ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Scalable Multi-level Preconditioners for Variable Viscosity Stokes Flow Problems Arising from Geodynamics

Dave May
Department of Earth Sciences
ETH Zürich, Switzerland
(dave.may@erdw.ethz.ch)

Geodynamics := ???

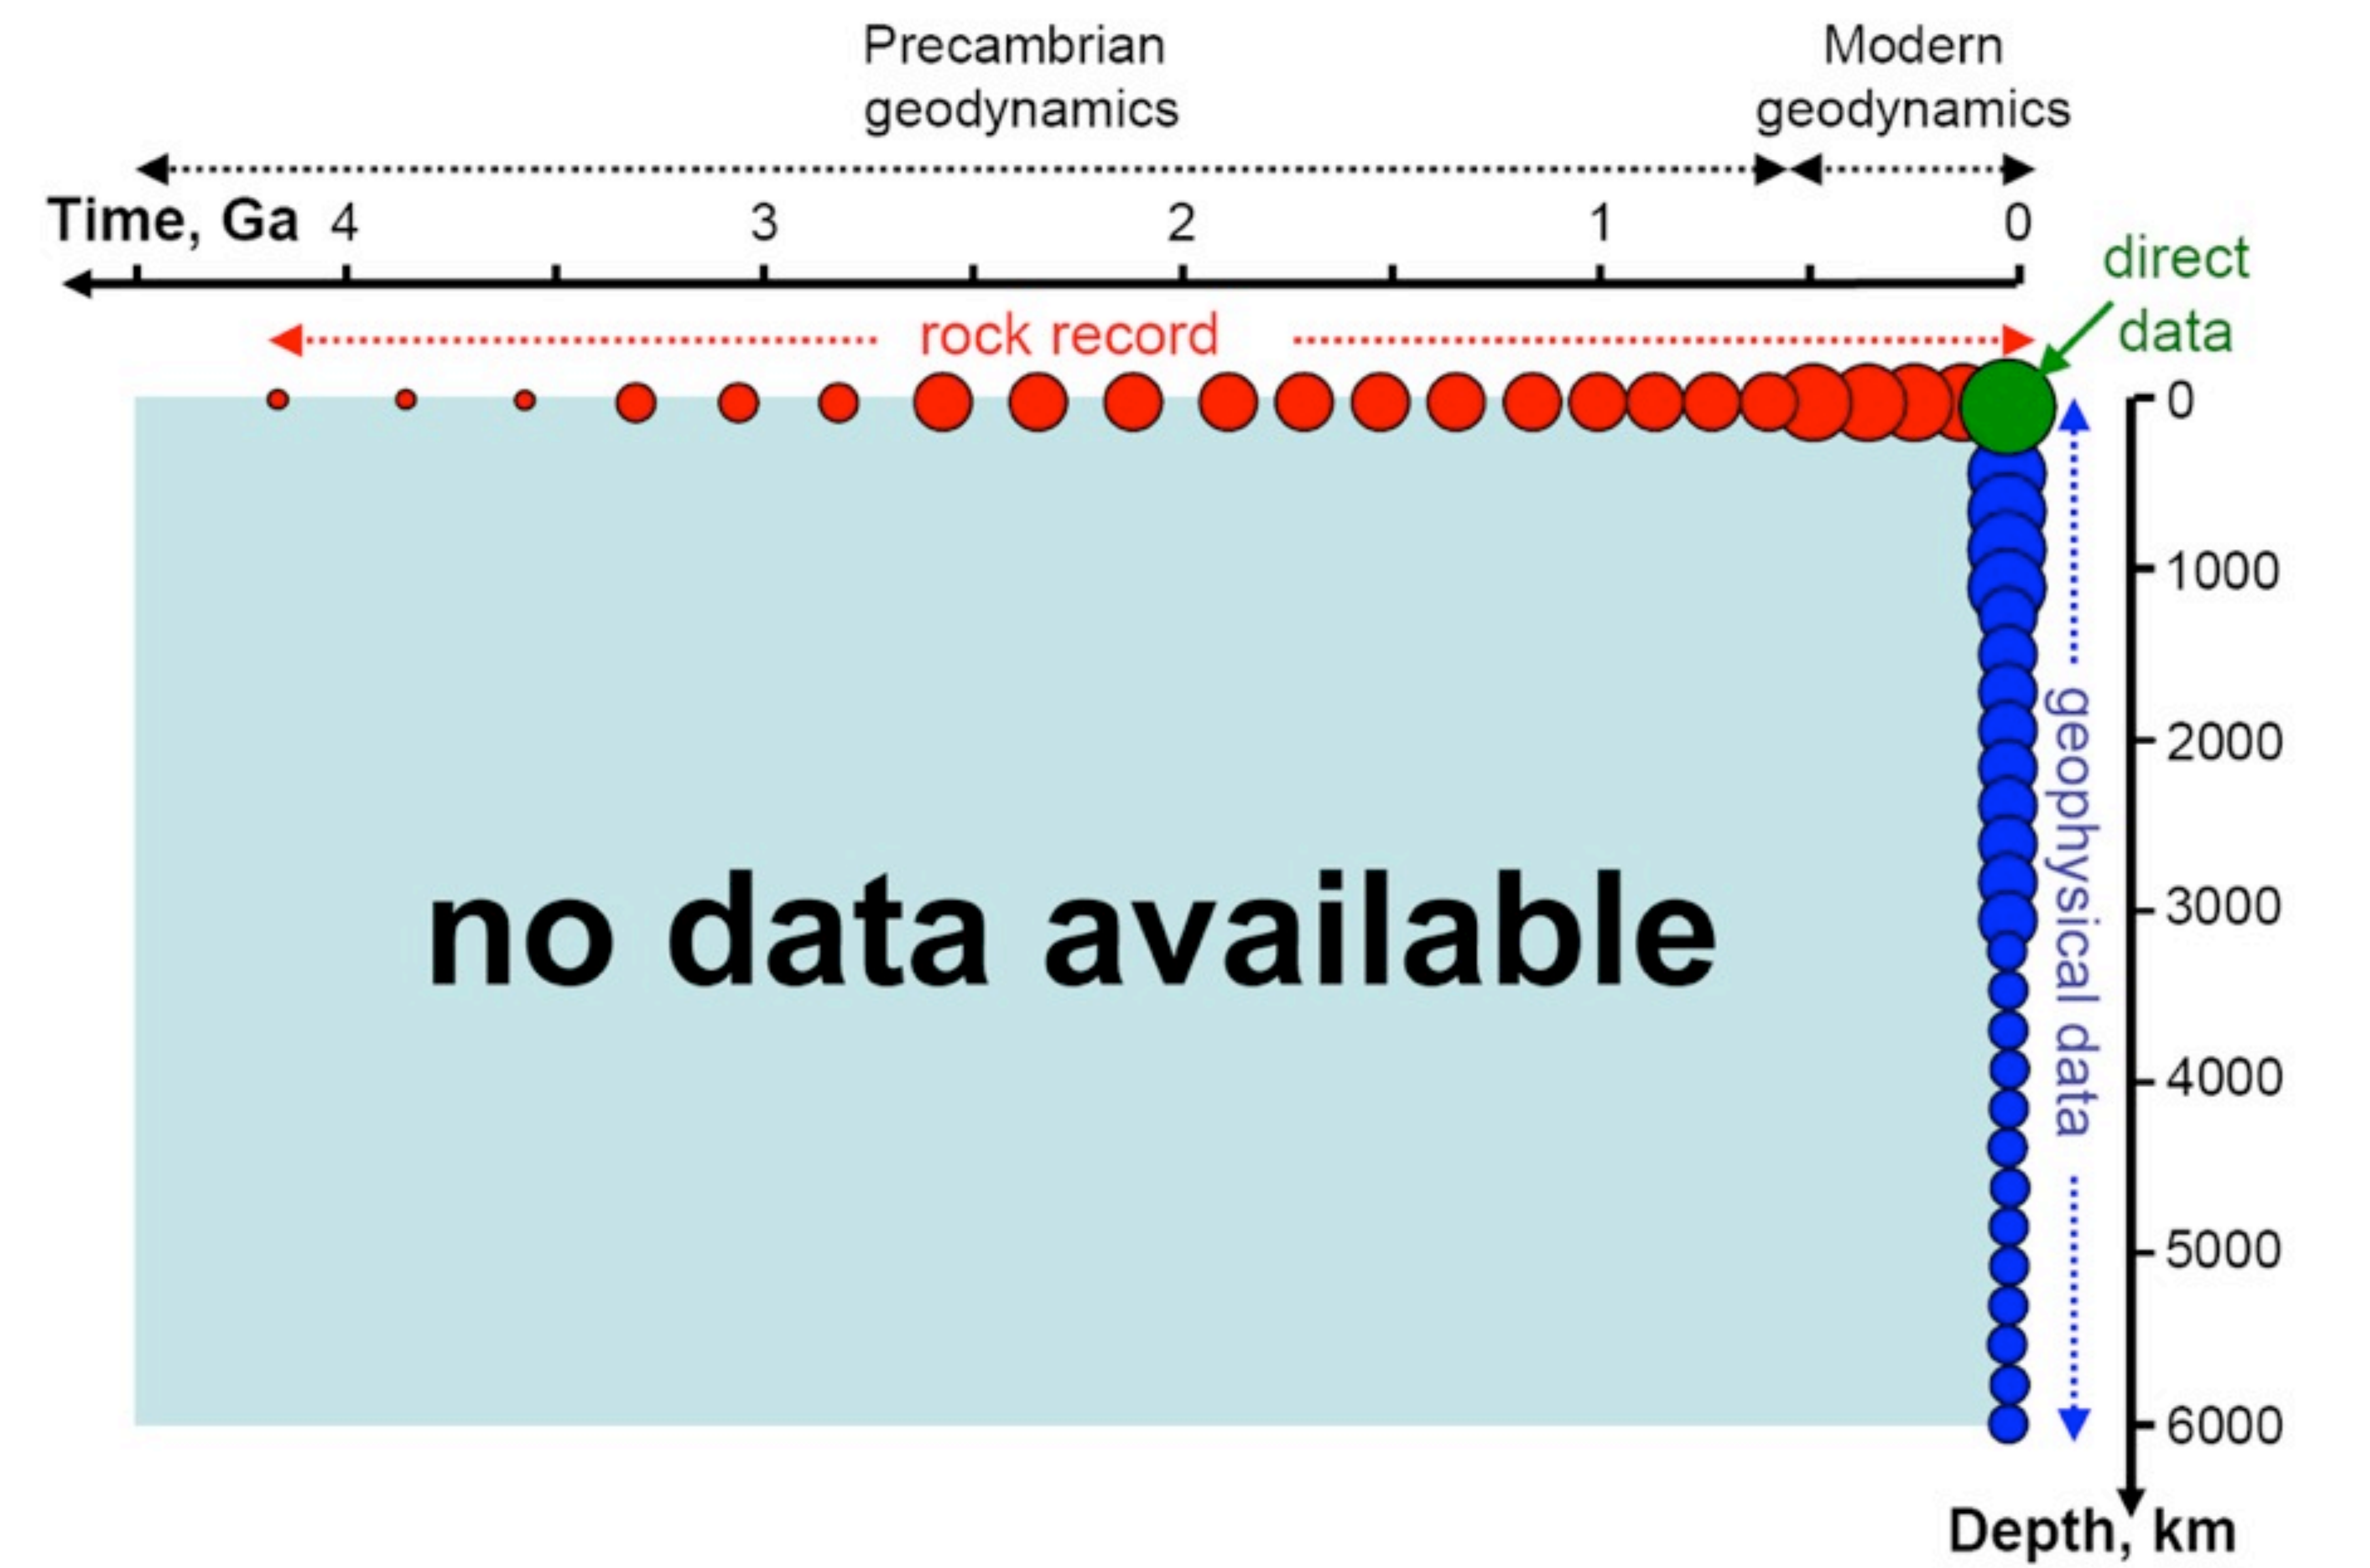

Geodynamics := ???

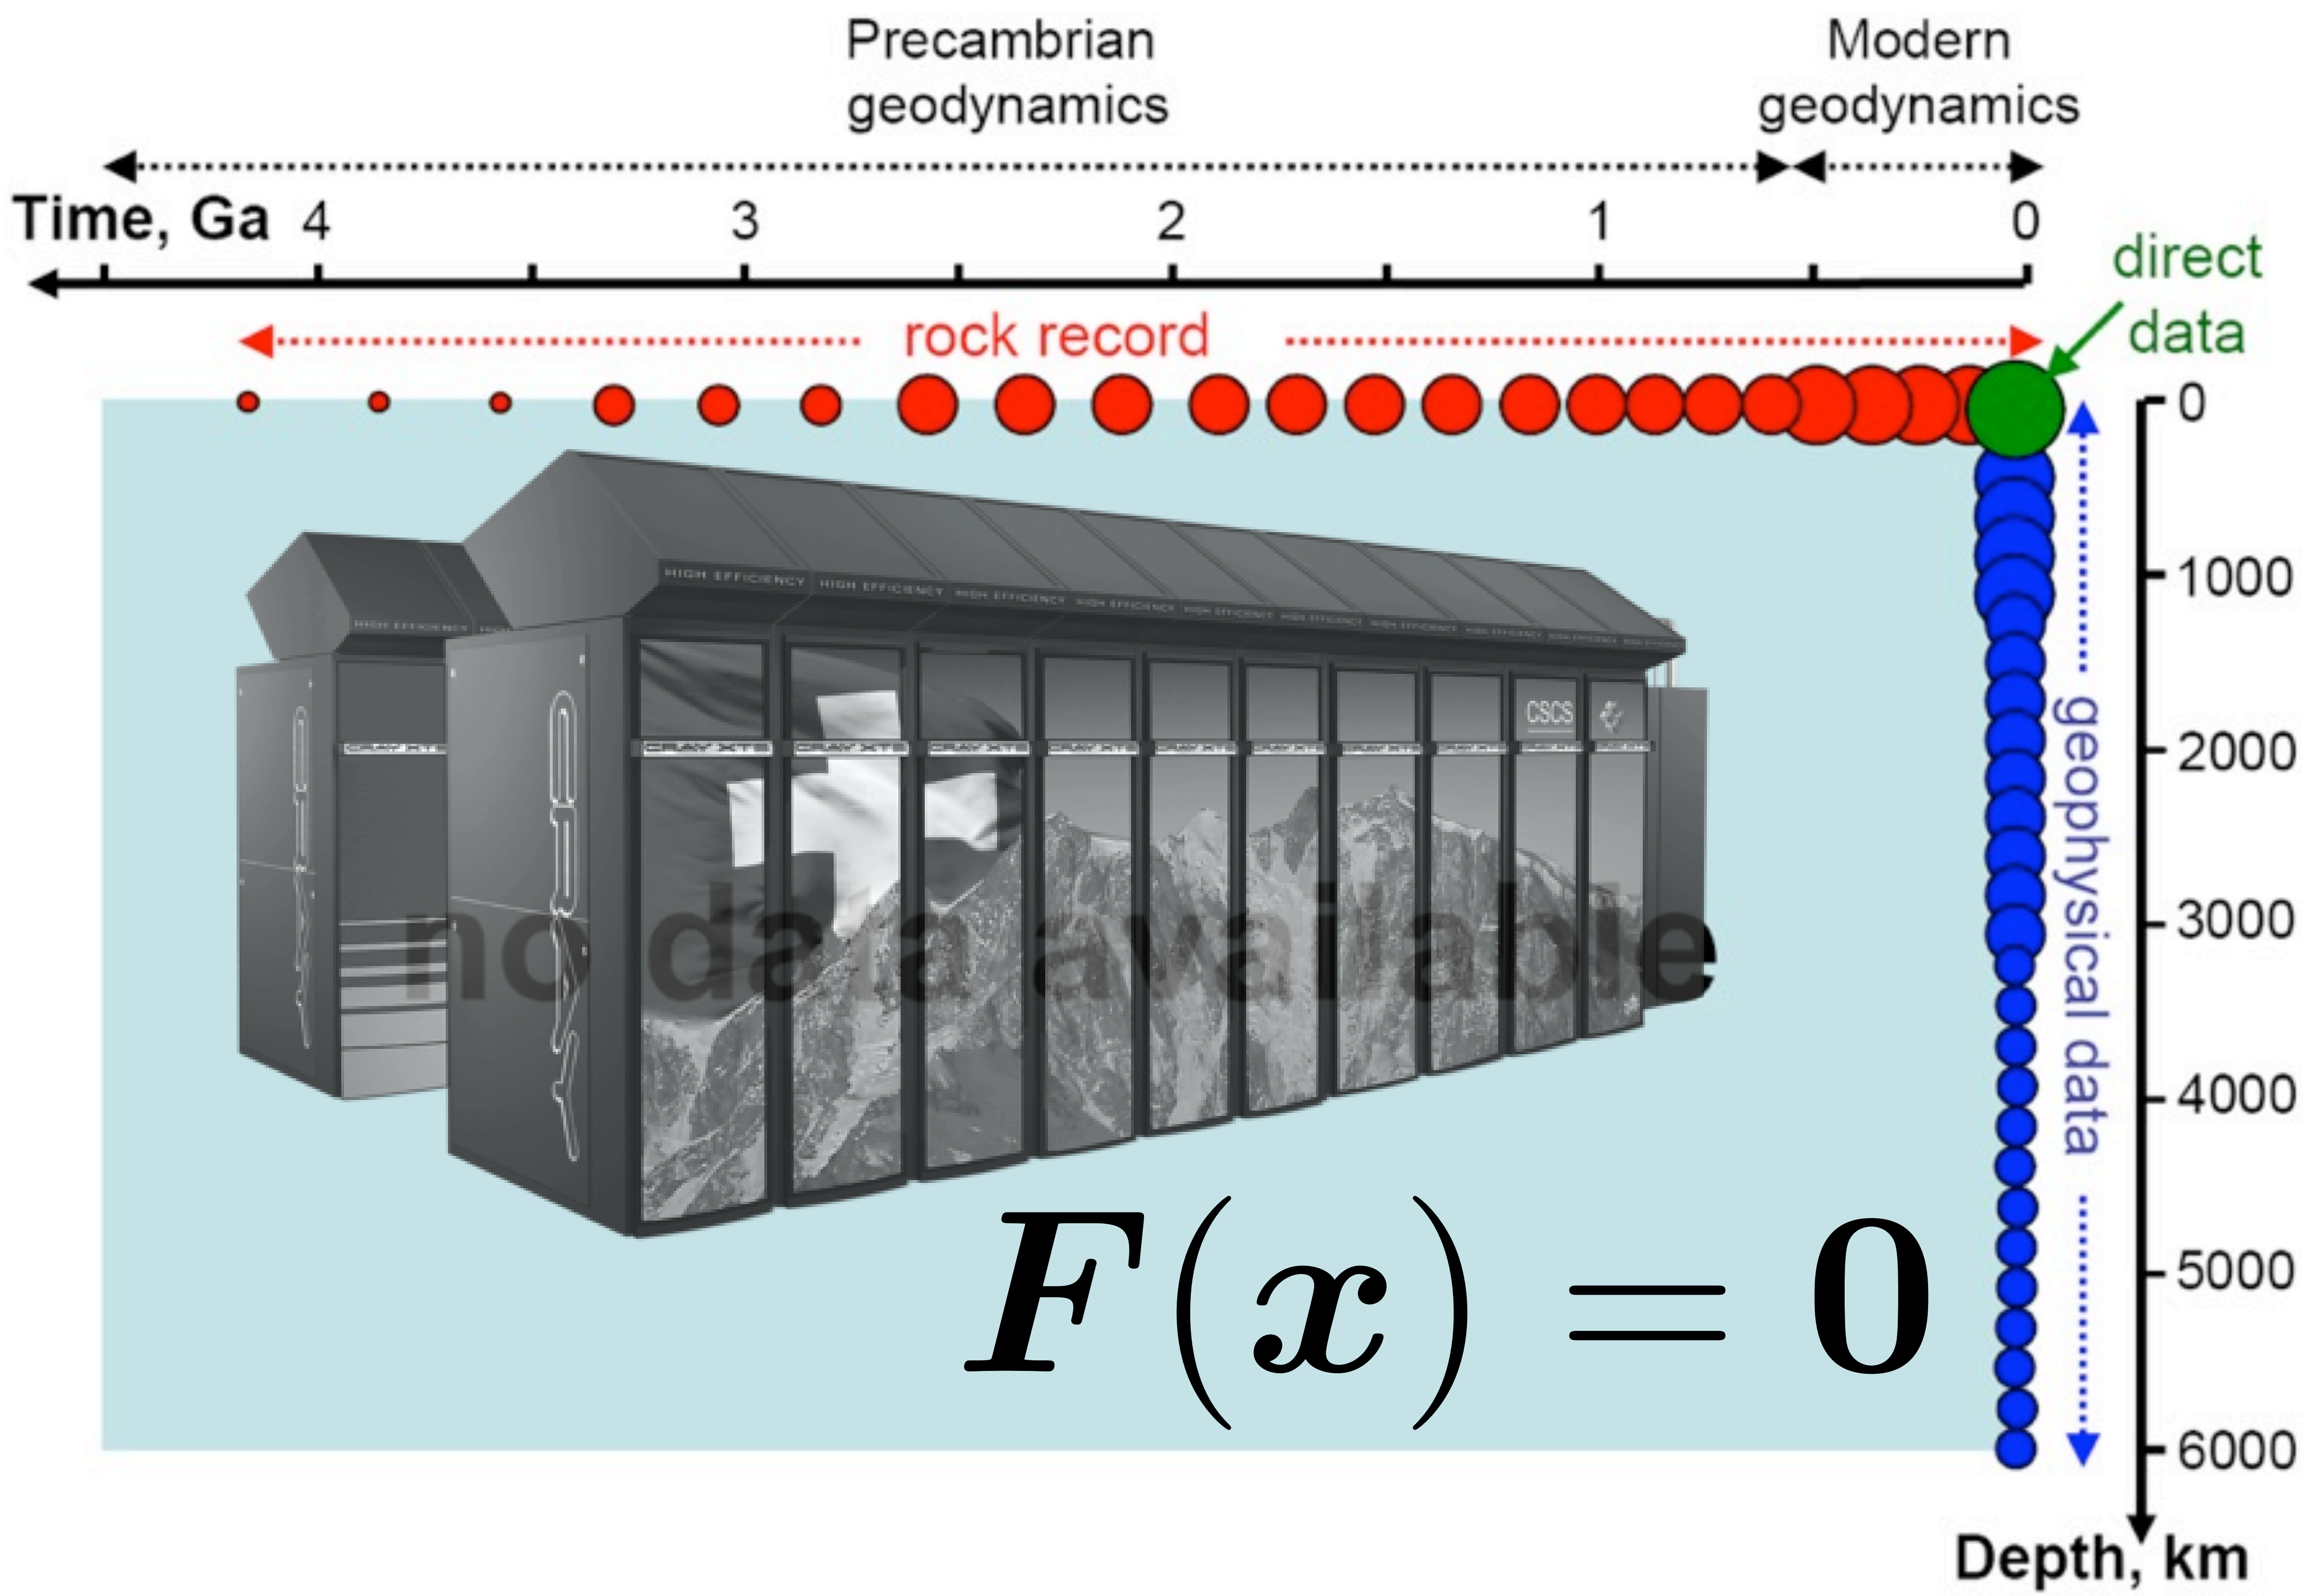

Outline

- Geodynamic background and motivations
- Spatial discretisations for long term geodynamic applications
- Newton formulation
- Scalable preconditioners for Saddle point problems
- Two-dimensional geodynamic examples
- Geometric multi-grid with matrix-free smoothers
- Three-dimensional geodynamic applications

Convective Engine of the Earth

http://asapscience.tumblr.com/post/50419005208/the-earths-center-is-out-of-sync-we-all-know

http://www.see.leeds.ac.uk/structure/dynamicearth/convection/models.htm

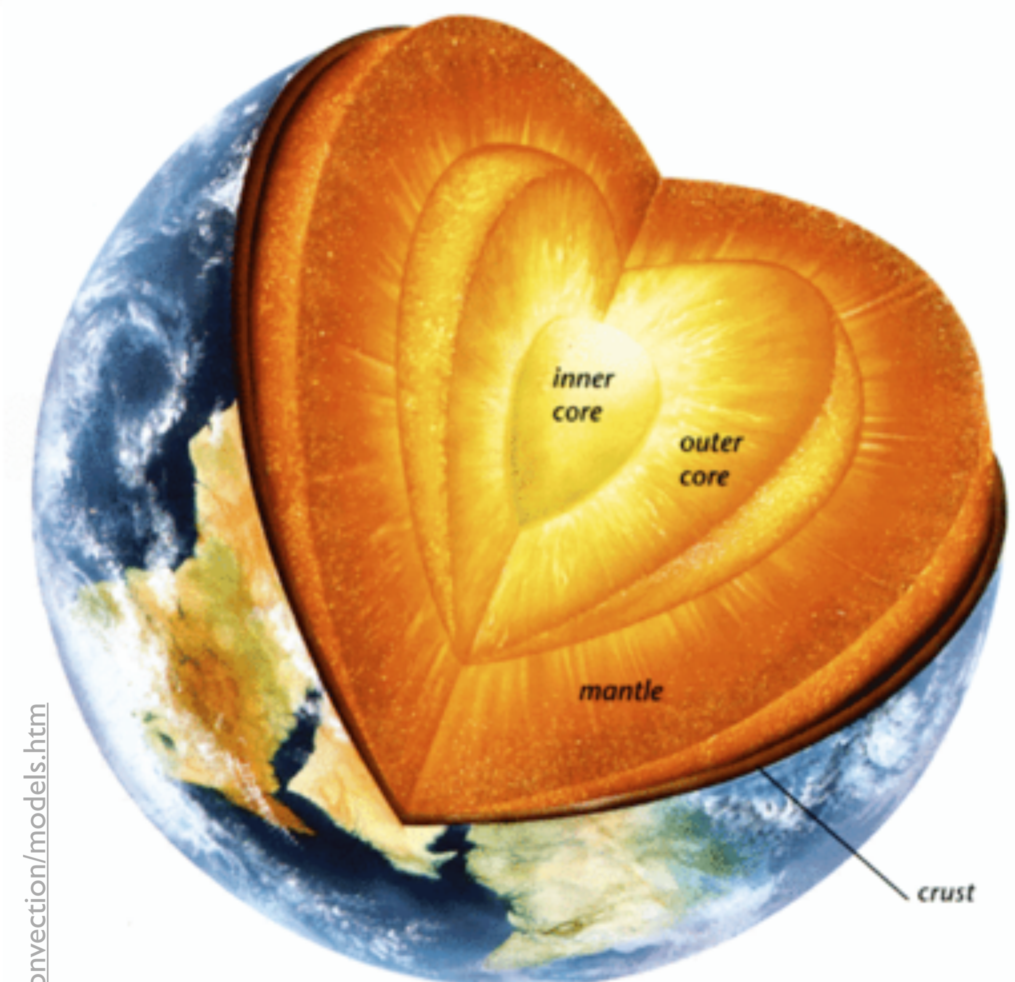

http://dreamtigers.wordpress.com/2011/05/11/plate-tectonic-metaphor-illustrations-cmu/

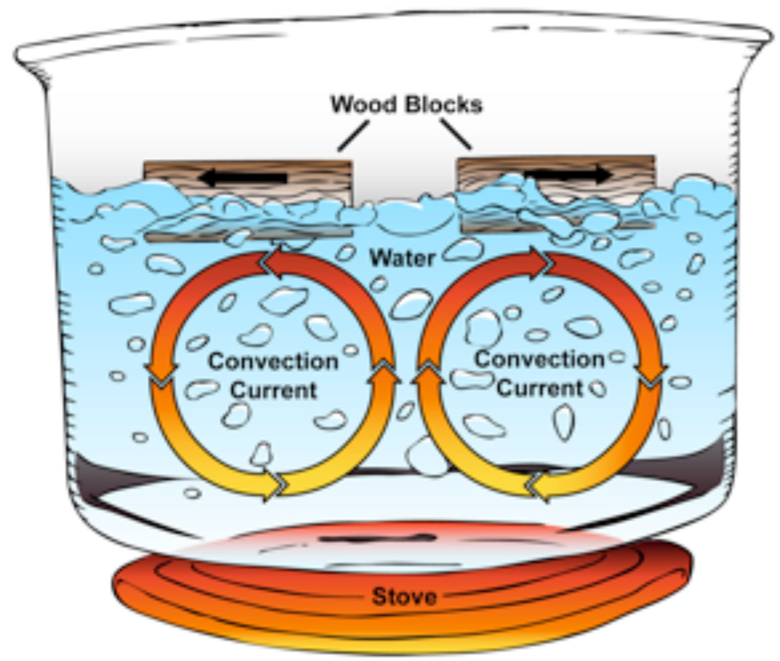

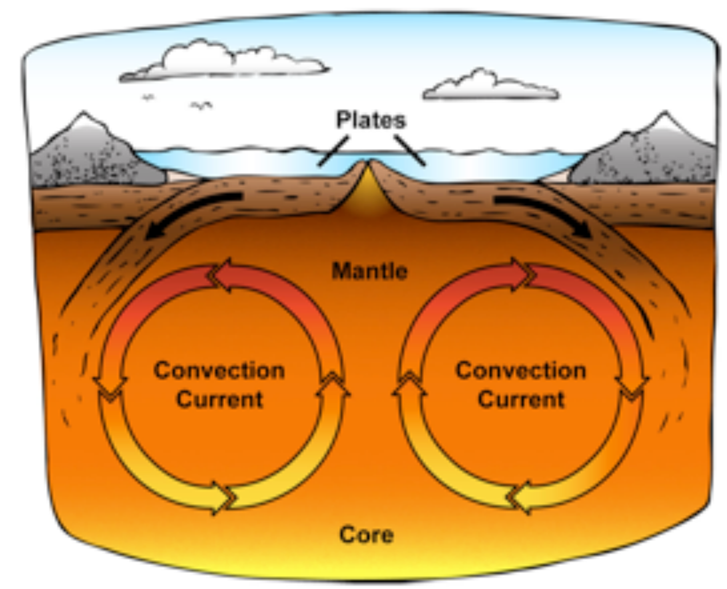

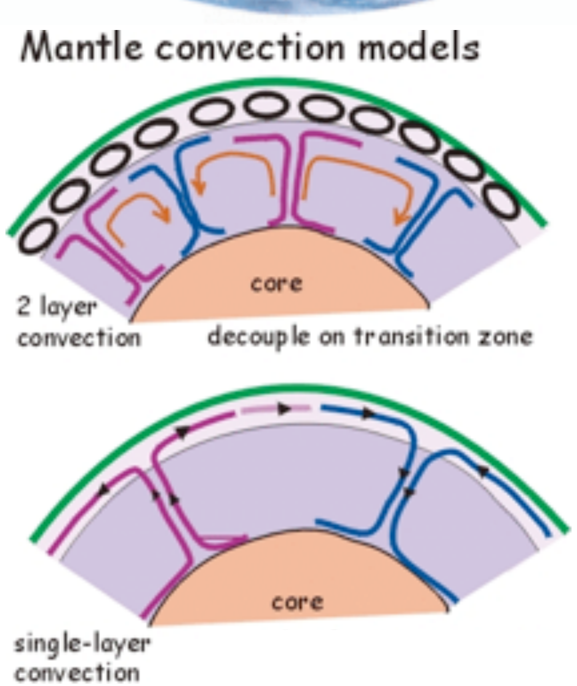

- Long time scale process. Very viscous, creeping flow regime
- Highly temperature dependent viscosity - large contrast in material properties (1e10)
- Stokes like Rayleigh-Bénard convection with strongly variable viscosity

Regional Geodynamics Processes

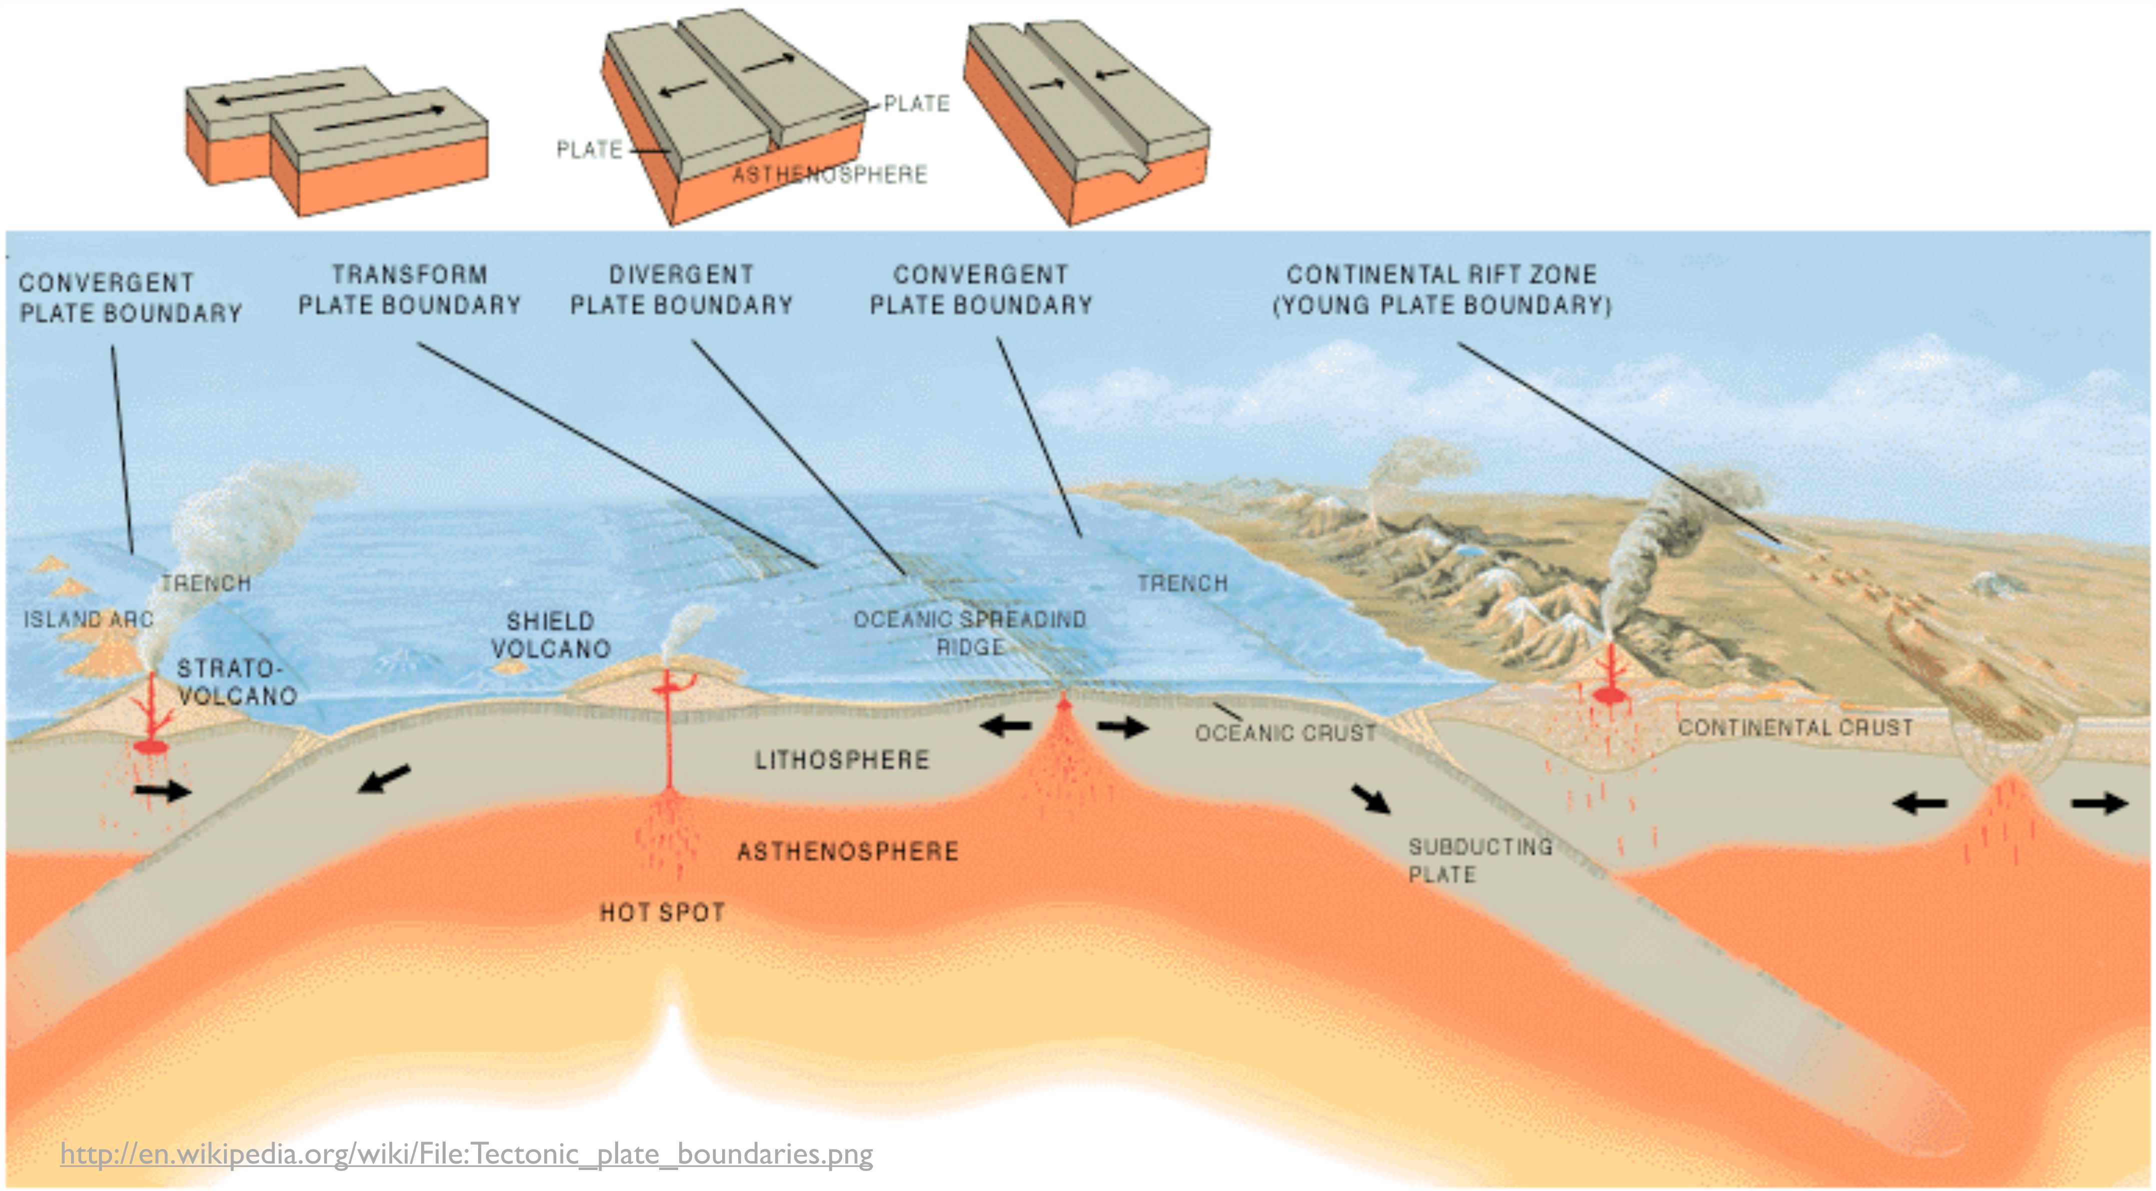

- Topography variations
- Large variation in length scales
- Presence of faults (material failure)
- Melting

- Complex constitutive behaviour
- Large deformation
- Deformation past the onset of material failure

