Textbook multigrid efficiency for hydrostatic ice flow

Jed Brown[1] Barry Smith[2]

[1]VAW, ETH Zürich

[2]MCS, Argonne National Lab

CCSM Land Ice Working Group 2010-02-17

Hydrostatic equations

- Valid in the limit $w_x \ll u_z$, independent of basal friction[1]
- Eliminate p and w by incompressibility:
  3D elliptic system for $\boldsymbol{u} = (u, v)$

$$-\nabla \cdot \left[\eta \begin{pmatrix} 4u_x + 2v_y & u_y + v_x & u_z \\ u_y + v_x & 2u_x + 4v_y & v_z \end{pmatrix} \right] + \rho g \nabla s = 0$$

$$\eta(\gamma) = \frac{B}{2}(\epsilon^2 + \gamma)^{\frac{1-\mathfrak{n}}{2\mathfrak{n}}}, \qquad \mathfrak{n} \approx 3$$

$$\gamma = u_x^2 + v_y^2 + u_x v_y + \frac{1}{4}(u_y + v_x)^2 + \frac{1}{4}u_z^2 + \frac{1}{4}v_z^2$$

and slip boundary $\sigma \cdot \boldsymbol{n} = \beta^2 \boldsymbol{u}$ where

$$\beta^2(\gamma_b) = \beta_0^2(\epsilon_b^2 + \gamma_b)^{\frac{\mathfrak{m}-1}{2}}, \qquad 0 < \mathfrak{m} \leq 1$$

$$\gamma_b = \frac{1}{2}(u^2 + v^2)$$

[1] C. Schoof and R. Hindmarsh, *Thin-film flows with wall slip: an asymptotic analysis of higher order glacier flow models*, 2010

Hydrostatic solver

- 3D elliptic system for (u, v)

$$-\nabla \cdot \left[\eta \begin{pmatrix} 4u_x + 2v_y & u_y + v_x & u_z \\ u_y + v_x & 2u_x + 4v_y & v_z \end{pmatrix} \right] + \rho g \nabla s = 0$$

- strong anisotropy and viscosity variation, shear important
- current solvers
  - Picard iteration
  - serial linear solve or 1-level ASM, high iteration counts
  - slow convergence with nonlinear sliding
- `petsc/src/snes/examples/tutorials/ex48.c`
  - conforming Q_1 finite element discretization, no σ-transform
  - multigrid Newton-Krylov
  - smoother respects strong coupling in vertical: rapid coarsening
  - globalized by grid sequencing
  - quadratic convergence including slip conditions
  - robust to discontinous sliding, high aspect ratio
  - high throughput: SSE2 integration, blocked matrix formats

Convergence with grid sequencing

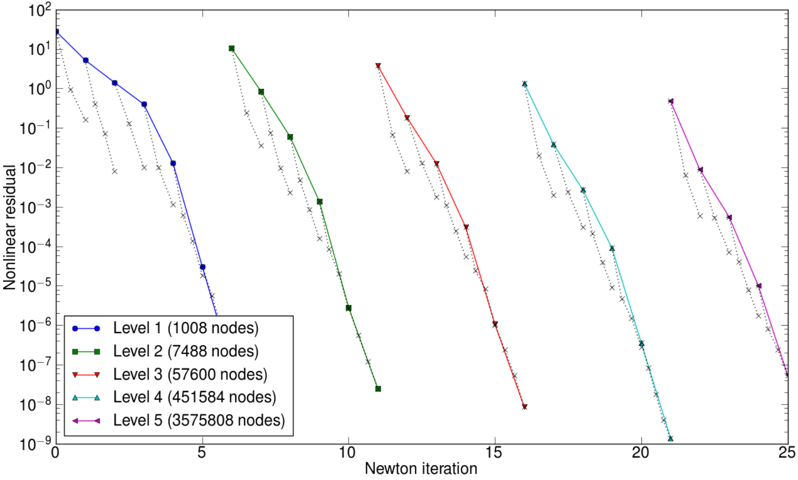

- colored points are nonlinear residuals
- black $\times$ marks are unpreconditioned linear residuals

