Progress Report:

Implementation and Validation of Multi-surface Plasticity in a Discontinuous Mesh Finite Difference Code

(SCEC Project 17162)

June 6, 2019

Investigators:

Daniel Roten[1] (PI)
Kim Olsen[1] (co-PI)
Yifeng Cui[2] (co-PI)

Institutional Affiliations:

1. San Diego State University
   Dept. Geological Sciences
   5500 Campanile Dr, MC 1020
   San Diego, CA 92182-1020

2. University of California San Diego
   San Diego Supercomputer Center
   9500 Gilman Dr, MC 0505
   La Jolla, CA 92093-0505

Abstract

SCEC project 18168 aims to implement an Iwan-type plasticity model in the discontinuous mesh version of the AWP-GPU finite difference (FD) code. A further scope of the project is to investigate the impact of Iwan-type plasticity on realistic earthquake ground motion scenarios, in particular a M7.8 (ShakeOut) earthquake scenario on the southern San Andreas fault. As of this writing, the implementation of multi-surface plasticity in AWP-GPU has not yet been fully completed, and verification of the method against the CPU code of AWP and other 1D or 2D wave propagation codes is still pending. This delay was mainly caused by migration from the Titan supercomputer to the newer system Summit (both at OLCF) as well as the decommission of Blue Waters (NCSA). Porting of the code to the newer Volta GPUs revealed several preexisting bugs that needed to be addressed before development could continue. In this preliminary report, we summarize the current state of code porting and development. We als present new results of ShakeOut scenario simulations using the Iwan model performed with the CPU version of AWP, and the sensitivity of ground motions to the choice of rock strength parameters.

Introduction

Strong ground motions during earthquakes are the result of wave propagation phenomena which occur over a wide range of scale lengths, from hundreds of kms for source directivity effects to a few tens of meters for local nonlinear site response. Although these phenomena can be predicted using simulations of dynamic rupture and wave propagation, the high computational costs frequently limit such simulations to the lower band (< 4 Hz) of the spectrum relevant for buildings (0–10+ Hz). On the other hand, high-frequency calculations which simulate the nonlinear response of soft soils usually only include the earth structure underlying a site of interest. In addition, simplifying assumptions are often made, such as a vertically incident wavefield or a horizontally layered soil structure, with effects of the source, path and site treated as decoupled processes. This traditional site response formalism neglects the important interplay between multi-dimensional and potentially nonlinear phenomena, such as long period-surface waves, finite source effects or fault zone plasticity.

In order to generate synthetic ground motion time histories and physics-based seismic hazard maps for the entire frequency range of engineering interest, a paradigm shift away from site response formalism has been proposed for SCEC5. This paradigm shift requires computer codes which are able to model ground motions as the response of a dynamical system, and include physics which represents inelastic effects both in the fault damage zone and in the shallow crust. Such codes must be able to resolve the short wavelengths in low-velocity sediments while simulating a domain that is large enough to include a finite source, which requires high computational efficiency as well as scalability.

One example of such a scalable code is the finite difference (FD) program AWP, which is being used in many SCEC projects, including CyberShake simulations (Graves *et al.*, 2011), southern California tomography (Lee *et al.*, 2014) and high-frequency ground motion simulation (Cui *et al.*, 2013; Withers *et al.*, 2015; Roten *et al.*, 2016). Two versions of the code are currently being maintained: a CPU version (Cui *et al.*, 2010) which supports kinematic and dynamic sources, and a more efficient and scalable GPU version which supports only kinematic sources. In particular, the GPU version includes support for a discontinuous mesh for shorter time to solution (Nie *et al.*, 2017; Roten *et al.*, 2018c,a). Both versions support plasticity using a Drucker-Prager (DP) yield condition (Roten *et al.*, 2016). Support for a multi-surface Iwan-type model has been implemented and verified in the CPU version of AWP within SCEC project 17162. The scope of SCEC project 18168 is to implement the Iwan model in the GPU version of AWP, and to include the capability to simulate nonlinearity (both Drucker-Prager and Iwan-type) inside the transition zone between discontinuous mesh resolutions.

Migration of AWP-GPU-DM to OLCF Summit

In the past, AWP-GPU-DM was tested and deployed on supercomputers equipped with NVIDIA GPUs, including Titan at OLCF[1] and Blue Waters at NCSA[2]. For example, the 4 Hz simulations of the Shakeout scenario (Roten *et al.*, 2016) where performed on both Titan and Blue Waters. Our recent simulations of a M9 megathrust earthquake in the Cascadia subduction zone, performed using a discontinuous mesh, were deployed exclusively on Titan (Roten *et al.*, 2019). However, both Blue Waters and Titan were recently decommissioned. Titan, which went out of production at the end of 2018, was replaced by the newer Summit system which became available in early 2019.

