Tutorial Manual for

Groundwater Vistas

Version 6

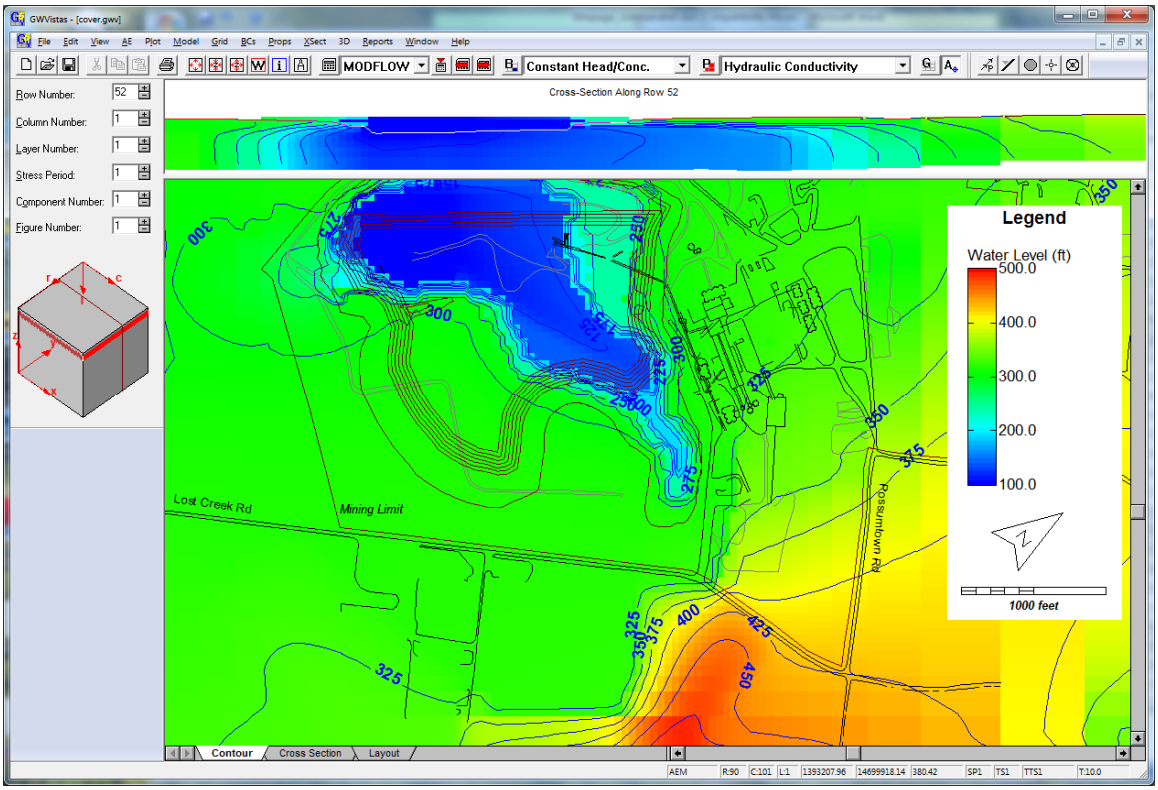

ESI Programming Team:
James O. Rumbaugh
Douglas B. Rumbaugh

Environmental Simulations, Inc.
P.O. Box 156
Reinholds, PA 17569
tel. (610) 670-3400
fax. (610) 670-9239

support@groundwatermodels.com
http://www.groundwatermodels.com

TABLE OF CONTENTS

Introduction

Groundwater Vistas Tutorials

The Groundwater Vistas tutorials used to be contained within the main GV manual. In Version 6 we have moved them here to their own manual to make it easier to use during Seminars and Webinars, as well as for use on your own as you explore the capabilities of the software. If you are new to Groundwater Vistas, it would be good to start with the Introductory Tutorial in the next chapter. That illustrates most of the functionality of the software on a simple model. The model used in the Introductory tutorial is also small enough that it can be used with the student version. That is, if you do not have a license for Groundwater Vistas you can still work through that first tutorial. The student version of GV is limited to 50 rows, 50 columns, and 4 layers.

The second tutorial called "Using GIS Data in Groundwater Vistas" is a more complicated model. This second model is more realistic and takes you through the steps in building a model based on data imported from a series of shapefiles. This tutorial also provides additional information on particle tracking and contaminant transport.

Later tutorials cover topics such as model calibration and optimization. The calibration tutorial is particularly extensive covering the use of Pest for advanced calibration techniques. As time goes on, more tutorials will be added.

Introductory Tutorial

The GV exercise, described below, introduces you to most of the important features of this software in a step-by-step example. You will be given very specific instructions to show how to use GV to design finite-difference models for MODFLOW. In a graphical user environment such as Windows, it is difficult to tell you exactly what to do during each step because many of the steps involve using the mouse. This demonstration provides several snap-shots of the GV screens to show you what your screen should look like, however, in case you miss a step. This tutorial also assumes that you are using ESI's MODFLOW^{win32} and MODFLOW2000^{win32}.

Starting a New Model

We will start this exercise by showing you how to create a new model using GV. First, double-click on the GV icon to start the program. You will see a small menu over a blank model design window. Select **File|New** from the main menu or click the new document button  . You will see a rather large dialog on your screen that asks you for basic information describing your model. These data are used to construct the initial model, which will have uniform row and column spacings and uniform layer thicknesses. All aquifer properties (hydraulic conductivity, storage, etc.) will initially be uniform (homogeneous). For simple modeling studies, you only need to add boundary conditions in order to have a complete model ready to run with MODFLOW.

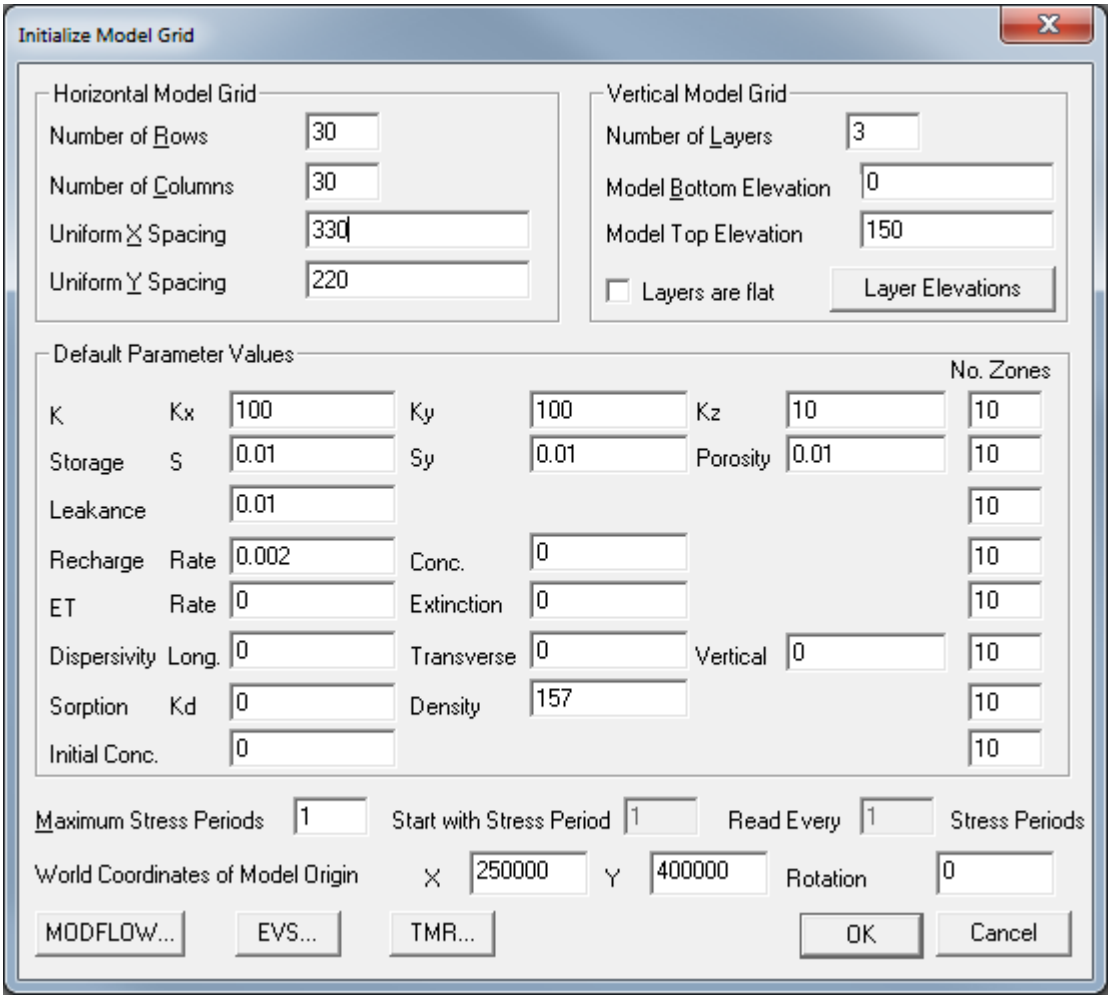

Now, fill in the dialog with the following information. When you are finished, press the OK button on the dialog. Before you choose OK, though, your screen should look like the one shown above.

Number of Rows	30
Number of Columns	30
X Spacing	330 ft
Y Spacing	220 ft
Number of Layers	3
Bottom Elevation	0.0 ft
Top Elevation	150.0 ft
Number of Stress Periods	1
Change Kz to	10.0 ft|d
Change Recharge rate to	**0.002 ft|d**

Enter World Coordinates of Model Origin: X= 250,000.0 Y= 400,000.0

You will notice that you can supply the world X and Y coordinates of the model origin (lower left corner) and an angle of rotation. In this case you should enter the real world coordinates of the lower left corner of the model grid. We will be importing maps from ArcView and this places the model into a real-world coordinate system. It is easier to import map files if you know in advance where the model should be placed on these maps.

After clicking the OK button, your screen should be similar to the one shown below. The model has uniform row and column spacings and the rows and columns are labeled.

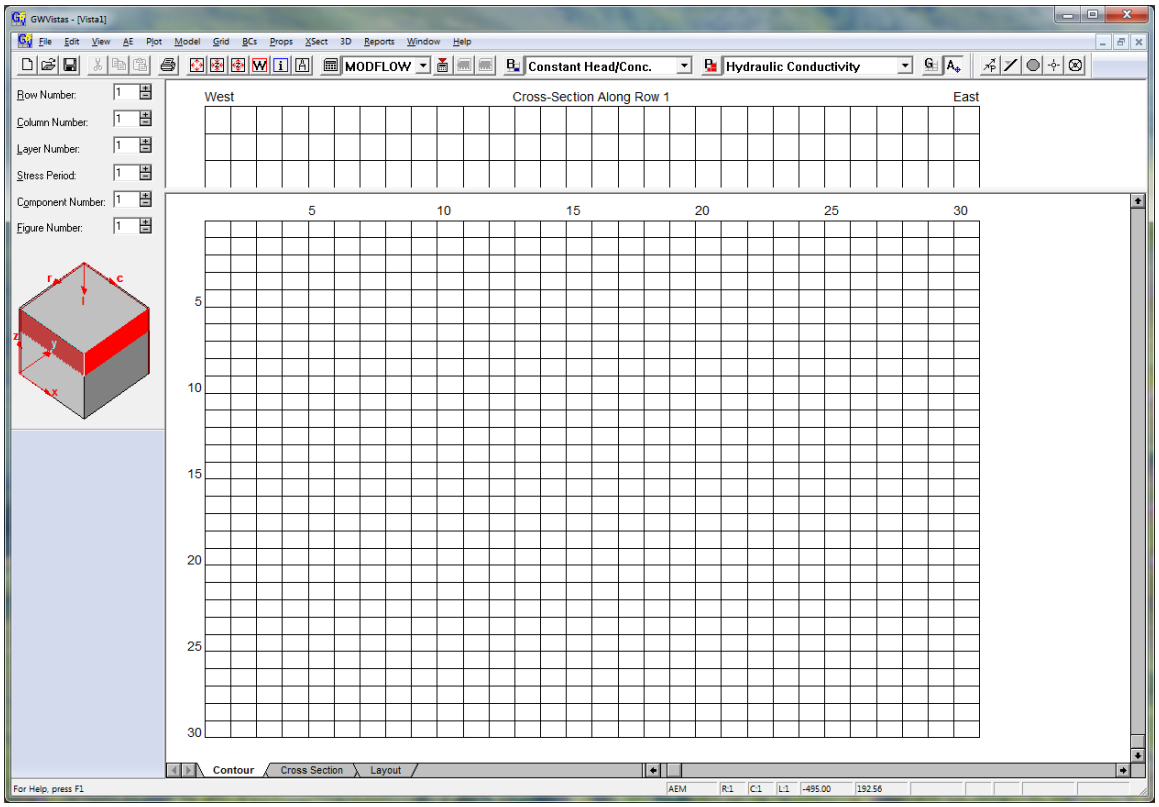

Now, let's change the font size used for the row and column labels. All text used to annotate the GV model may be modified in terms of font style and size. To change the font for the row and column labels, select **Grid|Options** and click the font button. Change the font size to 8 points and click OK to return to the Grid Options dialog. Click OK again to return to the GV menu.

You will now import three base map to display with the model. Select **File|Map|Shapefile**. A file open dialog will be displayed. Choose the shapefile called *river.shp*. Next, GV will ask you to enter a map file name. Enter *river.map* and click OK. Note that there is already a file called river.map in the tutorial directory. Either overwrite that map or use a different map name. Finally GV will ask you for a color to use on that shapefile. Choose blue. After clicking OK, you should see the river appear on the left side of your screen. Now repeat that process for the shapefiles called *site_bdy.shp* and *outcrop.shp* naming the maps *site_bdy.map* and *outcrop.map*. Your screen should look like the one below.

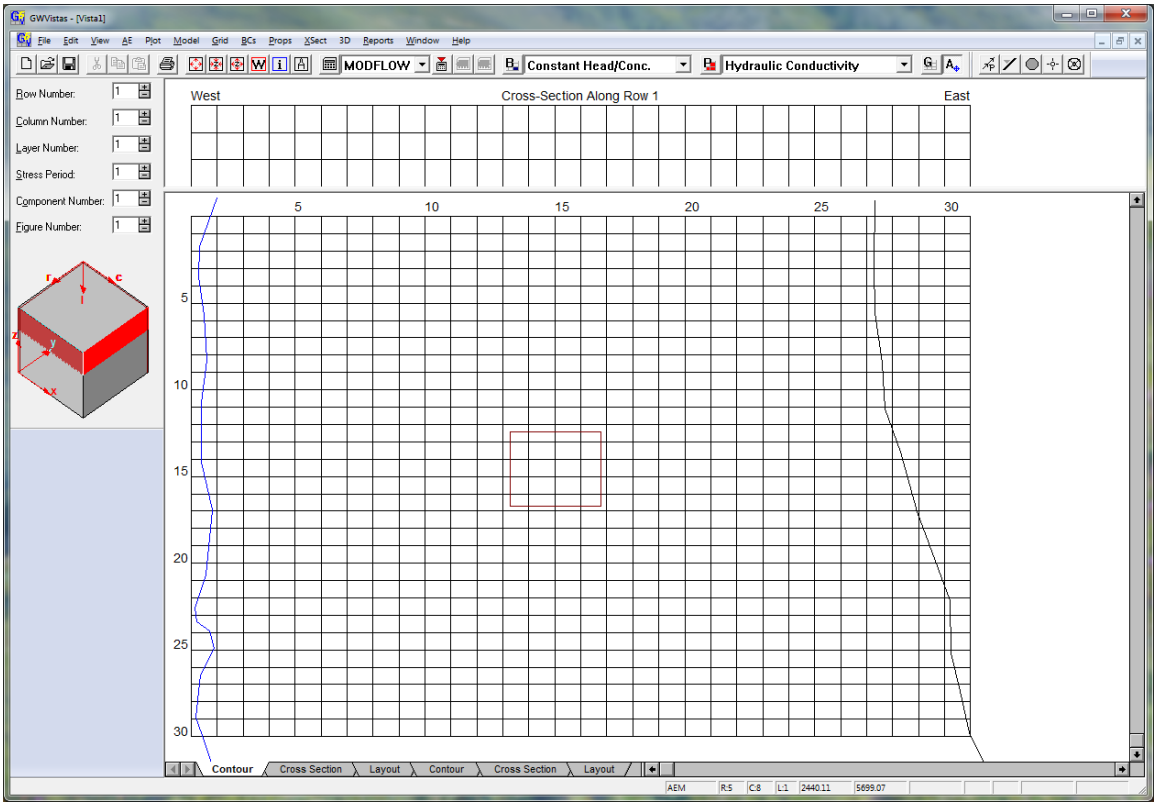

Adding Rows & Columns

GV has four different modes when designing the model. These include Analytic Elements, Grid, Boundary condition, and Property zones. The design operation that you may perform is determined on the Edit menu. Select **Edit** from the main menu. At the bottom of the pulldown menu you will see selections entitled Grid, Aquifer Properties, Boundary Conditions, and Analytic Elements. A check mark appears next to the option that is the current selection and the appropriate button is pushed down on the toolbar. The  button represents Analytic Elements,  represents Boundary Conditions,  stands for Property Zones, and  represents Grid operations. The Grid option allows you to add, delete, and move rows, columns, and layers. Aquifer Properties refers to assigning physical properties (e.g., hydraulic conductivity) to each cell in the model. Analytic Elements refers to the grid-independent boundary conditions in GV as well as annotation and calibration targets. You will see the buttons on the right side of the toolbar change depending upon which button is pressed down. This customization provides you with the most commonly used commands for each mode.

GV gives you the ability to insert, move, and delete rows, columns, and layers. In order to perform these operations, you must be in *grid* mode. This is accomplished by selecting **Edit|Grid** from the main menu or by pressing the  button on the toolbar. The word grid will appear in one of the panes of the status bar at the bottom of the GV window.

Once in grid mode, the cursor behaves differently than in other modes. When you are close to a row or column grid line, the cursor changes shape to either a left-right or up-down arrow. Pressing the left mouse button when this cursor appears allows you to slide the row or column line to a new position. You may not slide it beyond the adjacent row or column, however.

You may insert or delete rows, columns, and layers using the menu commands. These are fairly straightforward. Select the command (**Grid|Insert|Row** for example), move the cursor on the screen, and click the **left** mouse button. When deleting a row or column, the row or column closest to the cursor is deleted. Layers may be added above or below the current layer (the current layer is displayed as L:1,2,3,... on the status bar).

The right mouse button has a special use in GV. When you are in Grid mode, the **right** mouse button inserts a row or column into the model or deletes the nearest row or column without going through the menu as described above. The current grid operation (shown at the bottom of the Grid menu) determines what is added or deleted. To add rows or columns to the model, select **Grid|Insert Row** or **Grid|Insert Column** from the menu. To delete rows or columns from the model, select **Grid|Delete Row** or **Grid|Delete Column** from the menu. A check mark appears next to the type of action that GV will take when the right mouse button is clicked in Grid mode. The appropriate button is also pushed down on the toolbar ( to delete a row,  to insert a row,  to delete a column, and  to insert a column).

In this example, you will add two rows and two columns to the model. First, click the  button on the toolbar to enter Grid mode. Next, split row 15 into two new rows by placing the cursor anywhere within row 15 and click the right mouse button. Repeat this procedure for the next row to the south (Row 16 of the original model or new row number 17). When you insert a row or column, the default behavior is to split the current cell in half. You may change the way rows|columns are inserted by selecting **Grid|Options**.

Adding columns works the same way. Start by selecting **Grid|Insert Column** to place a check mark next to "Insert Column" on the Grid menu. Now split columns 14 and 15 just like you did for rows 15 and 16 above. Position the cursor within Column 15 and click the right mouse button. Repeat for the original column 14. Your screen should look like the one shown below:

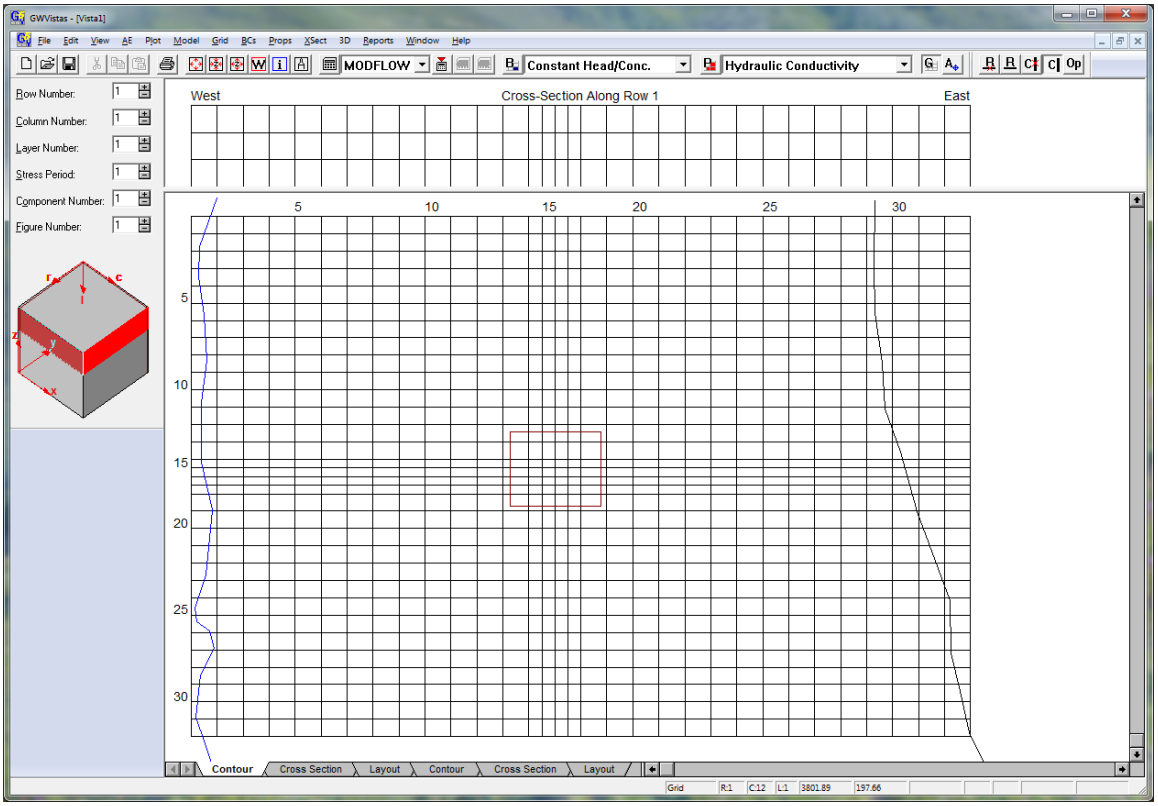

Adding Boundary Conditions

You will now select Boundary Conditions as the current design mode so that you can add boundary condition cells to the model design. Select **Boundary Conditions** on the Edit pulldown menu. You may also click the toolbar icon containing the large letter B ().

In this example, you will add a column of constant heads along the left edge of the model in layer 1. You will then add two wells in the bottom layer (layer 3) of the model.

The easiest way to set a large number of boundary conditions is to use the Window command. Select **BCs|Insert|Window** from the main menu (or  from the toolbar). The cursor will change shape and appear like a mini-finite-difference grid. Move the cursor to the upper left corner of the model (row 1, column 1) and press the left mouse button. Make sure the cursor is inside the grid before pressing the left mouse button. Hold the mouse button down and move the cursor to the lower left corner (row 32, column 1). Release the mouse button and a dialog appears to confirm the coordinates of the window that you just created. Simply press the OK button to accept these coordinates. Next, a constant head dialog appears. The only item that must be changed is the value of constant head. **Change this value to 150 ft**. A common mistake here is to forget to reset the head to 150 from the default of 0.0. This results in constant head values below the layer bottom and the whole model goes dry!

Notice that the boundary condition is *steady-state.* This means that the boundary cells will be active during the entire simulation. In MODFLOW, constant heads are active for the entire simulation and cannot be changed. Your screen should now look like the one shown below.

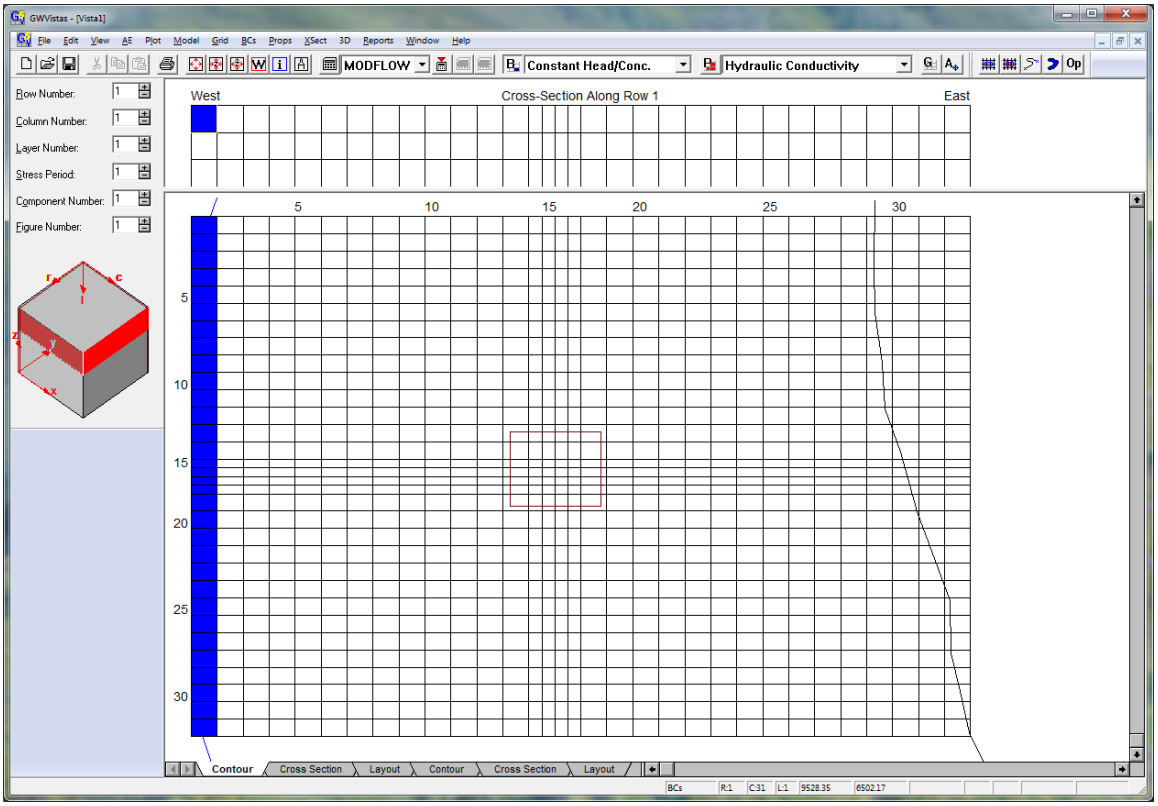

Now, move to layer 3 of the model. The easiest way to change layers is by clicking the "+" button next to "Layer" on the 3D cube (called the Reference Cube) that is on the left side of the screen. Click the "+" button twice to get to Layer 3. The model will be redrawn and the constant head cells will disappear. This happens because these constant head cells were defined in layer 1 (the top layer) and we are now viewing the bottom layer of the model. You should still see one constant head in the upper left corner of the cross-section view however.

Select **BCs|Well** from the main menu. This places a check mark next to the word "Well" indicating that we are now editing wells. You may also notice that GV now has a dropdown list of boundary conditions on the toolbar. You can simply select the BC type you want from that list as well.

Next, select **BCs|Insert|Single Cell**. Move the cursor to row 16, column 15 and click the left mouse button. (You could also add a well by simply moving the cursor to row 16, column 15 and clicking the right mouse button.) You will notice that as you move the cursor on the screen, the status bar at the bottom of the screen updates the current row and column. When you are at row 16, column 15 the status bar should read "R:16 C:15". After clicking the mouse button, a well edit dialog will appear. (Note that in MODFLOW, negative flow rates are used for production and postive rates for injection). Enter a flow rate of -30000 ft^3|d. GV assumes that units of length and time are consistent for all parameters. In this example, all length units are in feet and all time units in days. Therefore, hydraulic conductivity is in ft|d, well flow rates are in ft^3|d, recharge is in ft|d, etc. Enter a name for the well of *Well 1* as shown below:

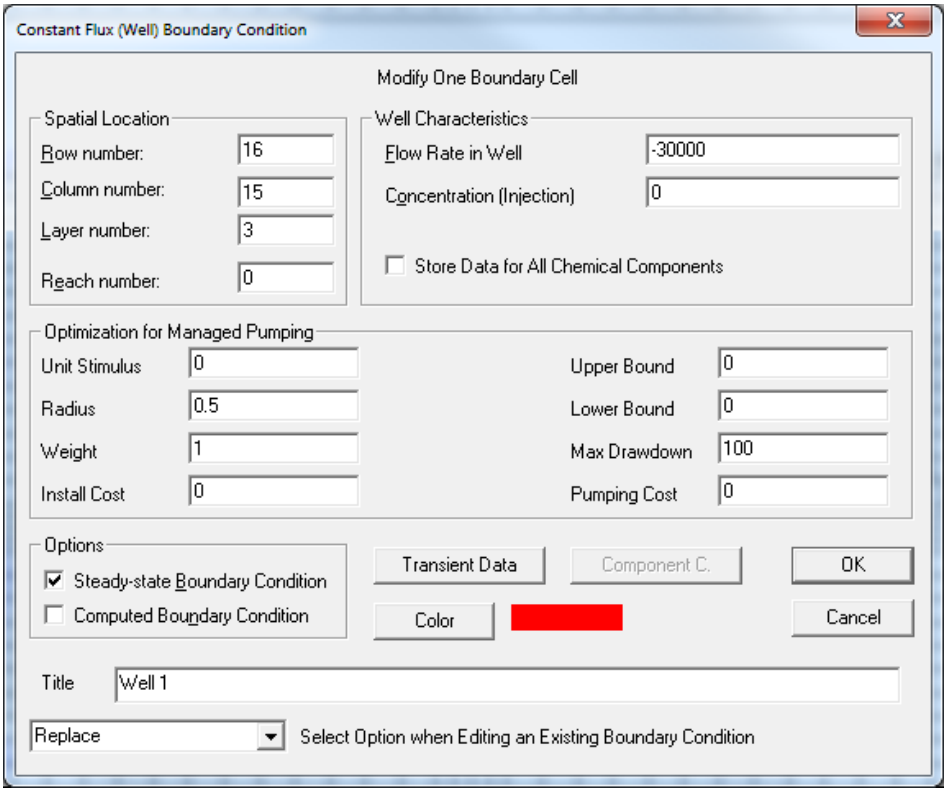

Add another well at row 8, column 16. This time, enter a flow rate of -40000. Use the same sequence of events that you used to insert the first well. Your screen should not look like the one below.

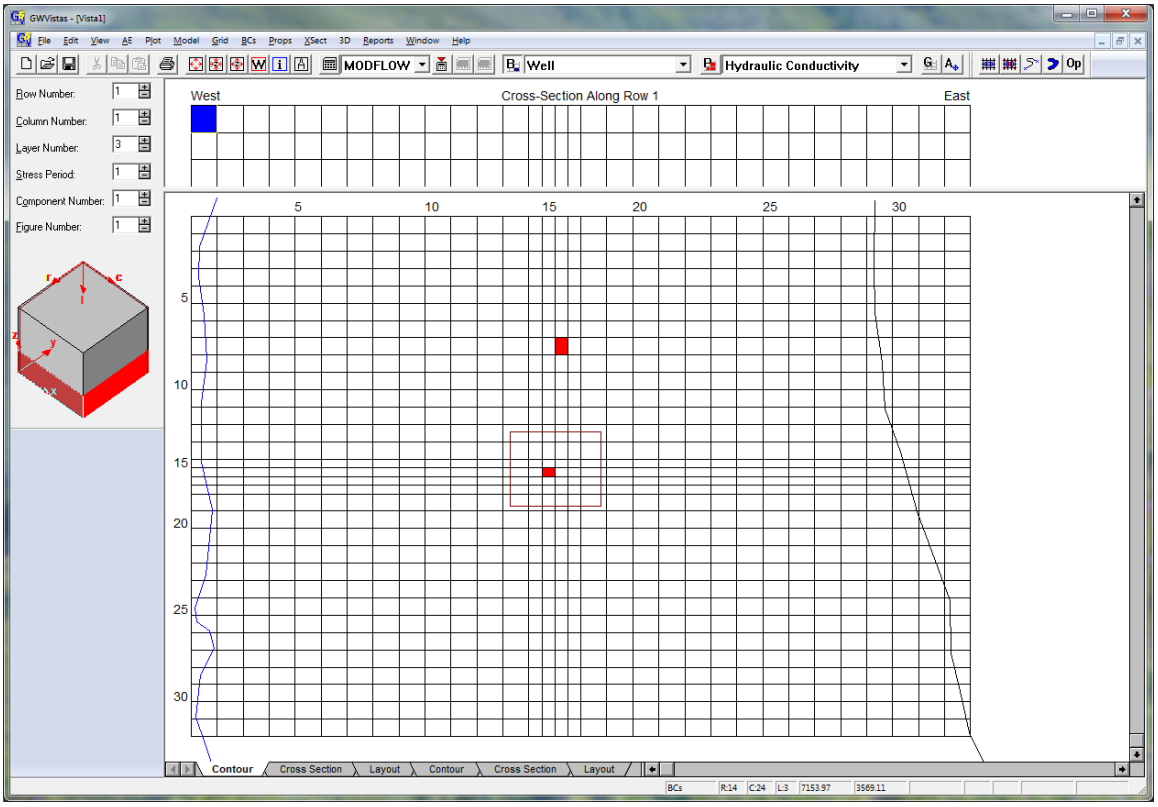

GV can display a text message next to boundary condition cells. You entered the names of the wells above. To display them, select **BCs|Options** and turn on *Plot BC Titles*. You should also change the font to 10 points. Now, you should see the well names next to the wells.

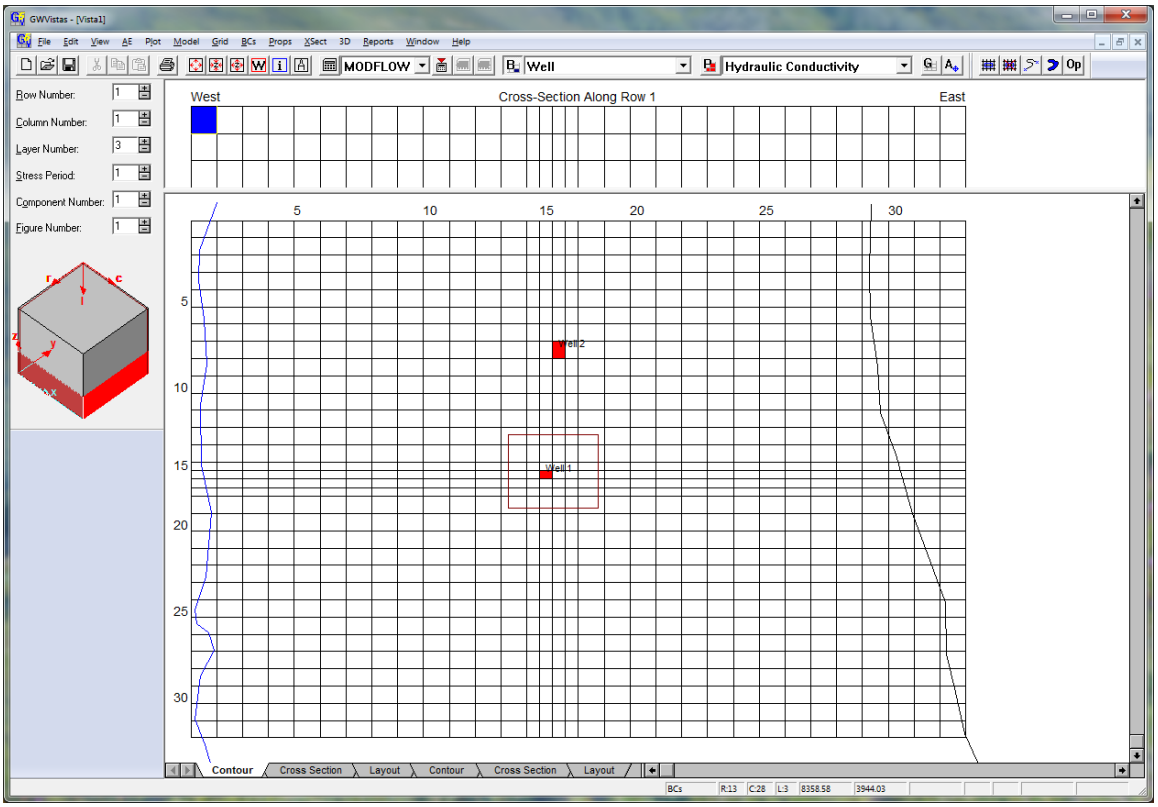

Changing Layer Elevations

Most models use varying layer elevations to follow some stratigraphic horizon. The easiest way to accomplish this is to import SURFER files. SURFER is a very nice program for contouring data and Groundwater Vistas can import SURFER grid files for any model property. We will first import a SURFER file for the top of layer 1. By default, Groundwater Vistas assumes that the bottom of layer 1 is the same as the top of layer 2, etc. That means that under most circumstances, you only need to worry about the top of layer 1 and then the bottoms of all other layers.

First, go back up to layer 1. Start by selecting **Props|Top Elevation**. Now, select **Props|Import|SURFER**. Find the file called *land_surface.grd*. Make sure to check the box labeled *SURFER file in Site Coordinates*, which means that the SURFER grid file is in the same coordinate system as your base maps.

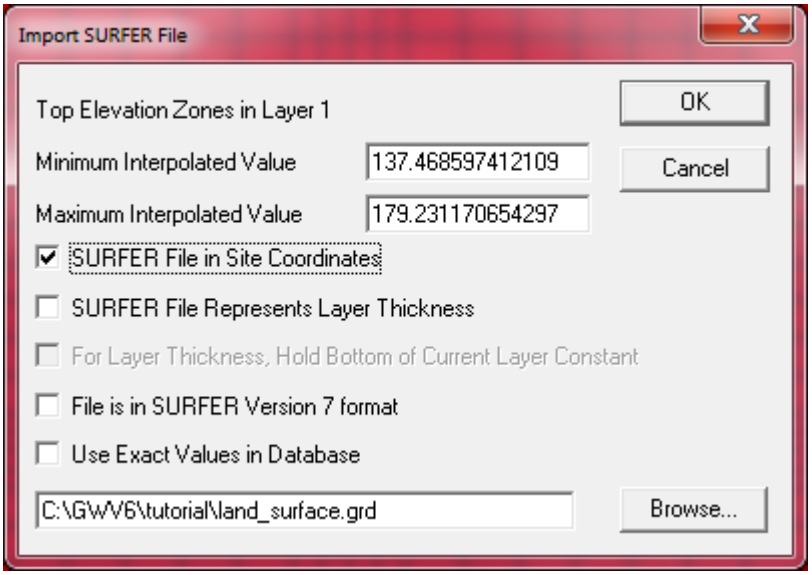

Click OK and you should see the top of layer 1 change to a curved surface in the cross-section window. To adjust the color of the grid cells for the new elevations, use **Props|Property Values|Reset Matrix Bounds**.

We will now import the bottom elevations of the other layers. Select **Props|Bottom Elevation** and then **Props|Import|Surfer**. The files are called *layer1_bottom.grd*, *layer2_bottom.grd*, and *layer3_bottom.grd*. Be sure to change to the correct layer before importing the SURFER files and also check the option labeled *SURFER File in Site Coordinates*. After importing bottom elevations for all layers, your screen should look like the one below:

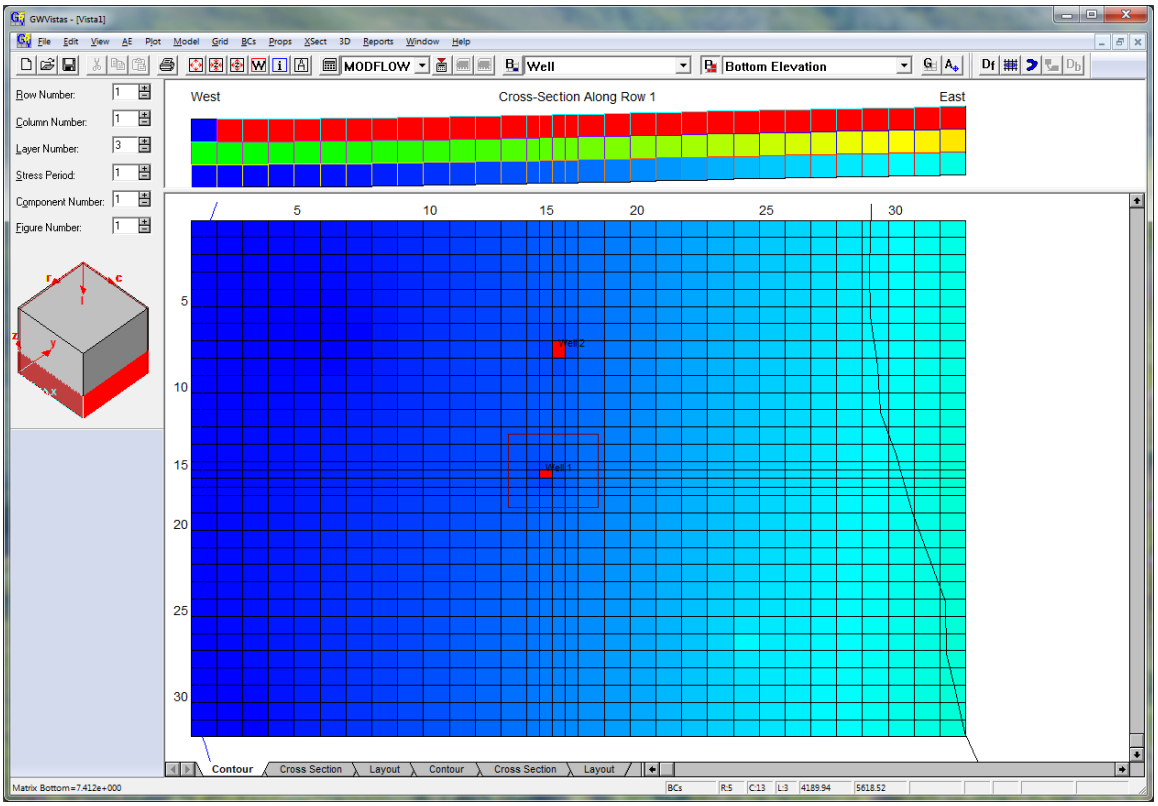

By default, Groundwater Vistas uses matrices of values for layer elevations and starting heads. This means that you no longer need to worry about a zone database for layer elevations. This makes working with

layer elevations much easier. If you want to edit the layer elevations, you simply select **Props|Property Values|Matrix Editor**. The matrix editor is like a spreadsheet that shows you the elevations in your model. The default format is exponential. Change this to "Fixed" next to Format at the bottom of the spreadsheet window. Do that now and your screen should look like the following one for layer 1.

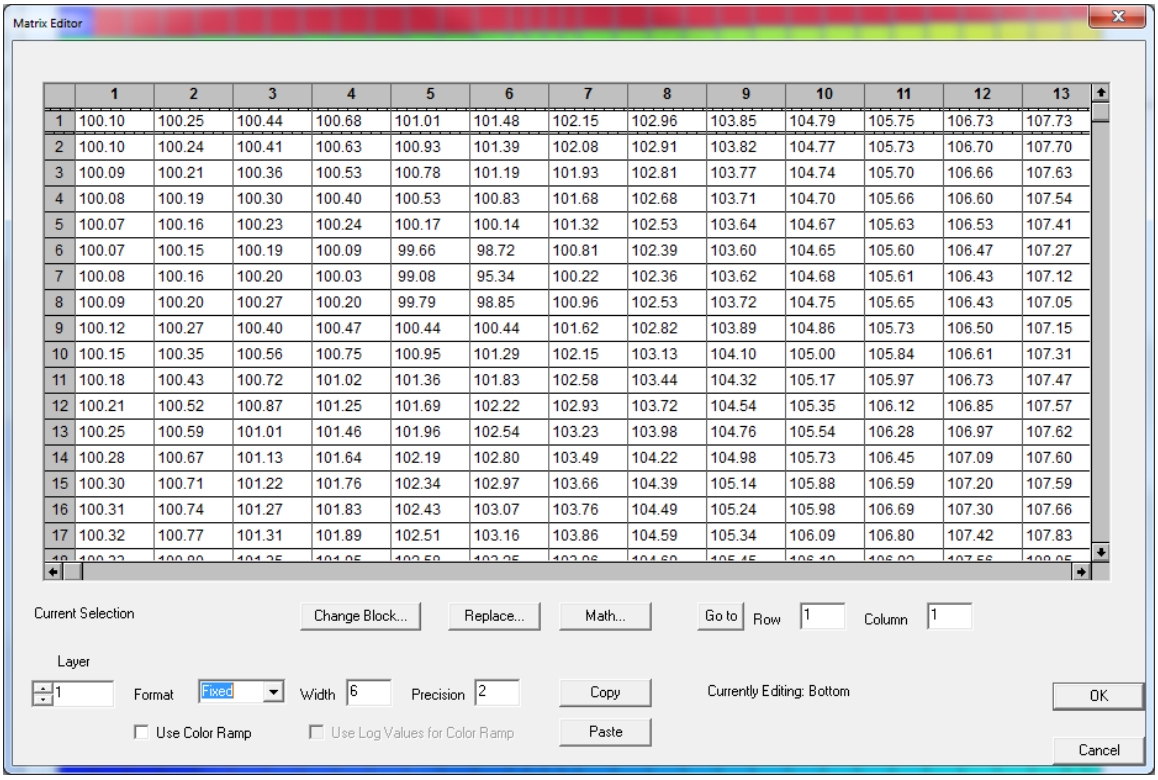

	1	2	3	4	5	6	7	8	9	10	11	12	13
1	100.10	100.25	100.44	100.68	101.01	101.48	102.15	102.96	103.85	104.79	105.75	106.73	107.73
2	100.10	100.24	100.41	100.63	100.93	101.39	102.08	102.91	103.82	104.77	105.73	106.70	107.70
3	100.09	100.21	100.36	100.53	100.78	101.19	101.93	102.81	103.77	104.74	105.70	106.66	107.63
4	100.08	100.19	100.30	100.40	100.53	100.83	101.68	102.68	103.71	104.70	105.66	106.60	107.54
5	100.07	100.16	100.23	100.24	100.17	100.14	101.32	102.53	103.64	104.67	105.63	106.53	107.41
6	100.07	100.15	100.19	100.09	99.66	98.72	100.81	102.39	103.60	104.65	105.60	106.47	107.27
7	100.08	100.16	100.20	100.03	99.08	95.34	100.22	102.36	103.62	104.68	105.61	106.43	107.12
8	100.09	100.20	100.27	100.20	99.79	98.85	100.96	102.53	103.72	104.75	105.65	106.43	107.05
9	100.12	100.27	100.40	100.47	100.44	100.44	101.62	102.82	103.89	104.86	105.73	106.50	107.15
10	100.15	100.35	100.56	100.75	100.95	101.29	102.15	103.13	104.10	105.00	105.84	106.61	107.31
11	100.18	100.43	100.72	101.02	101.36	101.83	102.58	103.44	104.32	105.17	105.97	106.73	107.47
12	100.21	100.52	100.87	101.25	101.69	102.22	102.93	103.72	104.54	105.35	106.12	106.85	107.57
13	100.25	100.59	101.01	101.46	101.96	102.54	103.23	103.98	104.76	105.54	106.28	106.97	107.62
14	100.28	100.67	101.13	101.64	102.19	102.80	103.49	104.22	104.98	105.73	106.45	107.09	107.60
15	100.30	100.71	101.22	101.76	102.34	102.97	103.66	104.39	105.14	105.88	106.59	107.20	107.59
16	100.31	100.74	101.27	101.83	102.43	103.07	103.76	104.49	105.24	105.98	106.69	107.30	107.66
17	100.32	100.77	101.31	101.89	102.51	103.16	103.86	104.59	105.34	106.09	106.80	107.42	107.83

Click OK when you are done exploring the matrix editor.

Assigning Aquifer Property Zones

Aquifer properties that do not use matrices are defined using zones of equal value. See the Concepts chapter for more information on zones. When defining a parameter with zones, you must first add a zone to the property database and then assign that zone number to a cell or group of cells in your model. We will illustrate this concept by adding a zone of lower hydraulic conductivity in this model.

Start by selecting **Props|Hydraulic Conductivity**. You should see that each cell is colored white and if you move the cursor around, you will see that the hydraulic conductivity is 100 ft|d, as shown on the left side of the status bar. You will now add a zone of hydraulic conductivity of 25 ft|d. First, you must define this zone in the database. Select **Props|Property Values|Database**. The database should look like the following:

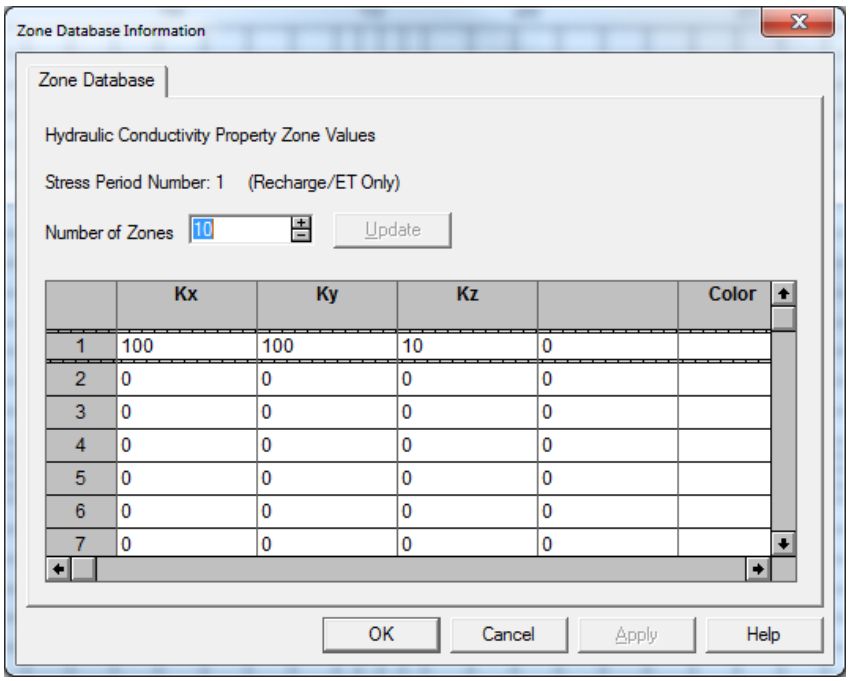

There are 10 zones defined in the database but only 1 is being used. Zone 1 has values of Kx, Ky, and Kz of 100, 100, and 10, respectively. Enter a value of 25 for Kx, Ky, and Kz in zone 2. Also, change the color to blue. You can change the color by double-clicking on the color column in the database next to zone 2. Your database should now look like the following:

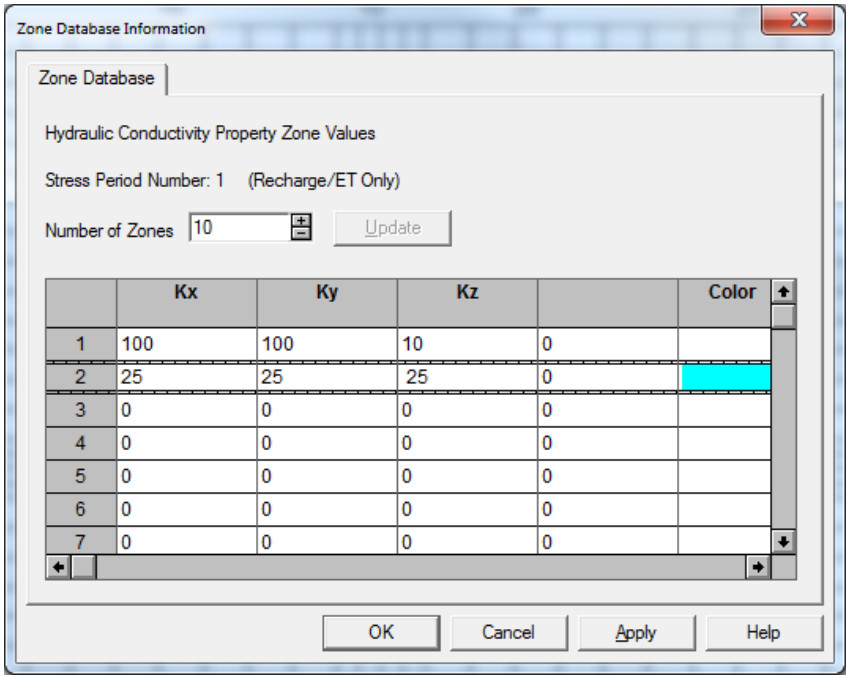

You now have defined zone 2 as having an isotropic conductivity of 25. No cell in the model has this value yet, however. Now, move to layer 3. Select **Props|Set Value or Zone|Window** and then drag a window covering the last eight columns of the model. After dragging the window, enter a zone number of 2 and your screen should look like the following:

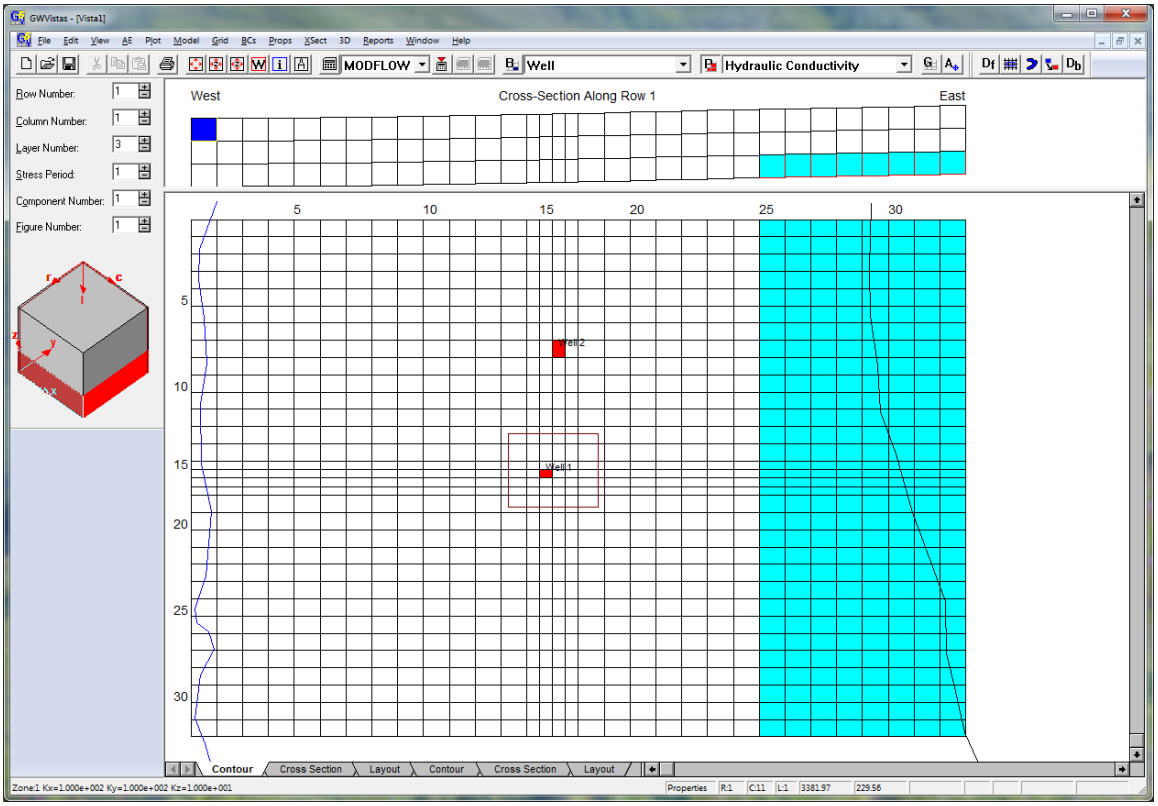

You should see that all rows in columns 25 through 32 turn blue. If you move the cursor over these cells, you will see 25 for Kx, Ky, and Kz on the status bar in the lower left corner of your Groundwater Vistas window. **Repeat this for layers 1 and 2 as well so that each layer has a block of zone 2 on the east side of the model**.

Setting No-Flow Boundaries

A no-flow boundary is a special kind of constant flux boundary where the flow is set to zero. This means that the cell is not part of the flow system. MODFLOW calls such cells *inactive*. You define no-flow boundaries in Groundwater Vistas on the **BCs** menu. You will now make cells east of the outcrop line no-flow.

Go to layer 1 and select **BCs|No Flow**. Now, select **BCs|Insert|Polygon**. Click the left mouse button as you move the cursor around the east side of the model and around the outcrop line you imported from the *outcrop.shp* ArcView shapefile. When you reach the point where you started digitizing, double-click to end the digitizing. Make sure to outline the whole area. For example, start at the northern end of the outcrop, digitize to the south, but then you also need to digitize east of the grid and back up to the north before double-clicking. You should see black cells in the upper right corner of your model, as shown below:

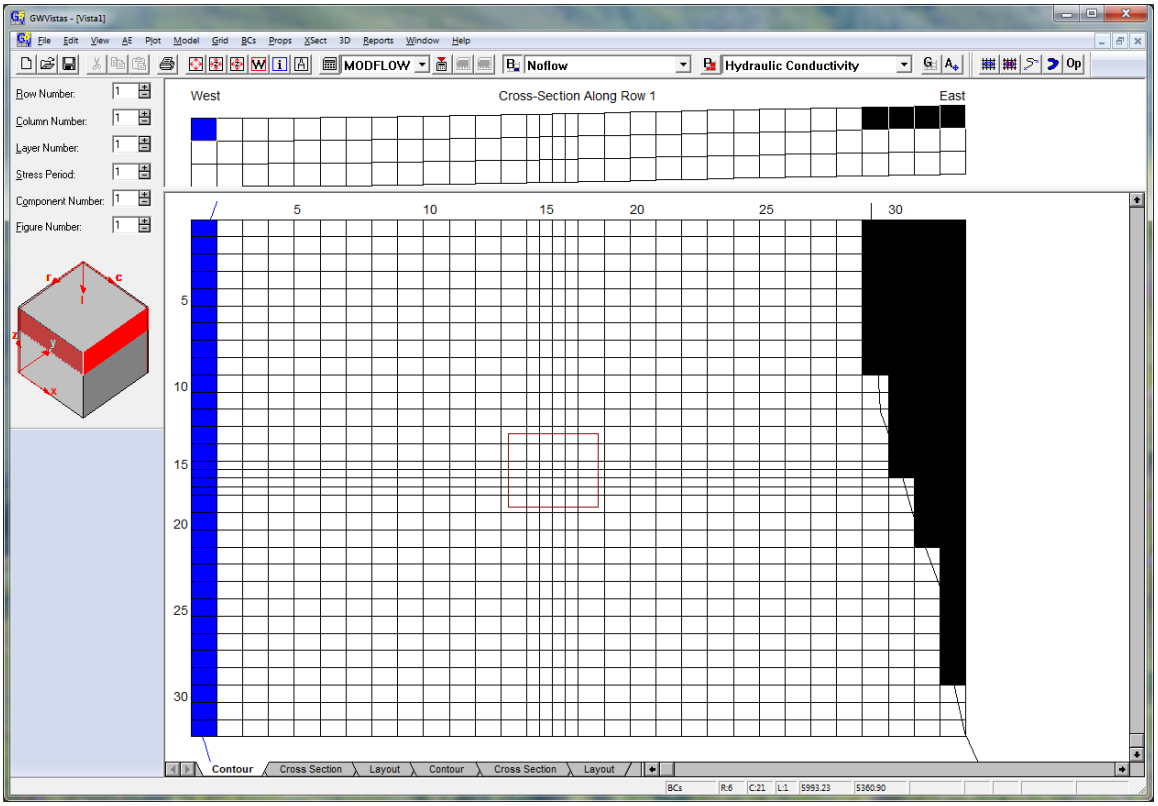

You could repeat this in layers 2 and 3 but there is an easier way. Go to layer 2 and select **BCs|Copy** and click OK (the defaults on the dialog are fine). This copies the boundary conditions of the current type (no-flow) from layer 1 to the current layer. Repeat this for layer 3 as well.

Creating MODFLOW Datasets

The example model is now complete. You will now create a MODFLOW data set, run the simulation, and analyze the results.

The preprocessing of model-specific options is accomplished using the **Model** selection on the main menu. Select **Model** and you will see options including MODFLOW, MODPATH, MT3D, etc. Select MODFLOW from the dropdown menu to reveal another menu with numerous selections. You will select several of these options before creating the data files.

First, select **Model|MODFLOW|Packages**. The dialog displays all of the MODFLOW packages and allows you to create only the ones necessary for the current model. The default selections do not normally need to be changed. The only change you need to make is the root file name. Change this parameter to *t2*. This means that the MODFLOW data files will be named T2.BAS for the BASIC package, T2.WEL for the well package, etc. The default MODFLOW version is **MODFLOW2000** which should be used for ESI's MODFLOW2000^{win32}. There are also selections for MODFLOW-SURFACT and MODFLOWT, as well as MODFLOW88|96. For now, leave this option at the default and select OK when you are finished.

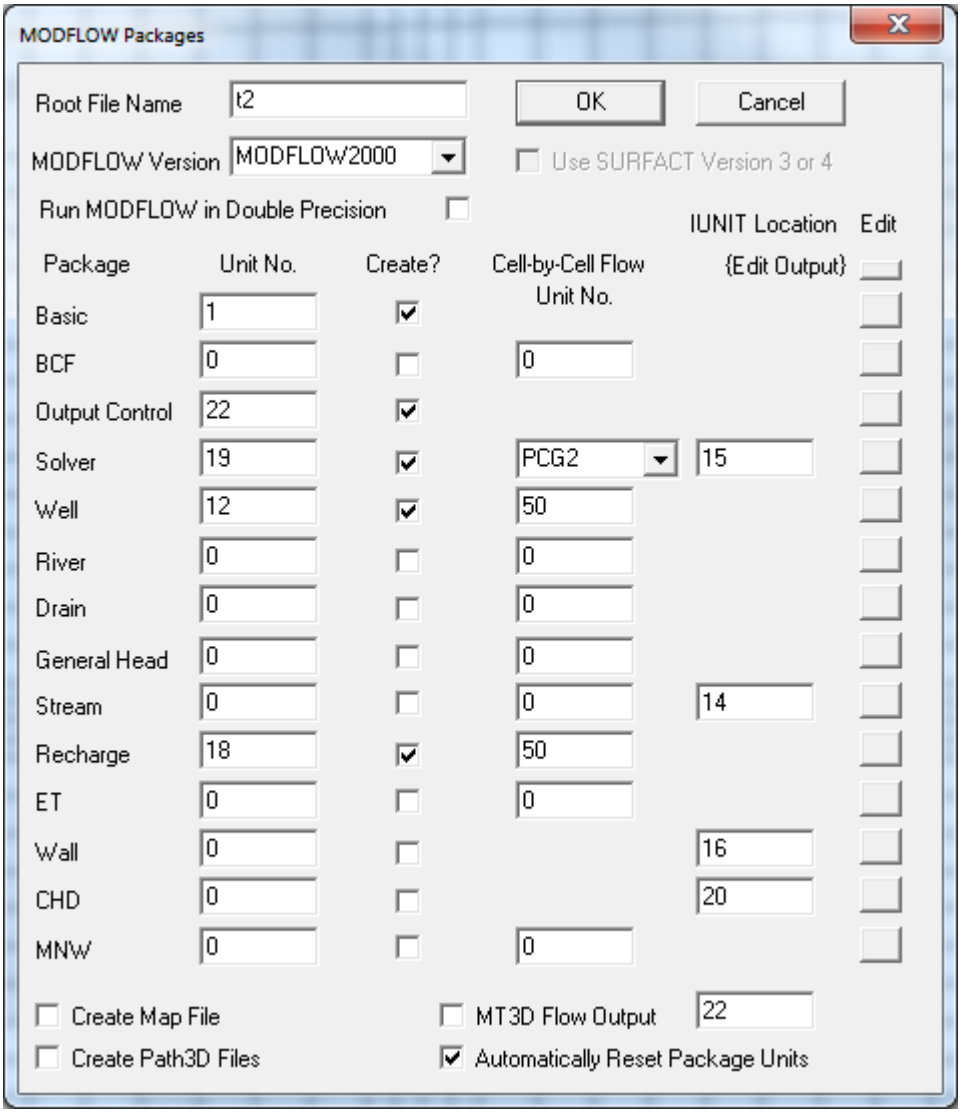

Choose **Model|MODFLOW2000|Options** and then click on the **BCF-LPF** tab to edit the layer types in the Layer Property Flow (LPF) Package. Change layers 2 and 3 from type 3 to type 0 (confined) layers. It is normally a good idea to make layers confined if you know that the water table will not drop into those layers. This speeds model convergence and can help avoid problems such as cells drying out. In this example, we know that the water table will always be in layer 1 so we make the lower layers confined.

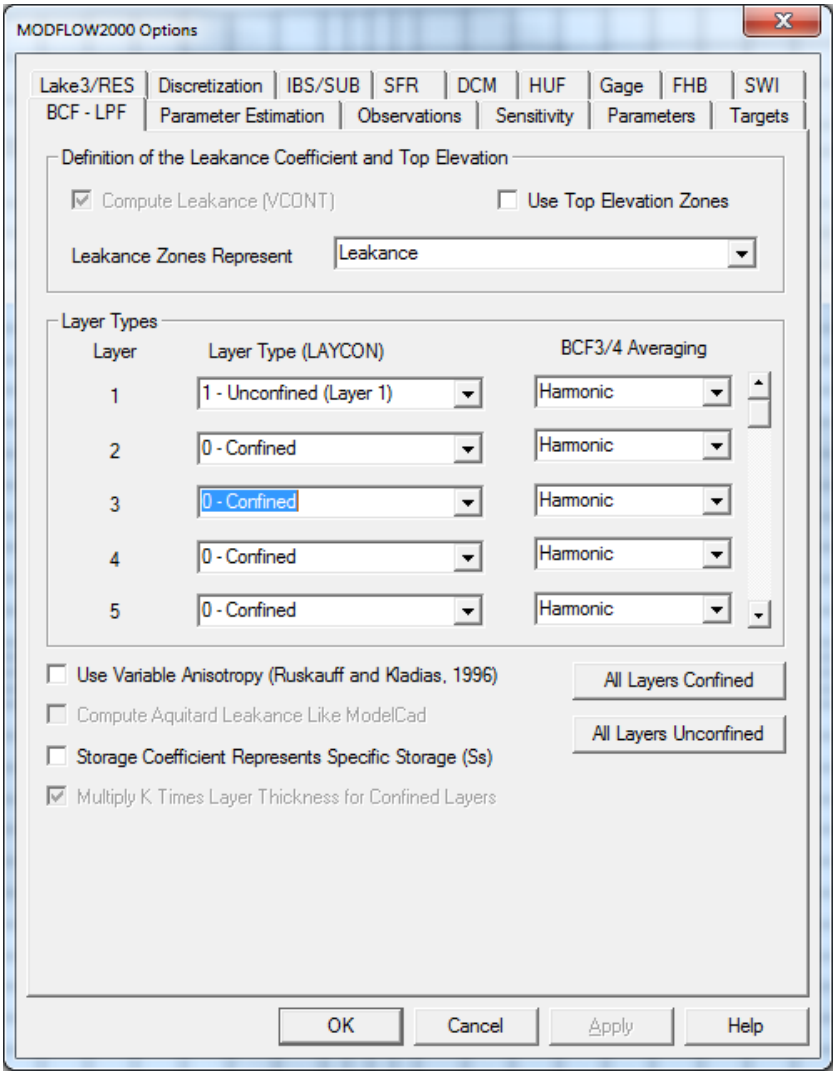

Now, select **Model|MODFLOW|Package Options** and click on the **Initial Heads** tab. Make sure the default starting heads for each layer are all set to 150. This is necessary because if the starting head is below the bottom of the layer (100 is below the bottom of most cells in layer 1, for example), then those cells will become dry. For steady-state models, you should make sure that your starting heads are always above the bottom of each cell in the model. An even better option is to set starting heads equal to the top of layer 1 from the dropdown list. Do that now and your screen should look like the one below.

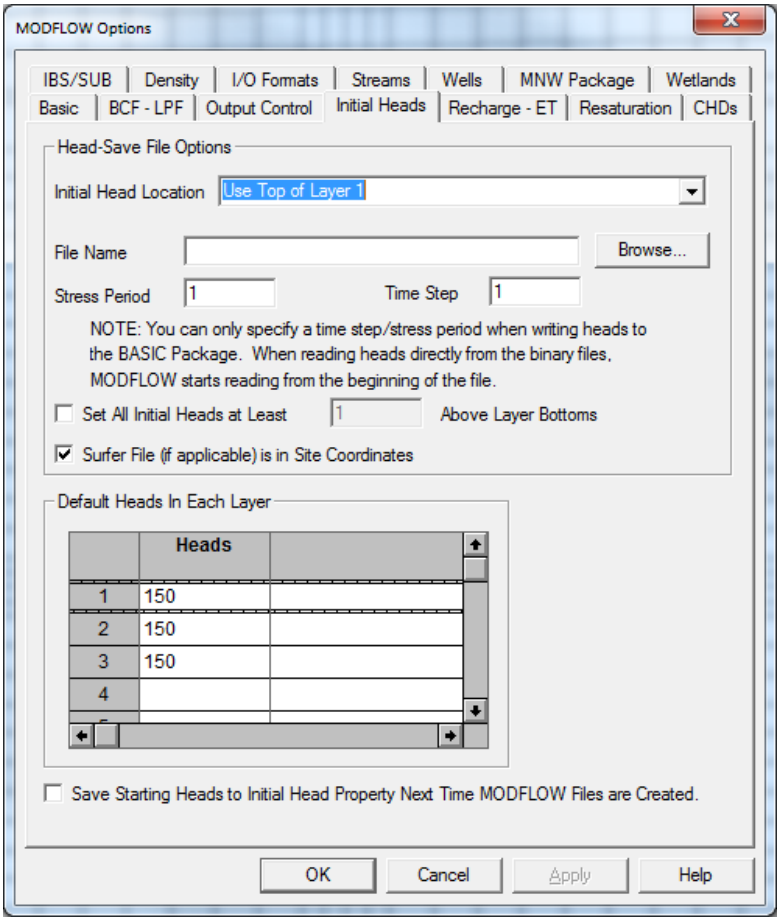

You may wonder why we have some options on the MODFLOW menu and some on the MODFLOW2000 menu when we are using MODFLOW2000. The logic behind the preprocessing menus on the Model menu is that any options that are common to all versions of MODFLOW are provided on the Models|MODFLOW menu. Initial heads are required by all versions of MODFLOW. The LPF Package is unique to MODFLOW2000 so we had to edit its options on the MODFLOW2000 menu.

Now, select **Model|MODFLOW|Solver Options**. Make sure the number of inner iterations is set to 25 for the PCG2 solver. It often speeds convergence to increase the number of inner iterations to between 25 and 50. At this point, all of the model options have been selected. You will now run the model and view the results.

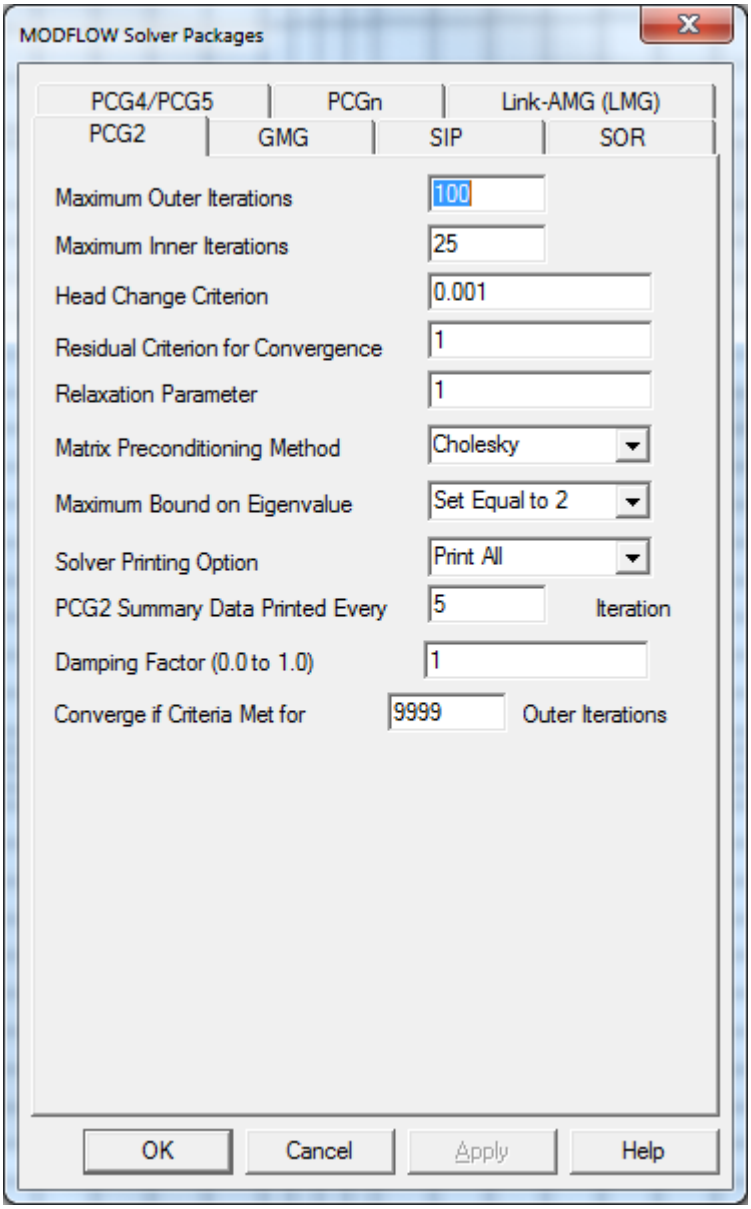

One last thing to change before we actually run the model is to reset the Working Directory. This directory is where GV will create the MODFLOW data files. It is also where MODFLOW will write the model results. By default, the working directory is the directory where GV was launched from (default: c:\gwv6). It is a good idea to have a special directory for your model runs that is separate from the GV program directory and from the GV file (*.gwv). In this example, we will use **c:\gwv6\tutorial\work**. If you installed GV in another location, then the directory will be **tutorial\work** under your GV directory. To change the working directory, select **Model|Paths to Models**. The working directory is the last item on the dialog. After changing it, click OK.

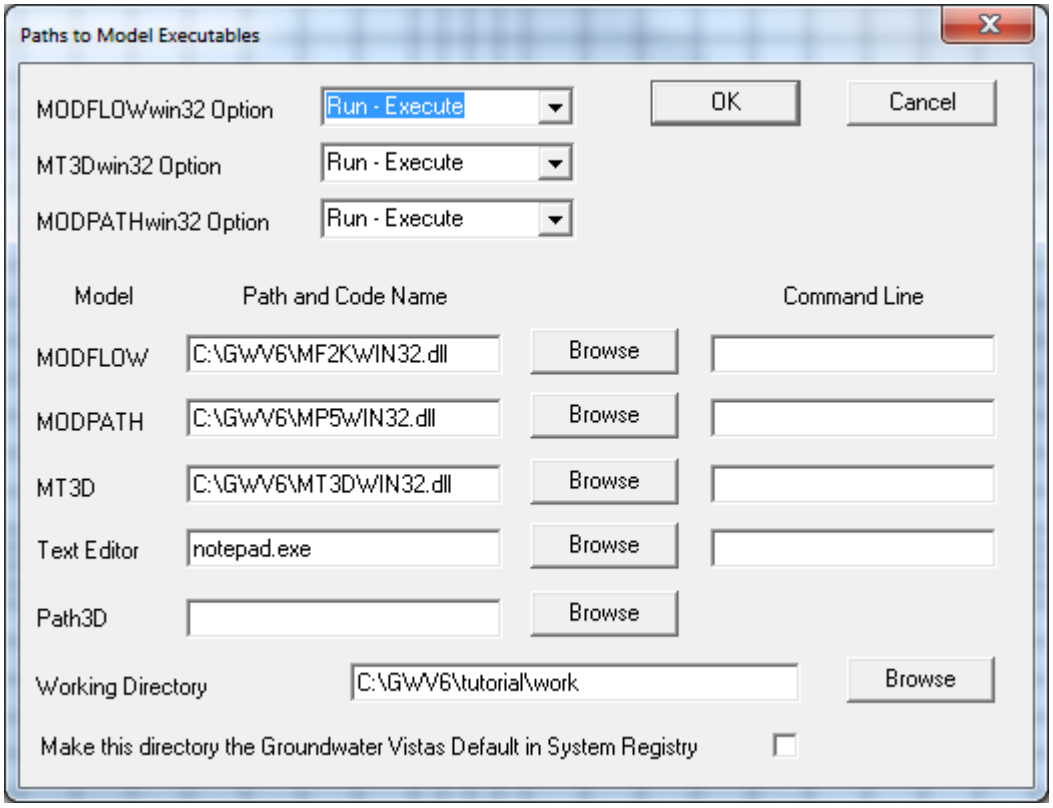

This is a good time to save your work. Choose **File|Save As** from the main menu. Use the file name T2.GWV for this first example.

You are now ready to create data files for a MODFLOW simulation. The simplest way to create data sets and run the model is to click the calculator button on the toolbar. A message first asks if you want to create datasets. Choose **YES** to do that now. A progress dialog will be displayed briefly as the files are created. You will then be asked if you want to view the error file. Choose **NO** for this example as there should be no errors to view. After clicking No, MODFLOW^{win32} will start to run. Clicking Cancel on this dialog will stop the run.

When MODFLOW is finished, a dialog will be displayed notifying you that the simulation is done and asks if you want to process the results. Choose YES to start the post-processing session. You will next see a dialog that allows you to specify what model results to analyze (the dialog is shown below). The default settings allow you to read the head file from the MODFLOW run. We will also look at the mass balance results so place a check mark next to *Cell-by-Cell Flow* as shown on the example dialog below. Click OK when you are done.

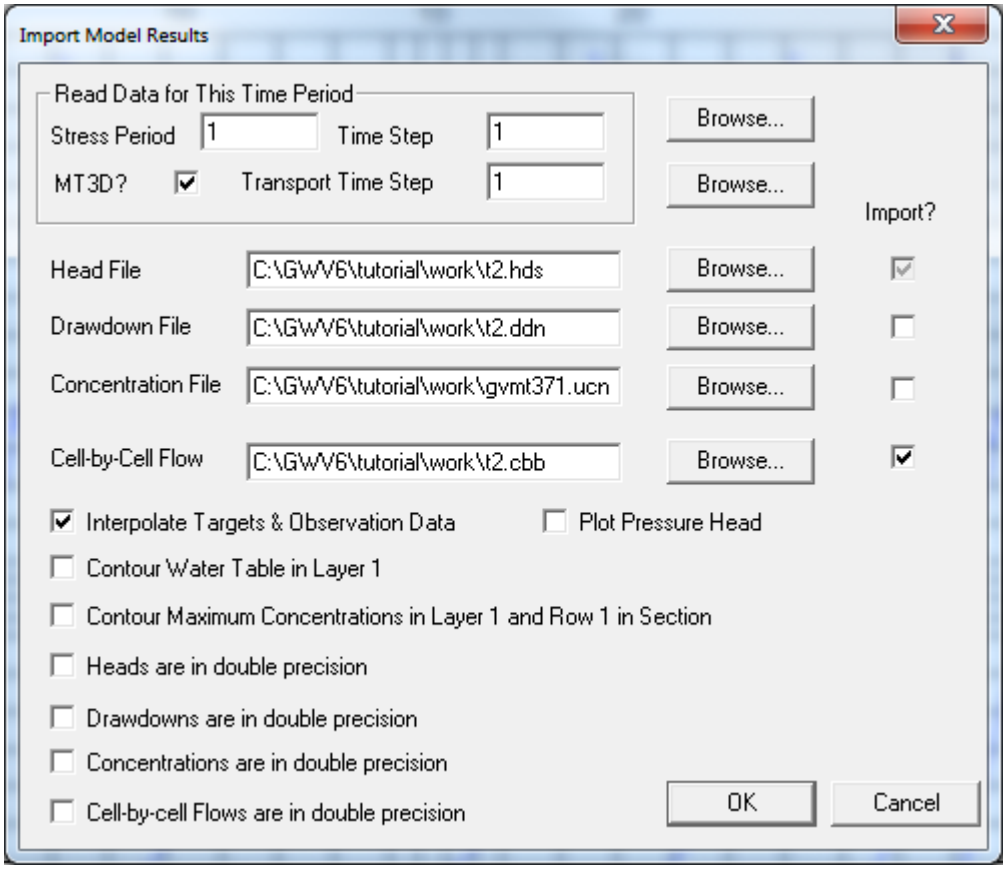

GV automatically contours the head results for the current layer and cross-section views. To change the contouring, select **Plot|Contour|Parameters(plan)** and change the contour interval to 0.5 ft and change the minimum and maximum values to whole numbers. The resulting contours for layer 3 and for the cross-section along row 1 are shown below. Your screen should look similar, unless you have changed to another layer.

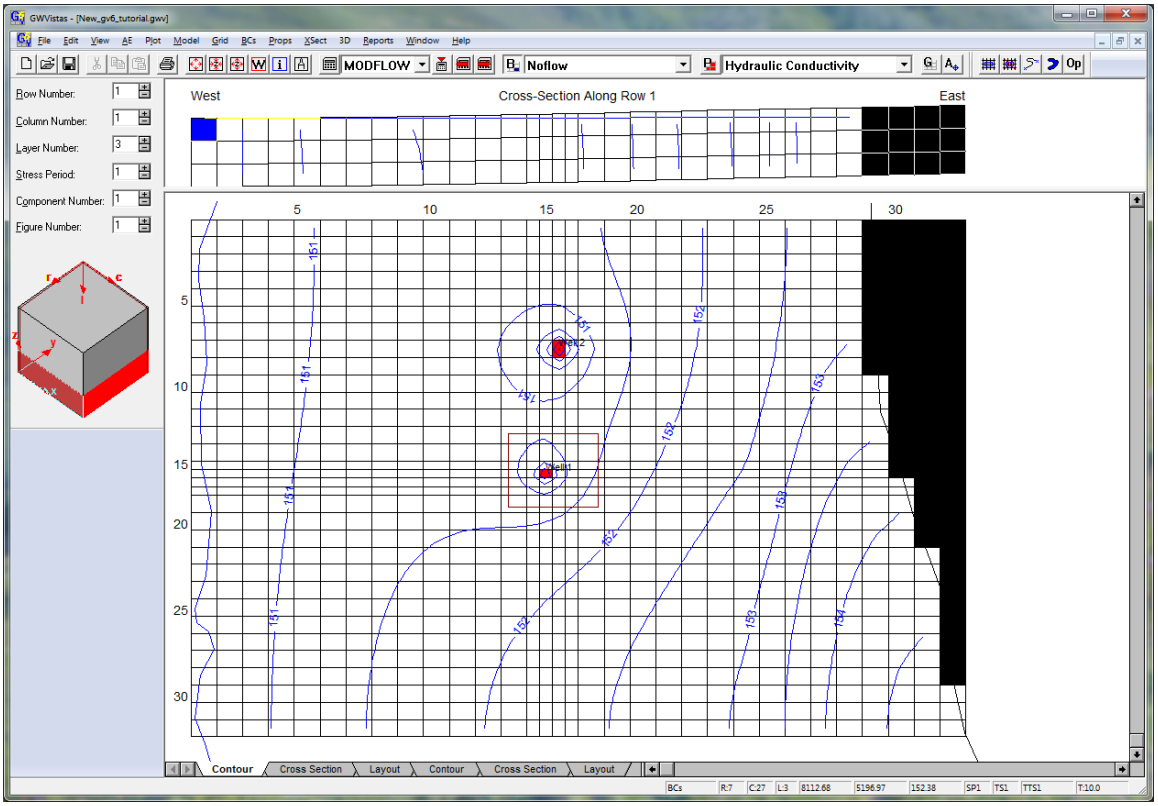

You may contour any layer or cross-section by simply changing the settings on the reference cube. For example, if you click the "-" button next to Layer on the cube, the layer above the current layer will be contoured and displayed. Similarly, if you select a new cross-section, it will also be recontoured. You may plot velocity vectors on the map and cross-section by selecting **Plot|What to Display…** from the main menu and placing a check mark next to Vectors on the dialog. You may also produce a color flood map by placing a check next to "Color Flood". You display all or none of these graphics using this dialog.

Other Types of Plots

GV provides you with the tools to create many types of maps and graphs that are useful in analyzing model results. These are accessed from the **Plot** menu. We will explore a few of these plots now.

Select **Plot|Mass Balance|Model Summary** from the main menu. A dialog appears summarizing the flow of water into and out of the entire model domain. The summary is similar to what MODFLOW writes in the output file. An example is shown below.

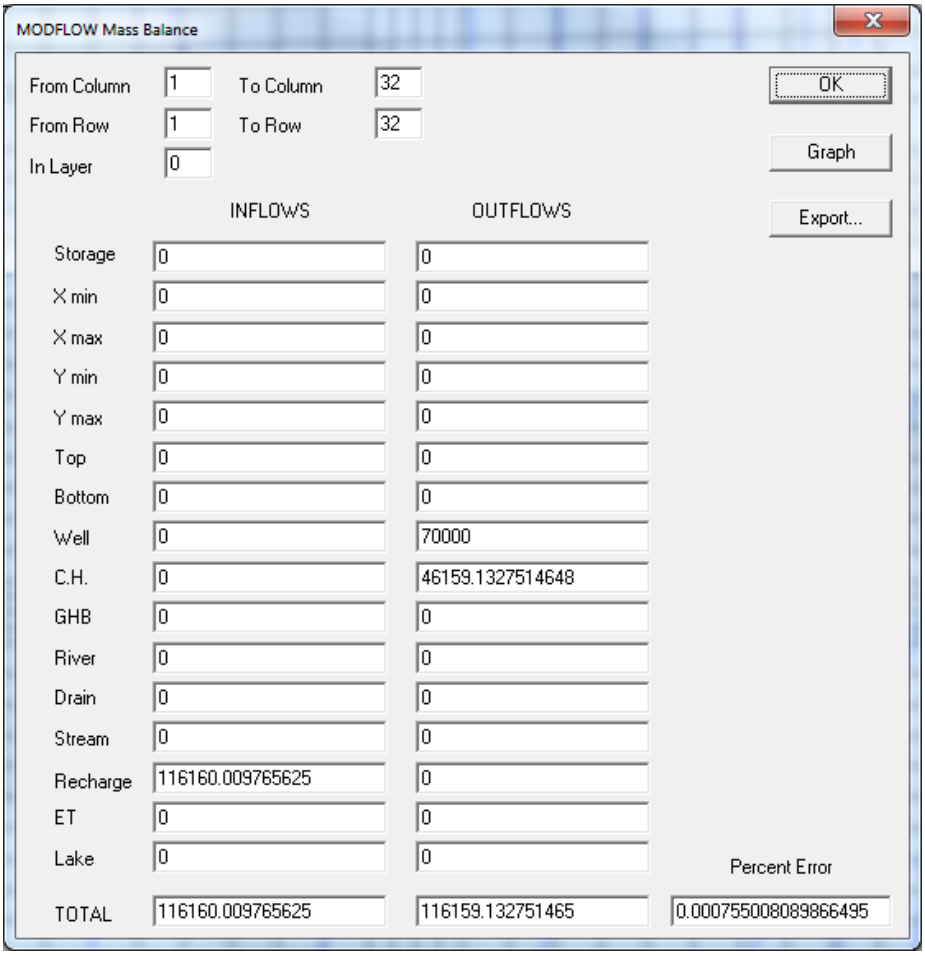

You may produce a bar chart of the mass balance results by simply clicking on the **Graph** button on the dialog. An example is shown below. Another useful mass balance feature of GV is that the flux rate of water into or out of a boundary cell is displayed on the status bar when the cursor moves over a boundary cell. Try this feature by moving the cursor over a constant-head cell in layer 1. The flow rate into the constant head for Row 1, Column 1 should be about -1329 $ft^3|d$ (NOTE: You must be in Boundary Condition mode to view fluxes; press the "B" button on the tool bar to enter BC mode). The negative sign is the MODFLOW convention for water being removed from the aquifer.

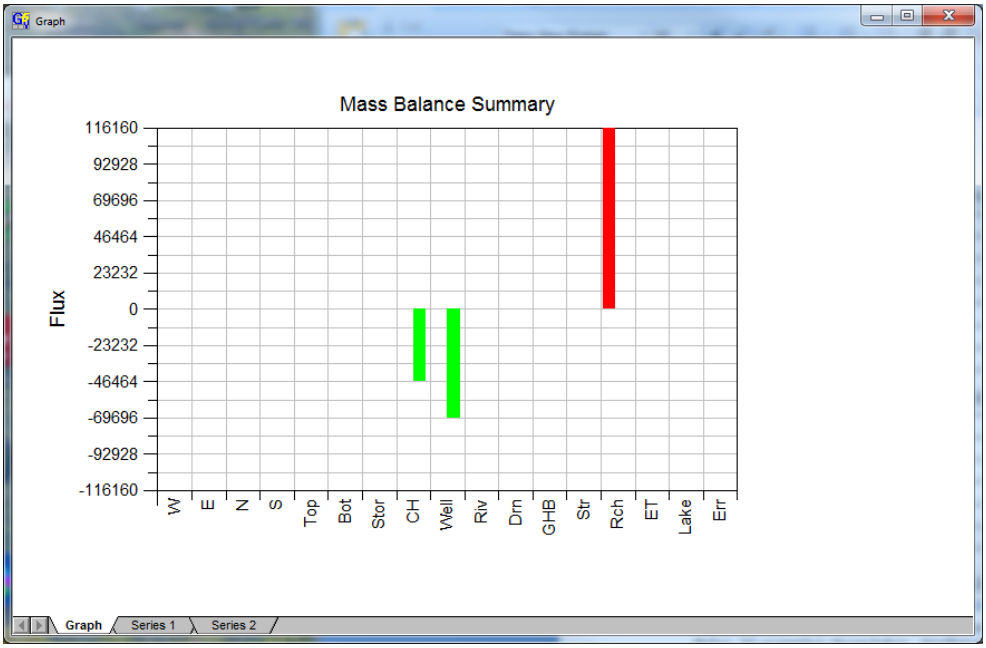

Another graph that shows head relationships between different layers of the model is the "profile" plot. A profile is simply a graph of head, concentration, drawdown, or water flux plotted versus distance along the current cross-section. Select **Plot**|**Profile**|**Head** from the main menu to display a profile of head in your example model. An example profile plot through Column 15 is shown below. To make this plot, click the "+" button next to **Column** on the reference cube until the number 15 appears. This changes the cross-section view to the row you specify and also determines which row or column will be plotted using the profile command. This plot shows the head relationships around a pumping well.

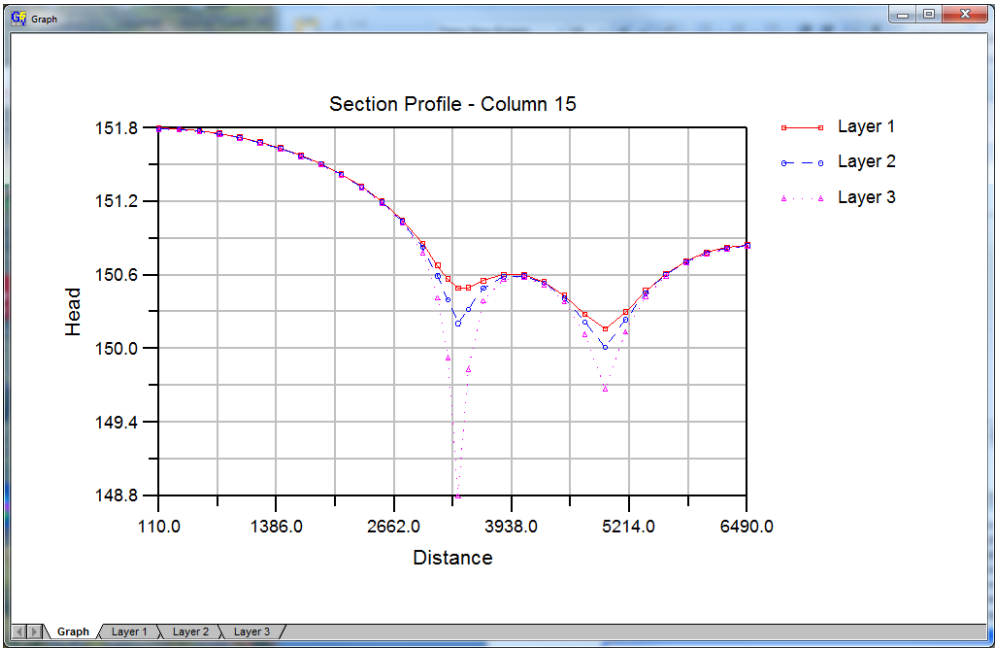

Particle-Tracking in GV

GV provides pre- and post-processing of particle tracking results for MODPATH (both the steady-state and transient versions) and for PATH3D. In the following example, we will use MODPATH[win32].

MODPATH[win32] is the same as the USGS MODPATH Version 5 but with a nicer interface. This version works with both MODFLOW[win32] and MODFLOW2000[win32].

Particle tracking is a form of transport modeling in which only the bulk movement of groundwater is investigated. Particle tracking neglects the effects of chemical reactions, dispersion, and diffusion. The results of a particle tracking simulation are displayed by plotting pathlines through the aquifer system. Travel times are often labeled on the pathlines.

