Terra-Neo

A two-scale approach for efficient on-the-fly operator assembly on curved geometries

Simon Bauer[1] Markus Huber[2] Marcus Mohr[1] Ulrich Rüde[3]
Jens Weismüller[1,5] Markus Wittmann[4] Barbara Wohlmuth[2]

[1]Dept. of Earth and Environmental Sciences, Ludwig-Maximilians-Universität München
[2]Institute for Numerical Mathematics (M2), Technische Universität München
[3]Computer Science 10, FAU Erlangen-Nürnberg
[4]Erlangen Regional Computing Centre (RRZE), FAU Erlangen-Nürnberg
[5]Leibniz Supercomputing Centre (LRZ)

SPPEXA Annual Plenary Meeting
21.03.2017

Geophysicist

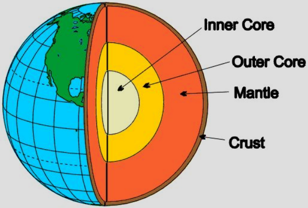

Earth Mantle represented as a thick spherical shell

Geophysicist

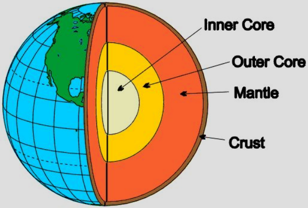

Earth Mantle represented as a thick spherical shell

Mathematician/Computer Scientist

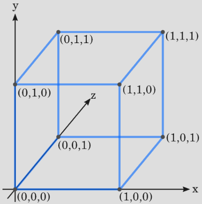

+ well-suited for matrix-free FE implementations

+ regular grid → constant access patterns

Geophysicist

Mathematician/Computer Scientist

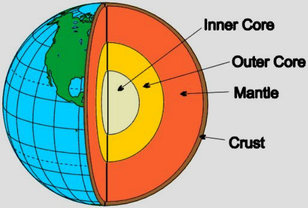

≠

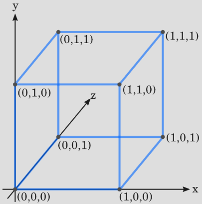

Earth Mantle represented as a thick spherical shell

+ well-suited for matrix-free FE implementations

+ regular grid → constant access patterns

Hierarchical Hybrid Grids (HHG)

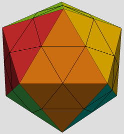

Hybrid approach

- initial coarse unstructured mesh (macro elements)
- regular refinement of macro elements
- constant access patterns (FE stencils) within each macro element for the refined grids
- matrix-free implementation
- compute stencils on-the-fly (only **ONCE per macro element**), reduces memory traffic
- excellent **parallel scalability** and node **performance** [Gmeiner et al., 2015]

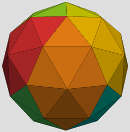

Projection of fine grid nodes

- accurate representation of geometry on all levels
- FE stencils are not constant within macro elements
- need to assemble stencil for **EACH fine grid node** → **SLOW**

Sequence of Mappings (2D sketch)

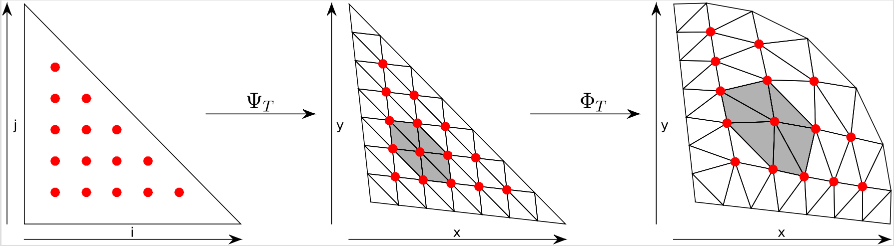

reference domain original HHG domain projected HHG domain

Sequence of Mappings (2D sketch)

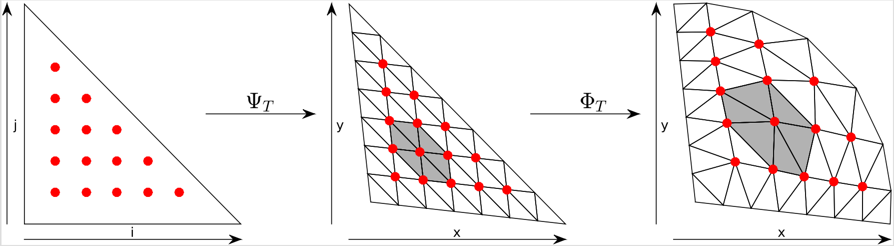

reference domain original HHG domain projected HHG domain

Can we get the performance of the original HHG also on the projected domain?

