

Simulations of a stretching bar using a plasticity model from the shear transformation zone theory

Chris H. Rycroft^{a,b}, Frédéric Gibou^c

^a*Department of Mathematics, University of California, Berkeley, CA 94720*

^b*Department of Mathematics, Lawrence Berkeley Laboratory, Berkeley, CA 94720*

^c*Departments of Mechanical Engineering and Computer Science, University of
California, Santa Barbara, CA 93106*

Abstract

An Eulerian simulation is developed to study an elastoplastic model of amorphous materials that is based upon the shear transformation zone theory developed by Langer and coworkers [1]. In this theory, plastic deformation is controlled by an effective temperature that measures the amount of configurational disorder in the material. The simulation is used to model ductile fracture in a stretching bar that initially contains a small notch, and the effects of many of the model parameters are examined. The simulation tracks the shape of the bar using the level set method. Within the bar, a finite difference discretization is employed that makes use of the essentially non-oscillatory (ENO) scheme. The system of equations is moderately stiff due to the presence of large elastic constants, and one of the key numerical challenges is to accurately track the level set and construct extrapolated field values for use in boundary conditions. A new approach to field extrapolation is discussed that is second order accurate and requires a constant amount of work per gridpoint.

Keywords: plasticity, numerical methods

1. Introduction

Plastic deformation in amorphous materials (*e.g.* glasses or granular media) displays a wide variety of complex behavior such as shear-localization

Email addresses: chr@math.berkeley.edu (Chris H. Rycroft),
fgibou@engineering.ucsb.edu (Frédéric Gibou)

and jamming, or shear history dependence. Particles in an amorphous material interact according to short-range forces on a very short timescale, and may undergo rapid random vibrations or collisions with each other, while keeping their overall configuration. Plastic deformation, in which the particles change their configuration, happens on a longer timescale. Experiments and simulations show that these configuration changes can happen in small localized regions that have been described as flow defects [2, 3, 4] and mean field approaches based upon this microscopic model have been developed.

Over the past decade, Langer and coworkers [1] have developed the shear transformation zone (STZ) theory, which can be viewed as extension to this work. The basic concept, introduced by Falk and Langer [5], is to say that an amorphous material contains a number of STZs, that represent localized regions of the material that are susceptible to plastic deformation. Typically, STZs may be present in multiple orientations; this could either be described with tensors, or by proposing a simplified model in which there are several populations aligned with each axis. Macroscopic plastic deformation is then the sum of events at STZs, each causing local configuration changes. After the configurational change, the STZ is no longer available for further transformations in the original shear direction. However, STZs are continually being created and annihilated, either mechanically during the shearing itself, or thermally in response to the background temperature. Depending on the rate factors used, this simple physical picture can be used to understand many of the features of amorphous materials. Jamming can be explained when the rates of STZ formation during shear are not enough to sustain the process. Shear localization happens when deformation creates more STZs in one particular area, making it progressively more amenable to flow. History-dependent effects occur when shear leads to STZs being preferentially present at certain orientations as opposed to others.

The STZ theory can be interpreted within a statistical mechanics framework [6], by making use of an effective temperature T_{eff} , that measures the amount of configurational disorder in a packing. At high energies, this is equivalent to the thermal temperature, but the two separate below the glass transition. Suppose the energy required to form an STZ is E_Z . Then the probability of a site being occupied by an STZ is proportional to the Boltzmann factor $\exp(-E_Z/k_B T_{\text{eff}})$. In situations where the real temperature is unimportant, an athermal version of the theory can be developed [7, 8] where the dynamics is driven by the effective temperature only.

From the microscopic principles of the STZ theory, a macroscopic elasto-

plastic continuum model can be derived which can be used to predict and test against real data. Manning et al. [9, 10] have simulated a variety of models to investigate shear-localization in a one-dimensional model, compared with analytic results. Pechenik [11] has derived a multi-dimensional elastoplastic model, and a two-dimensional model has been developed to study the necking instability in a stretched bar [12]. However, in this model, the width of the bar is described as function of horizontal position and time only, meaning that the numerics cannot approach the pinch-off point.

The aim of this paper is to develop a numerical framework for simulating and testing continuum elastoplastic STZ theories for complex deforming bodies. A variety of different numerical techniques could be employed, and in solid mechanics it is particularly common to employ the finite-element method with a mesh deforming with the object, such as those employed by the software package ABAQUS. However, in the STZ theory it is natural to make use of an Eulerian perspective, since it is never necessary to reference an initial undeformed state, and all history information can be encapsulated by the STZ densities themselves. We have therefore developed a finite difference framework, using the level set method to track the edge of bodies under consideration. The level set method makes it particularly easy to study topological changes and examine nonlinear behavior that would be hard to study analytically.

To simulate a deforming body, the continuum stress and velocity fields are only held for grid points that are within the body itself. Near the edge of the body, boundary conditions are applied by extrapolating the fields and constructing ghost values. Some of the key numerical challenges are to correctly handle this extrapolation, in the presence of large elastic constants that make the underlying equations stiff. In the following Section, the governing equations are presented, and are rescaled into a dimensionless form that is used within the simulation. In Section 3, the implementation of the level set method and extrapolation technique is discussed. The numerical methods for solving the elastoplastic model are presented in Section 4. In Section 5 several convergence studies are performed to test the numerical scheme, and a variety of simulations are carried out to investigate the different parameters in the STZ system.

2. Governing equations

This paper is based upon a two-dimensional elastoplastic model with a velocity $\tilde{\mathbf{u}} = (\tilde{u}, \tilde{v})$ and a stress tensor

$$\tilde{\sigma} = \begin{pmatrix} -\tilde{p} + \tilde{s} & \tilde{\tau} \\ \tilde{\tau} & -\tilde{p} - \tilde{s} \end{pmatrix}.$$

The angular velocity is defined as $\tilde{\omega} = \frac{1}{2}(\partial\tilde{v}/\partial\tilde{x} - \partial\tilde{u}/\partial\tilde{y})$. Here, tildes are used to signify dimensional quantities that are later replaced with rescaled dimensionless equivalents to be used within the simulation. In dimensional form, the elastoplastic system can be written as

$$\tilde{\rho}_0 \frac{d\tilde{u}}{d\tilde{t}} = -\frac{\partial\tilde{p}}{\partial\tilde{x}} + \frac{\partial\tilde{s}}{\partial\tilde{x}} + \frac{\partial\tilde{\tau}}{\partial\tilde{y}} + \tilde{\kappa}\nabla^2\tilde{u} \quad (1)$$

$$\tilde{\rho}_0 \frac{d\tilde{v}}{d\tilde{t}} = -\frac{\partial\tilde{p}}{\partial\tilde{y}} - \frac{\partial\tilde{s}}{\partial\tilde{y}} + \frac{\partial\tilde{\tau}}{\partial\tilde{x}} + \tilde{\kappa}\nabla^2\tilde{v} \quad (2)$$

$$\frac{d\tilde{p}}{d\tilde{t}} = \tilde{K} \left(\frac{\partial\tilde{u}}{\partial\tilde{x}} + \frac{\partial\tilde{v}}{\partial\tilde{y}} \right) \quad (3)$$

$$\frac{d\tilde{s}}{d\tilde{t}} = 2\tilde{\omega}\tilde{\tau} + \tilde{\mu} \left(\frac{\partial\tilde{u}}{\partial\tilde{x}} - \frac{\partial\tilde{v}}{\partial\tilde{y}} \right) - \frac{2\tilde{\mu}\tilde{s}\tilde{D}}{|\tilde{\sigma}_0|} \quad (4)$$

$$\frac{d\tilde{\tau}}{d\tilde{t}} = -2\tilde{\omega}\tilde{s} + \tilde{\mu} \left(\frac{\partial\tilde{u}}{\partial\tilde{y}} + \frac{\partial\tilde{v}}{\partial\tilde{x}} \right) - \frac{2\tilde{\mu}\tilde{\tau}\tilde{D}}{|\tilde{\sigma}_0|} \quad (5)$$

where \tilde{K} and $\tilde{\mu}$ are the elastic constants, $\tilde{\rho}_0$ is the density of the material, \tilde{D} is the plastic deformation rate, $\tilde{\kappa}$ is the viscosity, and $\tilde{\sigma}_0$ refers to the deviatoric part of the stress tensor. Here, $d/d\tilde{t}$ represents the convective derivative, so that for a given field f , $df/d\tilde{t} = \partial f/\partial\tilde{t} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} f$. In this model, we assume that elastic constants are large, so that the elastic deformations are small when compared to the plastic parts. This allows us to assume a constant density. The above system is based on the assumption that the rate-of-deformation tensor can be viewed as the sum of the elastic and plastic parts, referred to as hypoelastoplasticity. For small elastic deformations, this can be shown to be equivalent to the hyperelastoplastic model, in which the total deformation is written as a map $F = F_e \circ F_p$, where F_e is the elastic part and F_p is the plastic part.

In its most general form, the STZ theory specifies the plastic part of the rate-of-deformation tensor as

$$D_{ij}^{\text{pl}} = e^{-1/\chi} f_{ij}(\tilde{\sigma}_0, T),$$

where $\chi = T_{\text{eff}}k_B/E_Z$ is a dimensionless effective temperature, that characterizes the configurational disorder of the material, and affects its mechanical response. The tensor f_{ij} contains the details of the STZ dynamics and is a function of the deviatoric stress tensor and the bath temperature T . A number of different functional forms for D_{ij}^{pl} have been considered but here we concentrate on a specific athermal model discussed by Bouchbinder et al. [7, 8] where T plays no role and there is no detailed tensor dependence, so that

$$\tilde{D} = \frac{\epsilon_0}{\tau_0} e^{-1/\chi} \tilde{q}(|\tilde{\sigma}_0|)$$

where ϵ_0 is a dimensionless constant of order unity and τ_0 is a characteristic timescale at which the material responds to microscopic perturbations. The function \tilde{q} is given by

$$\tilde{q}(\zeta) = \begin{cases} \frac{\tilde{a}}{\zeta} (\zeta - s_Y)^2 & \text{if } \zeta > s_Y \\ 0 & \text{if } \zeta \leq s_Y \end{cases}$$

where \tilde{a} is an additional model parameter that has dimensions of inverse stress. Plastic deformation only occurs when the magnitude of the deviatoric stress tensor exceeds a specified yield stress s_Y . In this formalism the effective temperature χ evolves according to

$$\frac{d\chi}{dt} = \frac{2\epsilon_0}{\tau_0 c_0} e^{-1/\chi} \tilde{q}(|\tilde{\sigma}_0|) \frac{|\tilde{\sigma}_0|}{s_Y} (\chi_\infty - \chi),$$

where c_0 is a dimensionless parameter of order unity, so that as plastic deformation occurs, the configurational disorder of the material increases until it saturates at a value of χ_∞ .

2.1. Rescaling the equations

The simulations in this paper are carried out in rescaled dimensionless units. To begin, the yield stress is scaled out, so that

$$p = \frac{\tilde{p}}{s_Y}, \quad s = \frac{\tilde{s}}{s_Y}, \quad \tau = \frac{\tilde{\tau}}{s_Y}$$

$$K = \frac{\tilde{K}}{s_Y}, \quad \mu = \frac{\tilde{\mu}}{s_Y}, \quad \rho_0 = \frac{\tilde{\rho}_0}{s_Y}.$$

The first five of these quantities are dimensionless, while the last has units of $[\text{Time}]^2/[\text{Length}]^2$, an inverse speed squared. The transverse wave speed is given by considering the linearized equation for s ,

$$\frac{\partial^2 s}{\partial \tilde{t}^2} \sim \frac{\mu}{\rho_0} \frac{\partial^2 s}{\partial \tilde{x}^2},$$

from which it is natural to define a speed

$$c_s = \sqrt{\frac{\mu}{\rho_0}}.$$

Using this, lengths and times can be rescaled in terms of a characteristic system size L , so that $x = \tilde{x}/L$ and $t = c_s \tilde{t}/L$. The velocities are then rescaled according to $c_s u = \tilde{u}$ and $c_s v = \tilde{v}$ so that the transverse wave speed is unity. The function controlling the plastic deformation is rescaled according to $a = \tilde{a} s_Y$ and $a q(|\sigma_0|) = \tilde{q}(|\tilde{\sigma}_0|)$, so that

$$q(\zeta) = \begin{cases} \frac{1}{\zeta}(\zeta - 1)^2 & \text{if } \zeta > 1 \\ 0 & \text{if } \zeta \leq 1. \end{cases}$$

A dimensionless viscosity can be defined as $\kappa = \tilde{\kappa} c_s / L \tilde{\mu}$. Finally, many of the physical parameters in the STZ model can be scaled out by defining

$$\nu = \frac{2\epsilon_0 L a}{c_s \tau_0},$$

after which the system of equations can be written as

$$\mu \frac{du}{dt} = -\frac{\partial p}{\partial x} + \frac{\partial s}{\partial x} + \frac{\partial \tau}{\partial y} + \mu \kappa \nabla^2 u \quad (6)$$

$$\mu \frac{dv}{dt} = -\frac{\partial p}{\partial y} - \frac{\partial s}{\partial y} + \frac{\partial \tau}{\partial x} + \mu \kappa \nabla^2 v \quad (7)$$

$$\frac{dp}{dt} = K \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) \quad (8)$$

$$\frac{ds}{dt} = 2\omega\tau + \mu \left(\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y} \right) - \mu \nu e^{-1/\chi} \frac{s}{|\sigma_0|} q(|\sigma_0|) \quad (9)$$

$$\frac{d\tau}{dt} = -2\omega s + \mu \left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} \right) - \mu \nu e^{-1/\chi} \frac{\tau}{|\sigma_0|} q(|\sigma_0|) \quad (10)$$

$$\frac{d\chi}{dt} = \frac{\nu}{c_0} e^{-1/\chi} (\chi_\infty - \chi) |\sigma_0| q(|\sigma_0|). \quad (11)$$

The rescaled plastic deformation can then be defined as $D = \nu e^{-1/\chi} q(|\sigma_0|)$. The parameter ν controls the amount of plastic deformation, with a smaller value corresponding to more brittle behavior, and a larger value giving a more ductile response.

3. A level set implementation with rapid field extrapolation

The level set method is used to track the edge of the bar as it deforms in the simulation. This method has been extensively used in the literature and standard techniques are described in textbooks [13, 14]. However, in this section, we present a variation that avoids some problems specifically associated with simulating the PDE system introduced in the previous section. In many problems involving moving interfaces, particularly those involving multi-phase fluid flow, the continuum fields exist everywhere, and the level set boundary demarcates regions with different physical characteristics, such as changes in density or velocity. In the solid mechanics simulations presented here, all the field values for stress and velocity can only be defined for points within the bar itself. In order to apply boundary conditions and move the level set, it is necessary to be able to rapidly extrapolate field values to a small skin of grid points next to the bar. Aslam [15] introduced a method to extrapolate normally outward from the level set boundary by solving a sequence of separate PDE problems. If executed with the most efficient methods, this requires $O(N \log N)$ operations with a large prefactor. Here, we introduce a variation of the method that can carry out normal extrapolation with a single sweep, using $O(N)$ operations with a small prefactor, and achieving higher accuracy close to the interface.

For the simulations presented here, and in other situations, it is advantageous to maintain the level set as a signed distance function during the computation. However, as discussed elsewhere, transporting the level set according to the simulation velocity field will typically perturb this property. In the current simulations, if the signed distance function property is not maintained, then the normal field extrapolations will become inaccurate. In other situations, the signed distance function property is advantageous for accurately calculating forces applied on the interface, such as surface tension.

