Chapter 4

IMPROVEMENTS TO SAGE

4.1. INTRODUCTION

Chapter 3 described the implementation of coupled consolidation into SAGE, which was a significant addition to the program's capabilities. This chapter describes other important changes that were made to SAGE in the course of this research study. As mentioned previously, SAGE is a 2-D finite element program for static geotechnical engineering problems originally developed at Virginia Tech by Clark Morrison, under the supervision of Professors J. M. Duncan and George M. Filz (Morrison, 1995). The purpose of the enhancements made to SAGE by the writer was to increase the usefulness and flexibility of the program. The following sections describe these enhancements made to SAGE.

4.2. AXISYMMETRIC FORMULATION

The original version of SAGE implemented only plane strain analyses (Morrison, 1995). Although many geotechnical problems are suited to the plane strain assumption, many other problems fit the axisymmetric assumption. Therefore, SAGE was extended to use an axisymmetric formulation, in addition to the plane strain formulation. The implementation of axisymmetric conditions in SAGE was performed assuming the symmetry of both problem geometry and boundary conditions about the axis of revolution.

Four types of algorithms in SAGE were affected by the extension to axisymmetric conditions. The first was the generation of the element stiffness matrices. The second was the integration of element stresses to obtain nodal forces. The third was integration of strains to obtain stresses. The fourth was calculation of body loads and integration of tractions to obtain nodal loads.

SAGE was extended to axisymmetric conditions by creating two sets of subroutines for these algorithms. One set of subroutines is used for plane strain analyses and the other is used for axisymmetric analyses. For example, there are two subroutines in SAGE for the generation of the element stiffness matrices for 2-D elements. One subroutine generates the element stiffness matrix for plane strain conditions and the other generates the element stiffness matrix for axisymmetric conditions. The program determines which set of subroutines to use based on a flag in the input file, which identifies whether a plane strain or axisymmetric analysis is being analyzed.

The original version of SAGE included 2-D elements, beam-bar elements, zero-thickness interface elements, and Wilson elements (2-D + interface). The implementation of the axisymmetric formulation was performed for 2-D elements, but not for interface or Wilson elements. A shell element, the axisymmetric equivalent of a beam element, was implemented during the extension of SAGE to axisymmetric conditions.

4.3. ALGORITHM FOR CORRECTING CUMULATIVE ERROR PROBLEM IN SAGE

SAGE uses an incremental equilibrium approach for step-by-step analysis. Problems are modeled as a sequence of construction steps. Furthermore, each construction step is subdivided into substeps. This incremental approach allows complex time load histories to be modeled. Changes in body loads, tractions, boundary conditions, and removal and addition of material can be modeled in construction steps. Similar approaches are common to many geotechnical finite element programs.

The finite element equations, which in general are nonlinear because of material nonlinearities, are assembled and solved for each substep using the Newton-Raphson method or one of two quasi-Newton methods (line search or dogleg search). The loads and boundary conditions specified for each substep are increments of the changes in loads and boundary conditions for a construction step. The results of each substep calculation are incremental nodal displacements and changes in stress in elements. These incremental values are added to cumulative values from the end of the previous substep calculation to keep track of cumulative displacements and stresses.

The basic goal of the Newton-Raphson and quasi-Newton methods is to minimize the error in the solution. The primary measure of the error in the solution is the residual force vector, which is the difference between the applied loads and the loads due to the internal stresses in the finite elements at the end of each substep iteration. If the residual force vector is too large, the finite element equations are assembled again with the residual force vector as the load vector. This iteration continues until the residual force error is minimized to a user-specified tolerance. In the previous version of SAGE the residual forces were ignored at the end of each substep. This created a situation where errors could accumulate through a series of construction steps with many substeps.

Figure 4.1 illustrates an example SAGE analysis of a conventional drained triaxial test on a cylindrical clay sample. The Modified Cam Clay model was used to model the material response of the clay. The test consists of the following steps. First, a hydrostatic stress of 50 psf is specified as the initial condition. Next, tractions are applied in construction step 0 to represent the application of an all-around confining stress of 2032 psf. Finally, construction steps 1 through 15 model the loading of the sample by imposing a downward displacement of the top face of the cylindrical soil sample. Figure 4.1 also shows the final stresses in the soil from an analysis performed using the previous version of SAGE. The final lateral stress in the sample should be 2082 psf, but instead it is 2290.451 psf. The error in the lateral stress is due to cumulative error. Figure 4.2 is a plot of the lateral stress in the soil sample in construction steps 1 through 15. It can be seen that the lateral stress increases during each of these steps. Although the convergence criterion was met for every substep calculation in the entire analysis, the error in the lateral stress increases because the errors from each substep were ignored.

