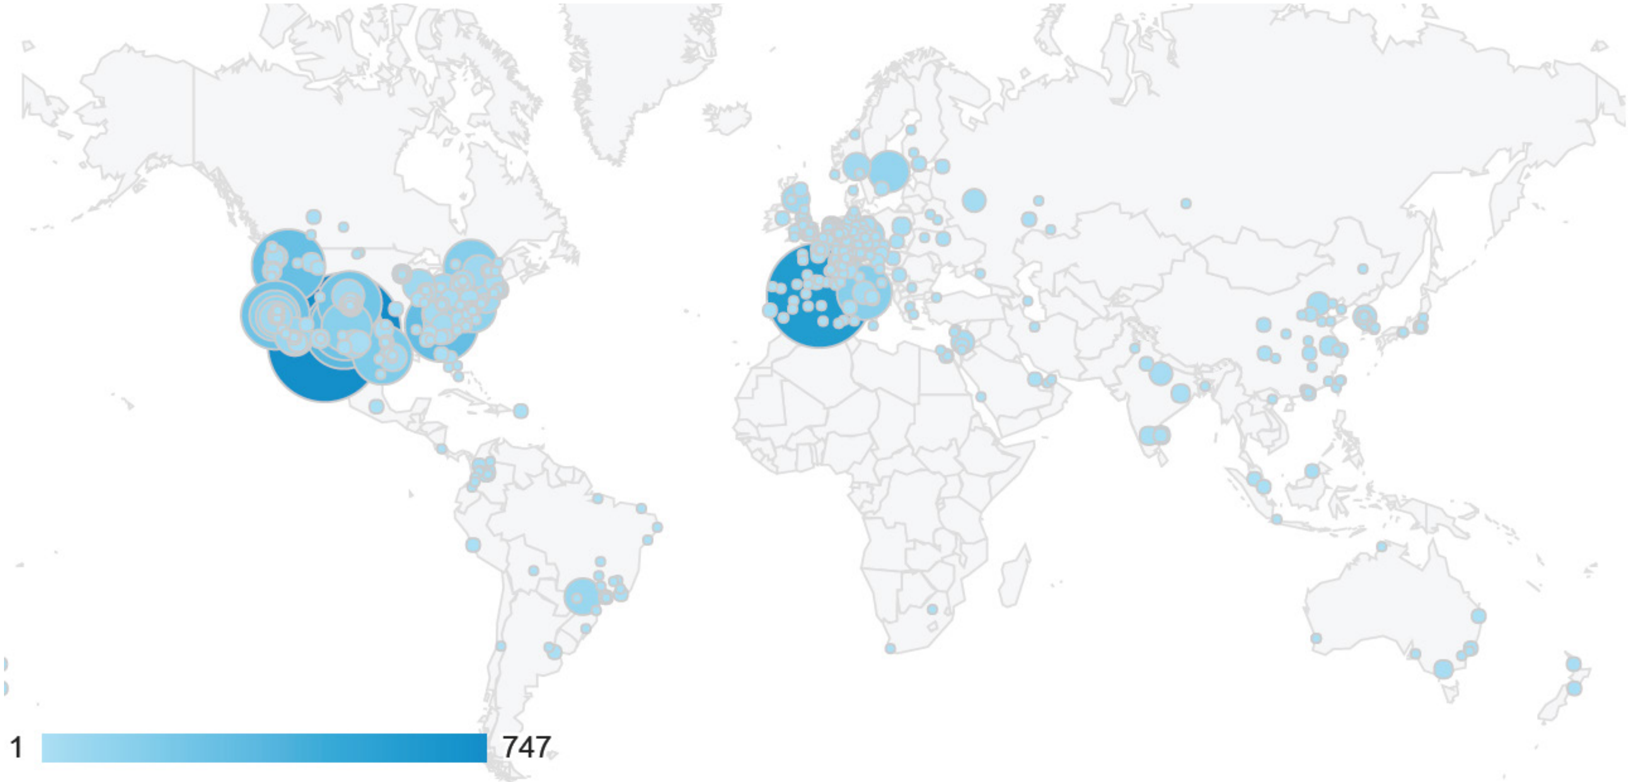

PFLOTRAN: Parallel Performance

- Scalability is good below ~10K processes
- Above ~10K cores, parallel performance degrades in the solvers and I/O
  - Parallel solvers
    - Breakdown in parallel preconditioners within Newton-Krylov solvers
    - Increasing communication cost (i.e. global reduction)
  - Parallel I/O
    - Bottlenecks in collective parallel I/O
- Future challenges
  - Transition to parallelism through accelerators (e.g. GPUs)

