

Response to Reviewer 2

Journal: Geoscientific Model Development

Manuscript Number: EGUSPHERE-2026-2922

Title: DynEarthSol v2.0: An Efficient Explicit Lagrangian Solver for Geodynamics, Surface Processes, and Earthquake-Cycle Dynamics

Authors: Chase J. Shyu, Sungho Lee, Xuesong Ding, Jaeyoon Keum, Eh Tan, Luc L. Lavier, and Eunseo Choi

Reviewer 2: Ylona van Dinther

Overall evaluation

This paper presents the second version of DynEarthSol, which becomes a 2D and 3D geodynamic modelling framework designed to solve the dynamics of geodynamic, surface process and earthquake cycle processes. Three technical breakthroughs justifying a second version of this model and this paper are (i) bitwise reproducibility, (2) coupling to surface processes, and (3) including state evolution of RSF in the existing rate-dependent formulation. The paper is mostly well-written, although it is often unclear what is exactly physically and mathematically implemented, such that several major points involve questions about that. Hence I think that the guidelines of GMD for reproducibility based on the text are not met. Nonetheless, remarkable technical has been done to get such a complex model with unstructured meshes and remeshing working fast across CPU-GPU architectures. I think this is important technical work that deserves publication in Geoscientific Method Development (GMD). However, two out of three technical challenges do not meet the standards I expect from GMD. Bitwise reproducibility is claimed as one of three novelties, but this is not demonstrated in the paper. RSF is claimed as “validated”, but the verification is limited to one local closed-form check of equations that does not verify many key processes of earthquake cycle dynamics. Since your RSF implementation is not fully physically sound in terms of an underlying assumptions on stress being equal to strength not being met and a physical process limiting rupture speeds that seems to be missing, it is important to either fix that and/or have a robust reproduction or closer evaluation of the resulting earthquake loading and release and large scale earthquake pattern. Since such robust technical proof is either absent or too limited, I recommend major revisions. If the editor thinks that some of the main points below on technical correctness are of high importance to GMD, more time might be needed for the technical implementation, such that reject and encouragement for resubmission would better allow for that time.

Response to the overall evaluation

We thank the reviewer for identifying the need to clarify the implementation and strengthen the evidence for reproducibility and earthquake-cycle behavior. The principal revisions are as follows.

(1) RSF formulation and recalculation. DES now uses an invariant strain-rate measure and both weakening and healing time-step limits for state evolution. All RSF benchmark, application,

and sensitivity calculations in this study were recomputed with this formulation, including all 74 parameter-study runs.

(2) Application-scale comparison. We added a heterogeneous-fault comparison with Herrendörfer et al. (2018), evaluating recurrence, slip partitioning, and rupture extent.

(3) Verification and sensitivity tests. The expanded evaluation includes local constitutive verification, zero-velocity healing, mesh resolution, and fictitious-density sensitivity. A 2000 yr RSF control produced bitwise-identical numerical states across the tested CPU configurations (Comment 5). Appendix D documents finite-speed disturbance propagation in the damped, mass-scaled system. Extensional tests show fault growth near the Arthur reference angle in EP and both EP–RSF cases, with stable sliding for velocity strengthening ($a > b$) and earthquake-like stick–slip for velocity weakening ($a < b$) (Section 5.3.2).

(4) Documentation and context. Section 5.2 now states the DES threshold-plasticity formulation and update sequence explicitly. Section 5.1 separates interface-based and continuum RSF implementations, expands the continuum literature, and clarifies the relation to Tong and Lavier (2018) and Lavier et al. (2021). Across the manuscript, rheology names were standardized, the earlier 3D/OpenCL implementation was acknowledged through Ta et al. (2015), and unsupported claims were removed. The responses below describe these changes point by point.

Major points

Implementation of RSF using plasticity

R2–Comment 1

Adding the state evolution to RSF and doing earthquake cycling is one of three main novelties. However, exactly how this is implemented is not well described and seems not in line with fundamental assumptions in RSF.

We agree that the original manuscript did not describe the constitutive update clearly enough. DES incorporates Dieterich–Ruina rate and state dependence into the friction coefficient used in the Mohr–Coulomb yield criterion (Section 5.2).

At each step, DES calculates $V^n = 2w\dot{\epsilon}_{II}^n$ from the total deviatoric strain rate, advances θ^n with the aging law, evaluates $\mu(V^n, \theta^{n+1})$ and the corresponding friction angle, and then performs the yield test and return mapping. Stress may therefore remain below strength while an element is elastic and equals strength only during yielding. The time step limits the fractional change in θ , and all reported RSF calculations were recomputed with this limit and the invariant rate measure. Sections 5.2.2 and 5.4.3, and our response to Comment 2, explain the consequences of the threshold behavior.

Changes in manuscript.

Section 5.2 was rewritten to give the friction and state laws, their integration into the Mohr–Coulomb return map, the invariant continuum rate measure, the update order, and the state-evolution time-step limit. Appendix G gives the application settings.

R2–Comment 2

RSF is derived under the assumption that fault slip occurs always - even at extremely low and negligible rates down to the -20, -30 or even -40 m/s -, such that stress and strength are always equal (e.g., Nakatani, JGR, 2001; Herrendorfer et al., JGR, 2018; Siroattanakul et al., Nature, 2025). In your plasticity formulation, stress and strength are per definition not equal, as you use strength as a threshold. In the interseismic period between earthquakes stress is smaller than (interface) strength, while we now know that ample of slip occurs below seismic slip rates as well, such as during slow slip events (hence the seismic coupling at faults is often less than half). Hence this leads to the question if, and how, your threshold plasticity formulation meets the fundamental assumption of RSF stating stress = strength, and how it then captures slip transients during the inter-seismic period. Note that this is the reason that Herrendorfer and related papers implement RSF as a highly non-linear viscous formulation, where slip can always occur - and can be very important interseismically - and stress and strength are always equal. In that paper interface strength defined by Nakatani, JGR, 2001 is used to locally approximate plasticity's threshold at which slip rates become noticeable.

DES uses the Dieterich–Ruina rate-and-state terms to set the friction coefficient in a Mohr–Coulomb yield strength. The fault zone can remain partly elastic during interseismic phases: below yield, elements deform elastically while θ continues to evolve. Plastic deformation resumes locally when stress reaches the yield strength, allowing elastic loading between slip transients. The viscoelasto-plastic model of Herrendörfer et al. (2018) also includes elasticity, but permits plastic flow at very low rates while maintaining stress equal to strength.

In the strike-slip comparison (Fig. R1), 64% of fault displacement at $X = 69$ km accumulates during interseismic phases and 36% during rapid episodes. This interseismic contribution is concentrated in short slip transients, as quantified in Appendix G1.

Changes in manuscript.

Section 5.2.2 specifies the stress–strength relation in the DES threshold formulation. Section 5.4.3 reports the interseismic slip and explains how DES and Herrendörfer et al. (2018) differ in their treatment of plastic deformation during interseismic phases.




Figure R1 (revised manuscript Figure 15). Comparison with the strike-slip configuration of Herrendörfer et al. (2018). (a) Fault-relative slip rate derived from paired stations on the two band faces, following their Figure 4b; $V_{\text{max}} = \max_i |v_{x,i}^+ - v_{x,i}^-|$ is the maximum fault-relative slip rate across the station pairs. (b) Cumulative-slip profiles along the band in the three rate classes of their Figure 4c; the rapid class is indexed by slip increment to compare the spatial accumulation of slip.

R2–Comment 3

Formulations in the abstract, conclusions and many other locations in the main text you write that “this approach enables the simulation of complex fault behaviors, from stable creep to stick-slip instabilities, “ (1.8-9), “allows stable creep, stick-slip instability, and dynamic rupture-like behavior to be represented within the same continuum framework” *(. 39-40) and the “full qualitative spectrum from .. stable creep to self-sustaining earthquake cycling and creep”. To me and most readers such texts and the full spectrum

implies that also slow slip events are simulated. However, also in light of the above, I wonder whether those are simulated correctly? Also since you never use the words to describe or list them. Could you show one panel in Figures 12-13 with non-decaying (similar magnitude) slow slip events? You could also plot stress and strength throughout a few earthquake cycles or slip transients.

- Also what means dynamic rupture-like? Please reformulate in relation to comments below

Four of the 74 calculations have intermediate-velocity responses. Two representative intermediate cases reach peak nodal speeds of 6.1×10^{-6} and $4.5 \times 10^{-7} \text{ m s}^{-1}$, respectively (Figs. R2 and R3). Their longest continuous intervals above ten times the loading rate last 67 and 89 days, respectively. Stress reaches the yield strength during the transients and falls below it during elastic unloading. These slow-slip-like transients recur with varying amplitudes.