Although OLCF Summit is somehow similar to Titan (i.e., both systems use NVIDIA GPUs to accelerate computations), differences in the deployed hardware, libraries and compilers led to a broad range of issues which initially prevented the group from running simulations on Summit. Because we plan to deploy AWP-GPU-DM on Summit, and we were no longer able to run production simulations on Titan, it was essential to address these issues before proceeding with the implementation and testing of the Iwan model.

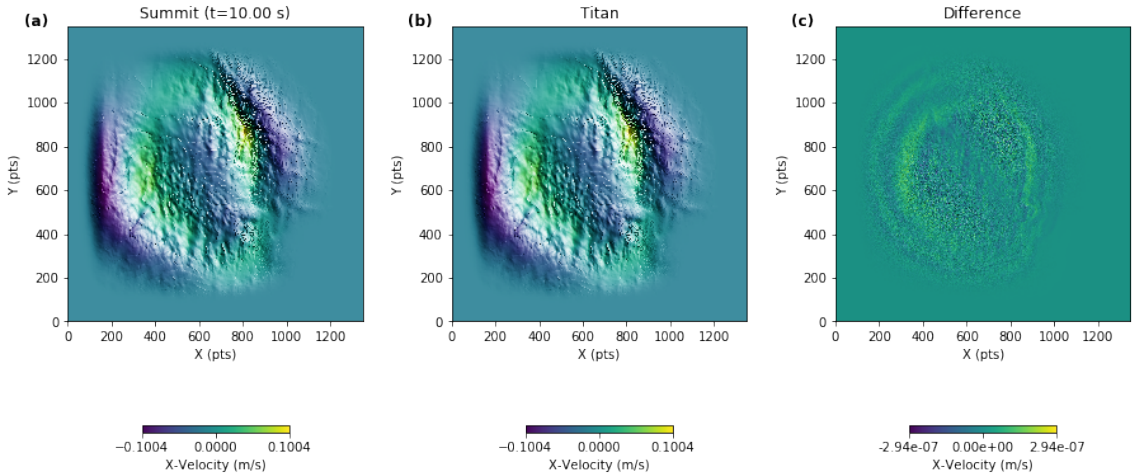

Figure 1: Surface snapshots of X-velocity obtained from simulation with AWP-GPU-DM using three mesh sizes on (a) OLCF Summit and (b) OLCF Titan. (c) Difference in surface velocity between the solutions obtained on the two systems. Horizontal axes show grid points (surface $\Delta h = 66\ ^2/_3$ m).

Porting issues that were identified and which have now been addressed were related to the type of MPI implementation, the GPU hardware and the version of NVIDIA's CUDA library. In particular, we found that the CUDA kernel which interpolates the wavefield from the coarse to the fine mesh in discontinuous mesh simulations consistently triggered memory access errors on the Volta GPUs used on Summit. These issues had not been encountered in the Kepler GPUs used on the older Titan system. The problem was identified early using the CUDA memcheck tool, but the initial solution produced incorrect results, and a correctly working version of the interpolation kernel has only been developed recently. Inconsistencies between simulation results obtained on Summit and Titan were also traced back to uninitialized memory on the GPU (i.e., memory which was allocated but not set to zero) at several locations in the code. We note that, although these incorrect initializations have been present in the application for a long time, they only started causing problems after migration to the Volta GPUs on Summit. A separate issue which caused inconsistencies between solutions computed on Summit and Titan was traced back to implementation differences between Spectrum's MPI

[1] Oakridge Leadership Computational Facility
[2] National Center for Supercomputing Applications

library used on Summit, and the MPICH library used on Titan (Ossian O'Reilly, personal communication, 2019).

All these technical concerns related to porting AWP-GPU-DM from Titan to Summit have now been addressed in the latest version of AWP-GPU-DM. Figure 1 compares snapshots of the surface wavefield obtained from simulations on Summit and Titan. Both systems predict consistent velocities, and differences are 6 orders of magnitude smaller than the amplitude of the propagating wavefield.

Implementation of Iwan model in AWP-GPU-DM

In contrast to the CPU version of AWP, which represents material parameters, stresses and velocities inside the medium using 3D Fortran arrays, the GPU version relies on flattened arrays to store these variables in GPU memory. AWP-CPU-Iwan extends the 3D arrays to 4D, where the 4th dimension represents the number of the yield surface in the Lamé parameters and stress tensors (with the overlay stress field and total shear and bulk moduli stored at index 0). The same strategy can not be applied on the flattened array used by AWP-GPU-DM.

For the implementation of the Iwan model, we settled on a memory layout where the stress (or material propery) associated with each yield surface increases fastest. For example, the shear stress τ_{xx} associated with each yield surface at grid position (i, j, k) would be stored as follows:

$\mathtt{xx}^1_{i,j,k}$	$\mathtt{xx}^2_{i,j,k}$	$\mathtt{xx}^3_{i,j,k}$	$\mathtt{xx}^4_{i,j,k}$	$\mathtt{xx}^5_{i,j,k}$	...	$\mathtt{xx}^{N-1}_{i,j,k}$	$\mathtt{xx}^N_{i,j,k}$	$\mathtt{xx}^1_{i,j,k+1}$	...

