S1) could not meet the incompressibility constraint and yielded no solution (Tables 3.1 and 3.2; Figure 3.2). The difficulty is likely due to the larger $f_{\eta e}$ values along the slopes of the low viscosity zone that result from the form of the sigma function (Figures 3.1b and 3.2b). For example, in Model S1, the maximum $f_{\eta e}$ is 5.5, whereas $f_{\eta e}$ is 3.16 in Model L1 (Figure 3.2b). The sigma weak zone also has more nodes at the lower viscosity values, because form of the the sigma function curve flattens at the top and base. This could affect the representation of the low viscosity zone on the lower levels (coarser meshes) of the multigrid, which in turn could affect model run time and convergence as will be described in the sections that follow.

3.4.2 Mechanisms to transfer viscosity from upper to lower level

Models L7-17 and S7-S15 vary the mechanism for transferring the viscosity structure from the finest mesh to the lower multigrid levels (Tables 3.1 and 3.2). For each of the viscosity transfer options, $O_{t\eta}$, the viscosity averaging scheme was varied to include the arithmetic, geometric, and harmonic means (Tables 3.1 and 3.2). The maximum viscosity for these models is 1e5. The viscosity within the weak zone varies by 5 orders of magnitude. The low viscosity zone spans 21 nodes, and results for models with the low viscosity zone generated from the linear versus sigma function are compared.

The choice for the mechanism by which viscosity on the finest grid is represented on the succesively coarser grids can result in a difference in model run time by greater than 8% and, in some cases, determine whether a solution can be obtained (Figure 3.3). In general, models that use the node-to-node injection scheme (Option 0) and the node-to-node projection scheme (Option 1) result in faster run times and better

convergence than models that use the element-to-element injection scheme (Option 2) and the element-to-element projection scheme (Option 3) (Tables 3.1 and 3.2; Figure 3.3). The node-to-node transfer options give the most smoothing because information has to go through more interpolations as it is passed from integration points, to elements to nodes and then back again (Appendix C). In contrast, element-to-element transfer options, have the least amount of inherent smoothing, because the information is interpolated fewer times. Thus, we find that the goal in choosing a transfer option is to preserve a representation of the low viscosity zone on the coarser meshes (lower FMG levels), but to have enough smoothing as the viscosity is passed down to the coarser meshes to keep the effective $f_{\eta e}$ value low.

The importance of smoothing across the low viscosity zone as it is passed down to

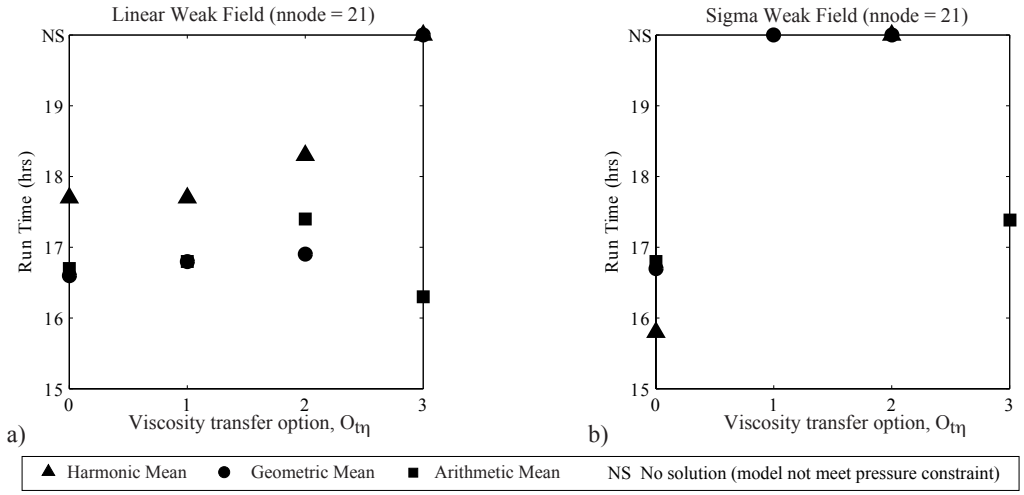

Figure 3.3: Run times for models that vary the viscosity transfer option and viscosity averaging schemes. Models with weak zone field generated from a linear function (a) and a sigma function (b). Option 0 uses a node-to-node injection scheme; Option 1 uses a node-to-node projection scheme; Option 2 uses an element-to-element injection scheme; and Option 3 uses an element-to-element projection scheme. Low viscosity zone spans 21 nodes.