Speed and displacement describe monitoring-node motion; fault-relative slip is calculated from paired nodes on the two band faces (Fig. R1). The earlier “full spectrum” and “dynamic rupture-like” claims have been removed.

Changes in manuscript.

Abstract, Section 5.1, Section 5.4, and Conclusions: the full-spectrum and “dynamic rupture-like” claims were removed. Section 5.4 and Figs. 13 and 14 describe the response types and their stress–strength histories, distinguishing nodal quantities from fault-relative slip. Appendix G gives the parameter-map classification.




Figure R2 (revised manuscript Figure 13). Responses at fixed $D_c = 2.5 \times 10^{-3}$ m, with a/b increasing from left to right. The third column contains intermediate-velocity transients reaching $6.1 \times 10^{-6} \text{ m s}^{-1}$. The fifth row compares the Mohr-circle radius R with the threshold R_y at which Eq. (5.5) gives $f_s = 0$; the two meet while the element yields and separate on unloading.




Figure R3 (revised manuscript Figure 14). Responses at fixed $a/b = 0.6$, with D_c increasing from left to right; panels and quantities follow Figure R2. Peak speed decreases across the four cases, and the second column has the intermediate-velocity response.

R2–Comment 4

You explain and justify your slip rate calculation in 5.3.2. You argue that principal stresses axes remain close to fault parallel and fault normal throughout the cycle. This might be a allowable approximation when the fault is defined as a planar structure, but this does not hold during fault growth. Preuss et al., JGR, 2019, shows the evolution of fault stress and friction during fault growth with an invariant RSF formulation, which are significant and can explain fault orientation mismatches in nature.

- To evaluate your slip rate calculation you could evaluate whether you can reproduce the Arthur angle as well.
- Note that Preuss 2019 uses the invariant formulation of van Dinther 2013 and Herrendorfer 2018 ($V_{\text{slip}} = \text{plastic strain rate} * 2 * dx$, where dx is grid size or fault width), which does not suffer from the same limitation in approximation V_{slip} as you suggest in l. 578.
- Ideally, check your implementation against the Arthur angle, but at least discuss your limitation more adequately and be aware when removing the definition of fault plane latter on.
- This computation is also computationally efficient and allows for the mesh-objectivity under plasticity, as discussed latter on. Hence one could consider adopting this one.
- Alternatively, this implementation could for a predefined planar fault be evaluated in a large-scale benchmark, but this will be discussed in the next point.

All earthquake-cycle calculations use the invariant, direction-independent rate measure (Section 5.2.3)

$$V = 2w \dot{\epsilon}_{II}, \quad \dot{\epsilon}_{II} = \sqrt{\frac{1}{2} \dot{\epsilon}'_{ij} \dot{\epsilon}'_{ij}},$$

with w the element minimum altitude ($w = 2A/l_{\text{max}}$ in two dimensions, $w = 3\Omega/A_{\text{max}}$ in three). This definition removes the previous stress-derived shear direction and is unchanged by rigid-body motion.

The extensional tests compare EP with velocity-strengthening ($a/b = 1.2$) and velocity-weakening ($a/b = 0.2$) EP-RSF using the invariant rate measure, with $\phi_0 = 30^\circ$ and $\psi = 0^\circ$ (Section 5.3.2, Fig. 11). At 12 kyr, shear bands extend across multiple elements and dip at 52.4° , 54.5° , and 51.2° , respectively, measured over depths 2–8 km. All three dips are close to the Arthur reference angle of 52.5° , evaluated at the initial friction angle because RSF friction evolves. The velocity-weakening ($a < b$) case shows stepwise displacement and repeated stress drops; the EP and velocity-strengthening ($a > b$) cases accumulate displacement more smoothly. Thus, adding RSF preserves fault growth near the Arthur reference angle under both stable sliding and earthquake-like stick-slip.

The invariant formulations cited by the reviewer scale plastic or viscoplastic strain rates. In DES, a plastic-only rate vanishes below yield, leaving the velocity floor to determine the direct-effect term during elastic loading. Using the total deviatoric strain rate lets V follow both elastic loading and plastic yielding. This V describes element-level deformation; fault-relative slip is evaluated separately.

The invariant definition removes orientation dependence but does not by itself make the local plasticity mesh objective: w is an element length, and V consequently scales with element size. Section 5.4.2 and Appendix G quantify this sensitivity.




Figure R4 (revised manuscript Figure 11). Spontaneous localization in EP and EP-RSF models with $\phi_0 = 30^\circ$ and $\psi = 0^\circ$. (a) Geometry, boundary conditions, mesh detail, and seed monitored in (e,f). (b–d) Accumulated plastic strain at 12 kyr minus its initial value, on a common linear scale, for EP, velocity-strengthening EP-RSF ($a/b = 1.2$), and velocity-weakening EP-RSF ($a/b = 0.2$). Dashed lines and angles show fits over depths 2–8 km. The Arthur, Coulomb, and Roscoe values use the initial angles. (e) Magnitude of seed-node displacement from its initial position. (f) Second invariant of deviatoric stress, τ_{II} , in the monitored seed element. Histories show 5–10 kyr; the legend in (e) applies to both panels.

Changes in manuscript.

Section 5.2.3 defines the invariant rate measure in Eq. (5.6), distinguishes it from the plastic-rate measures of van Dinter et al. (2013), Herrendörfer et al. (2018), and Preuss et al. (2019), and states its interpretation and mesh dependence. All reported earthquake-cycle calculations were recomputed with this measure; Section 5.4.2 and Appendix G quantify its resolution sensitivity.

Section 5.3.2 presents the spontaneous-localization comparison in Fig. 11; Appendix F gives the model configuration.

Noise and reproducibility

R2–Comment 5

One of the main technical points in the abstract, and one of three technical breakthroughs presented, is the bit-wise reproducibility. However, this is not demonstrated in the paper. Since your figures 12 and 13 suggest the presence of numerical noise I doubt that at least the earthquake cycle simulations are truly bitwise reproducible. Please repeat identical simulations across multiple runs and relevant parallel execution configurations and compare the complete numerical state using e.g., binary comparison. If reproducibility is not demonstrated or only partially for repeated runs using one fixed hardware and execution configuration, please narrow the claim accordingly.

- Could you also clarify which operation was previously nondeterministic and how the two-stage update removes this nondeterminism?

The previously nondeterministic operation was the element-to-node accumulation during force assembly. Direct parallel accumulation allows several elements to update the same nodal force, so both write conflicts and the order of floating-point addition can vary. In DES v1.x the elements were partitioned into $2T$ contiguous index bands, where T is the OpenMP thread count, and the even- and odd-numbered bands were traversed in two successive parallel loops so that concurrently active bands would not share nodes. Because the band boundaries are a function of T , the accumulation order at a given node changed with the thread count and differed again under serial execution. The required node-disjointness was moreover only verified by a node-index-range test that issued a warning rather than an error, so at coarse resolution or high thread count the assembly could still incur a genuine race, making results not merely irreproducible but potentially incorrect.

The two-stage update first writes every element–node contribution to a separate array location. A node-based loop then gathers the contributions in ascending element-number order, fixing the summation sequence.

We added a system-level reproducibility test using an RSF earthquake-cycle control spanning 2000 yr. The same calculation was run on three hosts using AMD EPYC 9654 and 9474F processors and 24 OpenMP threads. The complete monitor records and two checkpoint states were byte-identical. For all 18 saved states, every numerical payload was also byte-identical; the raw save containers differed only in the recorded wall-clock time. Thus the earthquake-cycle calculation was bitwise reproducible across the tested CPU hosts and processor models.

That test held the thread count fixed. We have since verified thread-count invariance separately, on the 3.9×10^4 -element benchmark configuration, where repeated CPU runs at differing OpenMP thread counts are bitwise identical to one another and repeated GPU runs are likewise bitwise identical to one another. CPU and GPU results are not bitwise identical to each other, which is expected: the two-stage update fixes the order of summation and only that, whereas the two paths are compiled separately and may contract multiply–add pairs or hold intermediate expressions at differing precision, altering the rounding of individual terms. We have therefore narrowed the manuscript claim to the tested CPU configurations and to invariance under repetition and thread

count, and we make no claim of bitwise identity across compilers, contraction settings, or accelerator platforms.