Coupled Regional / Global Processes

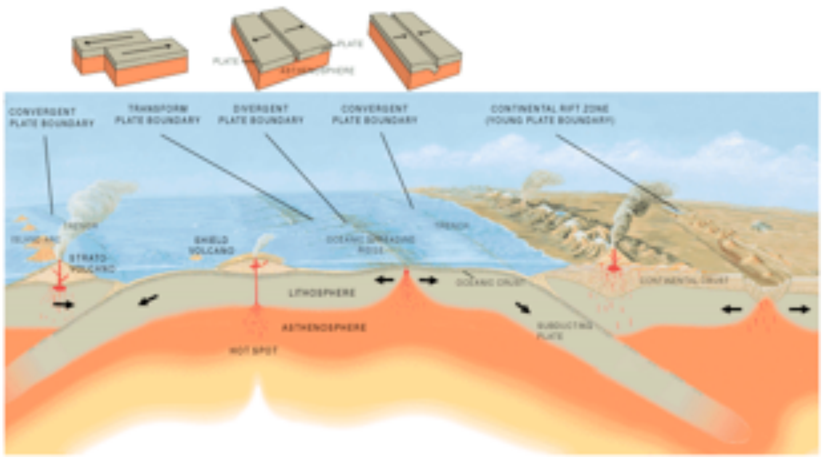

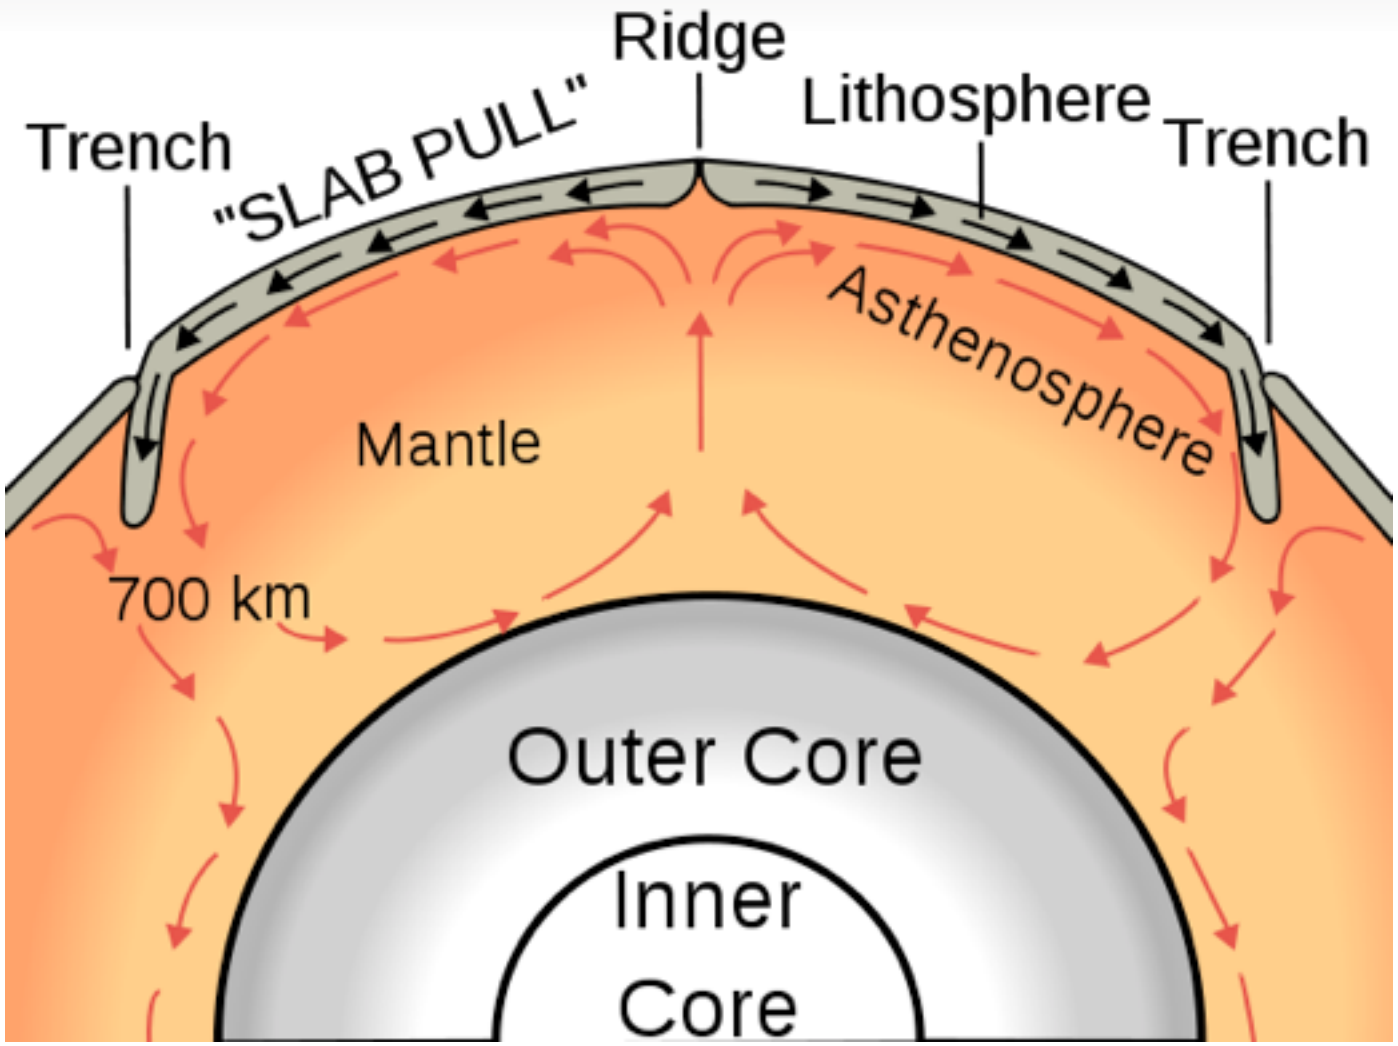

http://en.wikipedia.org/wiki/File:Oceanic_spreading.svg

Coupled Regional / Global Processes

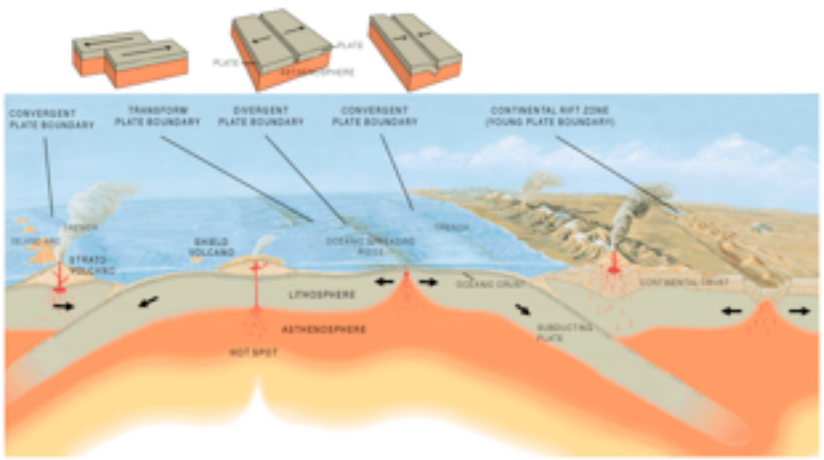

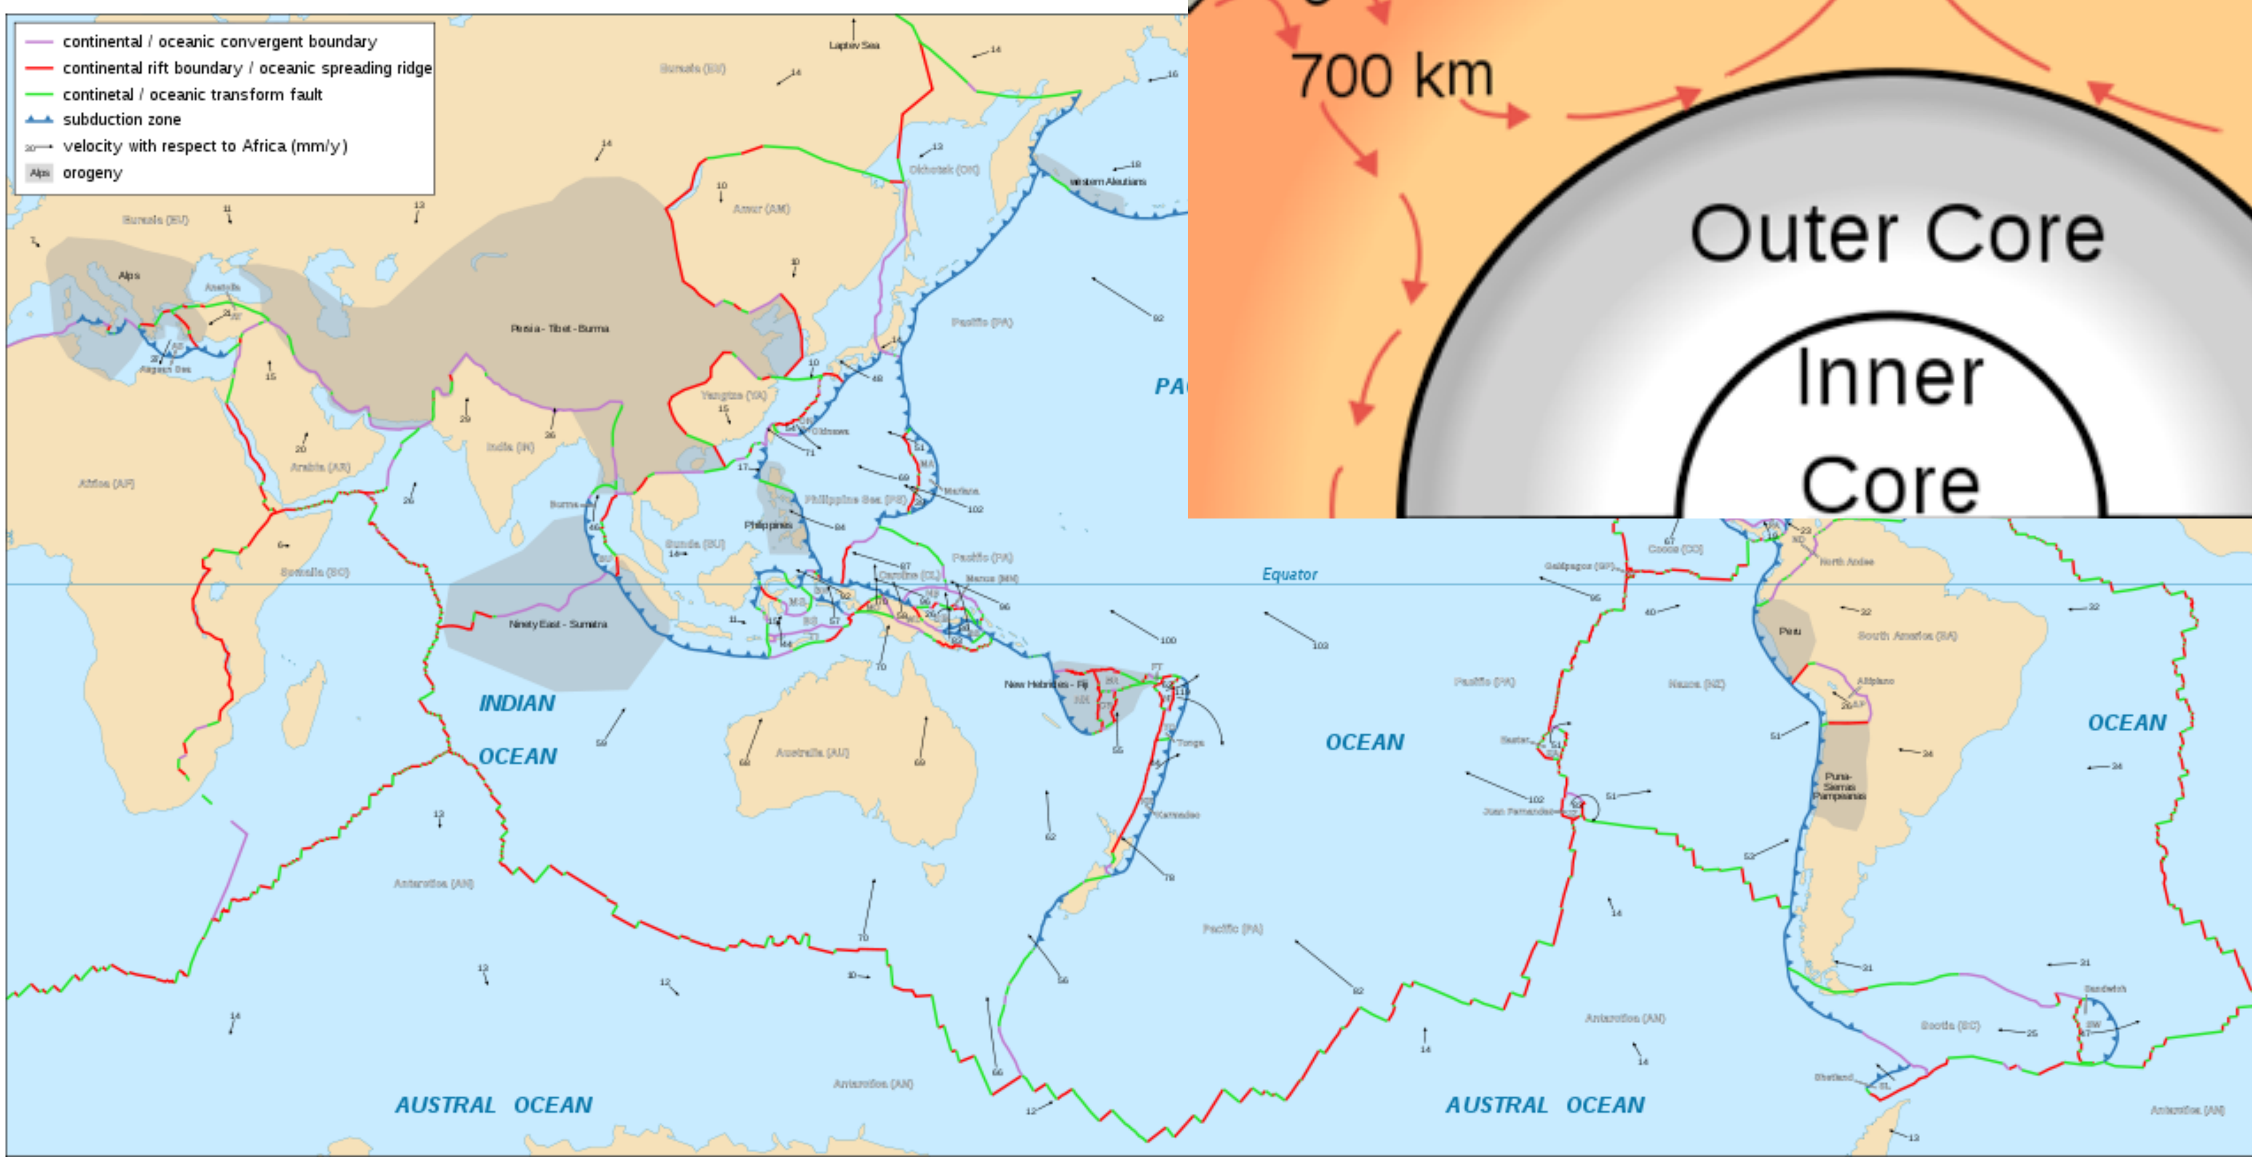

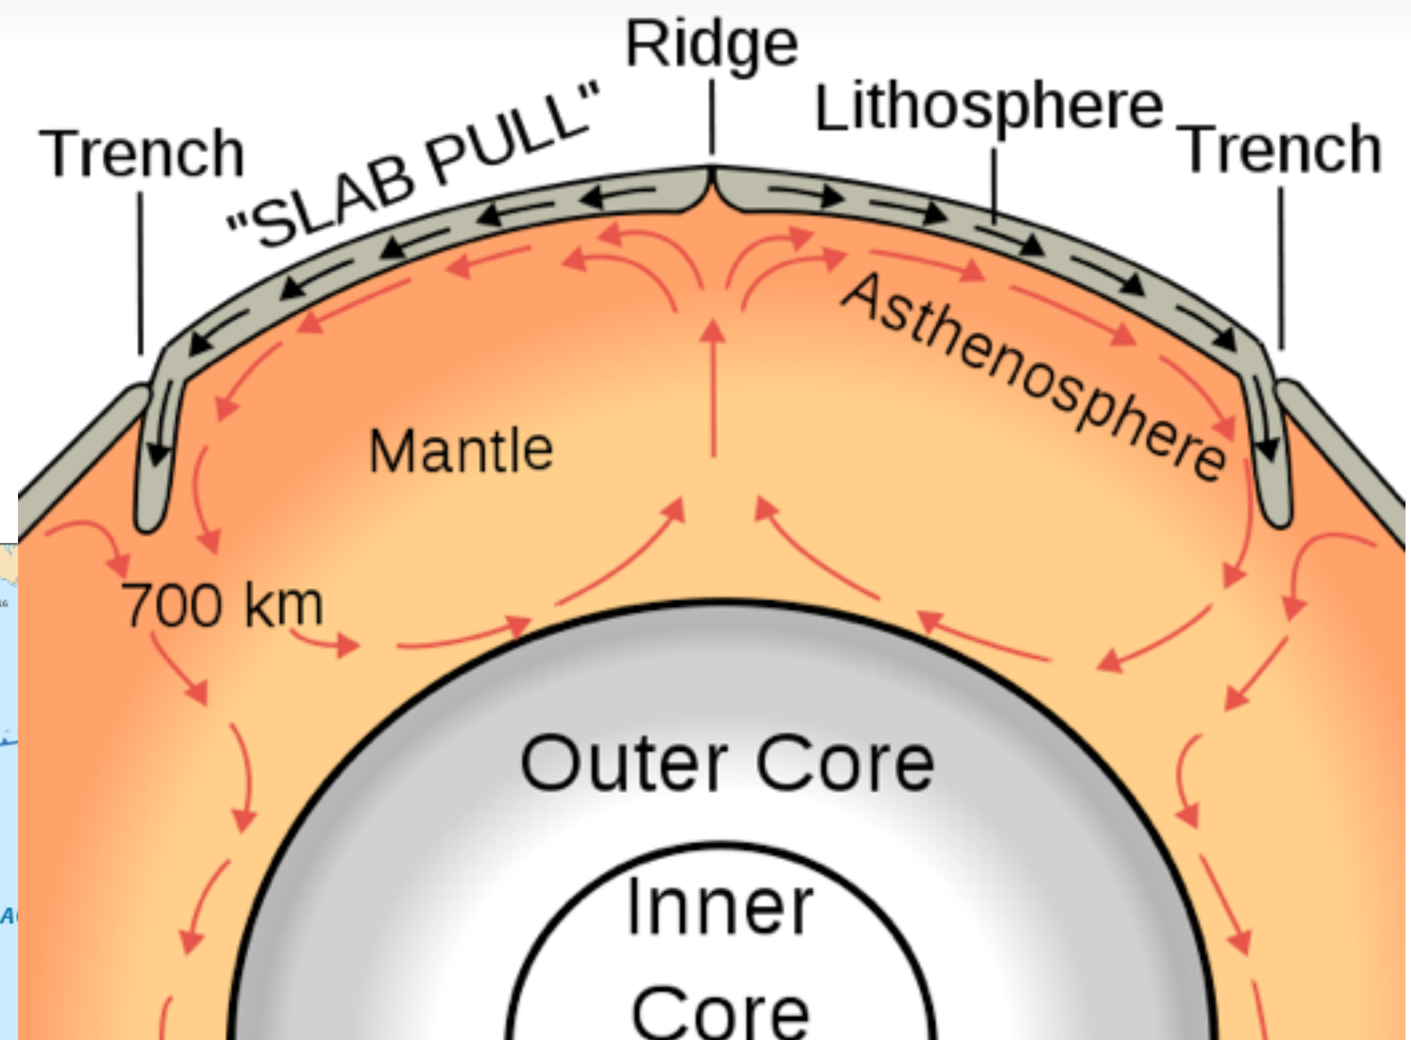

http://en.wikipedia.org/wiki/File:Oceanic_spreading.svg

Coupled Regional / Global Processes

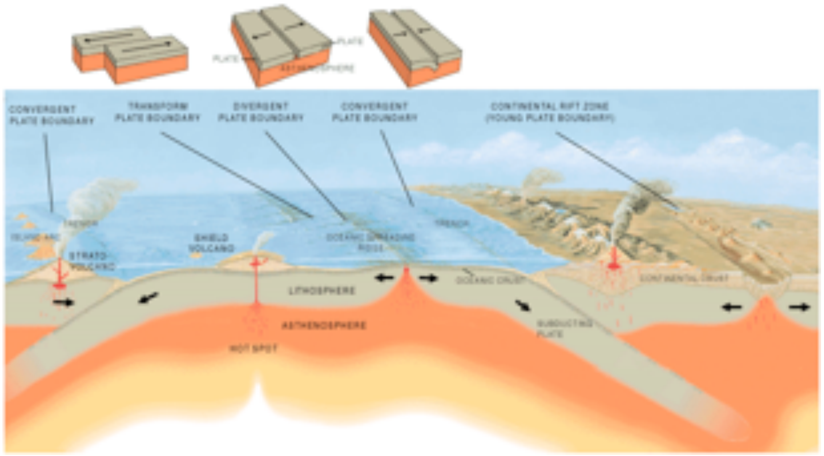

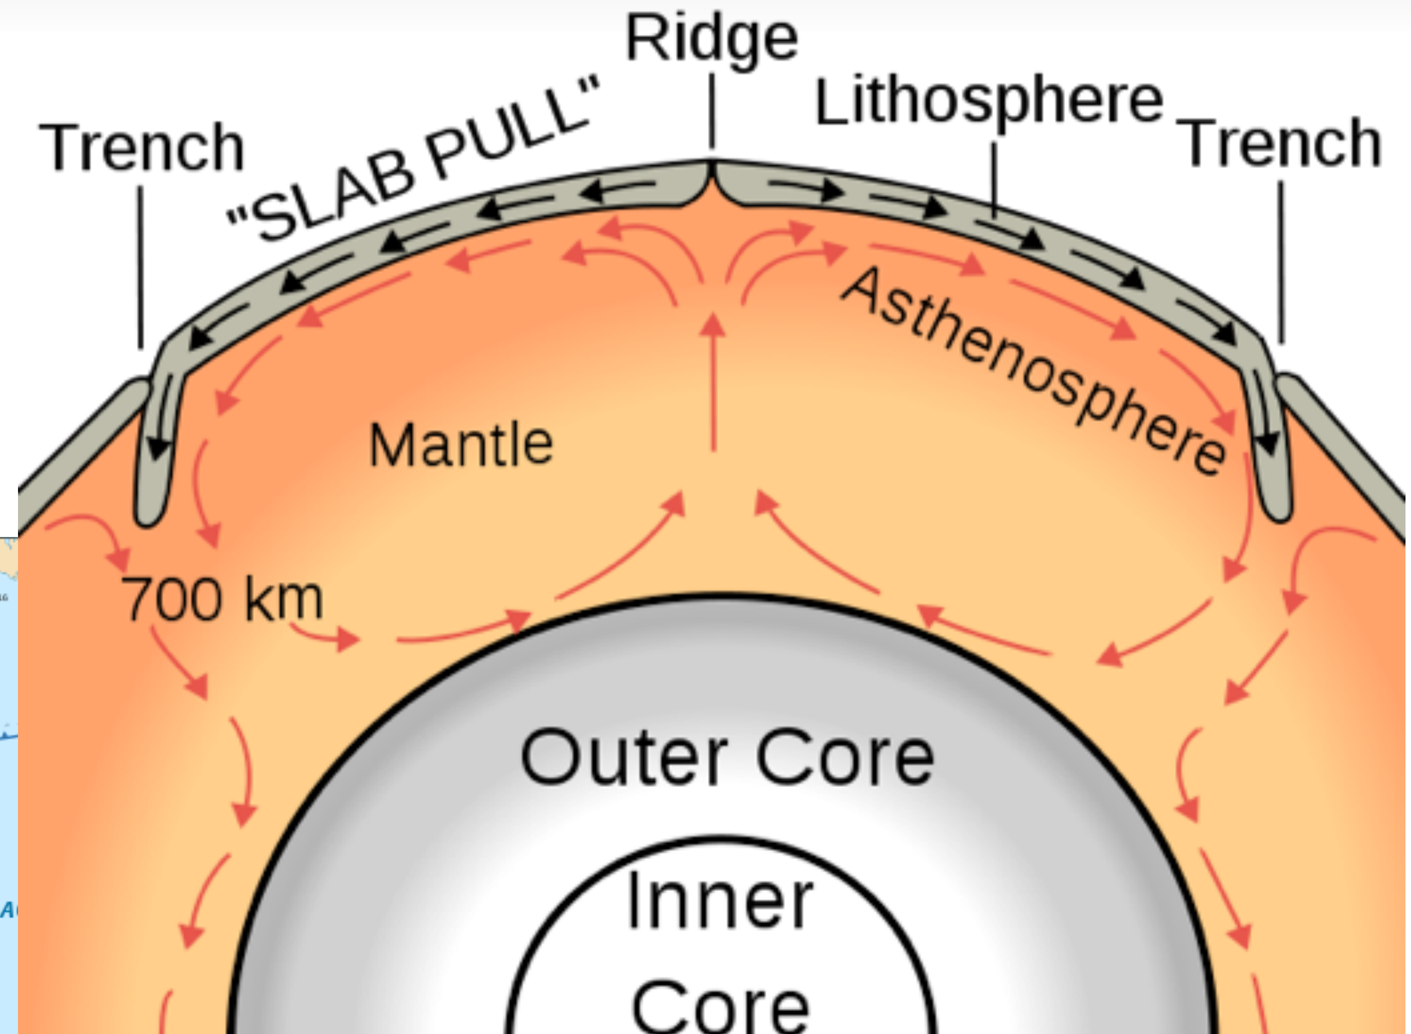

http://en.wikipedia.org/wiki/File:Oceanic_spreading.svg

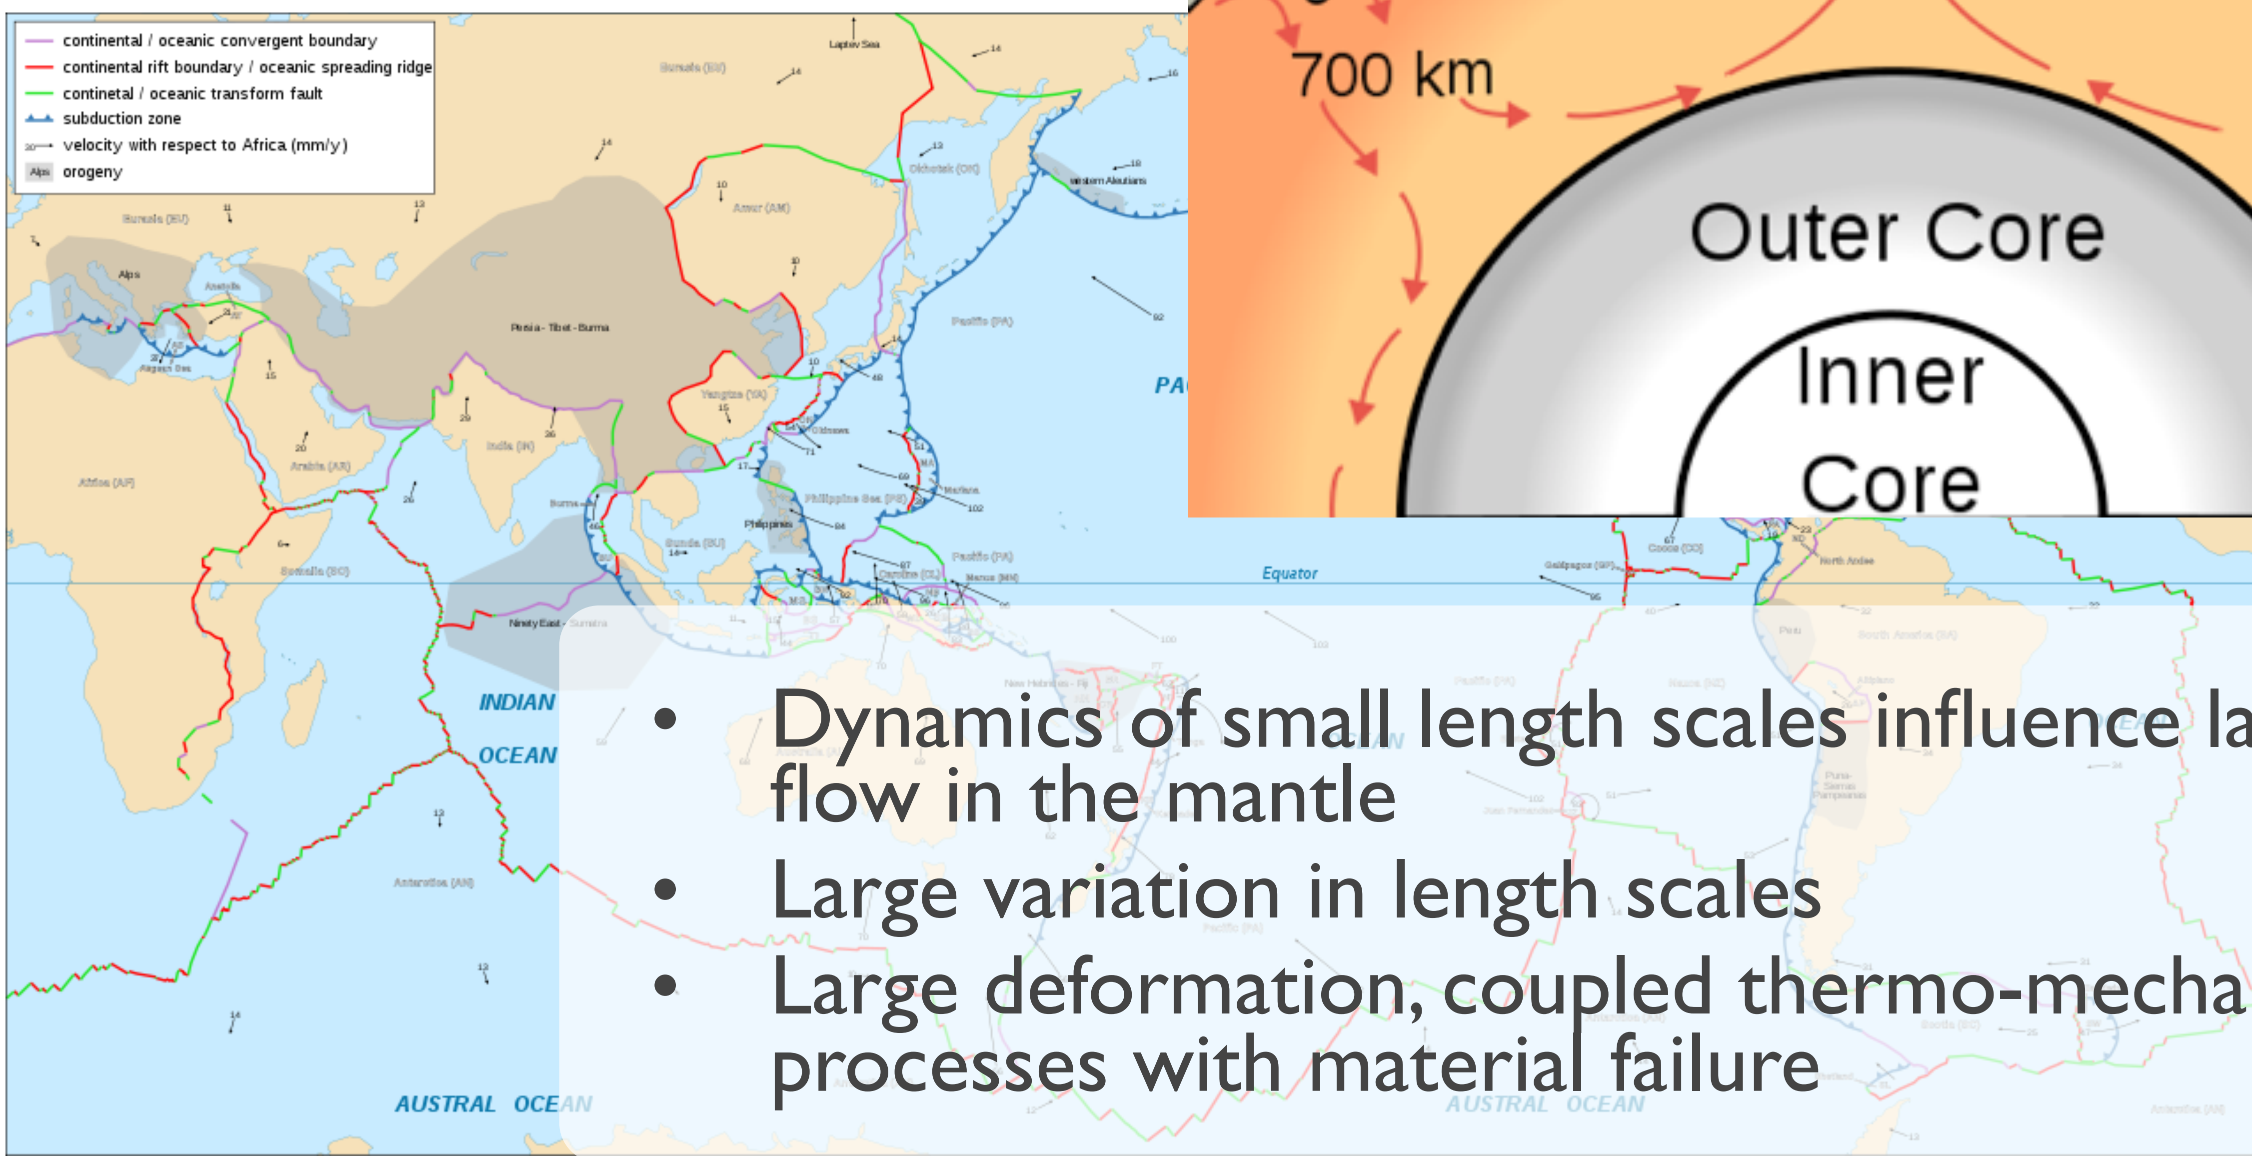

- Dynamics of small length scales influence large scale flow in the mantle
- Large variation in length scales
- Large deformation, coupled thermo-mechanical processes with material failure

Geology is Complex

- Zagros Mountains

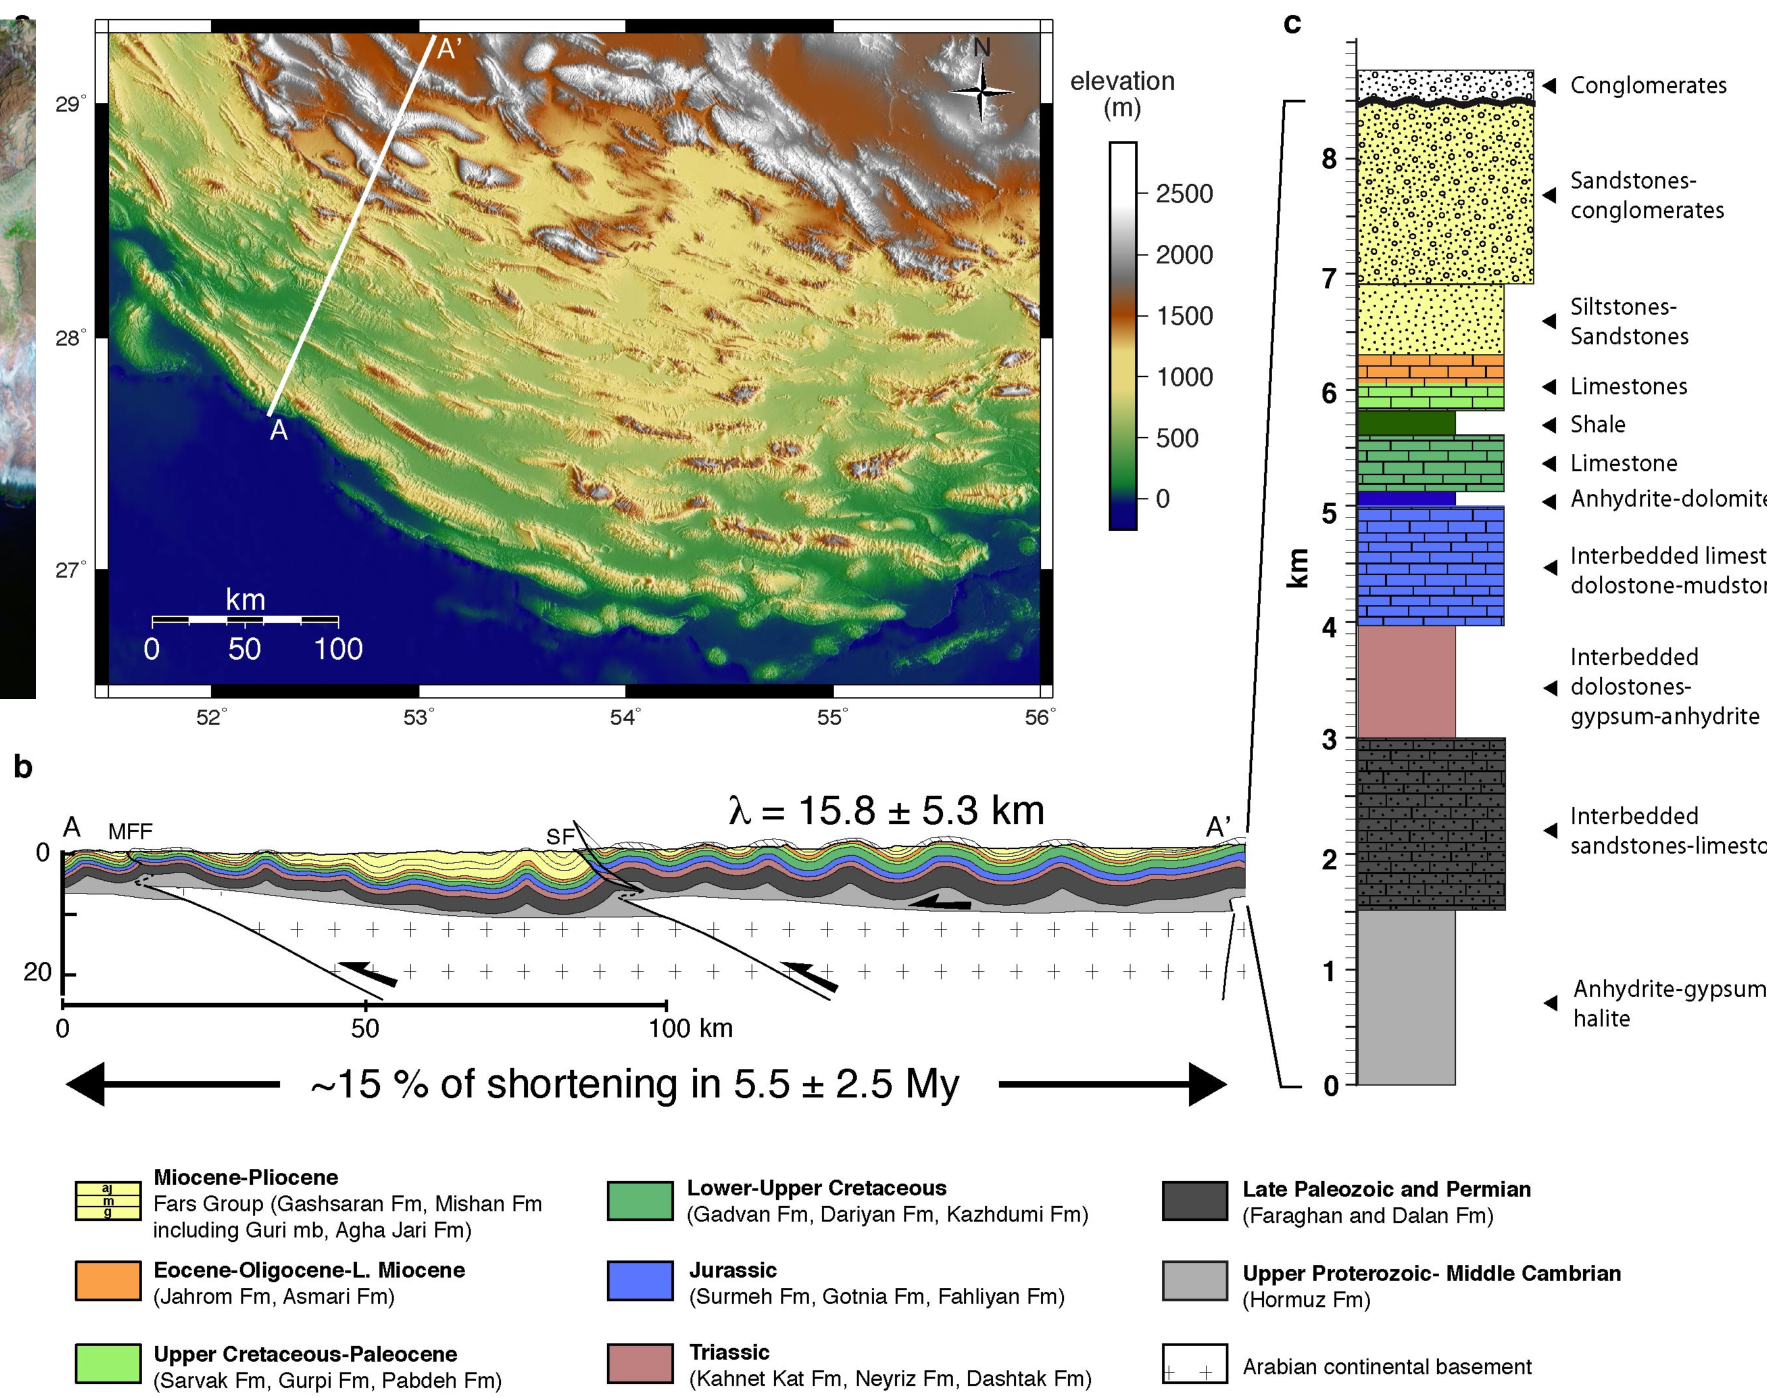

- Small and large amplitude ductile folding
- Discontinuous material properties
- Faulting
- Small length scales
- High aspect ratio