This tendency for cumulative error was obviously undesirable and posed a serious problem to the implementation of coupled consolidation. In a coupled consolidation analysis, each substep calculation represents an incremental time step. Coupled consolidation analyses could conceivably use very many incremental time steps, which meant that very large errors could accumulate and compromise the results of an analysis. Therefore, a solution that reduces the tendency for cumulative error was developed.

The solution implemented in SAGE was to add the residual forces from the end of a substep calculation to the substep load vector for the next substep. This simple change can be significant because it changes what the substep load vector represents. Instead of simply representing an increment of applied load as in the original SAGE, the substep load vector now represents the difference between the total applied force and the total force due to internal stresses at each substep. This is a more rigorous method of satisfying equilibrium, and hence produces a better solution, as can be seen in Figures 4.1 and 4.2.

4.4. PIEZOMETRIC LINES

A piezometric line is an input device that allows SAGE users to specify pore pressures in 2-D elements at different construction steps in an uncoupled analysis. Piezometric lines were implemented in SAGE to allow changes in piezometric levels (i.e. pore pressures) to be modeled in uncoupled analyses.

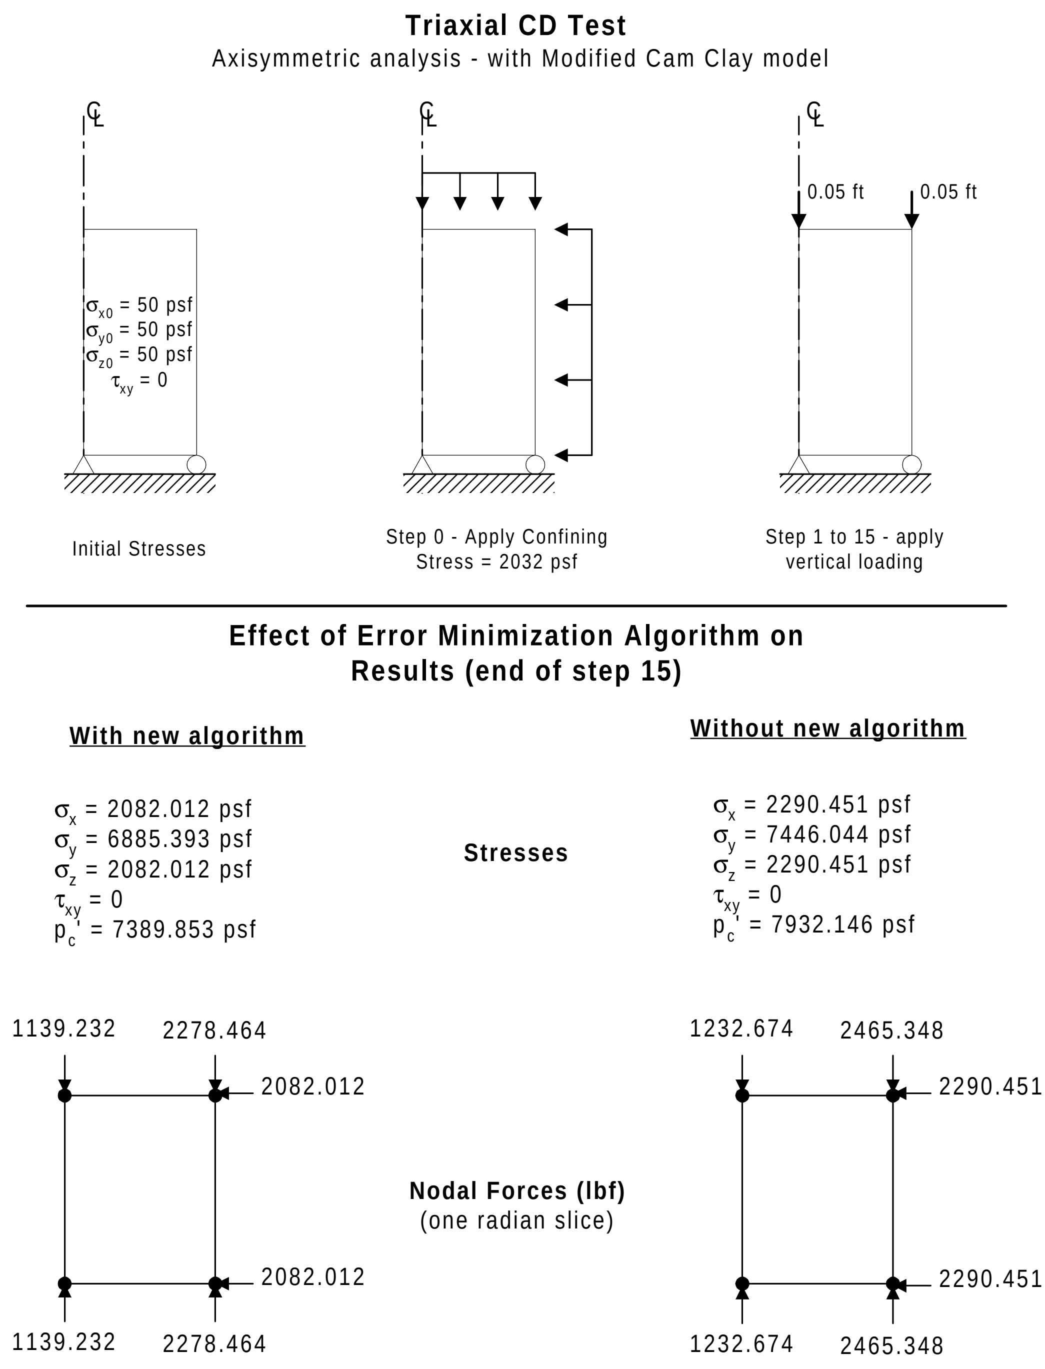

Figure 4.1. Example of cumulative error.

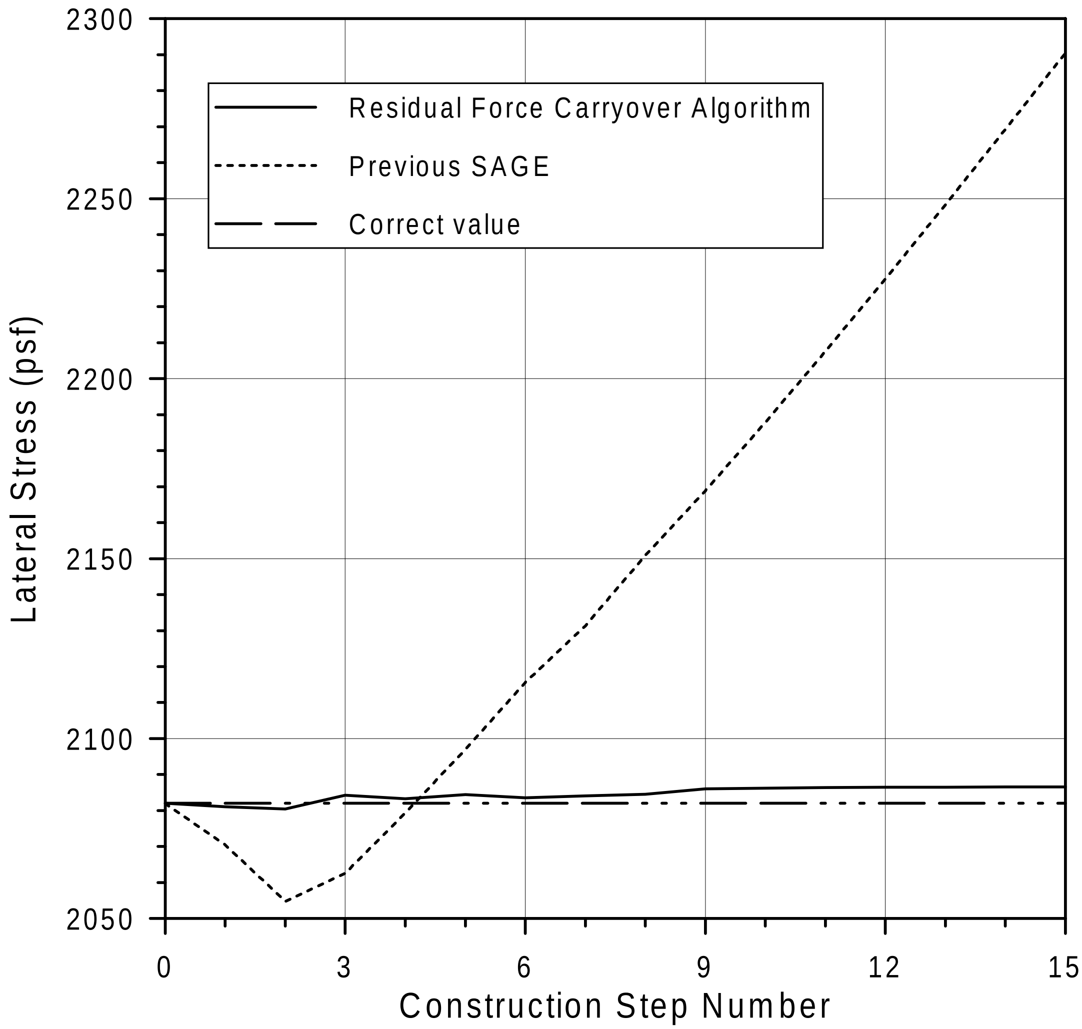

Figure 4.2. Effect of residual force carryover algorithm on example problem.