Several approaches have been discussed in the literature for maintaining the signed distance function property of the level set. One strategy is to periodically reinitialize the level set to be signed distance function. A popular

approach is that introduced by Sussman et al. [16], where the equation

$$\frac{\partial \phi}{\partial T} = \text{sign}(\phi)(1 - |\nabla \phi|)$$

is solved to a steady state by integrating forward in the fictitious time unit T . This PDE gradually improves the signed distance property, so that that $|\nabla \phi| \rightarrow 1$. This method was tried in the current simulations, but it has two drawbacks. Firstly, it was found to be computationally expensive, requiring a significant number of updates for each grid point. It also has a tendency to perturb the position of the boundary, $\phi(\mathbf{x}, t) = 0$, since this is not explicitly constrained. For the current study, which uses moderately stiff elastic constants with zero stress boundary conditions applied at the level set, these perturbations can lead to unpredictable results: small movements of the level set generate fictitious elastic waves. Fine features of the level set can be smeared out, with the initial notch in the bar becoming more rounded. Movement of the level set can be minimized by upwinding derivatives at adjacent gridpoints [17, 18] although this was not considered here.

A second approach to reinitialization is that described by Sethian [13]. In this method, the position of the interface is first located by linearly interpolating between adjacent gridpoints where ϕ changes sign. With the boundary located, a first order method can be used to update the level set values at these adjacent gridpoints. Two fast marching methods can then be employed to initialize the values at the remaining gridpoints, solving the Eikonal equation $|\nabla \phi| = 1$ forwards for positive gridpoints, and backwards for negative gridpoints. This method was also investigated, and typically runs faster, since the algorithm only sets each gridpoint once, rather than iteratively improving. It also does not smear out fine features of the level set boundary as much as the PDE approach. However, despite finding the boundary explicitly, the resetting of the adjacent level set values is only first order accurate, and can cause perturbations that lead to artificial elastic modes.

With either of these methods, a parameter must be introduced that sets the interval between successive reinitializations. This choice of interval is a trade-off, with a shorter interval creating more perturbations, and a longer interval leading to a less accurate level set. Having this additional parameter is undesirable, since it was found to have a weak effect on some of the physical results.

The method presented here maintains the signed distance function property throughout the calculation, and does not require specific corrections at certain times. It makes use of a combination of a direct update for points next to the boundary using the second order method of Chopp [19], and a second order fast marching method to update the remaining points. The level set is stored as a narrow band that is continually updated, rather than being rebuilt at specific intervals. The code keeps track of a list of points currently in the band, sorted according to ϕ values. The procedure for building an initial narrow band is described below, after which the method for moving the band is introduced. The field extrapolation method is presented in Subsec. 3.3 and tested in Subsec. 3.4.

3.1. Building the narrow band

To begin, a discretized function $\psi_{i,j}$ is provided whose zero contour is the interface to be tracked, but that may not be a signed distance function itself. To construct the narrow band level set function, all gridpoints (i, j) are found that are next to the interface, so that a neighboring gridpoint has an opposite sign.

For each of these points, the iteration described by Chopp [19] is carried out. To begin, a bicubic interpolation $\psi_b(\mathbf{x})$ is constructed from the $\psi_{i,j}$. The vector \mathbf{x}^0 is set to the position of the gridpoint (i, j) , and the following iteration is carried out:

$$\delta_1 = -\psi_b(\mathbf{x}^k) \frac{\nabla\psi_b(\mathbf{x}^k)}{\nabla\psi_b(\mathbf{x}^k) \cdot \nabla\psi_b(\mathbf{x}^k)} \quad (12)$$

$$\mathbf{x}^{k+1/2} = \mathbf{x}^k + \delta_1 \quad (13)$$

$$\delta_2 = (\mathbf{x}^0 - \mathbf{x}^k) - \frac{(\mathbf{x}^0 - \mathbf{x}^k) \cdot \nabla\psi_b(\mathbf{x}^k)}{\nabla\psi_b(\mathbf{x}^k) \cdot \nabla\psi_b(\mathbf{x}^k)} \nabla\psi_b(\mathbf{x}^k) \quad (14)$$

$$\mathbf{x}^{k+1} = \mathbf{x}^{k+1/2} + \delta_2. \quad (15)$$

This is a combination of two Newton–Raphson root finding methods. After the iteration has been carried out, the limiting point \mathbf{x}^∞ satisfies both $\psi_b(\mathbf{x}^\infty) = 0$ and $(\mathbf{x}^\infty - \mathbf{x}^0) \times \nabla\psi_b(\mathbf{x}^\infty) = 0$. Typically, three or four iterations are required in order to achieve double precision floating point accuracy.

In general, the above algorithm converges very reliably. Typically, the magnitude of the gradient $|\nabla\psi_b(\mathbf{x}^k)|$ is non-zero close to the zero contour, and thus the divisors in the equations for δ_1 and δ_2 are large. However, certain cases can cause problems, such as a local maximum, described by

$\psi_{i,j} = h/2$ and $\psi_{i\pm 1,j} = \psi_{i,j\pm 1} = -h/2$ where h is the grid spacing, in which case $\psi_b(\mathbf{x})$ is locally flat at (i, j) so that $|\nabla\psi_b(\mathbf{x}^k)|$ vanishes. In the simulation, these problem cases are recognized either when the number of iterations exceeds 100, or when the size of step taken $|\mathbf{x}^{k+1} - \mathbf{x}^k|$ exceeds $\sqrt{2}h$, the maximum step that can be taken within a single grid square. In those cases, the simulation falls back on the less accurate first order method described by Sethian [13].

Once the iteration is carried out, the level set function at (i, j) is initialized to be $\phi_{i,j} = \pm|\mathbf{x}^\infty - \mathbf{x}^0|$ using a negative sign for points inside the body and a positive sign for points outside. We write $c_\alpha = (i, j)$ to index a particular grid point, and then store two unsorted index lists, $L_1^- = \{c_{-\alpha}, \dots, c_{-1}\}$ and $L_1^+ = \{c_0, \dots, c_\beta\}$ for negative points and positive points respectively. After all points are considered, L_1^- and L_1^+ can be resorted so that the points are in ascending order by the corresponding value of ϕ .

The remainder of the level set narrow band can then be constructed using two fast marching methods to sweep upwards and downwards from these known points. We make use of the standard techniques, first creating a list of trial values for $\phi_{i,j}$ at gridpoints adjacent to those that are known. The values of the points are chosen via the second order discretization described by Sethian [13]. This is based upon a first order difference operator that switches to a second order scheme when the gridpoints are available and monotonic. In the negative x direction, this can be written as

$$E_{i,j}^{-x}\phi = \begin{cases} \frac{3\phi_{i,j} - 4\phi_{i-1,j} + \phi_{i-2,j}}{2h} & \text{if } \phi_{i-1,j} \text{ and } \phi_{i-2,j} \text{ are known and} \\ & \phi_{i-2,j} < \phi_{i-1,j} \\ \frac{\phi_{i,j} - \phi_{i-1,j}}{h} & \text{otherwise} \end{cases}$$

and similar expressions exist in the other directions. With this, the trial values of $\phi_{i,j}$ can be constructed by solving the quadratic

$$\max [E_{i,j}^{-x}\phi, -E_{i,j}^{+x}\phi, 0]^2 + \max [E_{i,j}^{-y}\phi, -E_{i,j}^{+y}\phi, 0]^2 = 1.$$

These are stored in a binary heap, that allows us to select the lowest one and fix its value. New values are then constructed for the $\phi_{i,j}$ at locations adjacent to this fixed value. The computation of the narrow band stops once the values of ϕ reach a cutoff.

During this computation, new index lists $L_2^- = \{c_{-\gamma}, \dots, c_{-1-\alpha}\}$ and $L_2^+ = \{c_{\beta+1}, \dots, c_\delta\}$ can be constructed using the negative and positive points respectively. Unlike L_1^\pm , these lists can be constructed in the ascending order

of the corresponding ϕ values without the need for subsequent resorting, since the fast marching method considers the points in order. From here, a master list of indices $\{L_2^-, L_1^-, L_1^+, L_2^+\} = \{c_{-\gamma}, \dots, c_\delta\}$ can be constructed. Each section of this list is sorted, but the maximum value of L_2^- may exceed the minimum of L_1^- , and the maximum value of L_1^+ may exceed the minimum of L_2^+ . A global ordering can be very quickly constructed by applying a single sweep of merge sort at these boundaries. We refer to $N = \delta + \gamma + 1$ as the total number of elements on this list.

3.2. Moving the narrow band

We now consider moving the band according to a given discretized velocity field $\mathbf{u}_{i,j}$, making use of techniques similar to those used in the initial band construction. Bicubic interpolations of the velocity $\mathbf{u}_b(\mathbf{x})$ and the level set function $\phi_b(\mathbf{x})$ can be defined, from which it is possible to define a normal velocity field

$$n(\mathbf{x}) = \mathbf{u}_b(\mathbf{x}) \cdot \nabla \phi_b(\mathbf{x}).$$

To update the position of the band, all gridpoints (i, j) next to the interface are first considered. Starting with \mathbf{x}^0 at the gridpoint, the Newton iteration given in Eqs. 12–15 is applied to find the closest point \mathbf{x}^∞ on the interface. The new value at the gridpoint is then given by

$$\phi_{i,j}(t + \Delta t) = \phi_{i,j}(t) + n(\mathbf{x}^\infty) \Delta t$$

where Δt is the timestep. Once all gridpoints next to the interface are considered, the remaining points can be initialized using two second order fast marching methods, in the same manner as described above. Similar routines can be used to construct a sorted list $\{c_{-\gamma}, \dots, c_\delta\}$ of the narrow band points.

3.3. Field extrapolation

For a scalar field $f(\mathbf{x})$ given within a region described by a level set, $\phi(\mathbf{x}) < 0$, it is possible to carry out a constant normal extrapolation by solving to steady state the PDE

$$\frac{\partial f}{\partial t} + H(\phi) \hat{\mathbf{n}} \cdot \nabla f = 0$$

where $\hat{\mathbf{n}} = \nabla \phi / |\nabla \phi|$. Here H is the Heaviside function, so that only values outside the region are modified. In practice, this procedure can be carried out using a fast marching method, constructing new level set values outside

the band, and then setting the field values by using $\hat{\mathbf{n}} \cdot \nabla f = 0$ with the second order operators $E^{\pm x}, E^{\pm y}$ introduced previously.

As discussed by Aslam [15], this approach can be generalized to carry out higher order field extrapolation. For example, for quadratic extrapolation, the sequence of PDEs

$$\begin{aligned}\frac{\partial f_{nn}}{\partial t} + H(\phi)\hat{\mathbf{n}} \cdot \nabla f_{nn} &= 0 \\ \frac{\partial f_n}{\partial t} + H(\phi)(\hat{\mathbf{n}} \cdot \nabla f_n - f_{nn}) &= 0 \\ \frac{\partial f}{\partial t} + H(\phi)(\hat{\mathbf{n}} \cdot \nabla f - f_n) &= 0\end{aligned}$$

must be solved. The first normal derivative is given by

$$\begin{aligned}f_n(\mathbf{x}) &= \hat{\mathbf{n}} \cdot \nabla f \\ &= \frac{\phi_x f_x + \phi_y f_y}{|\nabla \phi|}\end{aligned}$$

and the second normal derivative is

$$\begin{aligned}f_n(\mathbf{x}) &= \hat{\mathbf{n}} \cdot \nabla (\hat{\mathbf{n}} \cdot \nabla f) \\ &= \frac{1}{|\nabla \phi|^2} (\phi_x^2 f_{xx} + 2\phi_x \phi_y f_{xy} + \phi_y^2 f_{yy} + (\phi_x^3 f_y + \phi_y^3 f_x) \phi_{xy} \\ &\quad + \phi_x \phi_y (\phi_x (f_y (\phi_{yy} - \phi_{xx}) - f_x \phi_{xy}) + \phi_y (f_x (\phi_{xx} - \phi_{yy}) - f_y \phi_{xy}))).\end{aligned}$$

It can be seen from these expressions that in order to construct f_n , first derivatives in x and y are needed. These are calculated using centered differences, and thus they can only be calculated for points where the orthogonally adjacent points are also within the region. For the second normal derivative, cross derivatives ϕ_{xy} and f_{xy} are needed, which additionally requires that the diagonally adjacent gridpoints are also within the region. Because of these restrictions, Aslam makes use of a modified system of equations,

$$\begin{aligned}\frac{\partial f_{nn}}{\partial t} + H(\phi)\hat{\mathbf{n}} \cdot \nabla f_{nn} &= 0 \\ \frac{\partial f_n}{\partial t} + H(\phi - h)(\hat{\mathbf{n}} \cdot \nabla f_n - f_{nn}) &= 0 \\ \frac{\partial f}{\partial t} + H(\phi - \sqrt{2}h)(\hat{\mathbf{n}} \cdot \nabla f - f_n) &= 0,\end{aligned}$$

so that the derivatives f_n and f_{nn} are extrapolated from interior regions that contain only gridpoints for which the required neighbors have valid field values. These equations are then computed using three separate fast marching methods, extrapolating outwards from the boundaries $\phi(\mathbf{x}) = -\sqrt{2}h, -h, 0$ in sequence. In some cases, f_n and f_{nn} may exist outside the regions from which they are extrapolated, but these values are not made use of.

Here, a modified approach is introduced, that makes use of the sorted boundary index $\{c_{-\gamma}, \dots, c_\delta\}$ to extrapolate outwards normally in a single sweep with $O(N)$ operations. By making use of a bisection search, the first index α can be found so that $\phi(c_\alpha) > -\sqrt{2}h$, in $O(\log N)$ time. R is defined as the set of gridpoints that are inside the region $\phi(\mathbf{x}) < 0$. The extrapolation considers each point in the list in turn, according to the following iteration:

```

while  $\alpha \leq \delta$  do
   $(i, j) = c_\alpha$ 
  if  $(i, j)$  is within  $R$  then
    if  $(i \pm 1, j)$  and  $(i, j \pm 1)$  are all within  $R$  then
      if  $(i \pm 1, j \pm 1)$  are all within  $R$  then
         $\{(f, f_n, f_{nn})$  exist, no action needed $\}$ 
      else
        Extrapolate  $f_{nn}$ 
      end if
    else
      Extrapolate  $f_{nn}$ 
      Extrapolate  $f_n$ 
    end if
  else
    Extrapolate  $f_{nn}$ 
    Extrapolate  $f_n$ 
    Extrapolate  $f$ 
  end if
   $\alpha \rightarrow \alpha + 1$ 
end while

```

With this algorithm, each point is only considered once, and any values for f , f_n , and f_{nn} that can be directly constructed are used, instead of being extrapolated. Extrapolating a field at a particular gridpoint makes use of

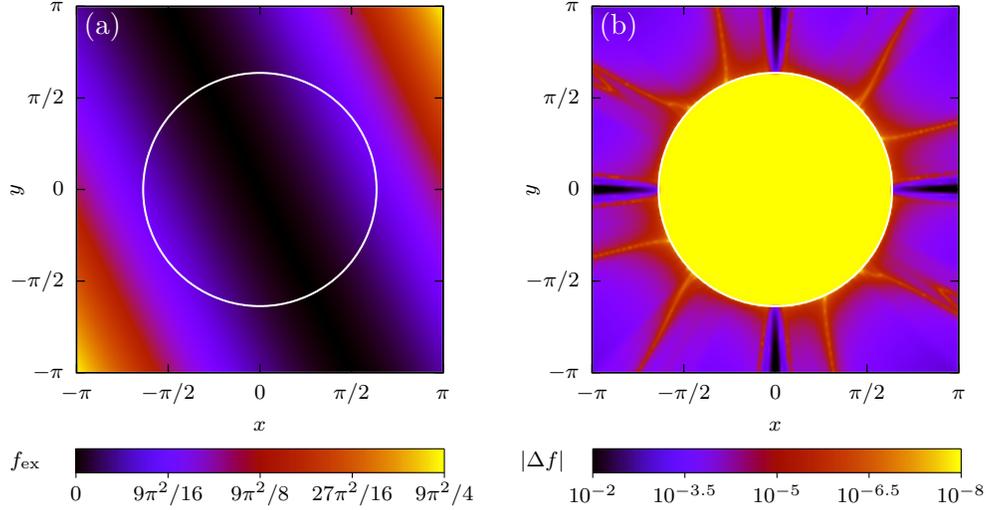


Figure 1: (a) A quadratic function f that is initially defined inside the circle $x^2 + y^2 = 4$, and is extrapolated to the region outside the circle. (b) Errors Δf between the original function f and the extrapolated function f_{ex} . These figures make use of a 256×256 grid.

the equation

$$\hat{\mathbf{n}} \cdot \nabla g = l, \quad (16)$$

where (g, l) can be $(f_{nn}, 0)$, (f_n, f_{nn}) , or (f, f_n) . The direction of extrapolation is computed by evaluating the derivatives of the existing level set field. In the x direction we compute

$$\left(\frac{\partial g}{\partial x} \right)_{i,j} = \begin{cases} E_{i,j}^{-x} g & \text{if } E_{i,j}^{-x} \phi > -E_{i,j}^{+x} \phi \\ E_{i,j}^{+x} g & \text{if } E_{i,j}^{-x} \phi \leq -E_{i,j}^{+x} \phi. \end{cases}$$

Here, when $E_{i,j}^{\pm x} g$ is evaluated, the monotonicity condition in the definition of $E_{i,j}^{\pm x}$ applies to the values of ϕ , since if ϕ is non-monotonic g should not be treated as smooth over this range. In a similar manner, $(\partial g / \partial y)_{i,j}$ and $\hat{\mathbf{n}}_{i,j}$ can be computed, after which the value of $g_{i,j}$ can be found by substituting into Eq. 16. For maximum efficiency, values of f_n and f_{nn} for interior points are not computed beforehand, but are computed dynamically and stored only when an extrapolation procedure at a neighboring gridpoint requires them. In cases where an extrapolation is attempted and there are no neighboring values to reference, the extrapolated value is set to zero.

