PROCEEDINGS, Thirty-Sixth Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, January 31 - February 2, 2011
SGP-TR-191

INVESTIGATION OF INJECTION-INDUCED SEISMICITY USING A COUPLED FLUID FLOW AND RATE/STATE FRICTION MODEL

Mark W. McClure and Roland N. Horne

Stanford Geothermal Program
Department of Energy Resources Engineering, 367 Panama Street
Stanford University, CA 94305-2220, USA
e-mail: mcclure@stanford.edu, horne@stanford.edu

ABSTRACT

This paper describes a numerical investigation of seismicity induced during injection into a single, isolated fracture. A model was developed and used that couples (1) fluid flow and (2) rate and state friction. Rate and state friction theory describes how friction on a fault depends on both sliding velocity and past sliding history. Rate and state friction is used widely in earthquake modeling. We investigated the effect of various factors, both geological and practical, that impact induced seismicity. Our modeling indicated that shear-induced pore dilation may be an important process that prevents injection from triggering slip events that propagate far from the injector. The effect of injection schedule was examined. It was found that decreasing injection pressure over time was a successful strategy for reducing the maximum event magnitude. The model predicted that significant seismicity should occur after injection ceases, which is consistent with observations from EGS stimulations. The post-injection events were caused by a redistribution of fluid pressure after injection. Production of fluid from the well immediately after injection inhibited post-injection events. In prior work, we investigated the impact of injection schedule using a simpler treatment of friction that we refer to as static/dynamic. The injection schedule findings described in this paper were consistent with our results from our static/dynamic modeling. Despite the apparent success of the static/dynamic modeling, it has important drawbacks, and we discuss some of those issues in this paper. Our eventual goal is to perform coupled fluid flow and rate and state simulation on a large network of fractures. We discuss various strategies we have implemented or are implementing to improve the efficiency of the simulations and make large scale simulations possible.

INTRODUCTION

Enhanced Geothermal Systems (EGS) are characterized by the use of hydraulic stimulation to enhance flow rate in high temperature, low productivity wells, typically located in crystalline basement rock. Water is injected at high pressure without proppant, and in most cases the increased fluid pressure triggers slip on pre-existing fractures. The process of using increased fluid pressure to trigger fracture slip is often referred to as "shear stimulation." When the fractures slip, their permeability is permanently enhanced, and well productivity can be improved by order of magnitude (Tester, 2007).

An important challenge for the deployment of EGS is that shear stimulation triggers microseismicity, very low magnitude seismic events that sometimes can be felt at the surface (Majer et al., 2007). Induced seismicity threatens the public acceptance of EGS, and the possibility of triggering a truly damaging seismic event, while seemingly remote, deserves careful consideration.

Shear stimulation directly impacts induced seismicity, the productivity of the EGS system, and the long term temperature decline, and so a practical need has developed for credible EGS stimulation modeling.

We have developed a model that couples (1) fluid flow and (2) rate and state friction for the purpose of simulating EGS stimulation. This paper summarizes the use of our model to investigate induced seismicity during injection into a single, isolated fracture.

The advantage of a model that couples rate and state friction and fluid flow is that it is a physically based technique that allows direct modeling of the frictional phenomena that cause induced seismicity.

Rate and state friction is widely used in earthquake modeling. It has been very successful at replicating

both laboratory observations of rock friction and large scale earthquake phenomenon (Dieterich, 2007, Segall, 2010). It relates friction to fracture sliding velocity and past sliding history, and provides a mechanism for both slow, aseismic slip and rapid, seismic slip.

Often in EGS modeling, friction is assumed constant, which makes it impossible to model seismicity. Seismicity occurs due to high rate slip caused by rapid weakening of friction. Another requirement for modeling seismic slip is that the heterogeneous stresses induced by slip must be described accurately. Seismicity is a process where slip nucleates at a given location and rapidly propagates due to induced stresses, and so accurate description of induced stresses is critical.

Some recent EGS modeling work has used a simplified treatment of friction in which sliding elements have a constant nonsliding resistance to slip, and once slip initiates an instantaneous drop in either friction (McClure and Horne, 2010) or stress (Baisch et al., 2010) is imposed. We refer to the approach in McClure and Horne (2010) as "static/dynamic friction."

Compared to static/dynamic friction, rate and state friction is a considerably more robust approach. Unlike static/dynamic it can model aseismic slip. Static/dynamic friction relies on the unrealistic assumption that all slip during a seismic event is simultaneous and instantaneous. Finally, in this paper we show how static/dynamic friction simulation can converge to an unrealistic result for an increasingly fine scale discretization.

We compared the modeling described in this paper to our earlier results from McClure and Horne (2010) and found that the results were similar in most respects. Despite the potential pitfalls of static/dynamic simulation, we conclude that because of its reasonable performance and relative efficiency, it may be appropriate for some applications.

Coupling between slip and fluid flow can occur for a variety of reasons. Slip can induce an increase in both fracture permeability and pore volume. Also, slip on one fracture can affect the normal traction on another, directly changing fluid pressure. Fluid flow affects slip because frictional strength depends on fluid pressure.

We examined the effect of shear-induced pore dilation, injection schedule, and the characteristic displacement scale d_c on the resulting seismicity.

Shear-induced pore dilation tended to damp out seismic events, resulting in more, lower magnitude events.

We studied a variety of factors related to injection schedule. A higher injection pressure resulted in a greater number of events. Decreasing injection pressure over time resulted in a significantly smaller maximum event magnitude than any other injection schedule. Seismicity was triggered at the periphery of the stimulated region after injection stopped due to redistribution of fluid pressure. Producing fluid from the well after injection inhibited post-injection seismicity. Our findings were consistent with Baisch et al. (2006) who predicted that post-injection seismicity should occur due to a redistribution of fluid pressure and that producing fluid after injection should inhibit post-injection seismicity. Our injection schedule results were consistent with our prior findings using static/dynamic friction in McClure and Horne (2010).

The effect of increasing characteristic displacement scale, d_c, was to damp out seismic slip, as predicted by rate and state theory.

Our eventual goal is to perform efficient coupled rate and state and fluid flow simulation on a complex network of multiple fractures. In the appendix we provide many of the techniques we are implementing to improve computational efficiency.

METHODOLOGY

Problem Definition

Our numerical model required simultaneous solution of six equations for six primary variables. The variables were velocity (v), state (θ), mass of fluid in a cell (m), shear traction (τ), normal traction (σ_n), and cumulative shear displacement at a location (D). The equations solved were unsteady-state fluid mass balance (with Darcy's law), frictional force equilibrium (with a radiation damping approximation term), two stress strain relationships that relate shear displacement to normal and shear traction, the aging law for state evolution, and the time integral relationship between slip velocity and cumulative shear displacement.

The problem was solved on a one-dimensional fracture embedded in a two-dimensional homogenous, isotropic medium. The code is capable of modeling multiple fractures, but it was used only for a single fracture in the work described in this paper. The two-dimensional stress/strain problems were solved using plane stress, which assumes the thickness of the medium in the third dimension is infinite. For some calculations, an infinite thickness fracture would lead unrealistic results. For example, an infinite width fracture would have an infinite flow rate. Therefore for calculations not involving stress and strain, the fracture width was set to be b_w. The

simulations were isothermal. The fluid was single phase liquid water.

The unsteady-state fluid mass balance equation in a fracture is (Aziz and Settari, 1979, with fracture aperture E replacing porosity):

$$\frac{\partial(\rho E)}{\partial t} = \nabla q + s \qquad (1)$$

where q is the mass flux rate, s is a source term, E is the void aperture (the pore volume per cross-sectional area of fracture), and ρ is the fluid density. Darcy flow is assumed. Mass flow across an area A in a direction x_i is (Aziz and Settari, 1979):

$$q = \frac{k\rho A}{\mu}\frac{\partial P}{\partial x_i} \qquad (2)$$

where P is fluid pressure, μ is fluid viscosity, and k is permeability. The permeability k is given by the "cubic law" and defined as (Jaeger et al., 2007):

$$k = \frac{e^2}{12} \qquad (3)$$

where e is hydraulic aperture, which is the effective aperture for flow in the fracture. Hydraulic aperture is equal to void aperture between two smooth plates, but can be lower than void aperture between rough surfaces as in a rock fracture. For flow in a one-dimensional fracture, the cross-sectional area A is b_w*e and so the mass flow rate is:

$$q = \frac{k\rho b_w e^3}{12\mu}\frac{\partial P}{\partial x_i} \qquad (4)$$

For a closed fracture, force equilibrium requires that shear traction be equal to the frictional resistance to slip. An additional term, $v*\eta$, called the radiation damping term, is included to approximate the damping effect of inertia on sliding at high velocities (Segall, 2010). The variable η is on the order of 20 MPa/(m/sec.), which means that the radiation damping term is small for $v << 1$. The frictional equilibrium equation is (Segall, 2010):

$$\tau - \eta v = \mu_f \sigma_n^{'} \qquad (5)$$

where μ_f is the coefficient of friction and $\sigma_n^{'}$ is the effective normal traction, defined as (Segall, 2010):

$$\sigma_n^{'} = \sigma_n - P \qquad (6)$$

where compressive tractions are taken to be positive. Following the rate and state friction law, the coefficient of friction is defined as a function of sliding velocity and state (Segall, 2010):

$$\mu_f = f_0 + a\ln\frac{v}{v_0} + b\ln\frac{\theta v_0}{d_c} \qquad (7)$$

where f_0, v_0, a, b, and d_c are material constants. The variable d_c is referred to as the characteristic displacement scale. The parameters a and b are ~0.01, much smaller than f_0, which is ~0.6. Their relatively small value is consistent with the observation that only a fraction of the stress borne by a fracture is usually released during a seismic event.

