Boundary element methods for earthquake modeling with realistic 3D geometries

T. Ben Thompson[1], **Brendan J. Meade**[1]

[1]Department of Earth and Planetary Sciences, Harvard University, Cambridge, MA, USA

Key Points:

- We develop a novel Galerkin boundary element method implementation with continuous displacement and slip.
- A lack of stress singularities enables earthquake cycle simulation on non-planar faults.
- A GPU-enabled fast multipole method allows solving million element problems in seconds.

Corresponding author: T. B. Thompson, `t.ben.thompson@gmail.com`

Abstract

Boundary element methods have become a foundational tool in earthquake science for the modeling of earthquake cycle kinematics. Despite their wide use and convenience typical rectangular and triangular constant slip dislocation methods produce stress singularities at the edges of every element rendering these models physically unrealistic. As we demonstrate, in an earthquake cycle simulation where the stress influences the fault slip through a friction relationship, these un-physical stress singularities manifest in severe numerical artifacts which limit their applicability to the calculation of on fault stresses and dynamic earthquake cycle modeling. To solve this problem, we develop a singularity free Galerkin boundary element method using continuous linear displacement and slip basis functions. We use Gaussian and Sauter-Schwab quadrature combined with a Stokes theorem based regularization approach in lieu of analytical formulae. In order to solve the large dense linear systems that emerge from boundary element methods, we use a fast multipole method to accurately approximate far-field element interactions. Combining these theoretical approaches with an optimized parallel implementation and GPU acceleration, we are able to solve one million element problems in seconds on a desktop computer.

1 Plain Language Summary

Earthquake scientists want to simulate and understand how faults work. To do this, we combine the physics of how the Earth bends and how faults stick and slide with our maps of faults. The current mathematical methods to do this have certain areas where the fault slip is vastly different between two points very close together. That produces infinite forces and is unrealistic. But, those infinite forces don't actually cause a problem when we are simulating a perfectly flat fault. We are presenting an improved mathematical method that doesn't have any jumps in fault slip and doesn't produce any infinite forces. This enables simulating the evolution of geometrically realistic faults over thousands of years. In addition to presenting the method, we implement it in a very computationally efficient way, using the fast multipole method. The fast multipole method is a very accurate and fast approximation of the force between far away portions of the fault. Combining the fast multipole method with parallel GPU code, we can run simulations many times faster than previous methods.

2 Introduction

A wide range of studies of fault slip and earthquake mechanics have demonstrated the effectiveness of boundary element methods for earthquake science. Boundary element methods have been so successful because they allow researchers to implicitly account for full three dimensional elasticity while limiting the solution domain to only the interfaces. Using half space Green's functions to further implicitly account for the free surface has allowed focusing on only the fault surface and eliminates the difficulty of creating an appropriate volumetric mesh. These faults are also the best mapped parts of the Earth's crust due to their illumination by seismic activity thus aiding in the creation of surface meshes. Boundary element methods are commonly applied to earthquake slip inversion [*Hudnut et al.*, 1994; *Reilinger et al.*, 2000; *Simons et al.*, 2011], earthquake cycle modeling [*Lapusta et al.*, 2000; *Liu and Rice*, 2005; *Segall and Bradley*, 2012; *Luo and Ampuero*, 2018] and dynamic rupture modeling [*Perrin et al.*, 1995; *Lapusta and Liu*, 2009].

A further reason for the success of the method is the simplicity with which a boundary element method (BEM) implementation can be assembled using closed form analytic solutions for constant slip rectangular or triangular dislocation elements [*Okada*, 1992; *Thomas*, 1993]. However, these analytical dislocation solutions produce unphysical singular stress fields due to the jump in slip at the at the boundary between elements. A consequance is unrealistic stress faults both near to an on fault surfaces themselves. This issue can be ignored in a slip inversion setting where there is no need to evaluate stresses near element edges. However, as we demonstrate, in a both quasi-static fault stressing calculations and earthquake cycle simulation where the stress influences the fault slip through a friction relationship, these unphysical stress singularities manifest in severe numerical artifacts and inaccuracy. This has limited BEM-based approaches for earthquake cycle simulation to evenly discretized planar faults with mostly constant frictional properties [*Liu and Rice*, 2005; *Liu and Rubin*, 2010; *Segall and Bradley*, 2012; *Luo and Ampuero*, 2018].

The solution is to adopt a discretization that maintains a continuous displacement and slip field. Such a discretization will not have stress singularities. However, despite the introduction of the constant basis displacement discontinuity BEM method in 1976 [*Crouch*, 1976; *Crouch and Starfield*], over the next 40 years, no BEM implementations using linear basis functions have been introduced in earthquake science. Why is this? The primary reason is that extending the analytical methods used to derive the constant slip dislocation ele-

ments to linear slip dislocation elements would be extremely challenging. Similarly, numerical integration of the hypersingular ($O(1/r^3)$) divergent integrals in most BEM formulations for cracks and faults is also exceptionally difficult. In a collocation-type BEM method where integral equations are enforced at specific points, the numerical integration remains partially unsolved. However, if we instead enforce the integral equation in a weighted sense using a Galerkin boundary element method, we can "spread" the singularity over another surface integral [*Nedelec*, 1982; *Balakrishna et al.*, 1994; *Sutradhar et al.*, 2008]. Combining the Galerkin BEM with a Stokes theorem based regularization approach [*Bonnet*, 1995; *Li et al.*, 1998; *Frangi et al.*, 2002], we can reduce the hypersingular integrals to weakly singular integrals and use a continuous linear basis for the displacement and slip field.

A second issue with common BEM implementations is the dense matrix of element interactions. Inverting and even storing this matrix becomes infeasible as problem sizes increase. We manage this limitation by approximating farfield element interactions with the fast multipole method [*Greengard and Rokhlin*, 1987]. This approximation can be very accurate due to the low rank of clustered farfield interactions between elements. The resulting algorithm scales linearly in both time and memory requirements as a function of number of elements and enables studying detailed regional or even global geometrically realistic models of earthquake behavior. The key parts of our BEM and FMM implementations are parallelized over many cores and partially implemented in CUDA to make use of GPU capabilities. The result is an algorithm that can perform a one million element BEM matrix vector product in 2.1 seconds on a standard desktop computer. We use this new tool to investigate the strongly nonlinear effects of nonplanar fault and Earth surface geometry on earthquake and slip modeling problems.

3 The problems with constant slip dislocation elements

Constant slip triangular and rectangular dislocation elements have stress singularities at their edges [*Kelvin*, 1848; *Okada*, 1992; *Jeyakumaran et al.*, 1992; *Thomas*, 1993; *Meade*, 2007; *Nikkhoo and Walter*, 2015]. This is due to their unphysical displacement field. At the edge of the element, there is an infinitesimal jump in slip. The displacement gradient is then infinite, resulting in infinite strains and stresses. Despite this shortcoming, dislocation elements have become the primary BEM approach in the earthquake science community because of their effectiveness for slip inversion problems [*Hudnut et al.*, 1994; *Reilinger et al.*, 2000; *Simons et al.*, 2011]. For slip inversion problems, the main elastic interactions that

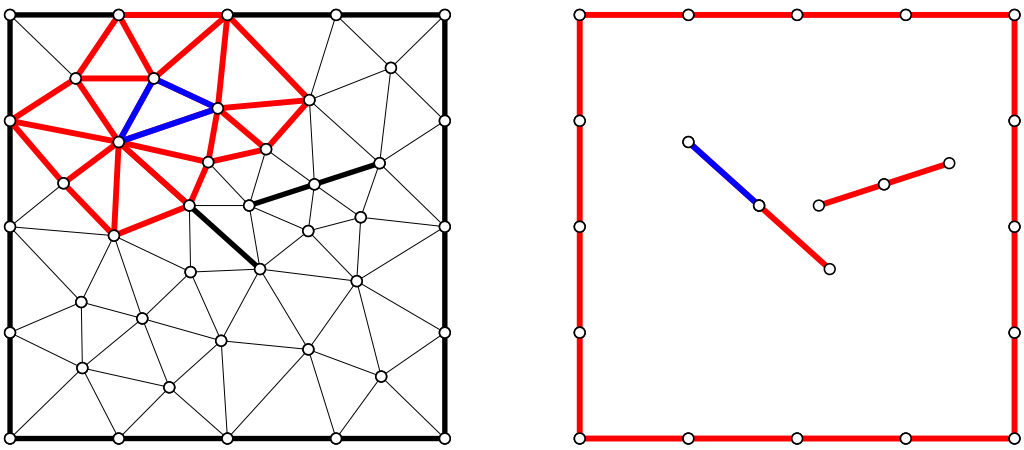

(a) (b)

Figure 1. a) A finite element discretization of a simple square mesh with two embedded faults. The red highlighted triangles are the triangles that interact with the blue source triangle in the finite element matrix. b) A boundary element discretization of the same geometry. The blue source element has a non-zero interaction with every other element in the mesh. This results in a dense matrix as opposed to the sparse matrix from the finite element method.

need to be calculated are the Green's functions between the fault elements and the locations of surface displacement observations. Surface displacement observations are almost always located far from the edges of the fault elements. As a result, the stress singularity can be ignored.

On the other hand, constant slip dislocation elements are also used for earthquake simulation [*Liu and Rice*, 2005; *Richards-Dinger and Dieterich*, 2012; *Shibazaki et al.*, 2012; *Qiu et al.*, 2016; *Luo and Ampuero*, 2018; *Yu et al.*, 2018]. At each time step in an earthquake cycle simulation, the stresses on the fault are calculated from the current slip and then a friction relationship is used to calculate the new fault slip rate. As a result, stresses are being evaluated at observation points on the fault plane. These observation points are very close to the singularities at the edges of the elements and as a result can be severely distorted. Distorted stresses then enter into frictional calculations and produce inaccurate slip rates.

However, in the common case of a planar fault mesh composed of rectangular elements that are all the same size, this singularity effect cancels out [*Bradley*, 2014]. This is because the stress is only evaluated at the observation point at the center of the element and the dis-

torting effect of the singularities at each of the four edges of the rectangle cancel each other out. When that very delicate symmetric singularity cancellation is disrupted, constant slip dislocation elements can no longer accurately model stresses on the fault surface itself. In particular, unlike rectangular dislocations, triangular dislocations fail to accurately model fault stresses on a planar fault for all the choices of observation points we tested. We demonstrate these failure modes of dislocation elements in Figure 2 where we evaluate the shear stress on a strike slip fault with mesh geometries using both rectangular dislocation and triangular dislocation. We also compare several choices of observation point.

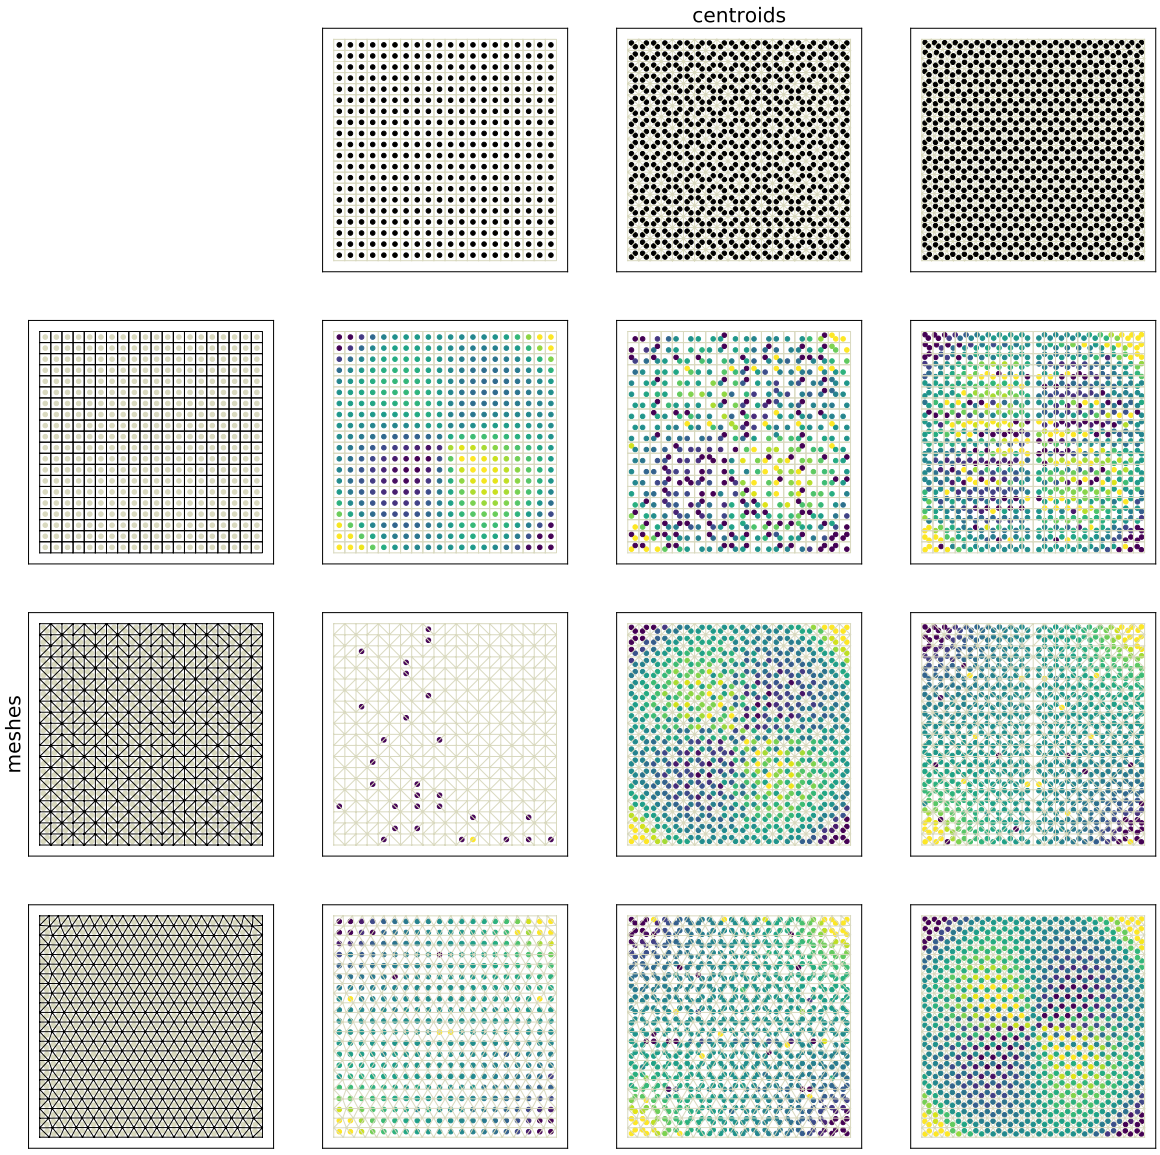

Figure 2. A comparison of the shear stress resulting from a Gaussian strike slip pulse at the center of a planar fault beneath a half space. The first row shows the evaluated shear stress using rectangular dislocations. The second row uses triangular dislocations. The third row uses triangular dislocations on a mesh composed of equilateral triangles to see if we can recover the singularity balance found in the rectangular case. The columns show the stresses evaluated at the centroids of the elements in the first, second and third meshes respectively. We see that when rectangular dislocations are used and the observation points are chosen to be the centers of those rectangles, we recover the correct shear stress distribution (row 1, column 1). We also see that using the centroids of the two triangular meshes gives a shear stress field that is close to correct, except with a erroneous oscillation overlaid (row 2, column 2 and row 3, column 3). Finally, as seen in all the other figures, evaluating at anything besides an element's centroid results in entirely incorrect shear stress field. In contrast, a boundary element method using linear slip variation can evaluate the shear stress at any observation point because there are no induced and unphysical slip jumps and stress singularities.

