Massively Parallel Adaptive Simulations Using PETSc for Turbulent Boundary Layer Flows

Michel Rasquin, Argonne National Labs, Dan A. Ibanez, Rensselaer Polytechnic Institute, Benjamin Matthews, University of Colorado Boulder, Cameron Smith, Onkar Sahni, and Mark S. Shephard, Rensselaer Polytechnic Institute, Jed Brown, Kenneth Jansen, University of Colorado Boulder

A stabilized finite element flow solver (PHASTA) has been shown to scale to 3.1 million processes with it native solver but effort is underway to incorporate FASTMath tools for adaptivity (MeshAdapt), partition improvement (ParMa and Chef) and linear equation solution (PETSc). In this study, performance of standard PETSc and an improved versions that caches the communication of global assembly are compared.

PHASTA Flow Solver

- Massively parallel Navier Stokes flow solver
- Models compressible or incompressible turbulent, unsteady flows.
- 3D finite element discretization in space => Complex geometries
- Fully implicit in time => Δt governed by the physics, not by the numerical scheme
- Mesh adaptivity =>
  - Grid matches physical scale
  - Anisotropic and transient
- Used to model air flow over airplane wing, two phase flow applied to nuclear reactor, blood flow, etc.

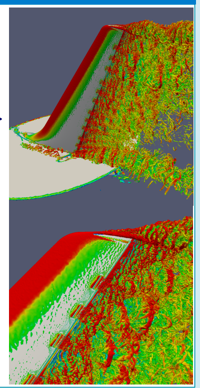

Native Solver Parallelization

Parallel strategy:

- Domain decomposition approach based on the elements
- Both compute stages operate off the same mesh partition
- Partition defines inter-part relations (part-to-part comm.)

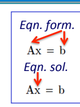

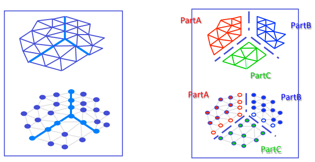

Locally, *incomplete* values (in ***b***, ***A***, ***q***, etc.) for shared *dofs*.

Apply communications to *complete* values/entries (in ***b***, ***q*** only NOT full A)

(a) during *Eqn. form.*

(b) during *Eqn. sol.*

Native Solver Scaling on Mira

92B element mesh (with resp. 8B and 64B dofs)

- Strong scaling: from 256K to 768K cores (Mira)
- 1, 2 and 4 MPI processes per core
- 3,145,728 MPI processes on full Mira system
- Scaling factor defined as (tbase × ncorebase)/(t × ncore):
  - For 1 MPI process per core: scaling factor 1 = perfect scaling
  - For 2 or 4 MPI processes per core: scaling factor means acceleration wrt. MPI process per core (Base 256k cores 1MPI)

PHASTA Native Solver Scaling on Mira

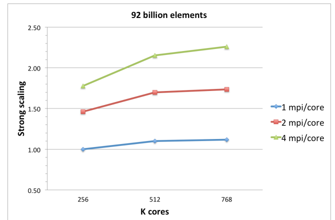

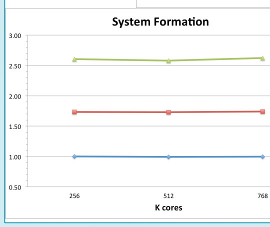

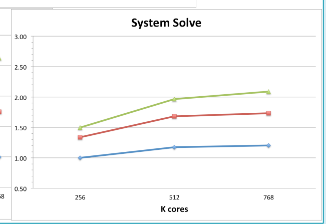

Comparing Native and PETSc Solvers

- Unsteady Navier-Stokes equations are highly non-linear.
- Implicit time integration requires reducing the non-linear PDE residual below a given threshold before accepting the current time step.
- Typically, this is accomplished by multiple linearizations and linear equation solves on each time step.
- Non-linear threshold depends to some degree on the problem being solved.
- Typically, 3-4 orders of magnitude reduction in non-linear residual yields solution statistics that are independent on solver threshold.
- Two solving levels where considered in this study
  - "Loose" $\log(R/R_0)=-32$ or 32 dB non-linear residual reduction
  - "Tight" $\log(R/R_0)=-45$ or 45 dB non-linear residual reduction
- Transonic flow over a wing, unstructured grid shock and boundary layer adaptation followed by uniform refinement (MeshAdapt)=>0.5 Billion element mesh.

Equation Solution Scaling

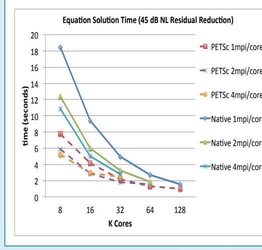

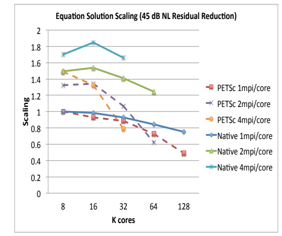

Global Assembly Improvement

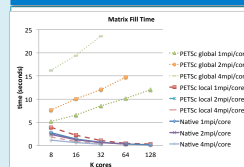

Standard Global Assembly

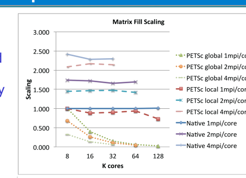

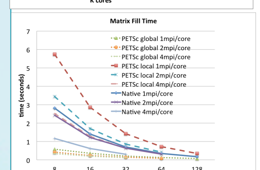

New, cached Global Assembly

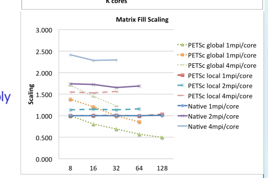

Summary and Future Plans

- New, cached Global Matrix Assembly has removed scaling bottleneck.
- Though still poorest scaling component, Global Assembly is now far smaller than local matrix assembly and KSP solve
- Setup time is roughly 3x subsequent use but still faster than standard way.
- Future work:
  - Extend study to higher core counts: Conversion to 64 bit PETSc completed
  - Evaluate other preconditioners available in PETSc
  - Coordinate with adaptive meshes.

More Information: http://www.fastmath-scidac.org or contact Lori Diachin, LLNL, diachin2@llnl.gov, 925-422-7130

SciDAC Scientific Discovery through Advanced Computing · Argonne · Berkeley Lab · Lawrence Livermore National Laboratory · Sandia National Laboratories · Rensselaer · SMU · Stony Brook University · Berkeley · The University of British Columbia · Wisconsin