Pore pressures were not considered during uncoupled analyses in the previous version of SAGE. The following paragraphs explain how a piezometric line is used in SAGE.

Three pieces of information are associated with piezometric lines in SAGE. One piece of information is the piezometric line, which is defined by a series of coordinate pairs. The line segments connecting the coordinate pairs define the piezometric line. The second is the rate at which pore pressures increase with depth below the piezometric line, S_w. The third is a list of the 2-D material types influenced by the line. A piezometric line can affect more than one 2-D material and more than one piezometric line can be used in a construction step. This gives SAGE users flexibility when specifying pore pressures.

The information from piezometric lines is used to calculate pore pressure at Gauss points and pore water head at nodes of 2-D elements. Figure 4.3 illustrates how a piezometric line is used to calculate pore pressures and heads. The effect of pore pressures is included in uncoupled analyses in SAGE by adding the buoyant forces due to pore pressure to the load vector. Because SAGE uses incremental analyses, buoyant forces are calculated by integrating the change in pore pressure from one construction step to the next using the expression given in Equation 4.1.

$$F_b^e = \int_{\Omega^e} [B]^T \{u_{new} - u_{old}\} d\Omega \qquad \text{Equation 4.1}$$

u_{old} is the pore pressure at the end of the previous construction step. The buoyant forces due to a piezometric line are added to the construction step load vector.

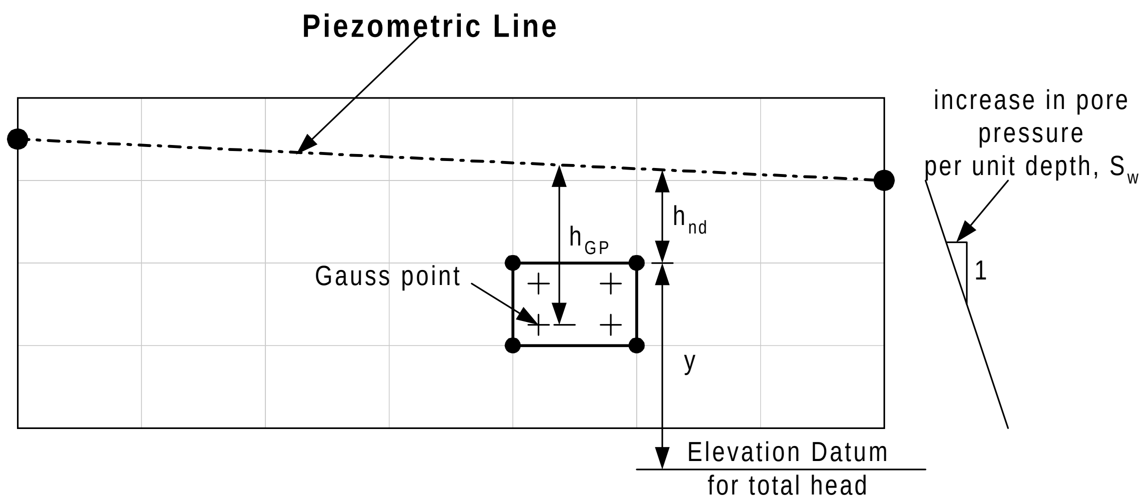

Pore pressure at Gauss point, $u_{GP} = h_{GP} * S_w$

Pore pressure at node, $u_{nd} = h_{nd} * S_w$

Head at node, $H_{nd} = u_{nd}/\gamma_w + y$

Note: A datum of y = 0 is usually convienent.

Figure 4.3. Piezometric lines.

Piezometric lines are also useful for establishing initial conditions for consolidation analyses. SAGE allows several construction steps to be performed using uncoupled calculations before a coupled analysis begins. Using piezometric lines with gravity turn-on calculations provides an easy way to specify initial conditions, because SAGE calculates initial effective stresses, pore pressures, and nodal heads. Without piezometric lines, gravity turn-on, and the initial uncoupled calculations, users would have to specify initial stress values at all Gauss points and nodal points in the input for coupled analyses.