You start the particle tracking analysis by adding particle starting locations to the model. Particle traces are then drawn either downgradient (forward tracking) or upgradient (reverse tracking) from these starting locations. First, move to layer 1 by clicking the "-" button next to layer on the cube until the number 1 appears in the Layer field. Select **AE|Particle|Line** from the GV main menu. Move the cursor to the northeast corner of the model and west of the no-flow boundary. Click the left mouse button and drag a line to the lower right corner, again keeping the line west of the no-flow area. Make sure the line of particles is within the bounds of the finite-difference grid. If the particles are outside the grid, they will not be tracked. Release the left mouse button to display a dialog. Change the number of particles along the line to 20 and click OK. Your particle line should look like the one below:

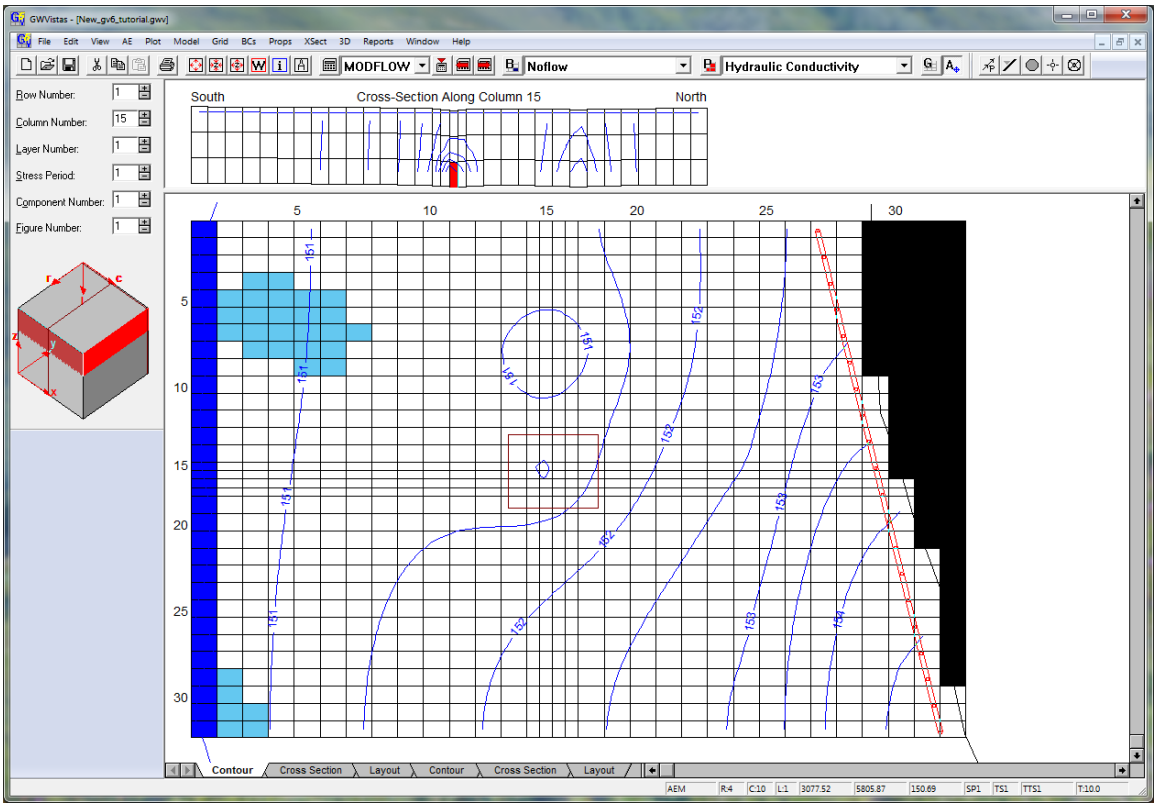

(You may notice the light blue cells depicted in layer 1 above. These are called "flooded" cells and represent cells where the water table is above the top of layer 1).

MODPATH has options just like MODFLOW only not as many. Select **Model|Modpath|Packages** and change the root file name for the MODPATH run to *t2mp*. It is not a good idea to use the same root name as MODFLOW because MODPATH uses some of the same file types with a slightly different format. Also, make sure that the MODPATH version is 5. Click OK when you are done.

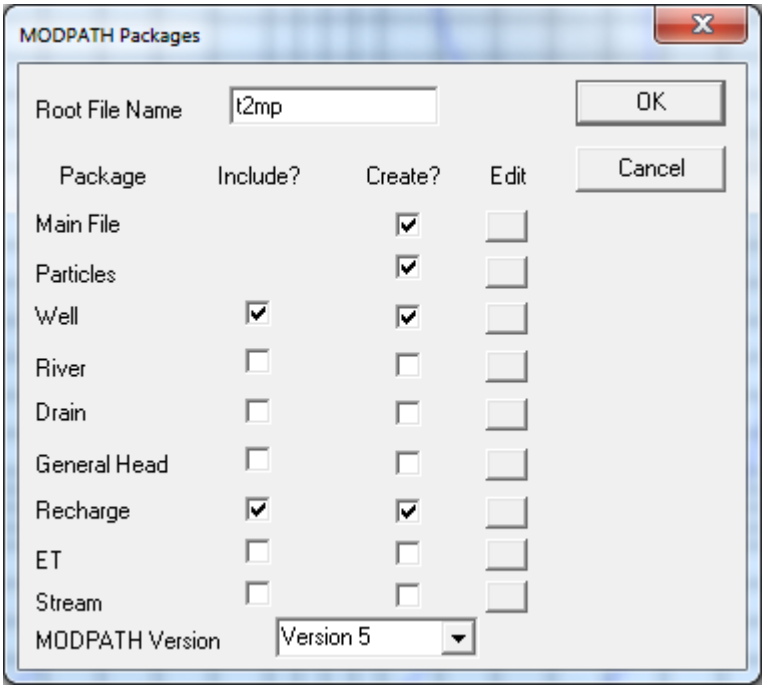

Now, select **Model|Use MODPATH** to make MODPATH the current model. A check mark should appear next to **Use MODPATH** near the bottom of the Model menu or use the dropdown list next to the calculator button on the toolbar. Click the calculator button on the toolbar and click the **Yes** button on the next dialog that asks to create MODPATH data sets. Select **CANCEL** from the next dialog that asks if you want to view the Error|Warning file. Now, MODPATH runs and produces results in a *pathline* file that GV will process. The MODPATH screen should flash on your screen. With this simple example, it should not take long. If you manage to see it, MODPATH should look like the following.

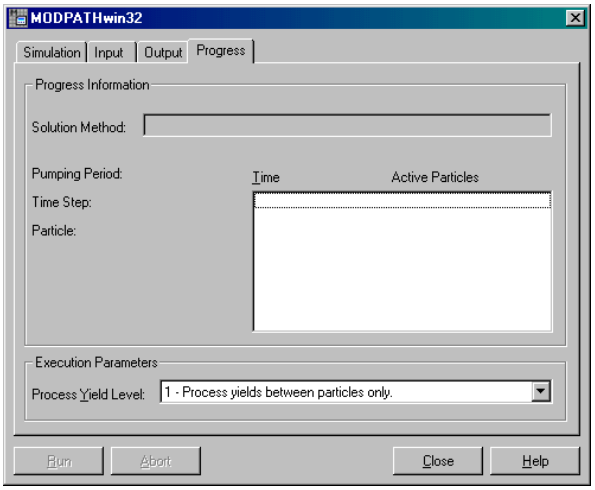

After finishing, a dialog asks whether you would like to display MODPATH results. Select **YES** and then find the file called *t2mp.ptl* in the subsequent dialog. Click OK when this file is selected. GV will now read the file and display the pathlines in both plan and cross-section views. Your screen should look similar to the one shown below.

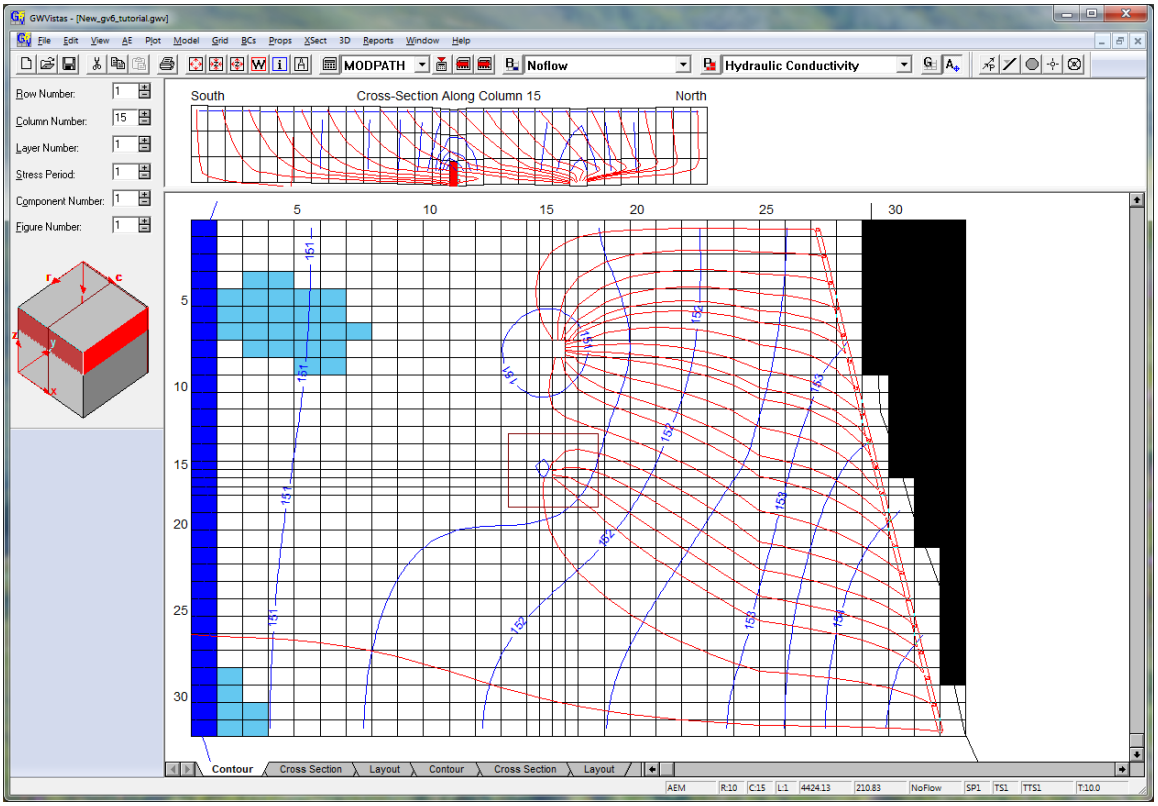

There are many options that can be used to modify the particle trace display. Usually you will want arrows on the traces. Most display options in Groundwater Vistas are found on the **Plot** menu. Select **Plot|Particles|Options** and click on the *Arrows* tab. Change the distance between arrows to 500 ft and the size of the arrows to 100 ft (these are both in model units of length). Also, check the box to display the arrows.

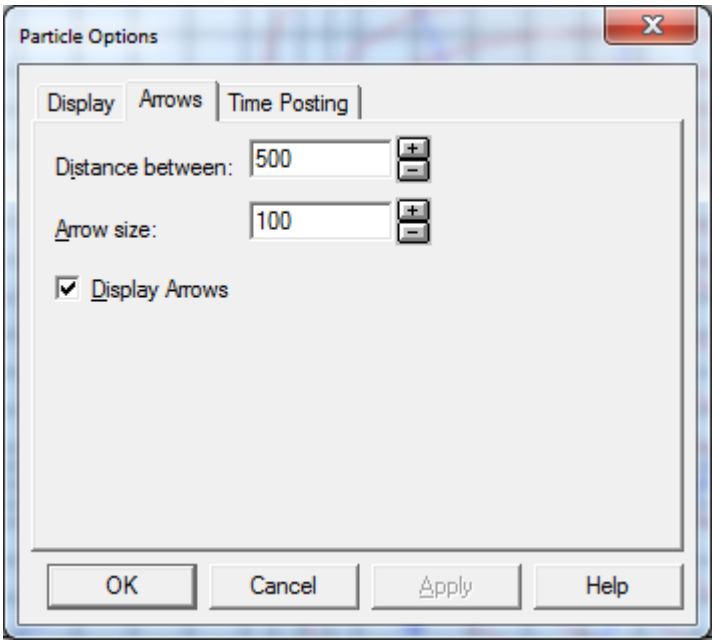

This method puts an arrow at a constant distance along the trace. Another method places the arrows at a constant time along the trace. Select **Plot|Particles|Options** and click on the **Time Posting** tab. Keep the default interval of 365 days but change the precision from 2 to 0 and change the multiplier from 1 to

0.00274. This will display an arrow every 365 days but the label will be in years (1|365 = 0.00274). Check the box to display time postings and click the font button. Change the font to 8 or 10 points.

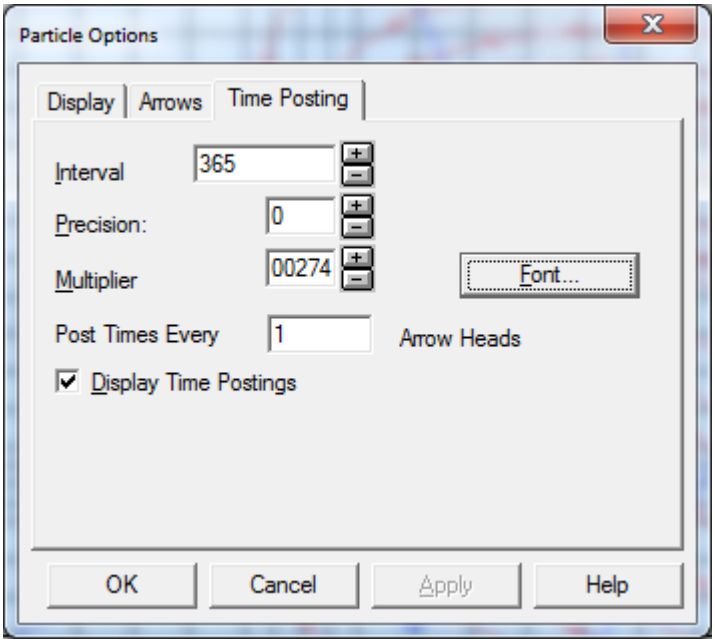

Now GV will redraw the traces with arrows spaced at equal time intervals instead of distance intervals. This is handy for looking at capture zones around wells for different time intervals.

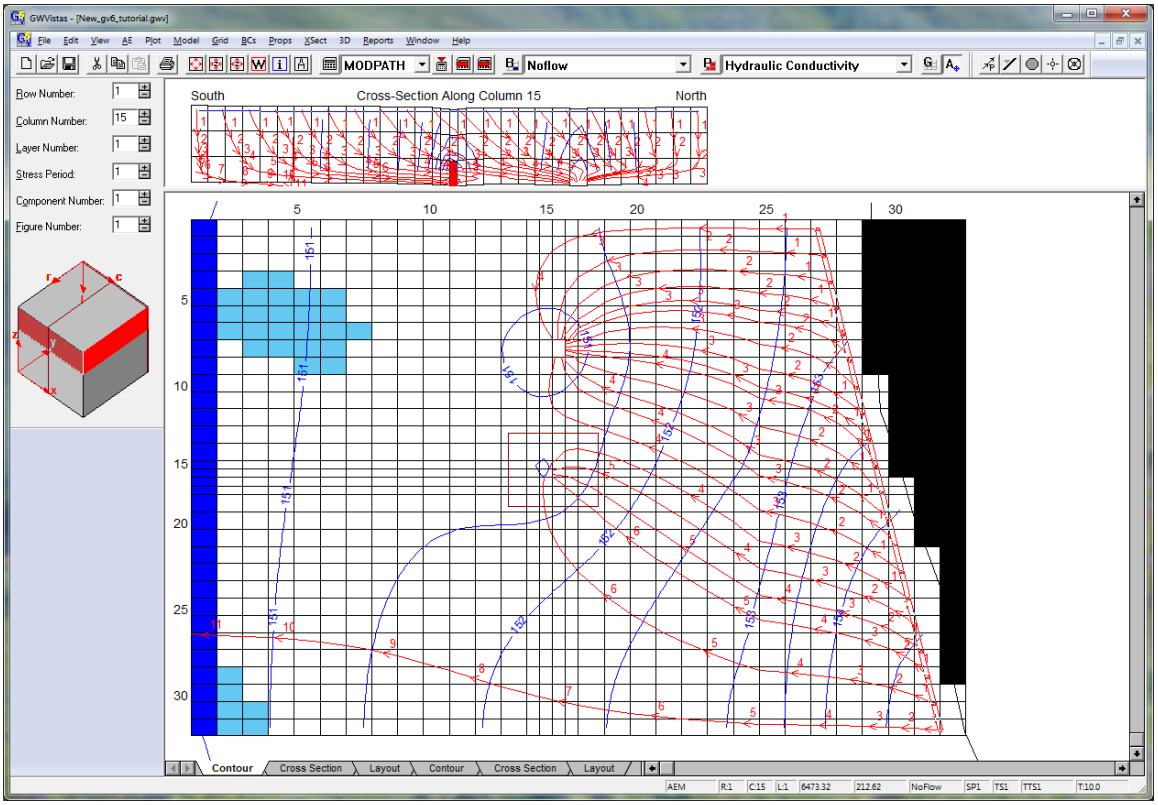

Model Calibration with GV

Calibration is one of the most complex parts of applying groundwater models. GV assists model calibration in three ways: (1) calculation of calibration statistics for head, drawdown, concentration, or flux, (2) automated parameter sensitivity analysis, and (3) automatic model calibration using a nonlinear least-squares technique built right into the GV interface or using external inverse models such as PEST, MODAC, and MODFLOW2000. All of these inverse models are included at no extra charge with Groundwater Vistas.

You start by adding calibration targets to the model. A calibration target is a point in the aquifer where a measurement of head, drawdown, concentration, or flux has been made. Calibration targets may be either steady-state or transient. When you run the model to compare against the target values, GV reads the model results and interpolates the model result in both space and time to compute an error or residual. Analysis of residual statistics is a powerful way of determining calibration quality and guiding further refinements to the model. The following exercise will illustrate the calculation of calibration statistics and automatic sensitivity analysis. We will then use the inverse model called Pest to calibrate the model.

We will start by defining 16 head target locations in our example model. Rather than type in the data manually, you will import a text file containing the target data. GV provides many data import features for calibration targets, boundary conditions, aquifer properties, and base maps. Select **AE|Import|Target Text File** from the main menu. Find the file called *targets.dat*, which should be in the tutorial directory (default is c:\gwv6\tutorial). Click OK when you have found it.

A dialog now prompts for the format of the file. The data in this file can be in any order as long as the data are delimited (separated by commas, spaces, or tabs) and each target contains the same number of data values. The only change needed is to skip 1 line. The rest of the information is correct. The numbers in the bottom part of the dialog denote columns (fields) in the text file containing the data described to the left side of the field. For example, the 1 next to “Name” means that the target name will be contained in column 1 of the file. Similarly the X coordinate of the target can be found in column 2. Click OK after changing the number of lines to skip at the top of the file.

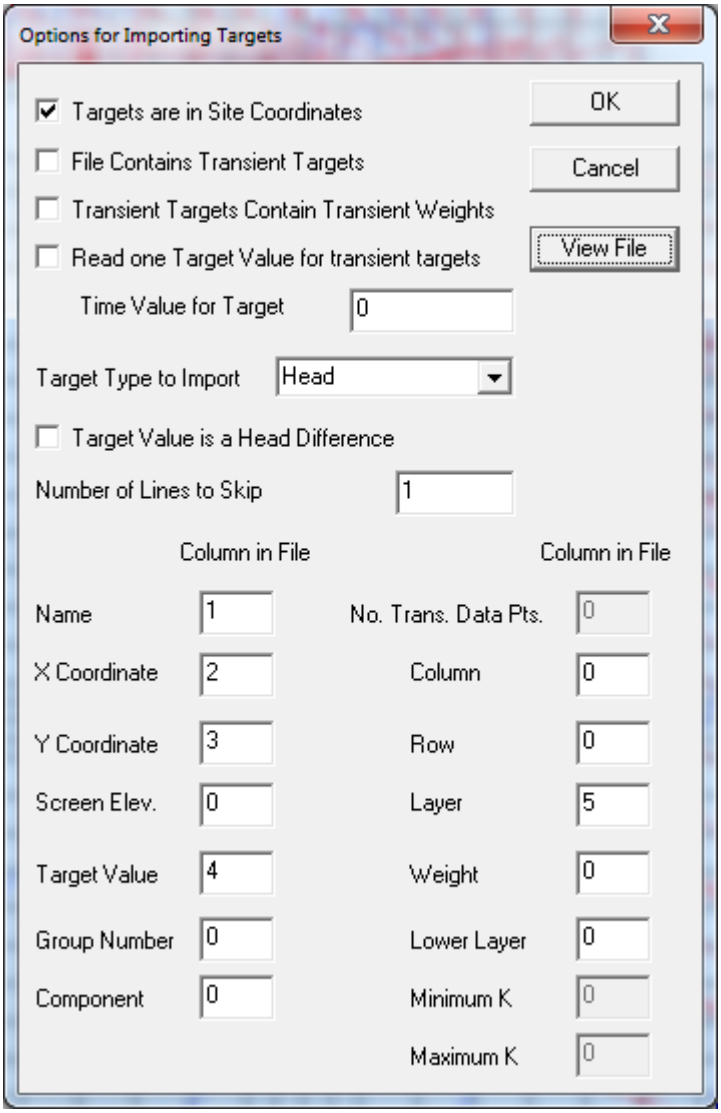

GV will report the number of targets successfully imported. In this example, there should be 16. The targets will appear on the plan view as small blue dots. Targets are only displayed for the layer in which they are defined. You should see 6 in layer 1, 5 in layer 2, and 5 in layer 3. You may edit target information by double-clicking on a target symbol (You must be in Analytic Element mode, however; click the  button on the tool bar to edit targets in this manner).

You will need to import model results again in order to compute the calibration statistics. Select **Plot|Import Results** and click OK on the dialog. All of the options should be set properly from the last run you made. You may view the calibration statistics by selecting **Plot|Calibration|Statistics|Plots…** A dialog is displayed that allows you to select the type of targets to use in the calculation (head, concentration, drawdown, or flux). You may also plot only selected ranges of layers. To view the statistics for this model, click the Statistics button. The residual sum of squares should be about 42.9 ft^2. Click Close to leave this dialog. Now click the *Plot Observed vs. Simulated* button. A graph of observed vs. computed heads is displayed. Targets are color coded by layer. Your plot should be similar to the one shown below.

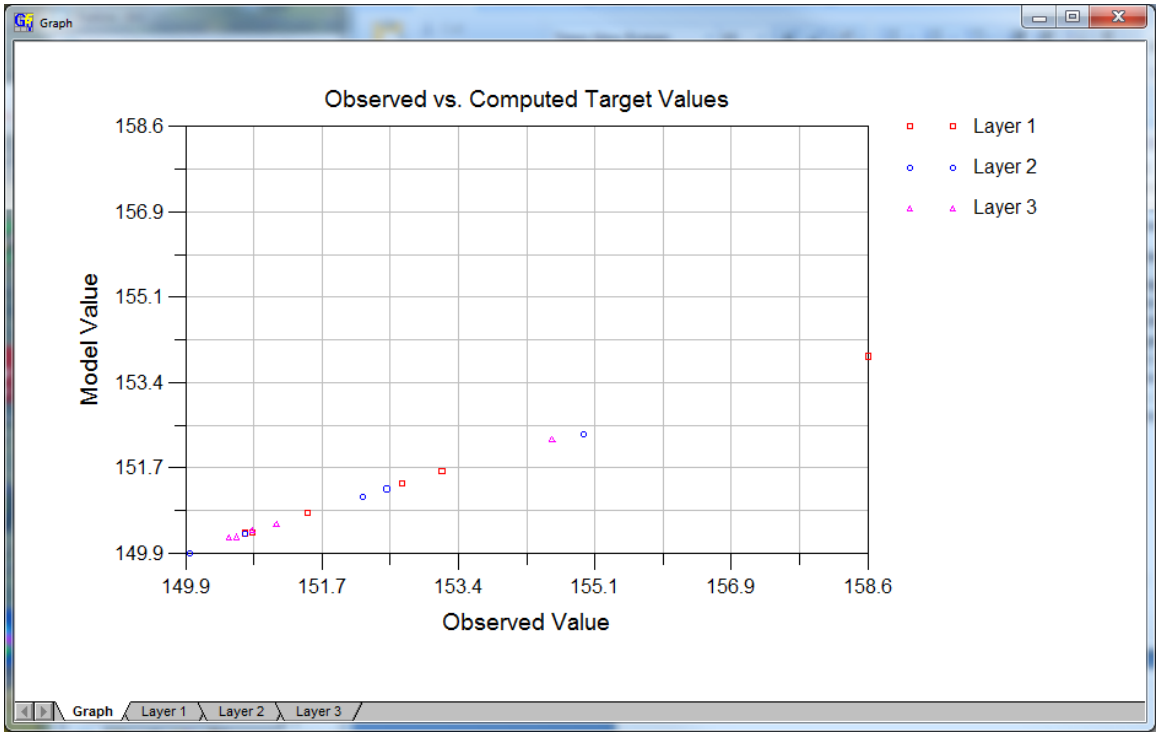

Ideally the plot shown above should be a straight line oriented at a 45-degree angle. This means that the observed value should equal the simulated value. In this example, the plot is a straight line but there is a strong bias such that higher heads are simulated too low.

Another way to view the target residuals (errors) is to post them on the contour map. You do this by selecting **Plot|Calibration|Post Residuals**. Target residuals are posted when a check mark is displayed next to *Post Residuals* on this menu. In this example, the residuals are too small to read. You may change the font size by selecting **Plot|Calibration|Options**. Click the font button to change the size or font. Select **View|Refresh** to redraw the window with the new font. The grid also makes viewing difficult sometimes. To turn off the grid display, select **Plot|What to Display** and uncheck the option for finite-difference grid. Your screen should now look something like the one shown below.

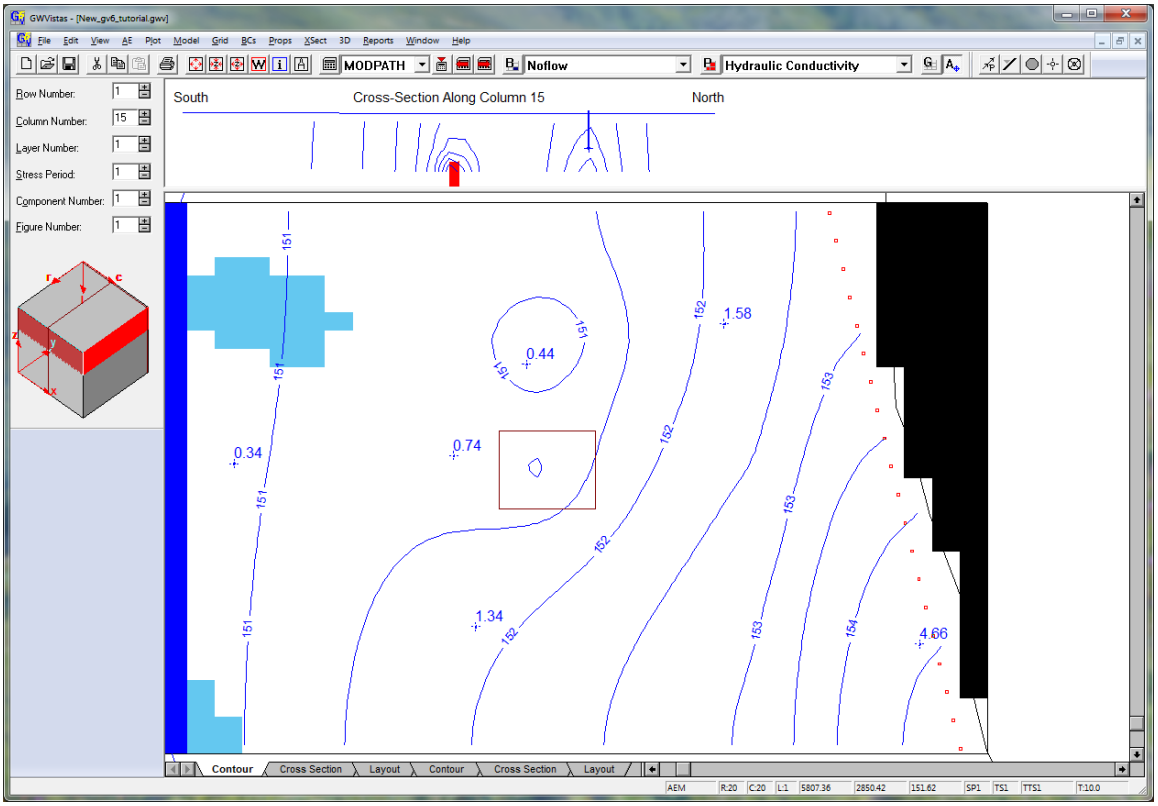

Sensitivity analysis is an integral part of model calibration. Sensitivity analysis is the process whereby model parameters or boundary conditions are altered slightly and the effect on model calibration statistics is observed. By producing a series of simulations with different values for a single model parameter, you get a feel for how a parameter may be modified in order to achieve a better calibration. This is a tedious process because many simulations are required for each parameter and there are often many parameters to analyze. GV provides you with an automated way of performing a sensitivity analysis that greatly improves the efficiency of the process. You simply choose a parameter type, the number of simulations, and a parameter multiplier for each simulation. GV then runs MODFLOW the desired number of times and produces a sensitivity plot. For each simulation, GV multiplies your initial parameter value by the multiplier you specify. After all of the simulations are finished, GV plots calibration statistics versus parameter multiplier to visually show the results of the analysis.

You start a sensitivity analysis by selecting **Model|Auto Sensitivity|Options**. Select Kx as the parameter to vary. Keep the default of Zone number 1 and change the number of simulations from 5 to 11. Click on the "Multipliers" button. Enter the values 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, and 1.5 for simulations 1 through 11, respectively.

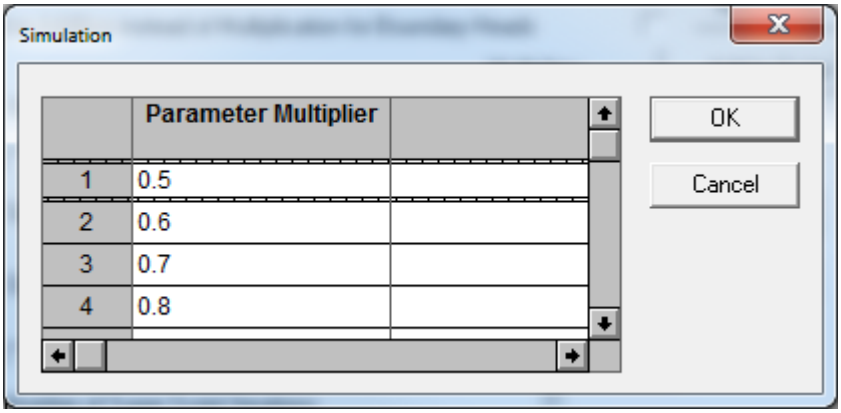

Normally you want to evaluate multiple parameters. In this case, we want to look at Kx and Kz in zones 1 and 2. The best way to do this is to run the automatic sensitivity analysis from a script file. GV can generate this script file for you. First, click the **New Script** button. GV will prompt for a file name. Make sure to place the file in the working directory (default: c:\gwv6\tutorial\work). Then GV asks for a range of zones. Enter 1 for the minimum zone and 2 for the maximum zone:

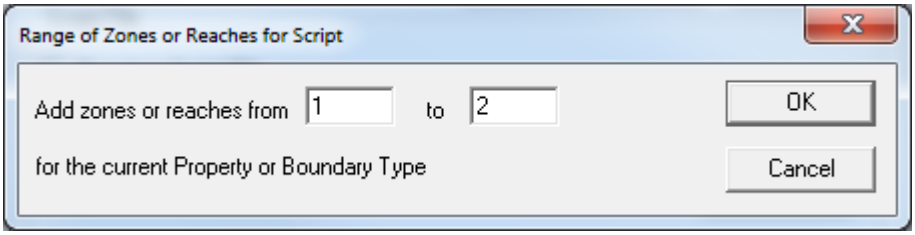

Repeat this for Kz by changing the parameter from Kx to Kz and click on the **Add to Script** button. Again, use zones 1 to 2 on the next dialog. To see the script file, select **Edit Script**. This will load the current script file into a text editor, as shown below.

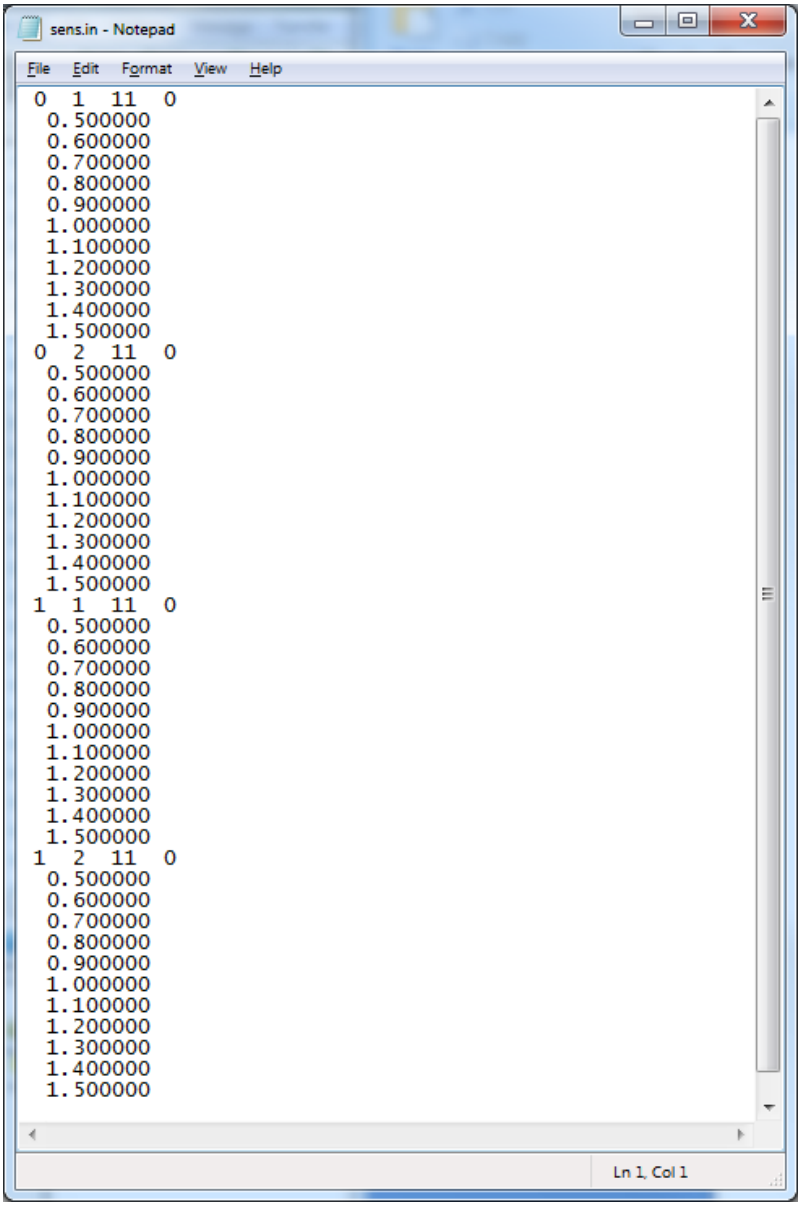

Each parameter is defined with one header line followed by the multipliers to use. The header contains the parameter type (first number on line), zone number, number of simulations, and the fourth number is not important. The GV manual describes what the parameter type numbers mean.

Now click the **Run from Script File** checkbox to instruct GV to use that file during the sensitivity analysis., as shown below.

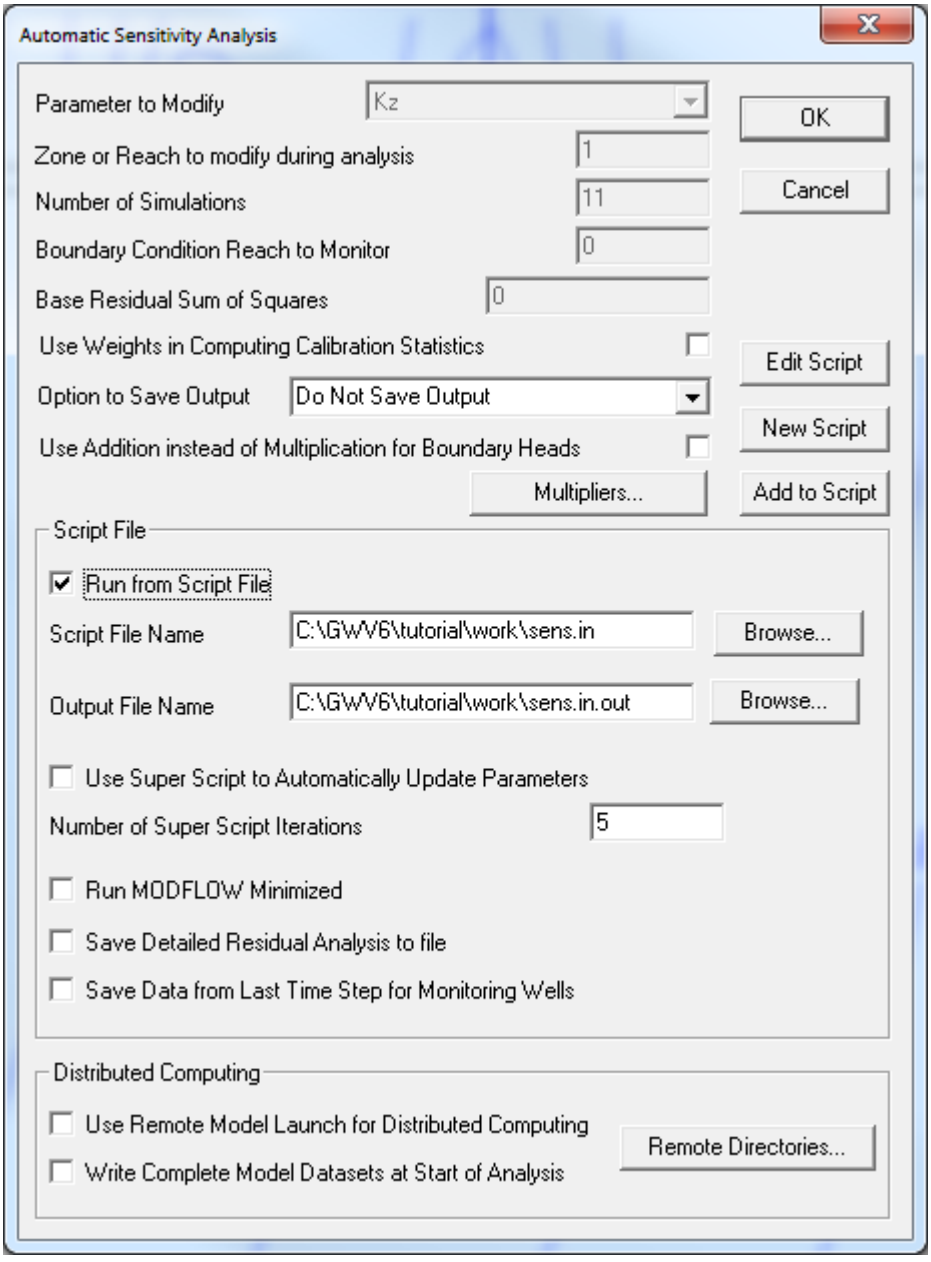

Click OK when you are done. Running the analysis is as simple as selecting **Models|Auto Sensitivity|Run Analysis** from the main menu. You will see MODFLOW flash on your screen numerous times and after the last simulation a dialog will ask whether you would like to see the results. Select "Yes" and you are asked to find the output file. The default is correct, so just select it. You are then asked to choose a variable to plot for each run. The choices include sum of squared residuals, residual mean, residual standard deviation, average drawdown, and total flux to a designated boundary condition reach. Choose the default, sum of squared residuals, and a plot of residual sum of squares vs. multiplier is displayed.. Your graph should look similar to the one shown below.

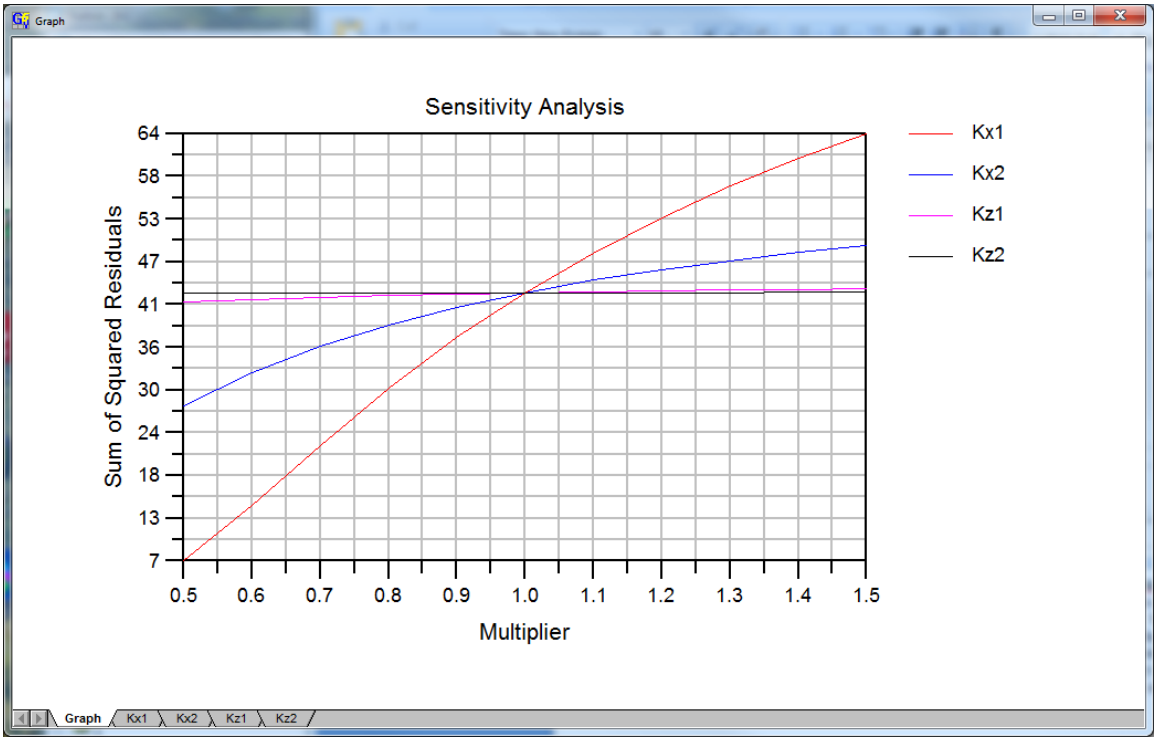

The way you interpret this graph is to look for the curve that has the lowest sum of squared residuals, because a lower value of that statistic means a better calibration. In this case, the red curve (for Kx zone 1) has a sum of squared residuals of 7 when the multiplier on Kx Zone 1 is 0.5. This means that if you multiply Kx (and Ky) by 0.5 and run the model, the sum of squared residuals will drop from its current value of around 44 to 7.

If you were trying to calibrate this model, you would go into the Hydraulic Conductivity database and multiply the Kx and Ky values for zone 1 by 0.5. Then you would repeat the sensitivity analysis and continue until the calibration was good. The trick here is to only change one parameter after each sensitivity run. You can learn more about this technique by following the tutorial in the Calibration chapter.

Automated Calibration with Pest

Groundwater Vistas offers several ways of automating the calibration process. GV has a very simple and easy-to-use automated calibration procedure (called inverse modeling) built into the GV interface (this is under Model|GV Calibration). GV also supports PEST, MODAC, and MODFLOW2000. In this tutorial, we will show you how to use Pest, which is the most sophisticated inverse model available for groundwater modeling. Many users are intimidated by Pest because it has a large number of options. However, we have set up Pest in GV with very good default values for most of these options. Thus, with just a few minor changes you should be able to use Pest on a variety of problems. This section shows how to adjust hydraulic conductivity in the simple example described above.

There are two basic things you need in any inverse or auto-calibration setup. The first is a set of targets (Pest calls these observations) that you compare to the model-simulated results. The second is to define a set of parameters to estimate. In the sensitivity analysis described in the previous section, you evaluated the sensitivity of horizontal and vertical hydraulic conductivity (Kx and Kz) in two different zones for a total of 4 parameters. We will now do the same thing with Pest.

Since you already defined the 16 head targets in the last session, we do not need to repeat that. We will start, then, with the parameters. Select Model|Pest|Parameters. This brings up a spreadsheet like the one below.

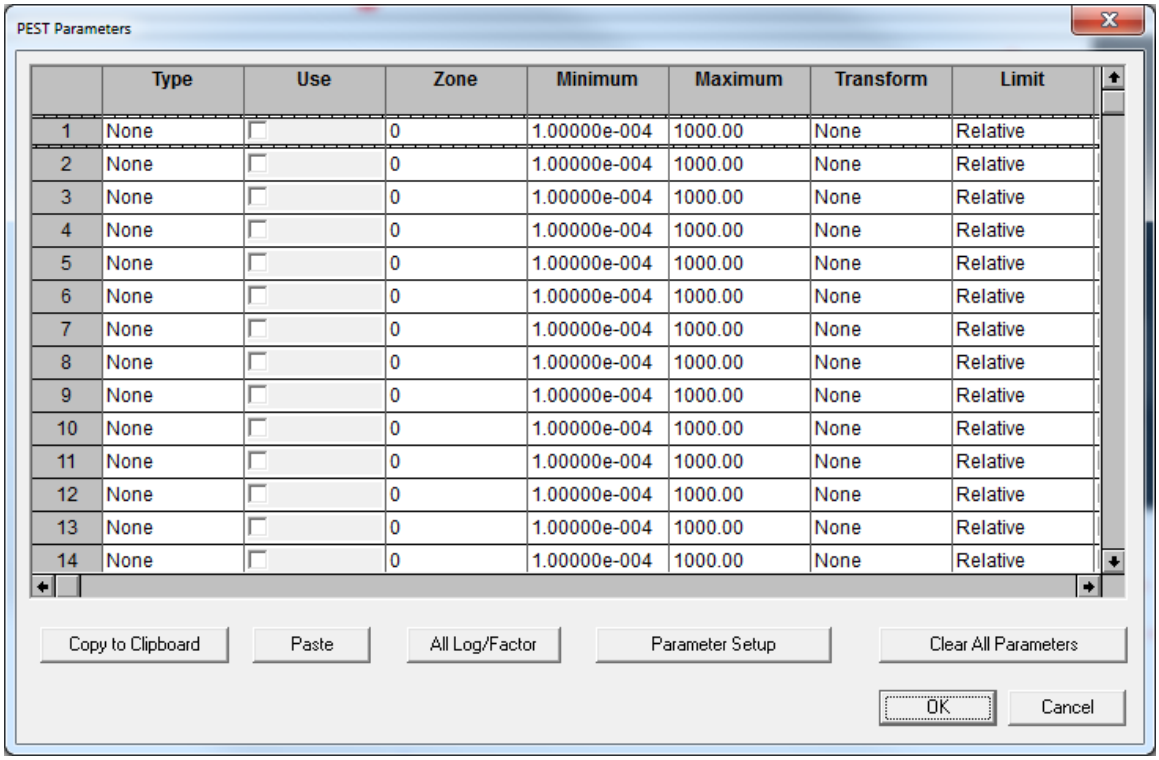

Initially there are no parameters defined. So you must add Kx and Kz in zones 1 and 2 just like for the sensitivity analysis. The easiest way to do this is to click the "Parameter Setup" button. You will do this twice, once for Kx and once for Kz. For some parameters (Kx, Kz, Storage, Specific Yield, and Recharge), you may use the "Add All Zones" checkbox to add them quickly. Also check the box called "Minimum|Maximum are multipliers…". Then enter 0.1 for minimum and 10 for maximum. This makes the minimum and maximum acceptable K values 10 times less or greater than the current values. Click OK when you are done.

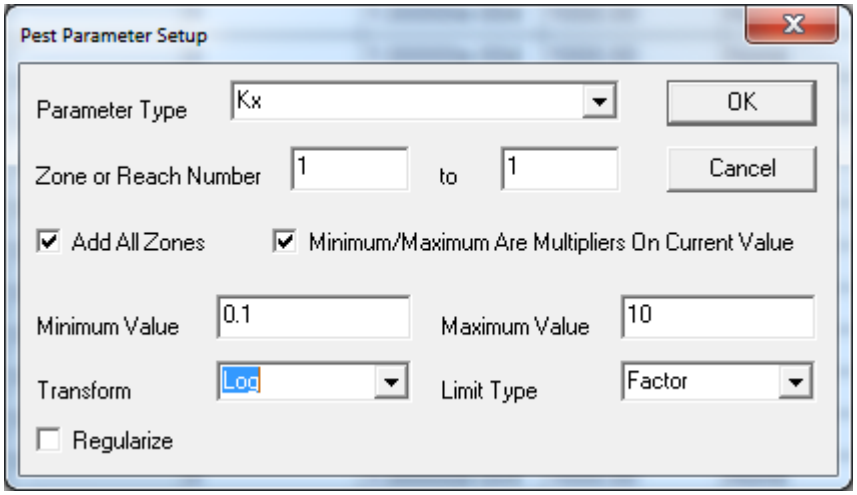

Now repeat for Kz, as shown below.

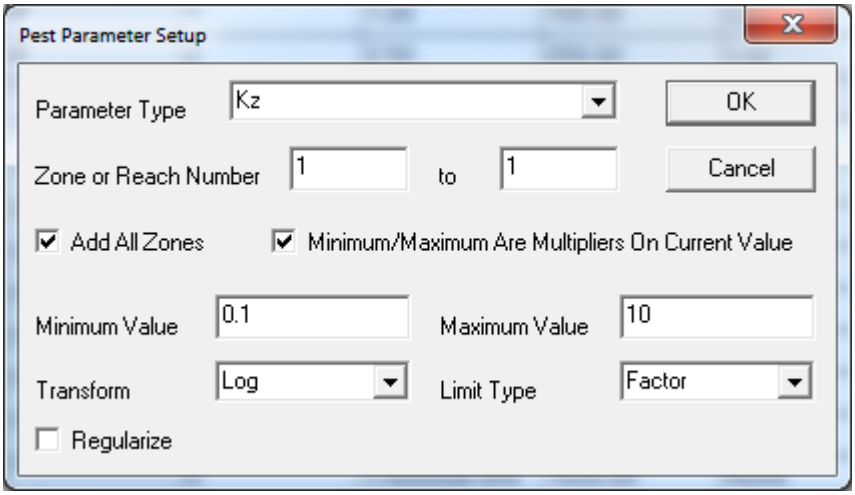

Your parameter spreadsheet should now have 4 entries like the example below:

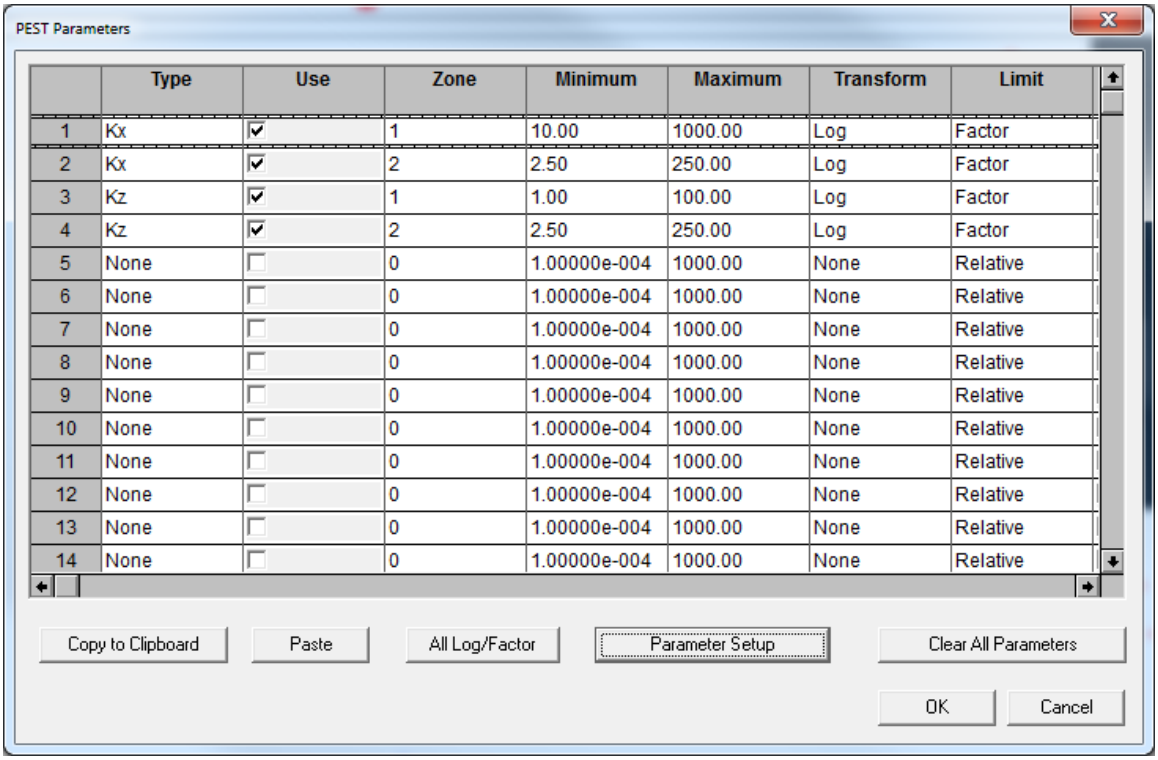

Note that each parameter is set to a log transform. This is generally a good idea when estimating hydraulic parameters in Pest. The minimum and maximum values for each parameter are also used by Pest to keep values in a valid range. If a parameter hits a bound after estimation, it usually means that the parameter is insensitive or the bounds you set are too restrictive. In the latter case, it could be there are other conceptual problems with the model. Anyway, click OK when your spreadsheet is filled out properly.

Now select Model|Pest|Options. This tab dialog is normally set up correctly for most problems. We do recommend, though, that you turn on the option to run models without screen output and use command line versions of models. The latter will keep the GV Windows versions of the models from popping up on the screen while Pest runs.

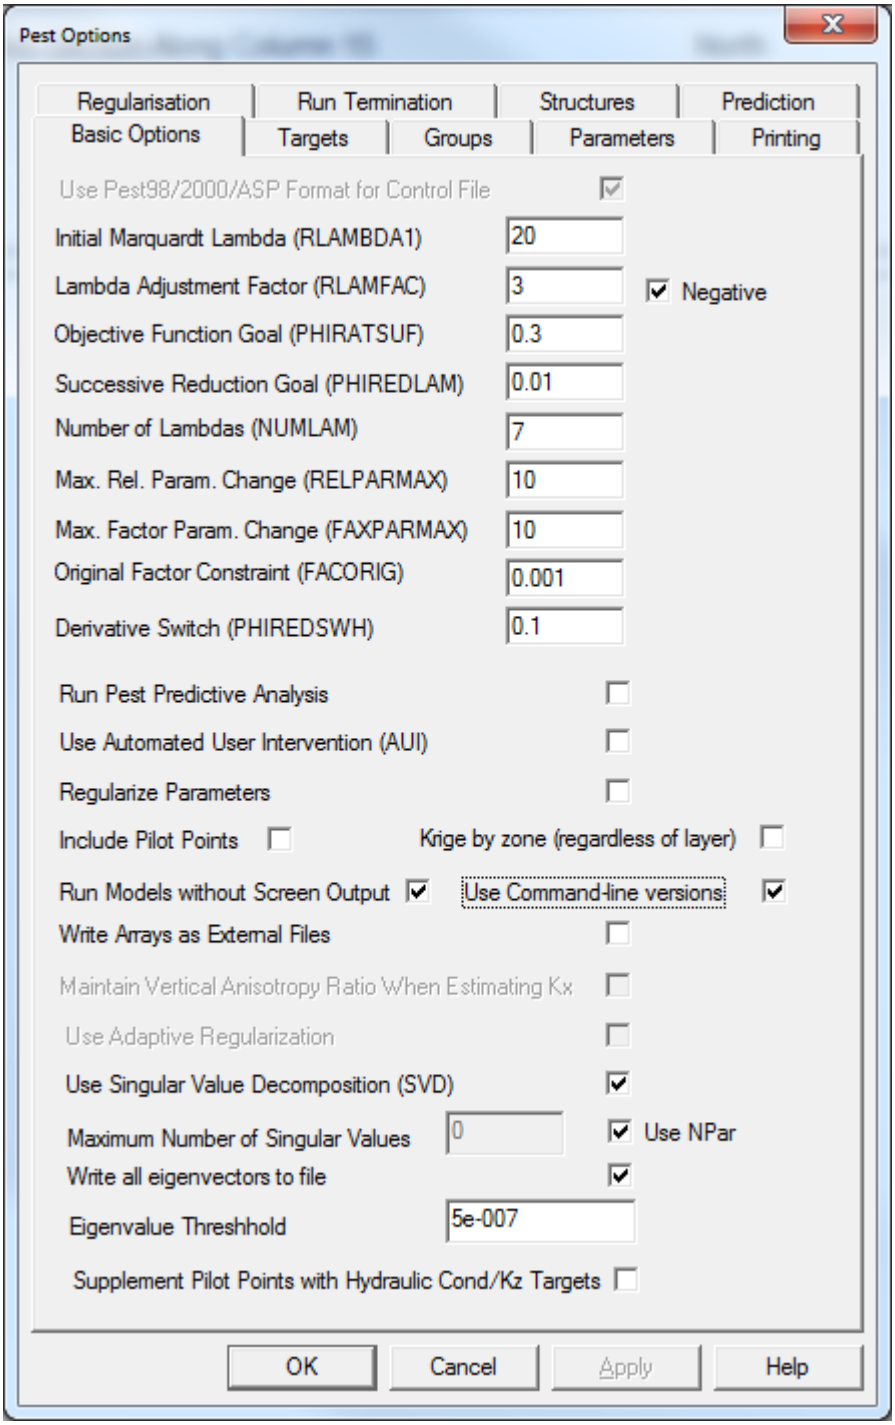

Click OK when you are done. That is all you need to do for a simple Pest problem. By "simple" we mean estimating just a few parameters in zones with only head targets. If you had other target types you would also need to activate those on the Target tab shown above.

Now just like MODFLOW you need to have GV create the Pest input files. Select Model|Pest|Create Datasets. Another important step with Pest is to run pestchek which looks for problems in the pest input files. Do this now using Model|Pest|Run Pestchek. GV will run Pestchek and put the messages in a text file, then display that text file in your text editor, as shown below.

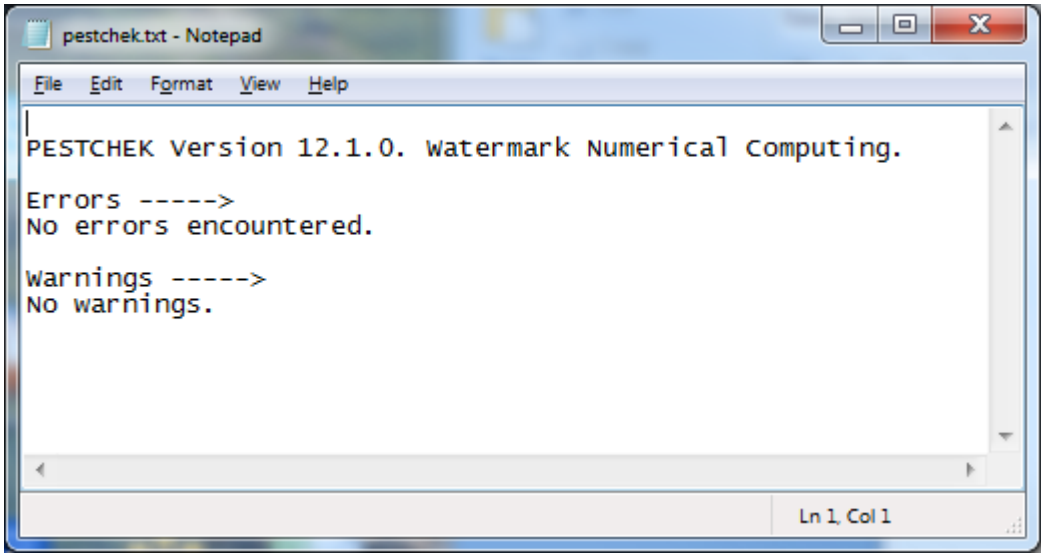

There must be no errors in order to continue. Warnings are generally not important but if you get some warnings and are not sure of their significance, feel free to send us the pestchek output. In this case you should get no errors or warnings.

The last step is to run Pest. Choose Model|Pest|Run Pest. A console (DOS) window should pop up on your screen like the one below.

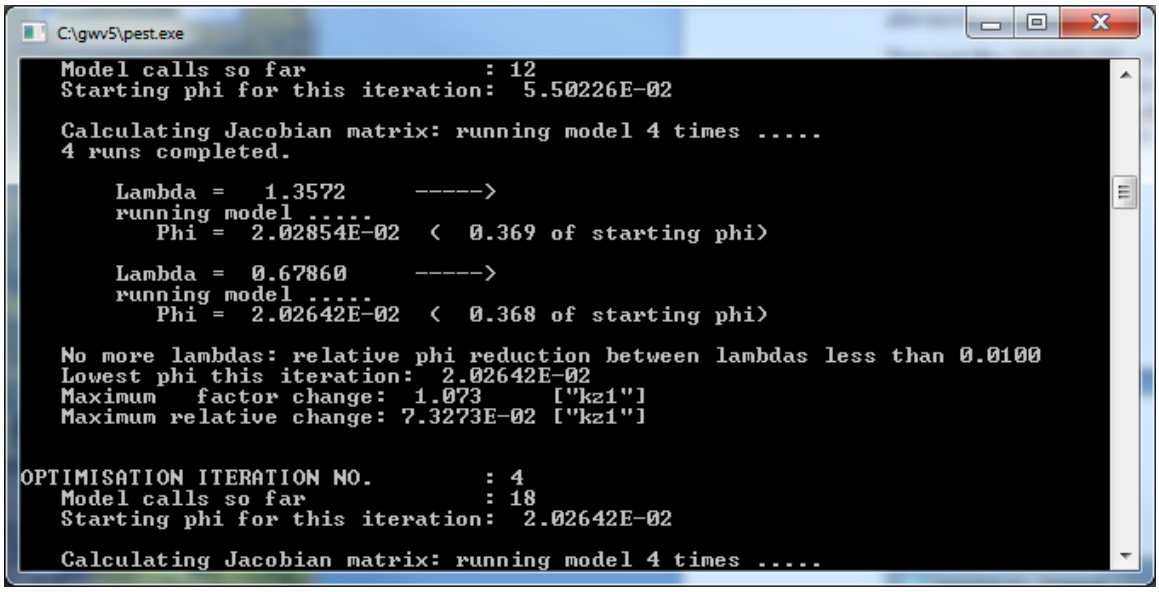

During each pest iteration, the model is run one or two times per parameter. Phi is the sum of squared errors, which in our case starts out at 42.9. As Pest runs, this value should decrease. In the example shown above, Phi is 0.02 at the 4th iteration. It quits soon after that.

When Pest finishes the parameter estimation run, it runs the model one time with the optimal parameter values. To see how it did, just use Plot|Import Results. Those results are the ones from that last simulation that Pest ran. You should see that the residuals are all quite low (< 0.1 ft).

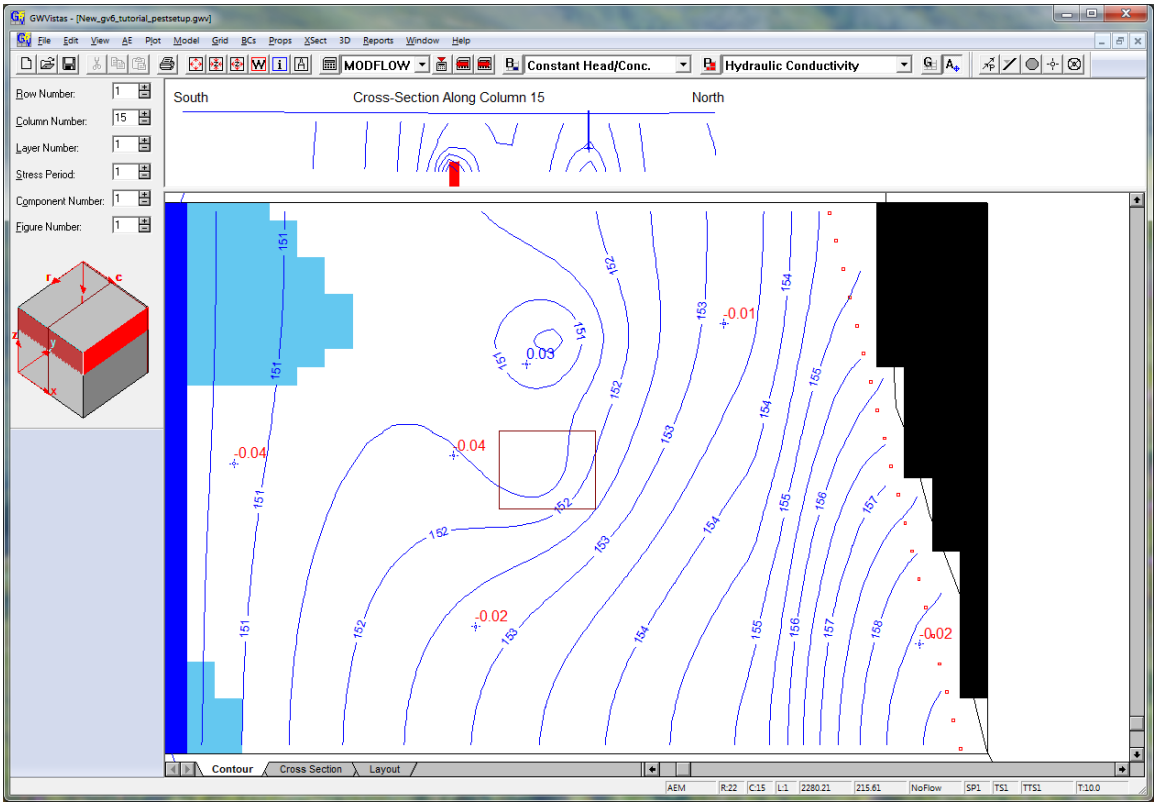

If you select Plot|Calibration|Statistics, you'll see that the sum of squared errors is quite low at about 0.02. To see the final parameter values, select Model|Pest|View Parameter Estimates. This shows the Pest par file output.

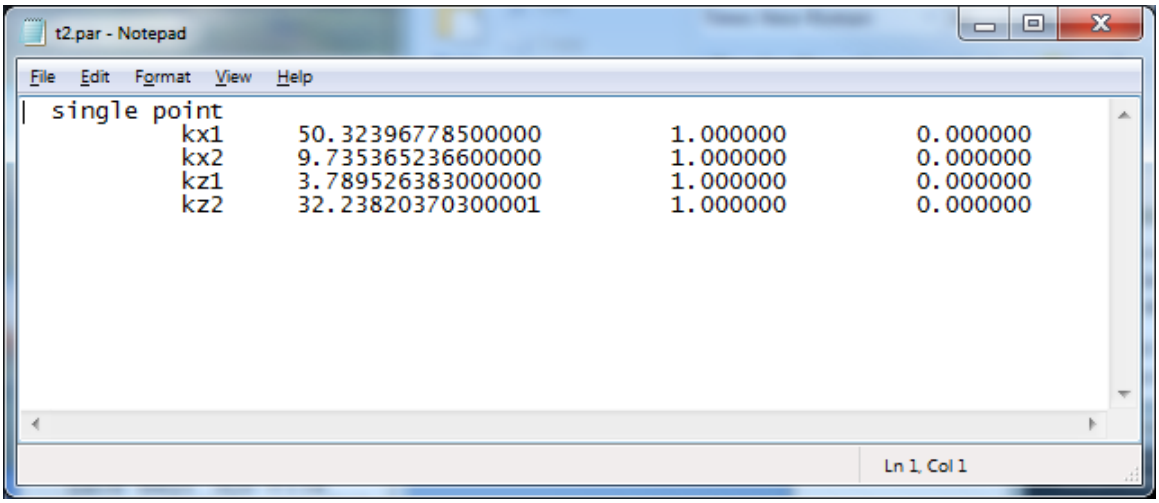

The first column contains the variable names, which should be obvious but they are also in the order presented under Model|Pest|Parameters if you have any questions. In this case, Pest decreased Kx1 from 100 to about 50, which is the correct answer. You will also notice that Kx2 is lower than Kz2. This happened mainly because Kz2 is a very insensitive parameter. You can see relative parameter sensitivity using Plot|Calibration|Pest|Final Sensitivities. A bar chart is created to illustrate the parameter sensitivity coefficients computed by Pest. The larger the bar, the more sensitive is the parameter.

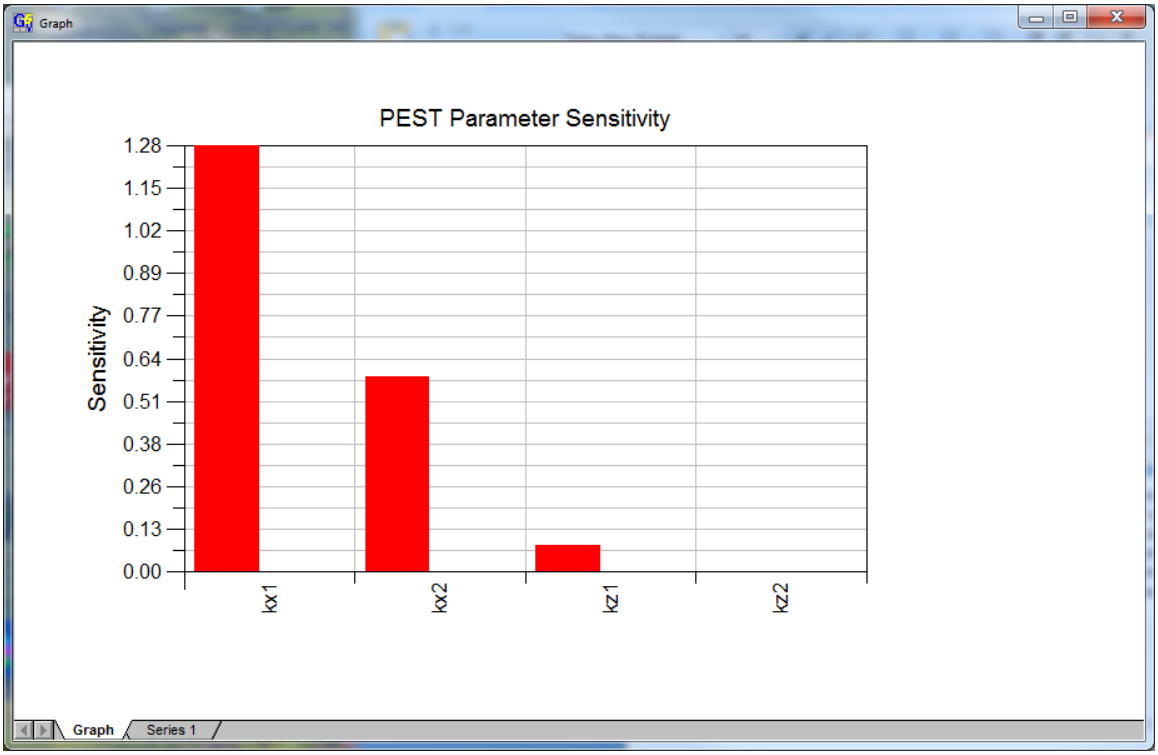

You can see from this that Kz2 is totally insensitive. The best way of dealing with this situation is to simply not estimate that parameter. You can try this by going back to the Pest parameter spreadsheet and removing the check next to Kz in zone 2. Then recreate the Pest input files and rerun Pest.

Once Pest is finished, it has estimated new values for model parameters and boundary conditions. However, those values are not part of your model design in Groundwater Vistas. They reside in the *.par file that Pest creates. To update the model with the latest Pest parameter estimates, select **Model|Pest|Update Parameters**. After using that function you should be able to run the model and get the same results that Pest did at the end of the estimation.

Setting Up a Transient Model

We will now load the previous model that you saved as T2.GWV. Select **File|Open** and select T2.GWV (or whatever you called the one before starting on MODFLOW2000).

You will now set up a model to evaluate the recovery of the aquifer after the pumping wells are shut off. The key aspects of this session will be to delete the 2 wells in layer 3 and set accurate initial conditions from the previous head-save file. First, to make sure that the proper head-save file exists we will rerun the T2 model. Simply click the calculator button  , create the MODFLOW data sets, and run the model. You do not need to import the results after the simulation.

The first step is to set up MODFLOW options to make a transient run. Select **Model|MODFLOW|Packages** and change the root file name to *T3*. Now select **Model|MODFLOW| Package Options** and uncheck the option labeled *Steady State Simulation*. Also, change the number of stress periods to 2. Click OK when you are done. Groundwater Vistas will ask if you want to copy recharge and ET to the new stress periods. Answer YES to these questions. Then GV will ask if you want to edit the stress period setup. Also answer Yes to this. We will have the first stress period be steady-state and the second stress period be transient. Do not change stress period 1 (1 time step of 10 days duration). Change the stress period length to 60.0 days for stress period 2. Change the number of time steps to 30 for stress period 2.

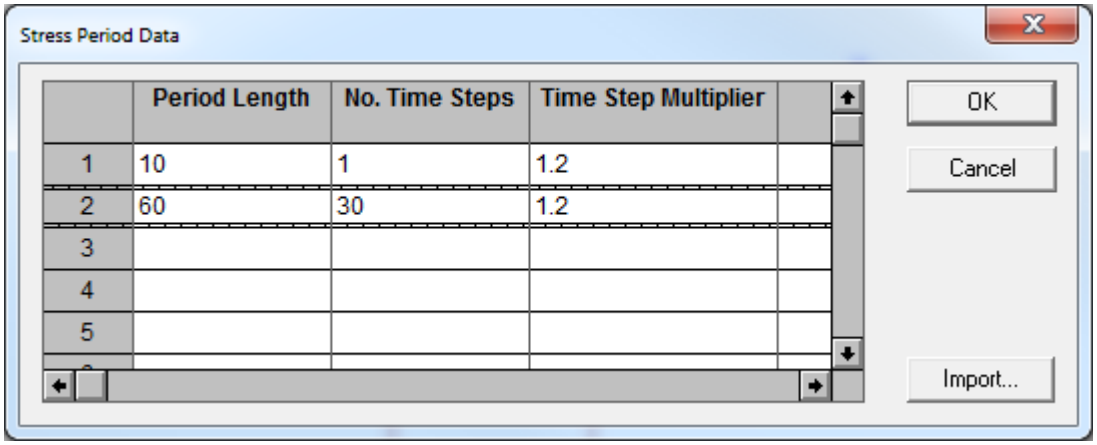

Since we are using MODFLOW2000, we also have to set the stress period type for each stress period. Groundwater Vistas knows this and will next display the stress period types spreadsheet. You could also get to this using **Model|MODFLOW2000|Stress Period Types|Edit Stress Period Types**. Make sure that the type is set to 0 for stress period 1 and 1 for stress period 2. A value of 1 means a transient stress period and zero is steady-state. Each stress period in MODFLOW2000 can be individually set to either steady-state or transient and you can mix them within any simulation. This is very handy, as we will see.

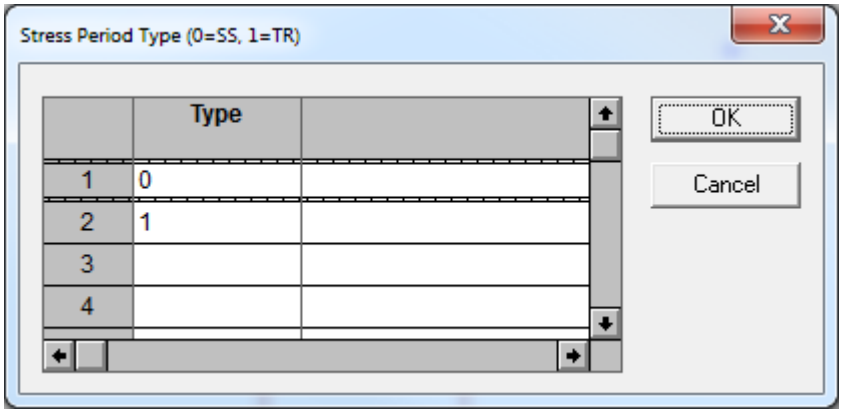

We now have a simulation with one steady-state stress period followed by a transient stress period. This is a very common type of simulation when performing a transient calibration for example. The next step is to turn off the wells in the second stress period to see how long it takes for the aquifer to recover.

We will start by editing the wells in layer 3. Click the + button next to "Layer" on the Reference Cube until you are in layer 3. Now, select **BCs|Well** to edit the wells. Move the cursor over one of the two wells (they are red squares on the screen) and double-click the left mouse button. Now, turn OFF the flag that says "Steady-state Boundary" and click the transient data button. In the first line of the spreadsheet enter a 1 for the starting and ending stress period. Enter -40000 for the pumping rate. You can then either leave line 2 all zeros or to be more explicit, enter a 2 for the starting and ending stress period and a zero for the pumping rate. This tells GV to have the well pump 40,000 cubic feet per day in stress period 1 and nothing in stress period 2.

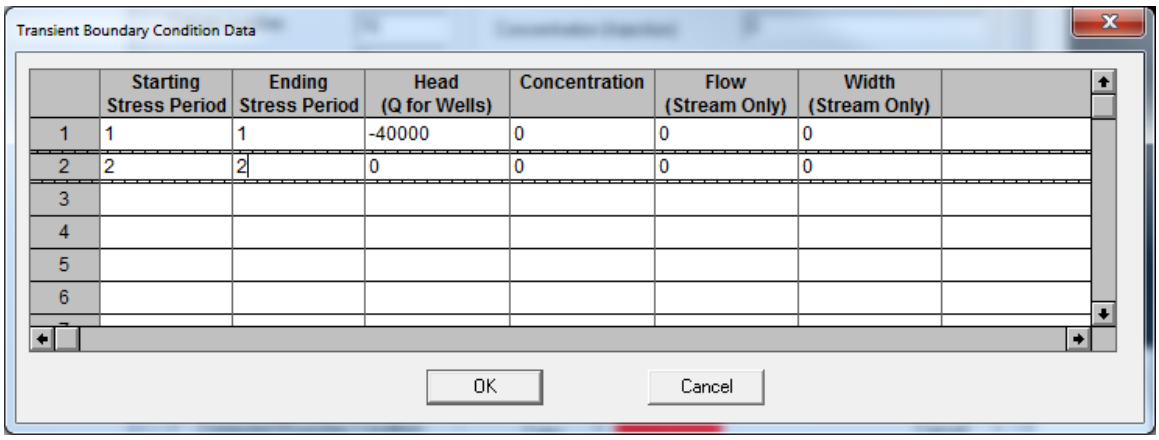

Repeat this procedure for the second well, which pumps -30000 cubic feet per day.

Steady-state models do not require accurate starting heads, however, transient simulations should start with accurate heads in all cells. By accurate, we mean that the heads should be consistent with properties and boundary conditions, otherwise the first few time steps are spent equilibrating the model to the boundary and parameter setup in the model. In this example, we really do not need to have accurate starting heads because the first stress period is steadystate. However, this is a useful thing to know how to do so we'll set up starting heads from the previous run anyway. In this example, we will use the heads from the T2 simulation. Select **Model|MODFLOW|Package Options** and click on the **Initial Heads** tab. Change the option at the top to Se*t Heads from Headsave, BASIC, Surfer, Matrix* and enter *t2.hds* next to file name for the head-save file.

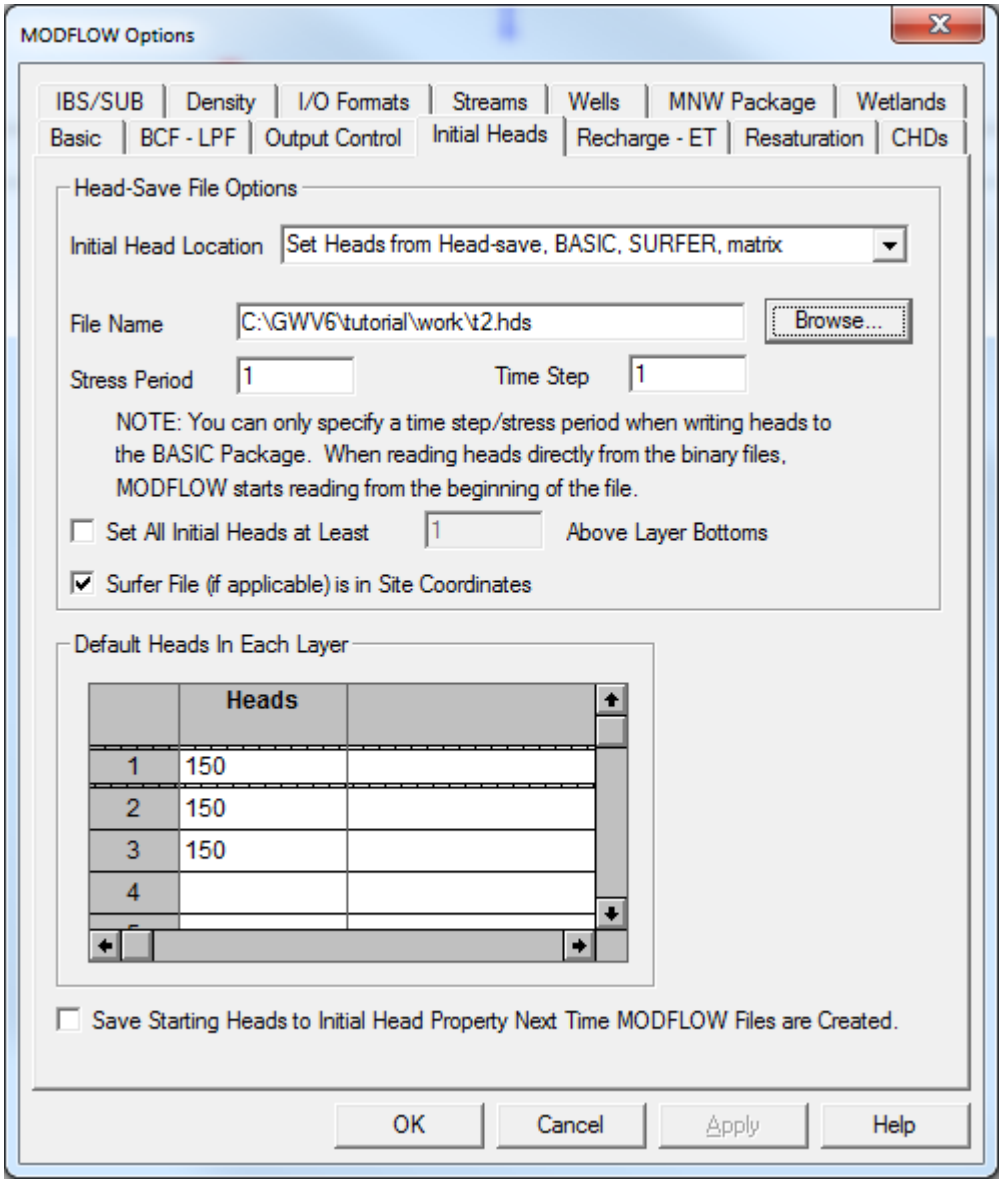

Transient runs require storage coefficients for each cell in the model. This has already been set when you set up the first run. Each cell in the model is assigned a storage value of 0.01. You can confirm this by selecting **Props|Storage**. Next click the database button on the toolbar  . Zone 1 should have specific yield and storage coefficient values of 0.01. When using MODFLOW2000 with the Layer Property Flow (LPF) Package, note that the first column in the database is specific storage. This is different from the BCF Package where this value would be the dimensionless storage coefficient. Change specific storage (Ss) in this example to 1.0e-05. Click OK when you are done.

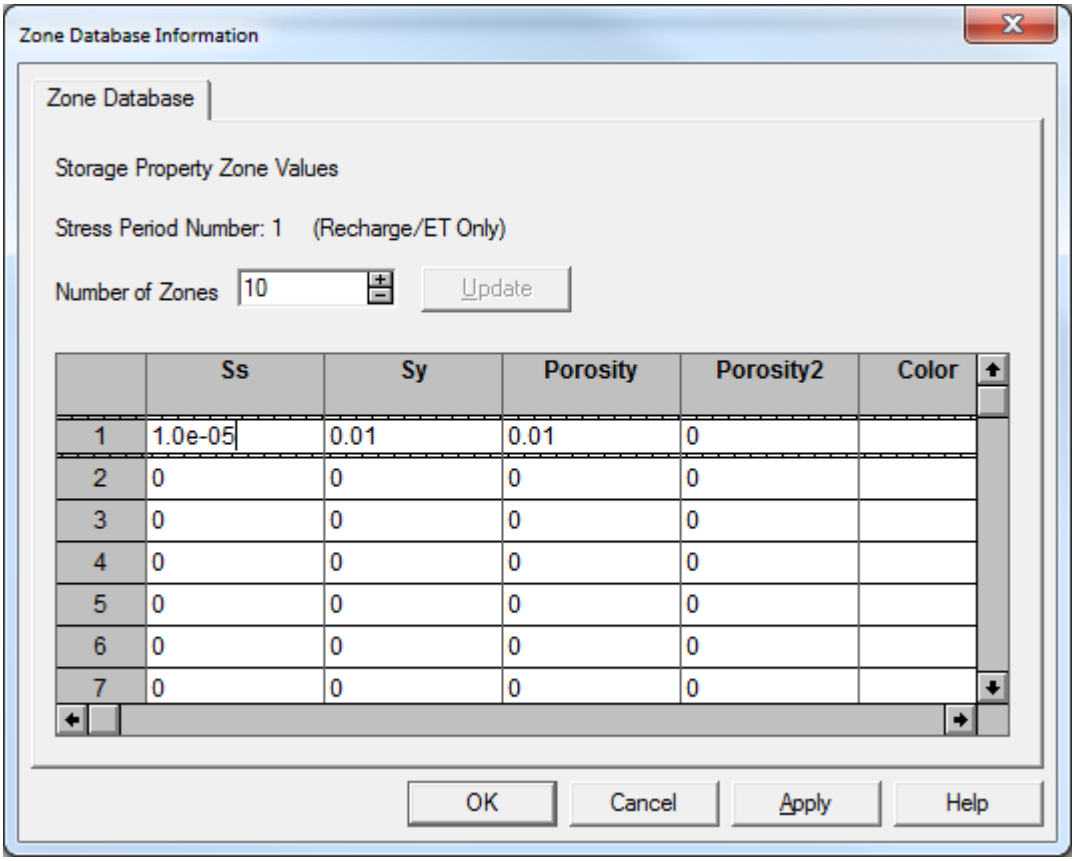

GV allows you to monitor head, drawdown, and concentration over time during a transient simulation.

You must first add observation wells, though to record the heads. Select  from the toolbar to enter Analytic Element mode. You add observation wells by selecting **Add|Well**. Move the cursor to Row 7, Column 15 and click the left mouse button. Make sure to change the pumping rate to 0.0 and place a check mark next to *Monitor Head|Concentration vs. Time*. Enter a 1 next to *Top Layer of Screen* and 3 next to *Bottom Layer of Screen*. When a monitoring well spans multiple layers you will get a hydrograph for each layer. Click the *Name* button and enter MW-1 for the name of this monitoring well. Click OK when you are done.

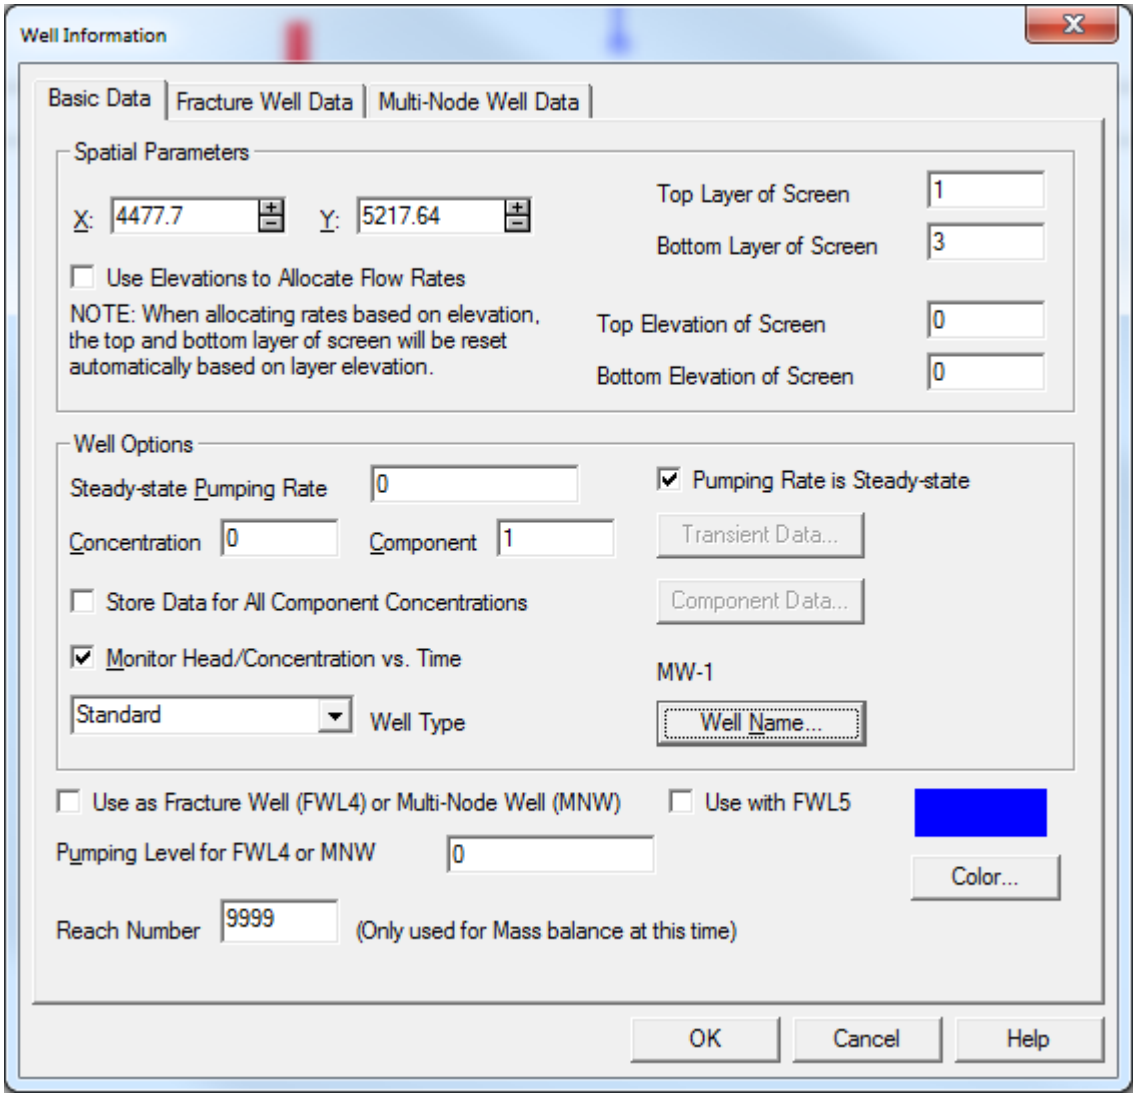

We will do one other thing that is very useful. Normally, MODFLOW computes drawdown in any given time step by subtracting the heads computed for that time step from the initial or starting heads of the simulation. It is often desirable, however, to have MODFLOW compute drawdown based on the heads computed for another time step. In our example, we would like drawdown to be computed based on the steady-state heads computed in stress period 1. MODFLOW2000 and MODFLOW2005 have an undocumented feature that allows this to happen. Select **Model|MODFLOW|Package Options** – Output Control tab and turn on the option for custom output control. Then enter 11 custom time steps.

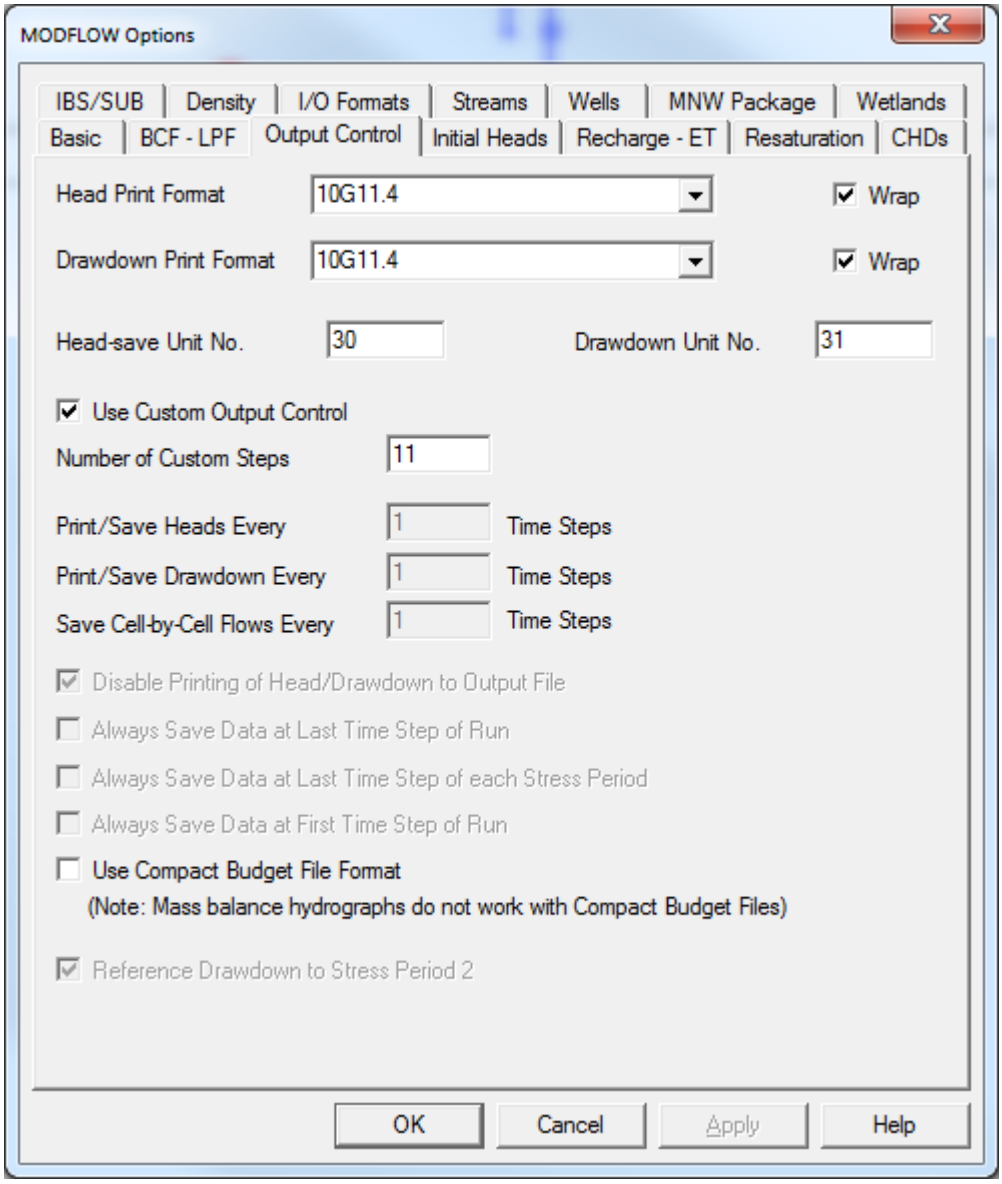
MODFLOW Options
IBS/SUB
Density
I/O Formats
Streams
Wells
MNW Package
Wetlands
Basic
BCF - LPF
Output Control
Initial Heads
Recharge - ET
Resaturation
CHDs
Head Print Format
10G11.4
Wrap
Drawdown Print Format
10G11.4
Wrap
Head-save Unit No.
30
Drawdown Unit No.
31
Use Custom Output Control
Number of Custom Steps
11
Print/Save Heads Every
1
Time Steps
Print/Save Drawdown Every
1
Time Steps
Save Cell-by-Cell Flows Every
1
Time Steps
Disable Printing of Head/Drawdown to Output File
Always Save Data at Last Time Step of Run
Always Save Data at Last Time Step of each Stress Period
Always Save Data at First Time Step of Run
Use Compact Budget File Format
(Note: Mass balance hydrographs do not work with Compact Budget Files)
Reference Drawdown to Stress Period 2
OK
Cancel
Apply
Help

Now, select **Model|MODFLOW|Custom Output Control**. Enter the information as shown below.