Solver Scalability

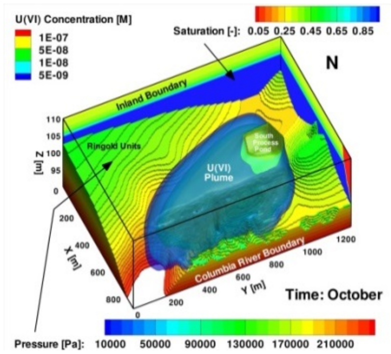

270 Million DOF Flow

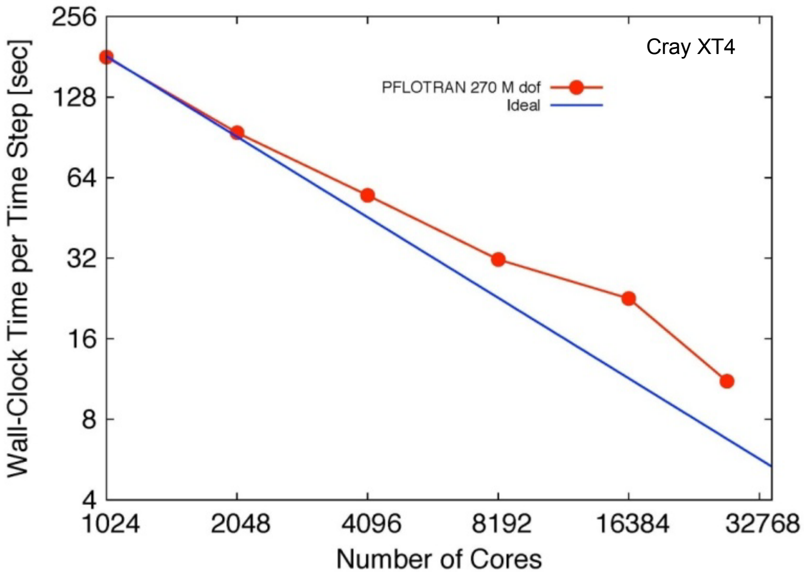

BCGStab w/ block Jacobi
R. Mills, 2008

1 Billion DOF Transport

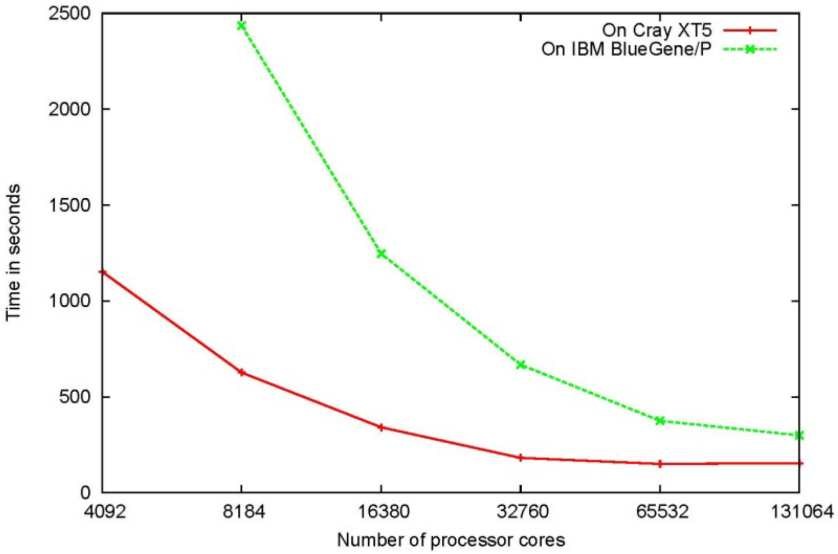

BCGStab w/ block Jacobi
Sripathi & Mahinthakumar, 2009

PFLOTRAN Performance on Jaguar (XT4)

270 Million Flow DOFs

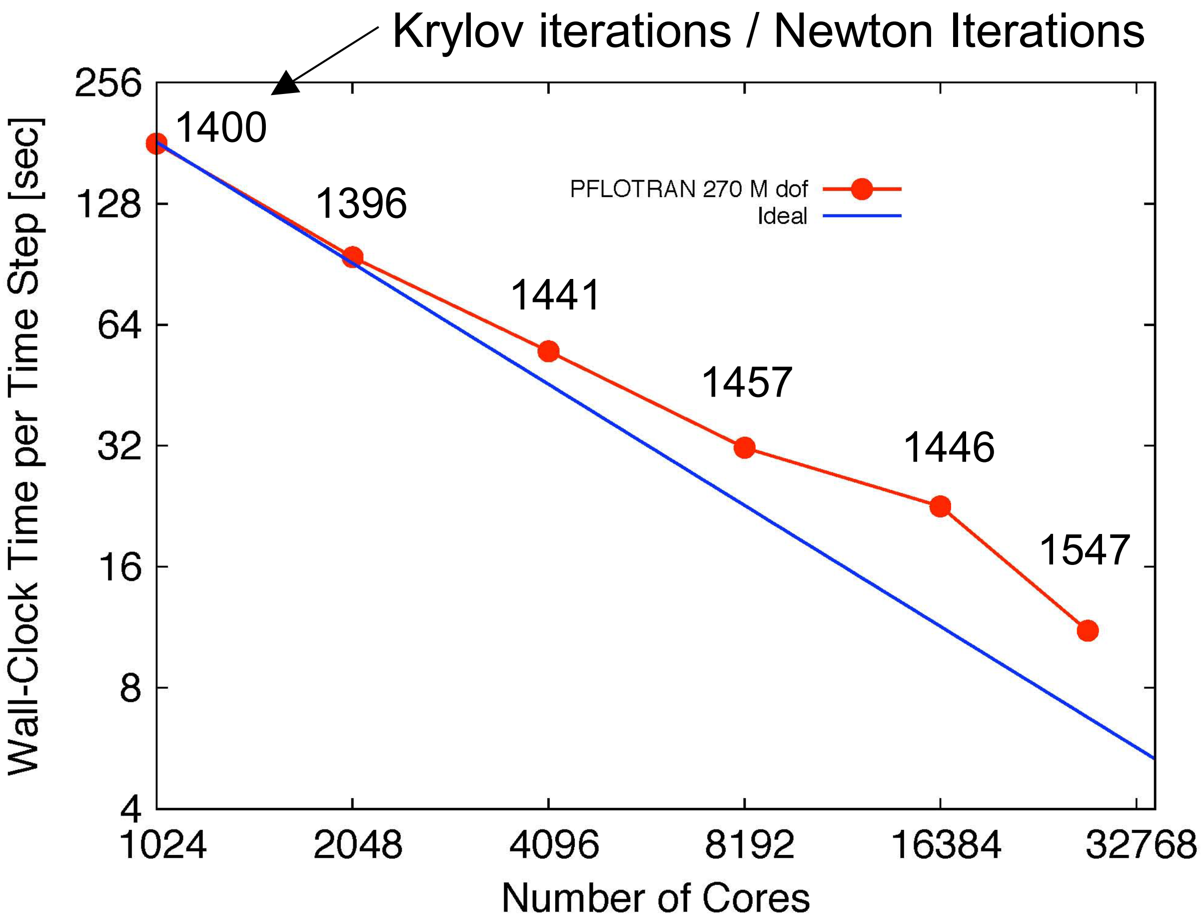

PFLOTRAN Performance on Jaguar (XT4)