4.5. STEADY STATE SEEPAGE MODULE

The addition of a steady state seepage module was another improvement made to SAGE. The steady state module is essentially a finite element seepage program within SAGE. The steady state module is used to solve confined and unconfined flow problems during uncoupled analyses with SAGE. The pore pressures calculated by the module are used to calculate buoyant forces acting on 2-D elements.

The steady state module was added to SAGE for two reasons. First, the module gives SAGE users the ability to solve unconfined and confined flow problems. Secondly, the steady state module aids in calculation of initial conditions for consolidation analyses. Using the steady state module establishes pore pressures, nodal pore water heads, and effective stresses consistent with the specified boundary conditions.

The development of the steady state seepage module is outlined in the following sections. First, the finite element formulation of the steady state flow problem used in SAGE is presented. Next, the implementation of the steady state seepage module in SAGE is discussed.

FINITE ELEMENT FORMULATION

Many finite element textbooks discuss the steady state flow problem and its finite element formulation (Zienkiewicz and Taylor, 1991; Griffiths and Smith, 1998; Reddy, 1993). The purpose of this section is to describe the governing equation and the finite element formulation used in SAGE.

The governing partial differential equation for steady state flow of a fluid through a rigid porous medium is the Laplace equation:

$$\nabla\left(k\nabla h\right)=0 \qquad \text{Equation 4.2}$$

where,

h = nodal head = $\dfrac{p_w}{\gamma_w}+\Omega$
k is the permeability of the soil
p_w is the pore water pressure
Ω is the elevation head

The boundary conditions for this problem are:

$\bar{h}=h$ on Γ_h
$\bar{q}=\vec{v}\cdot\vec{n}$ on Γ_v

The finite element formulation for the problem is given by:

$$\left[K_h\right]\{h\}=\{Q\} \qquad \text{Equation 4.3}$$

where,

$$\left[K_h\right]=\int_{\Omega}\left[B_h\right]^T\left[\kappa\right]\left[B_h\right]d\Omega$$
$$\{Q\}=\int_{\Gamma_v}\left[N_h\right]^T\{\bar{q}\}d\Gamma$$

$\left[B_h\right]$ = derivatives of $\left[N_h\right]$ (see Appendix A)
$\left[N_h\right]$ = vector of shape functions for nodal pore water head

$[\kappa]=\begin{bmatrix} k_{xx} & k_{xy} \\ k_{yx} & k_{yy} \end{bmatrix}$ is the permeability matrix

$k_{xx}=k_1\cos^2\theta+k_2\sin^2\theta$

$k_{xy}=k_{yx}=(k_1-k_2)\sin\theta\cos\theta$

$k_{yy}=k_1\sin^2\theta+k_2\cos^2\theta$

$\theta=$ the angle of anisotropy, referenced to x-axis; clockwise is positive

$\{h\}$ is a vector of pore fluid heads at the nodes

Permeability is assumed to be a function of pore pressure. This facilitates solution of unconfined flow problems. The three simple interpolation functions illustrated in Figure 4.4 are available in SAGE to specify how permeability varies with pore pressure. Two pairs of parameters are required for each function. One parameter pair is the unsaturated permeability, k_{unsat}, and the corresponding pore pressure, u_{unsat}. The second pair is the saturated permeability, k_{sat}, and the saturation pore pressure, u_{sat}.

The seepage forces, $\{F_u\}$, are calculated by integrating the difference in the pore pressures computed by the steady state module and the pore pressures at the end of the previous step.

$$\{F_u\}=\int[B]^T\{u_{new}-u_{old}\}d\Omega \qquad \text{Equation 4.4}$$

IMPLEMENTATION OF STEADY STATE SEEPAGE MODULE

The steady state module is invoked during uncoupled SAGE analyses when boundary conditions for pore fluid flow are changed. The steady state module is called during assembly of the construction step load vector. The steady state seepage module assembles the finite element seepage equations and solves them using direct (Picard) iteration. An iterative solution method is required because