Outline

1) **Introduce novel approach**

$$-\Delta u = f$$

2) **Numerical study and performance results**

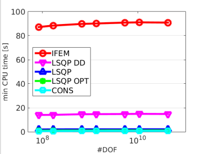

3) **Geophysical application**

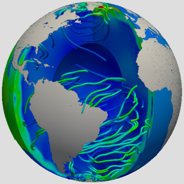

Novel Approach for Efficient Stencil Assembly on Curved Geometries

HHG Stencil Layout

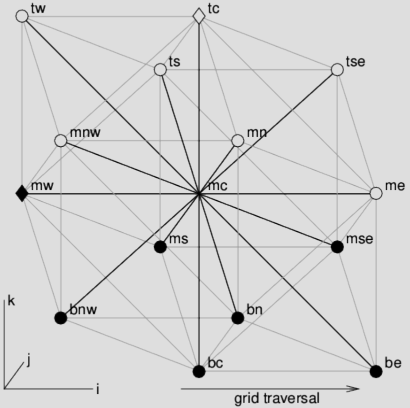

Layout of the 15pt stencil in HHG

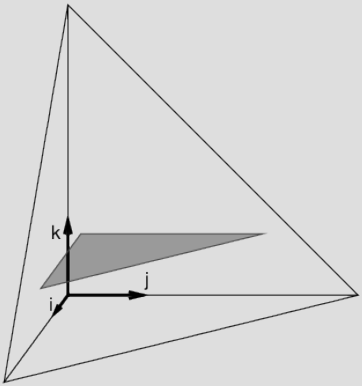

Selected interior points of a macro tetrahedron in 3D reference domain

Original (blue) vs. projected (red) stencil weights

weights (TC, . . .) of 15pt stencil for nodes in gray plane (previous slide)

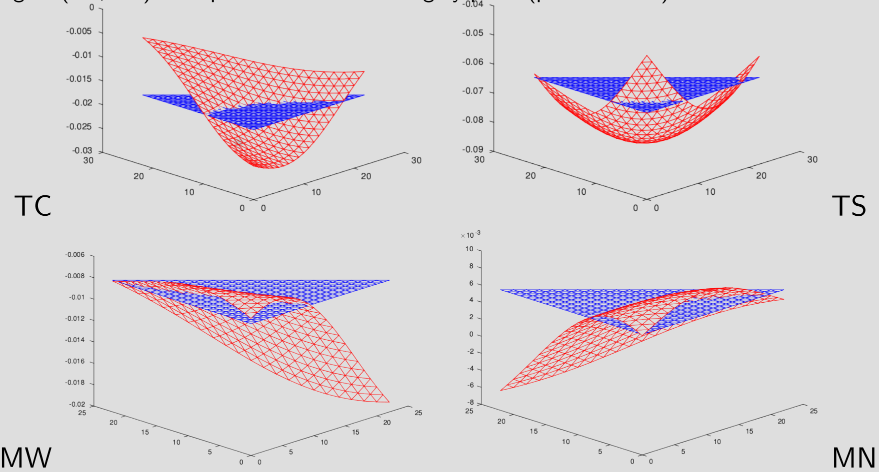

Novel Approach

Idea: Approximate stencils by low order polynomials

Replace expensive stencil assembly via evaluation of local element matrices by low-cost approximation

- Employ **quadratic** or **cubic** (or even higher order) polynomials
- Pre-compute polynomial coefficients in setup phase and store them
- Define one polynomial for each stencil weight at each level and for each macro element
- Fit polynomial coefficients by interpolation (IPOLY) or least-squares approach (LSQP)