Here, superscripts indicate the yield surface number, and N the total number of yield surfaces. Because the stress for the different yield surfaces is updated using an explicit for loop inside the stress and plasticity kernels, this layout reduces the memory access time.

We are developing new versions of the stress and plasticity kernels which compute the elastic and plastic stress updates individually for each yield surface (Fig. 2). These kernels will store the stresses inside additional arrays which are allocated only if the Iwan model is selected by the user. An additional overlay kernel is computing the sum of the stress field and stores the result inside the existing stress arrays employed in traditional elastic or elastoplastic calculations. This configuration allows us to leave the velocity and absorbing boundary kernels unchanged in computations involving the Iwan model.

Collaboration and coordination with USC/SCEC

Efforts to port AWP-GPU-DM from Titan to Summit and the development of the Iwan model were coordinated with researchers at USC/SCEC, in particular with Ossian O'Reilly, who is currently working on implementation of topography in AWP. The PIs participated in the weekly SCEC topography call to discuss issues encountered while using the code, and to coordinate the implementation of the features (Iwan model, DM) and porting to Summit with other users. Ossian O'Reilly visited SDSU on two occasions to discuss how to move forward. It was agreed to compile a series of test cases which will be used to make sure that new features do not change the expected behavior of the code.

Simulation of the ShakeOut scenario using the Iwan model

An important application of AWP with the Iwan model is the simulation of large scenario earthquakes, which dominate the seismic hazard, including realistic nonlinear attenuation in soft near-surface sediments. Previously (Roten *et al.*, 2018b, SCEC report 17162) we have performed one realization of the ShakeOut scenario, which stipulates a M7.8 earthquake rupture on the southern San Andreas fault (SAF) from southeast to northwest, using AWP-CPU with the Iwan model. Meanwhile, we have explored the sensitivity of ground motions in the Los Angeles basin to the strength of sedimentary deposits in more detail.

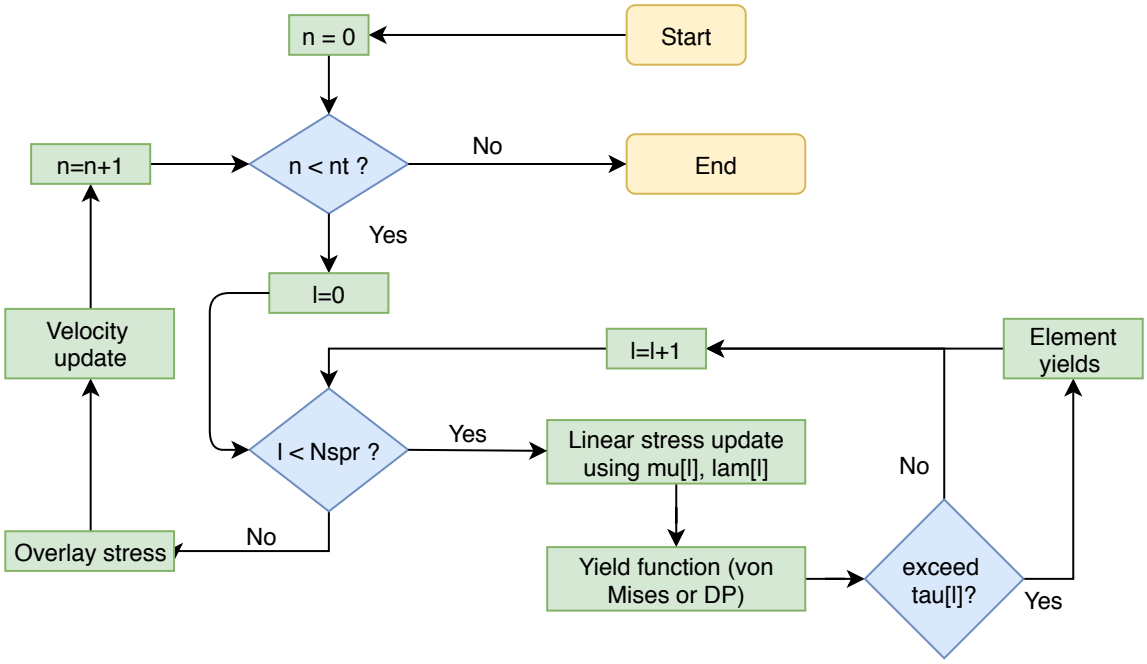

Figure 2: Flowchart indicating implementation of Iwan model in AWP. The counters `n` and `t` indicate the yield surface number and time step, respectively. `Nspr` indicates the number of yield surfaces and `nt` the number of time steps. The variables `mu`, `lam` and `tau` indicate the Lamé parameters μ and λ and the yield stress τ associated with the current yield surface.

