Algorithm Design for Large Scale Parallel FFT-Based Simulations on CPU-GPU Platforms

Anuva Kulkarni, Franz Franchetti, Jelena Kovačević

Carnegie Mellon University

Electrical & Computer ENGINEERING

Motivation

Extreme memory requirements and high communication overhead prevent scaling of large scale iterative simulations involving parallel FFTs to higher grid sizes, which is necessary for high resolution analysis.

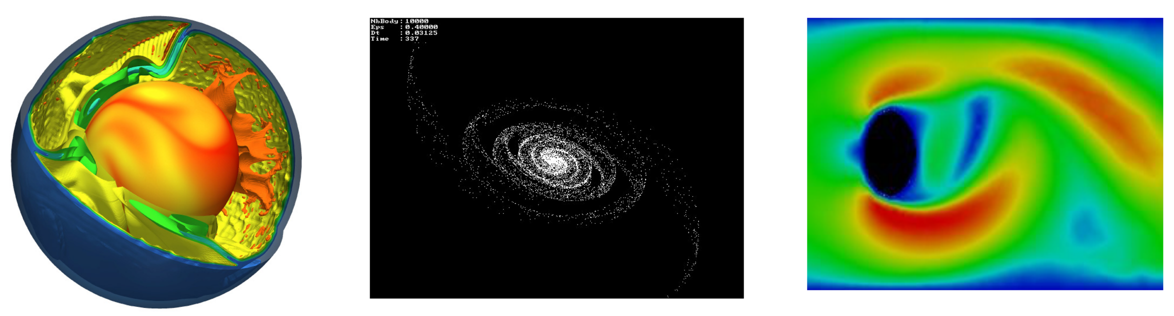

An example is that of Moulinec Suquet's Basic Scheme to compute local stress and strain fields in materials , a partial differential equation simulation that uses FFTs

Stress in crystals

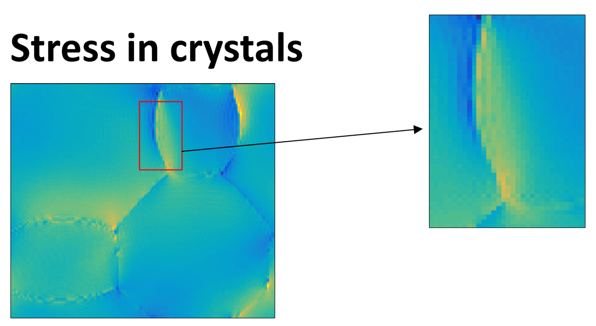

Boundaries are critical regions of interest

3D Hooke's law:

stress strain

$$\sigma_{ij} = C_{ijkl} : \epsilon_{kl}$$

3x3x3x3 stiffness tensor

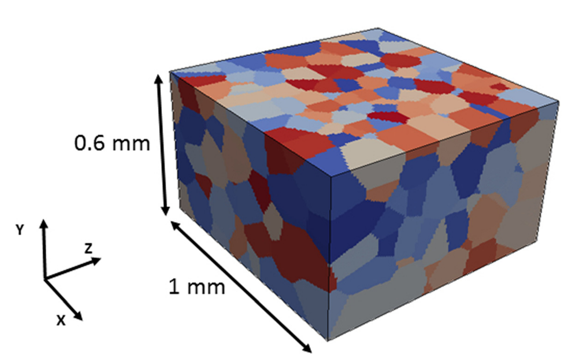

Elliptical PDE:

$$C^0_{ijkl} u_{k,lj}(\mathbf{x}) + \tau_{ij,j}(\mathbf{x}) = 0$$

MSC Basic Scheme is solved by convolution with Green's function using FFT.

Increasing grid resolution is desirable. However, larger problem sizes must be run with parallelized code. This requires large parallel FFT computations which means **high memory usage and all-all communication.**