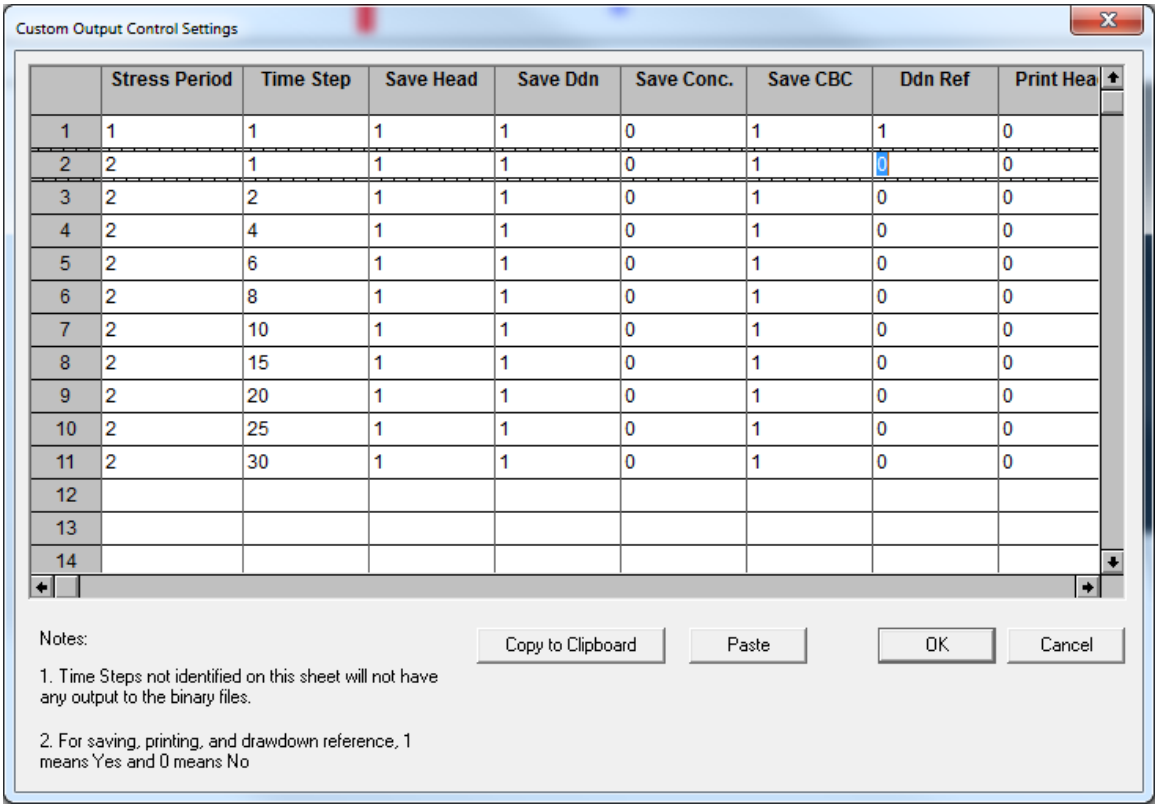

Do not forget to enter a 1 in the first row of the table under the heading Ddn Ref. This is the flag that tells MODFLOW to compute drawdown from the results of this time step.

You are now ready to run the two-stress period model. Click the  button to create data sets and run MODFLOW. When the run is done, select Yes to import the results. Click the *Browse* button at the top of the import dialog to see what time steps have been saved in the head-save file. Choose time step 30 in the second stress period. Make sure you are in layer 1 and your screen should look like the following:

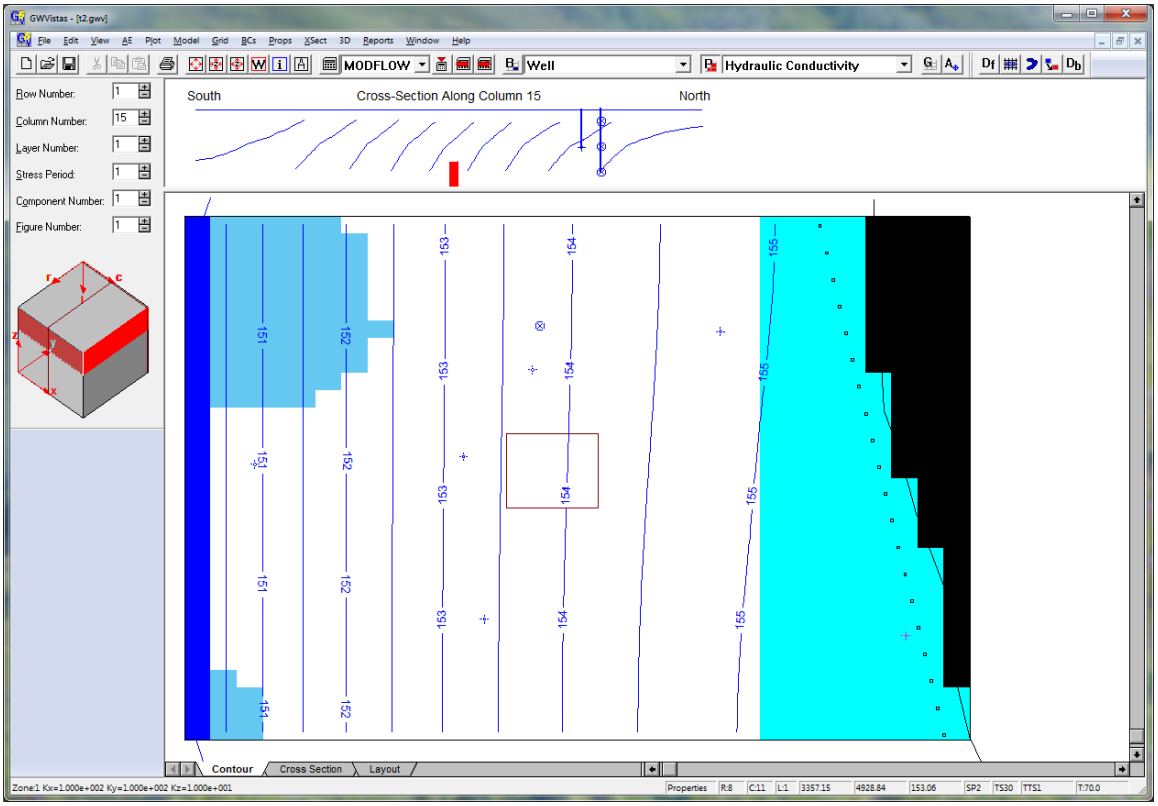

You can plot a hydrograph for the observation well you added by selecting **Plot|Hydrograph|Monitoring Well**. A dialog will show all of the observation wells in the model. In this case, there is only one. Click OK and your screen should look like the one below.

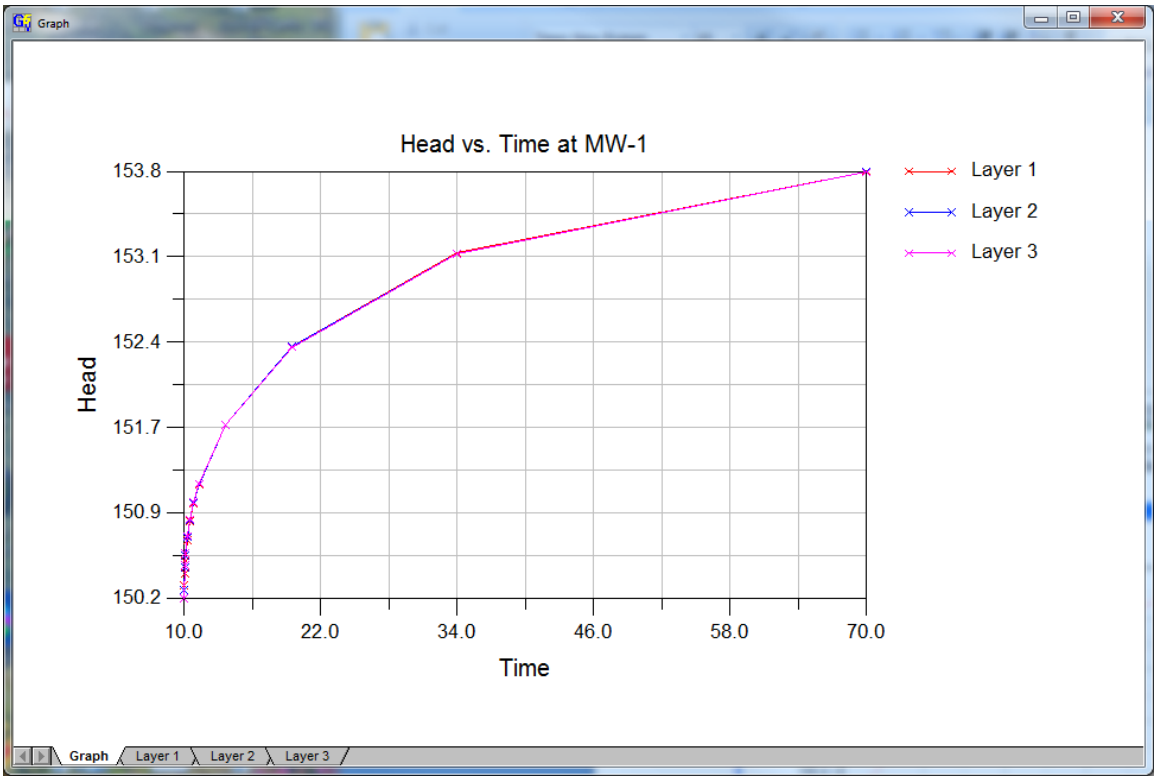

A new option in GV6 is the ability to plot small hydrographs on the plan view next to a well or target. Right-click the monitoring well you just added and select Hydrograph from the context menu. You should see a small hydrograph next to the well as shown below.

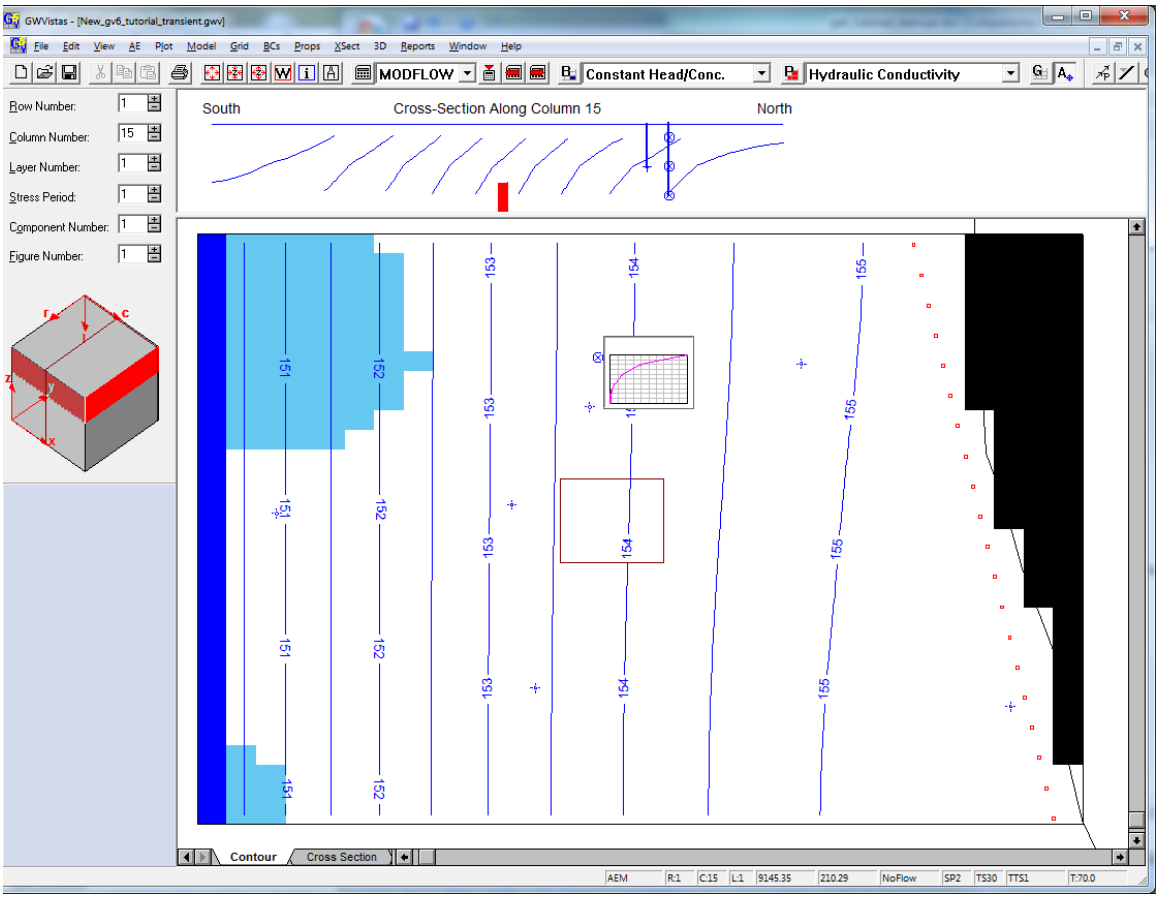

You can control the size of the hydrograph using **Plot|Hydrograph|Plot Options.**

This ends the transient simulation. Save this file for later use by selecting **File|Save As**. Enter the file name TR.GWV.

Transport Modeling with MT3D

We will now introduce you to transport modeling using MT3D. Groundwater Vistas comes with the public domain version of MT3D that has been compiled specifically to work with GV. We will use the new version called MT3DMS but we will just refer to it as MT3D in the following discussion.

MT3D simulations start with a MODFLOW simulation. MODFLOW creates a special flow file (similar to a cell-by-cell flow file) that MT3D uses to compute velocities and flow rates into or out of boundary cells. We will start by going back to the original T2 steady-state run. Close any open models in GV and select **File|Open**. Select the T2.GWV file.

We will now add a source of contamination in the transport model using a special recharge zone. This is a handy way to introduce a source of contamination and is appropriate for leaching of contaminants from the unsaturated zone. Select **Props|Recharge** and then select **Props|Property Values|Database**. Enter a recharge rate of 0.002 in zone 2 and a concentration of 1000.0. Change the color to a nice red! (To change the color, double-click the color field for zone 2).

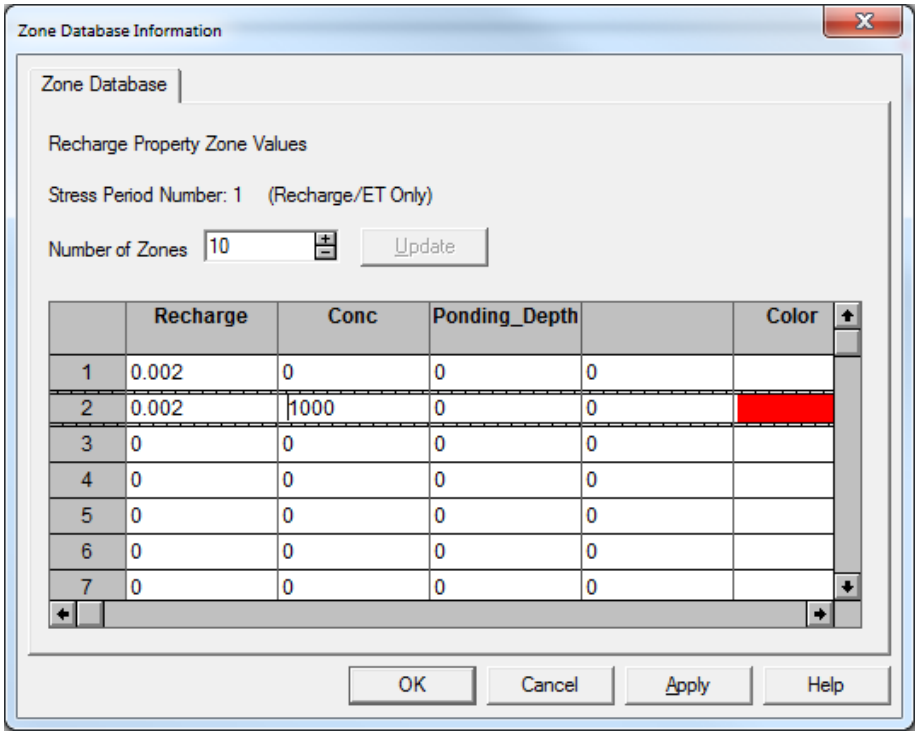

Now, as in the hydraulic conductivity example described above, select **Props|Set Value or Zone Numbers|Window** and drag a window on the east side of the model, similar to the example below. Make this window area zone 2.

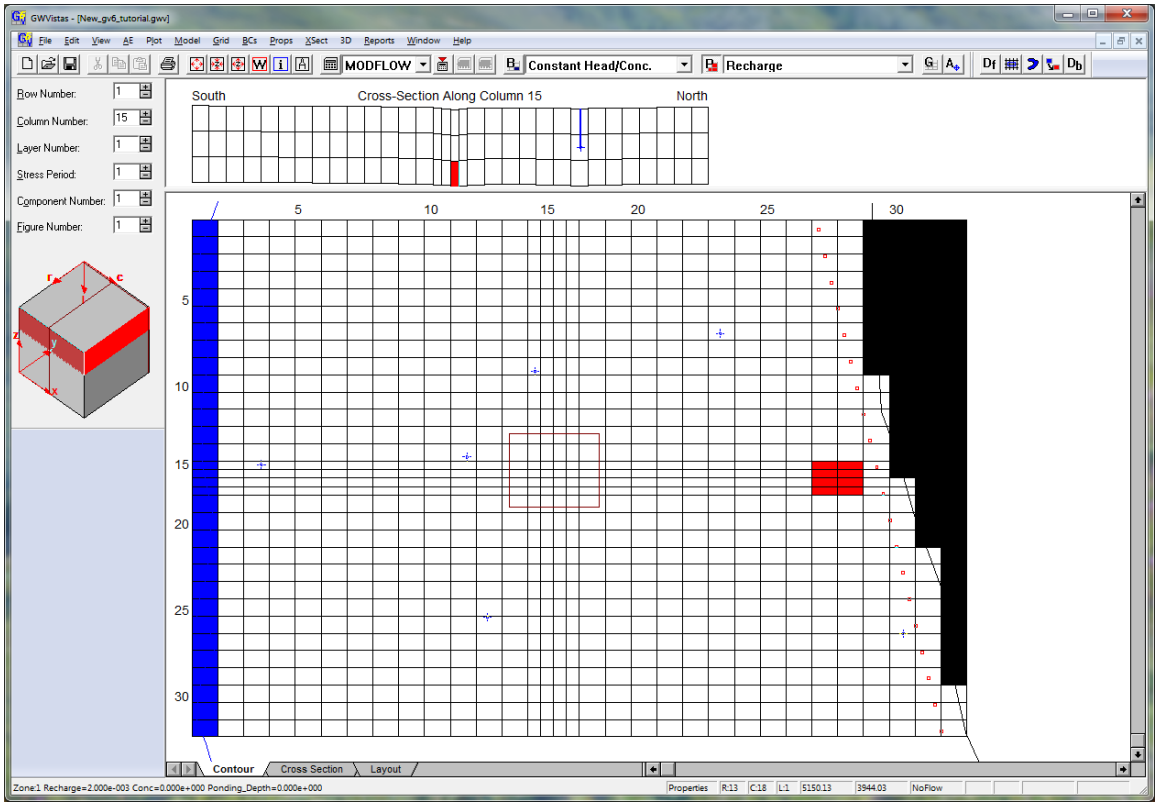

The basics of the contaminant transport and flow models are now configured. You need to first run MODFLOW to generate the flow file for MT3D. Select **Model|MODFLOW|Packages** and change the

root file name to T4. Place a check mark next to *Mt3D Flow Output* at the bottom of the dialog. The latter option creates the file that MT3D requires.

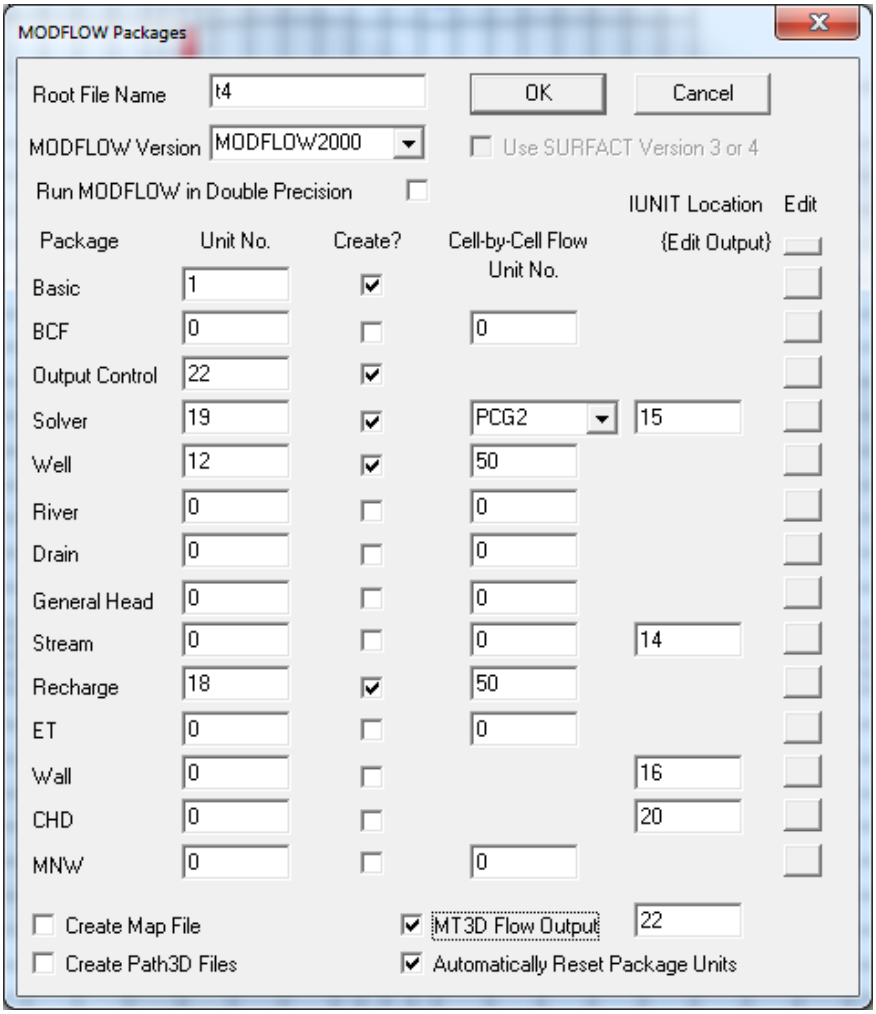

Select **Model|MODFLOW|Stress Period Setup** and change the length of the stress period to 100.0 days. The length of the MODFLOW simulation will then be the length of simulation for MT3DMS. Finally, click the  button to run MODFLOW.

After MODFLOW is done running, you need to set some MT3D options. Start by selecting **Model|MT3D|Packages**. Change the root name to **T4MT** and make sure the version is MT3DMS at the bottom of the dialog. Then turn on the GCG solver package and turn off reactions. We will assume a conservative tracer as a contaminant.

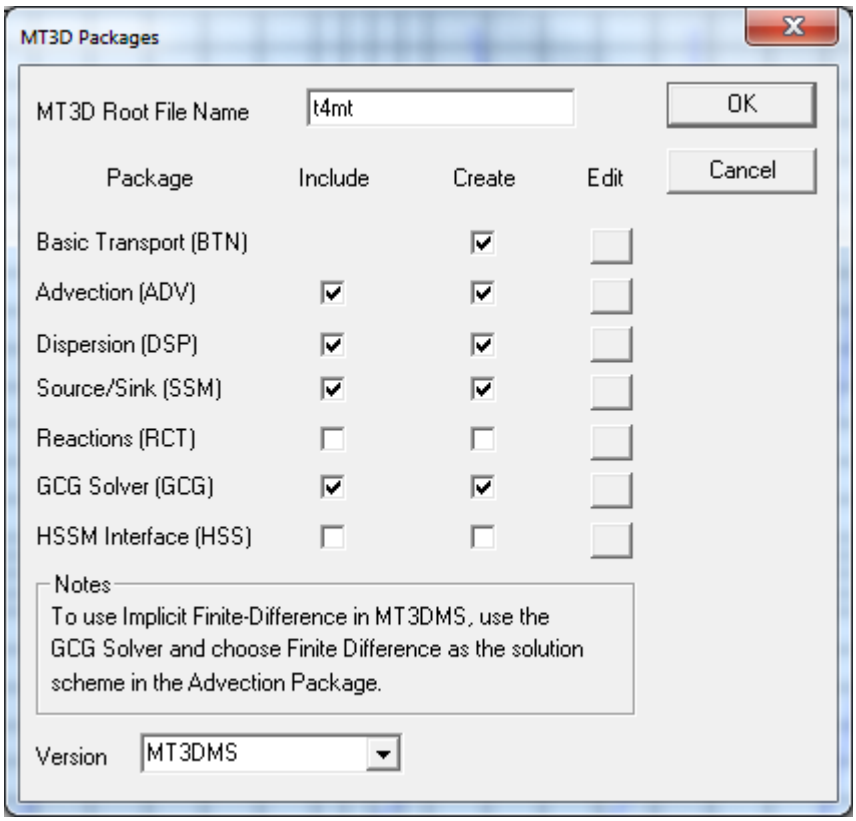

Now, select **Model|MT3D|General Options** and click on the **Advection** tab. The first option on the dialog is the solution scheme. Make sure this is set to *Finite Difference*. This scheme is the most stable of the MT3D methods and often gives the best mass balance, although TVD is quite good as well. The main problem with TVD is that run time is much greater than for finite difference.

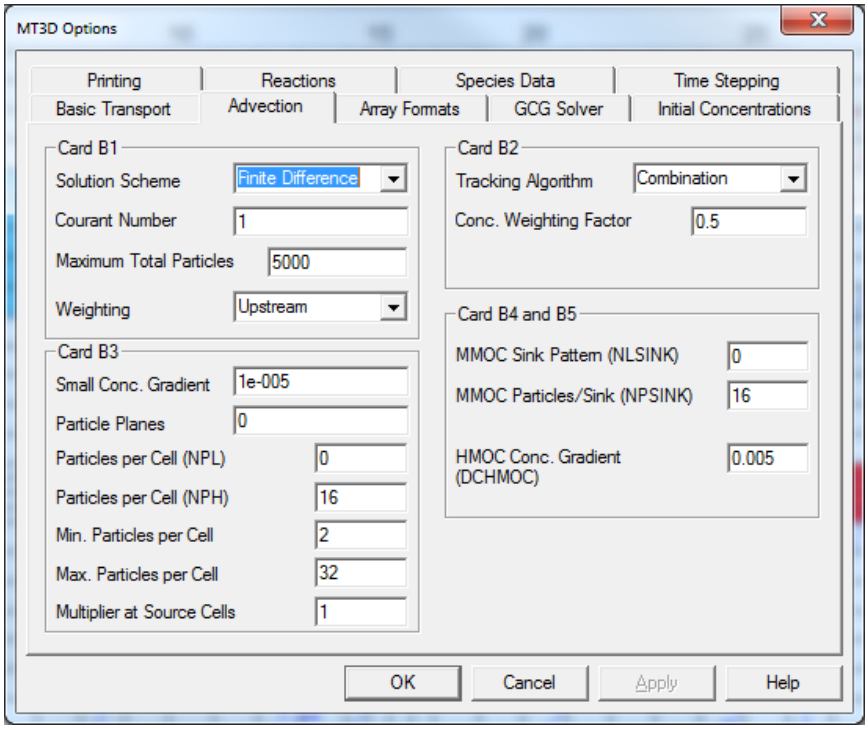

Next, click on the **Printing** tab. Place a check mark next to the item labeled *Save Concentration in Binary File*. Change the frequency of output to *Every N Time Steps* and the number to 2.

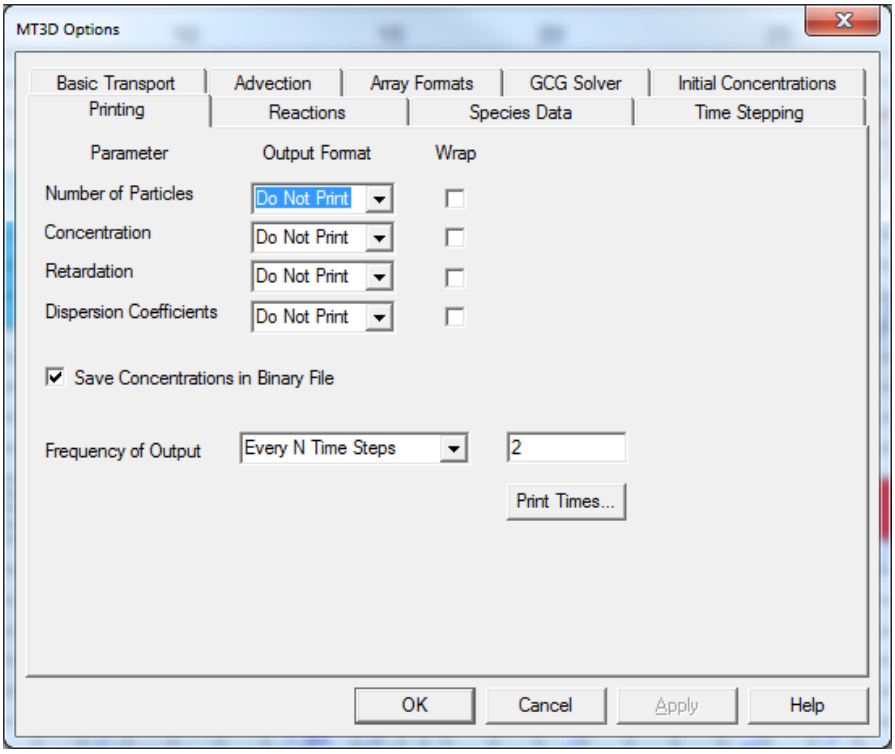

Click on the **Time Stepping** tab. Confirm that the initial time step size is 0.1 with a maximum step size of 500 days and a multiplier of 1.2. Although the MODFLOW run is steadystate, the MT3D run is transient. The same velocities will be used for each MT3D time step. It is generally a good idea to start the transport simulation using very small time steps and increase them gradually. Note that if you choose TVD, the time steps are uniform in length and MT3D will generally override your chosen time step size if it is too large.

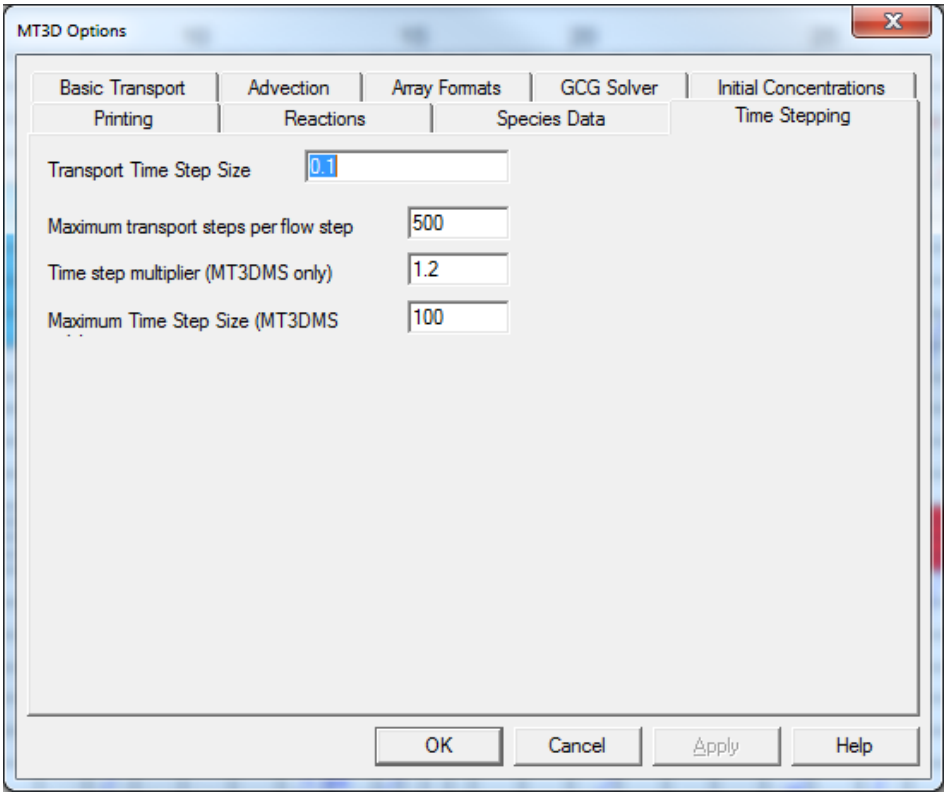

You are now ready to run MT3D and contour the concentrations. Select **Model|Use MT3D**. Now, when you click the  button, GV runs the MT3D model. Do this now. If MT3D does not run, select **Model|Path to Models** and confirm that the MT3D model is *MT3DWIN32.dll*.

After MT3D is done running, import the results. When the import results dialog is on the screen, click the Browse button next to the transport time step. This shows all of the times that concentration was saved. Choose the last time step at 100 days (should be step number 30). Click OK and GV will read heads and concentrations.

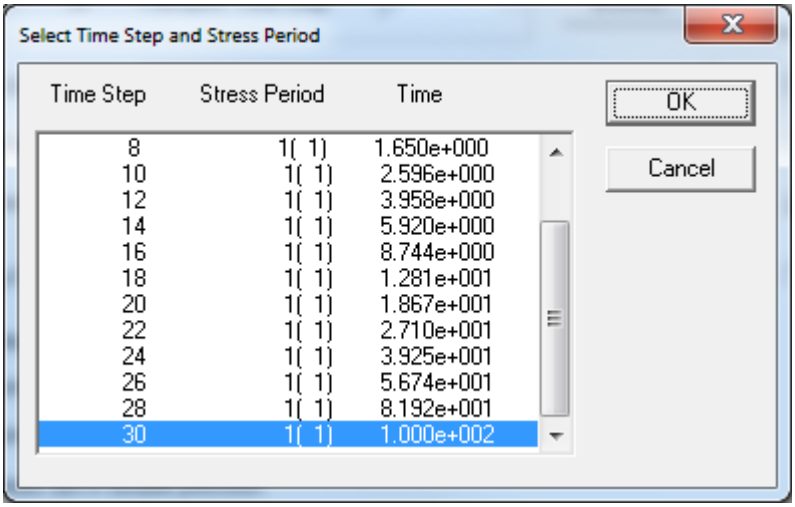

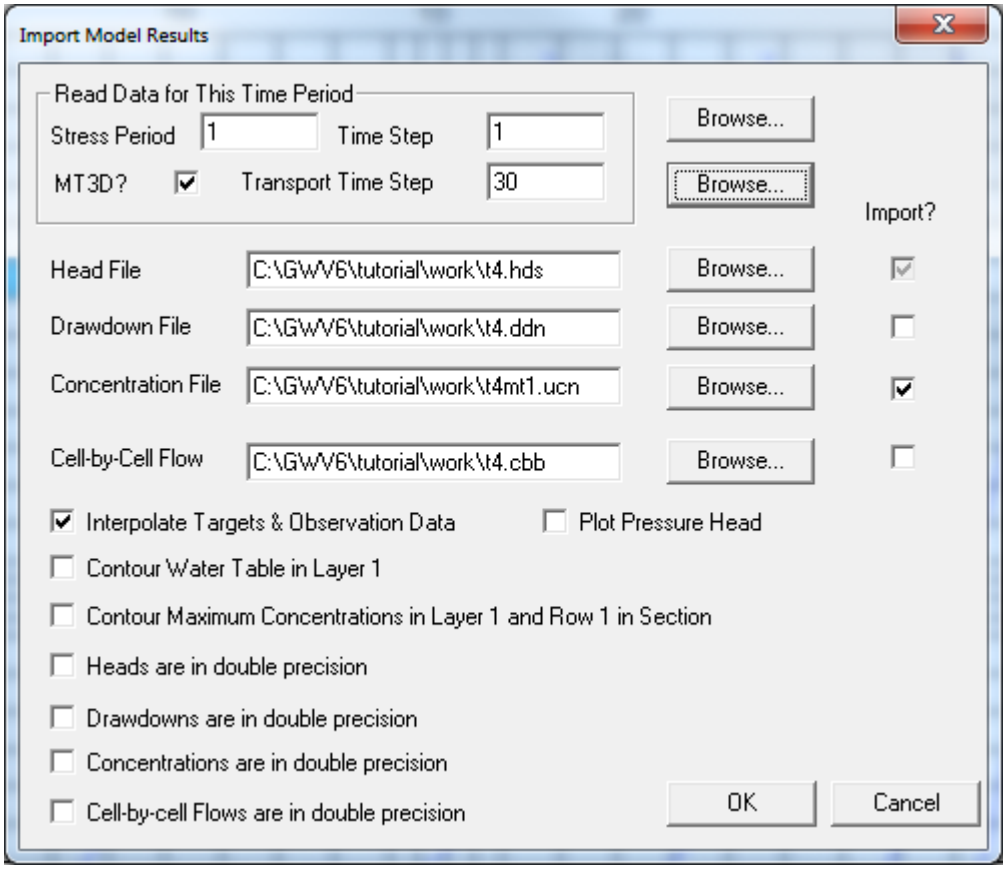

By default GV contours heads. To contour concentrations, select **Plot|What to Display**. Change the variable to contour from *Head* to *Concentration*. Click OK and GV will contour concentrations. Concentrations are contoured on log cycles. Select **Plot|Contour|Concentrations** if you want to change the starting log cycle or turn off the log contouring.

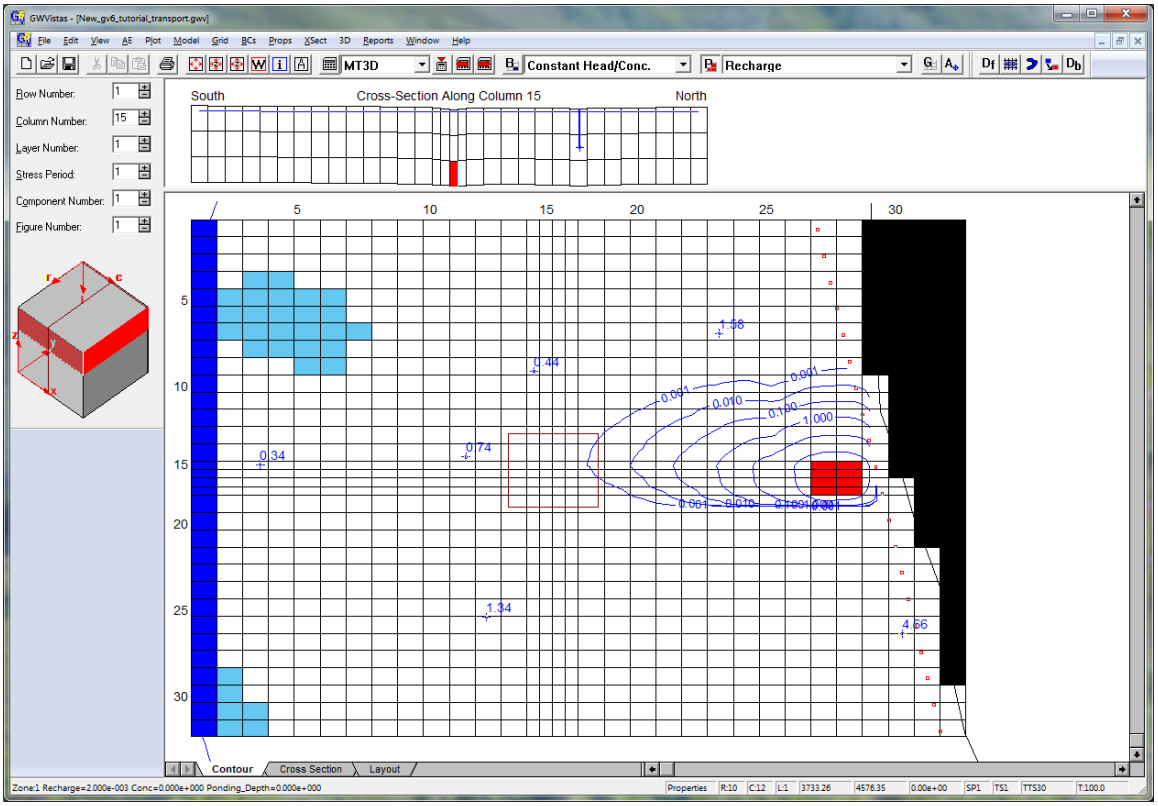

Stochastic MODFLOW

The official name of this product is *Stochastic MODFLOW* but it also includes Monte Carlo versions of MODPATH, MT3D, and Groundwater Vistas. *Stochastic* MODFLOW is the first modeling environment for the MODFLOW family of models that allows for the quantification of uncertainty.

The stochastic simulation method used in this software is called the Monte Carlo technique. This document is meant to be a practical guide to Monte Carlo simulation as it is applied to the groundwater models MODFLOW, MODPATH, and MT3D. No basic instruction is provided on using MODFLOW, MODPATH, MT3D, or on geostatistical simulation techniques. We assume that you are already familiar with these models.

Groundwater Vistas plays a vital role in using the Monte Carlo versions of the MODFLOW family of models. GV provides a convenient way of generating data files for the models, launches the models, and most importantly helps to make sense of the results. You will require a special version of Groundwater Vistas, called the **Advanced** version, in order to be able to run stochastic simulations. If you would like to run through this tutorial to see how the monte carlo versions work, simply remove the dongle from your parallel port. In demo mode, GV allows you to test out the features in the **Advanced** version.

Stochastic MODFLOW comes with Monte Carlo versions of MODFLOW, MODPATH, and MT3D. These are console applications that run in a DOS window. Windows versions will be available in the future. The stochastic models are located in the Groundwater Vistas directory (default: c:\gwv6) and are named:

Smodflow.exe Stochastic MODFLOW

Smodpath.exe Stochastic MODPATH

Smt3dms.exe Stochastic MT3DMS

This tutorial covers the basic operation of the Monte Carlo versions of MODFLOW, MODPATH, and MT3D. The tutorial assumes that you are already familiar with Groundwater Vistas and have run the general GV tutorial in the previous sections of this chapter.

You should first start Groundwater Vistas and open the GV file called *stutor.gwv* in the tutorial directory (default: c:\gwv6\tutorial). The following sections show how to configure the simulation, run the Monte Carlo versions of MODFLOW, MODPATH, and MT3D, and analyze the results.

You need to confirm that the working directory is set up correctly. Select **Model|Paths to Models** and make sure that the working directory (c:\gwv6\tutorial) exists. If this directory does not exist, please change it to one that is valid.

Stochastic MODFLOW

In a deterministic model, you have a fixed set of parameters and boundary conditions, usually from a calibrated model. When making predictions, you simply change the stresses in the model to simulate what will happen to the groundwater system in the future. You make one run and present the result in a report. In some cases, you may bracket the deterministic solution with a best and worst case. While this approach is generally accepted in our industry, it does not really address the issue of uncertainty in parameter distributions in the model and how that uncertainty effects our predictions.

In monte carlo modeling, on the other hand, you can make key model parameters *uncertain* by specifying a distribution type and associated statistical characteristics. Instead of making one simulation, you make hundreds or thousands of simulations. In each simulation, a different value is selected for the uncertain parameters. When processing the results of a monte carlo simulation, you look at the probability that something will happen by evaluating all of the hundreds or thousands of simulations. GV helps you perform this evaluation by organizing and summarizing the monte carlo simulations.

Open the file called **stutor.gwv** located in the gwv6\tutorial directory. We will start this example by selecting the parameter that will be uncertain. In this case, we will just use horizontal hydraulic conductivity in Zone 1. Select **Model|Stochastic|Parameters**. Select Kx as the parameter type. Use a normal distribution, zone number 1, standard deviation of 30.0 ft|d, minimum value of 10.0 ft|d, and maximum of 500.0 ft|d. Your screen should look like the following dialog.

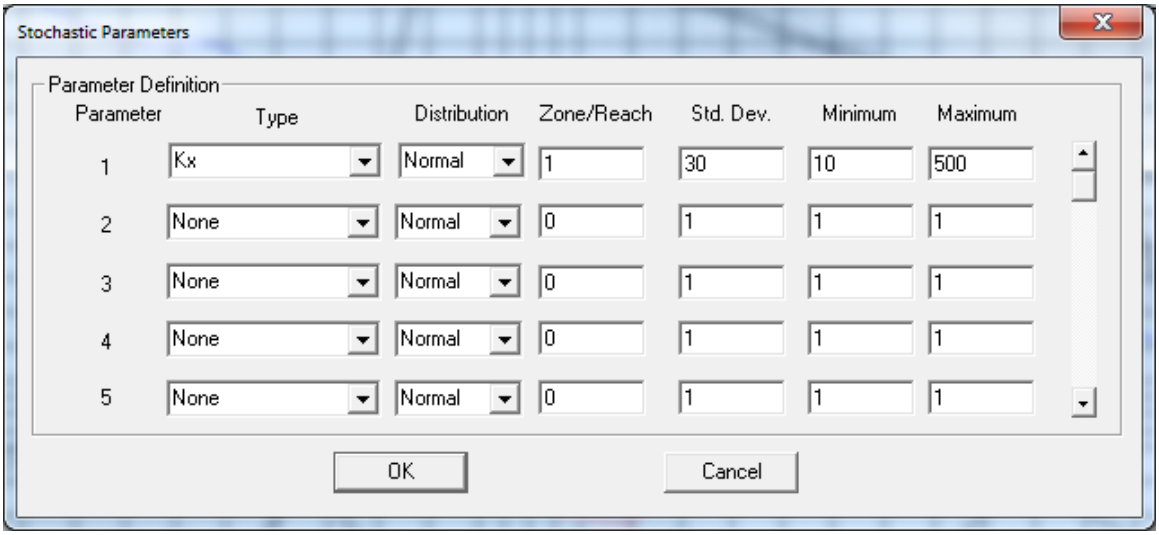

Next, select **Model|Stochastic|MODFLOW Options**. Use an output file name of *smod.out* and 100 realizations. In the jargon of monte carlo simulations, a realization is one MODFLOW run. An ensemble is the collection of all realizations, 100 in this case. Before running the model, change the root file name to *t2s*. This is done by selecting **Model|MODFLOW|Packages**.

We have now selected the parameters that will vary with each realization and how they will vary. In this case, *Stochastic* MODFLOW will use a normal distribution for sampling of K values and will limit the range to between 10 and 500 ft|d. We will also run 100 realizations and each realization will use a homogeneous K distribution.

Now run the simulation by selecting **Model|Stochastic|Create MODFLOW Datasets** and then **Model|Stochastic|Run Stochastic MODFLOW**. If *Stochastic* MODFLOW does not run, check **Model|Stochastic|Paths** to be sure that you have *smodflow.exe* for the *Stochastic* MODFLOW program and it is pointing to no directory or the directory containing the program (c:\gwv6 by default).

After the simulation is done, you may do some calculations and display some of the stochastic results. The first thing to do is to compute the mean head field and the standard deviation in heads. Select **Plot|Stochastic|Compute Mean Std.** On the dialog, enter the name of the head-save file (c:\gwv6\tutorial\t2s.hds) and also enter a name for the mean and standard deviation files (mean.hds and std.ddn are good choices). Click OK to compute these values.

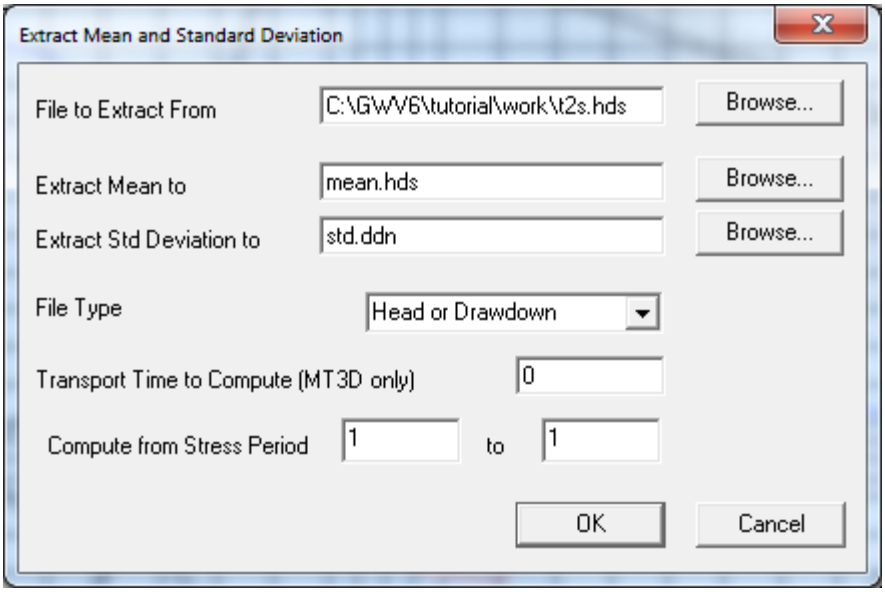

We will now import the mean and standard deviation into GV for contouring. Select **Plot|Import Results**. Browse to find the *mean.hds* file for the head file and the *std.ddn* file for drawdown. Even though the latter file is not really drawdown, it has the same format as drawdown. Your screen should look something like the following (you may need to modify the contour interval to get the exact same results).

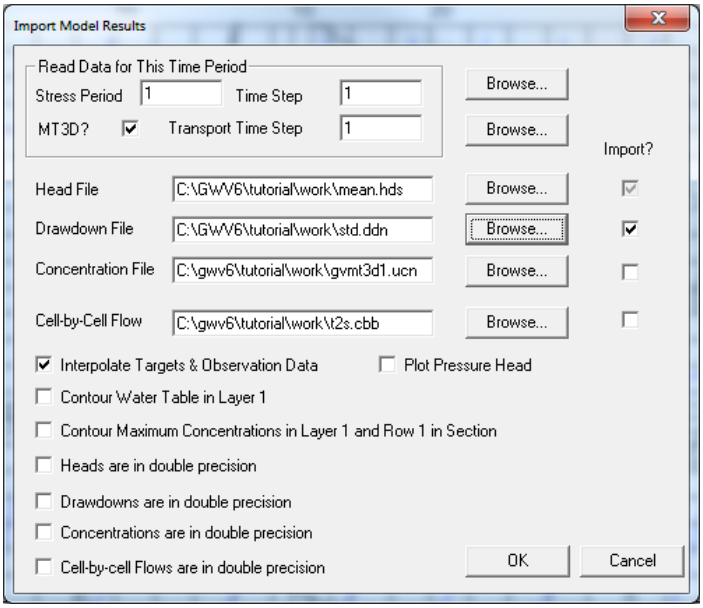

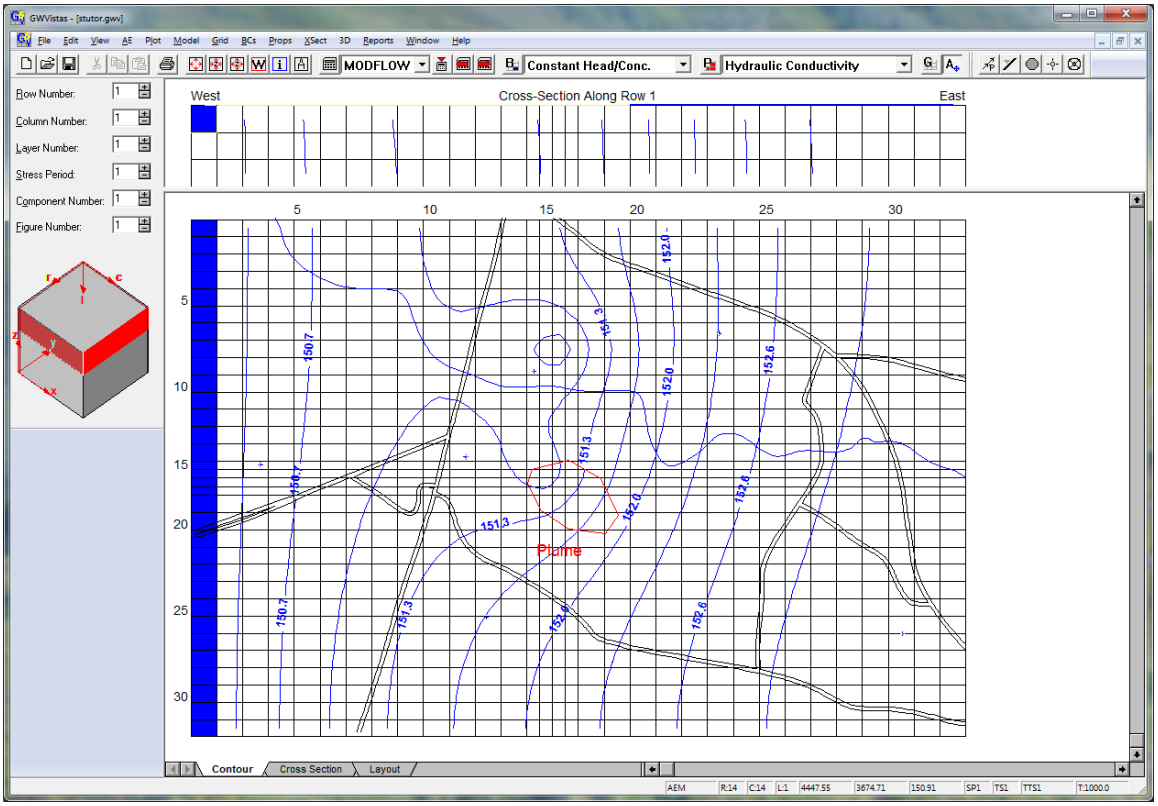

The contour map on your screen is the average of heads from the 100 realizations. To contour the standard deviation in head, select **Plot|What to Display**. Change the contoured variable from head to drawdown. Remember that drawdown in this case represents the standard deviation.

This shows the variability in head through the 100 realizations. Remember from basic statistics that about 68 percent of the heads in the simulation should fall within the one standard deviation of the mean and about 95 percent will fall within two standard deviations.

Now we will evaluate how the calibration results look for the various realizations. First select **Plot|Stochastic|Import Target Data** and choose the file *t2s.hds* as the input file. One way of looking at the results is to create a scatter plot of sum of squared residuals for each realization. Select **Plot|Stochastic|Graph|Calibration Statistics**. Simply click the *Scatter Plot for All Realizations* button and you should see the following graph:

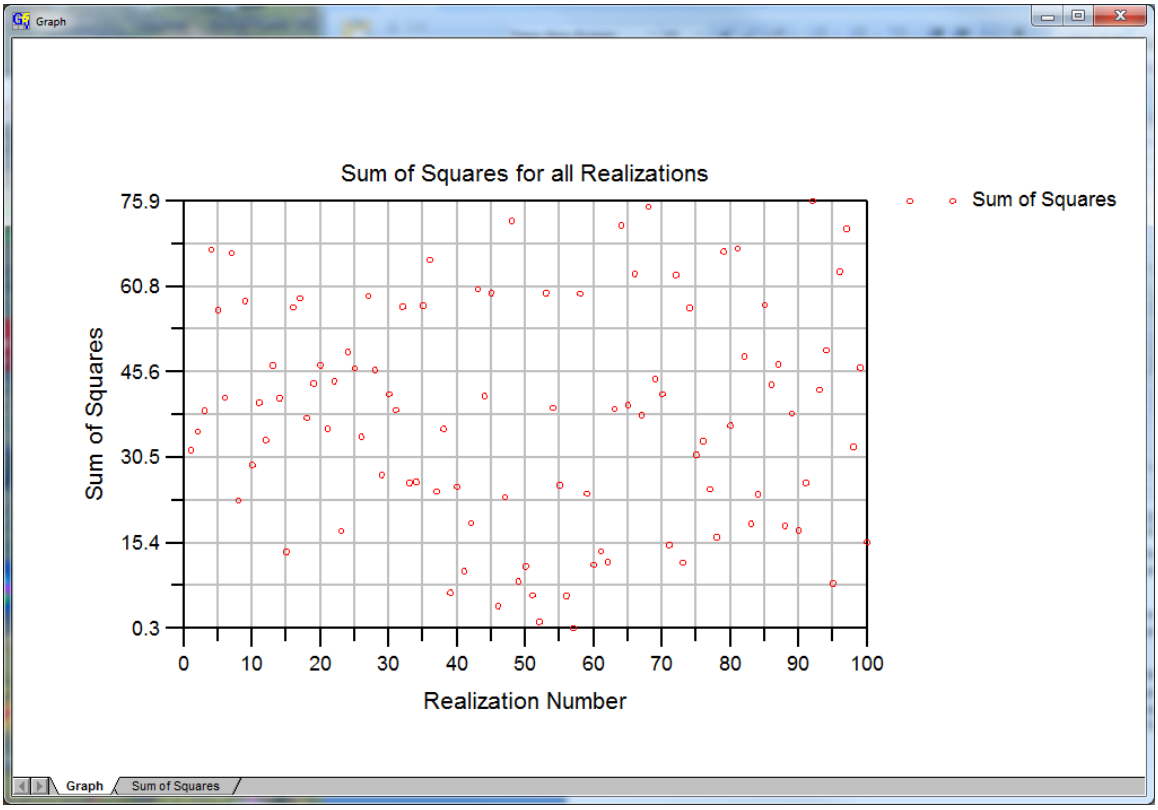

The base case in this simulation has a sum of squared residuals of 39.4 so you can see that there are many simulations that have a better calibration. You can see which ones are better by clicking on the tab below the chart called "Sum of Squares". If you scroll through the data spreadsheet, you will see that realization number 57 has a sum of squared residuals of 0.25 which appears to be the lowest. To see what K value was used in that realization, edit the file called *param.dat*. Scroll down to realization 57 and you will see that a K value of 52.26 ft|d was sampled. In fact, this is quite close to the theoretically perfect value of 50.0 ft|d for this simulation (that is, we created the targets in this simulation using a K value of 50 ft|d so we know what the *correct* value is in this case). While this is somewhat wasteful of computer resources, you can see that stochastic simulation could be used to augment model calibration (using the theory that *even a blind squirrel finds an acorn every once in a while*).

You can even see the head distribution from realization 57 by selecting **Plot|Stochastic|Extract Realization**. Select the file *t2s.hds* and enter a realization number of 57. Click OK to create a new file; call it r57.hds. To view the heads, select **Plot|Import Results** and use r57.hds for head.

Another thing you will often want to do is justify that enough realizations were simulated. You can get a handle on this by selecting **Plot|Stochastic|Graph|Calibration Statistics** as above and then press the *Cumulative Average* button. Your screen should look like the following:

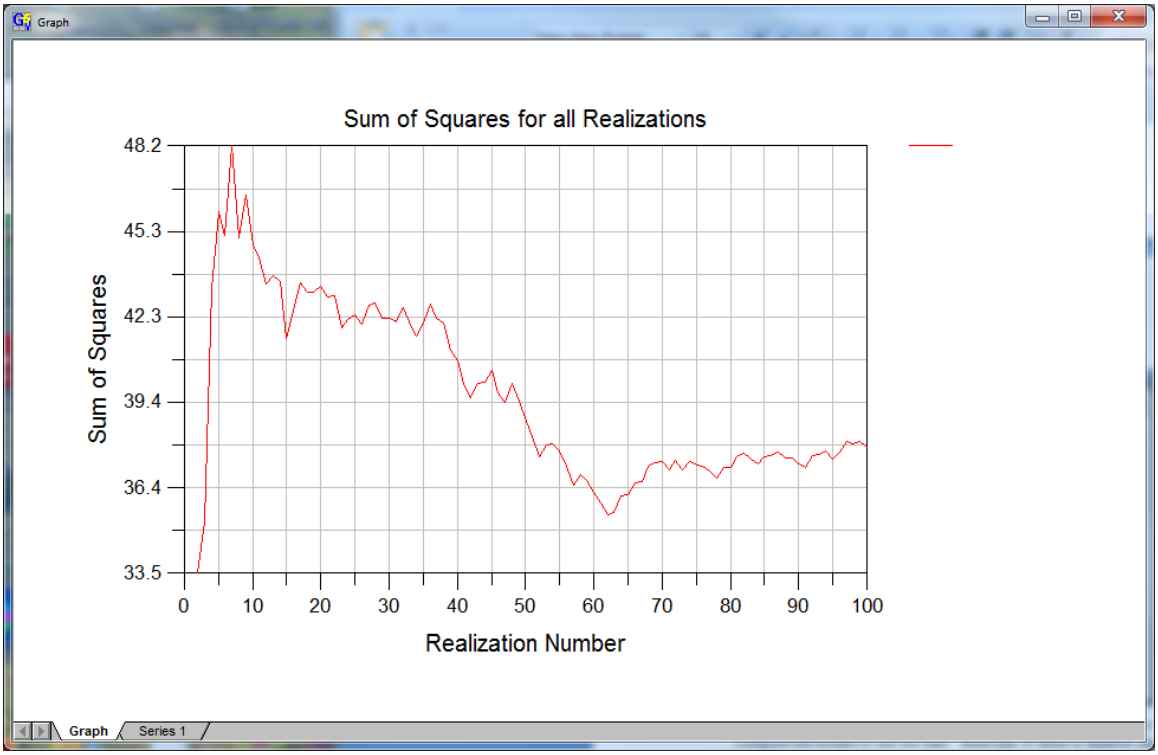

In this case, we cannot say conclusively that 100 was enough. We would need to run 150 or 200 and compare the results to see for sure. Basically if there is limited change in the cumulative average then you might conclude that you ran enough simulations. Another check is that the mean head field should be virtually identical to your steady-state calibrated model (assuming that all parameters used either a normal or log-normal distribution). If the mean head field and the steady-state calibration heads are the same, then you can at least conclude that you have sampled enough realizations to adequately characterize the mean response of the model.

Another common type of plot is the CDF (cumulative distribution function) which can be plotted for targets or observation wells. In the example model, we have 16 targets. You can create a CDF plot for a target by selecting **Plot|Stochastic|Graph|Target CDF**. Choose the first target on the list and click OK (actually it should already be chosen so just hit OK). Your plot should look like the one below.

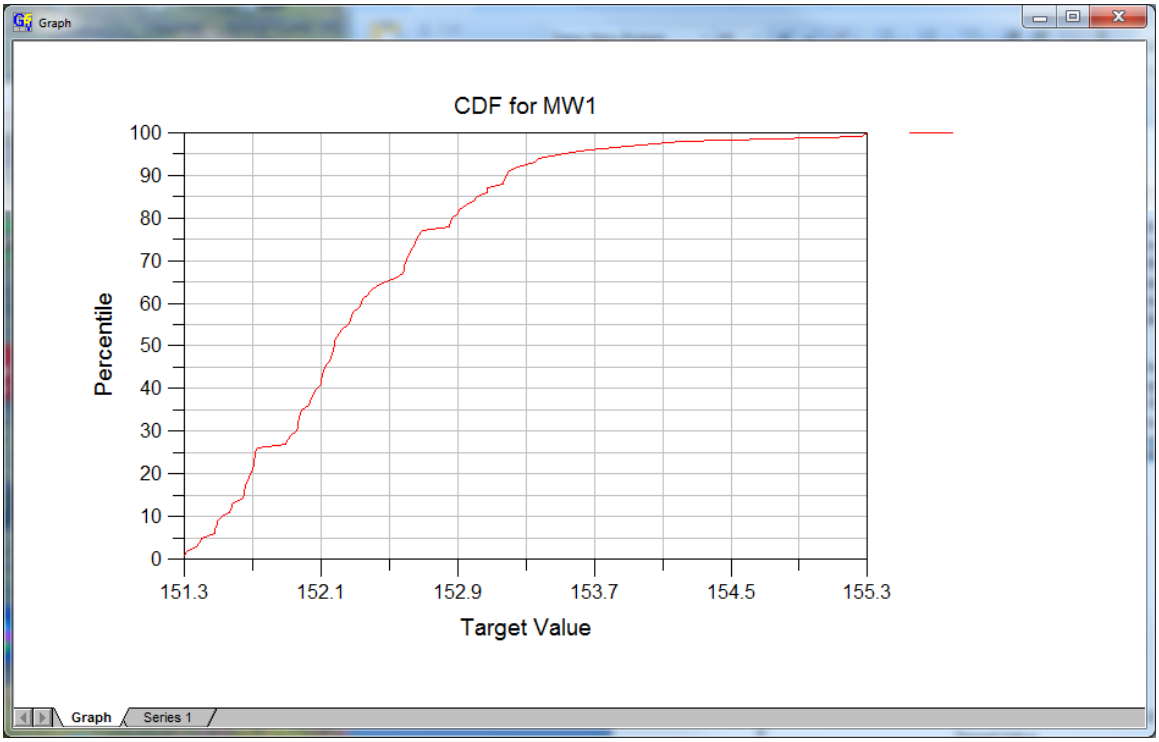

The CDF plot is interpreted by selecting a percentile (or probability) and stating that the there is a probability of X that the head value at this target is less than that value. In the example shown above, there is an 80 percent probability that the head at target MW1 will be less than 152.9 ft. The CCDF plot is the inverse of a CDF plot (i.e., the probability is that the value will be greater than a selected value). You create a CCDF plot just like a CDF plot but when you select the target you check the option for creating the CCDF plot as described in the last chapter.

Stochastic MODPATH

Stochastic MODPATH has only one use in the current version. It computes the probability that a particle will be captured by one or more boundary cells. These boundary cells are identified using a unique reach number. For example, to compute the probability of capture of three wells in a wellfield, you would first give these three wells a unique reach number in Groundwater Vistas. No other boundary cell (including other boundary types like drains and rivers) can have this reach number. This technique is ideal for wellhead protection studies and for evaluating the probability of failure of a pump & treat system.

We will now run *Stochastic* MODPATH to determine the probability of capture for the two wells in layer 3. These wells were previously assigned a reach number of 2 to distinguish them from the constant head cells in layer 1. Before running the Monte Carlo MODPATH simulation, select **Model|Stochastic|MODPATH Options** and enter a 2 for boundary reach and change the particle option to *One Particle Per Cell*. Your screen should look like the one below.

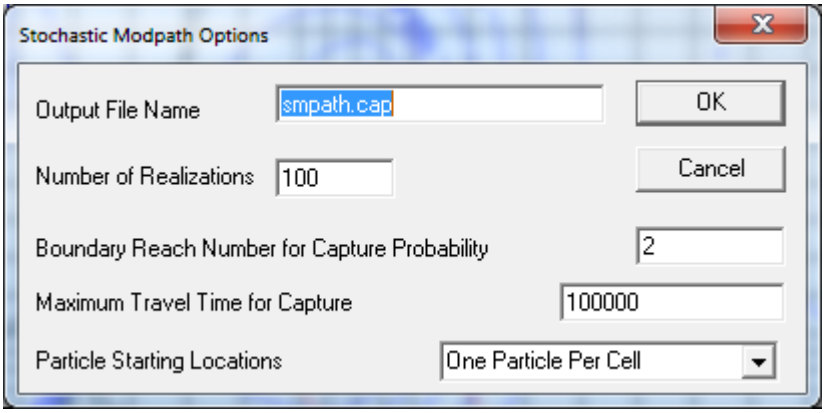

When you run *Stochastic* MODPATH, the probability of capture for each particle will be written to the file *smpath.cap*. After clicking OK on this dialog, select **Model|Stochastic|Create MODPATH Data Sets**. Now run the model by selecting **Model|Stochastic|Run Stochastic MODPATH.**

After MODPATH is done with the 100 realizations, select **Plot|Stochastic|Convert Capture to HDS**. First you identify the capture file (default is c:\gwv6\tutorial\work\smpath.cap but you may need to browse to the correct directory) in a standard file open dialog. Next, you will see a file save dialog where you enter the name of the head-save format file that will contain capture probability. Enter *capture.hds* here. Now select **Plot|Import Results** and browse to find capture.hds for the head file. Your screen should look like the one below.

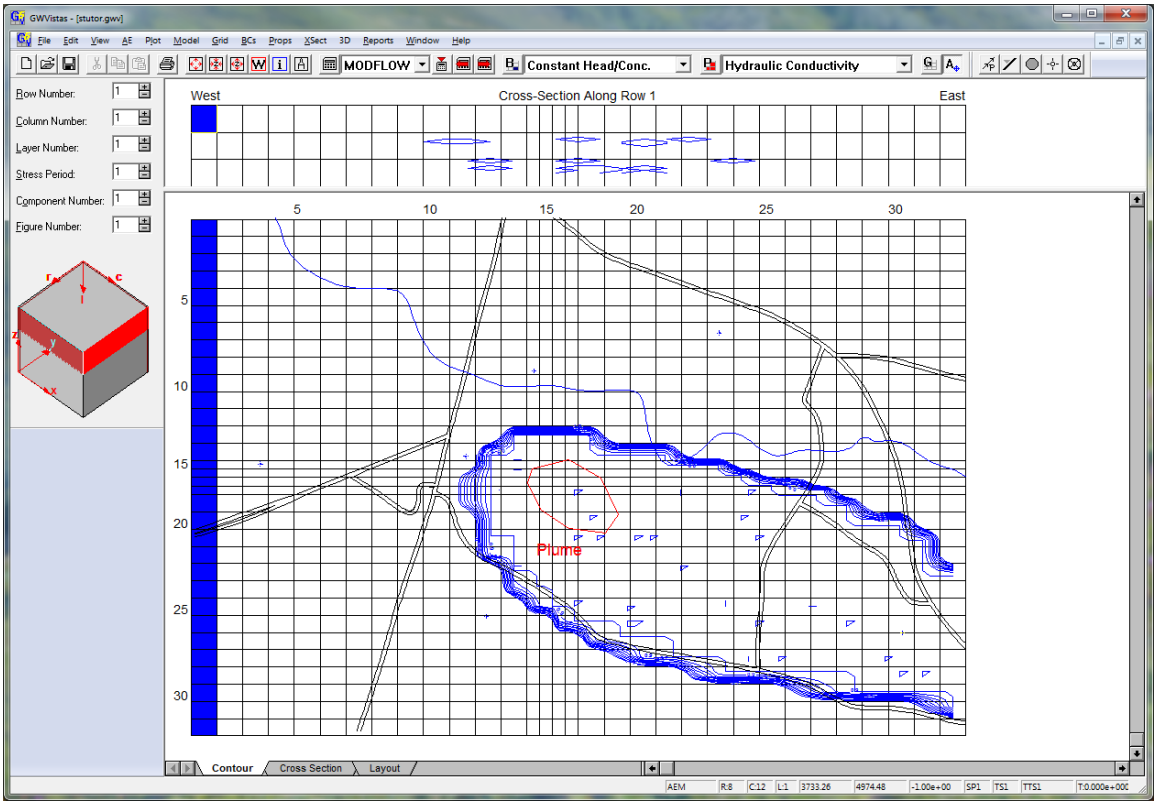

This plot is not very interesting because there is not a lot of variability in the capture probability for this model. The example does, however, show how you can get a capture probability map for any problem. Another useful technique is to use the color flood option (**Plot|What to Display**) for capture probability. Your might try it in this case just to see how it is done.

Stochastic MT3D

The third Monte Carlo model we will explore is MT3D. The example model has been configured with a second recharge zone that injects contaminant in the recharge water. This is supposed to simulate a leaking landfill or impoundment. We will add this parameter to the list of uncertain variables and then run *Stochastic MT3D*. Start by selecting **Model|Stochastic|Parameters**. Enter a second parameter as recharge concentration with a log normal distribution, a standard deviation of 1.0 (e.g., one order of magnitude), a minimum value of 1 and a maximum value of 10000. The mean value is the concentration currently in zone 2 which is 100. After entering this data, your screen should look like the one below.

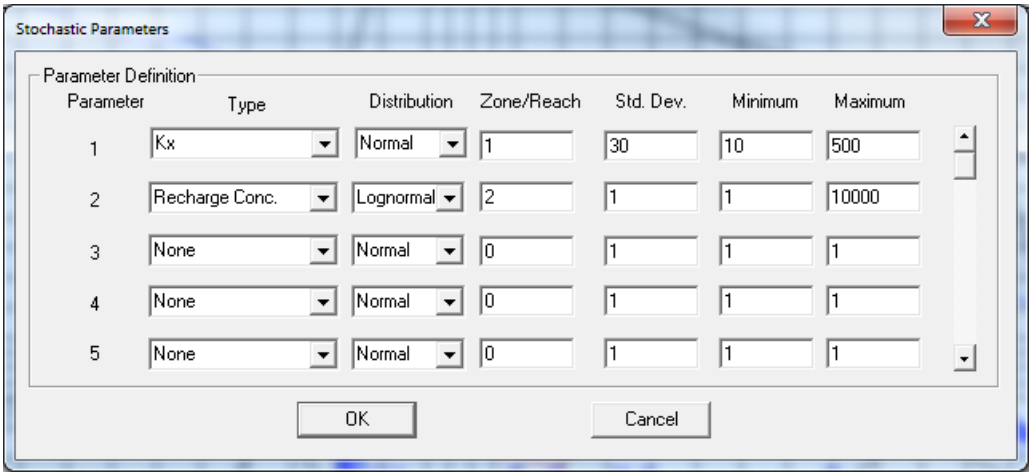

One other important concept to remember is that *Stochastic* MT3D does not use the observation well file created by the standard version of MT3D. Instead, GV reads data for observation wells and targets from the binary concentration file (ucn file). Because of this requirement, you must explicitly define what times to save concentrations in the binary file. This is done by selecting **Model|MT3D|General Options** and click on the **Printing** tab. In this example, we are only interested in the concentrations at the end of the run so make sure that the frequency of output is 1. Click the *Print Times* button and confirm that the printout time is 1000 days. When you run *Stochastic* MT3D on your own models, be sure to include as many times as you need in the binary file but not too many. These files can be quite large.

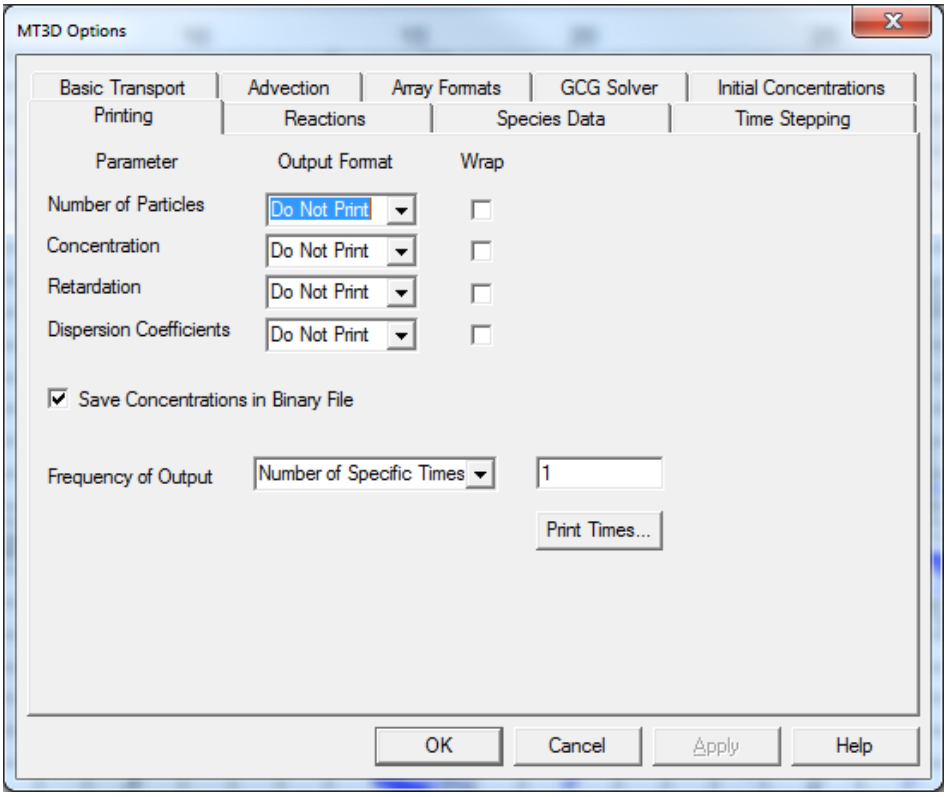

We now create MT3D data files by selecting **Model|Stochastic|Create MT3D Data Sets**. Even though this example is small, it may take an hour for *Stochastic* MT3D to run 100 realizations. Select **Model|Stochastic|Run Stochastic MT3D.**

Processing the results is similar to *Stochastic* MODFLOW. First, compute the mean and standard deviation for concentration at 1000 days. Select **Plot|Stochastic|Compute Mean & Std**. Fill in the dialog as shown below.

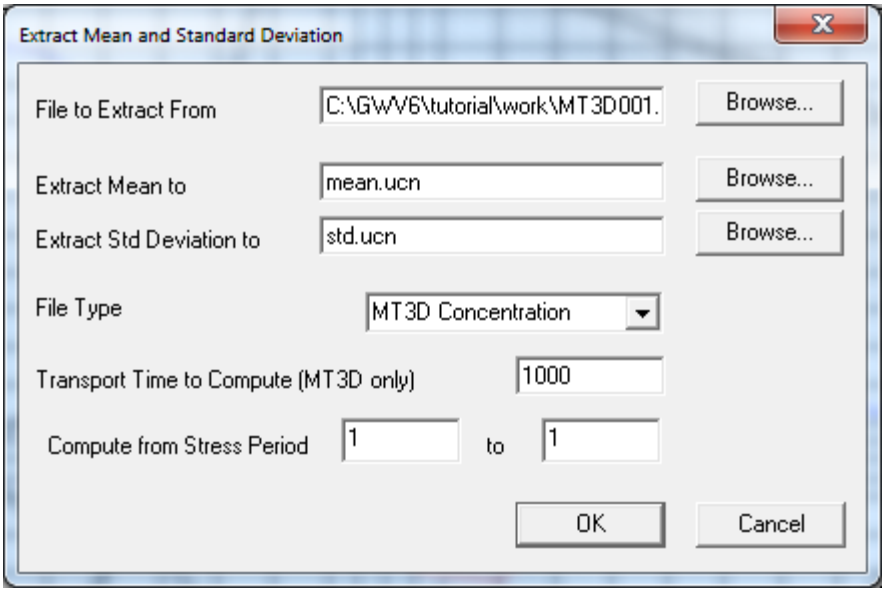

Be sure to specify the transport time of 1000 and that the file type is an *MT3D Concentration* file. You can now import the mean concentration and mean heads by selecting **Plot|Import Results**. The mean contaminant plume should look like the one shown below.

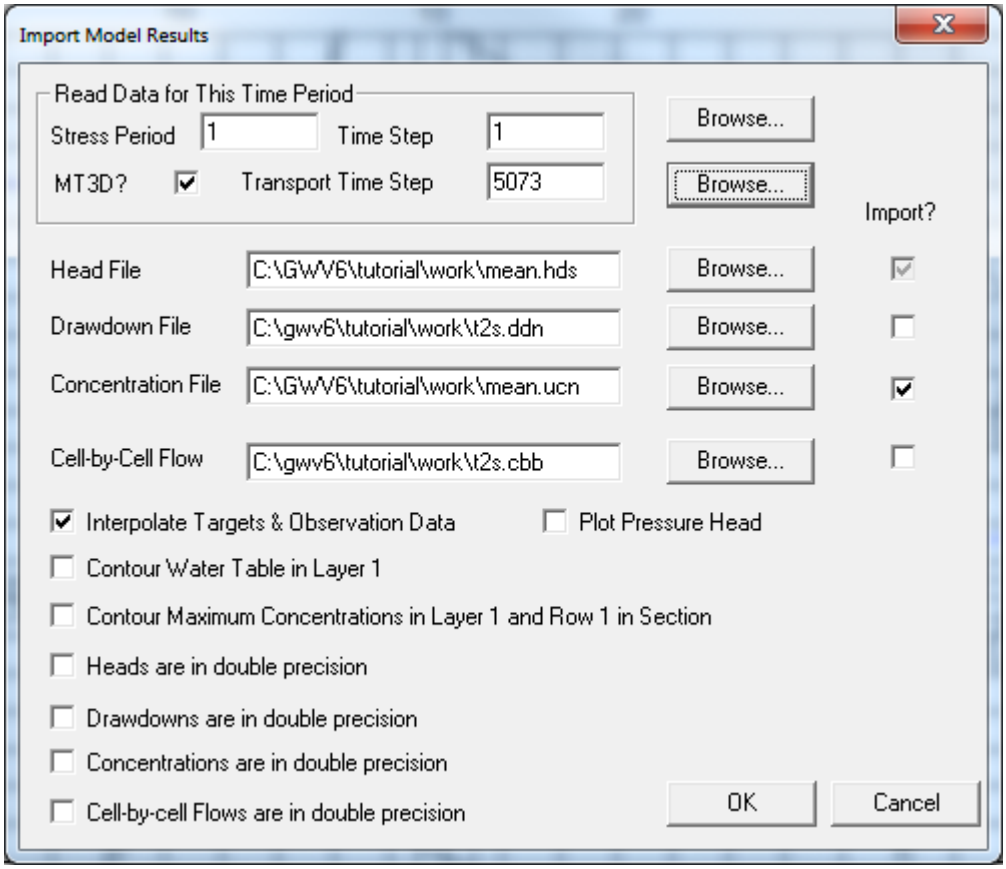
Import Model Results
Read Data for This Time Period
Stress Period
1
Time Step
1
MT3D?
Transport Time Step
5073
Browse...
Browse...
Import?
Head File
C:\GWV6\tutorial\work\mean.hds
Browse...
Drawdown File
C:\gwv6\tutorial\work\t2s.ddn
Browse...
Concentration File
C:\GWV6\tutorial\work\mean.ucn
Browse...
Cell-by-Cell Flow
C:\gwv6\tutorial\work\t2s.cbb
Browse...
Interpolate Targets & Observation Data
Plot Pressure Head
Contour Water Table in Layer 1
Contour Maximum Concentrations in Layer 1 and Row 1 in Section
Heads are in double precision
Drawdowns are in double precision
Concentrations are in double precision
Cell-by-cell Flows are in double precision
OK
Cancel

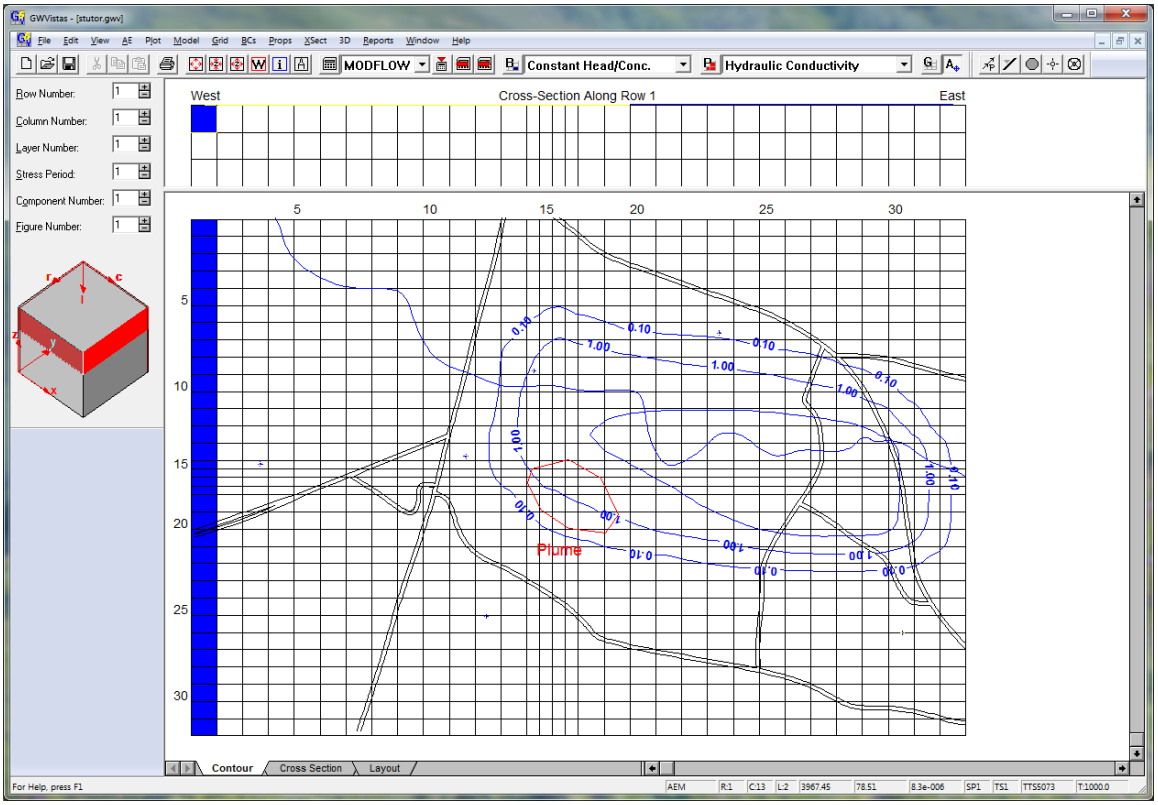
Constant Head/Conc.
Hydraulic Conductivity
Row Number:
Column Number:
Layer Number:
Stress Period:
Component Number:
Figure Number:
West
Cross-Section Along Row 1
East
Plume
Contour
Cross Section
Layout

One common plot you will want to make is a CCDF plot showing the probability that concentrations will exceed a certain value in a receptor well. First, we will add an observation well near the southern-most pumping well in layer 3. Make sure you are viewing layer 3 and that the **A** button (analytic elements) is pushed down on the toolbar and then select **Add|Well**. Move the cursor to the southern well and click the left mouse button. On the dialog, check the option to *Monitor Head|Concentration* and set the pumping rate to zero. You must also set the top and bottom screen layers to 3. Your dialog should look like the one shown below.

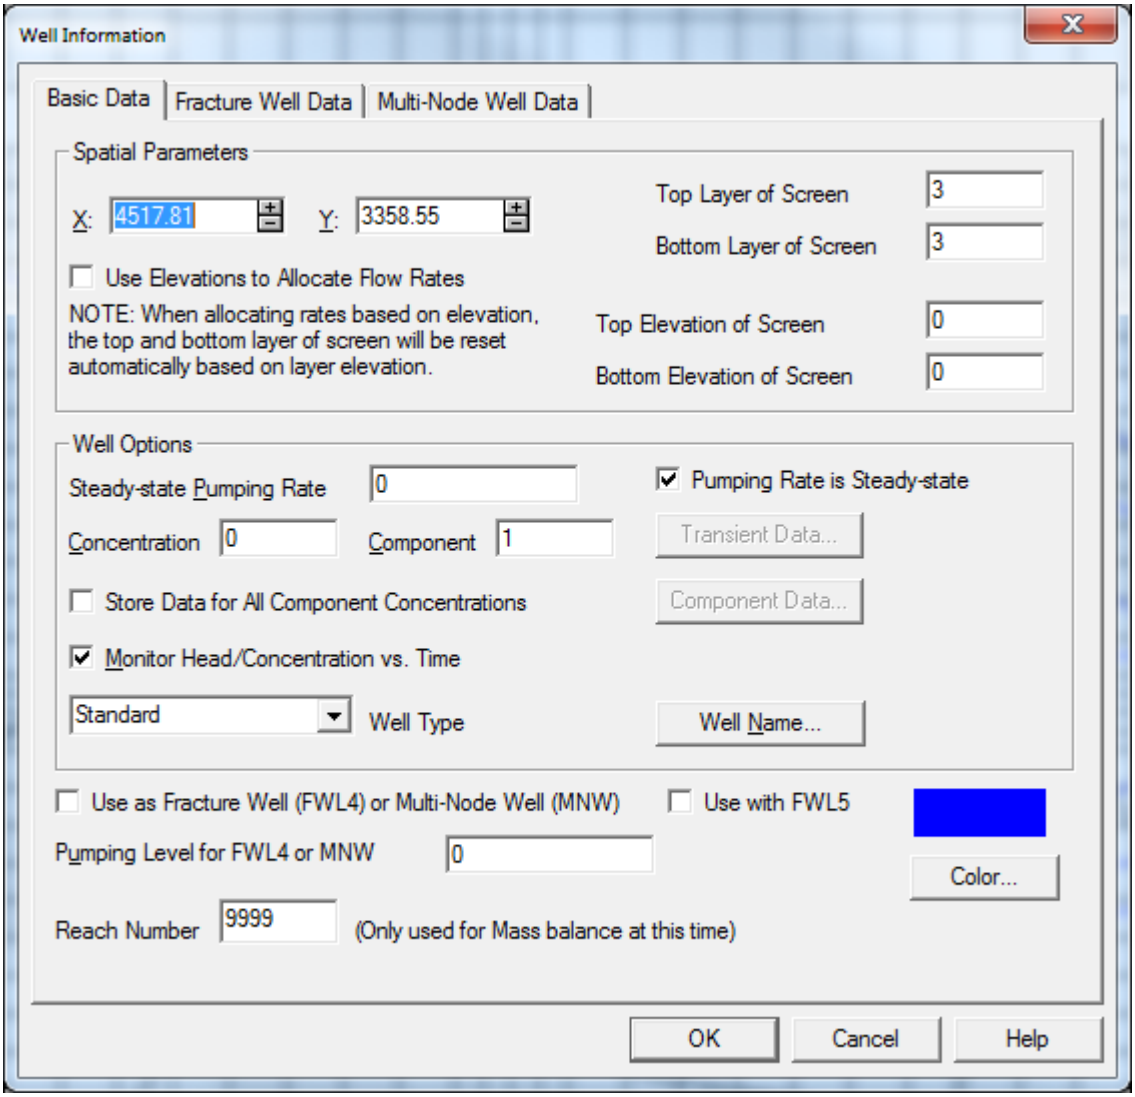

Now, we must import concentrations for observation wells. Select **Plot|Stochastic|Import Observation Data**. Browse to find the *mt3d.ucn* file from the Monte Carlo run and specify that the file type is concentration as shown below.

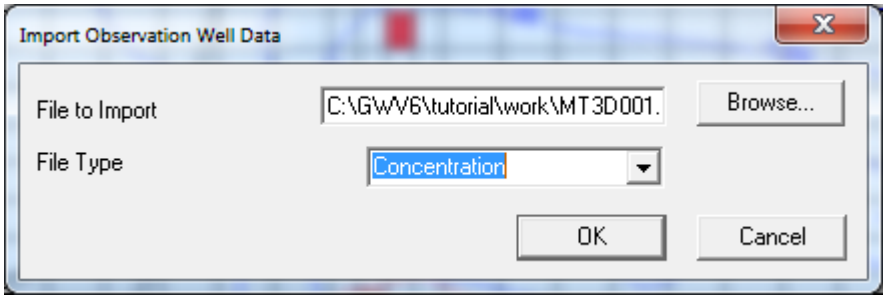

GV will report that 100 realizations have been imported. Now select **Plot|Stochastic|Graph|Observation Well CDF**. Select the one observation well, select concentration as the type, and check the option for a CCDF plot as shown below.

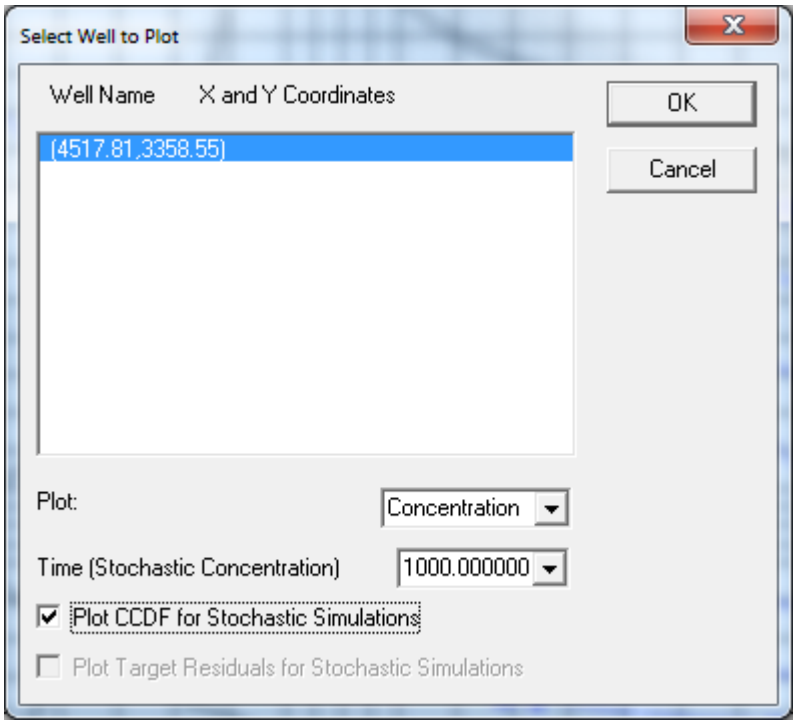

Your plot should look like the following.