Definition of Reference Strain

The strength of the sediments is controlled by the reference strain, γ_r, which is defined as the strain where the shear modulus reduction curve reaches half the low-strain value. In our previous simulations, we used the EPRI93 shear modulus reduction curves for Sand (Electric Power Research Institute, 1993) to define depth-dependent reference strains inside the sedimentary fill of basins, which were identified as regions where $v_s < 1{,}950$ m/s. Using the EPRI curves, the reference strain γ_r assigned inside basins ranges from 0.1% near the surface to 0.5% at depths of 300 m or more (Fig. 3).

Here, we use of an empirical relationship (Darendeli, 2001) to assign the reference strain as a function of depth. In the Darendeli (2001) relationship, the reference strain is assigned as a function of plasticity index PI, overconsolidation ratio OCR and mean effective confining pressure σ_0':

$$\gamma_r = \left(\phi_1 + \phi_2 \cdot \mathrm{PI} \cdot \mathrm{OCR}^{\phi_3}\right) \sigma_0'^{\phi_4}. \tag{1}$$

We used the tabulated mean values for the four soil groups analyzed by Darendeli (2001) to define the parameters ϕ_1 to ϕ_4. We assumed a value of 0 for the plasticity index (i.e., no clay content) and a value of 4 for the overconsolidation ratio, which corresponds to isotropic consolidation (i.e., coefficient of earth at rest K_0=1). It is noted that the OCR in the Darendeli relationship has no effect on γ_r if PI = 0. If we neglect pore water pressure effects for the computation of the effective vertical stress, the reference strain ranges from $\gamma_r = 0.8 \cdot 10^{-4}$ at the surface to $\gamma_r \approx 3 \cdot 10^{-3}$ at 4 km depth (Fig. 3). The Darendeli (2001) relationship predicts a generally slower increase in γ_r with depth compared to the EPRI curves. We also consider simulations using upper and lower bound models of the reference strain, which were obtained by adding or subtracting a standard deviation on the reference strain. At the reference strain, where $G/G_{\max} = 0.5$, the

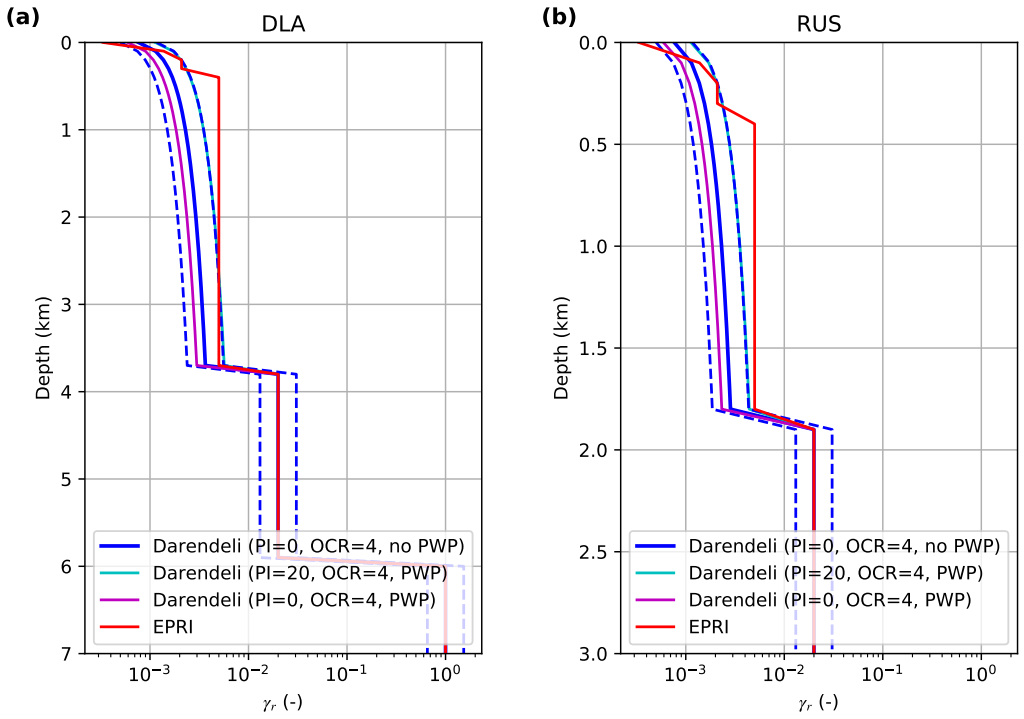

Figure 3: Reference strain γ_r as a function of depth at strong-motion sites (a) DLA and (b) RUS obtained with the Darendeli (2001) relationship using different assumptions for the plasticity index PI and with or without pore water pressure (PWP) effects, and using the EPRI (Electric Power Research Institute, 1993) shear modulus reduction curves. Dashed blue lines show reference strain obtained by adding or subtracting one standard deviation from the reference strain shown in the solid blue line. OCR = Overconsolidation ratio.