Geology is Complex

- European Alps

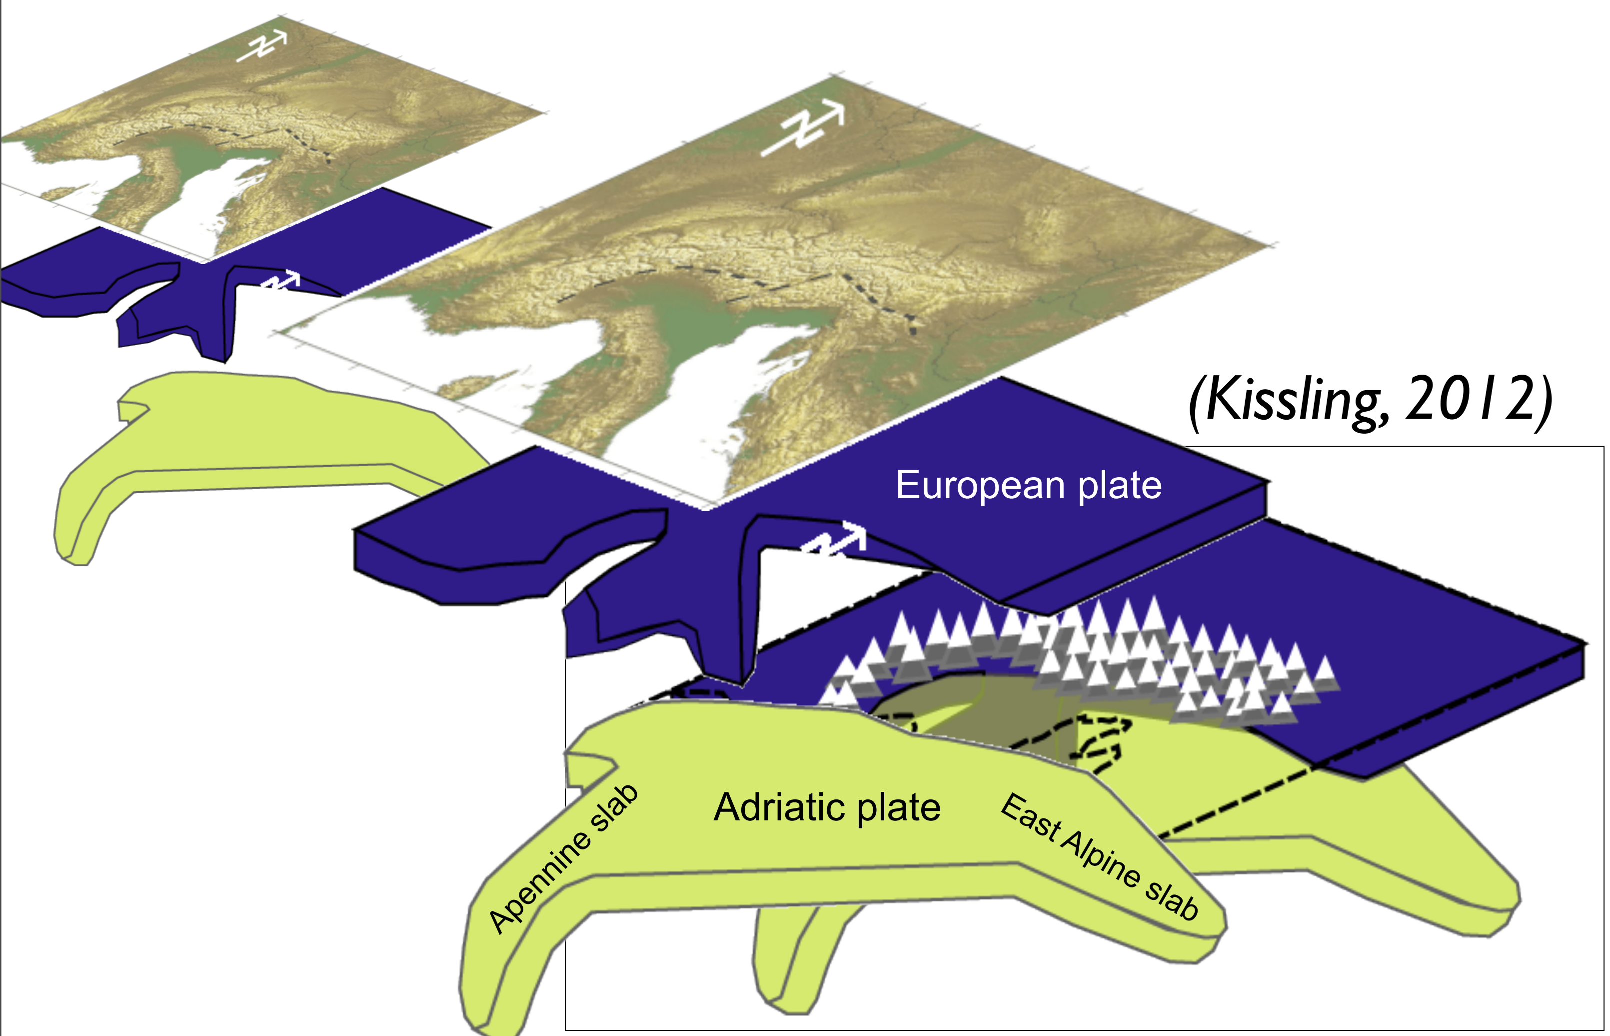

(Kissling, 2012)

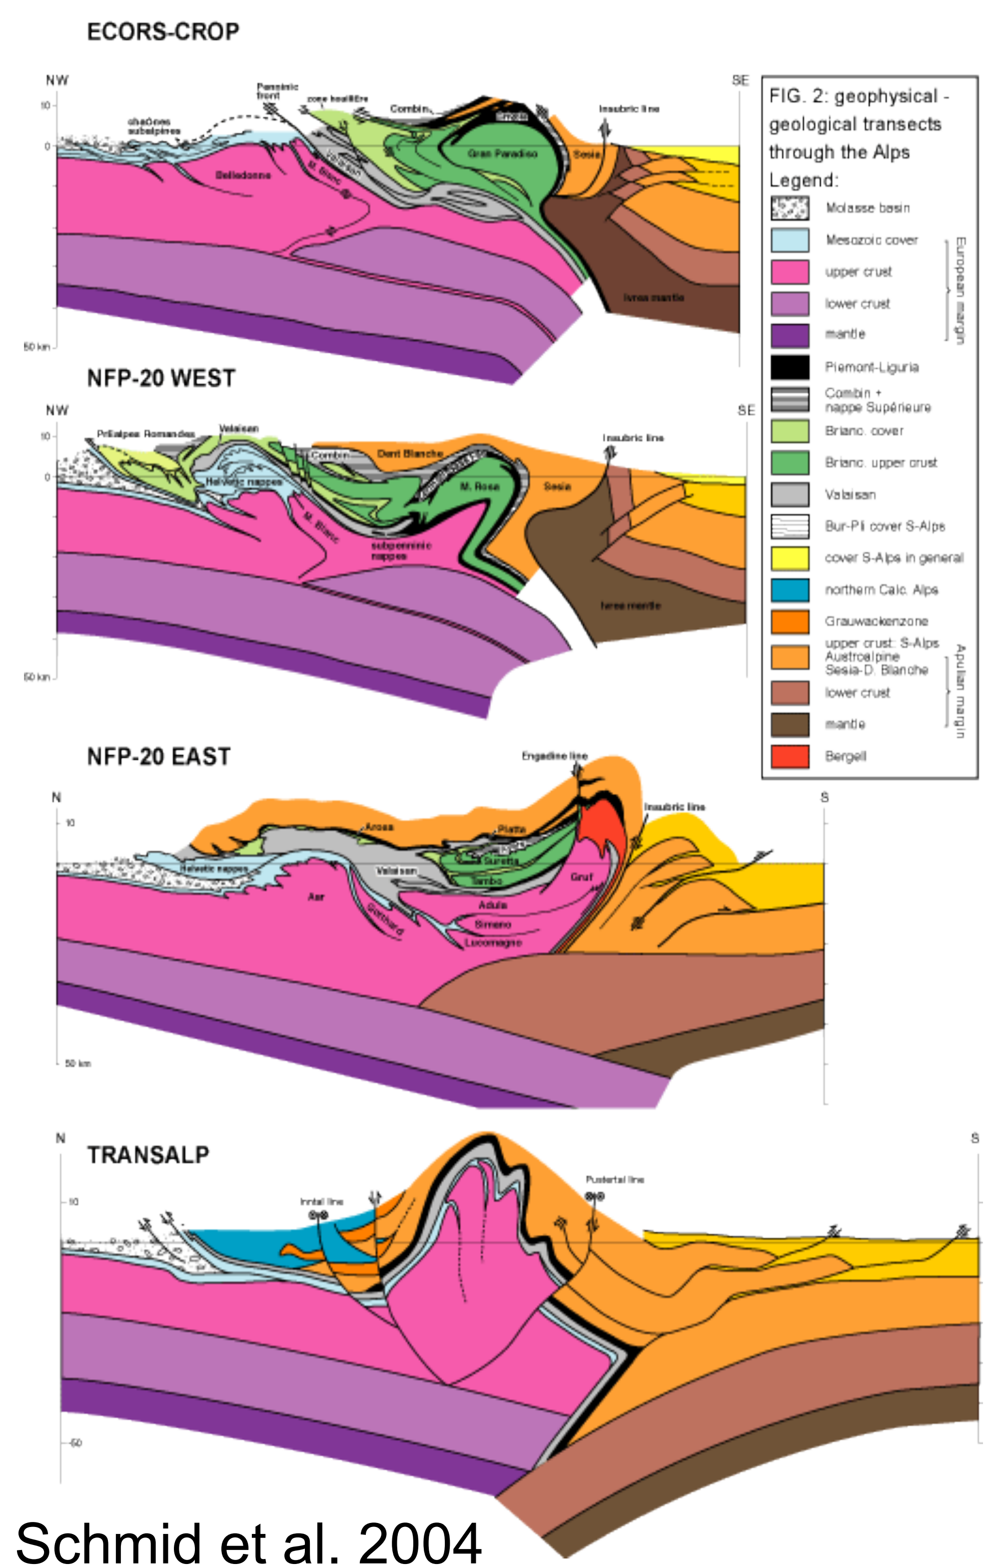

Schmid et al. 2004

- Inherently 3D
- Discontinuous properties
- Severe ductile folding + faulting
- Small length scales

"...a total mess" - even by geological standards

Geodynamic Motivations

Continental rifting

20 million years of evolution

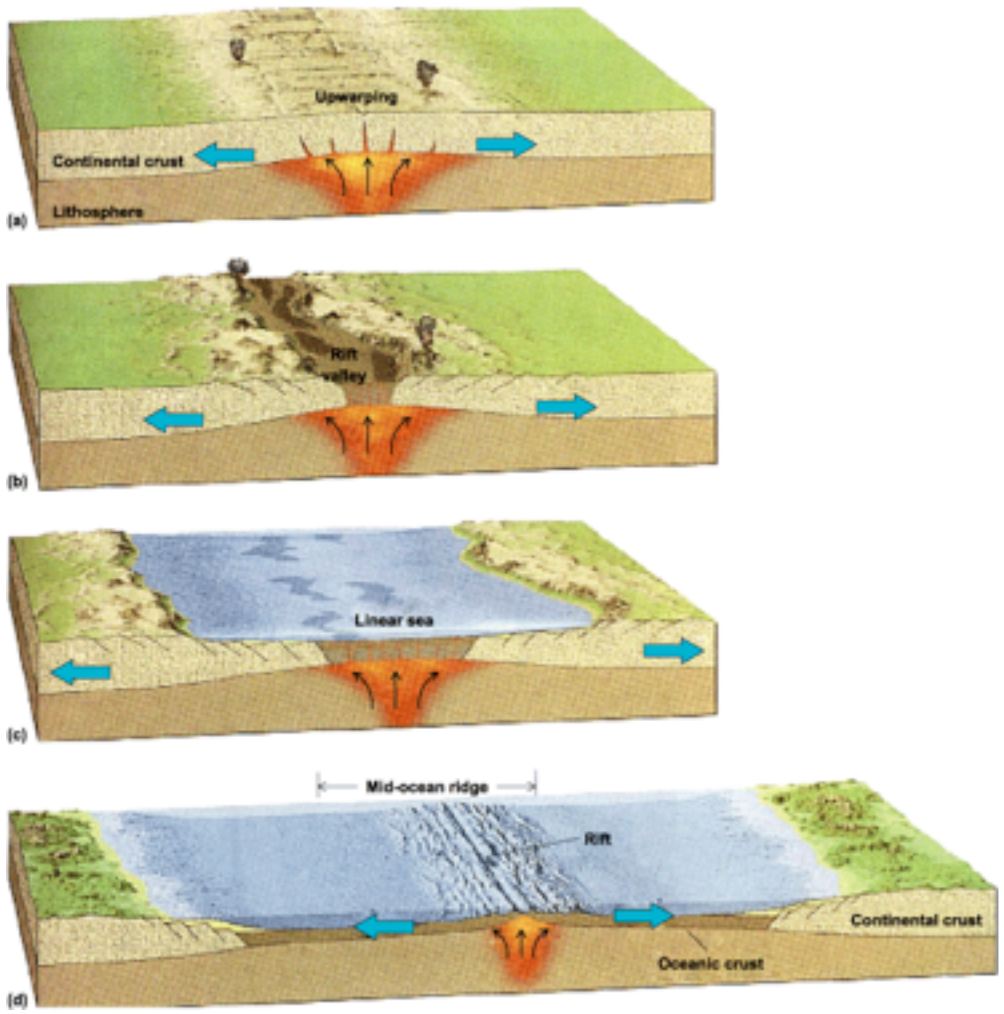

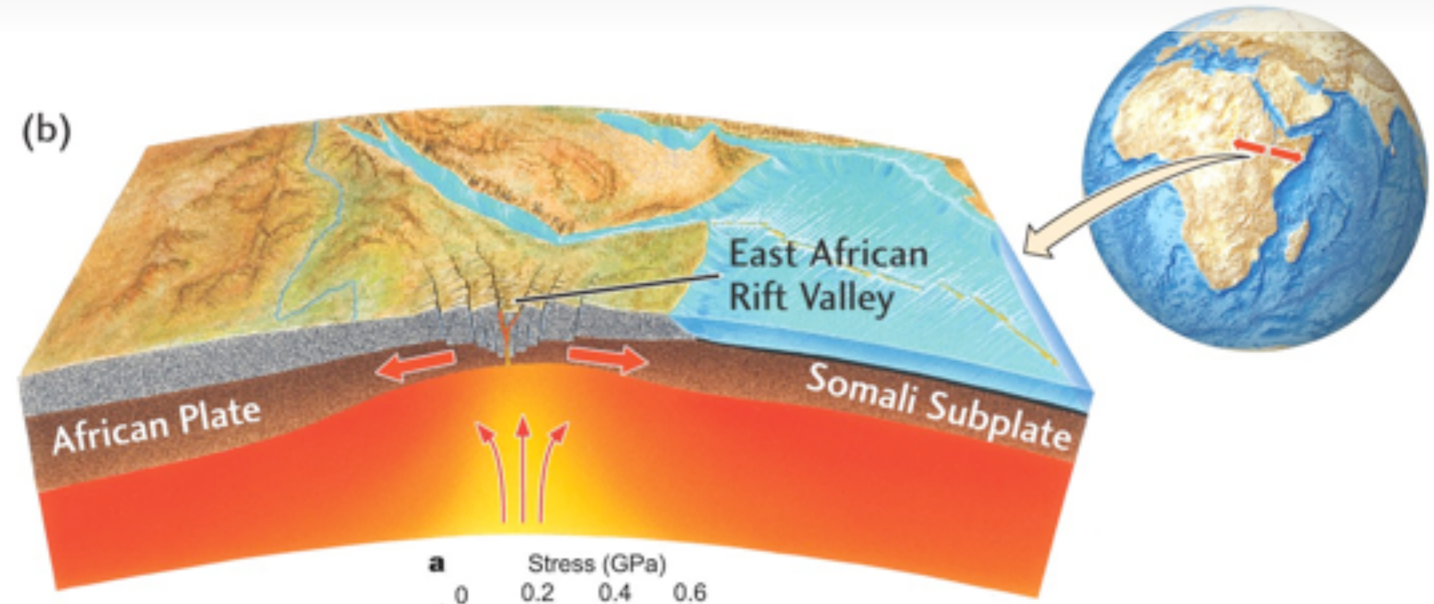

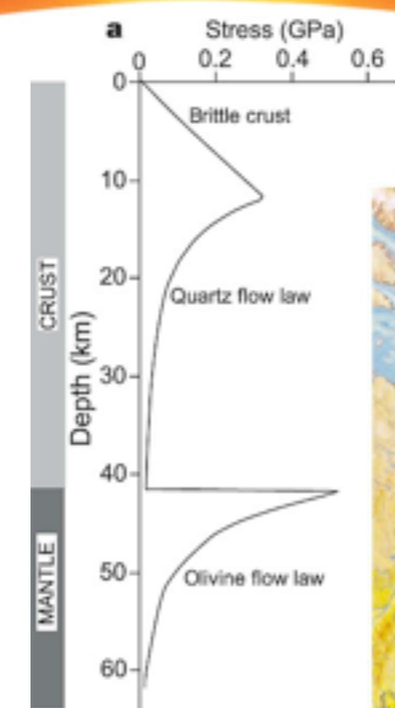

- Follow the **4D** evolution of rocks over millions year time spans ⟶ large deformation
- Complex constitutive laws
- Large contrast in material properties
- Deformation past the onset of material failure

Problem Statement

- Incompressible, Variable Viscosity (VV) Stokes:

$$\left[2\eta D_{ij}(\boldsymbol{u})\right]_{,j} - p_{,i} = f_i$$
$$u_{k,k} = 0$$

in Ω

$$u_i = \bar{u}_i \quad \text{on } \Gamma_D$$
$$\sigma_{ij} n_j = \bar{t}_i \quad \text{on } \Gamma_N$$

- **Non-linear constitutive behaviour**

- Evolution of coefficients

$$\frac{D\eta}{Dt} = 0, \qquad \frac{Df_i}{Dt} = 0$$

- Non-linear boundary conditions

- Conservation of Energy: $\frac{DT}{Dt} = \left[\kappa T_{,k}\right]_{,k} + Q$

Problem Statement: Coefficients

- Incompressible, Variable Viscosity (VV) Stokes:

$$\left[2\eta D_{ij}(\boldsymbol{u})\right]_{,j} - p_{,i} = f_i \qquad \text{in } \Omega$$
$$u_{k,k} = 0$$

- **Non-linear constitutive behaviour** (η)

Arrhenius [u,p,T dependence]

$$\eta = A(\sqrt{I_2'})^{\alpha} \exp\left(\frac{E+Vp}{nRT}\right) \qquad I_2' = \tfrac{1}{2} D_{ij} D_{ij}$$

Plasticity [u,p dependence]

$$F_s := \sqrt{J_2'} - \tau_{\text{yield}}^{VM}, \qquad \text{where } \tau_{\text{yield}}^{VM} := \text{const.} \qquad J_2' = \tfrac{1}{2}\tau_{ij}\tau_{ij}$$

$$F_s := \sqrt{J_2'} - \tau_{\text{yield}}^{DP}, \qquad \text{where } \tau_{\text{yield}}^{DP} := C_0 \cos(\phi) + p \sin(\phi),$$

$$\eta = \frac{\tau_{\text{yield}}}{2\sqrt{I_2'}} \qquad \text{if } \sqrt{J_2'} > \tau_{\text{yield}},$$

- **Boussinesq approximation** (f_i)

$$f_i = \rho_0 \left[1 - \alpha(T - T_0)\right] g_i$$

Problem Statement

- Incompressible, Variable Viscosity (VV) Stokes:

$$\left[2\eta D_{ij}(\boldsymbol{u})\right]_{,j} - p_{,i} = f_i \quad \text{in } \Omega$$
$$u_{k,k} = 0$$
$$u_i = \bar{u}_i \quad \text{on } \Gamma_D$$
$$\sigma_{ij} n_j = \bar{t}_i \quad \text{on } \Gamma_N$$

- Non-linear constitutive behaviour
- Evolution of coefficients

$$\frac{D\eta}{Dt} = 0, \qquad \frac{Df_i}{Dt} = 0$$

Use independent spatial discretisations for

(i) the flow variables (velocity, pressure)

$\longrightarrow$ Mixed FEM [Q2-P1]

(ii) coefficients (viscosity, density)

$\longrightarrow$ Lagrangian markers

(aka Material Point Method)

Spatial Discretisation (FE)

Discrete Variational Problem (with interpolated coefficients)

Seek $(\boldsymbol{u}_h, p_h) \in \boldsymbol{V}^h \times Q^h$ with $\eta^\lambda \in X$ and $\boldsymbol{f}^\lambda \in (X)^d$ such that

$$\left.\begin{aligned} A(\boldsymbol{u}_h, \boldsymbol{v}_h) + B(\boldsymbol{v}_h, p_h) &= F(\boldsymbol{v}_h) \\ B(\boldsymbol{u}_h, q_h) &= 0 \end{aligned}\right\} \quad \text{for all } (\boldsymbol{v}_h, q_h) \in \boldsymbol{V}_0^h \times Q^h.$$

$$A(\boldsymbol{u}, \boldsymbol{v}) = \int_\Omega \sum_{i,j=1}^{d} 2\eta^\lambda D_{ij}(\boldsymbol{u}) D_{ij}(\boldsymbol{v})\, dV,$$

$$B(\boldsymbol{v}, q) = \int_\Omega q \nabla \cdot \boldsymbol{v}\, dV,$$

$$F(\boldsymbol{v}) = \int_\Omega \boldsymbol{v} \cdot \boldsymbol{f}^\lambda\, dV + \int_{\Gamma_N} \boldsymbol{v} \cdot \bar{\boldsymbol{t}}\, dS.$$

$$\boldsymbol{V} := (V)^d = \left\{ \boldsymbol{v} \in (H^1(\Omega))^d \,|\, \boldsymbol{v} = \bar{\boldsymbol{u}} \text{ on } \Gamma_D \right\},$$

$$\boldsymbol{V}_0 := (V_0)^d = \left\{ \boldsymbol{v} \in (H^1(\Omega))^d \,|\, \boldsymbol{v} = \boldsymbol{0} \text{ on } \Gamma_D \right\},$$

$$Q := \left\{ q \in L_2(\Omega) \;:\; \int_\Omega q\, dV = 0 \right\},$$

$$X := \left\{ x \in L_2(\Omega) \right\},$$

- Reconstruct coefficients (viscosity, density) for the flow problem using material points

Spatial Discretisation (MPM)

- Reconstruct coefficients (viscosity, density) at quadrature points for the flow problem using material points

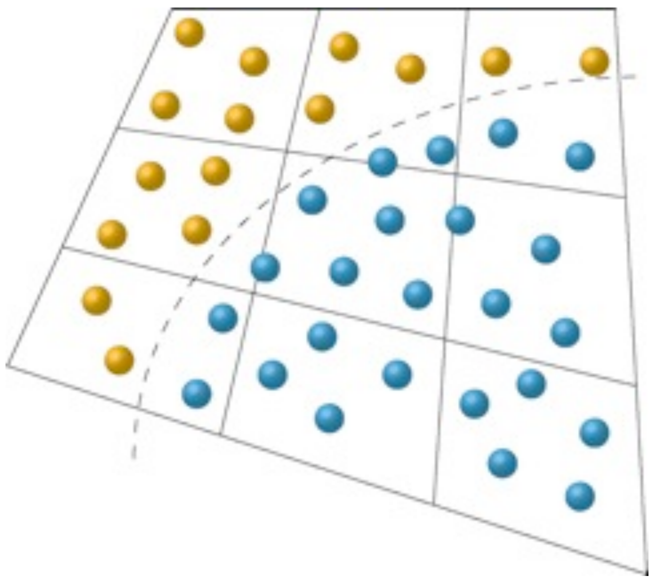

Particle In Cell (PIC) *Harlow & Welch, Phys. Fluids, (1965)*

Material Point Method (MPM) *Sulsky & Brackbill, JCP, (1991)*

$$A(\boldsymbol{u}, \boldsymbol{v}) = \int_\Omega \sum_{i,j=1}^{d} 2\eta^\lambda D_{ij}(\boldsymbol{u}) D_{ij}(\boldsymbol{v})\, dV,$$

$$F(\boldsymbol{v}) = \int_\Omega \boldsymbol{v} \cdot \boldsymbol{f}^\lambda\, dV + \int_{\Gamma_N} \boldsymbol{v} \cdot \bar{\boldsymbol{t}}\, dS.$$

[a] Local L2 projection (Q1)

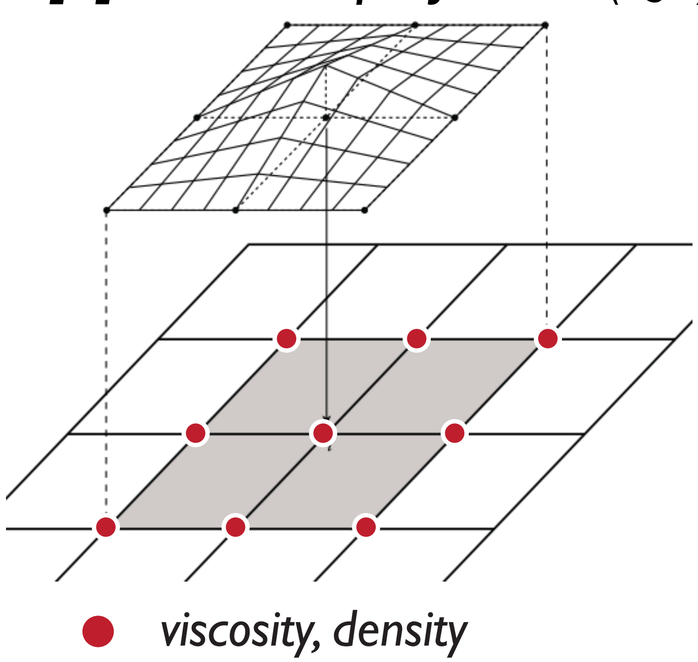

[b] Piecewise constant (P0)

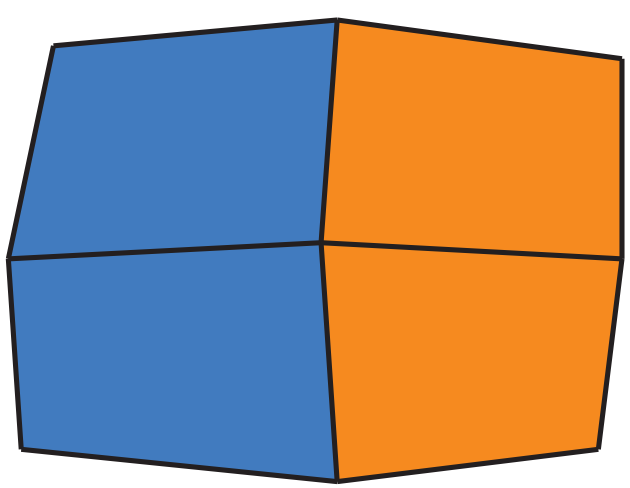

[c] Piecewise linear (P1)

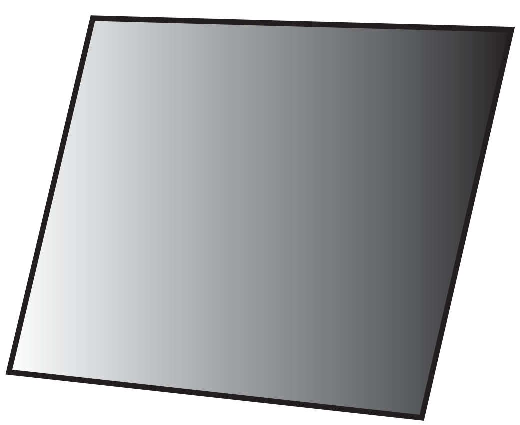

Newton Framework

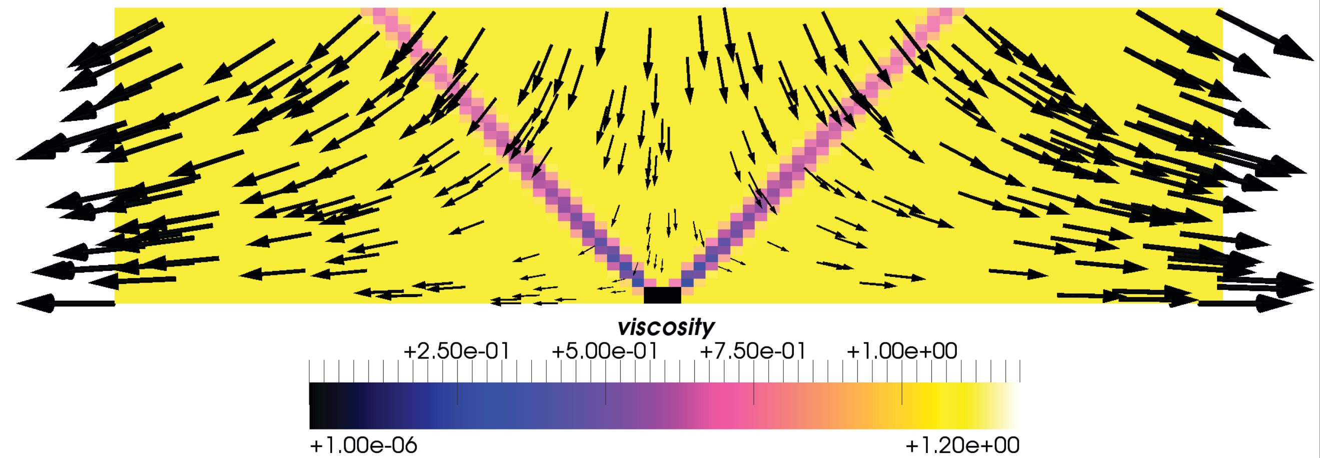

STOKES NON-LINEAR RESIDUALS

$$\hat{F}_{u_i} := \left[2\eta(\boldsymbol{u},p)D_{ij}(\boldsymbol{u})\right]_{,j} - p_{,i} - f_i(\boldsymbol{u},p)$$
$$\hat{F}_c := u_{k,k}$$

DISCRETISE

$$F_u = Au - Bp - f$$
$$F_c = B^T u$$

LINEARISE

$$\begin{bmatrix} A + \delta A & B + \delta B \\ B^T + \delta B^T & 0 \end{bmatrix} \begin{bmatrix} \delta u \\ \delta p \end{bmatrix} = - \begin{bmatrix} F_u \\ F_c \end{bmatrix}$$

STOKES JACOBIAN

$$\mathcal{J}_s = \begin{bmatrix} A + \delta A & B + \delta B \\ B^T + \delta B^T & 0 \end{bmatrix}$$

void FormFunction(Vec X,void *ctx) {

- `Extract u,p from X`
- `Update nonlinearities on markers`

$$f := \tau_{II} - \tau_y \leq 0$$

$$\tau_{II} := \sqrt{\tfrac{1}{2}\tau_{ij}\tau_{ij}}$$

$$\eta_{vp} = \begin{cases} \dfrac{\tau_y}{\sqrt{2\epsilon_{ij}\epsilon_{ij}}} & \text{if } \tau_{II} > \tau_y \\ \eta & \text{otherwise} \end{cases}$$

- `Project marker properties to QP`
- `Evaluate FE Stokes residuals`

$$F_u^e = A^e u^e - Bp^e - f^e$$
$$F_c^e = (B^e)^T u^e$$

}

Saddle Point Preconditioners

Newton update requires linear solve

$$\begin{bmatrix} A + \delta A & B + \delta B \\ B^T + \delta B^T & 0 \end{bmatrix} \begin{bmatrix} \delta u \\ \delta p \end{bmatrix} = - \begin{bmatrix} F_u \\ F_c \end{bmatrix}$$

$$\mathcal{A} \; x = b$$

$$\mathcal{A} = \begin{bmatrix} A + \delta A & B + \delta B \\ B^T + \delta B^T & 0 \end{bmatrix}$$

$$x = \begin{bmatrix} \delta u \\ \delta p \end{bmatrix} \qquad b = - \begin{bmatrix} F_u \\ F_c \end{bmatrix}$$

The ideal approach should be *optimal* in the sense that the convergence rate of method will be bounded independently of:

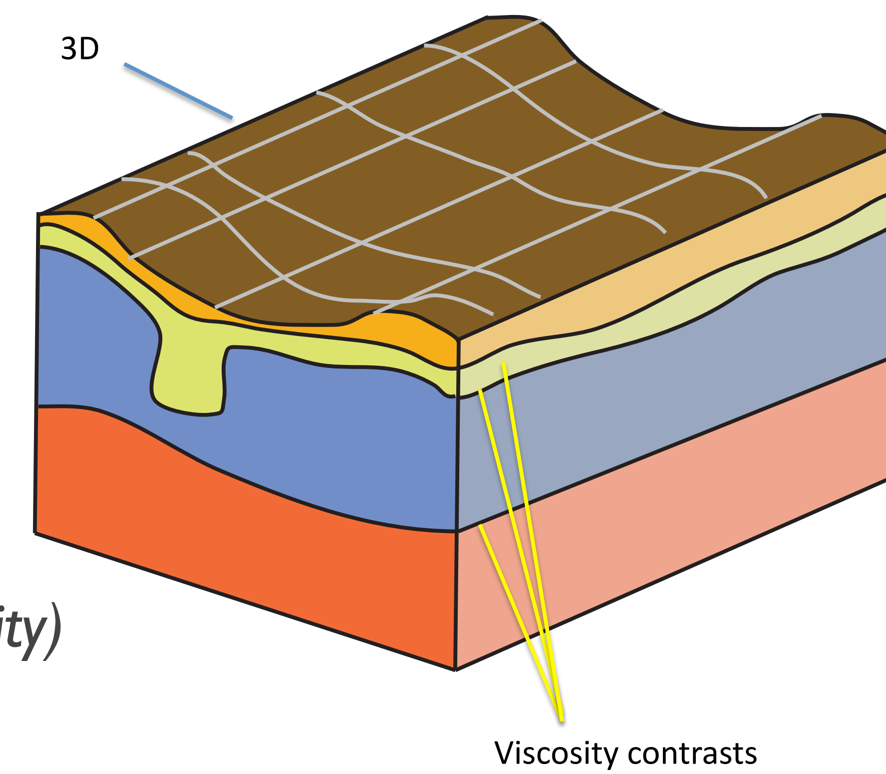

- the discretisation parameters *(e.g. grid resolution)*
- the constitutive parameters *(e.g. smooth vs. discontinuous viscosity)*
- the constitutive behaviour *(e.g. isotropic vs. anisotropic)*
- and we desire that the solution is obtained in *O(n)* time... i.e. multigrid

These are a challenging set of requirements

Newton MG for Saddle Point Systems

- Apply a Krylov method (e.g. FGMRES, GCR) directly to

$$\mathcal{A}x = b$$

right preconditioned with

$$\mathcal{B}_s = \begin{bmatrix} A' & B \\ 0 & -S^* \end{bmatrix} \quad \text{where} \quad S^* = \int_{\Omega_e} \frac{1}{\bar{\eta}_e} M_i M_j \, dV$$

$$S^* \approx S = B^T A^{-1} B$$

- Standard upper block triangular preconditioner, demonstrated to be effective for VV Stokes

See - Elman's book (2005) *- Burstedde, CMAME, (2009)*
- Geenen et al, G3, (2009) *- Grinevich & Olshanskii, SIAM J. Comput, (2009)*

Newton MG for Saddle Point Systems

- Apply a Krylov method (e.g. FGMRES, GCR) directly to

$$\mathcal{A}x = b$$

right preconditioned with

$$\mathcal{B}_s = \begin{bmatrix} A' & B \\ 0 & -S^* \end{bmatrix} \quad \text{where} \quad S^* = \int_{\Omega_e} \frac{1}{\bar{\eta}_e} M_i M_j \, dV$$

$$S^* \approx S = B^T A^{-1} B$$

- Standard upper block triangular preconditioner, demonstrated to be effective for VV Stokes

See - Elman's book (2005)
- Geenen et al, G3, (2009)
- Burstedde, CMAME, (2009)
- Grinevich & Olshanskii, SIAM J. Comput, (2009)

- Applying the action of the Stokes preconditioner on a vector *t*

$$s = \mathcal{B}_s^{-1} t$$

requires the action of

$$u = A'^{-1} v$$

→ Apply Algebraic MultiGrid (AMG) or Geometric MultiGrid (GMG) to ***A'***

Two-dimensional Examples

Prototype geodynamic processes

Crustal/lithospheric Problem
(Boundary driven processes)

strong upper crust
Layered material

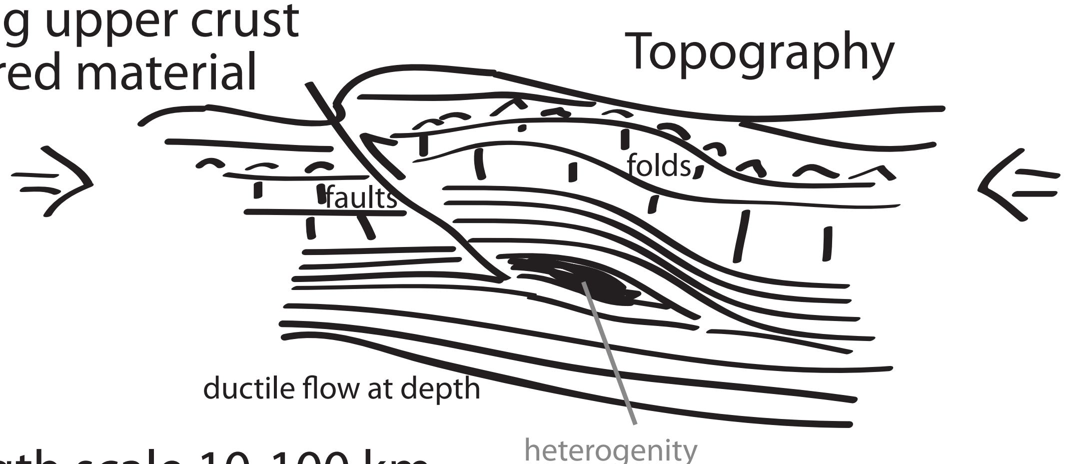

Length scale 10-100 km

weak density contrasts,
gravitationally stable
(density increases with depth)
thick high viscosity layer.

Upper Mantle Problem
(Body force driven processes)

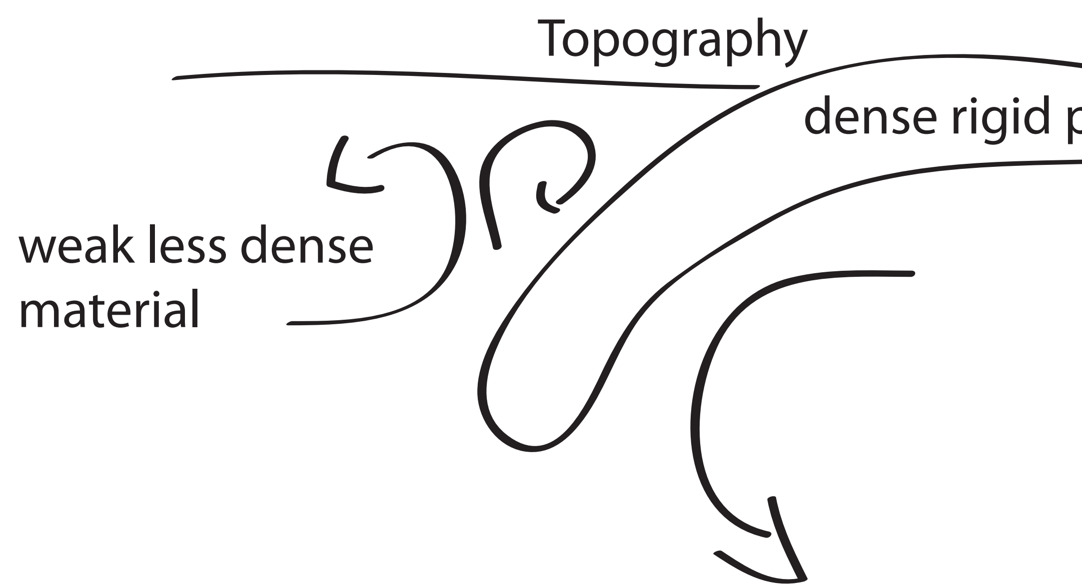

Length scale 100-1000 km

Reverted density gradient,
thick low viscosity layer

Numerical realisations

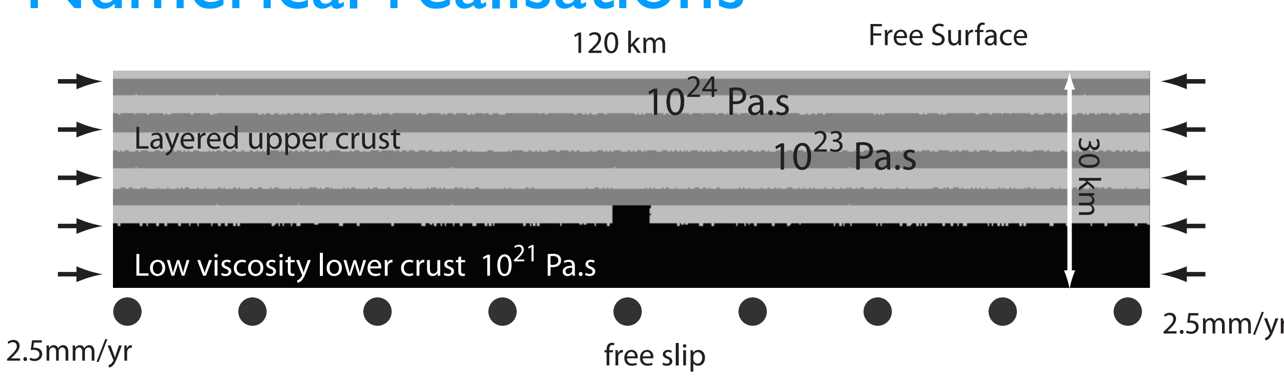

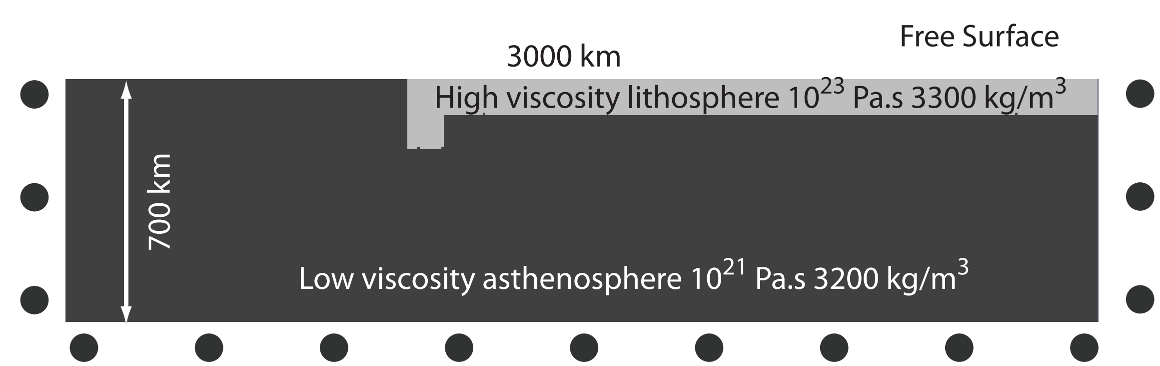

[May, Le Pourhiet, JCP, In prep.]