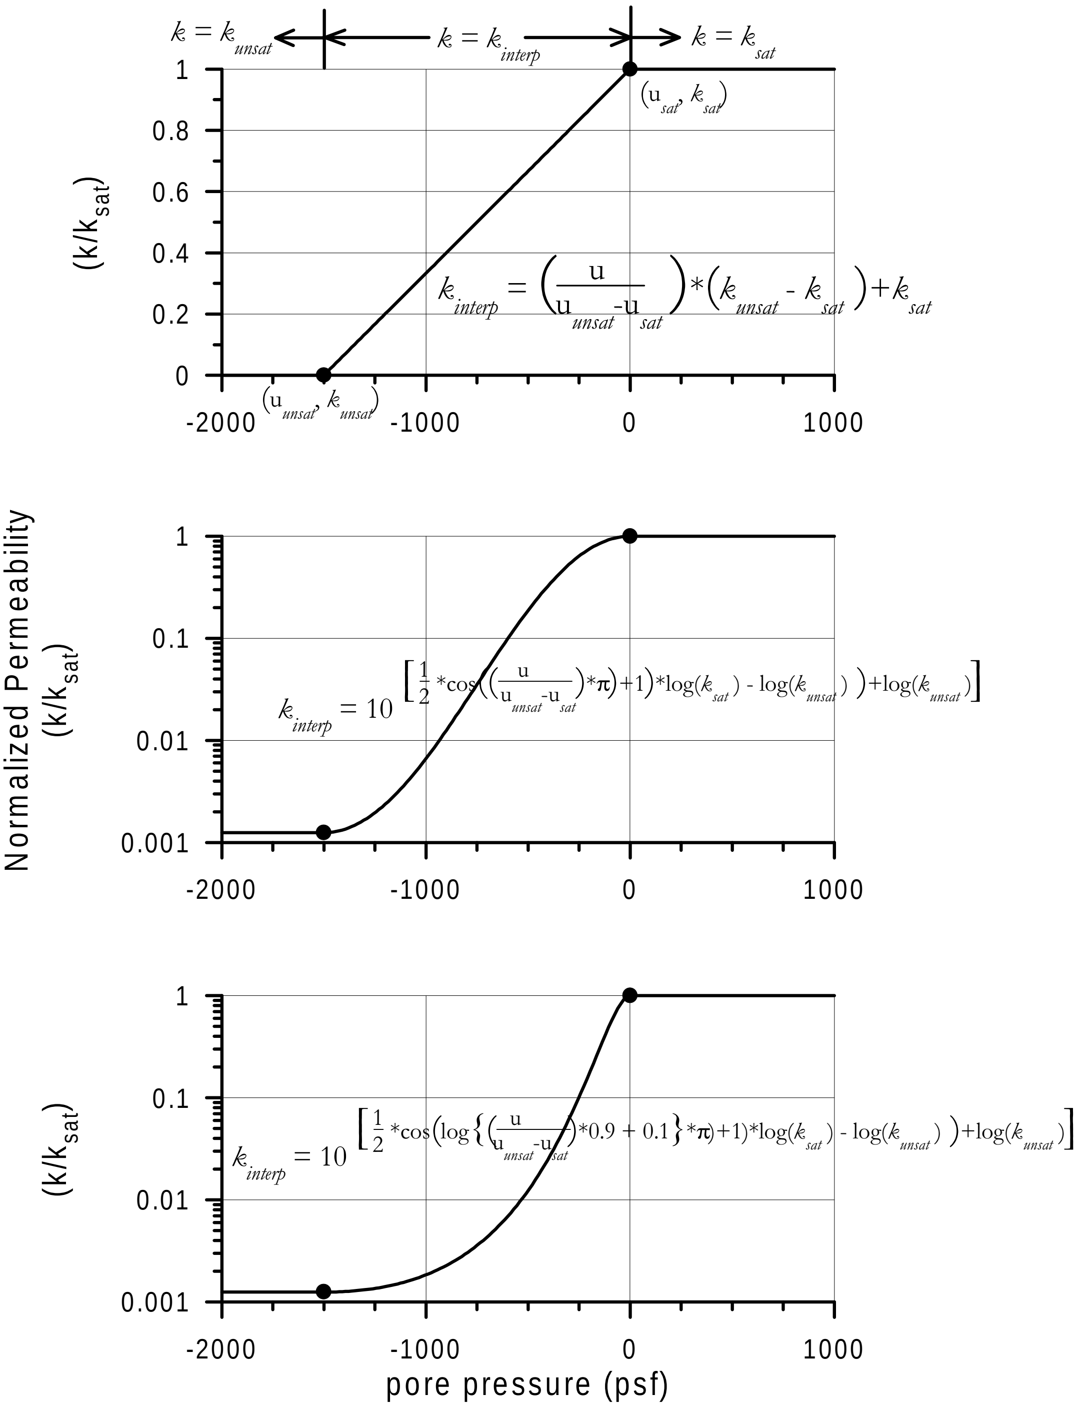

Figure 4.4. Permeability functions.

the seepage equations are generally nonlinear, because permeability is a function of pore pressure, which is directly related to nodal head, h. The euclidean norm of the nodal pressure head vector, $\left\{ p_w / \gamma_w \right\}$, is used in the convergence criterion. Convergence is achieved when the change in the pressure head norm between iterations is less than a user-specified tolerance. After convergence, seepage forces are calculated and added to the construction step load vector.