The solution to this problem is to use elements with linear slip variation. Then, by enforcing displacement and slip continuity constraints at the boundaries between elements, we can eliminate singular stresses. But, this approach requires new solutions for the displacement and tractions in the domain due to a linearly varying dislocation. Instead of tackling this challenging problem, we instead use numerical approaches to integrate the appropriate Green's functions. The result is a method that can accurately calculate stresses on the fault plane regardless of mesh density variations or nonplanar geometries.

4 Galerkin BEM

As opposed to the analytical dislocation approach to the BEM, there are two primary numerical approaches: collocation and Galerkin methods. The collocation approach imposes the boundary integral equation at many individual points on the surface mesh. The Galerkin approach, by contrast, imposes the boundary integral equation in the sense of a weighted sum over each basis function's support. We adopt the Galerkin BEM (GBEM) because we believe it is better suited to crack and fault modeling. The Green's function that calculates the stress in the domain resulting from a point slip source on a fault has a $O(1/r^3)$ behavior. As such, it is "hypersingular" and, unless carefully treated, integrals of this Green's function over a source element are divergent [*Sutradhar et al.*, 2008]. See Figure 4 for a summary of the singular behavior of the BEM kernels. Through integration by parts, the GBEM provides a simple way of regularizing this divergent behavior [*Bonnet*, 1995; *Li et al.*, 1998; *Frangi et al.*, 2002]. As an added benefit, the GBEM is generally an order of magnitude more accurate than the collocation method.

The potential downside is that the three dimensional GBEM requires evaluating four dimensional Green's function integrals for the interaction between every pair of elements. Four dimensional singular nearfield integrals can become very computationally expensive. This issue is minimized by the regularized integration process and by using specific quadrature rules tailored to the nearfield integrals.

4.1 The displacement boundary integral equation

The GBEM, like almost all boundary integral approaches to linear isotropic elasticity, is derived from the Somigliana identity:

$$u_k(\mathbf{x}) + \int_S T^*_{kj}(\mathbf{x},\mathbf{y})u_j(\mathbf{y})d\mathbf{y} = \int_S U^*_{kj}(\mathbf{x},\mathbf{y})t_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \tag{1}$$

where $S = \partial V$ (the boundary of V), $u_k(\mathbf{x})$ is k-th component the displacement field at $\mathbf{x}$, $t_k(\mathbf{x})$ is k-th component of the traction field and

$$U^*_{kj}(\mathbf{x},\mathbf{y}) = \frac{1}{16\pi\mu(1-\nu)r}\left[(3-4\nu)\delta_{kj} + r_{,k}r_{,j}\right] \tag{2}$$

$$T^*_{kj}(\mathbf{x},\mathbf{y}) = \frac{-1}{8\pi(1-\nu)r^2}\left[\{(1-2\nu)\delta_{kj} + 3r_{,k}r_{,j}\}\frac{\partial r}{\partial \mathbf{n}} - (1-2\nu)\{n_j r_{,k} - n_k r_{,j}\}\right] \tag{3}$$

are the fundamental elastic and traction Green's functions of elasticity, with δ_{ij} is the Kronecker delta, μ as the elastic shear modulus, ν as the poisson ratio, $\mathbf{n}$ is the normal vector to S at y, $r = \|\mathbf{x} - \mathbf{y}\|$ and $r_{,i} = \partial r/\partial x_i$.

At an intuitive level, this equation says that, if we know the displacement and traction on the boundary of an elastic body, then we have sufficient information to determine the displacement everywhere within the elastic body. For this reason, the equation is also sometimes known as the displacement boundary integral equation.

The Somigliana identity also gives a method for solving for all the boundary information from partial boundary information. If at every point on the boundary of the domain, either displacement or traction is known, then the other field can be solved for by inverting equation (1). In most practical applications, this process involves a numerical discretization of the Somigliana identity – the BEM.

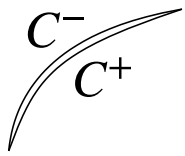

Figure 3. A schematic showing the approach of treating a crack as two infinitesimally separated surfaces with balanced forces and a displacement jump.

So far, this exposition ignores cracks or faults in the elastic body. However, for earthquake science, these are the most important surfaces in the problem. The standard approach to treat cracks is to separate the crack into two infinitesimally separated surfaces, C^+ and C^-

(see Figure 3). Then, introducing the displacement discontinuity $\Delta u = u^+ - u^-$, enforcing force balance across the crack, $t^+ + t^- = 0$, we get:

$$u_k(\mathbf{x}) + \int_F T^*_{kj}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y})d\mathbf{y} + \int_S T^*_{kj}(\mathbf{x},\mathbf{y})u_j(\mathbf{y})d\mathbf{y} = \int_S U^*_{kj}(\mathbf{x},\mathbf{y})t_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \tag{4}$$

where C is the crack surface. Note that the U^* integral over F on the right hand side has dropped out to the force balance assumption.

4.2 The traction boundary integral equation

Critically, the traction on the fault surface does not appear in equation (4). Solving for traction given fault slip is a critical step in many earthquake simulation problems. To solve this issue, another integral equation can be derived from the Somigliana identity by taking gradients and applying the elastic constitutive equations:

$$\sigma_{lk}(\mathbf{x}) - \int_S A^*_{lkj}(\mathbf{x},\mathbf{y})t_j(\mathbf{y})d\mathbf{y} = -\int_S H^*_{lkj}(\mathbf{x},\mathbf{y})u_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \tag{5}$$

with

$$A^*_{lkj}(\mathbf{x},\mathbf{y}) = \frac{1}{8\pi(1-\nu)r^2}\left[(1-2\nu)\{\delta_{lj}r_{,k} + \delta kjr_{,l} - \delta_{lk}r_{,j}\} + 3r_{,l}r_{,k}r_{,j}\right] \tag{6}$$

$$\begin{aligned} H^*_{lkj}(\mathbf{x},\mathbf{y}) = \frac{\mu}{4\pi(1\nu)r^3}\Big[&3\frac{\partial r}{\partial \mathbf{n}}(\{1-2\nu\}\delta_{lk}r_{,j} + \nu(\delta_{kj}r_{,l} + \delta_{lj}r_{,k}) - 5r_{,l}r_{,k}r_{,j}) \\ &+ (1-2\nu)(3n_jr_{,l}r_{,k} + n_k\delta lj + n_l\delta_{kj}) \\ &+ 3\nu(n_lr_{,k}r_{,j} + n_kr_{,l}r_{,j}) \\ &- (1-4\nu)n_m\delta_{lk}\Big] \end{aligned} \tag{7}$$

This equation is often called the traction boundary integral equation, because in contrast to equation (1), given the boundary conditions, we calculate the traction at any interior point instead of the displacement.

Just like the displacement boundary integral equation, we can treat cracks as two infinitesimally separated surfaces with a jump in displacement and balanced tractions. The resulting integral equation is:

$$\begin{aligned} \sigma_{lk}(\mathbf{x}) - \int_S A^*_{lkj}(\mathbf{x},\mathbf{y})t_j(\mathbf{y})d\mathbf{y} = \\ -\int_S H^*_{lkj}(\mathbf{x},\mathbf{y})u_j(\mathbf{y})d\mathbf{y} - \int_F H^*_{lkj}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \end{aligned} \tag{8}$$

While the fault traction still does not appear within any surface integrals, the point $\mathbf{x}$ can be chosen to lie on the surface F, giving us the stress and traction on the fault surface. As a result, in contrast to the displacement boundary integral equation, this integral equation can be used to solve for fault stress and traction and will be a fundamental component of our earthquake simulations.

4.3 Discretization

As an illustrative example I describe the BEM discretization for a fault with known slip underneath a topographic free surface. This is a common problem and retains all the challenges of the more general elastic BEM problem. I will start from the displacement boundary integral equation (4) and assume that the surface tractions are zero:

$$u_k(\mathbf{x}) + \int_S T^*_{kj}(\mathbf{x},\mathbf{y})u_j(\mathbf{y})d\mathbf{y} = -\int_F T^*_{kj}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \tag{9}$$

Because the slip on F is known, the right hand side of this equation can be fully calculated. In the next section, we will discuss how to perform this calculation. However, the surface integral on the left hand side is a functional of the unknown surface displacements. Our goal is to transform this integral equation into a linear system to allow solving for the surface displacement with standard numerical linear algebra.

First, we will approximate our surface, S, as a mesh composed of elements, S_i with $u_j(\mathbf{y})$ defined as a sum of basis functions on each element.

$$\begin{aligned} u_k(\mathbf{x}) + \sum_{i=0}^{N} \int_{S_i} T^*_{kj}(\mathbf{x},\mathbf{y}) \sum_r \phi_r(\mathbf{y})\tilde{u}_{rj}(\mathbf{y})d\mathbf{y} = \\ -\int_F T^*_{kj}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y})d\mathbf{y} \quad \forall \mathbf{x} \in V \end{aligned} \tag{10}$$

$$u_j(\mathbf{y}) = \sum_r \phi_r \tilde{u}_{r,j} \tag{11}$$

where $\tilde{u}_{r,j}$ are the unknown coefficients of the displacement basis expansion. Now, we have an integral equation that relates the displacement at an arbitrary point $\mathbf{x}$ to the unknown displacement coefficients.

The next step is to choose how to impose this integral equation. A traditional constant basis displacement discontinuity BEM will enforce equation (10) at the centroid point of

many rectangular or triangular elements. This is called a collocation method. By contrast, a Galerkin method enforces equation (1) in a weighted sense over the entire mesh. In particular, the weighting functions are chosen to be the same as the basis functions used for discretizing the displacement field. To be precise:

$$\int_S \phi_q(\mathbf{x})\left[u_k(\mathbf{x}) + \sum_{i=0}^{N} \int_{S_i} T^*_{kj}(\mathbf{x},\mathbf{y}) \sum_r \phi_r(\mathbf{y})\tilde{u}_{rj} d\mathbf{y}\right] d\mathbf{x} = \int_S \phi_q(\mathbf{x})\left[-\int_F T^*_{kj}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y}) d\mathbf{y}\right] d\mathbf{x} \quad \forall q \tag{12}$$

with the previous integral equation integrated against each basis function $\phi_q(\mathbf{x})$.

At this point, if there are N basis functions in the mesh, then we have $3N$ (N for each component of displacement) unknowns and $3N$ equations that form our linear system:

$$A_{IJ}U_J = b_I \tag{13}$$

$$U_J = \tilde{u}_{r(J)j(J)} \tag{14}$$

$$A_{IJ} = M_{IJ} + T_{IJ} \tag{15}$$

$$b_I = -\int_{S_{i(I)}} \int_F \phi_{q(I)}(\mathbf{x}) T^*_{k(I)j}(\mathbf{x},\mathbf{y})\Delta u_j(\mathbf{y}) d\mathbf{y} d\mathbf{x} \tag{16}$$

$$T_{IJ} = \int_{S_{i(I)}} \int_{S_{i(J)}} \phi_{q(I)}(\mathbf{x}) T^*_{k(I)j(J)}(\mathbf{x},\mathbf{y})\phi_{r(J)}(\mathbf{y}) d\mathbf{y} d\mathbf{x} \tag{17}$$

$$M_{IJ} = \int_{S_{i(I)}} \phi_{q(I)}(\mathbf{x})\phi_{r(J)}(\mathbf{x}) d\mathbf{x} \tag{18}$$

where $i(I)$, $r(J)$, $q(I)$, $k(I)$, $j(J)$ map from the matrix row and column indices I and J to the relevant element index, basis function index or displacement component index.