lower levels is also supported by the run time dependence on the averaging scheme within the viscosity transfer options (Tables 3.1 and 3.2; Figure 3.3). The fastest run times occur for models using either a geometric or arithmetic mean, which tend to weight larger magnitude values more during averaging and therefore would tend to smooth away the lower viscosity values at the center of the weak zone (Tables 3.1 and 3.2; Figure 3.3). In contrast, the harmonic mean averaging scheme, which weights lower magnitudes values more, generally results in an increase in run times, suggesting the preservation of the lower viscosities within the weak zone may lead to higher effective values of $f_{\eta e}$ on the coarser levels, making the problem more difficult to solve. It is interesting to note that while most models run with option 3 did not converge to a solution, models that used this option with the arithmetic mean averaging did converge. This is likely due to the fact that the arithmetic mean weights the higher magnitude viscosity values most during averaging, and therefore is still able to provide enough smoothing even with the element-to-element transfer option. More tests need to be done to determine if this is independent of the the sampling of the weak zone for this particular model set up.

Models that use the sigma shaped low viscosity zone are more sensitive to the viscosity transfer options than models that use a linear low viscosity zone (Tables 3.1 and 3.2; Figure 3.3). The higher sensitivity of the models with the sigma shaped low viscosity zone may be a result of how the nodes are distributed within the weak zone field. The weak zone field generated from the linear function results in a constant factor of viscosity change, $f_{\eta e}$, across the elements in the low viscosity zone (Appendix B). In contrast, the factor of viscosity change, $f_{\eta e}$, is variable for models with the low viscosity zone generated from a sigma function. Therefore, depending

on which viscosity values are sampled during the transfer down to the coarser grids, the effective $f_{\eta e}$ on the coarser meshes could be very large in models that use the sigma low viscosity zone.

3.4.3 Number of multigrid cycles and relaxations within FMG

Models L18-L31 and S18-S32 vary the choice of the V versus W cycle and the number of smoothing iterations within the V and W cycles (Tables 3.1 and 3.2). The specific smoothing parameters varied within the V (or W) cycle are: R_{ds}, the number of relaxations on the velocity residual as it is passed down to the lower levels; R_{ll}, the number of relaxations on the lowest level; R_{us}, the number of relaxations as the correction is passed up to the finer levels; and R_{ul}, the number of relaxations on the uppermost level. For these models, the low viscosity zone spans 21 nodes. Results for models with the low viscosity zone generated from the linear versus sigma function are compared.

For models with the sigma weak zone field, the use of the W cycle can determine whether a solution is obtained (Table 3.2). For example, Model S1 did not meet the incompressibility constraint and no solution was obtained. However, if the W cycle is used instead (Model S18) the model converges, and the run time is more than 5% faster than the equivalent model with a linear weak zone field (Model L18). In contrast, the use of the W cycle on the model with the linear weak zone field (Model L18) yielded a 6% slower run time than the equivalent model that used the V cycle (Model L1) (Table 3.1). The improvement in convergence with the use of the W cycle on the model with the sigma weak zone field suggests that the W cycle is the better choice in the FMG for such a model that is having difficulty converging to the incompressibility constraint, otherwise the V cycle is sufficient. For these models, the

C Appendix: Low Viscosity Zone and Pre-existing Functionality in CitcomCU

In contrast to Appendix B, which includes a description of new functionality added to CitcomCU to use arbitrarily shaped low viscosity zones along a plate interface, Appendix C provides a brief background of pre-existing functionality in CitcomCU that may affect the effectiveness of changes described in Appendix B.

C.1 Overview

A Brief History of Citcom

The original CITCOM code, a cartesian finite-element mantle convection code, was written by Louis Moresi (Moresi and Solomatov [1995], Moresi and Gurnis [1996]). Subsequently, a variety of mantle convection codes have been developed that build off the original CITCOM code. CitcomS (Zhong et al. [2000]) is a parallelized spherical mantle convection code that implements the full multigrid algorithm with a consistent projection scheme. Shijie Zhong subsequently developed CitcomCU (Zhong [2006]) which is a parallelized spherical and cartesian thermo-chemical mantle convection code that also implements the full multigrid algorithm. CitcomS and CitcomCU are maintained by the Computational Infrastructure for Geodynamics (CIG) and can be downloaded from the CIG website at http://www.geodynamics.org/cig/. At UC Davis, CitcomCU-1.0.2 has been modified to read in a scalar three-dimensional low viscosity zone and to use a composite, non-linear rheology similar to that used in CitcomT (Billen and Hirth [2007]). The model results that are discussed in this paper were run with this modified version of CitcomCU, referred to as CitcomCU_ucd.