4.6. PINNED CONNECTIONS FOR BEAM ELEMENTS

Another feature added to SAGE was pinned connections for beam-bar elements. When two or more beam-bar elements were joined at a common node in the previous version of SAGE, the connection was rigid, and transferred moments between beams. The addition of the ability to model pinned connections between beam-bar elements helped to make SAGE a more adaptable modeling tool.

A pinned connection in SAGE is formed by two or more nodes, which are connected so that they share the same global degrees of freedom for x and y displacement but have independent rotational degrees of freedom. The result is a group of nodes (a pin group), which always have the same x and y displacement but have independent rotation. These nodes are connected so that they all experience the same x and y displacements, but each node rotates independently of the other nodes.

The modeling of a pinned connection in SAGE involves two steps. The first step is to specify in the input file to SAGE which nodes form a "pin". The second step is accomplished when the connectivity of the beam-bar elements meeting at the pin is specified. If two beam-bar elements are connected to different nodes in a pin, the elements will be able to rotate independently of each other at the pin. If two beam-bar elements are connected to the same node in a pin, they will rotate together at that node.

Pinned connections were implemented in SAGE by using global degree of freedom mapping arrays. A global degree of freedom mapping array is an array that assigns global degrees of freedom for each active local degree of freedom for all nodes in a finite element mesh. When a group of nodes is identified as a pin group, all of the nodes in the pin group are assigned the same global degree of freedom numbers for the x and y displacement degrees of freedom, and unique global degree of freedom numbers for rotation.

4.7. SEKIGUCHI AND OHTA MODEL

Another improvement to SAGE was the implementation of the Sekiguchi and Ohta model. The Sekiguchi and Ohta model is an elastic-plastic critical state model that is similar to the original Cam Clay model. The difference between the Sekiguchi and Ohta model and the original Cam Clay model is stress invariant used to measure deviatoric stresses.

The yield surface for the Sekiguchi and Ohta model is shown in Equation 4.5.

$$f = \frac{\lambda - \kappa}{1 + e_0} \cdot \ln\left(\frac{p'}{p'_c}\right) + D\eta^* = 0$$

Equation 4.5

where,

$D = \dfrac{\lambda - \kappa}{M(1 + e_0)}$ = dilatancy

λ = virgin compression slope for void ratio log p'

κ = recompression slope for void ratio log p'

e_0 = initial void ratio

M = critical state parameter

$p' = \frac{1}{3} \cdot \sigma'_{ii}$

σ' = effectivestress

$$\eta^* = \sqrt{\frac{3}{2}\left(\frac{s_{ij}}{p'} - \frac{s_{ij_0}}{p'_0}\right)\left(\frac{s_{ij}}{p'} - \frac{s_{ij_0}}{p'_0}\right)}$$
$$s_{ij} = \sigma'_{ij} - p' \cdot \delta_{ij}$$

The failure condition or critical state is reached when

$$\frac{q}{p'} = M \qquad \text{Equation 4.6}$$

Where,

$$q = \sqrt{3 \cdot J_{2D}}$$
$$J_{2D} = \det[s] \cdot \mathrm{trace}[s]^{-1}$$

When the value of the function f is less than 0 the stress-strain behavior of the soil is elastic and the elastic stress strain matrix, $[D^e]$, is used. The stresses in the soil must always lie on or within the surface defined by Equation 4.5 (i.e. $f \leq 0$). The consolidation pressure, p'_c, is the hardening parameter for the Sekiguchi and Ohta model. p'_c determines the size of the yield surface. The stress strain behavior is governed by the elastic-plastic stress-strain matrix, $[D^{ep}]$, when $f = 0$ or the yield surface is expanding. Associative plasticity is assumed.

$$[D^{ep}] = \frac{[D^e]\left\{\frac{\partial f}{\partial \sigma}\right\}\left\{\frac{\partial f}{\partial \sigma}\right\}^T [D^e]}{\left\{\frac{\partial f}{\partial \sigma}\right\}^T [D^e]\left\{\frac{\partial f}{\partial \sigma}\right\} + \mathrm{trace}\left(\left\{\frac{\partial f}{\partial \sigma}\right\}[I]\right)} \qquad \text{Equation 4.7}$$

The implementation of the Sekiguchi and Ohta model in SAGE followed the implementation of the Modified Cam Clay model. Morrison (1995) describes the implementation of the Modified Cam Clay model.