Sandia National Laboratories

270 Million Flow DOFs

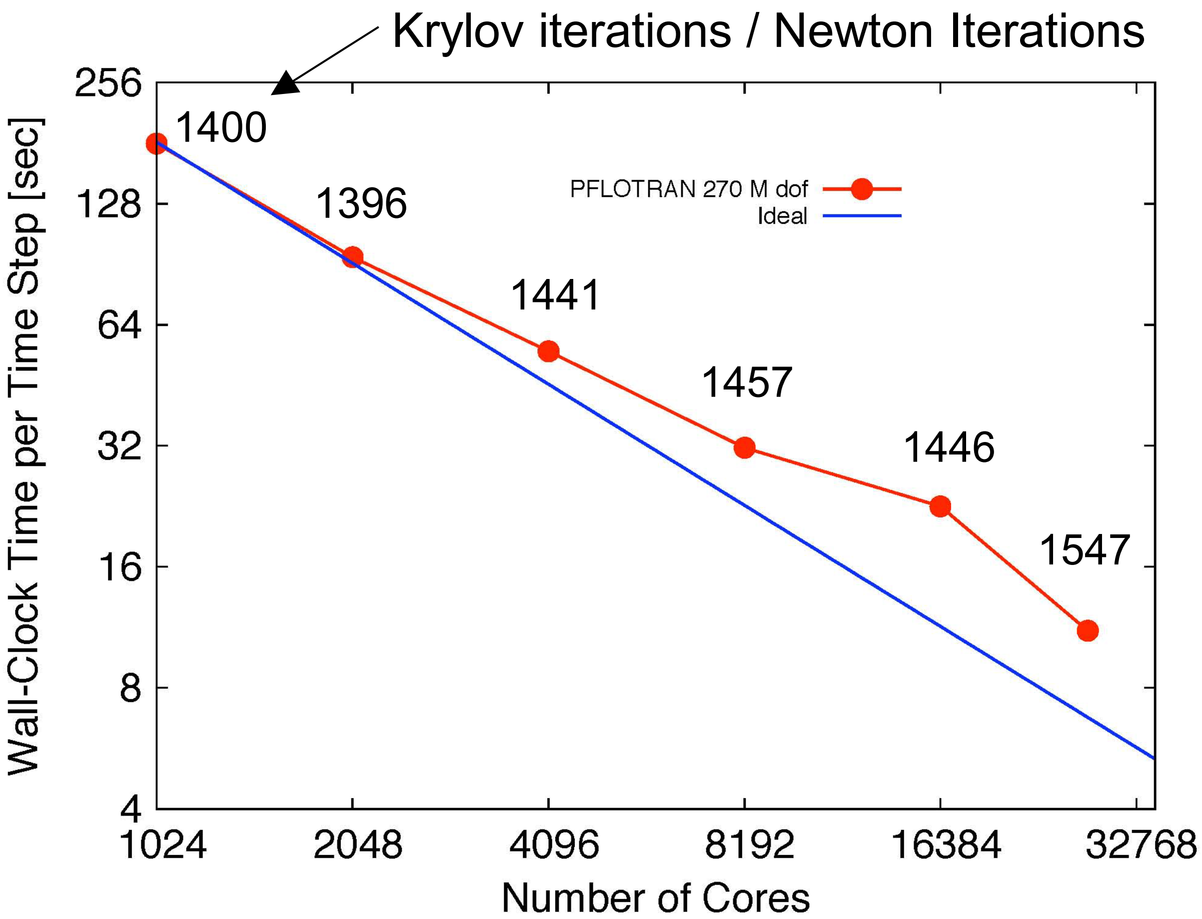

- Overall 60% efficient

PFLOTRAN Performance on Jaguar (XT4)

270 Million Flow DOFs

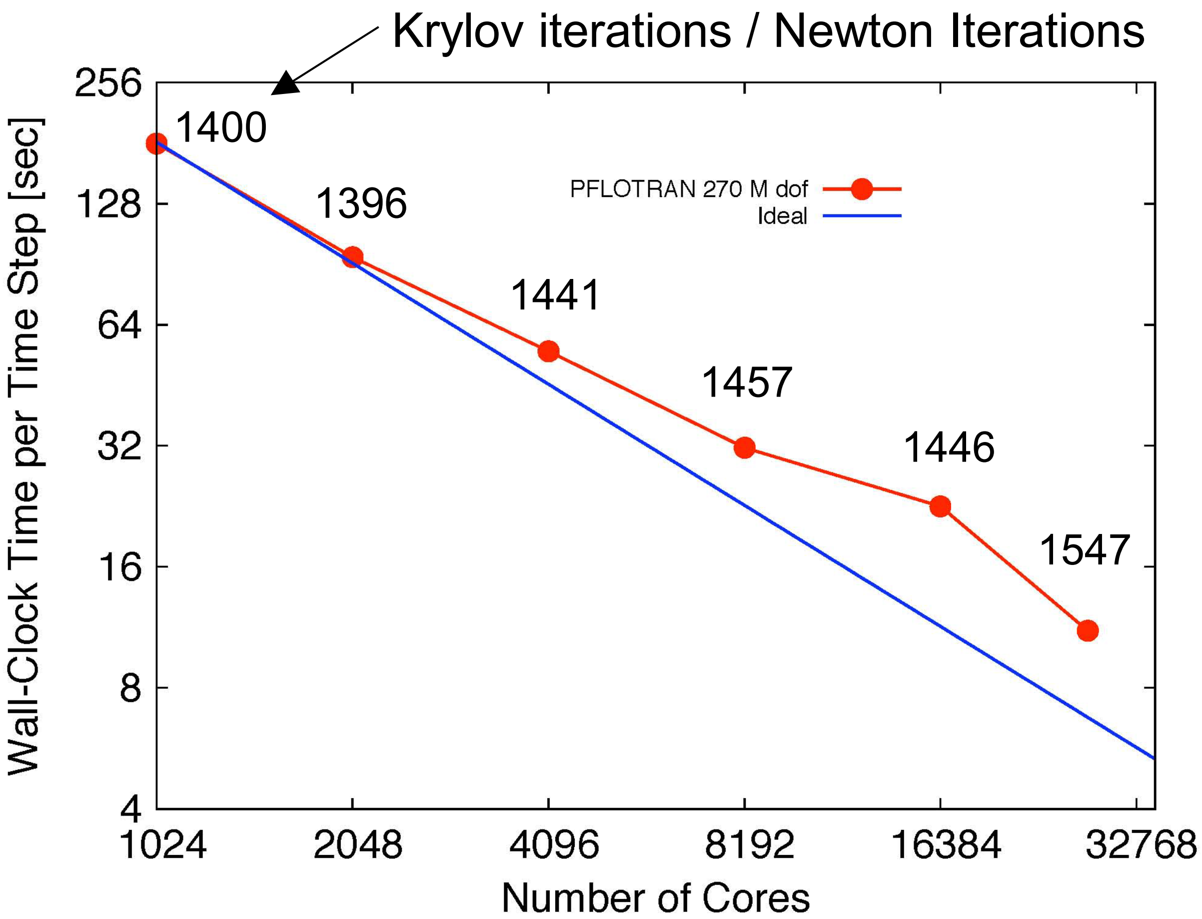

- Overall 60% efficient
- Science
  (1-2% of overall time)
  - 75% efficient
    (wall clock)

PFLOTRAN Performance on Jaguar (XT4)

