On Scaling Modes and Balancing Stochastic, Discretization, and Modeling Error T33E-2972

University of Colorado Boulder

Jed Brown jed@jedbrown.org CU Boulder & ANL

Argonne NATIONAL LABORATORY

Why compute?

Decision What to build? How to build it? Which experiment to perform? What to study next? Decisions usually require a finite number of quantities even though the simulation approximates functions.

When Many decisions are time-sensitive. Today's weather forecast is of no use if issued tomorrow.

Cost What is the cost of making the decision without access to simulation? Actual cost must not exceed this.

Relevant dimensions

Problem complexity Complexity of the physics. Sensitivity of the functionals of interest. Treatment of stochasticity.

Accuracy In a norm? Which quantities of interest?

Time Per problem instance. For the first instance. Compute time versus human time.

Cost Incremental cost. Human cost. Opportunity cost. Subsidized?

- Terms relevant to scientist/engineer
- Compute meaningful quantities – needed to make a decision or obtain a result of scientific value—not one iteration/time step
- No flop/s, number of elements/time steps

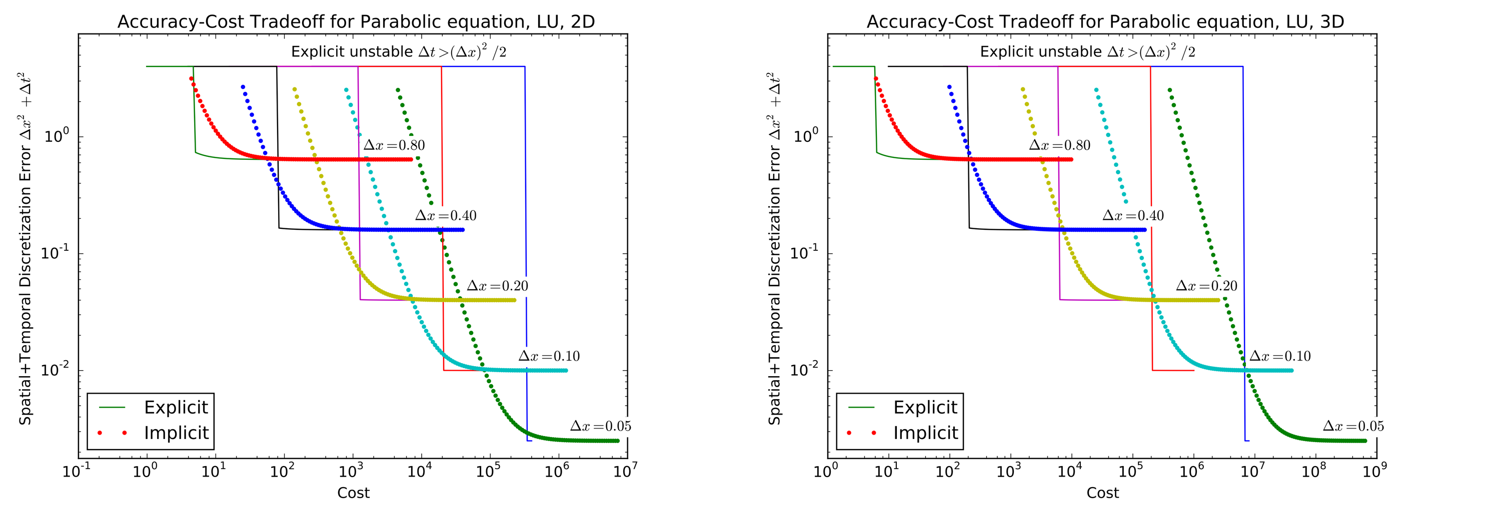

Figure: Accuracy vs Cost tradeoff when varying time step size for different choices of spatial resolution, each using a direct solver in 2D and 3D. The spatial and temporal resolution are both second order. Explicit time integrators must satisfy a stability criteria $\Delta t \le \frac{1}{2}(\Delta x)^2$. LU factorization is used as the solver, having a time cost of $\mathcal{O}(N^{3/2})$ and a storage cost of $\mathcal{O}(N \log N)$ in 2D, versus $\mathcal{O}(N^2)$ time and $\mathcal{O}(n^{4/3})$ space in 3D.