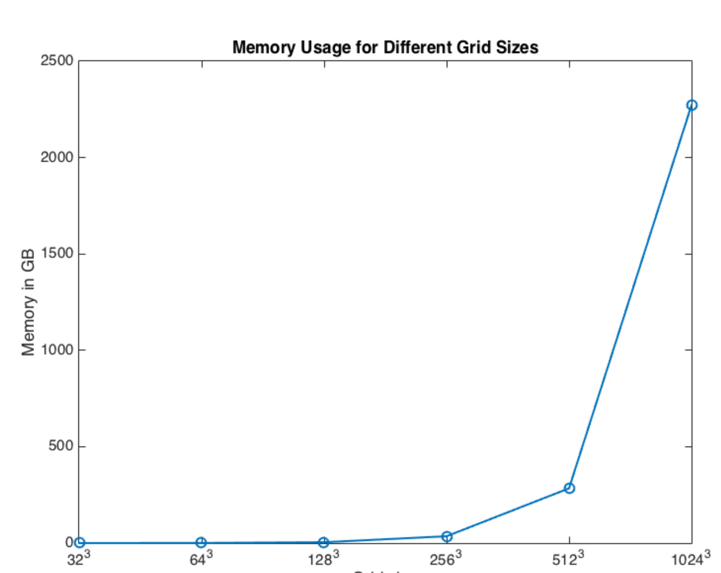

Memory requirement increases **32.4k** times!

Goal: **Overcome these limitations and run stress-strain simulations for larger problem sizes**

Our solution: An algorithm and software co-design for heterogeneous platforms using irregular domain decomposition and local FFTs.

Background

Algorithm 1 MSC Basic Scheme

1: **Initialize:**
$\epsilon^0 \leftarrow E$,
$\sigma^0_{mn}(\mathbf{x}) \leftarrow C_{mnkl}(\mathbf{x}) : \epsilon^0_{kl}(\mathbf{x})$
2: **while** $e_s > e_{tol}$ **do**
3: $\hat{\sigma}^i_{mn}(\xi) \leftarrow \mathrm{FFT}(\sigma^i_{mn}(\mathbf{x}))$
4: Check convergence
5: $\Delta\hat{\epsilon}^{i+1}_{kl}(\xi) \leftarrow \hat{\Gamma}_{klmn}(\xi) : \hat{\sigma}^i_{mn}(\xi)$
6: Update strain: $\hat{\epsilon}^{i+1}_{kl}(\xi) \leftarrow \hat{\epsilon}^i_{kl}(\xi) - \Delta\hat{\epsilon}^{i+1}_{kl}(\xi)$
7: $\epsilon^{i+1}_{kl}(\mathbf{x}) \leftarrow \mathrm{IFFT}(\hat{\epsilon}^{i+1}_{kl}(\xi))$
8: Update stress: $\sigma^{i+1}_{mn}(\mathbf{x}) \leftarrow C_{mnkl}(\mathbf{x}) : \epsilon^{i+1}_{kl}(\mathbf{x})$

Original method by Moulinec and Suquet

Stress and strain fields in grain interiors are smooth. Hence we can treat these as separate domains.

99% energy of the space-domain Green's function is concentrated at central peak. Hence, Green's function can be truncated before convolution.

Right:
N = 512
Slice of 3D component of space-domain Green's function

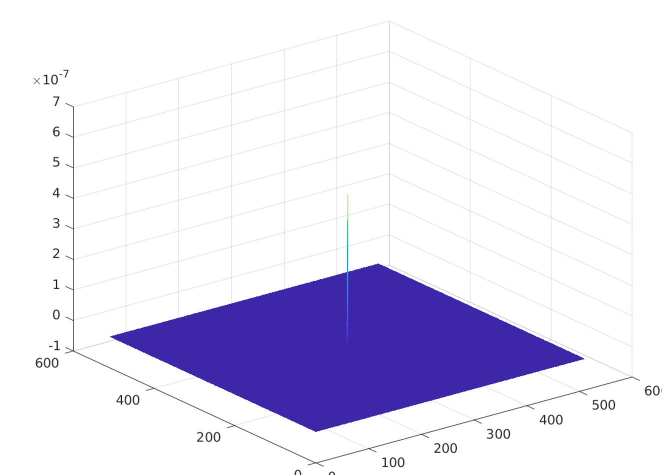

Proposed Method

The proposed MSC Alternate Scheme is a co-design of algorithm and software for heterogeneous platforms. It enables scaling of stress-strain simulations to large grids by overcoming high memory requirements and communication bottlenecks.

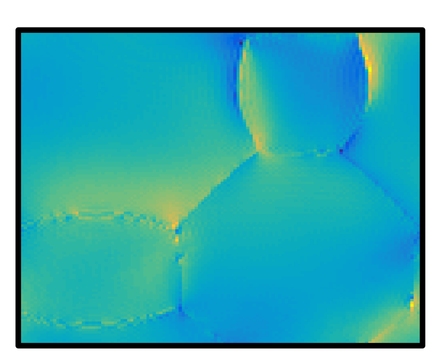

Smooth stress fields in grain interiors can be compactly represented with data models. This reduces communication overhead.
Grain boundaries are convolved at full resolution without approximations.