Sandia National Laboratories

270 Million Flow DOFs

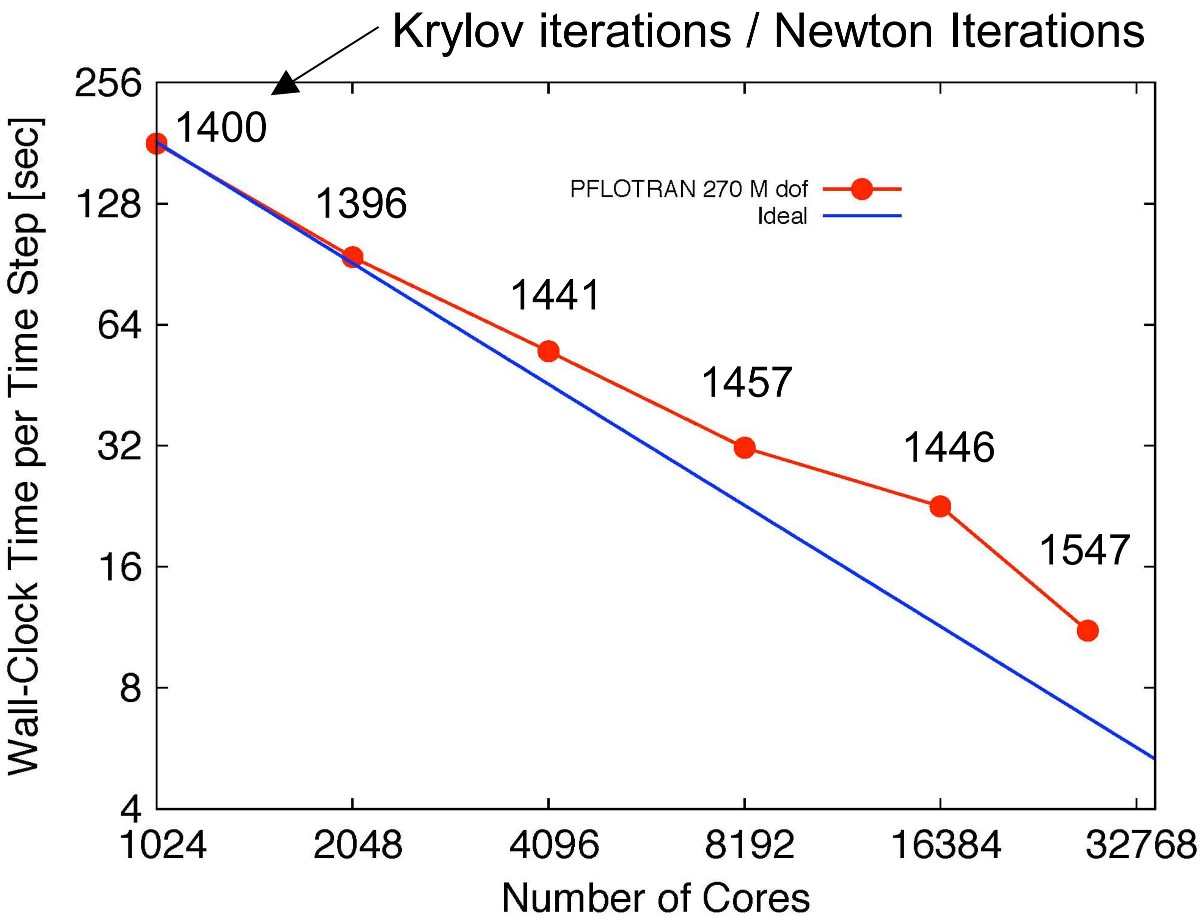

- Overall 60% efficient
- Science
  (1-2% of overall time)
  - 75% efficient (wall clock)
- Solvers
  (96% of overall time)
  - 90% efficient (iteration count)
  - 64% efficient (wall clock)

PFLOTRAN Performance on Jaguar

Processor Cores	4000	8000
Problem Size (dofs)	64.8 million	129.6 million
Problem Size / Processor Core	16200	16204
Krylov Iterations / Newton Iterations	418.4	462.6 (10.6% difference)
Nonlinear Solver Wall Clock Time (sec)	2531	2844 (12.4% difference)

PFLOTRAN Performance on Jaguar

Processor Cores	4000	8000
Problem Size (dofs)	64.8 million	129.6 million
Problem Size / Processor Core	16200	16204
Krylov Iterations / Newton Iterations	418.4	462.6 (10.6% difference)
Nonlinear Solver Wall Clock Time (sec)	2531	2844 (12.4% difference)

Breakdown in “weak” scalability

PFLOTRAN Performance on Jaguar

Processor Cores	4000	8000
Problem Size (dofs)	64.8 million	129.6 million
Problem Size / Processor Core	16200	16204
Krylov Iterations / Newton Iterations	418.4	462.6 (10.6% difference)
Nonlinear Solver Wall Clock Time (sec)	2531	2844 (12.4% difference)

Breakdown in "weak" scalability

Global reduction communication overhead

Geometric Multigrid (hypre PFMG)