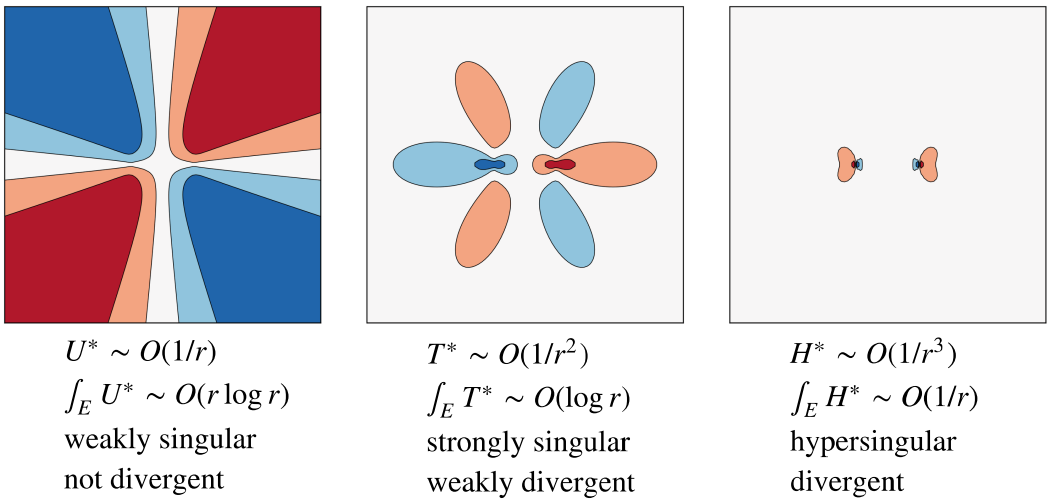

(a)

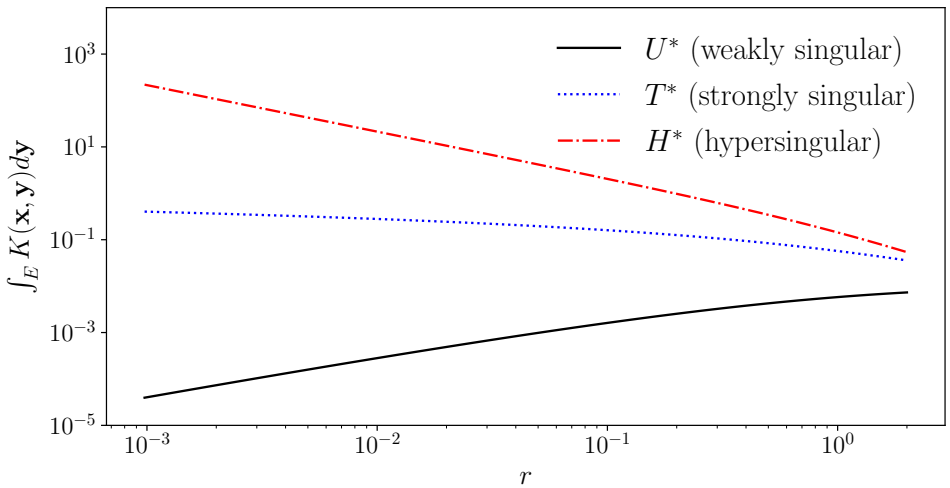

(b)

Figure 4. We compare the singular behavior of the U^*, T^* and H^* kernels. The A^* kernel has the same singular behavior as T^*. a) The filled contour plots show the value of the the respective kernel integrated over an element from -1 to 1 on the x-axis. The divergent behavior of the integral of the H^* kernel is apparent at the end points of the element. Below the contour plots, we give some information about the singular behavior of each kernel. b) The behavior of the kernel integrals in log-log space as the observation point approach the tip of the source element (i.e. $r \to 0$). We can see the divergence of the hypersingular kernel.

This linear system has three main components. Equation 16 is a double surface integral representing the effect of fault slip on the observation surface element, $S_{i(I)}$. Equation

(17) is a double surface integral representing the effect of surface displacement from a particular source surface element, $S_{i(J)}$ on the observation surface element, $S_{i(I)}$. Equation (18) is termed the mass matrix and differs from the other two components in that it has no inner Green's function integral.

With this linear system set up, the remaining task is to compute the entries of the matrix and right hand side.

By decomposing the slip, Δu, into a set of basis function on a fault mesh and dropping much of the index notation for the remainder of the discussion and unify the problem calculation of many different surface integrals into the basic problem of calculating:

$$\int_{E_1}\int_{E_2} \phi(\mathbf{x})K(\mathbf{x},\mathbf{y})\psi(\mathbf{y})d\mathbf{y}d\mathbf{x} \tag{19}$$

where K is one of U^*, T^*, A^* or H^* and $\phi(\mathbf{x})$ and $\psi(\mathbf{y})$ are the basis functions of interest on elements E_1 and E_2. Calculating this element pair integral is the fundamental task of assembling a boundary element matrix.

4.4 Singularity and Regularization

We have reduced the problem of assembling a non-singular, non-planar, elastic boundary element matrix to calculating equation (19). If E_1 and E_2 are disjoint, this is fairly straightforward. However, in equation (17), there are element pair integrals where $E_1 = E_2$ (coincident elements) or E_1 shares an edge or vertex with E_2 (edge adjacent or vertex adjacent). In these cases, there are points for which $\mathbf{x} = \mathbf{y}$ and the integrand in the element pair integral is undefined.

The behavior of these singular integrands is very important. The U^* kernel has an $O(1/r)$ behavior, termed weakly singular. This means if the integral is evaluated for $\mathbf{x} \in S$ that while the integrand is singular, the integral itself is well defined. The T^* kernel has an $O(1/r^2)$ form that is termed strongly singular which means the integral is actually divergent. However, the integral can be interpreted in a physically meaningful way using its Cauchy principal value excluding an infinitesimal ball around the source point $\mathbf{y}$. The kernel A^* has an $O(1/r^2)$ behavior, bringing the same challenges as the Cauchy principal value integral of T^* does in the displacement boundary integral equation.

However, the kernel H^* suffers from a more extreme $O(1/r^3)$ hypersingular behavior. Individual element pair hypersingular integrals that are coincident or edge adjacent are divergent. Fortunately, the divergence terms drop out of the final equations. Due to the nature of the discretization process, at an edge between two elements, there will be two separate displacement values; one from the sum of basis functions on each of the touching elements. But, displacement should be continuous at this boundary. After imposing this continuity constraint, when using a Galerkin method, the opposite sign divergent terms cancel. We leave out the complex proof of this fact and refer to *Sutradhar et al.* [2008]. This divergence cancellation is one of the main reasons why we use the Galerkin discretization as opposed to collocation. With a collocation method, the hypersingular integrals in the collocation method are truly divergent, while the divergent terms resulting from the hypersingular integrals in the Galerkin method drop out.

Despite this theoretical divergence calculation, the numerical calculations of the strongly singular and hypersingular element pair integrals of T^*, A^* and H^* is extremely difficult.

To work around these difficult singular integrals, we make use a regularized form of the integrals [*Frangi et al.*, 2002]. The regularization approach uses a form of Stokes theorem to integrate (19) by parts and reduce the order of the singularity, transfering derivatives to the basis functions. Focusing on the hypersingular kernel, H^*:

$$\int_{E_1}\int_{E_2} R_i(\phi(\mathbf{x}))B_{iksj}(\mathbf{x},\mathbf{y})R_s(\psi(\mathbf{x}))d\mathbf{y}d\mathbf{x} \tag{20}$$

where R_i is the "surface rotor" defined as:

$$R_i(\phi(\mathbf{x})) = e_{bci}n_b(\mathbf{x})\frac{\partial\phi}{\partial x_c}(\mathbf{x}) \tag{21}$$

and

$$\begin{aligned} B_{iksj} =& \frac{-\mu}{8\pi r}(\delta_{eg} - r_{,e}r_{,g})e_{iep}e_{kgr} \\ &\times\left[\frac{2\nu}{1-\nu}\delta_{ps}\delta_{rj} + \delta_{pr}\delta_{sj} + \delta_{pj}\delta_{sr}\right] \end{aligned} \tag{22}$$

is the regularized hypersingular kernel, where e_{ijk} is the Levi-Civita tensor. Importantly, it has $O(1/r)$ before and is weakly singular. As a result, the integral (20) is well defined. A similar regularized version of the strongly singular integrals of T^* and A^* can be derived. For the hypersingular kernel, this regularization is only possible in a Galerkin formulation because it relies on transfering a derivative to both the source basis function, $\psi(\mathbf{y})$ and the

observation basis function $\phi(\mathbf{x})$. In a collocation method, no observation basis function exists.

4.5 Numerical quadrature of boundary element integrals

The remaining task is to calculate the surface integrals and solve the linear system. The traditional analytical dislocation approach will be exceptionally difficult for these double surface integrals. As a result, we use numerical quadrature methods.

In the previous section, we left the mathematical discussion agnostic to the exact form of the basis functions and the shape of the mesh elements. For here on, we will focus on triangular mesh elements and linear basis functions. Then, numerical quadrature of the mass matrix (18) is simple since the product of the two linear basis functions is a quadratic over the observation triangle. A Gaussian quadrature with three points will integrate such a function exactly.

Relationship	Geometry	Method	# Integrals
Coincident		Change of variables	81 million
Edge Adjacent		Change of variables	243 million
Vertex Adjacent		Change of variables	729 million
Nearfield		5th order Gauss	6.5 billion
Farfield		2nd order Gauss	81 trillion

Figure 5. The five types of integrals required for building a Galerkin boundary element matrix and the methods we use for each types. We include the number of integrals of each type that must be computed for a BEM problem with one million elements.

Then, we will separate the integration of (20) into four situation depending on the relationship between the two elements, E_1 and E_2: non-touching, vertex adjacent, edge adjacent and coincident. Despite the regularization process, a weak singularity remains at $\mathbf{x} = \mathbf{y}$ for all the kernels. So, the quadrature rules must take that behavior into account to converge quickly.

For the non-touching case, the integrand is never singular and behaves smoothly. As a result, Gaussian quadrature methods are ideal. We use the tensor product of two triangular Gaussian quadrature rules [*Zhang et al.*, 2009]. The order of the quadrature rule depends

on the distance between the elements. Element pair integrals for elements that are far from each other can be approximated very accurately with just two or three quadrature points. For closer elements, we often use a fifth order Gauss quadrature rule. For the three touching cases, we use different forms of the Sauter-Schwab quadrature rules [*Duffy*, 1982; *Sauter and Schwab*, 2010]. These rules make use of changes of variables to ameliorate the $O(1/r)$ singularity. Then, Gaussian quadrature is used on the resulting non-singular integrand. These integrals can be accurately integrated to four decimal places with eight quadrature points per dimension. The Sauter-Schwab rules split the integration domain into two to eight subdomains. With four dimensions in each subdomain due to the double surface integral, that means that 8^5 are necessary for the coincident integrals.

5 Fast multipole method

The boundary element matrix resulting from discretizing any of the elastic kernels is a dense matrix. To store a dense matrix requires $O(N^2)$ memory and to invert a dense matrix requires $O(N^3)$ floating point operations. As a result, traditional boundary element implementations have been limited to less than 50,000 elements. However, several extremely effective methods for producing a sparse approximation to these dense matrices have been developed [*Hackbusch and Nowak*, 1989; *Bebendorf and Rjasanow*, 2003; *Liu and Nishimura*, 2006].

One of these approaches, the fast multipole method (FMM) is an approximate method for calculating the farfield interactions in an n-body problem or integral equation [*Greengard and Rokhlin*, 1987]. Nearby elements are grouped together and their effects on farfield elements are approximated using a spherical harmonic expansion. In terms of the linear algebra, the FMM is based on the realization that any block of a n-body or BEM matrix that does not contain the diagonal of the matrix can be decomposed and accurately represented by only a small subset of its eigenvalues. In other words, the off diagonal blocks, which correspond to farfield interactions, of the matrix are low rank. The FMM is a physically motivated approach to take advantage of this low rank property.

To explain the FMM, we first define which element interactions are farfield. We separate the non-touching element pair integrals discussed in the previous section into nearfield and farfield integrals based on the ratio of the distance between the elements to the radii of the minimum bounding spheres of those elements. The chosen minimum ratio is called the

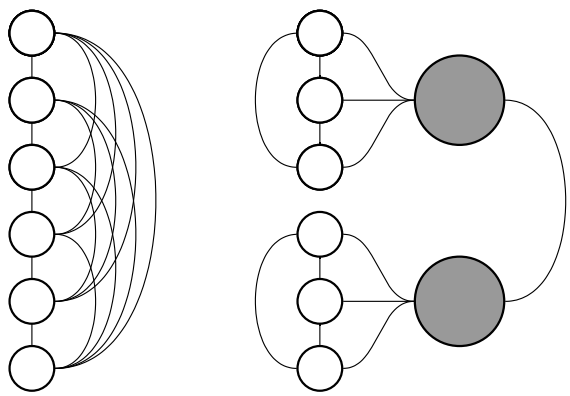

Figure 6. A schematic comparison between a direct calculation (left) and the FMM (right). In the direct calculation, all element pair integrals must be explicitly calculated. In the FMM, a tree structure is imposed and only nearfield calculations are performed directly. Farfield calculations are performed through "multipole expansions" that represent many elements through a single spherical harmonic expansion.

multipole acceptance criteria or MAC. So, an interaction between elements i and j is farfield if:

$$\text{dist}(i, j) < MAC(R_i + R_j) \tag{23}$$

These farfield element interaction integrals are all approximated using the FMM.

Most presentations of the FMM focus on point to point interactions of the form

$$F_i = \sum_j K(\mathbf{x}_i, \mathbf{y}_j) G_j \tag{24}$$

where G_j are known source coefficients and $K(\mathbf{x}, \mathbf{y})$ is a kernel function that normally decays rapidly with the distance between $\mathbf{x}$ and $\mathbf{y}$. However, using the FMM as a tool in the Galerkin boundary element method, we have a slightly different form involving source and observation surface integrals. For example, rearranging (19):

$$F_i = \int_{E_i} \phi_i(\mathbf{x}) \sum_j \Big[\int_{E_j} K(\mathbf{x}, \mathbf{y}) \sum_r \big[\psi_r(\mathbf{y}) G_r\big] d\mathbf{y} \Big] d\mathbf{x}) \tag{25}$$

But, by approximating the surface integrals over E_i and E_j with a quadrature rule, we can return to the summation form of the point to point FMM in (24):

$$F_i \approx \sum_q w_q \phi_i(\mathbf{x}_q) \sum_j \Big[\sum_p w_p K(\mathbf{x}_q, \mathbf{y}_p) \sum_r \big[\psi_r(\mathbf{y}_p) G_r\big] \Big] \tag{26}$$

Although there are several more summations, the fundamental form is the same. For this transformation from surface integrals to a point to point interaction sum to work, the same

quadrature rule must be used for every source element and for every observation element. Since we are approximating the farfield entries in the BEM matrix, we generally use a 1st or 2nd order Gauss quadrature rule. However, this quadrature rule does not apply for the nearfield matrix entries, which are included in the sum (26). These nearfield matrix entries are calculated directly using higher order and specialized quadrature rules as discussed in the previous section. Because these integrals are approximated in both the nearfield and farfield portions of the matrix, we have:

$$A_{ij} = A_{ij}^{near} + A_{ij}^{far} + \widetilde{A}_{ij}^{near} \tag{27}$$

This is incorrect because a second, low accuracy nearfield term is being added to the matrix. For ease of implementation, we directly cancel this term when constructing the nearfield matrix. This solution is similar to the precorrected FFT scheme [*Nie et al.*, 2002]. Now that we have established that the surface integral BEM terms can be transformed into a point to point sum, the remaining discussion of the FMM will focus on the simpler point to point setting.