Picard with ASM overlap 1/ICC(1), 8 subdomains

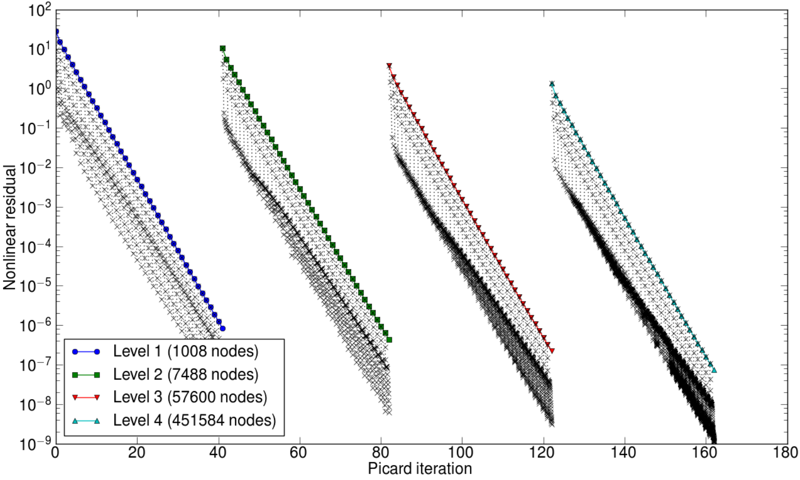

- $\mathfrak{n} = 3, \mathfrak{m} = 1/3, \frac{\eta_{\max}}{\eta_{\min}} \approx 380$
- geometry of ISMIP-HOM test A, "dimpled sombrero" sliding
- slow nonlinear convergence, over 100 iterations for rtol 10^{-2}

Picard with multigrid preconditioning, 512 subdomains

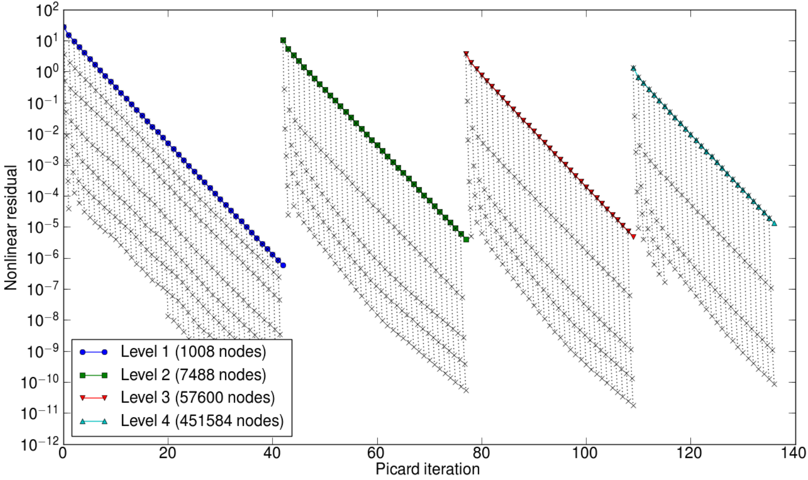

- slow nonlinear convergence
- very fast linear solves (note rtol is now 10^{-5})