These tests establish deterministic repetition, but do not establish that every repeated fluctuation is physical; the origin and resolution sensitivity of the fluctuations are addressed in the next response.

Changes in manuscript.

Abstract, Introduction, and Section 3.1: the general claims of bitwise reproducibility across hardware and accelerator platforms were narrowed to the tested CPU configurations. Section 3.1 now identifies the previously nondeterministic element-to-node operation and describes the DES v1.x thread-count-dependent band colouring that was its specific origin. It notes that band disjointness was checked by warning rather than enforced and explains how ascending element traversal of the support list fixes the gather order. The section reports both the complete-state comparison for the 2000 yr RSF control and the separate thread-count comparison on the 3.9×10^4 -element benchmark. It also explains why CPU and GPU results are not bitwise identical to each other.

R2–Comment 6

Your figures 12 and 13 with earthquake cycle results show considerable noise, whose amplitudes can vary from 0.5-3 orders of magnitude. A large part of this may well be numerical in nature.

- A part of this in the left and right “cree” panels may be physical as well, hence please show the left and right panels for much short times/displacements, such that about 10 cycles are seen.
- Other parts in the central panel are likely numerical still, as oscillations are prominent in the inteseismic and at the end of coseismic as well. Such oscillations are not present in e.g., the comparable work of Herrendorfer et al., JGR, 2018, which - like many boundary element earthquake cycle simulations (Rice et al. JGR 1993) - show very clean sequences with fully periodic recurrence of characteristic earthquakes, after the first or first few (spin-up) earthquakes.
- Additionally, the small variability in recurrence rate and slip you note, also suggests that there might be something that is not fully appropriately solved.
- Often reducing the grid size can help to achieve a smooth, noiseless evolution. Does that, or anything else, help you to reduce the noise in your simulations?
- Also, is this noise bitwise reproducible?

A 2000 yr earthquake-cycle control reproduced its checkpoint and saved numerical states bitwise across the tested CPU configurations, ruling out run-to-run parallel nondeterminism for that test.

Figs. R2 and R3 now pair each full time history with detailed views of selected intervals, including speed, displacement, friction, state and stress relative to strength, allowing the small oscillations to be examined directly.

The fine, reference and coarse meshes place about nine, four to five and two elements across the fault zone. All three produce rapid episodes. The first-episode time changes by less than 1%, and mean recurrence changes by -9.6 to $+2.7\%$. Recurrence intervals remain variable after refinement, with coefficients of variation of 0.06 on the finest mesh and 0.02 on the other two. No remeshing occurs in these calculations. In the density-floor test (Comment 9), removing the residual floor reduces the coefficient of variation of recurrence intervals from 0.16 to 0.05 and changes event-scale

properties. Together, these tests support the response class and first-episode timing while showing that finer-scale variability remains sensitive to numerical settings.

Changes in manuscript.

Figs. 13 and 14 pair full time histories with detailed views of selected intervals. Appendix G reports the three-resolution comparison over the simulated intervals, including mean recurrence and its variability, and states that no remeshing occurred.

R2–Comment 7

Reproducibility of plasticity is also affected by mesh-objectivity, which is only discussed in one sentence in relation to another paper with a different implementation and a focus only at rate-dependent plasticity (Goudarzi et al., 2023). However, mesh-objectivity for your long-term geodynamic simulation coupled to surface processes would not be helped by that rate-dependence, yet the fault locations and intensity and maybe even surface expressions could be affected. In various geodynamic codes this is dealt with by adding Kelvin viscoelasticity to regularise the fault width (Duretz et al., GGG, 2019). ASPECT includes a similar type of regularisation, such that this difference could explain a distinct part of the difference between your and ASEPECT’s simulation, which is distinctly smoother and has less faults.

- It is not clear whether your model is reproducible for different grid sizes. Could you clarify what spatial or temporal regularization controls shear-zone width, and demonstrate mesh convergence or discuss the limitations for your surface processes and earthquake-cycle simulations?
- Since you write here and there, e.g., in L. 551, that your “shear zones are explicitly resolved” I would expect to see that demonstrated following the previous question, or such formulations implying “resolvedness” need to be updated.

On the shear zone width in DES

Our plasticity implementation does not include an explicit regularization of the kind formulated by Duretz et al. (2019). The width of the shear bands is therefore controlled by the element size. We note that the regularization of Duretz et al. (2019) introduces a length scale through a Kelvin viscoplastic element. We are pursuing a comparable regularization and have expanded the plan for future improvement in the Conclusions section.

On the mesh convergence in DES

We have tested the resolution dependence of the DES model with three series of simulations: three-dimensional models with adaptive refinement at $h_{\min} = 1.0, 0.5$ and 0.25 km; three-dimensional models with uniform resolution at $h = 2.0, 1.0$ and 0.5 km; and two-dimensional models with uniform resolution at $h = 1.0, 0.5$ and 0.25 km (Fig. R5). As expected for an implementation without an internal length scale, the shear band width scales with the element size, and the peak plastic strain within the bands scales inversely, since a comparable fault offset is accommodated across a mesh-set band width. The faulting pattern itself is reproduced at $h = 1.0$ km and finer: The location, dip and surface expression of the four faults (primary, footwall-snapping, and the two in the hangingwall) are unchanged across this range. At $h = 2.0$ km, a different faulting mode appears and strain localization is much less intense. We further note that the uniform and adaptively

refined three-dimensional models agree at matched resolution (0.5 and 1.0 km), indicating that the faulting pattern is controlled by the resolution rather than by the refinement strategy.




Figure R5. Resolution tests of the DES model. Plastic strain at 0.48 My for three series of simulations. Left: two-dimensional models with uniform resolution at $h = 1.0$, 0.5 and 0.25 km. Center: three-dimensional models with uniform resolution at $h = 2.0$, 1.0 and 0.5 km. Right: three-dimensional models with adaptive mesh refinement at $h_{\min} = 1.0$, 0.5 and 0.25 km. Three-dimensional results are shown on a cross section at 1 km from the front face of the domain. All panels share the color scale shown.

The two-dimensional models with a uniform resolution reproduce the primary fault and the footwall-snapping secondary fault over the same range of resolutions, but do not develop the fault that cuts off the hangingwall block, even at $h = 0.25$ km (Fig. R5). We therefore interpret that the formation of rider block is promoted in the three-dimensional models given the rheological parameters.

The resolution test results are now a new section in Appendices.

Implications for tectonic-surface process coupling

Neither the DES nor the ASPECT model presented here carries a regularization length scale of the kind discussed by Duretz et al. (2019), by construction in DES and by choice in ASPECT, and yet both produced the same primary fault, consistent with Lavier et al. (2000), at comparable minimum element sizes. Also, within DES, the full faulting pattern is stable across meshing strategy (i.e., adaptive or uniform) and across a factor of four in resolution.

The three-dimensional faulting pattern is reproduced at resolutions of 0.25–1.0 km. We speculate that the remaining difference between DES and ASPECT in secondary faulting is more likely attributable to element type and mesh structure: irregular linear tetrahedra in DES versus locally structured, higher-order (Q2–Q1) elements in ASPECT.

Implications for earthquake cycle modeling

As detailed under Comment 6 and in Appendix G3 Décollement mesh-resolution sensitivity, the three-resolution earthquake-cycle comparison shows that the response class and first-episode time are much less sensitive to resolution than the element-scale rate measure. The fault-zone width remains 100 m in all three runs, whereas the width of localized deformation within it remains mesh dependent. The interseismic rate measure V spans a factor of 4.8, as Eq. (5.6) implies, but the interseismic friction coefficient changes by only about 0.01. These runs quantify resolution sensitivity rather than establish convergence.

Shear-band orientation provides a separate test of the plasticity formulation (Kaus, 2010; Choi and Petersen, 2015). In the extensional tests, EP and both velocity-strengthening and velocity-

weakening EP–RSF develop shear bands across multiple elements at dips of 51.2° – 54.5° , close to the 52.5° Arthur reference evaluated at the initial friction angle (Section 5.3.2; Comment 4). This agreement holds under both stable sliding and earthquake-like stick–slip. These tests use a single mesh resolution and do not establish mesh convergence.

On the expression, “shear zones are explicitly resolved”

We have revised the statement at L. 551 and similar expressions elsewhere. "Explicitly resolved" overstated what our discretization supports; we now describe the shear zones as represented at the element scale, with a width set by the local element size rather than by an internal length scale of the constitutive model.

Changes in manuscript.