Multigrid and Full Multigrid Methods

CitcomCU uses the FMG to expedite convergence of the residual in the approximation to the velocity solution (Zhong et al. [2000]). The multigrid method (MG) is a numerical approach commonly used to accelerate convergence of the discretized form of an elliptic partial differential equation, such as Laplace's equation ($\nabla^2\Psi = \mathbf{0}$) (Press et al. [1997], Briggs et al. [2000], Wesseling [2004]). A telescoping set of successively coarser grids, referred to as multigrid levels, are used. The approximation to the solution is obtained through the multilevel correction scheme and by relaxation at each level. This combined procedure is implemented through V, or W, cycles, where the W cycle is equivalent to performing two V cycles for each level. At the top of the V, i.e. on the finest mesh, an initial approximation to the solution (usually zero) is used, and the residual is calculated. During the downward stroke of the V, the residual is passed down to the lower levels. At the bottom of the V, i.e. at the coarsest mesh, the residual equation is solved by direct methods, such as Gaussian elimination, or by an approximation method, such as the Gauss-Seidel or weighted-Jacobi method. During the upward stroke of the V, the error that was solved for on the coarsest mesh is passed up to the higher levels and is used to correct the original approximations on the finer meshes. By relaxing at each level during this process, both the short and long wavelength errors are reduced. The net result is to efficiently reduce the residual and expedite the convergence of the numerical approximation.

In the full multigrid method (FMG), the problem is defined on all of the grids, as opposed to only on the finest grid as in the MG. In addition, rather than starting with a guess on the highest level, i.e. the finest mesh, as was the case for the MG, in the FMG, an approximation to the solution is first calculated on the coars-

est mesh. This initial approximation is then used as a starting point for a series of nested telescoping V (or W) cycle iterations that work sequentially upwards to the finer mesh (Figure 3.7). A solution to the problem is determined for each level during this iterative process. In CitcomCU, the Gauss-Seidel method is used to relax on the downward and upward strokes of the V cycle. Gauss-Seidel is also used, instead of a direct method, to solve for the solution on the coarsest grid.

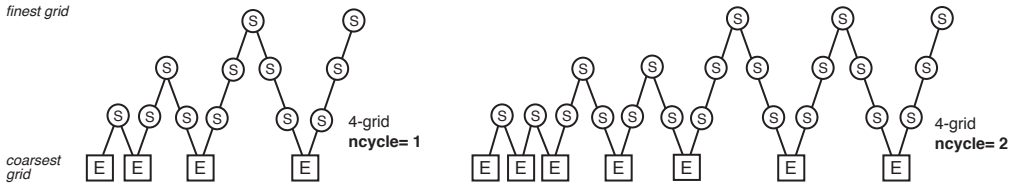

Figure 3.7: Illustration of the full multigrid algorithm of a four level system. The coarsest of the four grids is on the lower level and the grids become increasingly fine toward the highest level. E refers to exact solution, because a direct method is often used to solve on the coarsest grid. S denotes smoothing, such as by Gauss-Seidel. The full multigrid algorithm proceeds starting at the coarsest mesh and ending on the finest mesh (from left to right in each of these figures). *Left* Structure for a choice of one V cycle within the full multigrid algorithm. *Right* Structure for a choice of one W cycle (i.e., 2 V cycles) within the full multigrid algorithm. Image modified from Press et al. [1997].

C.2 Representation of Viscosity Field on Multigrid Levels

Project Viscosity Algorithm

We define the weak zone field, A_{wk}, *a priori* on the finest mesh of the finite element model, i.e., on the upper most multigrid level. During code execution, the viscosity structure, η_{eff}, is calculated and blended with the low viscosity zone, η_{wk}, also on the finest mesh. The viscosity on the finest mesh is transferred to the coarser meshes of the lower multigrid levels via the `project_viscosity` subfunction (Moresi and

Solomatov [1995]; Moresi and Gurnis [1996]; Zhong et al. [2000]).