3.4. Tests of the algorithm

The algorithms to build the level set and carry out field extrapolation have been tested in a square domain, $x \in [-\pi, \pi]$, $y \in [-\pi, \pi]$, for a variety of $n \times n$ grid sizes. The initial boundary is a circle defined by $\psi(x, y) = \sqrt{x^2 + y^2} - 2$, and the function to be extrapolated is

$$f(x, y) = \left(x - \frac{y}{2}\right)^2$$

which is initially defined in the region $\psi(x, y) < 0$. This function serves as a good test of the algorithm, since it is quadratic along any normal ray from the circle's surface, so the extrapolated field $f_{\text{ex}}(x, y)$ should converge exactly to it.

The algorithms were tested on fifteen different grid sizes from $n = 32$ to $n = 4096$. A typical output, for $n = 256$, is shown in Fig. 1(a), where the initial field $f(x, y)$ is plotted within the circle, and the extrapolated field $f_{\text{ex}}(x, y)$ is plotted outside. A variety of different running times were analyzed for the different values of n and are shown in Fig. 2. The large grid sizes require a significant amount of memory allocation (*i.e.* more than 100 Mb), and times are reported for both a Mac Pro system and a Gentoo Linux system, to examine possible differences in computer architecture and cache utilization. The first time reported is the initial grid scan of the ψ field that is required in order to locate the interface. While this is a very rapid operation, it is necessary to consider every gridpoint, and therefore it scales according to $O(n^2)$.

The next two times make use of a narrow band over the range $-5\sqrt{2}h < \phi < 2.5\sqrt{2}h$, for a grid spacing of $h = 2\pi/n$. Here, we would expect the time to build the level set function to scale according to $O(n \log n)$, but we see slightly worse scaling, that can be fit to $O(n^{1.4})$ on both machines – this is likely to be explained by a lowering of memory cache performance when dealing with a very large grid. The field extrapolation is significantly faster than building the level set function, running between four to ten times faster. Theoretically, we expect the field extrapolation to have $O(n)$ complexity as each gridpoint requires a constant amount of time, although here, the results fit well to $O(n^{1.2})$. The remaining two times in the figure are based upon computing the complete level set function and extrapolating f over the entire grid. We expect at least $O(n^2)$ scaling as each gridpoint must be considered. The times are approximately $O(n^{2.1})$, although some small differences become apparent at large grid sizes between the Linux and Mac Pro systems. For

both the narrow band and the complete grid, it can be seen that the field extrapolation is at least several times faster than building the level set.

Errors in the two algorithms were considered for the case when the complete fields were constructed. We examined errors in the L_1 , L_2 , and L_∞ norms, and we also considered the L_∞ norm restricted to the gridpoints immediately outside the circle that border the interface. Figure 3(a) shows the differences between the built ϕ field and the initial ψ field. Since the ψ field is a perfect distance function to the interface we would expect these two fields to perfectly match. The scaling in the L_1 , L_2 , and L_∞ norms is approximately $O(n^{-1.9})$, $O(n^{-1.8})$, and $O(n^{-1.0})$ respectively. This is consistent with the second order fast marching method employed here. If the residual is plotted, as shown in Fig. 1(b), it can be seen that the vast majority of error occurs close to $x = 0$ and $y = 0$, which are the regions where the fast marching method makes use of the first order scheme when there are not enough gridpoints available. The errors also become larger for areas that are further away from the boundary. If the L_∞ norm is restricted only to the gridpoints adjacent to the boundary, the errors are very small and scale according to $O(n^{-3.0})$.

The errors between the extrapolated field f_{ex} and the original field f are shown in Fig. 3(b), and follow a very similar behavior. The scaling in the L_1 , L_2 , and L_∞ norms is approximately $O(n^{-2.0})$, $O(n^{-1.5})$, and $O(n^{-1.0})$ respectively. For gridpoints adjacent to the boundary, the L_∞ norm scales according to $O(n^{-2.0})$. In the following simulations, only extrapolated values gridpoints that are close to the boundary are needed, and this scaling is satisfactory.

4. Numerical methods

The simulations are carried out in the fixed rectangular domain $-1 < x < 1, -0.5 < y < 0.5$. Vertical walls are located at $x = \pm x_{\text{wall}}$, with an initial value of $x_{\text{wall}} = 0.73$. A rectangular bar occupies the region $|y| < 0.25$ between the walls, and an initial notch of width n_w and height n_h is made in the bar by removing the region $|x| < n_w, y > 0.25 - n_h(1 + \cos(x\pi/n_w))/2$. A finite difference discretization on a $2n \times n$ rectangular grid is employed. In the x direction, the gridpoints cover the range from -1 to 1 , while in the y direction, they cover the range from $-(n-1)/(2n-1)$ to $(n-1)/(2n-1)$ to ensure an equal grid spacing in the two coordinate directions of $h = 2/(2n-1)$. A level set $\phi_{i,j}$ is initialized as the signed distance function to

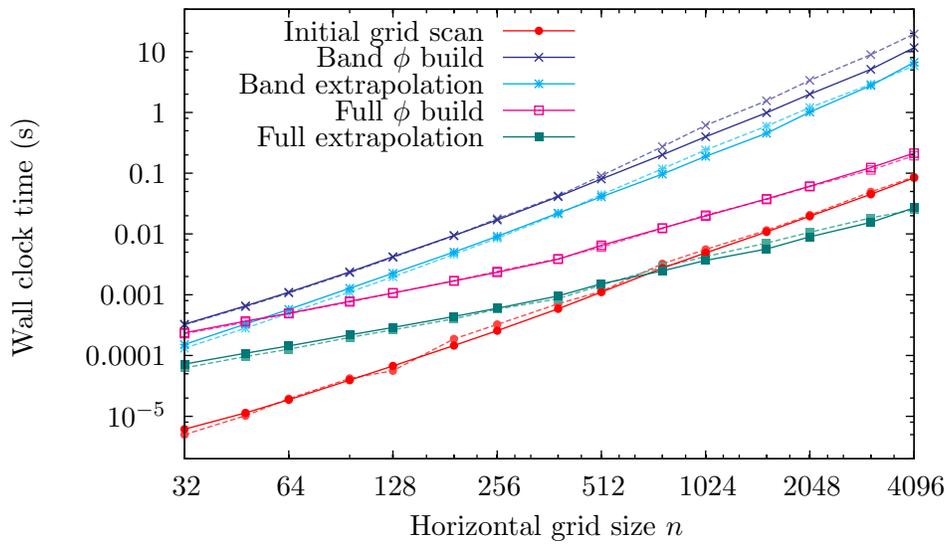


Figure 2: Running times for the level set algorithm and the quadratic extrapolation routine for different grid sizes. Solid lines were carried out using a single thread on a Mac Pro system with 5 Gb of memory and a dual-core 2.66 GHz Intel Xeon processor. Dashed lines were carried out using a single thread on a Gentoo Linux system with 16 Gb of memory and quad-core 2.60 GHz AMD Opteron processor.

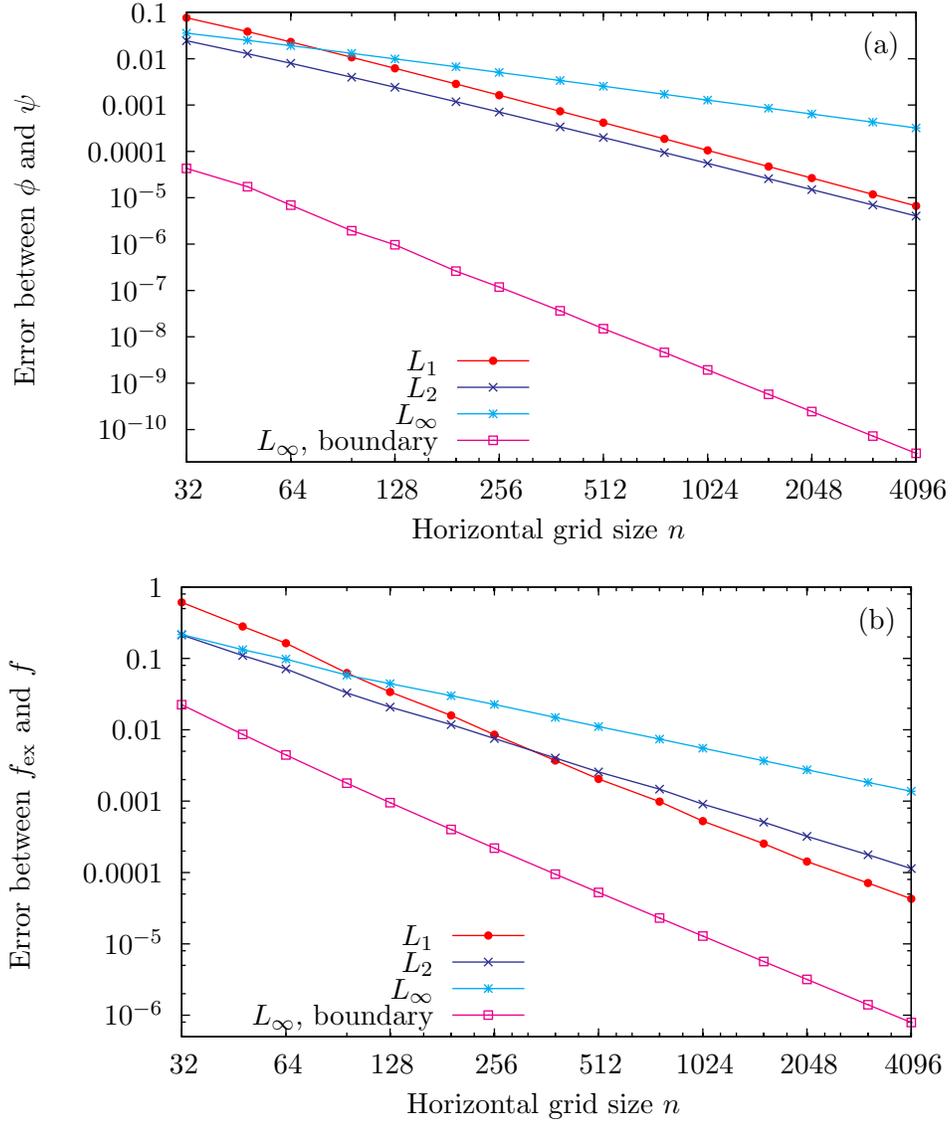


Figure 3: Errors in (a) the level set field ϕ and (b) the extrapolated field f_{ex} for the test problem, for four different error measures.

the edge of the bar. The interior of the bar is defined by $\phi_{i,j} < 0$, and for these gridpoints the velocity and stress are initialized to be zero, while the initial effective temperature is set to χ_0 .

The simulations are carried out using a first order semi-implicit integration with a fixed timestep Δt . For each step, the following procedure is taken:

1. Extrapolate the velocity, stress, and effective temperature to a skin of gridpoints outside the bar.
2. Move the level set boundary using the velocity field.
3. Carry out a finite difference update for all points within the bar.

Parts 1 and 2 make use of the level set methods discussed Sec. 3, using a narrow band with the same size as those considered in the numerical tests. Although the extrapolation method can carry out quadratic extrapolation, we made use of linear extrapolation, since this provided more stable results with the finite difference schemes that were employed. The finite difference update and boundary conditions for part 3 are discussed below.

4.1. Implementation of the derivatives

The finite difference update is carried out using a second order accurate spatial discretization that makes use of a combination of centered differencing and the essentially non-oscillatory (ENO) scheme developed by Shu and Osher [20]. We write the first and second order centered difference approximations of a given field f as

$$\begin{aligned} [f_x]_{i,j} &= (f_{i+1,j} - f_{i-1,j})/2h \\ [f_{xx}]_{i,j} &= (f_{i+1,j} - 2f_{i,j} + f_{i-1,j})/h^2. \end{aligned}$$

In the horizontal direction, the ENO derivative is then given by

$$\left\{ \frac{\partial f}{\partial x} \right\}_{i,j} = \frac{1}{2h} \begin{cases} -f_{i+2} + 4f_{i+1,j} - 3f_{i,j} & \text{if } u_{i,j} < 0 \text{ and} \\ & |[f_{xx}]_{i,j}| > |[f_{xx}]_{i+1,j}| \\ 3f_{i,j} - 4f_{i-1,j} + f_{i-2,j} & \text{if } u_{i,j} > 0 \text{ and} \\ & |[f_{xx}]_{i,j}| > |[f_{xx}]_{i-1,j}| \\ f_{i+1,j} - f_{i-1,j} & \text{otherwise.} \end{cases}$$

The scheme switches between a second order forwards and backwards derivative and a centered difference, depending upon which group of three points

has a lower second derivative. This procedure reduces short range oscillations and tolerates rapid changes in the field values. Below, a derivative evaluated using centered differencing is written within square brackets, and a term evaluated using the ENO method is written within braces.

Two different discretization methods have been investigated. In method A, the derivatives appearing on the right hand side of Eqs. 6 to 11 are calculated with centered differences, while the advective derivatives that are part of the left hand side are computed using the ENO method. The velocity and pressure are updated using an explicit Euler step, so that

$$u_{i,j}^{n+1} = u_{i,j}^n + \Delta t \left(-u\{u_x\}_{i,j}^n - v\{u_y\}_{i,j}^n + \frac{[\tau_y - p_x + s_x + \mu\kappa\nabla^2 u]_{i,j}^n}{\mu} \right) \quad (17)$$

$$v_{i,j}^{n+1} = v_{i,j}^n + \Delta t \left(-u\{v_x\}_{i,j}^n - v\{v_y\}_{i,j}^n + \frac{[\tau_x - p_y - s_y + \mu\kappa\nabla^2 v]_{i,j}^n}{\mu} \right) \quad (18)$$

$$p_{i,j}^{n+1} = p_{i,j}^n + \Delta t \left(-u\{p_x\}_{i,j}^n - v\{p_y\}_{i,j}^n + K[u_x + v_y]_{i,j}^n \right). \quad (19)$$

The deviatoric stress components are handled using a semi-implicit method in time, that can better handle rapid plastic deformation. In certain situations, the terms involving $\nu e^{-1/\chi} q(|\sigma_0|)$ in Eqs. 9–11 may grow large if $|\sigma_0|$ takes a large step above the normalized yield of 1. These terms should then cause the deviatoric stress to decay. However, if an explicit scheme is used, this decay may overshoot, and become unbounded. We therefore compute the plastic deformation as

$$D_{i,j}^n = \nu e^{-1/\chi_{i,j}^n} q((|\sigma_0|)_{i,j}^n)$$

and make use of

$$s_{i,j}^{n+1} = \frac{s_{i,j}^n + \Delta t (-u\{s_x\}_{i,j}^n - v\{s_y\}_{i,j}^n + 2\omega_{i,j}^n \tau_{i,j}^n + \mu([u_x - v_y]_{i,j}^n))}{1 + (\frac{s}{|\sigma_0|})_{i,j}^n \mu D_{i,j}^n \Delta t} \quad (20)$$

$$\tau_{i,j}^{n+1} = \frac{\tau_{i,j}^n + \Delta t (-u\{\tau_x\}_{i,j}^n - v\{\tau_y\}_{i,j}^n - 2\omega_{i,j}^n s_{i,j}^n + \mu([u_y + v_x]_{i,j}^n))}{1 + (\frac{\tau}{|\sigma_0|})_{i,j}^n \mu D_{i,j}^n \Delta t} \quad (21)$$

$$\chi_{i,j}^{n+1} = \frac{\chi_{i,j}^n + \Delta t (-u\{\chi_x\}_{i,j}^n - v\{\chi_y\}_{i,j}^n + \chi_\infty (|\sigma_0|)_{i,j}^n D_{i,j}^n \Delta t / c_0)}{1 + (|\sigma_0|)_{i,j}^n D_{i,j}^n \Delta t / c_0}. \quad (22)$$

With this method, larger plastic deformations just create increasingly larger denominators, meaning that the deviatoric stress more rapidly decays to zero.