Under a rate and state framework, all fractures are slipping at all times. If a fracture bears little shear traction, it will have a tiny, but nonzero slip velocity. Very tiny velocities, ~10^{-15} m/sec. or less, are physically meaningless, but this is not a practical problem because these fractures behave approximately as if they were locked.

The state variable can be interpreted as the average contact time of asperities on the fault. The "aging law" of state evolution is (Segall, 2010):

$$\frac{\partial\theta}{\partial t} = 1 - \frac{\theta v}{d_c} \qquad (8)$$

A consequence of the aging law is that at very low velocities, friction gradually strengthens, which is consistent with laboratory results. One possible explanation for this effect is that over time, asperities deform plastically into one another, increasing the total frictional contact surface area. Another possible explanation is that surface adhesion increases with time as impurities diffuse away from the contact (Segall, 2010).

The stresses induced by fracture slip were calculated according to the equations of quasistatic equilibrium in a continuum assuming that body forces are equal to zero. They are given by the vector equation (Jaeger et al., 2007):

$$\nabla^T \mathrm{T} = 0 \qquad (9)$$

where **T** is the stress tensor.

Linear elasticity in an isotropic, homogeneous body was assumed, which means that the relationship between stress and strain is given by Hooke's law (Jaeger et al, 2007):

$$\mathrm{T} = \frac{2G\nu_p}{1-2\upsilon_p} trace(\varepsilon) I + 2G\varepsilon \qquad (10)$$

where I is the unit vector, ε is the strain tensor, ν_p is Poisson's ratio, and G is the shear modulus.

The cumulative displacement at any point is equal to the time integral of velocity:

$$D = \int v dt \qquad (11)$$

Other constitutive relationships were used. Both void and hydraulic aperture are related to effective normal traction and cumulative displacement. There is not a universally accepted equation in the literature for the relationship between these variables. We followed Willis-Richards et al. (1996), Rahman et al. (2002), Kohl et al. (2007), and others by using:

$$E = \frac{E_0}{1+9\sigma_n^{'}/\sigma_{Enref}} + D\tan\frac{\phi_{Edil}}{1+9\sigma_n^{'}/\sigma_{Enref}} + E_{res} \quad (12)$$

where E_0, σ_{Enref}, E_{ref}, and φ_{Edil} are material constants. We allowed these constants to be different between hydraulic aperture, e, and void aperture E.

We defined a D_{emax} and a D_{Emax}. These parameters are a maximum displacement at which point further displacement does not contribute to further increase in aperture. The fracture could still experience shear displacement, but further displacement did not affect aperture beyond these values. Such a property was not recognized in early laboratory testing of shear displacement and aperture coupling in granite (Barton et al., 1985). More recent laboratory work has observed this phenomenon. Esaki et al. (1999) and Lee and Cho (2002) both found that for a shearing fracture in granite, permeability increased rapidly at first, but permeability did not increase further after 5-10 mm of slip. Esaki et al. and Lee and Cho both observed an increase in mechanical aperture with slip beyond 10 mm of slip. It is not clear whether or not void aperture continued to increase after 10 mm of slip. In the base case simulation, D_{emax} was used but D_{Emax} was not.

Fluid density and viscosity are related to fluid pressure (and temperature, but the simulations are isothermal). Values were interpolated from a large table of properties generated using the freeware Matlab code XSteam 2.6 by Magnus Holmgren (2007).

A microseismic event was considered to have begun when the maximum velocity on the fracture exceeded 5 mm/sec. A slip event was considered finished when the highest velocity on the fracture dropped below 2.5 mm/sec. The typical duration of a slip event was around one second.

The total amount of displacement on the fracture during the event was correlated to seismic magnitude. The seismic moment M_0 is a measure of the size and energy release of an earthquake (Stein and Wysession, 2003). M_0 is defined as:

$$M_0 = GD_s A_s \quad (13)$$

where D_s is the average displacement, and A_s is the area of slip. From Hanks and Kanamori (1979), the seismic moment scale defines the moment magnitude M_w as:

$$M_w = \frac{\log_{10} M_0}{1.5} - 10.73 \quad (14)$$

where M_0 is defined in dyne-cm. As for fluid flow, the dimension of the fracture out of the plane (in the third dimension) was taken to be b_w.

We neglected elastodynamic transfer of stress. Dynamic stress transfer may have some effect on the results but is extremely expensive to calculate. Lapusta (2001) found for a single fracture case, calculations neglecting dynamic stresses could be made consistent with dynamic stress calculations simply by using a lower value of η. Lapusta suggested that for geometries more complex than a single fracture, dynamic stress transfer may play a bigger role.

Spatial Discretization

The fracture was discretized into constant length elements.

The force equilibrium and stress/strain relations were solved with the two-dimensional linear elastic boundary element method (BEM) from Crouch and Starfield (1983). The problem reduces to finding the induced stresses $\Delta\sigma_n$ and $\Delta\tau$ at each element caused by the cumulative shear displacements from each element. Stresses and displacements are linearly related so that:

$$\Delta\tau_i = \sum_{j=1}^{n} A_{ij} D_j \quad (15)$$

$$\Delta\sigma_{n,i} = \sum_{j=1}^{n} B_{ij} D_j \quad (16)$$

where A and B are matrices of interaction coefficients calculated according to Crouch and Starfield (1983). For the case of a single linear fracture, B is zero.

We neglected stresses induced by fracture normal displacement. The fracture in our simulations was never "open" because it never experienced tensile stress. In reality, there is some slight fracture normal displacement of a closed fracture due to loading or unloading, but these displacements are small (Barton et al., 1985), ~0.1 mm, and their effect should be slight compared to the effects of pore pressure change caused by injection and shear-induced dilation.

The mass balance equation was solved using the finite volume method. The flow between two adjacent elements in a linear fracture can be calculated directly from Darcy's law. The transmissibility between two elements was calculated using the harmonic average. Flow between fracture elements at an intersection can be calculated according to the method of Karimi-Fard et al. (2004).

Time Discretization

The issue of solving mechanical and flow equations simultaneously has been discussed at length in the literature of poroelasticity. One way to solve the problem is to use implicit Euler time-stepping on every equation simultaneously and solve the entire problem as a large coupled system of equations. This is a "fully coupled" scheme (Kim et al., 2011). The fully coupled strategy is stable and accurate but is expensive computationally. We used an "explicitly coupled" scheme in which a rate and state time-step was taken, and then the time-step was repeated for the flow problem using the rate and state result as a boundary condition.

The rate and state time-step was taken with an explicit, third-order Runge-Kutta scheme (Abramowitz and Stegun, 1972). In this time-step, the state, shear traction, normal traction, and cumulative displacement were updated. Next a flow time-step was taken using implicit Euler to find m^{n+1}. The frictional equilibrium equation, Equation (5), is an algebraic constraint, not a differential equation. This equation was solved last, on an element by element basis with sliding velocity as the unknown. The flow equations and the frictional equilibrium equation were solved using Newton-Raphson iteration.

Figure (1) summarizes the coupling strategy.

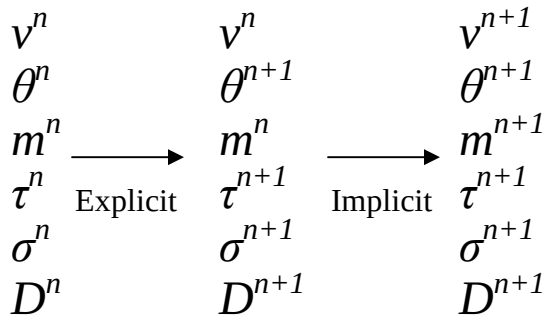

Figure 1: Explicit coupling strategy.

The advantage to splitting the problem is that different parts of the problem are most appropriately solved in different ways. The implicit Euler scheme is always numerically stable and is necessary to solve flow equations such as Equation 1. However the implicit Euler scheme requires solving a large system of equations. It would be impractical to attempt implicit Euler with the (BEM) equations from the Crouch and Starfield (1983) method because the BEM uses a dense matrix of interaction coefficients. The matrix inversion requirements would be very large. Explicit time-steps require only multiplications of the BEM matrix.