standard deviation as defined in the shear modulus reduction curves by Darendeli (2001) simplifies to

$$\sigma = \exp(\phi_{13}) + \sqrt{\frac{0.25}{\exp(\phi_{14})}} = 0.09638. \quad (2)$$

Subtracting one standard deviation from the reference strain obtained without pore water pressure effects (PI=0, OCR=4) results in a reference strain that is slightly lower than the reference strain obtained with effective confining pressure (including pore water effects, Fig. 3). Adding one standard deviation to the reference strain obtained with PI = 0 results in virtually the same reference strain obtained using PI = 20 with pore water pressure effects. These results show that reference strains including plus or minus one standard deviation represent viable choices for the sediments of the San Bernardino and San Gabriel basins. Outside basins, a value of γ_r = 2% for rock was assigned (Schnabel *et al.*, 1972). At depths of 6 km or more, we set a reference strain of 100%, effectively prohibiting nonlinearity (Fig. 3a).

Sensitivity of Ground Motions to Reference Strain in Sediments

As shown in many previous studies (e.g. Olsen *et al.*, 2008, 2009; Roten *et al.*, 2016), dynamic simulations of the ShakeOut scenario predict strong long-period ground motions inside the waveguide connecting the SAF

to the Los Angeles basin (LAB). Here, spectral accelerations at a period of 3 seconds (3s-SAs) reach values of more than 1g in the linear case (Fig. 4a). If nonlinearity based on the Iwan model is taken into account, using the reference strain predicted by the Darendeli relationship (PI = 0, no pore water pressure effects), 3s-SAs are limited to values of 0.5g or less inside the waveguide (Fig. 4b).