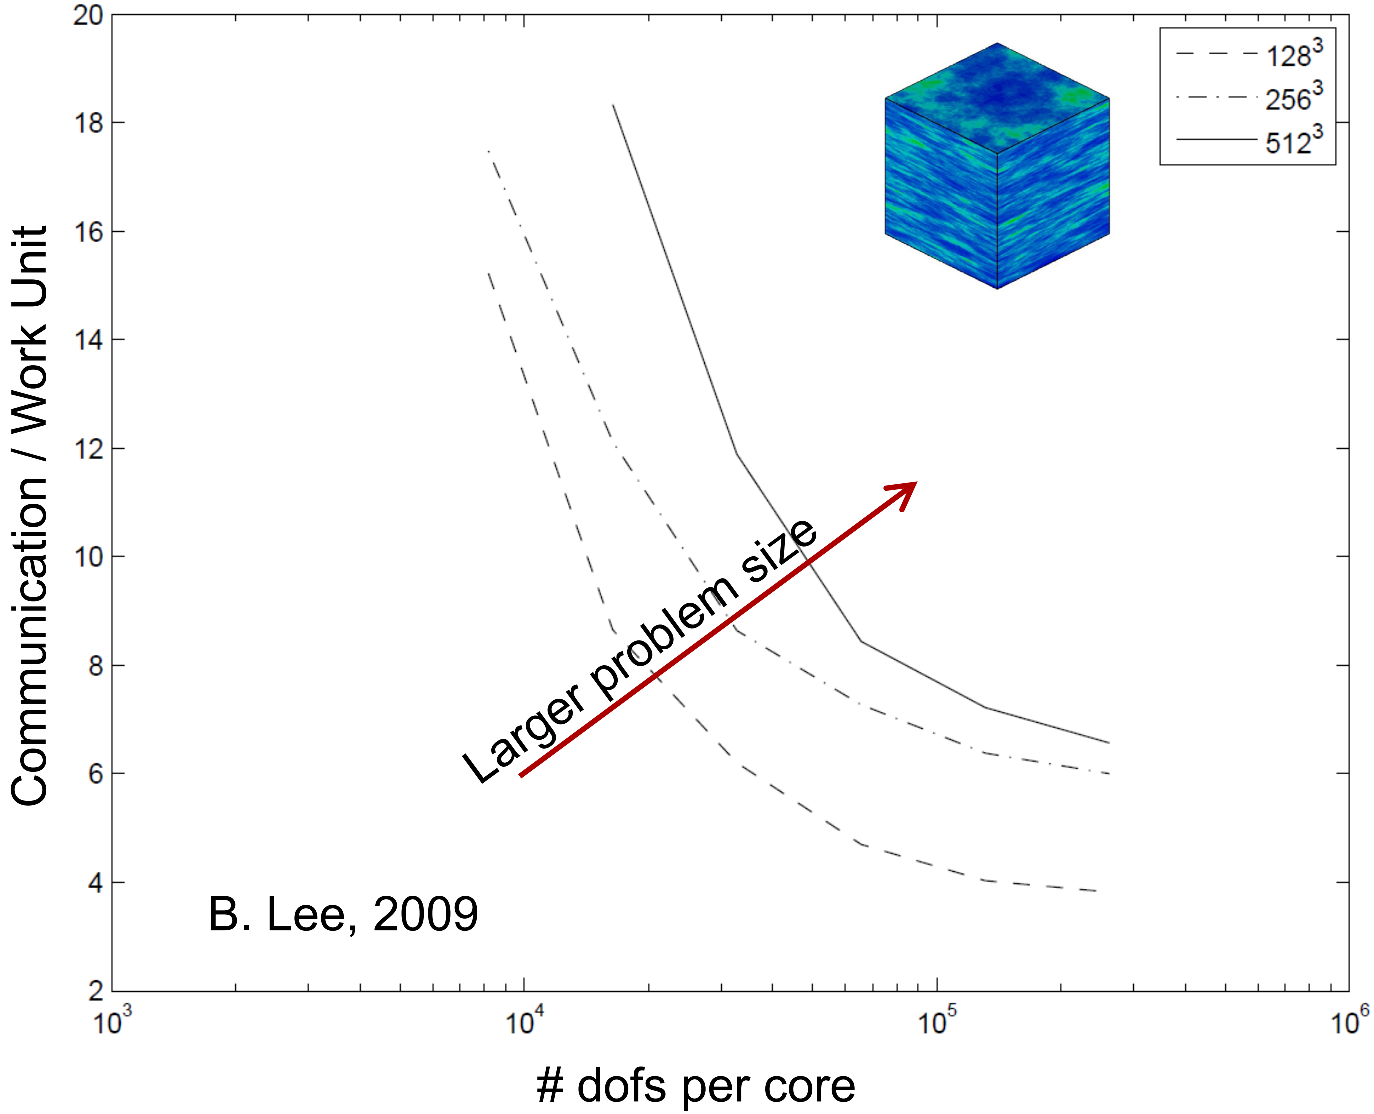

2M dofs
17M dofs
134M dofs

Geometric Multigrid (hypre PFMG)