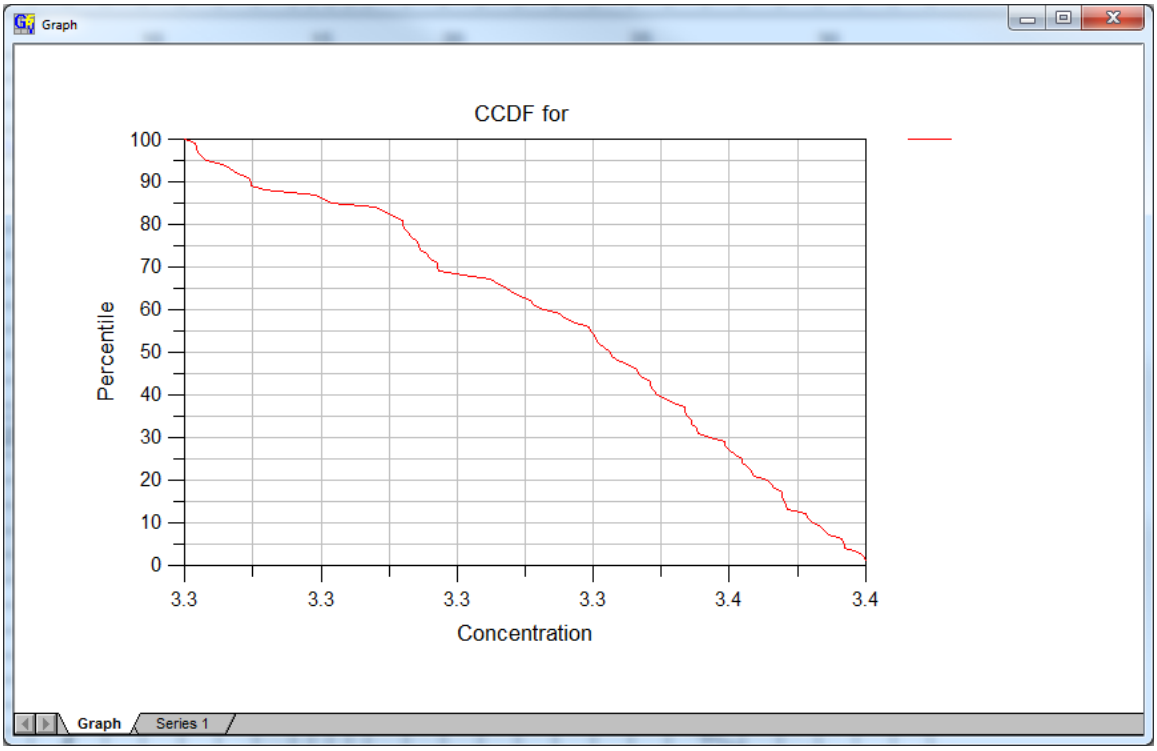

The plot indicates that there is a 50 percent probability that concentrations will exceed 3.35 ppb.

This concludes the tutorial for the Monte Carlo versions of MODFLOW, MT3D, and MODPATH. You have now seen most of the capabilities of these powerful simulators. You should now be ready to try out some stochastic simulation on your own models!

Using GIS Data in Groundwater Vistas

Building the Model

This exercise will build a more realistic model than the simple examples in the last session. The focus of this section will be on integrating Geographic Information Systems (GIS) with the MODFLOW model. Additional features for model design within Groundwater Vistas will also be explored. The model built in this section will be used in the next several sessions for particle-tracking, contaminant transport, mine dewatering, and analysis of water supply scenarios.

The model consists of a bedrock aquifer overlain in some areas by an alluvial aquifer system. The aquifer extends from a topographic divide to a large river. A three-dimensional depiction of the model area is shown below.

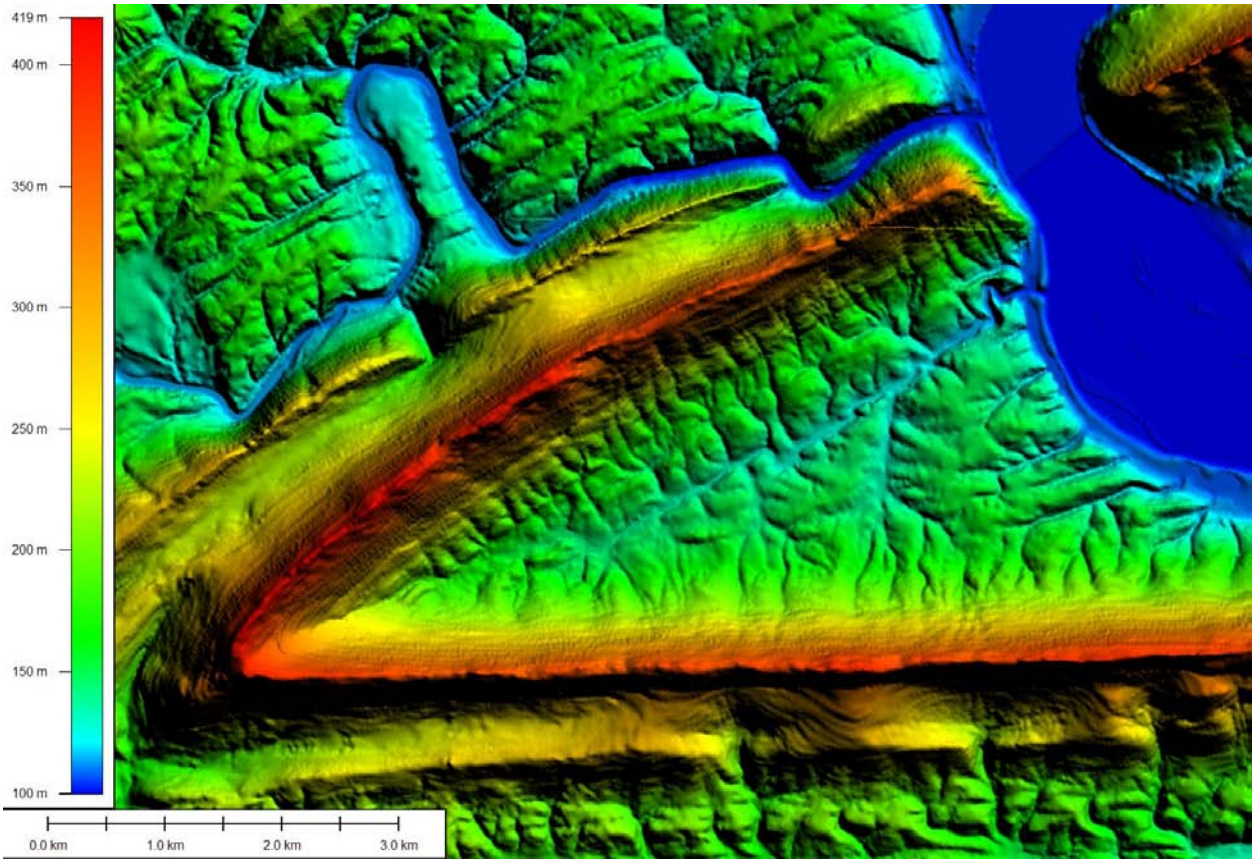

(Note: the picture above was produced using GlobalMapper and digital elevation model files)

The alluvial aquifer occupies the areas near the stream channels in the low-lying parts of the model. Groundwater flows generally from the topographic divide to the streams and also towards the major river to the east.

A GIS is a very convenient way of constructing the various components of the model design, including topography, hydrologic features (rivers and streams), geologic units (e.g. alluvial units in this case), well locations, model extent, etc. The simple GIS for this model is shown below.

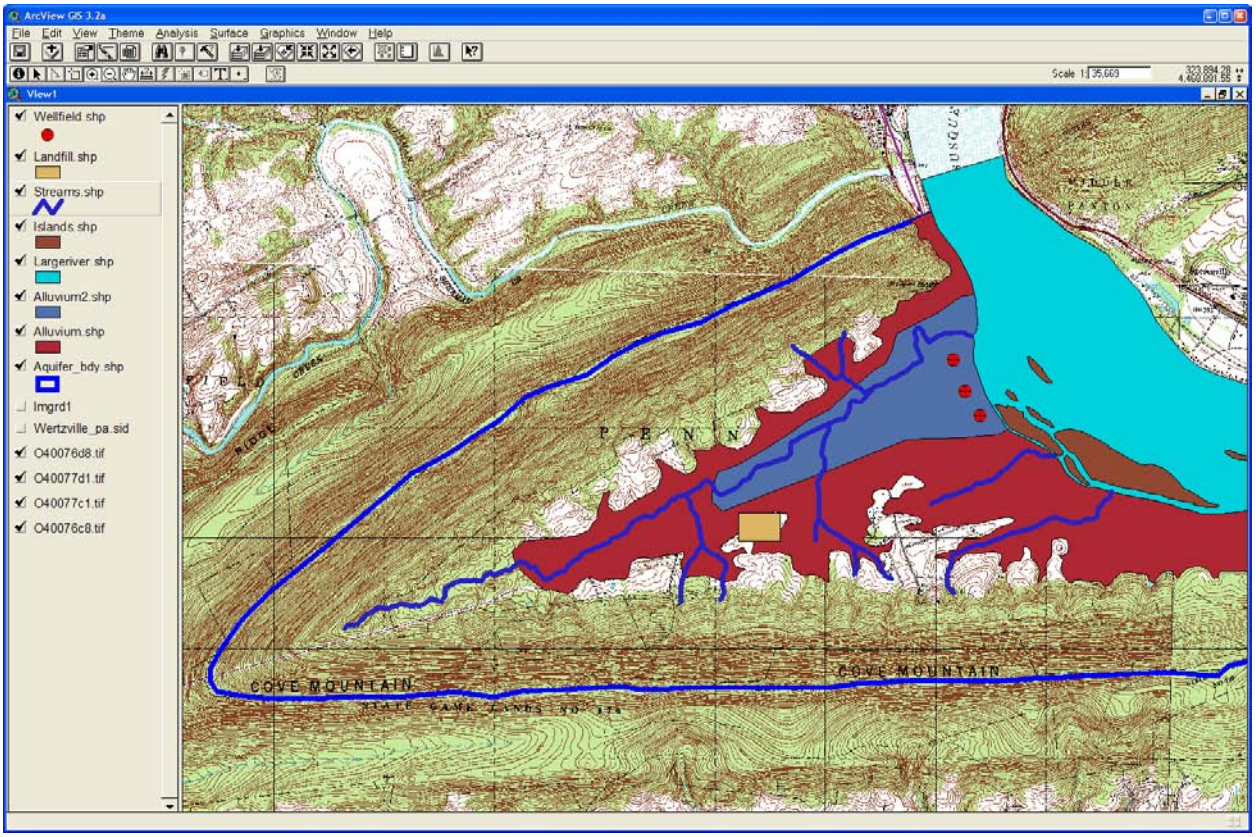

The interface between the GIS and the model design is through a series of shapefiles shown in the GIS legend above. The files we will use for this model include:

Wellfield.shp	Pumping well locations (points)
Landfill.shp	Location of a leaking landfill (polygon)
Streams.shp	Polylines of Tributary streams to the main river
Largeriver.shp	Polygon of the major river to the east
Islands.shp	Location of islands in the major river
Alluvium.shp	Shallow alluvial unit
Alluvium2.shp	Deep alluvial unit
Aquifer_bdy.shp	Boundary of the aquifer (polygon)
Surface_topography.grd	Surfer grid file of surface topography (not shown)

The only aspect of the model not provided in the files listed above is the design of the finite-difference grid. Everything else about this model will come from the GIS. All of these shapefiles were digitized in the GIS from aerial photography and topographic maps in a relatively short time. In many cases, though, such data can be obtained from public sources and modified to suit the needs of the model.

The topographic data came from publicly available digital elevation models (DEM). The program ***GlobalMapper*** (from www.globalmapper.com) was used to stitch together several DEMs of the area and then export them to SURFER format, which Groundwater Vistas can read. The GlobalMapper software is very easy to use and inexpensive. It is very useful at converting between many different file formats and coordinate systems.

Starting the New Model

Start the new model by selecting **File|New** in Groundwater Vistas. Change the number of rows to 50 and the number of columns to 150. Keep the grid spacing for both rows and columns as 100 m. We will use 3 layers in the model. Change the default hydraulic conductivity (Kx and Ky) to 1 m|d. Set the vertical hydraulic conductivity (Kz) to 0.01 m|d. Enter a recharge rate of 0.001 m|d. Your screen should look like this one.

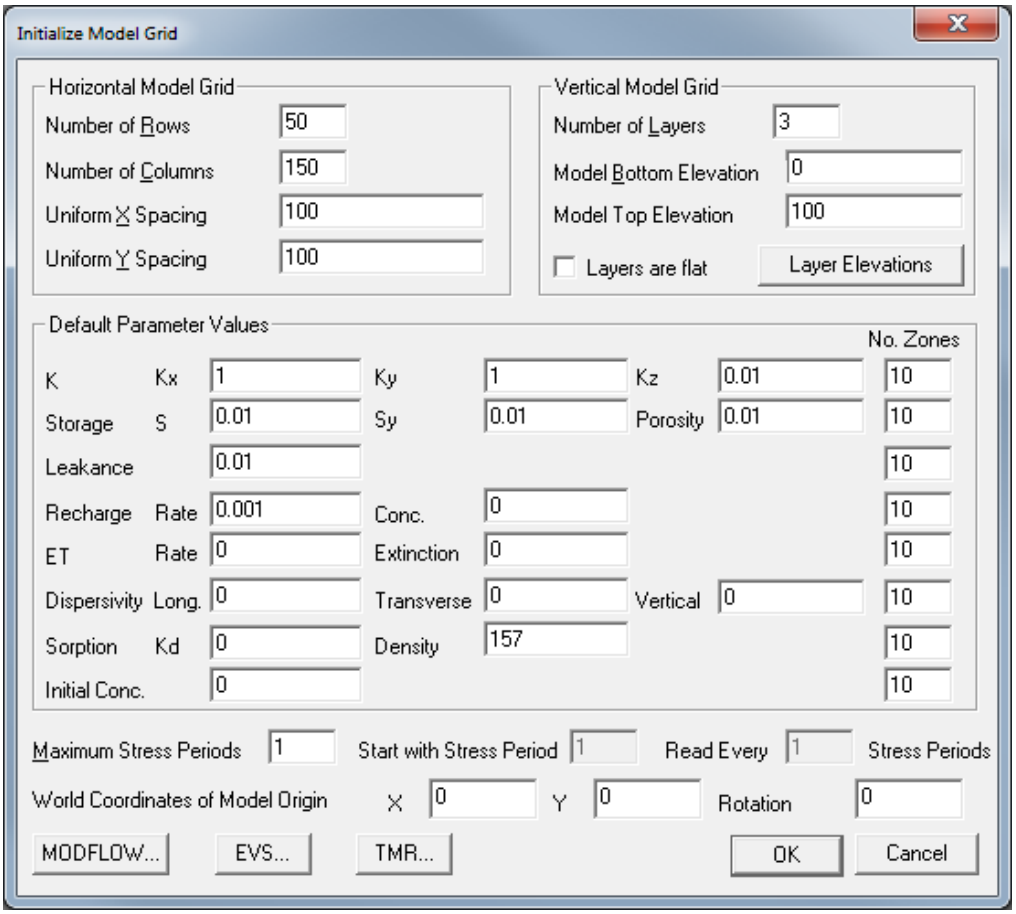

Now, to register the model with a real-world coordinate system, we will import some of the shapefiles listed above for base maps. The first will be the model boundary. Select **File|Map|Shapefile** and browse to find the shapefile called ***aquifer_bdy.shp*** located in the c:\gwv6\tutorial directory. After selecting this file, another dialog is displayed asking you for a name for the map file. GV will read the shapefile and convert it to a

Groundwater Vistas map file. Use the name ***aquifer_bdy.map***. After importing this map, you model should look like the one below.

You may wonder how Groundwater Vistas figured out the mapping between the model grid and the real-world coordinate system. If you do not specify a non-zero value for the X and Y coordinates of the lower left corner on the initialization dialog (shown above), GV chooses the minimum X and Y coordinates on the first map file you overlay on the model as the real-world origin. In this case, the origin of the aquifer outline is the correct real-world origin for the model. If you create such an outline for each of your models, then matching your model to a world coordinate system will be relatively simple.

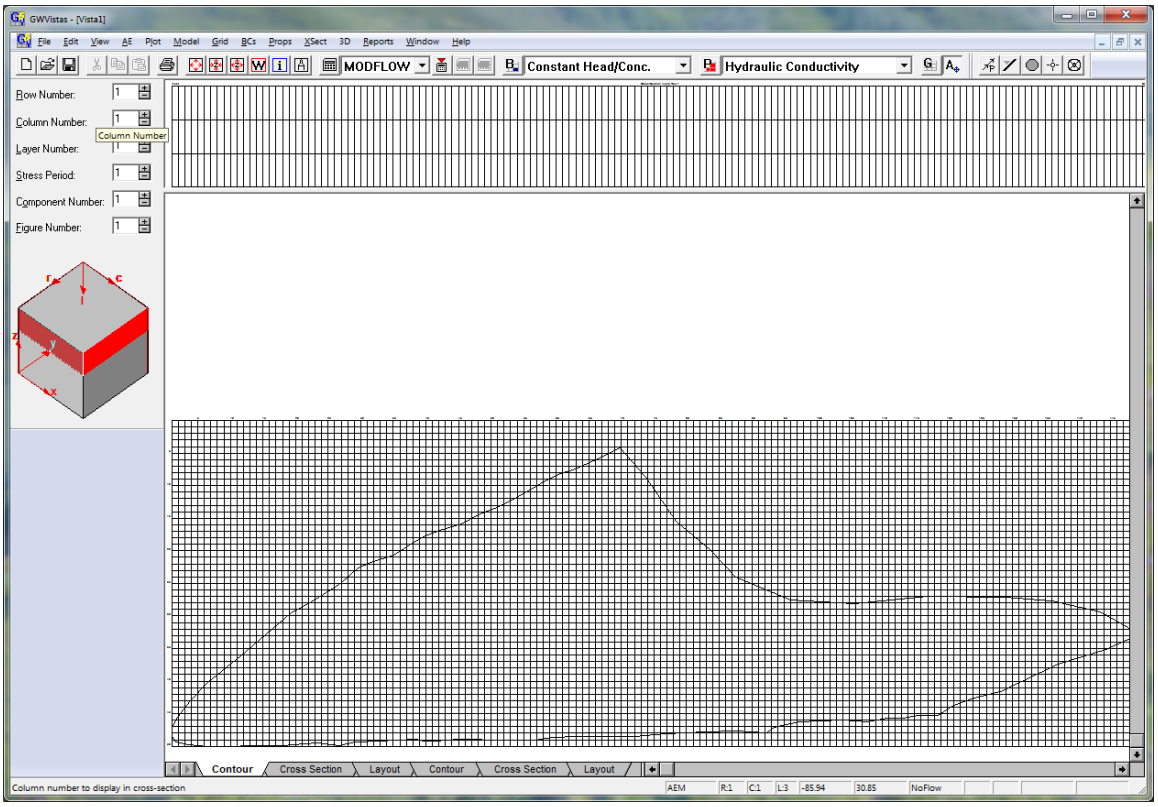

Note that there are several rows and columns that go beyond the edge of the model domain. In order to conserve memory, you will delete these rows and columns.

To delete the rows along the northern side of the model, there are two approaches. First option is to select **Grid|Delete|Row**, then move the cursor just below the top row in the model, and click the left mouse button. The second option is to press the **G** button on the toolbar and then press the **R** button that has a red "X" next to it on the toolbar (immediately to the right of the "A" button). Then move the cursor just below the top row and click the ***right*** mouse button. Use either method and repeat it 4 times.

Use the same basic approach to delete the two eastern-most columns.

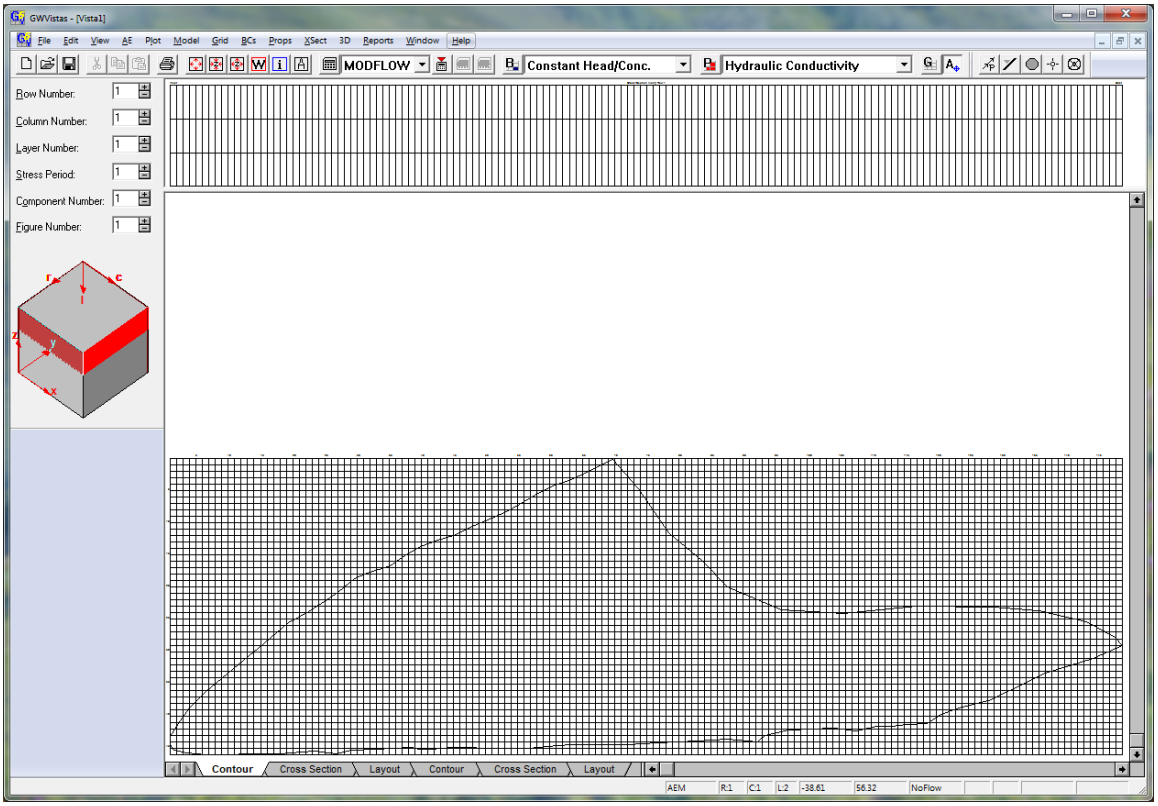

Importing Inactive Cells

The aquifer we are going to simulate lies within the polygon you imported as a base map. All cells outside that boundary should be inactive (i.e., no-flow cells). We will import the same shapefile we used as a map to automatically fill in those cells with no-flow boundary conditions. This is done by selecting **BCs|No-flow** and then **BCs|Import|Shapefile**. Browse to find the same shapefile (aquifer_bdy.shp). GV will then ask you "Set no-flow cells inside polygon?". Answer **NO** to this prompt and GV will make cells outside the polygon inactive. GV will then ask if you want to delete existing boundary conditions. Answer NO to this one as well (it actually makes no difference whether you say yes or no to this one as there are no other boundary conditions in the model).

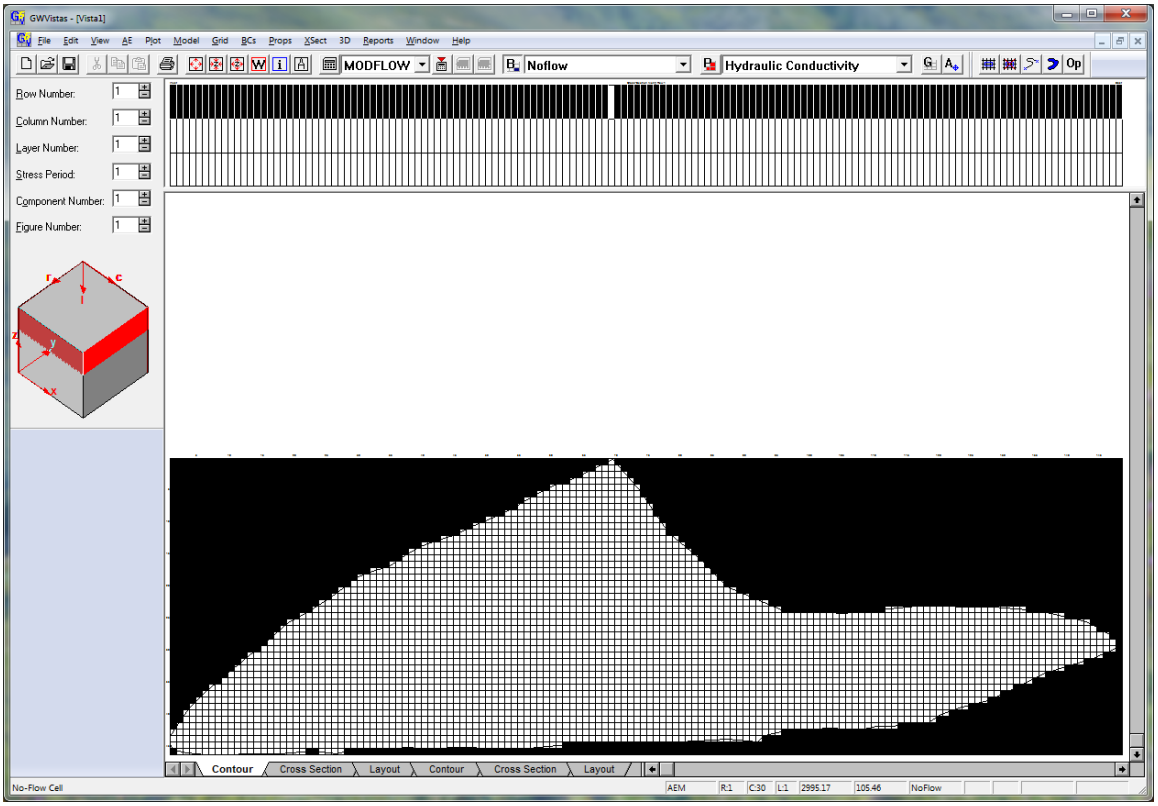

Remember that this is a 3-layer model. You imported the aquifer boundary shapefile into layer 1. You also need to do the same in layers 2 and 3. Another option, though, is to simply copy these no-flow boundary cells to the other layers. Go to layer 2 and select **BCs|Copy** and click OK. Repeat for layer 3.

Before you add any other features to the model, use the **File|Map|Shapefile** menu to display several additional maps. Import the shapefile on the left to create the map on the right. Your model should then look like the one below.

Landfill.shp	landfill.map
Streams.shp	streams.map
Largeriver.shp	river.map
Islands.shp	islands.map

[Note: the map files listed above are already in the tutorial directory so they can be used with the pre-built version of this model. If you get an error that these files cannot be overwritten, then just rename them from the ones listed above.]

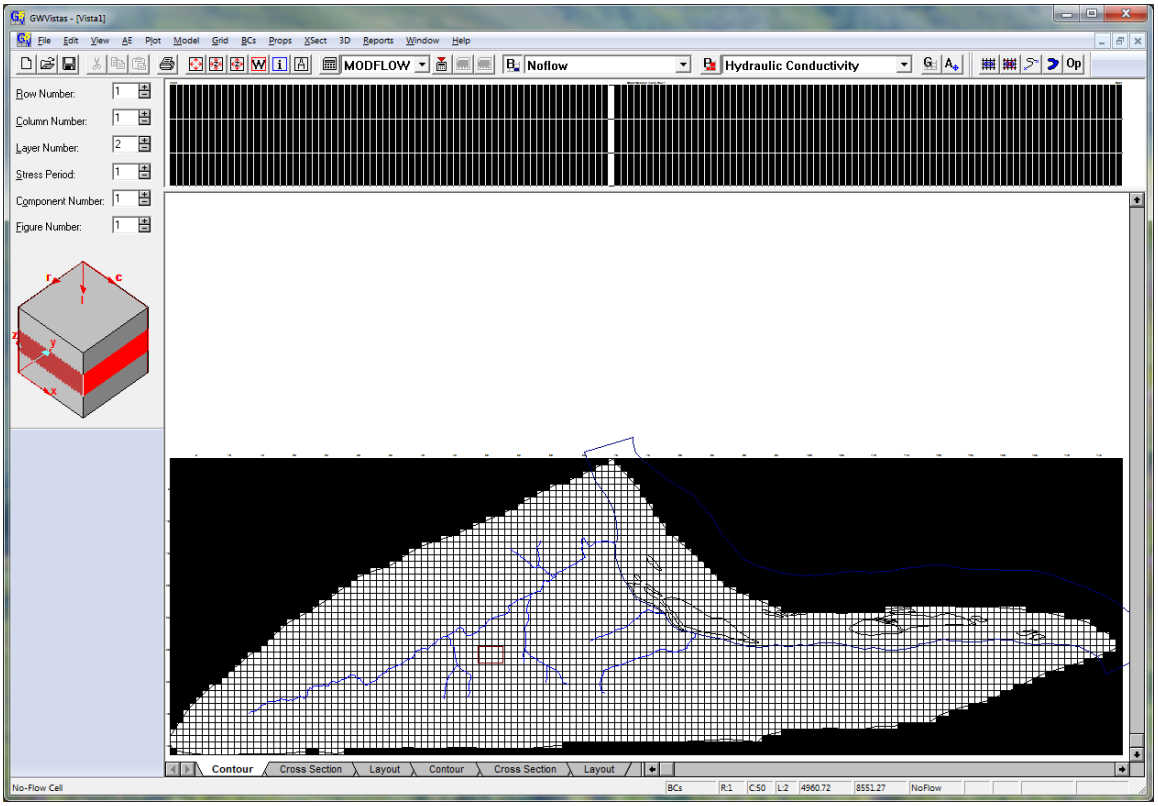

Refining the Grid

Before we add any additional boundary conditions or aquifer properties, we will add some additional rows and columns from the landfill area to the river. The easiest way to do this is to press the **G** button on the toolbar (for Grid mode) and then the **R** button (the one without the red x) to insert rows. Start with row 33 (just south of the landfill) and work your way to the north, ending with row 14. Move the cursor within the row and click the right mouse button. Repeat for each row.

Now insert new columns by pressing the **C** button on the toolbar and right-clicking from column 76 (western edge of the largest island in the river) to column 48 (just west of the landfill). After doing this, your model should look like the one below.

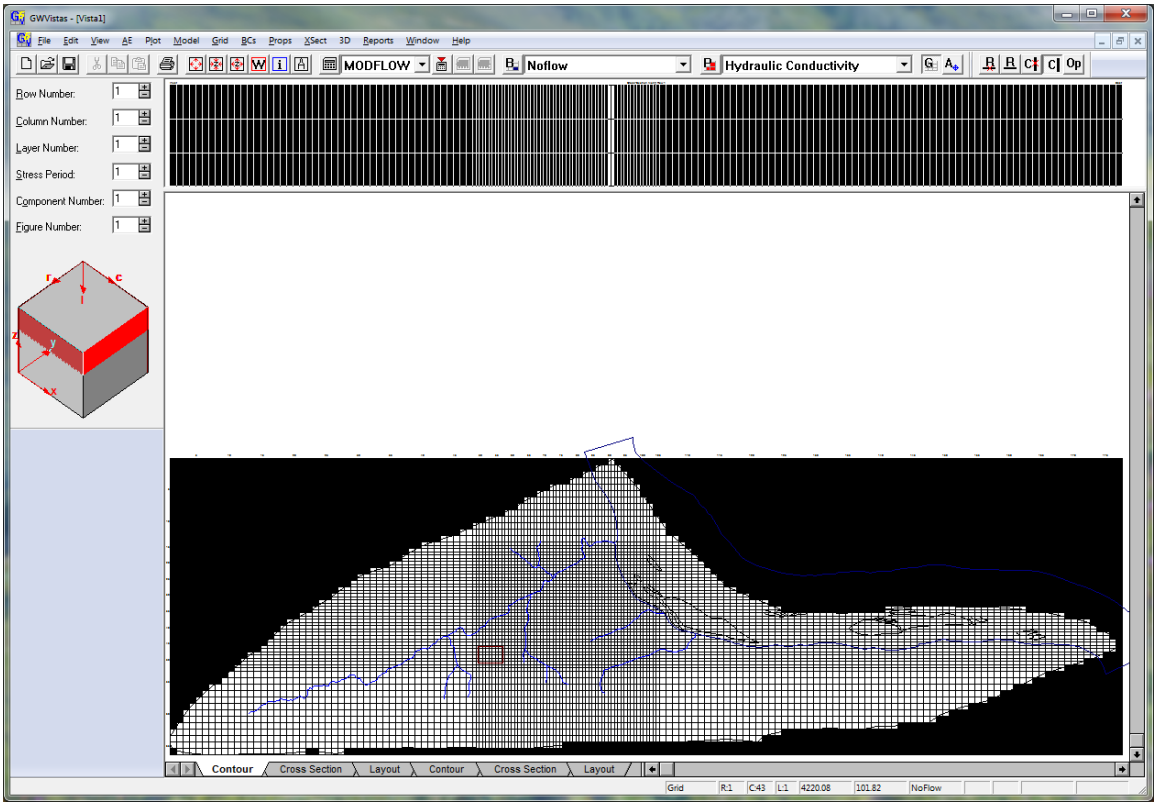

The "rule-of-thumb" in finite-difference modeling is that the change in grid spacing between adjacent rows and columns should not be larger than 50%. In the model we have designed so far, there is a jump of 100% at the boundary of the refined grid. To smooth this transition, select **Grid|Edit|Smooth Grid**. Click OK to keep the default transition of 1.5 (50%) between adjacent cells.

Now is a good time to save your work. Select **File|Save As** and give the model a name you will remember. In Version 6, Groundwater Vistas asks you if you want to change root file names whenever you use the "Save As" command. This allows you to automatically choose a new name for the simulation each time you save the model under a new name.

Defining Layer Elevations

The next step is to adjust the layer elevations based on the digital elevation model (DEM) in the file ***surface_topography.grd***. **First, make sure you are viewing layer 1.** Select **Props|Top Elevation** and then **Props|Import|SURFER** and browse to find this file in gwv6\tutorial. Check the option that the file contains site coordinates (e.g., the same coordinates as the base maps).

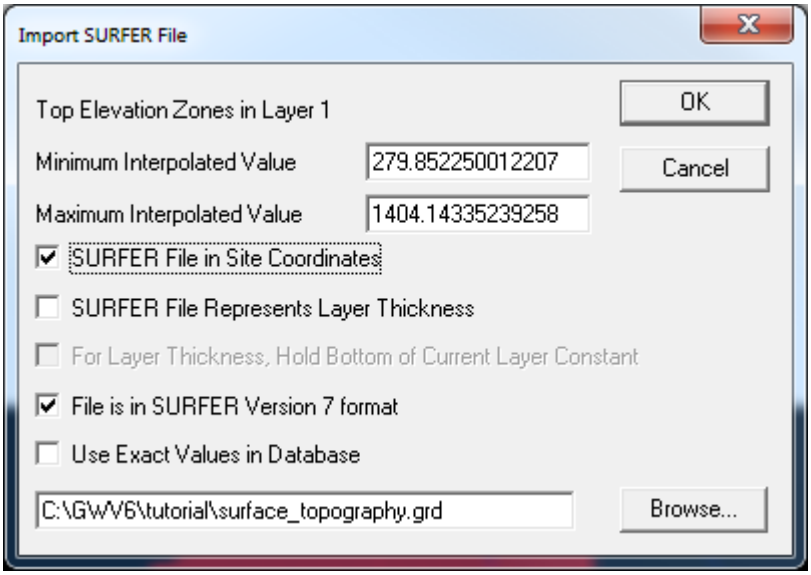

Click OK and GV will import the SURFER file and interpolate a top elevation for layer 1 from the DEM. The screen will probably appear a uniform red. To remap the colors, select **Props|Property Values|Reset Matrix Bounds.** Now the elevation colors should look like the following.

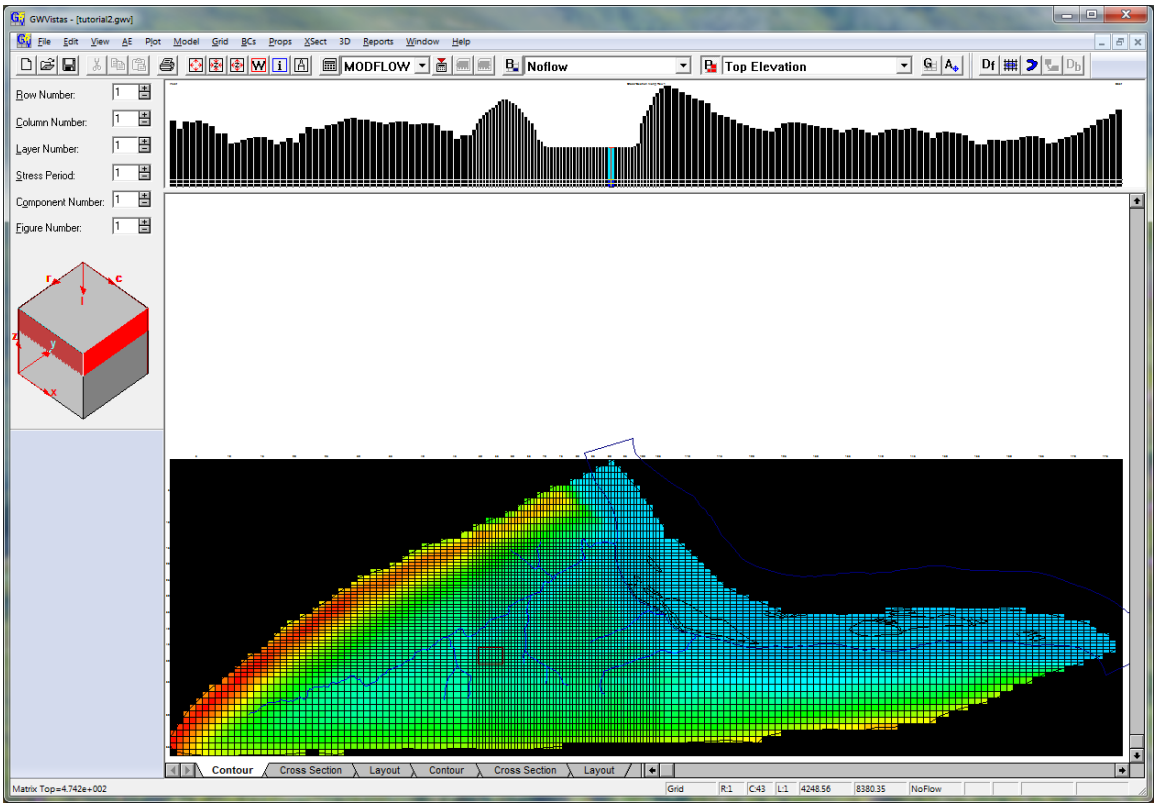

In a fully 3D model, you need to define the top elevation for layer 1 and then the bottom elevations for every layer. The DEM you just imported is the top of layer 1. There are many ways of defining layer elevations. You can import SURFER grid files, as we did above. The elevations can also be interpolated from known points by importing a text file of X, Y, and Z (elevation) data. GV also interfaces with 3D Visualization tools like EVS and Leapfrog Hydro, each of which can export data from a block 3D model to define layers in the MODFLOW model.

In the current model, we will be using uniform layer elevations for the bottoms of all 3 layers. The conceptual model in this fractured rock terrain is that hydraulic conductivity decreases with depth. There is no definable bottom to the aquifer – it's just that the hydraulic conductivity decreases to a point where it becomes negligible. That depth is at about 200 to 500 meters.

You will make the bottom elevation 250 m above sea level for layer 1, 200 meters for layer 2, and 100 meters for layer 3. Start by selecting **Props|Bottom Elevation**. Make sure you are viewing layer 1. To set the bottom elevation of layer 1 equal to 250 m, select **Props|Set Value or Zone|Clear** and enter 250 in the dialog. Repeat this for layers 2 and 3. Go to each of those layers and use the same command. The bottom elevation for layer 2 will be 200 m. The bottom elevation for layer 3 is 100 m.

Select **File|Save** to save your modifications.

Defining Boundary Conditions

Boundary conditions in this model consist of the small streams and the larger river to the east. The large river is a polygon in the GIS because it is so wide. The smaller streams are polylines because they are not very wide relative to a model grid cell. Polylines are also useful when a linear feature changes in water level significantly over its length. GV will interpolate river stage from the water level data entered into the GIS for both lines of each stream segment.

The only tricky part about defining these features is the desire to keep the islands in the large river as active cells. One way to do this is to temporarily make these islands no-flow cells and then delete the no-flow cells after importing the rivers.

Make sure you are viewing layer 1. Then start by selecting **BCs|No-flow** and then **BCs|Import|Shapefile**. Choose the ***islands.shp*** file in the gwv6\tutorial directory. When GV asks if you want to set no-flows inside the polygons, choose "Yes". The islands should now appear as no-flow cells.

Next, use **BCs|River** and **BCs|Import|Shapefile** and import the ***largeriver.shp*** file. The procedure when importing shapefiles with attribute data is to match up field names in the shapefile database with model features. The following dialog is displayed for the large river polygon. Change the field called "Exclude cells which have less than…" to 0.5.

That means that a river will be defined in a cell only if at least half the cell within the river polygon. Then change the river stage field to “WLE” (stands for water level elevation). Change the bottom elevation field to “BOTTOM”, thickness to “THICKNESS”, and hydraulic conductivity to “KV”, as shown below.

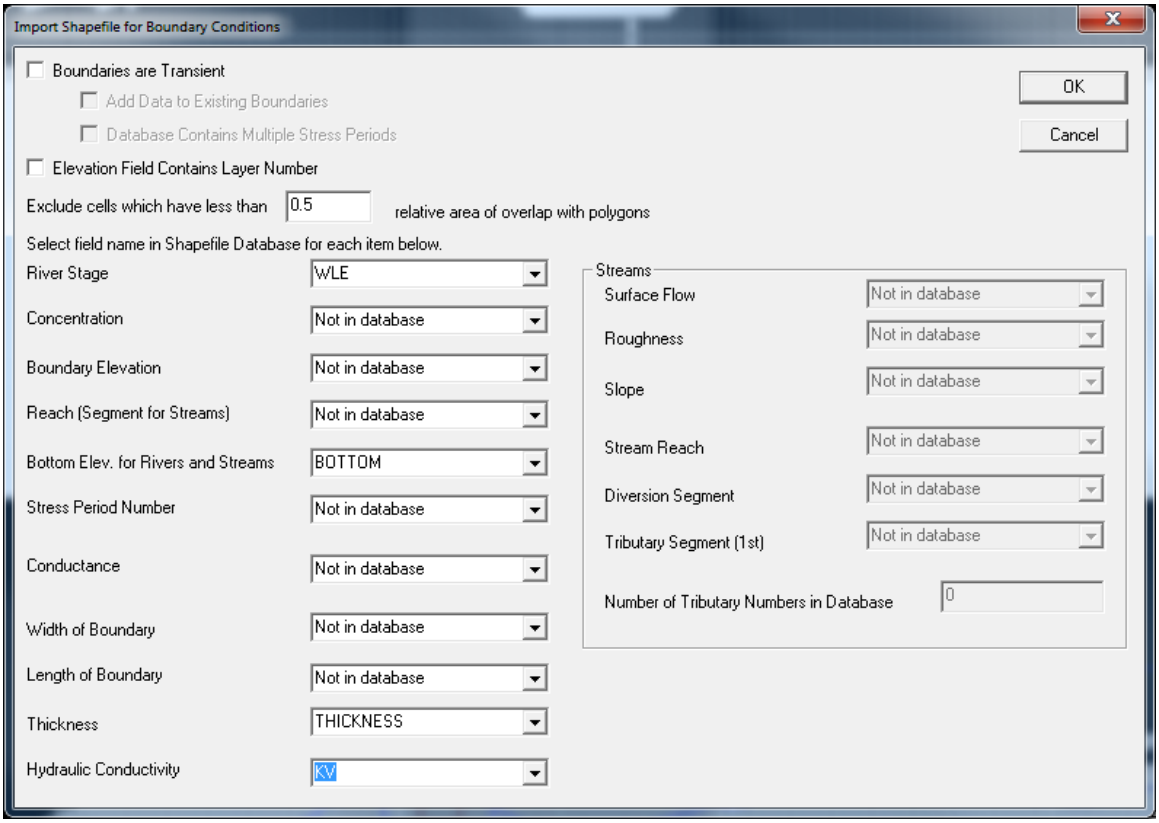

Click OK when done and the green river cells should appear within the area of the large river.

We need to convert the islands back to active cells. Select **BCs|No-Flow** and **BCs|Delete|Polygon**. Digitize a polygon around the islands. Digitizing is done by clicking the left mouse button to set points on the polygon and double-clicking to end the polygon. If you make a mistake, you can right-click the delete the last point you entered.

You can digitize a polygon that goes around all of the islands. Just make sure the polygon does not cross into any no-flow areas that you want to keep. After deleting these no-flow islands, your model should look like the one below.

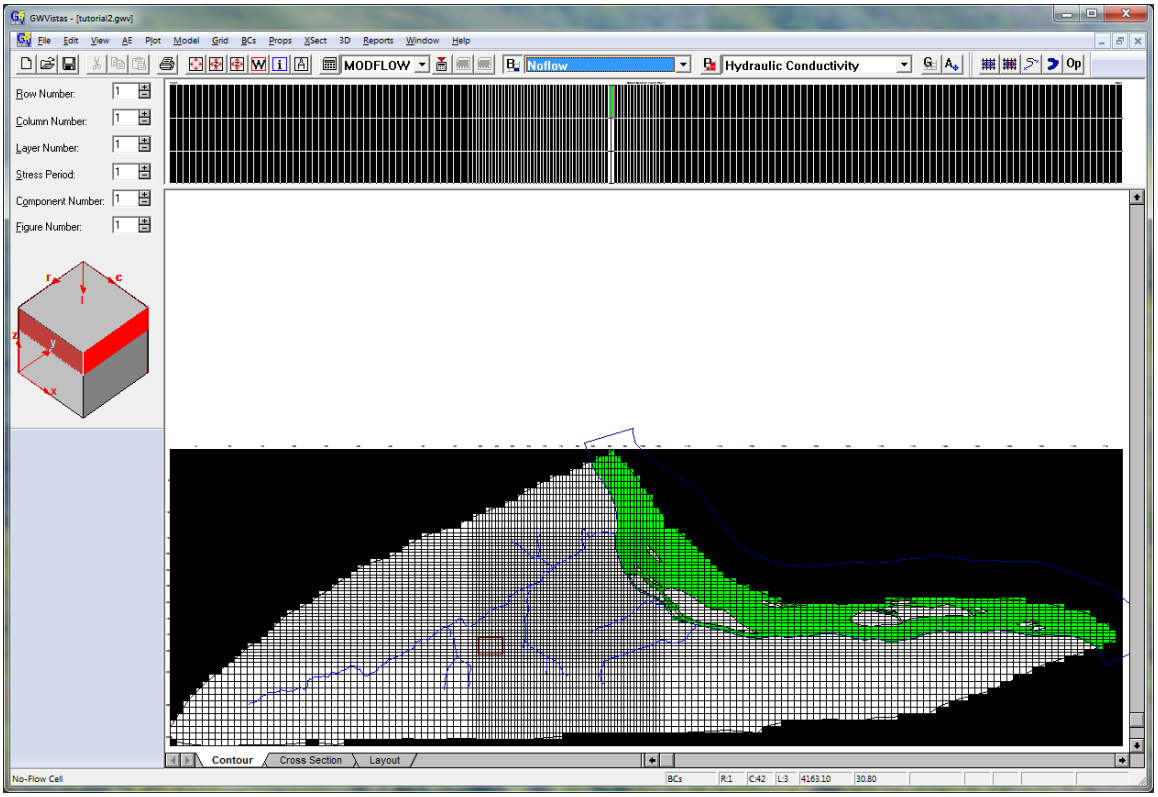

The smaller streams are defined by the ***streams.shp*** file. Select **BCs|River** and then **BCs|Import|Shapefile** and find that shapefile in the tutorial directory. GV will ask you whether data have been defined for each end of the polylines. Answer Yes to that question. Change the field names as shown below.

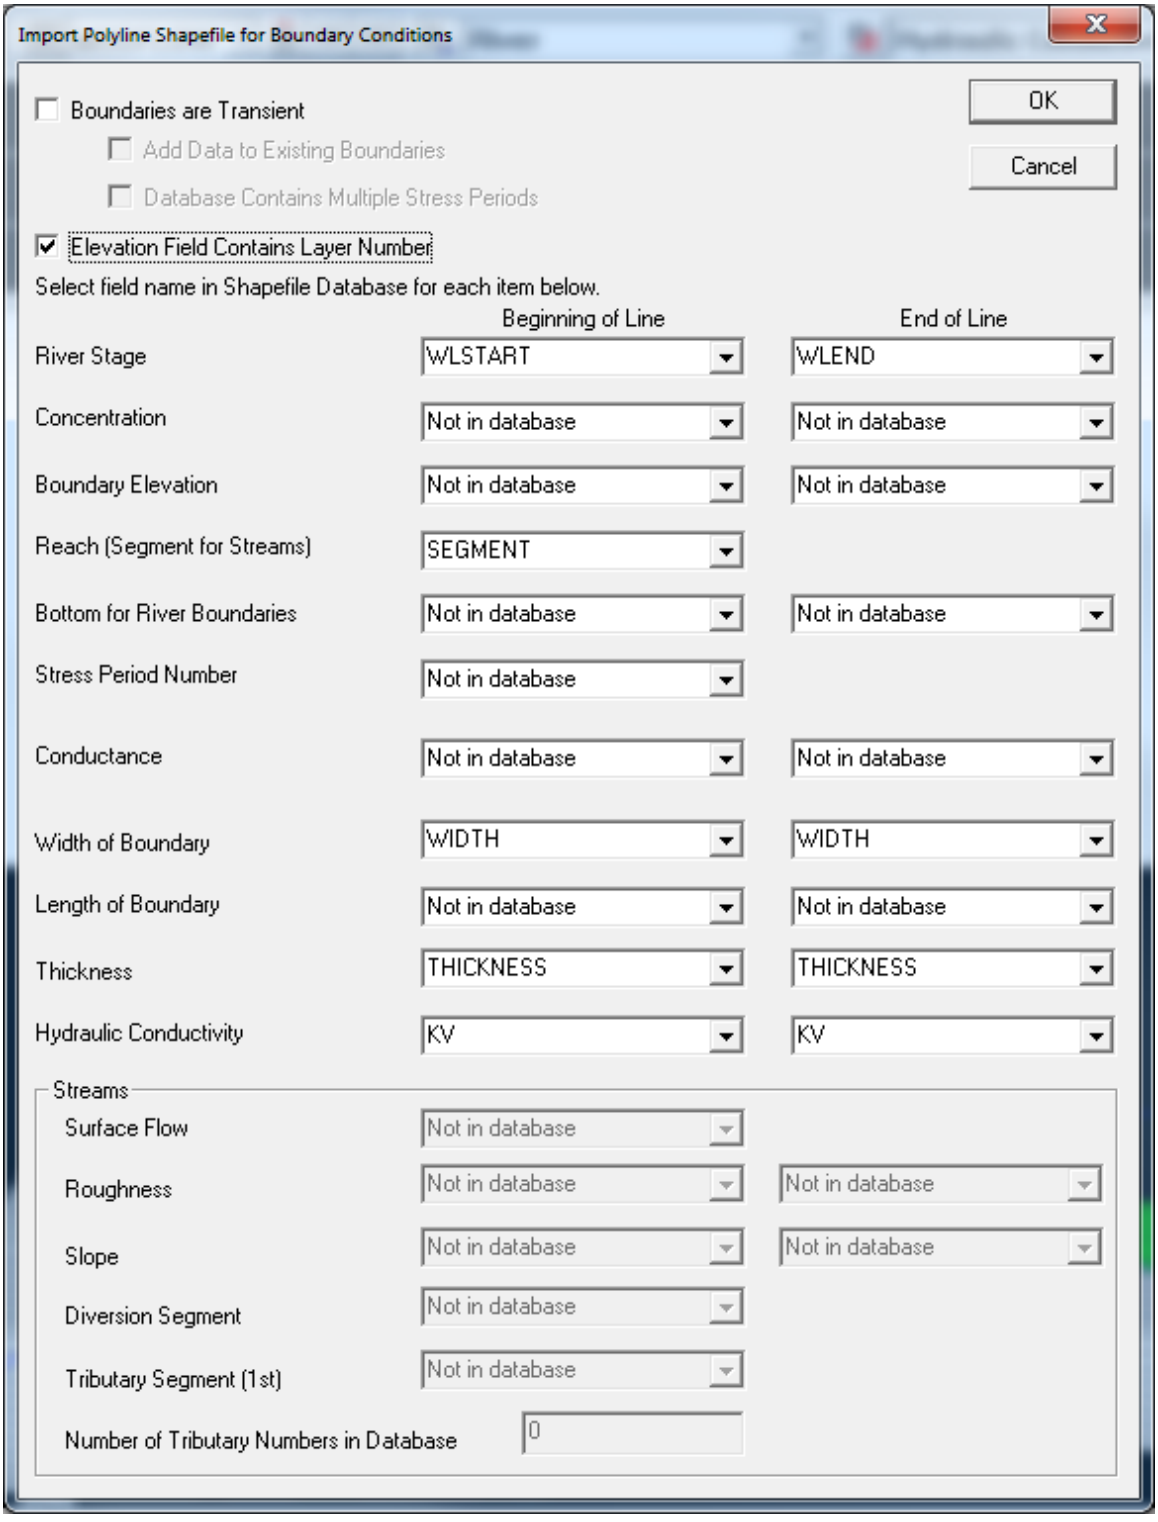

You may wonder why the length field is not defined above while the other elements of the conductance calculation (width, Kv, bed thickness) are defined. When importing polylines, GV will compute the length for you if you do not have it in the database.

Similarly for polygons like the large river imported in the previous step, both the length and width are computed by GV. Your model should now look like the one below.

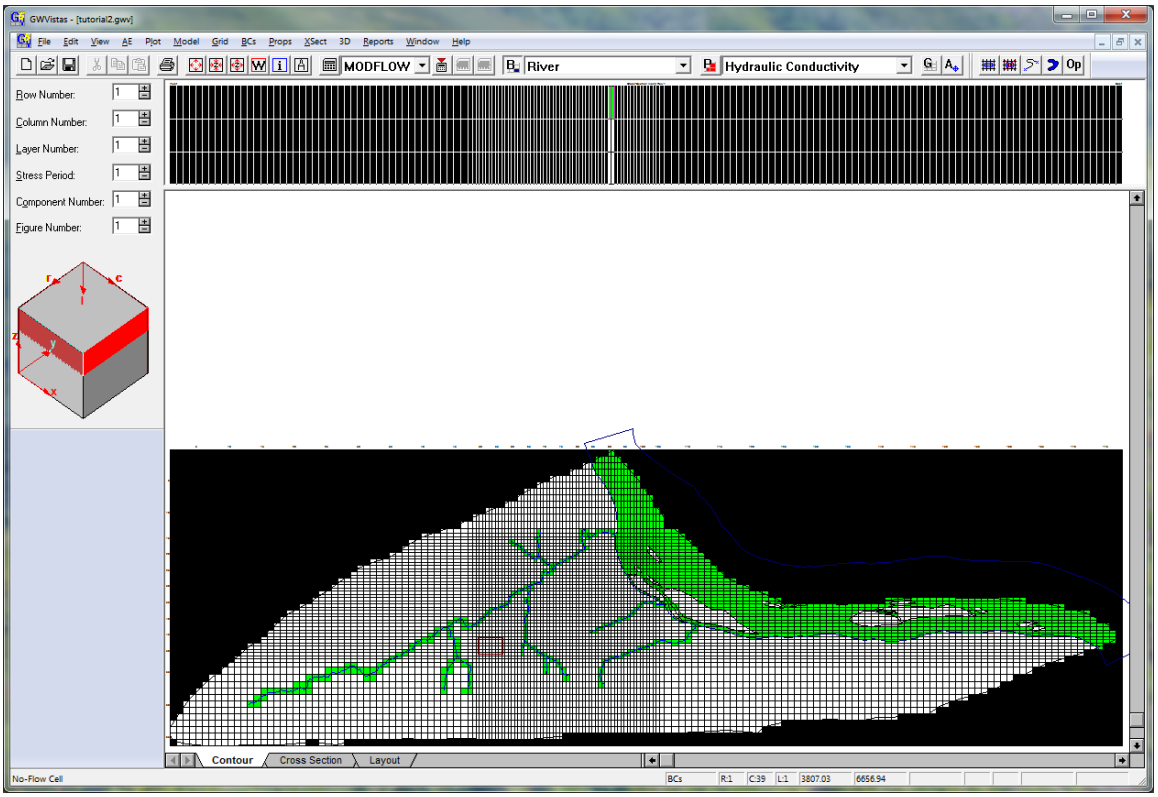

The river boundary condition should only be used for streams that are always flowing. Ephemeral streams are better simulated by drains because they will go dry if the water table drops below the stream channel. In this model, the upper portions of the tributary streams are ephemeral but we have defined them as rivers from the shapefile you just imported. There is a relatively simple way of making them drains, though.

Select **BCs|Modify|Change BC Type**. Change the options so that you are converting rivers to drains and only in reach 1. That is the most upstream portion of the main tributary channel.

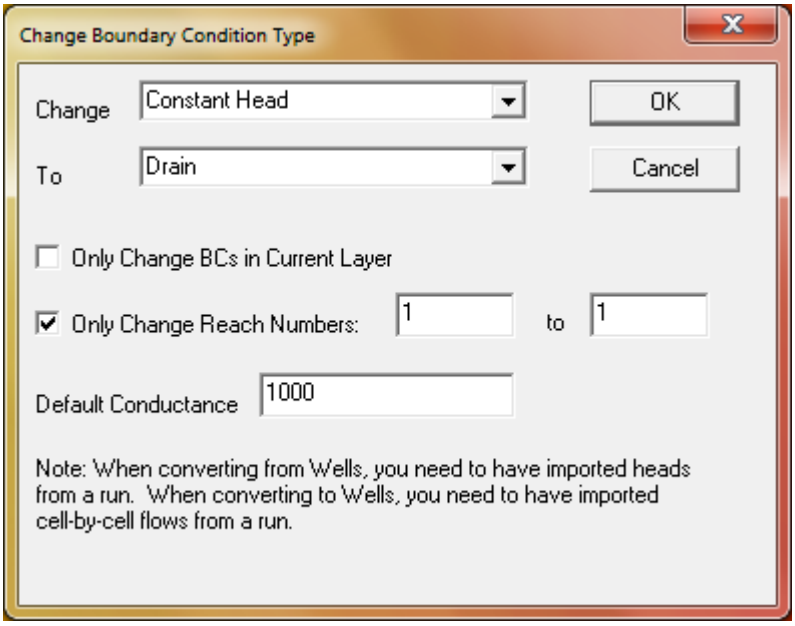

Click OK, and those cells should turn yellow, indicating they are no drains. Repeat this procedure for reaches 6 through 12, 14 through 15. After all of those changes, your model should look like the following.

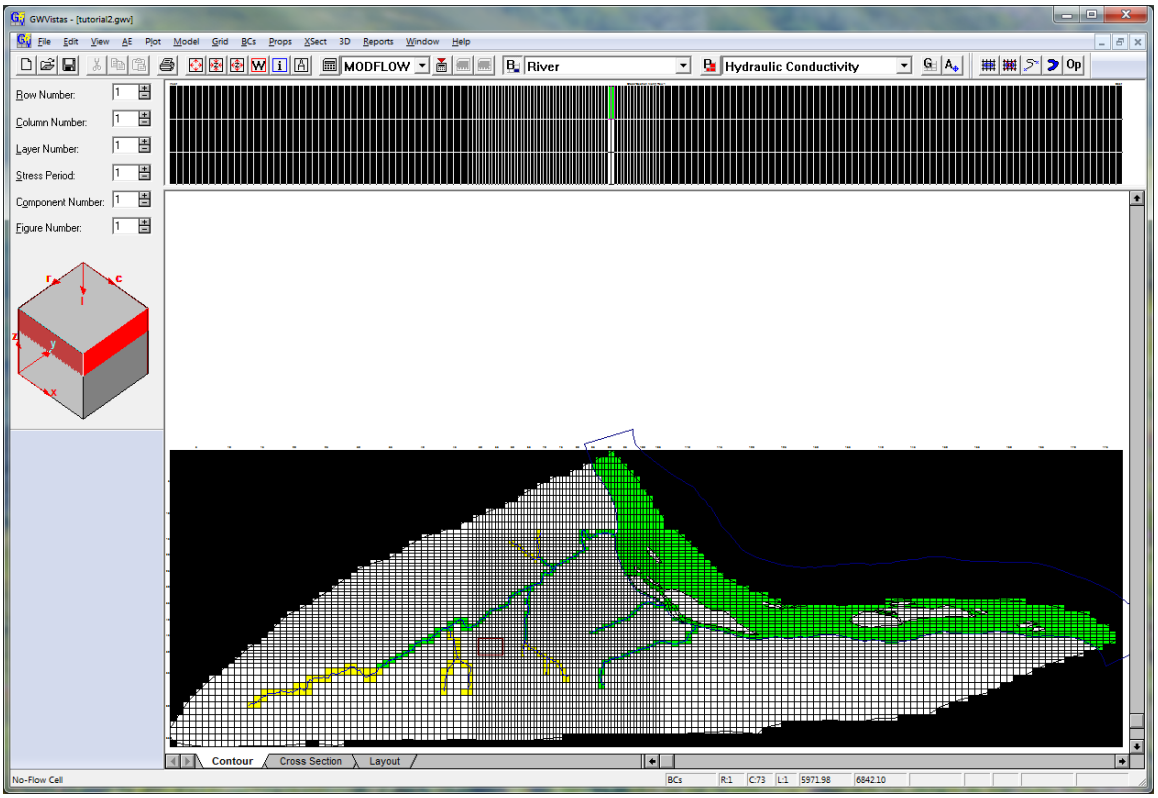

Adding Pumping Wells

Wells can be simulated in Groundwater Vistas using either a Boundary Condition or an Analytic Element. The only real difference is that the Analytic wells can span more than one layer. You will import 3 pumping wells from a shapefile called ***wellfield.shp***. Select **AE|Import|Shapefile** and browse to find that file in the gwv6\tutorial directory. Modify the import dialog as shown below.

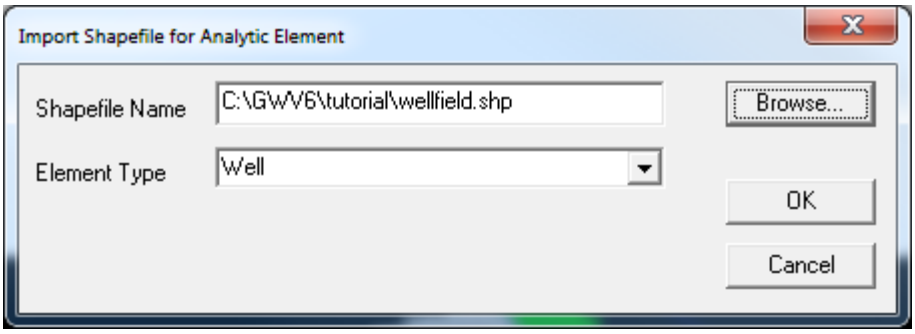

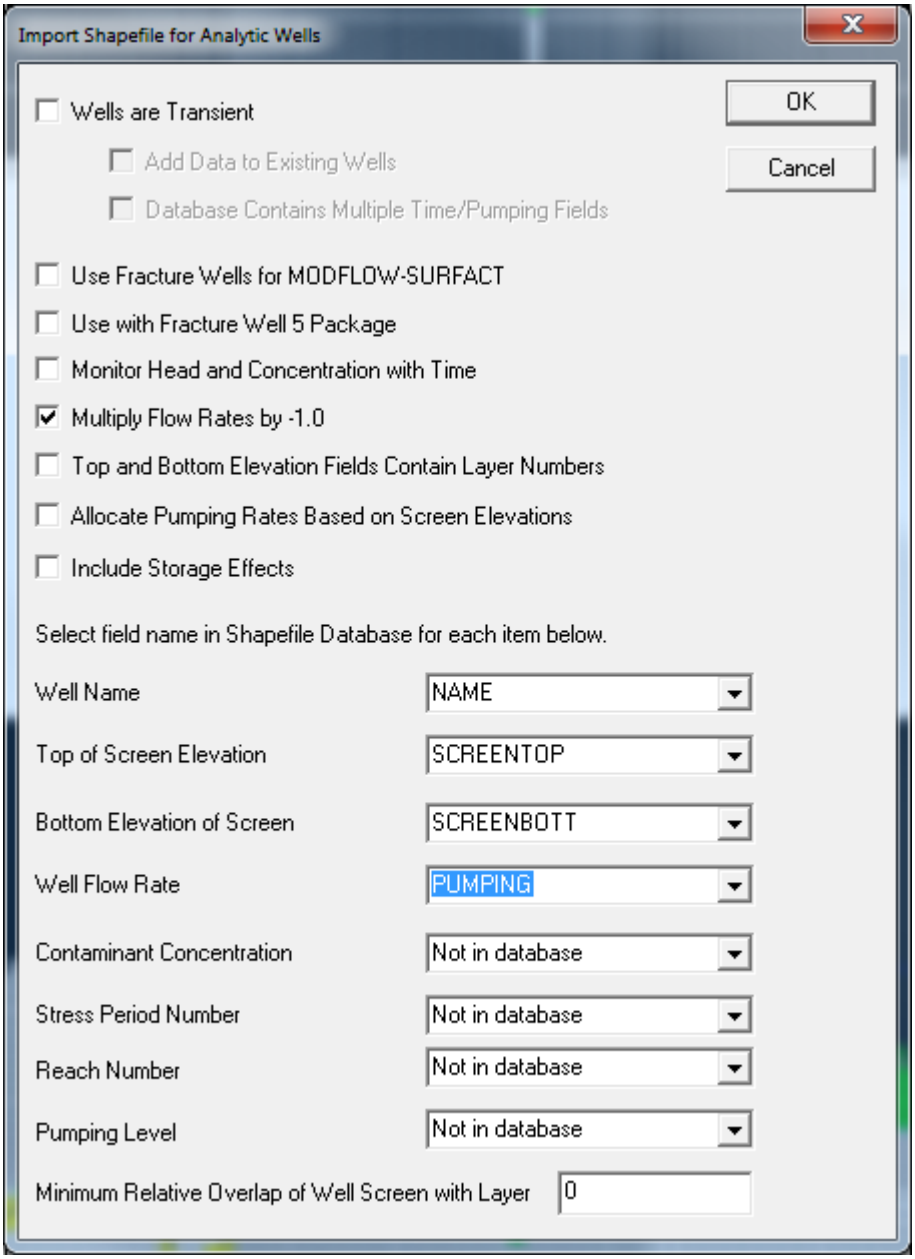

You should get a message that 3 wells were imported. You may only see one well, however, in layer 1. The other 2 wells are screened only in layer 2.

Setting Hydraulic Conductivity

We will set up the model with 5 different zones of hydraulic conductivity. The first three are in bedrock, with zone 1 assigned to layer 1, zone 2 to layer 2 and zone 3 to layer 3. Zone 4 is the alluvium in layer 1 and zone 5 is the alluvium in layer 2. The first step is to enter the hydraulic conductivity values for each zone. Select **Props|Hydraulic**

Conductivity and then **Props|Property Values|Database**. Enter the Kx, Ky, and Kz values as shown below. Also set a unique color for each of the five zones.

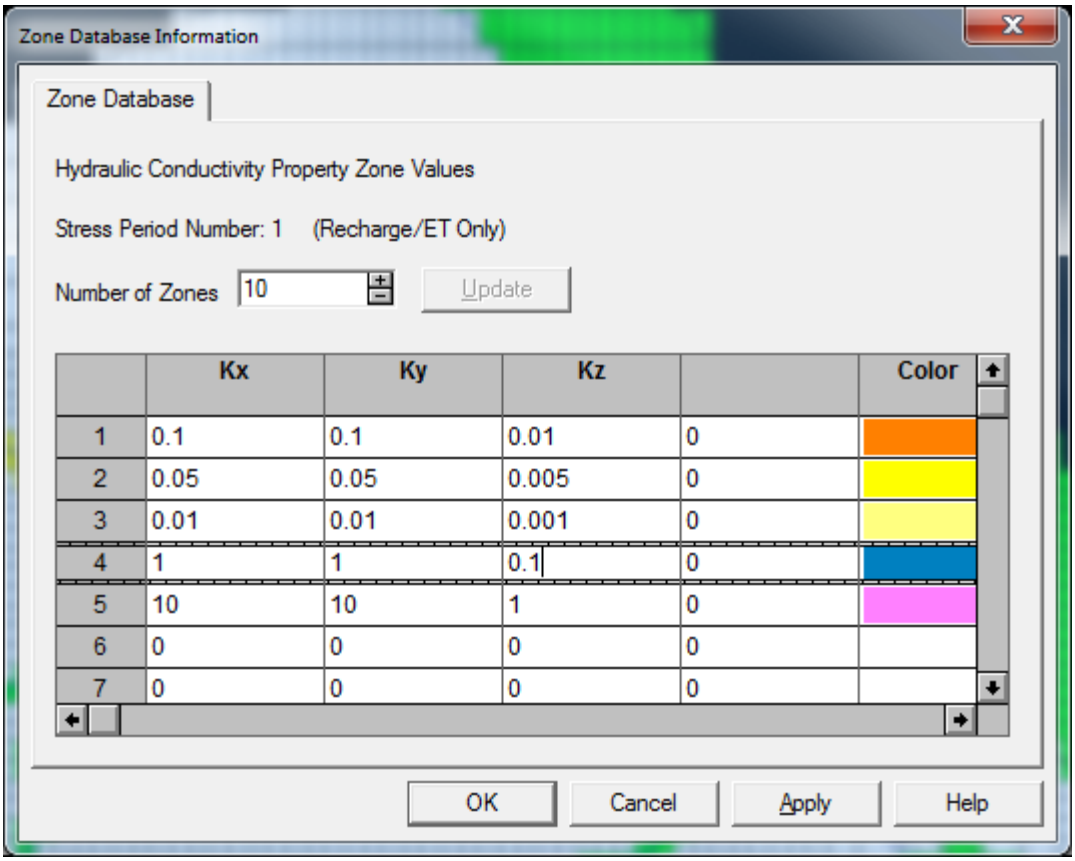

Now we need to assign zone numbers to model cells. By default when you start a new model, all cells in the model in all layers are assigned zone 1. Therefore, we do not need to do anything in layer 1 which will contain zone 1. Go to layer 2 and choose **Props|Set Zone or Value|Clear** and enter 2 for the zone number. Repeat this in layer 3, clearing the layer to zone 3.

The alluvial units are imported from shapefiles called ***alluvium.shp*** and ***alluvium2.shp*** for layers 1 and 2, respectively. First, go to layer 1. Then select **Props|Import|Shapefile** and browse to find the first file. The only thing we are importing from this shapefile is the zone number. Set up the dialog as shown below.

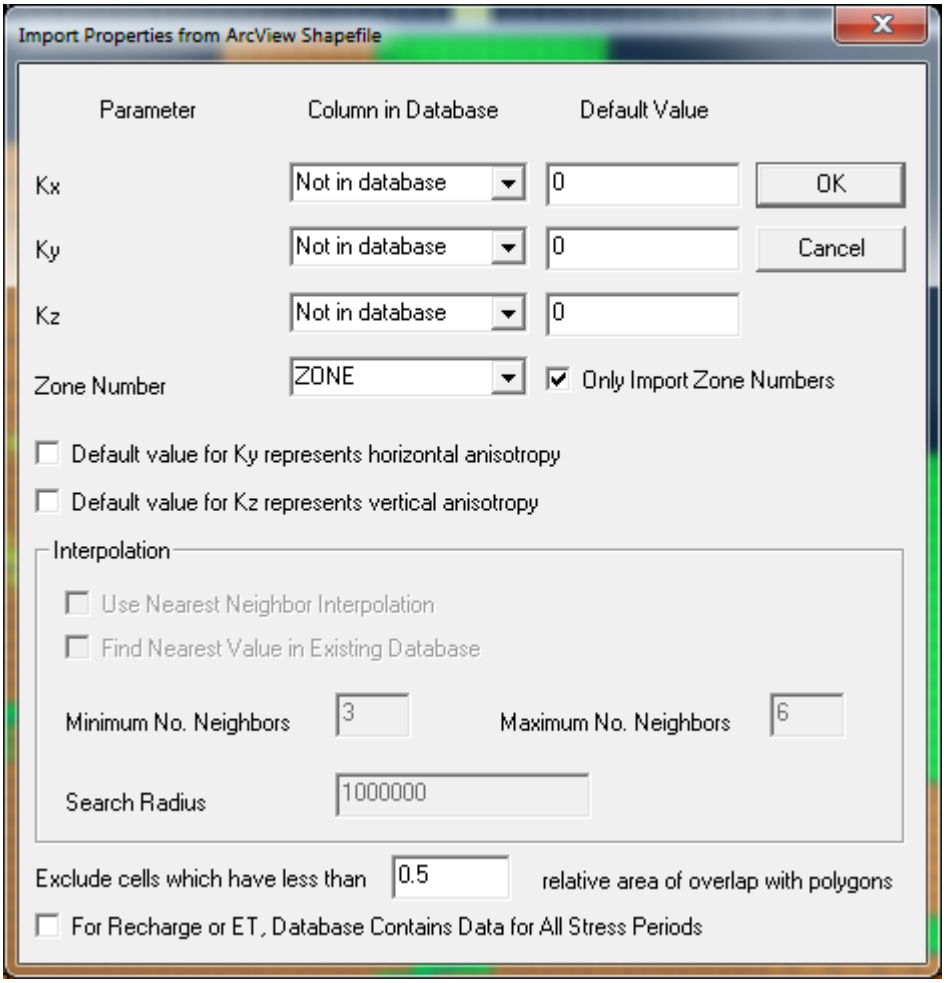

Go to layer 2 and repeat this for the second alluvium shapefile. After importing both files, your model should look like the one below in layer 1.

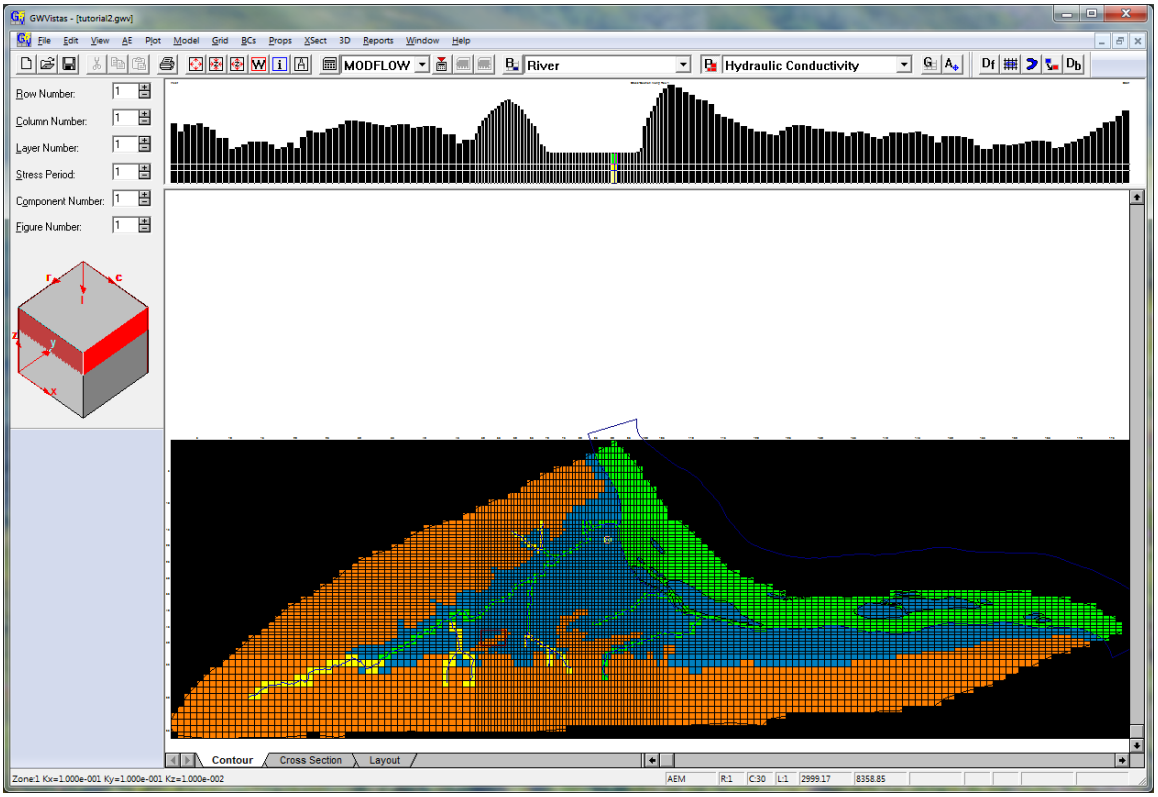

You are finished setting up the model. Note that most of the work was done in the GIS and then the data simply imported into the model. The advantage of this approach is that you can quickly redesign the model grid and then reload the GIS information.

Running the Model

The model is ready to run. Select **Model|MODFLOW|Packages** and change the root name to ***model1***. Note that Vistas figured out that you have rivers, recharge, drains, etc. and set up those packages for you. We also need to set up a few other MODFLOW options. Select **Model|MODFLOW|Package Options** and click on the **Initial Heads** tab. Change the initial heads option to use the top of layer 1. This is a good option to use at first until the model is stable.

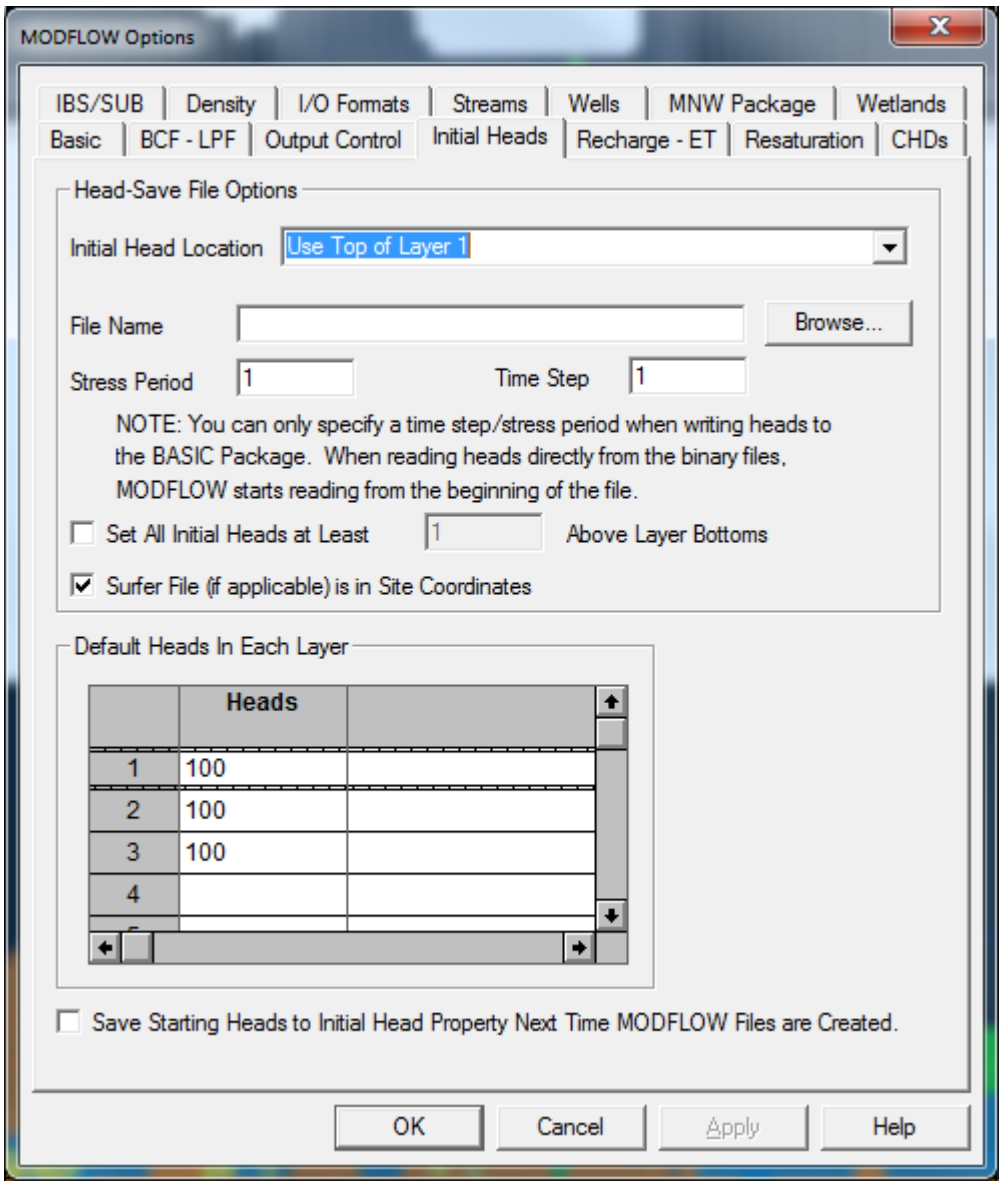

Run the model by clicking the calculator button and answering Yes to create datasets. After the run, import model results. Select **Plot|Contour|Parameters (plan)** and change the contour interval to 5 m. Your model results should look like the following.

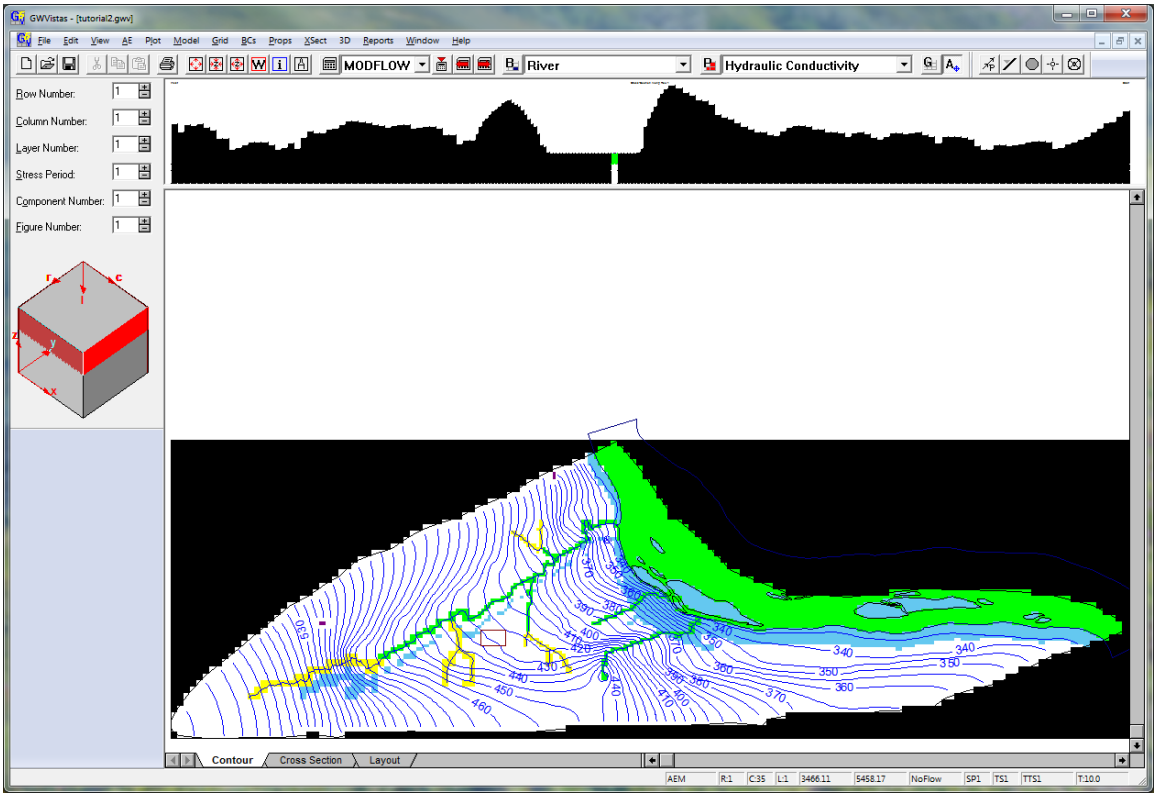

We will be using this model in several more exercises. Please save it and give it a unique name that you will recognize so you can use it later.

Particle Tracking with MODPATH

There are two ways of determining if an area of interest (the landfill in this case) is within the capture zone of a well. One uses reverse particle tracking. That is, particles are released around the pumping wells and tracked upgradient to define the portions of the aquifer contributing flow to the wells. The second approach is to place particles in the landfill area and track them downgradient to see if the wells capture them. We will start with the reverse particle-tracking strategy.

Using MODPATH for particle-tracking is relatively simple. Once the MODFLOW model has been run, you only need to add a couple features to the model. One is to define porosity, which MODFLOW does not require. The second is to add particle starting locations.

Start by defining the porosity of the bedrock and alluvial aquifers. Open your Groundwater Vistas file for this model and select **Props|Storage** – porosity is a part of the storage property in Groundwater Vistas. Since the bedrock and alluvial aquifers will have differing porosity values, we will copy the hydraulic conductivity zones to the storage property. Select **Props|Set Value or Zone|Copy**. The default property to copy from is already set to hydraulic conductivity (K). Simply check the option to copy all layers, as shown below.

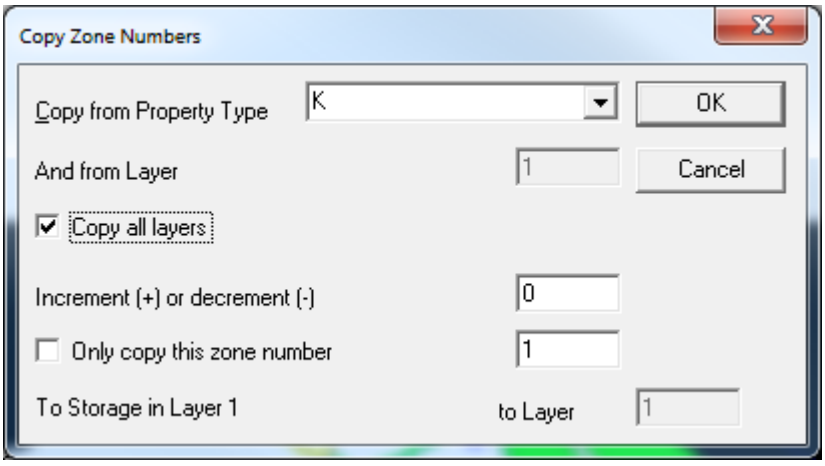

Now select **Props|Property Values|Database** and enter a porosity value of 0.01 for the bedrock zones (1 to 3) and a value of 0.2 for the alluvial zones (4 and 5), as shown below.

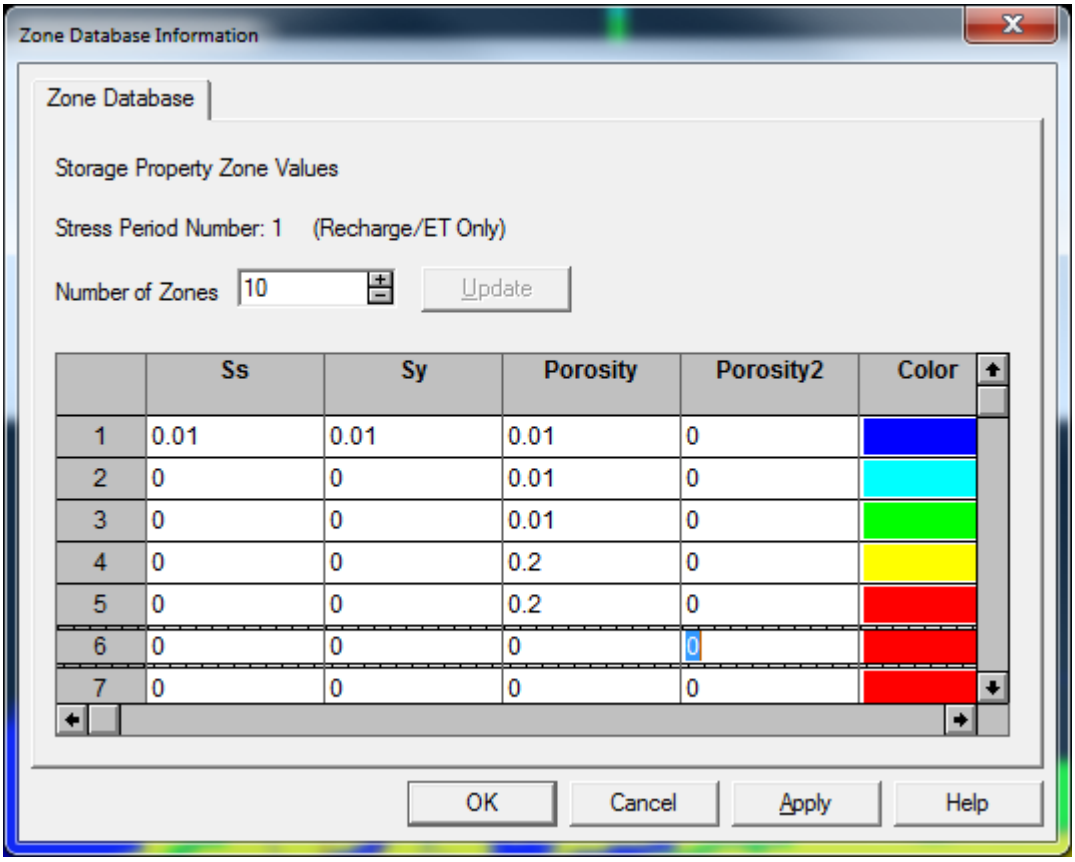

The second step is to add particle starting locations. Go to layer 2, which is the lower layer of the well screens. Zoom in on the wellfield area using View|Window or some other option on the View menu. You might also want to display the finite-difference grid if you turned it off before. This is helpful because the wells are not located at the center of the grid cells. However, the reverse tracking is most useful when the particles are

spread evenly around the edge of the cell containing the well. Select **AE|Particle|Circle** and drag a circle around one of the wells. Release the left mouse button and change the number of particles to 12. Also, change the Z Offset to zero. The Z Offset is MODPATH terminology for the vertical position of the particle within the layer. A value of 0.0 means the bottom of the layer and a value of 1.0 means the top.

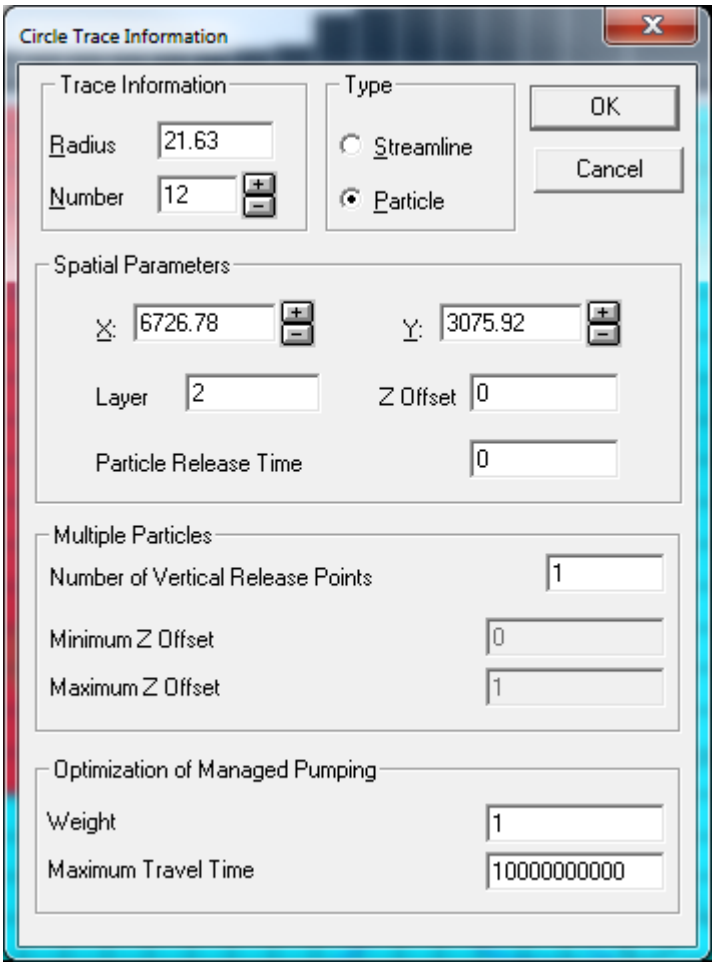

Repeat this procedure for the 2 other wells. You should then have 3 circles of particles around each well, as shown below.

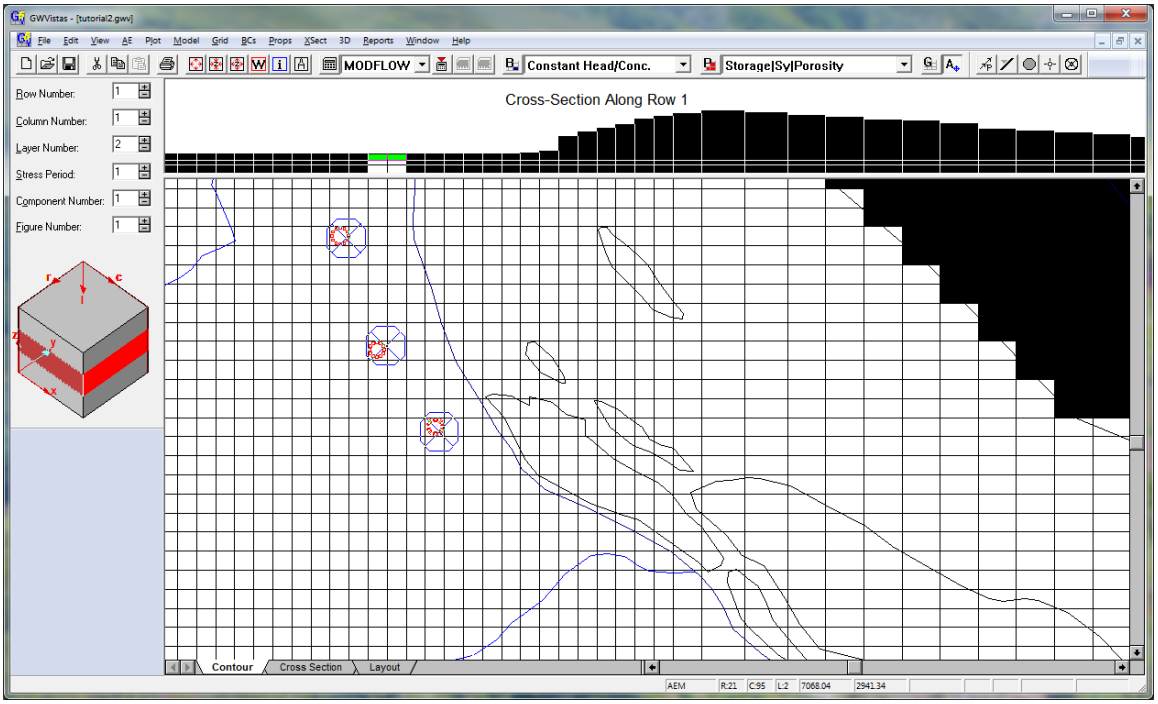

The only other thing to do is set MODPATH for reverse tracking. Select **Model|MODPATH|Particle Options** and change the tracking to "Reverse".