In Sect. 2.2 Rheology framework,

“The Mohr–Coulomb return mapping with strain weakening is local and carries no internal length scale, so the width of localized deformation is set by the local element size.”

In Sec. 4.3 Test case: Gaussian fault evolution

“The initial mesh is refined in the central region such that the element size is about 1.2 km there and about 2.5 km outside (Fig.5b). During the simulation, adaptive refinement further reduced the element size to a minimum ($h_{\min}$) of about 0.5 km within the developing shear bands.”

“This suggests that in this particular test case our coupled surface processes were not significant enough to fundamentally alter the crustal deformation due to limited surface deposition (i.e., surface load redistribution).

The faulting pattern described above is reproduced across a factor of four in element size and under both uniform and adaptive meshing, although the shear band width scales with the element size (see Appendix C).”

At the end of Sec. 4.3 Test case: Gaussian fault evolution

“Neither model carries a regularization that sets an internal length scale: the plastic damper available in ASPECT (Bangerth et al., 2024) was not enabled in the calculation shown here, and DES has no equivalent. The difference in secondary faulting is therefore unlikely to reflect regularization in one code and not the other. We speculate instead that it is attributable to the element type and mesh structure: i.e., irregular linear tetrahedra in DES against the locally structured, higher-order (Q2–Q1) elements in ASPECT, which represent localized shear deformation as a smoother band.”

In Sections 5.2 and 5.4

Sections 5.2 and 5.4 state the absence of intrinsic length regularization and report the resolution sensitivity, citing Duretz et al. (2019). The phrase “explicitly resolved” is removed where it implied mesh objectivity.

In Conclusions

“Because the plasticity model in DES carries no internal length scale, the width of localized deformation remains set by the local element size. Incorporating an internal length scale through regularization is a priority for future development.”

A new section in Appendices

Physical implementation: issues and lack of clarity

R2–Comment 8

Your description of the physical equations is not clear everywhere, such that the reader can not understand what you did or what is implemented. This is evident from my questions above, but also for inertia:

We rewrote Sections 2.1 and 5.2.2–5.2.5 to state the physical momentum equation, Mohr–Coulomb yield function and stress–strength relation, invariant rate measure, mass scaling and damping, state integration, and time-step rule. Detailed responses are given under Comments 1, 2, 4, 9 and 10; Appendix G provides the application settings and sensitivity tests.

Changes in manuscript.

Sections 2.1 and 5.2.2–5.2.5 revised; see the changes listed under Comments 1, 2, 4, 9 and 10.

Inertia: description and implementation issues with unboundedness at large slip rates

R2–Comment 9

It is not clear if, and how, inertia in a physical form is implemented. Some formulations suggest so “Navier-Stokes equations (l. 57)” or “propagation of elastic waves “ (l. 675). However, others do not.

- Equation 2.1 uses a fictitious mass-scaled density, and the method was originally designed to damp inertial forces and recover quasi-static equilibrium. Please clarify whether $\rho_f = \rho$ during coseismic events, whether dynamic relaxation remains active and how this is tuned in relation to adaptive time stepping, and whether the resulting P- and S-wave speeds are physical. If not, the terms “dynamic,” “dynamic rupture,” and “elastic-wave propagation” should be accurately reformulated. The present formulation may be better described as a dynamically relaxed or mass-scaled explicit approximation.
- Also if you had inertia, you would not need an additional bound as a novel criteria in your adaptive time stepping, which you do in form of a local shear wave speed maximum (l. 586). You explain that you do need that to ensure that the fictitious density remains bounded when slip velocities approach shear wave speeds. However, physically, slip rate remains bounded due to inertia. Hence it is not clear why you need this.
- Without it is not fully dynamic, hence you should not claim earthquake-cycle dynamics is solved (e.g. in the title). And you would also need another energy sink to limit unbounded slip rate growth, which is in the absence of inertia typically done using a quasi-dynamic approximation (Rice et al., JGR, 1993). However, I do not see anywhere you implement that?
- I suspect that you cutoff maximum slip rates with your additional shear wave speed bound in your time stepping, rather than implementing a physical mechanism to limit rupture growth. This numerical trick may well impact your solution, may also contribute to the noise you observe towards the end of the event, and inhibits a proper verification and benchmarking of your solution.

- Since lines 675 and others discuss the simulation of waves and claim issues due to a linear approximation of waves, I wonder what your waves look like. And how you could simulate them without physical inertia. Could you show plots of the (underresolved) wave field in the supplements? And clarify how you simulate those?

DES retains the second-order acceleration term throughout the earthquake cycle. Equation (2.1) gives the momentum balance with material density; mass scaling enters through the discrete nodal mass, following Choi et al. (2013) and Tong and Lavier (2018). Componentwise FLAC damping remains active with a coefficient of 0.8 and modifies the assembled nodal force without removing acceleration. It differs from the radiation-damping approximation used in quasi-dynamic models (Rice, 1993).

Mass scaling permits longer explicit time steps during slow loading. The décollement and strike-slip models use scaling coefficients $c = 10^5$ and 10^6 , respectively. Their minimum fictitious density is 1.67ρ , including during rapid episodes, so the scaled compressional and shear wave speeds reach at most $\sqrt{G/K} = 0.775$ of their material-density values. The shear-speed ceiling sets this density floor; it does not cap nodal velocity, the constitutive rate measure V , or slip rate. The separate time-step rule in Eq. (5.9) combines the elastic and state-evolution limits, with other active processes able to impose a tighter global limit.

To isolate the effect of the density floor, we changed the apparent-speed ceiling so that ρ_{fict} returns to ρ wherever the ceiling applies. The mesh and all other parameters are unchanged in this 2000 yr comparison (Appendix G).

	reference $\rho_{\text{fict}} \geq \rho K/G$	modified $\rho_{\text{fict}} \rightarrow \rho$	ratio
First episode (yr)	745.2	745.2	1.00
Peak nodal speed, median (m s^{-1})	0.199	0.372	1.87
Peak nodal speed, maximum (m s^{-1})	0.491	0.564	1.15
Episode duration, median (s)	15.0	19.6	1.30
Stress drop ΔR , median (MPa)	4.87	6.53	1.34
Recurrence (yr)	47.8 ± 7.8	50.4 ± 2.6	1.05

The first episode remains at 745.2 yr and mean recurrence changes by 5%, whereas peak speed, duration, stress drop, and recurrence variability remain sensitive to the density floor.

Appendix D diagnoses compressional and shear disturbances using velocity divergence and curl (Fig. R6). A rapid episode produces a four-lobed compressional pattern, an expanding shear front, and progressive arrival along a transect. The fraction of the squared signal envelope ahead of a cone falls to approximately 10^{-3} at the mass-scaled compressional speed. These finite-speed arrivals demonstrate inertial propagation. The title retains “Earthquake-Cycle Dynamics” because the second-order equation governs the entire cycle.

Changes in manuscript.

Equation (2.1) now uses material density, with fictitious density introduced only in the discrete nodal mass. Sections 2.1 and 5.2.4 describe the retained acceleration, FLAC damping, mass scaling, and characteristic speeds; Section 5.2.5 gives the state update and time-step rule. Section 5.4 and Appendix G report the density-floor sensitivity. The terms “dynamic rupture”, “dynamic rupture-like”, and “elastic-wave propagation” were removed. New Appendix D and Fig. D1 demonstrate finite-speed propagation in the mass-scaled system.




Figure R6 (revised manuscript Figure D1). Propagation of the disturbance radiated by one rapid episode of the model in Sect. 5.4, at 950.91 yr, using the compressional-sensitive diagnostic $\nabla \cdot \mathbf{v}$ and the shear-sensitive diagnostic $\nabla \times \mathbf{v}$, the out-of-plane component in this 2D section; both are signed. (a–d) Snapshots 2.0 s apart during the rapid episode. Each field uses one color scale across the two times, set from the bulk, so values inside the 100 m fault zone are clipped, and the vertical exaggeration is 2.4. (e) The $|\nabla \times \mathbf{v}|$ envelope along a line 2.5 km above the fault, taken as the analytic-signal magnitude after a 2 s running mean is removed, normalized by its global maximum. Dashed lines are cones drawn from the nucleation patch at the mass-scaled speeds $c_p^{\text{app}} = 4472$ and $c_s^{\text{app}} = 2582$ m s⁻¹. (f) Fraction of the squared envelope of each field arriving ahead of a cone moving at the speed on the horizontal axis. Dashed lines mark the mass-scaled speeds; dotted lines mark the material-density $c_s = 3333$ and $c_p = 5774$ m s⁻¹.