Adaptive time-stepping was used. The time-steps were chosen based on four criteria. The first was a built in error estimation from the third-order Runge-Kutta method. The second was the change in fluid pressure during the previous time-step. The third was the number of iterations used by the flow simulator in solving the non-linear system of equations. The fourth was the relative amount of velocity change for each of the elements at the previous time-step. Each of the four criteria had a target value and an adjustment factor (either up or down) was calculated that would move each value toward its target based on the result from the previous time-step. Of the four criteria, the adjustment that resulted in the most conservative time-step was used. If any of the criteria exceeded four times the target, the entire time-step was discarded and repeated with a smaller time interval. During seismic events when slip was very rapid, very small time-steps on the order of microseconds were necessary. In between seismic events, time-steps on the order of seconds, minutes, or hours were taken.

PROBLEM SETUP

Simulations were performed for injection into the center of a single, isolated one-dimensional fracture embedded in a two-dimensional whole space. The code is capable of simulating multiple fracture but simulations were limited to the single fracture case in this paper for the sake of simplicity.

The fracture was 500 m long and oriented 20° clockwise from the vertical y-axis. The two-dimensional problem could be interpreted as viewing a strike-slip fault in plan view, a normal fault in side view, or a reverse fault in side view, rotated 90°. The fracture was discretized into 2000 25 cm elements. The rock around the fracture was assumed to have zero permeability.

The parameters of the base case simulation (Case 1) are given in Table 1.

Table 1:List of simulation base case parameters.

P_{init}	40 MPa	D_{emax}	5 mm
T_{init}	200°C	D_{Emax}	-
θ_{init}	2589760 sec.	E_{snref}	95 MPa
σ_{xx} *remote*	60 MPa	e_{snref}	95 MPa
σ_{yy} *remote*	100 MPa	f_0	0.6
σ_{xy} *remote*	0 MPa	d_c	.05 mm
a	0.011	v_0	.001 mm/sec.
b	0.02	φ_{Edil}	0.5°

G	10000 GPa	φ_{edil}	1º
ν_p	0.1	q_{inj}	0.2 kg/sec.
E_0	1 mm	p_{injmax}	55 MPa
e_0	.01 mm	T_{inj}	200ºC
E_{res}	.002 mm	b_w	100 m
e_{res}	.0002 mm		

Injection was performed at the constant rate q_{inj} unless the injection pressure required to maintain q_{inj} exceeded p_{injmax}, in which case injection was performed at constant pressure p_{injmax} until q_{inj} could be achieved at a pressure below p_{injmax}. Once the outer elements of the fracture had slipped by 0.1 mm, injection was ceased. The simulation was continued after injection stopped for a period equal to 20 times the duration of injection.

The frictional parameters a, b, and d_c deserve some discussion. In order for unstable slip to occur, a must be smaller than b (Segall, 2010). This is because in order to achieve runaway velocity acceleration, the friction weakening effect of state decrease must be greater than the friction strengthening effect of velocity increase. The parameter d_c controls the minimum size of a patch of slip that can slip unstably and cause seismicity (Segall, 2010). d_c also limits the size of the spatial discretization. The element size must be significantly smaller than a characteristic length scale related to a, b, d_c, and σ_n', otherwise the result is numerically unstable (Lapusta, 2001).

A total of twelve simulations were carried out. The base case is Case (1). The other runs used identical parameters except for the changes shown in Table 2.

Table 2: Parameters for Cases 1-12. In each case, variables not explicitly changed in this table are the same as given in Table 1.

1	-
2	E_{phidil} = 2.5º
3	D_{Emax} = 5 mm
4	D_{Emax} = 5 mm, E_{phidil} = 2.5º
5	E_{phidil} = .0001º
6	Slide list
7	d_c = 5 mm
8	Constant pressure inj. at 55 MPa
9	Constant pressure inj. at 45 MPa
10	Constant pressure inj. at a pressure linearly declining from 55 MPa to 45 MPa over 10,000 sec., then constant at 45 MPa
11	Constant pressure inj. at 50 MPa
12	Constant pressure production at 40 MPa after the end of injection

RESULTS

Plots of injection rate, injection pressure, and event magnitude versus time for Cases (1) through (12) are shown in Figures 2-13. Table (3) gives summary data for the twelve simulations.

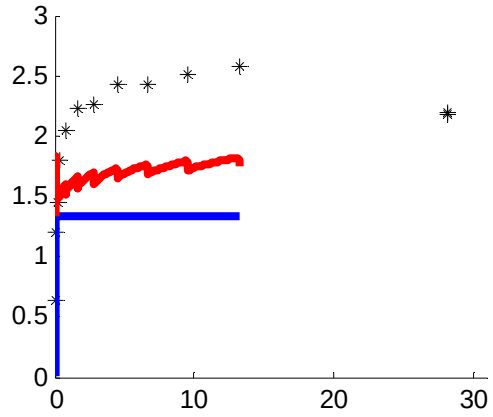

Figure 2: Case 1 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

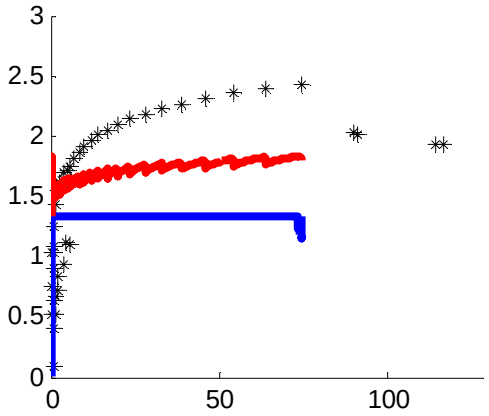

Figure 3: Case 2 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

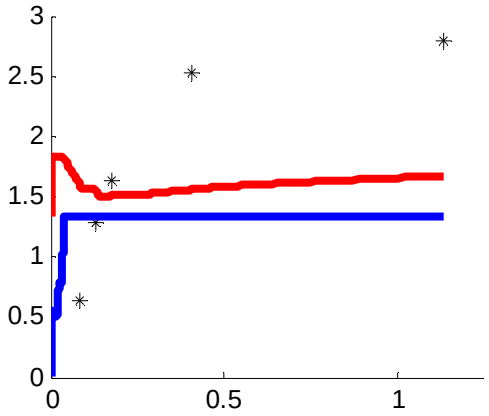

Figure 4: Case 3 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

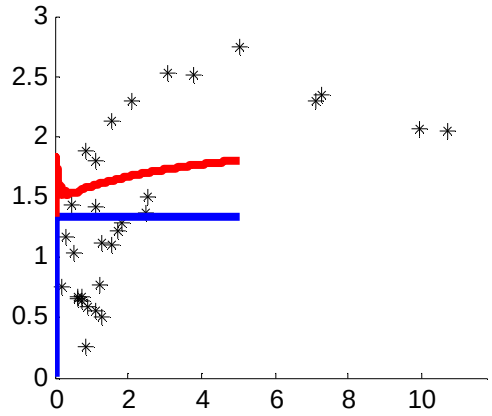

Figure 5: Case 4 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

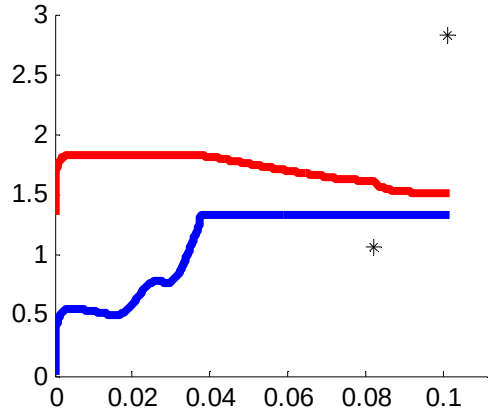

Figure 6: Case 5 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

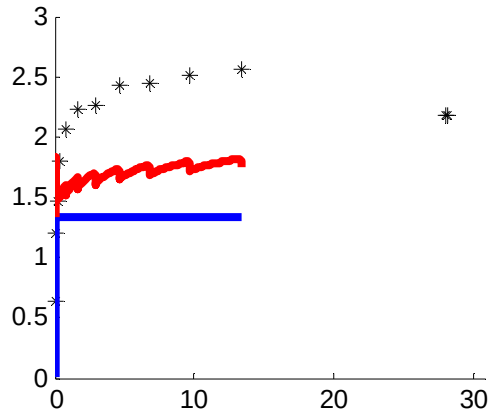

Figure 7: Case 6 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

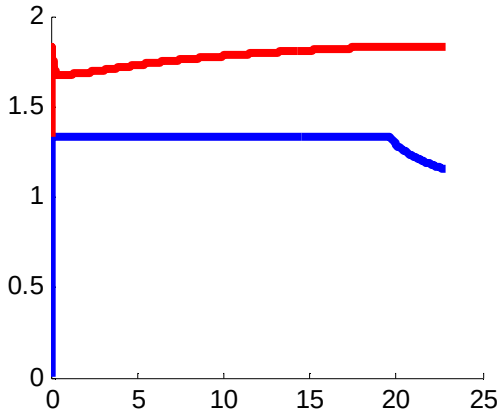

Figure 8: Case 7 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

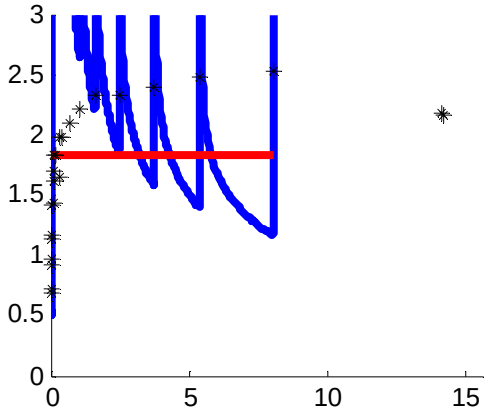

Figure 9: Case 8 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

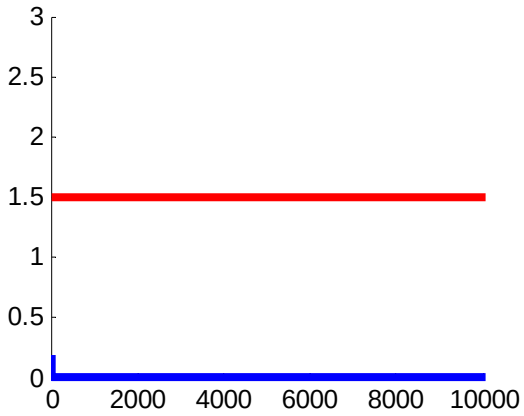

Figure 10: Case 9 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours). In this simulation, the fracture was not fully stimulated, and no events occurred.