Visco-plastic Shortening

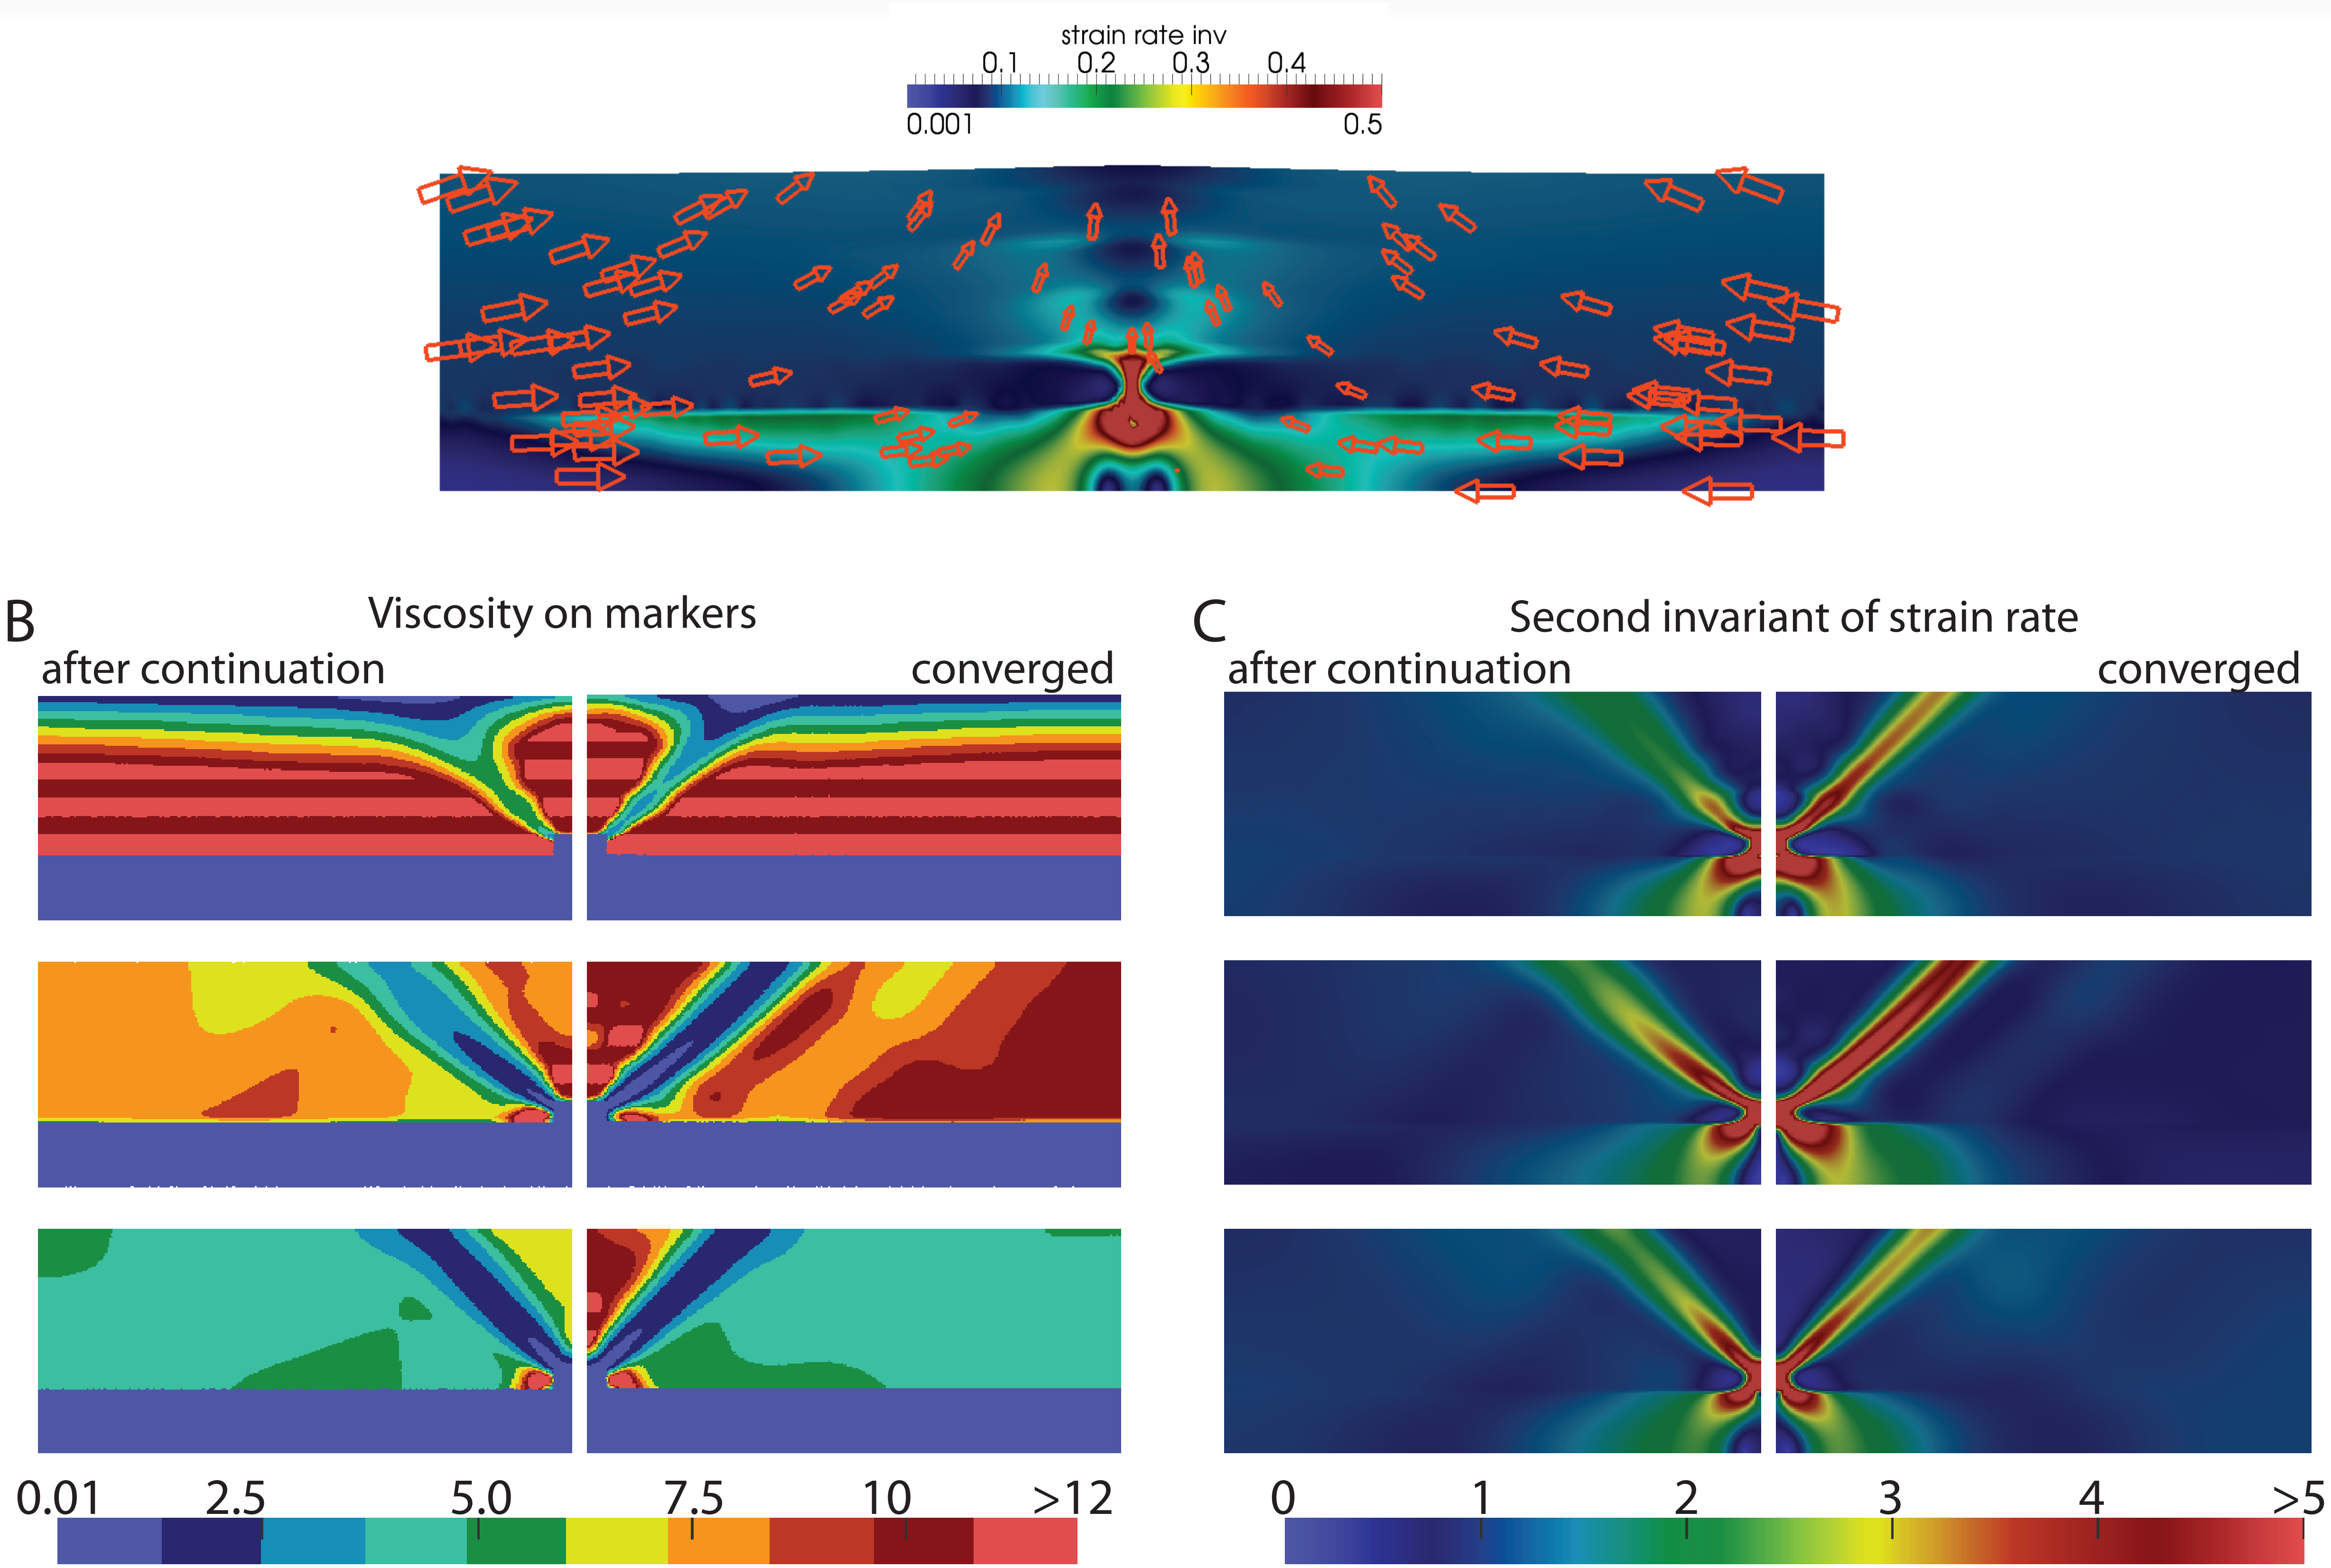

Visco-plastic Shortening

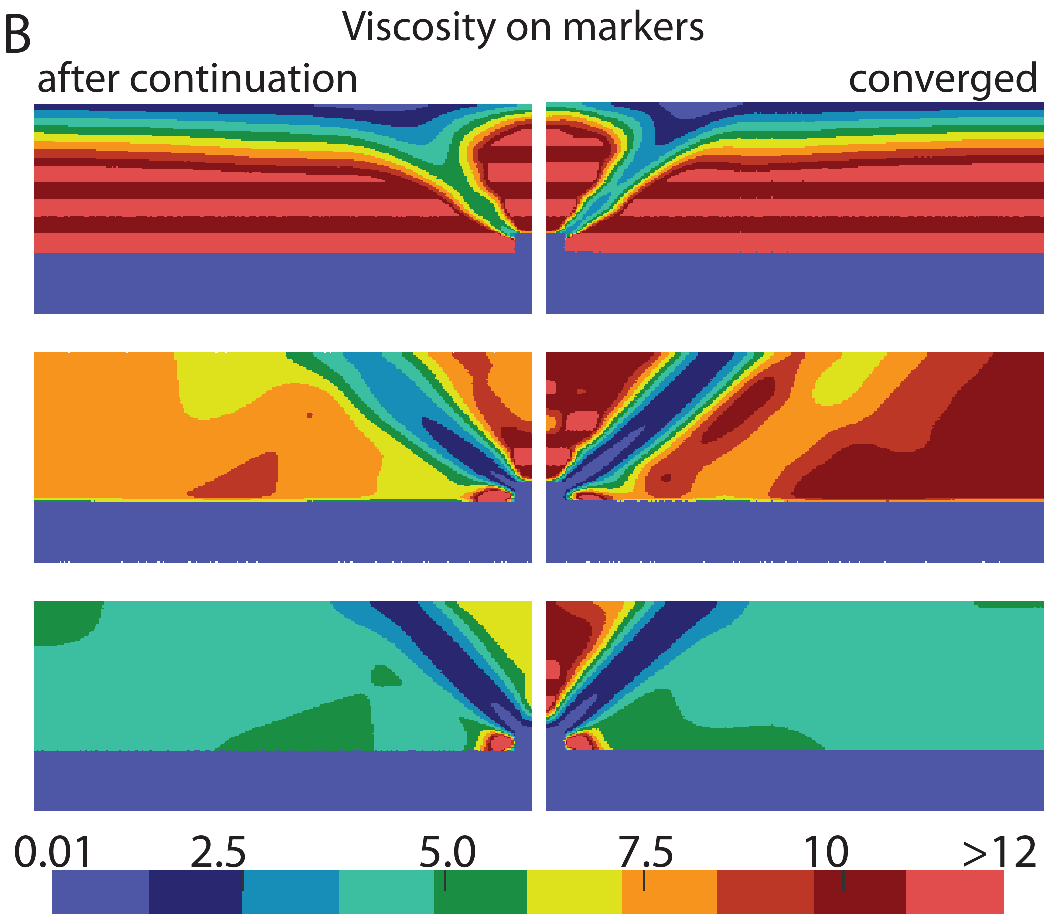

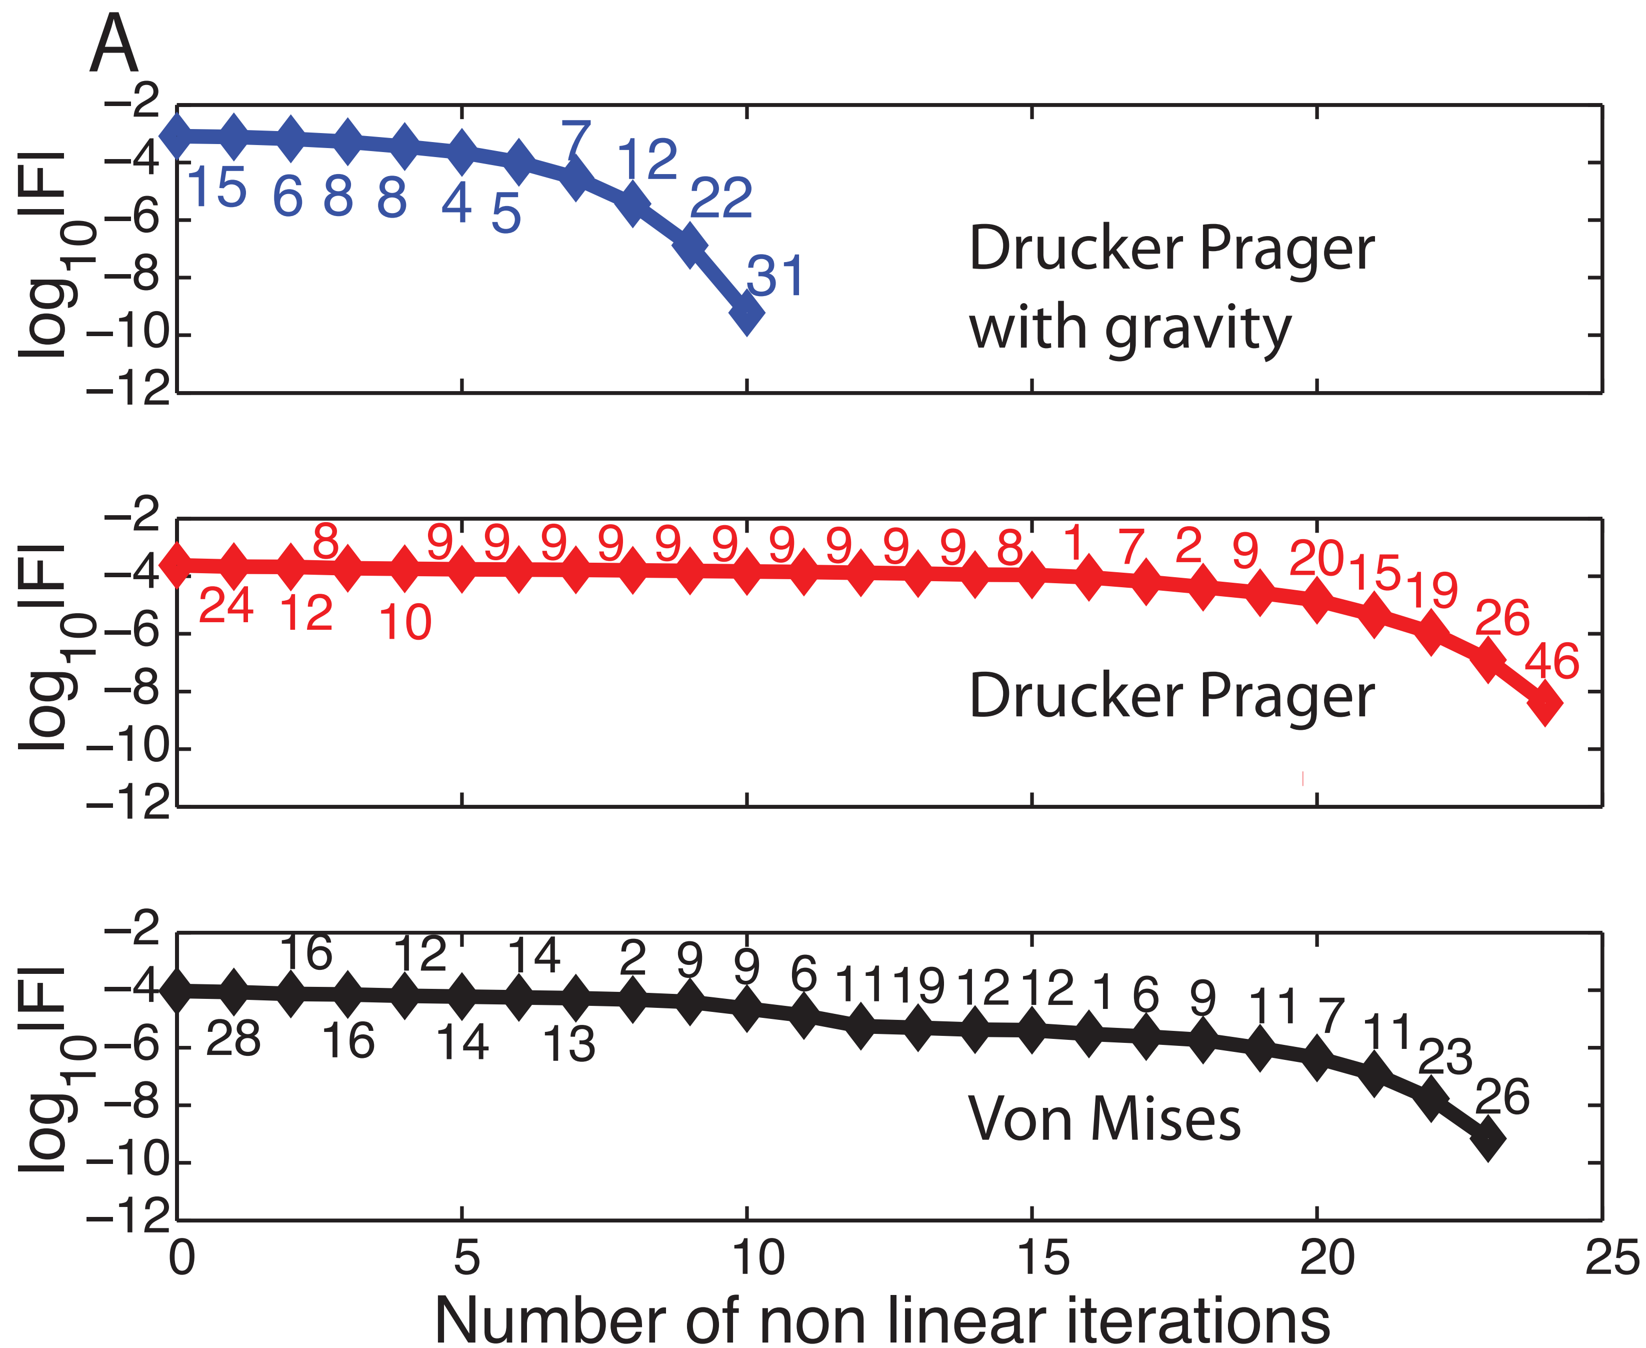

Subduction: Coordinate Evolution

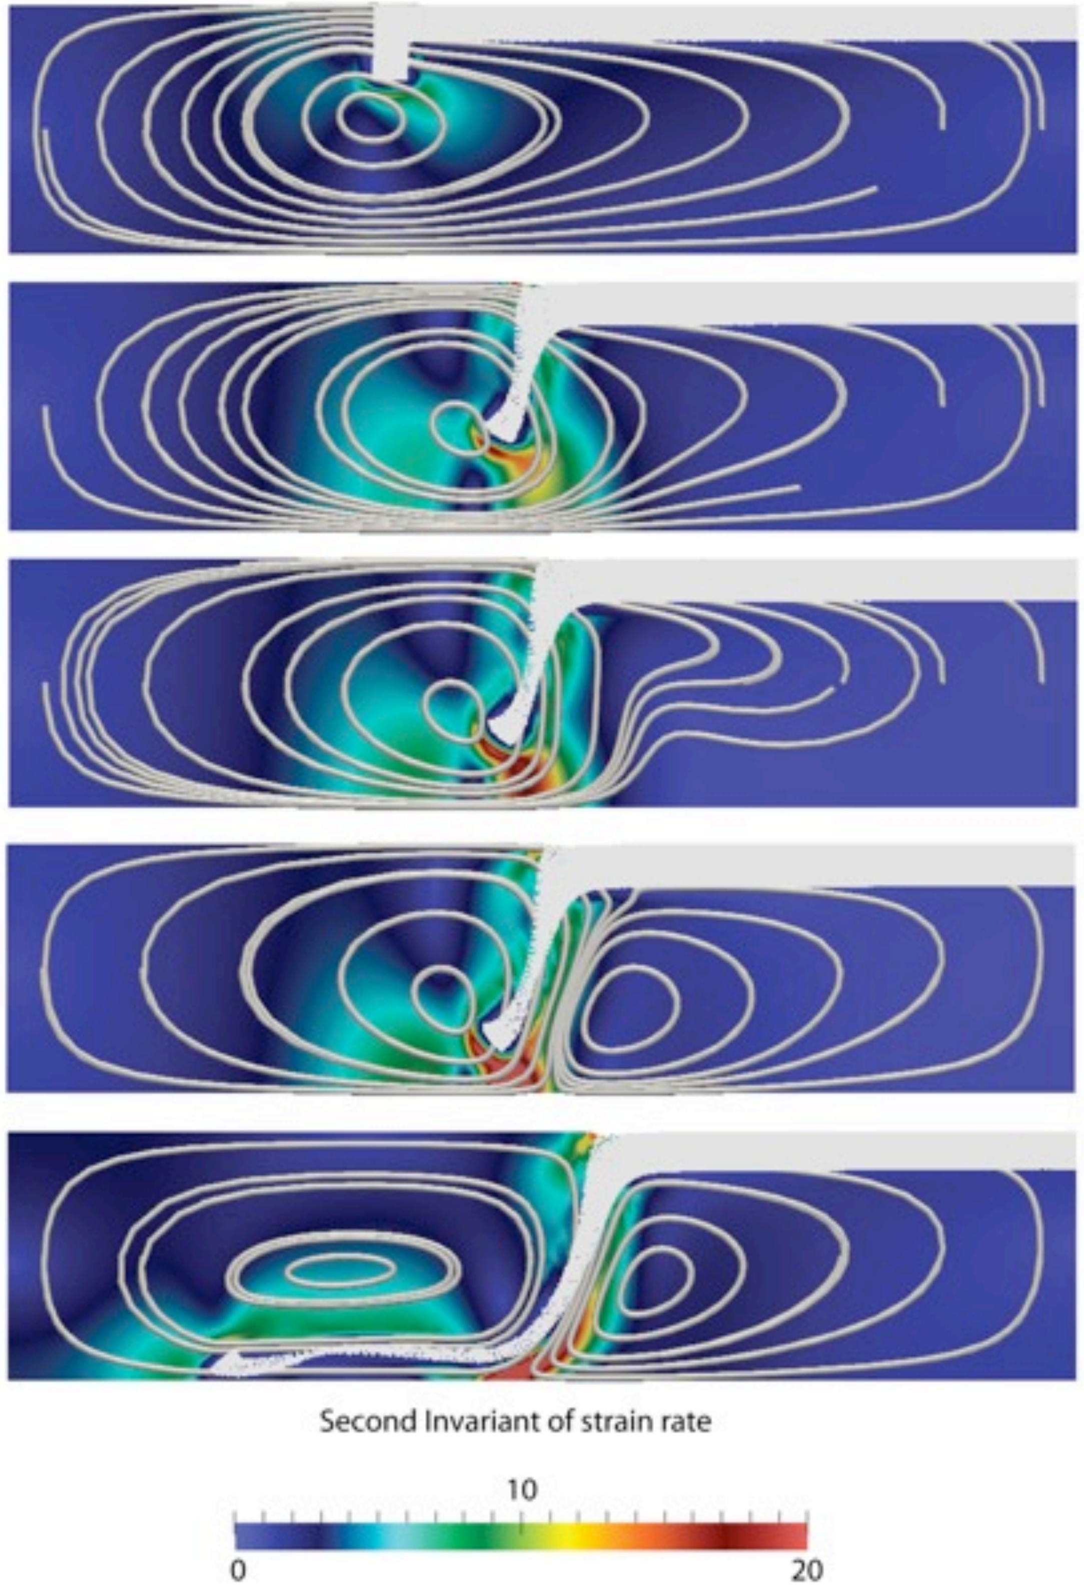

STOKES FLOW
+ COORDINATE EVOLUTION

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{u}$$

$$\hat{F}_{u_i} := \Big[2\eta(\boldsymbol{u}, p, \boldsymbol{x}')D_{ij}(\boldsymbol{u}, \boldsymbol{x}')\Big]_{,j} - p_{,i}(\boldsymbol{x}') - f_i(\boldsymbol{u}, p, \boldsymbol{x}')$$