Picard without grid sequencing

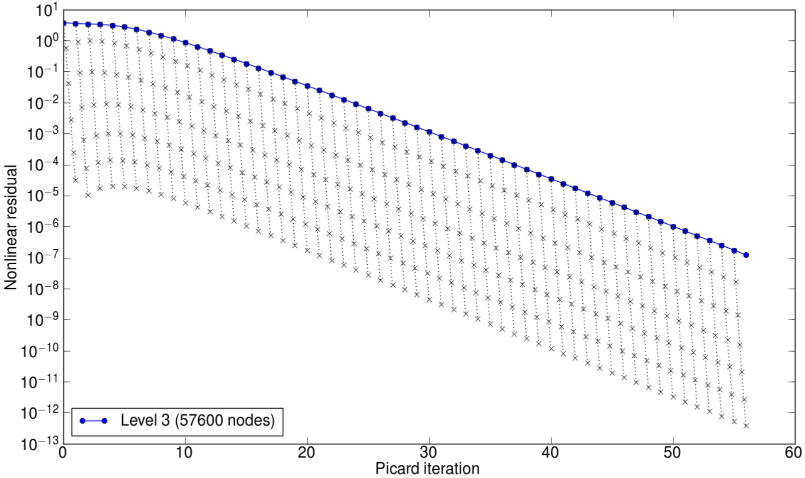

- slow initial convergence

Picard for rheology, Newton-linearized sliding

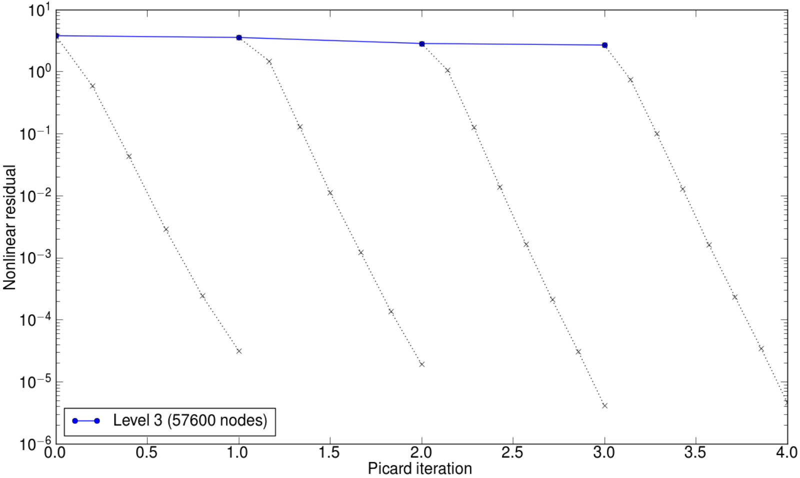

- cannot Newton-linearize only sliding, step is not descent direction

Switch to Newton

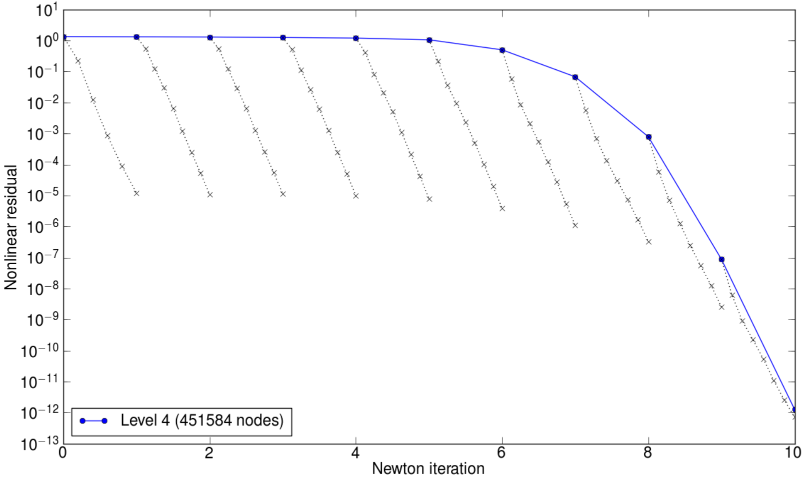

- *much* faster nonlinear convergence
- linear systems slightly more difficult
- initial nonlinear convergence is slow