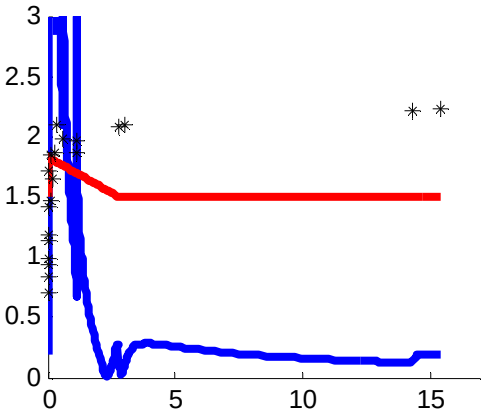

Figure 11: Case 10 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

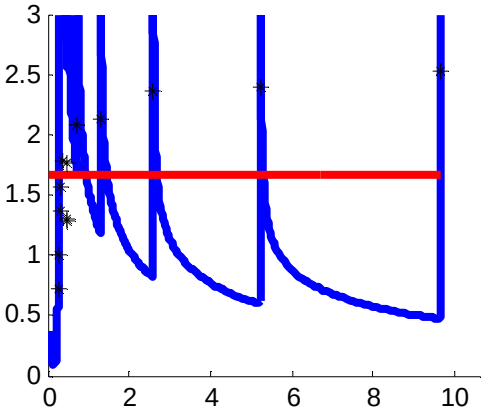

Figure 12: Case 11 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours).

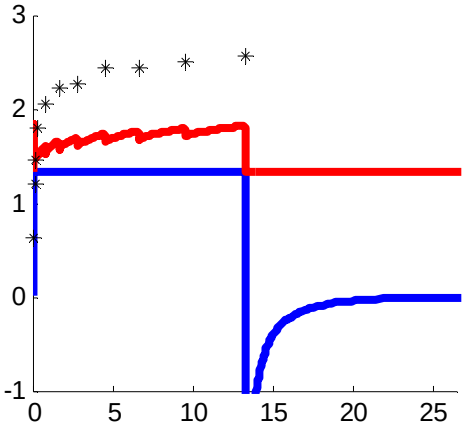

Figure 13: Case 12 injection rate (blue, kg/sec./0.15), injection pressure (red, MPa/30), and event magnitude (black) versus time (hours). Negative flow rate indicates fluid production and positive flow rate indicates injection.

Table 3: Summary data for Cases 1-12. The first column gives the case number. The second column gives the maximum event magnitude during injection. The third column gives the maximum event magnitude after injection. The fourth column gives the number of events magnitude less than 2.0. The fifth column gives the number of events between 2.0 and 2.5. The sixth column gives the number of events magnitude greater than 2.5. The seventh column gives the duration of injection until the entire fracture had been stimulated.

	Max. Mag. During	Max. Mag. After	< 2.0	<2.5, >2.0	>2.5	Time to Complete (hours)
1	2.57	2.19	4	7	2	13.25
2	2.43	2.02	30	13	0	74.56
3	2.80	-	3	0	2	1.13
4	2.74	2.34	23	6	3	5.04
5	2.82	-	1	0	1	0.10
6	2.57	2.19	4	7	2	13.38
7	-	-	0	0	0	22.87
8	2.53	2.16	15	8	1	8.08
9	-	-	0	0	0	-
10	2.23	-	15	5	0	15.39
11	2.52	-	8	4	1	9.64
12	2.571	-	4	5	2	13.25

In Cases (7) and (9), no seismicity occurred. In Case (9), the stimulation was never completed, because after several years of simulation time, the outer elements of the fracture still had not slipped.

DISCUSSION

A number of issues are discussed in the following subsections. A few general observations are made first. Subsequent subsections discuss the effect of slip-induced void dilation, a comparison of rate and state friction results to static/dynamic results, a discussion of issues associated with static/dynamic simulation, the effect of injection schedule on seismicity, and the effect of the characteristic displacement scale, d_c.

General Comments

At the beginning of injection, events were relatively low in magnitude, but over time, they grew larger. Magnitude increased over time because magnitude was related directly to the surface area of fracture that slipped. At later times, more fracture was available to slip. A similar effect was observed in numerical simulations by McClure and Horne (2010) and Baisch et al. (2010).

In the simulations, slip tended to be confined to areas that had already slipped. During each slip event, there tended to be some growth of the slipped region into previously unstimulated fracture, but that growth died out after some distance. Such confinement of the slipping region was partly caused by shear-induced dilation, as discussed later, and partly because the fracture was not initially close to slipping.

It is possible to imagine a scenario in which confinement would not occur. For example, if injection were carried out into a fault that was about to slip anyway, the event might be triggered, and it could propagate far from the injector well. A runaway slip event occurred in Case (5), as discussed later. An unconfined slip event has never occurred during an actual EGS stimulation. It is unclear how likely such an event would be, but it is a subject worth consideration.

As long as the fracture was not naturally on the verge of slipping, it would be easier to induce slip on a region of fracture that had already been induced to slip than on a new region of fracture. Under these conditions, slip would remain relatively confined. As the stimulated region got larger, the potential for larger events would grow because there would be more previously stimulated fracture available to slip. The largest events would occur across the entire fracture.

In many simulations, there were post-injection events. These events generally occurred in pairs very close together several hours after injection ceased. The pairs corresponded to events on either side of the one-dimensional fracture.

The post-injection events occurred because of pressure redistribution. During injection, there was a pressure gradient away from the wellbore. After injection stopped, the pressure redistributed to become uniform everywhere. The redistribution lowered pressure near the injector and increased pressure away from the injector. A similar effect was described in McClure and Horne (2010), Baisch et al. (2010), and Baisch et al. (2006). A schematic diagram of this effect is shown in Figure 14.

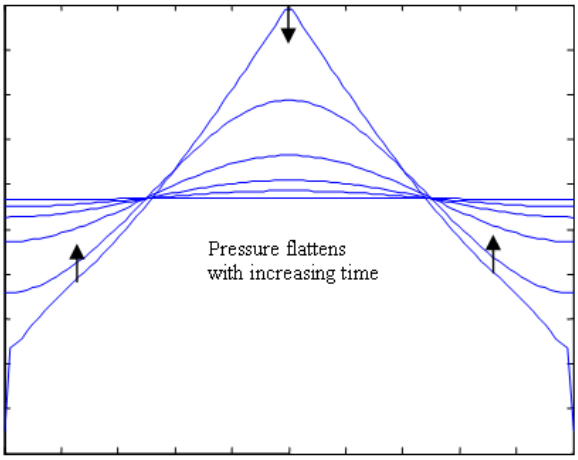

Figure 14: Schematic of the change in pressure distribution along a fracture after injection stops. The x-axis is distance along the fracture, and the y-axis is fluid pressure. The injector is located at the center. Reproduced from McClure and Horne (2010).

The post-injection events were of significant magnitude but smaller in magnitude than the largest events during injection. In actual EGS stimulation, post-injection events are often larger than events during injection. The discrepancy may be partly an artifact of our use of a one-dimensional fracture. For a two-dimensional planar fracture, the periphery would have relatively more surface area than in the one-dimensional case. The post-injection events, which occur at the periphery, would be relatively larger. In the one-dimensional simulations, peripheral events tended to happen in closely spaced pairs because the center region, which did not slip, prevented slip on one side from triggering slip on the other. For a two-dimensional, planar fracture, the periphery would be a ring-shaped region, and there would be no barrier to prevent slip anywhere in the periphery from inducing slip everywhere else in the periphery. Baisch et al. (2010) simulated induced seismicity in a two-dimensional planar fracture and observed larger events post-injection than during injection. They modeled a fracture of infinite size and noted that pressure redistribution caused the

stimulated region continued to grow even after injection stopped.