$$\hat{F}_c := u_{k,k}(\boldsymbol{x}')$$

$$\boldsymbol{x}' = \boldsymbol{x} + \Delta t\, \boldsymbol{u}$$

Linear solve with dt = 6e-5 : typical free surface instability

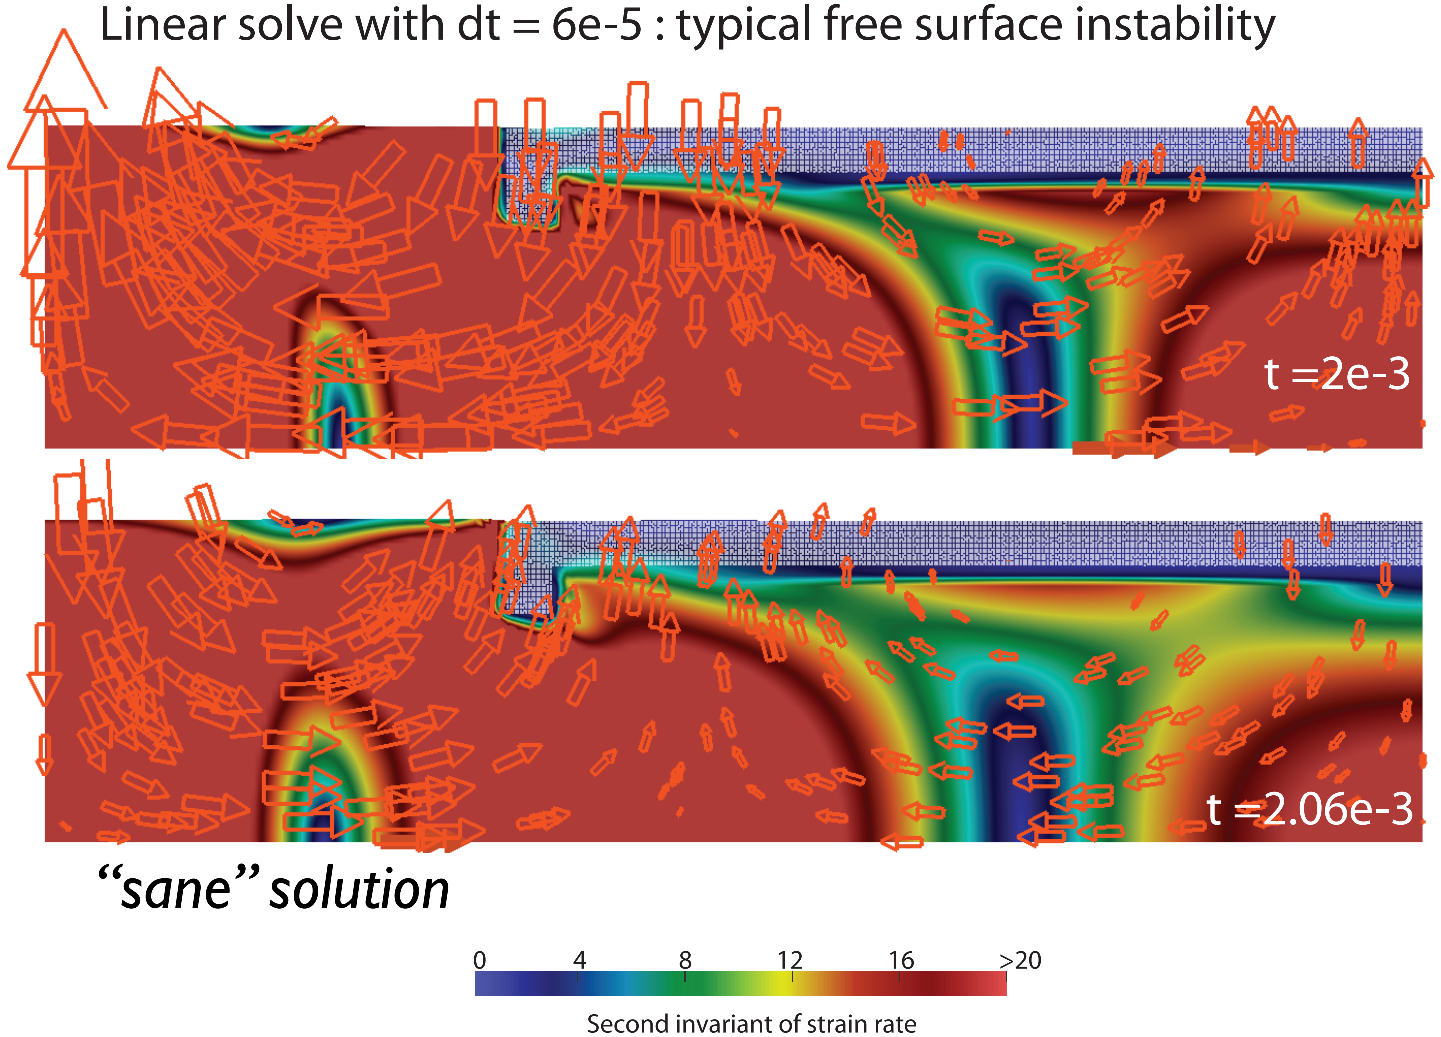

“sane” solution

Subduction: Coordinate Evolution

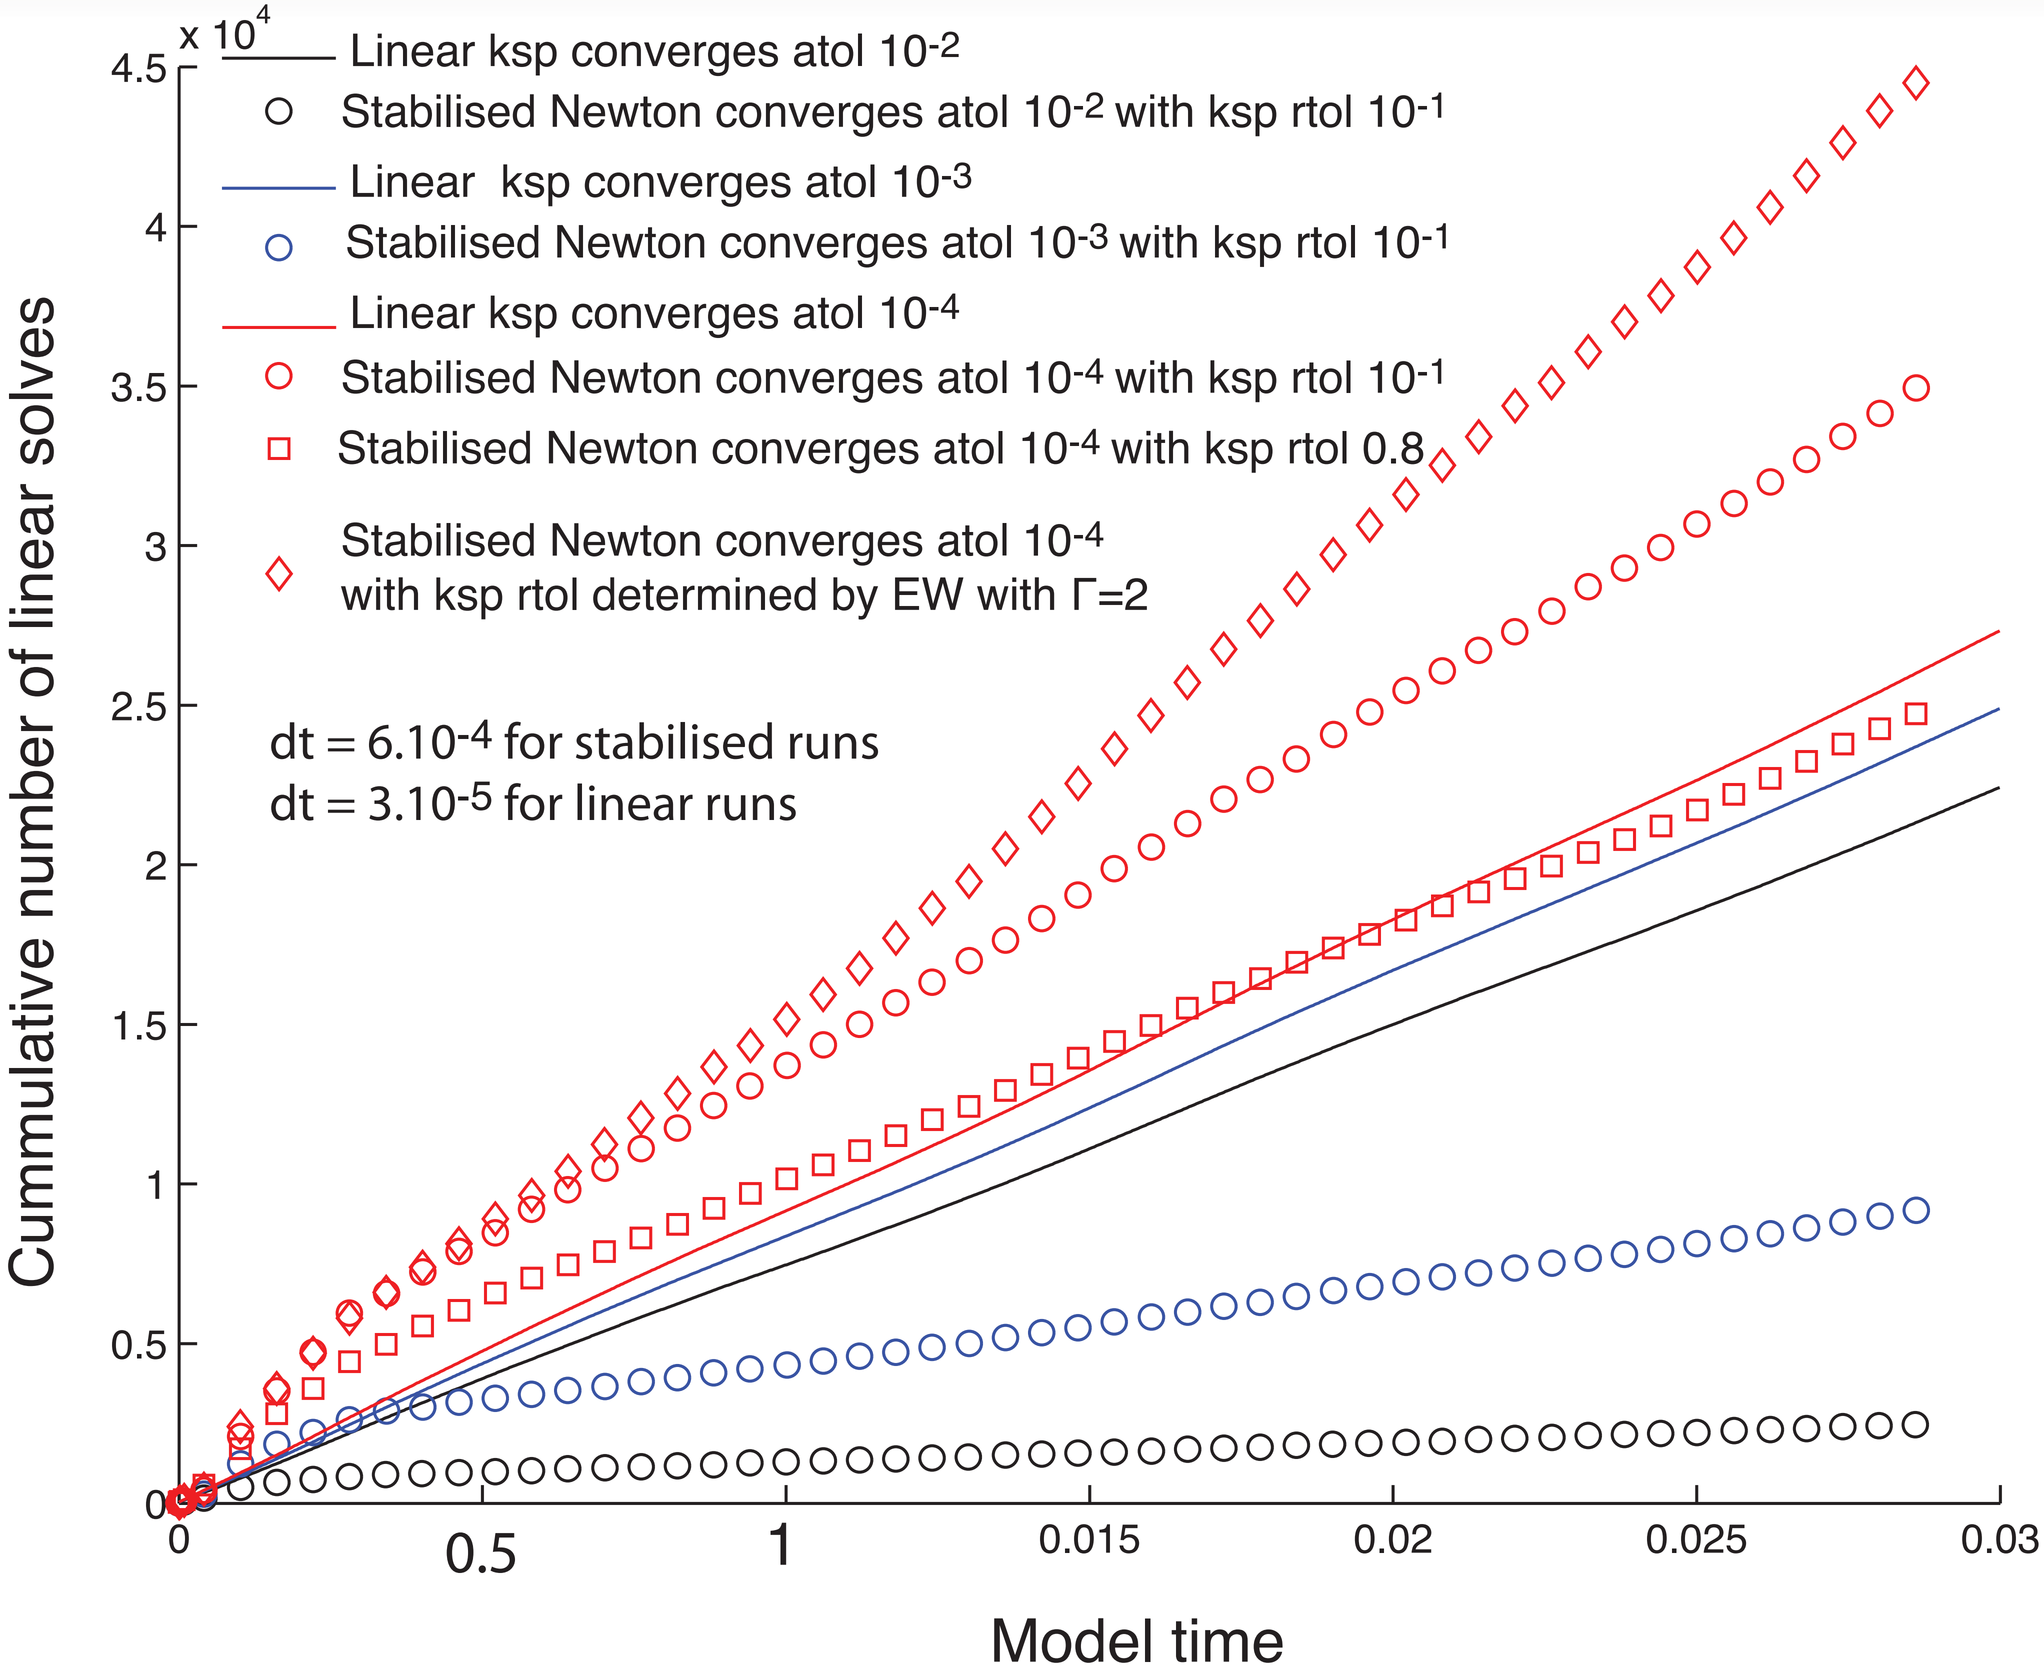

Subduction: Coordinate Evolution

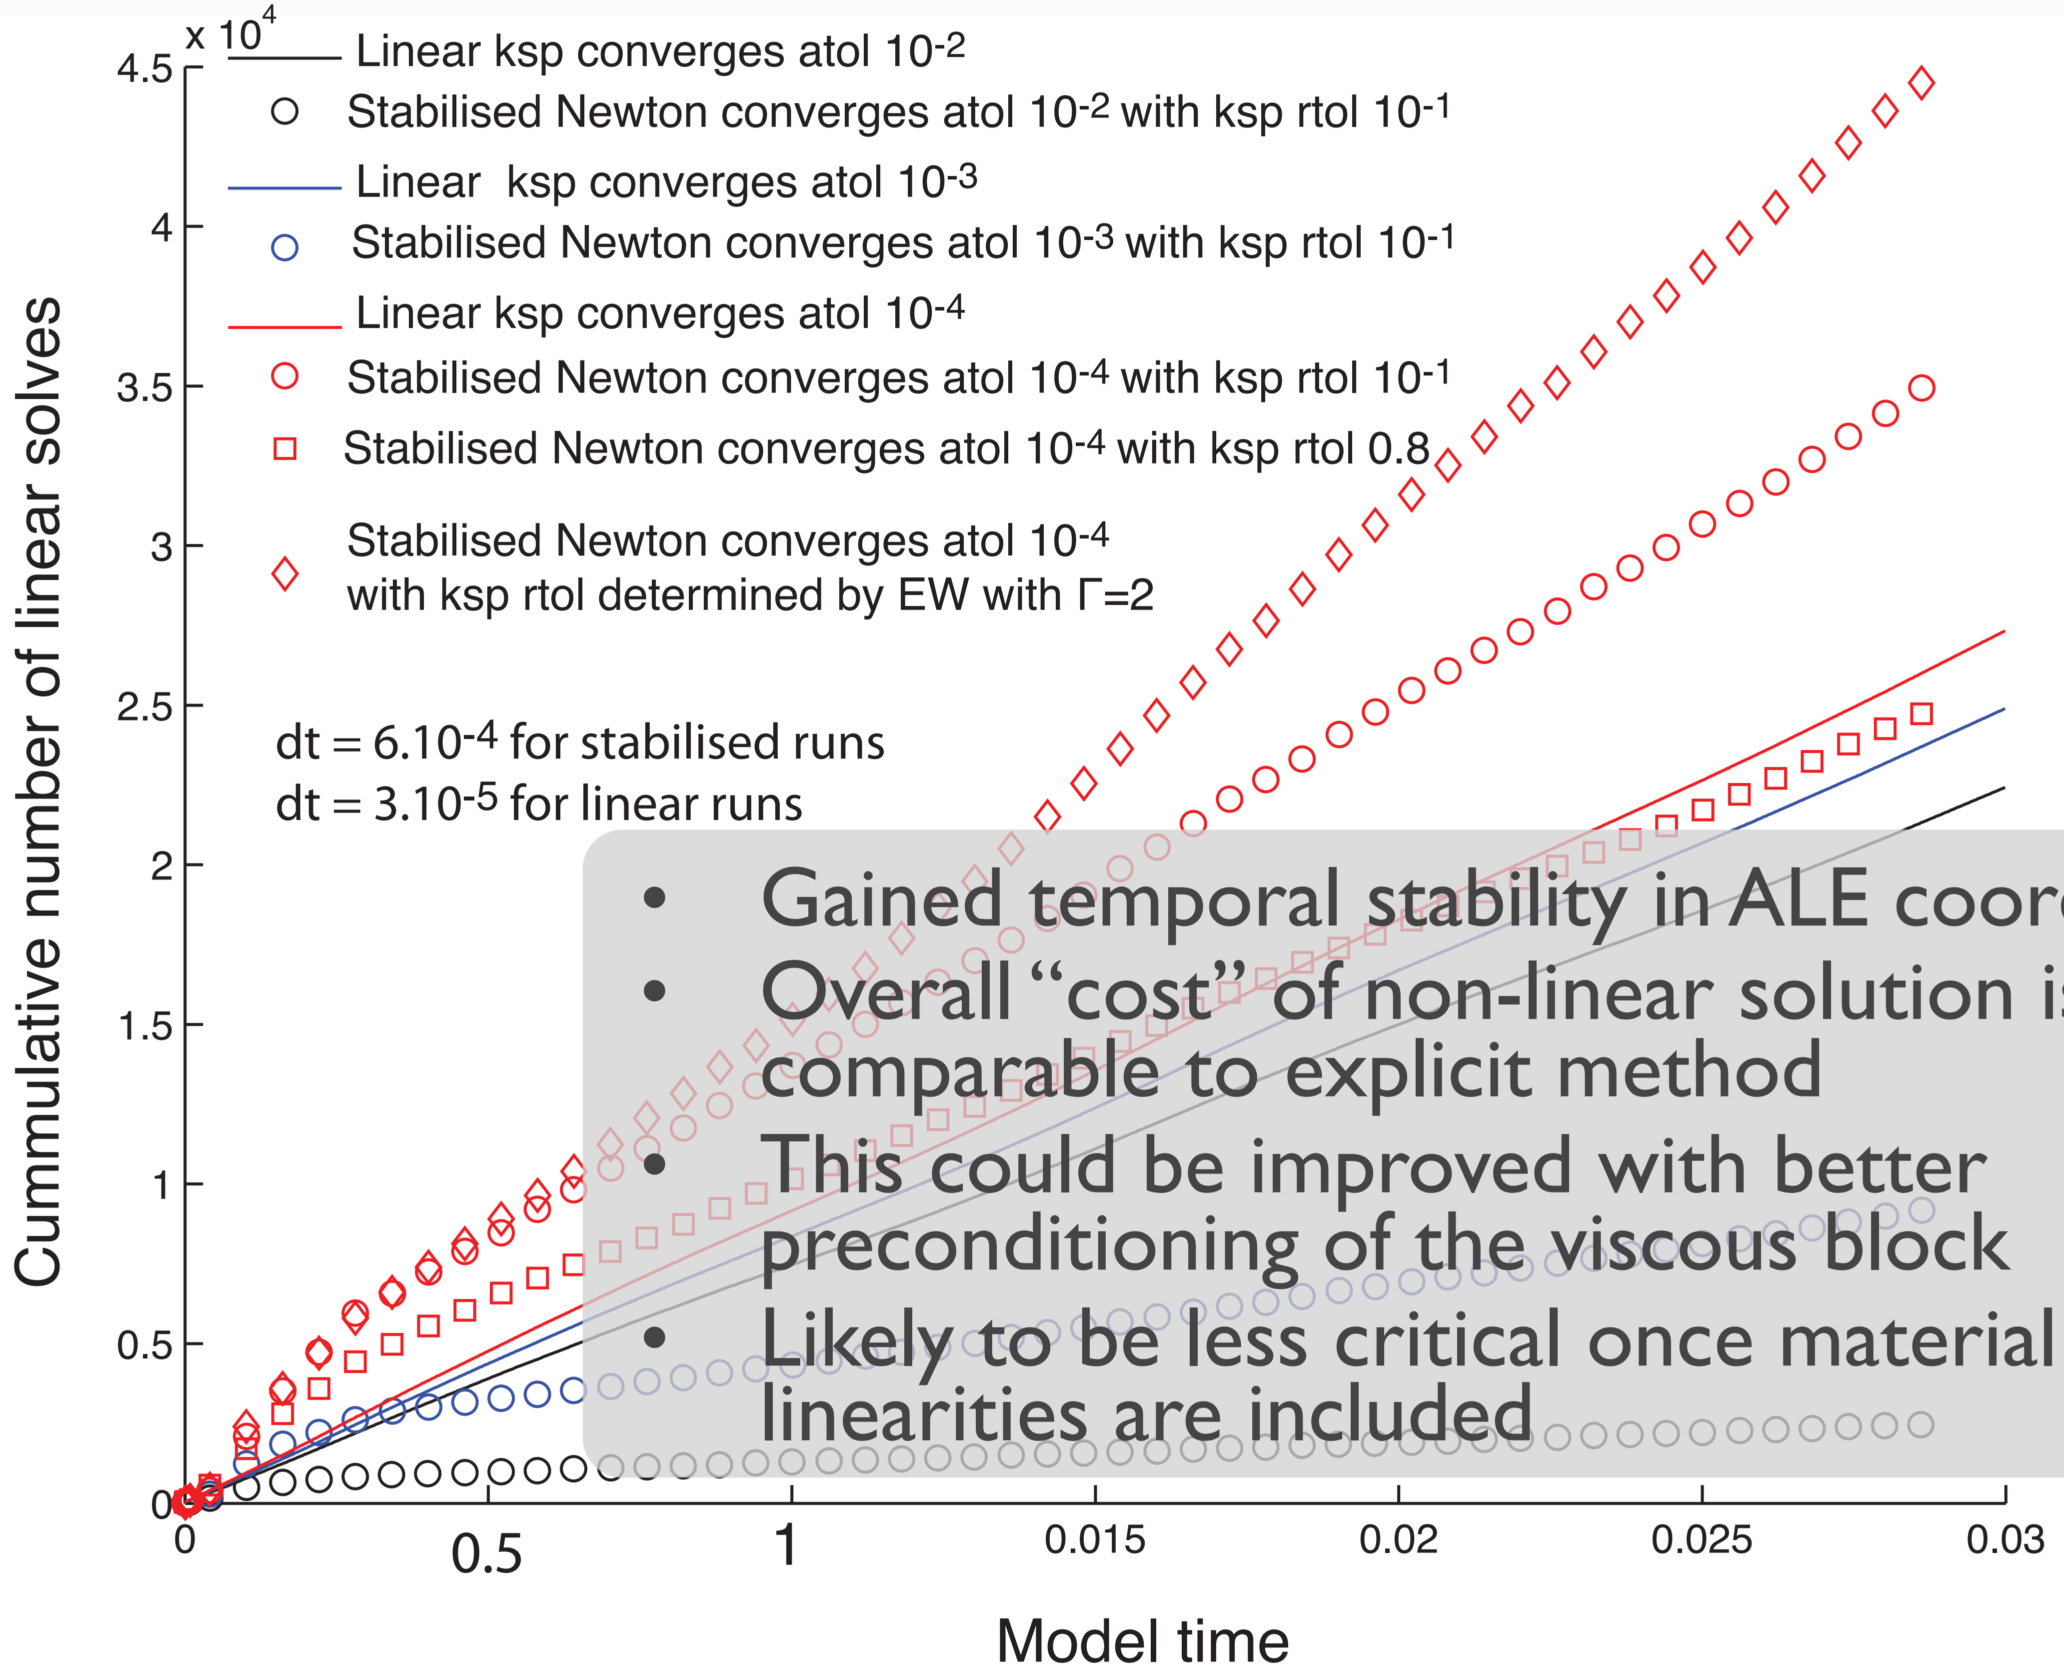

- Gained temporal stability in ALE coordinate update
- Overall "cost" of non-linear solution is less, or comparable to explicit method
- This could be improved with better preconditioning of the viscous block
- Likely to be less critical once material non-linearities are included

Moving to 3D: pTatin3d

["p" stands for PETSc, pragmatic and pedantic]

- ***Open source project with the following features***
  - Fully parallel (flat MPI), robust and scalable 3D FE-MPM discretisation and solvers for non-linear variable viscosity Stokes
  - Physics is extensible
  - Flexible solver design (defer as many choices as possible to run time)
  - Low memory to maximize numerical resolution, maximize resources and permit wide usability to geodynamic community without massive HPC access
  - Employ algorithms which exploit modern multi-core architectures.
    Target hardware; IBM BG/Q, Cray XE6

Parallel algebra support provided by PETSc (www.mcs.anl.gov/petsc)

Performance Issues

"Strong" smoothers require assembling operators (e.g. CG/ICC - GMRES/ILU)

STORAGE IS EXPENSIVE

***A*:** (Q2) 64 x 64 x 64 ~ 19.3 GB

3 x A_{ii}: (Q2) 64 x 64 x 64 ~ 6.4 GB

+ temporary vectors for the solver
+ whatever else you might need...
  e.g. markers, quadrature point fields...

Performance Issues

"Strong" smoothers require assembling operators (e.g. CG/ICC - GMRES/ILU)

STORAGE IS EXPENSIVE

A: (Q2) 64 x 64 x 64 ~ 19.3 GB

3 x ***A****ii*: (Q2) 64 x 64 x 64 ~ 6.4 GB

+ temporary vectors for the solver
+ whatever else you might need...
  e.g. markers, quadrature point fields...

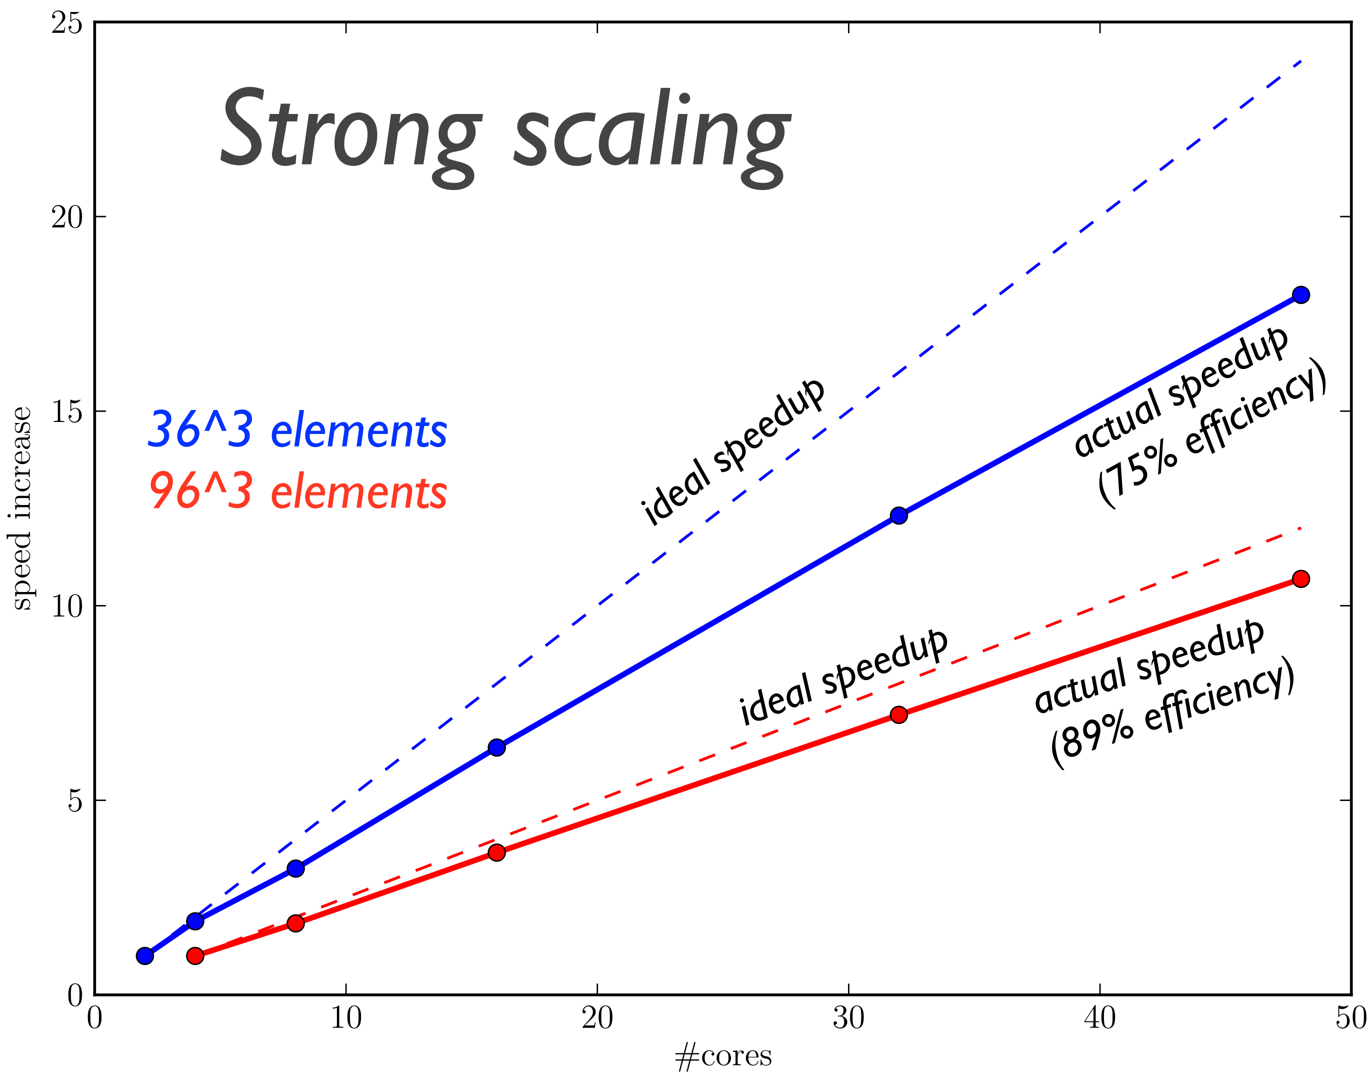

Weak scaling

$M_x \times M_y \times M_z$	n_p	m_x	CPU time (sec)	
36 × 36 × 36	1	36	1.1066e+00	
72 × 36 × 36	2	36	1.1409e+00	(100%)
72 × 72 × 36	4	36	1.1981e+00	(95%)
72 × 72 × 72	8	36	1.3956e+00	(82%)
72 × 72 × 144	16	36	1.3991e+00	(82%)
144 × 144 × 72	32	36	1.4139e+00	(81%)
144 × 144 × 108	48	36	1.4249e+00	(80%)
18 × 18 × 18	1	18	1.3685e-01	
36 × 18 × 18	2	18	1.4262e-01	(100%)
36 × 36 × 18	4	18	1.5121e-01	(94%)
36 × 36 × 36	8	18	1.7698e-01	(81%)
36 × 36 × 72	16	18	1.7826e-01	(80%)
72 × 72 × 36	32	18	1.8086e-01	(79%)
72 × 72 × 54	48	18	1.8214e-01	(78%)

"ulysse": *[SGI Altix UV 100]* 6 nodes; 8 x Xeon E7-8837 (2.67GHz); 8 GB RAM/core

Performance of Matrix-Free (MF) SpMV [weak scaling]

"hexagon": [Cray XE6] 696 nodes; 2x16 AMD Interlagos (2.3GHz); 1 GB RAM/core

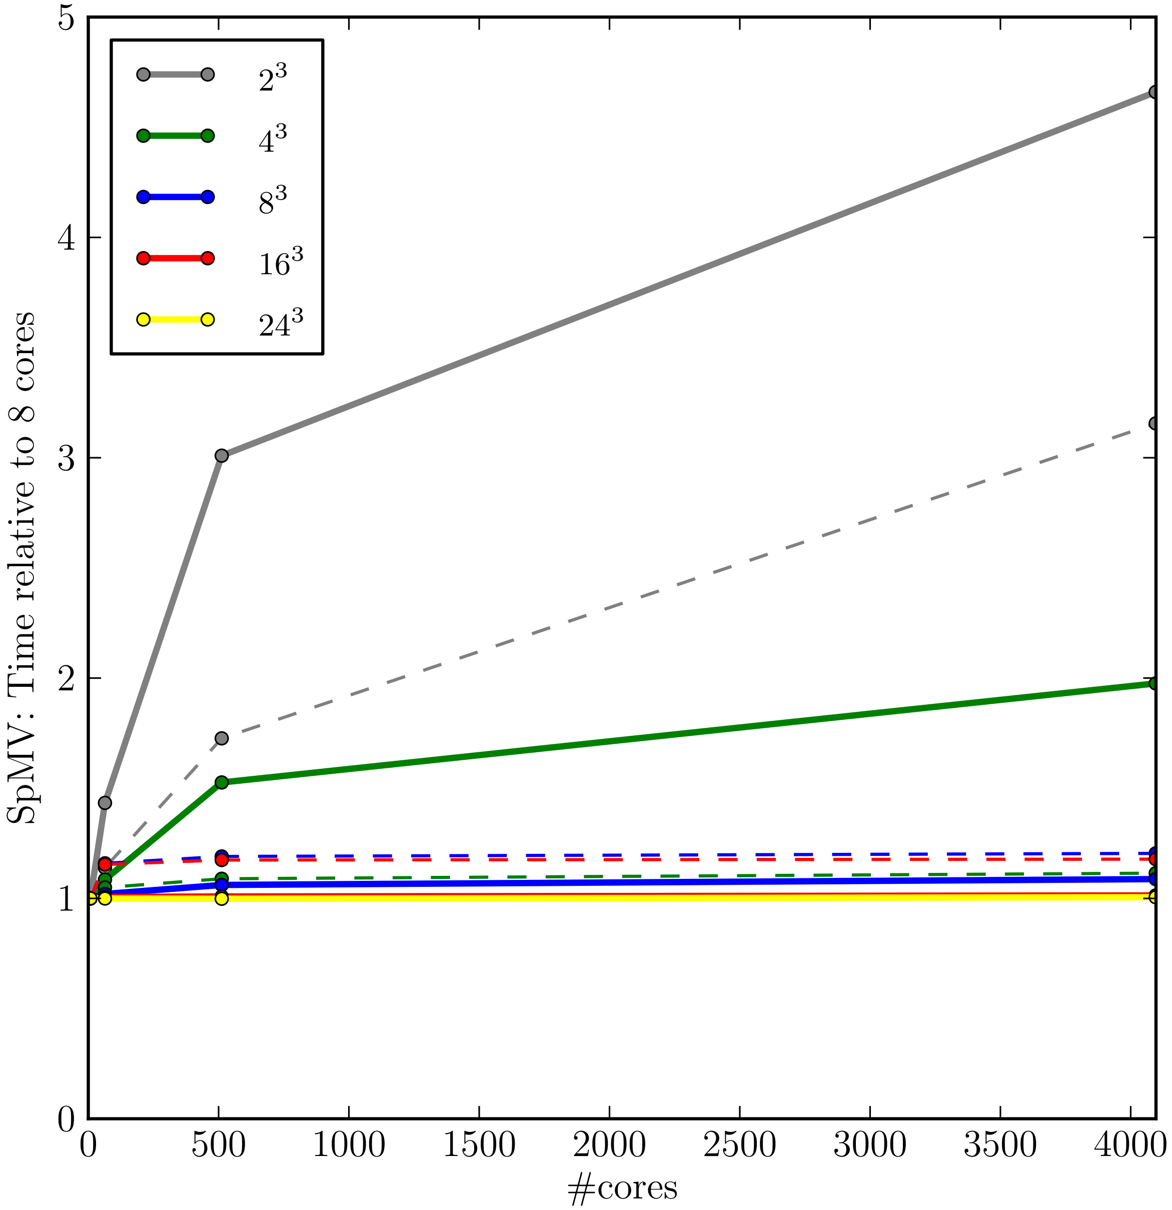

Performance of Matrix-Free (MF) SpMV [weak scaling]

"hexagon": [Cray XE6] 696 nodes; 2x16 AMD Interlagos (2.3GHz); 1 GB RAM/core

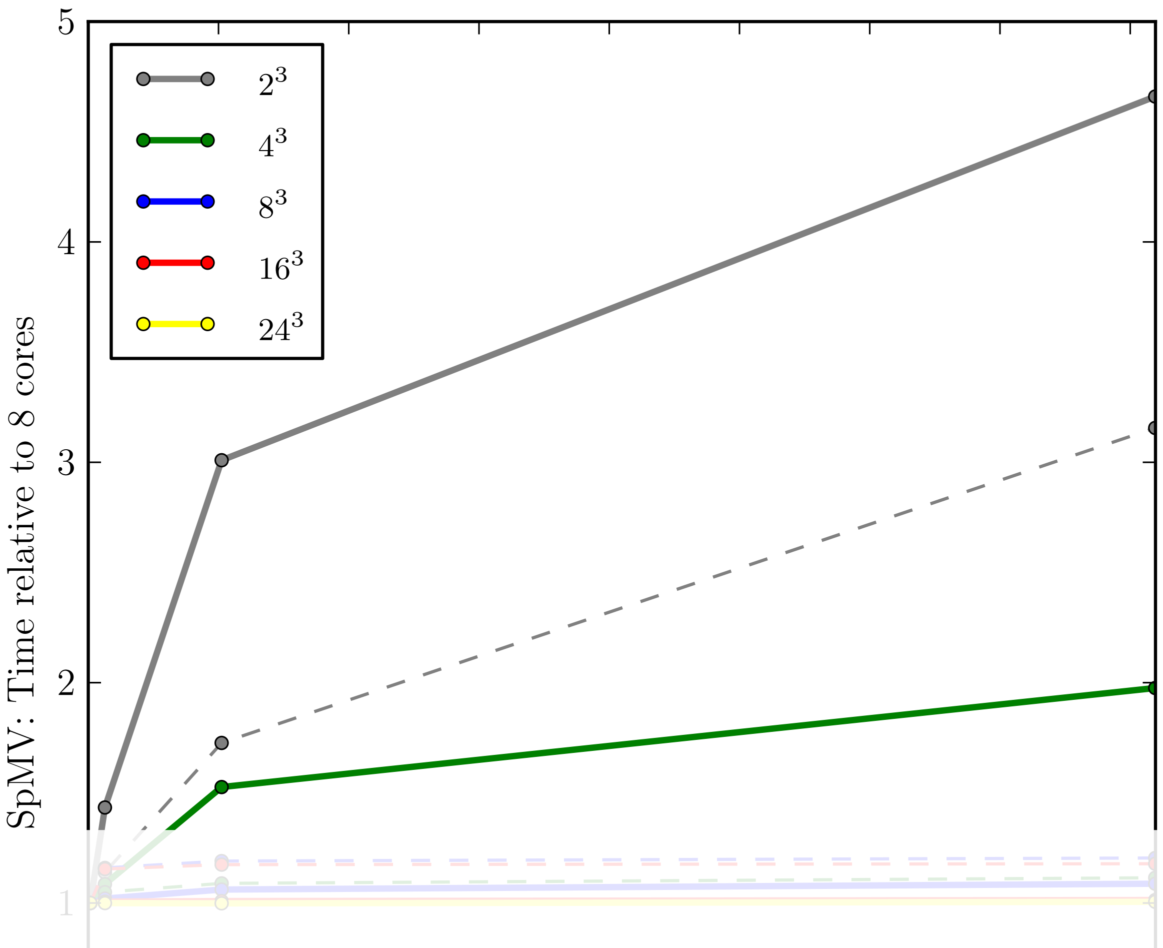

Elements per core	Elements on 4096 cores	Efficiency ASM	Efficiency MF
2^3	32^3	0.32	0.21
4^3	64^3	0.90	0.51
8^3	128^3	0.83	0.92
16^3	256^3	0.85	0.99
24^3	384^3	*	0.99

Performance of Matrix-Free (MF) SpMV [weak scaling]

"hexagon": [Cray XE6] 696 nodes; 2x16 AMD Interlagos (2.3GHz); 1 GB RAM/core

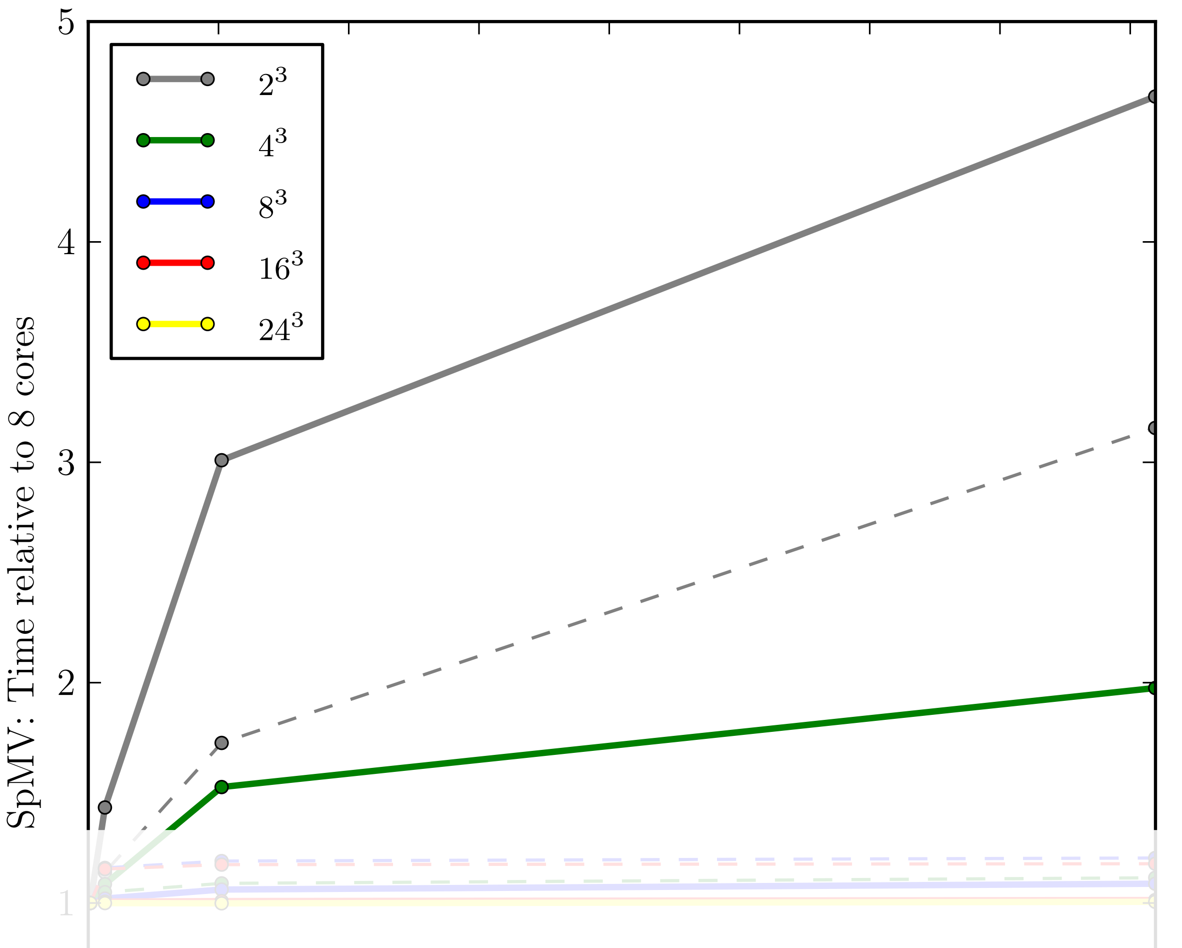

Elements per core	Elements on 4096 cores	Efficiency ASM	Efficiency MF
2^3	32^3	0.32	0.21
4^3	64^3	0.90	0.51
8^3	128^3	0.83	0.92
16^3	256^3	0.85	0.99
24^3	384^3	*	0.99

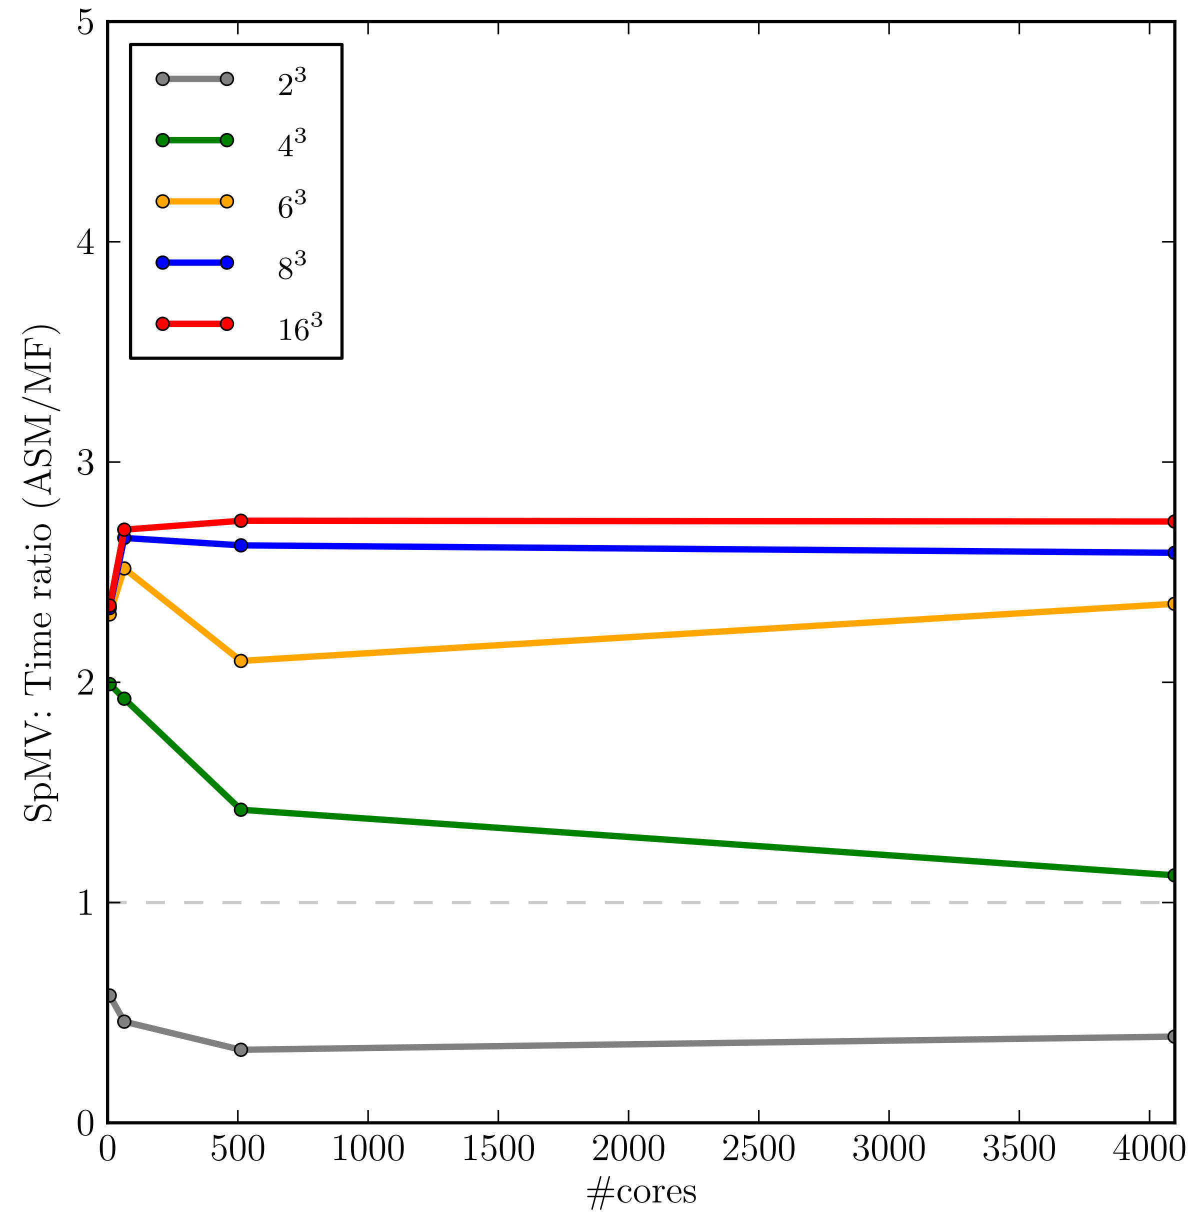

MF is faster than ASM when number of elements per core is larger than 8. Typical scenario on fine levels.