Ideal model characteristics

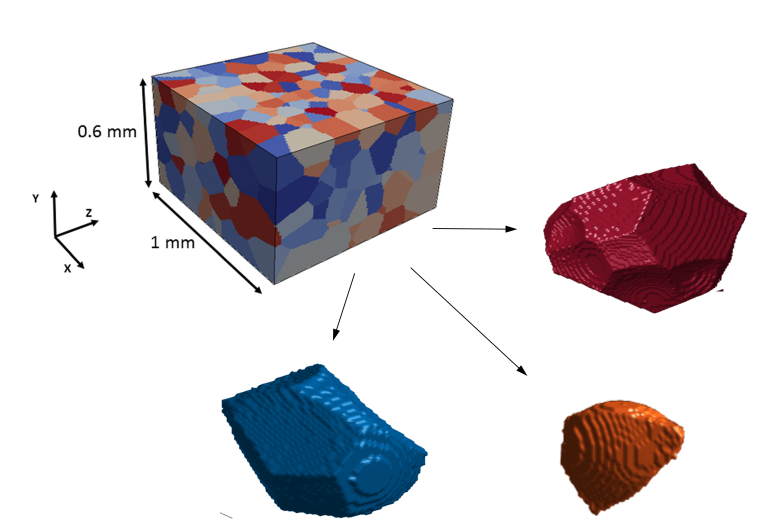

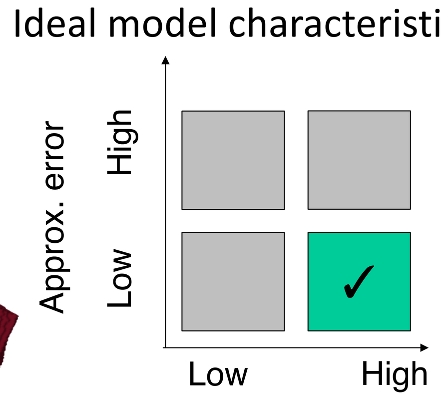

Each grain (domain) is assigned to a GPU. For small grains, single GPU can process multiple grains. Distribution will be done using appropriate load balance. Local FFTs are performed on the GPU side.