In discretization method B, the spin ω is computed using centered differencing, but all remaining spatial derivatives are handled using the ENO scheme. This requires carrying out an eigenvector decomposition on the system of equations to rewrite it in terms of advective components. Since the effective temperature only depends on derivatives of itself, the non-trivial part of the decomposition depends on the velocity and stress only. Equations 6 to 10 can be rewritten as

$$\begin{pmatrix} u_t \\ v_t \\ p_t \\ s_t \\ \tau_t \end{pmatrix} = \begin{pmatrix} \kappa \nabla^2 u \\ \kappa \nabla^2 v \\ 0 \\ 2\omega\tau \\ -2\omega s \end{pmatrix} + \begin{pmatrix} -u & 0 & -\mu^{-1} & \mu^{-1} & 0 \\ 0 & -u & 0 & 0 & \mu^{-1} \\ -K & 0 & -u & 0 & 0 \\ \mu & 0 & 0 & -u & 0 \\ 0 & \mu & 0 & 0 & -u \end{pmatrix} \begin{pmatrix} u_x \\ v_x \\ p_x \\ s_x \\ \tau_x \end{pmatrix} + \begin{pmatrix} -v & 0 & 0 & 0 & \mu^{-1} \\ 0 & -v & -\mu^{-1} & -\mu^{-1} & 0 \\ 0 & -K & -v & 0 & 0 \\ 0 & -\mu & 0 & -v & 0 \\ \mu & 0 & 0 & 0 & -v \end{pmatrix} \begin{pmatrix} u_y \\ v_y \\ p_y \\ s_y \\ \tau_y \end{pmatrix} - \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{\mu s D}{|\sigma_0|} \\ \frac{\mu \tau D}{|\sigma_0|} \end{pmatrix}.$$

For the x component, the eigenvectors and corresponding eigenvalues are

$$\begin{aligned} \mathbf{A} &= (\mu p_x + K s_x, \mu v_x + \tau_x, -\mu v_x + \tau_x, \mu \eta u_x - p_x + s_x, -\mu \eta u_x - p_x + s_x) \\ \alpha &= (u, u - 1, u + 1, u - \eta, u + \eta) \end{aligned}$$

and for the y component, the eigenvectors and corresponding eigenvalues are

$$\begin{aligned} \mathbf{B} &= (-\mu p_y + K s_y, \mu u_y + \tau_y, -\mu u_y + \tau_y, -\mu \eta v_y + p_y + s_y, \mu \eta v_y + p_y + s_y) \\ \beta &= (v, v - 1, v + 1, v - \eta, v + \eta) \end{aligned}$$

where $\eta = \sqrt{1 + K/\mu}$. The scheme works by evaluating the quantities in the vectors \mathbf{A} and \mathbf{B} using the ENO scheme based upon the velocity given by the corresponding eigenvalue. Once these quantities are found, matrix inversions can be carried out to find expressions for the derivatives $u_x, v_x, \dots, s_y, \tau_y$, which can then be used with the same discretization given in Eqs. 17–22.

Computing the derivative at a point (i, j) requires referencing nearby gridpoints. Using the boundary conditions described in the following section, ghost values can be constructed for the orthogonally adjacent points if real values are not available. The ENO discretization may additionally make use of values that are two gridpoints away. In cases where real values are not available, the simulation makes use of a first order one-sided derivative.

4.2. Boundary conditions

When a discretized derivative operator in the integration schemes given above references a gridpoint that is not within the bar, a ghost value for that gridpoint is constructed that is consistent with the boundary conditions. Here we consider a point (i, j) within the bar, and a point (i', j') that is orthogonally adjacent and outside the bar.

We first consider the boundary conditions at the edge of the bar that is described by the level set function. Before beginning, we note that during the stage when boundary conditions are used, we have already carried out a normal extrapolation, and field values are available at (i', j') even though it is outside the level set. For the velocity and stress components, which are unconstrained at the boundary, these extrapolated values are used directly.

For the stress tensor, the normal–normal and normal–tangential components must vanish, so that $\sigma_{nn} = \sigma_{nt} = 0$, while the tangential–tangential component σ_{tt} is left unconstrained. This condition is effectively applied at the point P on the line from (i, j) to (i', j') where linear interpolation of the level set vanishes. The fraction of the distance along this line where this happens is given by

$$z = \frac{\phi_{i,j}}{\phi_{i,j} - \phi_{i',j'}}.$$

From this, a stress tensor σ_P at this point can be constructed as $\sigma_P = (1 - z)\sigma_{i,j} + z\sigma_{i',j'}$. Similarly, the derivatives ϕ_x and ϕ_y can be linearly interpolated to P , after which an outward-pointing normal vector $\mathbf{n}_P = (n_P^x, n_P^y)$ can be constructed. Let θ be the angle the \mathbf{n}_P makes with the x axis in a counter-clockwise sense. Then the stress tensor at P , in the frame rotated by this angle, is given by

$$\sigma'_P = \begin{pmatrix} -p + s \cos 2\theta + \tau \sin 2\theta & \tau \cos 2\theta - s \sin 2\theta \\ \tau \cos 2\theta - s \sin 2\theta & -p - s \cos 2\theta - \tau \sin 2\theta \end{pmatrix}.$$

From the given values of p , s , and τ , the tangential–tangential component of the stress tensor can be calculated as

$$\sigma_{tt} = -p - s \cos 2\theta + \tau \sin 2\theta,$$

making use of the identities $\sin 2\theta = 2n_P^x n_P^y$ and $\cos 2\theta = 1 - (n_P^y)^2$. From this, ghost values at P that satisfy the stress boundary conditions can be

computed using

$$\begin{aligned} p_g &= -\frac{1}{2}\sigma_{tt} \\ s_g &= -\frac{1}{2}\sigma_{tt} \cos 2\theta \\ \tau_g &= -\frac{1}{2}\sigma_{tt} \sin 2\theta. \end{aligned}$$

From here, ghost values at (i', j') can be constructed using linear interpolation. This is done using the values at point P and at point $(i'', j'') = (2i - i', 2j - j')$, so that

$$\sigma_{i',j'} = \frac{2\sigma_g - (1 - z)\sigma_{i'',j''}}{1 + z}.$$

This worked better than interpolating using P and (i, j) . In that case, the denominator in the interpolating formula is z , which can give unpredictable results when P is very close to (i, j) and z is small.

A similar procedure is used to apply boundary conditions at the end walls. We consider constructing a ghost value at a point (i', j) that is horizontally adjacent to point (i, j) , and in the same manner write (i'', j) to be the point on the opposite side of (i, j) so that $i'' = 2i - i'$. We write z to be the proportion of the distance from (i, j) to (i', j) where the wall is located. Some of the fields are left unconstrained, and some are fixed to a certain value. For a particular unconstrained field f , the ghost value is constructed as

$$f_{i',j} = 2f_{i,j} - f_{i'',j}.$$

If a field needs to be fixed to a certain value f_{fix} , then the ghost value is

$$f_{i',j} = \frac{2f_{\text{fix}} - (1 - z)f_{i'',j}}{1 + z}.$$

For similar reasons as discussed above, linear interpolation is done using point (i'', j) , to avoid stability problems when the wall position is very close to the (i, j) .

For all cases, u is constrained to $\pm u_{\text{wall}}$ and p , s , and χ are unconstrained. For the majority of simulation runs, sliding boundary conditions were used in which v is unconstrained, and τ is fixed to zero. However, fixed boundary conditions were also considered in which v is fixed to zero, and τ is unconstrained.

5. Numerical results

Figures 4–6 show snapshots of χ , p , and $|\sigma_0|$ respectively, from a typical run of the simulation using a grid size of 600×300 and initial conditions set by $\chi_0 = 0.074$, $n_w = 0.15$, and $n_d = 0.1$. The walls move at a speed $u_{\text{wall}} = 0.005$, the elastic parameters used are $K = 60$, $\mu = 30$, and the STZ plasticity parameters are set to $\nu = 10^8$, $c_0 = 1$, and $\chi_\infty = 0.015$. Discretization method A is used, with sliding boundary conditions at the end walls.

From $t = 0$ (a) to $t = 1$, a period of elastic deformation takes place. The effective temperature remains unchanged, but a negative pressure builds up within the bar, that is focused largely under the initial notch. Coupled with this is an increase in deviatoric stress, that begins to approach the plastic yield value of 1 at around $t = 1.5$ (b). A pair of shear bands of plastic deformation then nucleates from below the notch, that are coupled with an increase in the effective temperature. At $t = 3$ (c) these bands stretch halfway across the bar, and by $t = 3.5$ (d) they begin to reach the opposite side of the bar. By $t = 7.5$ (e) the bands are well developed, and small secondary reflected bands can be seen emanating from $x = \pm 0.3$.

The shear bands continue to grow wider and more plastic deformation takes place. By $t = 23$ (f) only a small triangle of undeformed material is left between two large shear bands, and the initial notch has become larger as the bar becomes thinner. At $t = 34.3$ (g), the bar is on the point of separating and all material near the center of the bar has become heated, up to a value of $\chi = 0.13$. The effective temperature does not reach the saturation value of $\chi_\infty = 0.15$. The bar splits and the two parts then move away with no further deformation, to reach the final state at $t = 50$ (h). Some remnants of deviatoric stress are visible, and the two parts undergo small elastic vibrations.

Examining the force exerted on the end walls during the stretching process is a useful method of comparing simulations with different parameters. The force on the walls at $x = \pm x_{\text{wall}}(t)$ can be calculated according to

$$F^\pm(t) = \int_{y_l^\pm(t)}^{y_u^\pm(t)} \sigma_{xx}(\pm x_{\text{wall}}(t), y, t) dy$$

where $y_l^\pm(t)$ and $y_u^\pm(t)$ are the lower and upper vertical positions where the bar meets each wall. Numerically, this can be computed using the trapezoidal

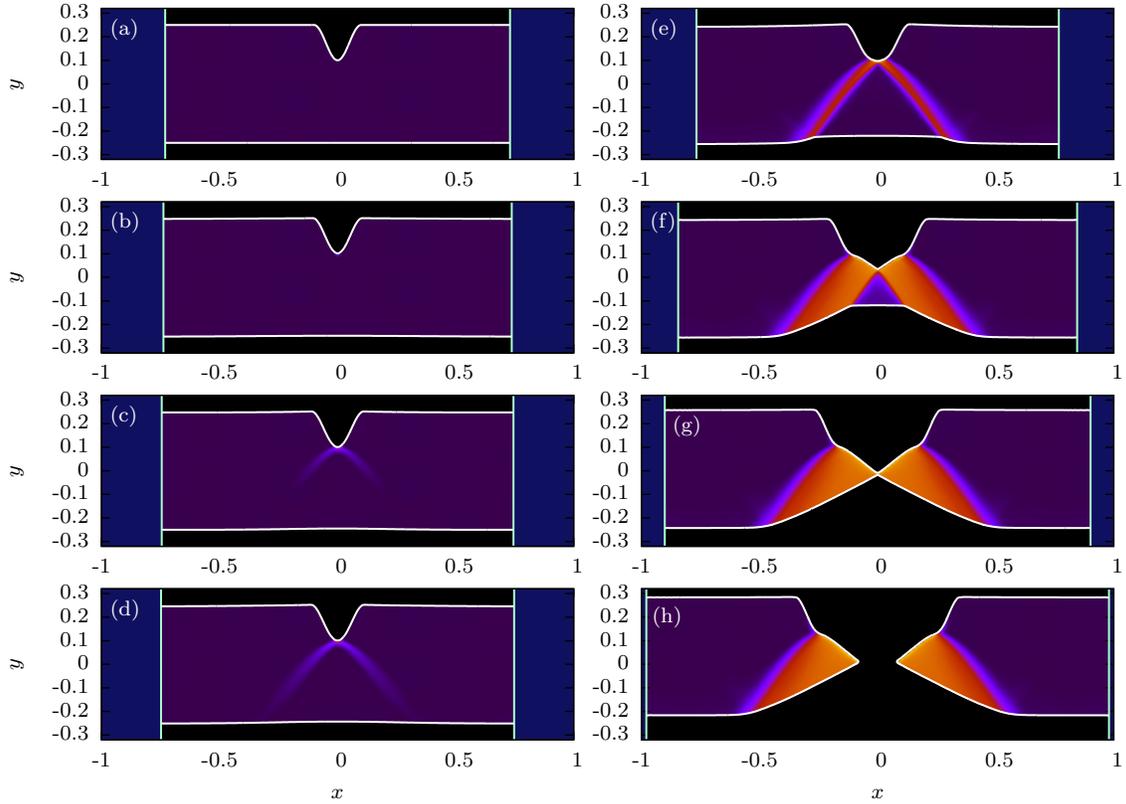


Figure 4: Snapshots of the effective temperature field χ for a stretching bar simulation using a 600×300 grid and initial conditions set by $\chi_0 = 0.074$, $n_w = 0.15$, and $n_d = 0.1$. The wall moves at a speed $u_{\text{wall}} = 0.005$, the elastic parameters used are $K = 60$, $\mu = 30$, and the STZ plasticity parameters are set to $\nu = 10^8$, $c_0 = 1$, and $\chi_\infty = 0.015$. Figures (a)–(h) correspond to $t = 0, 1.5, 3.0, 3.5, 7.5, 23.0, 34.3, 50.0$ respectively. See Fig. 7 for a key to the color gradient.