We use a version of the FMM based on spherical harmonics to approximate the elastic interaction integrals [*Greengard and Rokhlin*, 1987; *Liu and Nishimura*, 2006; *Pham et al.*, 2012]. The main expansion is

$$\frac{1}{\|\mathbf{x}-\mathbf{y}\|} = \sum_{n=0}^{\infty}\sum_{m=-n}^{n} \overline{S_{n,m}}(\mathbf{x}-\mathbf{y}_c)R_{n,m}(\mathbf{y}-\mathbf{y}_c), \quad \|\mathbf{y}-\mathbf{y}_c\| < \|\mathbf{x}-\mathbf{y}_c\| \tag{28}$$

where $\mathbf{y}_c$ is the center of the spherical harmonic expansion and is assumed to be close to $\mathbf{y}$, the bar indicates the complex conjugate, and

$$R_{n,m}(\mathbf{x}) = \frac{1}{(n+m)!}P_n^m(\cos\theta)e^{im\phi}r^n \tag{29}$$

$$S_{n,m}(\mathbf{x}) = (n-m)!P_n^m(\cos\theta)e^{im\phi}\frac{1}{r^{n+1}} \tag{30}$$

are the terms in the spherical harmonic expansion, where (ρ, θ, ϕ) are the spherical coordinates of $\mathbf{x}$ and P_n^m are the associated Legendre functions defined as

$$P_n^m = (1-x)^{m/2}\frac{d^m}{dx^m}P_n(x) \tag{31}$$

in terms of the Legendre polynomial of degree n, $P_n(x)$. In practice, the function $R_{n,m}$ and $S_{n,m}$ are not calculated directly but instead via a recurrence where $R_{n,m}$ depends on the previous entries in the sequence and likewise for $S_{n,m}$.

The key component of the "multipole expansion" (28) is the linear separation of the component of $1/r$ related to the source points ($R_{n,m}$) and the component related to the observation points ($\overline{S_{n,m}}$). To make this more concrete, suppose we have N sources located in a

cluster far from M observation points and we would like to compute the n-body sum

$$F_i = \sum_{j}^{N} \frac{1}{\|\mathbf{x}_i - \mathbf{y}_j\|} \quad \forall i \in 1, ..., N \tag{32}$$

The default approach of directly computing the sum for each observation point has a cost $O(NM)$. However, because the sources and observation points are far away, it's possible substitute in a truncated multipole expansion from equation (28), giving:

$$F_i = \sum_{j}^{N} \sum_{n=0}^{P} \sum_{m=-n}^{n} \overline{S_{n,m}}(\mathbf{x}_i - \mathbf{y}_c) R_{n,m}(\mathbf{y}_j - \mathbf{y}_c) \quad \forall i \in 1, ..., N \tag{33}$$

And because only the $R_{n,m}(\mathbf{y}_j - \mathbf{y}_c)$ term depends on the summation over j, we can swap the order of summation to get:

$$F_i = \sum_{n=0}^{P} \sum_{m=-n}^{n} \overline{S_{n,m}}(\mathbf{x}_i - \mathbf{y}_c) \Big[\sum_{j}^{N} R_{n,m}(\mathbf{y}_j - \mathbf{y}_c) \Big] \quad \forall i \in 1, ..., N \tag{34}$$

The innermost summation over j is independent of i and thus can be calculated once for all i. As a result, using the harmonic expansion gives an approximate approach to computing the sum (32) that requires $O(MP^2)$ operations to calculate the innermost sum and $O(NP^2)$ operations to calculate the F_i. If $P^2 \ll N$ and $P^2 \ll M$, then this fast multipole method is a more efficient approach to calculating the sum. In Figure 7, we compare the FMM to direct calculation for the elastic U^* kernel, demonstrating that the approximation effectively reproduces the correct results. The approximate nature of the FMM is not a significant impedient due to the rapid convergence of the expansion (28). With $P > 10$, we can easily achieve machine precision while still accelerating the computation of an n-body summation. In practice, since the approximation is only used for farfield interactions, increasing the multipole acceptance criteria for a more stringent definition of the farfield will also improve the accuracy. The tradeoff between improving accuracy through increasing the order or increasing the MAC is shown in (FIGURE!).

Two issues remain before the multipole expansion can be successfully implemented to acclerate the elastic boundary element method. First, several of the elastic kernels include terms like $\frac{\mathbf{x}_i - \mathbf{y}_i}{\|\mathbf{x}-\mathbf{y}\|^2}$. These terms cannot be directly approximated using equation (28). However, if we recognize that

$$\frac{d}{dx_i}\left[\frac{1}{\|\mathbf{x} - \mathbf{y}\|}\right] = \frac{\mathbf{x}_i - \mathbf{y}_i}{\|\mathbf{x} - \mathbf{y}\|^2} \tag{35}$$

and that the derivatives of R_n, m and S_n, m are easily calculated, then the method can be extended to all the elastic interaction kernels.

Second, the example summation above only considered a pair of source and observation points that were well separated. In reality, some observation and source points will be

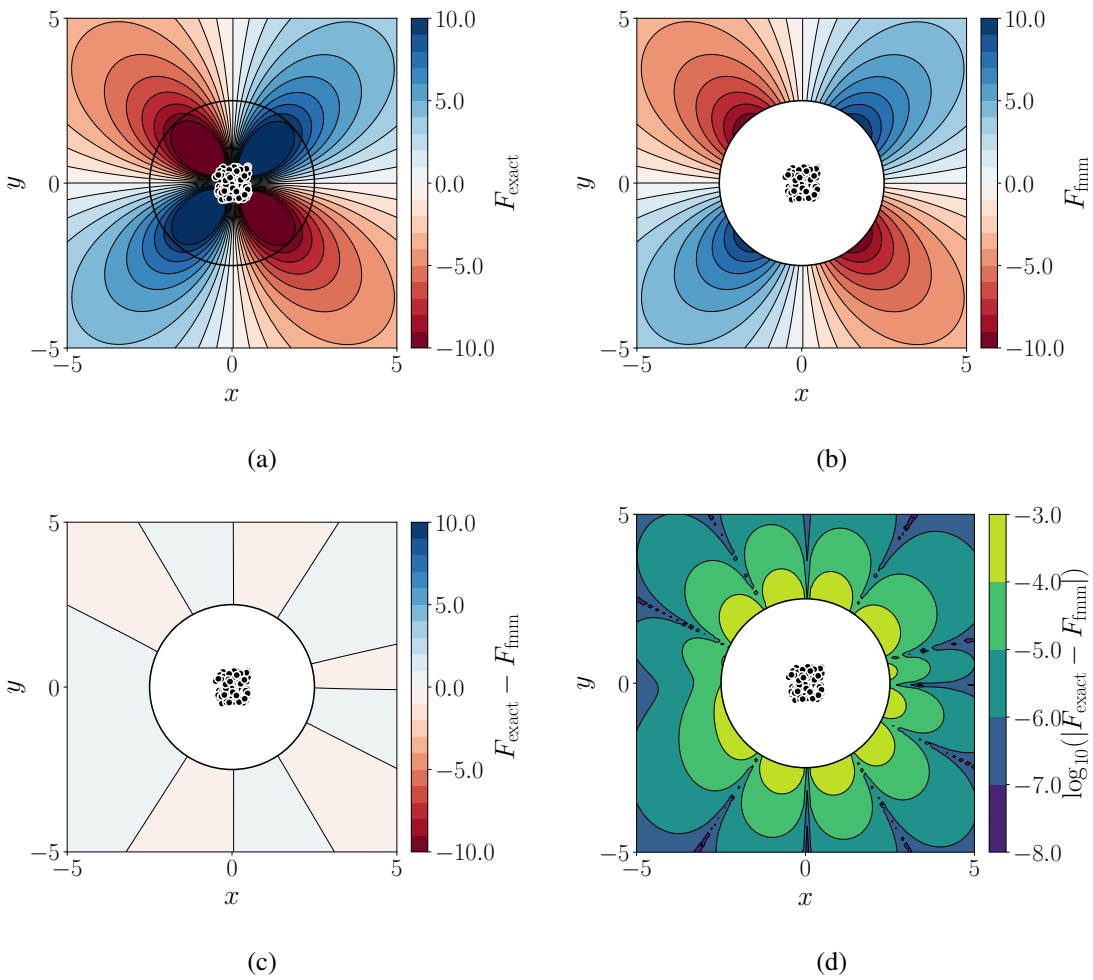

Figure 7. a) The result from performing a direct n-body sum using the U^* kernel and the sources shown by the black dots. The black circle shows the MAC of 2.5 chosen for this example. b) The result from performing an FMM approximate sum with $P = 5$. The inner circle is excluded because the FMM approximation is invalid near the sources. c) The difference between the FMM result and the direct result. d) The base 10 log of the difference. Note the maximal difference of 10^{-3} and the decay in error further from the sources.

close together and some will be well separated. To deal with this, we build a hierarchical tree of both the source and observation points. Then, several operators are used to build up the nearfield exact calculation and farfield approximation from the individual cells of the tree. These operators are usually named with a convention X2Y where X and Y are the input and output type respectively which can be "P" for the input points, "M" for the multipole expansion and "L" for the local expansion. The P2M operator takes the input field and calculates the multipole expansion for a cell. The M2M operator takes the multipole expansions for several child cells and computes the multipole expansion for the parent cell. The M2L op-

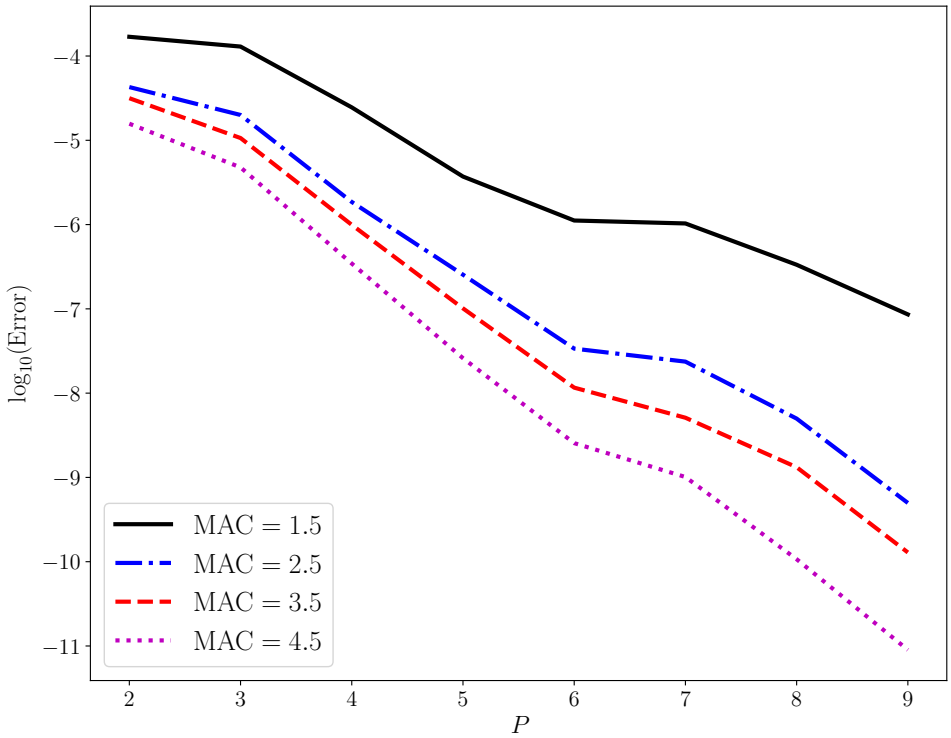

Figure 8. The error in the FMM multipole approximation as a function of both MAC and order, P. We can see that increasing the MAC improves the error directly and also increases the slope of the error curve as a function of P. Also, the error decreases to a negligible level for reasonably small order.

erator computes the local expansion at an observation cell from the multipole expansion at a source cell. The L2L operator computes the local expansions for each child cell from the local expansion for a parent cell. The L2P operator computes the full sum from the local expansion in a given observation cell. Finally, the P2P operator computes the exact interaction between a source and an observation cell. We use an extra operator that is not normally included in the FMM, the M2P operator, which computes the sum at an observation cell from a source multipole expansion. This M2P operator is equivalent to L2P(M2L), but can be more efficient when the number of points in the observation cell is small [*Yokota*, 2013].

The performance of the FMM depends heavily on well constructed hierarchical trees. We follow a variant of the dynamic octree construction and traversal methods from *Yokota* [2013] that uses an octree where each cell is represented by a sphere rather than a box. The root cell is sized to tightly fit all the points in a sphere. Tree cells are then recursively divided at their center until no cell has more than $N_{max} \gg P$ points. At each cell division, the new

tree cells are resized to tightly fit the points they contain. Compared to a static octree, this approach maximizes the portion of the interactions that can be evaluated with the multipole approximation.

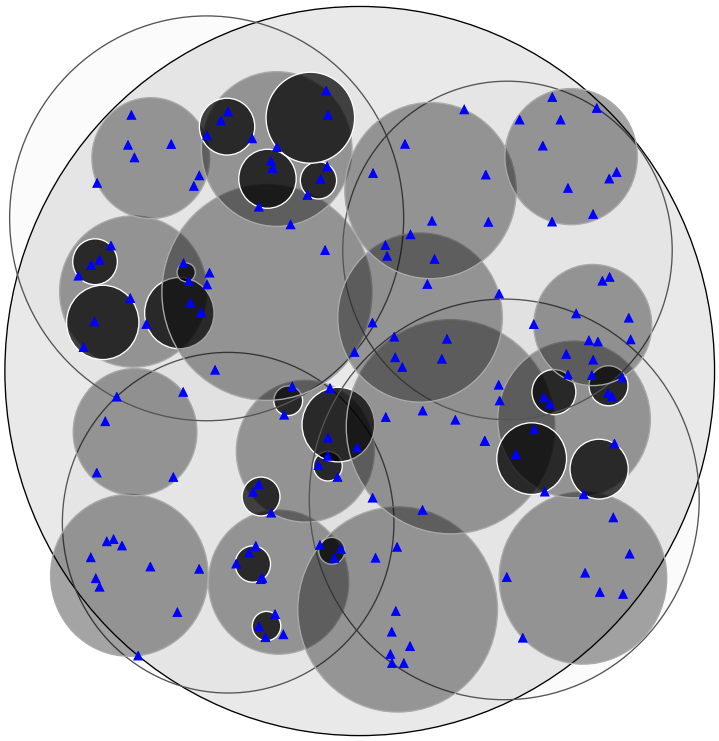

Figure 9. An example of a spherical octree for the FMM. The circles are the cells in the tree. Each successive level down the tree, the cells are shaded more darkly. The points contained within the tree are shown as blue triangles. Note that the tree construction algorithm is adaptive and only divides a parent cell if there are sufficient points within that parent cell. In this division only occurs if there are more than 10 points within a cell. In some areas, the cells overlap. This does not present a problem to the FMM algorithm and every point is arbitrarily assigned to one of the cells.