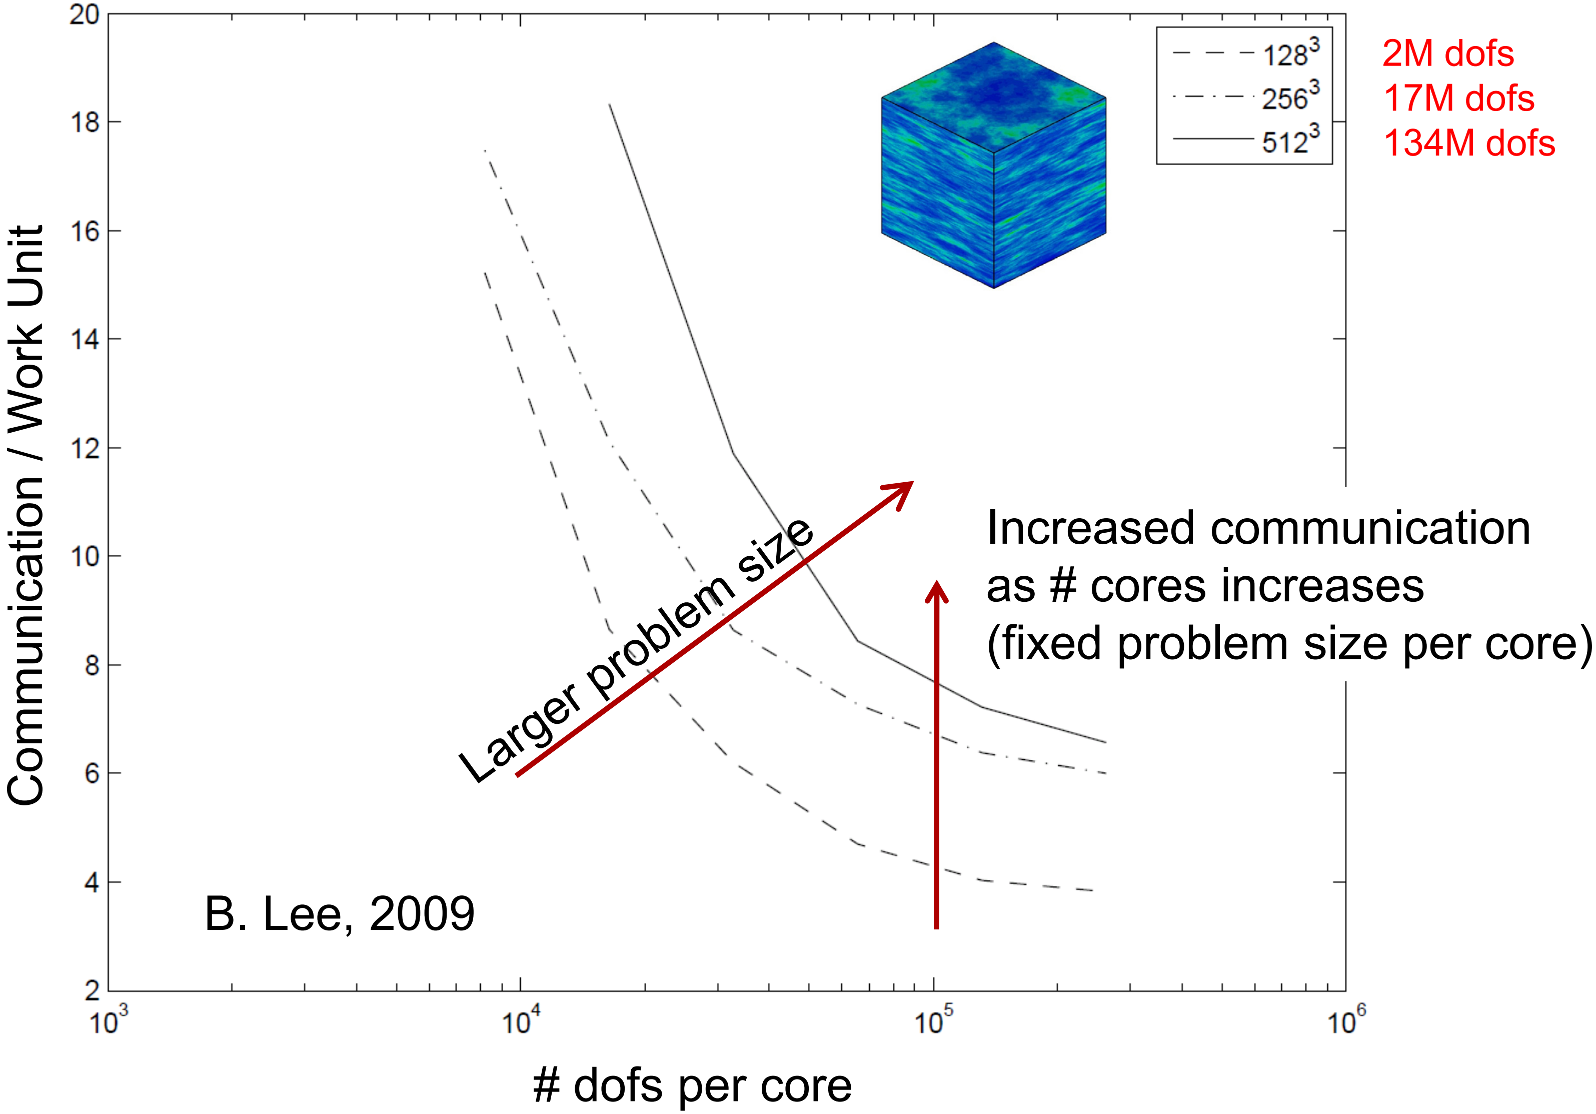

Geometric Multigrid (hypre PFMG)