Interpretation

- For a fixed spatial resolution, explicit methods cannot deliver crude accuracy fast (unstable).
- If we choose the best spatial resolution for each accuracy, explicit methods are a good choice for crude accuracy (coarse grids).
- Not shown: explicit methods are sensitive to coefficient contrast, stretched meshes, domain shape, etc.
- Explicit methods keep pace with implicit/LU in 3D. The cost of the solve grows too fast for implicit's larger allowed steps to pay off.
- For an $\mathcal{O}(N)$ solver such as multigrid, implicit rapidly wins for stricter accuracy.

On Modeling Error versus Discretization Error

The 1986 Editorial Policy Statement for the Journal of Fluids Engineering [1] states

> *A professional problem exists in the computational fluid dynamics community and also in the broader area of computational physics. Namely, there is a need for higher standards on the control of numerical accuracy. [...] It [is] impossible to evaluate and compare the accuracy of different turbulence models, since one [cannot] distinguish physical modeling errors from numerical errors related to the algorithm and grid. This is especially the case for first-order accurate methods and hybrid methods.*

The statement goes on to mandate systematic evaluation of numerical errors. This policy has been influential in defining the standards of the computational engineering community, for whom validation is readily available because products are manufactured and tested. Models that do not ascribe to this standard are widely viewed as less reliable and lower quality.

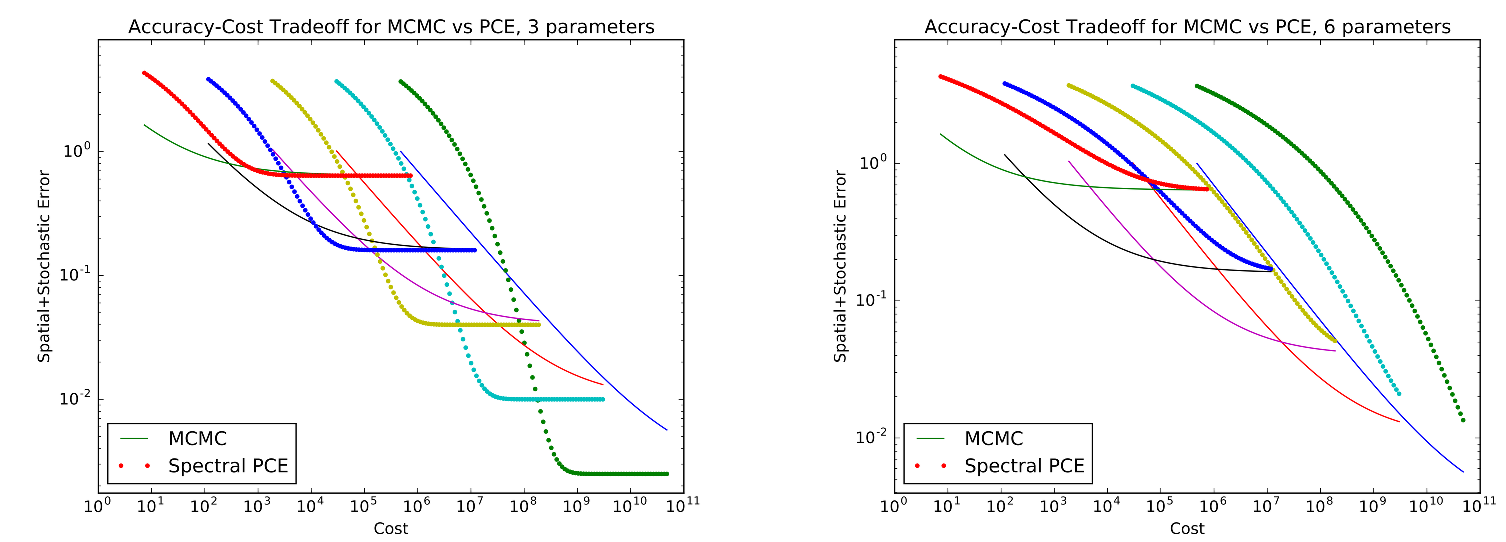

Figure: Accuracy versus Cost for computing the mean for a 3D stochastic PDE. PCE converges more rapidly than MCMC, but suffers the curse of dimensionality. The spatial resolution and sampling technique should be chosen to balance errors – there is no point computing with a very accurate spatial discretization when stochastically under-resolved, or vice-versa.