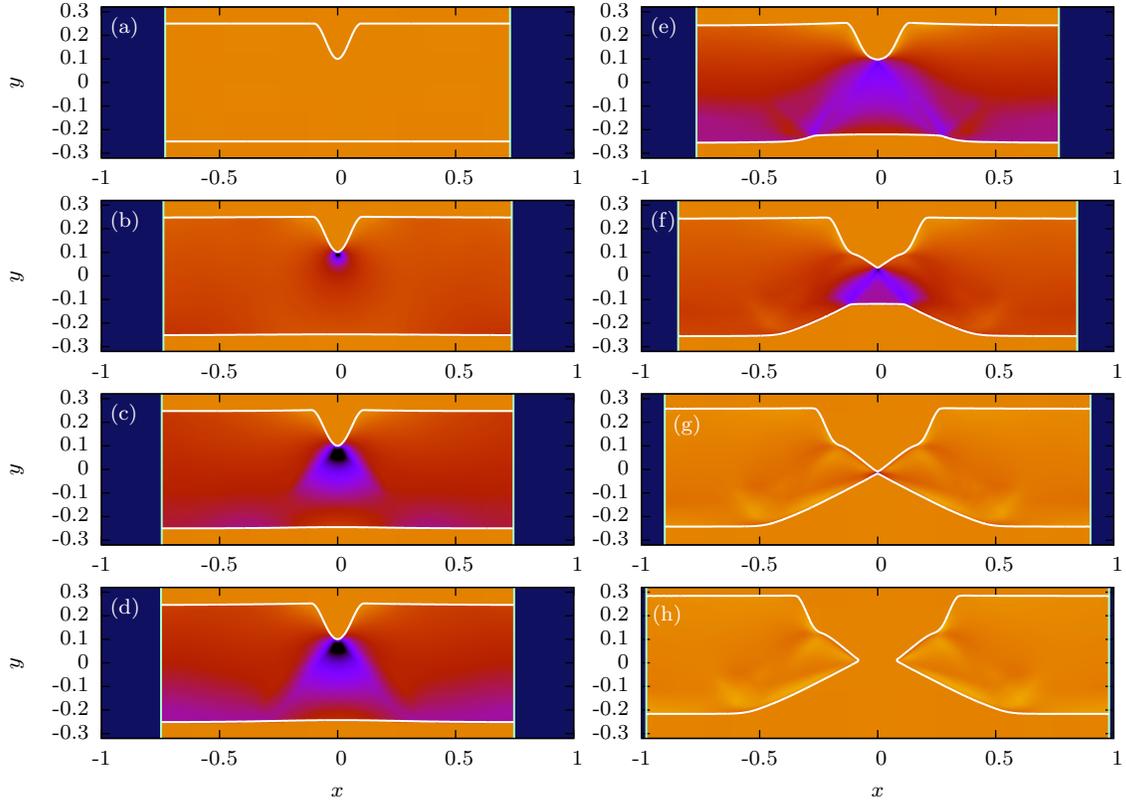


Figure 5: Snapshots of the pressure p for the same simulation shown in Fig. 4. Figures (a)–(h) correspond to $t = 0, 1.5, 3.0, 3.5, 7.5, 23.0, 34.3, 50.0$ respectively. See Fig. 7 for a key to the color gradient.

rule, and bilinear interpolations of the fields. For the graphs in this paper, the average force $F(t) = (F^-(t) + F^+(t))/2$ is then calculated, although due to symmetry $F^+(t)$ and $F^-(t)$ agree to a high degree of accuracy – the maximum differences seen are less than 5×10^{-3} , and occur after pinch-off when the details of the pinch-off point may introduce a small amount of asymmetry.

This physical process, whereby the bar is split into two by a pair of shear bands nucleating from the initial notch, happens for many of the simulations considered in this paper. In the following subsections, we analyze this behavior in more detail. We begin by examining the numerical accuracy in Subsec. 5.1, and then proceed to look in detail at the shear band structure in Subsec. 5.2. In Subsec. 5.3, the effect of the elastic and STZ parameters

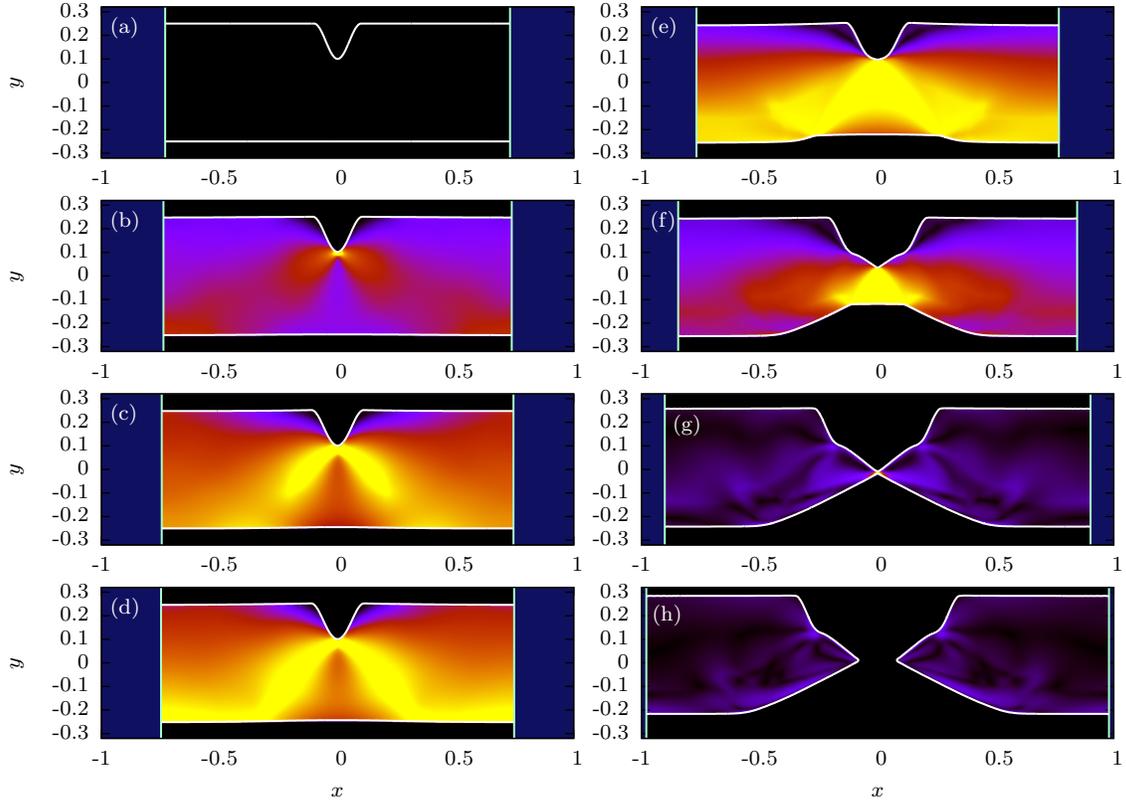


Figure 6: Snapshots of the magnitude of deviatoric stress $|\sigma_0|$ for the same simulation shown in Fig. 4. The material begins to yield plastically at $|\sigma_0| = 1$. Figures (a)–(h) correspond to $t = 0, 1.5, 3.0, 3.5, 7.5, 23.0, 34.3, 50.0$ respectively. See Fig. 7 for a key to the color gradient.

is investigated, followed by an examination of the role of viscosity and small perturbations in Subsec. 5.4. Alternative integration methods and models are considered in Subsec. 5.5.

5.1. Grid size and numerical convergence

The PDE system under consideration is strongly nonlinear, and we are not able to construct an analytic solution to the current problem that would allow us to directly examine numerical convergence. Instead, we consider how the results differ when the grid is refined and the timestep is altered. To correctly model the viscous term in the PDE system, the timestep is chosen to satisfy $\Delta t = \lambda h^2$. For the rest of this paper, the timestep is set using

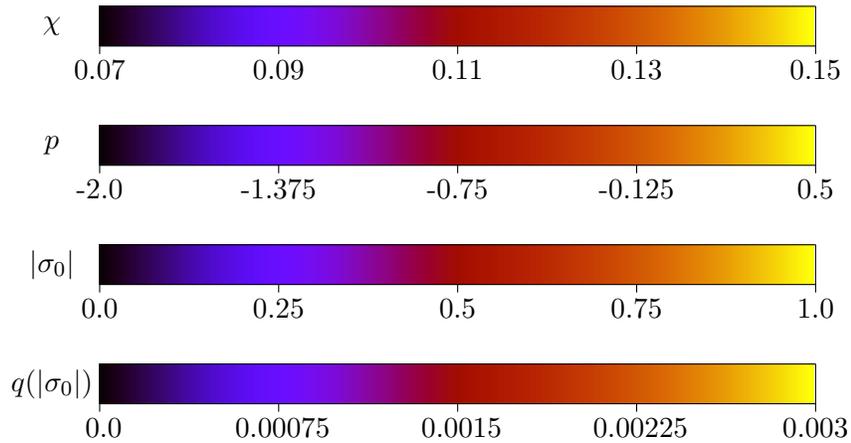


Figure 7: Keys to the color gradients used in the figures, for effective temperature χ , pressure p , the magnitude of deviatoric stress $|\sigma_0|$, and the plastic deformation function $q(|\sigma_0|)$.

Grid size	Timestep	Pinch-off time	Computation time
80×160	3.164×10^{-4}	31.12	13.5 min
100×200	2.020×10^{-4}	30.89	29.9 min
150×300	8.948×10^{-5}	32.61	1.59 hr
200×400	5.025×10^{-5}	33.32	6.08 hr
300×600	2.230×10^{-5}	34.63	46.9 hr

Table 1: The grid sizes considered to examine the numerical convergence. For each grid size, the simulation timestep and the of value t for which the bar pinches off into two are reported. The final column lists the wall clock times of the simulation runs up to $t = 50$ on a single thread of a Mac Pro system with 5 Gb of memory and a dual-core 2.66 GHz Intel Xeon processor, that scale according to the fourth power of the grid size, as expected from timestep restrictions.

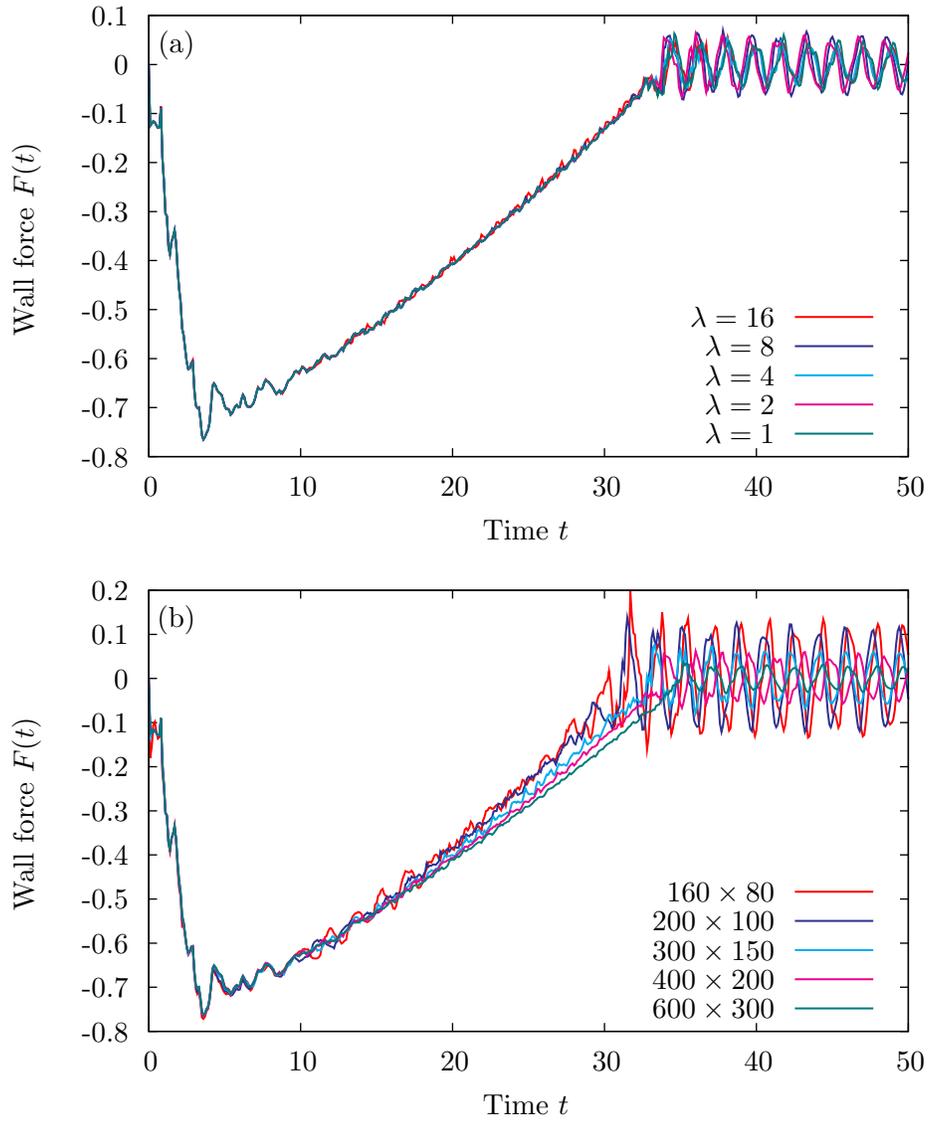


Figure 8: Plots of the force exerted on the end wall, using the same simulation parameters given in Fig. 4, for five different timesteps using a grid size of 400×200 , and (b) five different grid sizes using $\lambda = 2$.

$\lambda = 2$, but here we consider $\lambda = 1, 2, 4, 8, 16$ for a grid size of 200×400 using the same simulation parameters as presented above. For the values of λ considered, the size of the timestep plays a very limited role on the results, as can be seen in Fig. 8(a), where the forces exerted on the wall are almost identical. By using $\lambda = 2$ we are well below the limit where temporal errors could be important.

Errors associated with the spatial discretization are more significant, and Tab. 1 shows the details of a sequence of runs that were performed with different grid sizes. As would be expected, the simulation time scales approximately with the fourth power of the number of horizontal gridpoints, with the largest grid size taking approximately two days to compute. Plots of the wall force are shown in Fig. 8(b). While the curves have different random fluctuations, they all exhibit the same trend, and the results for the two finest grids are very similar.

One of the largest possible sources of error is the due to the boundary conditions applied at the level set boundary, particularly in regions of high deformation. Figure 9 shows the initial position of the boundary, and its position at two later times, for the sequence of grid sizes that were considered. In this plot, it was necessary to account for vertical drift. The sliding boundary conditions do not constrain the vertical velocity, and thus small random variations can cause it to move position. To correct for this, the point y_w where the top of the bar intersects the wall is found, and then a drift $\Delta y = y_w - 0.25$ is computed. The graph shows that we have good convergence as the grid size is increased, with the curves being very close at $t = 150$ and still similar at $t = 300$. The main point where the curves differ is at the notch tip, which becomes progressively sharper as the resolution is increased. It is not surprising that this is the most difficult point to simulate, since the previous simulation snapshots suggest that the shear bands emanate from here. It is also a difficult point at which to apply the field extrapolation, since outward-pointing normals around the notch will cross. This does not cause problems for the extrapolation routine itself, but may mean that an extrapolated value used to construct a derivative may be less consistent.

5.2. Shear band structure

Here, we examine more closely the formation and structure of the shear bands seen in Figs. 4–6, by looking at cross-sections of the simulated fields. Figure 10(a) shows the effective temperature in the bar at $t = 3.0$, for six

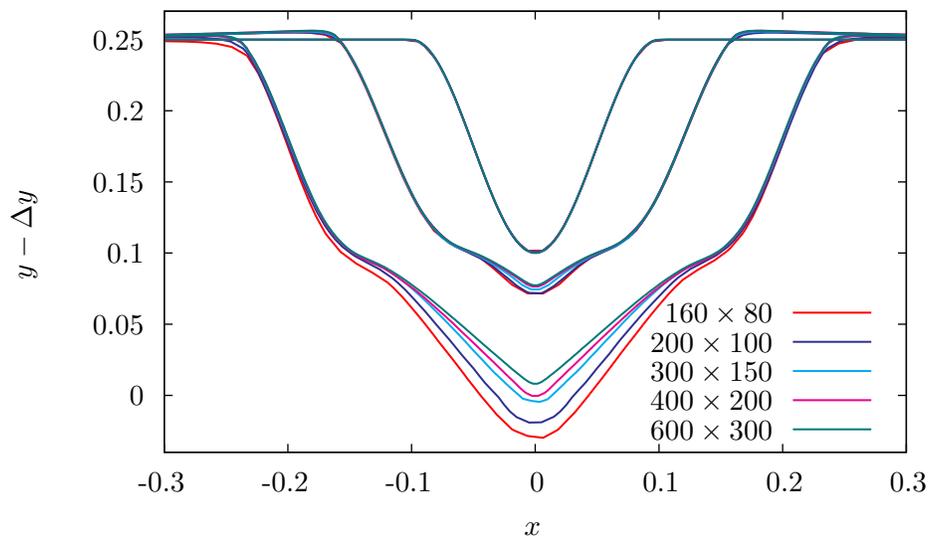


Figure 9: Plots of the deformation of the initial notch for five different grid sizes. The highest family of curves correspond to the initial notch position, at $t = 0$. The middle and bottom families correspond to the notch position at $t = 150$ and $t = 300$ respectively. To counteract bar drift due to the sliding boundary conditions, a displacement Δy is added to the curve positions.

different values of y . At this time, the shear band has just formed. The increases seen in χ are small, and are approximately Gaussian and symmetric. The increases are progressively larger for higher values of y that are closer to the initial notch. At the later time of $t = 12.5$, shown in Fig. 10(b), the shear band is much more developed, with χ approaching values of 0.13. By this point, the structure of the shear band is much more even throughout the bar, with the curves from $y = 0.05$ through to $y = -0.2$ looking very similar in form.