R2–Comment 10

The rheological model description varies throughout the paper. Sometimes correctly, but also sometimes less completely. For example, in the abstract you write about RSF being implemented in a MC elastoplastic structure, but latter I also see visco-elasto-plastic or elasto-visco-plastic and full implementations. Could you check all your rheology model statements and try to clarify what you do where?

DES supports EP, EVP, EP–RSF, and EVP–RSF. The performance benchmark uses EVP and the surface-process application uses strain-weakening EP. The constitutive tests in Section 5.3.1 and the localization comparison in Section 5.3.2 use EP and EP–RSF, whereas the décollement and strike-slip applications use EP–RSF. EVP–RSF is available but is not used in this study.

Changes in manuscript.

Section 2.2 and the workflow figure list the supported rheologies consistently. The performance-benchmark table now identifies EVP, Section 4 identifies strain-weakening EP, and the opening of Section 5 identifies the EP and EP–RSF cases used in the RSF tests and applications.

Benchmarking

R2–Comment 11

The above issues illustrate the importance of benchmarking your solution also in an earthquake cycle setup.

Quickly reading your text that references to SCEC (State-wide California Earthquake Center) benchmarks of Sequences of Aseismic and Seismic Slip (SEAS) across several pieces of text makes the not careful reader think that you executed them, but you have not. First, could you not refer to those benchmarks if you did not execute them?

We agree. We did not run the SCEC/SEAS benchmarks and have therefore removed all references to that benchmark suite from the manuscript.

Changes in manuscript.

References to the SCEC/SEAS benchmark suite were removed from Sections 5.1, 5.3, and 5.4 and from the Conclusions.

R2–Comment 12

Instead you compared whether you would get nucleation, if friction increases and decreases, what maximum slip rates are, that you get recurrence and conclude that “cycle statistics are broadly in line with quasi-dynamic RSF sequence result (Lapusta.2000; Erickson.2020)”. I have seen quite some simulations that reproduce these very general features, but in those there still could be many smaller and larger issues. Hence quantification and closer examination matters here.

At the start of your benchmarking section 5.4 you indeed describe that you “must reproduce physically reasonable nucleation, rupture and healing behaviour” (L. 593/4). However, you do not verify of any of the properties whether these are physically reasonable.

I understand that reproducing SCEC benchmark designed for non-volume-based codes using a volume-based code is an endeavour you would not like to start on, and is it not realistic to demand in this paper. However, several aspects can be tested independently and analytically. This involves inter seismic healing and post seismic relaxation, such as addressed in Figure 10 of Herrendorfer et al., (2018).

The zero-velocity test verifies the aging-law healing limit examined by Herrendörfer et al. (2018) in their Fig. 10b. With all nodal velocities held at zero, θ grows almost linearly from θ_0 , and an independent iteration of Eq. (5.8) reproduces the DES state evolution over 12.70 yr.

In the décollement calculations, μ recovers to within 1% of its interseismic value after 33 yr at $a/b = 0.2$ and 4.4 yr at $a/b = 0.5$. Both recovery times are shorter than the corresponding median recurrence intervals of 225 and 50 yr (Section 5.4.2).

Postseismic relaxation is not independently verified here. Comment 13 presents the comparison with their Figs. 4b and 4c.

Changes in manuscript.

Section 5.3.1 reports the zero-velocity healing verification, and Section 5.4.2 quantifies frictional recovery between rapid episodes. The general claims about “physically reasonable” nucleation, rupture, and healing were replaced by specific comparisons.

R2–Comment 13

In addition, your recurrence pattern can be compared to Herrendorfer et al., (2018)’s Figure 4c. Since your setup is very similar to that one, such a comparison should be rather straightforward and very meaningful. Other more detailed evaluations for the nucleation and propagation are available in there as well, which would be useful for evaluating those processes, but that can be regarded as beyond the scope of this study.

Several authors of volume-based codes have been discussing to setup a separate, new benchmark that is designed such that it can be robustly simulated and defined for volume-based codes. The idea is to start from a setup similar to Herrendorfer et al., (2018) and yours. I think it would be useful for your model to join this endeavour. However, until then and in any case, the large-scale cycle pattern of Herrendorfer et al., (2018) can be reproduced independently of the progress on that front.

In the comparison with Figs. 4b and 4c of Herrendörfer et al. (2018), DES produces a median recurrence interval of 12.1 yr, close to the reference 12.7 yr (Fig. R1). DES cycles are more variable, with intervals of 5.3–22.8 yr ($CV = 0.33$). Rapid slip remains mainly within the velocity-weakening patch, whereas rupture in the reference model enters the velocity-strengthening regions. Rapid episodes account for about 36% of fault displacement at $X = 69$ km, with the remainder accumulating during interseismic phases; the reference model releases nearly all accumulated slip deficit during rupture.

DES divides the reference fault-zone width among three elements and assigns $D_c = L/3 = 3.33$ mm from the reference one-cell distance $L = 0.01$ m (Section 5.4.1; Fig. R7). The element-level D_c is not directly interchangeable across band discretizations (Section 5.4.3). Because band discretization and D_c differ between the models, this comparison cannot isolate the effect of the constitutive formulation. We agree that a benchmark designed for volume-based codes would be valuable and would be glad to participate.

Changes in manuscript.

Section 5.4.3 (“Strike-slip: comparison with Herrendörfer et al. (2018)”) and Fig. 15 compare the slip-rate history, cumulative-slip profiles, interseismic slip, and the resolution dependence of the element-level D_c . Section 5.4.1 gives the comparison setup, with the geometry in Fig. 12b.




Figure R7 (revised manuscript Figure 12b). (b) The DES comparison model based on Herrendörfer et al. (2018). Unlike the décollement parameter-study model, this is a map view of a strike-slip fault, so gravity is zero, the confining stress is the uniform traction P_B on the two sides, and X and Z are horizontal coordinates. The 250 m band is discretized with three elements across its thickness and 250 m along strike, in a mesh that coarsens outward to 8 km at the domain boundaries. The band is drawn with an exaggerated thickness; the magnified stretch at right gives its true proportions, with the monitor stations on its two faces.

R2–Comment 14

Another way to evaluate whether you capture essential earthquake-cycle physics and reproduce its nucleation is via comparing / demonstrating that your simulations nucleate a “dynamic” rupture once the nucleation size is reached. How do your nucleation size and the theoretical nucleation size mentioned in L. 643 compare?

- L. 641-644 describes that you resolve the nucleation size because you use 4-5 elements across the fault width of 100m. However, in all other smaller width fault zone applications, nucleation size ALONG the fault is measured, such that ruptures from then on can propagate ALONG the fault. -> also holds? Update text
- Note also that if it were the across fault width that mattered, I would not term your simulation as “well resolved” as it is only just or not even resolved.

The relevant resolution comparison is along the fault. For $a/b = 0.2$, the low- a/b interface solution of Rubin and Ampuero (2005) gives fixed nucleation half lengths of about 103 and 51 m for $D_c = 5 \times 10^{-3}$ and 2.5×10^{-3} m, respectively, under the same G -based modulus convention used in Appendix G2. These reference lengths span about four to five and two to three local element lengths along the fault (Fig. R8). The internal nucleation structure remains unresolved on this mesh.

Across the parameter study, L_∞ gives the clearest overall separation of rapid and non-rapid responses, correctly classifying 61 of 69 runs (Appendix G2; Fig. R9). Both constant- L_∞ contours cross from rapid to non-rapid behavior as a/b increases, showing that a/b also affects the response at fixed L_∞ .

Changes in manuscript.

Section 5.4.1 compares the reference nucleation half length with the along-fault element length and

states that the internal nucleation structure remains unresolved. Appendix G2 and Fig. G1 report the response classification and the constant- L_∞ comparisons.




Figure R8 (revised manuscript Figure 12a). (a) Domain, materials, mesh, and boundary conditions for the décollement parameter study: a vertical cross section with gravity. Four to five elements with local lengths of about 20–25 m span the 100 m fault-zone width.




Figure R9 (revised manuscript Figure G1). Response classification in the friction-parameter study. (a) The a/b - D_c parameter grid and the 14 calculations along $L_\infty = 325$ and 450 m. Marker colors distinguish rapid, intermediate-velocity, recurrent low-velocity, and near-loading-rate responses. (b) Peak nodal speed against a/b along the two constant- L_∞ contours, with line colors and dashes as in (a). (c) Peak nodal speed against L_∞ for all runs with finite L_∞ . Horizontal dotted lines mark the peak-speed class boundaries, and the dashed line marks the loading rate. The asymptotic nucleation interpretation of L_∞ applies at larger a/b .