Performance of MF-SpMV [strong scaling]

"hexagon": [Cray XE6] 696 nodes; 2x16 AMD Interlagos (2.3GHz); 1 GB RAM/core

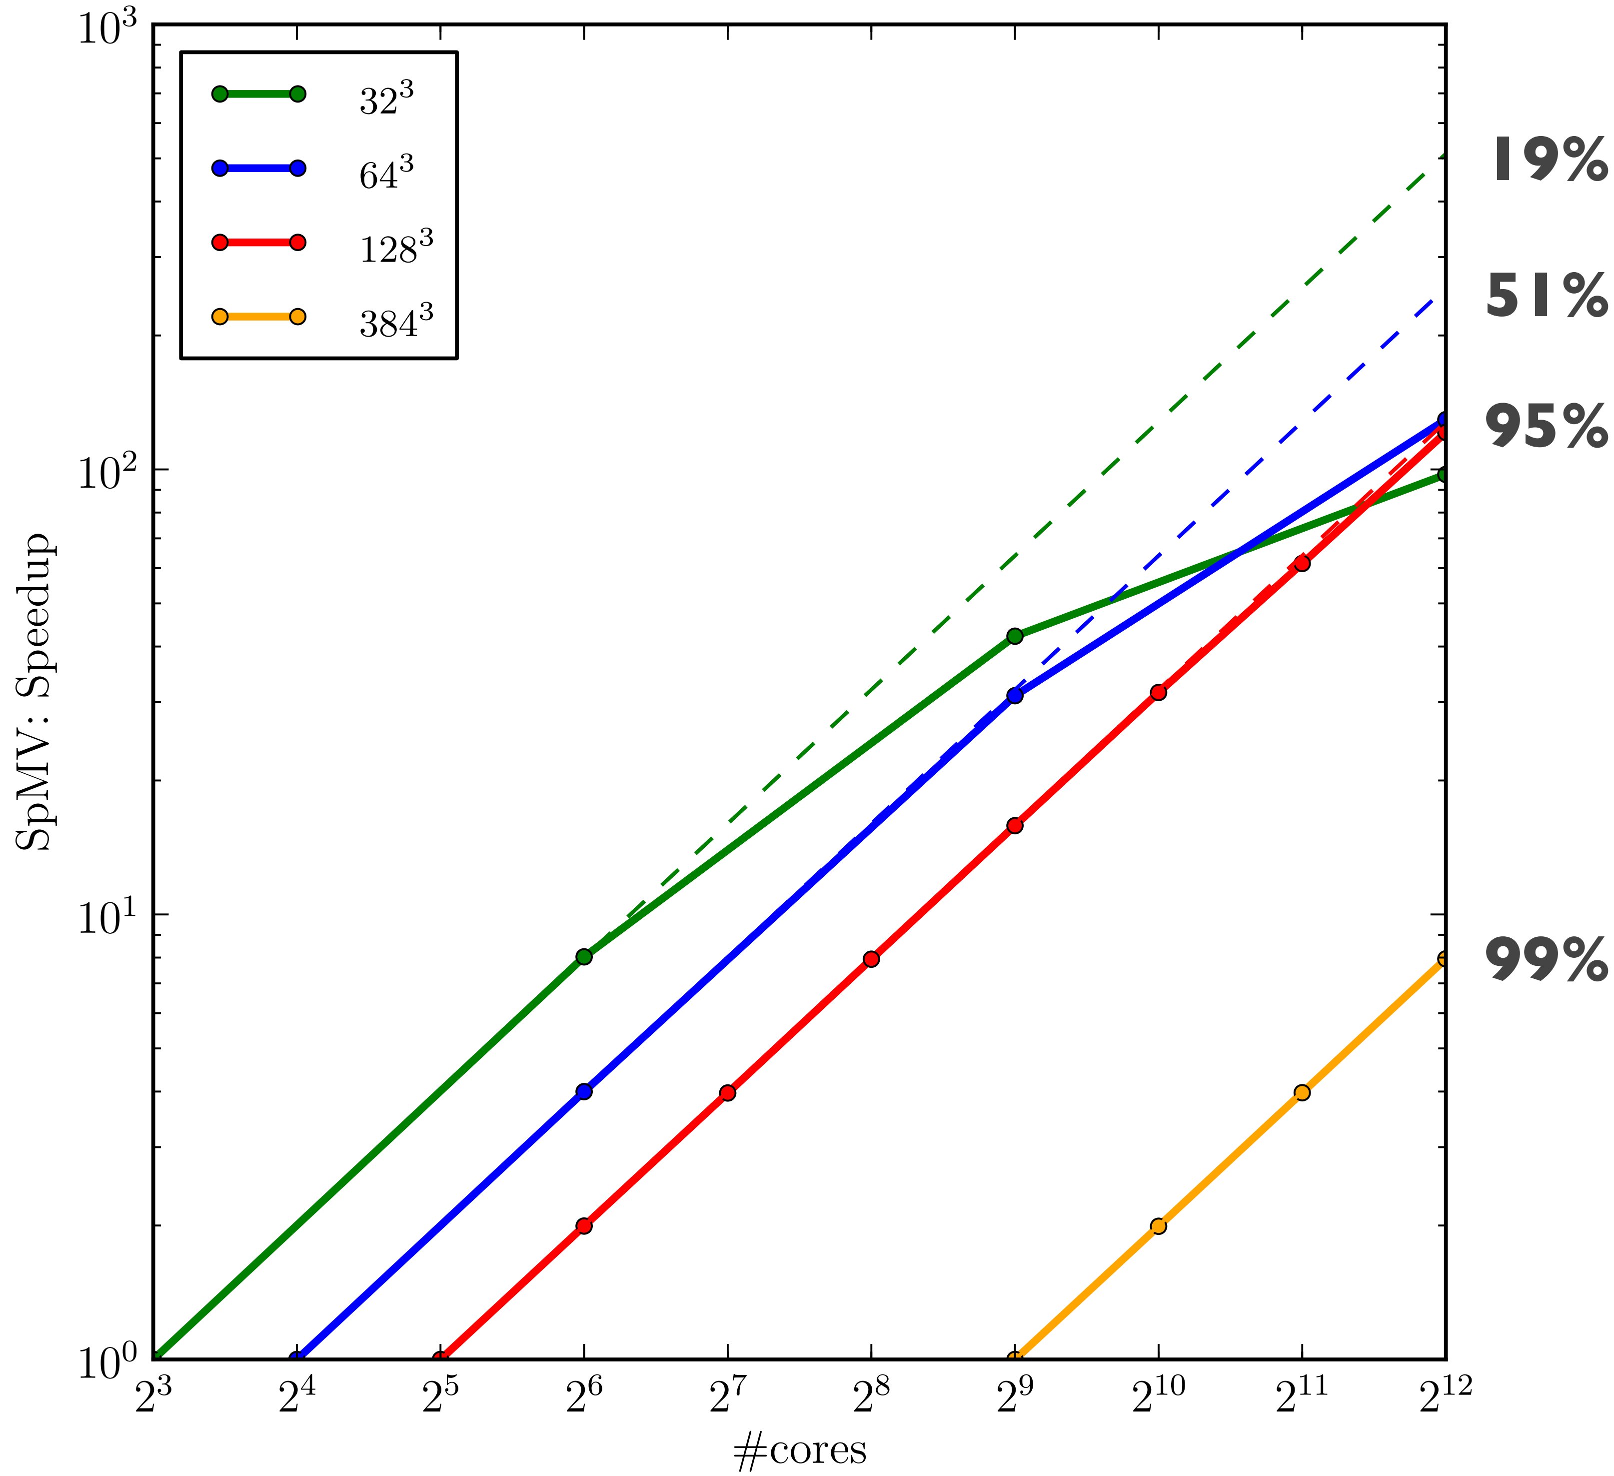

Excellent strong scaling when using more than 8 Q2 elements per core

Hybrid Multi Level Strategy for Au=v

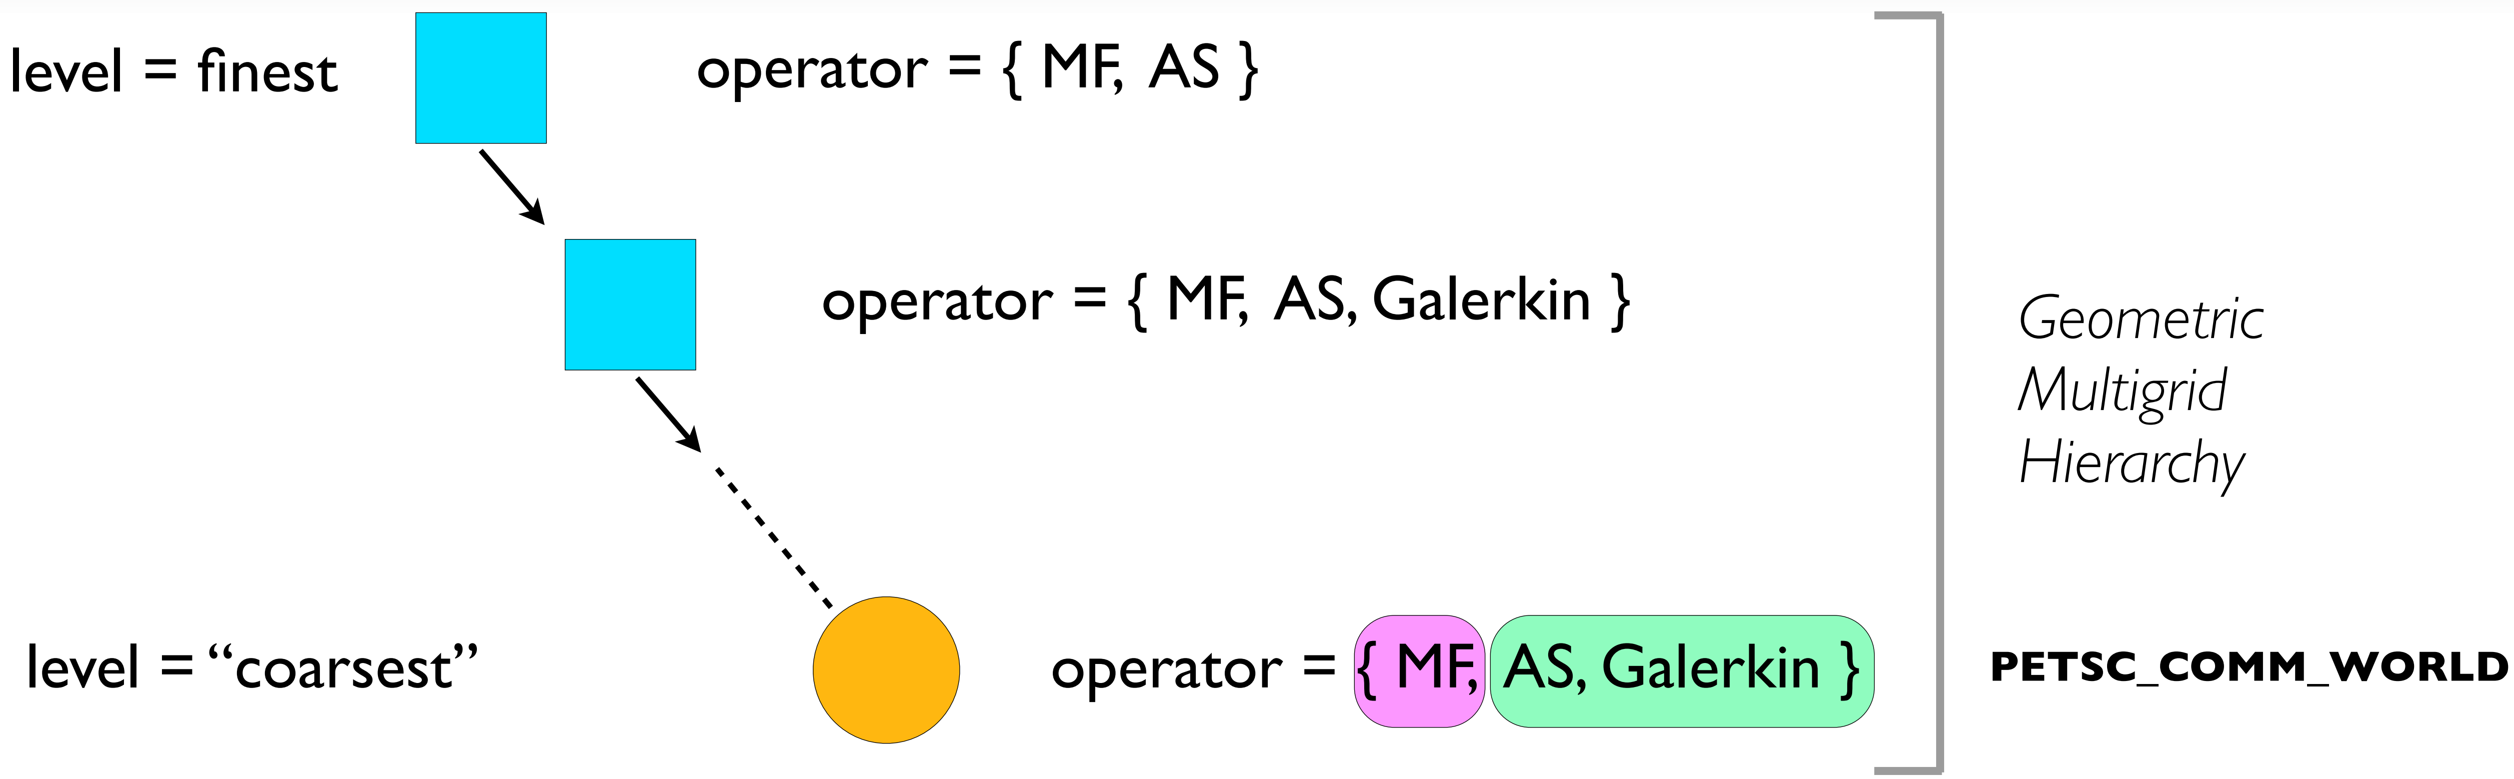

Hybrid Multi Level Strategy for Au=v

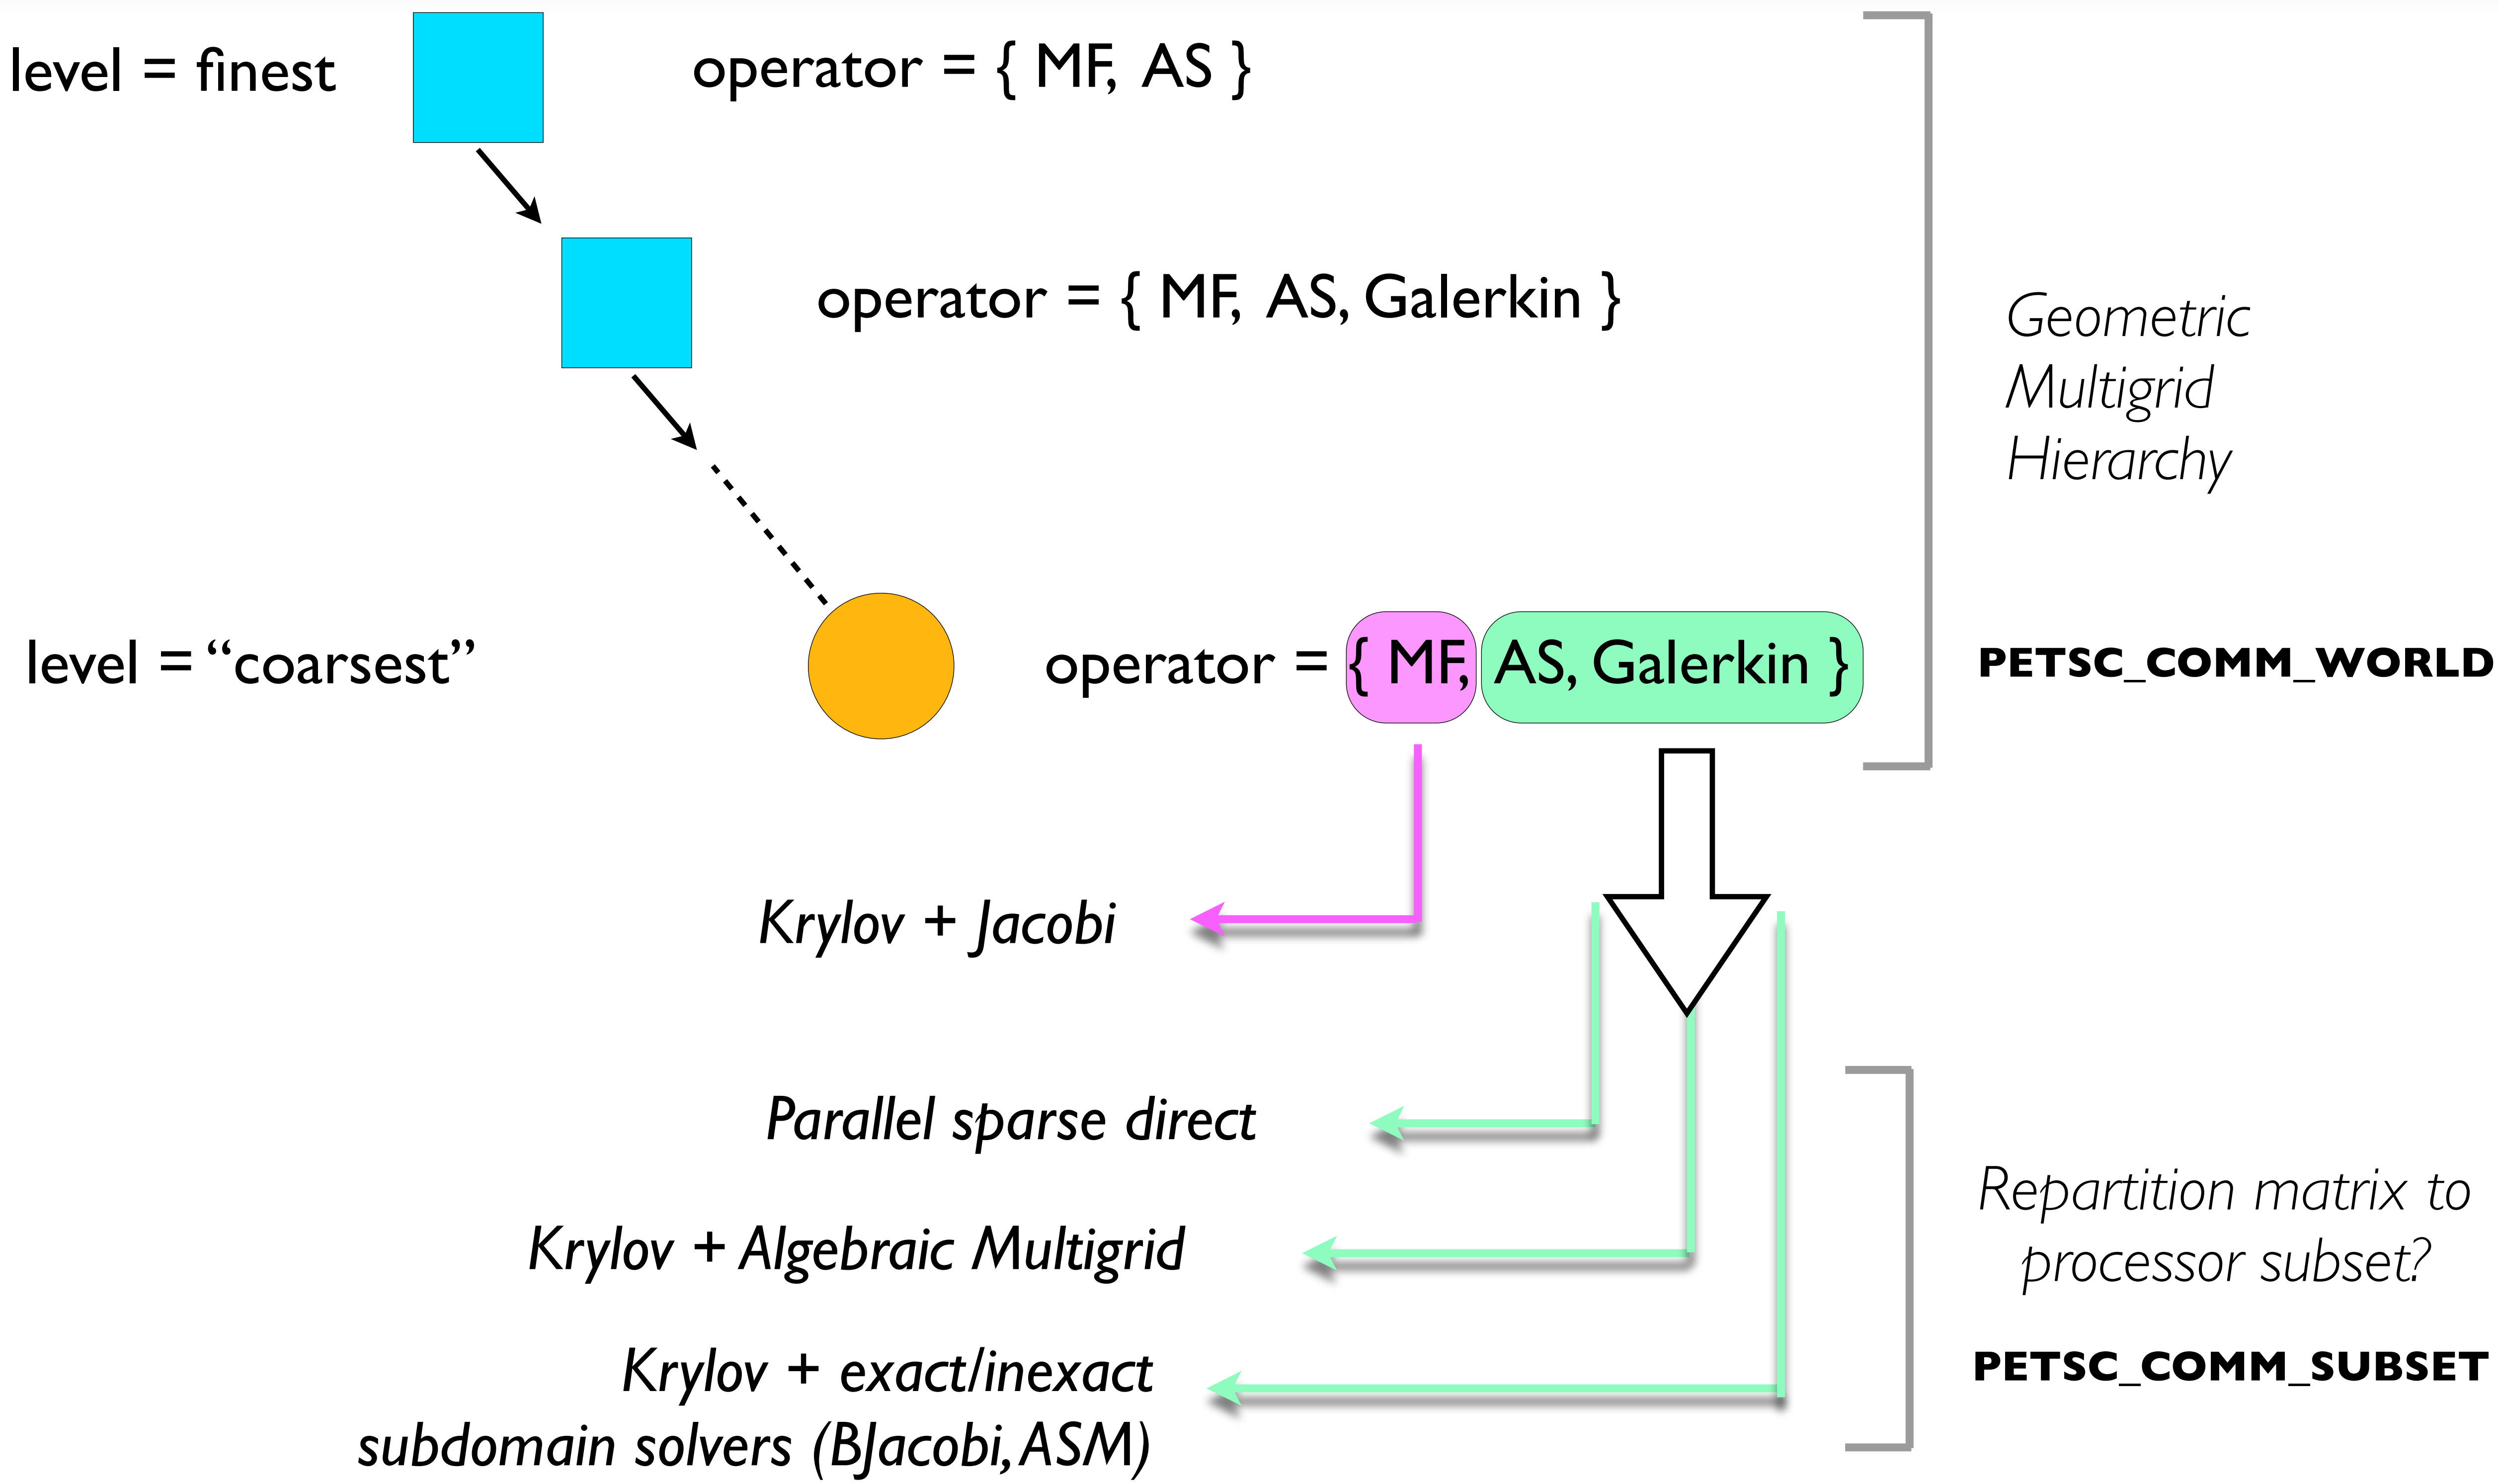

Hybrid Multi Level Strategy for Au=v

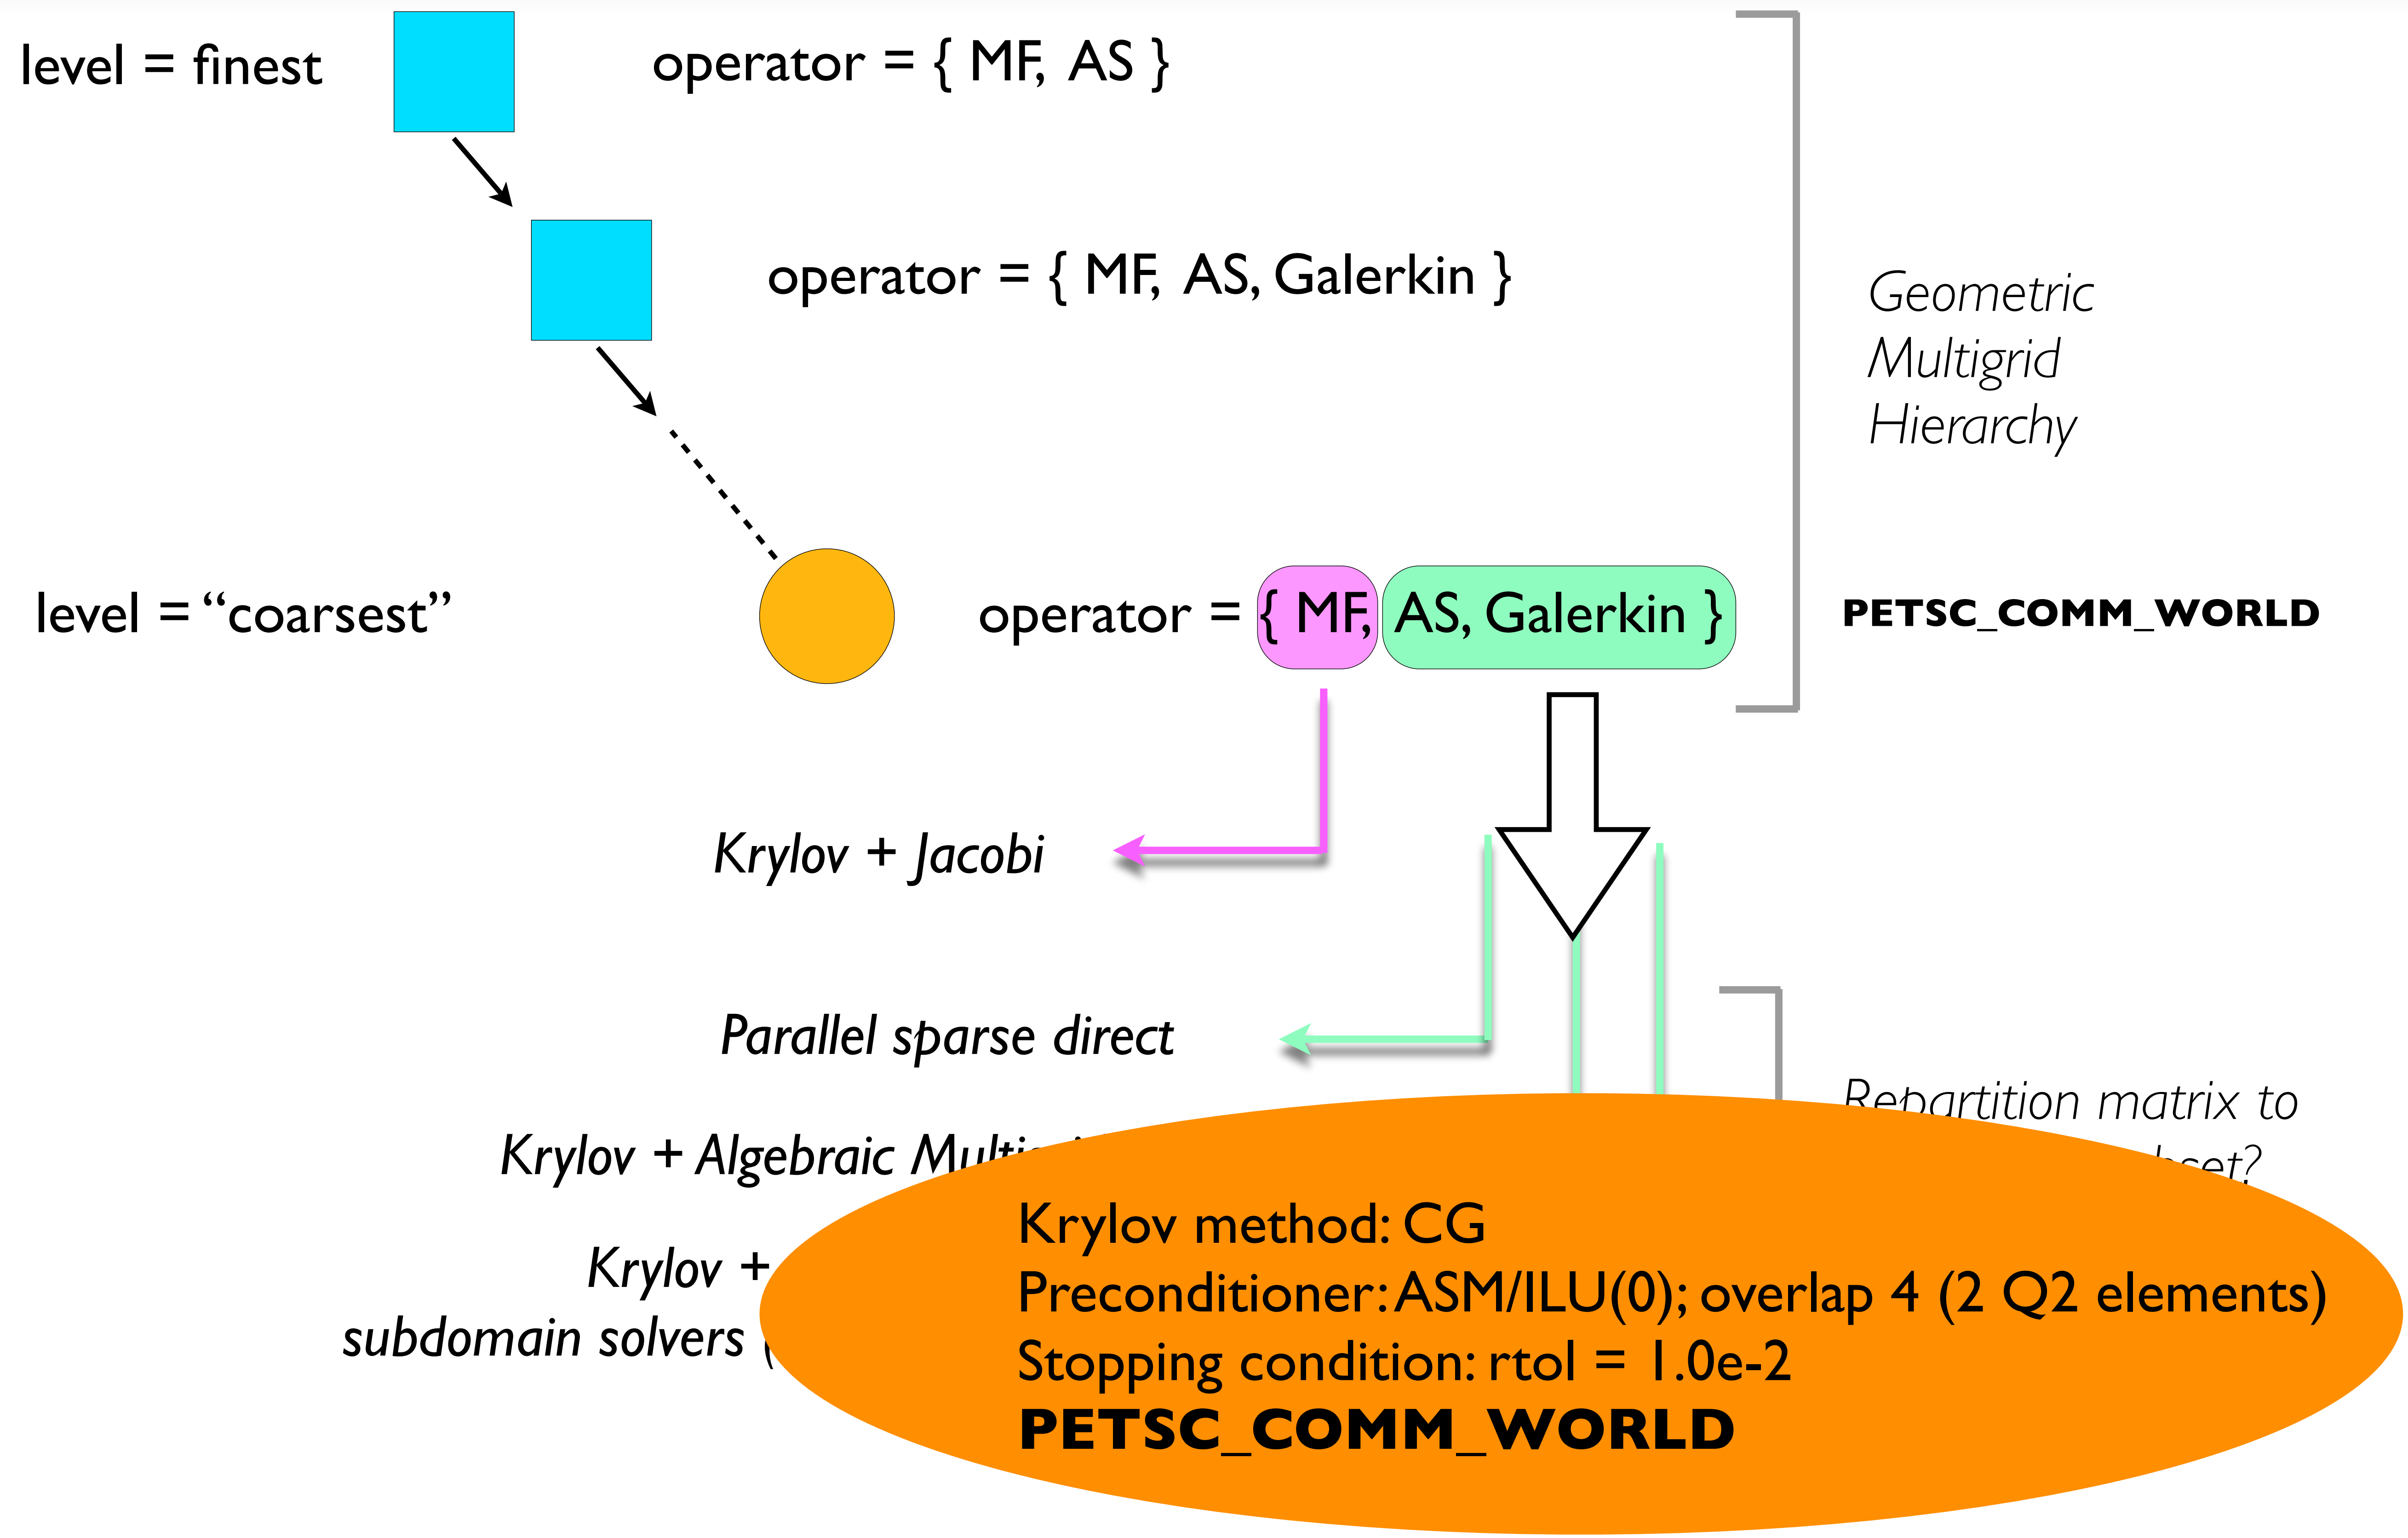

Convergence History: Stokes

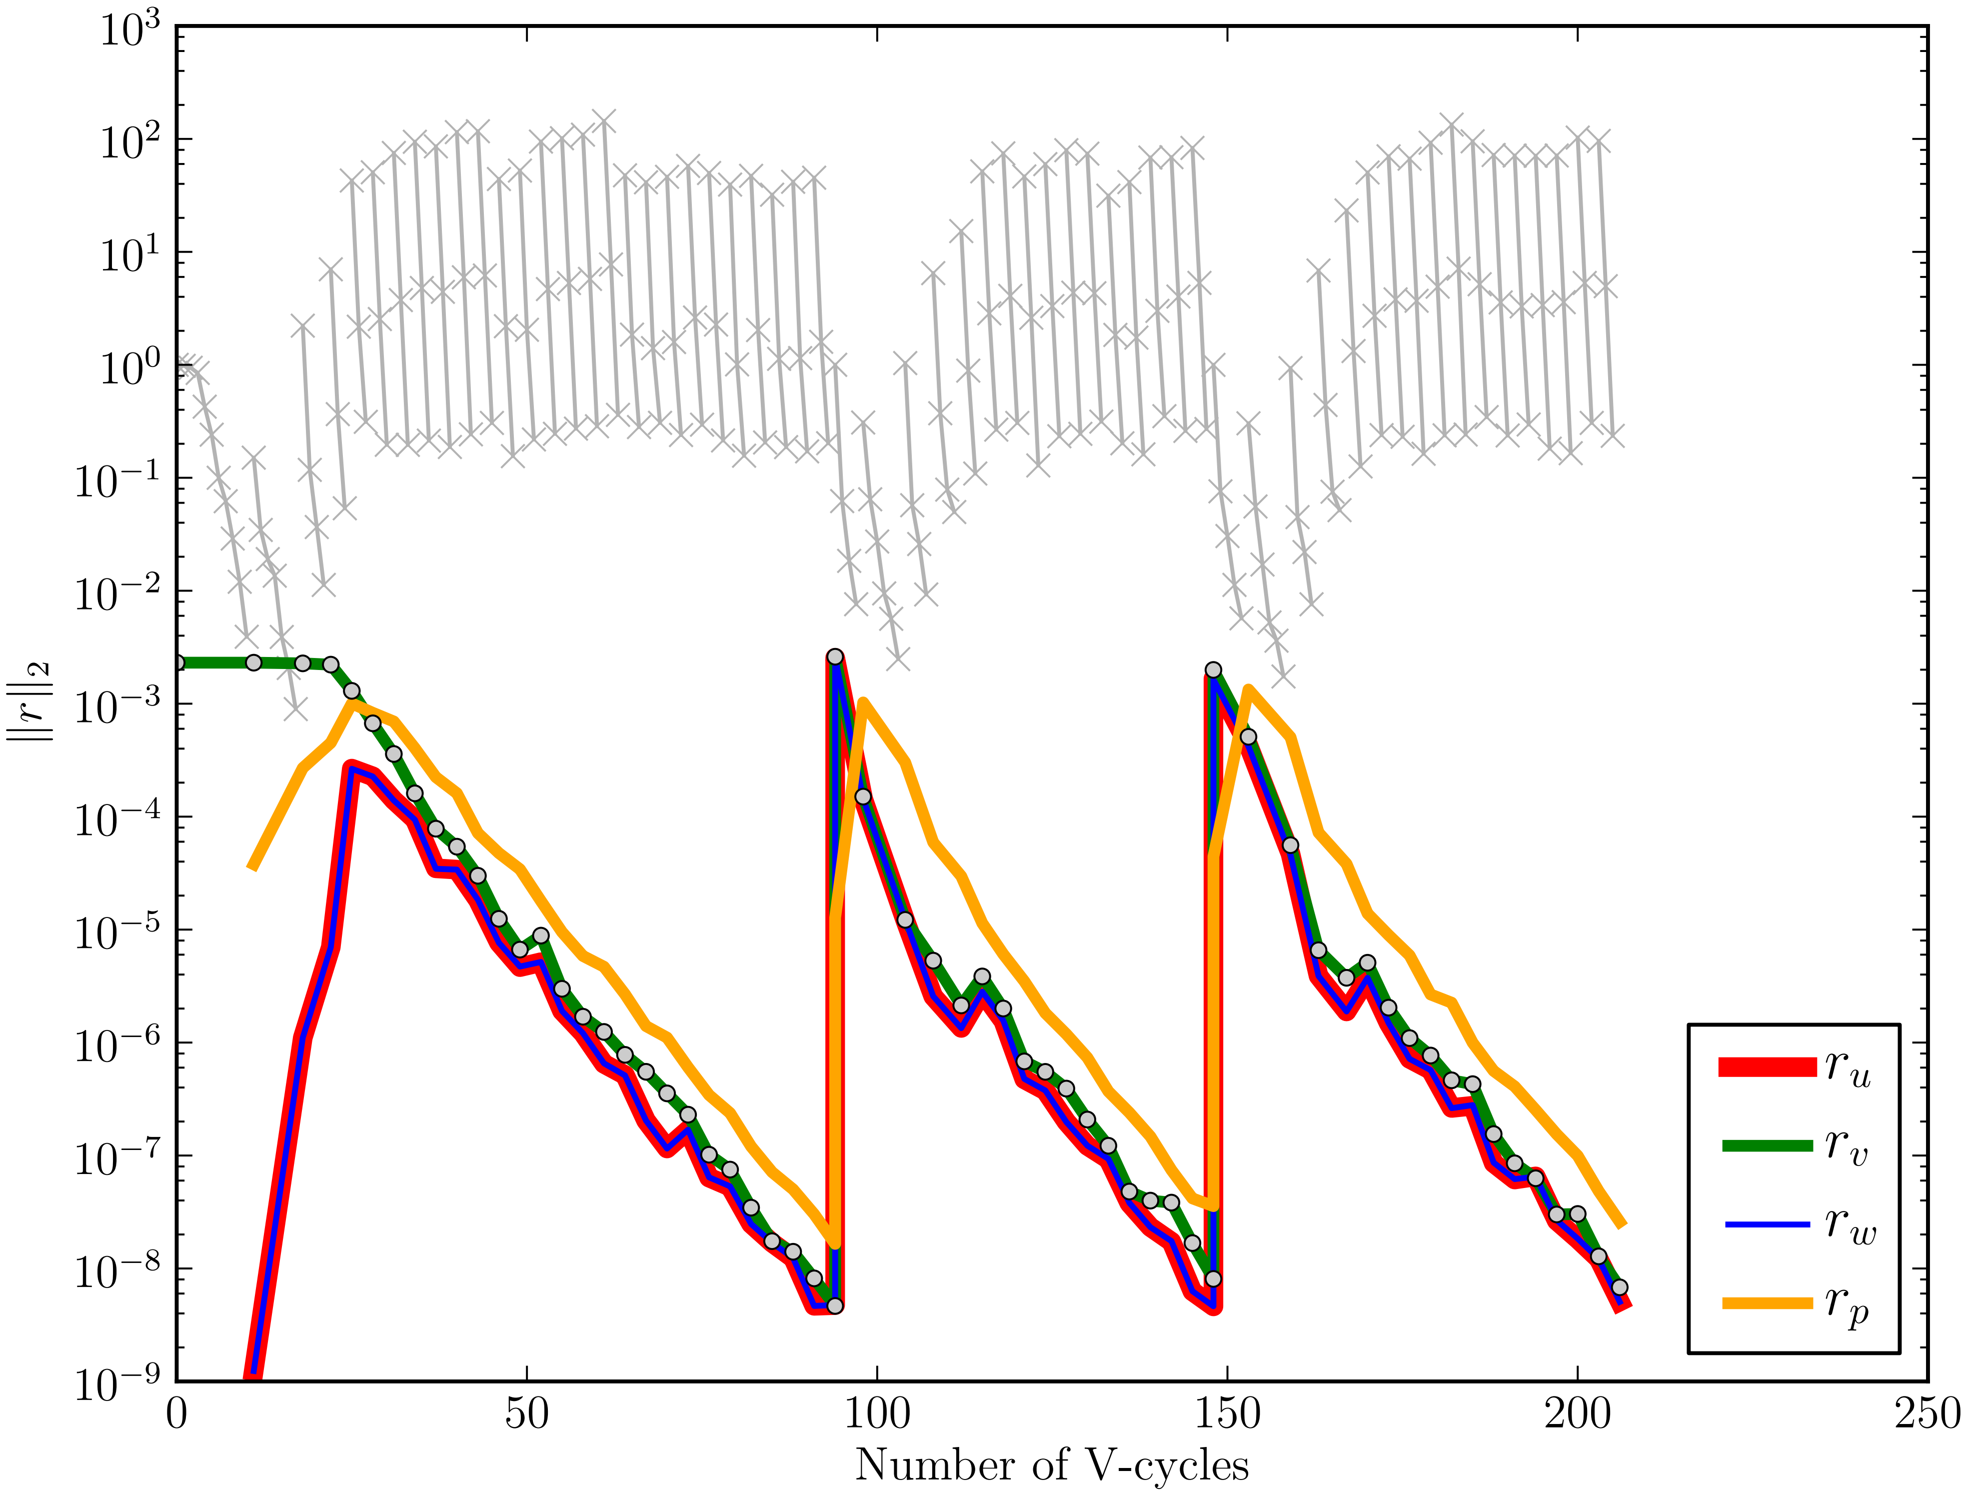

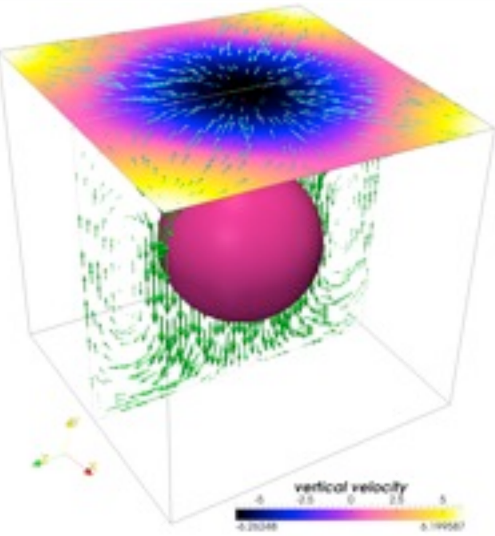

- Sedimenting sphere
  $R = 0.25$
  $\Delta\eta = 10^4$
- 3 time steps
- 32^3 elements
- 3 levels
- Cheby(4) + Jacobi smoother
- LU coarse