It can also be seen that the shear band has become asymmetric, rising faster on the left side, and having a long tail on the right side. To determine the cause of this, it is useful to examine a close-up of a single shear band, shown at $t = 12.5$ in Fig. 11. Figure 11(b) shows a plot of $q(|\sigma_0|)$, the function that controls the plastic deformation, where it can be seen that all of the plastic deformation occurs on the left side of the shear band. This behavior suggests that the region of increased χ is progressively widened from the left side as the bar is stretched. It is worth noting that the values of $q(|\sigma_0|)$ plotted are very small. This is typical for the entire stretching process, where the maximum value of $|\sigma_0|$ is only just above 1, pointing to a very fast relaxation of stress by plastic deformation for these particular simulation parameters – this is explored in more detail in the following section.

5.3. Physical effect of the elastic and STZ parameters

We now turn attention to how the various parameters in the PDE system in Eqs. 6–11 affect the stretching process. Figure 12 shows a sequence of snapshots taken at $t = 15$ for three different values of ν , the parameter that controls the rate of plastic deformation. The plots look very similar for $\nu = 10^8$ and $\nu = 10^9$, and although not shown here, a simulation with $\nu = 10^{10}$ also looks identical. Despite a factor of 100 difference in ν , these simulations give similar results because the timescale on which the plastic response occurs happens very rapidly, meaning that $|\sigma_0|$ can never achieve values significantly higher than 1, since any plastic behavior happens immediately. This can be confirmed by looking at the amount by which $|\sigma_0|$ exceeds 1. For $t = 12.5$ with $\nu = 10^8$, a typical maximum value of $q(|\sigma_0|)$ is 6×10^{-4} , while for $\nu = 10^{10}$, a typical maximum value is 6×10^{-6} . Thus, the overall plastic response, which is proportional to $\nu q(|\sigma_0|)$, is very similar for the two simulations. Figure 13(a) shows the wall force for these three simulations, and as would be expected, the $F(t)$ curves look almost identical.

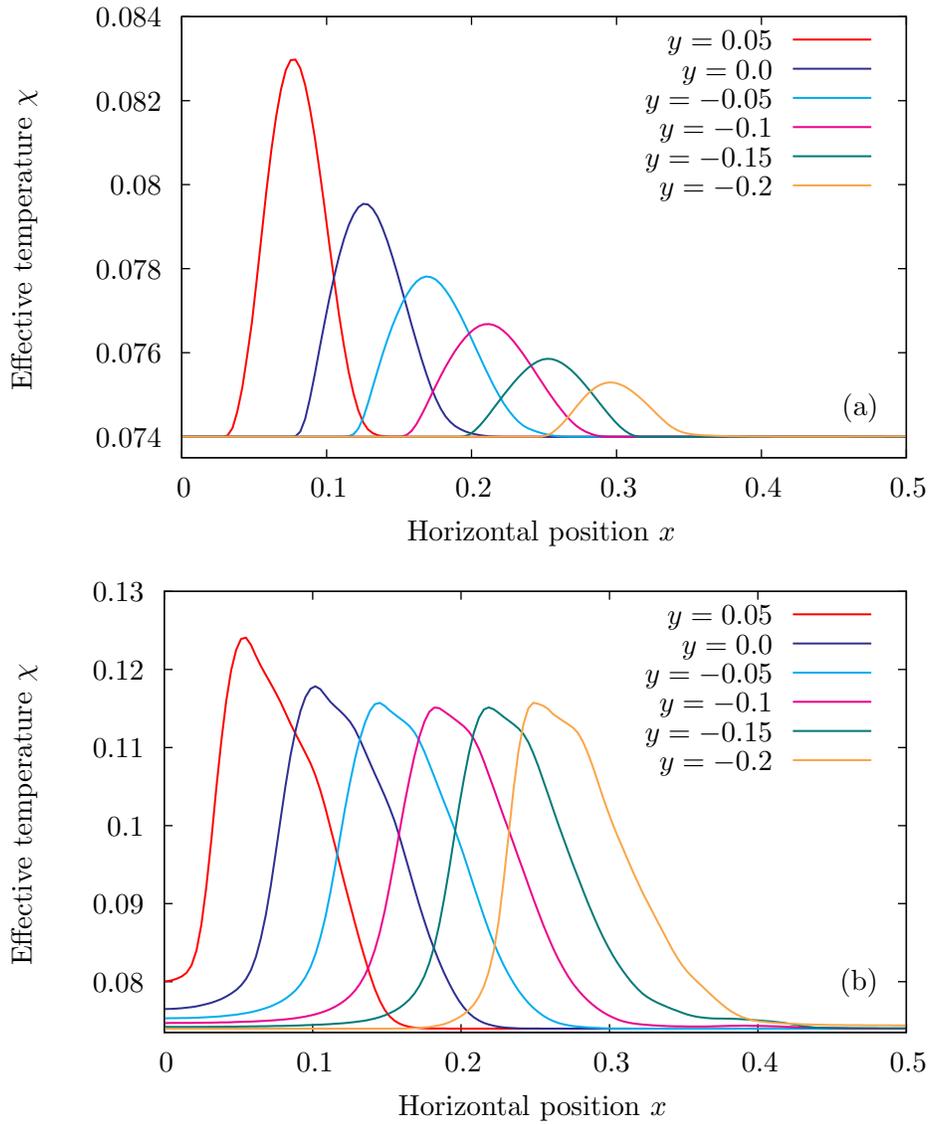


Figure 10: Cross-section plots of the effective temperature for several different values of y , at (a) $t = 3.0$ and (b) $t = 12.5$, for the same simulation run presented in Fig. 4.

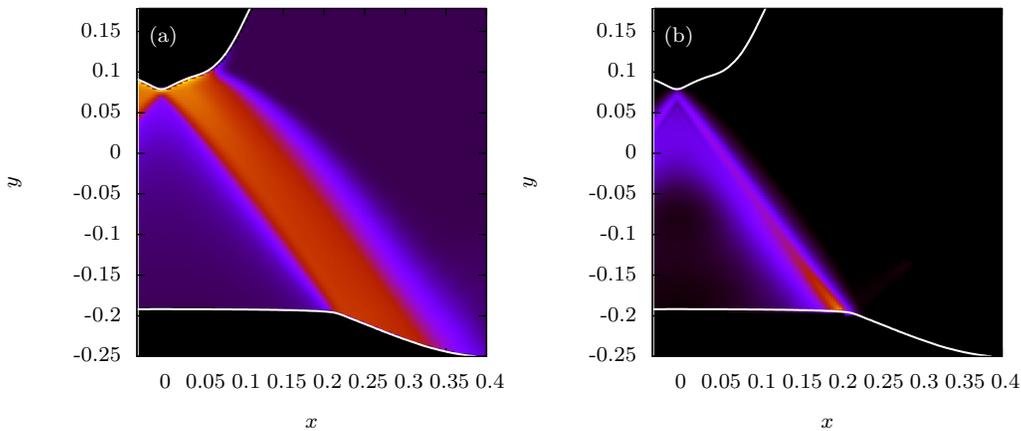


Figure 11: Close-up of the shear band at $t = 12.5$ after it is well-developed, showing (a) the effective temperature, (b) the plastic deformation function $q(|\sigma_0|) = (|\sigma_0| - 1)^2/|\sigma_0|$ for regions where $|\sigma_0| > 1$. See Fig. 7 for a key to the color gradients.

Also shown in Fig. 12 is a simulation using $\nu = 10^7$, where some differences become apparent. While the shear bands form in a similar manner, but the angle that the notch makes as it grows is shallower. This behavior seems to be primarily due to the plastic response being slower in the vicinity of the notch tip. The bar takes longer to split for this case, and the wall forces in Fig. 13(a) decrease at a slower rate.

Fig. 14 shows a sequence of three snapshots for $\nu = 10^6$, where a qualitatively different behavior can be seen. Here, diagonal shear bands of higher χ form, but yielding at the notch tip does not occur immediately, causing two notches to be visible at $t = 10.0$ (a). At $t = 17.0$ (b), a second pair of shear bands appears, which then begin to widen in a similar manner to the other simulations, becoming large by $t = 24.0$ (c).

Two simulations were also carried out to investigate the effect of the elastic moduli. Figure 15 shows a snapshot at $t = 24.0$ of a simulation where $K = 100$ and $\mu = 50$. Overall, there is little qualitative difference with the runs using $K = 60$, although the wall force, shown in Fig. 13(b), more rapidly reaches its peak value. For lower elastic moduli of $K = 20$ and $\mu = 10$, a sequence of snapshots shown in Fig. 16 exhibit different behavior. The initial notch, shown at $t = 6.5$ (a) undergoes a significant amount of elastic deformation before plastic behavior occurs. Once the plastic deformation occurs, the angle of the notch, shown at $t = 24.0$ (b) is sharper, meaning

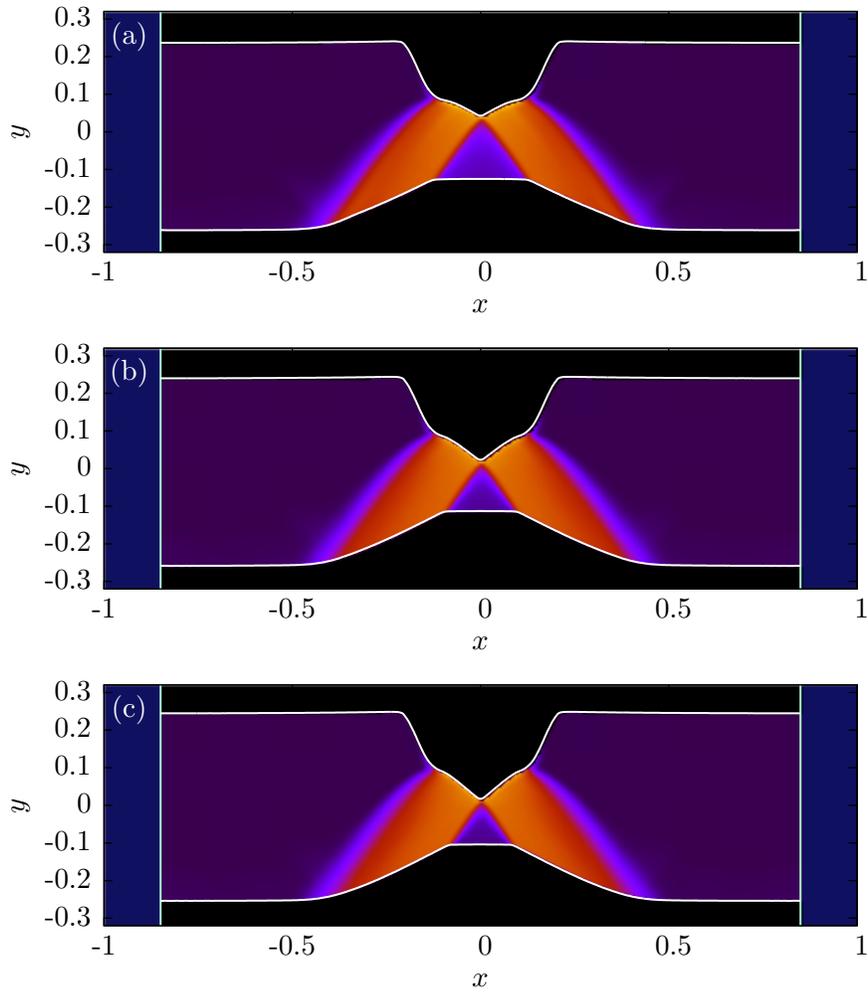


Figure 12: Simulation snapshots of the effective temperature χ at $t = 24.0$ for (a) $\nu = 10^7$, (b) $\nu = 10^8$, and (c) $\nu = 10^9$. See Fig. 7 for a key to the color gradient.

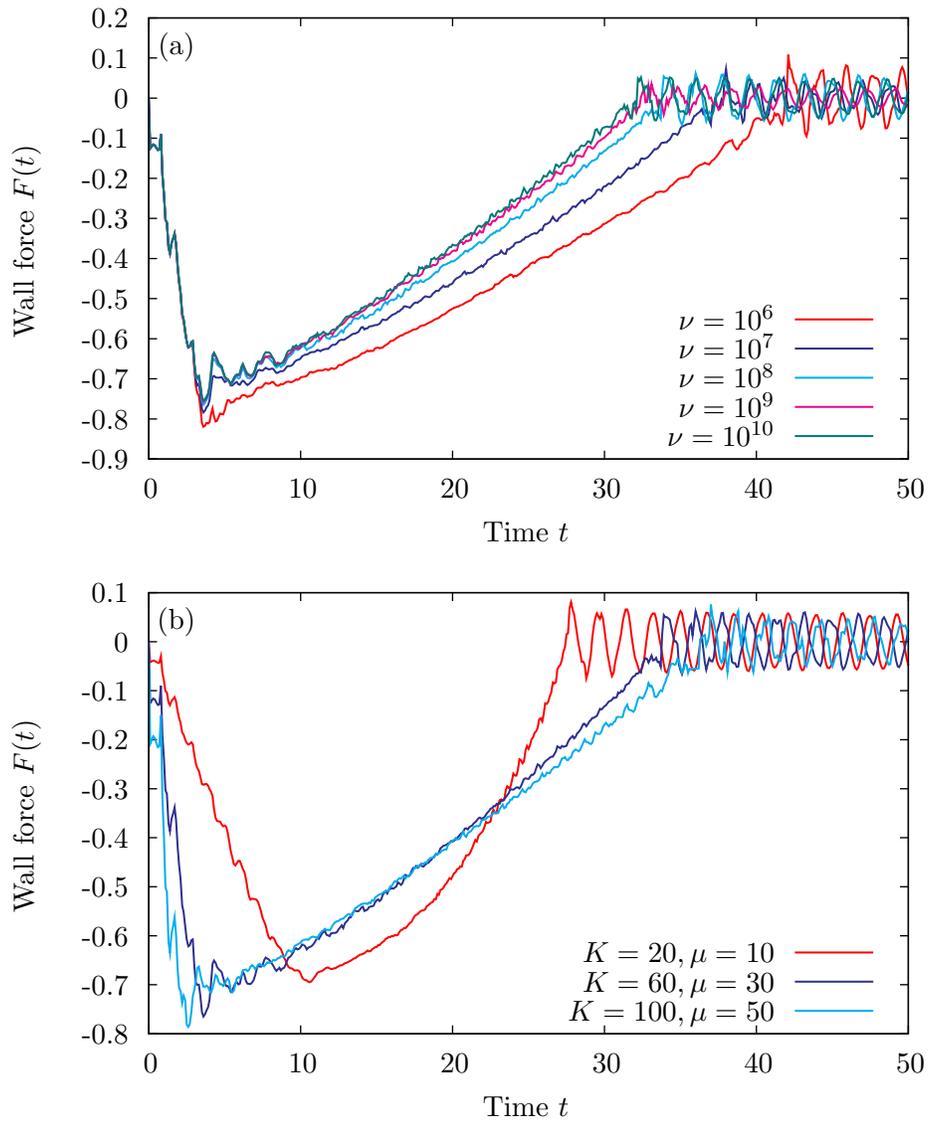


Figure 13: Plots of the wall force for (a) several different values of ν , and (b) several different elastic moduli.