R2–Comment 15

I appreciate the new ‘shear-box’ benchmark of a single element that you introduce, which can be used by others as well. However, the aspects discussed above are not tested in there. Also you impose the velocity field, which makes the element-level slip rate projection known in advance. This predefines the slip rate, but you argue in l. 621-622 that you test and validate the nodal-to-element velocity interpolation and the projection onto the maximum-shear-stress direction, i.e., the way in which you calculate the slip rate. I do not think you can verify your slip rate equation in this predefined slip rate benchmark with fixed stress orientation. You test one simple case where key things to test are defined.

The simple-shear benchmark verifies the local constitutive response at a known rate (Section 5.3.1). The simple-shear velocity field gives both elements the same $\dot{\epsilon}_{II}$, w , and V (Fig. R10). Equation (5.6) uses this invariant measure without a stress-derived slip direction.

Evolving fault kinematics are examined in the extensional tests, where shear bands develop from a localized weak seed in EP and both EP–RSF cases (Section 5.3.2; Comment 4).

Changes in manuscript.

Section 5.3.1: the claim that the benchmark validates the interpolation and projection was removed; the scope was restated as local constitutive verification under known kinematics. Fig. 10a was redrawn to mark the element length scale w used by Eq. (5.6).




Figure R10 (revised manuscript Figure 10). Simple-shear constitutive verification. (a) Two-element geometry and minimum element altitude w in $V = 2w\dot{\epsilon}_{II}$. (b–d) Shear stress and relative error for elastoplastic, steady-state RSF, and aging-law RSF cases. Lines denote reference solutions and symbols DES results; the inset compares state evolution.

R2–Comment 16

Also I do not think based on this benchmark you can claim that “All three benchmark problems show excellent agreement between numerical and analytical results (Fig. 10b–d), validating the RSF implementation.” (L. 620). I agree that this benchmark verifies the local RSF constitutive update for prescribed constant-velocity simple shear. However, “validating the RSF implementation” seems too broad. The test does not verify the coupled system-level behaviour under evolving slip rates, including spontaneous nucleation, rupture propagation, arrest, healing, or earthquake recurrence. I suggest replacing “validating” with “verifying the local constitutive implementation” and clearly distinguishing this test from validation of the full earthquake-cycle model.

We revised the statement to identify what the benchmark verifies: the elastoplastic return mapping and RSF stress and state updates under simple shear (Section 5.3.1). Sections 5.3.2 and 5.4 separately evaluate coupled fault growth and earthquake-cycle behavior, including the comparison with Herrendörfer et al. (2018).

Changes in manuscript.

Section 5.3.1 describes the benchmark as verification of the elastoplastic return mapping and RSF stress and state updates under simple shear. Sections 5.3.2 and 5.4 present the localization tests and earthquake-cycle applications separately.

R2–Comment 17

You study four models with RSF parameters to generate different types of sequences. All your cases have the fault being one property; either velocity-weakening (VW) or velocity-strengthening (VS). However, you do not have the transition between VW and VS then. It is at that transition where in nature nucleation is preferred, so if you want to study and claim reasonable nucleation you need that. Additionally, that transition is numerically harder to simulate, so you need to show you can do that well with one case. This is also done in all SCEC benchmarks.

If not a full SCEC or a targeted benchmark, then expect at least to verify different aspects as mentioned above. The state-of-the-art in method developments has reached that stage several years ago. Also since you in your abstract write that you build a robust platform, I think robustness should be demonstrated.

The strike-slip model contains the requested weakening–strengthening transitions (Section 5.4.3, Fig. 15). Following Herrendörfer et al. (2018), the model contains a velocity-weakening patch over $X = 32\text{--}108$ km, with $a = 0.011$ and $b = 0.017$, flanked by velocity-strengthening regions with $b = 0.001$ across 4 km transitions. The calculation produces spontaneous rapid episodes, and the cumulative-slip profiles show rapid slip concentrated mainly within the weakening patch and decreasing across the transitions. Comment 13 compares its recurrence, slip partitioning, and rupture extent with the reference model.

The décollement study isolates the effects of uniform a/b and D_c , while Section 5.3.2 compares fault growth under uniform velocity-strengthening and velocity-weakening RSF. The abstract no longer claims a robust platform.

Changes in manuscript.

New Section 5.4.3 and Fig. 15 add a heterogeneous fault with velocity-weakening to velocity-strengthening transitions and compare its recurrence, slip partitioning, and rupture extent with Herrendörfer et al. (2018). Section 5.4.1 distinguishes this case from the uniform décollement parameter study.

Medium points

R2–Comment 18

The abstract and many places throughout the manuscript write about “validation”, but you are not checking that the model reproduces the physical system adequately. You do not compare to any laboratory or natural observations or discuss in that direction. You are doing “verification”, i.e., whether your equations are implemented and solved correctly.

We use “verification” for the local constitutive comparisons with analytical and semi-analytical solutions under known kinematics (Section 5.3.1). The earthquake-cycle applications evaluate coupled behavior at model scale; they do not provide validation against laboratory or natural observations.

Outside the RSF section, bitwise reproducibility is described as a requirement for debugging and checking numerical consistency, the DES–goSPL coupling is described as tested, and the correspondence between the goSPL river network and the DES incision pattern is described as qualitative support.

Changes in manuscript.

Manuscript-wide: “validation” and its variants were removed. Section 5.3 now distinguishes local constitutive verification from the model-scale applications of Section 5.4. Sections 3.1 and 4.3 and the Conclusions now use evidence-specific terms such as “bitwise reproducibility”, “tested”, and “qualitative support”.

R2–Comment 19

Your abstract states “across temporal and spatial scales” and your introduction “to bridge the gap between long-term tectonic deformation and earthquake-cycle dynamics”. However, your simulations are done at fixed time frames; geodynamics OR earthquake cycle, so they do not bridge them by simulating both in the same framework. Could you update your formulations throughout and clarify the link with the latest paper as well? Koelzer et al., (2026) is under review like yours (<https://doi.org/10.22541/essoar.15007647/v1>), and there geodynamic and earthquake-cycle time scales are bridged and the functioning of earthquake-cycles under in a model that also in the same application shows tectonic deformation is clear.

We thank the reviewer for bringing the study by Koelzer et al. (2026) to our attention. We extended the Arthur-angle tests from EP to EP–RSF, including both velocity-strengthening and velocity-weakening cases. These simple extensional models reproduce fault growth near the Arthur reference angle, accompanied by stable sliding or repeated earthquake-like stick-slip. The results support using DES to model fault growth together with earthquake-cycle behavior. We revised the Abstract and Introduction to describe these fault-growth and slip-cycle results and to distinguish the separate tectonic–surface-process and earthquake-cycle applications. Applying this approach to subduction zones and more complex tectonic settings, as explored by Koelzer et al. (2026), is a promising direction that we plan to investigate in future work, beyond the scope of the present study.

Changes in manuscript.

Abstract and Section 1: the “across temporal and spatial scales” and “bridge the gap” formulations were removed, and the main applications are identified as separate. Section 5.3.2 adds the coupled fault-growth and slip-cycle calculation.

R2–Comment 20

You do not describe the change from 2D to 3D as one of the main novelties (you only note it briefly at line 54). However, I think this is an important novelty, which was facilitated by your technical HPC work. Why was that not added as a key novelty?

We appreciate this point. The omission of Ta et al. (2015) made the development history unclear: DES was already a three-dimensional unstructured solver, so the 2D-to-3D extension is not a novelty of DES v2.0. We added the missing reference and corrected the Abstract, Introduction, Methods, and Conclusions.

The advance in DES v2.0 is instead the modernization of the existing 3D solver through a single-source OpenACC CPU–GPU implementation, deterministic two-stage updates, and GPU-resident time integration and remapping. The earlier study used OpenCL for selected kernels on tightly coupled CPU–GPU hardware and excluded remeshing from its performance analysis.

Changes in manuscript.

Abstract, Sections 1, 2, and 3.2, and Conclusions: the 2D-to-3D extension is no longer presented as new to DES v2.0. Ta et al. (2015) is now cited for the earlier 3D and OpenCL implementation, and the DES v2.0 contribution is framed around its single-source OpenACC implementation, deterministic updates, and GPU-resident remapping.