Phase I: Successes

Computational aspects for prototype development: MATLAB-FORTRAN workflow

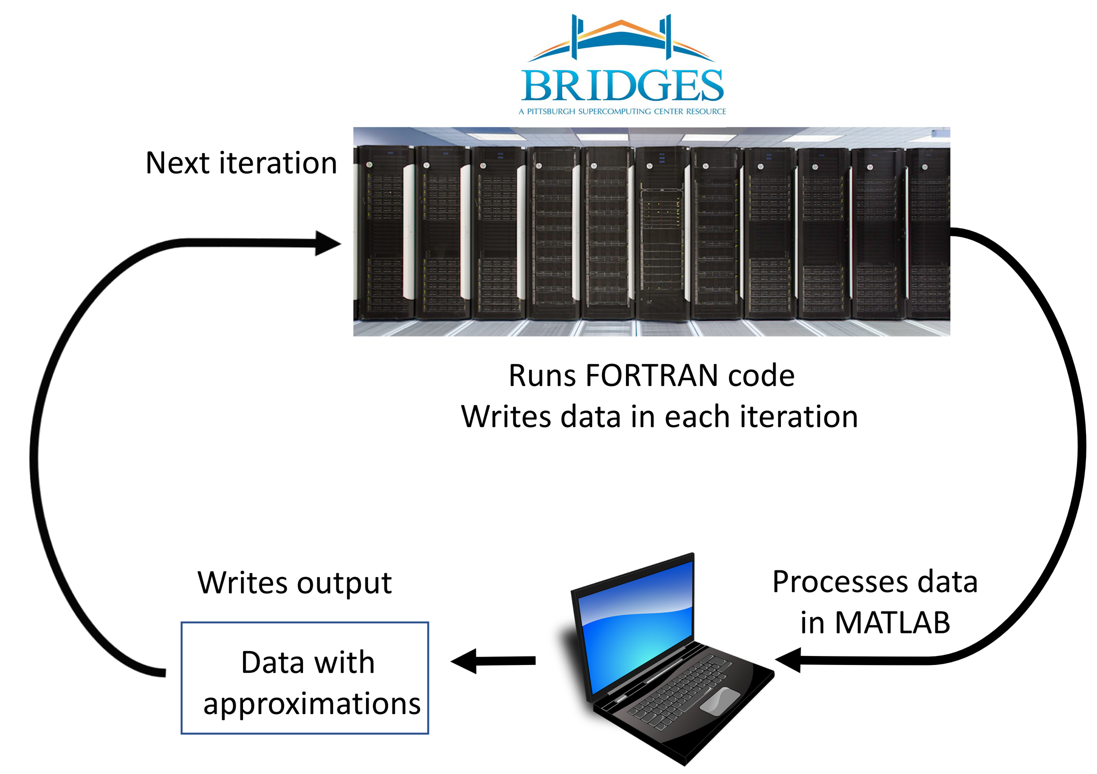

Smooth regions for data modeling:

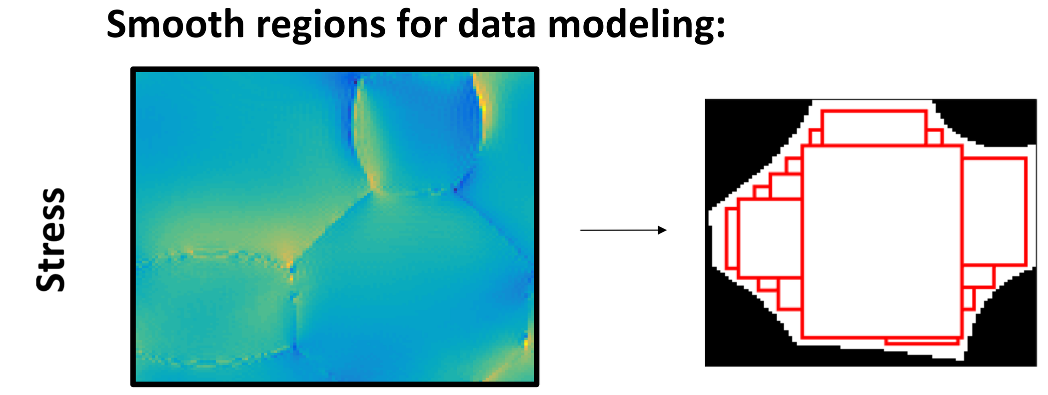

Convergence of our method:

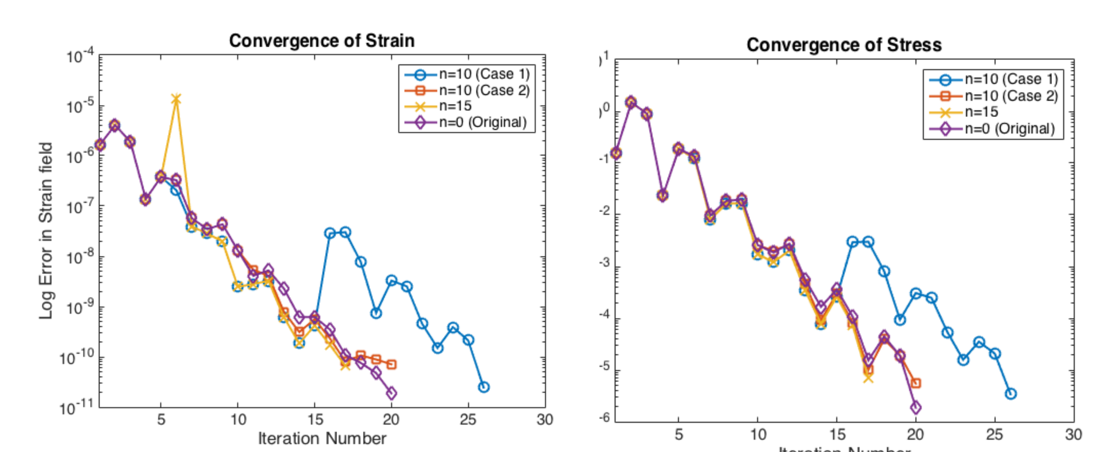

Mismatch in stress field in each iteration compared to MSC Basic Scheme:

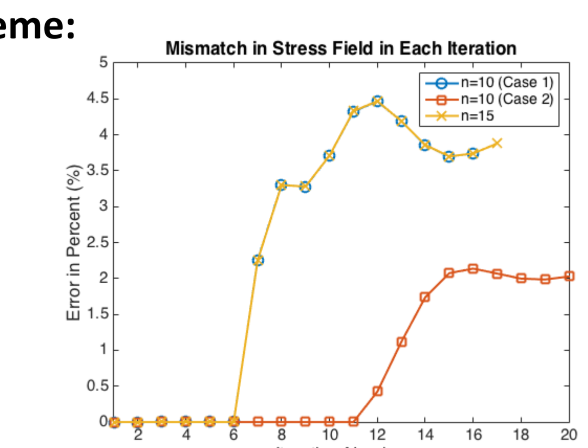

Stress field slice in MSC Basic Scheme:

Stress field slice in MSC Alternate Scheme:

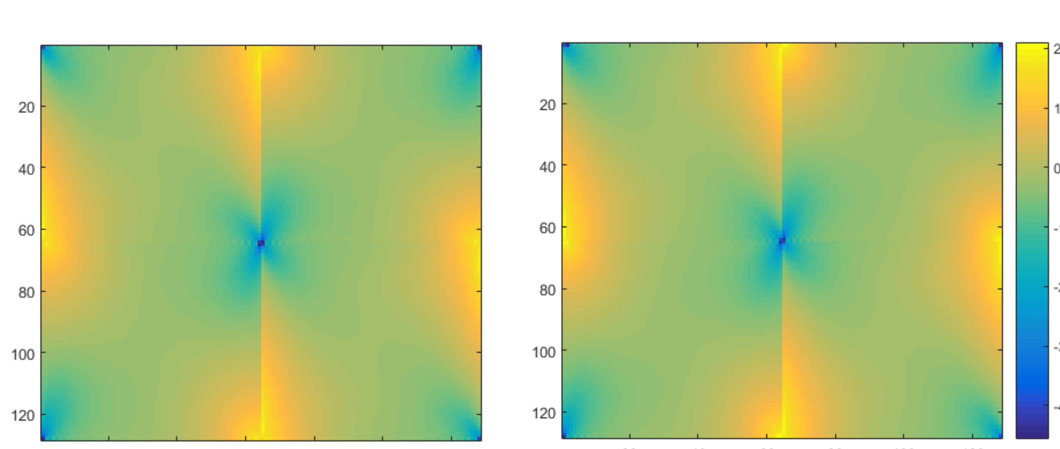

Phase II: Plans

Thrust 1: Algorithm improvement

- De-noise convolution output
- Reduce model error
- Extend to different datasets with irregular grains