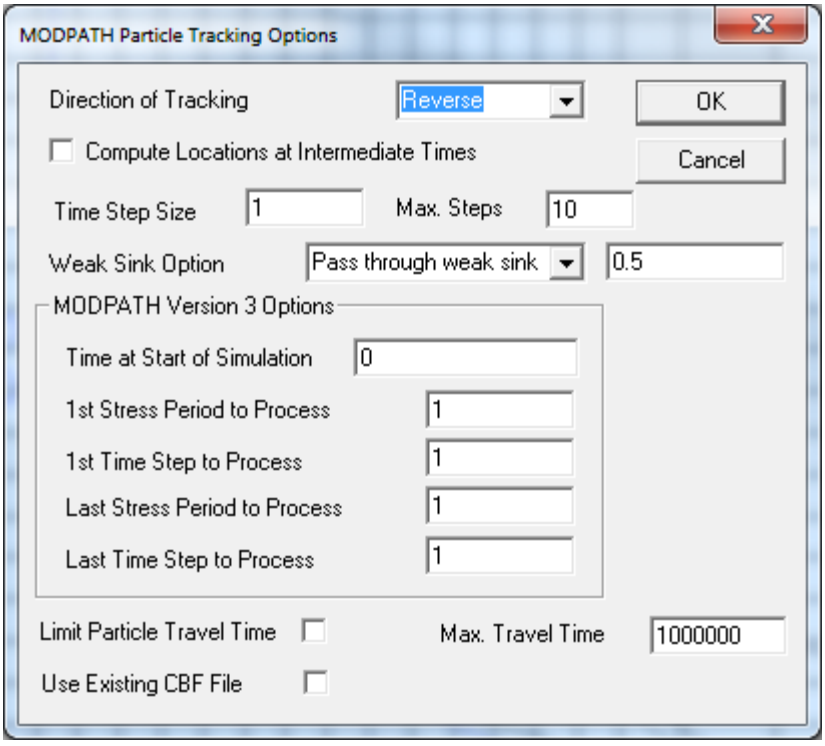

It is a good idea to give the simulation a unique MODPATH root file name. Select **Model|MODPATH|Packages** and enter a name you will remember. Now, to run MODPATH, select **Model|Use MODPATH** and click the calculator button. Answer Yes

to create datasets. MODPATH should run very quickly. After the run, import the modpath results (*.ptl)

The particle traces do seem to indicate that the landfill is within the capture zone of the central well, as shown below.

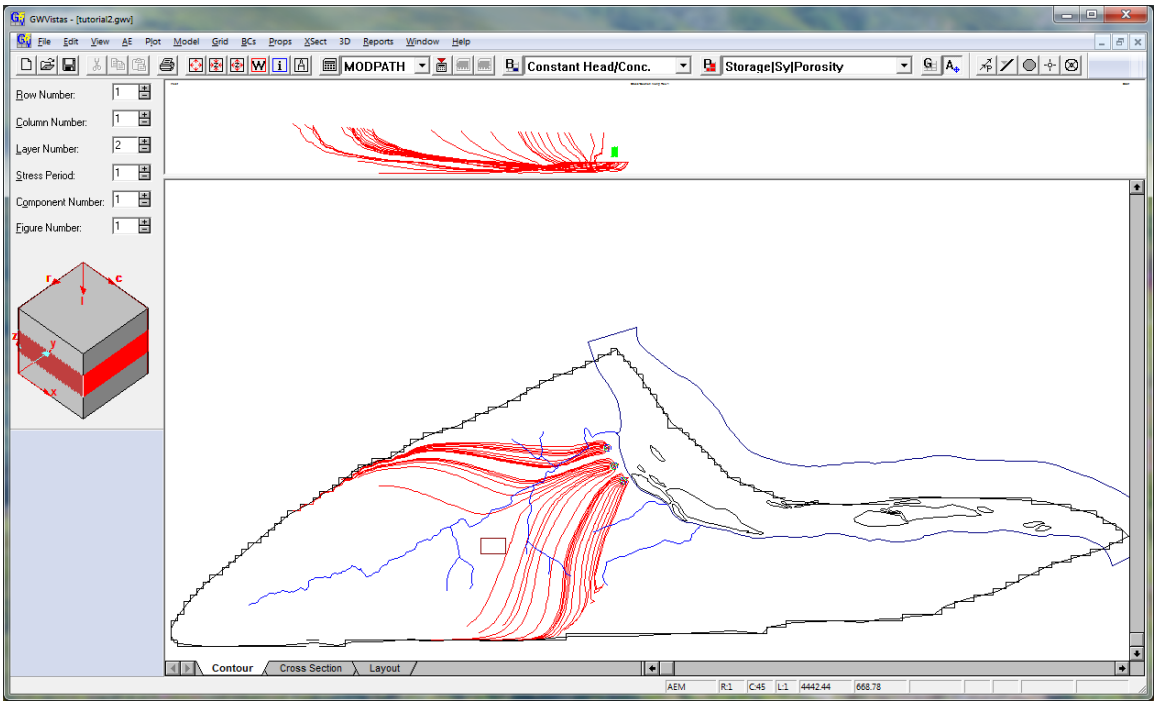

One question is how long the potential contaminants will take to reach the well. Select **Plot|Particle|Options** and click on the time posting tab. Turn on time postings and use 3650 (10 years) as the interval between arrows. Change the multiplier to 0.00274. This will round the posted time values to integer years (1|365).

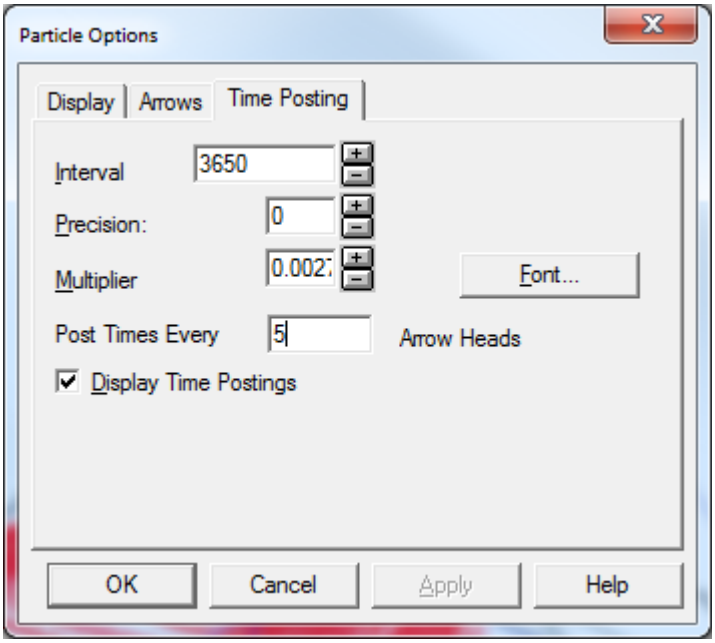

Obviously the wellfield would be in no immediate danger!

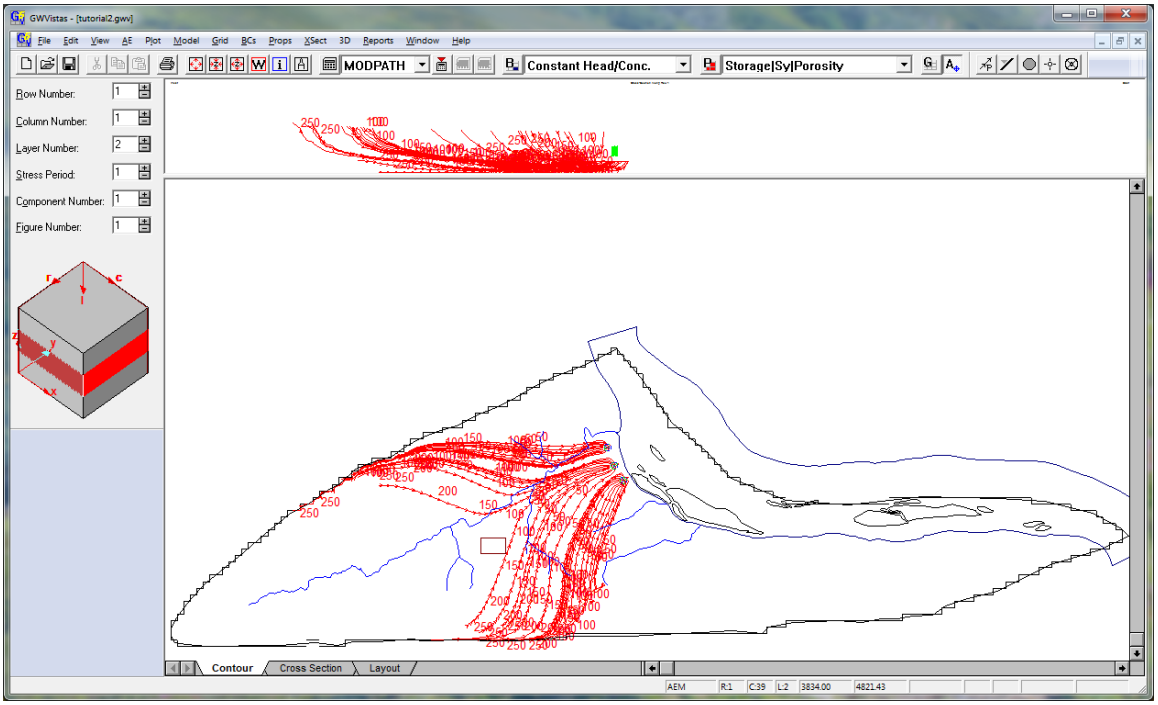

Reverse Particle Tracking

To confirm these results, we will add particles in the landfill area and track them toward the wellfield. First, though, we will delete the existing particles. The easiest way of doing this is to select **Edit|Select All|Particles** and then **Edit|Delete**.

Add the particles using **AE|Particle|Window** and drag a window round the landfill. Make sure the layer number you enter for the particles is 1. After you see the following dialog, GV will prompt for the number of particles per cell. Just keep it at the default of 1.

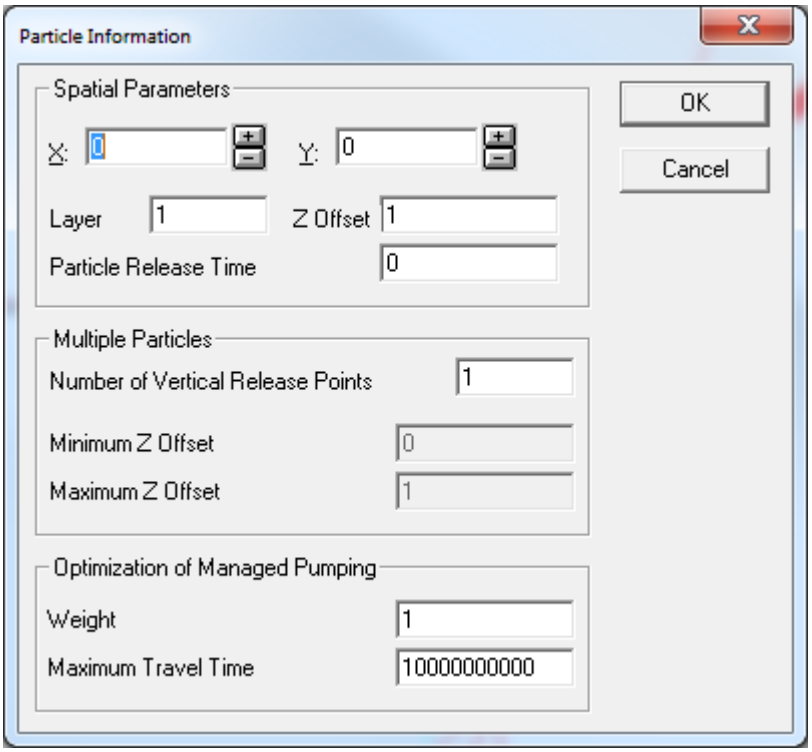

Select **Model|MODPATH|Particle Options** and change from "reverse" to "forward". Rerun MODPATH and import the results. The results should show that all particles get captured by two wells and not just one. Also, there are too many particles to adequately see the travel time results. We can change the display to only show one particle. Zoom in on the landfill and record the row and column number of one of the particles. Then select **Plot|Particle|Options** and turn on the option to only display particles from a particular row|column|layer location. Now, the plot shows a travel time that is much shorter than the reverse tracking indicated. Why?

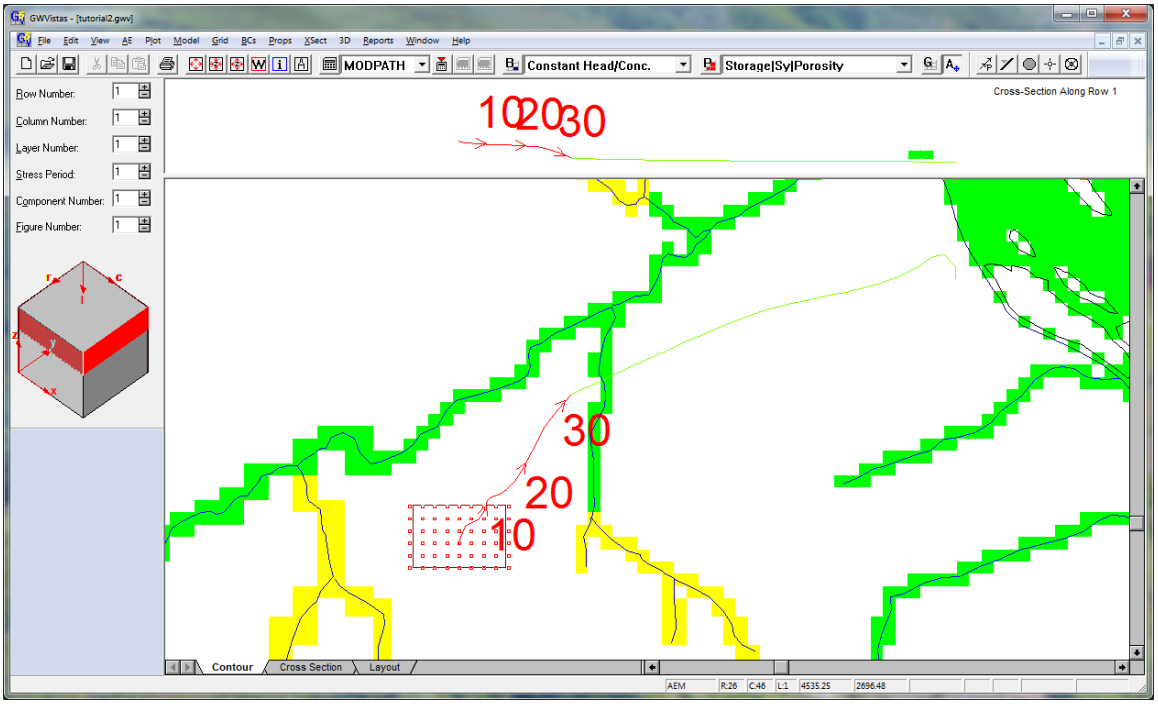

In the picture above, we turned on the option to display a different color for each layer where the particle travels. This indicates that the particle starts in layer 1 and ends in layer 2 (the high K alluvium).

If you go back and do the same with the previous simulation using reverse tracking you will see that the particles travel in layer 3, the low-K bedrock. So obviously you need to pay attention to the vertical dimension when interpreting particle tracks. In this case, releasing particles in layer 1 around the wells would have also been helpful.

Transport Modeling with MT3DMS

In the last exercise, we used MODPATH to determine whether a leaking landfill would be within the capture zone of a wellfield. The analysis indicated that the landfill was within the capture zone and that it would take about 50 to 60 years to reach the wells. This type of simulation, though, does not give any indication of the concentration of contaminants that might reach the wells. In order to do that a contaminant transport model, like MT3DMS, is required. We will take the same model from the last exercise and run MT3DMS to determine the concentration of contaminants that might reach the wells.

Setting Up the MT3DMS Simulation

There are several steps to go through in setting up a typical contaminant transport simulation. These will usually include:

- Rerun the MODFLOW simulation to save a flow-link file for MT3DMS

- Define transport properties (dispersivity, porosity, and chemical reaction data)
- Assign contaminant sources
- Assign initial concentrations
- Set some MT3DMS-specific options

The first step is to rerun MODFLOW2000 and save the flow-link file. This file is not saved by default because it can be large and because not all models involve a transport simulation. Open the MODPATH model you created in the last session (or open *gwv6\tutorial\bedrock_example_pt1.gwv* if you did not do the last section). Select **Model|MODFLOW|Packages** and check the box labeled "MT3D Flow Output" in the lower right side of the dialog, as shown below.

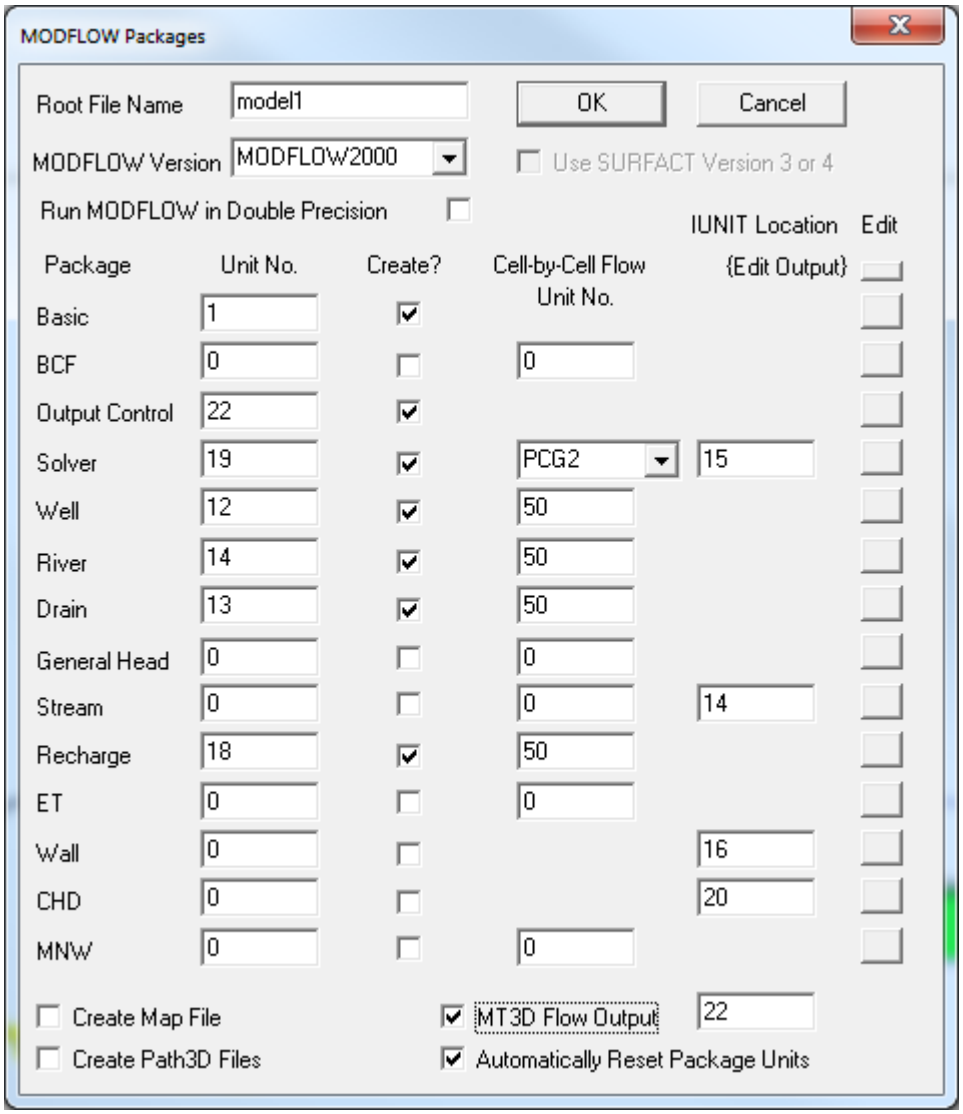

We also need to change the length of time the MT3D model will simulate. In a steady-state MODFLOW simulation, the stress period length does not have any meaning. However, when using MT3DMS with a steady-state flow run, the stress period length is

the length of time that transport will be simulated. Select **Model|MODFLOW|Stress Period Setup** and change the stress period length to 10950 days (30 years).

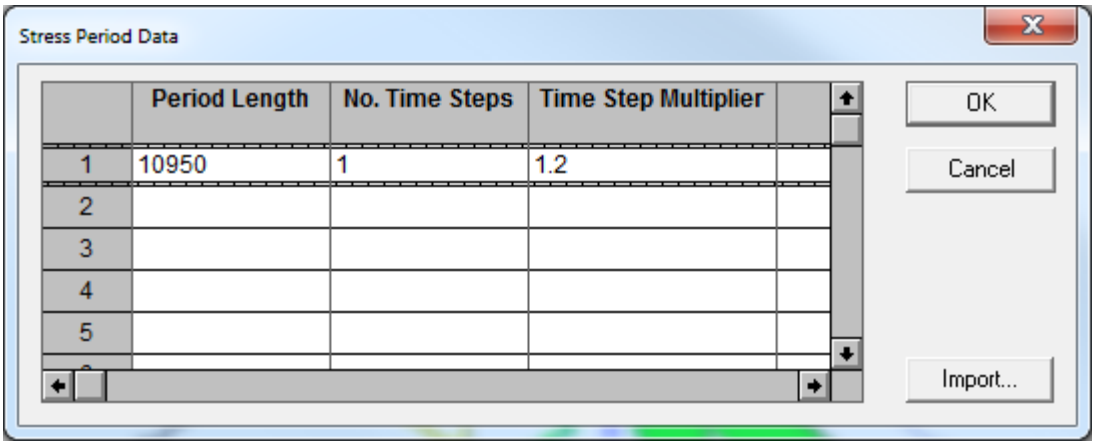

Now run MODFLOW2000 and create new input datasets.

Several properties related to contaminant transport are required by MT3DMS. These include porosity, dispersivity (longitudinal, transverse, and vertical), contaminant half-life, and distribution coefficient (for computing a retardation coefficient). In the last session you defined the porosity for use in MODPATH so we do not need to do that again. Dispersivity is entered by selecting **Props|Dispersivity** and **Props|Property Values|Database**. Enter the dispersivity values as shown below.

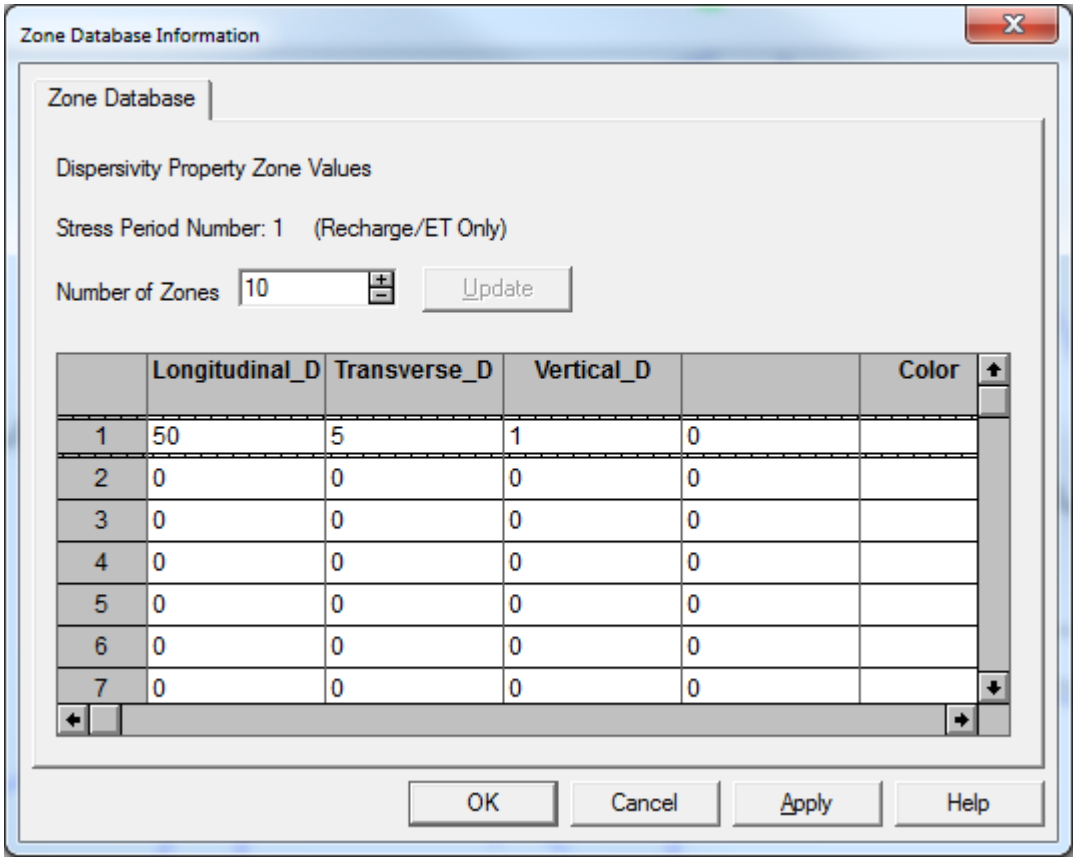

In the first simulation, we will assume that the contaminant is a conservative tracer. That is, it does not degrade nor adsorb onto the aquifer material. Thus, no reaction properties will be entered.

We will add a source of contaminant in the landfill. There are two common ways of defining a contaminant source. The first is to enter a concentration in recharge because many contaminant sources start above the water table in the vadose zone. The second option is to define a constant concentration boundary condition that holds the contaminant at a specified concentration. The latter is most appropriate for non-aqueous phase contaminants where the constant concentration is the solubility limit of the contaminant. In this case, we will use a recharge source.

Select **Props|Recharge** and **Props|Property Values|Database**. Define a second recharge zone with the same recharge rate as the first and with a concentration value of 1000.0. Give the second zone a different color.

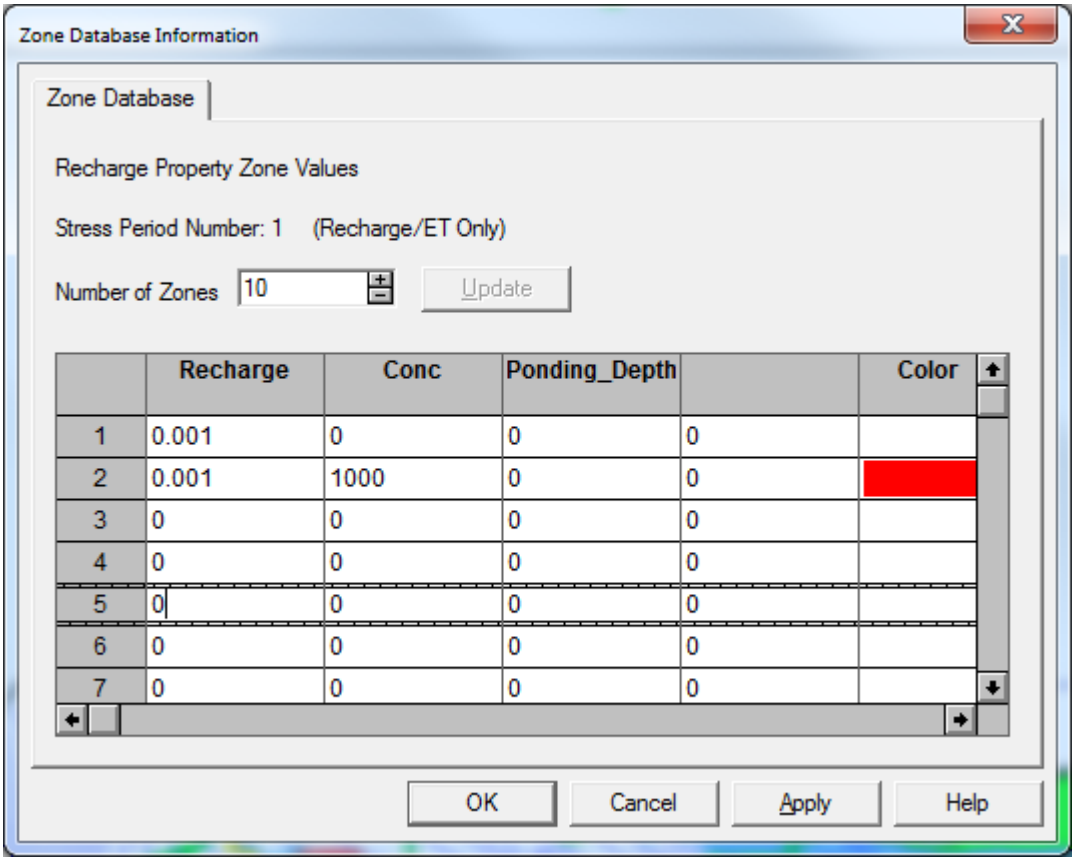

The database now contains the data for the second zone. To enter it in the model, select **Props|Set Zone or Value|Window** and drag a window around the landfill. Enter 2 for the zone number when prompted. Your model should look like the following.

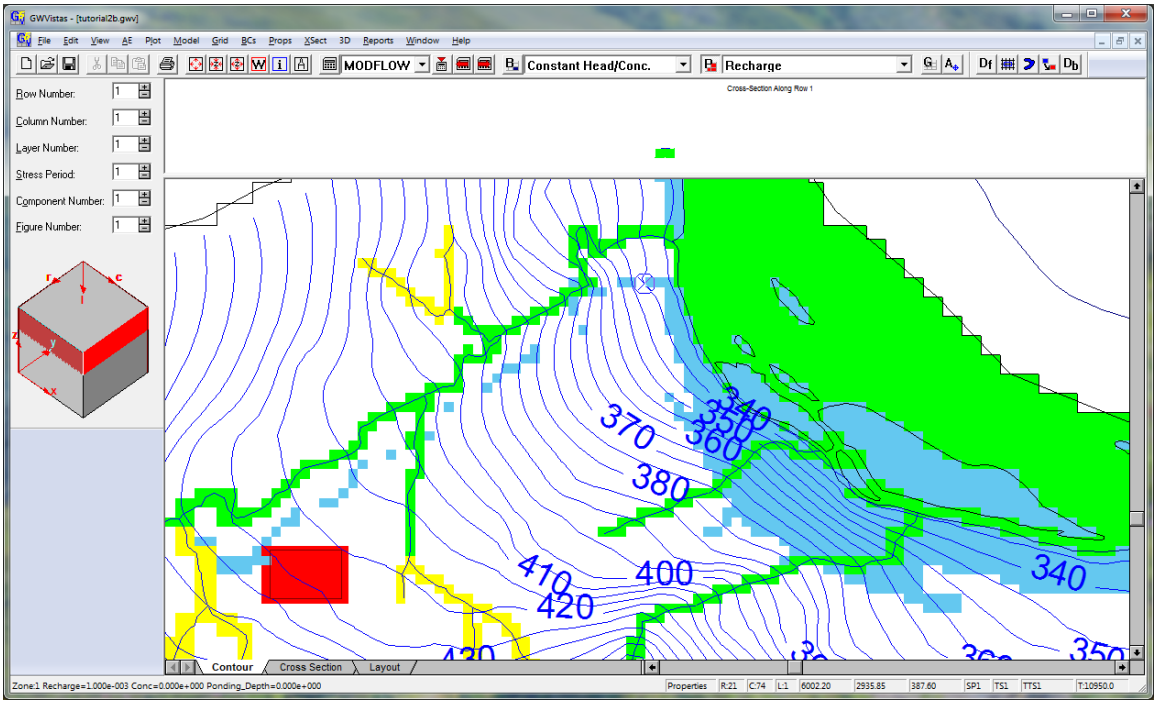

Since we are starting from the beginning of the landfill operations we will not define an initial concentration. The default initial concentration is zero, meaning that there are no contaminants in the aquifer to start the simulation.

There are several MT3DMS options to set before running the transport model. The first is to give the run a root name. Select **Model|MT3D|Packages** and enter a unique run name (e.g. mt3d1). Also, uncheck the 2 columns next to the reaction package. We are simulating a conservative tracer so not reactions are required.

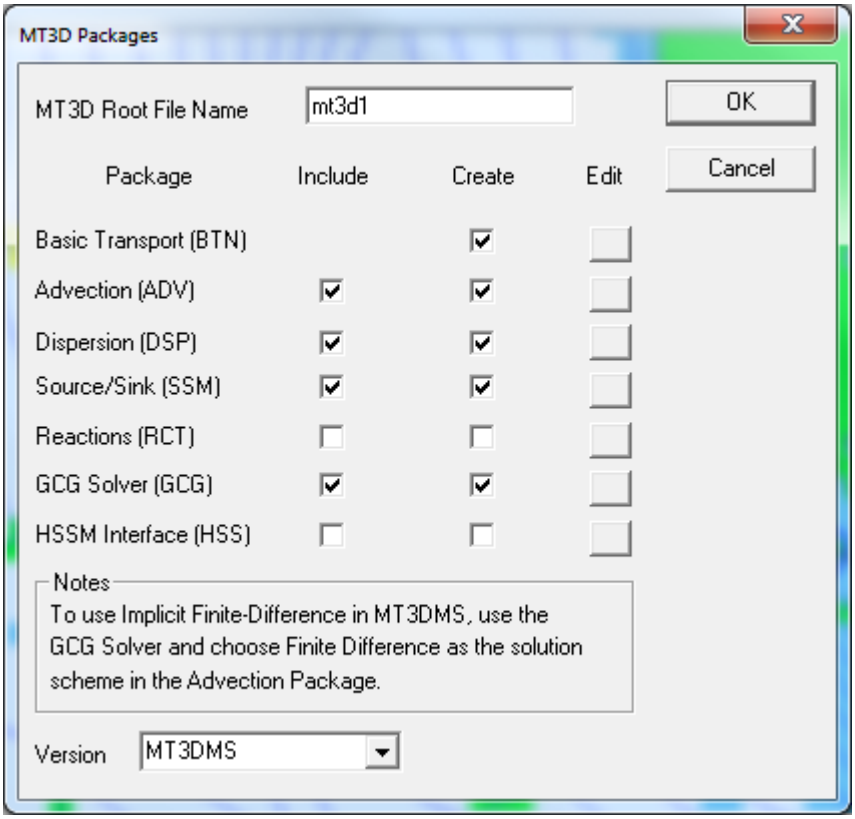

Next, select **Model|MT3D|General Options – Time Stepping tab**. Change the number of transport steps per flow step from 500 to 5000. Click on the **Printing** tab and change from saving every 2 time steps to every 5 time steps.

Running MT3D

You are ready to run the transport model. Select **Model|Use MT3D** and then click the calculator button. Create datasets and after the model is finished import results for the last time step. Make sure the concentration file is called MT3D11.UCN (or the mt3d root name you entered with a "1" attached). If not click the browse button next to the concentration file name. Put a check next to the concentration file. Click the browse button next to the transport step number to choose the last time step, as shown below.

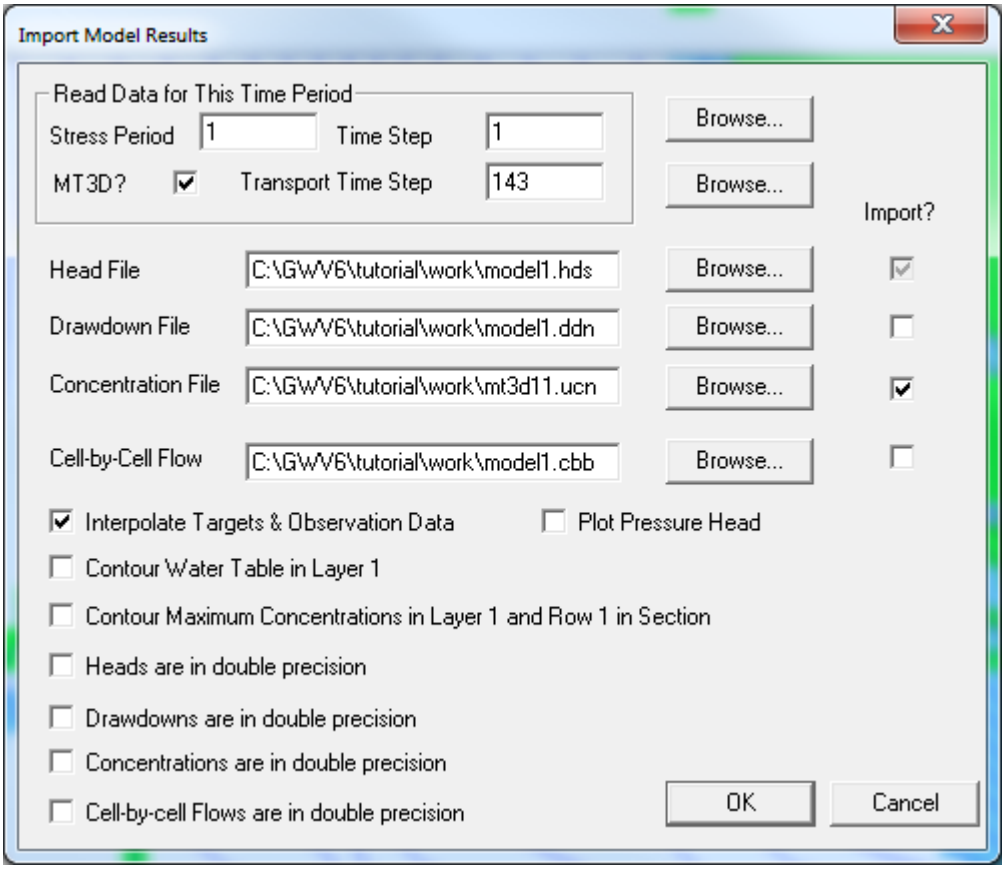

Contour the concentration results by selecting **Plot|What to Display** and changing the contour variable from "Head" to "Concentration". Click OK and go to layer 2. Your results in layer 2 should look like the following:

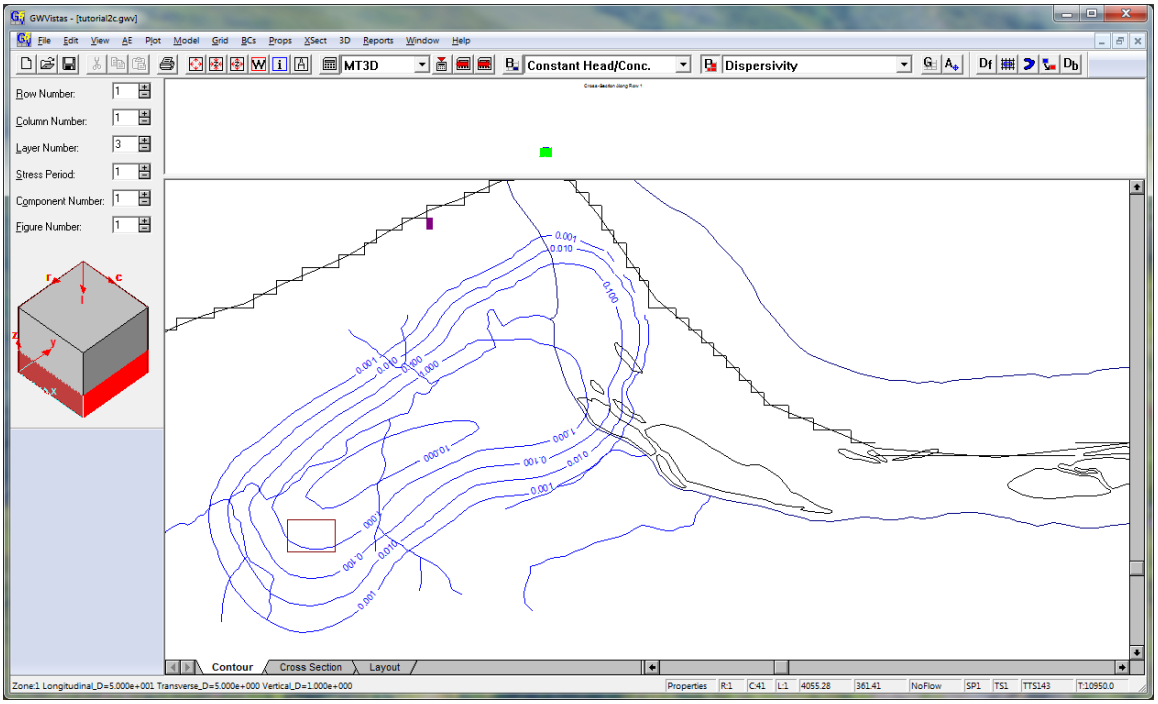

Notice that the highest concentration contour is 10 even though our source concentration was 1000. You might think that the highest concentration should be closer to 1000. However, since we are using a recharge source, the 1000 concentration is in the recharge water. That water is diluted as soon as it hits the first layer in the model. The more groundwater flows through that first layer, the model dilution occurs.

Concentration Break-Through Curves

The purpose of this type of model is to determine whether concentrations will migrate to some receptor above some regulatory limit. In this example, the receptors of concern are the 3 pumping wells. To determine whether concentrations get above some limit, the best thing to do is plot concentration break-through curves for each well. To do that, push down the **A** button on the toolbar. Double-click on each well and put a check in the option labeled "Montitor Head and Concentration vs. time". Now, select **Plot|Import Results** and click OK. Monitoring data are stored in each well during the import of results. Now select **Plot|Hydrograph|Monitoring Well**. The hydrograph for the southern well is shown below.

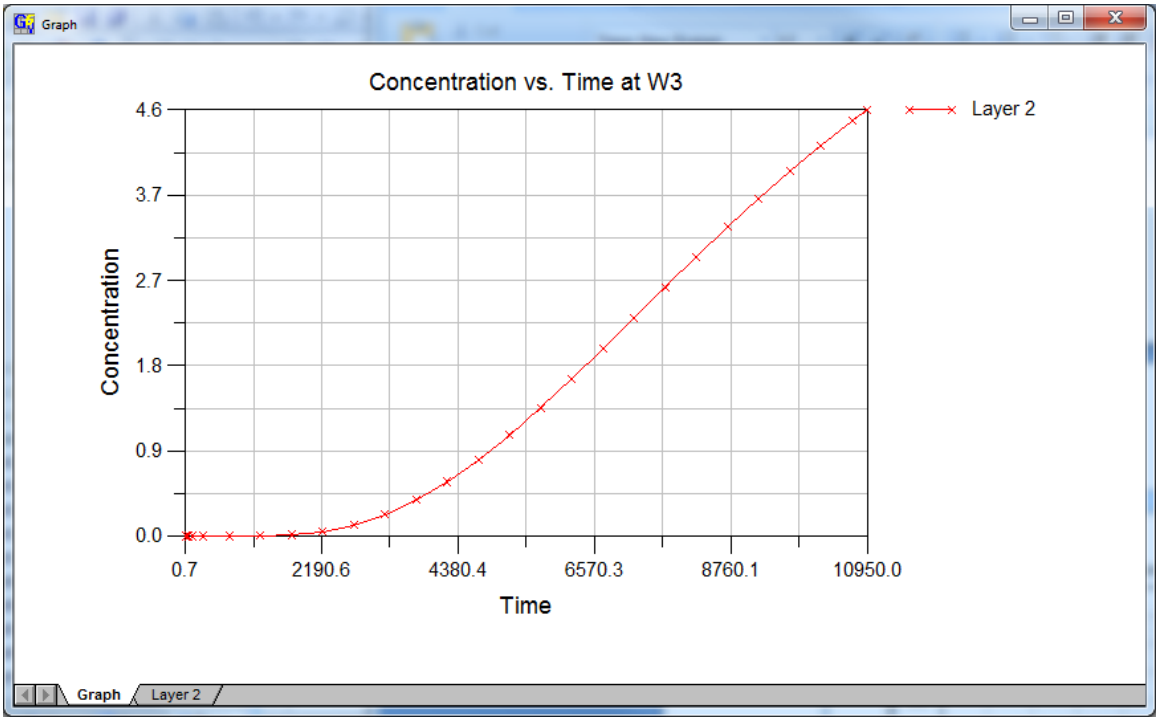

If the regulatory limit for this contaminant was 1.0, then there would be a problem. If the limit is higher than 4.6 then at least for the first 30 years there would not be a problem. However, the concentration has not peaked by 30 years. You might try running this model longer to see what the maximum concentration will be and how long it will take to get there. This is done by simply selecting **Model|MODFLOW|Stress Period Setup** and increasing the stress period length and rerunning MT3DMS (Note that you do not need to rerun MODFLOW2000 even if you change the period length).

Using the TVD Transport Scheme

You may notice that the concentration contours indicate quite a bit of smearing of the plume. This is caused both by the dispersion values you entered above but also by numerical dispersion from the finite-difference approximation. Another transport scheme called TVD is sometimes a better choice because it limits numerical dispersion. Run the previous example again using TVD. Select **Model|MT3D|General Options – Advection tab**. Change the scheme at the top from finite-difference to TVD. Also click on the **Time Stepping** tab and change the initial time step size to 10. Rerun MT3DMS.

One thing you will see right away is that the model run time has increased dramatically. This is, unfortunately, one of the problems with TVD – it requires much smaller time step sizes. After the run is finished, import the last time step again and compare to your previous result. You will probably see that the plume looks significantly different!

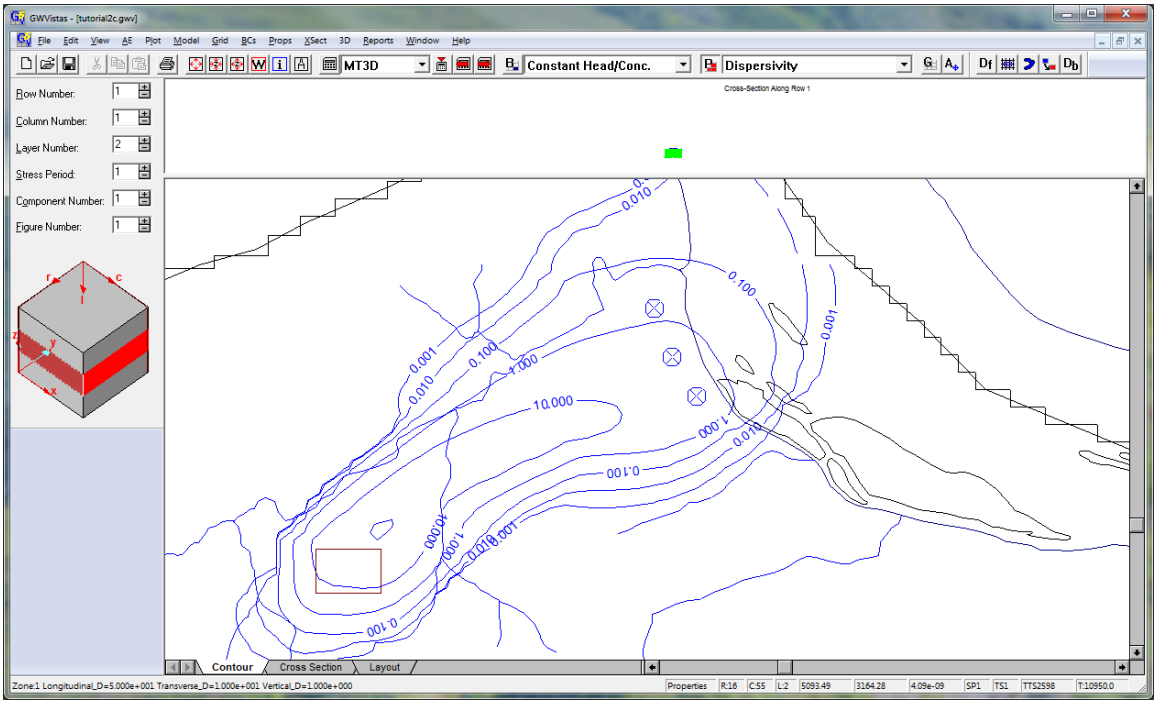

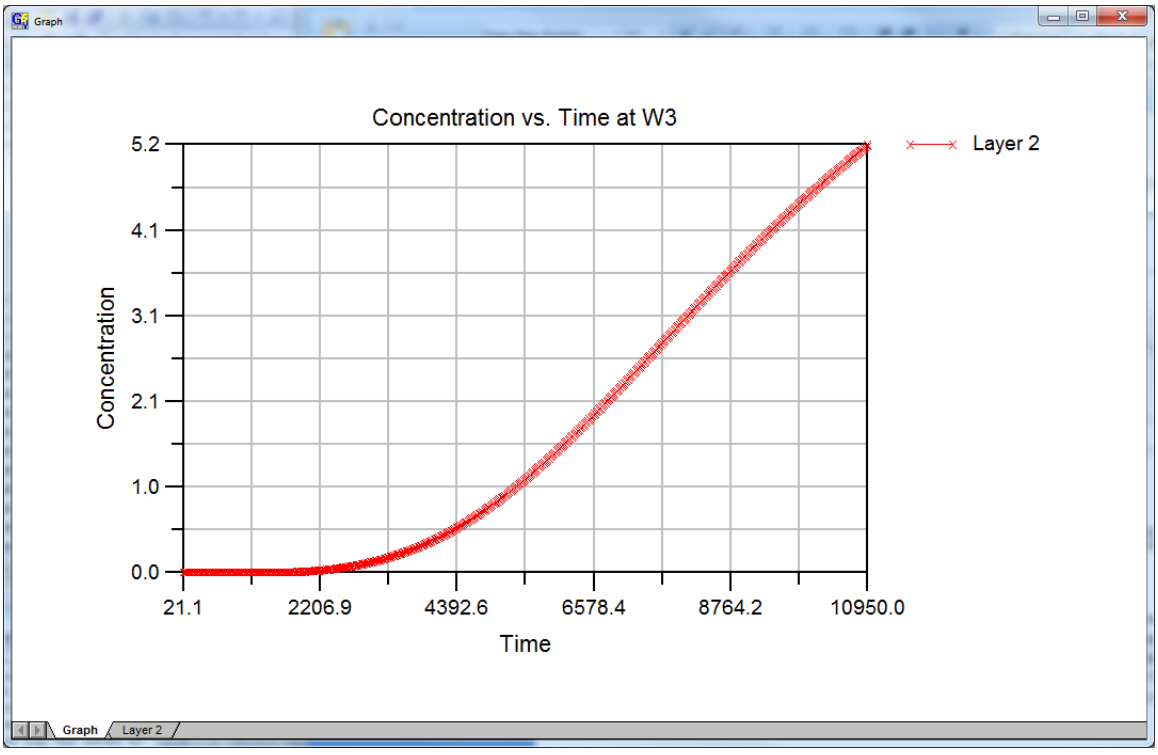

In the TVD simulation, the plume is narrower with less lateral dispersion. The concentrations near the center of the plume are higher and the concentrations at the wellfield are also higher. So, which simulation is correct? Since the finite-difference solution introduces a lot of numerical dispersion that was not called for by the parameters

you specified in the run, you would have to say that the TVD scheme is the more accurate solution.

Special Applications

Mine Dewatering

The first special application of MODFLOW is to mine dewatering. The issue in mine dewatering usually revolves around how much impact there will be from the dewatering of an open pit mine, how much water must be pumped by the mine over time, and then how long will it take the pit to fill back up once mining ceases.

Mine dewatering is generally simulated in MODFLOW using a sequence of drains places within the mine pit. The drain elevation (water level) is set to the elevation of the mine floor over time as the mine gets deeper and deeper. We will start from the flow model you saved before the particle tracking example. If you did not save that model, you can use **gwv6\tutorial\bedrock_example.gwv**. Open that model now.

In this example, the first stress period will represent steady-state conditions prior to mining. Mining will then following a sequence of 7 stress periods of 1-year duration each. During each year, the mine will get deeper by about 30 meters. This is not designed to be realistic but to just show you how a mine dewatering plan would be implemented in MODFLOW.

Start by selecting **Model|MODFLOW|Packages** and change the MODFLOW root file name to ***mine1***. Next select **Model|MODFLOW|Package Options** and remove the steady-state check-box and change the number of stress periods from 1 to 8, as shown below. Click OK when you are done.

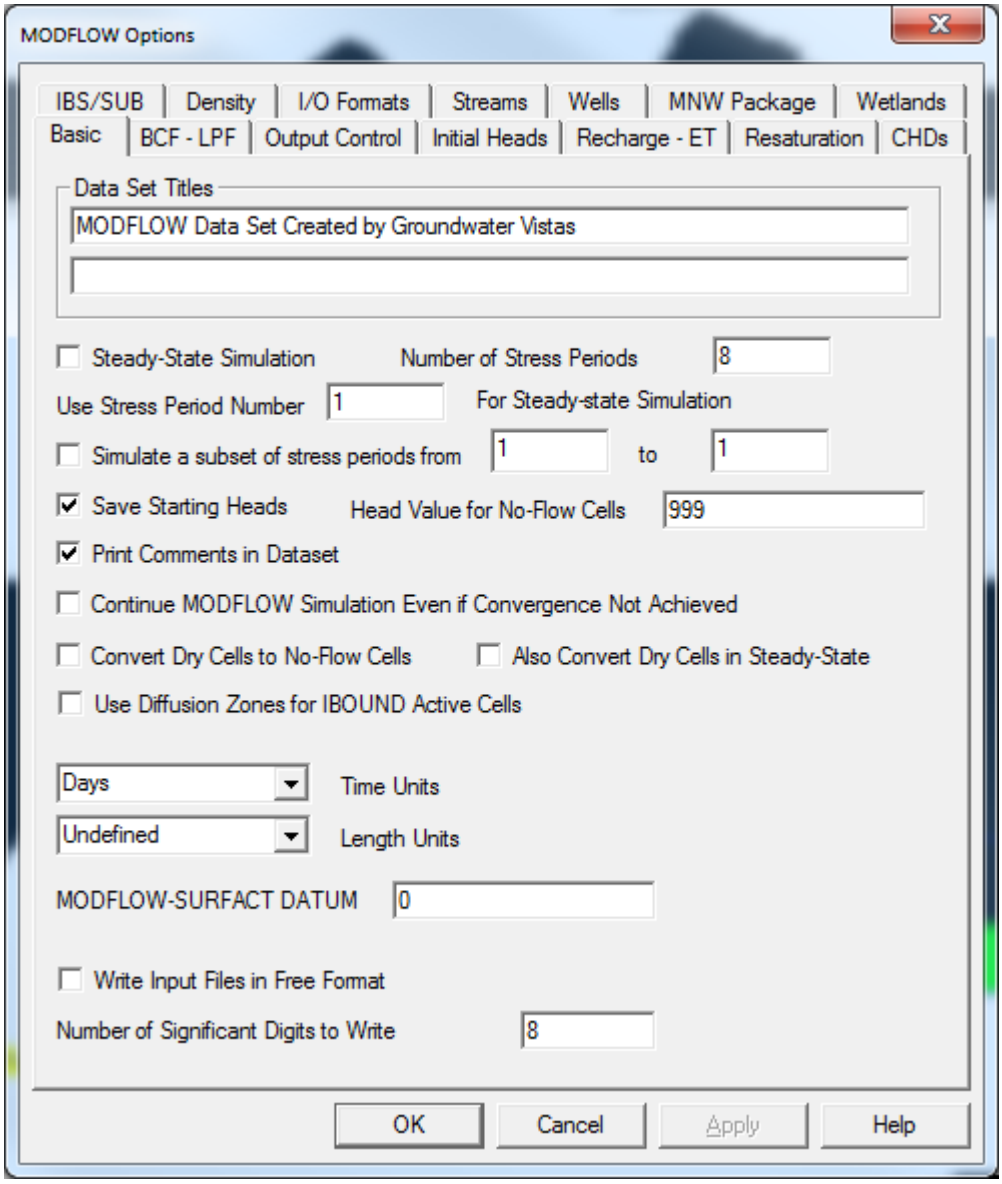

Groundwater Vistas will then ask if you want to copy Recharge from stress period 1 to the other stress periods. Answer Yes to this and all other prompts. The last question asks if you want to edit the stress period setup. Enter 365 for each stress period length in periods 2 through 8. Use 10 time steps in periods 2 through 8 as well, with a multiplier of 1.2. The multiplier of 1.2 is used to start each stress period with a small time step.

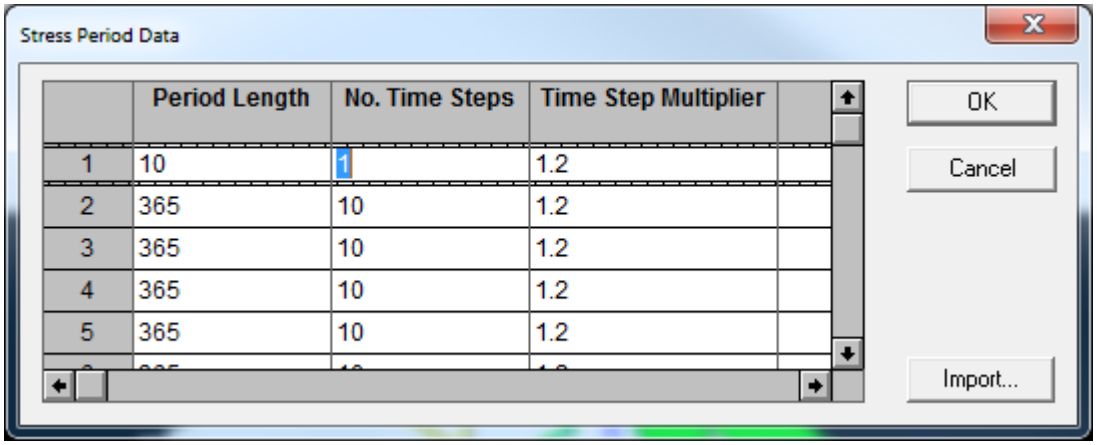

MODFLOW2000 has a very useful feature in that each stress period can be steady-state or transient. In this example, the first stress period is steady-state and the rest are transient. GV will automatically show this dialog to you when you change the number of stress periods. In this case, it will also correctly set the stress period types. You can edit the stress period types manually by selecting **Model|MODFLOW2000|Stress Period Types|Edit Stress Period Types**. You should see a zero in stress period 1 (representing steady-state) and 1 for all the other stress periods.

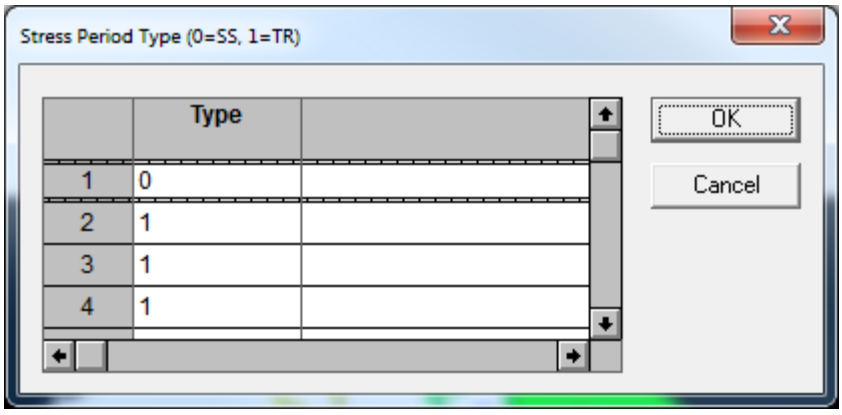

We will now take advantage of an undocumented feature in MODFLOW2000 to compute drawdown from the results of a particular time step. In previous versions of MODFLOW, drawdown was always computed from the starting heads for the simulation. In the current model, we would like to compute drawdown due to mine dewatering from the results of the steady-state stress period. To do this, select **Model|MODFLOW|Package Options – Output Control tab**. Turn on the option for custom output control and enter 8 for the number of custom steps.

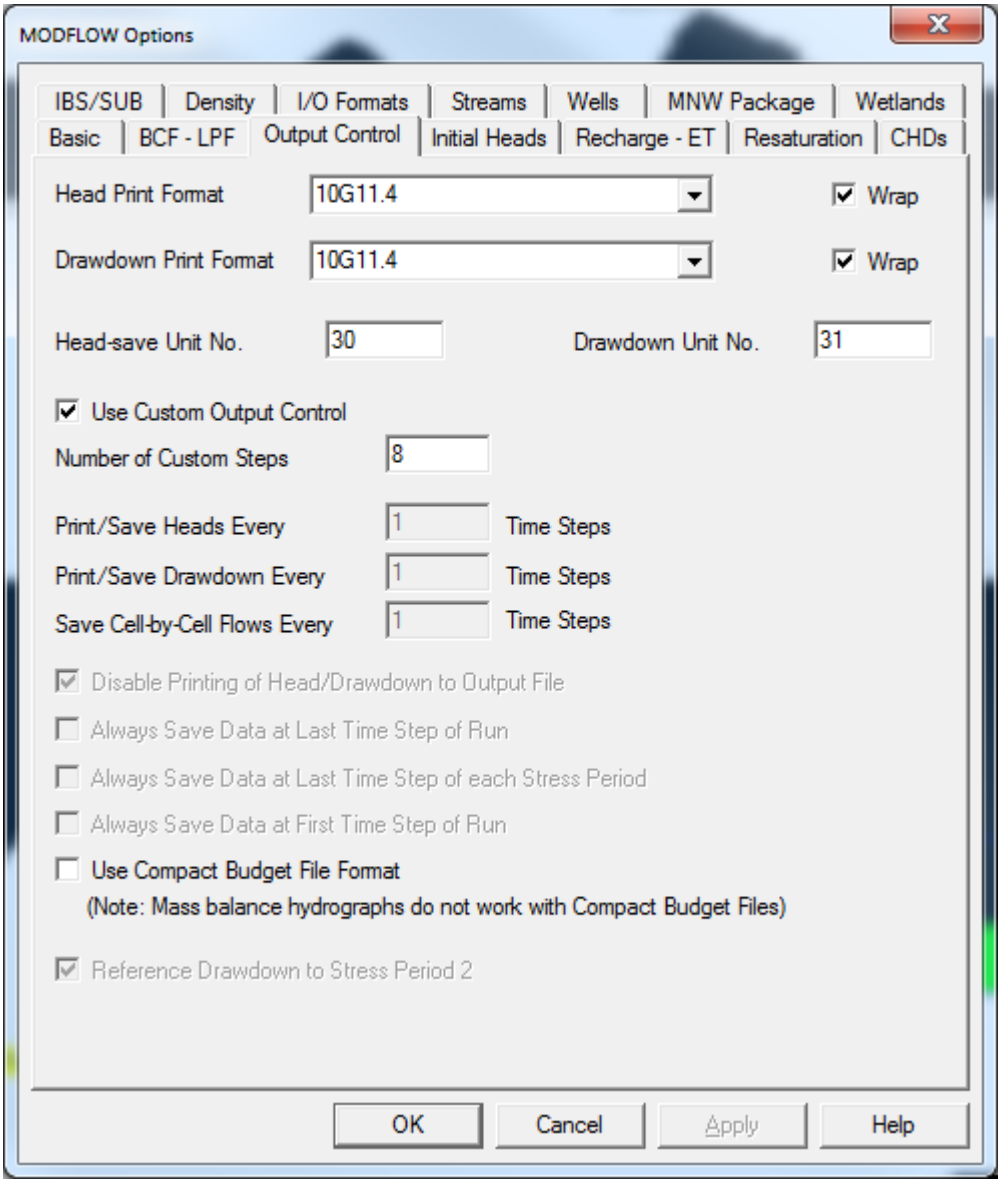

Now, select **Model|MODFLOW|Custom Output Control**. Fill out the table as shown below. In particular make sure to put a 1 in the first row beneath the heading "Ddn Ref". This option places the key word ***DDREFERENCE*** in the output control file in the first time step, which in turn causes MODFLOW to compute drawdown from that first stress period.

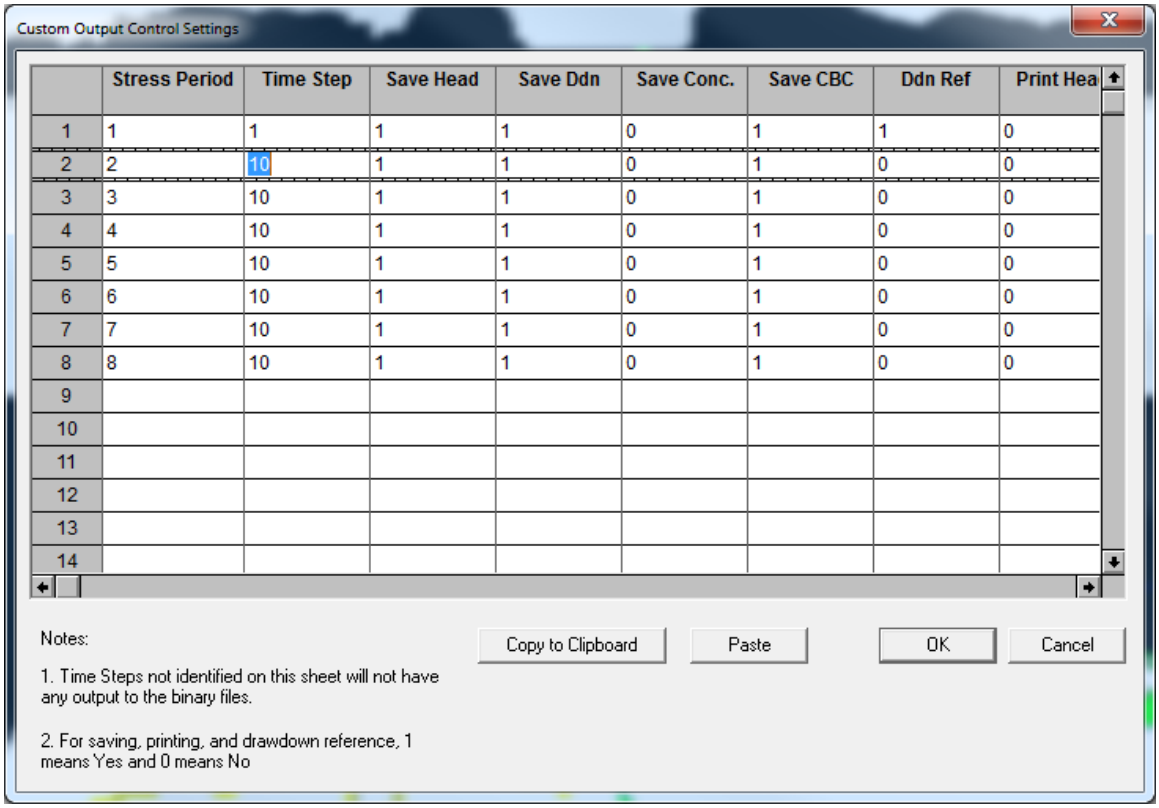

Because the simulation will have transient stress periods, we need to enter values for specific storage and specific yield. Select **Props|Storage** and **Props|Property Values|Database**. Enter 1.0e-5 for specific storage and 0.01 for specific yield, as shown below.

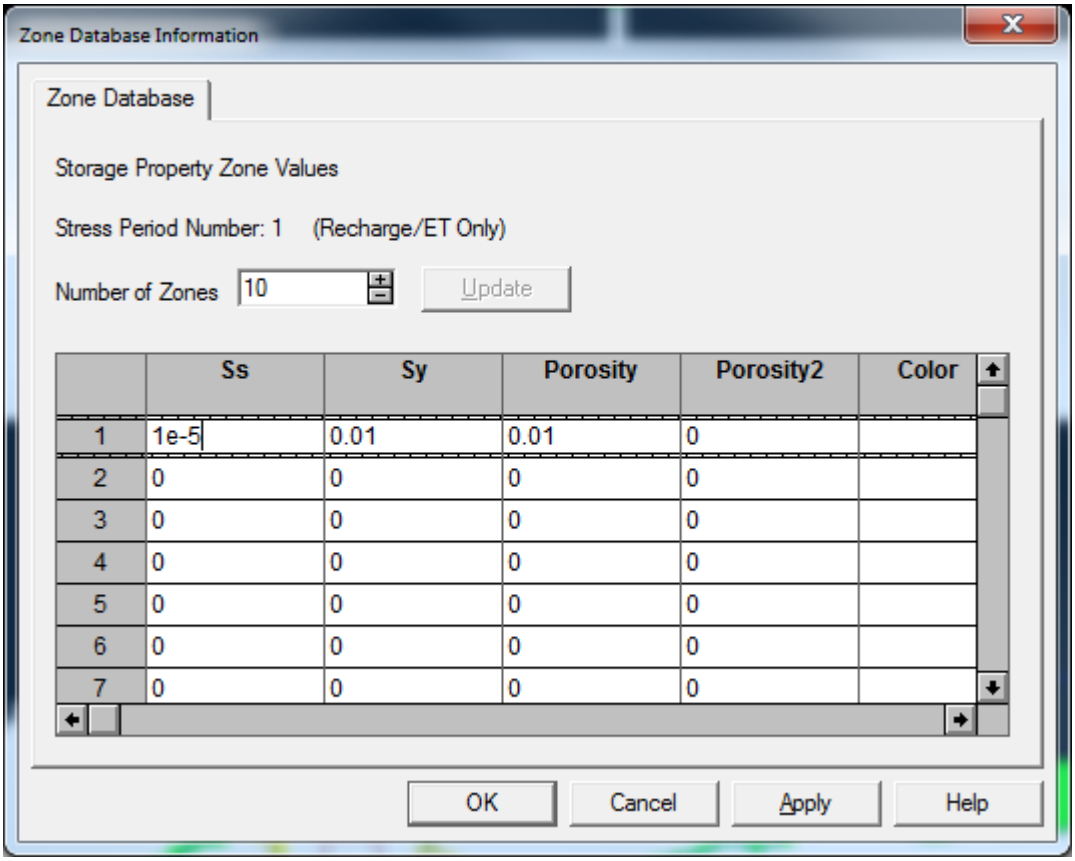

You will now place drain cells within the mine, along with a transient schedule of mine floor elevations. Zoom in on the red rectangle we used in previous examples as a "landfill". We will assume for this exercise that it is a mine pit. Make sure you are viewing layer 1 of the model. Select **BCs|Drain** and **BCs|Insert|Window**. Drag a window around this red rectangle and a drain dialog appears. Uncheck the "steady-state" option and enter 50 for the length and width of the drain (cell dimensions). You may also want to change the color of these mine drains to distinguish them from other drains in the model. Also, if you change the reach number to something unique (e.g. 99), you can evaluate mass balance in these drains more easily.

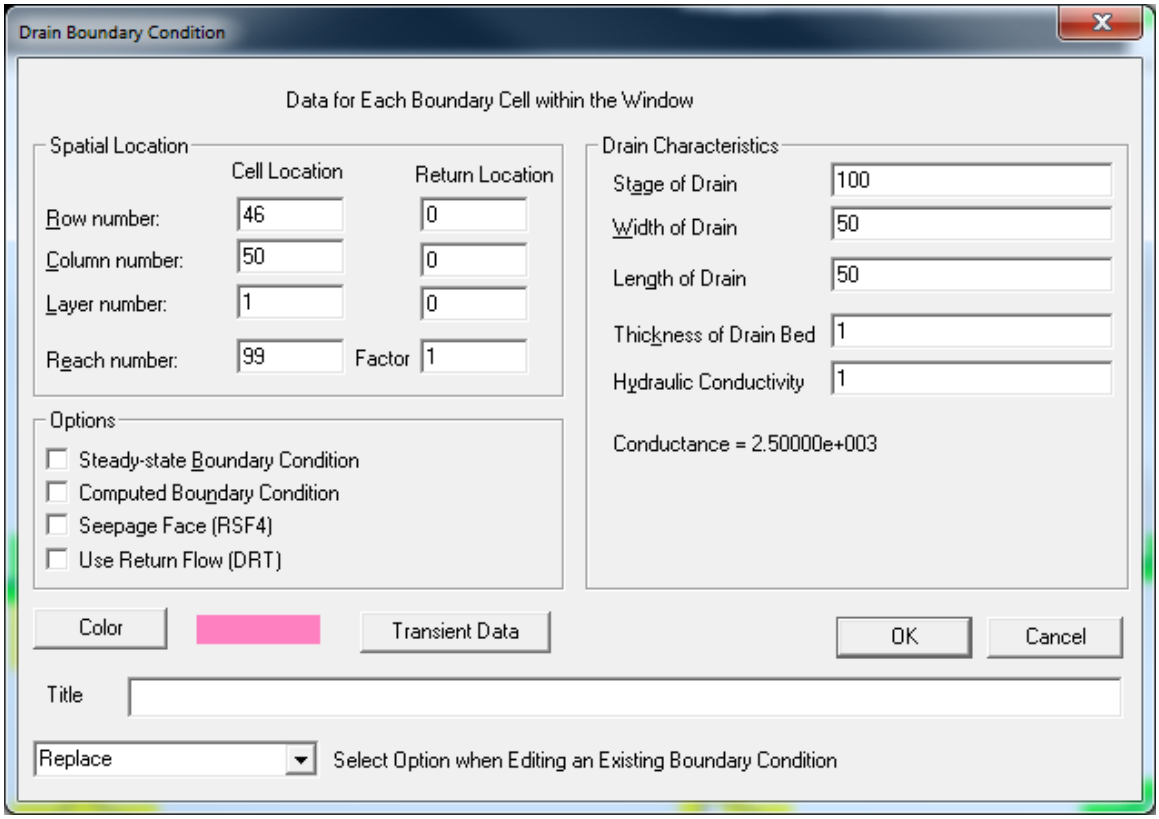

Click the "Transient Data" button to enter the drain elevations over time. The first and second columns are the starting and ending stress period for each drain Head (elevation). Start with stress period 2 and increment by 1 to stress period 8 as shown below. The head for stress period 2 is 480 and decreases by 30 m for each period. Also set the hydraulic conductivity to 1 (Kv) for each period.

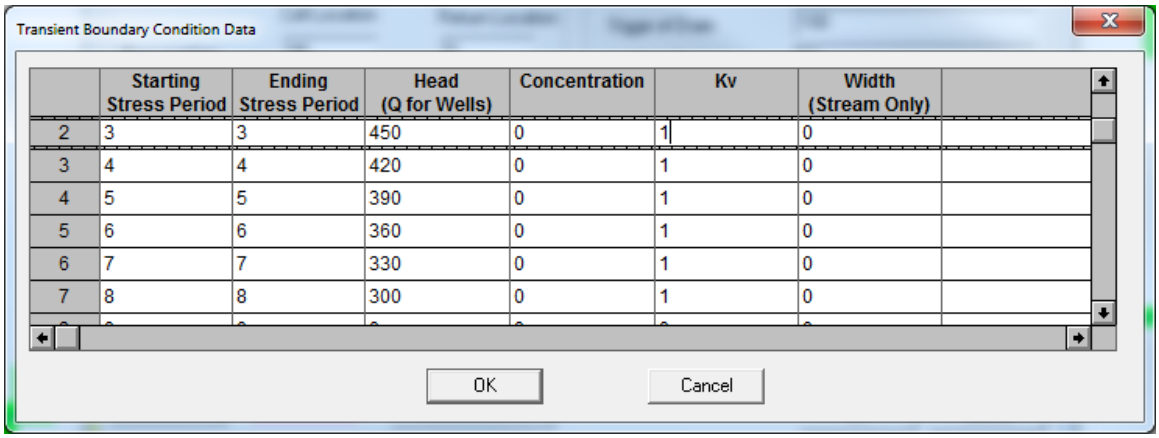

	Starting Stress Period	Ending Stress Period	Head (Q for Wells)	Concentration	Kv	Width (Stream Only)
2	3	3	450	0	1	0
3	4	4	420	0	1	0
4	5	5	390	0	1	0
5	6	6	360	0	1	0
6	7	7	330	0	1	0
7	8	8	300	0	1	0

You can now run the model. Click the calculator button and answer yes to create datasets. After the run, import the results of the last time step (default). Put a check next to drawdown on the right so that we can view the impact of the mine. Also, put a check next to the cell-by-cell flow file (*.cbb) so we can view mass balance information. After importing results, select **Plot|What to Display** and change the contour variable from

Head to Drawdown. The drawdown should look like the following (after changing the contour interval to 10 m).

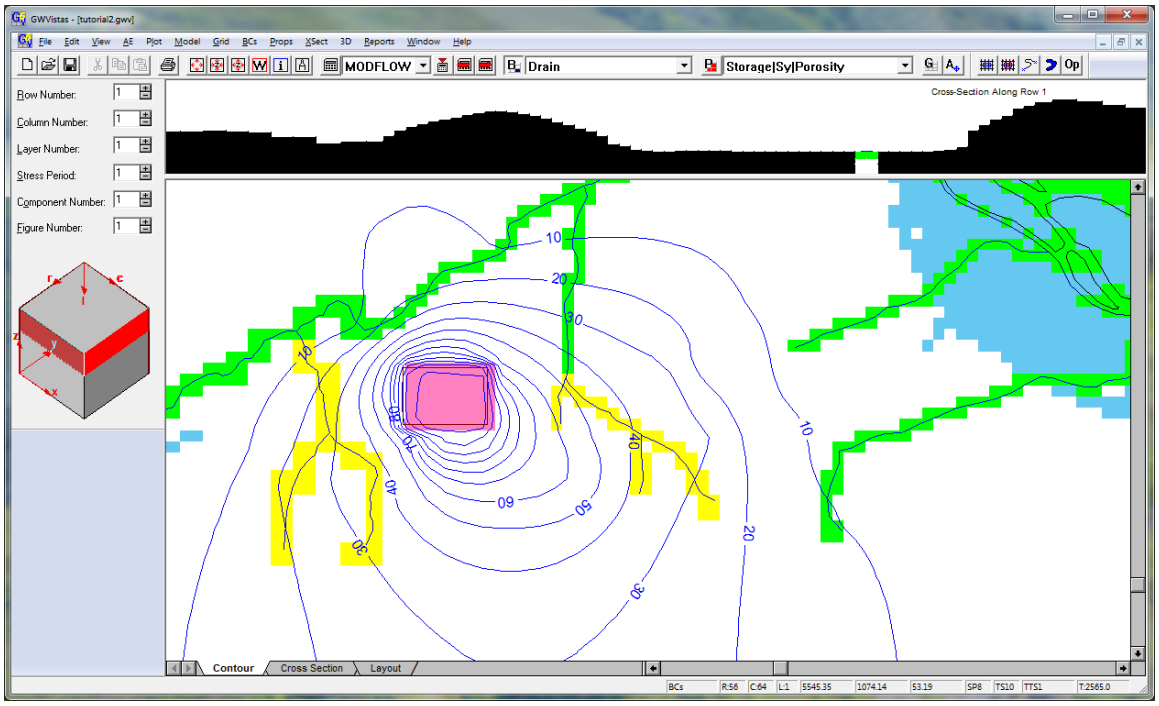

To display the amount of water pumped over time from the mine, select **Plot|Hydrograph|Mass Balance|Model Summary**. The graph below has been modified to make the drain flow a bold line (double-click the graph to modify it). You can see that the mine does not start pumping significant amounts of water until at least half-way through the life of the mine.

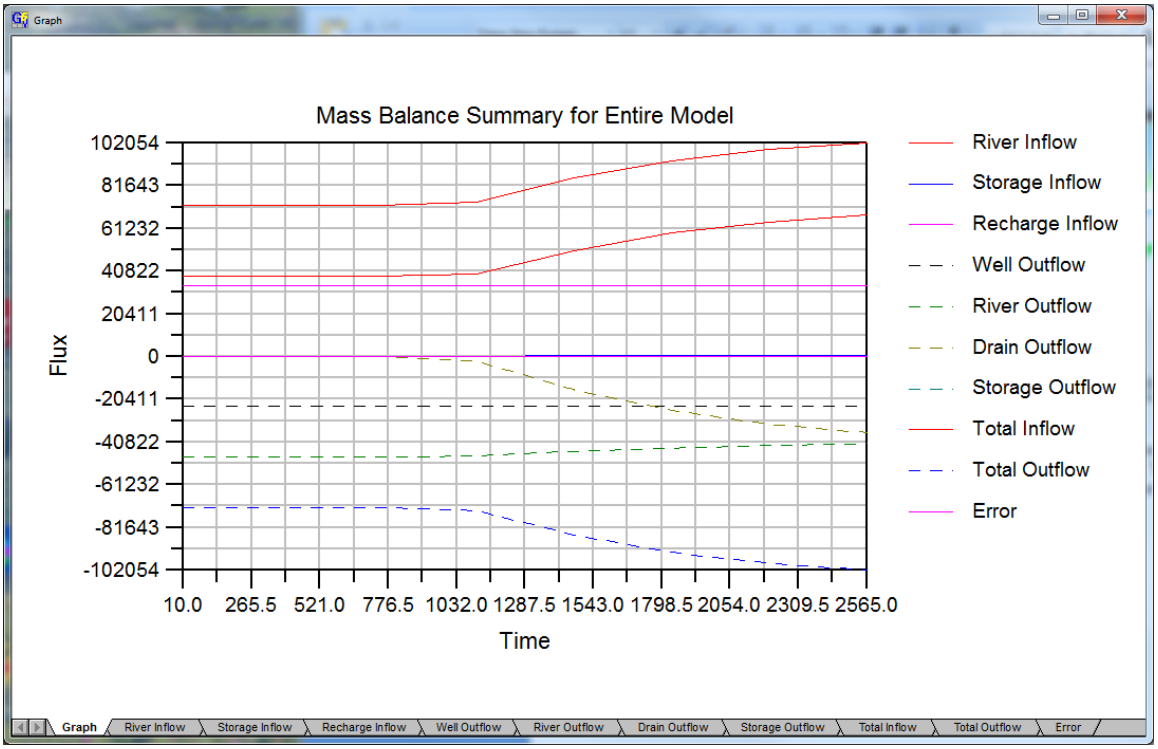

Now we will create a simulation that evaluates the length of time required to fill the mine pit after mining stops. This is most easily accomplished using the MODFLOW2000 Lake Package (LAK3). Before we do this, save the model you just created using a name like ***bedrock_example_mine1.gwv***.

The first step is to delete the drains you added to represent the mine. Select **BCs|Drain** and then **BCs|Delete|Window** and drag a window around the mine drains. After they have been deleted, add lakes using **BCs|Lake** and **BCs|Insert|Window**. Enter the initial stage at 300 m (the bottom of the mine), a minimum stage of 290 and maximum stage of 500. Change the Lake Number to 1.

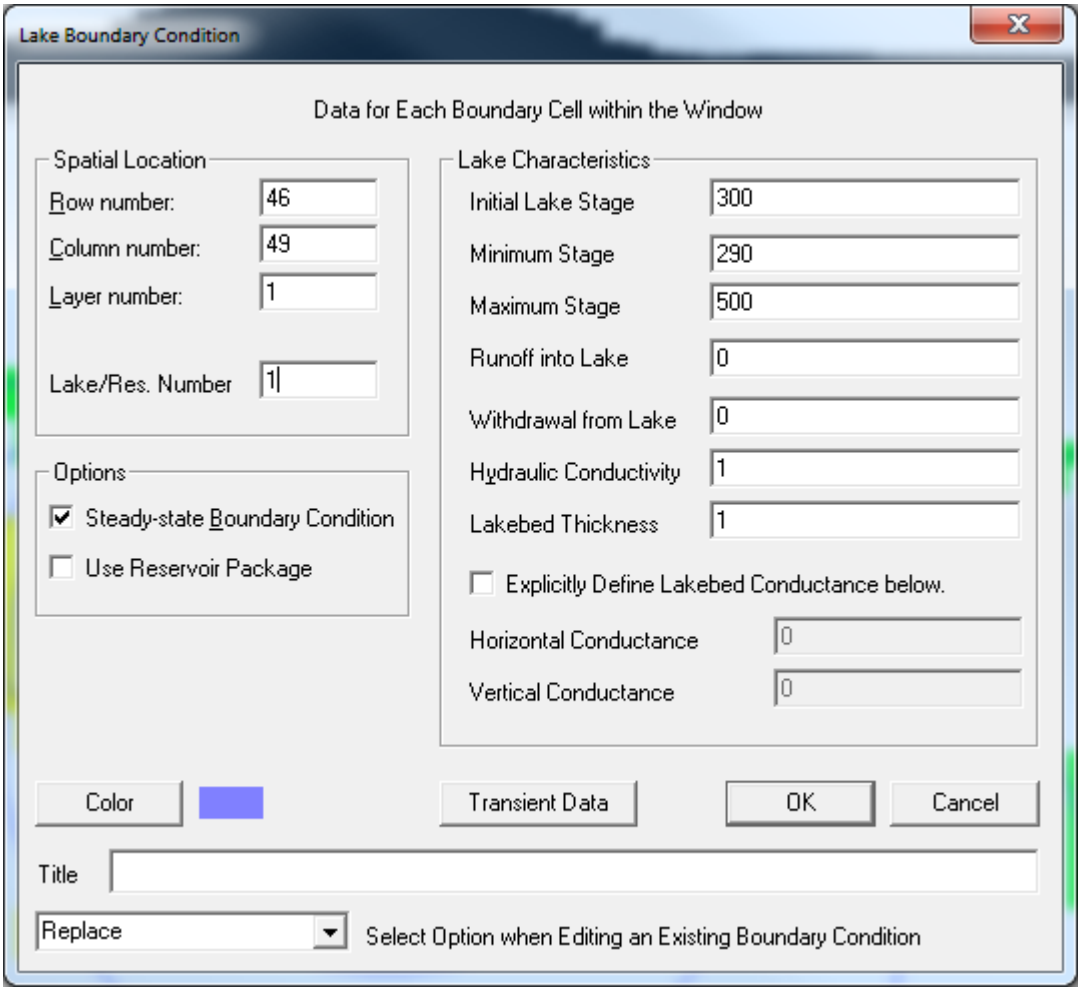
Lake Boundary Condition
Data for Each Boundary Cell within the Window
Spatial Location
Row number: 46
Column number: 49
Layer number: 1
Lake/Res. Number 1
Options
Steady-state Boundary Condition
Use Reservoir Package
Lake Characteristics
Initial Lake Stage 300
Minimum Stage 290
Maximum Stage 500
Runoff into Lake 0
Withdrawal from Lake 0
Hydraulic Conductivity 1
Lakebed Thickness 1
Explicitly Define Lakebed Conductance below.
Horizontal Conductance 0
Vertical Conductance 0
Color
Transient Data
OK
Cancel
Title
Replace
Select Option when Editing an Existing Boundary Condition

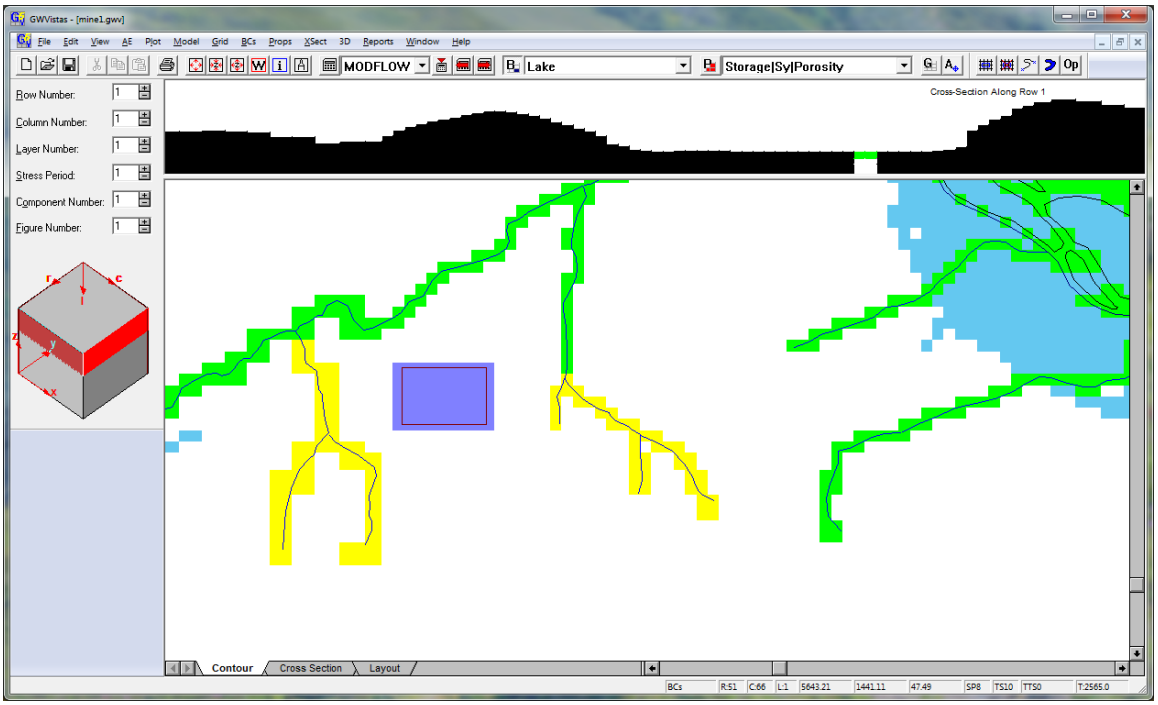
MODFLOW
Lake
Storage|Sy|Porosity
Row Number:
Column Number:
Layer Number:
Stress Period:
Component Number:
Figure Number:
Cross-Section Along Row 1
Contour
Cross Section
Layout

This second run will only have one stress period. Select **Model|MODFLOW|Packages** and give it a new root name (like ***mine2***). Select **Model|MODFLOW|Package Options** and change the number of stress periods to 1.

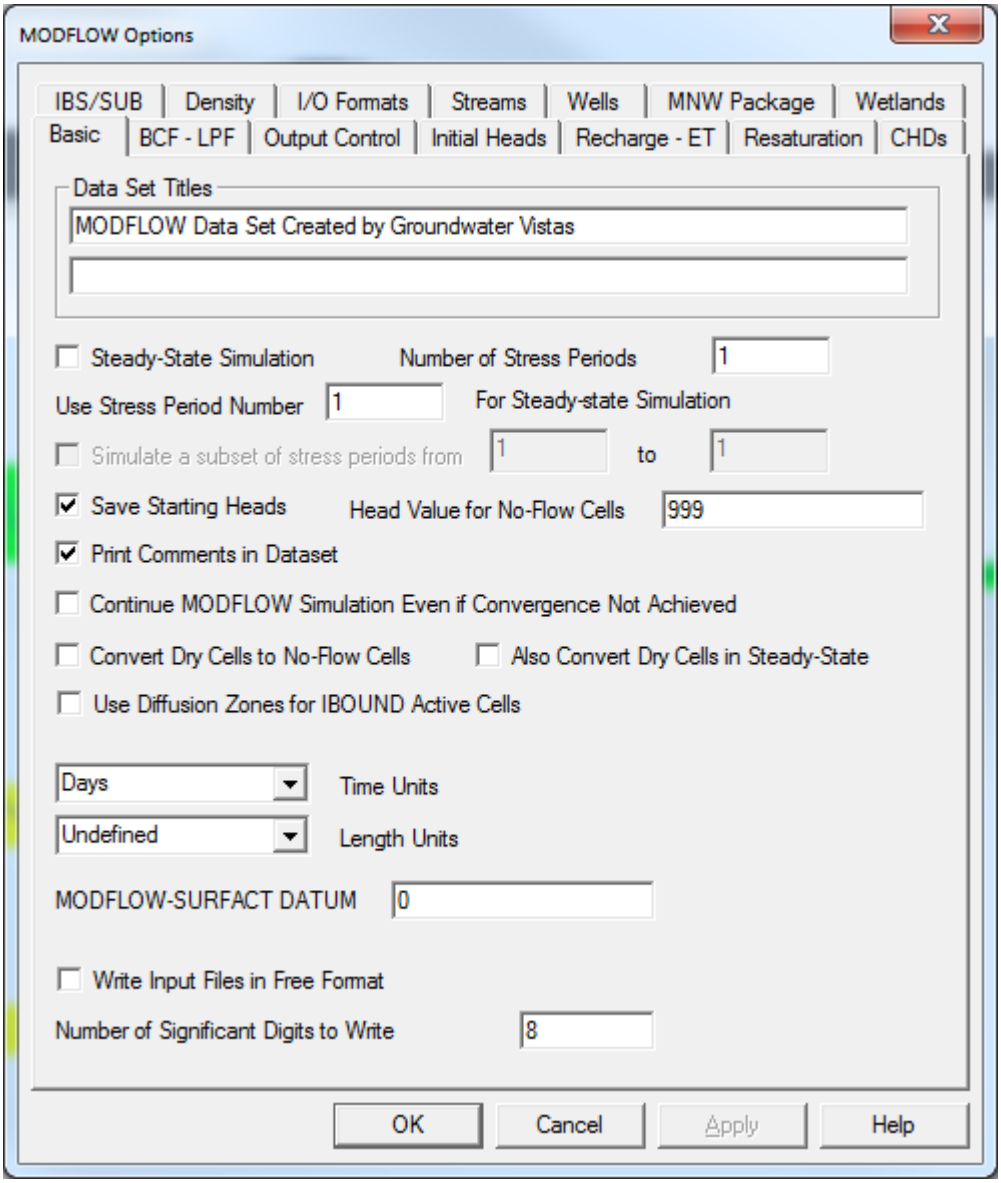

Click on the **Initial Heads** tab and change the initial heads option from "Top of Layer 1" to "Set Heads from Head-save…". Then click the browse button and find the results of the last simulation (e.g. mine1.hds). Enter stress period 8 and time step 10 beneath the head file name.

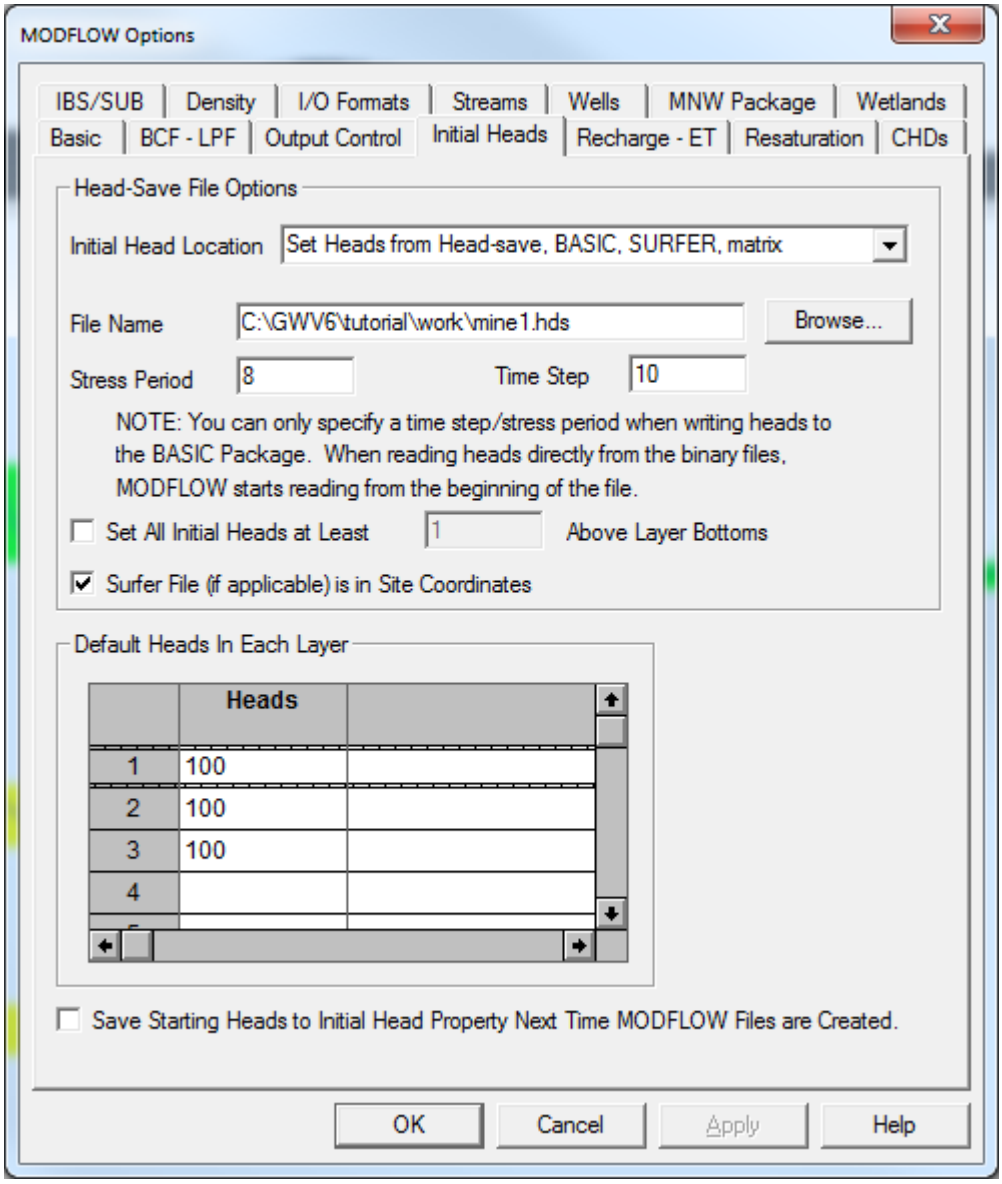

Click the **Output Control** tab and turn off the custom output control.