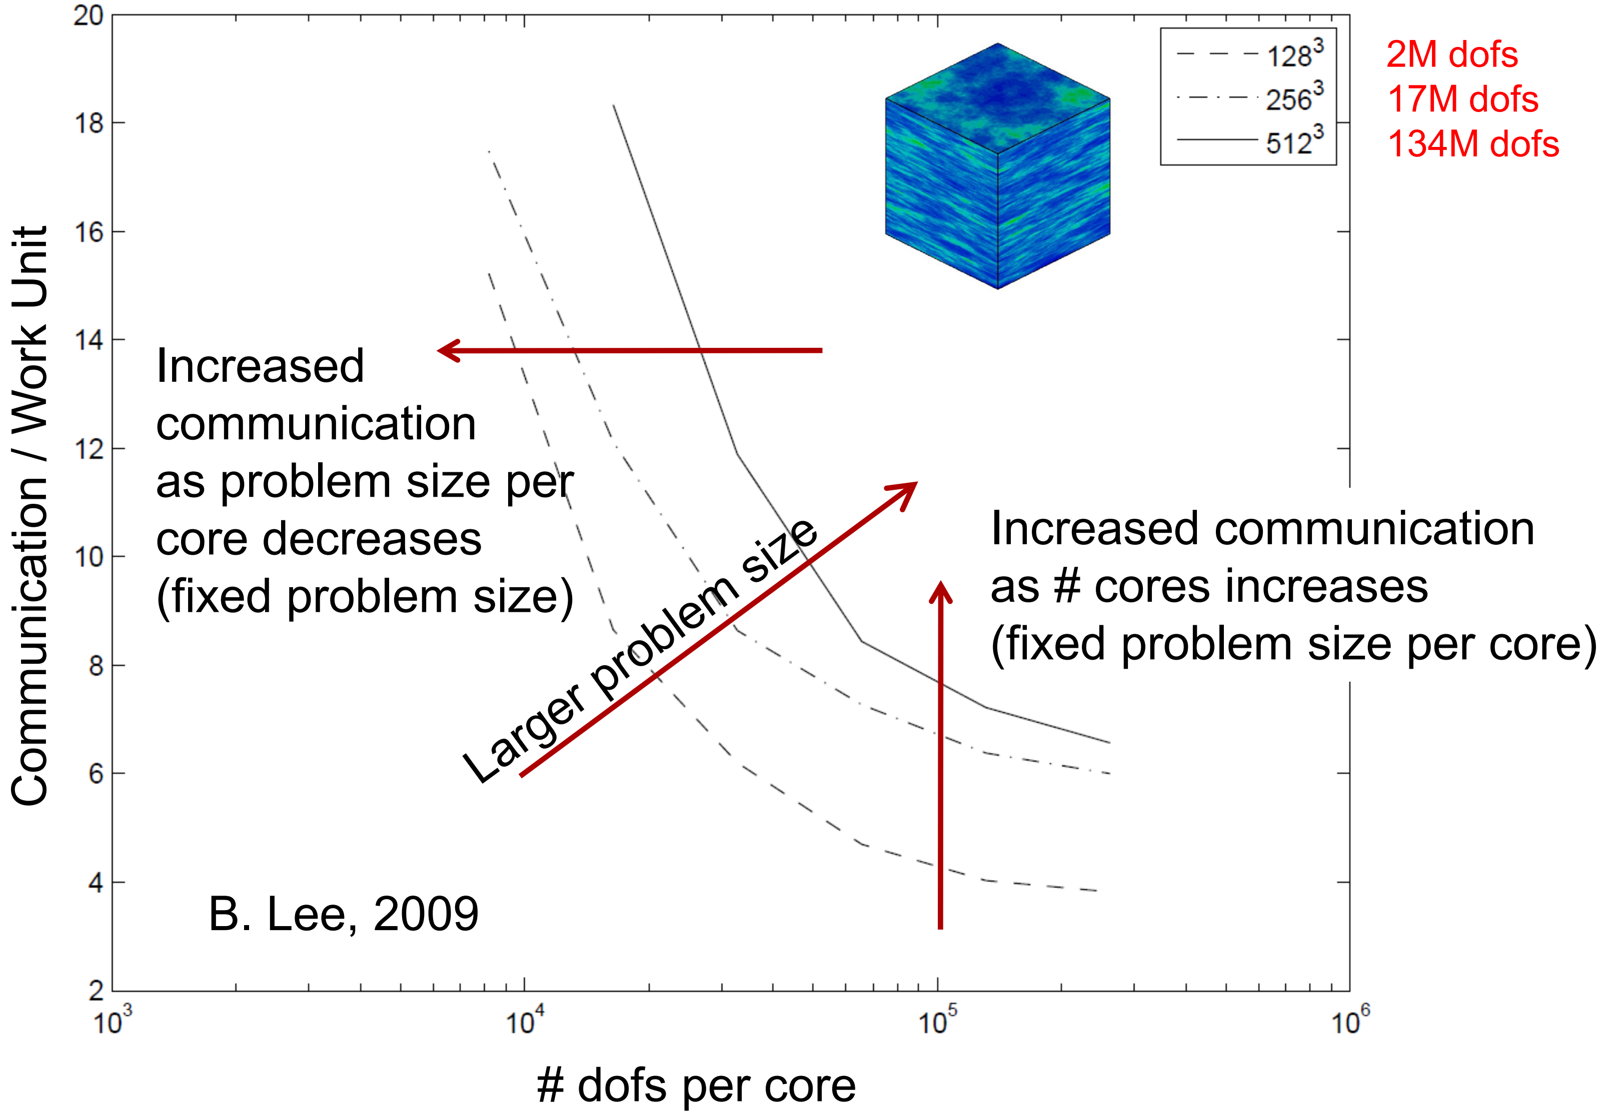

2M dofs
17M dofs
134M dofs

Scalable Parallel I/O through HDF5

Collective parallel I/O for jobs with
< ~10,000 cores

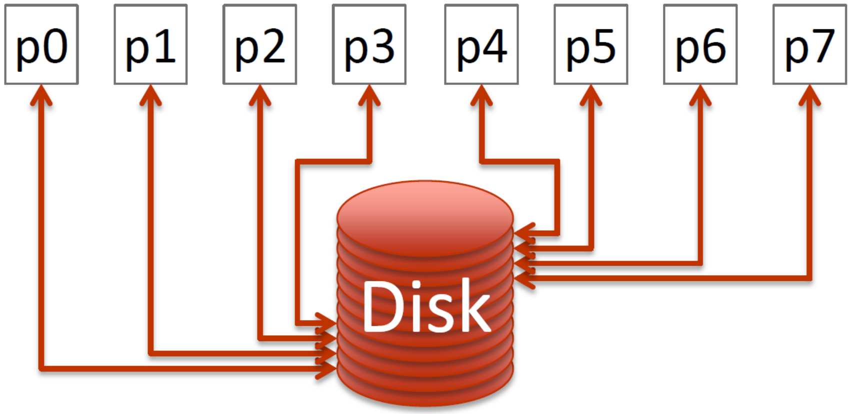

Two-phase parallel I/O for jobs with
> ~10,000 cores

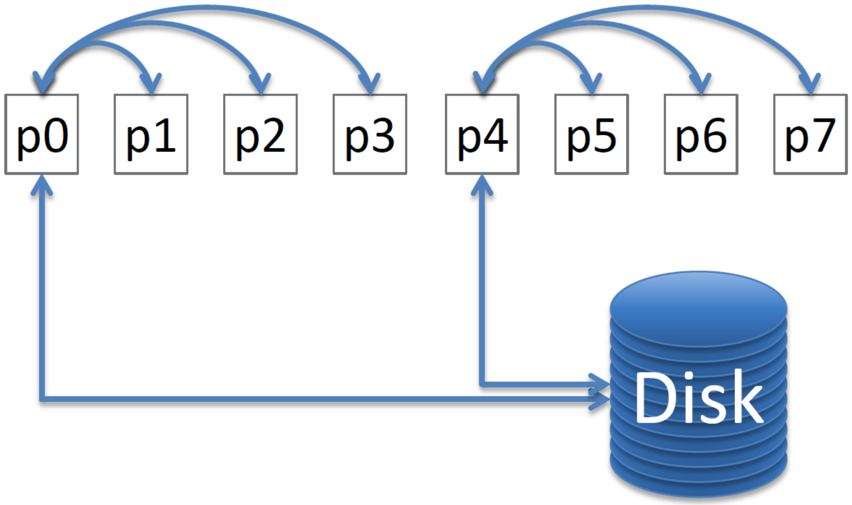

Scalable Parallel I/O through SCORPIO

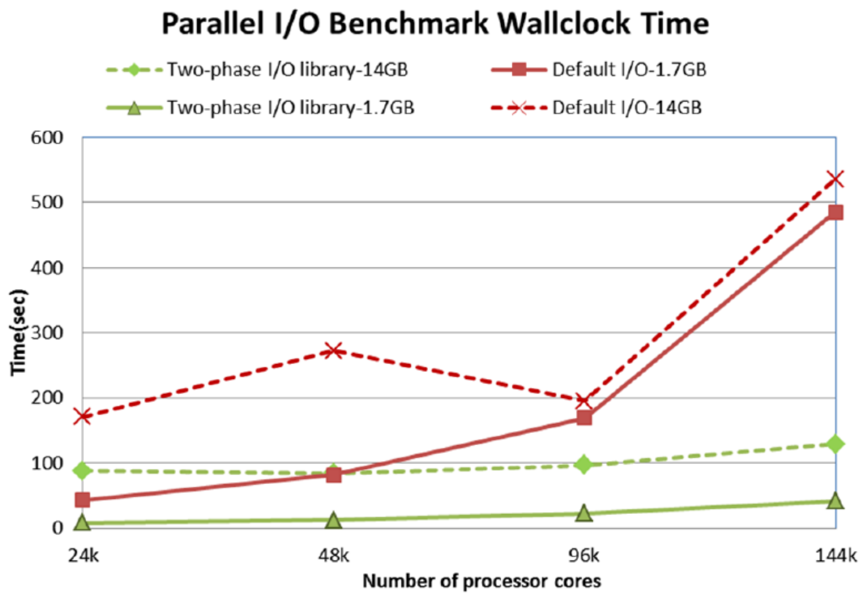

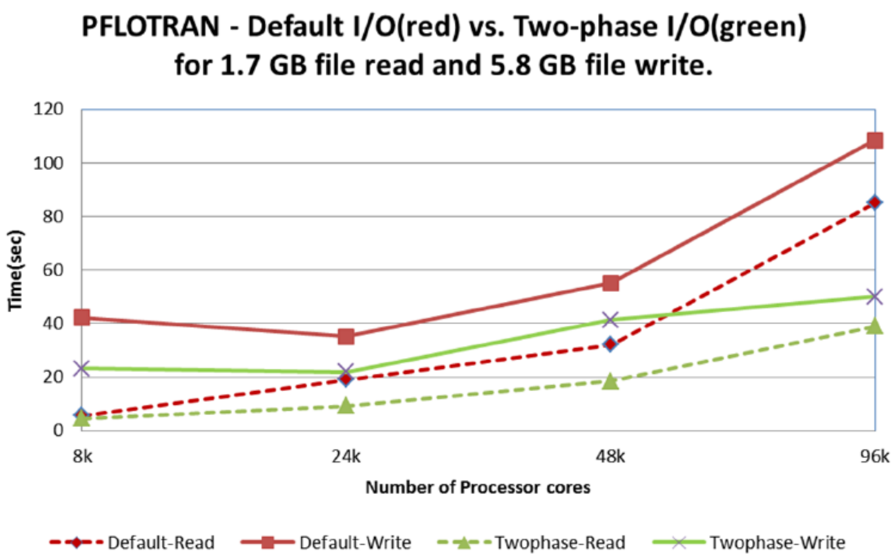

SCORPIO: A Scalable Two-Phase Parallel I/O Library With Application To A Large Scale Subsurface Simulator