In CitcomCU, there are four predefined options within the `project_viscosity` subfunction for how to transfer the viscosity, η_f, from a higher to lower multigrid level (Table 3.3). All four of the options begin on the higher level with the viscosity values located on the integration points. Then while still on the higher level, an average viscosity is calculated for each element (top two rows of Table 3.3). The four options then differ in whether the averaged elemental viscosity is partitioned onto the nodes on the higher level, whether it is further averaged while it is passed down to the lower level, and whether it is passed to the nodes, integration points, or elements on the lower level. In all cases, the final result is to have the viscosity on the integration points on the lower level.

Level	Description	Option 0	Option 1	Option 2	Option 3
Upper	Start with η_f on integration points	X	X	X	X
Upper	Average η_f for each element	X	X	X	X
Upper	Pass η_f to nodes of element	X	X	–	–
Transfer	Average η_f from nodes	–	X	–	–
Transfer	Average η_f from elements	–	–	–	X
Transfer	Transfer η_f to nodes on coarser mesh	X	X	–	–
Transfer	Transfer η_f to elements on coarser mesh	–	–	–	X
Transfer	Transfer η_f to integration points on coarser mesh	–	–	X	–
Lower	Pass η_f from nodes onto integration points	X	X	–	–
Lower	Pass η_f from elements onto integration points	–	–	–	X
Lower	End with η_f on integration points	X	X	X	X

Table 3.3: Mechanisms for transferring viscosity, η_f, from an upper to lower multigrid level. Options 0, 1, 2, and 3 correspond to the four choices in the project_viscosity subfunction. Option 0 uses a node-to-node injection scheme; Option 1 uses a node-to-node projection scheme; Option 2 uses an element-to-element injection scheme; and Option 3 uses an element-to-element projection scheme. The Level listed in the leftmost column refers to the multigrid level on which the action in the subfunction occurs. The X denotes the action taken by the respective option.

Viscosity Averaging Schemes

Several of the subfunctions that are called by `project_viscosity` in CitcomCU involve averaging of the viscosity, η_f, and use the arithmetic mean by default. In the arithmetic mean averaging scheme, the larger order of magnitude numbers will dominate (Beyer [2000]). This implies that the smaller order of magnitude values in a localized low viscosity zone may not be preserved after an averaging scheme is applied that uses the arithmetic mean. It may be worthwhile to preserve the smaller viscosity values in the low viscosity zone and to therefore use either the harmonic or geometric means. The harmonic mean better preserves lower magnitude values in a list (Beyer [2000]). The geometric mean will return values between that of the arithmetic and harmonic means, although the values returned from a geometric mean averaging scheme tend to be closer to those returned from using the arithmetic mean. To test the influence of the geometric and harmonic means on the projection of the viscosity to the lower multigrid levels for models that incorporate a localized low viscosity zone, we added these averaging schemes to the `visc_from_gint_to_nodes` and `visc_from_gint_to_ele` subfunctions located in the Nodal_mesh.c source file and to the `project_scalar` and `project_scalar_e` subfunctions located in the Solver_multigrid.c source file.

For reference, the formulas for the arithmetic mean, X_a, the geometric mean, X_g, and the harmonic mean, X_h are:

$$X_a = \frac{1}{n}(x_1 + x_2 + ... + x_n) \tag{3.21}$$

$$X_g = (x_1 x_2 \ \ ... \ \ x_n)^{\frac{1}{n}} \tag{3.22}$$

$$X_h = \frac{n}{\frac{1}{x_1} + \frac{1}{x_2} + ... + \frac{1}{x_n}} \tag{3.23}$$

where $X_a > X_g > X_h$ (Beyer [2000]). As an example, for the following set of values $\{10^{20}, 10^{21}, 10^{22}, 10^{23}, 10^{24}$, and $10^{25}\}$, the respective averages are: $X_a = 1.85 \times 10^{24}$, $X_g = 3.16 \times 10^{22}$, and $X_h = 5.40 \times 10^{20}$.

Three-Dimensional Lithosphere and Mantle Dynamics: Models of the Subduction-Transform Plate Boundary System in Southern Alaska

By

Margarete Ann Jadamec
B.S. (University of Connecticut) 1999
M.S. (University of Alaska) 2003

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

GEOLOGY

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Magali I. Billen

Donald L. Turcotte

Louise H. Kellogg

Committee in Charge

2009