Sampling Points for Determination of Polynomial Coefficients

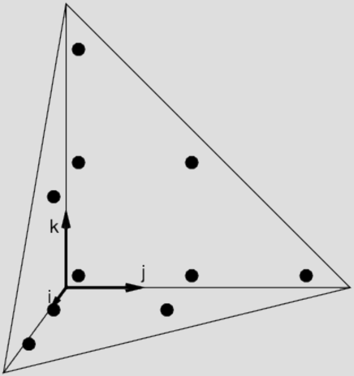

IPOLY: (quadratic) 10 sampling points → interpolation polynomial

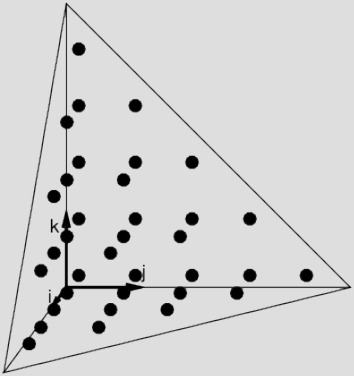

LSQP: employ more sampling points and solve least-squares problem

Polynomial Evaluation

1D: a quadratic polynomial
$p_i := p(z_i) = a_2 z_i^2 + a_1 z_i + a_0$
can be evaluated incrementally with only two flops

$$p_{i+1} = p_i + \delta p_i, \quad \delta p_{i+1} = \delta p_i + \delta k$$

where

$$\delta p(z_i) := p'(z_i)h + \frac{p''(z_i)}{2}h^2$$
$$\delta k := 2ah^2$$

Polynomial degree $q > 2$: Formula can be generalized to higher polynomial degrees (based on Taylor series expansion)

3D: ansatz transfers directly to higher dimensions

Numerical Study and Performance Results

Model Problem

$-\Delta u = f \text{ in } \Omega,\ u = 0 \text{ on } \partial\Omega$

Exact solution:

$u(x,y,z) = g(x,y,z) \cdot h(x,y,z)$

with

$$g(x,y,z) = (r - r_{\min})(r - r_{\max})$$
$$h(x,y,z) = \sin(10x)\sin(4y)\sin(7z)$$

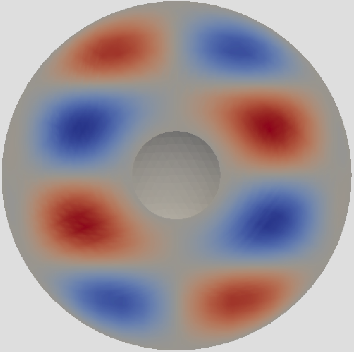

Numerical Experiments (1/2)

- Employ geometric multigrid solver with V(3,3)-cycle and Gauss-Seidel (GS) smoother
- Test different numerical strategies:

CONS	original HHG (non-projected)
IFEM	projected HHG (assemble FE stencils exactly)
IPOLY	employ interpolation polynomial(quadratic)
LSQP	employ approximation polynomial (quadratic)
LSQP (q=3)	employ approximation polynomial (cubic)
DD (Double Discretization)	use IFEM for residual and IPOLY/LSQP for the smoother

Numerical Experiments (2/2)

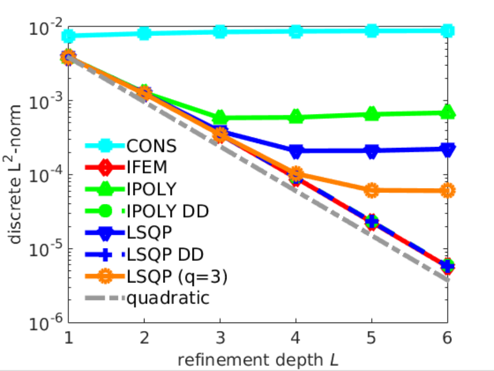

Discretization error for different levels of refinement

- **LSQP** better than **IPOLY**, but both deteriorate after some levels
- **cubic** more accurate than **quadratic** polynomials (but more expensive)
- **DD** (Double Discretization) schemes are **exact** for all levels
- critical level L_C depends on macro element size H, **for smaller H we get larger L_C**
- Disc. error is in $\mathcal{O}(h^2) + \mathcal{O}(H^{q+1})$
  q : polynomial degree