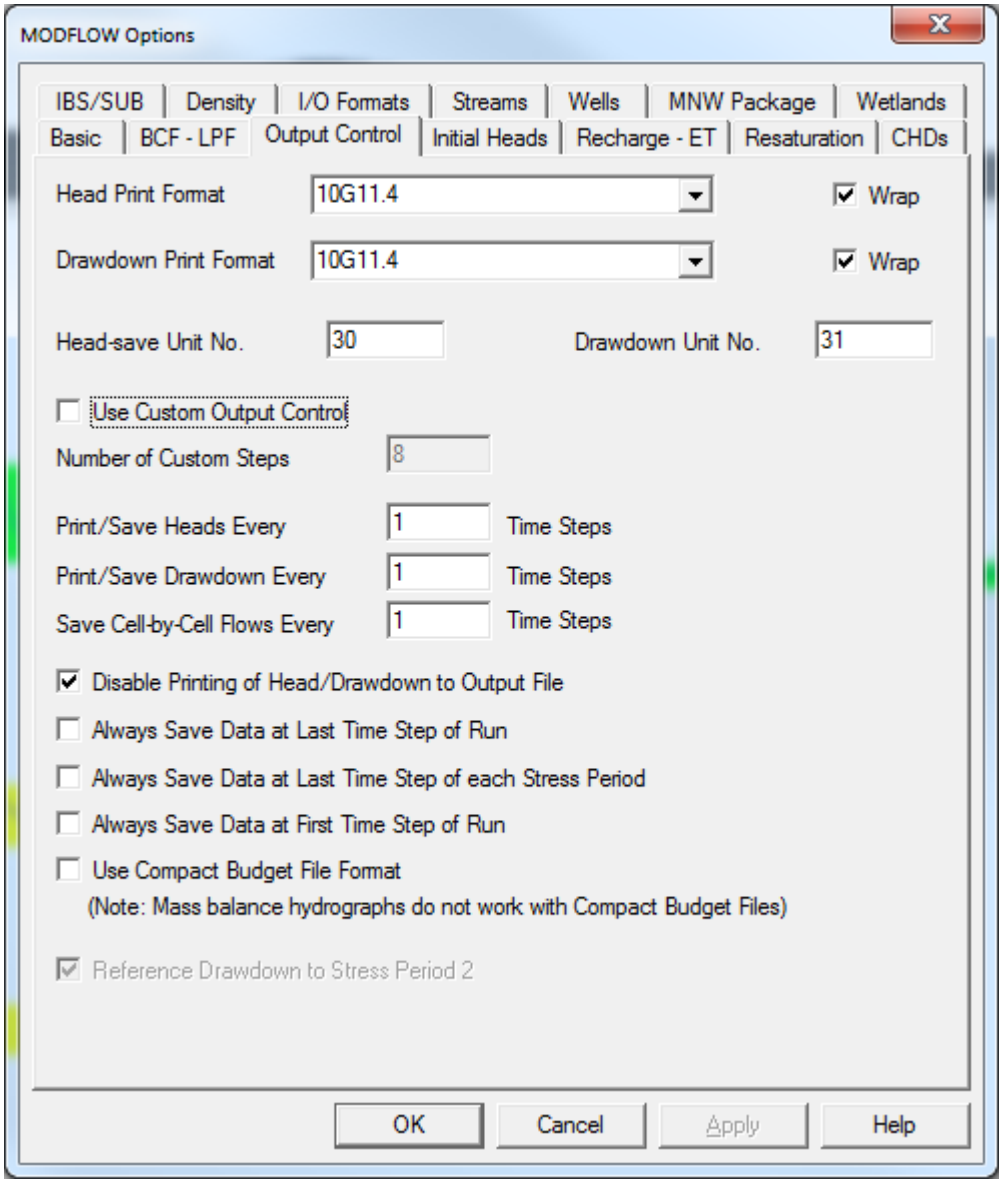

Now you need to enter a period length for the one transient stress period. Select **Model|MODFLOW|Stress Period Setup**. Use a period length of 1,825 d (5 years) and 25 time steps with a multiplier of 1.2.

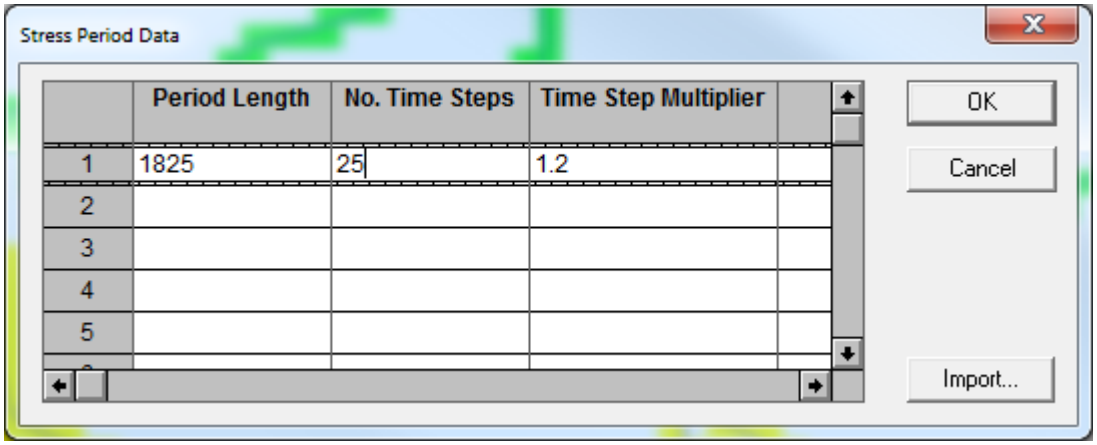

Remember that we had set up the previous simulation so that the first stress period was steady-state. We need to change that to transient. Select **Model|MODFLOW2000|Stress Period Types|Edit Stress Period Types** and make the period type 1.

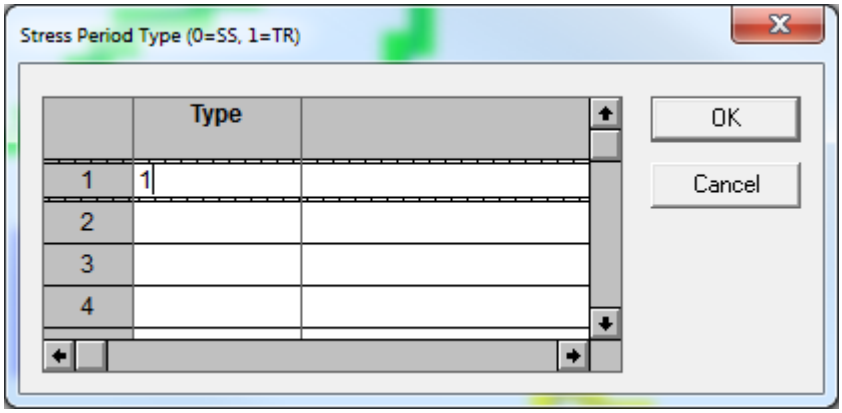

Finally, we want to be able to plot the change in lake level with time. Add a monitoring well inside the lake. Select **AE|Well** and put the well anywhere inside the lake cells. Give it a name and make the pumping rate zero. Check the option to monitor head|concentration over time.

Now run the model you just created in MODFLOW2000. Import the results of the last time step. To view the change in lake stage over time, select **Plot|Hydrograph|Monitoring Well** and pick the one you just created above. You should see something like the following.

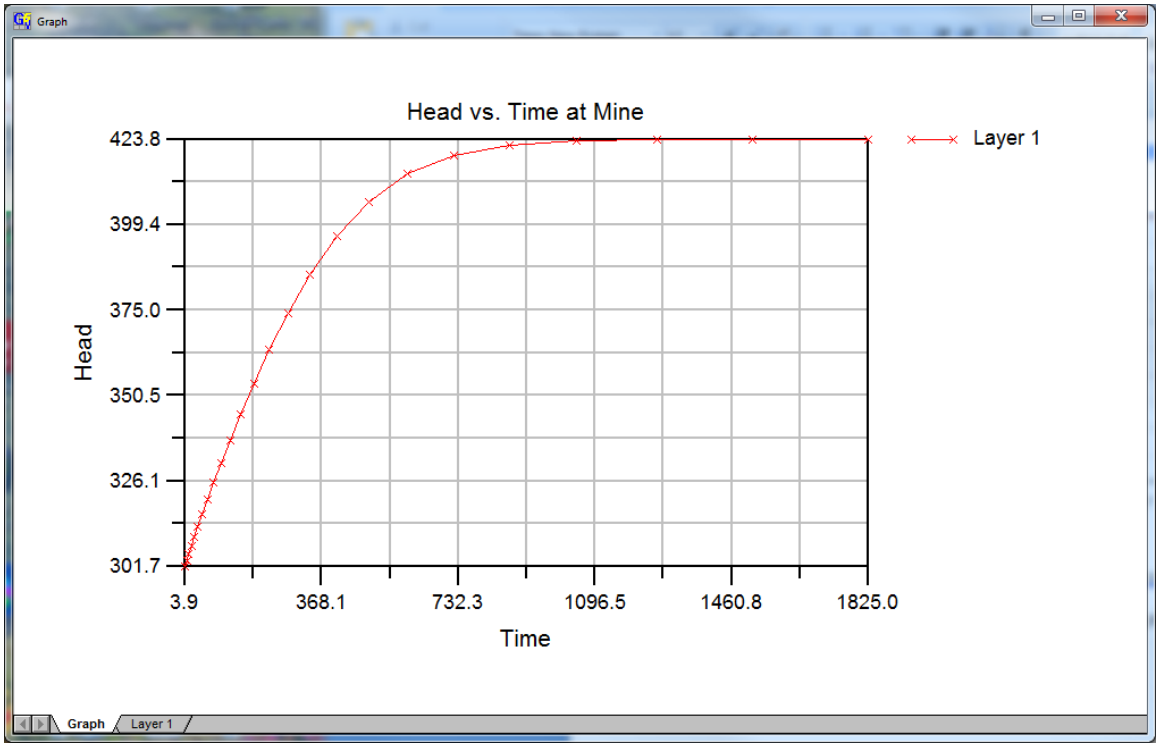

This graph shows that it will take about 3 years for the mine pit to fully recover after mining ceases.

Impacts from Water Supply Development

Evaluating the impacts from developing new groundwater supplies is similar in concept to the mining example we just simulated. The goal is to determine the impact on nearby wells and also on surface water.

We will start from the first model you created. If you did not save that model, you can use gwv6\tutorial\bedrock_example.gwv. Open that model now.

In this example, the first stress period will represent steady-state conditions prior to additional pumping. New wells will then be added in a 1-year stress period. Start by selecting **Model|MODFLOW|Packages** and change the MODFLOW root file name to ***well1***. Next select **Model|MODFLOW|Package Options** and remove the steady-state check-box and change the number of stress periods from 1 to 2. Click OK when you are done.

Groundwater Vistas will then ask if you want to copy Recharge from stress period 1 to the other stress period. Answer Yes to this and all other prompts. The last question asks if you want to edit the stress period setup. Enter 365 for the second stress period. Use 10 time steps, with a multiplier of 1.2. The multiplier of 1.2 is used to start each stress period with a small time step.

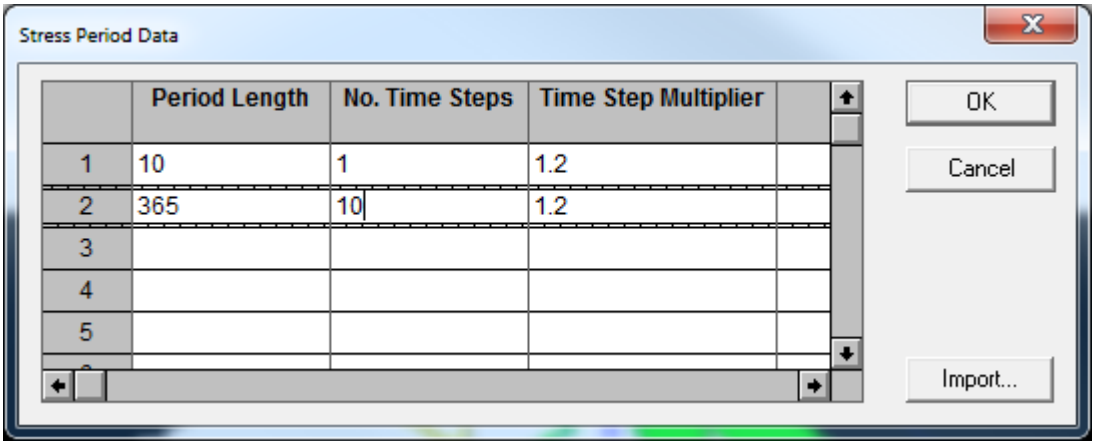

MODFLOW2000 has a very useful feature in that each stress period can be steady-state or transient. In this example, the first stress period is steady-state and the second one is transient. GV should set this properly in the next dialog it displays. The first period should be zero (steadystate) and second 1 (transient).

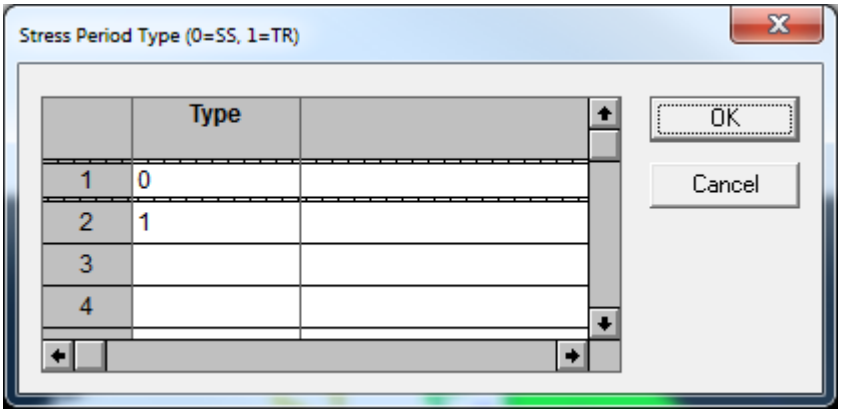

We will now take advantage of an undocumented feature in MODFLOW2000 to compute drawdown from the results of a particular time step just as we did in the mine dewatering example in the last section. In previous versions of MODFLOW, drawdown was always computed from the starting heads for the simulation. In the current model, we would like to compute drawdown from new pumping wells from the results of the steady-state stress period. To do this, select **Model|MODFLOW|Package Options – Output Control tab**. Turn on the option for custom output control and enter 11 for the number of custom steps.

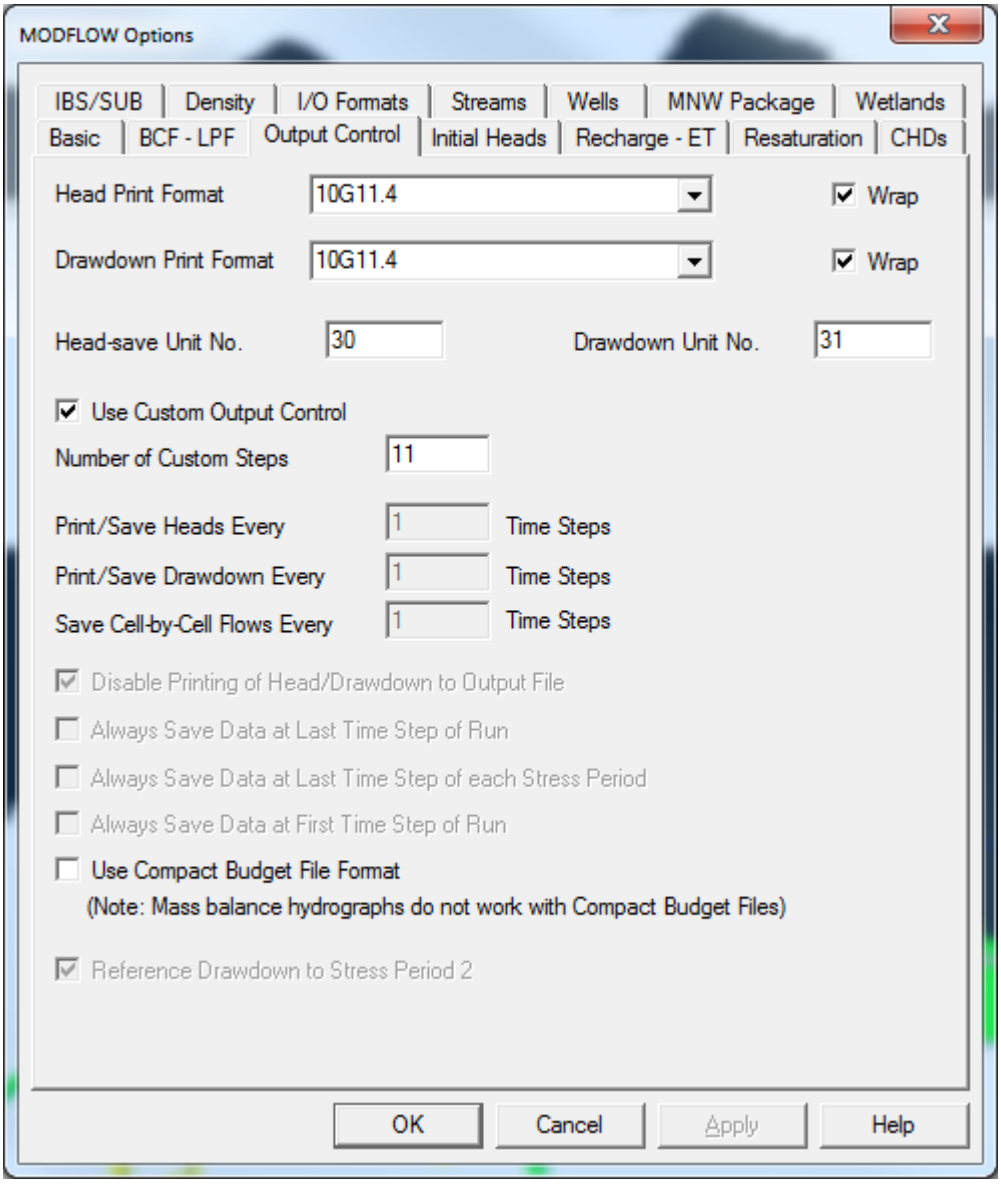

Now, select **Model|MODFLOW|Custom Output Control**. Fill out the table as shown below. In particular make sure to put a 1 in the first row beneath the heading "Ddn Ref". This option places the key word ***DDREFERENCE*** in the output control file in the first time step, which in turn causes MODFLOW to compute drawdown from that first stress period.

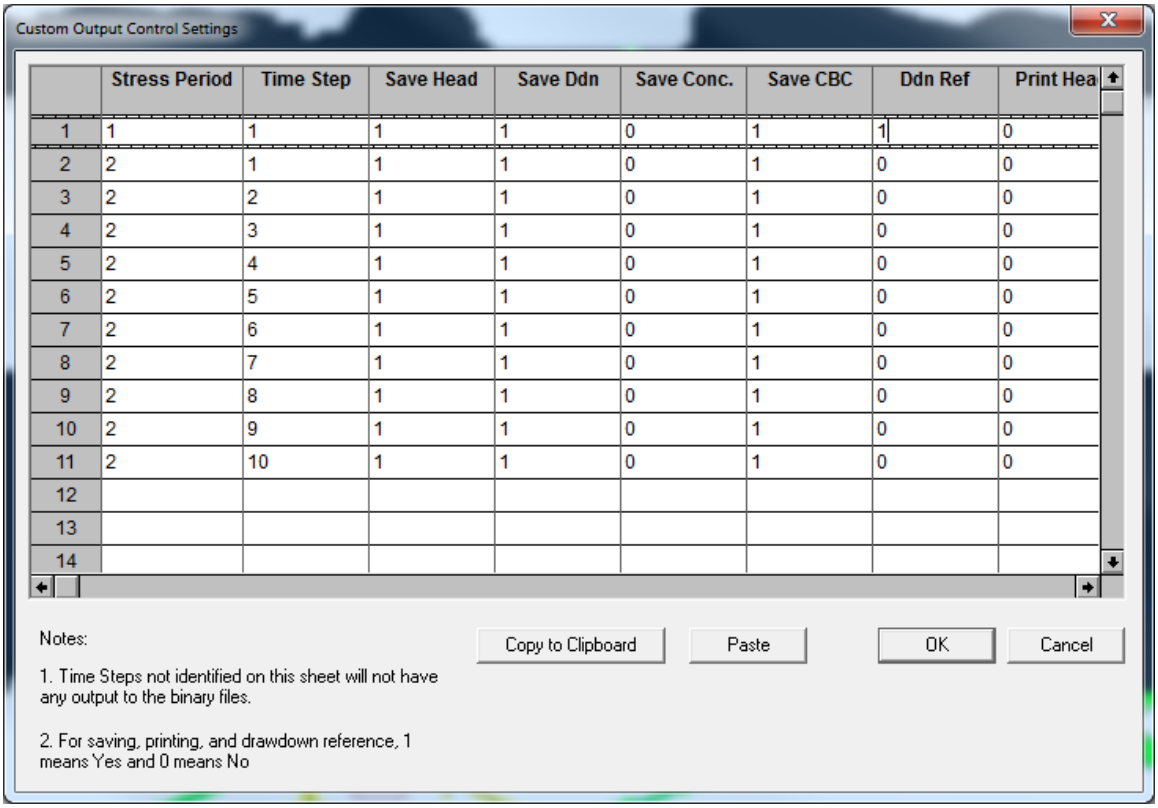

Because the simulation will have transient stress periods, we need to enter values for specific storage and specific yield. Select **Props|Storage** and **Props|Property Values|Database**. Enter 1.0e-5 for specific storage and 0.01 for specific yield, as shown below.

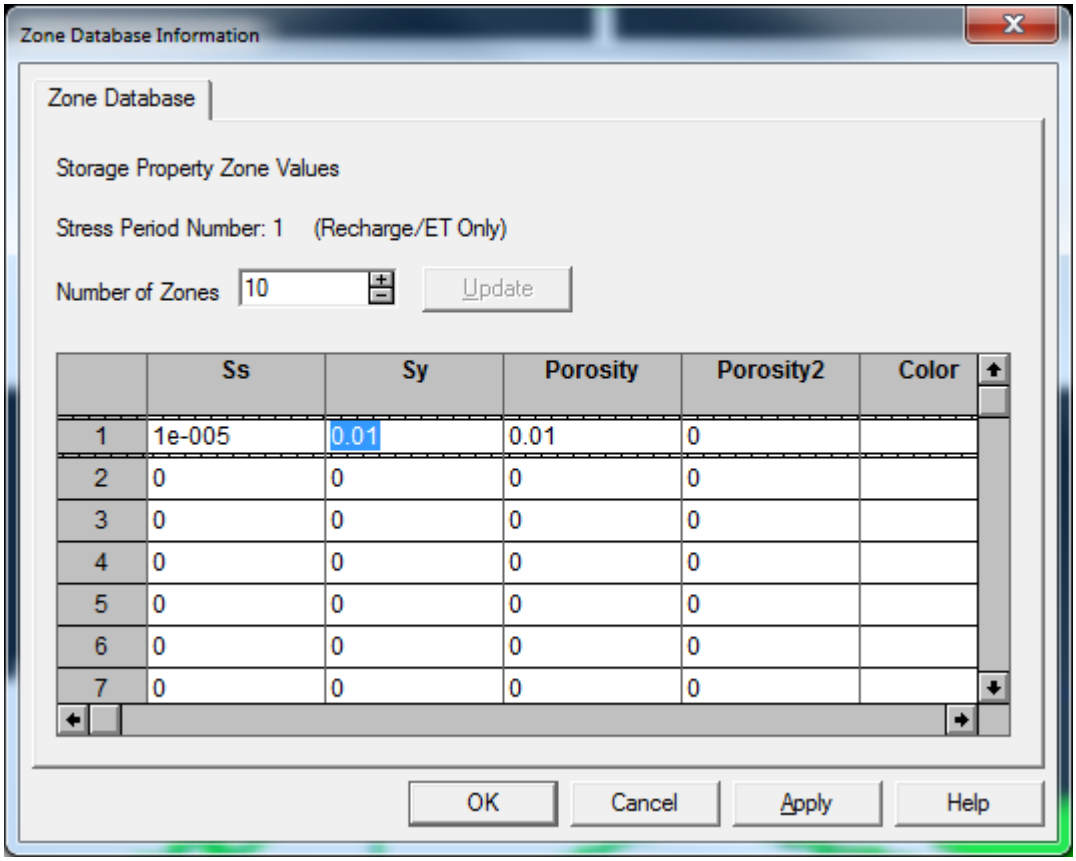

Now we will enter a new well that will only pump in stress period 2. Select **AE|Well** and put the well near row 32, column 85. Pump from layers 1 and 2. Give the well a name, such as ***New Well***. Uncheck the steady-state flag and then click the ***transient data*** button. Enter a rate of -15,000 m^3|d for the second stress period.

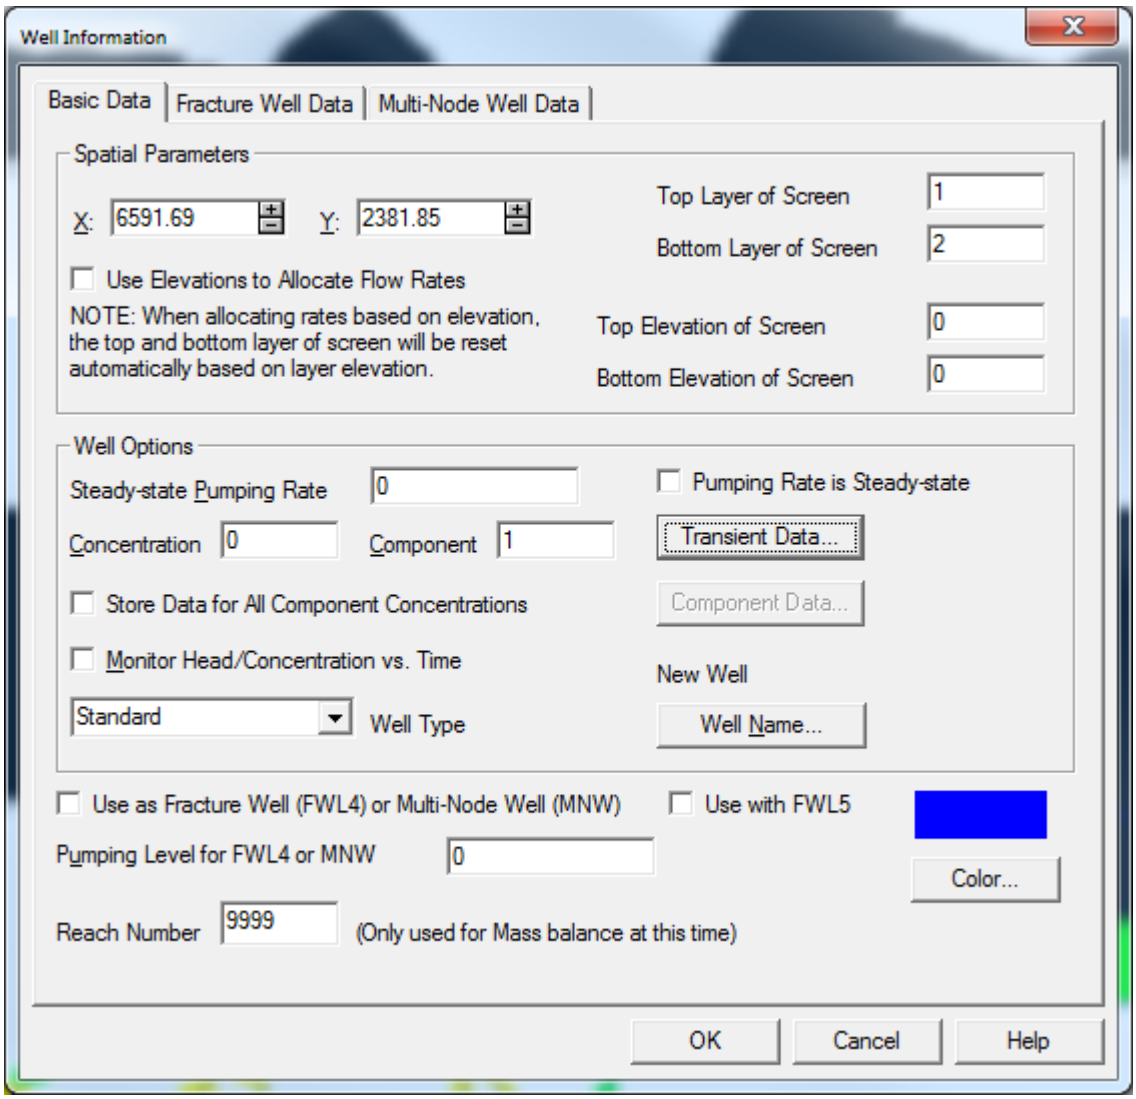

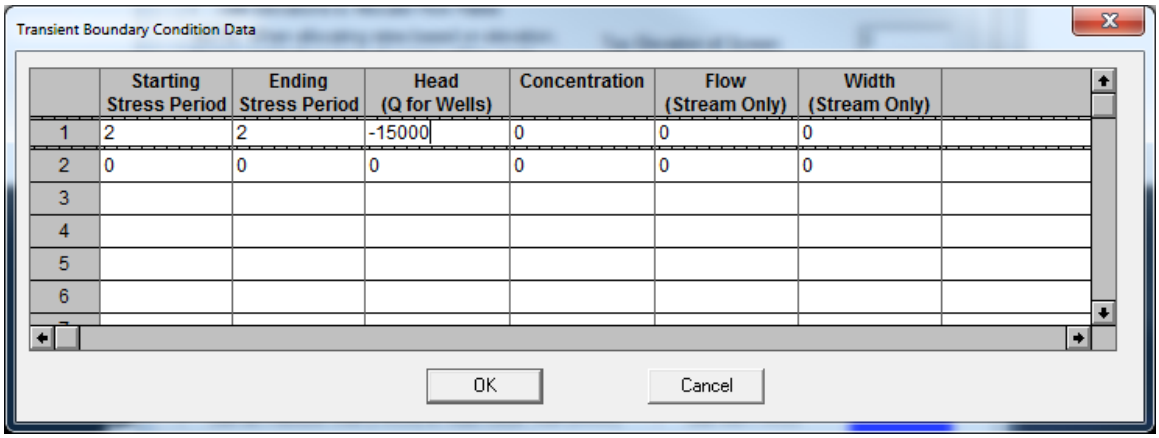

You can now run the model. Click the calculator button and answer yes to create datasets. After the run, import the results of the last time step (default). Put a check next

to drawdown on the right so that we can view the impact of the mine. Also, put a check next to the cell-by-cell flow file (*.cbb) so we can view mass balance information. After importing results, select **Plot|What to Display** and change the contour variable from Head to Drawdown. The drawdown should look like the following (after changing the contour interval to 2 m).

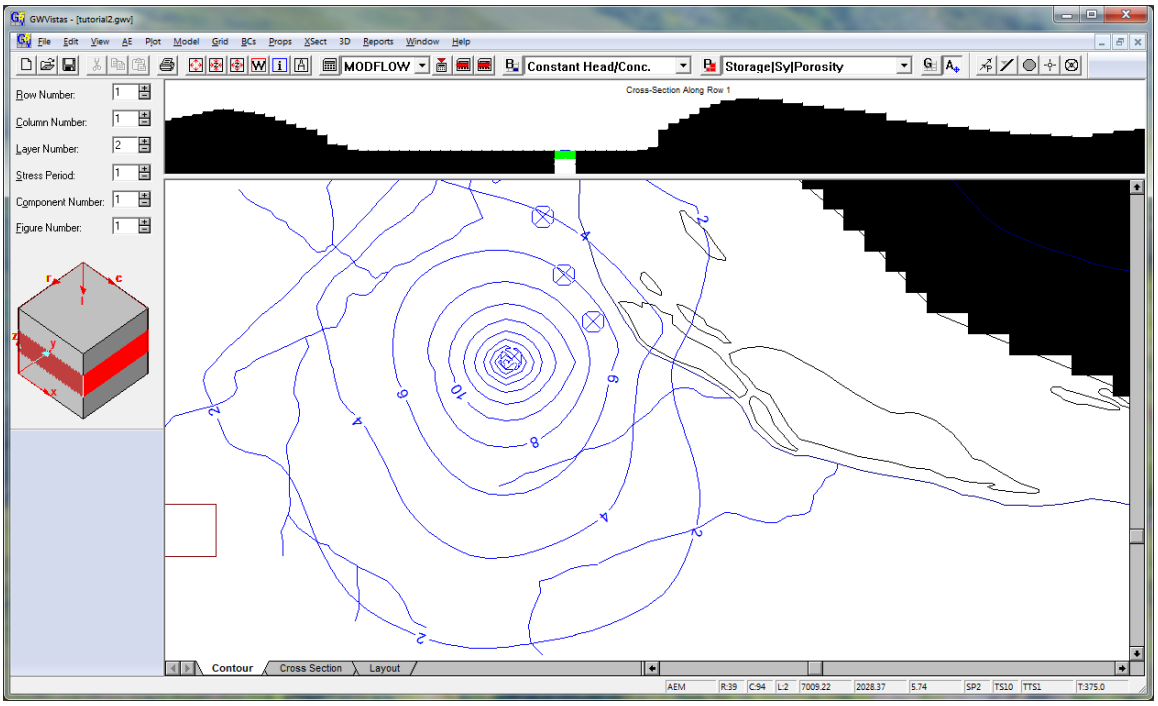

The drawdown predicted by the model is about 4 to 7 meters at the existing wells. In many parts of the USA, that would be too much impact to allow the well to operate. Another issue that often causes problems is the amount of water removed from surface water.

To display the amount of water removed from the river by this new well, select **Plot|Hydrograph|Mass Balance|Model Summary**. The graph below has been modified to make the river flow a bold line (double-click the graph to modify it). You can see that about half of the new well pumping comes from the river (bold dashed line is the river outflow, lighter dashed line above it is the well outflows). Depending upon the amount of surface flow, that might be too much to allow the well to pump.

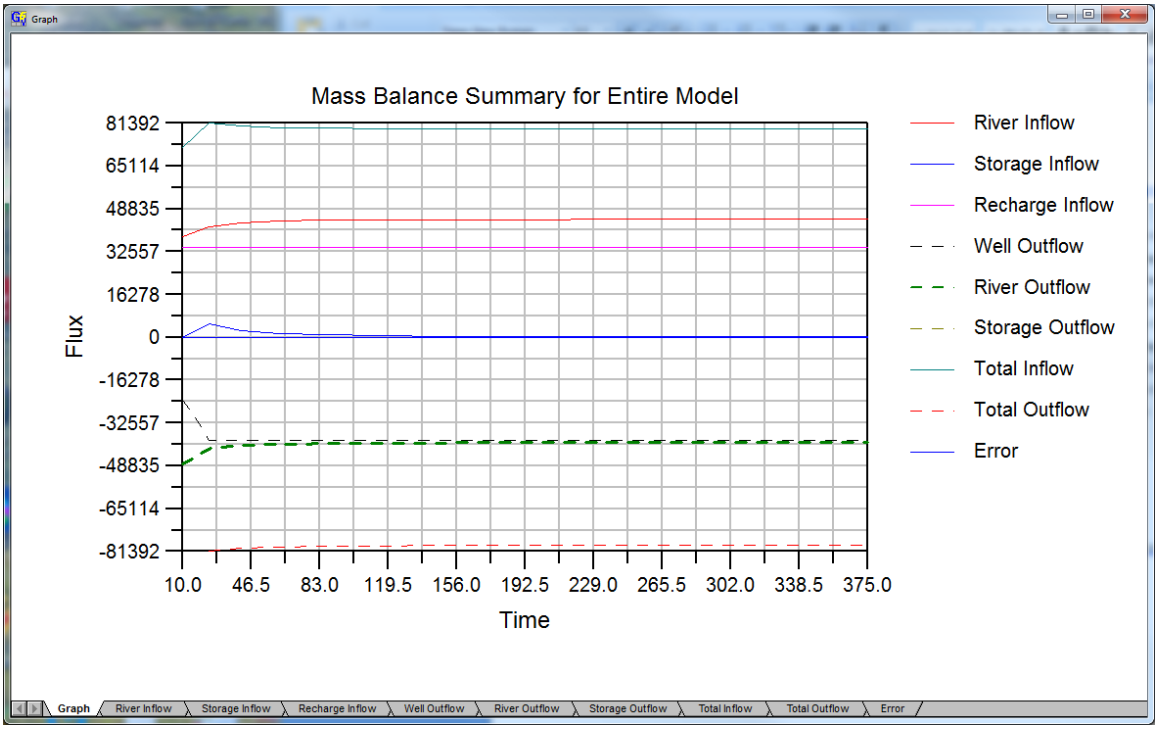
Graph
Mass Balance Summary for Entire Model
81392
65114
48835
32557
16278
0
-16278
-32557
-48835
-65114
-81392
Flux
10.0
46.5
83.0
119.5
156.0
192.5
229.0
265.5
302.0
338.5
375.0
Time
River Inflow
Storage Inflow
Recharge Inflow
Well Outflow
River Outflow
Storage Outflow
Total Inflow
Total Outflow
Error
Graph
River Inflow
Storage Inflow
Recharge Inflow
Well Outflow
River Outflow
Storage Outflow
Total Inflow
Total Outflow
Error

Effect of Boundary Condition Selection

Introduction

This exercise illustrates how different boundary conditions can be used to create the same model with the same computed water levels. Three different models will be designed and the same predictive scenario will be run with each design. You will see that although the models all produce the same results in an unstressed state, the predictions when the model includes a pumping stress are all very different. This is one example of model non-uniqueness. As part of the exercise, you will also learn new tools within Groundwater Vistas for designing groundwater models.

The model will be a simple confined aquifer with a uniform gradient from east to west. If this were a real-world study, the model could be thought of as a small piece of a regional aquifer where there is limited data outside the area of interest and no regional hydrologic boundaries close to the site. This situation is actually quite common in modeling where you only have data in a very small part of a much larger system. In that case, the modeler will often keep the model small and create a design that allows groundwater to flow in one side of the model and out the other based on an observed flow direction and gradient. The figure below shows conceptually what is being simulated.

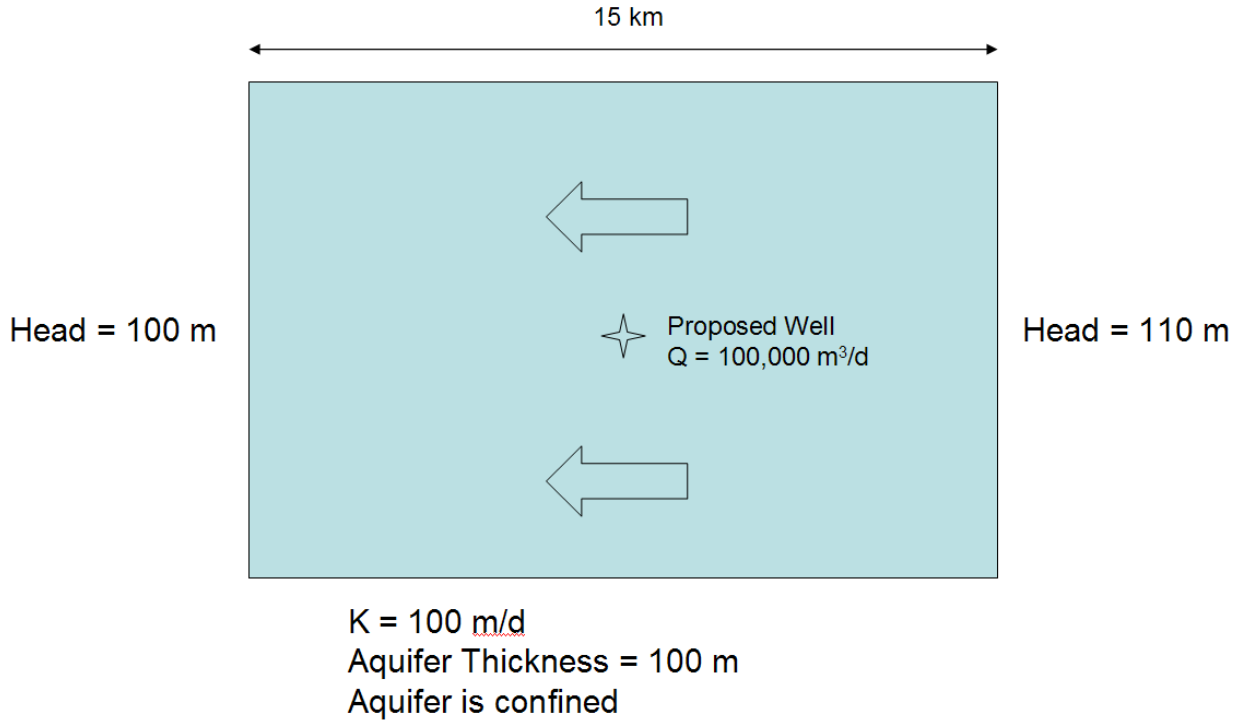

Model 1: Upgradient and Downgradient Constant Heads

Start by selecting **File|New** from the Groundwater Vistas menu. Change the number of columns to 75 and the X and Y spacing to 200 m. The remaining options can be left at their default settings. This will create a one-layer model with K=100 m|d, aquifer thickness of 100 m, 50 rows, and 75 columns.

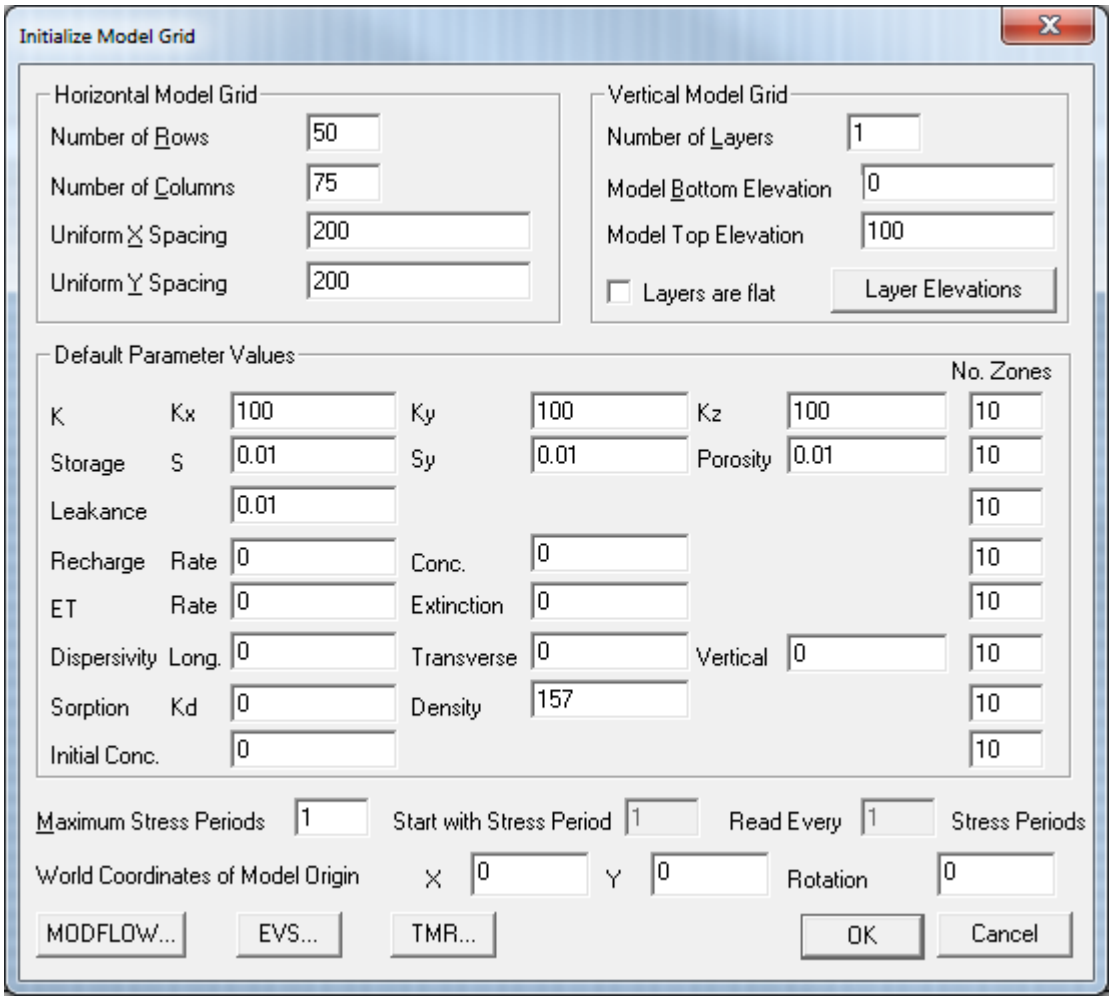

Click OK when you are finished entering the data described above.

In the first version of this model you will add constant head cells on the east and west sides of the model. The head on the west is 100 m and the head on the east is 110 m. Start by selecting **BCs|Constant Head** and then select **BCs|Insert|Window** or click the  button on the toolbar. Now, move the cursor somewhere inside the upper left (northwest) cell in the model, hold down the left mouse button, and drag a window so the end point is inside the lower left (southwest) cell in the model. Release the mouse button and the following dialog appears.

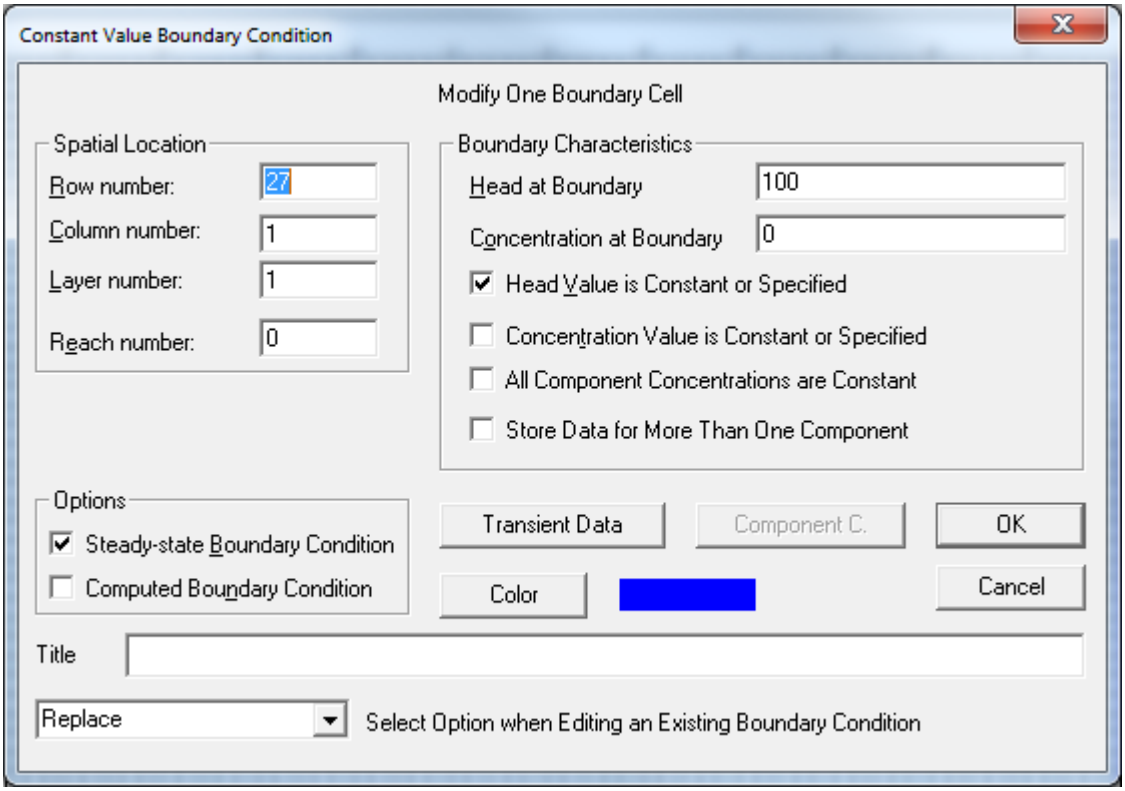

Enter a value of 100 m in the dialog in the field labeled “Head at Boundary”, as shown above. Click OK when you are done. You should now see a line of blue constant head cells along the western side of the model.

Repeat this procedure along the right (eastern) edge of the model and use a head value of 110 m.

The model should look like the one below.

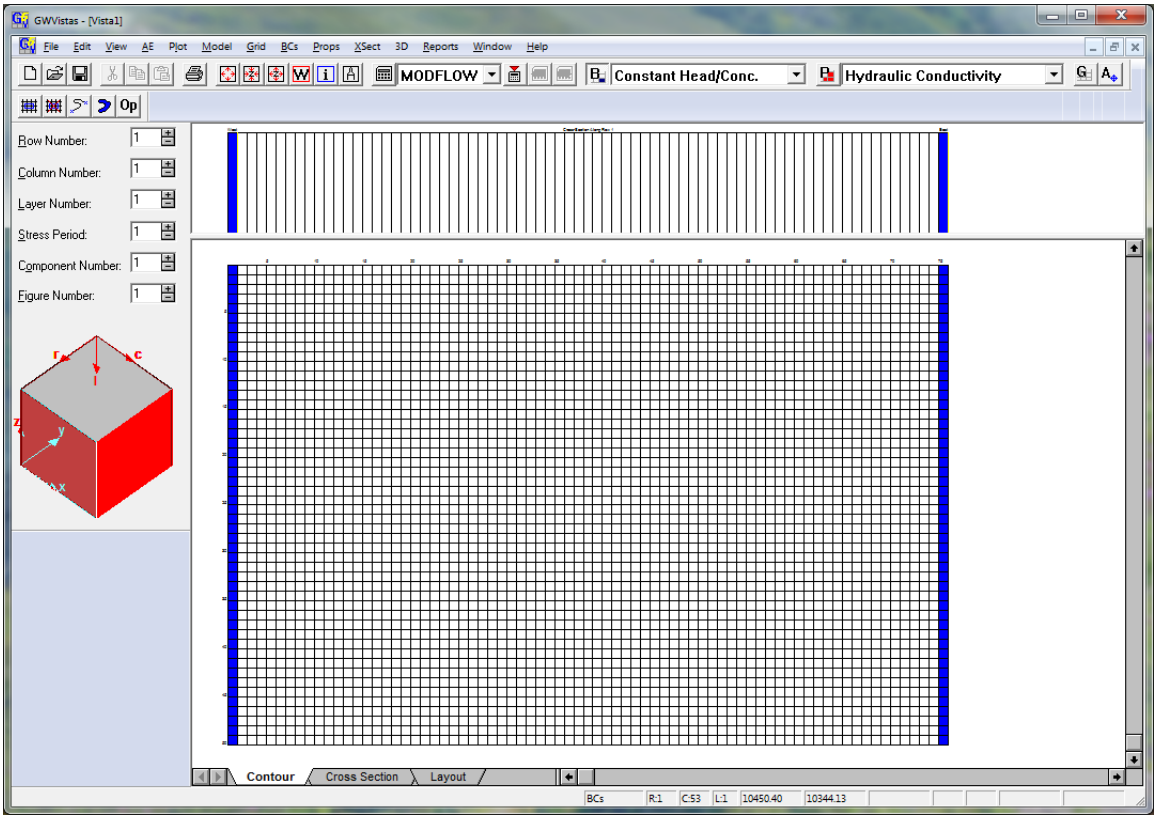

Since the aquifer properties are homogeneous the only thing we have to do to complete the model is give it a root file name (for the MODFLOW input files) and change the aquifer type to confined. First, select **Model|MODFLOW2000|Options**. Change the layer type for layer 1 from 1 (unconfined) to 0 (confined).

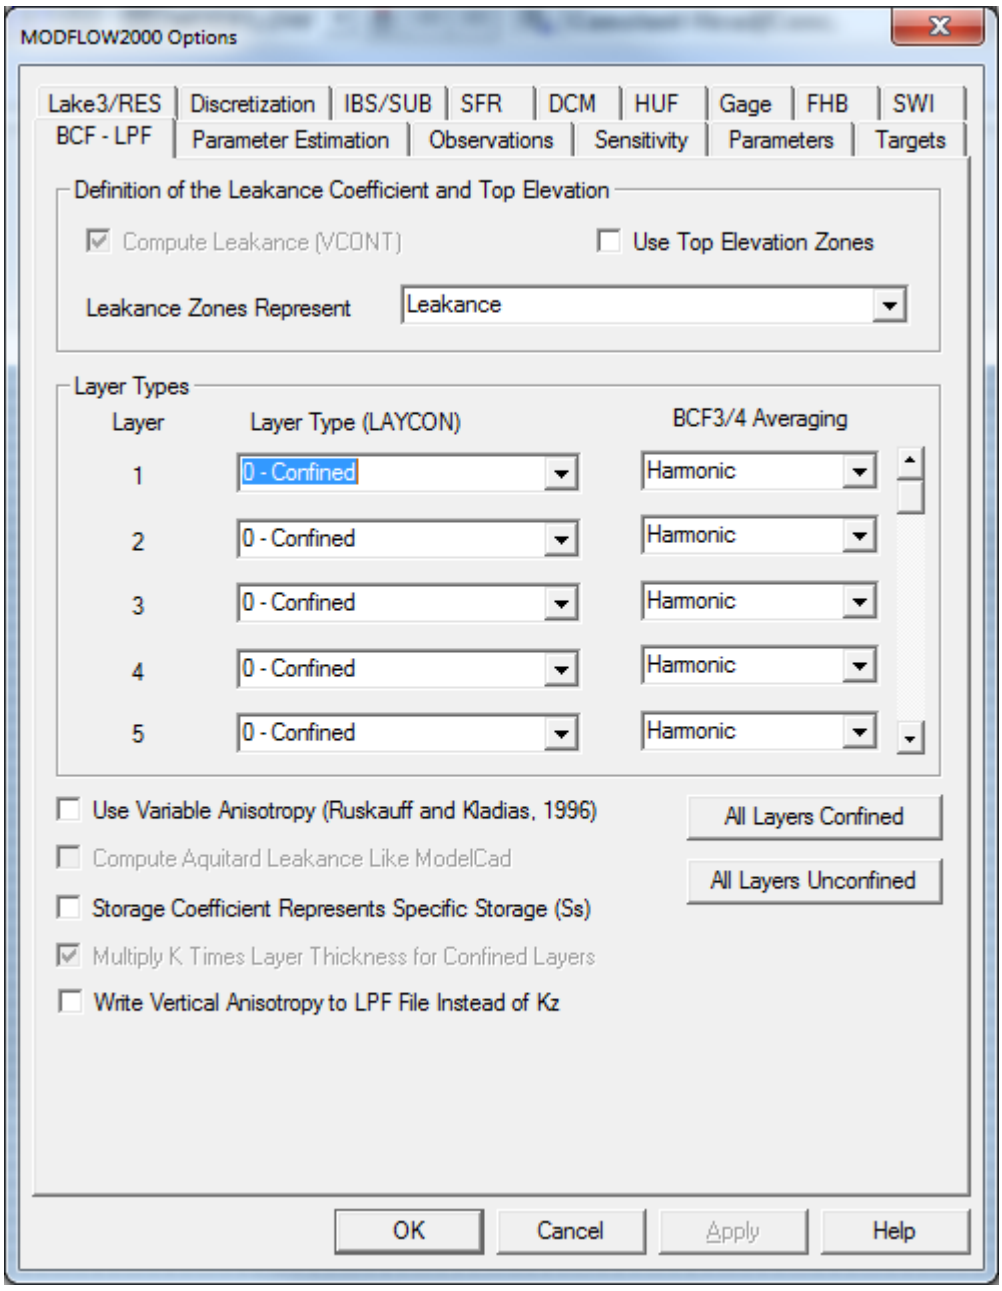

Now, select **Model**|**MODFLOW**|**Packages** and change the root name at the top to ***bctest1***. Click OK when you are done.

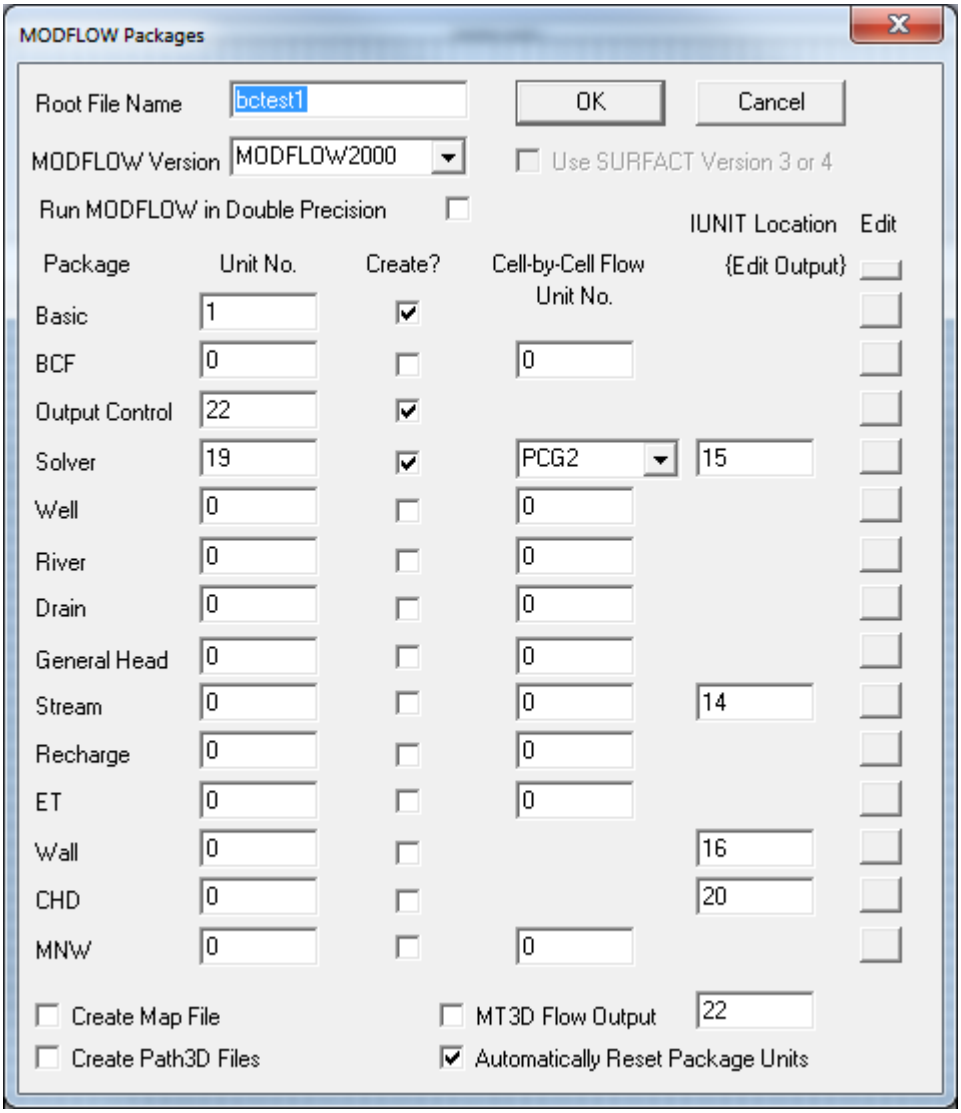

The model is ready to run. Simply click the calculator button on the toolbar  , and answer Yes to create datasets. Then you can answer No to viewing the error|warning file. MODFLOW should run very fast. GV will then ask if you want to import results. Select Yes and click OK to just accept the default results. Select **Plot|What to Display** and turn off display of the grid so you can see the contours better. Also, select **Plot|Contour|Parameters (plan)** and change the minimum, maximum, and contour interval to 100, 110, and 1, respectively. You may also need to increase the font size to 8 points. After all that, your model should look like the one below.

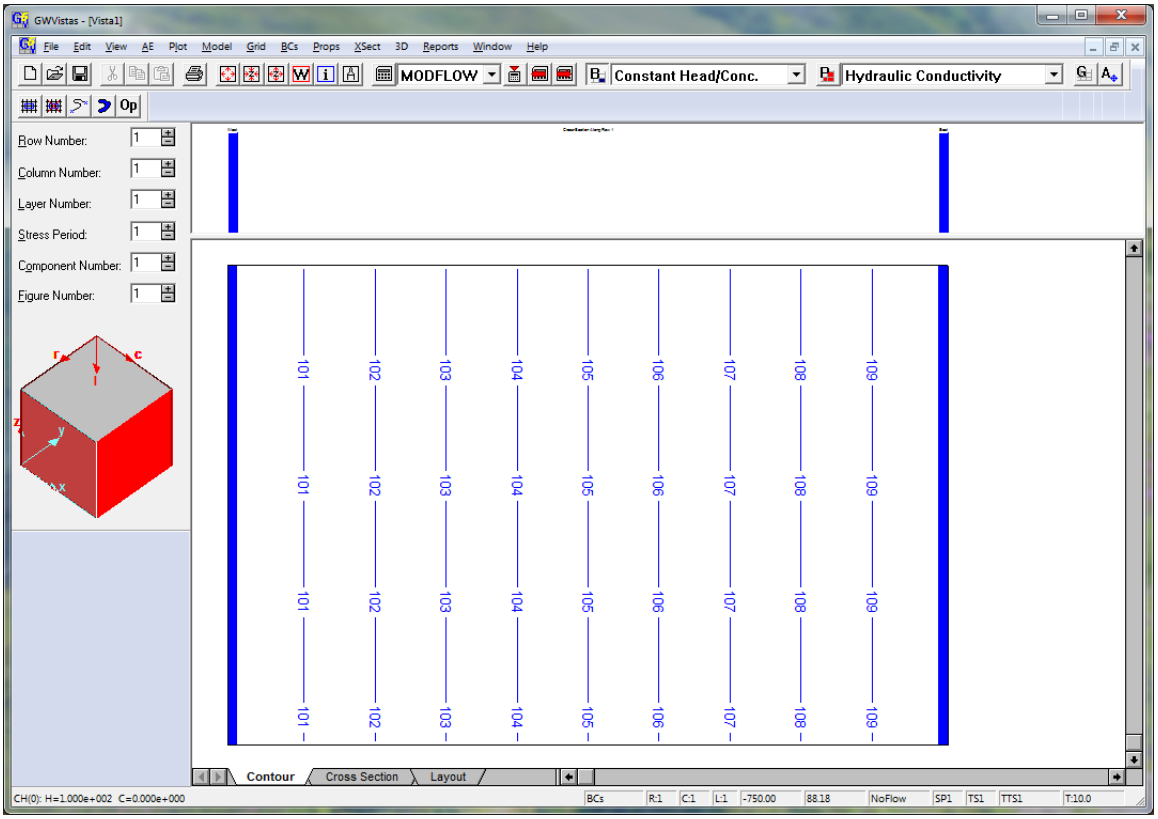

You should see a uniform gradient from west to east, as shown above. Save this model design using **File|Save As** and give it a name like ***BCTEST1.GWV***. We will come back to this model design later.

Model 2: Constant Heads Surrounding the Model

The second model will have constant heads surrounding all sides of the model grid. While this is not the best approach to creating a model, it is a common design strategy. Use the procedure from the last model to add another line of constant heads on the north side of the model and also on the south side. When you are done, you should have blue cells surrounding your model. Do not worry about the head value you enter for these constant heads. We will be using a special feature of Groundwater Vistas to apply a uniform gradient across all constant head cells after you have entered them. Your model should now look like the one below.

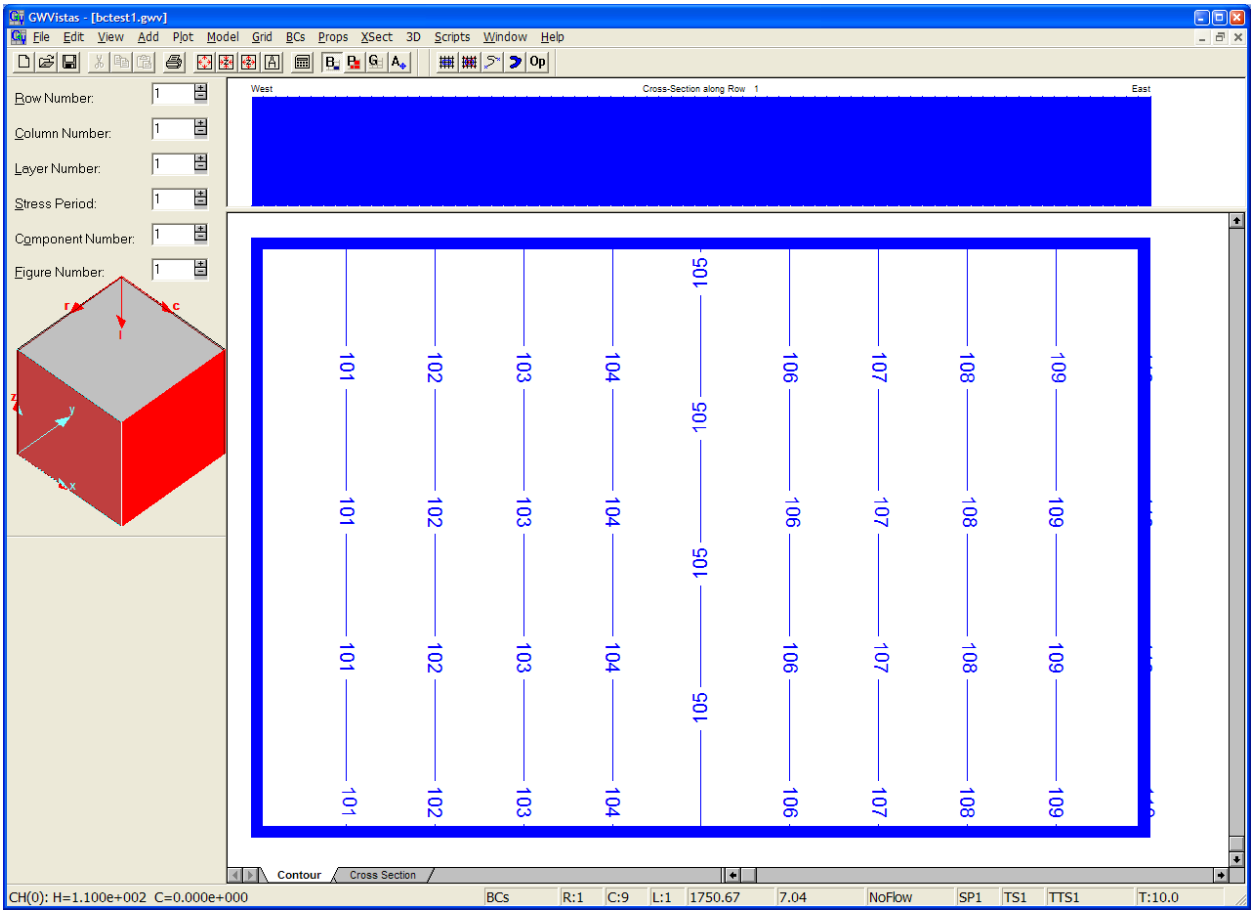

To achieve the same uniform gradient, you will now use the "Gradient Fill" option in Groundwater Vistas. Select **BCs|Modify|Gradient Fill**. Make sure the Fixed Head Value field says 100. Change the gradient to 0.000675. Leave everthing else at the default value. The data on this dialog tells Groundwater Vistas that the head value in row 1, column 1 (northwestern-most cell) is 100 m. The gradient is 0.000675 m|m (dimensionless). The fact that the gradient is positive means that heads increase to the east. Click OK when you are done.

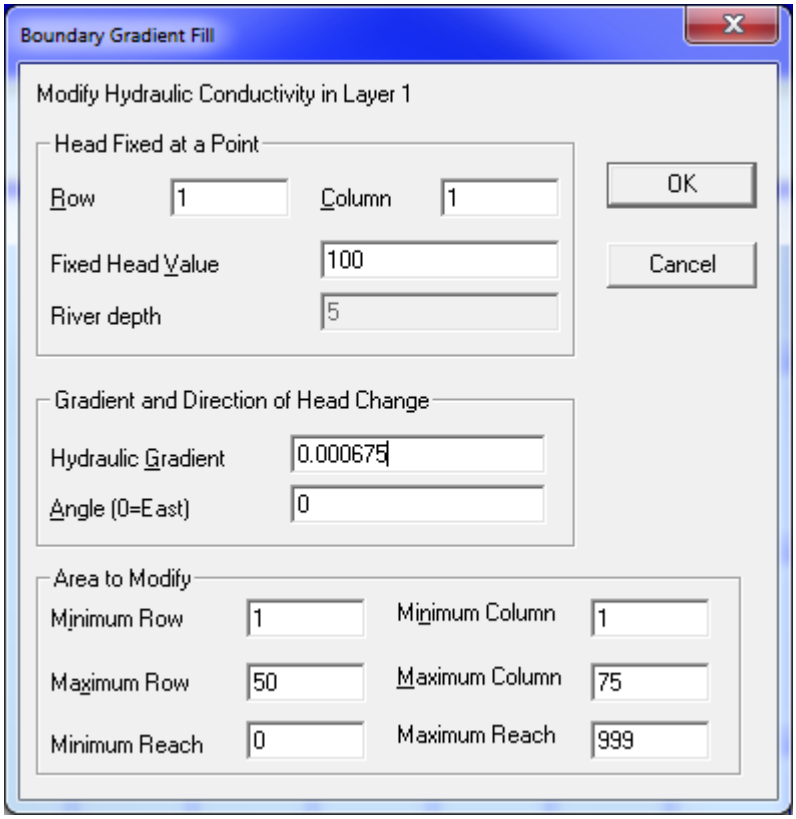

Now, if you double-click a constant head cell along the northern or southern boundary, you will see that the head has been changed from whatever you entered to a new value that is designed to give you a uniform gradient.

Use **File|Save As** to save this new model as **BCTEST2.GWV**. GV will ask if you want to change root file names. Enter bctest2 for the new modflow root file name.

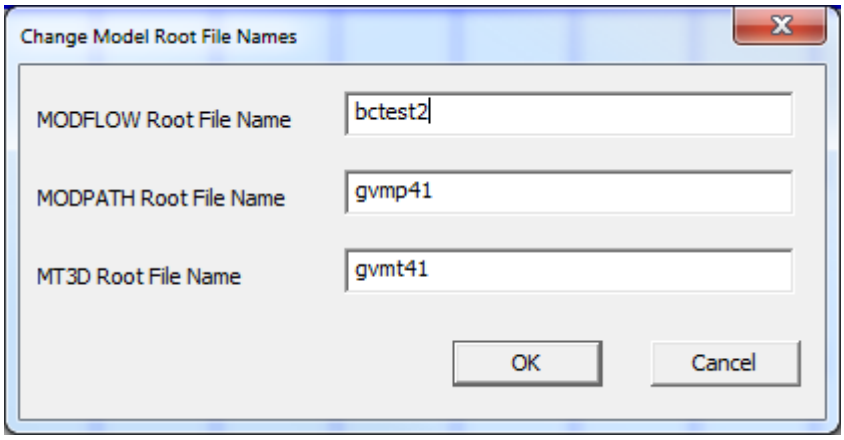

Now, run the model by clicking the calculator as you did in the first model and import the results. The water level contours should look exactly as before (see below).

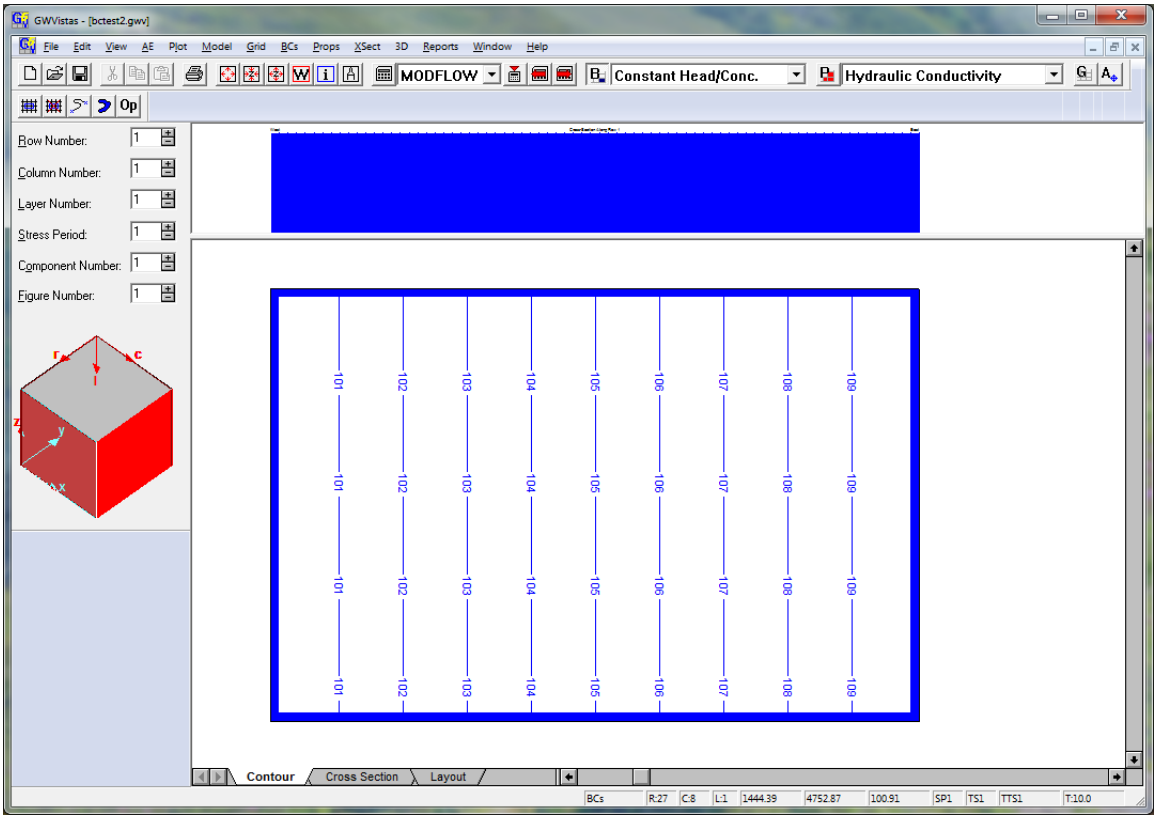

Just to be sure that we have the same water level contours, you will use another feature of Vistas that allows you to compare two binary MODFLOW output files. Select **Plot|File Operations|Compare Two Files**. Click the browse button next to the "Subtract" field and go find the first head-save file, which should be called ***bctest1.hds***. This will probably be located in the ***c:\gwv6\models*** directory. Next, click the Browse button next to the "From" field and find the results of the second run (***bctest2.hds***). Finally, enter the file name ***dif.hds*** in the "Results in" field. To actually make the comparison, click the **View** button. Vistas will then compare each layer and time step in the files to see what the differences are. It then displays the maximum head difference for each layer and time step in the window above the file names. If you were successful, that number should be about 0.01 meters, indicating that the results are very close.

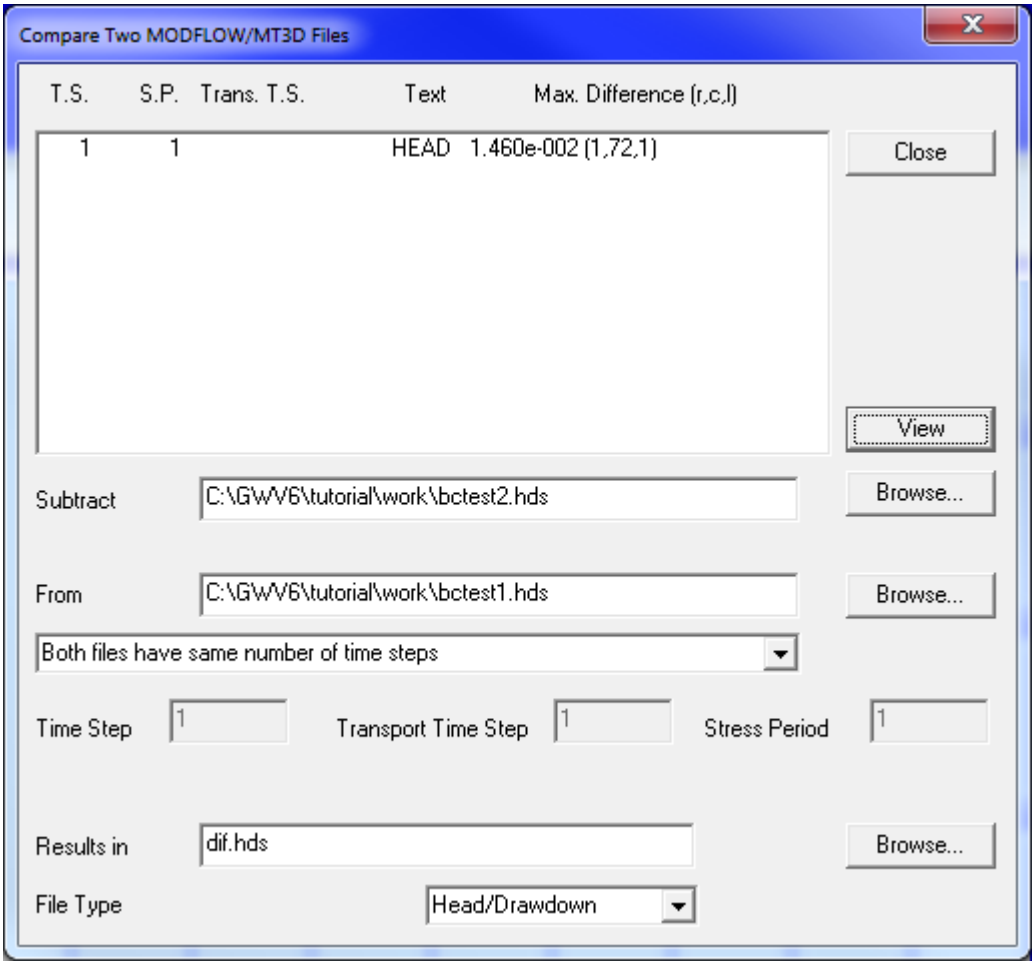

Model 3: Constant Head and Constant Flux

Make sure that you saved the second run to a file called ***bctest2.gwv*** using **File|Save As**. Open the first model you created in this session, ***bctest1.gwv*** using **File|Open**. This is the model with constant heads on the east and west but NOT on the north and south. You will change the eastern constant head cells to constant flux (wells).

In order to set the appropriate flow rate (flux) in these cells, we need to know what was computed in the first model. Run that model again and import the results, including the cell-by-cell flows. To do this, put a check to the right of the *.cbb file on the import dialog, as shown below.

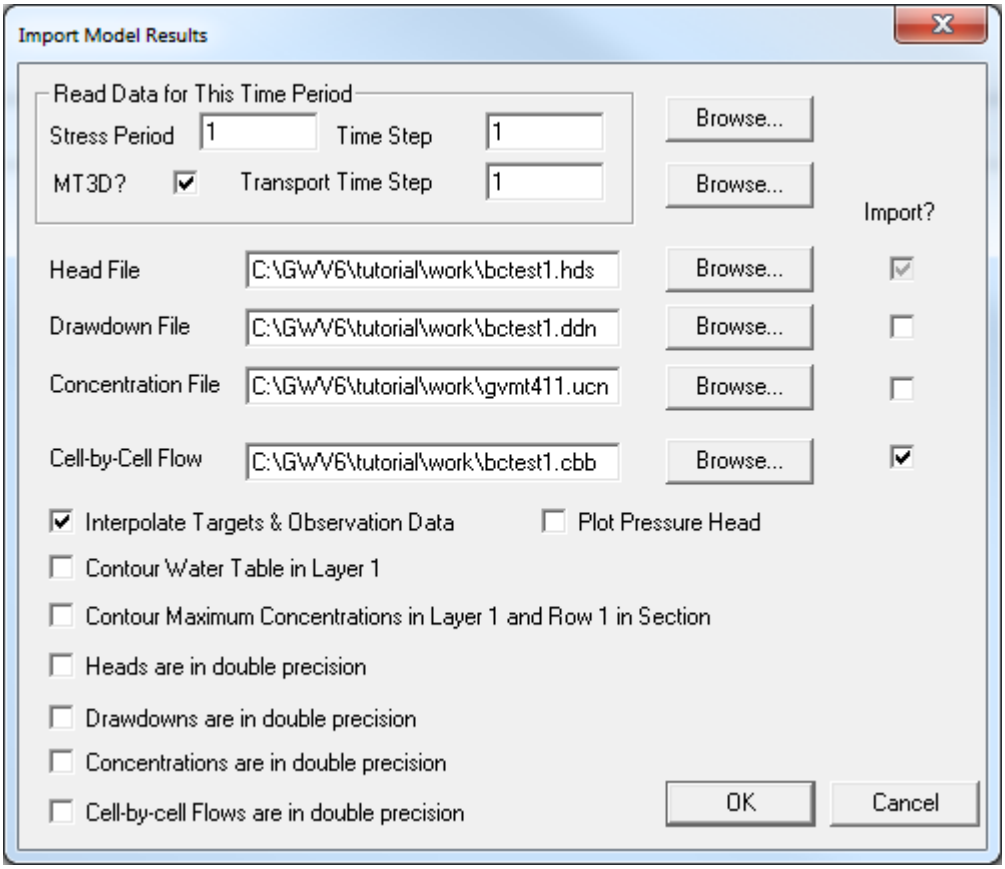

Push down the **B** button on the toolbar to enter Boundary Condition mode. Move the cursor over one of the constant head cells on the right and the flux rate should be shown on the status bar as "F=1.351e+03", which means that the cell is injecting 1,351 m^3/d into the model through the constant head cell. You could also compute this value from the gradient (0.000675) times the cell width (200 m), times the hydraulic conductivity (100 m|d), times the aquifer thickness (100 m).

Now we will put the constant flux cells in the model. First, delete the constant head cells on the right (east) by selecting **BCs|Constant Head** and then **BCs|Delete|Window**. Drag a window that encloses the eastern cells and Vistas will delete them. Now, use **BCs|Well**. Select **BCs|Insert|Window** or click the  button on the toolbar. Now, move the cursor somewhere inside the upper right (northeast) cell in the model, hold down the left mouse button, and drag a window so the end point is inside the lower left (southeast) cell in the model. Release the mouse button and a well dialog appears. Change the flow rate from 0.0 to 1351.0.

Select Model|MODFLOW|Packages and rename the root name as ***bctest3*** so that it will be separate from the previous two runs. Also, use **File|Save As** and rename the Vistas file to ***bctest3.gwv***. Now run the model and your screen model should now look like the one below.

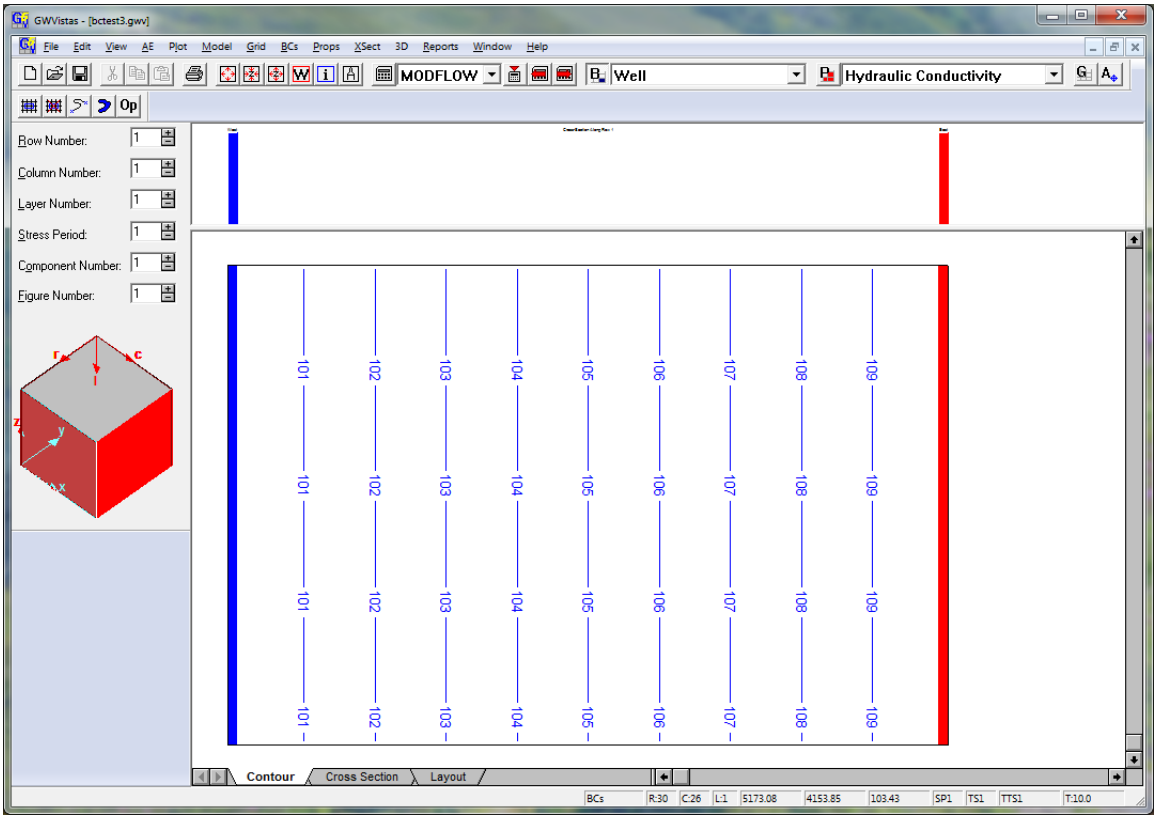

The red cells on the east are now wells that inject 1,351 m^3|d in each cell. Notice that the contours look exactly like the previous two model simulations. We now have three models with three different boundary condition designs that all yield the exact same solution. Which one is right? (Sorry, rhetorical question…)

Predicting Impact from a Well

To test the predictions from these three different models we will add one pumping well in the center. Select **AE|Well** and move the cursor to row 25, column 38 and click the left mouse button. An analytic well dialog is shown on your screen. Enter the flow rate as -100000.0 (100,000 m^3/d) – in MODFLOW a negative flow rate means that the water is pumped out of the model.

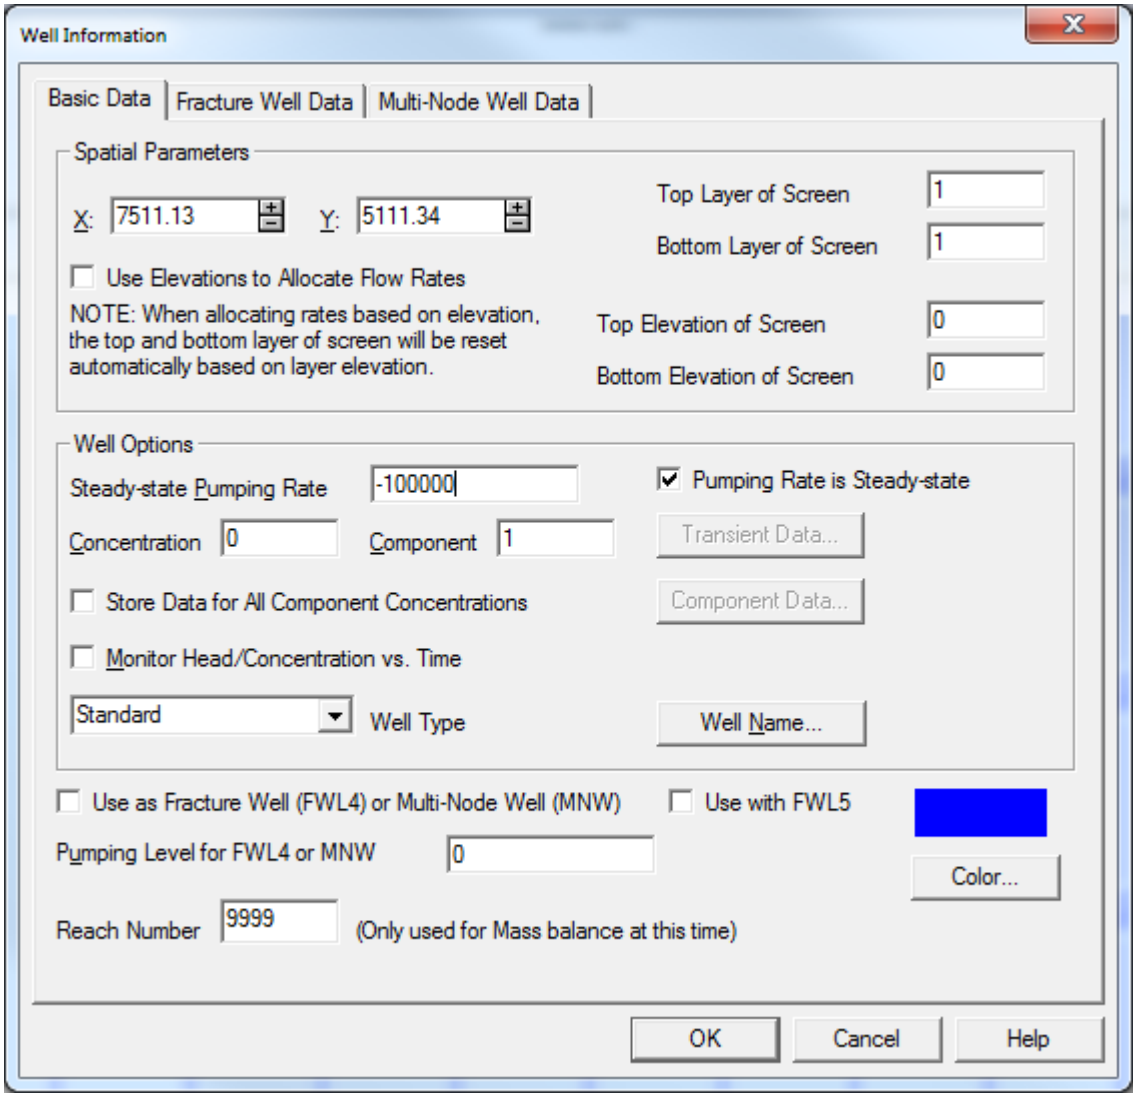

You should also select **Plot|Map and AE Options|Well Options** and change the well size to 200 so you can see it better. Your model should look like the one below. Note that the contours are still straight lines because we have not rerun the model yet.

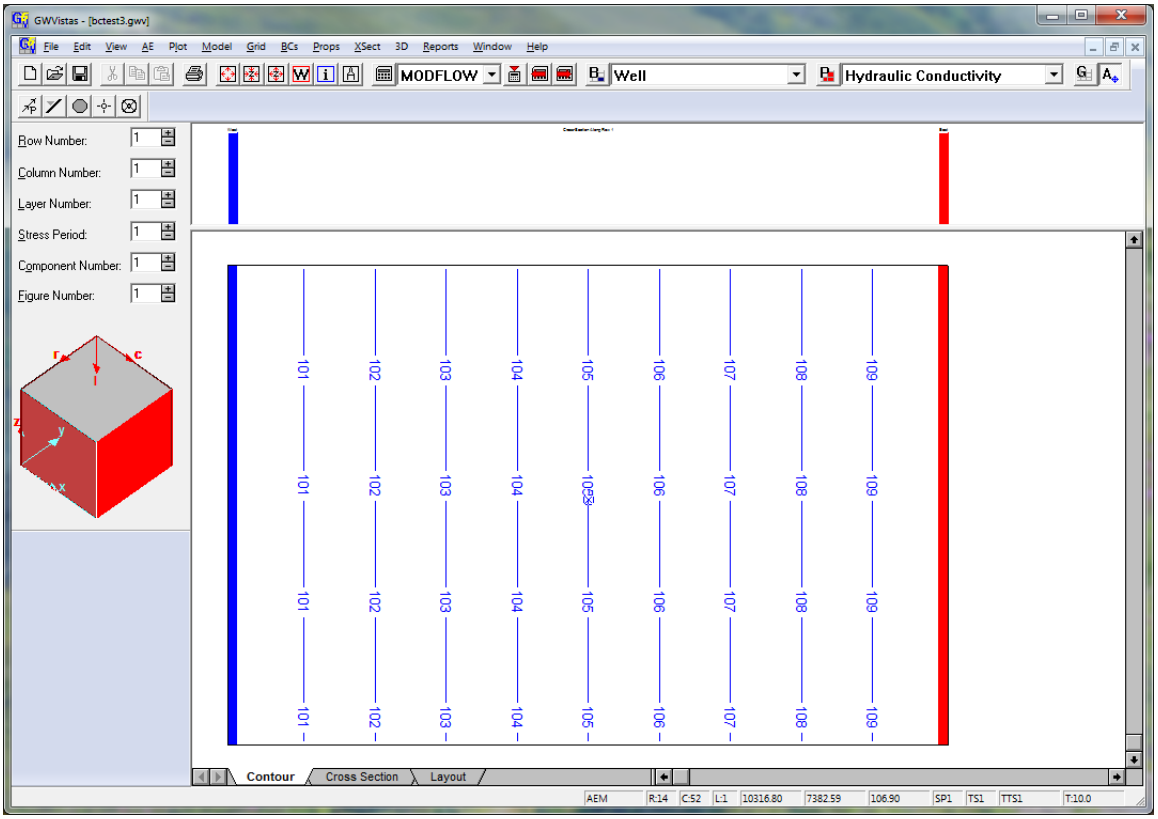

Now, click the calculator button to run the model with the well. After the simulation is over, import results as before and you should see the cone of depression around the well, as shown below.