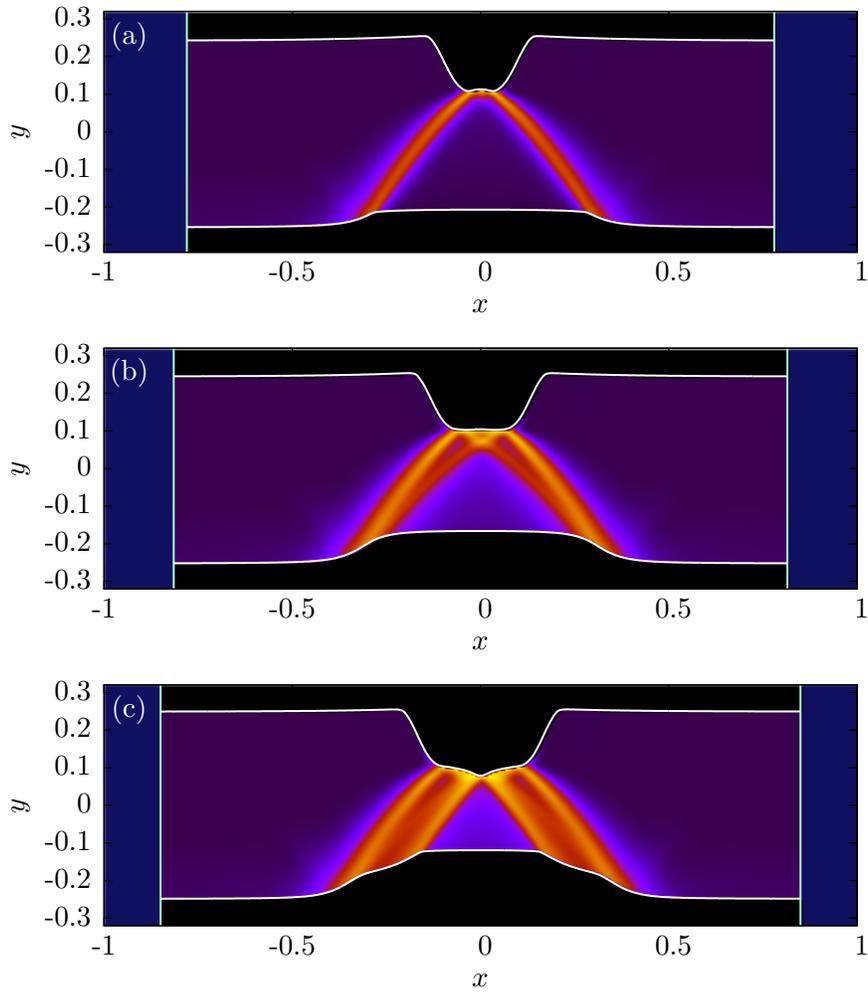


Figure 14: Simulation snapshots of the effective temperature χ for $\nu = 10^6$ at (a) $t = 10.0$, (b) $t = 17.0$, and (c) $t = 24.0$. See Fig. 7 for a key to the color gradient.

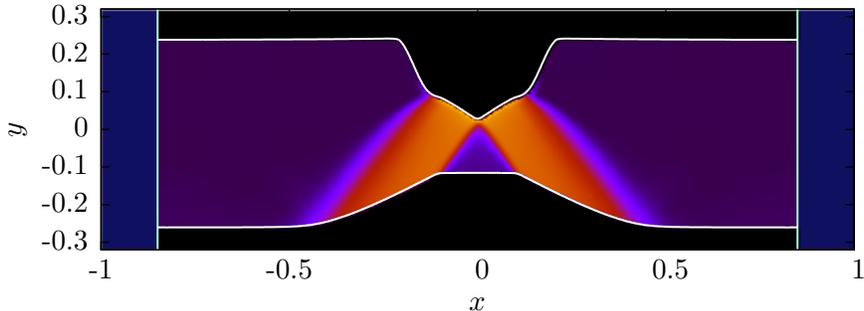


Figure 15: Simulation snapshot of the effective temperature χ at $t = 24.0$ for higher elastic moduli of $K = 100, \mu = 50$. See Fig. 7 for a key to the color gradient.

that the two remaining halves (c) have flatter ends. The forces on the wall for this simulation run, shown in Fig. 13(b), behave very differently to the other simulations, and take a long time to reach a maximum value.

5.4. Viscosity and small perturbations to χ

The viscous term in the PDE system is needed to damp out elastic sound waves. If no viscosity is used with the centered difference scheme employed here, then the simulation rapidly breaks down, with elastic waves growing in magnitude. While it is physically reasonable to include a viscous damping, with sound waves in real materials being attenuated, the value of $\kappa = 0.002$ used in the previous sections was chosen to be small, so that viscous effects are relatively less important when compared the plastic deformation. Here, we examine whether the value of κ plays a significant role in the deformation process. Figure 17 shows simulation snapshots at $t = 24.0$ for $\kappa = 0.001, 0.004$ keeping all other simulation parameters the same. While the shear bands are similar in width, the angle of the notch varies. This suggests that the most significant effect of the viscosity is to alter the dynamics near the notch, which appears reasonable since the velocity fields vary the most around this point. As shown in Fig. 18(a), higher viscosities lead to a slightly lower decay in the wall force.

The initial value of the effective temperature can also play a role in the dynamics. Since the plastic deformation has a factor of $e^{-1/\chi}$, even a relatively small change in this field can have a significant effect on the material behavior. Figure 19(a) shows a snapshot of a simulation using an initial temperature of 0.068, with all other simulation parameters kept the same. We

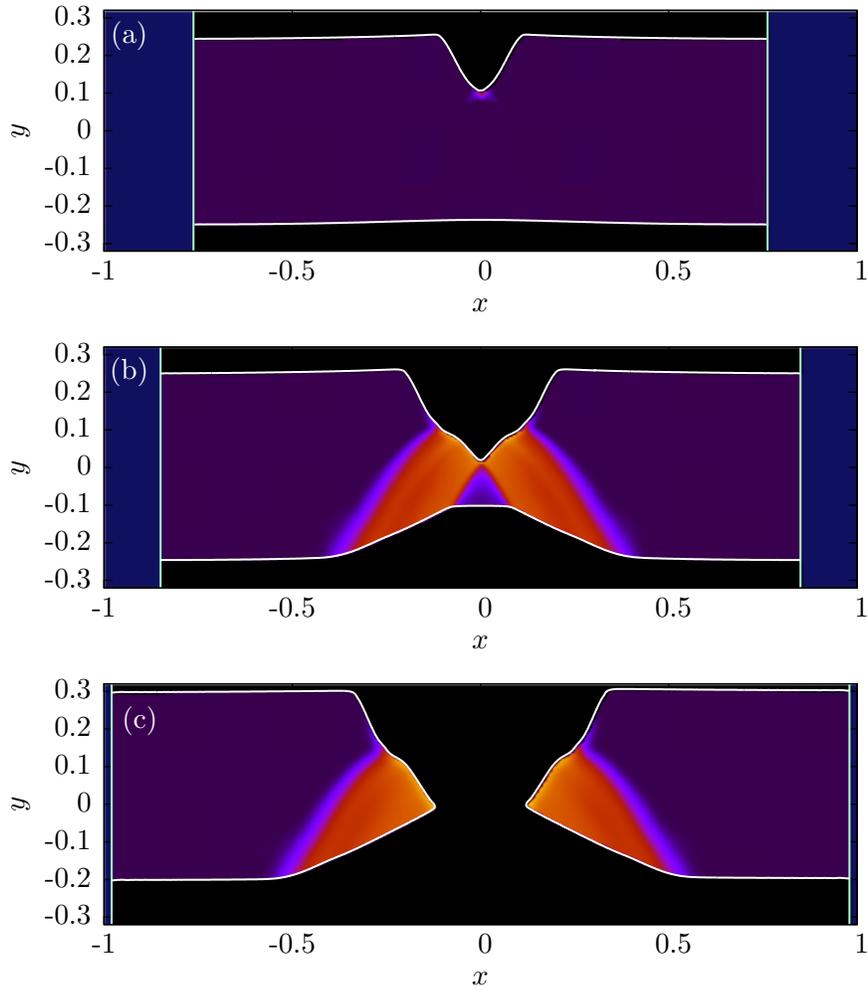


Figure 16: Simulation snapshots of the effective temperature χ for lower elastic moduli of $K = 20, \mu = 10$ at (a) $t = 6.5$, (b) $t = 24.0$, and (c) $t = 50.0$. See Fig. 7 for a key to the color gradient.

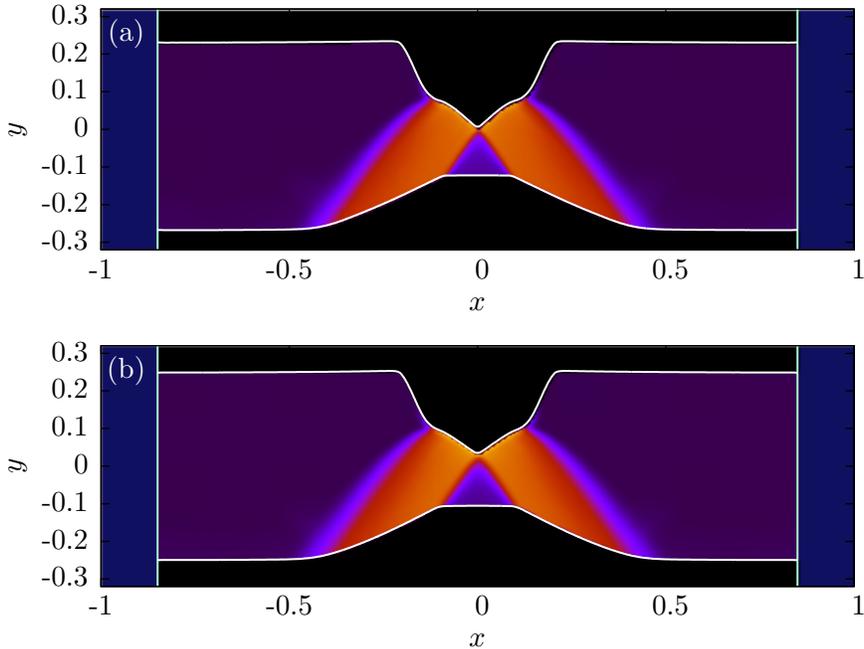


Figure 17: Snapshots of the effective temperature χ at $t = 24.0$, for the different values of viscosity of (a) $\kappa = 0.001$ and (b) $\kappa = 0.004$. See Fig. 7 for a key to the color gradient.

see that the plastic deformation has happened slightly more slowly than in Fig. 12(b), although there are no large qualitative changes. Correspondingly, the wall force for this simulation, shown in Fig. 18(b), takes a slightly longer time to decay.

It is interesting to ask whether small variations in the effective temperature would have a significant impact on the stretching process. To investigate this, a grid $U_{i,j}$ was initialized where each point is chosen to be uniformly distributed from 0 to 1. From here, a smoothed random field $R_{i,j}$ was computed according to

$$R_{i,j} = \frac{\sum_{k=0}^m \sum_{l=0}^n e^{-0.3((i-k)^2+(j-l)^2)} U_{k,l}}{\sum_{k=0}^m \sum_{l=0}^n e^{-0.3((i-k)^2+(j-l)^2)}}.$$

For the 400×200 grid used here, this corresponds to smoothing on a scale of $h/\sqrt{2} \times 0.3 = 0.0118$. With this, the effective temperature is initialized so that $\chi_{i,j} = 0.5 + 0.4R_{i,j}$, which gives an initial average of $0.5 + 0.4/2 = 0.7$. Figure 19(b) shows a snapshot of this simulation at $t = 24.0$. While there are some small random variations in the shear bands, there is no large breakage

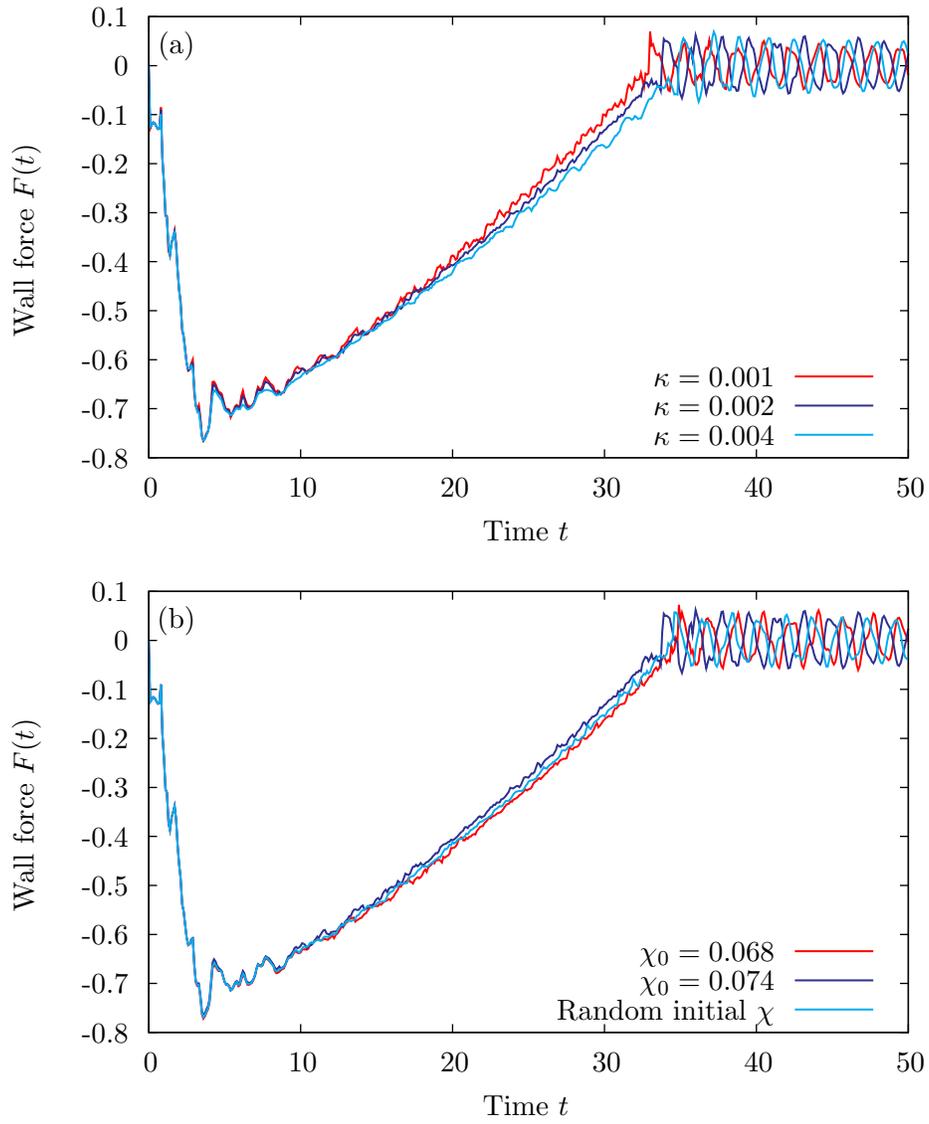


Figure 18: Plots of the wall force for simulations using (a) varying values of viscosity and (b) using varying initial effective temperatures.

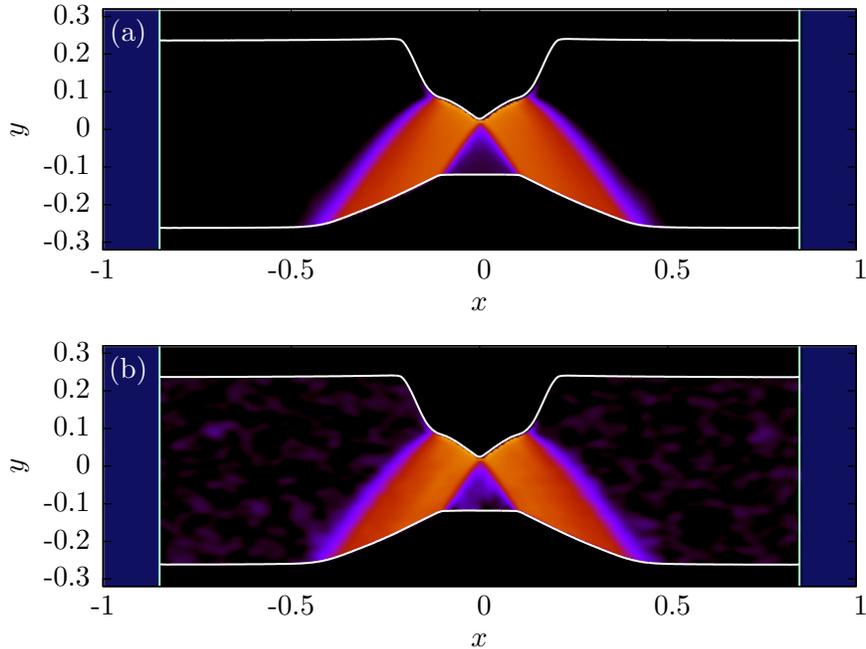


Figure 19: Snapshots of the effective temperature χ at $t = 24.0$ for simulations using different initial conditions for χ . In (a), the effective temperature is initialized to 0.068. In (b), the effective temperature is randomly initialized according to the process described in the text, with a mean value of 0.07. See Fig. 7 for a key to the color gradient.

of symmetry, and the position of the edges of the bar is almost identical to the unperturbed case shown in Fig. 12(b).