Sreepathi et al., 2013

PFLOTRAN Performance: Accelerators

- In 2011 timeframe, PFLOTRAN was slated as one of six Tier 1 apps at the Oak Ridge Leadership Computing Facility to be ported from Jaguar to Titan for acceptance of Titan from Cray.
  - "The proposal is to do a one-time modification of Pflotran source code for the future. This would make the code high performing and flexible in supporting a variety of highly parallel processors and different accelerator programming models." - Cray
- The plan was to begin with point-wise kernels and then move on to parallelize the global solves across the grid.
- The chemical reaction kernel within PFLOTRAN was refactored to accommodate GPUs through OpenMP/CUDA directives instrumented by Cray/NVIDIA engineers.

PFLOTRAN Performance: Accelerators

Before:

```
  allocate(auxvar%total(reaction%naqcomp,option%nphase))
```

After:

```
  allocate(auxvar%total(chunk_size,num_threads, &
                        reaction%naqcomp,option%nphase))
do i = 1, ncomp
  icomp = reaction%eqcplxspecid(i,icplx)
  auxvar%total(ichunk,ithread,icomp,iphase) = &
    auxvar%total(ichunk,ithread,icomp,iphase) + &
    reaction%eqcplxstoich(i,icplx)* &
    auxvar%sec_molal(ichunk,ithread,icplx)
enddo
```

PFLOTRAN Performance: Accelerators

- Final Speedup
  - ~14× for an exaggerated biogeochemical system composed of ~30 chemical species.
  - <4× on a more realistic biogeochemical system (10-15 species)

PFLOTRAN Performance: Accelerators

- Final Speedup
  - ~14× for an exaggerated biogeochemical system composed of ~30 chemical species.
  - <4× on a more realistic biogeochemical system (10-15 species)

Amdahl's Law

$$T_p = s\,T_1 + \frac{(1-s)T_1}{p} \Rightarrow T_p = s\,T_1$$

T_p = parallel time
T_1 = serial time
p = # processes
s = sequential or non-GPU fraction

PFLOTRAN Performance: Accelerators

- Final Speedup
  - ~14× for an exaggerated biogeochemical system composed of ~30 chemical species.
  - <4× on a more realistic biogeochemical system (10-15 species)

Amdahl's Law

$$T_p = s\,T_1 + \frac{(1-s)T_1}{p} \Rightarrow T_p = s\,T_1$$

$s = 0.2$

$T_p = 0.2\ T_1$

Maximum Speedup = 5x

T_p = parallel time
T_1 = serial time
p = # processes
s = sequential or non-GPU fraction

Conclusion

- Presentation support provided by the IDEAS (Interoperable Design of Extreme-scale Application Software) Project funded by the United States Department of Energy Office of Science Advanced Scientific Computing Research and Biological and Environmental Research Programs

PFLOTRAN Multi-Physics Capability

- Flow
  - Single phase, variably-saturated
  - Multiphase gas-liquid
  - Interchangeable constitutive models and equations of state
- Energy
  - Thermal conduction and convection
- Multi-Component Transport
  - Advection
  - Hydrodynamic dispersion
- Chemical Reaction
  - Aqueous speciation
    - Ion activity models
  - Mineral precipitation-dissolution
  - Sorption
    - Isotherm-based
    - Ion exchange
    - Surface complexation
      - Equilibrium
      - Kinetic / multirate kinetic
  - Microbiological
    - Biomass
    - Inhibition
  - Radioactive decay with daughter products

Governing Equations

Multiphase Air-Water-Energy

Mass Conservation

$$\frac{\partial}{\partial t}\phi\left(s_l\rho_l X_a^l + s_g\rho_g X_a^g\right) = -\nabla\cdot\left(\rho_l X_a^l \boldsymbol{q}_l + \rho_g X_a^g \boldsymbol{q}_g + \boldsymbol{J}_a^l + \boldsymbol{J}_a^g\right) + q_a^G$$

$$\frac{\partial}{\partial t}\phi\left(s_l\rho_l X_w^l + s_g\rho_g X_w^g\right) = -\nabla\cdot\left(\rho_l X_w^l \boldsymbol{q}_l + \rho_g X_w^g \boldsymbol{q}_g + \boldsymbol{J}_w^l + \boldsymbol{J}_w^g\right) + q_w^G$$

Energy Conservation

$$\frac{\partial}{\partial t}\left(\phi\left[s_l\rho_l U_l + s_g\rho_g U_g\right] + (1-\phi)C_p^{\text{rock}}\rho_{\text{rock}}T\right) = -\nabla\cdot\left(\rho_l H_l \boldsymbol{q}_l + \rho_g H_g \boldsymbol{q}_g - \kappa_{\text{eff}}\nabla T\right) + q_e^G$$

Darcy Flux

$$\boldsymbol{q}_\alpha = -\frac{k k_\alpha}{\mu_\alpha}\nabla(p_\alpha - W_\alpha\rho_\alpha g z), \quad (\alpha = l,\ g)$$

Aqueous Diffusion

$$\boldsymbol{J}_a^l = -\tau\phi s_l D_l \rho_l \boldsymbol{\nabla} X_a^l$$

Gas Diffusion

$$\boldsymbol{J}_a^g = -\tau\phi s_g D_g^0 \left(\frac{T}{T_K}\right)^\theta \frac{p_0}{p_g}\rho_g\nabla X_a^g$$

Governing Equations

Biogeochemical Transport

Mass Conservation

$$\frac{\partial}{\partial t}\left(\phi s \Psi_j\right) + \nabla \cdot \left(\boldsymbol{q} - \phi s \boldsymbol{D} \boldsymbol{\nabla}\right) \Psi_j = Q_j - \sum_r \nu_{jr} I_r$$

$$\Psi_j = C_j + \sum_{i=1}^{N_X} \nu_{ji} K_i \prod_{j'=1}^{N_c} C_{j'}^{\nu_{j'i}}$$