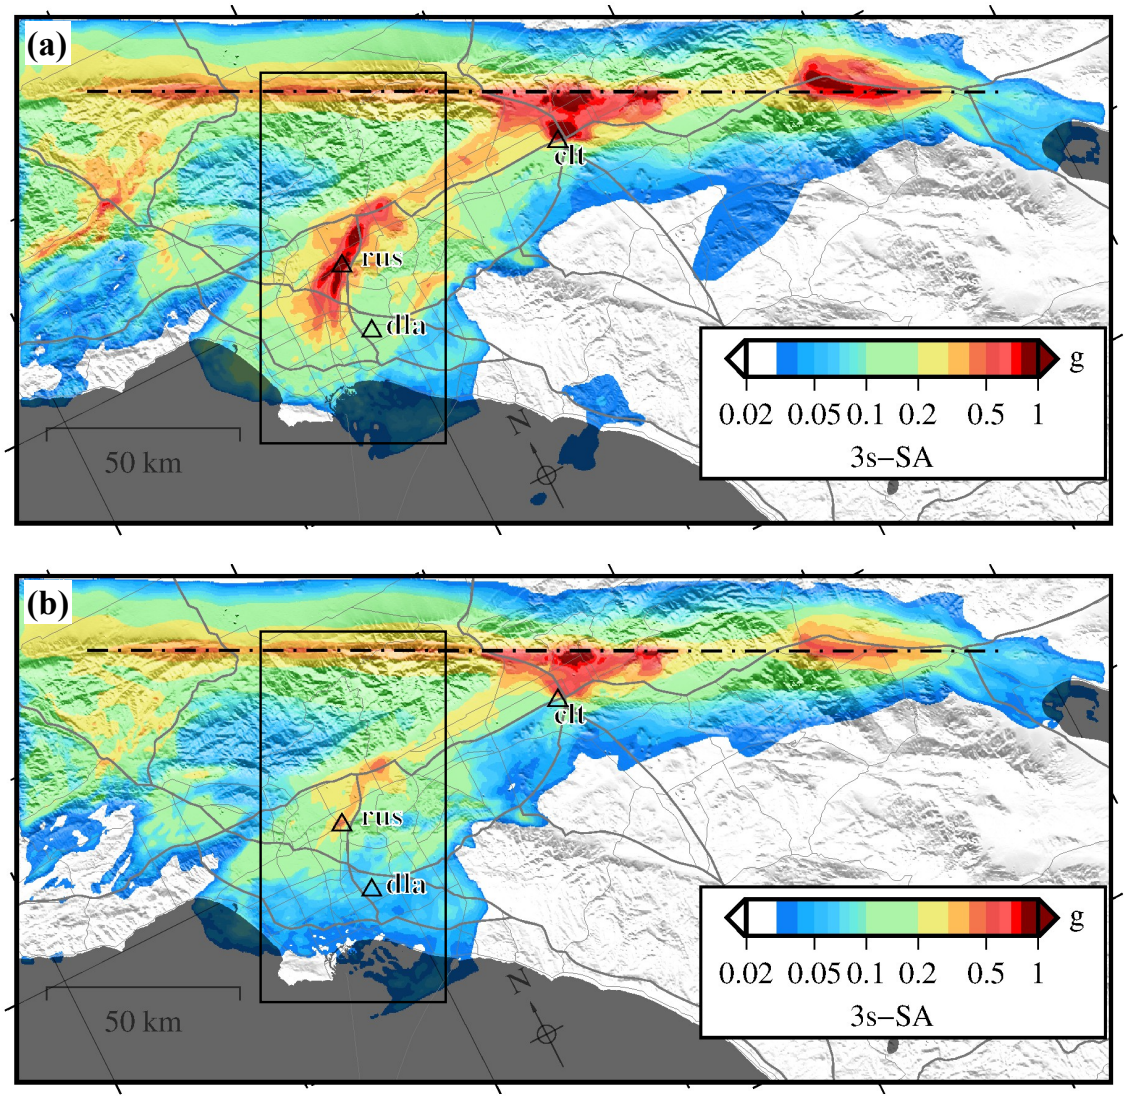

Figure 4: Spectral accelerations at 3 seconds (3s-SAs) obtained from dynamic simulation of the ShakeOut scenario for (a) the linear case and (b) the nonlinear case using the Iwan model and reference strains defined by the Darendeli (2001) relationship (PI=0, OCR=4, no PWP; Fig. 3). The black rectangle shows the extent of the map region shown in Figure 5.

Figure 5 compares ground motions obtained from different simulation assumptions inside a rectangular region surrounding the largest patch of waveguide amplification inside the Los Angeles and San Gabriel basins. Adding one standard deviation to the reference strain given by the Darendeli relationship increases 3s-SAs to values above 0.4g inside large parts of the waveguide. Long-period ground motions obtained using

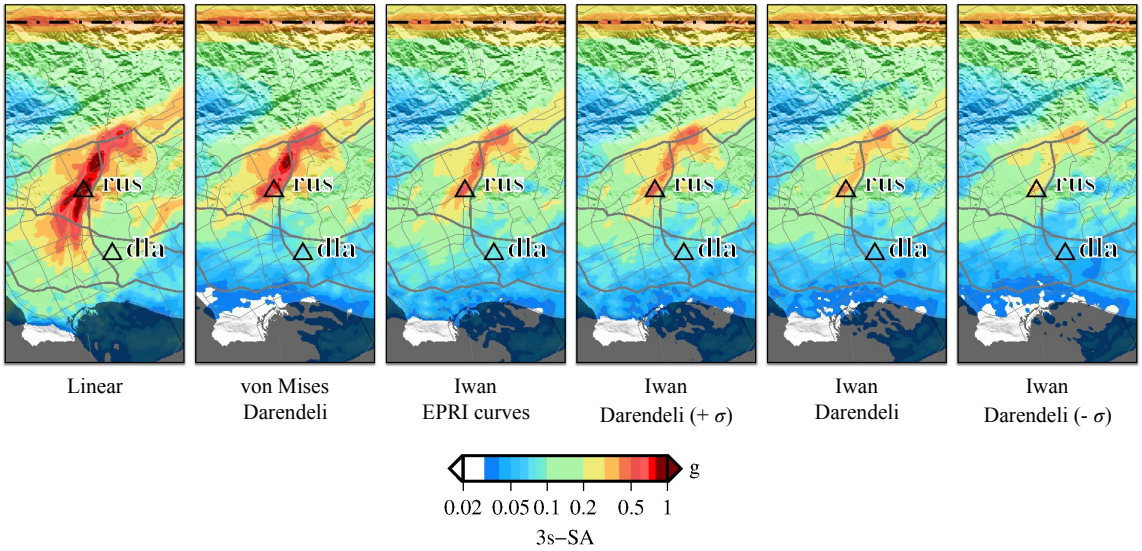

Figure 5: 3s-SAs inside rectangular area include main waveguide amplification patch (Fig. 4) obtained in the linear case, the nonlinear case using a single von Mises yield surface, and a multi-surface Iwan model using different definitions of the reference strain.

the EPRI curves are slightly larger than those from the Darendeli relationship plus one standard deviation, but remain mostly below $0.5g$. If we subtract one standard deviation from the reference strain obtained from Darendeli's relationship, the waveguide almost disappears from the ground motion maps, with 3s-SAs mostly below $0.3g$.

To assess how the use of an oversimplified nonlinear model affects ground motions, we also carried out one nonlinear simulation using a single von Mises yield surface (simulations using the Iwan model were carried out using 10 yield surfaces following a von Mises yield criterion). The yield stress τ_{max} of the single yield surface was defined as the stress at the reference strain in the hyperbolic model, which corresponds to $\tau_{\text{max}} = 0.5\ \gamma_r\ G$, with the reference strain computed from Darendeli's relationship (PI = 0, no PWP). 3s-SAs obtained using a single von Mises yield surface are lower than those obtained in the linear case, but still higher than the values obtained using the EPRI curves.