For the simulations at constant injection rate, injection pressure rose with time. There were decreases in injection pressure immediately after slip events due to enhanced permeability. For constant pressure injection, injection rate tended to decrease with time with spikes in injection rate after slip events. In actual EGS stimulation, a variety of possible relationships between rate and pressure have been observed, but overall, injection pressure tended to stay roughly constant with constant injection pressure.

The relationship between the injection rate and pressure is dependent on the relationship of slip and fluid pressure to fracture permeability. In general, increasing permeability with time should tend to prevent injection pressure from rising at during constant rate injection and rate from falling at constant pressure. The relationship was investigated in further detail in McClure and Horne (2010).

It is sometimes assumed in the EGS literature is that the advance of the stimulated region is caused by diffusion of pressure into the unstimulated fracture (Shapiro et al., 1999 and Bruel, 2007). But that was not the case in our simulations. Slip events nucleated in regions where pressure was increased but spread into parts of the fracture that had not yet slipped. These unstimulated regions had low permeability and low pressure. The spread of the slip event was caused by the increase in shear traction induced by the slip.

There are some aspects of our simulations that did not agree with observations from EGS projects. Our simulations had few small events (magnitude one or less). In reality, EGS stimulations trigger a large number of small events. Our simulations failed to capture small events because of our choice of characteristic displacement scale d_c, and because our fracture did not have heterogeneous properties.

The parameter d_c directly affects the minimum size of seismic events that can occur. Unstable slip cannot occur on a patch of fracture below a minimum patch size that is related to the d_c variable. Smaller patches slip slowly, and aseismically. The cause of the minimum patch size can be understood through a stability analysis of the rate and state system of equations (Segall, 2010).

In Case (7), a large d_c was chosen and the entire fracture was stimulated without causing a single seismic event. Slip occurred across the fracture, but it was slow and aseismic.

Smaller values of d_c would allow smaller events to be modeled. The drawback to using a smaller d_c is that it would become necessary to use a finer spatial discretization, which would be computationally expensive. Discretizations above a certain threshold related to d_c are numerically unstable (Lapusta, 2001).

Another reason our simulations had few small events is that our fracture was planar and had homogenous properties. Irregularity would lead to more stress heterogeneity and a greater number of smaller events.

A final reason our simulations had few small events is that we used a one-dimensional fracture. We assumed that the fracture width was 100 m, which meant that even a relatively "short" slip patch on the one-dimensional fracture had a significant slip area due to its width.

Effect of Pore Dilation

Cases (1), (2), (3), (4), and (5) examined the effect of slip induced void aperture dilation on seismicity.

Overall, pore dilation had the effect of damping out seismicity. This can be seen clearly from the case in which there was not any void dilation, Case (5). In Case (5) only two events occurred. The second nucleated close to the injector but spread until it reached the edge of the fracture, causing a magnitude 2.8 event, one of the largest in any of the simulations. The slip event would have propagated even further had the fracture been larger. The entire fracture had slipped six minutes after the initiation of injection.

In Cases (1) and (2), void dilation was stronger and events had a much harder time propagating across the fracture, especially into areas that had not yet slipped substantially. Comparing Cases (1), (2), and (5), it can be seen that the maximum magnitude went down as void dilation increased. More events occurred as shear dilation increased because the simulation lasted longer. More pumping and time was required to stimulate the entire fracture. More injection was required because the dilation increased the void volume of the fracture and led to an increase in the fluid injection required to increase its pressure.

Pore dilation damped out slip events because it caused a decrease in fluid pressure during slip. During rapid slip, fluid flow did not have time to occur, and so the mass of fluid at a given location was nearly fixed. Water density is relatively insensitive to pressure, so to conserve mass, the void aperture at a given location had to be nearly constant during slip. As slip tried to dilate the void aperture, the only way to keep void aperture constant was to decrease fluid pressure, increasing the effective normal traction on the element and compressing the

fracture slightly. The higher effective normal traction caused by the decrease in fluid pressure strengthened friction and tended to inhibit slip from occurring.

In Cases (3) and (4), a threshold D_{Emax} was set at which point additional shear displacement did not increase void aperture further. Not surprisingly, the displacement limit led to larger magnitude events. Once D_{Emax} was reached, events were able to propagate across the previously stimulated fracture unimpeded by pore pressure decrease. The limit on void aperture dilation did not affect the damping effect of void dilation on propagation of slip into virgin fracture, and so the runaway slip event seen in Case (5) was not observed.

Despite what might seem to be suggested from Case (5), runaway fracture slip does not necessarily have to occur if there is not void dilation. Void dilation is not the only phenomenon that might inhibit spread of seismic slip into previously unstimulated fracture. Nonplanar faults or multiple faults would make it more difficult for slip to spread by making the induced shear tractions less direct. If a fault was not naturally at a stress state that put it close to slip, there would be a natural resistance to slip that the seismic event would have to overcome. The value of the rate and state parameters *a* and *b*, are also very important, with a lower *b* relative to *a* leading to less tendency for seismic slip. The value of *b* chosen for this work is quite large relative to *a*, leading to conditions that are favorable for runaway slip. In the limiting case of *b* less than or equal to *a*, seismic slip is impossible. In that case, slip must be aseismic.

Effect of d_c

Case (7) demonstrated that a larger value for d_c inhibits seismicity. In Case (7), the entire fracture was stimulated, but the slip was uniformly slow and aseismic. This result is predicted by the theory of rate and state friction (Segall, 2010) which says that the value of d_c affects the minimum size of a patch that can slip unstably. Lapusta (2001) observed a similar effect with elastodynamic earthquake simulation. If the fracture in Case (7) had been large enough, the region of increased fluid pressure would have eventually gotten large enough that it would have slipped unstably.

Differences in the rate and state parameters *a*, *b*, and d_c in nature account for why sometimes fractures slip seismically, and sometimes they slip aseismically. Before initiating an injection experiment at a given location, characterization of the parameters *a*, *b*, and d_c could be useful for predicting seismic hazard.

Comparison of Rate/State Friction to Static/Dynamic Friction

Previously, in McClure and Horne (2010), we modeled injection to into a one-dimensional isolated fracture using static/dynamic friction instead of rate and state friction. A related approach was used by Baisch et al. (2010), who modeled injection induced seismicity in a two-dimensional planar fault. In spite of the significant differences between the two treatments of friction, the overall conclusions from McClure and Horne (2010) were consistent with the results from the rate and state modeling described in this paper. Simulations were qualitatively similar to the results of Baisch et al. (2010).

In the rate and state results from this paper, in McClure and Horne (2010), and in Baisch et al. (2010), event magnitudes increased with time as the stimulated region grew larger. Post-injection events occurred because of redistribution of pressure as described above.

Examples of static/dynamic results are given in Figures 15 and 16.

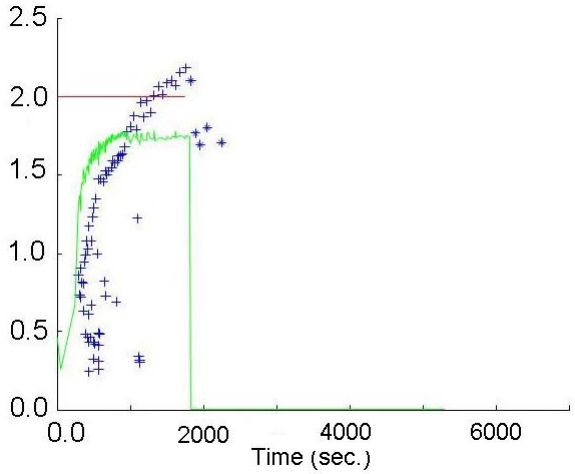

Figure 15: Reproduced from McClure and Horne (2010). Injection into a single, isolated fracture using static/dynamic friction. Injection pressure (red, MPa/30), injection rate (green,kg/sec./4), and event magnitude (blue).

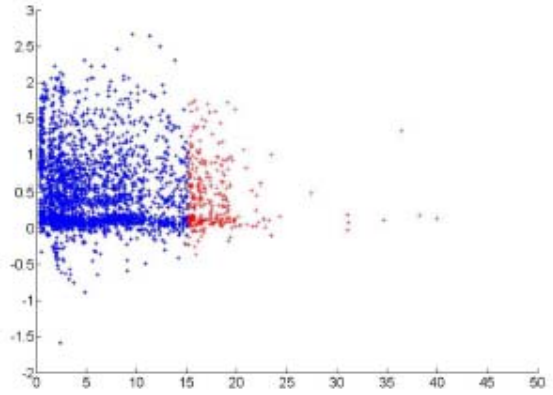

Figure 16: Reproduced from McClure and Horne (2010). Event magnitudes during injection into a network of many fractures using static/dynamic friction. The x-axis is time(hours).Blue events occurred during injection and red events occurred after injection.

From Figure 16, it is evident that when stimulation was performed into a network of many fractures, event magnitudes did not as grow with time as clearly as in the single fracture case. However, the largest observed events were later in the simulation. The difference is that with injection into a single fracture, slip events can occur across the entire stimulated region, but that is not as likely to happen in a large network of fractures because slip on one fracture does not necessarily increase shear traction on another.