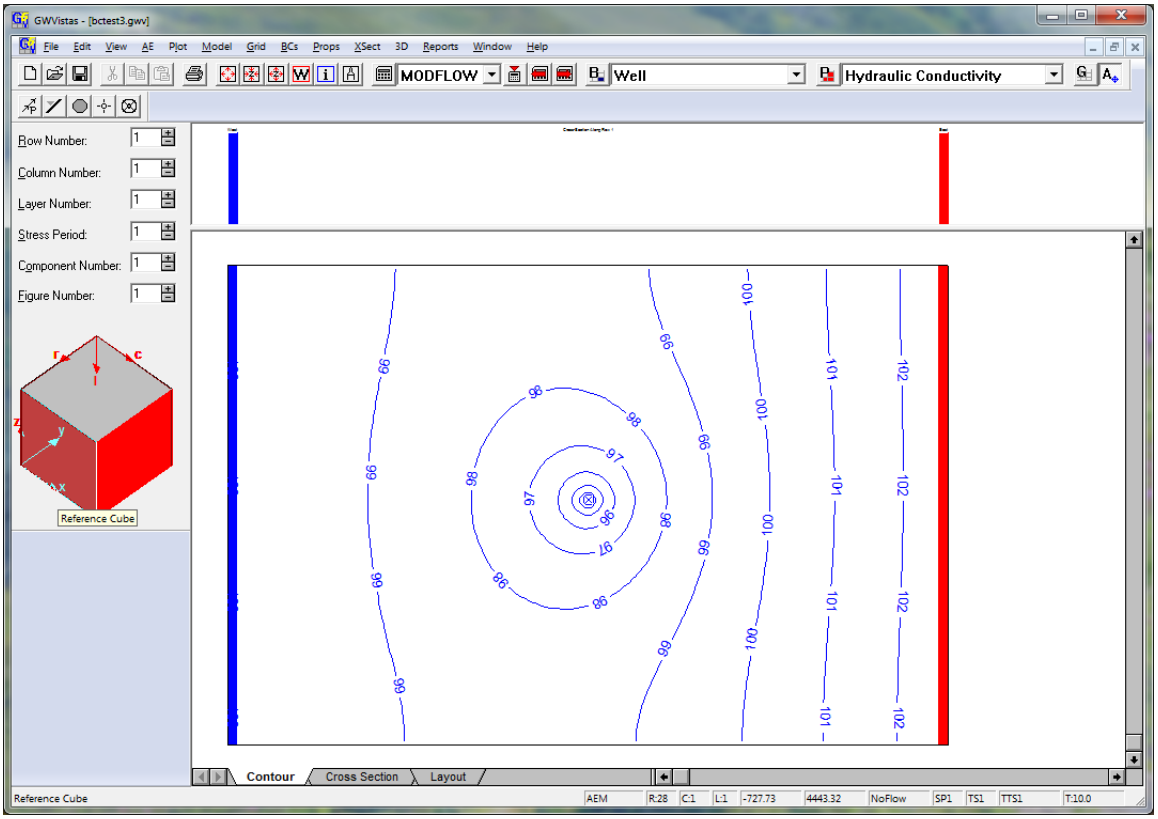

Now, repeat this with the other 2 models. An easy way to do this is to press the **A** button on the toolbar to enter analytic element mode. Click the left mouse button on the well you added so it turns solid blue. Now press **Ctrl-C** to copy the well to the clipboard. Use **File|Open** to open one of the other 2 models (e.g. ***bctest1.gwv*** or ***bctest2.gwv***) and press **Ctrl-V** to paste the well into the other model. Run that model and then repeat this procedure with the 3rd model.

The other two model runs should look like the ones below.

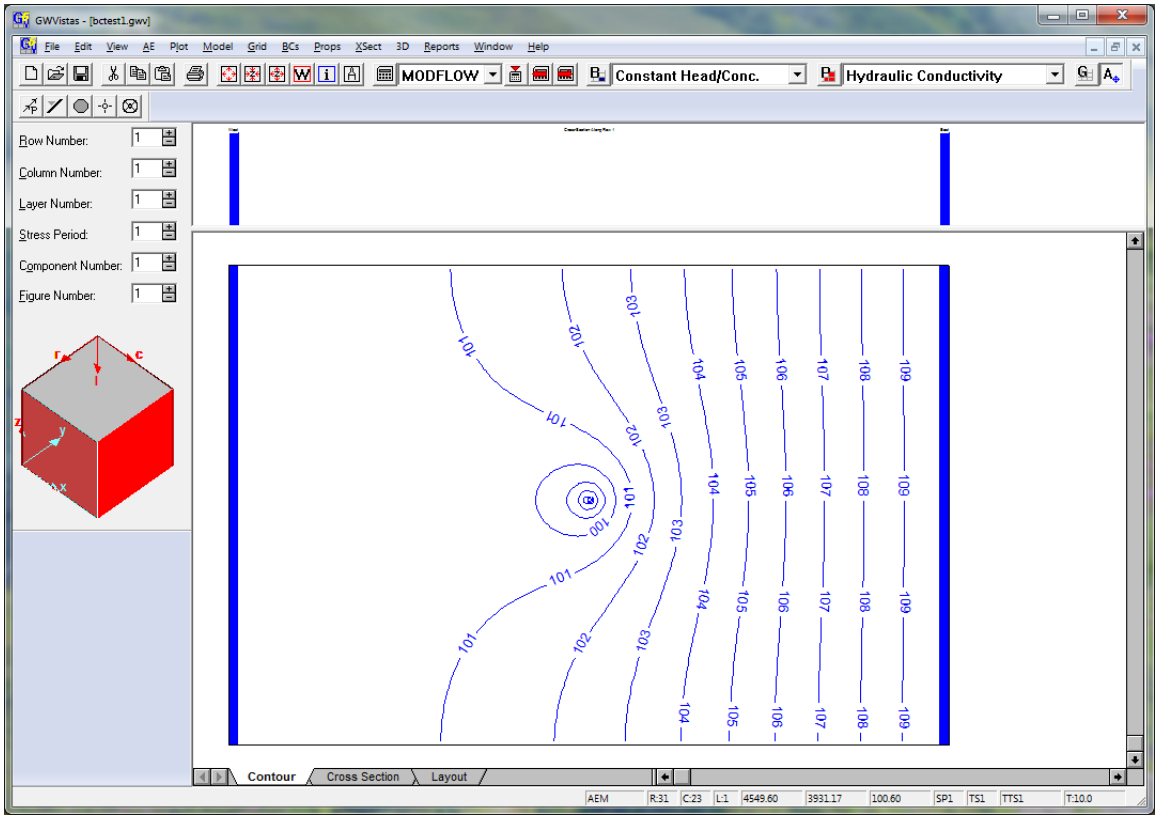
GWVistas - [bctest1.gwv]
File Edit View AE Plot Model Grid BCs Props XSect 3D Reports Window Help
MODFLOW
Constant Head/Conc.
Hydraulic Conductivity
Row Number: 1
Column Number: 1
Layer Number: 1
Stress Period: 1
Component Number: 1
Figure Number: 1
Contour
Cross Section
Layout

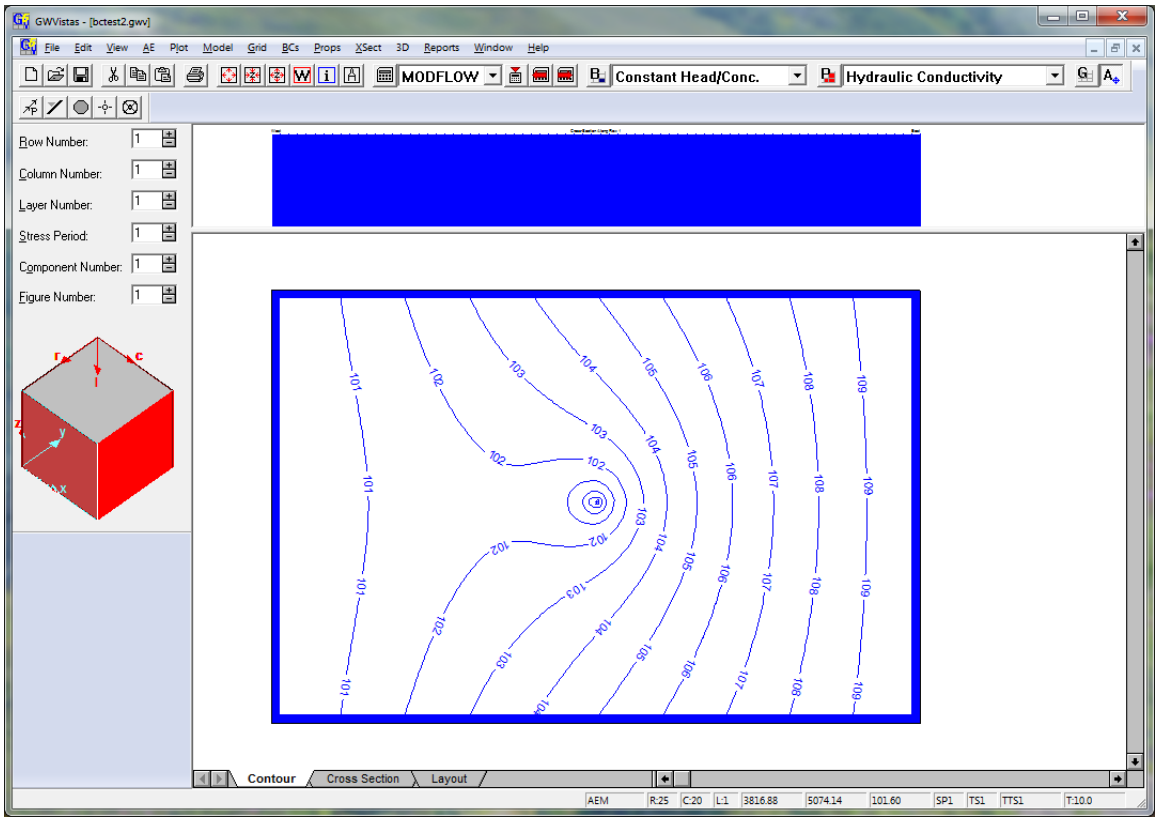
GWVistas - [bctest2.gwv]
File Edit View AE Plot Model Grid BCs Props XSect 3D Reports Window Help
MODFLOW
Constant Head/Conc.
Hydraulic Conductivity
Row Number: 1
Column Number: 1
Layer Number: 1
Stress Period: 1
Component Number: 1
Figure Number: 1
Contour
Cross Section
Layout

You can see that the results are quite different. The largest cone of depression (largest amount of drawdown) is for the model that has the one boundary set to a constant flux. This happens because the model cannot allow any extra water to enter the model domain in response to pumping. The next largest cone of depression is with the first model that had 2 sides of the model set to constant heads. The least amount of drawdown was for the model that has all sides as constant head. The latter does not allow the head to change on the model boundary and so limits the amount of impact from the well by supplying additional water to the model from all sides. Which one do you think is the most realistic approach?

Model Calibration in Groundwater Vistas

Introduction

The following tutorial focuses on Groundwater Vistas' tools for model calibration and uses a synthetic model to illustrate the importing and manipulation of calibration targets and using automated sensitivity analysis. Defining a set of calibration targets is a fundamental first step in applying any inverse model, such as Pest. Groundwater Vistas has several different types of targets, including head, drawdown, flux, concentration, gradient (magnitude & direction), and subsidence. In this simple example, we will import head targets into a model and explore the different ways of displaying and manipulating targets in preparation for using Pest.

To start, open the model called c:\gwv6\tutorial\CalEx0.gwv. Run the model by clicking the calculator button and import the results after the run is finished. Your screen should look like the one below:

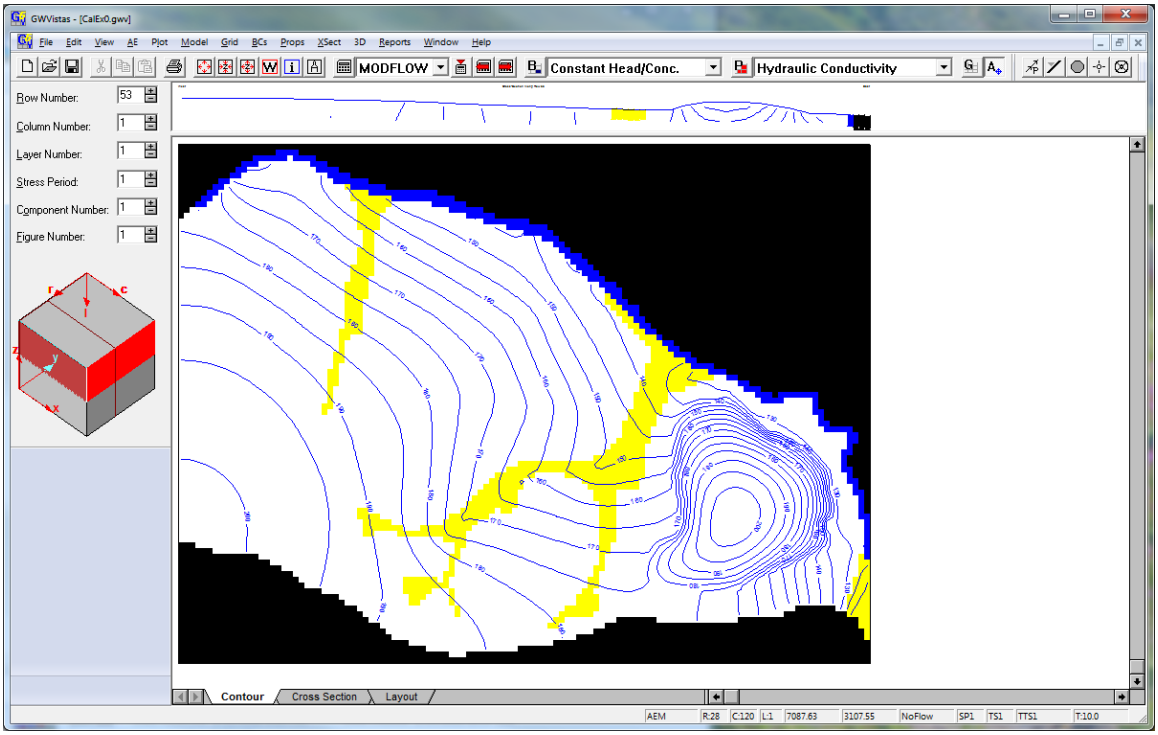

Preparing for Calibration

The first step in preparing for automated sensitivity analysis and Pest calibration runs is to make sure that the model will run reliably when parameters are changed. A good way of testing this is to set up an automated sensitivity analysis in Groundwater Vistas. There are two things that have been added to Groundwater Vistas' version of MODFLOW that are not standard but assist in stabilizing runs. The first is to continue model simulations even if the model has not converged. While this can lead to bogus simulations, it will allow Pest to keep going. To activate this option, choose **Model|MODFLOW|Package Options – BASIC tab,** as shown below.

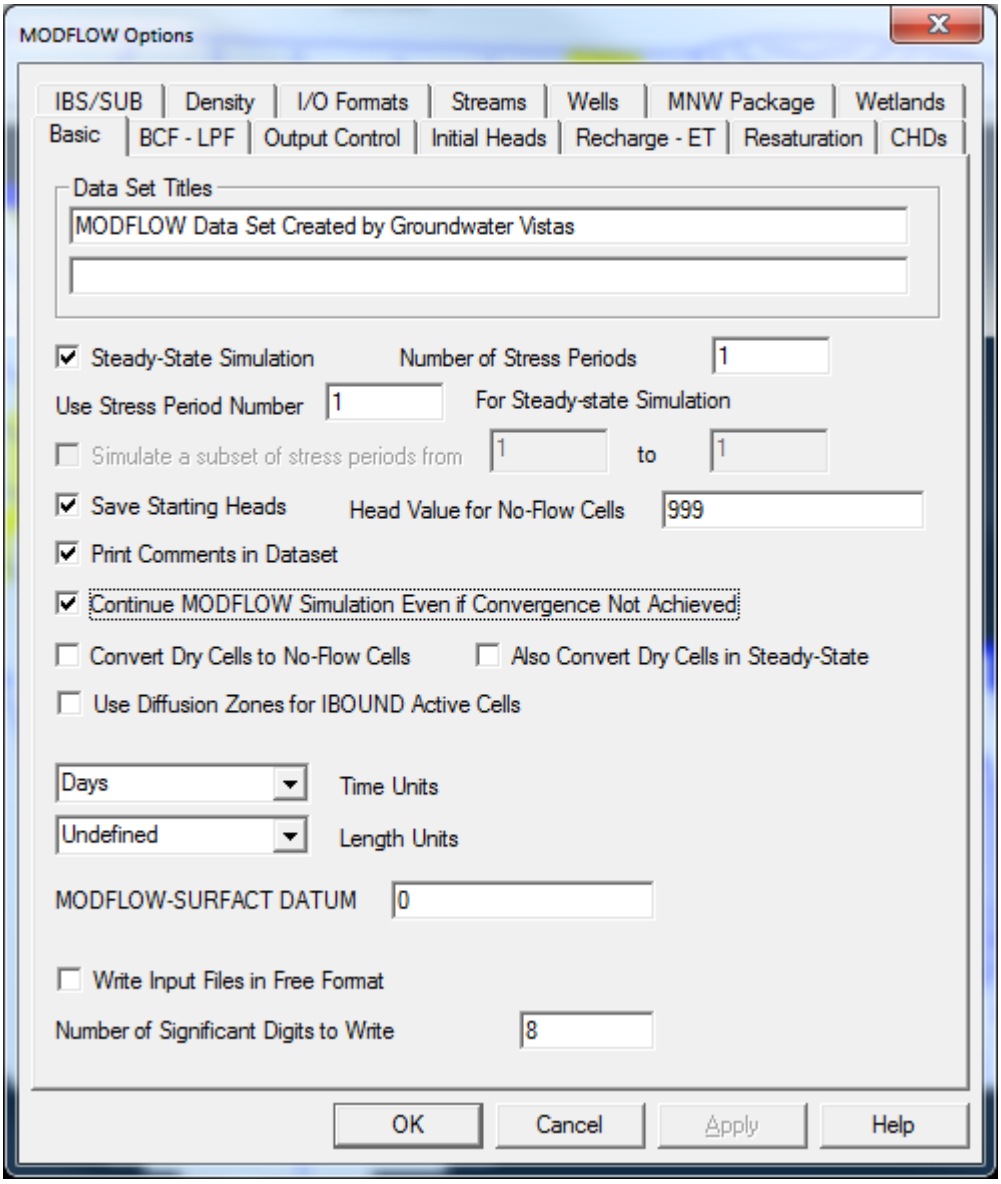

The option is called “Continue MODFLOW Simulation Even if Convergence Not Achieved”. You should make a habit of always turning this on when using automated sensitivity analysis and Pest. Note, however, that this is NOT a standard USGS option (although we have made several requests of the USGS to add this as an option). If you are always using Groundwater Vistas, it is not a big deal. It only becomes a problem if your simulations rely on this option and you transfer the model to someone using another preprocessor or version of MODFLOW.

The second option that helps with slightly unstable model simulations is related to the PCG2 solver. The PCG2 solver has a bad habit of often getting to the end of the run and then just slightly oscillating between outer iterations without ever converging fully. We added an option to stop iterating if the two convergence criteria (head and residual) are met in N successive outer iterations. Usually 5 to 10 is a good number. If you rely on this feature, make sure that the mass balance error is low (< 1 %). As with the previous option, this is not a USGS standard option. To activate this switch, use **Model|MODFLOW|Solver Options – PCG2 tab**. It is the last field on the property sheet.

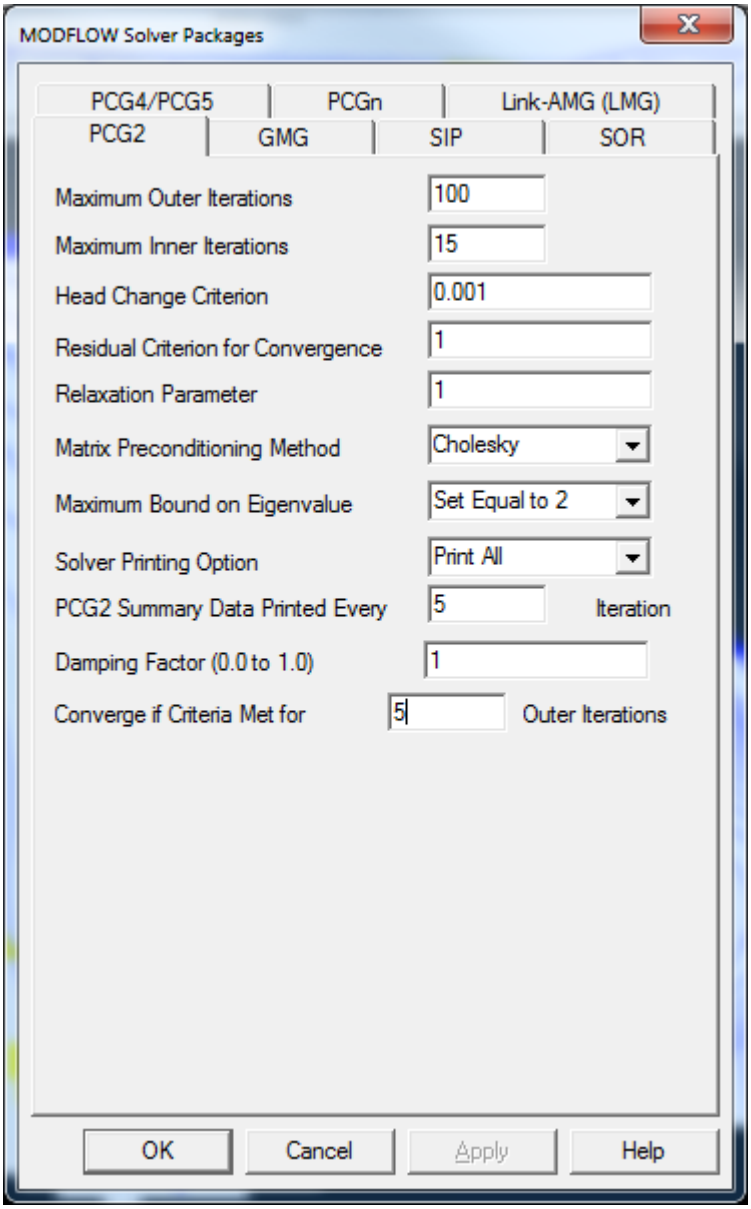

Adding Calibration Targets

Calibration targets can be added to the model manually by selecting **AE**|**Target** and left-clicking where you want the target to go. You can also click the  button on the toolbar and then left-click to add the target in the same manner. Usually, however, there are far too many targets to insert them all manually. The more common way of adding targets is either to import from a text file (e.g. comma-delimited Excel spreadsheet) or from a GIS shapefile. We will illustrate each method.

Start by importing a set of targets from a text file. The file is called ***CalEx0_targets.csv*** and should be in the gwv6\tutorial directory. You can open this file in Excel to view the contents.

	A	B	C	D	E	F	G
1	Name	X	Y	Head	Layer	Weight	
2	MW-1	1381607	394208.1	223.9	1	1	
3	MW-2	1381662	396499.3	197.5	1	1	
4	MW-3	1382469	395729.5	199.5	1	1	
5	MW-4	1382716	396746.8	168.8	1	1	
6	MW-5	1382121	396783.4	183.6	1	1	
7	MW-6	1384329	396499.3	152.8	1	1	
8	MW-7	1384999	395372	157.8	1	1	
9	MW-8	1383385	395335.4	195.7	1	1	
10	MW-9	1383046	393896.4	201.1	1	1	
11	MW-10	1383981	396105.2	174.4	1	1	
12	MW-11	1383880	393465.7	196.2	1	1	
13	DW-1	1382130	396636.8	186.4	2	1	
14	DW-2	1383890	395683.6	178.3	2	1	
15	DW-3	1382423	394703	207.2	2	1	
16	DW-4	1384320	393694.8	186.2	2	1	
17	DW-5	1383734	396783.4	162.1	2	1	
18							
19							
20							
21							

There are 16 targets with names, coordinates, target head values, layers, and weights. This is the generally the minimum information to define a target in Vistas, although layer and weight can be omitted. If layer is omitted, then the targets are imported into the current layer. Weights are assigned a value of 1.0 if they are not in the file.

Import this file by selecting **AE|Import|Target Text File** and browse to find this file. Then enter the column number of the data next to the corresponding label on this dialog. For example, Name is in column 1, X is in column 2, etc. You also need to enter a 1 for number of lines to skip. By default Vistas assumes that the X and Y coordinates are "site coordinates". That is, they are the same coordinates as your base map (see first option on the dialog below). Your dialog should look like the one below.

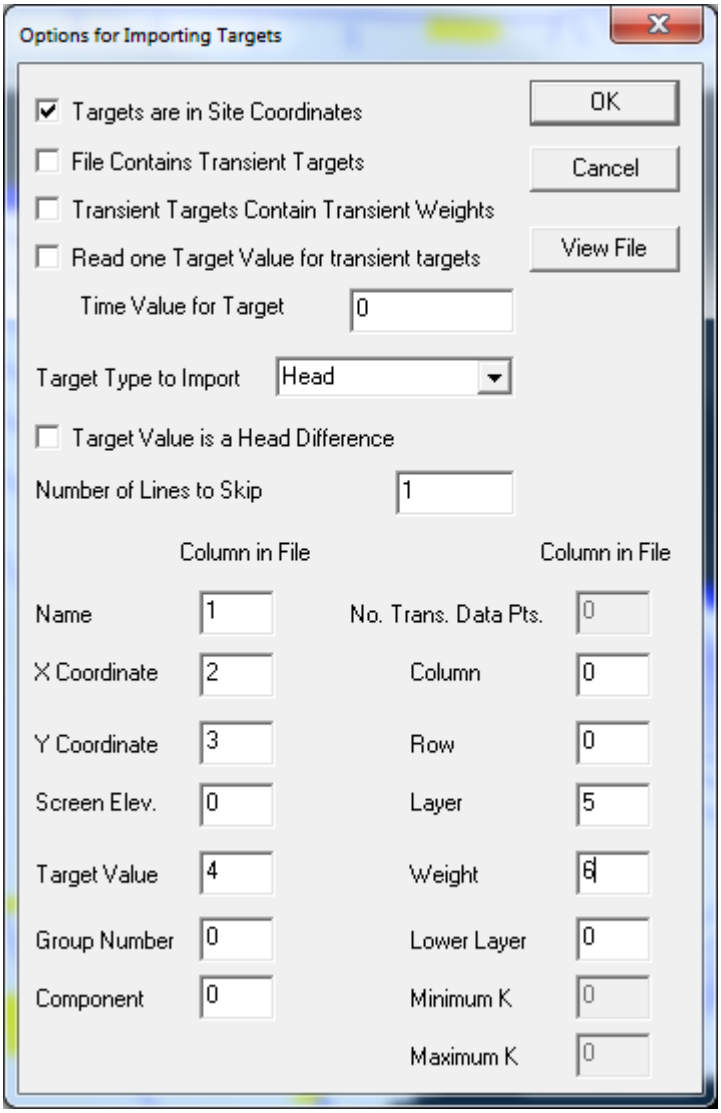

Click OK and Vistas should report that 16 targets were imported. If you do not get that number imported then there was something wrong on the import dialog above. You can confirm that you have 16 head targets by selecting **Model|Model Information**. This is a handy place to get a summary of all boundary conditions and analytic elements in the model.

Another option for adding targets is from a shapefile. The shapefile was invented by ESRI for ArcView and ArcInfo (now ArcGIS) as a standard file type for displaying and manipulating spatial data. They can be composed of either points, lines, or polygons. A shapefile is actually a collection of files. Groundwater Vistas uses 3 of these files: (1) the *.shp file contains the coordinates of each entity, (2) the *.dbf file contains the attributes of each entity (e.g. Name, target head, etc.), and (3) the *.shx file is tells which shape goes with which item in the attribute database. Groundwater Vistas further assumes that the coordinates are the same as your base map.

To import targets from a shapefile, use **AE|Import|Shapefile**. Browse to find the file CalEx0_targets.shp and change the element type from Well to Target.

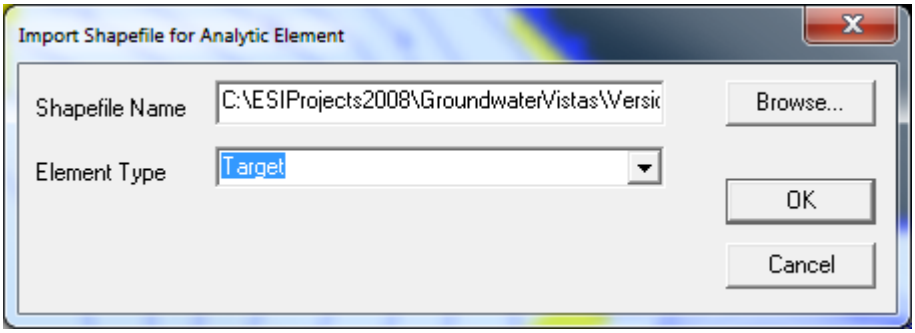

Groundwater Vistas will now open the database file and pull out the field names. The following dialog is displayed. Your job is to select the database field that corresponds to the data types needed for the target. If a particular data type is not in the database, then just leave the field as "Not in Database". Fill out the dialog as follows, making sure to check the option that the layer elevation field contains layer numbers:

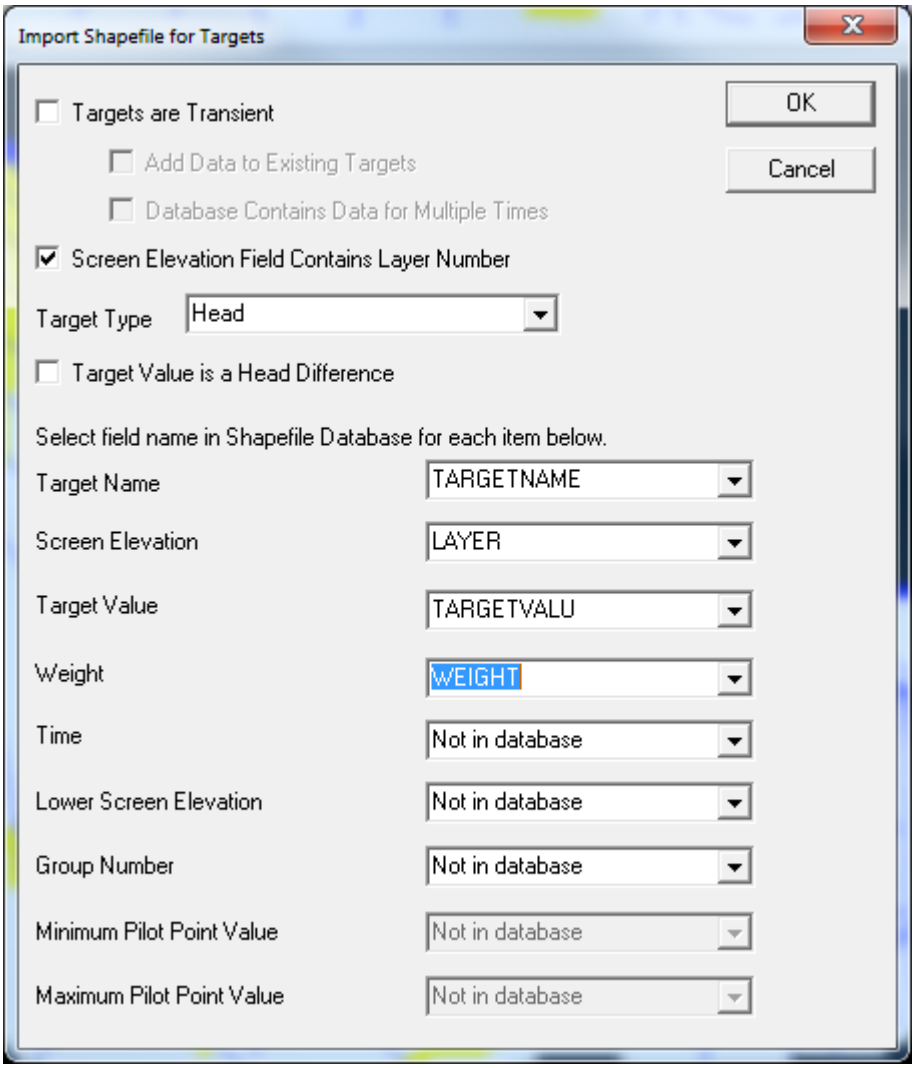

Note that you never tell GV which field contains coordinates. This is because the coordinates are obtained automatically from the shapefile. Click OK when you are done and Groundwater Vistas should again report that 16 targets were imported.

You will now see that a problem has occurred. Your targets have disappeared. This happened because we imported the targets twice, once from the text file and once from the shapefile. Groundwater Vistas does not delete duplicate targets, but rather simply keeps adding them to the model as you import them. If two targets lie identically at the same location, they essentially "white out" each other and nothing is displayed.

This is a rather common problem, which is why we set up this computer session the way we did. You can rectify the problem in a couple of ways. Generally the easiest approach is to delete all targets and import again. This is done by selecting **Edit|Select All|Targets** followed by **Edit|Delete** (or just hit the Del key). Another approach since we only have a few targets is to select **Edit|Analytic Element List** and click on the duplicate targets and hit the Delete button, as shown below.

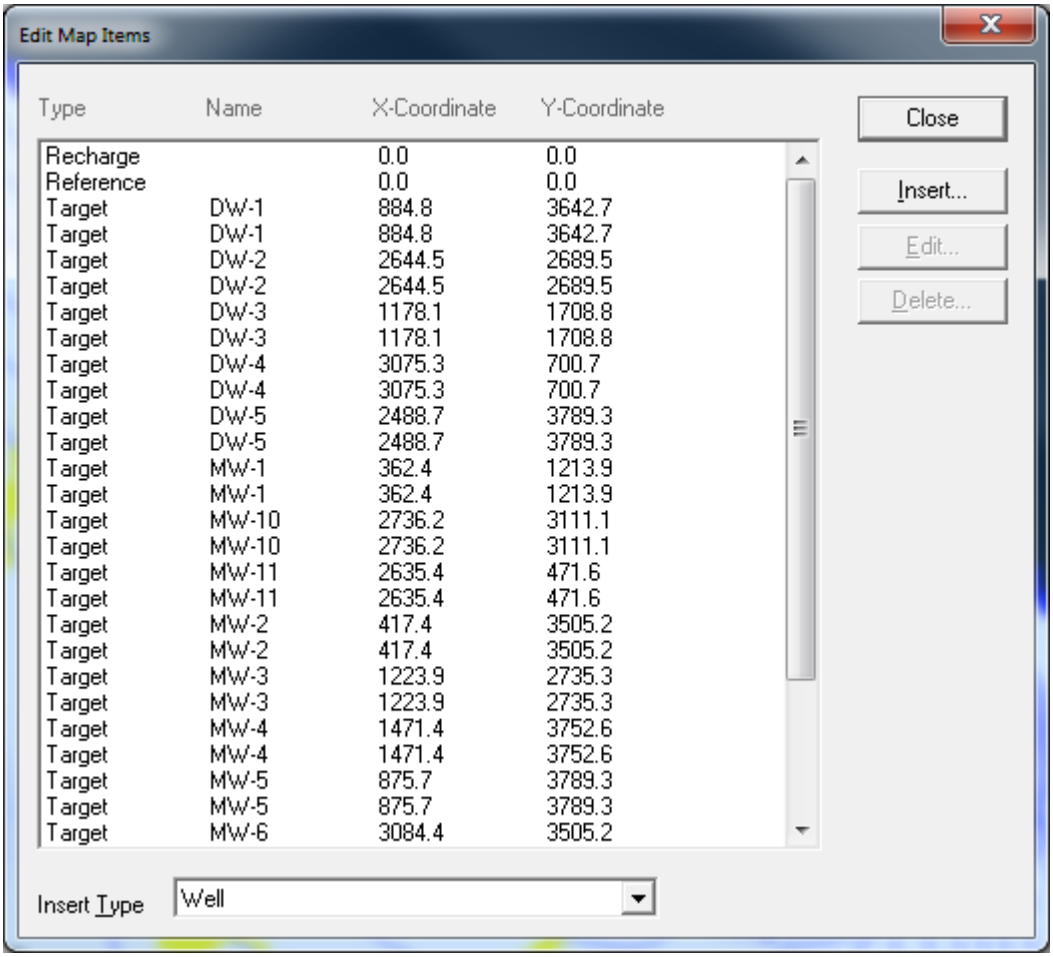

Pick a method outlined above to get back to one set of 16 calibration targets.

Calibration Statistics and Plots

The reason for using calibration targets is to be able to generate statistics and plots that show how well the model matches reality. The first thing you usually do is select **Plot|Calibration|Statistics-Plots**. If you do that now, you will get an error message:

Groundwater Vistas computes the residual (error) at each target when the results are imported. If you import targets after importing results, the residuals have not been computed. Simply select **Plot|Import Results** and click OK to update the target calculations.

There are many ways of analyzing residuals at target locations. The most common is to get the statistical summary and make basic charts. This is done through **Plot|Calibration|Statistics-Plots.** Do that now and then click the statistics button. You should get a summary that looks like the following:

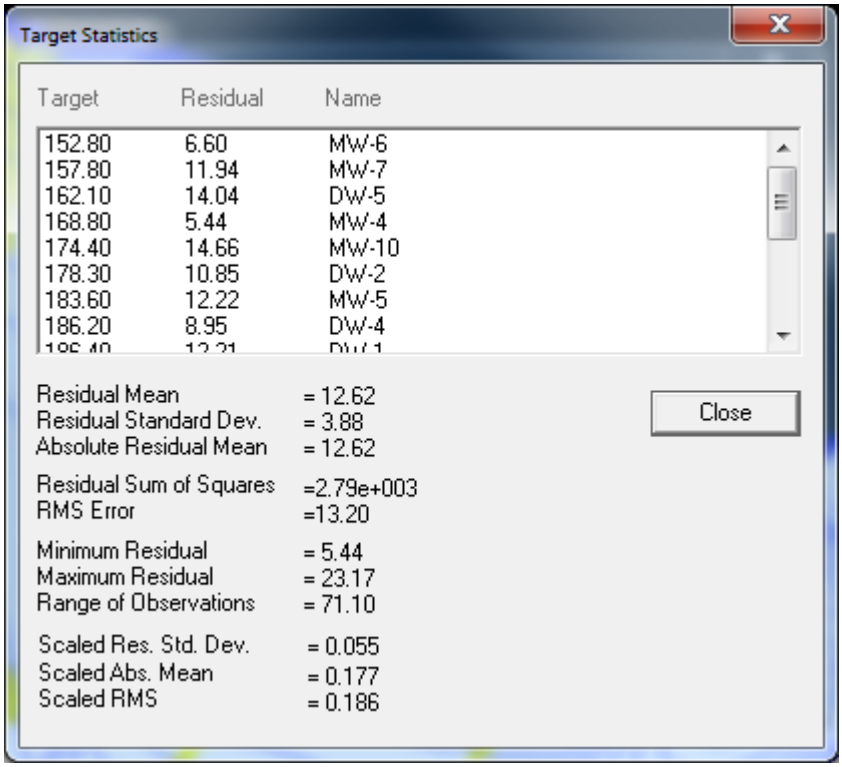

If you want a spreadsheet of the details for these targets, click Close and then click the Export button. The file that is exported is a comma-delimited file you can open in Excel. Various plots can also be created such as computed versus measured target values, computed versus residual, and a cumulative sum of squared residuals plot. The latter is very useful when using Pest because a few very large residuals can cause convergence problems in Pest. If you create the cumulative sum of squares plot now you will see the following:

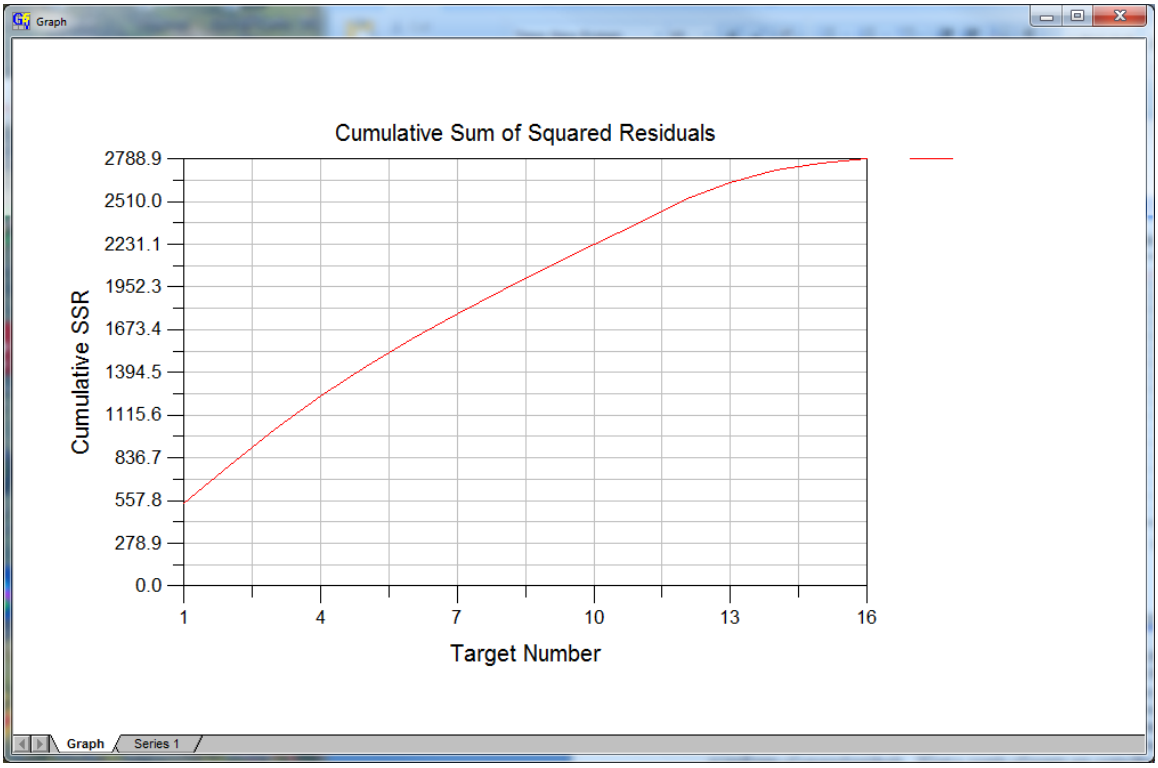

This shows that almost half of the total sum of squared errors (objective function in Pest) is controlled by just 4 targets and 20% of the sum of squares is the one worst target (558 ft^2 for target 1 in the chart above). This is not a significant problem early in the calibration. However, after you have been working on the calibration this might be an indication that there are a few targets (maybe just 1 target) that is a significant part of the overall objective function. In that case, you may consider using a weight less than 1.0 to remove any adverse influence of this one target on the calibration. Remember that Pest is trying to lower the overall sum of squared residuals. If just a couple of targets are controlling that value, then Pest will be spending most of its effort on those few targets.

Other common displays you can create for the target set is to post the residual at each target and/or draw a residual circle at each target. This is done using **Plot|Calibration|Post Residuals** and **Plot|Calibration|Post Residual Circles** (or just use Plot|What to Display). The options for how these postings are created are controlled using **Plot|Calibration|Options**. Try posting residuals and residual circles for the current model. Your plot may look something like the one shown below.

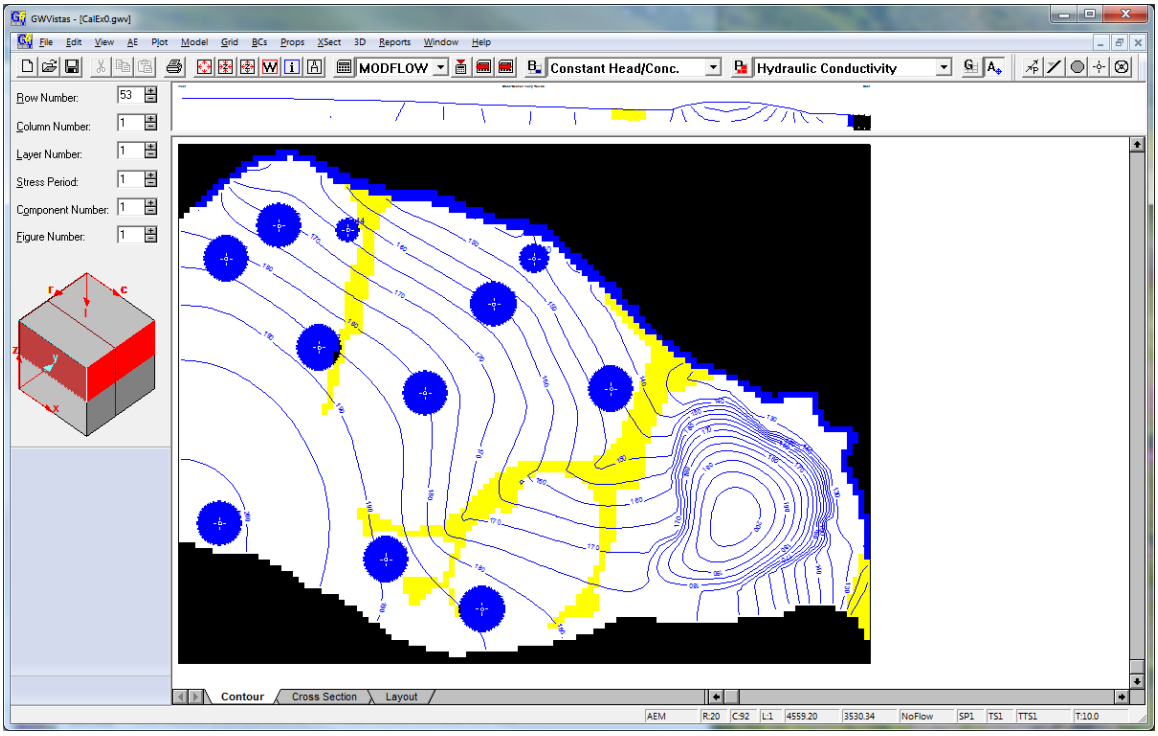

The colors used for residuals and residual circles are blue for positive residuals (model-computed values are too low) and red for negative residuals (model-computed values are too high). You can use the options menu described above to change those colors.

Editing Calibration Targets

Targets are edited in Vistas by either double-clicking individual targets (make sure the "A" button is pushed down on the toolbar) or using the **AE**|**Modify**|**Targets** menu.

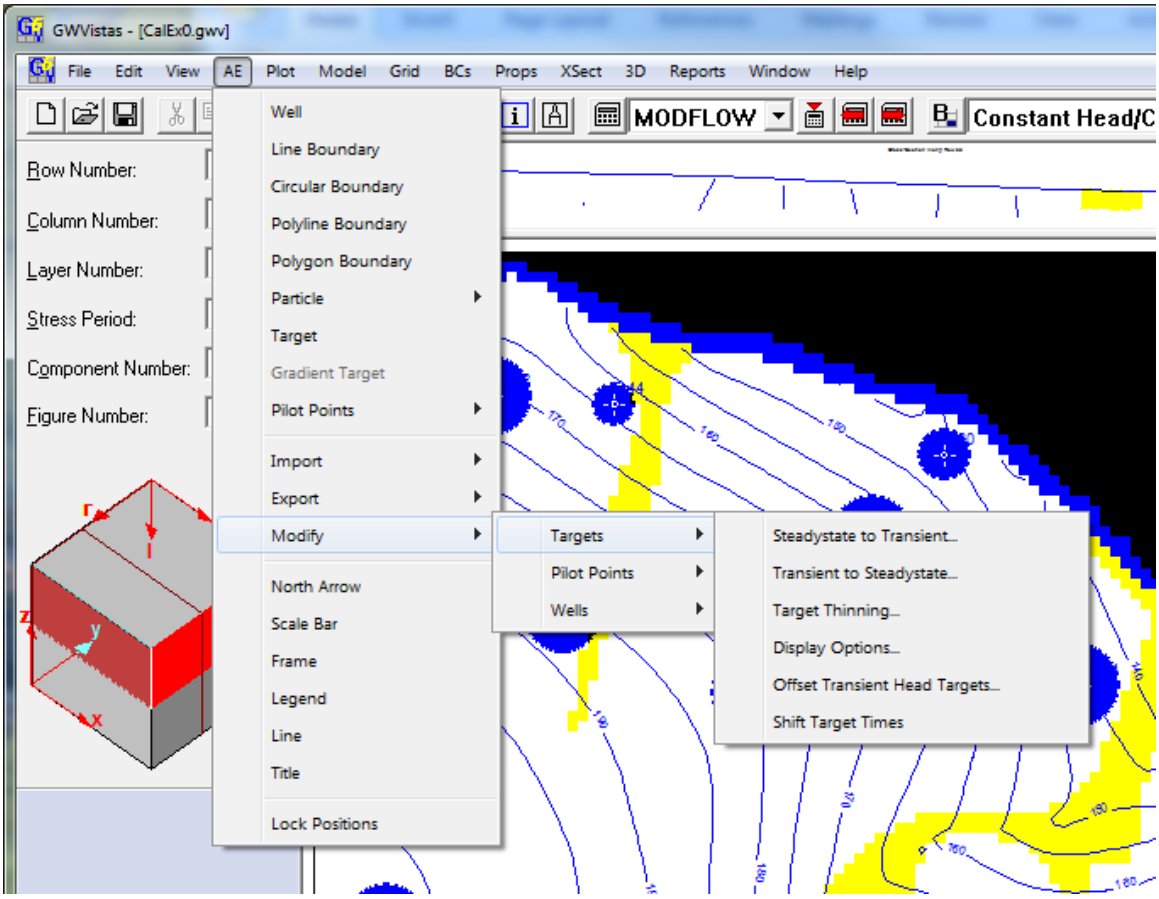

You can convert steadystate targets to transient or vice versa. Note that if there is more than one stress period and|or time step, even though the run is steadystate, all targets must be transient. The "Target Thinning" option is used when you have highly clustered targets. This option can automatically delete targets based on number per cell, distance, etc.

Calibration Target Report

Another type of export for target residuals is accessed from **Reports|Calibration|Target Residuals**. This report offers more options than the standard export under **Plot|Calibration|Statistics-Plots**. The target residuals report options are shown below.

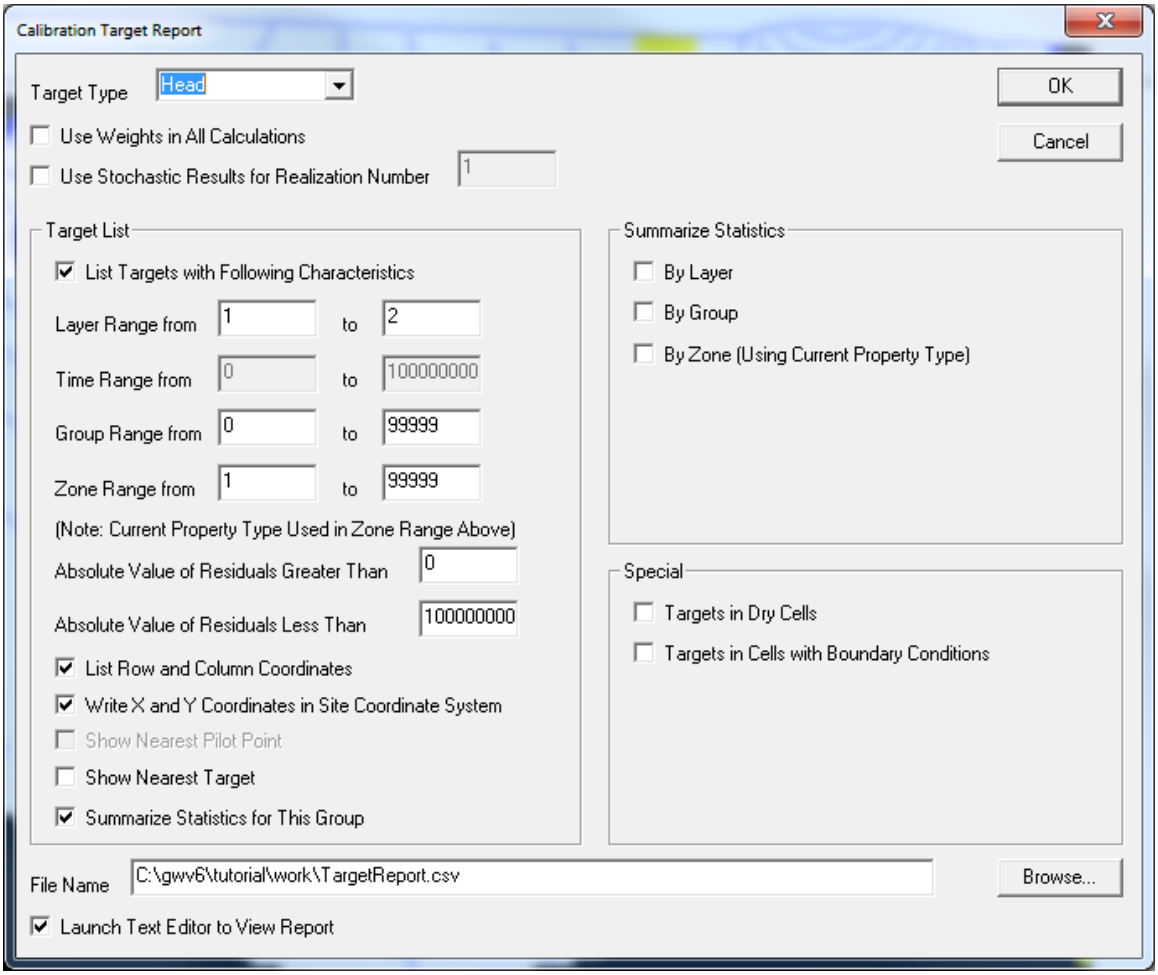

These options give you a lot of flexibility in creating residual tables and statistics and looking for problems such as targets in dry cells or extreme target values. Try creating a report for this model and open the file in Excel.

Automated Sensitivity Analysis

The purpose of the automated sensitivity analysis in Groundwater Vistas is to explicitly show the effect of small changes to model parameters on calibration statistics. Curves are plotted for each parameter tested and these curves help you decide which parameter to adjust and by how much. It is a computationally intensive operation and not as efficient as a Pest run. However, it does have the advantage of confirming model stability and provides good information on overall parameter sensitivity. We believe it is a good precursor to a Pest run.

The easiest way to run an automated sensitivity analysis through Groundwater Vistas is to use the new script file creation routines. Select **Model|Auto Sensitivity|Options**.

Change the parameter to Kx=Ky at the top of the dialog and confirm that the zone is set to 1. Note that when changing the parameter type at the top, do not type in the parameter name – simply use the dropdown list of names provided. In this case, Kx=Ky means that we are adjusting both Kx and Ky by the same amount (i.e. isotropic conditions).

Confirm that the number of multipliers to 11. Click the multipliers button and confirm that the multipliers are 0.5, 0.6, 0.7. 0.8, 0.9, 1.0, 1.1, 1.3, 1.5, 1.7, 2.0.

Once you have defined the first parameter, click the *New Script* button and put the script file in c:\gwv6\tutorial\work. Call it *sens1.in*. Note that GV automatically will change the output file name to sens1.in.out so you know which automatic sensitivity analysis results corresponds to which input file.

After giving the file name, GV will prompt for a range of zones to use. This is a quick way of adding many consecutive zones without having to do it manually. Enter 1 for the first zone and 2 for the second zone. We have two K zones of interest in this model.

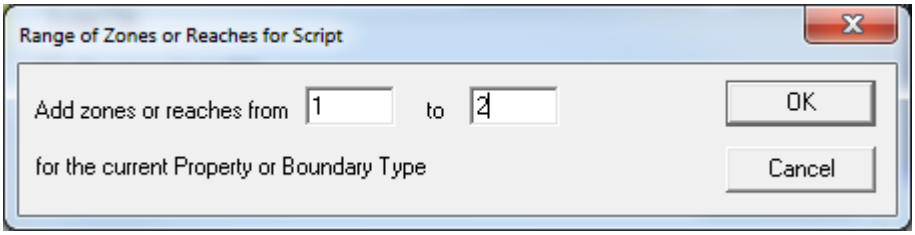

Now change the property type from Kx=Ky to Kz.

Click the *Add to Script* button and enter the zone range of 1 to 2 again.

You now have 4 parameters in the script file, Kx1, Kx2, Kz1, and Kz2.

Now change the parameter type to Drain Conductance.

Click the *Add to Script* button again and enter a range of drain reaches from 0 to 1.

You now are done creating the script file unless you want to add recharge. To look at the script file, click the *Edit Script* button. The file should look like the following one:

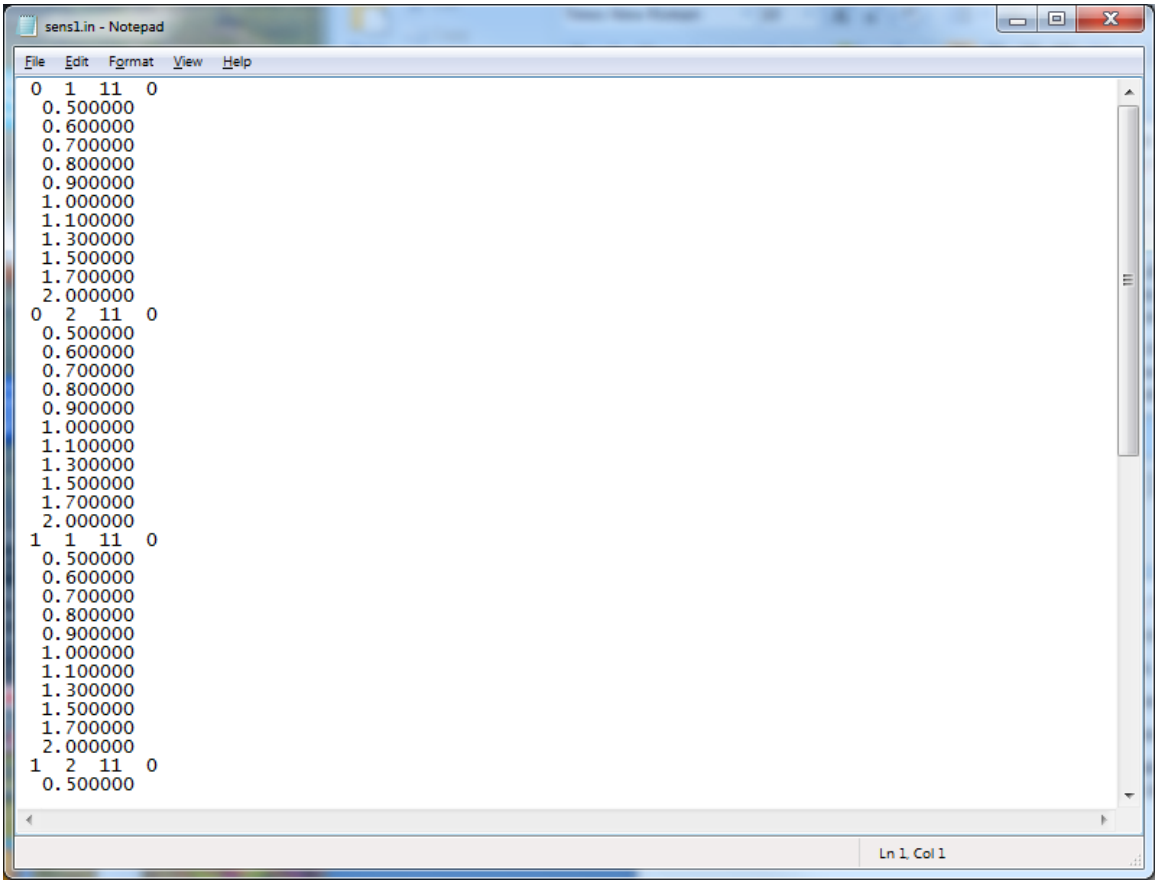

The format of the script file is documented in the Groundwater Vistas manual. The first number is the parameter number (0=Kx, 1=Kz, etc.). The second number is the zone or boundary reach and the third number is the number of simulations to run. Following the parameter definition line are the multipliers in order from lowest to highest.

To run the script file, click the *Run from script file* button. Be sure that the script output file is located in the same directory as the input file (normally it will be because GV will automatically put it in the right location of you use the "New Script" button. In this case, it should be but if you put it somewhere else, change it now.

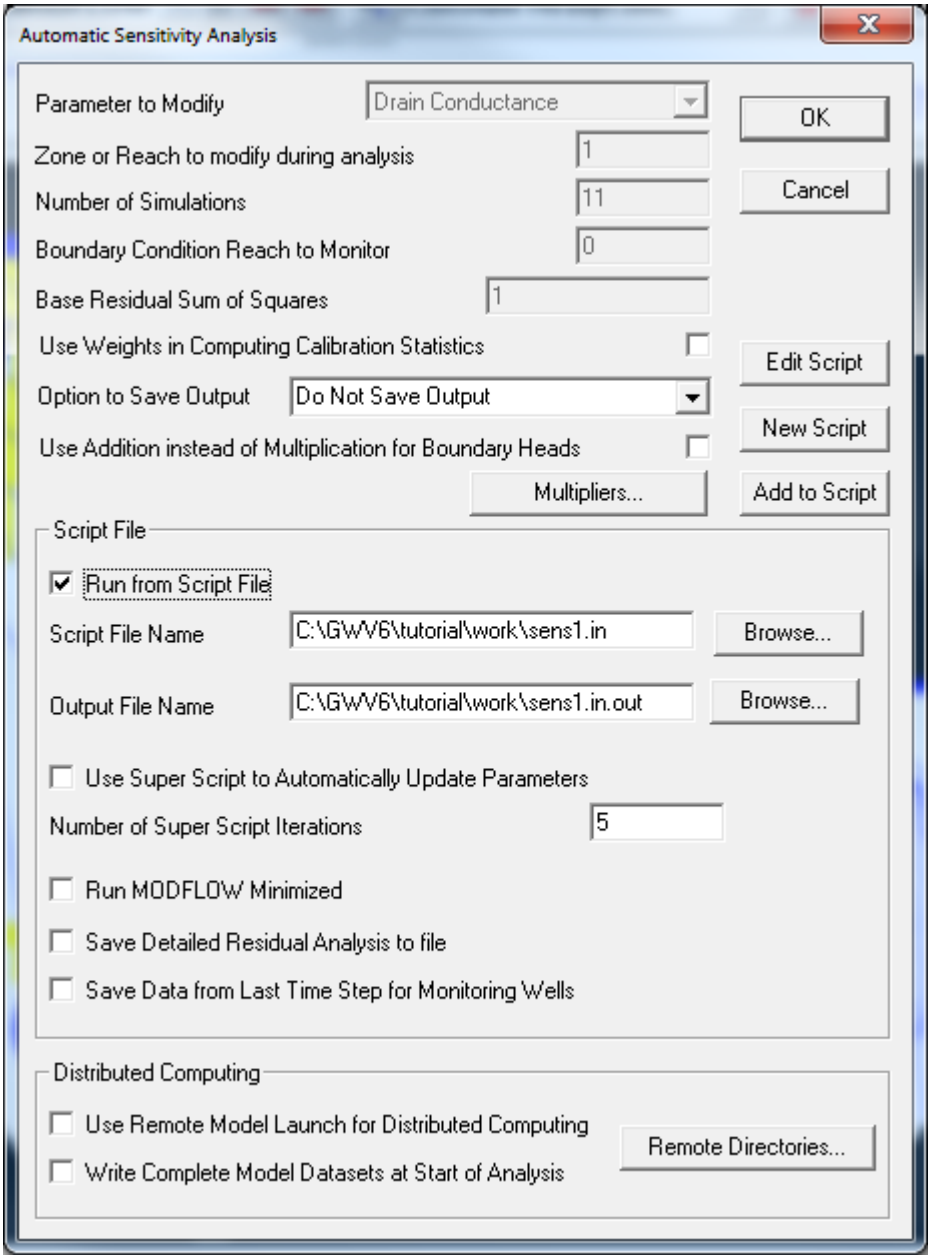

Now, start the sensitivity analysis by selecting **Model|Auto Sensitivity|Run Analysis**. The model will run 66 times (6 parameters x 11 runs per parameter), which should not be too long since this is a relatively simple model.

When the model is done running, you will be asked if you want to create a plot. Select Yes, and then just click OK on the next dialog that prompts for the file name. GV puts the results of the analysis in the script output file (see dialog above – sens*1.in.out*). This is a comma-delimited file that you can view in a text editor or in Excel. Next click OK on the dialog asking which statistic to plot. Although you may choose from a variety of statistics, the sum of squared errors is the best one to use. Your first plot should look like the following:

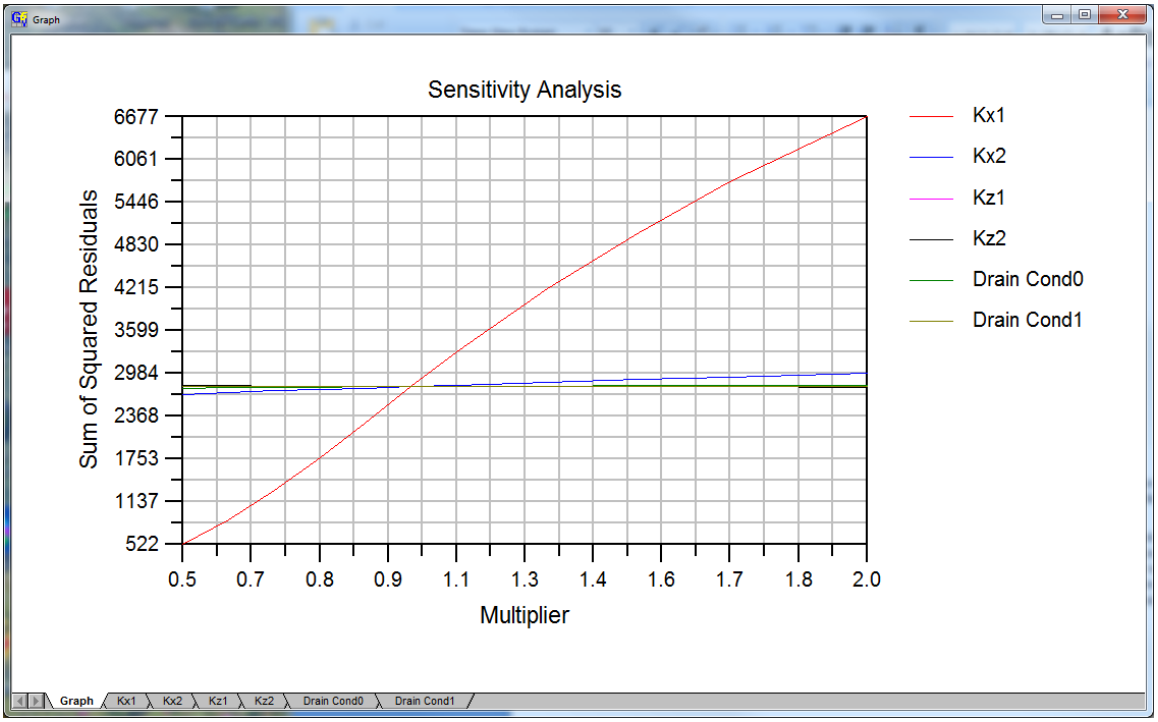

This curve shows that reducing Kx1 by a factor of 0.5 would substantially reduce the sum of squared errors and therefore improve the quality of the calibration. The initial sum of squared residuals (multiplier = 1.0) was about 2,790 in this case. If you look at the Kx1 curve above, you can see that at a multiplier of 0.5, the sum of squared residuals drops to 525, a substantial improvement. The trick to reading this graph is to look for the curve that hits the lowest point on the Y-axis (sum of squared residuals). A low sum of squares equates to a better model, assuming the parameter change makes sense. Once you find the curve that has the lowest sum of squared residuals, look at the parameter type and multiplier. If that value seems to be justified, change the parameter in Groundwater Vistas and run the model again.

Make that change now by selecting **Props|Property Values|Database**.

Multiply the Kx and Ky values for zone 1 by 0.5 and enter those values in the database.

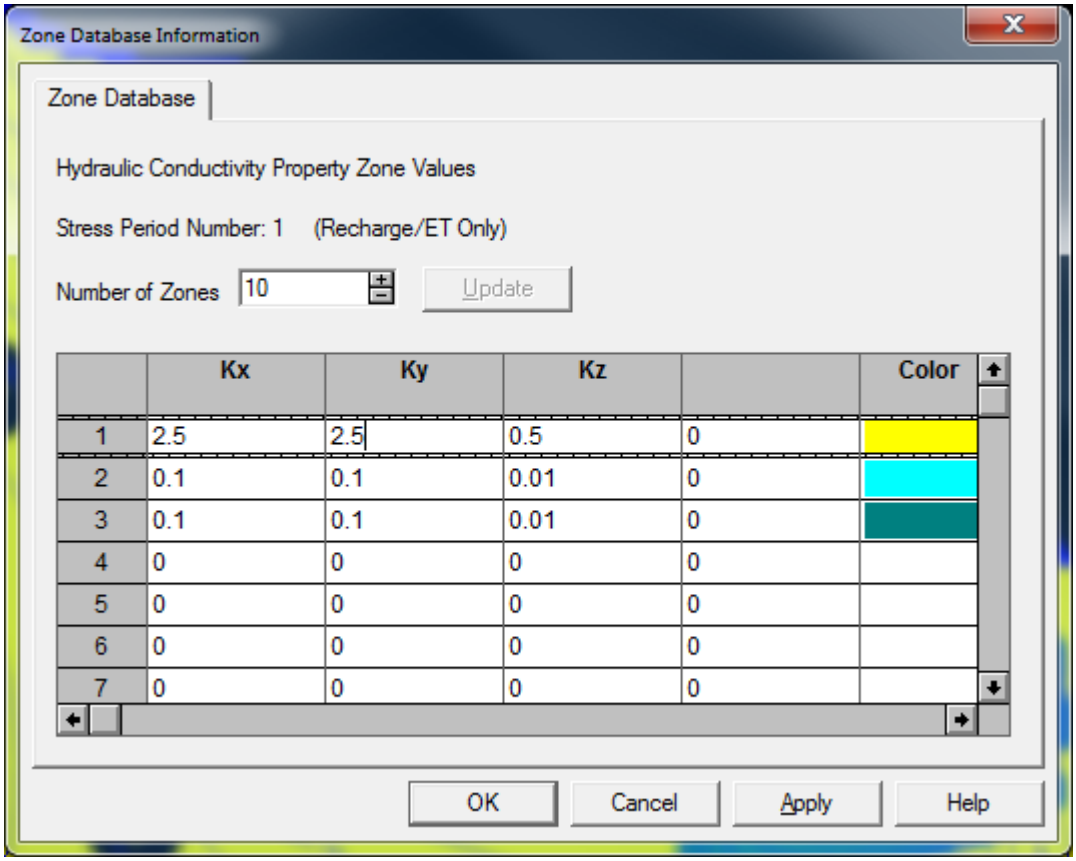

After making the change, run the model one time to confirm the results and make another pass through the automatic sensitivity analysis.

Repeat this analysis 4 more times. If you want the computer to do all the work for you, select **Model|Auto Sensitivity|Options** and turn on the Super script option at the bottom. In Super script runs, GV decides which parameter to change, makes that change, and runs another script analysis. Change the number of super script runs from 5 to 4 and try it. In this type of analysis, you don't get a graph but rather a summary of which parameters were changed and by how much. The changes are permanent in your model so if you don't like what happened, you can use the super script output file to go back and reset to the original values.

If you make 5 sets of auto-sensitivity runs (your initial one plus 4 more), you should be able to get the sum of squared errors down to about 128 and the calibration should look reasonably good. If you want to compare your progress to the theoretically "correct" answer, here are the four additional runs.

In run 2, the Kx1 curve now reaches a minimum value at a multiplier of 0.7. This brings the sum of squared residuals from 525 to about 293. If you are unsure of the exact multiplier, just click on the Kx1 tab at the bottom of the graph window. This will show a spreadsheet of multiplier versus sum of squared residuals so you can see the raw data that went into the graph for this parameter. Change Kx and Ky in zone 1 from 2.5 to 1.75.

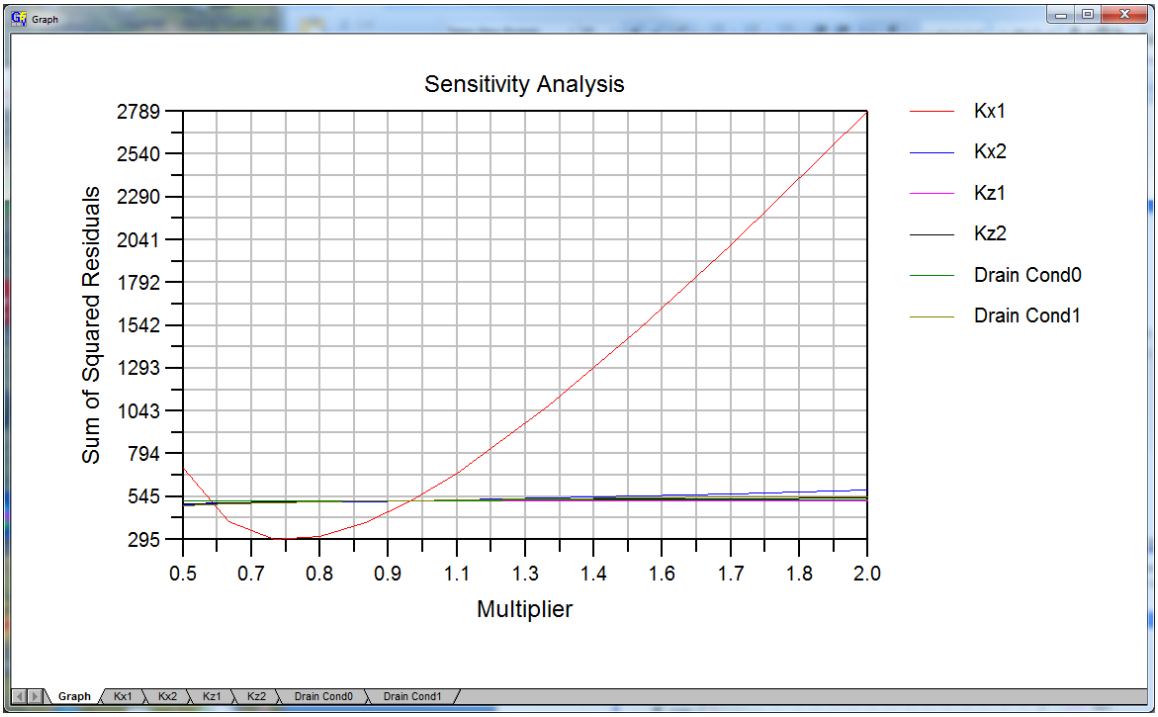

In run 3, the results are starting to look like they are not getting much better. This is partially because of the scaling on the Y-axis where there are some large sum of squared residual values. The graph right after the run looks like this.

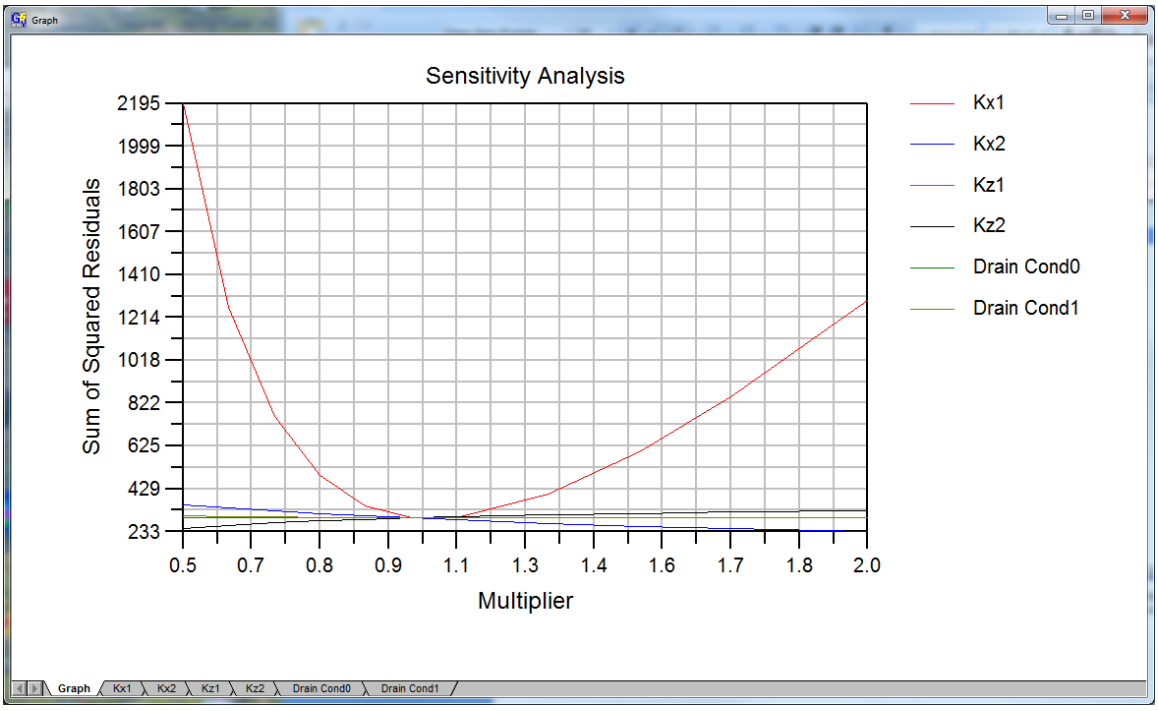

In order to better see what is going on, we need to change the scaling of the Y-axis. Right click the graph and choose properties. On the Y-axis tab, change the maximum Y to 500. Now the graph is easier to read, as shown below.

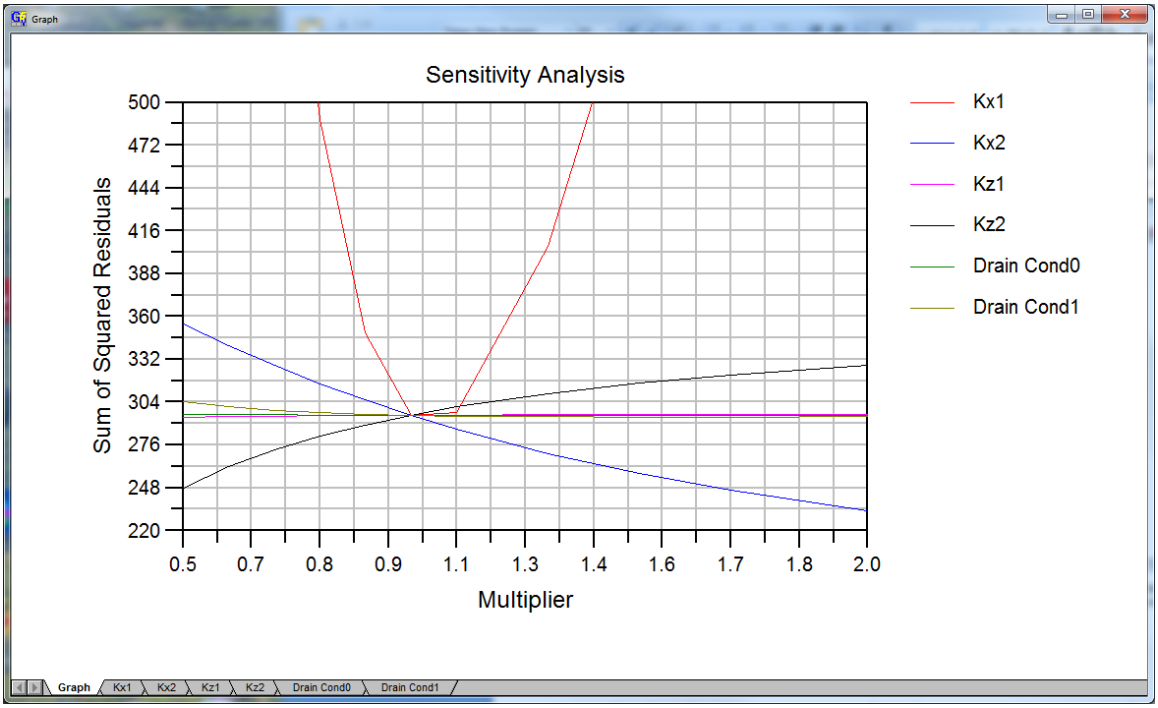

It is easier to see now that multiplying Kx2 by 2.0 will lower the sum of squared residuals from 293 to about 233. Go to the property database and make that change to both Kx and Ky.

The results of run 4 need to be scaled as in run 3. After doing that, it is clear that multiplying Kz in Zone 2 by 0.5 will reduce the sum of squared down to about 178. Go to the database and change Kz2 to 0.005.

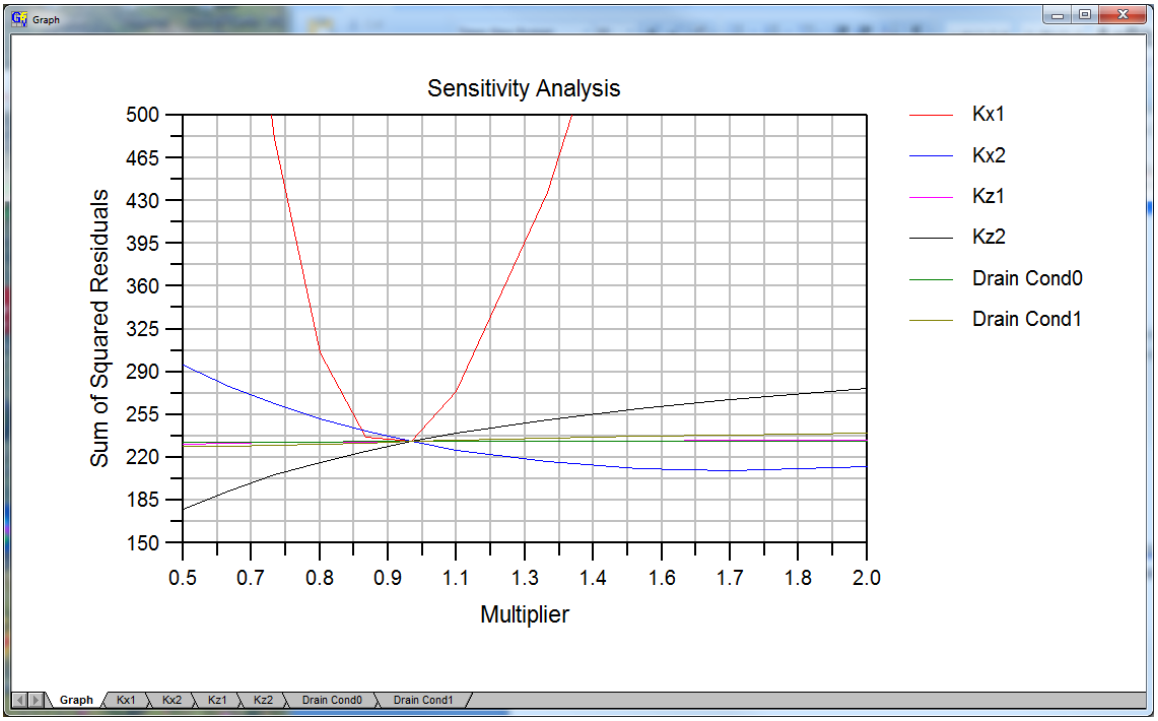

After run 5, the results are like run 4, as shown below. Reducing Kz by another factor of 0.5 will lower the sum of squared residuals to about 123.

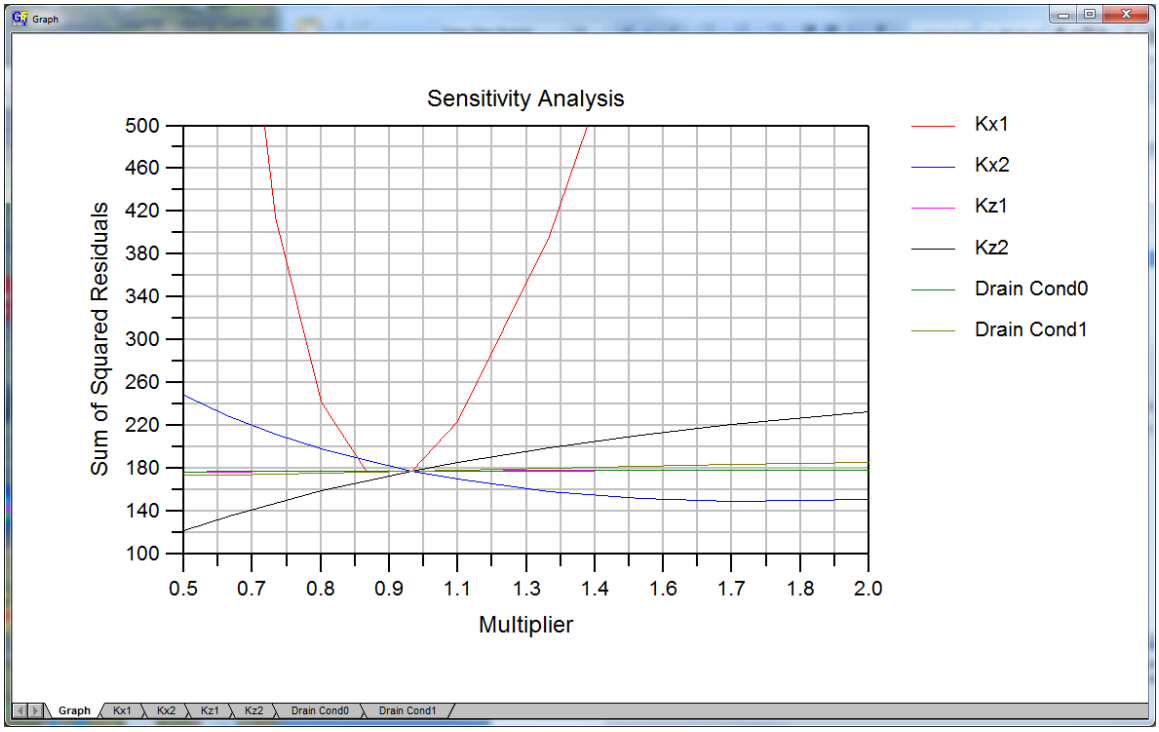

In summary, we made 5 sets of 66 runs each. After each set, we changed one parameter value and lowered our sum of squared residuals substantially from 2,790 to about 123. In addition, the process is logical and can be easily explained and documented.

Introduction to PEST

The following tutorial continues on from the previous session and introduces you to the Pest inverse model. You will start with a steadystate calibration and then add a transient simulation.

Select **File|Open** and find the file *CalEx1.gwv*. Click the calculator button to run the model one time just to make sure all files are up to date.

PEST has its own set of menus under **Model|Pest**.

First select **Model|Pest|Options - Basic Options**. This sets the main options for the Pest run.

Put a check in the option to put the BCF Package arrays in a separate file ("Write arrays as external files"). This is technically not necessary for MODFLOW2000 with the LPF Package but it is a good habit to get into as many times you will need to do this with PEST and it does not hurt to do it even if it is not necessary.

Put a check in the option to run models without screen output (this allows you to see the Pest screen output). Finally, put a check in the option to "Use Command-line versions". The latter will run Pest without the ESI versions of MODFLOW flashing up a window each time the model is run. Your dialog should look like the following:

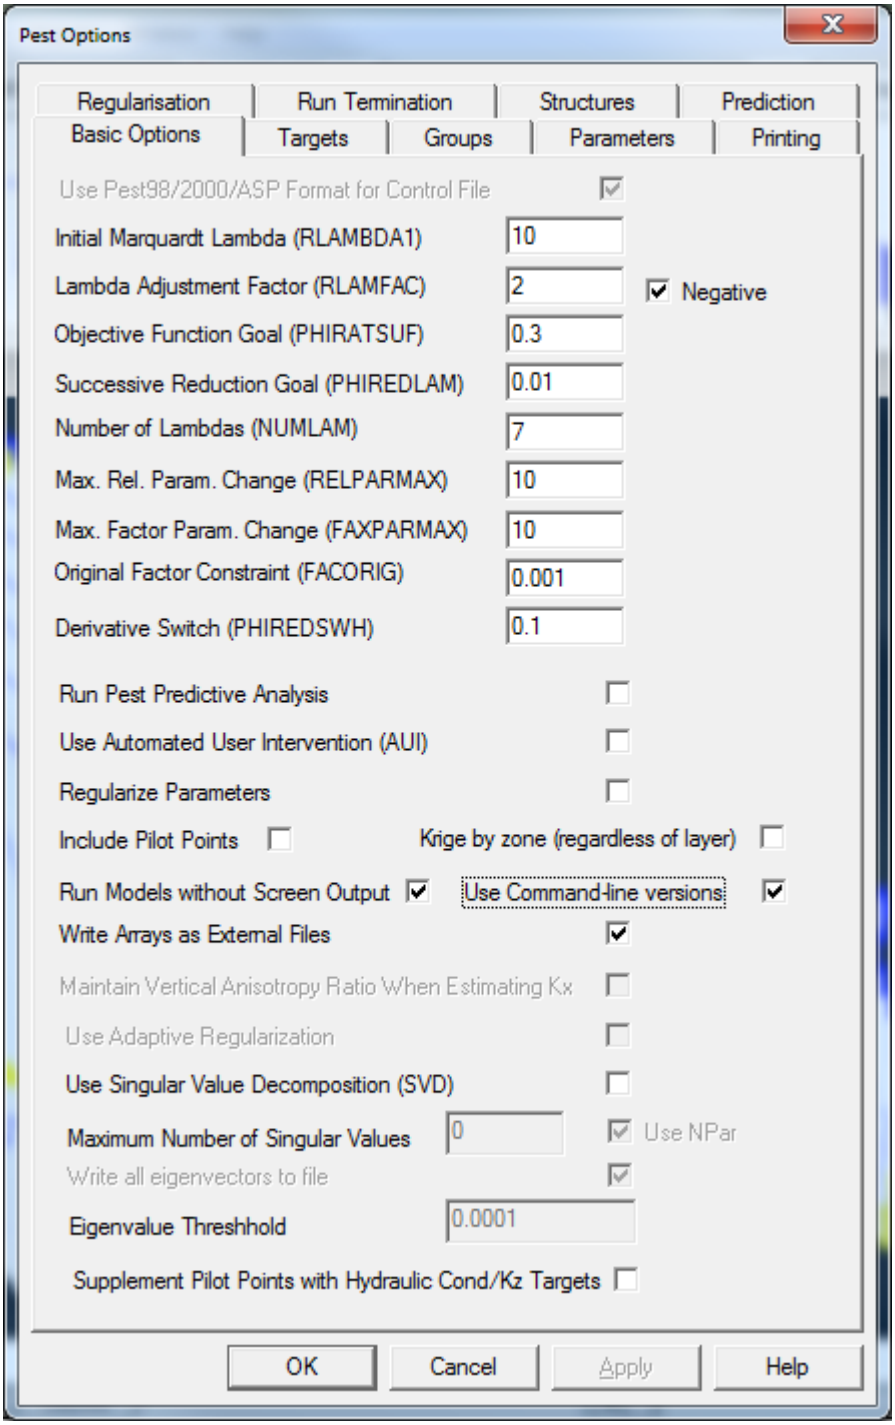

Next set up the Pest parameters by selecting **Model|Pest|Parameters**.

Change the first one to Kx Zone Number

Enter zone 1 in the field to the right of the parameter name.

Set the transformation to “log” and the limit type to “factor”.

Repeat this for Kx in Zone 2, Kz in Zone 1, and Kz in Zone 2.

An easier way to do this is to use the **Parameter Setup** button. If you use this approach, you would enter the following for Kx zones:

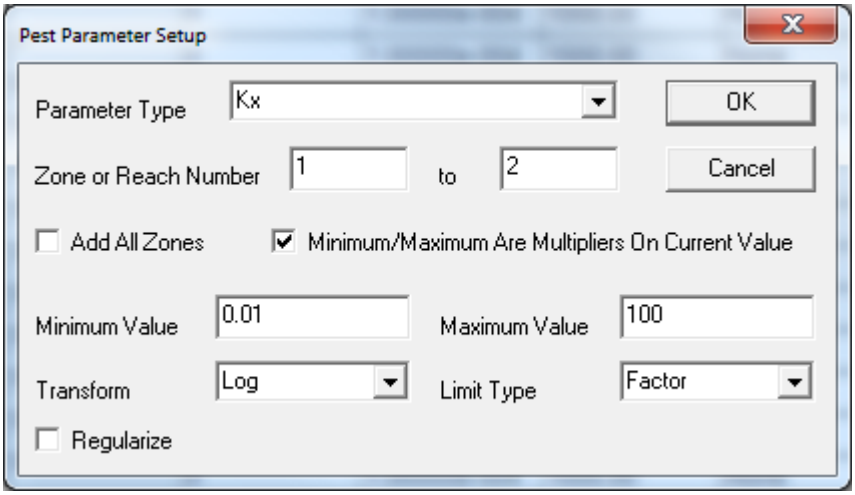

Note that by checking the option that the minimum and maximum values are multipliers you don't need to remember what the current values are in these zones. Also, for Kx, Kz, and storage, you can check the "Add All Zones" option and Vistas will put all K zones in the parameter spreadsheet. Note, however, if you do that for this problem you need to uncheck the "Use" column for zone 3.

We will also be estimating drain conductance. Add Drain conductance in reach 0 and Drain conductance in Reach 1. Use the "log" and "factor" options for all of them. Also, make sure to increase the *Parameter Upper Bound* for the drain conductances to 1e+05.