After this tree is constructed, a upwards tree traversal over the source tree is performed from the leaves to the root, calculating the multipole expansion at each cell using the P2M and M2M operators. Next, a dual tree traversal over the source and observation tree is performed. FIGURE describes this traversal in pseudocode. Finally, a downward tree traversal

over the observation tree is performed from the root to leaves, calculating the local expansion at each non-leaf node with the L2L operator and using the L2P operator at the leaves to calculate the final sum. This particular variant of the FMM is particularly straightforward to implement and has been demonstrated to be extremely efficient.

Algorithm 1 Dual tree traversal

```
procedure DTT(CELL S, CELL T)
    if S.Radius + T.Radius < MAC * Distance(S.Center, T.Center) then
        if S.NumPoints < PointThreshold then
            M2P(S, T)
        else
            M2L(S, T)
    else if S.IsLeaf and T.IsLeaf then
        P2P(S, T)
    else if S.Level ≤ T.Level then
        for C in S.Children do
            DTT(C, T)
    else
        for C in T.Children do
            DTT(S, C)
```

6 Linear system solution

Tectosaur is designed as the sum of three components: 1) nearfield matrix assembly, 2) farfield fast multipole method, 3) linear system solution. We have already discussed the construction of the nearfield matrix and the fast multipole method in the previous two sections. Here, we describe the iterative approaches we use to solve the linear systems resulting from the BEM discretization.

We solve most linear systems using iterative Krylov subspace methods because the large size makes direct inversion impractical. For example, due to the $O(n^3)$ scaling of matrix inversion, directly inverting a 1,000,000 $\times$ 1,000,000 matrix would require months of computing power even on a supercomputer. On the other hand, performing an FMM accelerated BEM matrix vector product can take less a second on a large shared memory machine

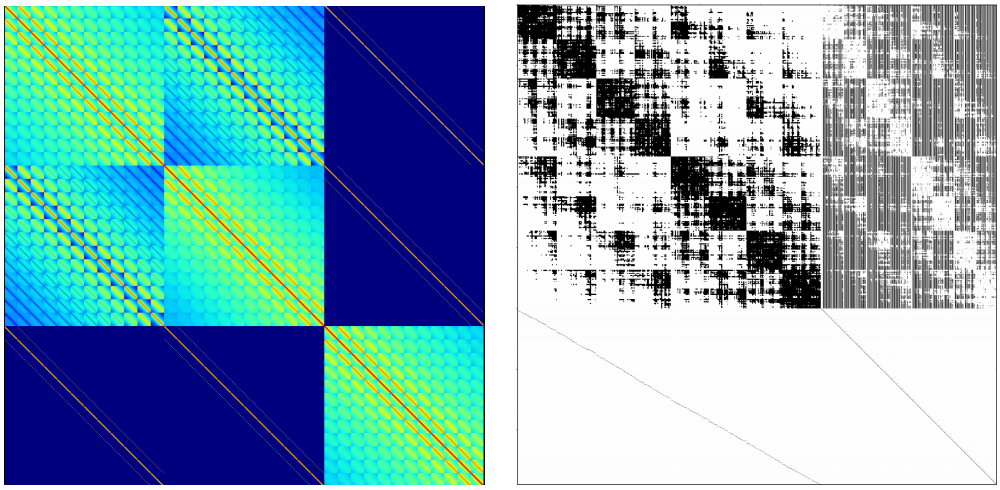

(a) (b)

Figure 10. a) The logarithm of the value of the entries in the dense matrix produced by the BEM for the T^* kernel for a triangulated planar mesh. Note the 3x3 grid of matrix blocks. These correspond to the three components of the displacement input. b) The sparsity pattern for a sparse matrix produced by approximating the same BEM matrix using the fast multipole method.

and a linear solve might require 10-100 matrix vector products, resulting in a total time to solution of less than two minutes. For symmetric and positive definite matrices, we use the conjugate gradient method for optimal speed of convergence. However, many BEM matrices are not positive definite. In that case, we use GMRES [*Saad and Schultz*, 1986]. When helpful, we precondition with the ILU implementation in SuperLU.

The matrices produced by calculating the BEM terms (19) lack global information about the displacement and traction fields. In particular, the displacement field should be continuous across element boundaries and drop to zero at the edges of a mesh. The displacement field should also have a jump in value anywhere a surface intersects a fault. Critically, this means that anywhere that a fault intersects another surface (e.g. the Earth's surface), the two meshes must be conforming, sharing triangle vertices and edges.

We use a general purpose system for handling any linear constraint of the form:

$$\sum_i c_i x_{J(i)} = r \tag{36}$$

where c_i are some coefficients multiplying the elements of the solution vector, $x_{J(i)}$, indexed by $J(i)$, and r is the non-homogeneous component of the constraint. There are many approaches for imposing constraints like these on a BEM (or finite element method) matrix.

We use an approach called Guyan reduction [*Guyan*, 1965]. The key idea is to rearrange the constraints so that the degrees of freedom (DOFs) are grouped into "master" DOFs and "slave" DOFs, where the slave DOFs can be calculated once the master DOFs are known. After doing this rearrangement, we can write the original DOFs in terms of a reduced set of DOFs as:

$$x = Cy + R \tag{37}$$

$$\text{where} \qquad x = \begin{bmatrix} x_u \\ x_m \\ x_s \end{bmatrix} \qquad C = \begin{bmatrix} I & 0 \\ 0 & I \\ 0 & \widetilde{C} \end{bmatrix} \qquad y = \begin{bmatrix} x_u \\ x_m \end{bmatrix} \qquad R = \begin{bmatrix} 0 \\ 0 \\ \hat{R} \end{bmatrix} \tag{38}$$

and x_u are the unconstrained DOFs, x_m are the master DOFs, x_s are the slave DOFs, $\widetilde{C}$ is the matrix of constraint coefficient derived by rearranging the constraints of form (36), and $\widetilde{R}$ is the vector of constraint inhomogeneities. Once the constraints are assembled in this form, by substituting in the constrained representation for x, we can solve the reduced and constrained linear system.

$$C^T A(Cy + R) = C^T b \tag{39}$$

The premultiplication by C^T is necessary for the process to result in a square matrix true to the original problem and can be thought of similarly to the way in which rows and columns are both added to a matrix when constraining it using Lagrange multiplier.

To make this process more concrete, we discuss a small example involving a free surface with four elements and a fault with known slip on one element that intersects that free surface. The form of the continuity and slip constraints are in Figure 11.

The first and third constraints are already in a form where a slave DOF is identifiable. The second constraint should be rearranged to:

$$u_3 = u_4 + s_1 \tag{40}$$

Returning to the example problem. Now, we choose u_1, u_3, u_5 as the slave DOF and build a constraint matrix that maps between the small vector of unconstrained and master DOFs and the full vector.

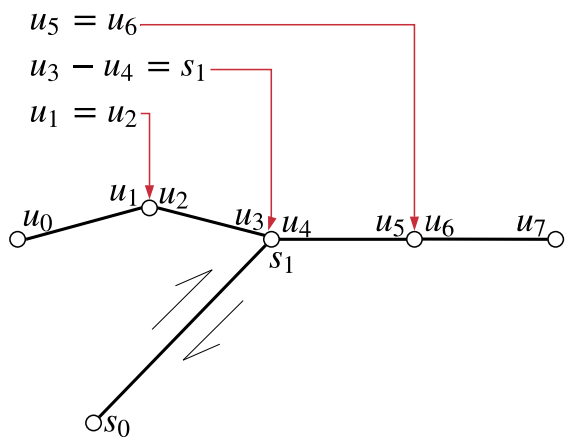

Figure 11. A small example two-dimensional boundary element mesh with the degrees of freedom identified and the constraints specified. The first and third constraints maintain displacement continuity at the junction between elements. The second constraint enforces the slip condition where the fault intersects the surface. Note that the fault intersects the surface at a vertex in the surface mesh. This is necessary to properly enforce the slip condition.

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \\ u_6 \\ s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_2 \\ u_4 \\ u_6 \\ s_1 \\ s_2 \end{bmatrix} \tag{41}$$

Guyan reduction allows us to treat constraints separately from constructing the main unconstrained BEM matrix. This allows substantially more modularity in the software design and enables a matrix free approach to the farfield computation. Furthermore, Guyan reduction avoids many of the matrix ill-conditioning issues that arise when using Lagrange multipliers or penalty methods.

7 Implementation

Tectosaur is available at `https://github.com/tbenthompson/tectosaur`, where there are several concrete examples of usage. Tectosaur is implemented in Python, C++,

CUDA and OpenCL. This mix of tools allows for rapid development in Python with the efficiency of highly parallel compiled CPU or GPU code when needed. We develop inefficient, but correct, code first in Python. Then we profile to determine the bottlenecks and migrate those portions of the code to C++ or CUDA. We GPU accelerate with both CUDA and OpenCL backends. All optimizations are done for the CUDA backend. However, the less efficient and poorly optimized OpenCL backend enables developing and testing on machines without NVidia GPUs, including most laptop computers. Most of the differences between CUDA C and OpenCL can be accomodated by simply replacing one set of syntax for the other. We hide these differences behind a layer of C macros.

The resulting implementation is highly efficient. We have run benchmarks on an Intel Xeon E7-8891 v4 CPUs with 10 cores and a NVidia P100 GPU. In Figure 12, we show the time required per matrix vector product with a BEM matrix produced by the T^* kernel on this machine as a function of the number of elements for both Tectosaur using the FMM and Tectosaur using a direct farfield calculation. Both methods are using heavily optimized CUDA code. We first notice that for small numbers of elements, the direct approach is faster, but both approaches are taking less than a fiftieth of a second. The FMM approach also shows some stochastic overhead at these small problem sizes due to parallelization overhead. Second, for more than approximately 1000 elements, the FMM is faster due to it's linear slope compared to the quadratic slope of the direct method. Third, for both methods, the time required is sublinear in the number of elements until approximately 10,000 elements. This is due to the parallelization overhead of approximately 0.05 seconds. An implementation could be developed that was more tailored to these smaller problems. However, we have focused more on medium to large problems of 100,000 to 1,000,000 elements.

Focusing on the largest problem in Figure 12, we see that running one matrix vector product with the matrix produced on a one million element mesh requires 2.1 seconds or 476 billion triangle-triangle interactions per second. In comparison, the direct method required 461 seconds. Thus, the FMM is providing a speedup for 230x.

Much of this speedup comes from the reduction in total number of operations from the FMM tree structure. However, a significant portion of the speedup also comes from the use of a matrix free algorithm as opposed to storing the matrix in memory. We can't compare directly to a stored dense matrix implementation due to memory limitation. Each triangle in this mesh has nine degrees of freedom because there are three basis functions per triangle

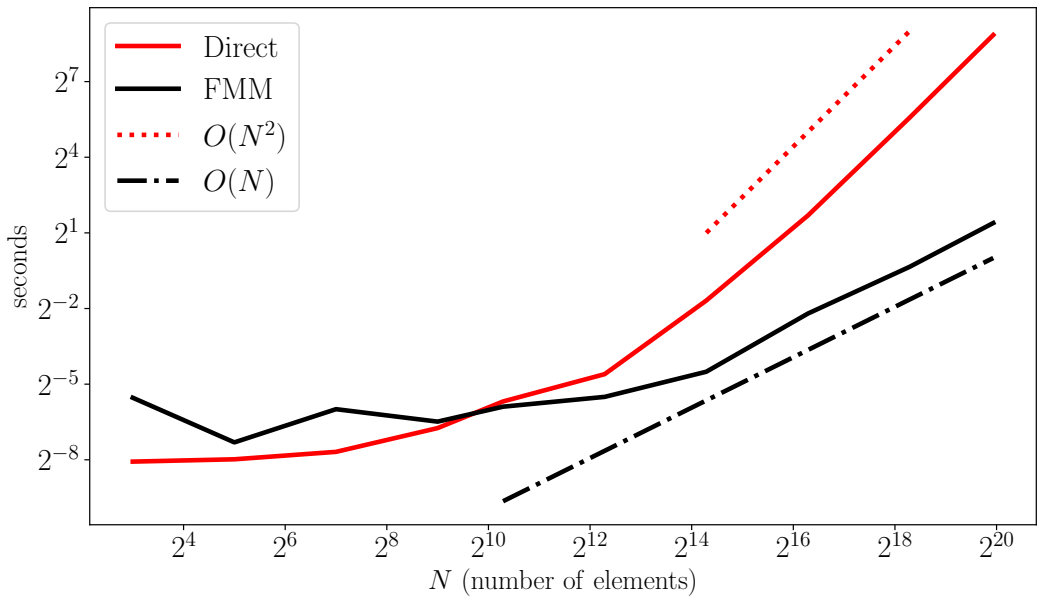

Figure 12. A log-log plot of runtime against number of elements for a single matrix vector product of the T^* BEM matrix. We show the performance of both a direct calculation and a calculation using the FMM. Dashed lines showing the slope for linear and quadratic scaling are included, demonstrated the linear growth in run time of the FMM.

and three components in a three-dimensional vector field. As a result, we are computing with a dense matrix with nine million rows and columns. To store such a matrix at 32-bit floating point precision would require 324 terabytes of memory. Storing a matrix that size would require thousands of nodes on some of the largest supercomputers. Even if we were able to store the matrix, we would be severely limited by memory bandwidth. On most modern computer hardware, about four floating point operations can be performed in the same time that one byte is loaded from RAM. As a consequence, when the entries of a matrix can be recalculated with only a few operations, it can be much more efficient to avoid storing the entries and recalculate them every time they are needed. With such a "matrix-free" method, the primary limitation will be floating point operations and as a result these methods are normally well suited to GPU acceleration. Because our FMM implementation is flop-limited, we are able to reach 5.3 teraFLOPS on one NVidia P100 GPU.