A difference between the rate and state simulations in this work and the results from both Baisch et al. 2010) and McClure and Horne (2010) is that there were fewer low magnitude events in the rate and state simulations. The reasons for the small number of minor events are discussed above. The basic reason is that rate and state friction can model aseismic slip, which static/dynamic cannot do, and this work used a value of d_c which unrealistically favored aseismic slip for small events.

Issues Associated with Static/Dynamic Friction

With static/dynamic friction, fractures can have only two possible coefficients of friction, a static μ_s and a dynamic μ_d. Fractures are considered either "slipping" or "static." A fracture is static if it has a frictional resistance to slip greater than its shear traction. If traction exceeds static resistance, the coefficient of friction is instantaneously lowered to μ_d and displacements are calculated so that the element is at equilibrium based on the new, lower dynamic coefficient of friction. Often slip of one element causes the adjacent elements to slip, which triggers yet more elements to slip, and a chain reaction can occur in which a large number of elements slip. Iteration is required to determine which elements should be considered slipping or stationary. The slip is assumed to be simultaneous. In a sense the calculations mimic the physical process of seismic nucleation because slip on one patch triggers slip on another and so on.

The static/dynamic approach is problematic for a number of reasons. One issue is that time is not present in the equations for slip. Displacements are calculated from the simultaneous solution to the coupled equations of quasi-static equilibrium for the slipping elements. The implication is that all slip during a single slip event is simultaneous and instantaneous. There is no causality in such a calculation and this is intuitively unsatisfying.

Another problem with static/dynamic friction is that it cannot distinguish between seismic and aseismic slip. With static/dynamic friction all slip is seismic because it enforces an instantaneous change in friction and an instantaneous change in displacement (which requires an infinite slipping velocity).

Another issue is that with static/dynamic friction, fractures are unstable in the limit of a very fine discretization.

Take for example a one-dimensional fracture of length S loaded with 6 MPa of shear traction and 50 MPa of effective normal traction with $\mu_s = 0.6$ and $\mu_d = 0.55$. The frictional strength of this fracture is 30 MPa, much greater than its shear traction. The fracture is discretized into N elements of length S/N. Now assume that the effective normal traction along one meter of the fracture is dropped to 9.9999 MPa. The elements on the one meter patch will slip. Their friction will drop to 0.55, and the resulting shear traction drop will be $(\mu_s - \mu_d)*\sigma_n' \sim 0.5$ MPa.

The displacement field along the one meter patch can be calculating using the Crouch and Starfield (1983) BEM method and turns out to be ellipsoidal with a maximum at the center and tapering to zero at the edges. It is also possible to calculate the induced shear traction on the fracture element directly adjacent to the one meter patch of slip. The induced traction on the adjacent element is $\Delta\tau_a$. It turns out that $\Delta\tau_a \sim N^{1/2}$, which means that as N goes to infinity, $\Delta\tau_a$ goes to infinity. This makes sense because as N increases, the next adjacent element gets closer and closer to tip of the slipping patch. The slip patch can be thought of as a crack responding to a uniform load, and according to the analytical solution for stress induced by a crack, there is a stress singularity at a crack tip equal to the inverse of the square root of the distance from the tip (Pollard and Fletcher, 2005). The consequence is that regardless of how great the frictional strength of the element next to the slipping region is (or regardless of low its initial shear traction

is) it can be induced to slip with a fine enough discretization.

When the adjacent element slips, it will experience a stress drop of $(\mu_s - \mu_d)*\sigma_n'$, which in the case of our example would be ~ 2.5 MPa. Note that the stress drop on the adjacent element will be greater than the stress drop on the original patch. The result is an even larger crack undergoing an even bigger stress drop than before. As a result, the stress induced on the next element adjacent to this new, larger slip patch will be even greater than on the first, and it too will be induced to slip, and so on until every element in the entire fracture has slipped. In practice, slipping patches of fractures do not always cause an entire fracture to slip in static/dynamic friction but that turns out to be a convenient discretization error!

It seems that while static/dynamic simulations have significant disadvantages, their overall results are qualitatively reasonable, at least in some applications. Static/dynamic simulations are far more efficient computationally than rate and state simulations, because they allow far greater time-steps and coarser spatial discretizations.

We conclude that static/dynamic friction seems to give results that are qualitatively similar to the more rigorous rate and state calculations. Static/dynamic friction could be a reasonable choice for some applications, especially for very large problems. Perhaps some intermediate treatment of friction could be found that would be more efficient than rate and state but would resolve the issues of static/dynamic.

Effect of Injection Schedule

Cases (8), (9), (10), (11), and (12) investigated the effect of injection pressure. In contrast to the other simulations, these simulations were run at constant injection pressure rather than flow rate. Flow rate changed with time, experiencing a spike after a seismic event and gradually decreasing with time between events. Overall, the flow rate was decreasing over time in all four cases, even taking into account the rate spikes after events. These cases can be compared to Case (1), in which injection was at constant rate, and injection pressure gradually increased with time until the maximum injection pressure was reached late in simulation, at which point injection was continued at constant pressure.

In McClure and Horne (2010), we examined the effect of injection pressure schedule using the static/dynamic friction model. We found that higher injection pressure led to more seismicity, but that starting at a high pressure and decreasing over time as in Case (10) reduced the maximum magnitude of events.

Our McClure and Horne (2010) results were confirmed using the rate and state simulations. The maximum magnitude during Case (10), which used a declining pressure over time was only 2.2, significantly lower than any other case. The maximum magnitude during both Case (8), injection at high pressure, and Case (11), injection at lower pressure, was about 2.5. The number of events greater than magnitude 2 was nine for Case (8) and five for Case (11). The maximum magnitude in Case (1), which had the same physical parameters but constant injection rate, was 2.4. Seven events greater than magnitude 2 occurred in Case (1).

The physical explanation is that relatively large events only occur when a large surface area of fracture slips. The ability for a slip event to propagate far into a region of unstimulated fracture is limited by void dilation (or some other effect, as discussed earlier). Therefore the maximum size of an event is limited roughly by the size of the region of fracture that has already slipped. At early time, the stimulated region is small, and the potential for a large event is low. It is safe to use a higher injection pressure at earlier time. Inducing more slip early on releases more of the shear traction on the fracture near the injector, which reduces that region's ability to slip later and contribute to the larger events.

Baisch et al. (2006) postulated that producing fluid from the well after injection would reduce post-injection seismicity. Both our simulations in this paper and our static/dynamic simulations in McClure and Horne (2010) found that to be the case. In Case (12), the well was put on production after injection, and there was no post-injection seismicity.

Case (9) used a very low injection pressure, 45 MPa. During this injection simulation, seismic slip never occurred. Very slight aseismic slip occurred in the region around the injector, but most of the fracture remained essentially locked. The injection rate went to zero over time. Even after years of injection, (at very low rate), the pressure remained near the initial pressure at the periphery of the fracture because the unstimulated permeability was very low and the pressure neither able to neither trigger enough slip to enhance permeability nor diffuse into the unstimulated fracture.

CONCLUSIONS

Modeling that couples rate and state friction and fluid flow simulation is a promising approach for simulation of EGS stimulation because it allows a direct, physically based simulation of the process of induced seismicity. Rate and state friction theory is widely accepted in the field of earthquake modeling, which makes it an attractive choice for EGS modeling.

A drawback to rate and state friction modeling is that it is relatively expensive computationally because it requires fine scale discretization in both time and space. In the appendix, we summarize various techniques that could be used to improve efficiency.

We applied rate and state modeling to investigate a number of issues related to EGS stimulation.

We verified our earlier static/dynamic friction results from McClure and Horne (2010) that reducing injection pressure over time should reduce the maximum magnitude of events, even compared to using a low injection pressure for the full duration of stimulation. We also verified our earlier result that producing fluid from the well immediately after injection should inhibit post-injection seismicity, as suggested by Baisch et al. (2006).

Shear-induced pore dilation damped seismicity, resulting in more, smaller events.

ACKNOWLEDGEMENTS

We gratefully acknowledge the Precourt Institute for Energy at Stanford University for funding this research. This work has benefitted greatly from interactions with a number of people at Stanford. Thank you very much to Andrew Bradley for his insight into fast boundary element solutions as well as the coupling of rate and state and fluid flow simulation. Thank you to Prof. Eric Dunham for his help and for lending us a rate and state simulation code which we used to check the accuracy of our simulator. Also thank you to Prof. Paul Segall and Prof. Hamdi Tchelepi for providing valuable insight.

REFERENCES

Abramowitz, M. and Stegun, I. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, U. S. Government Printing Office, Washington D. C., Tenth Printing, 1972.

Aziz, K, and Settari, A., *Petroleum Reservoir Simulation*, Blitzprint Ltd., Calgary, Alberta, 1979.

Baisch, S., Vörös, R., Rothert, E., Stang, H., Jung, R., and Schellschmidt, R., (2010), "A Numerical Model for Fluid Injection Induced Seismicity at Soultz-sous-Forêts," *International Journal of Rock Mechanics and Mining Sciences*, **47**, 405-413.

Baisch, S., Weidler, R., Vörös, R., and Jung, R., (2006), "A Conceptual Model for Post-Injection Seismicity at Soultz-sous-Forêts," *GRC Transactions*, **30**, 601-605.

