TOPS Contributions to PFLOTRAN

Barry Smith Matthew Knepley

Mathematics and Computer Science Division
Argonne National Laboratory

TOPS Meeting at SIAM PP '08
Atlanta, Georgia
March 14, 2008

Outline

1. Performance
2. Future Work

PFLOTRAN Project

Modeling Multiscale-Multiphase-Multicomponent Subsurface Reactive Flows using Advanced Computing

LANL Peter Lichtner (PI), Chuan Lu, Bobby Philip, David Moulton

ORNL Richard Mills

ANL Barry Smith, Matthew Knepley

PNNL Glenn Hammond, Steve Yabusaki

U IL Al Valocchi, Kalyana Babu

Project goals:

- Develop a next-generation code (PFLOTRAN) for simulation of multiscale, multiphase, multicomponent flow and reactive transport in porous media.
- Apply it to field-scale studies of
    - Geologic CO2 sequestration,
    - Radionuclide migration at Hanford site, Nevada Test Site,
    - Others. . .

Hanford Problem

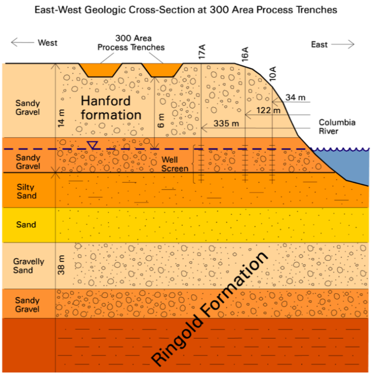

- U(VI) plumes continue to exceed drinking standards.
- Calculations predicted cleanup by natural attenuation years ago!
- Models with constant K_d do not account for slow release of U(VI)
- Kd approach implies behavior opposite to observations!

Hanford Problem

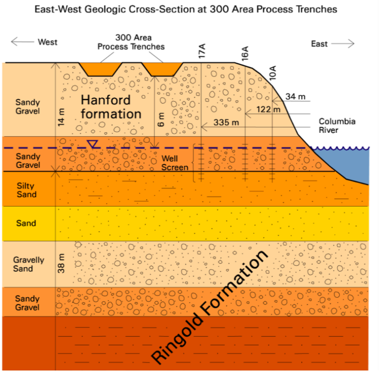

- 1 Gu
- DA: 2048 $\times$ 4096 $\times$ 128
- 1,073,741,824 grid points
- Runs on the Cray XT3

Jaguar: ORNL Cray XT3/4

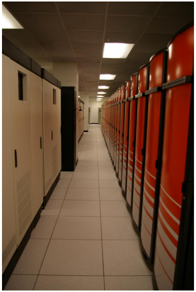 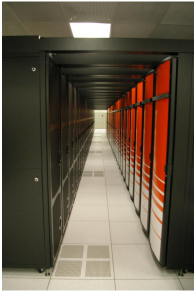

- Distributed memory, commodity processors
- Proprietary interconnect.
- 10^5 to 10^6 processors

Jaguar: ORNL Cray XT3/4

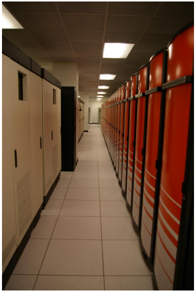

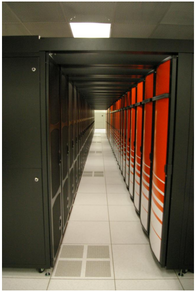

- 11706 dual-core 2.6 GHz Opteron nodes (23412 CPU cores)
- 119 Teraflops theoretical peak performance
- 46 Terabytes aggregate RAM
- 600 Terabytes of parallel disk storage
- 4.5 μs best-case network latency (vs. 35 μs for Gigabit Ethernet)

Convergence

Cores	Linear Its	Nonlinear Its
1024	1551	16
2048	1607	16
4096	1664	16

Table: Convergence of nonlinear and linear solvers for the Hanford problem.

Can also run well on 8192–16384 cores,
but the machine is currently being reconfigured

Scaling

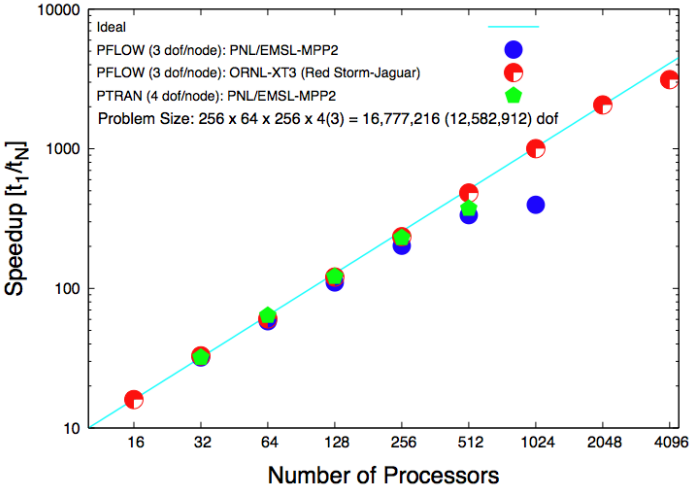
Ideal
PFLOW (3 dof/node): PNL/EMSL-MPP2
PFLOW (3 dof/node): ORNL-XT3 (Red Storm-Jaguar)
PTRAN (4 dof/node): PNL/EMSL-MPP2
Problem Size: 256 x 64 x 256 x 4(3) = 16,777,216 (12,582,912) dof
Speedup [t1/tN]
10000
1000
100
10
16
32
64
128
256
512
1024
2048
4096
Number of Processors

Scaling

Cores	SNESSolve Time (s)	SNESSolve MFlops/s	% of time reductions	% of time comm	% of time SNESSolve
1024	1124.6	85,719	28	3	96
2048	517.4	178,455	30	3	93
4096	267.6	364,540	26	2	82

Table: Performance of nonlinear solver for the Hanford problem.