Hybrid Multi Level Strategy for Au=v

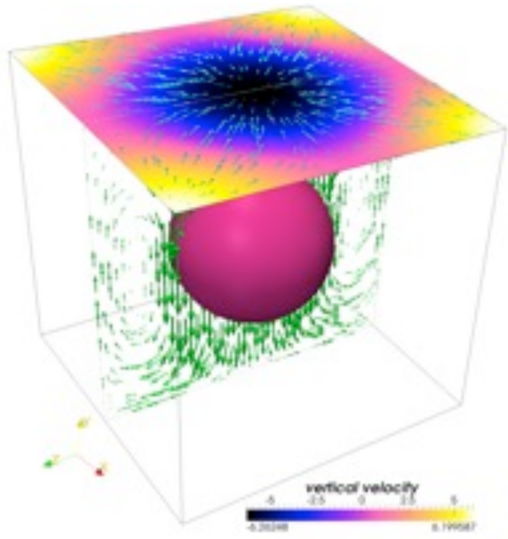

Single CPU test

* 32 x 32 x 32 : 4 MG levels
* Single iteration of Stokes solve
* *A u = v* terminated when initial residual reduced by 1e6
* Smoother: Chebychev/Jacobi - 6 iterations
* Coarse grid: LU

Hybrid Multi Level Strategy for Au=v

Single CPU test

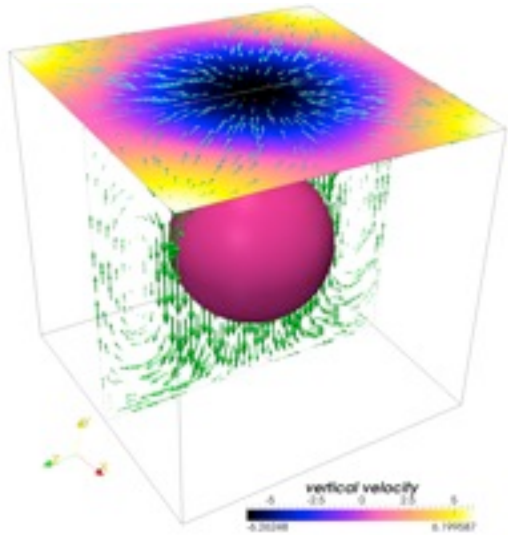

* 32 x 32 x 32 : 4 MG levels
* Single iteration of Stokes solve
* *A u = v* terminated when initial residual reduced by 1e6
* Smoother: Chebychev/Jacobi - 6 iterations
* Coarse grid: LU

Solve time (sec) (#number of iterations)

Coarse level configuration	$\Delta\eta = 10^0$	$\Delta\eta = 10^2$	$\Delta\eta = 10^6$	$\Delta\eta = 10^{10}$	**Mem. (GB)**
A, M, M, M	51 (#5)	123 (#13)	532 (#60)	1605 (#179)	0.7
G, A, M, M	51 (#5)	114 (#12)	185 (#20)	185 (#20)	0.8
G, G, A, M	51 (#5)	87 (#9)	120 (#13)	130 (#14)	1.2
G, G, G, A	43 (#5)	51 (#6)	73 (#9)	80 (#10)	4.4

coarse

G = Galerkin : A = Assembled : M = Matrix-free

Hybrid Multi Level Strategy for Au=v

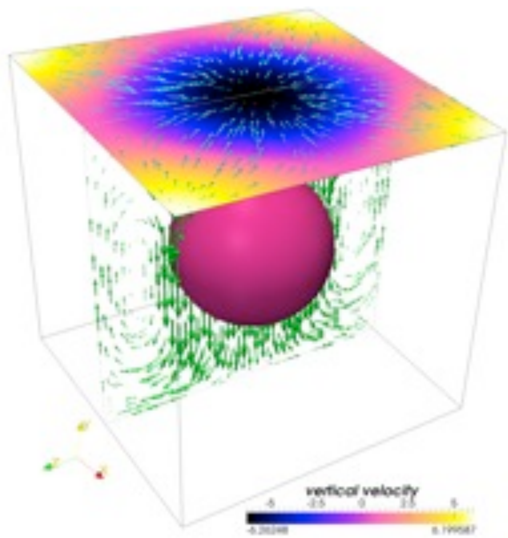

Single CPU test

* 32 x 32 x 32 : 4 MG levels
* Single iteration of Stokes solve
* *A u = v* terminated when initial residual reduced by 1e6
* Smoother: Chebychev/Jacobi - 6 iterations
* Coarse grid: LU

Solve time (sec) (#number of iterations)

Coarse level configuration	$\Delta\eta = 10^0$	$\Delta\eta = 10^2$	$\Delta\eta = 10^6$	$\Delta\eta = 10^{10}$	**Mem. (GB)**
A, M, M, M	51 (#5)	123 (#13)	532 (#60)	1605 (#179)	0.7
G, A, M, M	51 (#5)	114 (#12)	185 (#20)	185 (#20)	0.8
G, G, A, M	51 (#5)	87 (#9)	120 (#13)	130 (#14)	1.2
G, G, G, A	43 (#5)	51 (#6)	73 (#9)	80 (#10)	4.4

coarse

G = Galerkin : A = Assembled : M = Matrix-free

- Significant gains obtained from using "strong" coarse grid operators - memory increase is minimal
- Assembled Galerkin is 38% faster HOWEVER uses 3.7 times more memory

Convergence History: Viscous Block

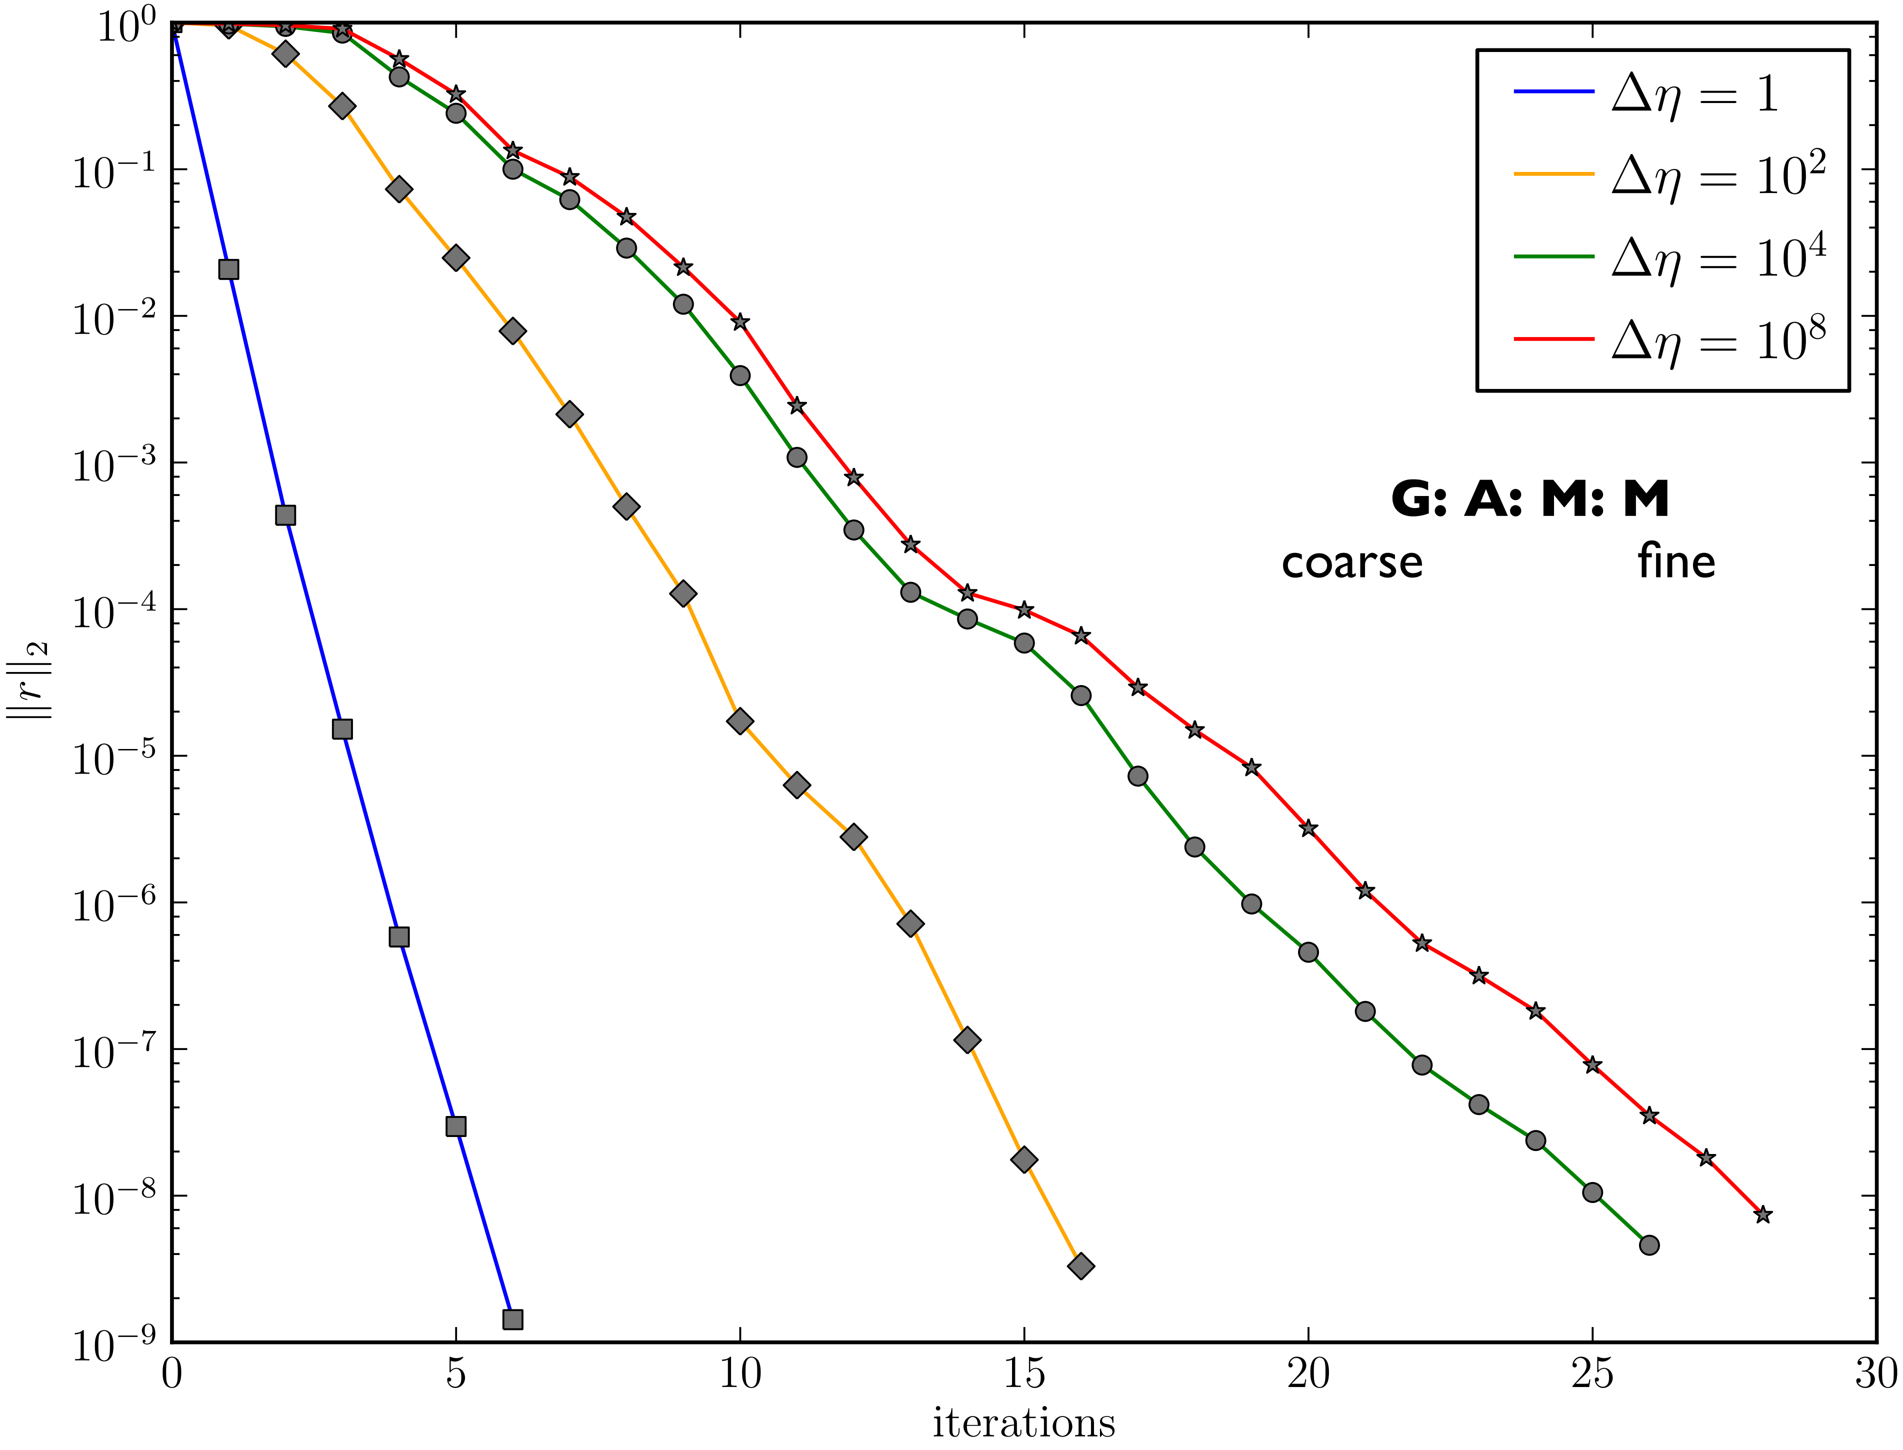

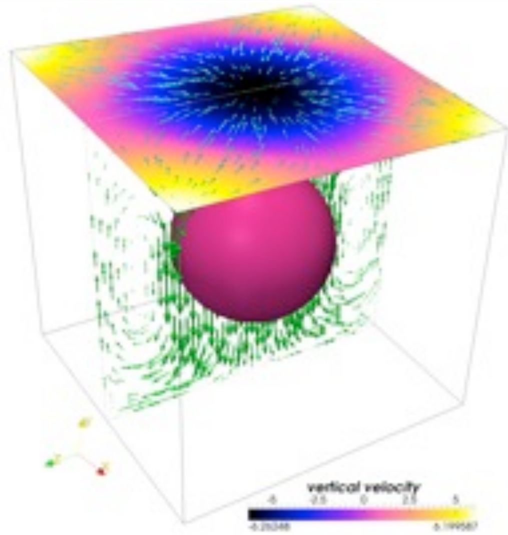

- Sedimenting sphere

- 32^3 elements
- **4 levels**
- Coarse Galerkin
- Cheby(4) + Jacobi smoother
- LU coarse

Parallel Performance

Cores	64	512	4096
Event			
MGSmooth Coarse	1.5907e+02	2.9855e+01	8.8030e+00
MatSolve*	8.6882e+01	1.8791e+01	3.4849e+00
MGSmooth Fine	5.4153e+02	6.8653e+01	9.1118e+00
MatMult*	8.5636e+02	1.1249e+02	1.6264e+01
VecDot*	4.8386e+00	1.1783e+00	1.3234e+00
VecMDot	2.0199e+00	3.8708e-01	2.0188e-01
VecNorm*	1.1429e+01	2.5994e+00	4.4239e-01
KSPGMRESOrthog	7.9125e+00	1.5767e+00	1.5306e+00
KSPSolve	9.6860e+02	1.2980e+02	2.1507e+01
J KSP #	24	24	23
A KSP #	100	101	98
MGCoarse KSP #	347	399	495

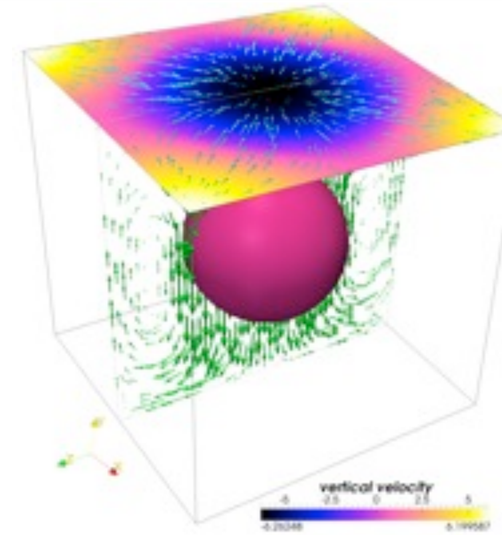

- Sedimenting sphere

$$R = 0.25$$
$$\Delta\eta = 10^4$$

- 96^3 elements
- 3 levels
- Cheby(10) + Jacobi smoother
- Krylov coarse grid solver

Stopping conditions

$$\delta_J^{\text{rel}} = 10^{-5}$$
$$\delta_A^{\text{rel}} = 10^{-2}$$
$$\delta_{A_C}^{\text{rel}} = 10^{-2}$$

Excellent strong scaling.
70% efficiency from 64 to 4096 cores

Oblique Rifting

in collaboration with L. Le Pourhiet (UPMC, Paris)

Major oil reservoirs have been discovered within the last 10 years in the Equatorial Atlantic. These oil fields were not explored before as companies had classified such “oblique continental margins” as having very low oil potential - an assumption which was largely based on *state-of-the-art 2D modeling*.

3 branches of oblique rifting co-exist

Oblique Continental Margins

- Oceanic crust
- Thin continental crust
- Uplifiting Area
- Subsiding Area

One branch is abandoned. The most oblique is favoured

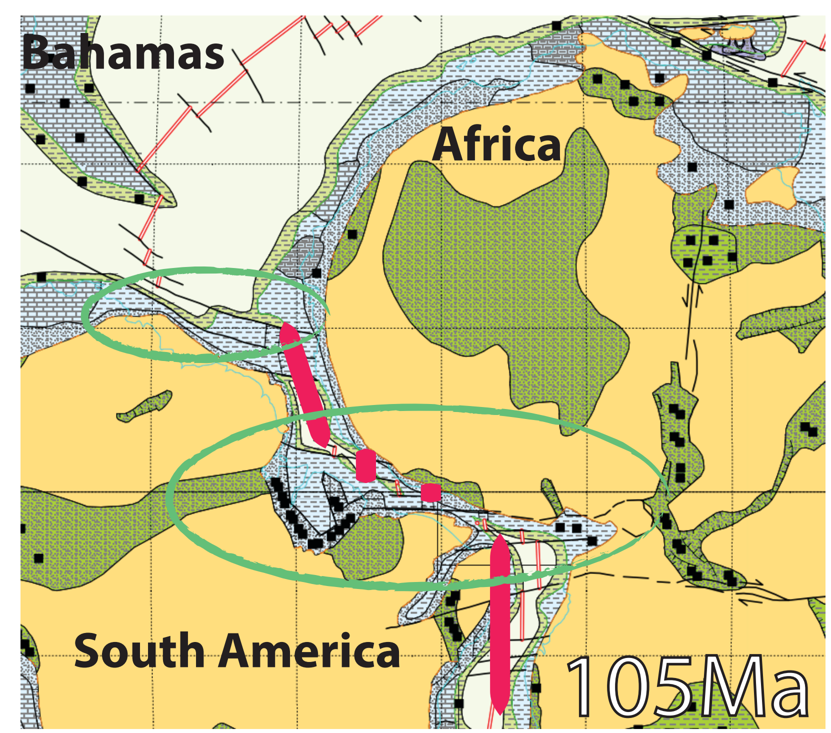

Oblique Rifting

in collaboration with L. Le Pourhiet (UPMC, Paris)

Major oil reservoirs have been discovered within the last 10 years in the Equatorial Atlantic. These oil fields were not explored before as companies had classified such “oblique continental margins” as having very low oil potential - an assumption which was largely based on *state-of-the-art 2D modeling*.

3 branches of oblique rifting co-exist

Strike slip systems in oblique settings have not been self-consistently modelled before.

Oblique Continental Margins

- Oceanic crust
- Thin continental crust
- Uplifiting Area
- Subsiding Area

3D modelling of continental rifting and break-up is numerically challenging as it requires;

- large domains, 4000 km x 4000 km x 300 km
- simulations to be performed over large time spans, > 30 million years
- resolving the influence of strongly non-linear material behaviour and large viscosity contrasts (1e6) between thin layers (< 10 km)

One branch is abandoned. The most oblique is favoured

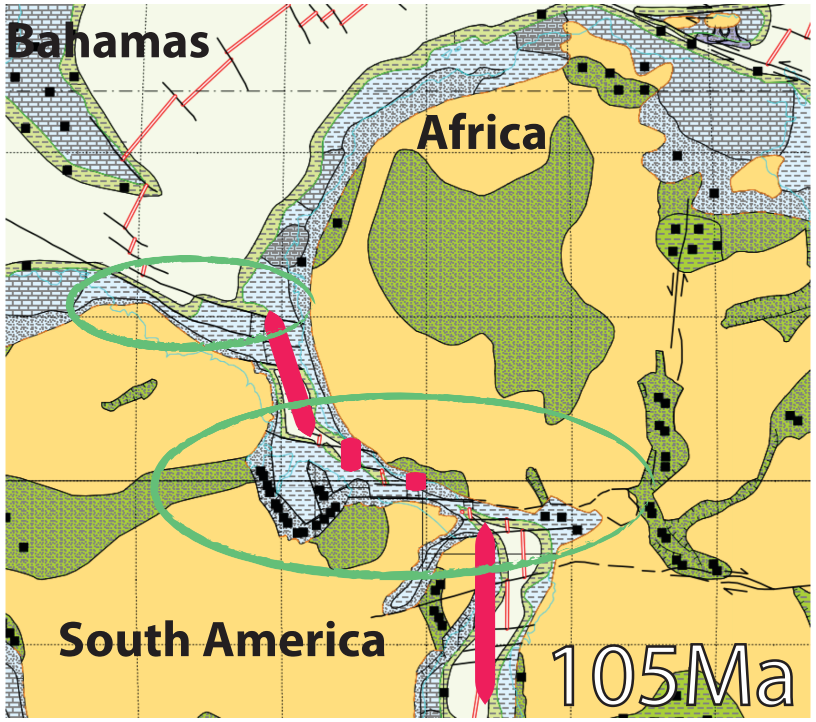

Understanding Origins of Obliquity

Non-linear, Arrhenius + Drucker Prager rheology (brittle / ductile)

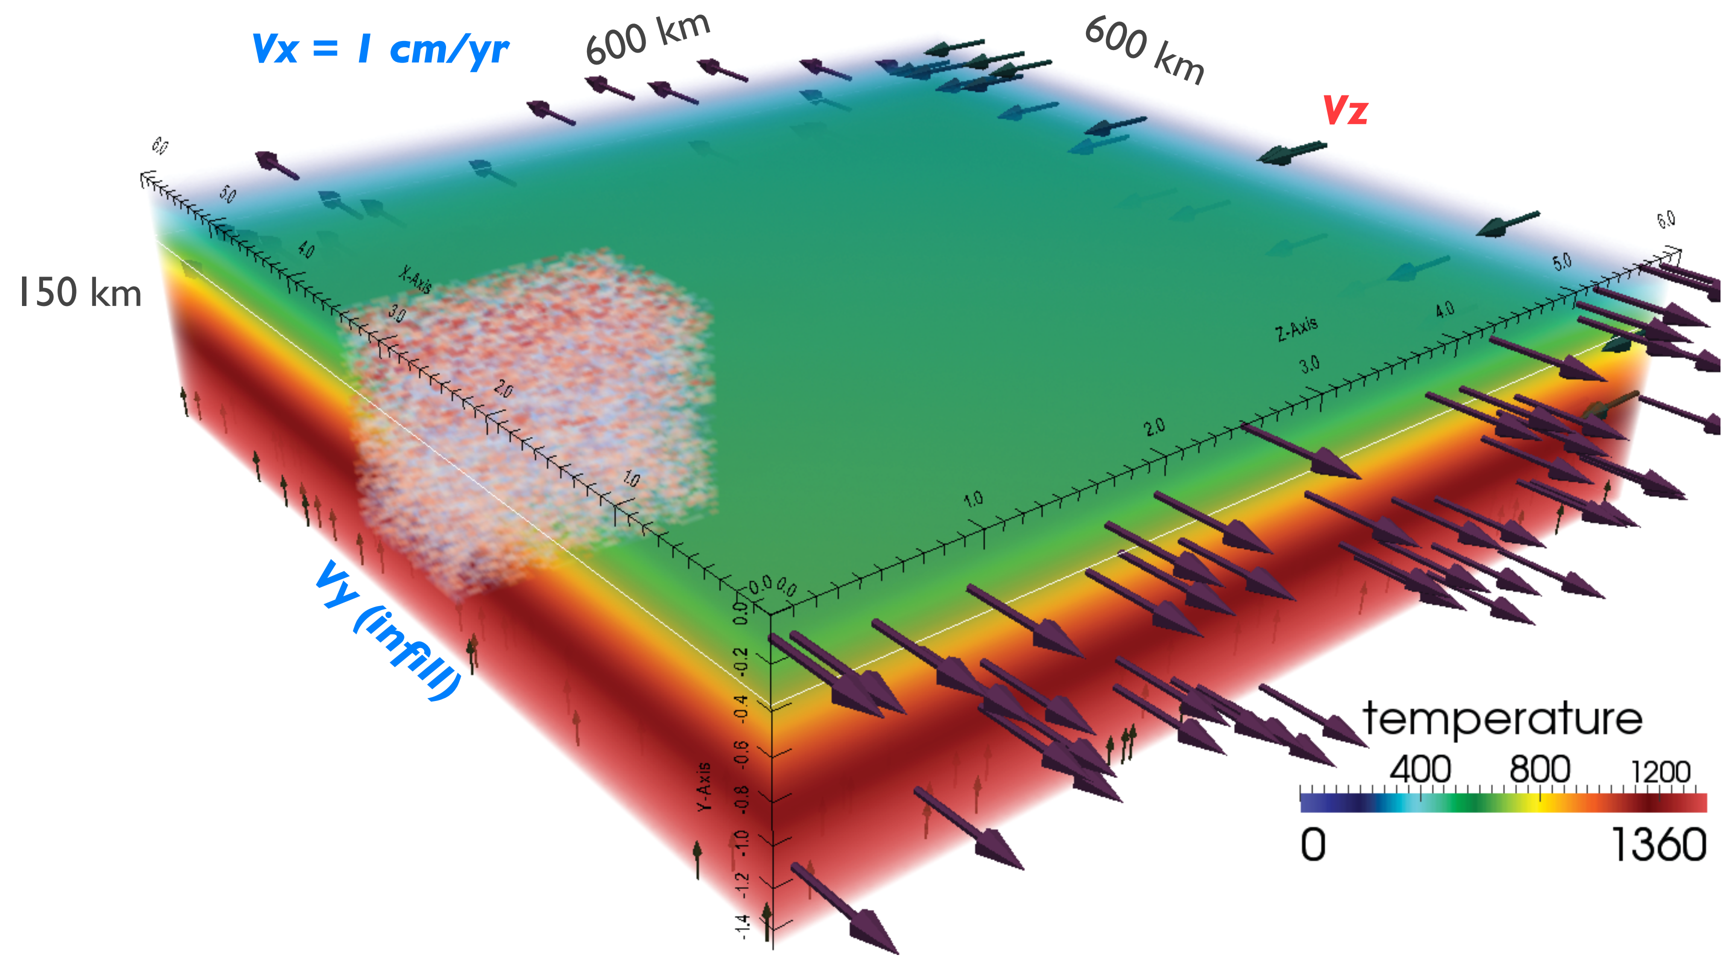

Understanding Origins of Obliquity

Rifting phase (5-10 Myr)

Only models with shortening normal to extension develop oblique branches

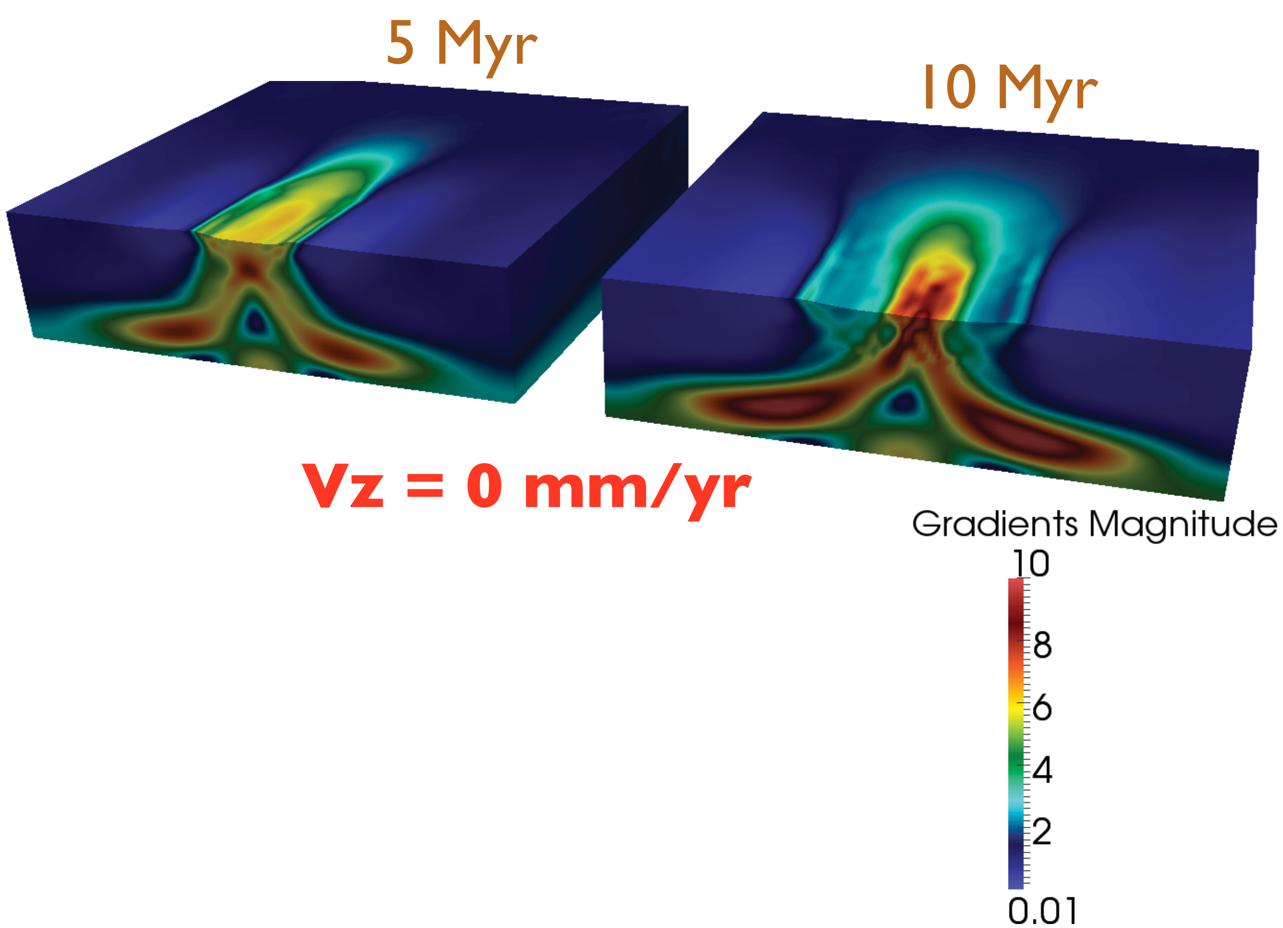

Non-linear, Arrhenius + Drucker Prager rheology (brittle / ductile)

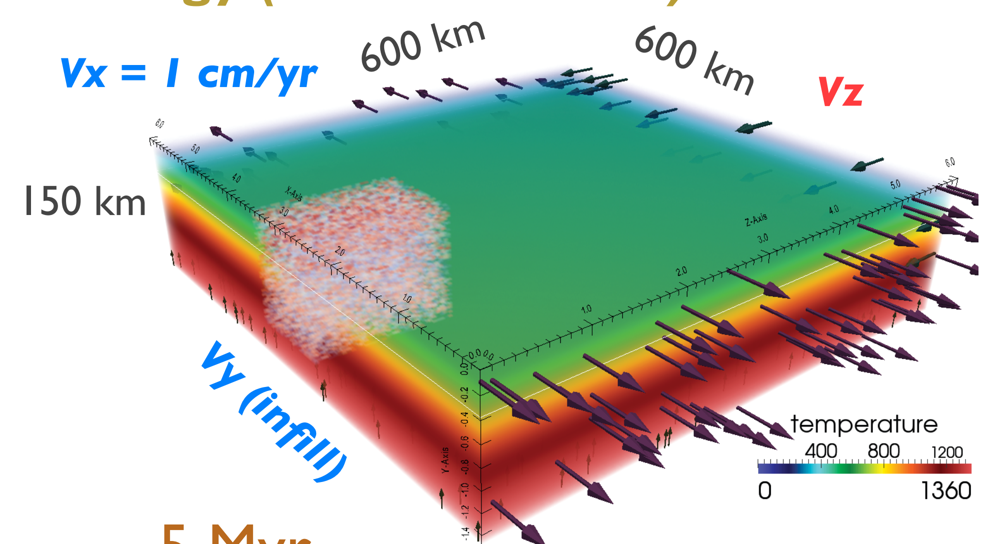

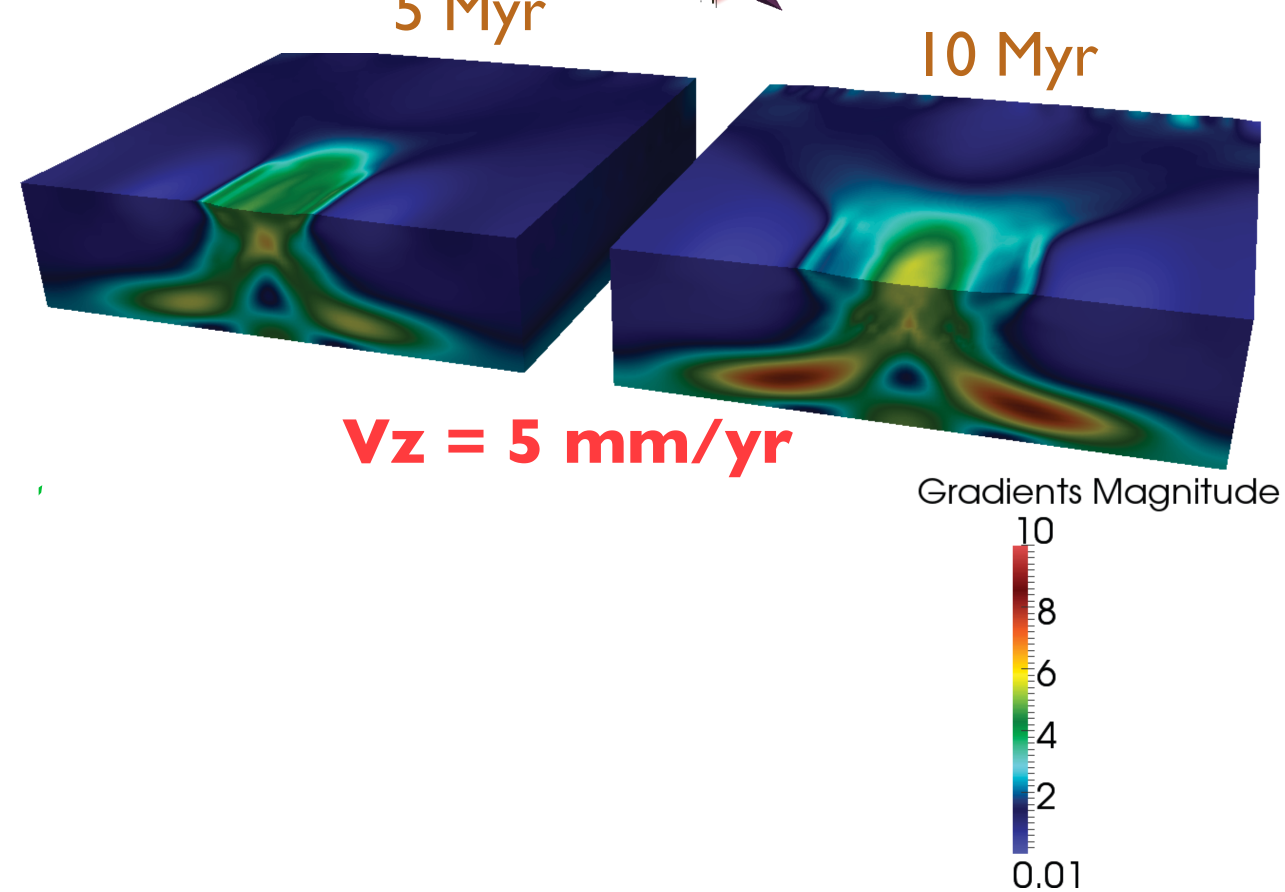

Understanding Origins of Obliquity

Rifting phase (5-10 Myr)