8 Planar fault slip

We first study a fault slip problem that we can benchmark against the well known analytic solution for the slip on a rectangular dislocation in a half space [*Okada*, 1992]. We model a planar vertical strike slip fault beneath a planar free surface. The free surface is finite but many extends away from the fault for NNN fault lengths. While, we will compare with the infinite free surface (halfspace) from the analytic solution, the difference in the solution due to the finite free surface is small. We impose a cosine slip pulse at the upper edge of the fault (see Figure 13, and solve the displacement on the free surface. The fault is discretized in a 48 x 48 grid of rectangles. These rectangles are used directly in the Okada dislocation approach to calculate surface displacements. For the BEM, the rectangles are further split into two triangles. The free surface is also triangulated in the BEM case and extends five fault lengths from the center. The free surface has 115,200 elements.

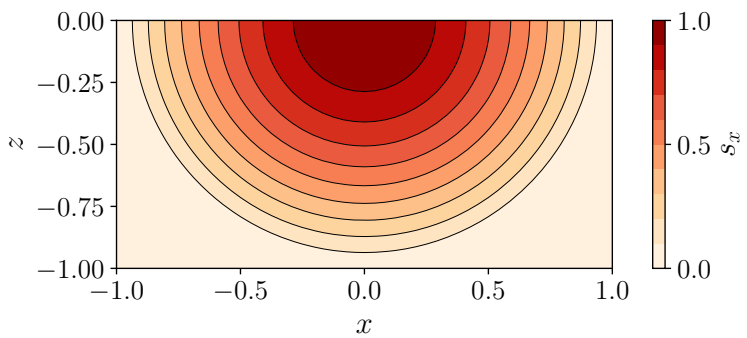

Figure 13. The input strike slip field on the fault.

In Figure 14, we present a comparison of the displacement field on the free surface between the BEM solution and the analytic Okada dislocation solution. Both the BEM and Okada approach suffer from error in discretizing the slip pulse. The BEM approach discretizes the smooth slip pulse onto the linear basis functions of the triangular fault mesh, losing some of the curvature. The Okada approach discretizes onto a set of constant basis functions on each rectangle. As a result, the discretization error for the slip pulse in the Okada approach is substantially larger than for the BEM approach. We exclude a single element on either side of the fault in the comparison in Figure 14 because the step function in slip between each adjacent element creates severe error. On the other hand, the Okada approach computes an exact surface displacement field from its discretized slip field while the BEM incurs further error when calculating the surface displacement field. In Figure 15, we show

that the two solutions converge rapidly as the discretization of both the fault and the surface gets finer.

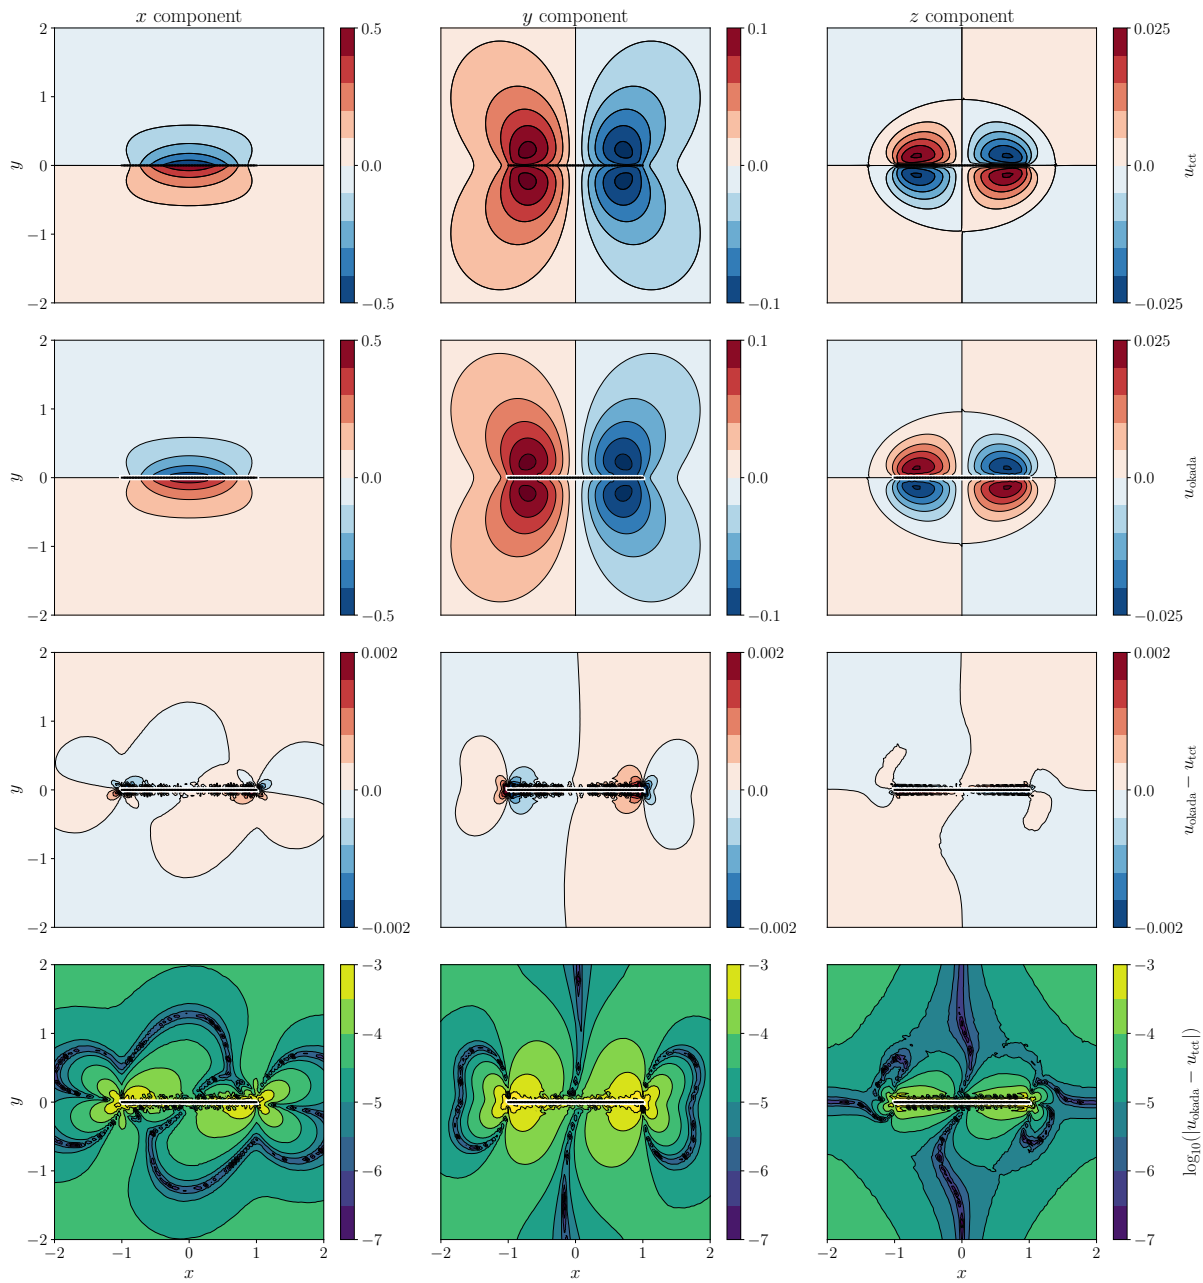

Figure 14. The first row shows the components of the displacement vector on the free surface resulting from the the slip field in Figure 13 as calculated by Tectosaur. The second rows shows the displacement vector as calculated by Okada dislocations. The third row shows the difference between the two calculations. The fourth row shows the base 10 logarithm of the difference between the two demonstrating that the differences are largest nearest the fault.

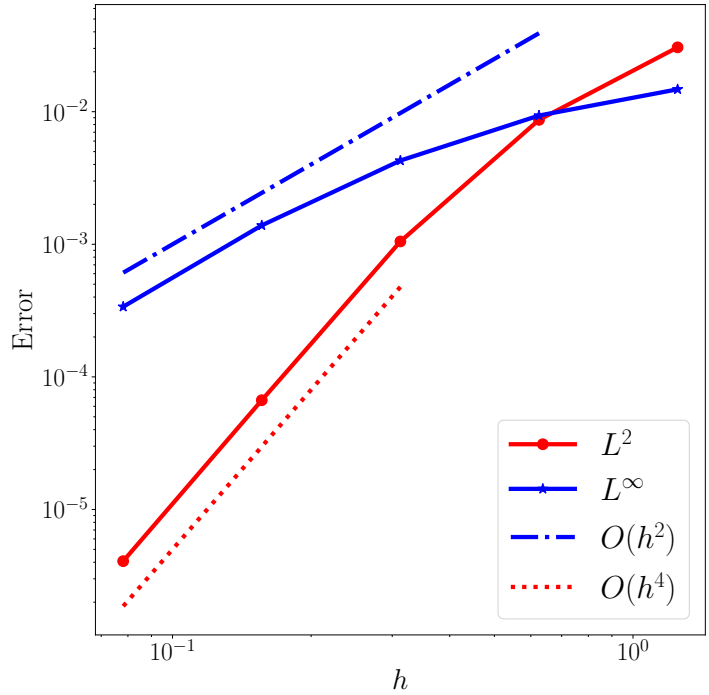

Figure 15. The convergence in both L^2 and L^∞ norm of the BEM-derived surface displacement as compared to the analytical Okada dislocation. Convergence is quadratic in the L^∞ norm and quartic in the L^2 norm.

9 Slip on a curved fault

As discussed in Section 2, dislocation elements fail to accurately model the stress field on the fault surface except under the special circumstance where a planar fault is tesselated with equal size rectangles. We demonstrate this failing and show that Tectosaur is able to accurately model stresses on a non-planar fault surface. We model strike slip on a curved fault dipping at 45 degrees. Figure 16 shows a map view of a low resolution mesh of the fault surface. The gaps and overlaps in the rectangular mesh demonstrates a major failing of using rectangles to mesh a non-planar surface even for displacements and stresses not on the fault surface. We do not compare to triangular dislocation elements because, as demonstrated in Figure 2, triangular dislocations fail to even model fault stresses on a planar fault.

To simplify the comparison, we model the fault surface in a full space. We choose a Gaussian strike slip pulse in the center of the fault. Then, we show in Figure 17 the resulting x-component shear stress field on the fault surface calculated with both the Tectosaur

and rectangular dislocations. On the rectangular dislocations, the shear stress is evaluated at the center of every element. While the two methods produce identical stress fields on a planar fault, on the curved fault, the dislocation approach produces an anomalous oscillation parallel to the direction of curvature. This oscillation has a magnitude six times greater than the true value of the shear stress. This oscillation is due to the effects of the singularities present at the edges of every dislocation element. Because the edges of adjacent elements are no longer equidistant to the center of a element as in the planar case, the effects of the edge singularities do not cancel out. As a result, the shear stress field produce by the constant dislocation approach is incorrect.

This geometry is designed to have a significant curvature and produce large errors from the rectangular dislocation approach. In some real geometries, the fault curvature might be substantially smaller compared to the discretization length scale and the stress field might appear closer to correct while still having large errors. Further, despite the erroneous results, the shear stress field produced by the constant slip dislocation approach is smooth. When embedded in a earthquake simulation, slip velocities would remain reasonable and the model would show typical stick-slip behavior. The effect is that the errors might go undetected unless the shear stress itself were carefully studied. So, this sort of modeling error can be particularly insidious and may be present in some published earthquake simulations on nonplanar fault geometries [*Shibazaki et al.*, 2012; *Li and Liu*, 2016; *Yu et al.*, 2018].

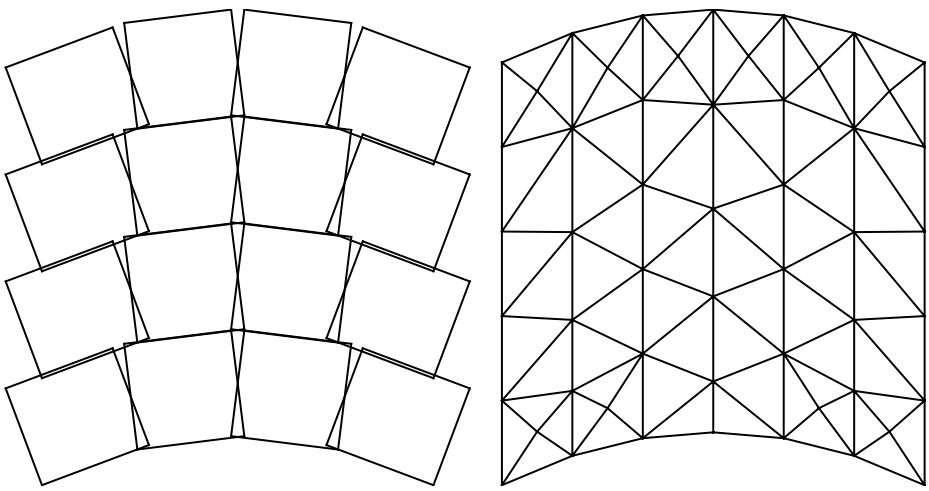

Figure 16. A map view comparison of a triangulation (right) and a rectangulation (left) for a curved, dipping fault. Note the gaps and overlaps in the rectangular mesh. The rectangles are chosen to minimize the size of these gaps and overlaps. For the shear stress calculations we perform, a much finer resolution mesh is used.