R2–Comment 21

Since your off-fault behaviour would always weaken (or strengthen) with slip rate changes, can you still simulate static off-fault plasticity as is studied in dynamic rupture studies?

Yes. RSF and plastic-strain weakening are assigned independently by material. The fault can use RSF while the host rock uses rate-independent plasticity, with cohesion and friction angle either fixed or evolving with plastic strain. The décollement and strike-slip applications use $C = 10^{50}$ Pa to suppress host-rock yielding (Appendix G1). With finite cohesion and plastic-strain weakening, the EP, velocity-strengthening EP–RSF, and velocity-weakening EP–RSF tests develop shear bands from a localized weak seed (Section 5.3.2, Fig. 11; Comments 4 and 19). Thus, the assigned rheology and strength parameters determine whether the surrounding material remains elastic or undergoes plastic deformation and localization.

Changes in manuscript.

Appendix G1 specifies the wall-rock cohesion used to suppress off-fault yielding in the earthquake-cycle applications. Section 5.3.2 presents the finite-strength localization tests.

R2–Comment 22

Since you have a complex workflow with certain parts of it operating on CPU’s and others on GPU’s, are your timings and performance affected by availability of either resource? Does that mean that your code performs better on certain nodes that have an optimal balance between CPU and GPU architectures, such as available on the Grace Hopper?

Yes, the timings depend on both resources. The time-integration kernels run on the GPU, whereas mesh-topology generation and output remain CPU-bound, so as GPU execution becomes faster these CPU stages take a larger share of the total wall time. Section 3.2.5 now quantifies this: on the H100 configuration the GPU-resident kernel step falls from 94.3% of wall time at the smallest problem size to 69.5% at the largest, while the CPU-side share rises to between 6.0% and 27.5%. The reported H100 and A100 results are therefore system-level measurements for the H100–Grace and A100–EPYC 7763 node configurations rather than isolated GPU comparisons, and Section 3.4 now says so.

A well-balanced system should benefit this hybrid workflow, and we note that the H100 configuration is in fact a Grace–Hopper node, which is the balance the reviewer asks about. We are nonetheless careful not to overclaim. The two configurations differ in host CPU, host core count, and GPU generation simultaneously, so they do not isolate the effect of CPU–GPU balance and we cannot identify an optimal ratio from them. Establishing one requires varying the host thread count against a fixed device, which we plan as a separate scaling test.

Changes in manuscript.

Section 3.4 now names the host CPU and core count of each configuration, identifies the cluster each was run on, and states that because remeshing and I/O retain CPU-bound stages the host CPU contributes to their measured performance, so comparisons of the hybrid stages between the H100 and A100 configurations reflect node-level rather than GPU-level differences. Section 3.2.5 quantifies the shift in the CPU-side share of wall time with problem size.

R2–Comment 23

Your adaptive time stepping claims to resolve the state variable’s evaluation (l. 585), but you only have a time step criteria on the slip rate. Herrendorfer et al., JGR, 2018 suggests that also two criteria for are needed to resolve state weakening and healing.

We implemented both state-evolution criteria and reran all earthquake-cycle and sensitivity calculations. For the forward-Euler update, DES limits the fractional change in θ through $\Delta t \leq \Delta\theta_{\max} D_c/V$ for slip-induced state decay and $\Delta t \leq \Delta\theta_{\max} \theta$ for healing, adapted from Herrendörfer et al. (2018). These limits are evaluated over all elements and combined with the elastic and other active time-step limits.

All 74 parameter-study runs and the resolution and density-floor tests use $\Delta\theta_{\max} = 0.2$.

Changes in manuscript.

Section 5.2.5 now presents the forward-Euler state update and both state-evolution time-step limits in Eqs. (5.8)–(5.9). All earthquake-cycle and sensitivity calculations were rerun with these limits.

R2–Comment 24

Your references and state-of-the-art description are missing another geodynamic earthquake cycle model; Sobolev et al., GGG, 2017. They also use a finite element model with plasticity (albeit Drucker-Prager), where only the friction coefficient is extended with RSF and a state evolution to simulate elastoviscoplastic earthquake cycles. Since this is the most similar model to yours (also see points below), could you clarify the novelty w.r.t. that model/paper?

We thank the reviewer for highlighting Sobolev and Muldashev (2017). Their model couples RSF to Drucker–Prager plasticity and estimates slip velocity from plastic shear strain rate and element size. DES uses a Mohr–Coulomb return mapping and a rate measure based on the total deviatoric strain-rate invariant and local element length.

DES integrates momentum explicitly on unstructured Lagrangian meshes with remeshing capability, using scaled nodal masses and FLAC damping. Our rapid-slip calculations demonstrate propagating compressional and shear disturbances with finite-speed arrivals. The calculations of Sobolev and Muldashev (2017) focus on event onset and cumulative coseismic slip without resolving detailed rupture propagation. Our extensional tests further show fault growth across elements near the Arthur angle, without a predefined fault path, together with stable sliding for $a > b$ and repeated earthquake-like stick–slip for $a < b$. These results demonstrate coupled fault growth and earthquake-cycle behavior within the explicit DES framework.

Changes in manuscript.

Section 5.1.2 now includes Sobolev and Muldashev (2017) among the prior continuum RSF models and states the constitutive and rate-measure differences from DES.

Minor points

R2–Comment 25

At various locations you write about state-dependent restrengthening (e.g. 1.72), but state evolution also weakens during earthquakes. -> clarify

We agree. State evolution contributes to both weakening and restrengthening: decreasing θ lowers the state-dependent contribution to friction, whereas increasing θ raises it.

Changes in manuscript.

Sections 2.2 and 5.2.1 now describe both weakening and restrengthening produced by state evolution.

R2–Comment 26

Why do you compare the Grace CPU and EPYC CPU in Figure 3, whereas the A100 and H100 performed better in Figure 2? -> Add justification

Figures 2 and 3 address different questions. Figure 2 compares CPU and GPU performance as model size increases. Figure 3 examines scaling with CPU core count, measuring speedup against a

single-core baseline on each platform. This comparison has no directly equivalent GPU core-count baseline. It identifies where parallel efficiency is lost within a compute node, including the poor scaling of remeshing.

The hardware configurations reflect the resources available through our TACC allocations. The Grace CPU and H100 benchmarks were run on Vista, and the EPYC 7763 and A100 benchmarks on Lonestar6. The CPU–GPU pairings follow the compute-node hardware configurations.

Changes in manuscript.

A paragraph was added at the start of Section 3.4.4 stating that the subsection examines scaling with respect to CPU core count against a single-core baseline, that this independent variable has no direct GPU analogue so the analysis is restricted to the two CPU platforms, and that its purpose is to locate lost parallel efficiency rather than to rank hardware. Section 3.4 also now names the two clusters on which each configuration was run.

R2–Comment 27

L. 474: The references cited here do not all quantify off-fault plastic work in finite-thickness fault-zone models. Several instead concern fault-interface slip or field-scale hydromechanical evolution. Limit to those actually using finite-thickness fault zones or “plastic work” only if energy dissipation is explicitly calculated. Petrini et al., *Tectonophysics*, 2020 is worthwhile in this finite thickness - off-fault plasticity - earthquake context.

We did not quantify energy dissipation in this study and have therefore removed the term “plastic work”. Section 5.1 now describes plastic deformation and shear localization in DES, and cites Petrini et al. (2020) for a finite-thickness continuum model that couples deformation with pore-pressure evolution, compaction, and strength recovery.

Changes in manuscript.

Section 5.1 describes plastic deformation and shear localization in DES and cites Petrini et al. (2020) for coupled hydromechanical evolution. The term “plastic work” was removed.

R2–Comment 28

L. 476: In this introduction 5.1, is it fair to here only cite your own work?

We agree that the discussion needed broader coverage of previous work. Section 5.1 now places DES within the literature on interface-based and continuum earthquake-cycle models, including the continuum formulations of van Dinther et al. (2013), Sobolev and Muldashev (2017), and Herrendörfer et al. (2018). Tong and Lavier (2018) and Lavier et al. (2021) are identified as the direct precursors within DES.

Changes in manuscript.

Section 5.1 was reorganized into Sections 5.1.1 and 5.1.2 on interface-based and continuum RSF implementations, respectively, and the continuum RSF literature was expanded.