Numerical Experiments - Take Home Message

Disc. error is in $\mathcal{O}(h^2) + \mathcal{O}(H^{q+1})$
q : polynomial degree

Task: Identify (h,H)-parameter regime such that "$\mathcal{O}(H^{q+1}) \leqslant \mathcal{O}(h^2)$"
Large Scale: $\mathcal{O}(10^4)$ MPI-processes $\rightarrow$ lower bound for number of macro elements $\rightarrow$ small H $\rightarrow$ $\mathcal{O}(H^{q+1})$ is also small

LSQP (q=2) is accurate for at least 6 levels of refinement.

This allows a global resolution of the Earth mantle of 1km.

Node Level Optimization

Execution-Cache-Memory (ECM) model (single core)

Reported cycles from IACA for the 3 loops normalized to 8 stencil updates. (maximum values of each port of a certain category)

loop	total	arithmetic	data	type	comment
Loop 1	22	20	22	throughput	vectorized
Loop 2	56	56	36	throughput	vectorized
Loop 3	40	4	12	throughput	scalar w/ recursion

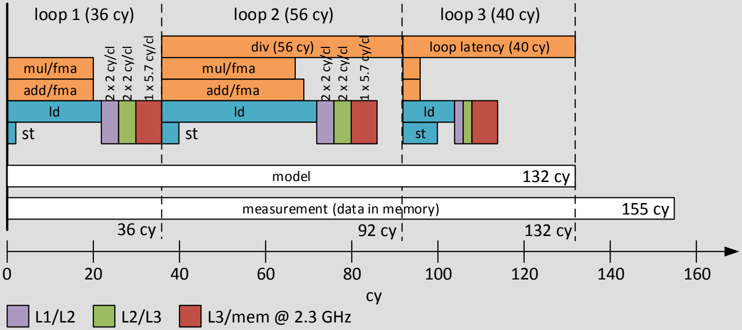

Weak Scaling on SuperMUC Phase 2

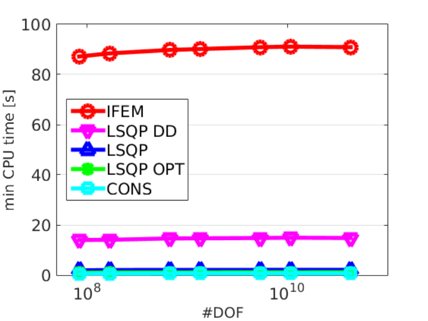

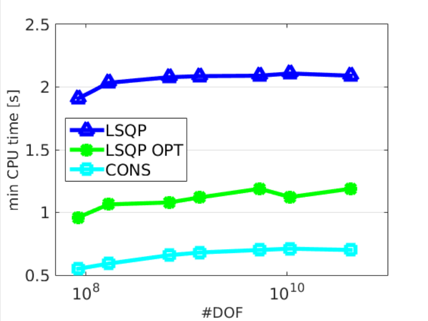

Minimal CPU time [s] for one V(3,3)-cycle. 16 macro elements are assigned per core with $3.3 \cdot 10^5$ DOFs per macro element.

Geophysical Application

Mantle Convection: Physical Model

Conservation of **momentum**, **mass** and **energy**:

$$\operatorname{div} \boldsymbol{\sigma} + \varrho \mathbf{g} = 0 \tag{1a}$$

$$\partial_t \varrho + \operatorname{div}(\varrho \mathbf{u}) = 0 \tag{1b}$$

$$\partial_t(\varrho e) + \operatorname{div}(\varrho e \mathbf{u}) + \operatorname{div} \mathbf{q} - H - \boldsymbol{\sigma} : \dot{\boldsymbol{\varepsilon}} = 0 \tag{1c}$$

with

$$\boldsymbol{\sigma} = 2\mu \left(\dot{\boldsymbol{\varepsilon}} - \frac{\operatorname{tr} \dot{\boldsymbol{\varepsilon}}}{3} \mathbf{I} \right) - p\mathbf{I} \ , \quad \dot{\boldsymbol{\varepsilon}} = \frac{1}{2} \left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T \right)$$

and an equation of state $\varrho = \varrho(p, T)$.