Barton, N., Bandis, S., and Bakhtar, K., (1985), "Strength, Deformation and Conductivity Coupling of Rock Joints," *International Journal of Rock Mechanics and Mining Sciences*, 22(3), 121-140.

Bruel, D., (2007), "Using the Migration of the Induced Seismicity as a Constraint for Fractured Hot Dry Rock Reservoir Modeling," *International Journal of Rock Mechanics and Mining Sciences*, 44, 1106-1117.

Crouch, S. L., and Starfield, A. M., *Boundary Element Methods in Solid Mechanics*, George Allen & Unwin Ltd., London, 1983.

Dieterich, J. H, (2007), "Applications of Rate- and State-Dependent Friction to Models of Fault Slip and Earthquake Occurrence," *Treatise on Geophysics*, vol. 4, chapter 4, ed. Kanamori, H., Elsevier, Amsterdam and Boston, 107–129.

Esaki, T., Du, S., Mitani, Y., Ikusada, K. and Jing, L., (1999), "Development of a Shear-Flow Test Apparatus and Determination of Coupled Properties for a Single Rock Joint," *International Journal of Rock Mechanics and Mining Sciences*, **36**, 641-650.

Hanks, T. and Kanamori, H., (1979), "A Moment Magnitude Scale," *Journal of Geophysical Research*, **84(B5)**, 2348-2350.

Holmgren, M., (2007), *XSteam: Water and Steam Properties According to IAPWS IF-97*, <www.x-eng.com>.

Jaeger, J.C., Cook, N.G.W, and Zimmerman, R.W., *Fundamentals of Rock Mechanics*, Blackwell Publishing, Malden, MA, 2007.

Karimi-Fard, M., Durlofski, L. J., and Aziz, K., (2004), "An Efficient Discrete-Fracture Model Applicable for General-Purpose Reservoir Simulators," paper SPE 88812, presented at the 2003 SPE Reservoir Simulation Symposium, Houston, TX.

Kim, J., Tchelepi, H. A., and Juanes, R., (2011), "Stability, Accuracy, and Efficiency of Sequential Methods for Coupled Flow and Geomechanics," paper SPE 119084, presented at the 2009 SPE Reservoir Simulation Symposium, The Woodlands, TX, (SPE Journal, in press).

Kohl, T. and Mégel, T., (2007), "Predictive Modeling of Reservoir Response to Hydraulic Stimulations at the European EGS Site Soultz-sous-Forêts," *International Journal of Rock Mechanics & Mining Sciences*, **44**, 1118-1131.

Lapusta, N., *Elastodynamic Analyses of Sliding with Rate and State Friction*, PhD Thesis, Harvard University, Cambridge, MA, 2001.

Lee, H. S., and Cho, T. F., (2002), "Hydraulic Characteristics of Rough Fractures in Linear Flow under Normal and Shear Load," *Rock Mechanics and Rock Engineering*, **35(4)**, 299-318.

Majer, E., Baria, R., Stark, M., Oates, S., Bommer, J., Smith, B., and Asanuma, H., (2007), "Induced Seismicity Associated with Enhanced Geothermal Systems," *Geothermics*, **36**, 185-222.

McClure, M., and Horne, R. N., (2010), "Numerical and Analytical Modeling of the Mechanisms of Induced Seismicity During Fluid Injection," *GRC Transactions*, **34**, 381-396.

Morris, J., and Blair, S., (2000), "An Efficient Displacement Discontinuity Method Using Fast Multipole Techniques," presented at the Fourth North American Rock Mechanics Symposium, Seattle, WA, 2000.

Pollard, D., and Fletcher, R., *Fundamentals of Structural Geology*, Cambridge University Press, Cambridge, UK, 2005.

Rahman, M. K., Hossain, M. M., and Rahman, S. S., (2002), "A Shear-Dilation Based Model for Evaluation of Hydraulically Stimulated Naturally Fractured Reservoirs," *International Journal for Numerical and Analytical Methods in Geomechanics*, **26**, 469-497.

Rjasanow, S. and Steinback, O., *The Fast Solution of Boundary Integral Equations*, Springer Science+Business Media, LLC, New York, 2007.

Segall, P., *Earthquake and Volcano Deformation*, Princeton University Press, Princeton, New Jersey, 2010.

Shapiro, S., Audigane, P., and Royer, J.-J., (1999), "Large Scale *in-situ* Permeability Tensor of Rocks from Induced Microseismicity," *Geophysical Journal International*, **137**, 207-213.

Stein, S and Wysession, M., *An Introduction to Seismology, Earthquakes, and Earth Structure*, Blackwell Publishing, Ltd., Malden, MA, 2003.

Tester, J. (ed), *The Future of Geothermal Energy: Impact of Enhanced Geothermal Systems (EGS) on the United States in the 21st Cen*tury, MIT Report, 2007.

Willis-Richards, J., Watanabe, K., and Takahashi, H., (1996), "Progress Toward a Stochastic Rock Mechanics Model of Engineered Geothermal Systems," *Journal of Geophysical Research*, **101(B8)**, 17,481-17,496.

APPENDIX: STRATEGIES TO IMPROVE MODEL EFFICIENCY

The simulations in this paper used 2000 fracture elements to simulate a single 500 m fracture. As an example, Case (1) required three hours of computation to complete, using a conventional desktop PC. The ultimate goal of this work is to be able to perform coupled fluid flow and rate and state simulations on large networks of fractures involving 100,000+ fracture elements. To simulate much larger systems, strategies need to be identified to increase efficiency. Four strategies to improve efficiency are given in the following subsections. Two have been implemented in the current work, and two are still being developed.

Efficient Matrix Multiplication

A strategy to increase the speed of the dense matrix multiplication BEM was implemented in the code used for this paper. We followed an approach similar to that outlined in Rjasanow and Steinbach (2007).

The underlying concept is that parts of the dense BEM matrix contain redundant information and can be approximated. The method of approximation we used was singular value decomposition, but there are other methods available for matrix approximation. For example Rjasanow and Steinback suggest Adaptive Cross Approximation. Another method is the Fast Multipole Method, which was applied to the Crouch and Starfield two-dimensional BEM method by Morris and Blair (2000).

In singular value decomposition, an (m x n) matrix A is decomposed into three matrices:

$$A = U\Sigma V^T$$

The middle matrix Σ is diagonal and the values can be ordered from largest to smallest along the diagonal. In a poorly conditioned matrix, a few values in Σ are much larger than the others. In that case, A can be well approximated by discarding the small values of Σ and the corresponding columns and rows of U and V^T and keeping only s singular values from Σ.

Multiplying A with a vector x requires $m*n$ multiplications. With singular value decomposition, the operation can be split into two multiplications: $V^T x$, which requires $s*n$ operations and yields an (s x 1) vector, and $(U\Sigma)(V^T x)$, which requires $s*m$ multiplications (as long as U and Σ are pre-multiplied). The full multiplication requires $s*(m + n)$ operations.

Attempting singular value decomposition (SVD) on a full matrix of interaction coefficients would not yield much gain in efficiency (possibly a reduction in

efficiency), because full interaction matrices are well-conditioned and as a result their singular values are close together. However blocks of the full matrix are poorly conditioned. SVD can be used very effectively on the poorly conditioned parts of the matrix. Matrices of longer range interactions tend to be more poorly conditioned and are prime targets. The full matrix multiplication can thus be split into many smaller problems most of which are then approximated using SVD.

We used a series of grids to divide up a realization of fractures (or a single fracture). The coarsest grid split the problem in half; the second split it into quarters, the third into eighths and so on. The blocks were not uniform in size. They were chosen so that every block at a given level of refinement contained nearly the same number of elements. An example of the hierarchy of grids is shown in Figure 17.

The interaction between any two blocks at a given grid level was described by a matrix of interaction coefficients. There was a trade off. Interactions between finer scale blocks required fewer singular values. However the efficiency gain from SVD was greater if coarser blocks were used. The challenge was to find optimal grouping of elements: a greater number of fine scale SVD's with fewer singular values or a smaller number of coarse scale SVD's. In general, larger blocks were used for long range interactions and smaller blocks were used for close range interactions.