R2–Comment 29

L. 477: Here, but also at other locations, the novelty of the current RSF implementation w.r.t. Tong and Lavier, 2018 and Lavier et al., 2021 is not always clear. -> add a few words on adding state evolution here and maybe elsewhere

We agree. Tong and Lavier (2018) and Lavier et al. (2021) are the direct precursors within DES. Both studies use the steady-state form of RSF, with plastic strain history serving as a proxy for frictional state through plastic hysteresis. DES v2.0 adds a separate state variable θ that evolves independently of accumulated plastic strain and represents time-dependent interseismic healing. The zero-velocity test verifies this local healing update (Section 5.3.1).

Changes in manuscript.

Section 5.1 distinguishes the earlier steady-state RSF formulation from the independent state evolution in DES v2.0. Section 5.2.1 gives its evolution law.

R2–Comment 30

L. 492: You may want to discuss finite and heterogeneous fault zone in subduction zones (e.g. in relation to your old work) and whether fault geometries are well-known at depth here.

Fault-zone geometry and internal structure at seismogenic depths are poorly constrained in active subduction zones. The observations from exhumed faults cited in Section 5.1 document heterogeneous fault cores and surrounding fractured rock; they motivate representing deformation across a fault zone.

Changes in manuscript.

Section 5.1 describes the heterogeneous structure of exhumed faults and its representation in continuum models.

R2–Comment 31

L. 517: van Dinther et al., (2013)’s model is also elastic, not only viscoplastic; how else can earthquakes stick and slip?

We agree. Our earlier description was too narrow. van Dinther et al. (2013) use a viscoelasto-plastic seismo-thermo-mechanical model that includes elastic, viscous, and plastic deformation. Section 5.1.2 now uses that description.

Changes in manuscript.

Section 5.1.2: the description of van Dinther et al. (2013) was corrected from “viscoplastic” to “viscoelastoplastic.”

R2–Comment 32

L. 518-522: Actually, Herrendorfer (2018)’s formulation is also a highly non-linear viscous one, because that allows it to match fundamental RSF assumptions with stress = strength. This is apparent in its supplements, but it is better explained in Koelzer et al., (2026).

We agree. Our earlier description did not make this constitutive distinction clear. In Herrendörfer et al. (2018), a nonlinear viscoplastic viscosity enforces $\tau_{II} = \sigma_{\text{yield}}$, so plastic deformation remains active even at very low rates. Sections 5.1.2 and 5.2.2 now distinguish this formulation from the DES threshold-plasticity update, in which stress can remain below strength without plastic deformation.

Changes in manuscript.

Sections 5.1.2 and 5.2.2 now describe the nonlinear viscoplastic formulation of Herrendörfer et al. (2018) and contrast its continuous stress–strength relation with the threshold-plasticity update used in DES.

R2–Comment 33

L. 549: Goudarzi et al. (2023) does not use only a single-cell-wide zone, but also has multiple cell shear zones, where the paper argues under which conditions these are resolved and how different implementations thereof compare.

We agree. Our earlier statement incorrectly characterized Goudarzi et al. (2023) as using only a single-cell-wide fault zone. Section 5.1.2 now describes their comparison of plasticity-based, Kelvin-viscoplastic, and phase-field representations, including fixed-thickness fault zones spanning multiple cells, and their evaluation of mesh independence.

Changes in manuscript.

Section 5.1.2: the description of Goudarzi et al. (2023) was corrected to include their multiple-cell fault zones and comparison of continuum fault representations.

References

- Bangerth, W., Dannberg, J., Fraters, M., Gassmoeller, R., Glerum, A., Heister, T., Myhill, R., and Naliboff, J.: ASPECT: Advanced Solver for Planetary Evolution, Convection, and Tectonics, User Manual, <https://doi.org/10.6084/m9.figshare.4865333>, 2024.
- Choi, E. and Petersen, K. D.: Making Coulomb angle-oriented shear bands in numerical tectonic models, *Tectonophysics*, 657, 94–101, <https://doi.org/10.1016/j.tecto.2015.06.026>, 2015.
- Choi, E., Tan, E., Lavier, L. L., and Calo, V. M.: DynEarthSol2D: An efficient unstructured finite element method to study long-term tectonic deformation, *Journal of Geophysical Research: Solid Earth*, 118, 2429–2444, <https://doi.org/10.1002/jgrb.50148>, 2013.
- Duretz, T., de Borst, R., and Le Pourhiet, L.: Finite thickness of shear bands in frictional viscoplasticity and implications for lithosphere dynamics, *Geochemistry, Geophysics, Geosystems*, 20, 5598–5616, <https://doi.org/10.1029/2019GC008531>, 2019.

- Goudarzi, M., Gerya, T., and van Dinther, Y.: A comparative analysis of continuum plasticity, viscoplasticity and phase-field models for earthquake sequence modeling, *Computational Mechanics*, 72, 615–633, <https://doi.org/10.1007/s00466-023-02311-0>, 2023.
- Herrendörfer, R., Gerya, T., and van Dinther, Y.: An invariant rate- and state-dependent friction formulation for viscoelastoplastic earthquake cycle simulations, *Journal of Geophysical Research: Solid Earth*, 123, 5018–5051, <https://doi.org/10.1029/2017JB015225>, 2018.
- Kaus, B. J.: Factors That Control the Angle of Shear Bands in Geodynamic Numerical Models of Brittle Deformation, *Tectonophysics*, 484, 36–47, <https://doi.org/10.1016/j.tecto.2009.08.042>, 2010.
- Koelzer, A. J., de Vos, M., Gerya, T., and van Dinther, Y.: Bridging the gap between millions of years and milliseconds: Modeling earthquake sequences and splay fault rupture in subduction zones, *ESS Open Archive*, <https://doi.org/10.22541/essoar.15007647/v1>, version 1, posted 19 August 2026, 2026.
- Lavier, L. L., Buck, W. R., and Poliakov, A. N. B.: Factors controlling normal fault offset in an ideal brittle layer, *Journal of Geophysical Research: Solid Earth*, 105, 23 431–23 442, <https://doi.org/10.1029/2000JB900108>, 2000.
- Lavier, L. L., Tong, X., and Biemiller, J.: The mechanics of creep, slow slip events, and earthquakes in mixed brittle–ductile fault zones, *Journal of Geophysical Research: Solid Earth*, 126, e2020JB020 325, <https://doi.org/10.1029/2020JB020325>, 2021.
- Petrini, C., Gerya, T., Yarushina, V., van Dinther, Y., Connolly, J., and Madonna, C.: Seismo-hydro-mechanical modelling of the seismic cycle: Methodology and implications for subduction zone seismicity, *Tectonophysics*, 791, 228 504, <https://doi.org/10.1016/j.tecto.2020.228504>, 2020.
- Preuss, S., Herrendörfer, R., Gerya, T., Ampuero, J.-P., and van Dinther, Y.: Seismic and aseismic fault growth lead to different fault orientations, *Journal of Geophysical Research: Solid Earth*, 124, 8867–8889, <https://doi.org/10.1029/2019JB017324>, 2019.
- Rice, J. R.: Spatio-temporal complexity of slip on a fault, *Journal of Geophysical Research*, 98, 9885–9907, <https://doi.org/10.1029/93JB00191>, 1993.
- Rubin, A. M. and Ampuero, J.-P.: Earthquake nucleation on (aging) rate and state faults, *Journal of Geophysical Research: Solid Earth*, 110, B11 312, <https://doi.org/10.1029/2005JB003686>, 2005.
- Sobolev, S. V. and Muldashev, I. A.: Modeling seismic cycles of great megathrust earthquakes across the scales with focus at postseismic phase, *Geochemistry, Geophysics, Geosystems*, 18, 4387–4408, <https://doi.org/10.1002/2017GC007230>, 2017.
- Ta, T., Choo, K., Tan, E., Jang, B., and Choi, E.: Accelerating DynEarthSol3D on tightly coupled CPU–GPU heterogeneous processors, *Computers & Geosciences*, 79, 27–37, <https://doi.org/10.1016/j.cageo.2015.03.003>, 2015.
- Tong, X. and Lavier, L. L.: Simulation of slip transients and earthquakes in finite thickness shear zones with a plastic formulation, *Nature Communications*, 9, 3893, <https://doi.org/10.1038/s41467-018-06390-z>, 2018.

van Dinther, Y., Gerya, T. V., Dalguer, L. A., Mai, P. M., Morra, G., and Giardini, D.: The seismic cycle at subduction thrusts: Insights from seismo-thermo-mechanical models, *Journal of Geophysical Research: Solid Earth*, 118, 6183–6202, <https://doi.org/10.1002/2013JB010380>, 2013.