5.5. Alternative methods and models

The fixed boundary conditions were also considered, and a simulation snapshot is shown in Fig. 20 at $t = 24.0$. These boundary conditions have the advantage that the y coordinate is constrained, so that the bar does not drift vertically as discussed in Subsec. 5.1. However, there is an additional complication that the stress field in an elastic bar can form weak singularities at corners [21, 22]. Here, because of viscosity and plastic yield, we do not encounter singular fields in the simulation, but additional increases in effective temperature are visible at the corner where the bar is fixed to the wall. The forces exerted on the wall in this simulation are almost identical to those using the sliding boundary.

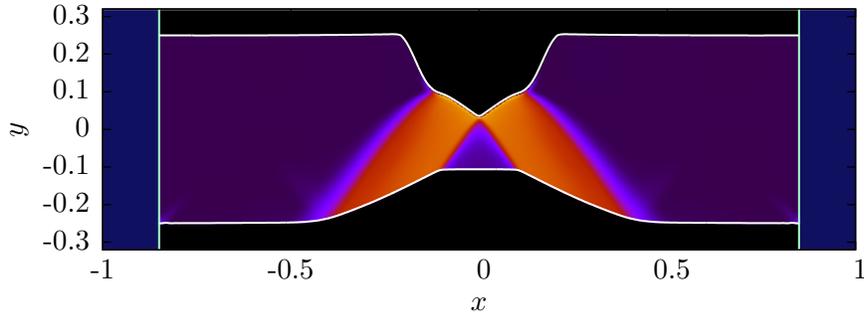


Figure 20: Snapshot of the effective temperature χ at $t = 24.0$, for a simulation making use of the fixed boundary conditions at the wall. Small increases in χ are visible at the interface between the wall and the lower edge of the bar. See Fig. 7 for a key to the color gradient.

Figure 21(a) shows a simulation snapshot using discretization method B. This technique has the advantage of being able to tolerate more rapid changes in field values, and as shown in Fig. 21(b) it can be run without requiring any viscosity, using $\kappa = 0$, which results in the bands in effective temperature being slightly sharper. This simulation technique has the disadvantage of being anisotropic, so that rapid changes in field values are handled well in the orthogonal directions, but not in the diagonal directions. As an example, consider the snapshot of the pressure shown in Fig. 21(c), which shows rapid horizontal changes that are much stronger than with centered differencing: smaller, rapid vertical changes are also visible below the initial notch. This causes a different behavior near the initial notch, leading to a small disturbance in the level set boundary at the point where the shear band starts. This feature does not appear to be significant, and a refinement study using the same grid sizes employed in Fig. 9 shows that the results converge, with this disturbance becoming progressively less visible. Switching between orthogonally and diagonally aligned stencils between timesteps may circumvent this issue.

6. Conclusion

We have developed a numerical method for studying elastoplastic models from the STZ theory, that makes use of the level set method and rapid field extrapolation. The scheme has been shown to converge well as the timestep and grid size are decreased. Using this method, ductile fracture

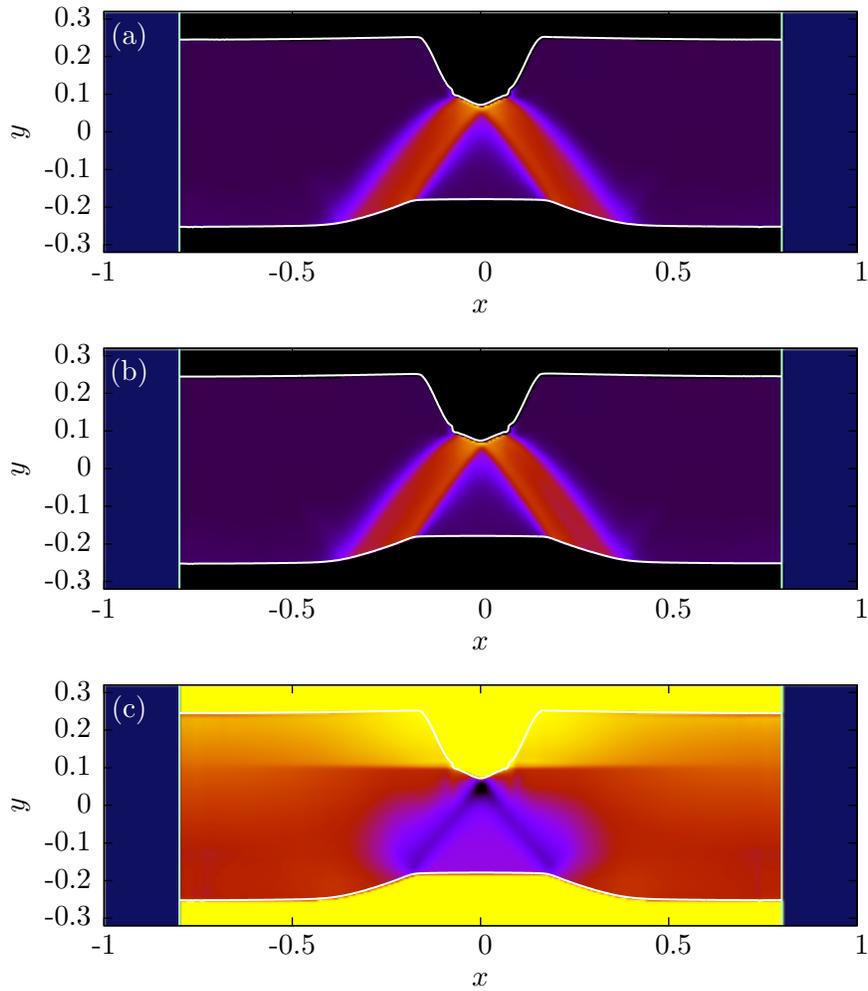


Figure 21: Simulation snapshots making use of discretization method B, at $t = 14.0$. (a) A snapshot of χ for a simulation using a viscosity of $\kappa = 0.002$. (b) A snapshot of χ for a simulation with no viscosity. (c) A snapshot of p for $\kappa = 0.002$, highlighting rapid horizontal and vertical changes in the field. See Fig. 7 for a key to the color gradients.

has been studied in a bar being stretched. The numerical method we use has allowed us to simulate the entire failure process, whereby shear bands of higher effective temperature form that gradually widen, until the bar splits into two. We have investigated the effect of the STZ parameters, the elastic constants, and the initial conditions on the stretching process, and we have considered the effects of using alternative discretization methods.

The simulation method presented can be used as a basis for further study. It can be readily adapted to a wide variety of geometries, and we have written a preliminary version in polar coordinates for computing STZ elastoplasticity problems involving circular and elliptical holes in an infinite medium, for which a variety of exact and asymptotic solutions have been examined [21, 23]. We also aim to study some newer versions of the STZ theory [24, 25, 26], and hope to incorporate orientational effects when a non-isotropic STZ density is employed.

7. Nomenclature

Physical quantities

\tilde{u} – Horizontal velocity (L/T)

\tilde{v} – Vertical velocity (L/T)

\tilde{p} – Pressure (M/T²)

$\tilde{s}, \tilde{\tau}$ – Shear stresses (M/T²)

$\tilde{\rho}_0$ – Density (M/L²)

$\tilde{\kappa}$ – Viscosity (M/T)

\tilde{K} – Bulk modulus (M/T²)

$\tilde{\mu}$ – Shear modulus (M/T²)

\tilde{D} – Plastic deformation (T⁻¹)

ρ_0 – Rescaled density (T²/L²)

c_s – Transverse wave speed (L/T)

L – Characteristic system size (L)

$\tilde{\sigma}$ – Stress tensor (M/T²)

$\tilde{\sigma}_0$ – Deviatoric stress tensor (M/T²)

s_Y – Yield stress (M/T²)

$\tilde{\omega}$ – Angular velocity (T⁻¹)

ϵ_0, c_0 – STZ model parameters

\tilde{a} – STZ model parameter (T²/M)

τ_0 – Molecular rearrangement timescale (T)

Rescaled quantities

u – Dimensionless horizontal velocity

v – Dimensionless vertical velocity

p – Dimensionless pressure

s, τ – Dimensionless shear stresses

κ – Dimensionless viscosity

K – Dimensionless bulk modulus

μ – Dimensionless shear modulus

D – Dimensionless plastic deformation

σ – Dimensionless stress tensor

σ_0 – Dimensionless deviatoric stress tensor

ω – Dimensionless angular velocity

a, ν – Dimensionless STZ model parameters

χ – Dimensionless effective temperature

χ_0 – Dimensionless initial effective temperature

χ_∞ – Dimensionless limiting effective temperature

Additional quantities

ψ – Initial signed distance function

ϕ – Level set function

x_{wall} – Dimensionless wall position

u_{wall} – Dimensionless horizontal wall speed

n_w, n_h – Dimensionless notch sizes

h – Dimensionless grid spacing

Bibliography

- [1] J. S. Langer, Shear-transformation-zone theory of plastic deformation near the glass transition, *Phys. Rev. E* 77 (2) (2008) 021502. doi:10.1103/PhysRevE.77.021502.
- [2] M. H. Cohen, D. Turnbull, Molecular transport in liquids and glasses, *J. Chem. Phys.* 31 (5) (1959) 1164–1169. doi:10.1063/1.1730566.
- [3] F. Spaepen, A microscopic mechanism for steady state inhomogeneous flow in metallic glasses, *Acta Metallurgica* 25 (4) (1977) 407–415. doi:10.1016/0001-6160(77)90232-2.
- [4] D. Turnbull, M. H. Cohen, On the free-volume model of the liquid–glass transition, *J. Chem. Phys.* 52 (6) (1970) 3038–3041. doi:10.1063/1.1673434.
- [5] M. L. Falk, J. S. Langer, Dynamics of viscoplastic deformation in amorphous solids, *Phys. Rev. E* 57 (6) (1998) 7192–7205. doi:10.1103/PhysRevE.57.7192.
- [6] J. S. Langer, M. L. Manning, Steady-state, effective-temperature dynamics in a glassy material, *Phys. Rev. E* 76 (5) (2007) 056107. doi:10.1103/PhysRevE.76.056107.
- [7] E. Bouchbinder, J. S. Langer, I. Procaccia, Athermal shear-transformation-zone theory of amorphous plastic deformation. I. Basic principles, *Phys. Rev. E* 75 (3) (2007) 036107. doi:10.1103/PhysRevE.75.036107.

- [8] E. Bouchbinder, J. S. Langer, I. Procaccia, Athermal shear-transformation-zone theory of amorphous plastic deformation. II. Analysis of simulated amorphous silicon, *Phys. Rev. E* 75 (3) (2007) 036108. doi:10.1103/PhysRevE.75.036108.
- [9] M. L. Manning, J. S. Langer, J. M. Carlson, Strain localization in a shear transformation zone model for amorphous solids, *Phys. Rev. E* 76 (5) (2007) 056106. doi:10.1103/PhysRevE.76.056106.
- [10] M. L. Manning, E. G. Daub, J. S. Langer, J. M. Carlson, Rate-dependent shear bands in a shear-transformation-zone model of amorphous solids, *Phys. Rev. E* 79 (1) (2009) 016110. doi:10.1103/PhysRevE.79.016110.
- [11] L. Pechenik, Dynamics of shear-transformation zones in amorphous plasticity: Nonlinear theory at low temperatures, *Phys. Rev. E* 72 (2) (2005) 021507. doi:10.1103/PhysRevE.72.021507.
- [12] L. O. Eastgate, J. S. Langer, L. Pechenik, Dynamics of large-scale plastic deformation and the necking instability in amorphous solids, *Phys. Rev. Lett.* 90 (4) (2003) 045506. doi:10.1103/PhysRevLett.90.045506.
- [13] J. A. Sethian, *Level Set Methods and Fast Marching Methods*, Cambridge University Press, 1996.
- [14] S. J. Osher, R. P. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces*, Springer-Verlag, New York, 2003.
- [15] T. D. Aslam, A partial differential equation approach to multidimensional extrapolation, *J. Comp. Phys.* 193 (1) (2004) 349–355. doi:10.1016/j.jcp.2003.08.001.
- [16] M. Sussman, P. Smereka, S. J. Osher, A level set method for computing solutions to incompressible two-phase flow, *J. Comp. Phys.* 114 (1) (1994) 146–159. doi:10.1006/jcph.1994.1155.
- [17] G. Russo, P. Smereka, A remark on computing distance functions, *J. Comp. Phys* 163 (1) (2000) 51–67. doi:10.1006/jcph.2000.6553.
- [18] A. du Chéné, C. Min, F. Gibou, Second-order accurate computation of curvatures in a level set framework using novel high-order reinitialization schemes, *J. Sci. Comput.* 35 (2) (2008) 114–131. doi:10.1007/s10915-007-9177-1.

- [19] D. L. Chopp, Some improvements of the fast marching method, *SIAM Journal on Scientific Computing* 23 (1) (2001) 230–244. doi:10.1137/S106482750037617X.
- [20] C.-W. Shu, S. Osher, Efficient implementation of essentially non-oscillatory shock-capturing schemes, *J. Comp. Phys.* 77 (2) (1988) 439–471. doi:10.1016/0021-9991(88)90177-5.
- [21] N. Muskhelishvili, *Some Basic Problems in the Mathematical Theory of Elasticity*, Noordhoff, Groningen, 1953.
- [22] A. H. England, On stress singularities in linear elasticity, *International Journal of Engineering Science* 9 (6) (1971) 571–585. doi:10.1016/0020-7225(71)90039-5.
- [23] E. Bouchbinder, T.-S. Lo, I. Procaccia, E. Shtilerman, Stability of an expanding circular cavity and the failure of amorphous solids, *Phys. Rev. E* 78 (2) (2008) 026124. doi:10.1103/PhysRevE.78.026124.
- [24] E. Bouchbinder, J. S. Langer, Nonequilibrium thermodynamics of driven amorphous materials. I. Internal degrees of freedom and volume deformation, *Phys. Rev. E* 80 (3) (2009) 031131. doi:10.1103/PhysRevE.80.031131.
- [25] E. Bouchbinder, J. S. Langer, Nonequilibrium thermodynamics of driven amorphous materials. II. Effective-temperature theory, *Phys. Rev. E* 80 (3) (2009) 031132. doi:10.1103/PhysRevE.80.031132.
- [26] E. Bouchbinder, J. S. Langer, Nonequilibrium thermodynamics of driven amorphous materials. III. Shear-transformation-zone plasticity, *Phys. Rev. E* 80 (3) (2009) 031133. doi:10.1103/PhysRevE.80.031133.

Acknowledgments

The authors are grateful to J. S. Langer who provided the initial motivation for the study and much support throughout, and to M. L. Manning and E. Bouchbinder for useful discussions. C. H. Rycroft was supported in part

by the National Science Foundation under grants DMS-0410110 and DMS-070590, and by the Director, Office of Science, Computational and Technology Research, U.S. Department of Energy under Contract Numbers DE-AC02-05CH11231. The research of F. Gibou was supported in part by the Department of Energy under grant agreement DE-FG02-08ER15991, the National Science Foundation under grant agreement DMS-0713858, and by the Institute for Collaborative Biotechnologies through Contract No. W911NF-09-D-0001 from the U.S. Army Research Office.

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.