4.8. PRE-CONSOLIDATION PRESSURE PROFILE LINES

Several of the material models in SAGE keep track of the maximum past stress state. For example, the critical state Modified Cam clay and Sekiguchi and Ohta models store the pre-consolidation pressure, p_c'. There was no control over the initial pre-consolidation pressure values in the previous version of SAGE. p_c' was assigned automatically using the rules of each material model after initial stresses had been defined with a gravity turn-on analysis or explicit assignment of element stresses. The purpose of implementing pre-consolidation pressure profile lines into SAGE is to provide more flexibility in defining initial conditions.

A pre-consolidation pressure profile line is defined by a series of coordinate pairs. Elevation and pre-consolidation pressure are the variables that make up each coordinate pair. The sequence of points defines an elevation profile of pre-consolidation pressure. SAGE uses the profile of pre-consolidation pressure to interpolate pre-consolidation pressure at Gauss points of 2-D elements. Figure 4.5 illustrates a profile line and how it is used to interpolate pre-consolidation pressures at Gauss points.

4.9. RESTART ANALYSIS

Another addition to SAGE was the capability to perform restart analyses. In a restart analysis, the initial condition is based on the displacements, stresses, and forces from the results of another analysis using the same mesh. Restart analyses are useful in two ways. First, a large analysis with many steps can be broken into two or more smaller analyses. Second, restart analyses expedite running analyses with the same beginning (initial sequence of steps), but with different endings. For example, one analysis would be run from the beginning step to the last step that the analyses have in common. The remaining analyses would be restarted from the end of the first analysis.

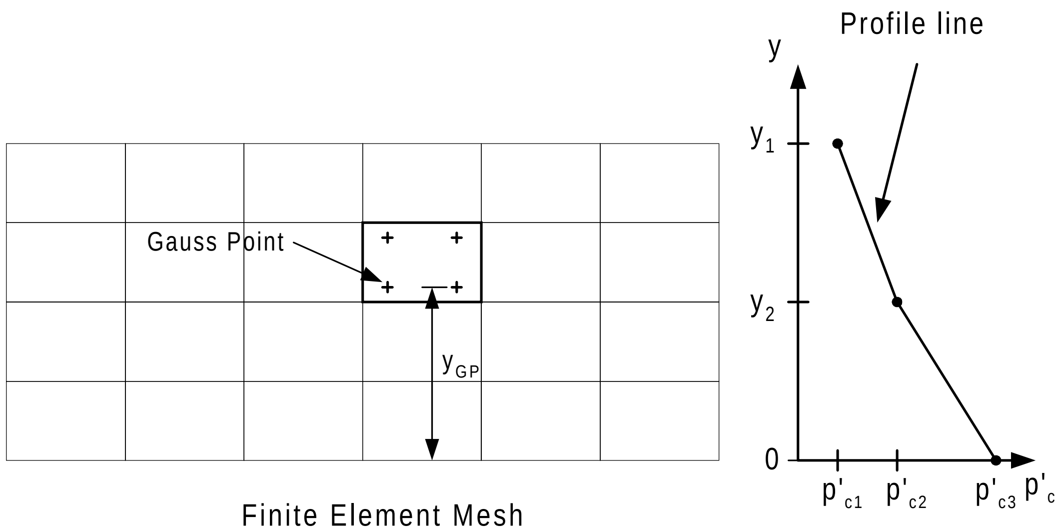

linear interpolation of p'_{cGP}:

$$p'_{c_{GP}} = \left(\frac{y_{GP} - y_1}{y_2 - y_1}\right)\left(p'_{c2} - p'_{c1}\right) + p'_{c1}$$

Figure 4.5. Profile lines.

The information required for a SAGE restart analysis includes:

1. Nodal displacement and head values
2. Nodal forces
3. Element forces
4. Element stresses
5. Element strains
6. Element state variables
7. Element attendance roster for the previous analysis step.
8. The global degree of freedom roster for the previous analysis step

Adding the restart capability involved creating two subroutines. One subroutine writes a file containing the information needed for a restart analysis, and the second reads this file and uses it to initialize the necessary arrays and variables.

A restart information file can be written after any construction step in an analysis. The user instructs the program to create the restart file by including a keyword (see Appendix D) in the SAGE input file. To run a restart analysis, the user includes a keyword and the name of the information file (see Appendix D) in the restart analysis input file.

4.10. CONCLUSION

The improvements described in this chapter enhance the capabilities of SAGE as a tool for modeling geotechnical problems. A significant amount of work was also performed to debug SAGE and improve its robustness. One of the principal challenges of having so many types of analyses and modeling tools in one program is to make them work together in a logical and reliable fashion. This challenge made some of the improvements, which would seem easy, difficult. Dealing with all element types in SAGE during the implementation of any of the improvements described in this chapter is the best example.