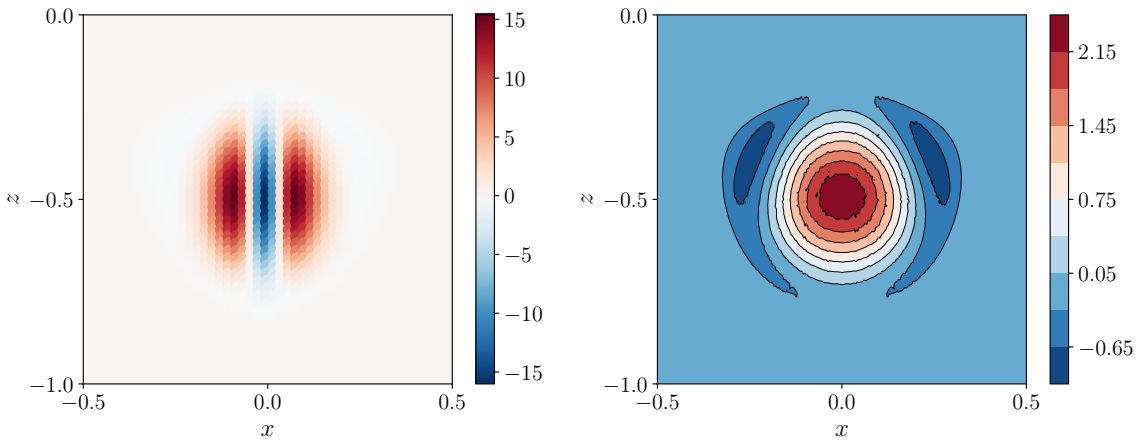

Figure 17. The along strike component of shear stress produced by a Gaussian strike slip pulse on a curved dipping fault (see Figure 16) as calculated by rectangular dislocations (left) and Tectosaur (right). The shear stress is projected into the x-z plane for plotting. The large negative shear stresses near $x = 0$ in the dislocation shear stress is clearly incorrect as it coincides with peak slip.

10 Quasidynamic earthquake simulation

Quasidynamic earthquake simulation is a powerful tool for investigating the frictional behavior of faults over many earthquake cycles without having to invest the numerical resources required for fully dynamic rupture modeling [*Rice*, 1993; *Liu and Rice*, 2005; *Thomas et al.*, 2014]. We implement quasidynamic simulation on arbitrary three-dimensional geometries using Tectosaur, enabling the study of real fault geometries rather than planar analogues.

The quasidynamic approximation is a first order approximation of inertial wave effects with a "radiation damping" term. The quasidynamic shear stress on the fault surface is

$$\tau_{\text{qd}} = \tau_{\text{static}} - \frac{\mu}{2c_s} V \tag{42}$$

where μ is the shear modulus, c_s is the shear wave speed, and V is the local fault slip velocity. The advantage of this approximation is that the shear stress can be calculated using static elastic numerical methods and then adjusted by the slip velocity.

To complete the system, we need to a friction law that relates shear stress to slip velocity. A common framework is rate-state friction where the strength of friction is related to both the rate of slip and a state variable the evolves during fault slip. The state variable evolution law can take various forms. Here, we present the aging law. Then, rate-state friction

takes the form

$$\tau_{\mathrm{qd}} = a\sigma_n \sinh^{-1}\left(\frac{V}{2V_0}e^{\Psi/a}\right) \tag{43}$$

$$\frac{d\Psi}{dt} = \frac{bV_0}{D_c}\left(e^{(f_0-\Psi)/b} - \frac{V}{V_0}\right) \tag{44}$$

where σ_n is the normal stress, Ψ is the state variable, f_0 is a the friction coefficient at a steady state slip velocity of V_0, a and b are dimensionless parameters determining the strength of velocity and state changes respectively on the evolution of friction and D_c is the state evolution length scale. In our implementation, fault slip is always parallel to the shear stress vector and can be in any direction on the fault plane.

Several approaches have been used for quasidynamic earthquake cycle simulation. Fourier domain convolution methods are extremely efficient for planar faults with a uniform discretization [*Rice*, 1993; *Lapusta et al.*, 2000]. However, these methods break down in the face of any non-planarity. Similarly, boundary element methods using rectangular or triangular dislocations [*Liu and Rice*, 2005; *Segall and Bradley*, 2012] have the aforementioned stress singularity issues, especially on nonplanar faults. Both Fourier and dislocation approaches struggle with rheologies beyond linear elasticity. In comparison, finite difference methods or finite element methods can successfully model a much wider range of rheologies including non-uniform material properties [*Erickson and Dunham*, 2014] or viscoplasticity [*Allison and Dunham*, 2018]. However, finite difference methods still have difficulty with complex nonplanar fault geometries. In comparison, our boundary element methods can model arbitrary nonplanar fault geometries.

We build a quasidynamic earthquake cycle simulator on top of Tectosaur. We track the current slip deficit and state variable at every degree of freedom on the fault surface. Then, at each time step, we calculate the traction on the fault surface from the slip deficit field using Tectosaur to solve the static elastic equations. The friction equations are then solved for the current slip velocity using Equation (43). The state derivatives are calculated using Equation (44). Finally, having both the slip deficit derivatives and state derivatives, we integrate in time. The algorithm is flexible to the particular time integration method. A popular method has been to use a time step dependent on the fastest slip velocity on the fault [*Lapusta and Liu*, 2009]. We follow *Erickson and Dunham* [2014] in using an adaptive Runge-Kutta algorithm.

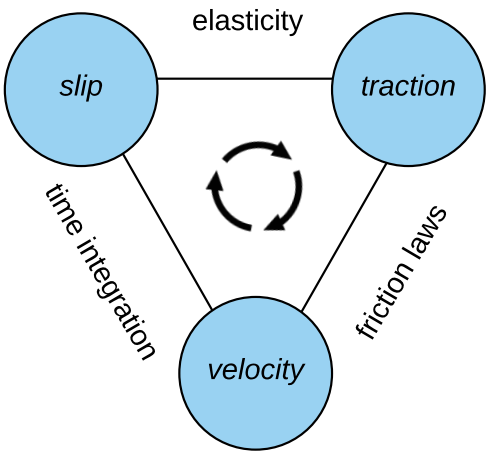

Figure 18. A schematic demonstrating the tools required to implement a rate-state quasidynamic earthquake simulator. An elastic solver is required to calculate the traction from the slip on the fault. The friction laws are solved to calculate the slip velocity from the fault traction. Finally, the velocities are integrated in time to compute the new slip field.

Beyond the typical quasidynamic implementation, there are some issues that arise when using a linear basis for the elastic solver. First, at a corner or sharp bend in the mesh, the normal vector on the sides of the bend or corner will be different. This implies that degrees of freedom for different triangles that are located at the same point in the mesh will almost certainly have different traction values. This is not a problem in and of itself and we ensure that these different traction values are consistent with a single underlying stress field. However, when the friction law is solved with two different traction values, we produce two different slip velocity values for the same point in the mesh. This violates the continuity of displacement implying that rate-state friction requires a C^1 mesh representation and is fundamentally ill posed at a sharp corner. We solve this issue by simply averaging the multiple velocities at a single mesh point. However, future research should explore this issue more.

We also suffer from the need to directly model a free surface as opposed to dislocation methods that already implicitly account for a half space. To solve the elastic equations exactly at each time step would require a full linear solve involving the T^* integral term to determine the surface displacement field. This linear solve is expensive, requiring many iterations with a Krylov subspace method. For most problems, we use the GMRES algorithm [*Saad and Schultz*, 1986]. By comparison, in a full space, only a single matrix vector product

is required to calculate the traction from the slip. To avoid the full linear solve, we note that the surface displacement field evolves very slowly compared to fault slip and the surface displacement for one time step is extremely close to the surface displacement for the next time step. As a result, we adopt a method whereby we perform a single iteration of a Jacobi iterative solver for the surface displacement at each time step. Because the surface displacement varies slowly, this approach is accurate while reducing the computational expense to a single matrix vector product.

Here, we first demonstrate our method on a square planar fault in a full space. We focus on this simple example to demonstrate that our approach is fundamentally sound. In a companion paper, we make use of the tools developed here to analyze a model with a real 3D geometry of the Cascadia subduction zone. In this example, the fault mesh is a two meter by two meter square covered by 200 x 200 grid of points and triangulated into 79,202 triangles. We initialize the fault with zero slip deficit, and a plate rate of 31.5 mm/yr. As a result, the entire fault is locked at the outset of the simulation. We choose $\mu = 3x10^{10}$ Pa, $\nu = 0.25$, $\rho = 2670\,\mathrm{kg/m^3}$, $a = 0.01$, $b = 0.015$, $V_0 = 10^{-6}$ m/s, $f_0 = 0.6$, $D_c = 0.000002$ m, and an ambient normal stress of $5x10^7$ Pa. We model the evolution of the fault for 40,000 time steps or 28.7 days. In Figure 19, we can see that the time step size varies over eight orders of magnitude. This emphasizes the adaptive time integration as a critical component of the algorithm [*Erickson and Dunham*, 2014].

In Figure 19, we plot the maximum slip rate over time on as well as the minimum state. The minimum state track the maximum velocity closely because when slip rates highest, the state variable is driven lower by the velocity weakening friction law. Because, we initiated the fault in a zero stress state, for the first 16 days, the fault is mostly locked and stress accumulates. As can be seen in the maximum slip rate, there are slow slip events during this time period that grow in size and slip rate each time they occur. We see a slip rate field that is mostly locked with some creep around the edges of the fault (Figure 20a). Eventually, at approximately day 16, the fault has gained enough stress to have its first rupture. The rupture nucleates once a sufficiently fast slow slip event reaches the critical slip patch size (Figure 20c) and then propagates across the whole fault (Figure 20d). Not all the stress is released, resulting in another rupture two days later after some small slow slip events (Figure 20b) with slip rates two to five times the plate velocity. The cycle continues with almost perfectly periodic earthquakes every two days.

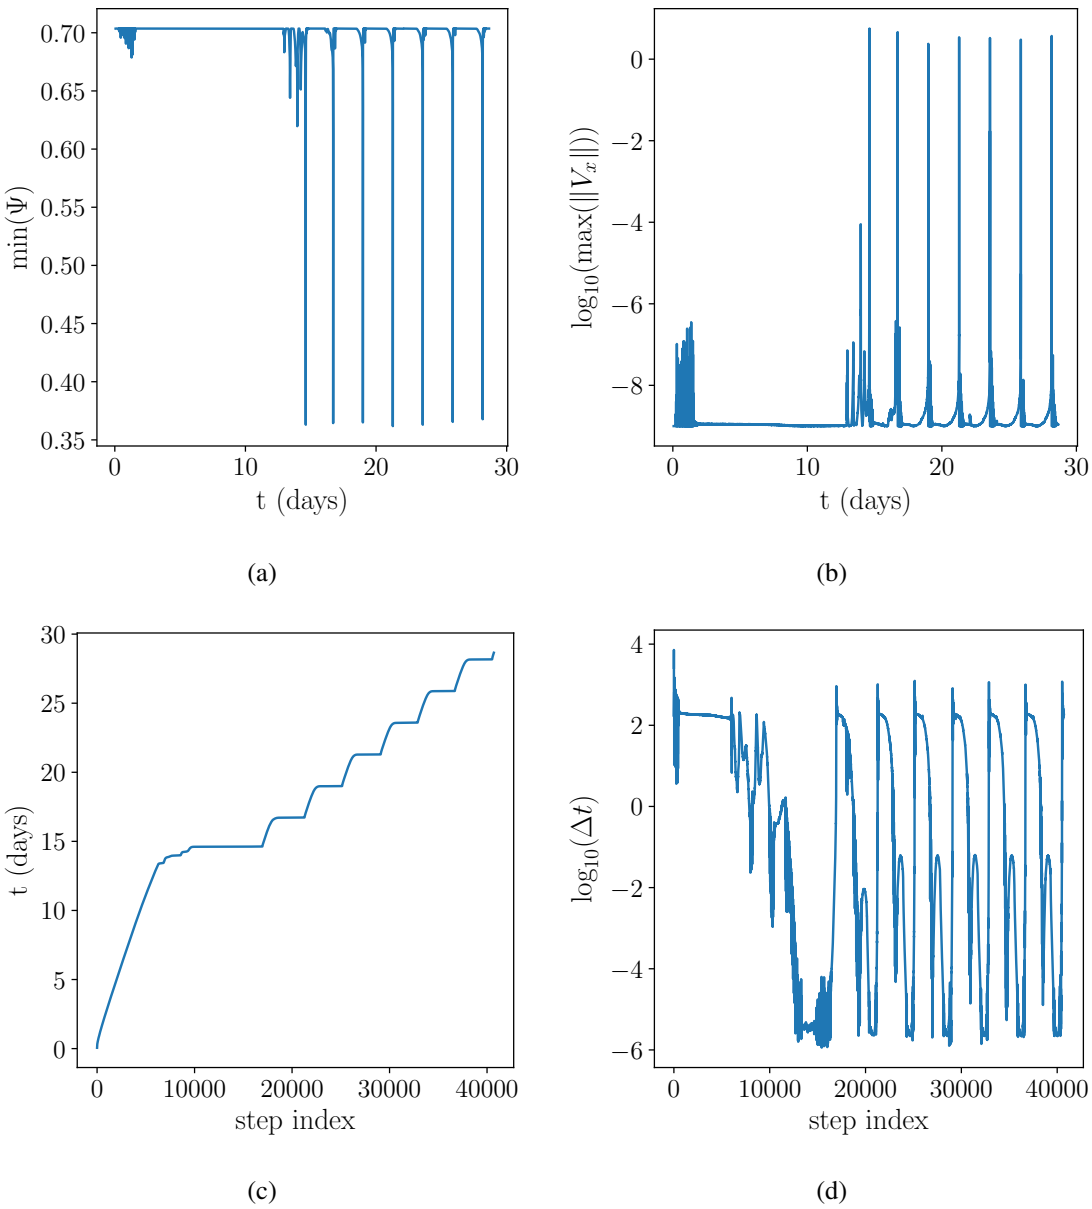

Figure 19. a) The minimum state parameter on the fault as a function of time. b) The maximum slip rate on the fault as a function of time. c) The model time as a function of time step index. d) The base ten logarithm of the time step size as a function of the time step index.