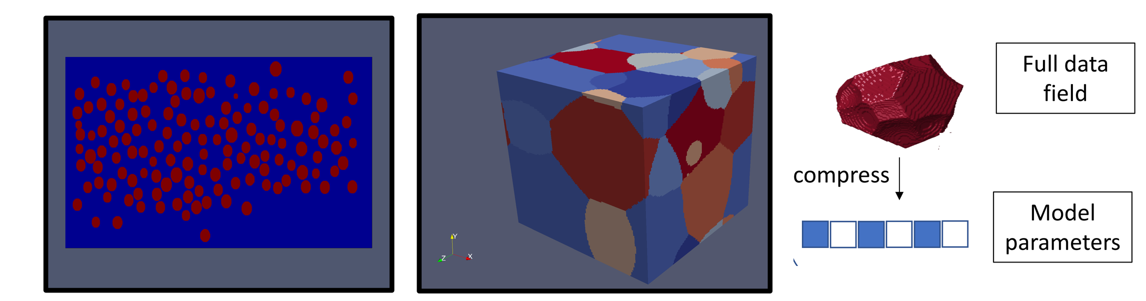

Thrust 2: Deployment on heterogeneous system with Tesla K80

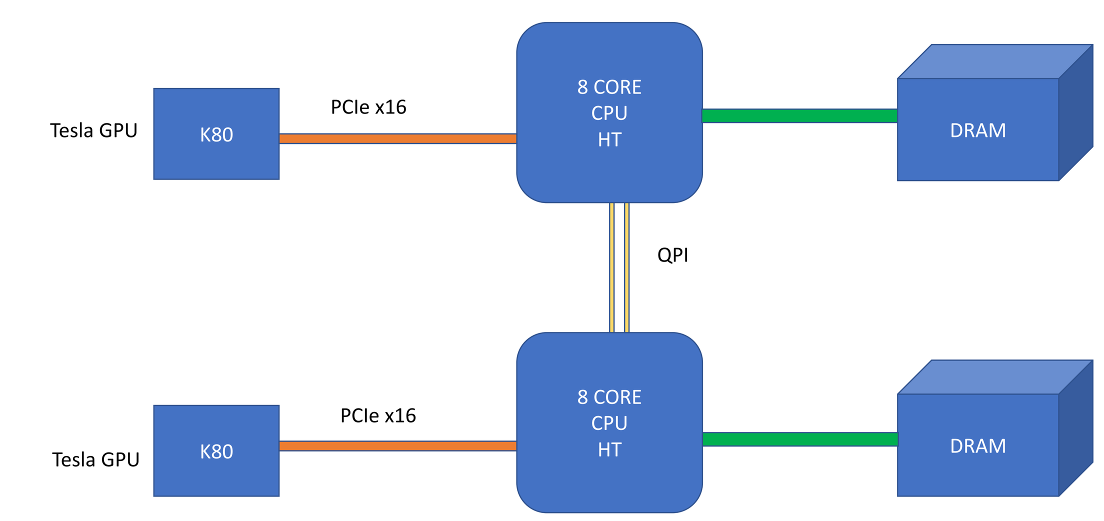

Thrust 3: Support and adapt to different accelerators and machines

Thrust 4: Extend to visco-plastic code [2] which includes deformation of crystals and studies cracking and fracture formation

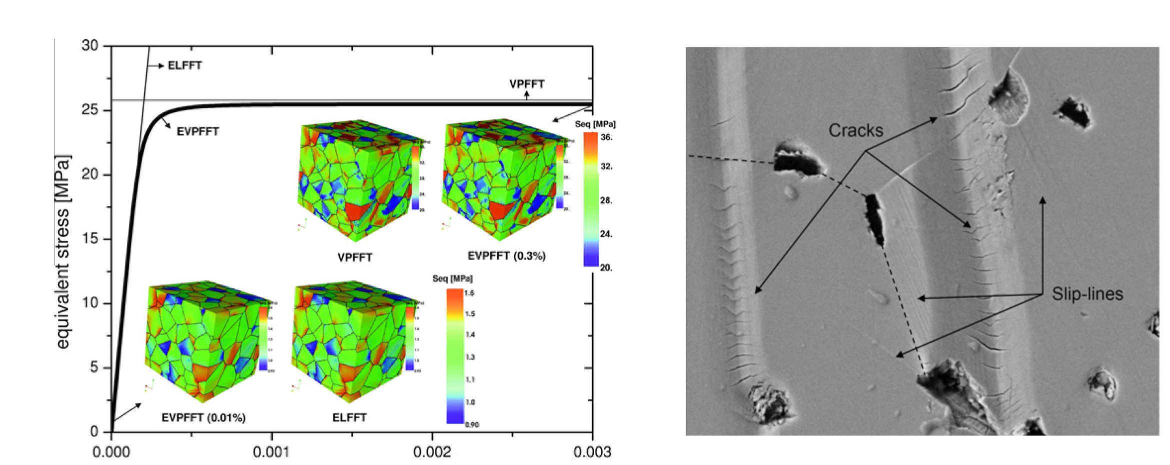

Acknowledgements

The authors would like to thank Dr. Anthony Rollett, Dr. Vahid Tari at CMU and Dr. Anirban Jana and Dr. Roberto Gomez at Pittsburgh Supercomputing Center for all their assistance and collaboration with this project.

References

[1] H. Moulinec and P. Suquet. 1998. A numerical method for computing the over- all response of nonlinear composites with complex microstructure. Computer methods in applied mechanics and engineering 157, 1-2 (1998), 69–94.
[2] R. A. Lebensohn. 2001. N-site modeling of a 3D viscoplastic polycrystal using fast Fourier transform. Acta Materialia 49, 14 (2001), 2723–2737.

Images:
http://www.ethlife.ethz.ch/archive_articles/120223_Simulation_subduktion_su/1202XX_einseitigeSubduktion_L1-hires.jpeg
http://public.lanl.gov/lebenso/2012-IJP-RLetal-EVPFFT.pdf

https://www.researchgate.net/profile/Tamas_Csanadi/publication/311421026/figure/fig2/AS:489562741972997@1493732261970/ Effect-of-load-and-crystal-orientation-on-crack-formation-during-scratch-test-b.ppm