Uncertain Data and Stochastic Resolution

- Geophysical observations are inherently sparse and inexact.
- Heavy preprocessing is often performed to produced gridded data, perhaps introducing numerical artifacts and inconsistencies.
- Spatial and temporal correlation abounds.
- We need models for the uncertainty.
- Gaussian assumptions are needed for linear scaling, but often invalid.
- Markov-Chain Monte Carlo: adaptive sampling technique robust to number of parameters, but slowly converging.
- Polynomial Chaos Expansion: spectrally accurate technique, suffers curse of dimensionality.
- Dimension reduction methods seek to use these expensive techniques only in a low-dimensional subspace.
- Always more expensive: higher order statistical moments, rare events.
- No point converging beyond accuracy of stochastic assumptions and physical model.

HPGMG and the Role of Versatility

- Is hardware moving in a direction that will help applications? Which ones?
- Model fidelity: resolution, multi-scale, coupling
  - Transient simulation is not weak scaling: $\Delta t \sim \Delta x$
- Analysis using a sequence of forward simulations
  - Inversion, data assimilation, optimization
  - Quantify uncertainty, risk-aware decisions
- Increasing relevance $\implies$ external requirements on time
  - Policy: 5 SYPD to inform IPCC
  - Weather, manufacturing, field studies, disaster response
- "weak scaling" [...] will increasingly give way to "strong scaling" [2]