Nonlinear convergence

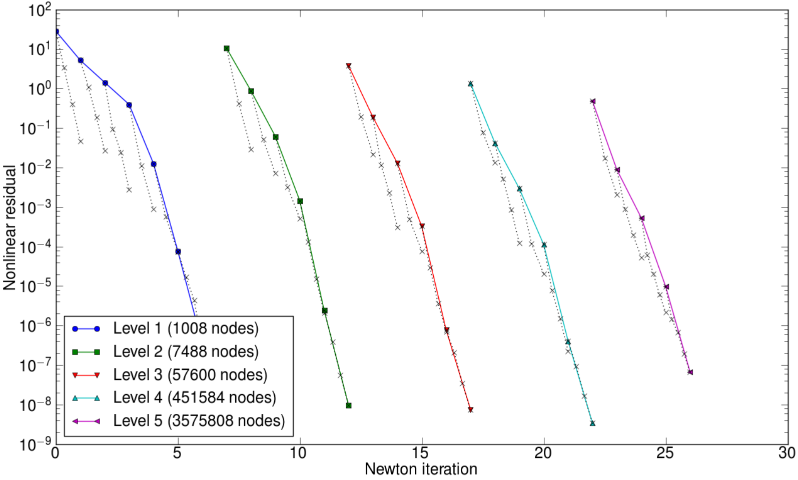

- ASM/ICC(0) smoothers, isotropic coarsening
- 512 subdomains on fine levels

Avoid oversolving

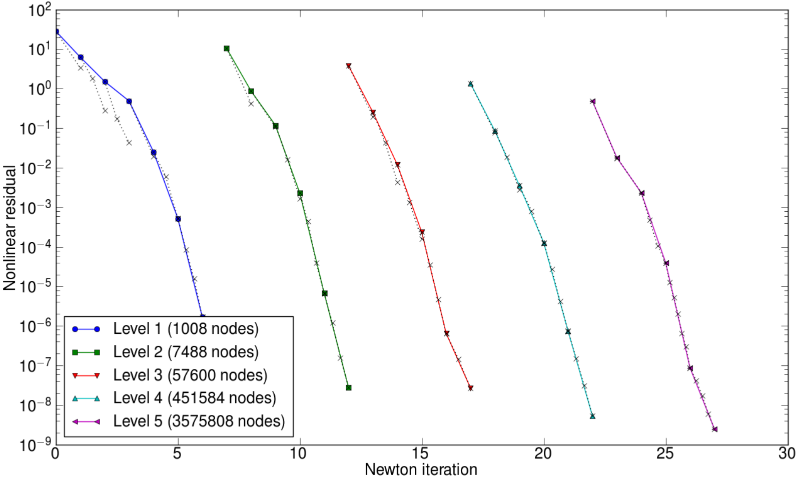

- Luis Chacon's variant of Eisenstat-Walker
- almost equivalent nonlinear convergence

Linear solve performance

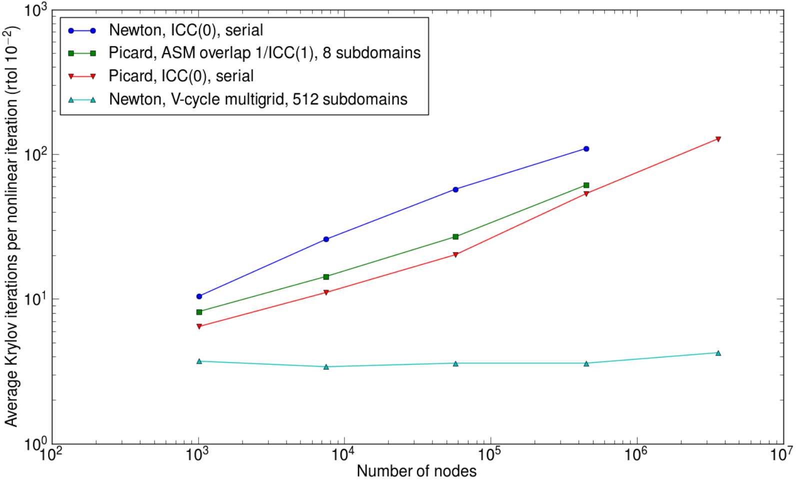

Aggressive coarsening

- coarsest level is responsible for providing global coupling
- subdomains very small on coarse levels: bottleneck is latency
- approaches to anisotropy
  - semi-coarsening (z-only): many intermediate levels, expensive
  - line smothers: allow rapid coarsening
- order unknowns so that ICC(0) performs exact column solves
  - allows isotropic coarsening at moderate aspect ratios
  - allows semicoarsening by factors ≥ 8