Only models with shortening normal to extension develop oblique branches

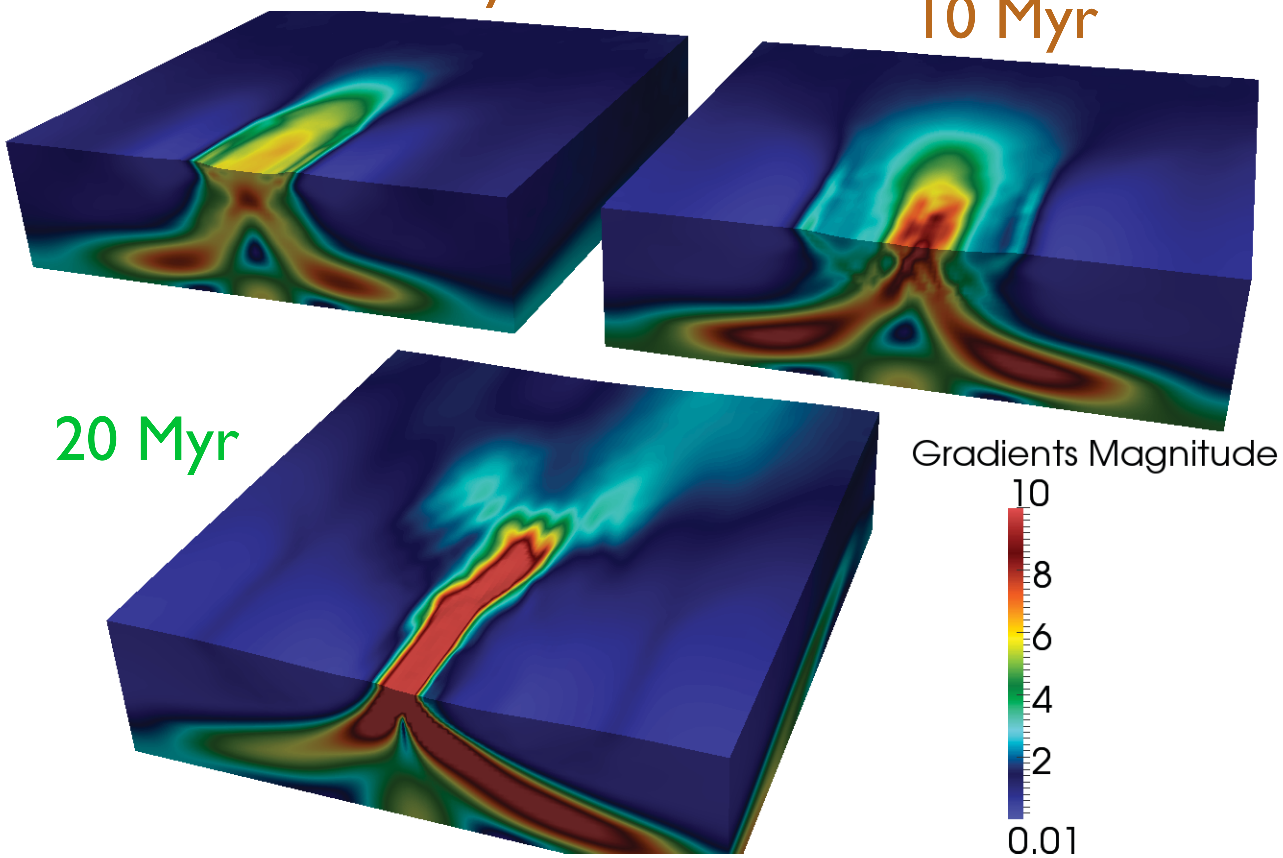

Non-linear, Arrhenius + Drucker Prager rheology (brittle / ductile)

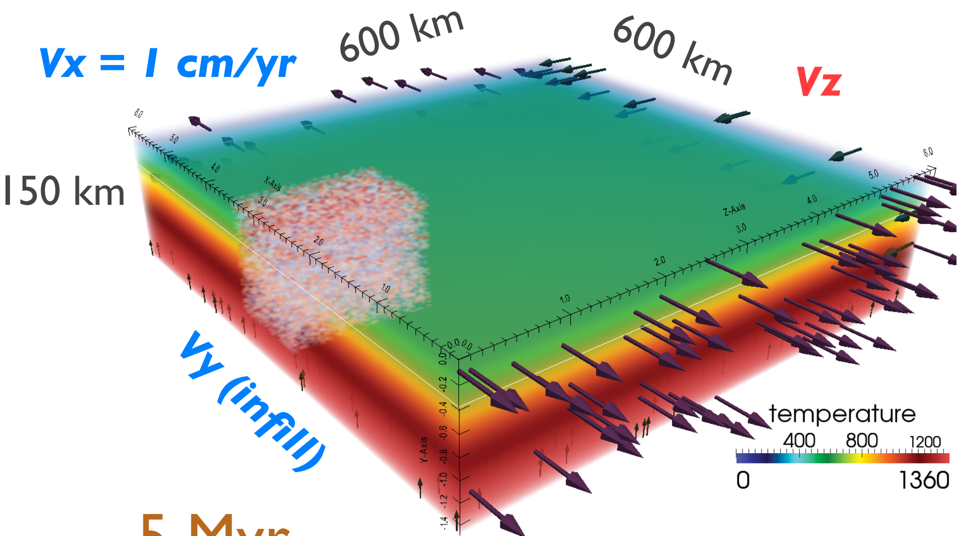

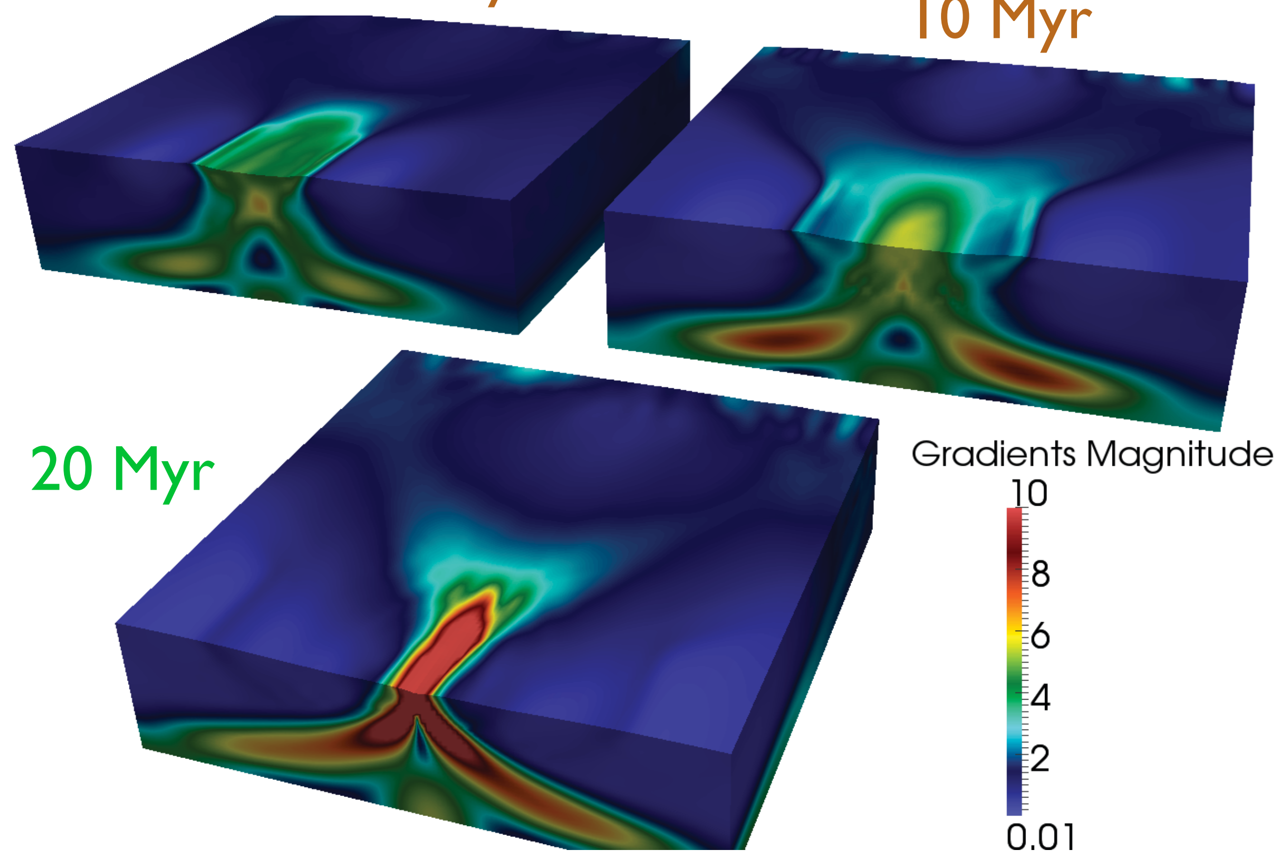

Break-up phase (~20 Myr)

Once continental break-up occurs (in the south), model with no shortening develops oblique rifting branches, self consistently. The timing when obliquity occurs may help constrain boundary conditions.

Understanding Origins of Obliquity

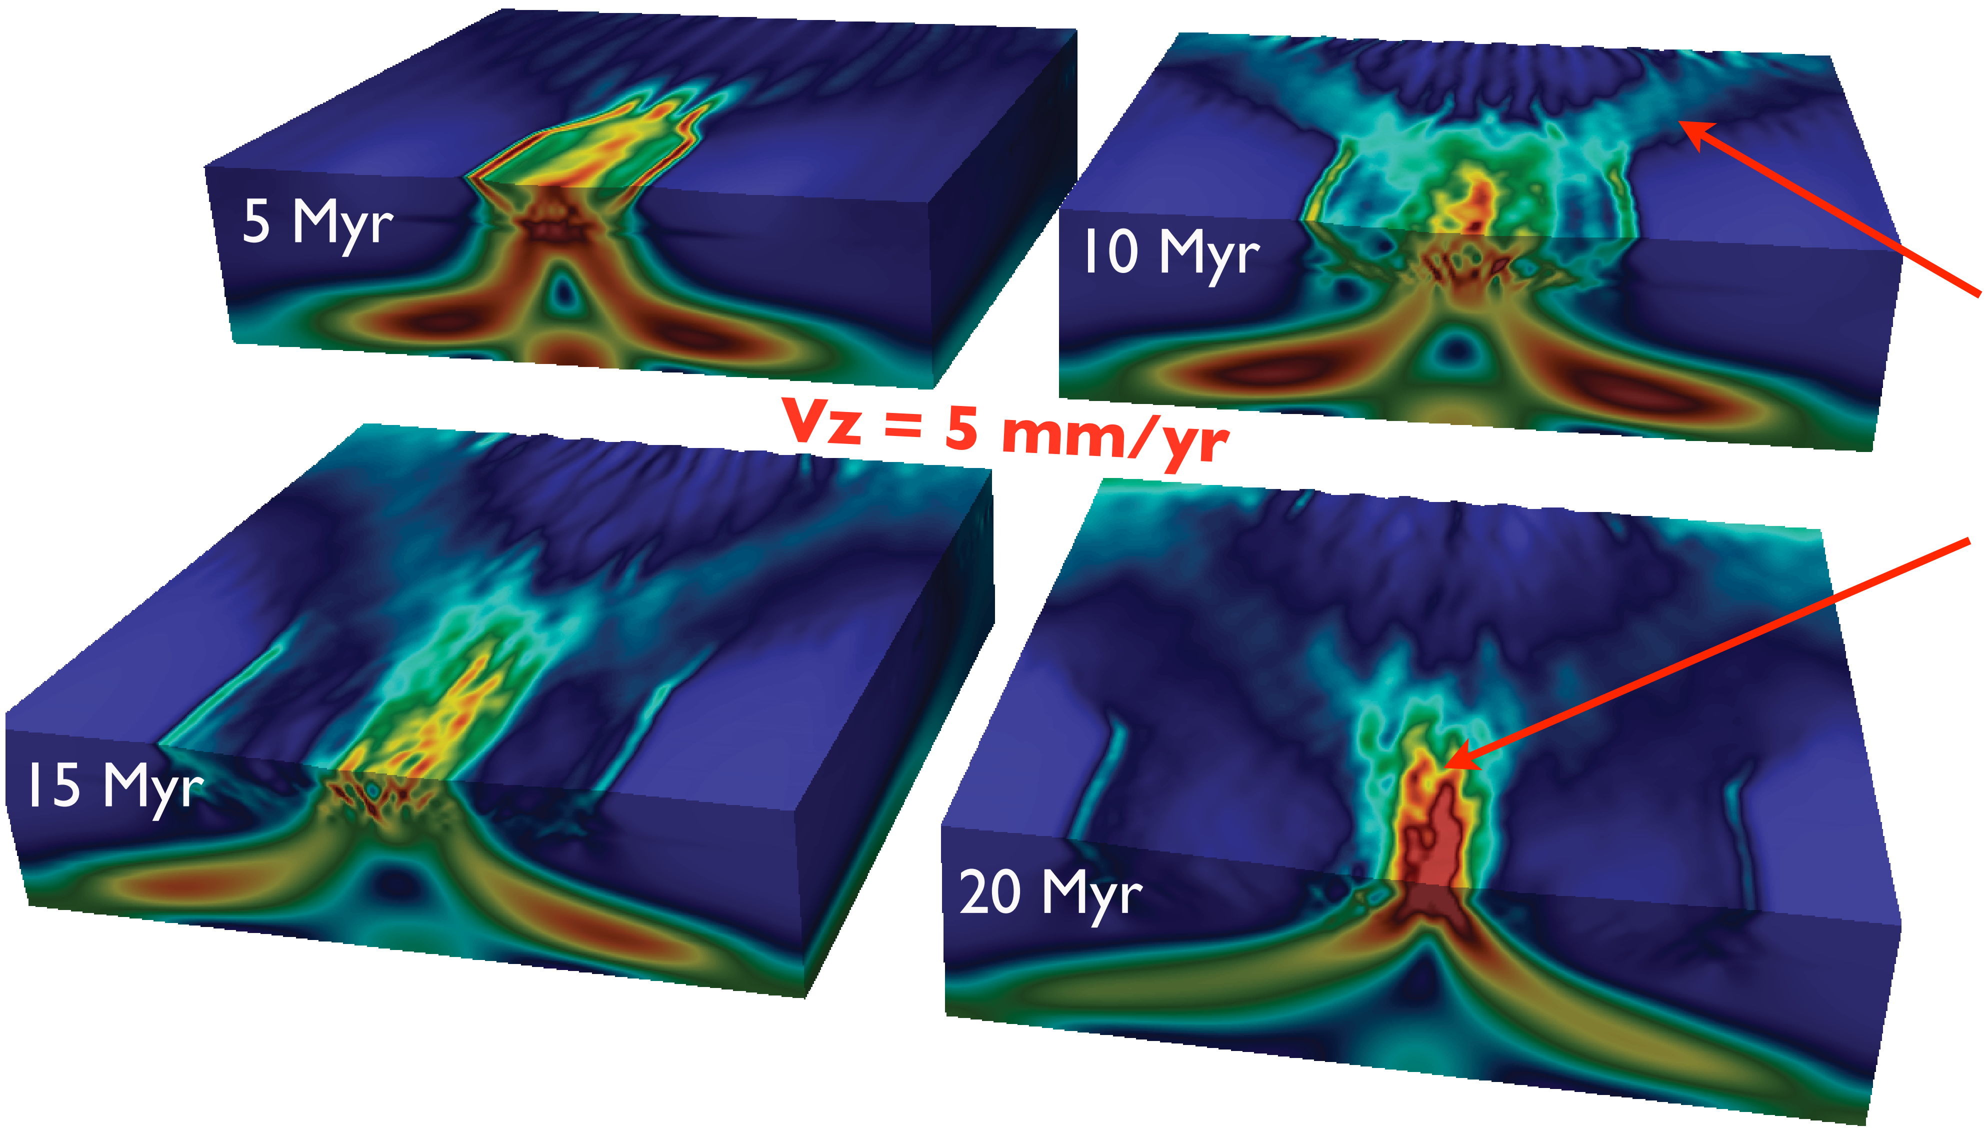

Oblique margin cross-cut by faults aligned with spreading direction

Once break-up occurs, it propagates by small segments with step over

Understanding Origins of Obliquity

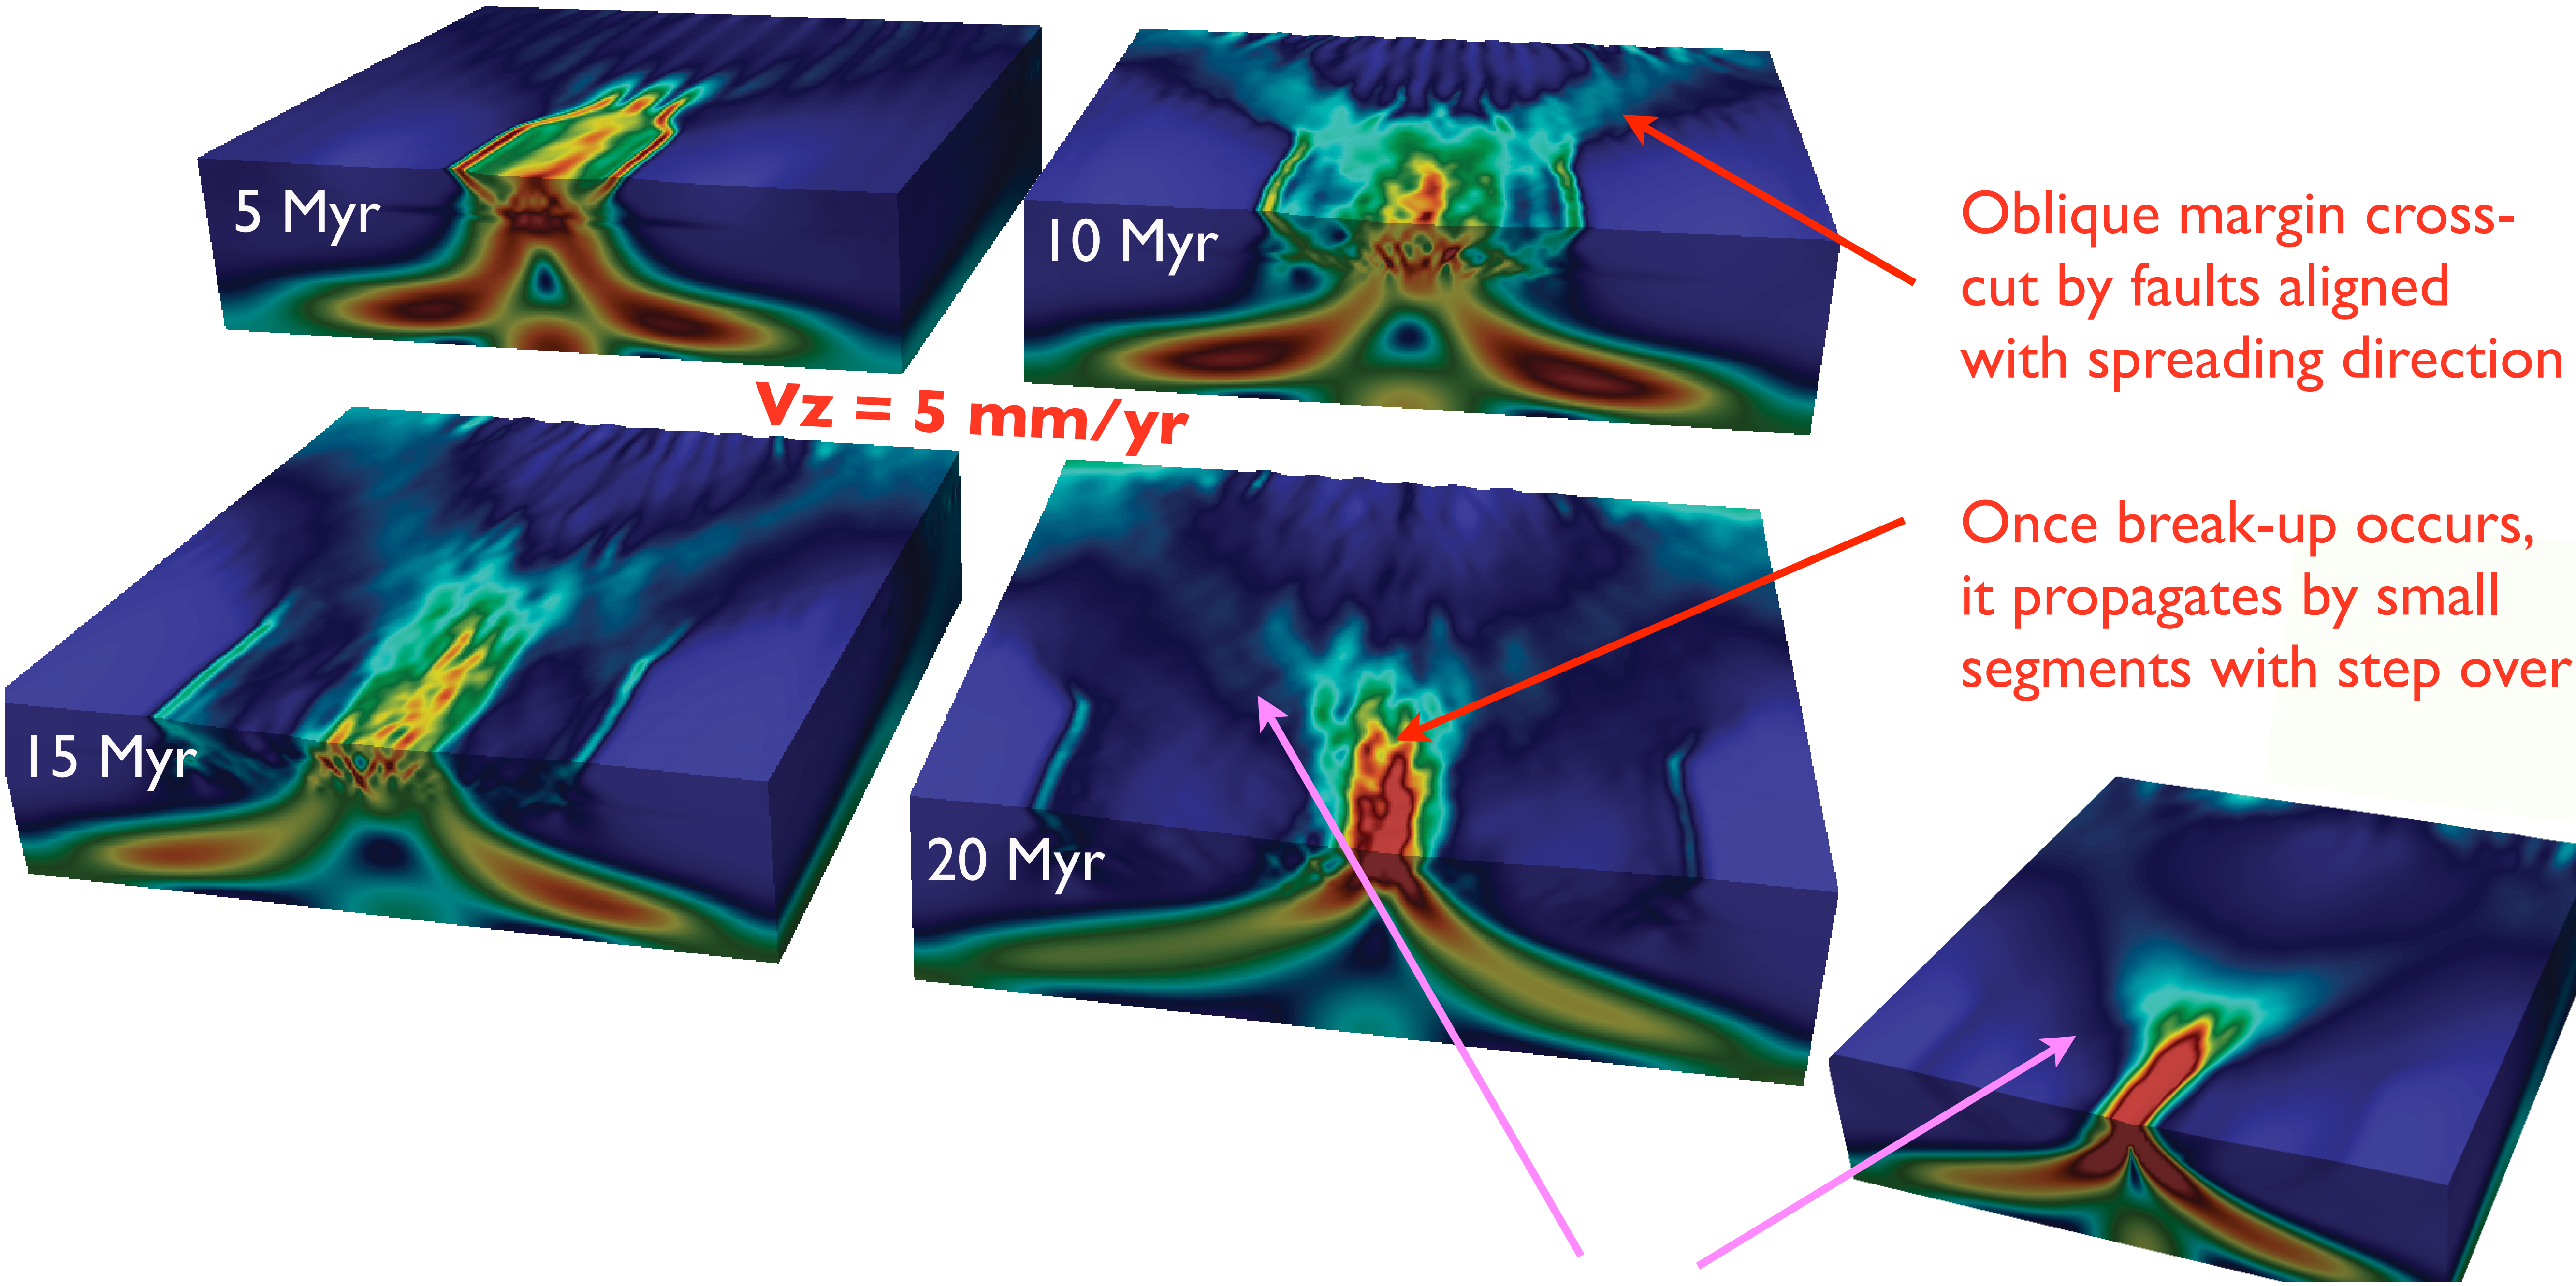

At lower resolution, the asymmetry of propagation, the small faults cutting the oblique margin are not resolved. Such structures are necessary to valid models from field geology

Understanding Origins of Obliquity

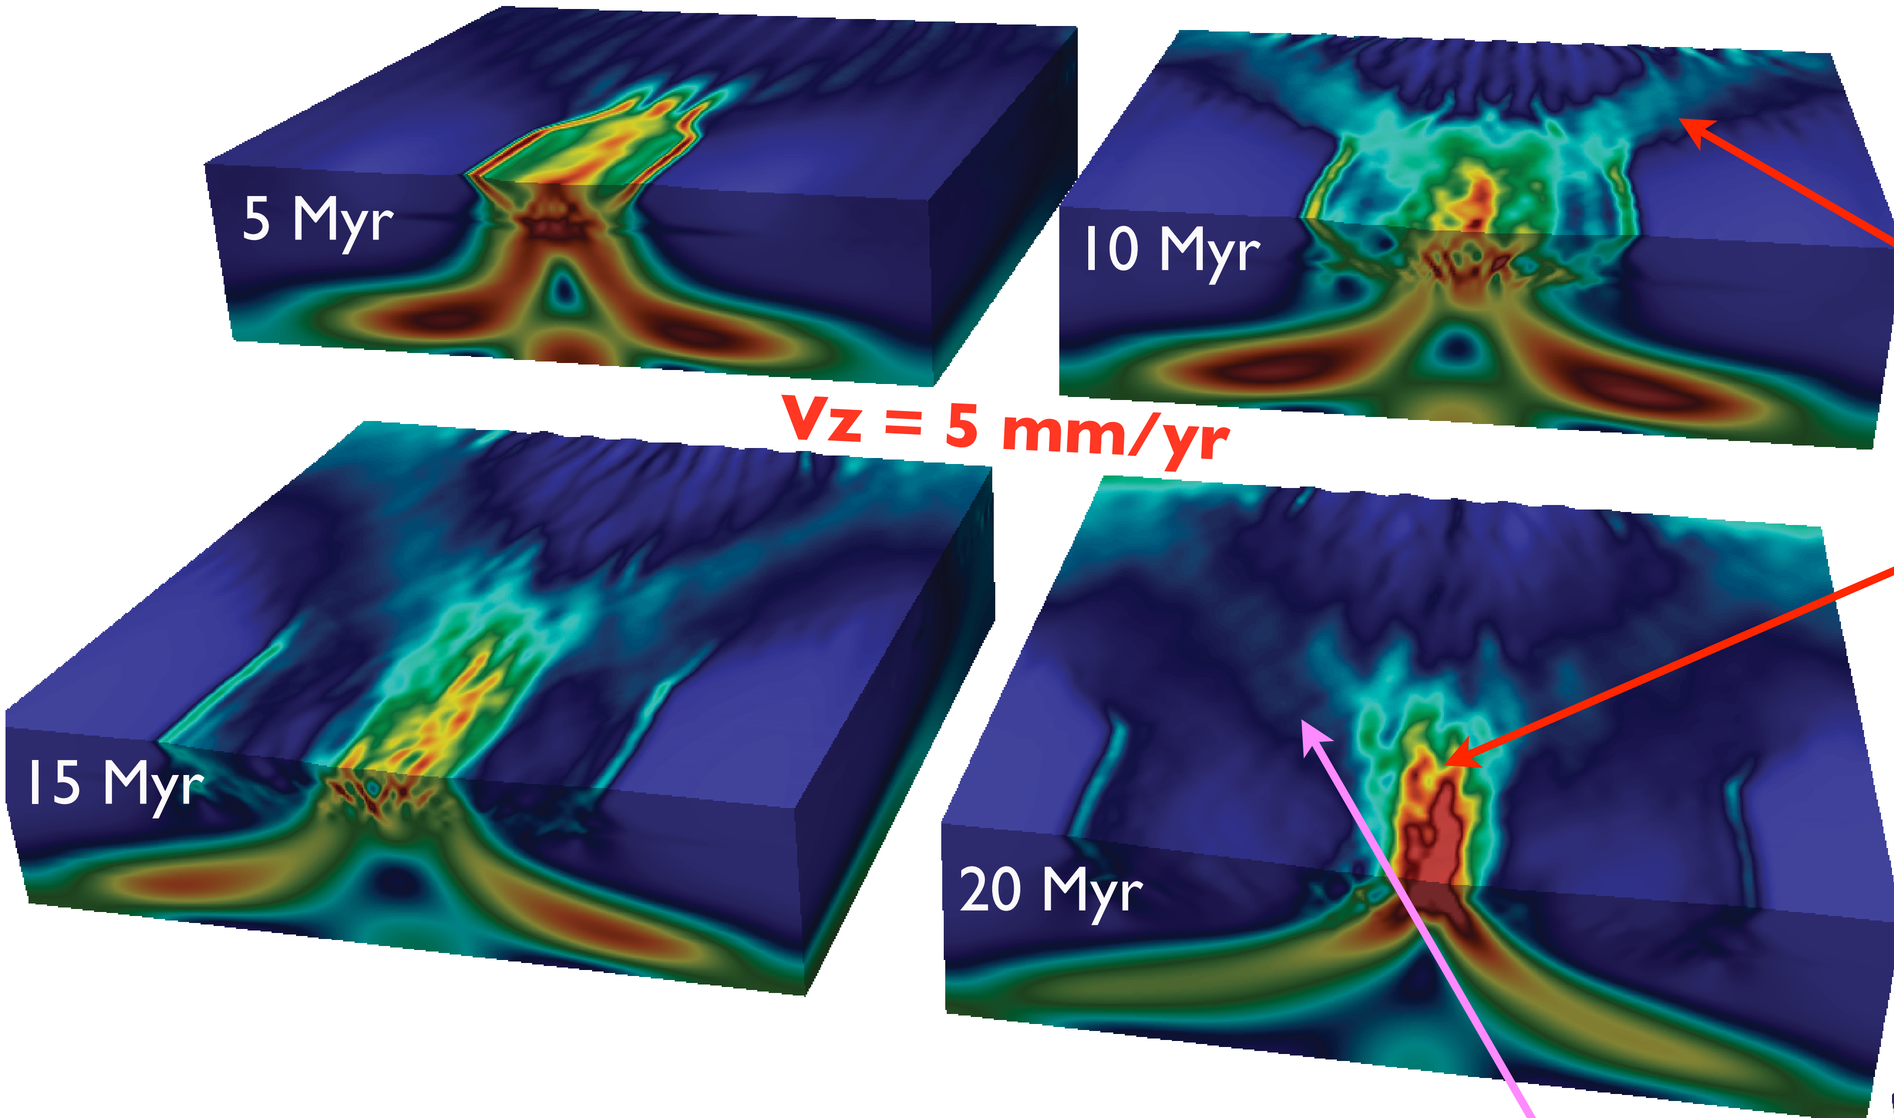

Oblique margin cross-cut by faults aligned with spreading direction

Once break-up occurs, it propagates by small segments with step over

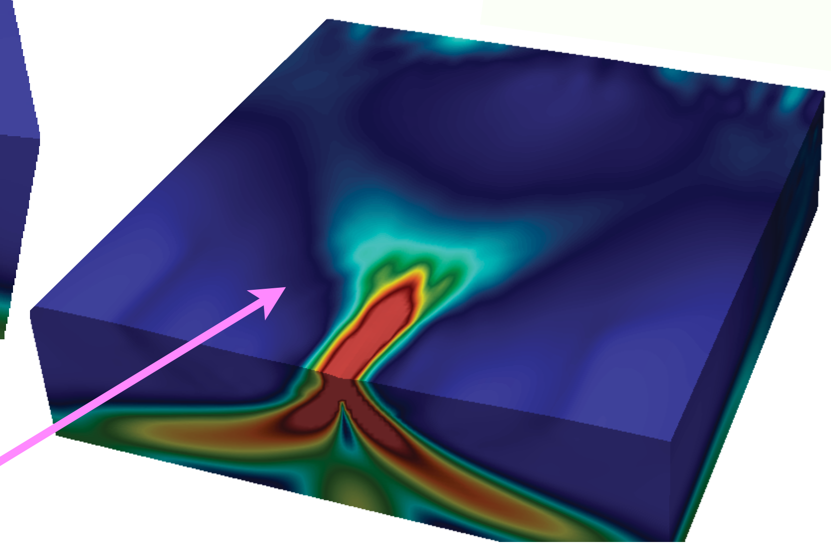

* 128 x 64 x 128 Q2 elements, 20 Myr (1800 time steps, full non-linear solve) 64 cores with ~1 GB/core in ~11 days

* Same model runs on 18 hours on 1024 cores [Cray XE6] due to good strong scaling capabilities

At lower resolution, the asymmetry of propagation, the small faults cutting the oblique margin are not resolved. Such structures are necessary to valid models from field geology

Rifting at Scale

- Geometry aspect ratio: 12 x 1.5 x 6
- Three level MG hierarchy
  - Viscosity gradients largest in y (coarsening less frequently)
  - Coarsen aggressively in directions with high aspect ratio
- Use Chebyshev + Jacobi smoothers
- Krylov coarse grid solver

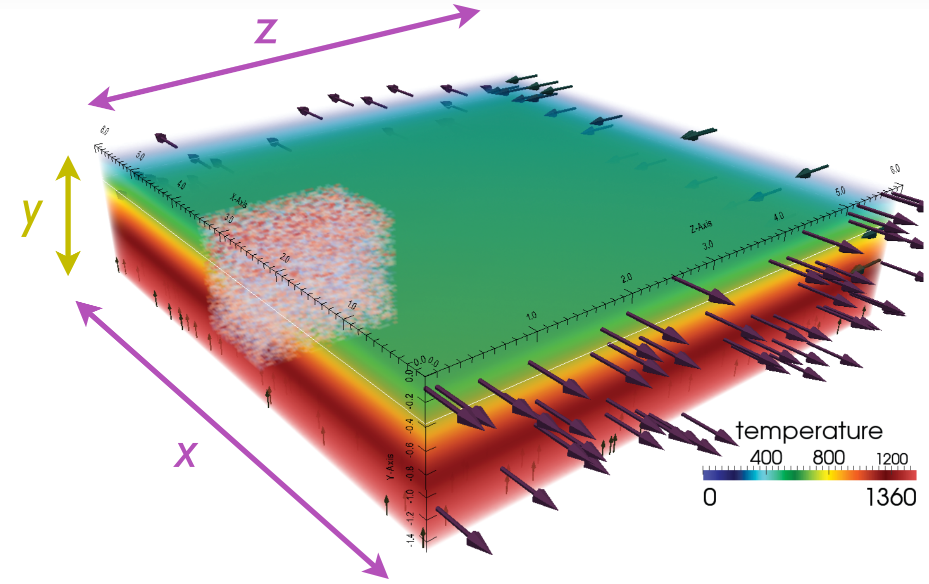

Mesh 1	***Mesh 2***	***Mesh 3***
256 x 32 x 128	512 x 64 x 256	1024 x 128 x 512
64 x 16 x 32	128 x 32 x 64	256 x 64 x 128
32 x 16 x 16	64 x 32 x 32	128 x 64 x 64
30 million DOFs	*237 million DOFs*	*1.9 billion DOFs*

- Results presented were performed using "Kraken" [Cray XT5]

Rifting at Scale

Mesh 1	*Mesh 2*	*Mesh 3*
256 x 32 x 128	512 x 64 x 256	1024 x 128 x 512
30 million DOFs	*237 million DOFs*	*1.9 billion DOFs*

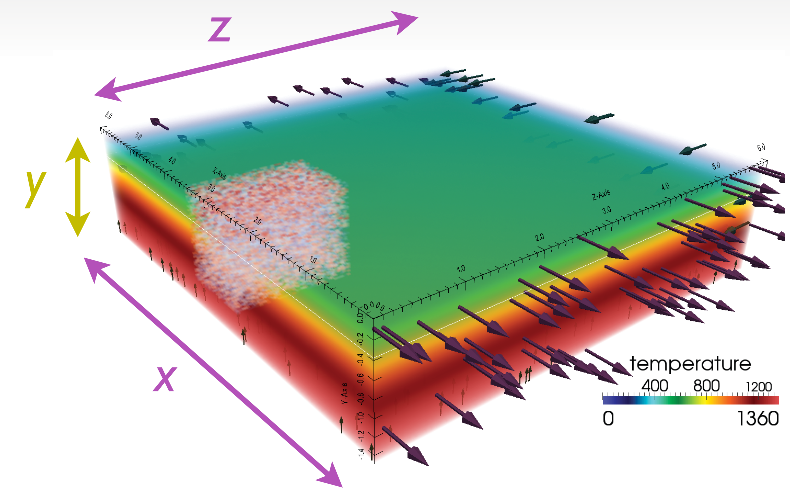

CPU time (sec) per *iteration* of Stokes problem,
15-20 iterations required per Newton step

	cores	Linear	Picard
Mesh 1	*512*	3.57	3.88
Mesh 2	*4096*	8.48	5.90
Mesh 3	*32786*	7.86	7.12

weak scaling

cores	Linear	Picard
2048	2.04	2.22
16384	4.56	4.71

strong scaling

Rifting at Scale

Mesh 1
256 x 32 x 128
30 million DOFs

Mesh 2
512 x 64 x 256
237 million DOFs

Mesh 3
1024 x 128 x 512
1.9 billion DOFs

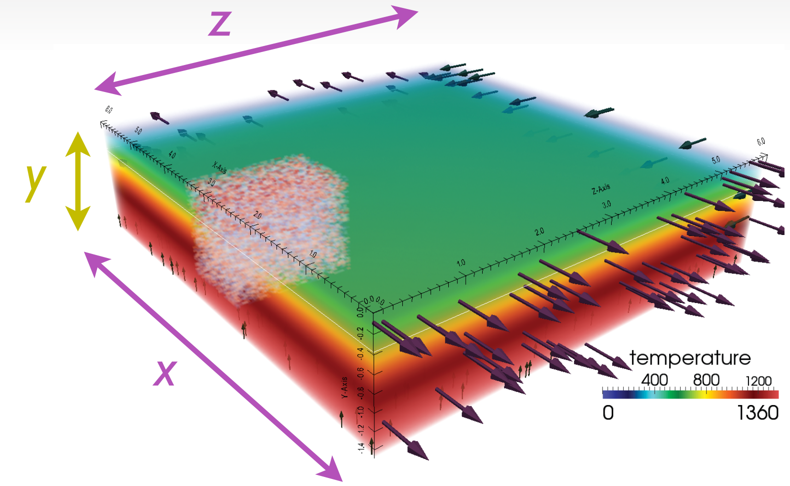

CPU time (sec) per *iteration* of Stokes problem, 15-20 iterations required per Newton step

	cores	*Linear*	*Picard*
Mesh 1	*512*	3.57	3.88
Mesh 2	*4096*	8.48	5.90
Mesh 3	*32786*	7.86	7.12

weak scaling

cores	*Linear*	*Picard*
2048	2.04	2.22
16384	4.56	4.71

strong scaling

- Fixed three level hierarchy obviously prohibits possibility of perfect weak scaling, however 8 times larger problems requires < 2 times more CPU time.
- Strong scaling is in the range of 20-43% efficiency. Strongly influenced by massively parallel, small sub-domain coarse grid solver.
- **Performing non-linear, high resolution 3D thermo-mechanically coupled visco-plastic (non-trivial) simulations is still challenging.**

Conclusions and Outlook

1. *Matrix Free (MF) SpMV kernels for Q2 elements reduce memory foot print, avoid the memory bandwidth bottleneck and thus scale well on parallel multi-core architectures*

2. *MF operators combined with Chebyshev/Jacobi can result in robust and efficient parallel MG smoothers for VV Stokes*

3. *Hybrid coarsening strategies can yield significant speed gains for "hard" problems*

4. *The distributed coarse grid solver is a scalability bottleneck. Further experimentation using (i) semi-redundant solves, (ii) AMG coarse grid solvers and (iii) Krylov methods with non-blocking global reductions needs to be conducted*

5. *Usage of hybrid MPI-OpenMP parallelism needs to be investigated*

Thanks for your attention... questions?