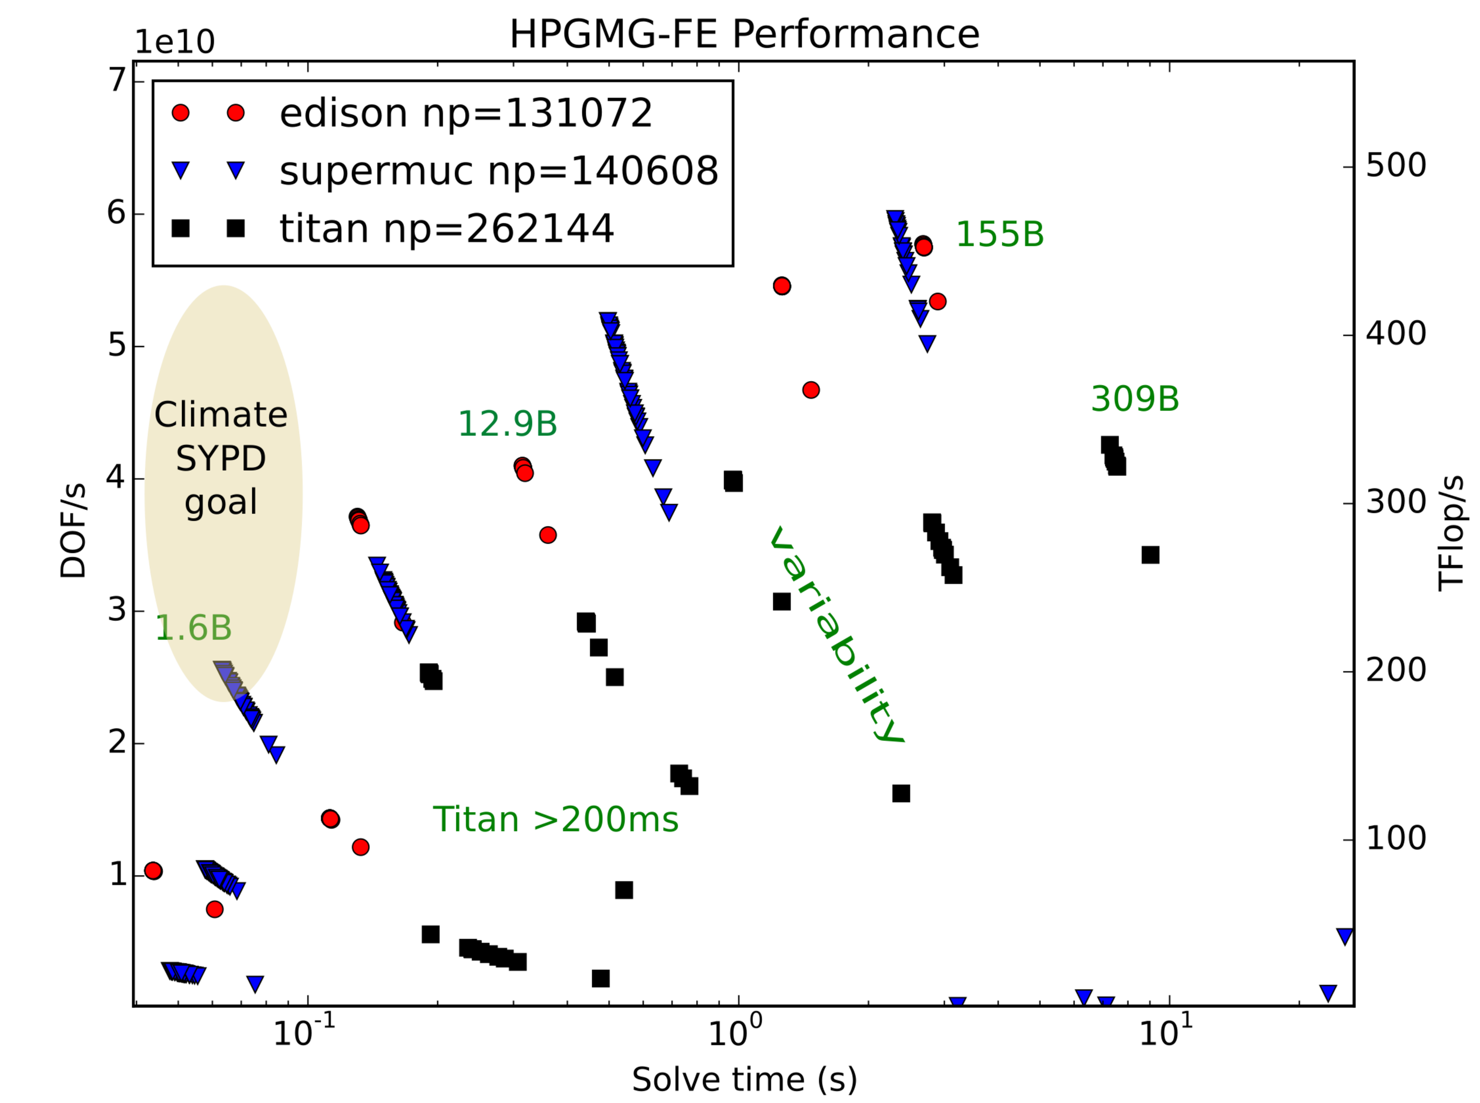

Figure: HPGMG-FE efficiency versus execution time on Edison, SuperMUC, and Titan.

Outlook

- Balancing errors affects the memory and time scaling regime.
- Sometimes hardware doesn't move in the direction the science needs.
- The asymptotics considered here represent an extremely idealized case. Real scenarios are much less clean and require comparison and competition among methods.

References

- P.J. Roache, K.N. Ghia, and F.M. White. Editorial policy statement on the control of numerical accuracy. *Journal of Fluids Engineering*, 108:2, 1986.
- Jack Dongarra, Pete Beckman, Terry Moore, et al. The international exascale software project roadmap. *International Journal of High Performance Computing Applications*, 25(1):3–60, 2011.
- Houman Owhadi, Clint Scovel, Tim Sullivan, et al. Brittleness of bayesian inference under finite information in a continuous world. *Electronic Journal of Statistics*, 9:1–79, 2015.

This work was supported by the Department of Energy Office of Advanced Scientific Computing Research. Thanks to NERSC for compute time on "Edison", ORNL for time on Titan, and LRZ for time on SuperMUC.