We already reported previously (Roten *et al.*, 2018b) that using a bilinear yield criterion does not result in underprediction of ground motion compared to using a multi-surface yield model. However, the analysis shown here (Fig. 5) allows a one-by-one comparison of the von Mises and Iwan models, as we used the same reference strain to calibrate the yield stress of the yield surface(s) in both cases. These results confirms that while the von Mises yield criterion does not reproduce plasticity as accurately as a multi-surface yield model, it still represents a viable first-order approximation that should be preferred over a fully elastic model.

Simulated seismograms extracted for the strong motion sites **clt** (Colton), rus (Whittier-Narrows) and downtown Los Angeles (Fig. 6 illustrate how peak ground velocities decrease with decreasing values for the reference strain in the sediments. Effects of plasticity are especially pronounced close to the SAF, at site clt, and inside the main waveguide at site rus. Howver, Iwan nonlinearity does not only affect ground motion amplitudes, but also the arrival times strong motion phases. For example, the largest peak in direction N180°E reaches LA about ∼0.5 later in the Iwan model compared to the linear solution (Fig. 6). This shift is a direct consequence of the shear modulus reduction inside the low-velocity sediments of the basins during strong shaking, which is not captured using a single yield surface.

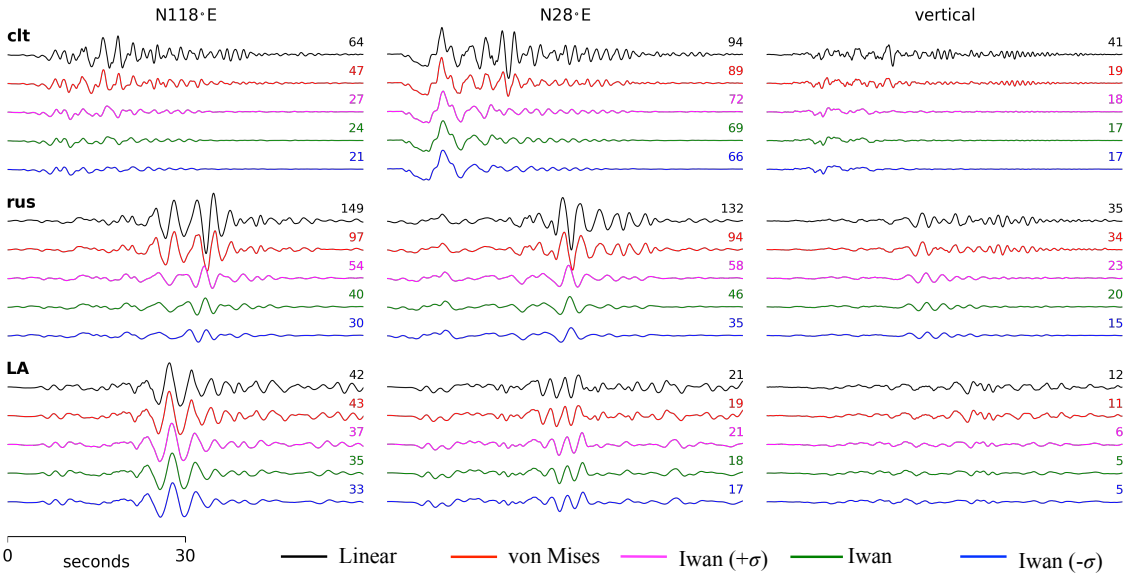

Figure 6: Synthetic seismograms computed at the strong motion sites clt, rus and in downtown Los Angels (LA) from dynamic simulation of M7.8 ShakeOut scenario using a linear model, a single von Mises yield surface and the Iwan model using different assumptions for the definition of the reference strain in sediments. Numbers above traces indicate peak velocity in cm/s.

Summary and Outlook

AWP-GPU-DM has successfully been ported to Summit, and we are currently completing the implementation of Iwan nonlinearity in the code. The development and optimization of the code will continue during SCEC project 19128 "Development and Optimization of Iwan Model in Discontinuous Mesh Finite Difference Code AWP-ODC-GPU". The implementation will be verified against the CPU version of AWP-Iwan and against established 1D and 2D nonlinear wave propagation codes, as done previously for the development of the Iwan model in the CPU code. In the framework of SCEC project 19056, "Verification and Validation of 3D Nonlinear Physics-based Ground Motion Simulations: Phase I" (PI D. Assimaki), the code will be verified against other 3D nonlinear wave propagation codes (such as Hercules, SEM3d and SPEED), and validated against observations made at the Garner Valley downhole arrays. These verification and validation exercises will initially be performed for 1D cases, but extended to 3D benchmarks in a later phase.