- Speedup is 4.2[a]
  - Good implementation efficiency
- All operations are scaling
  - `VecAXPY()`, `VecDot()`, `VecNorm()`
  - `MatMult()`
  - `PCApply()`
  - `DAGlobalToLocal()`

[a] very good

Scaling

Cores	SNESSolve Time (s)	SNESSolve MFlops/s	% of time reductions	% of time comm	% of time SNESSolve
1024	1124.6	85,719	28	3	96
2048	517.4	178,455	30	3	93
4096	267.6	364,540	26	2	82

Table: Performance of nonlinear solver for the Hanford problem.

- Reductions on the Cray are not particularly good
  - Badness is not increasing with more cores
- Something outside of the SNESolve is not scaling so well
  - I/O: Printing during the run disrupts timing on the Cray
  - Inactive cells at river boundary and above ground surface (10-15%)

[a] very good

Outline

PETSc Additions

- Linear Solvers
  - BiCG with fewer reductions
  - Geometric Multigrid
  - “Physics-based” preconditioners
- Nonlinear solvers
  - Phase transitions: Possibility of non-smooth Jacobians
  - Better handling of inequality/complementarity constraints

Sieve Additions

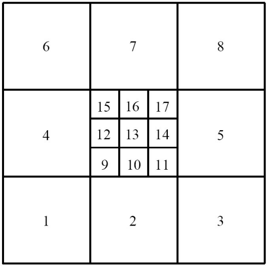

Unstructured mesh upgrade uses PETSc Sieves:

- Sieve allows more general mesh input:
  - Adjacency
  - Cell centers and volumes
  - Face centers and areas
- Sieve will mitigate load imbalance
- Automate distribution (partition at any depth)
  - Can partition faces (hypergraph w/ Zoltan)
- Arbitrary dimension and cell shape

Sieve Additions

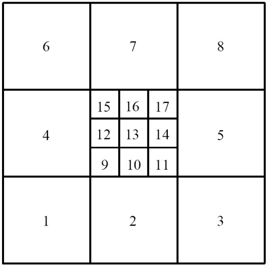

Unstructured mesh upgrade uses PETSc Sieves:

- Sieve allows more general mesh input:
  - Adjacency
  - Cell centers and volumes
  - Face centers and areas
- Sieve will mitigate load imbalance
- Automate distribution (partition at any depth)
  - Can partition faces (hypergraph w/ Zoltan)
- Arbitrary dimension and cell shape

Sieve Additions

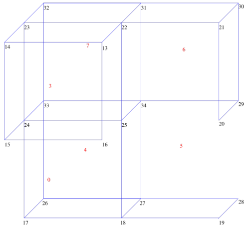

Process 0

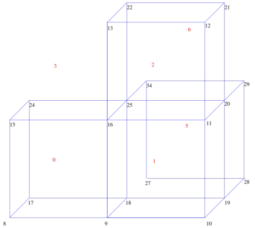

Process 1

Unstructured mesh upgrade uses PETSc Sieves:

- Sieve allows more general mesh input:
  - Adjacency
  - Cell centers and volumes
  - Face centers and areas
- Sieve will mitigate load imbalance
- Automate distribution (partition at any depth)
  - Can partition faces (hypergraph w/ Zoltan)
- Arbitrary dimension and cell shape

Sieve Additions

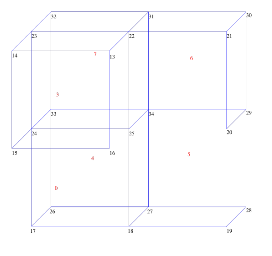

Process 0

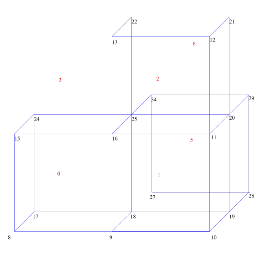

Process 1

Unstructured mesh upgrade uses PETSc Sieves:

- Sieve allows more general mesh input:
  - Adjacency
  - Cell centers and volumes
  - Face centers and areas
- Sieve will mitigate load imbalance
- Automate distribution (partition at any depth)
  - Can partition faces (hypergraph w/ Zoltan)
- Arbitrary dimension and cell shape

Barry's understanding of PFLOTRAN

PetscInitialize()

```
do some stuff (read files, ...)
loop over timesteps
    maybe do something
```

SNESSolve()

```
    maybe do something
maybe do something
```

PetscFinalize()

Barry's understanding of PFLOTRAN

- On 1024 cores, 4% of the time is spent doing something
- On 4096 cores, 18% of the time is spent doing something
- If this is load imbalance, why doesn't `SNESSolve()` suffer from it?
- Need more logging

Cores	Total (s)	SNESSolve (s)	do something (s)
1024	1170.9	1124.6	46.3
4096	324.8	267.6	57.1
	setup indices (s)	read materials (s)	unknown (s)
	1.4	16	28.9
	5.8	15	36.3

Table: Performance of nonlinear solver for the Hanford problem.