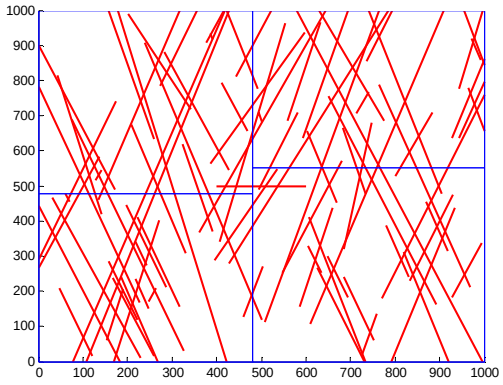

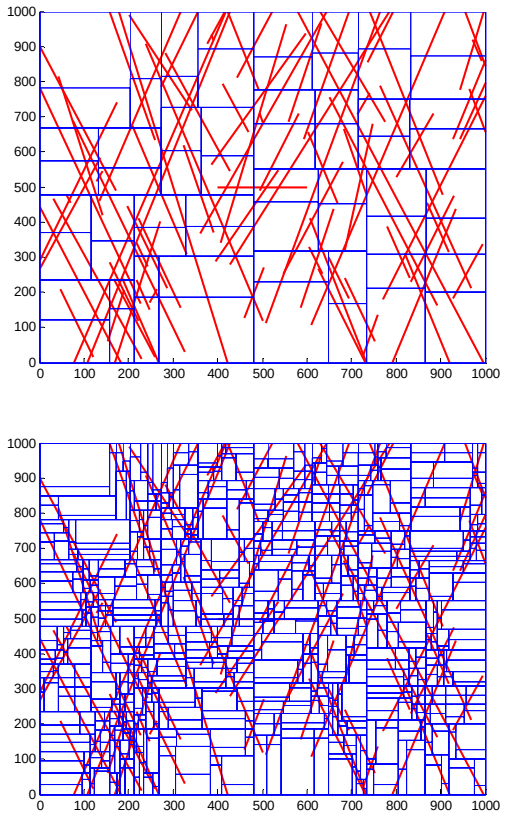

Figure 17: Splitting an arbitrary network of fracture elements into a series of blocks containing nearly equal numbers of elements. Different levels of refinement are shown.

A recursive algorithm was used to find the optimal groupings. Starting at the coarsest grid, SVD was applied to the matrix of interaction coefficients between each possible pair of blocks (including blocks to themselves) at that level of refinement. Each of the blocks in the pair was then decomposed to the next finer level, giving two smaller blocks from each original. For every pair of elements, SVD was repeated for the resulting 16 possible interactions between the four new child blocks. The number of singular values retained for each SVD was determined based on a predefined ratio between the largest singular value and the smallest singular value retained. If the coarser grid SVD would require fewer operations to perform multiplication than the sum of the 16 finer scale subproblems, the coarse SVD was accepted. If the finer SVD's required fewer operations, then the algorithm continued recursively by splitting them into smaller elements and comparing again. The recursion ended with interactions between individual elements. The individual element interactions not included in any of the SVD's were computed directly. These interactions tended to be between elements that were very close together (including all the interactions of

elements to themselves). The result of the algorithm was a list of smaller sub-problems and their singular value decompositions and a list of interactions that needed to be calculated directly. The sum of those subproblems was a close approximation to the original problem. The algorithm has a much better complexity than the n^2 complexity of standard matrix multiplication.

Computing singular value decompositions is very computationally intensive, but because the fracture geometry did not change with time in our simulations, SVD's only needed to be computed once. We used a separate program to create the SVD decompositions and save them into a text file.

Figure 18 shows a plot of number of elements versus multiplications required for four discretizations of the single fracture problem: one with 100, 500, 2000, and 10,000 elements, corresponding to 5 m, 1 m, 0.25 m, and 0.05 m.

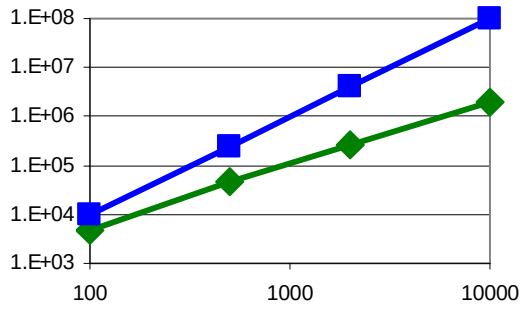

Figure 18: Matrix multiplication requirements for a matrix of interaction coefficients for a single, linear fracture using the 2D Displacement Discontinuity Method. The x-axis shows the number of elements and the y-axis shows the number of multiplications required. Shown are standard matrix multiplication (blue) and the hierarchical SVD decomposition method (green).

An accuracy comparison was performed for the case of a single straight 500 m fracture using discretizations with different levels of refinement. A displacement of 0.1 m was applied across the middle 20% of the fracture and 10^{-8} m everywhere else on the fracture. Singular values within 10^{-5} of the largest singular value were kept. Comparing the BEM to SVD methods, the maximum relative difference, defined as abs((BEM-SVD)/BEM), ranged from 8×10^{-5} for the 5 m element case to .0006 for the 0.05 m element case. The accuracy can be controlled by changing the ratio of largest singular value to smallest singular value that is kept. Using more singular values results in better accuracy but more computations are required.

The efficiency gain from the SVD method was not as great for the multiple fracture case because for the same level of accuracy more singular values were needed. The single fracture case described here required mostly less than five singular values per decomposition but for the multiple fracture case, tens of singular values were usually required.

For problems with hundreds of thousands of elements, SVD is not practical. The problem is that SVD has a cubic complexity, and so to perform the required singular value decompositions is too expensive, even though the SVD's only need to be performed once. Two possible alternatives are Adaptive Cross Approximations (Rjasanow and Steinbach, 2007) and the Fast Multipole Method (Morris and Blair, 2000). These methods follow an approach roughly similar to the one outlined above. We plan to implement all three methods in future work and compare the results.

Slide List Strategy

Another strategy to improve efficiency is to only update stresses caused by fractures that are slipping at non-tiny velocities. We call this the "slide list" strategy.

Fracture elements distant from the injector may have very low slipping velocity for much or all of the simulation. For example, in the above simulations, the initial slip velocity was around 10^{-14} m/sec. It is wasteful to constantly update stresses on these elements and to calculate the effect of their displacements on other elements. In the slide list strategy, we keep lists of elements that are sliding at greater than some minimum velocity (slide elements), elements that are close to sliding at that minimum velocity (check elements), and elements that are effectively not sliding (inactive). Only the effect of slide elements on slide and check elements is calculated at every time-step. At every time-step, the equilibrium equation residual is calculated for each check element to see whether that element should be recategorized as a slide element. At intervals, the stresses on the inactive elements are updated based on the accumulated slip of the sliding elements since the last update, and the inactive elements are checked to see if they should be recategorized as check elements.

We implemented the slide list strategy in Case (6). Case (6) was otherwise identical to Case (1). The results were very similar but not quite identical. In both cases, there were 13 seismic events, each at nearly the same time and of the same magnitude. The moment at which the outermost element in the fracture was stimulated was 13.25 hours in Case (1) and 13.38 hours in Case (6). The discrepancy may

have occurred because we did not allow state to evolve with time for the inactive elements. Another possibility is that the stress updating of the active and check elements was not integrated into the Runge-Kutta scheme as they were in the sliding elements. In future work, we will investigate the cause of the discrepancy and investigate the optimal way to determine when to update stresses on inactive elements.

Loose Coupling

We are working on the implementation of a "loose coupling" strategy. Loose coupling in poromechanical simulation refers to the strategy of taking multiple time-steps in one part of the problem while periodically taking longer time-steps for another part of the problem (Kim et al., 2011). This strategy can be effective when different processes are taking place on different time scales.

In poromechanical reservoir simulation of a process such as subsidence, the mechanical deformation is gradual over time, and so multiple fluid flow time-steps can be taken between mechanical updates.

In our problem, the situation is reversed. Small time-steps (as short as microseconds) are needed to simulate very rapid slip. Very little fluid flow occurs on such short time scales. An appropriate loose coupling strategy would be to take relatively long fluid flow time-steps and repeat that time-step with multiple rate and state sub-time-steps. Figure 19 illustrates the scheme.

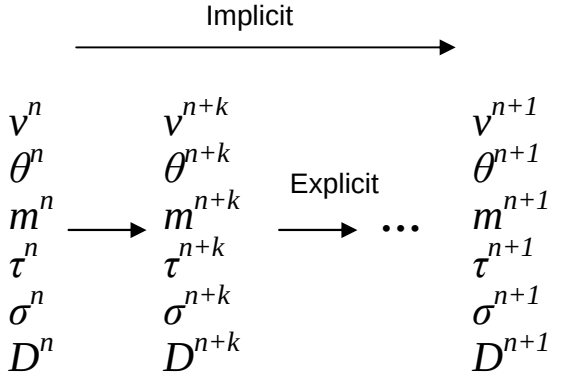

Figure 19: Proposed loosely coupled scheme.

From the implicit scheme, the mass in each element is known at time-step n+1. As time-step n+1 is repeated with sub-time-steps, the mass of fluid in each element is linearly interpolated from m^n and m^{n+1}. Shear displacement during the sub-time-steps change fracture aperture. If pressure were allowed to remain constant as shear dilation occurred, mass would not be conserved. To conserve mass, the fluid pressure of each grid block must be calculated at the end of each sub-time-step.

Separate Discretizations

We are experimenting with using different spatial discretizations for the flow and rate and state parts of the problem. Because the rate and state and the flow parts of the problem have different requirements for their spatial discretizations, they could be solved on different discretizations and values could be interpolated between the two.

The spatial discretization must be very fine for the rate and state calculations to ensure numerical stability. The flow simulations can get accurate solutions with a coarser discretization. Also, the Crouch and Starfield method requires almost equally sized elements for accuracy, but the flow simulations do not. Separate discretizations would allow the use of nonuniform spatial discretizations in the flow problem.