References

Cui, Y., Olsen, K.B., Lee, K., Zhou, J., Small, P., Roten, D., Ely, G., Panda, D.K., Chourasia, A, Levesque, J., Day, S.M., and Maechling, P. 2010. Scalable Earthquake Simulation on Petascale Supercomputers. *In: Proceedings of SC10, November 13-19, New Orleans, LA*.

Cui, Y., Poyraz, E., Olsen, K.B., Zhou, J., Withers, K., Callaghan, S., Larkin, J., Guest, C., Choi, D., Chourasia, A., *et al.* 2013. Physics-based seismic hazard analysis on petascale heterogeneous supercomputers. *Page 70 of: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. ACM.

Darendeli, Mehmet B. 2001. *Development of a new family of normalized modulus reduction and material damping curves*. Ph.D. thesis, The University of Texas at Austin.

Electric Power Research Institute. 1993. Guidelines for determining design basis ground motions. *Electric Power Research Institute Technical Report EPRI TR-102293*.

Graves, R., Jordan, T. H., Callaghan, S., Deelman, E., Field, E., Juve, G., Kesselman, C., Maechling, P., Mehta, G., Milner, K., Okaya, D., Small, P., and Vahi, K. 2011. CyberShake: A physics-based seismic hazard model for southern California. *Pure Appl. Geophys.*, **168**(3-4), 367–381.

Lee, En-Jui, Chen, Po, Jordan, Thomas H, Maechling, Phillip B, Denolle, Marine AM, and Beroza, Gregory C. 2014. Full-3-D tomography for crustal structure in southern California based on the scattering-integral and the adjoint-wavefield methods. *Journal of Geophysical Research: Solid Earth*, **119**(8), 6421–6451.

Nie, Shiying, Wang, Yongfei, Olsen, Kim B, and Day, Steven M. 2017. Fourth-Order Staggered-Grid Finite-Difference Seismic Wavefield Estimation Using a Discontinuous Mesh Interface (WEDMI). *Bull. Seism. Soc. Am.*, **107**(5), 2183–2193.

Olsen, K. B., Day, S. M., Minster, Y. A., Cui, Y., Chourasia, A. J., Okaya, D., and Maechling, P. 2008. Terashake2; Spontaneous rupture simulations of Mw 7.7 earthquakes on the southern San Andreas Fault. *Bull. seism. Soc. Am.*, **98**(3), 1162–1185.

Olsen, K. B., Day, S. M., Dalguer, L. A., Mayhew, J., Cui, Y., Zhu, J., Cruz-Atienza, V., Roten, D., Maechling, P., Jordan, T., Okaya, D., and Chourasia, A. 2009. ShakeOut-D: Ground Motion Estimates Using an Ensemble of Large Earthquakes on the Southern San Andreas Fault With Spontaneous Rupture Propagation. *Geophys. Res. Lett.*, **36**, L04303.

Roten, D., Cui, Y., Olsen, K.B., Day, S.M., Withers, K., Savran, W., Peng, W., and Mu, D. 2016. High-frequency Nonlinear Earthquake Simulations on Petascale Heterogeneous Supercomputers. *In: 2016 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC'16), November 13-18 2016, Salt Lake City, UT*. in press.

Roten, D., Olsen, K.B., Nie, S., and Day, S.M. 2018a. High-frequency Nonlinear Earthquake Simulations on Discontinuous Finite Difference Grid. *In: Proceedings of the 11th National Conference in Earthquake Engineering*. Earthquake Engineering Research Institute, Los Angeles, CA. 2018.

Roten, D., Olsen, K.B., and Cui, Y. 2018b. *A multi-surface plasticity model for 3D wave propagation simulation using AWP*. Final Report, SCEC Project 17162. Southern California Earthquake Center.

Roten, D., Olsen, K.B., and Takedatsu, R. 2018c. Numerical Simulation of M9 Megathrust Earthquakes in the Cascadia Subduction Zone. *In: Workshop on Best Practices in Physics-based Fault Rupture Models for Seismic Hazard Assessment of Nuclear Installations: issues and challenges towards full Seismic Risk Analysis, Cadarache-Château, France, May 14–16, 2018*.

Roten, D., Olsen, K.B., and Takedatsu, R. 2019. Numerical Simulation of M9 Megathrust Earthquakes in the Cascadia Subduction Zone. *Pure and Applied Geophysics*, **in press.**

Schnabel, P., Seed, H. Bolton, and Lysmer, J. 1972. Modification of seismograph records for effects of local soil conditions. *Bull. Seism. Soc. Am.*, **62**(6), 1649–1664.

Withers, K. B., Olsen, K. B., and Day, S. M. 2015. Memory-Efficient Simulation of Frequency-Dependent Q. *Bull. Seism. Soc. Am.*