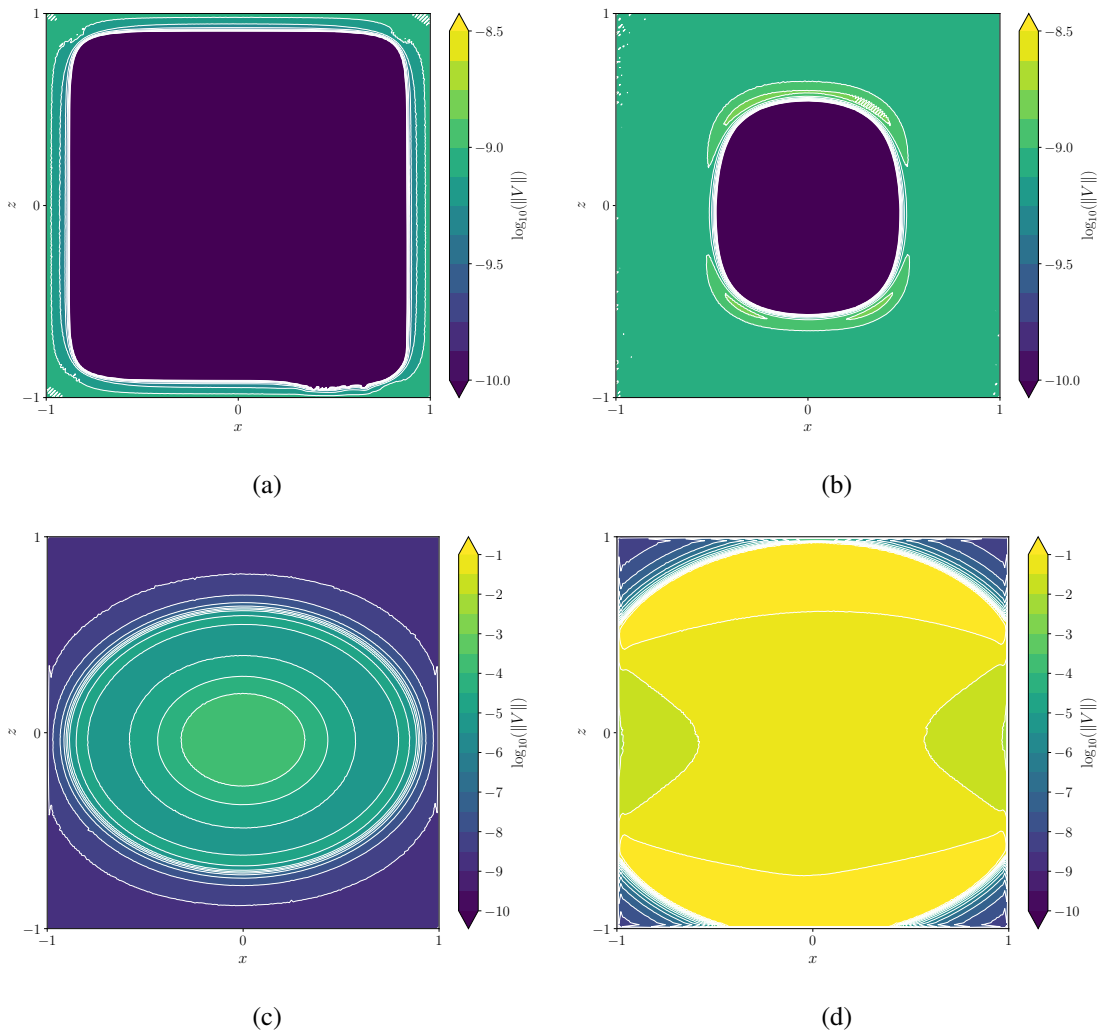

Figure 20. Four snapshots of the evolution of slip rate on a fault in a fullspace. The base ten logarithm of slip rate is plotted because slip rate varies by more than ten orders of magnitude between interseismic locked periods and ruptures reaching up to 0.1 m/s. Note that the upper two plots during the interseismic phase use a different color scale than the lower two plots showing rupture phase velocities. a) The interseismic period with most of the fault locked. b) A slow slip event during the interseismic period. c) The nucleation of a rupture. Rupture nucleation occurs once a critically large region begins slipping much faster than the plate rate. Here, we can see most of the the fault is slipping between 1 μm/s and 1 mm/s with slip rate reaching 100 mm/s in the rupture itself. d) The rupture propagating up the fault. The lower half of the fault has already ruptured and has slowed to postseismic creep rates.

11 Conclusions and outlook

In a field where the primary object of interest (the fault) is a boundary, it is natural to expect that a numerical method that focuses on the boundary would be ideal. Such methods

have dominated the fields of slip inversion and earthquake cycle modeling for decades. However, the next step will need to be capable of analyzing the wide range of complex behavior caused by realistic geometries in fault systems. Tectosaur enables studying that behavior efficiently and at high resolution.

The BEM in its simple form is limited to solving static linear elastic problems with piecewise constant elastic properties, but is able to do so for extremely complex boundary geometries. While that still enables a huge range of research, it is worth considering the long-term potential of the method. There are finite difference or volumetric integral equation approaches that allow a non-conforming volumetric mesh of the interior when combined with a BEM-based surface solution [*Mayo*, 1984; *Biros et al.*, 2004]. As a result, we can view the BEM as a very powerful approach for applying boundary conditions that enables mesh-free treatment of the interior for problems including wave propagation, dynamic rupture, nonlinear plasticity and smoothly varying elastic property variations. From this perspective, Tectosaur is one component of a larger toolbox that will be able to study almost any earthquake science modeling problem without ever constructing a volumetric mesh.

Acknowledgments

The data and source code for this work is available at https://github.com/tbenthompson/tectosaur. T. Ben Thompson appreciates the support of the Department of Energy Computational Science Graduate Fellowship.

References

Allison, K. L., and E. M. Dunham (2018), Earthquake cycle simulations with rate-and-state friction and power-law viscoelasticity, *Tectonophysics*, *733*, 232–256.

Balakrishna, C., L. Gray, and J. Kane (1994), Efficient analytical integration of symmetric galerkin boundary integrals over curved elements; elasticity formulation, *Computer methods in applied mechanics and engineering*, *117*(1-2), 157–179.

Bebendorf, M., and S. Rjasanow (2003), Adaptive low-rank approximation of collocation matrices, *Computing*, *70*(1), 1–24.

Biros, G., L. Ying, and D. Zorin (2004), A fast solver for the stokes equations with distributed forces in complex geometries, *Journal of Computational Physics*, *193*(1), 317–348.

Bonnet, M. (1995), Regularized direct and indirect symmetric variational bie formulations for three-dimensional elasticity, *Engineering Analysis with Boundary Elements*, *15*(1), 93–102.

Bradley, A. M. (2014), Software for efficient static dislocation–traction calculations in fault simulators, *Seismological Research Letters*, *85*(6), 1358–1365.

Crouch, S. (1976), Solution of plane elasticity problems by the displacement discontinuity method. i. infinite body solution, *International Journal for Numerical Methods in Engineering*, *10*(2), 301–343.

Crouch, S. L., and A. Starfield (), *Boundary element methods in solid mechanics: with applications in rock mechanics and geological engineering*.

Duffy, M. G. (1982), Quadrature over a pyramid or cube of integrands with a singularity at a vertex, *SIAM journal on Numerical Analysis*, *19*(6), 1260–1262.

Erickson, B. A., and E. M. Dunham (2014), An efficient numerical method for earthquake cycles in heterogeneous media: Alternating subbasin and surface-rupturing events on faults crossing a sedimentary basin, *Journal of Geophysical Research: Solid Earth*, *119*(4), 3290–3316.

Frangi, A., G. Novati, R. Springhetti, and M. Rovizzi (2002), 3d fracture analysis by the symmetric galerkin bem, *Computational Mechanics*, *28*(3-4), 220–232.

Greengard, L., and V. Rokhlin (1987), A fast algorithm for particle simulations, *Journal of computational physics*, *73*(2), 325–348.

Guyan, R. J. (1965), Reduction of stiffness and mass matrices, *AIAA journal*, *3*(2), 380–380.

Hackbusch, W., and Z. P. Nowak (1989), On the fast matrix multiplication in the boundary element method by panel clustering, *Numerische Mathematik*, *54*(4), 463–491.

Hudnut, K. W., Y. Bock, M. Cline, P. Fang, Y. Feng, J. Freymueller, X. Ge, W. Gross, D. Jackson, M. Kim, et al. (1994), Co-seismic displacements of the 1992 landers earthquake sequence, *Bulletin of the Seismological society of America*, *84*(3), 625–645.

Jeyakumaran, M., J. Rudnicki, and L. Keer (1992), Modeling slip zones with triangular dislocation elements, *Bulletin of the Seismological Society of America*, *82*(5), 2153–2169.

Kelvin, L. (1848), Note on the integration of the equations of equilibrium of an elastic solid, *Cambridge and Dublin Mathematical Journal*, *3*, 87–89.

Lapusta, N., and Y. Liu (2009), Three-dimensional boundary integral modeling of spontaneous earthquake sequences and aseismic slip, *Journal of Geophysical Research: Solid Earth*, *114*(B9).

Lapusta, N., J. R. Rice, Y. Ben-Zion, and G. Zheng (2000), Elastodynamic analysis for slow tectonic loading with spontaneous rupture episodes on faults with rate-and state-dependent friction, *Journal of Geophysical Research: Solid Earth*, *105*(B10), 23,765–23,789.

Li, D., and Y. Liu (2016), Spatiotemporal evolution of slow slip events in a nonplanar fault model for northern cascadia subduction zone, *Journal of Geophysical Research: Solid Earth*, *121*(9), 6828–6845.

Li, S., M. Mear, and L. Xiao (1998), Symmetric weak-form integral equation method for three-dimensional fracture analysis, *Computer Methods in Applied Mechanics and Engineering*, *151*(3-4), 435–459.

Liu, Y., and N. Nishimura (2006), The fast multipole boundary element method for potential problems: a tutorial, *Engineering Analysis with Boundary Elements*, *30*(5), 371–381.

Liu, Y., and J. R. Rice (2005), Aseismic slip transients emerge spontaneously in three-dimensional rate and state modeling of subduction earthquake sequences, *Journal of Geophysical Research: Solid Earth*, *110*(B8).

Liu, Y., and A. M. Rubin (2010), Role of fault gouge dilatancy on aseismic deformation transients, *Journal of Geophysical Research: Solid Earth*, *115*(B10).

Luo, Y., and J.-P. Ampuero (2018), Stability of faults with heterogeneous friction properties and effective normal stress, *Tectonophysics*, *733*, 257–272.

Mayo, A. (1984), The fast solution of poissonâĂŹs and the biharmonic equations on irregular regions, *SIAM Journal on Numerical Analysis*, *21*(2), 285–299.

Meade, B. J. (2007), Algorithms for the calculation of exact displacements, strains, and stresses for triangular dislocation elements in a uniform elastic half space, *Computers & Geosciences*, *33*(8), 1064–1075, doi:10.1016/j.cageo.2006.12.003.

Nedelec, J. (1982), Integral equations with non integrable kernels, *Integral equations and operator theory*, *5*(1), 562–572.

Nie, X.-C., L.-W. Li, and N. Yuan (2002), Precorrected-fft algorithm for solving combined field integral equations in electromagnetic scattering, *Journal of Electromagnetic Waves and Applications*, *16*(8), 1171–1187.

Nikkhoo, M., and T. R. Walter (2015), Triangular dislocation: an analytical, artefact-free solution, *Geophysical Journal International*, *201*(2), 1119–1141.

Okada, Y. (1992), Internal deformation due to shear and tensile faults in a half-space, *Bulletin of the Seismological Society of America*, *82*(2), 1018–1040.

Perrin, G., J. R. Rice, and G. Zheng (1995), Self-healing slip pulse on a frictional surface, *Journal of the Mechanics and Physics of Solids*, *43*(9), 1461–1495.

Pham, A. D., S. Mouhoubi, M. Bonnet, and C. Chazallon (2012), Fast multipole method applied to symmetric galerkin boundary element method for 3d elasticity and fracture problems, *Engineering Analysis with Boundary Elements*, *36*(12), 1838–1847.

Qiu, Q., E. M. Hill, S. Barbot, J. Hubbard, W. Feng, E. O. Lindsey, L. Feng, K. Dai, S. V. Samsonov, and P. Tapponnier (2016), The mechanism of partial rupture of a locked megathrust: The role of fault morphology, *Geology*, *44*(10), 875–878.

Reilinger, R., S. Ergintav, R. Bürgmann, S. McClusky, O. Lenk, A. Barka, O. Gurkan, L. Hearn, K. Feigl, R. Cakmak, et al. (2000), Coseismic and postseismic fault slip for the 17 august 1999, m= 7.5, izmit, turkey earthquake, *Science*, *289*(5484), 1519–1524.

Rice, J. R. (1993), Spatio-temporal complexity of slip on a fault, *Journal of Geophysical Research: Solid Earth*, *98*(B6), 9885–9907.

Richards-Dinger, K., and J. H. Dieterich (2012), Rsqsim earthquake simulator, *Seismological Research Letters*, *83*(6), 983–990.

Saad, Y., and M. H. Schultz (1986), Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on scientific and statistical computing*, *7*(3), 856–869.

Sauter, S. A., and C. Schwab (2010), Boundary element methods, in *Boundary Element Methods*, pp. 183–287, Springer.

Segall, P., and A. M. Bradley (2012), Slow-slip evolves into megathrust earthquakes in 2d numerical simulations, *Geophysical Research Letters*, *39*(18).

Shibazaki, B., K. Obara, T. Matsuzawa, and H. Hirose (2012), Modeling of slow slip events along the deep subduction zone in the kii peninsula and tokai regions, southwest japan, *Journal of Geophysical Research: Solid Earth*, *117*(B6).

Simons, M., S. E. Minson, A. Sladen, F. Ortega, J. Jiang, S. E. Owen, L. Meng, J.-P. Ampuero, S. Wei, R. Chu, et al. (2011), The 2011 magnitude 9.0 tohoku-oki earthquake: Mosaicking the megathrust from seconds to centuries, *science*, *332*(6036), 1421–1425.

Sutradhar, A., G. H. Paulino, and L. J. Gray (2008), *Symmetric Galerkin Boundary Element Method*, Springer.

Thomas, A. L. (1993), Poly3D: A three-dimensional, polygonal element, displacement discontinuity boundary element computer program with applications to fractures, faults, and cavities in the Earth's Crust, Ph.D. thesis, Stanford University.

Thomas, M. Y., N. Lapusta, H. Noda, and J.-P. Avouac (2014), Quasi-dynamic versus fully dynamic simulations of earthquakes and aseismic slip with and without enhanced coseismic weakening, *Journal of Geophysical Research: Solid Earth*, *119*(3), 1986–2004.

Yokota, R. (2013), An fmm based on dual tree traversal for many-core architectures, *Journal of Algorithms & Computational Technology*, *7*(3), 301–324.

Yu, H., Y. Liu, H. Yang, and J. Ning (2018), Modeling earthquake sequences along the manila subduction zone: Effects of three-dimensional fault geometry, *Tectonophysics*, *733*, 73–84.

Zhang, L., T. Cui, and H. Liu (2009), A set of symmetric quadrature rules on triangles and tetrahedra, *Journal of Computational Mathematics*, pp. 89–96.