ϱ	density	$\mathbf{u}$	velocity	p	pressure
$\mathbf{g}$	gravitational acceleration	μ	dynamic viscosity	$\boldsymbol{\sigma}$	Cauchy stress tensor
$\dot{\boldsymbol{\varepsilon}}$	rate of strain tensor	T	temperature	e	internal energy
$\mathbf{q}$	heat flux per unit area	H	heat production rate		

Mantle Convection: Generalised Stokes System

Re-casting in terms of key quantities velocity, pressure and temperature, we get a generalised Stokes problem for the quasi-stationary flow (1a),(1b) (with an elliptic *viscous operator* L)

$$\mathrm{L}(\mathbf{u};\mu) - \nabla p = \boldsymbol{F}(T)$$
$$\mathrm{div}(\mathbf{u}) = 0$$

$$\mu = \mu(\mathbf{x}, T, \mathbf{u})$$

coupled to the energy equation (1c)

$$\partial_t T + \mathbf{u} \cdot \nabla T + \mathrm{div}(\kappa \nabla T) - \frac{1}{c_p \varrho}(H - \boldsymbol{\sigma} : \dot{\boldsymbol{\varepsilon}}) = 0$$

only via the buoyancy term $\boldsymbol{F}(T)$.

c_p	specific heat capacity	κ	heat conductivity

LSQP for Stokes System

Mixed Finite Element discretisation:

$$\begin{bmatrix} \mathbf{A}(\mu) & \mathbf{B}^T \\ \mathbf{B} & -\mathbf{C} \end{bmatrix} \begin{bmatrix} u \\ q \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix}$$

Challenges:

1) Curved geometry $\rightarrow$ employ LSQP approach
   - system of PDEs $\rightarrow$ approximate stencils of each operator block individually with a quadratic polynomial
2) How do we deal with non-constant μ?

LSQP with Variable Coefficients

Remember: stencil weight $a_{k,m}^{I,J}$ is computed via:

$$a_{k,m}^{I,J} = \sum_{t \in \mathcal{N}(I) \cap \mathcal{N}(J)} \bar{\mu}_t E_t^{I,J} \quad (2)$$

- Approximate stencil weight including μ ("classical" LSQP approach) $\rightarrow$ only for "sufficiently smooth" μ
- Approximate local element matrices $\rightarrow$ replace expensive $E_t^{I,J}$ with polynomial approximation (LSQP LocEl)
  - $+$ works for general μ
  - – more expensive, but still much faster than computation of $E_t^{I,J}$

I, J	node indices	$\mathcal{N}(I)$	neighbourhood of node I
k, m	operator component, $1 \leqslant k, m \leqslant 3$	$E_t^{I,J}$	contribution of local element matrix
t	local element	$\bar{\mu}_t$	averaged viscosity on element t

Application - Velocity Streamlines

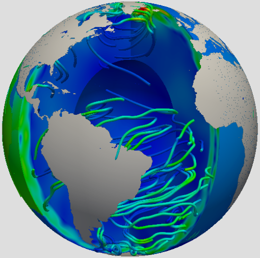

classic FE assembly

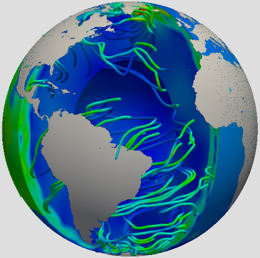

LSQP LocEl

Viscosity model [Davies et al. 2012] with jump at $d_a = 660/6371$

$$\mu(\mathbf{x}, T) = \exp\left(4.61 \frac{1 - \|\mathbf{x}\|_2}{1 - r_{\text{cmb}}} - 2.99\,T\right) \begin{cases} 1/10 \cdot 6.371^3 d_a^3 & \text{for } \|\mathbf{x}\|_2 > 1 - d_a, \\ 1 & \text{else.} \end{cases}$$