Nonlinear convergence: $\mathfrak{n} = 3, \mathfrak{m} = 1, \eta_{\max}/\eta_{\min} = 3800$

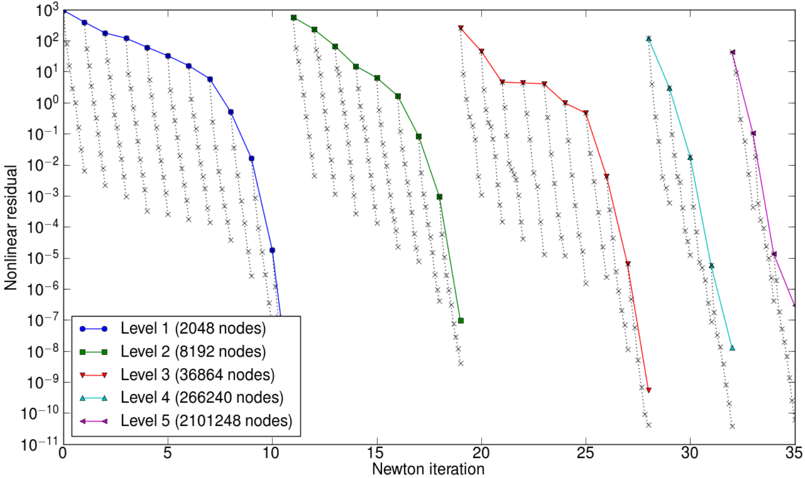

- element aspect ratio 640
- discontinuous sliding, strong shear zone (0.5 m/a to 46 km/a)
- rapid semi-coarsening, then isotropic quasi-2D coarsening

Same problem, rtol 10^{-2}

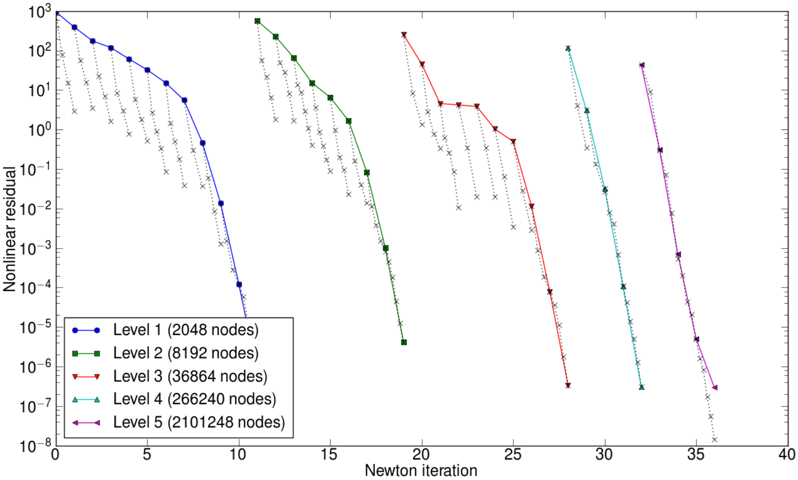

- almost equivalent convergence

Breakdown of incomplete factorization

- Jacobian is symmetric positive definite
- full Cholesky always has positive pivots
- incomplete Cholesky can generate negative pivots
- increasing fill can make the problem worse
- standard practice is to “shift” diagonal
  - may still be an adequate local smoother
  - destroys nonlocal properties of preconditioner
  - increasing overlap often makes it worse
- additive Schwarz with incomplete Cholesky is unreliable
- additive Schwarz with direct solves is reliable but slow

Conclusions

- Newton is almost 10 times faster than Picard
- deficiencies of serial incomplete factorization is more evident with Newton than Picard
- incomplete factorization is poor with 1-level domain decomposition
- multigrid preconditioning is ~ 50 times faster at $\sim 1M$ nodes
- grid sequencing is important for globalization
- preconditioner can be lagged with true Jacobian applied matrix-free