Present day temperature field T [Grand et. al 1997, Stixrude and Lithgow-Bertelloni 2005], plate velocities [GPlates] on surface and free-slip BC's on core-mantle boundary (cmb)

Application - Velocity Streamlines

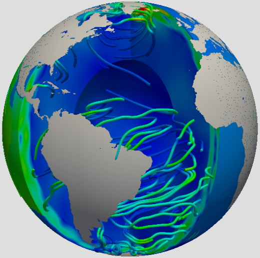

classic FE assembly

$\Longrightarrow$

speedup

5.9x

(excl. coarse grid solver)

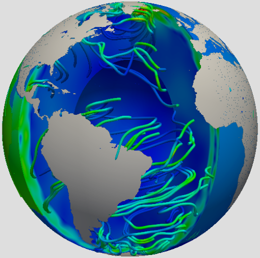

LSQP LocEl

Viscosity model [Davies et al. 2012] with jump at $d_a = 660/6371$

$$\mu(\mathbf{x}, T) = \exp\left(4.61 \frac{1 - \|\mathbf{x}\|_2}{1 - r_{\text{cmb}}} - 2.99\,T\right) \begin{cases} 1/10 \cdot 6.371^3 d_a^3 & \text{for } \|\mathbf{x}\|_2 > 1 - d_a, \\ 1 & \text{else.} \end{cases}$$

Present day temperature field T [Grand et. al 1997, Stixrude and Lithgow-Bertelloni 2005], plate velocities [GPlates] on surface and free-slip BC's on core-mantle boundary (cmb)

Thank you for your attention

Original (blue) vs. projected (red) stencil weights

weights (MC, . . .) of 15pt stencil for nodes in gray plane (previous slide)

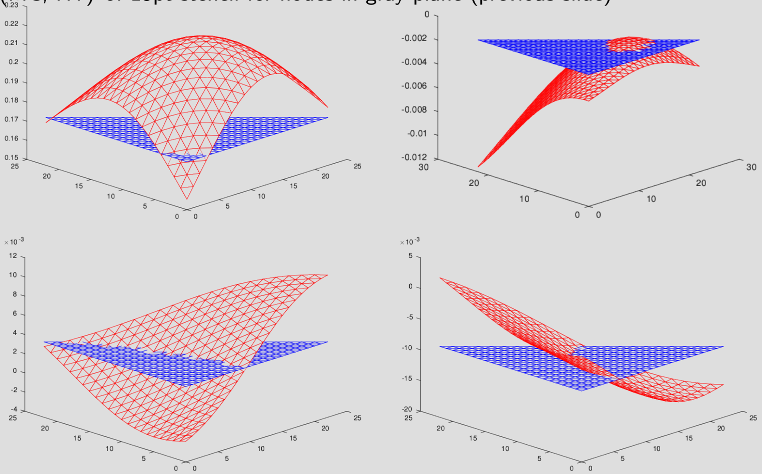

MC

MSE

BNW

BE

Numerical Experiments

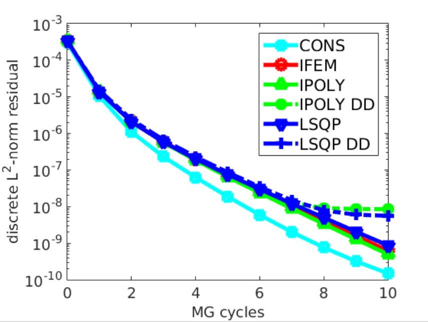

Residual in discrete L_2-norm

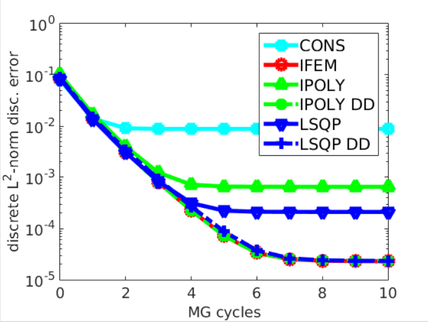

Discretization error in discrete L_2-norm