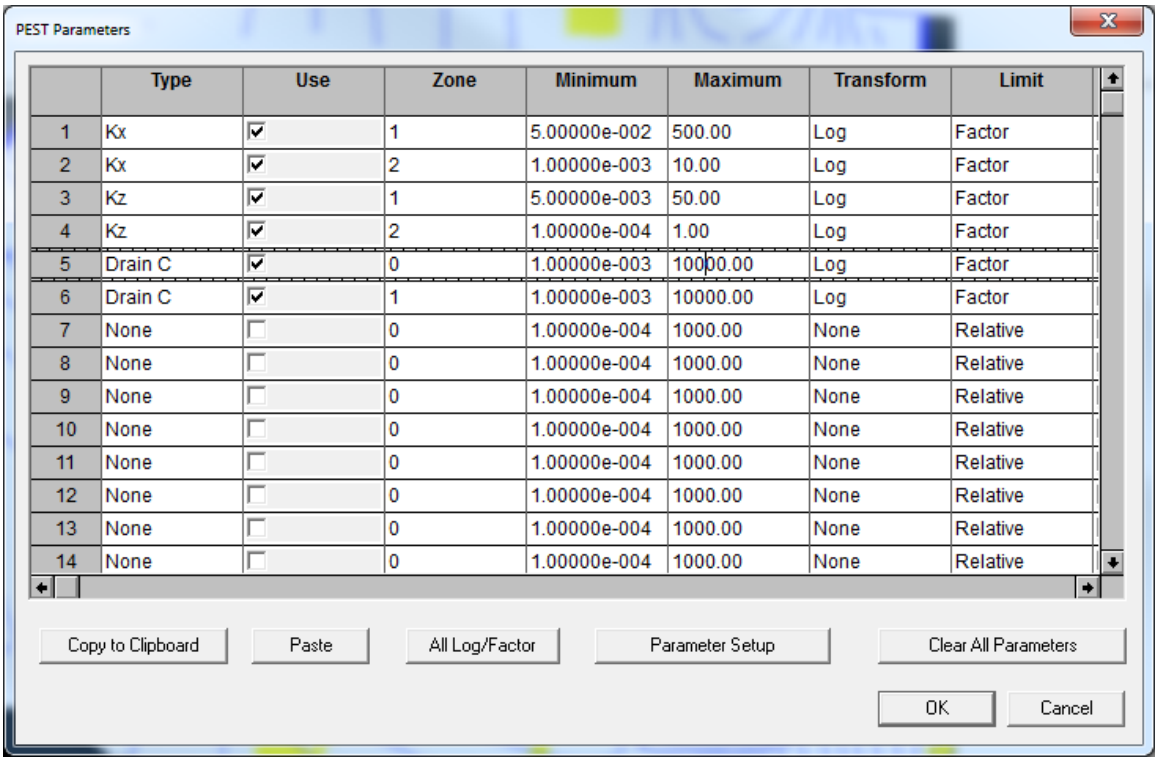

	Type	Use	Zone	Minimum	Maximum	Transform	Limit
1	Kx	☑	1	5.00000e-002	500.00	Log	Factor
2	Kx	☑	2	1.00000e-003	10.00	Log	Factor
3	Kz	☑	1	5.00000e-003	50.00	Log	Factor
4	Kz	☑	2	1.00000e-004	1.00	Log	Factor
5	Drain C	☑	0	1.00000e-003	10000.00	Log	Factor
6	Drain C	☑	1	1.00000e-003	10000.00	Log	Factor
7	None	☐	0	1.00000e-004	1000.00	None	Relative
8	None	☐	0	1.00000e-004	1000.00	None	Relative
9	None	☐	0	1.00000e-004	1000.00	None	Relative
10	None	☐	0	1.00000e-004	1000.00	None	Relative
11	None	☐	0	1.00000e-004	1000.00	None	Relative
12	None	☐	0	1.00000e-004	1000.00	None	Relative
13	None	☐	0	1.00000e-004	1000.00	None	Relative
14	None	☐	0	1.00000e-004	1000.00	None	Relative

Now, we will create the files necessary to run PEST.

First, select **Model|MODFLOW2000|Create Datasets** to write out the external arrays for the BCF package.

Then select **Model|Pest|Create Datasets**.

It is always a good idea with PEST to run the PESTCHEK utility. PESTCHEK looks at your PEST input file (root.pst) to see if there are any problems. To run it from GV, select **Model|Pest|Run Pestchek**. GV then runs pestchek and puts the screen output into a text file and then displays the text file. You should see no errors but two warnings. (The warnings are because gradient targets are turned on but there are no gradient targets in your model. This won't hurt anything but you can turn them off by selecting **Model|Pest|Options – Targets tab**).

You can now run Pest by selecting **Model|Pest|Run Pest**. Pest will start by running the model with the current parameter values. Pest will then display the starting sum of squared residuals (phi) and tell you how many times it will run the model during the first iteration, as shown below.

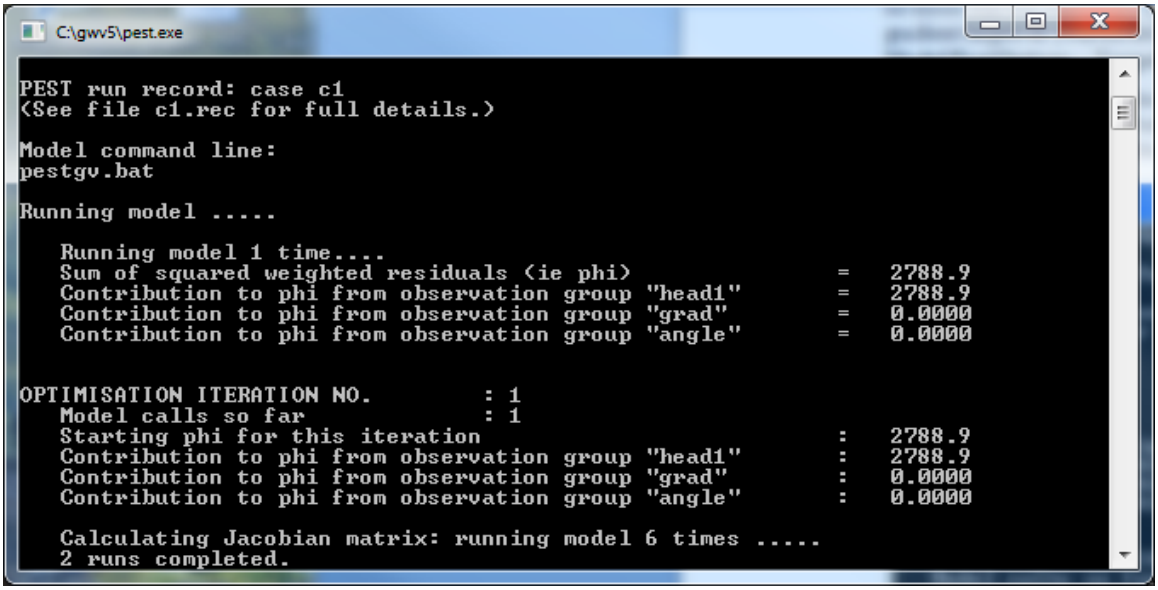

The starting phi (2788.9 ft^2 in this case) should closely match the value you get by running the model one time and selecting **Plot|Calibration|Stats**. If the Pest phi does not match the sum of squared residuals reported by Vistas, then something is wrong.

Now Pest will run the model one time for each parameter (6 times in this example). After these first 6 runs, Pest will estimate a new set of parameters using the starting Lambda value. Next, Pest will test multiple Lambda values to see if it can do a better job with more Lambda values.

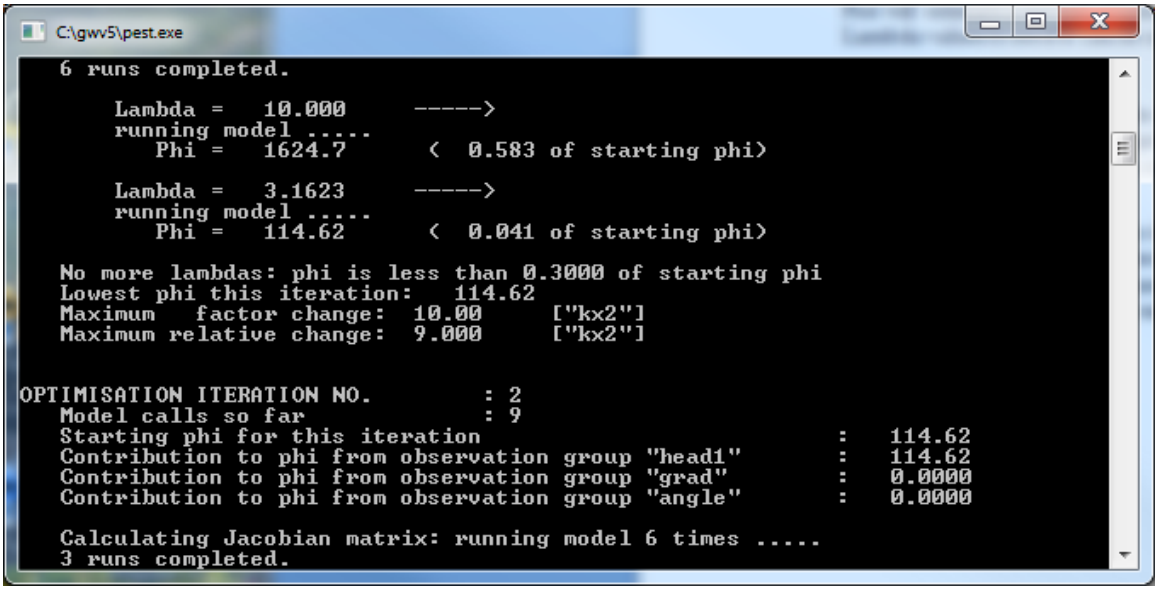

In this first iteration, Pest used the default Lambda value and reduced phi to 1,624.7. It then reduced Lambda from 10 to 3.16 and got a huge improvement in phi, down to 114.62. This was such a huge step (less than 0.3 of starting phi) that it moves on to the next iteration (Note that you can control that 0.3 value under the various Pest options).

After a few iterations, the phi has dropped to the level where it is not changing much between iterations. Pest will now start running the model twice for each parameter. This was done because the "switch" option is the default option that Vistas uses. It is generally a good idea to use this approach because it is very common to make good progress early in the Pest run and then as the model gets closer to the best fit, the change in phi slows.

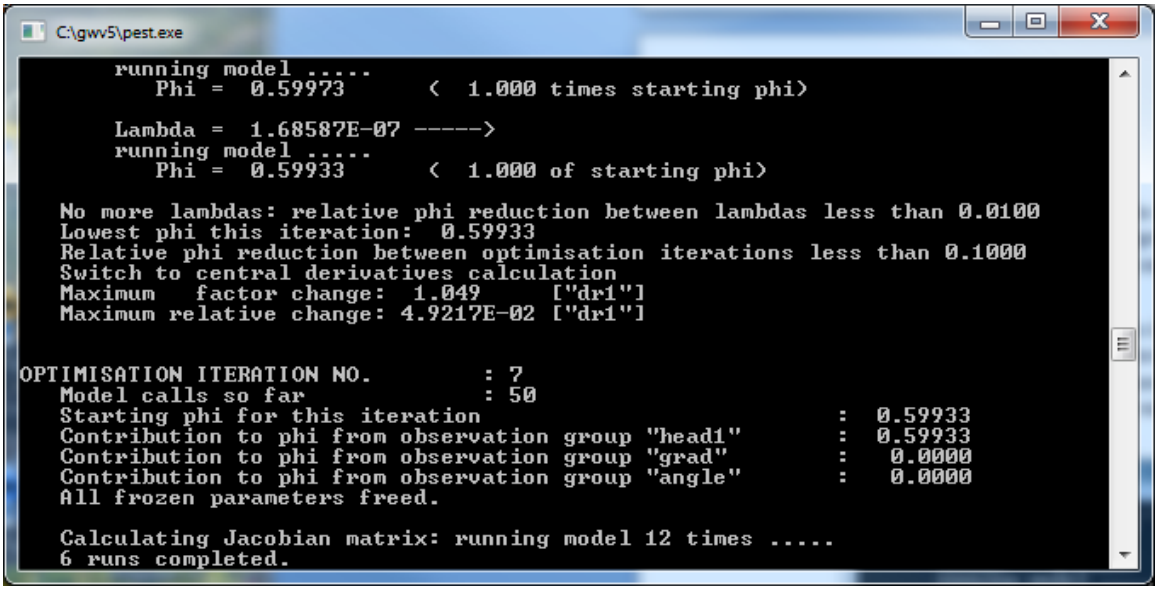

So in this example, the number of model runs eventually jumps from 6 runs to 12 runs. Phi has dropped down to less than 1, which is several orders of magnitude better than the initial phi. When this happens, you know you are close to a solution.

Eventually Pest will get to a point where phi is not changing very much at all. Pest will then report it is done and run the model one last time with the best estimates. This is handy because you can then simply choose **Plot|Import Results** in Vistas to see the calibrated head field.

Import results after your Pest run and your screen should look like the one below. Note that the residuals posted next to each target are very low numbers. You should also check the calibration statistics to confirm that your sum of squared residuals is close to the final Pest value (0.599 in this case).

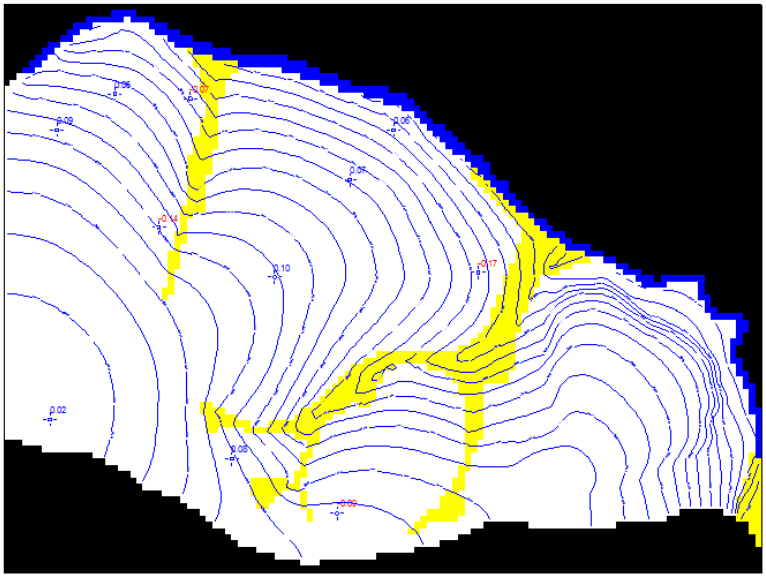

You can also view the Pest text output when you are done by selecting **Model|Pest|View Main Output**. If you plot the final sensitivities for the Pest run (**Plot|Calibration|Pest|Final Sensitivities**), you can see that Kx in zone 1 is the most sensitive parameter and drain reach zero is the lowest in terms of sensitivity.

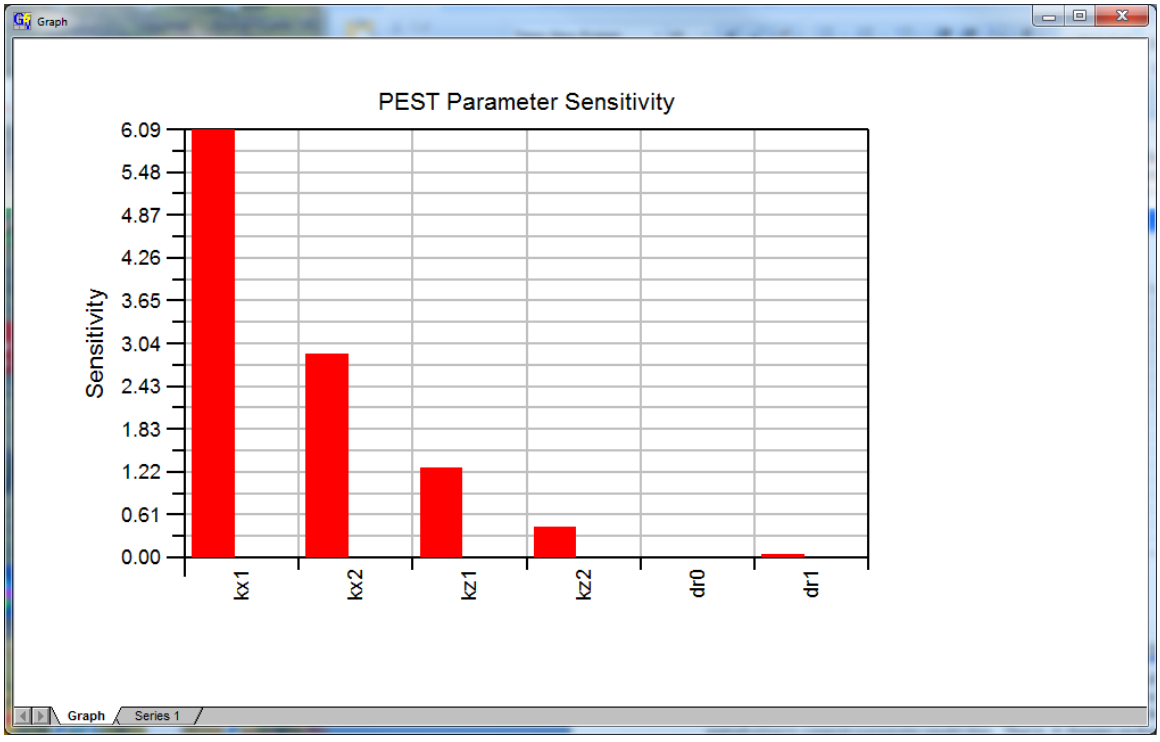

In the past we also used MODFLOW2000 in this calibration exercise and found that MODFLOW2000 cannot get nearly as good a fit as Pest unless Drain reach 0 is turned off as a parameter. This is an interesting result because MODFLOW2000 and PEST compute sensitivities in a different way. Pest uses perturbation to compute parameter sensitivities. That is, it changes each parameter slightly and then compares the change in error at each target to the base case. MODFLOW2000 computes sensitivities analytically which the authors at the USGS claim yields much more accurate results. Obviously this is not the whole story since Pest is so much better, at least in this case.

Before we do anything else, use **File|Save As** and give this version of the model a new name, such as CalEx1_PestSetup.gwv.

If you like the results of the Pest calibration run, the next step is to take the final parameter estimates from Pest, enter those in the current model, and run it again. Groundwater Vistas offers a simple way of doing this. Simply select **Model|Pest|Update Parameters** and answer Yes to the next prompt. Before doing anything else, it is a good time to save the model under a new name (e.g. CalEx1_PestResults.gwv) using **File|Save As**. It is good practice to save at least two versions of each calibration run. One is the setup of the run which we saved in the previous paragraph and the second is the result, which we save now. That way if you ever want to go back and reproduce what you did, you can do that or you can test the effectiveness of other Pest options and rerun it.

PEST and Pilot Points

Introduction

Written by

John Doherty
Watermark Computing

Adapted for Groundwater Vistas by

Jim Rumbaugh
Environmental Simulations, Inc.

This tutorial demonstrates some of the advanced and exciting opportunities for model calibration available through PEST-ASP and supplementary software available through the Groundwater Data Utilities. By the time you have finished this exercise you will have:

- used pilot points as a means of characterising the spatial distribution of an aquifer hydraulic property;
- used PEST's advanced regularisation functionality in conjunction with geostatistically-based regularisation constraints;
- used Groundwater Vistas to assist in PEST input file preparation.

This exercise is based on the Groundwater Visas (GV) graphical user interface (GUI), which simplifies the use of PEST-ASP and some of the utility programs related to pilot point calibration.

Why not use MODFLOW-2000?

At first site, using PEST to calibrate a MODFLOW model may seem to be an old-fashioned way of doing things; after all, parameter estimation functionality is available through MODFLOW-2000. Why not use that?

There are a number of reasons why the parameter estimation capabilities of MODFLOW-2000 could not be used in the present circumstances. They are as follows:-

MODFLOW-2000 parameter estimation functionality does not extend to the calibration of a composite model comprised of MODFLOW and MT3D. Thus borehole concentration measurements cannot be used in a MODFLOW-2000 parameter estimation run.

MODFLOW-2000 has no predictive analysis or regularisation functionality. As the present exercise shows, both of these can be extremely useful in many modelling contexts.

Many cases have been encountered where PEST's inversion engine is stronger than that of MODFLOW-2000, especially where many parameters require simultaneous estimation.

The "parameter estimation" process of MODFLOW-2000 does not work well where hydraulic conductivity or transmissivity parameterisation within a model domain is based on pilot points, and where the distribution of these properties within the model domain is best described by a log variogram.

If parameter estimation was required only for a MODFLOW model, and MT3D was thus not included in the parameter estimation and predictive analysis processes documented herein, then these problems could have been overcome by carrying out the following procedure:-

Using Groundwater Vistas, set up a MODFLOW-2000 model based on a simple parameterisation of the model domain.

Generate a set of MODFLOW-2000 input files for that problem.

Use programs PPK2FAC and FAC2MF2K from the Groundwater Data Utilities to create a set of MODFLOW-2000 input files in which the simplified parameterisation of the model domain is replaced by a more complex parameterisation based on pilot points.

Use the MODFLOW2000-to-PEST translator MF2PEST to build a set of PEST input files based on the enhanced MODFLOW-2000 input dataset.

Use program PPKREG from the Groundwater Data Utilities to incorporate geostatistically-based regularisation information into the PEST input dataset.

Run PEST to calibrate the model. However instead of using the normal version of MODFLOW-2000 in conjunction with PEST, use a special version of MODFLOW-2000 (supplied with PEST) known as MF2KASP. This overcomes the problems associated with log variograms mentioned above.

Strategy used in the Present Exercise

In the present exercise a slightly different strategy will be used. This strategy is a little simpler in some ways than that described above, and easily incorporates the inclusion of MT3DMS into the calibration process. Neither the *sensitivity* nor *parameter estimation* processes of MODFLOW-2000 will be required, for MODFLOW will only be run as a "forward model"; ie. it will only be run to calculate outputs - in this case head outputs. Furthermore, spatial interpolation of MODFLOW-generated heads to the sites of bores will not be undertaken using the MODFLOW-2000 *observation* process; rather it will be undertaken using Groundwater Vistas' targpest program. In some contexts, as is demonstrated below, this facilitates preparation for a PEST run. Furthermore, the same utilities can be used to interpolate MT3DMS concentration outputs to the sites of bores. (There is no *observation* process within MT3DMS to do this, so the use of such external spatial interpolation software is a necessity.)

In the present exercise the spatial distribution of transmissivity (rather than of hydraulic conductivity) within the model domain will be estimated on the basis of borehole head and concentration measurements. Estimation of transmissivity is entirely appropriate in the present case because we will be simulating steady state water levels in a heterogenous aquifer in which the spatial distribution of bedrock depth is only poorly known. Because parameter estimation is restricted to transmissivity only, the MODFLOW *bcf* package will be used in preference to the MODFLOW-2000 *lpf* or *huf* packages. Use of this strategy allows us to set the MODFLOW LAYCON variable to zero for the single layer comprising the model; this allows MODFLOW to run quickly (an advantage in the present circumstances), and removes the possibility of any numerical problems being encountered through the drying and rewetting of cells.

For the present exercise the MODFLOW, rather than MODFLOW-2000, version of MODFLOW is used. However use of MODFLOW-2000 would be almost identical.

Description of Study Area

Figure 1 shows a MODFLOW grid set up over a (theoretical) study area. Grid cells are 40m wide, except for a band of cells which cross in the centre of the study area which are 20m wide.

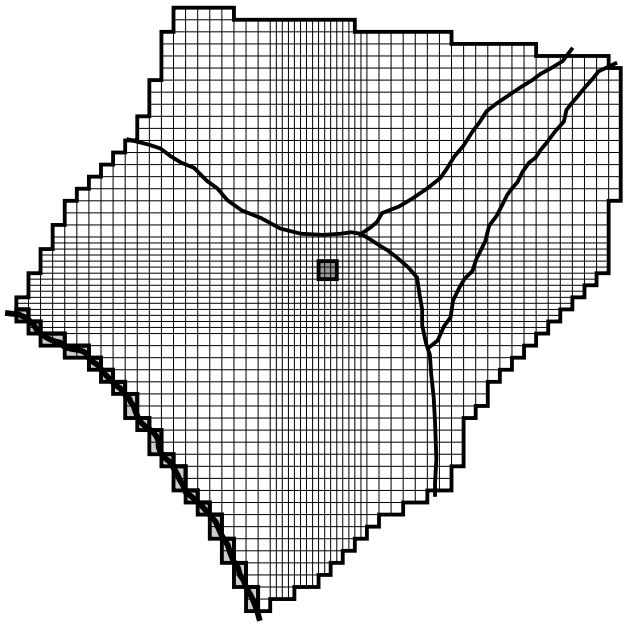

Figure 1. Map of finite-difference grid, river, source of contamination and alluvial boundaries. Scale: 1cm=200m.

A river runs along the south-western boundary of the study area, simulated by a line of fixed head cells in the model; the head in the river is 20m. To the north-east of the river, and running parallel to it, is a broad band of alluvium. Not many holes have been drilled all of the way through this alluvial material into the rocks which underlie it; hence information on basement depth is only sporadic. Nevertheless, based on those holes that have been drilled to date, alluvial thickness seems to average about 30m. An old creek channel, oriented in a north-easterly direction, exists in the north of the study area. The alluvium in this old creek channel is not quite as thick, and not quite as hydraulically conductive, as that in the main alluvial band running parallel to the river. To the east and west of this old creek channel are basaltic rocks. Though not as permeable as the alluvium, these rocks are nevertheless able to conduct water through fractures and vesicles, the latter occurring especially at the interface between successive flows.

Because the present study area is actually a figment of my imagination, I am able to disclose to you the "real" distribution of transmissivity within the geological units which underlie it; see Figure 2. (This was generated with the help of a stochastic field generator based on Mejia's algorithm.) Within the main alluvial unit, transmissivity varies between about 65 m^2|day and 140 m^2|day, being generally higher downstream and closer to the river. Within the old creek system, transmissivity varies between 60m^2|day and 80m^2|day. Transmissivity is uniformly 20m^2|day in the western basalt and uniformly 30m^2|day in the eastern basalt.

Effective porosity of all alluvial material is 0.07.

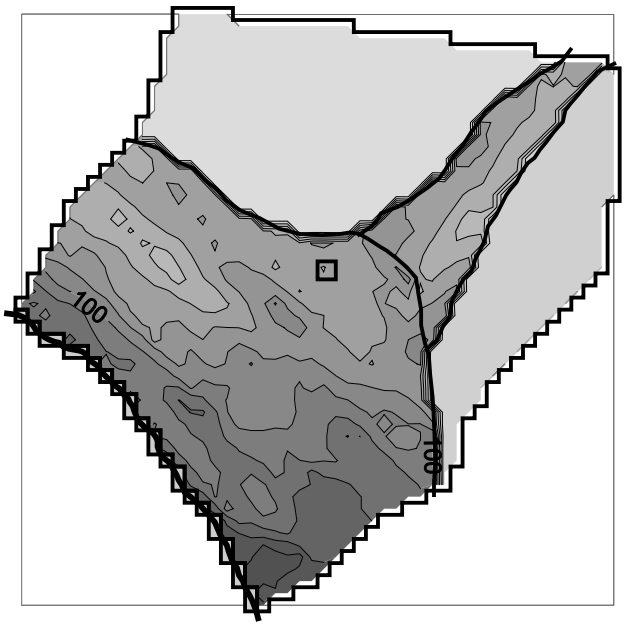

Figure 2. Transmissivity distribution within the study area. Contour interval is 10m^2|day; darker shade indicates greater transmissivity. The site where contaminant is introduced to the groundwater system is also shown.

The north-west and south-east boundaries of the model domain have been chosen to coincide with groundwater flow lines. The north and north-eastern boundaries coincide with a groundwater divide.

Recharge to the groundwater system is different within each of the geological units prevailing in the study area. Within the broad alluvial system associated with the river, recharge is 7×10^{-4}m|day (about 255 mm|yr), while within the smaller alluvial unit comprising the old creek system, recharge is 5×10^{-4}m|day (about 180 mm|yr). In the basalt unit to the west of this, recharge is 1×10^{-4}m|day (about 36 mm|yr) whereas in the basalt unit to the east of the old creek system, recharge is 2×10^{-4}m|day (about 72 mm|yr).

The focus of interest in the present study is a site whose dimensions are about 60m ×60m situated within the main alluvial unit near the place where it is joined by the old creek system; see Figure 2. Drilling and water sampling within the main alluvial unit to the south-west of that site has revealed a contaminant plume originating from under the site. A study of management records at the site has revealed that leakage has probably been occurring for about 1500 days; however as soon as it was detected, measures were immediately taken to prevent any further leakage. The contaminant leakage rate was such that there was minimal effect on the overall groundwater recharge rate under the site. However recharge waters entering the groundwater system under the site over the 1500 days during which leakage took place contained contaminant at a concentration of 100 units|m^3.

Figure 3 shows the steady state piezometric surface within the study area, as well as the disposition of the contaminant plume at the time its presence was first detected, ie. after 1500 days of continuous leakage. As is the case for the geological map pictured in Figure 2, when we come to calibrate a model for this study area shortly, we will not have the luxury of knowing either the head or contaminant distribution in anything like the detail pictured in Figure 3.

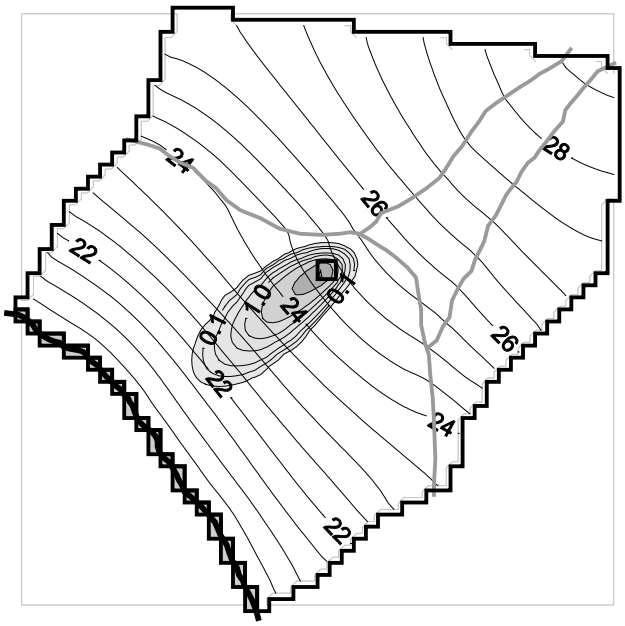

Figure 3. Piezometric and contaminant concentration contours 1500 days after commencement of leakage. Heads are in metres. In this, and similar diagrams, a logarithmic contour interval is used for contaminant concentration, the lowest contour level being 0.05 units|m^3.

Figures 4 and 5 depict the fate of the contaminant over time if no remedial action is taken. Over a period of 5000 days, starting at the time of commencement of leakage into the groundwater system, 324112 units of contaminant will enter the river at the lower end of the groundwater system. Only about 15 of these units entered the river during the 1500 days of leakage. A total of 378000 units of contaminant actually entered the groundwater system under the leaking site during the 1500 days over which leakage took place.

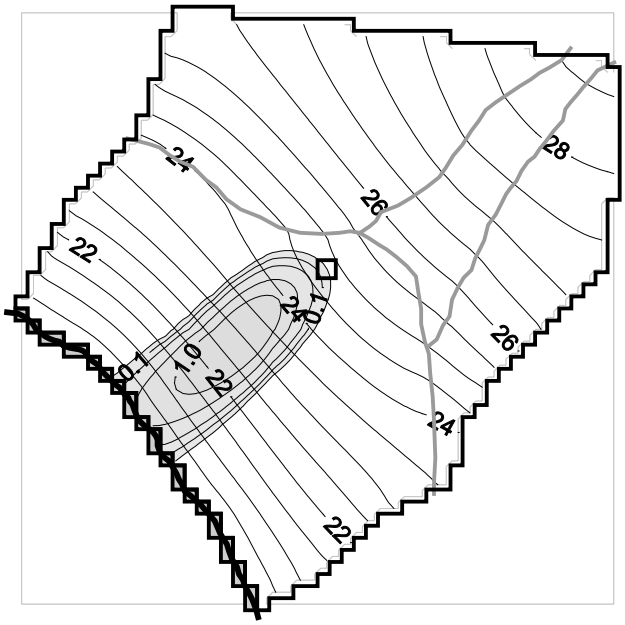

Figure 4. Piezometric and contaminant concentration contours 1500 days after cessation of leakage. Heads are in metres. A logarithmic contour interval is used for concentration, the lowest contour level being 0.05 units|m^3.

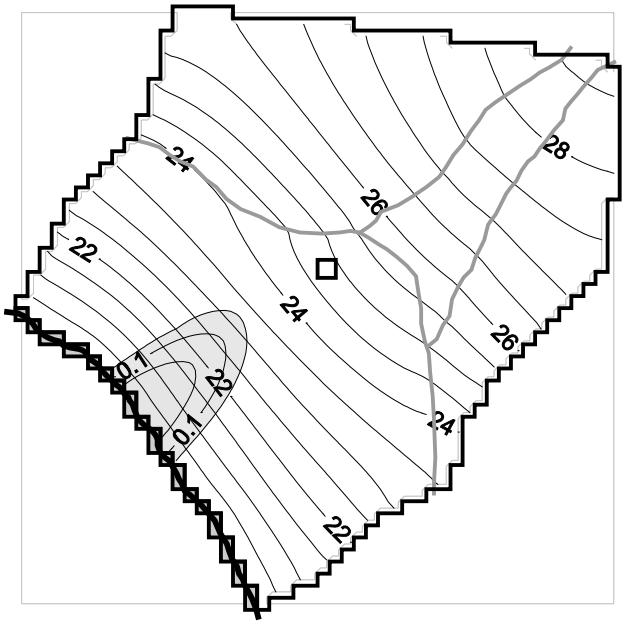

Figure 5. Piezometric and contaminant concentration contours 3500 days after cessation of leakage. Heads are in metres. A logarithmic contour interval is used for concentration, the lowest contour level being 0.05 units|m^3.

Using a Model

No model should ever be built without a clearly defined purpose in mind. In the present case the reason is clear. We have been asked to test the efficacy of extracting water from a single location within the contaminant plume, and of treating the water to remove contaminant from it. The treated water will then be released elsewhere. For the present purposes it will be assumed that it does not make its way back into the groundwater system, at least not within the model area. The aim of the remediation system is to intercept the contaminant before it emerges in the river which crosses the base of the study area. If no remediation system is in place, all of the 378000 units of contaminant that entered the groundwater system will find their way into the river; as was explained above, 324112 units will have entered the river by the time 5000 days have elapsed since leakage began. Let us assume that a regulatory requirement has been set specifying that total contaminant outflow to the river up until this time must not exceed 10000 units.

The location of the single extraction bore is shown in Figure 6 superimposed on the plume when its presence was first detected 1500 days after leakage into the groundwater system began. The bore will pump continuously at a rate of 300m^3|day (about 3.5 l|sec).

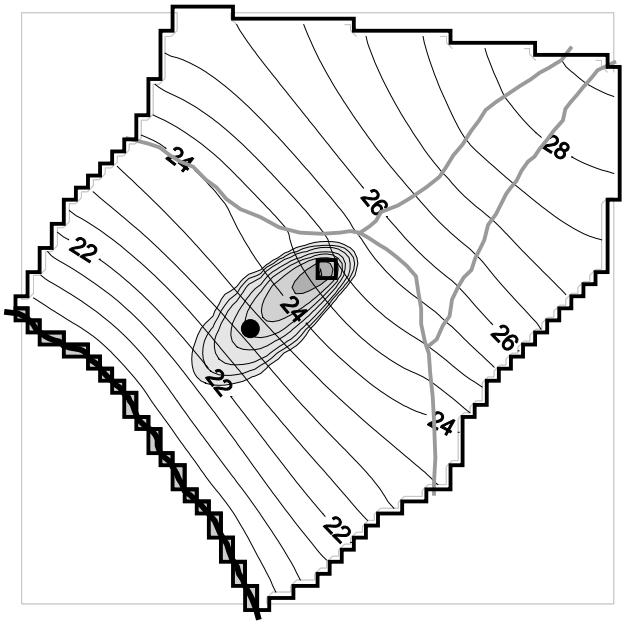

Figure 6. Location of extraction bore.

If we had complete knowledge of the hydraulic properties of the system it would be a simple matter to simulate the efficacy of the proposed remediation strategy on a computer, and thereby ensure that the regulatory requirement is met. Unfortunately, many modelling efforts are conducted under the pretence that such perfect knowledge of the system exists, for they do not take account of the fact that uncertainties in model parameterisation can lead to large uncertainties in model predictions. For the purposes of the present practical exercise we will assume that we know a lot about the system - more than we would probably know in real life; we will assume that our knowledge of recharge, the source concentration of the contaminant, the disperssivity and effective porosity of the subsurface, and the amount of time over which the contaminant was leaking, are all perfect. However we will acknowledge that we do not know the details of the geologically complex medium pictured in Figure 2. Hence we will parameterise our model in a way that accommodates the *existence* of such geological heterogeneity, and explore the predictive uncertainty arising from our lack of knowledge of the *details* of this heterogeneity.

But before we do any of that, let us find the “answer at the back of the book” by running a MODFLOW|MT3DMS model on the basis of the “true” geological conditions.

Running the Model

The "true" model is provided in Groundwater Vistas format as *ex1.gwv*. Launch Groundwater Vistas and open this file. You should see the following on your screen:

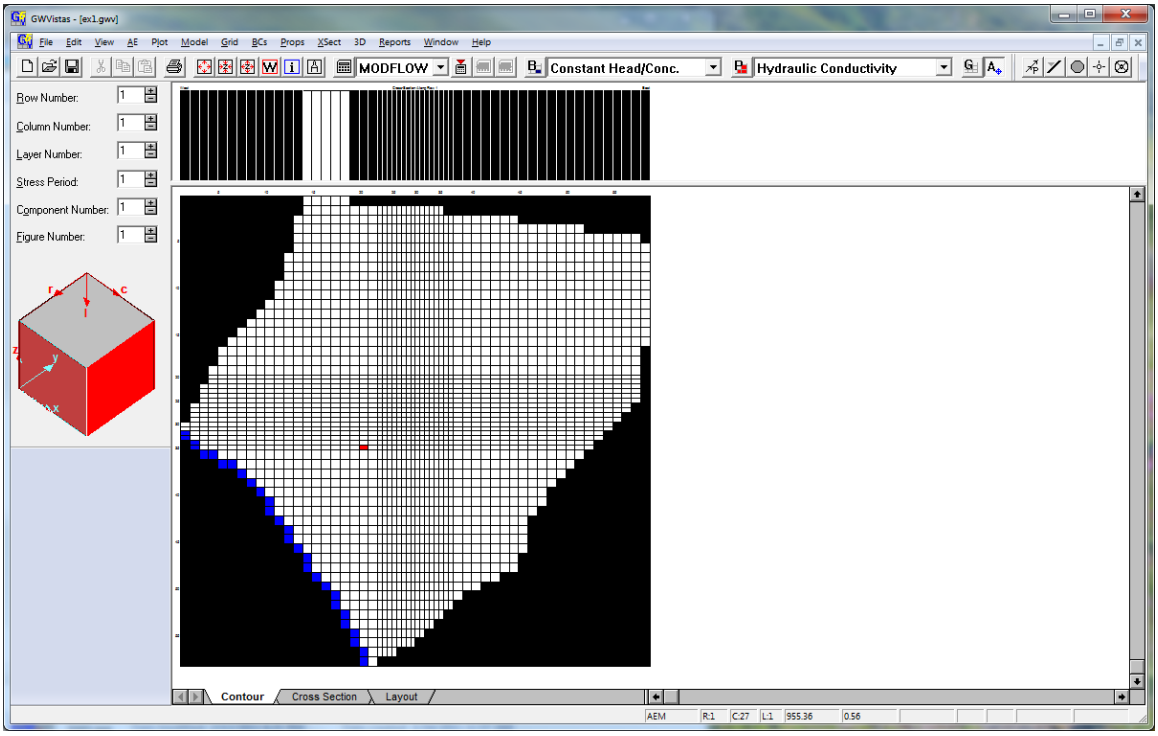

The model is run over two stress periods, the first of 1500 days duration and the second of 3500 days duration; during the second stress period the extraction bore operates continuously. Steady state conditions are assumed to prevail over both stress periods. (This is a slight violation of reality as some time must elapse before water levels adjust to the introduction of pumping; however as the adjustment time is not large in comparison to the length of the stress period, the errors incurred by this assumption are likely to be minimal. Use of this strategy has the beneficial effect that file transfer between MODFLOW and MT3D is minimised, for only a steady state flow field needs to be passed between the two models during every stress period.)

The model possesses only a single layer, three-dimensional aspects of the flow and transport processes being neglected for the present purposes.

Run the model now from Groundwater Vistas by pressing the calculator button on the toolbar.

Create datasets when prompted.

Import the results of the simulation for the first stress period and first time step. The default stress period is 2 when you are prompted to import the results. You must change it to 1 in order to have your plot look like the one below.

Your contours should be the same as those shown below:

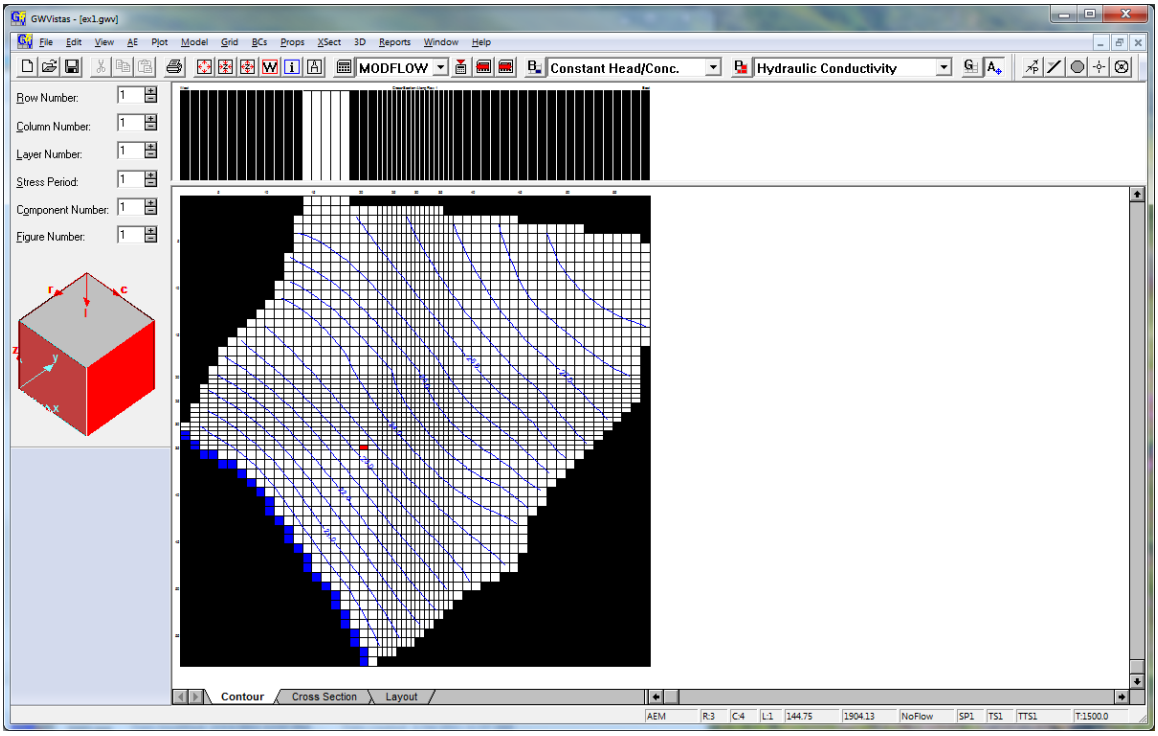

Now, close the *ex1.gwv* model. We will start with another version for the calibration exercise. Open the file *ex1a.gwv* and use that model for the remaining sections.

Calibrating the Model using Water Level Data Only

As a first step in the calibration process we will attempt to estimate the distribution of transmissivity within the model domain using borehole water level data only. It will be shown that, even where there is a relatively high borehole density, the use of water level data alone in the calibration process does not provide sufficient information to resolve geological fine detail through the calibration process. In many cases of model deployment, such as water supply applications, this fine detail does not matter. However it is often of crucial importance when predicting the movement of a contaminant through the subsurface.

When carrying out the model runs documented above, "true" hydraulic properties of the model domain were used in the model. However in the present instance our task is to attempt to infer hydraulic properties of the system by calibrating the model, and then to use the properties determined during this calibration process to make predictions using the model. As mentioned above, we will assume that recharge is perfectly known; we will also assume that the boundaries between the alluvial material and the basalt, and between the river and creek alluvia, are well mapped. So the aim of the calibration process is to infer the distribution of transmissivity within each of the four different zones within our study area, these being:-

the river alluvium (referred to as "zone 1");

the creek alluvium (referred to as "zone 2");

the western basalt (referred to as "zone 3"); and

the eastern basalt (referred to as "zone 4").

Figure 7 depicts the zonation of our model domain.

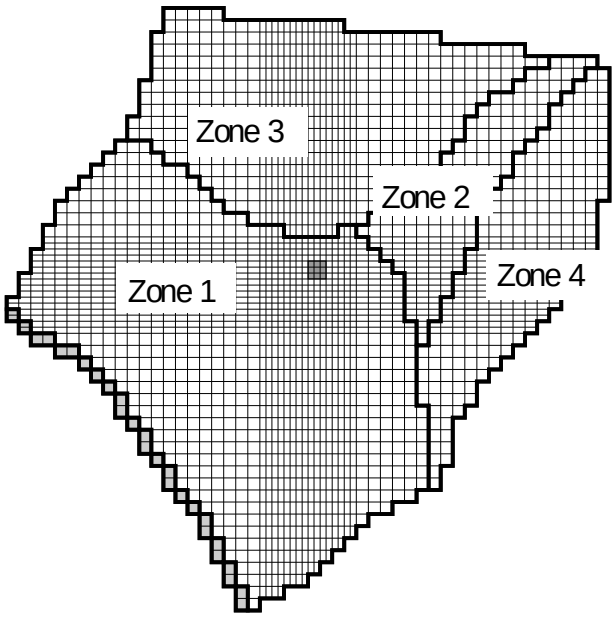

Figure 7. Zonation of the model domain.

You can see these zones by opening the Groundwater Vistas model called ex1a.gwv. Select Props|Hydraulic Conductivity and your screen should look like the one below:

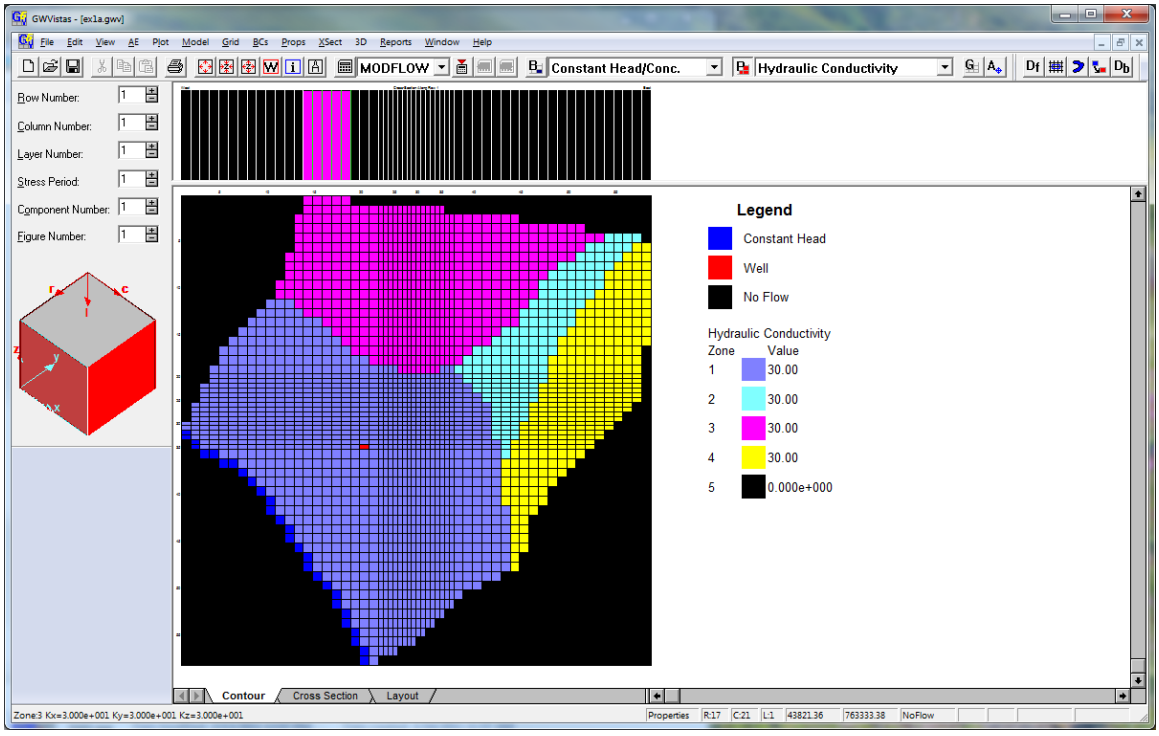

Pilot Points

Using conventional model calibration technology, the calibration process would probably proceed by defining four parameters, (viz. the transmissivity associated with each of the above zones), and adjusting these parameters until the fit between model outcomes and field observations is as good as possible. If the goodness of fit obtained on the basis of these four zones was not acceptable, then extra zones would be introduced into the model domain at locations where the modeller felt that they would "do the most good"; the parameterisation process would then be repeated with the new parameters included. If a good fit was still not obtained, more zones would then be introduced. The process would continue until the fit between model outcomes and field observations was acceptable.

There are a number of shortcomings associated with this approach. These include the following:

The procedure is quite laborious and slow.

In a case such as the present one, geological mapping provides no guidance on where to put extra zones; hence the final distribution of zones that the modeller achieves using this process is likely to be subjective and non-unique.

Characterisation of geological heterogeneity in the present study area by zones of piecewise uniformity is not in harmony with the nature of the alluvial material; therefore any zonation pattern that is finally decided upon will not "look right"; it will be defensible only on the basis that it is better to employ such a zonation scheme than to ignore geological heterogeneity altogether.

Piecewise constancy as a method of characterising geological heterogeneity lacks the flexibility required to explore the effects of small scale geological variability on model predictive uncertainty.

To overcome these problems, the distribution of transmissivity within the model domain will be described by a set of pilot points. A number of these pilot points will be introduced to the model domain and PEST will be asked to estimate the transmissivity of the aquifer at each such point. These "point-transmissivities" will then be spatially interpolated to all of the active cells within the model domain using kriging. Hence in estimating transmissivity values at pilot points, PEST will effectively be assigning parameter values to the whole of the model domain.

Individual pilot points can be assigned to different zones within the model domain. Only those points assigned to a particular zone can be used in calculating transmissivity values throughout that zone using the kriging interpolation procedure. Furthermore, the variogram upon which kriging is based can be different in each zone, reflecting differences in the geology, or in the level of heterogeneity, expected within each geological unit. Note that if only one pilot point is assigned to a particular zone, then that zone is characterised as being uniform. This will be done for zones 3 and 4 in our present study.

The locations of pilot point used in the present calibration exercise are plotted in Figure 8. The pilot points are provided to you in a text file called *c:\gwv6\tutorial\pilots.dat*. Import this file into GV as described below:

Select **AE**|**Import**|**Target Text File** (GV uses calibration targets to store pilot point information – pilot points are a type of target in GV). Locate the file called *c:\gwv6\tutorial\pilots.dat.*

Change the Target Type at the top from "Head" to "Pilot Point (Kx)" by selecting from the drop-down list of types.

Enter a value of 0 next to Layer.

Enter a value of 5 next to Target value (this means that the initial Transmissivity is in column 5).

Enter a 6 next to Minimum K and 7 next to Maximum K

Leave the other numbers on the dialog at default values.

Fill out the dialog in GV as shown below:

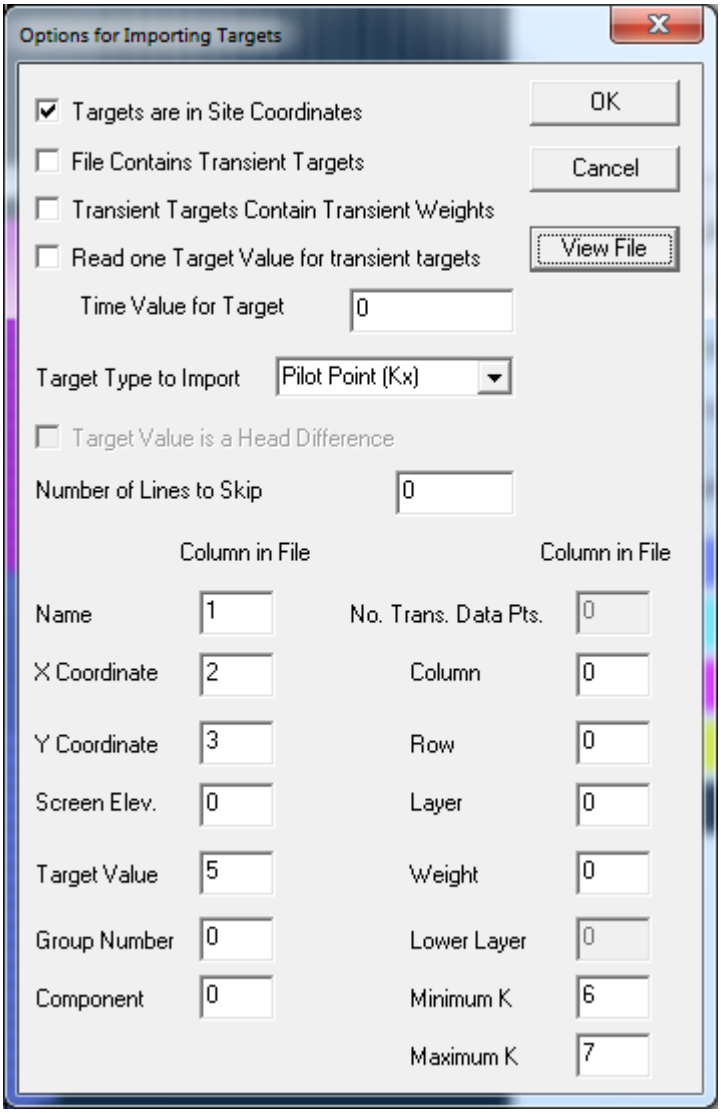

After clicking OK, GV should report that 35 targets (pilot points) were imported. Just to make sure that you imported these pilot points correctly, double-click on one of the target symbols (a circle inside a cross). **(Note: You will need to push down the A button on the toolbar in order to edit pilot points)** You should see a target dialog like the one below. Make sure the target value is 30 and the target type is *Pilot Point (Kx).*

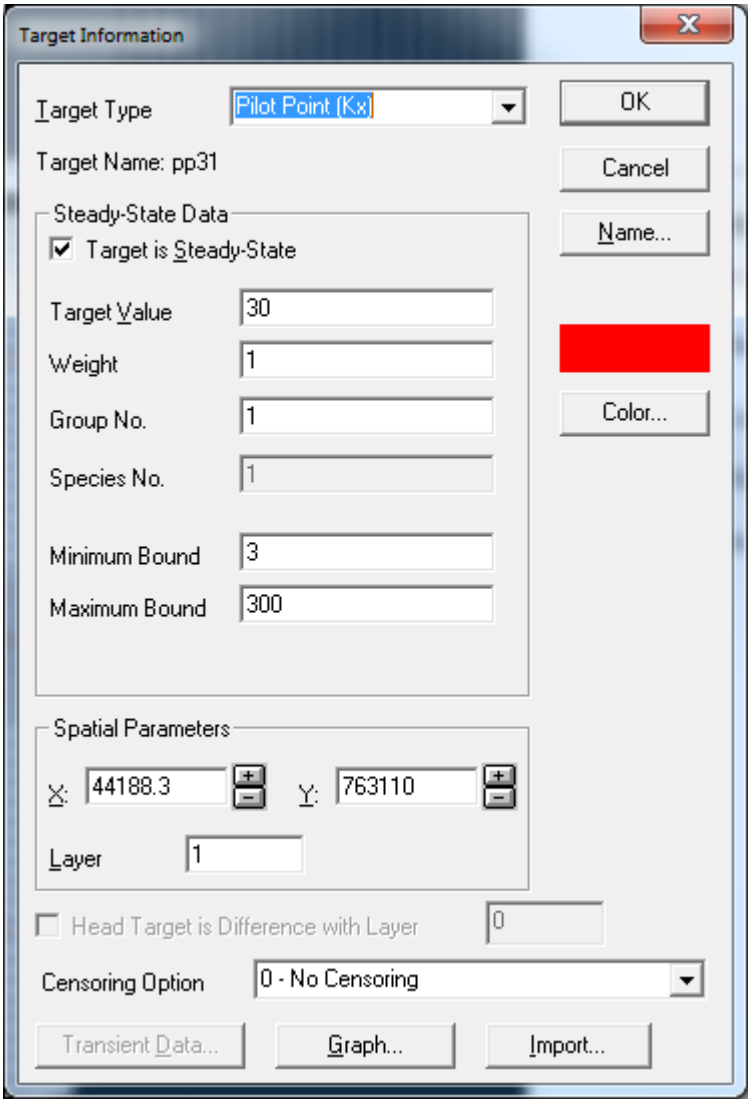

Now we will import another set of targets that are our water level calibration targets. These data are provided in a file called *c:\gwv6\tutorial\Pesttargets.dat*.

Select **AE|Import|Target Text File** and located the file called *c:\gwv6\tutorial\Pesttargets.dat.*

These targets are transient, representing the end of the first stress period (elapsed simulation time = 1500 days). Check the option that says *File Contains Transient Targets*

Also check the option that says Read One Target Value for Transient Targets.

Enter a value of 1500 for the target time value.

Enter a value of 0 next to Layer.

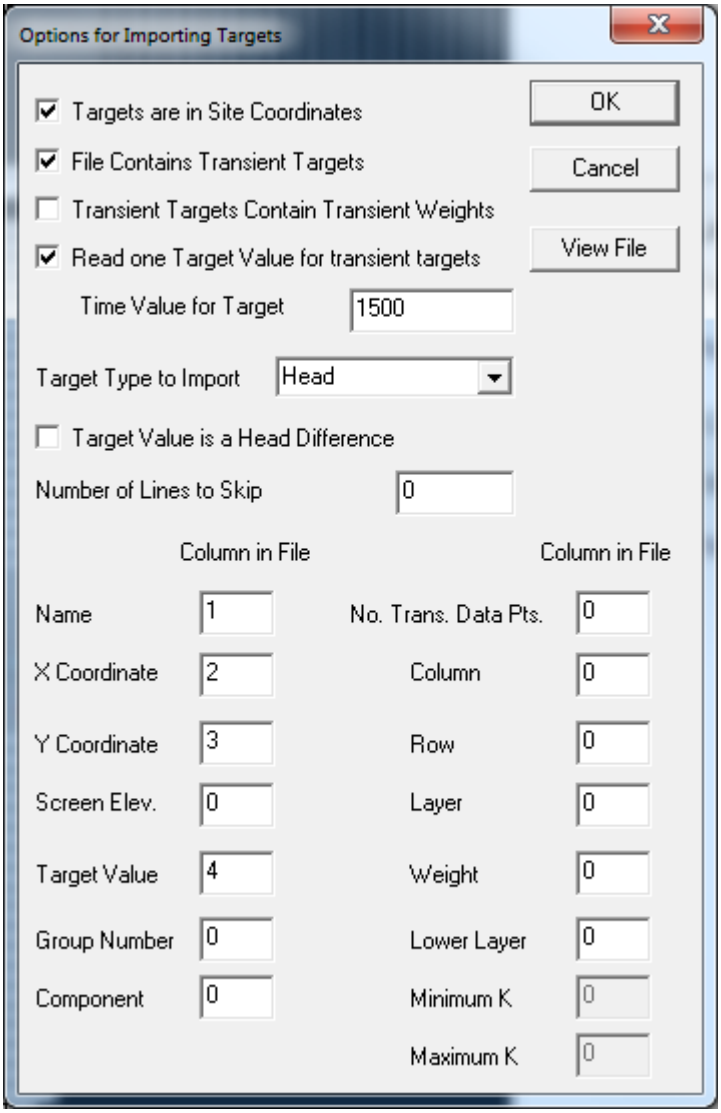

GV should report that 21 targets were imported.

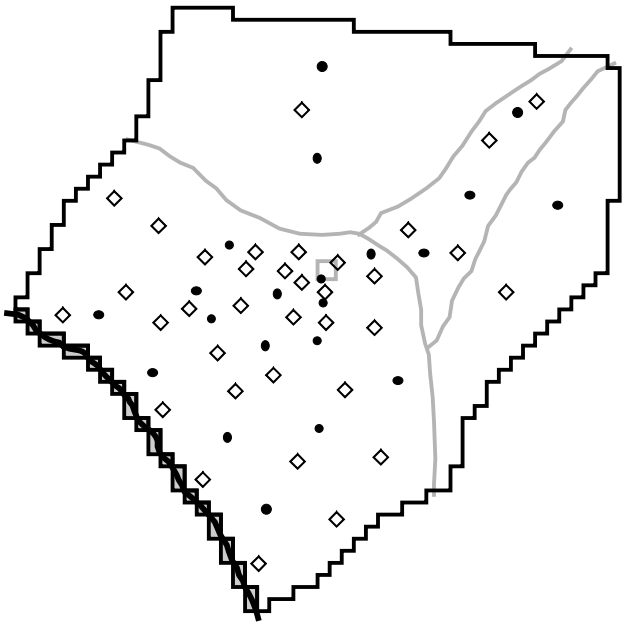

Figure 8. Locations of pilot points (open diamonds); observation bore locations are shown as closed circles.

The placement of pilot points at the locations shown in Figure 8 is somewhat subjective. Nevertheless a strategy was followed in choosing their locations:-

One point was assigned to each of zones 3 and 4 to make them effectively uniform, as described above.

Where more than one point was assigned to a zone, the points were distributed pervasively through the zone, including right up to the zone boundaries (especially the downstream boundary, as this gives PEST the chance to implant continuous geological structures in the subsurface which may have a pronounced effect on the movement of contaminants within it).

Pilot point density was increased where there is greater density of observation points.

Conventional wisdom dictates that the number of parameters involved in a parameter estimation exercise should be kept to a minium. However when using pilot points in conjunction with PEST's new "regularisation" mode, the opposite is often true. PEST's regularisation functionality prevents the onset of numerical instability, which often accompanies attempts to solve over-parameterised inverse problems. Furthermore, a superfluity of pilot points allows PEST to achieve a much better fit between model outcomes and field observations. PEST's regularisation functionality will be discussed shortly. For the moment, we will use pilot points without regularisation in parameterising our model.

Spatial Interpolation using Pilot Points

The use of pilot points in characterising the spatial distribution of a hydraulic property must be accompanied by a mechanism whereby hydraulic property values assigned to pilot points are spatially interpolated to the cells of the finite difference grid. If using programs provided with the Groundwater Data Utilities, spatial interpolation is accomplished using the kriging algorithm. Kriging is a method of spatial interpolation based on geostatistics. The cornerstone of geostatistics is the variogram; a variogram

describes the extent to which hydraulic property values (or any other type of data) pertaining to any two points are likely to be different from each other as a function of the distance between those points. See Section 5 of Part A of the Groundwater Data Utilities manual, or any standard text on geostatistics, for more details.

One of the benefits of using kriging as a basis for spatial interpolation is that the factors by which hydraulic properties at pilot points are multiplied before summation to obtain the hydraulic property value at a particular grid cell are independent of the actual hydraulic property values at the pilot points. Hence a set of "kriging factors" pertaining to each of the cells of the finite difference grid can be calculated in advance of the actual interpolation process. As the latter is undertaken again and again as the model is run repeatedly by PEST, the fact that it is not necessary to repeat calculation of the kriging factors on each occasion that the model is run can result in large savings in the time required to complete the overall parameter estimation process.

Calculation of kriging factors is undertaken by program PPK2FAC, which is run automatically from GV when you use pilot points. GV will create the necessary files for PPK2FAC and run the program. In some cases, though, it may be necessary to run the program from the command line. We will show you how to do that below.

Variograms upon which these kriging factors are based are supplied to PPK2FAC in a "structure file". For full details of the specifications of this file, see the documentation to the Groundwater Data Utilities. GV can create the structure file by first selecting **Model|PEST|Options – Structures Tab**. There are **next** and **previous** buttons at the bottom of the dialog to allow you to scroll through each structure definition.

Three geostatistical structures are defined for the present example. Each of these structures cites one variogram (though it could cite up to five). "Structure1" will be used to characterise zone 1 of our model domain (ie. the river alluvium), "structure2" will be used to characterise zone 2 (ie. the creek alluvium) whereas "structure3" will be used to characterise zones 3 and 4. Note that the variogram assigned to these latter zones is quite unimportant; because there is only one pilot point assigned to each of them, all cells within these zones will be assigned the one interpolated value (same as the respective pilot point) irrespective of the variogram.

The variogram assigned to "structure1" represents the river alluvium. This is an exponential variogram (VARTYPE=2) with a range of about 1500m (three times the value of the "A" variable appearing in the variogram specification). In contrast to variogram2, which is assigned to "structure2" (representing the creek alluvium), there is no anisotropy associated with "vario1". However for "vario2" the anisotropy value is 2.0, with the direction of anisotropy coinciding with the direction of the creek. Alignment in the direction of the creek is based on the premise that channel structures within this old creek valley will make it more likely for hydraulic property similarity to prevail in this direction than in a direction at right angles to it.

These structures have been defined for you in the *ex1a.gwv* model. Select **Model|PEST|Options - Structures Tab** and look at each of the three structures. The first is shown below.

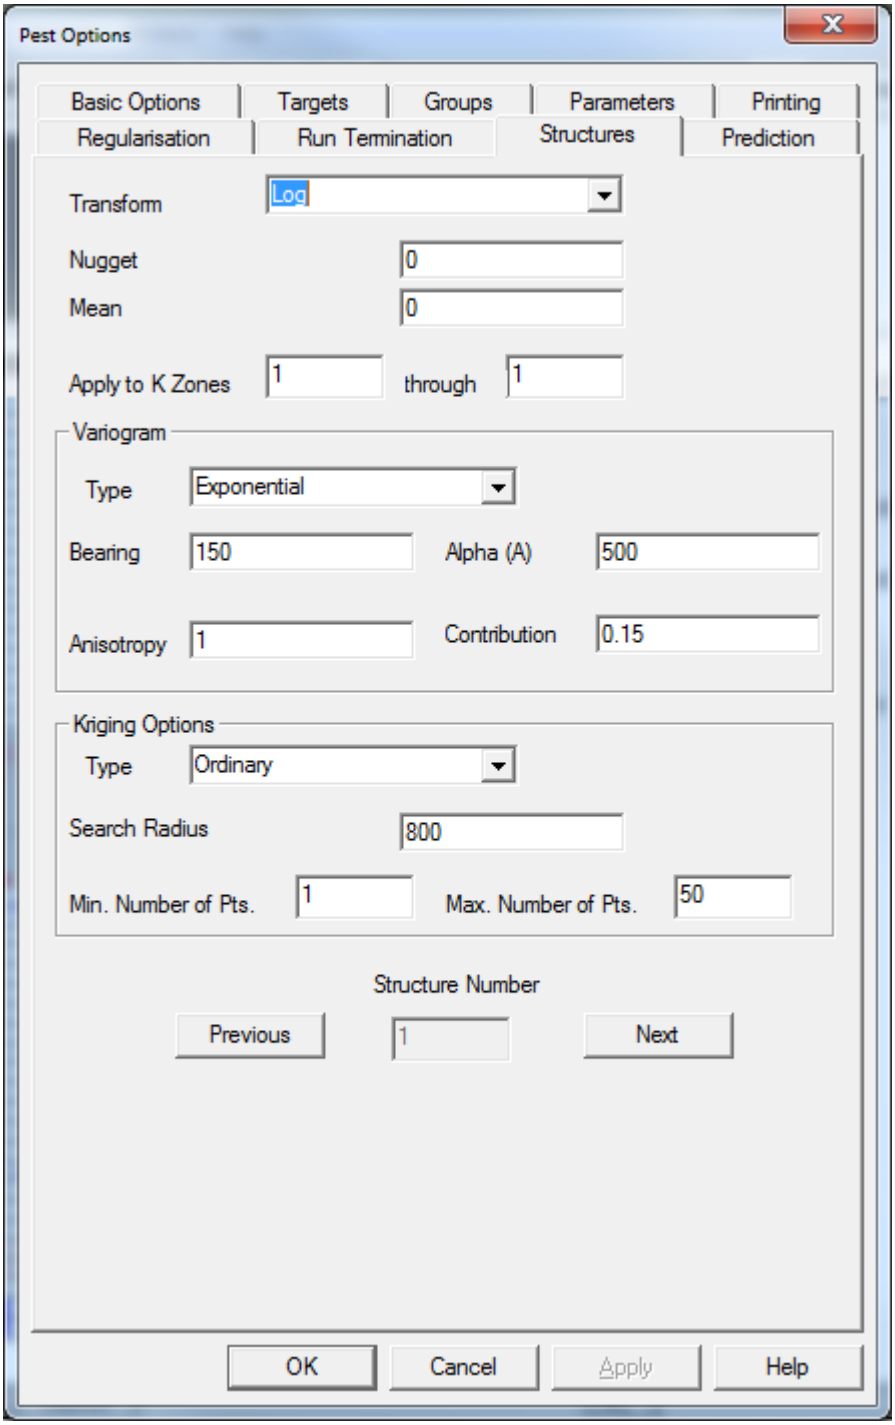

Note: When filling out this dialog for your own model, the key items are the following:

- Use Exponential variogram
- The anisotropy should be 1
- The Alpha parameter should be about 20% of the maximum length or width of your model domain
- Search radius should be high enough that you will find several pilot points from any point within the area being estimated.

- Set contribution to 1.
- All other fields can be left at default levels.

Note also that for all of the structures the TRANSFORM variable is set to “log”. Thus any variogram cited in each of these structures must pertain to the spatial distribution of the logarithm of the pertinent hydraulic property. This is in accord with the fact that most studies cited in the groundwater literature which treat transmissivity and|or hydraulic conductivity as a regionalised variable, indicate that its distribution is better described by a log variogram, rather than a variogram based on native property values.

To implement pilot points in GV, you just need to define the pilot point locations as we did above.

Select **Model|PEST|Options** - Basic Options tab

Check the option to include pilot points as shown below.

You also need to check the option to write arrays as external files. This is necessary to allow the use of some of PEST’s utility software to generate the transmissivity arrays.

Also check the option to run the models without screen output. This allows you to see the Pest output on the screen.

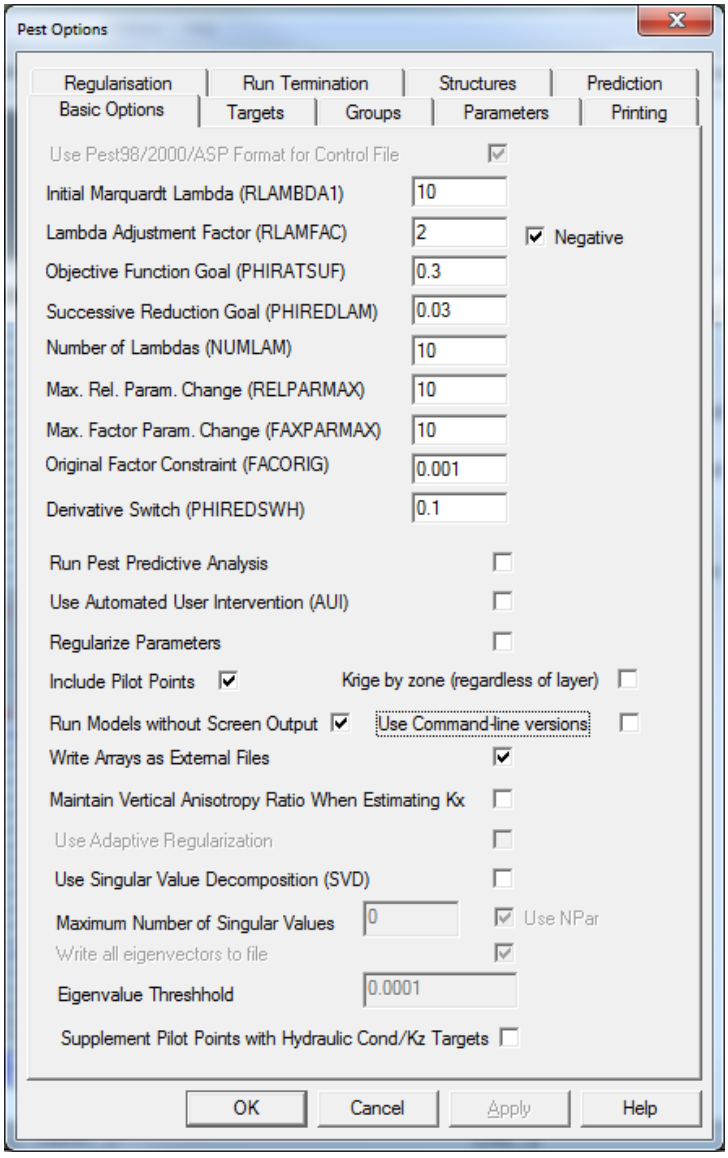

Confirm that these settings have been turned on and then select **Model|PEST|Create Data Sets**. You should see some DOS windows come up as GV runs the PEST utility software.

Program PPK2FAC will now be run in order to generate a set of kriging factors by which the transmissivity at each cell centre can be calculated from the transmissivities assigned to the pilot points.

After PKKFAC has run to completion three new files will exist in your working directory. *Factors1.dat* contains kriging factors for all active cells in the model domain. *Sd1.ref* contains a MODFLOW-compatible array in which the square root of the "kriging variance" is recorded for every active cell. This is a measure of the uncertainty associated with the spatial interpolation process, calculated from the geostatistical structure assigned to each zone. Finally, file *reg1.dat* contains information that will allow program PPKREG to add a series of "regularisation prior information equations" to the parameter estimation process later; see below.

Figure 9 shows that the "kriging uncertainty" increases with distance from the nearest pilot point; if a pilot point falls exactly at the location of a cell centre, then the kriging uncertainty at that cell centre is zero. Also, the uncertainties in zone 2 are elongated along the direction of the creek in accordance with the anisotropic nature of the geostatistical structure assigned to that zone. The kriging uncertainty pattern for zones 3 and 4 is circular due to the fact that there is only one pilot point assigned to each of these zones.

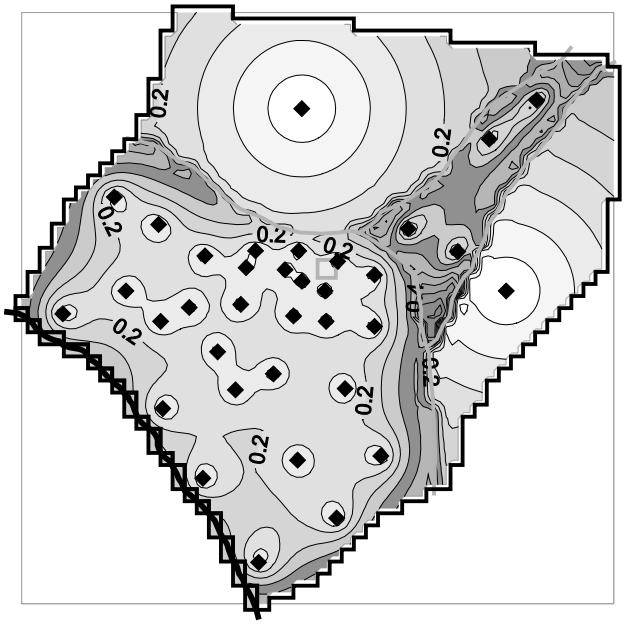

Fig. 9 "Kriging uncertainty" calculated by PPK2FAC on the basis of the geostatistical structures contained in file struct.dat. Darker colours indicate higher uncertainty.

Interpolation of hydraulic property values assigned to pilot points to the centres of cells comprising the finite difference grid is undertaken using program FAC2REAL. In a pilot point file, the value assigned to each point comprises the fifth element of each line. GV creates this file and calls it *points#.dat* where # is the layer number. In this example, there is only one layer so the file is called *points1.dat*.

FAC2REAL writes a MODFLOW-compatible array of real numbers to a file named *root#._kx* where # is the layer number and root is your root file name for the MODFLOW run (ex1a in this example). The elements of this array are comprised of numbers interpolated from the sites of pilot points to the centres of finite-difference cells comprising zones 1, 2, 3 and 4 of the model domain. Spatial interpolation is carried out by applying the kriging factors calculated by PPK2FAC to the pilot point values supplied in the fifth

column of the pilot points file. It is important to note the following aspects of the interpolation process carried out by FAC2REAL.

FAC2REAL can apply upper and lower limits to interpolated values at cell centres. This can be very useful, for interpolated values can sometimes be higher or lower (even negative) than any of the values assigned to the pilot points; this is more likely to happen when pilot points are close together and when the Gaussian variogram is used. In the present instance the lower interpolation limit is set uniformly at 10^{-6} and the upper interpolation limit is set uniformly at 10^{6}. However you can supply these limits on a cell-by-cell basis (this is not supported by GV).

If any of the geostatistical structures cited in the structure file supplied to program PPK2FAC have a TRANSFORM value of "log", this information is recorded in the kriging factor file written by PPK2FAC. FAC2REAL, when calculating a spatially interpolated value on the basis of pilot points to which such a geostatistical structure applies, takes the log of each pilot point value before applying the respective kriging factor. After the log of the interpolated hydraulic property value is calculated by summation of kriging factors times respective log-transformed pilot point values, the resultant value is raised to the power of 10 to "untransform" it before it is recorded in FAC2REAL's output file.

If there are any cells to which interpolation does not take place (either because they lie within a zone for which no kriging factors were calculated, or because they are located at a greater distance from any pilot point than the search radius supplied to PPK2FAC), then FAC2REAL records a user-supplied dummy value for these cells. In the present case this dummy value is 10^{35}. Since MODFLOW would probably object to a K or T value of 10^{35}, Groundwater Vistas runs a program called *kmakr.exe* after FAC2REAL runs. The *kmakr* program substitutes K or T values into the matrix when it finds a 10^{35} number. The K or T value is the default value in your model for the zone containing that cell.

GV creates the necessary information for FAC2REAL and creates a new model batch file called *pestgv.bat* which will run the FAC2REAL program before running MODFLOW. As stated in the last bullet above, GV also runs a program called *kmakr* right after the FAC2REAL program.

Preparing for a PEST Run

Initially our model will be calibrated on the basis of water levels only. PEST input file preparation will be carried out with assistance Groundwater Vistas. Files prepared in this way will then be amended in order to carry out more advanced analysis of the measurement dataset using PEST in the following sections.

When using pilot points to parameterise a groundwater model, the parameters to be estimated are the values (in this case transmissivity values) assigned to the pilot points. Pilot points are assigned values in a pilot points file; so our first task in preparing for a PEST run is to build a template file based on a pilot points file. GV does this automatically, however, when you created Pest input files above.

MODFLOW produces a plethora of output data in many different files. However from the point of view of calibrating the current model, the only output quantities that are important to us are MODFLOW-generated heads spatially interpolated to the sites of the observation bores depicted if Figure 6. The task of spatially interpolating MODFLOW outputs to discrete points within the model domain is carried out by a program called *targpest.exe* supplied with Groundwater Vistas. Groundwater Vistas creates a file called *targpest.dat* that supplies the necessary target information to targpest.exe. The advantage of using targpest is that it reads the MODFLOW head-save binary file directly which means that you get maximum precision in the computed head values. If you inspect the file *pestgv.bat* in a text editor, you will see that the targpest command immediately follows the MODFLOW run.

Just so to make sure that we have all of the files we need, select **Model|MODFLOW|Create Datasets**. That makes the MODFLOW input files for a single run.

Now select **Model|PEST|Create Datasets** which writes all of the PEST input files and runs some of the PEST utility programs.

PEST requires three types of input file, viz. one or more template files, one or more instruction files and a PEST control file. Groundwater Vistas builds all of these files for you so there is little effort in running PEST from GV. The GV approach, though, assumes that you will be using the *targpest* program to

interpret MODFLOW output. If you want to use the Groundwater Data Utilities that come with PEST, then you would need to do a lot of editing of the PEST files in a text editor.

Even though GV takes care of your PEST files, there can sometimes be problems that will cause the input files to be incorrectly formatted. For this reason, it is usually a good idea to run the PEST checking utility called *pestchek*. To do this, select **Model|Pest|Run Pestchek**.

PESTCHEK should report no errors or inconsistencies.

Running PEST

You should now run PEST directly from GV by selecting **Model|PEST|Run PEST**. If nothing happens then the PEST software is not in the current path for your computer.

After a few iterations you will notice that PEST has lowered the objective function to an incredibly low value. You can either stop PEST then, or let it run to completion. To stop PEST, select **Model|PEST|Stop Pest**. It is apparent from an inspection of the run record file *hcal.rec* that PEST is able to match every model output to every field observation almost exactly. Thus, in spite of the fact that there are often serious numerical problems encountered in cases where observations outnumber parameters, PEST was able to achieve a perfect solution to the present inverse problem.

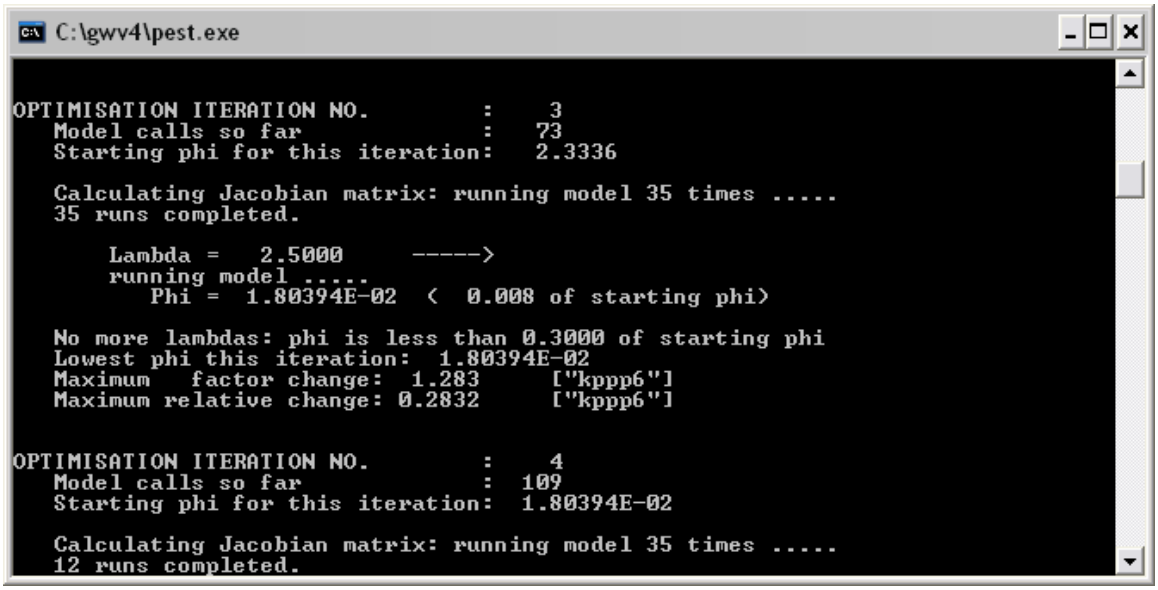

After a PEST run is complete, optimised parameter values are listed in the run record file (in this case *ex1a.rec*). They are also recorded in a "parameter value file", in this case *ex1a.par*. You can look at these files from GV by selecting **Model|PEST|View Main Output File** or **Model|PEST|View Parameter Estimates**. GV launches a text editor (notepad by default) to show these output files.

As usual, MODFLOW-generated heads are to be found in *ex1a.hds*. You can select **Plot|Import Results** and click OK to view the calibrated heads. Select **Plot|Calibration|Stats-Plots** to see the calibration statistics. If you click the button to plot observed versus simulated heads, your plot should look similar to the following:

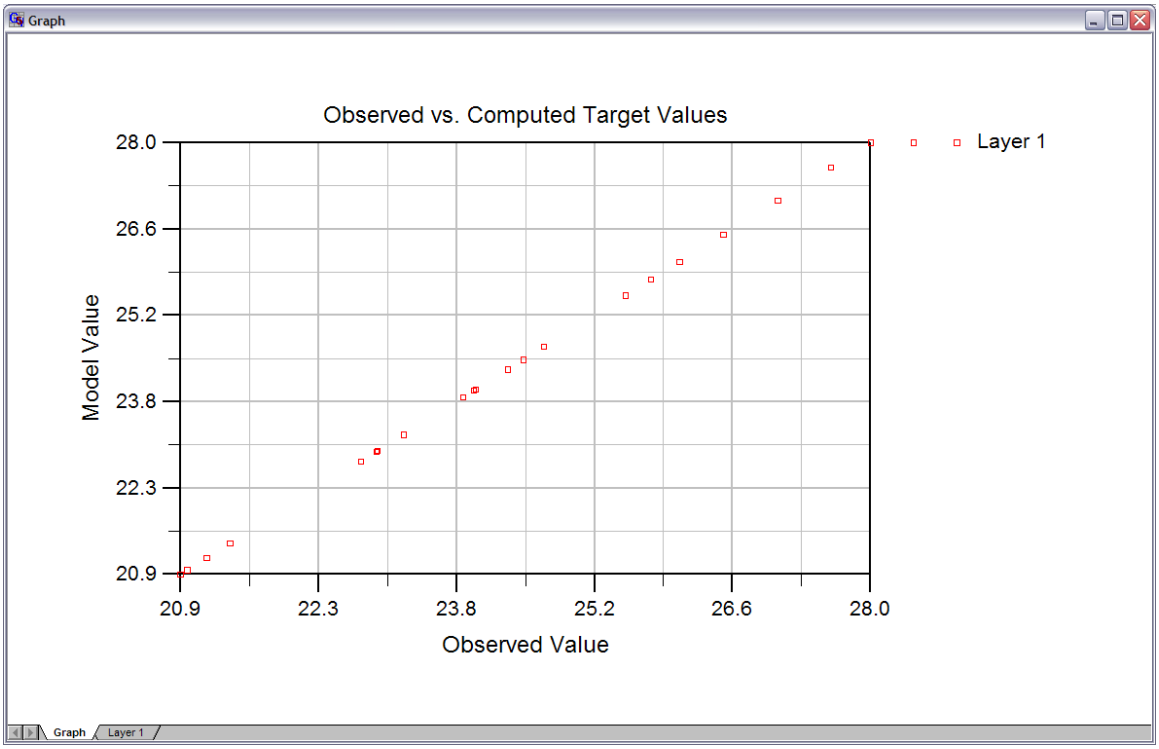

This is remarkably good!

Also, the transmissivity distribution used by MODFLOW (spatially interpolated from the locations of the pilot points) is found in file *ex1a1._kx*. If you imported this array into GV for graphical display you would see that the optimised transmissivity distribution is as shown in Figure 10.

You can import the estimated transmissivities by selecting **Props|Leakance** (we don't need this array and can use it to view the estimated T's – If you wanted to actually incorporate these values into your model, you would use the Hydraulic Conductivity property). At this point many of you will be confused as to why we are importing transmissivity into the leakance property. Essentially we are just using the leakance property to view the transmissivity values. Leakance is not used in this model so that importing T values into Leakance does not change the MODFLOW run in any way. Thus it is just a convenient way of inspecting the calibrated T values without accepting them at this point. To import these T values, select **Props|Import|Matrix**. Click the browse button and find the file c:\gwv6\tutorial\work\ex1a1._kx. Next, select Props|Property Values|Reset Matrix Bounds to reset the colors used to color flood the data.

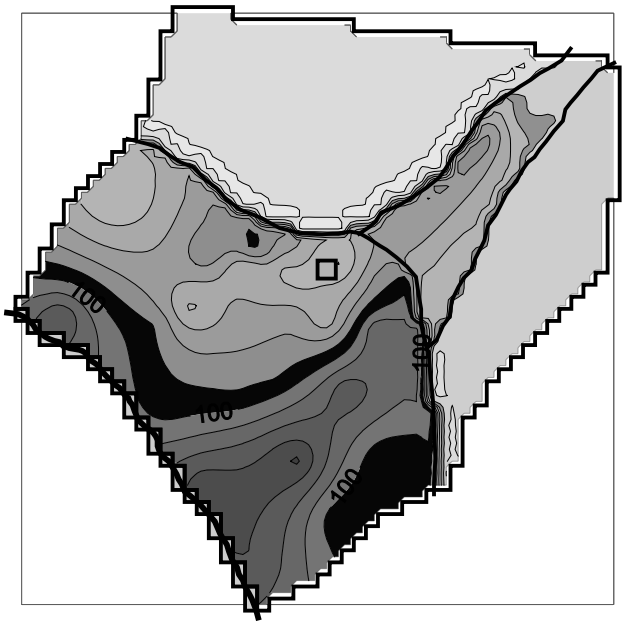

Figure 10. Transmissivity distribution within study area determined through calibration against heads. Contour interval is 10m^2|day; darker shade indicates greater transmissivity.

A comparison of Figure 10 with Figure 2 (which shows the "true" transmissivity distribution) reveals that although the "calibrated" transmissivity field has certain things in common with the true field (for example, the tendency of transmissivities to increase toward the eastern end of the river alluvium), there are nevertheless significant differences between the two transmissivity fields. Remember, however, that both transmissivity fields result in a perfect fit between model-generated heads at boreholes and their respective field-measured counterparts.

The difference between the "true" and calibrated transmissivity fields is depicted in Figure 11. Differences as high as ±40m^2|day can be seen within the river alluvium. It is thus apparent that the transmissivity distribution within the river alluvium cannot be uniquely determined through calibrating the model against water levels measured in available bores. In fact, the transmissivity distribution shown in Figure 10 is just one of an infinite number of transmissivity fields that would also have resulted in a perfect, or near-perfect, fit between model outcomes and field measurements. The big question is, given the uncertainty in the transmissivity field determined through the calibration process, what is the uncertainty in model predictions of contaminant movement, and hence in its predictions regarding the efficacy of the remediation system?

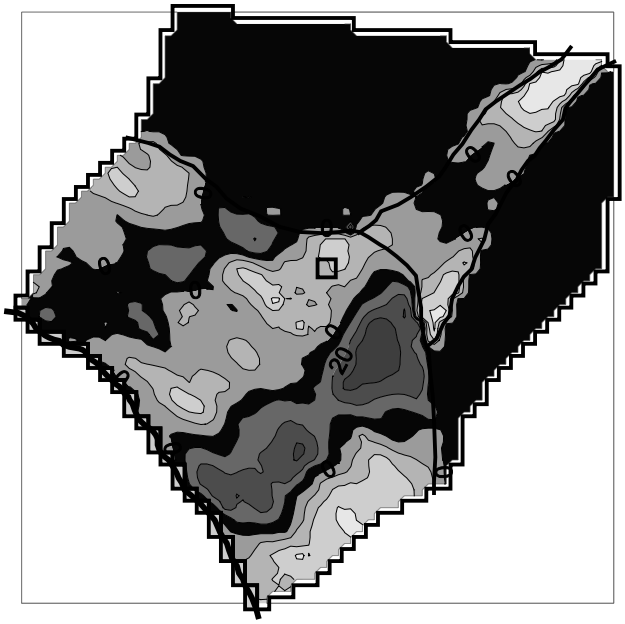

Figure 11. Difference between calibrated and true transmissivities. Contour interval is 10m^2\|day.

If MT3DMS is now run on the basis of the transmissivity field shown in Figure 10 in order to evaluate the efficacy of the remediation system, it is found that the amount of contaminant entering the river over the 5000 day simulation time is 2167 units.

Regularisation with Pilot Points

It has now become abundantly obvious that it is beyond our capacity to uniquely determine the transmissivity distribution of our model domain by calibrating the model on the basis of water level measurements taken in a number of bores distributed through that domain. Unfortunately, when simulating transport processes, it is often the “geological fine detail” that will determine the outcome of these processes. However the introduction of fine detail means the introduction of many parameters into the parameter estimation problem. This brings with it nonuniqueness, and, more often than not, numerical instability. We were lucky to get away with what we did in our previous calibration exercise in which we estimated 14 more parameters than there were observations (thus breaking one of the "golden rules" of parameter estimation). However, as mentioned above, one of the great advantages of using pilot points is that we can distribute a superfluity of those points throughout the model domain and then ask PEST to find for itself those regions within the study area where transmissivity must be greater or less than average in order to ensure that there is good agreement between model outputs and field measurements. If we had based our parameterisation of the model domain solely on zones, we might not have placed those zones in the correct position for the calibration process to properly infer the existence or extent of such heterogeneity.

The introduction of regularisation into the calibration process serves two purposes. Firstly it brings a high degree of numerical stability to a parameter estimation problem which would otherwise be highly susceptible to the deleterious effects of a singular normal matrix (you might have noticed when inspecting *ex1a.rec* that PEST was not able to calculate any parameter statistics due to singularity of the normal matrix.) Secondly, if regularisation constraints are appropriately defined, model calibration can proceed with a “homogeneous unless proven otherwise” philosophy; that is, in spite of the number of parameters at its disposal, PEST will make each zone within the model domain as uniform as it can in terms of the distribution of the estimated hydraulic property, introducing heterogeneity into a zone only where this is necessary in order to allow a good of fit between model outputs and field data to be achieved. Hence any heterogeneity which is introduced as an outcome of the calibration process is “there because it has to be

there". In many modelling contexts this philosophy of model calibration has a large intuitive appeal, allowing a modeller to use zones to characterise the distribution of some hydraulic property within a model domain while, at the same time, removing the inflexibility that accompanies the characterisation of a model domain by areas of piecewise parameter constancy.

In the present case we will use the same geostatistical information that was used by program PPK2FAC in calculating kriging factors, to introduce regularisation constraints pertaining to the relationships between different pairs of parameters into the inverse problem solved by PEST. If the "default condition" for each of our hydrogeological zones is one of hydraulic property uniformity then, for any two pilot points lying within the same zone, the preferred value for the difference between transmissivity values is zero. Relationships between pairs of parameter values can be introduced into the calibration process as prior information equations. The weight assigned to each of these prior information equations can be the same. Alternatively, if the weight is proportional to the inverse of the square root of the variogram calculated for the distance between the respective pilot points, then it can be shown that this is in harmony with the geostatistical characterisation of the area as encapsulated in the variogram. What this characterisation says, in short, is that "the closer are two points together, the more likely are the hydraulic properties at those points to be the same". By calculating weights on the basis of the inverse of the variogram, we are enforcing the "zero difference" condition more strongly for points which are closer together than for those which are farther apart.

The PPKREG Utility introduces regularisation prior information to a PEST control file. It does this on the basis of information written to a "regularisation information file" by program PPF2FAC; PPK2FAC writes this file (it was called *reg1.dat* in our case) when calculating kriging factors. It contains a matrix of variogram values calculated between pairs of pilot points. PPKREG uses this information to calculate the weights assigned to the regularisation prior information equations. However PPKREG does not formulate prior information equations which link parameter values pertaining to pilot points assigned to different zones; thus the role of zone boundaries in identifying preferred locations of hydraulic property discontinuity remains intact. Also, PPKREG does not link pilot points which are "too far apart" in terms of the search radius supplied to PPK2FAC when it calculated kriging factors.

Groundwater Vistas can now run PPKREG directly. To do this, select **Model|Pest|Options - Basic Options tab** and put a check next to *Regularize Parameters*

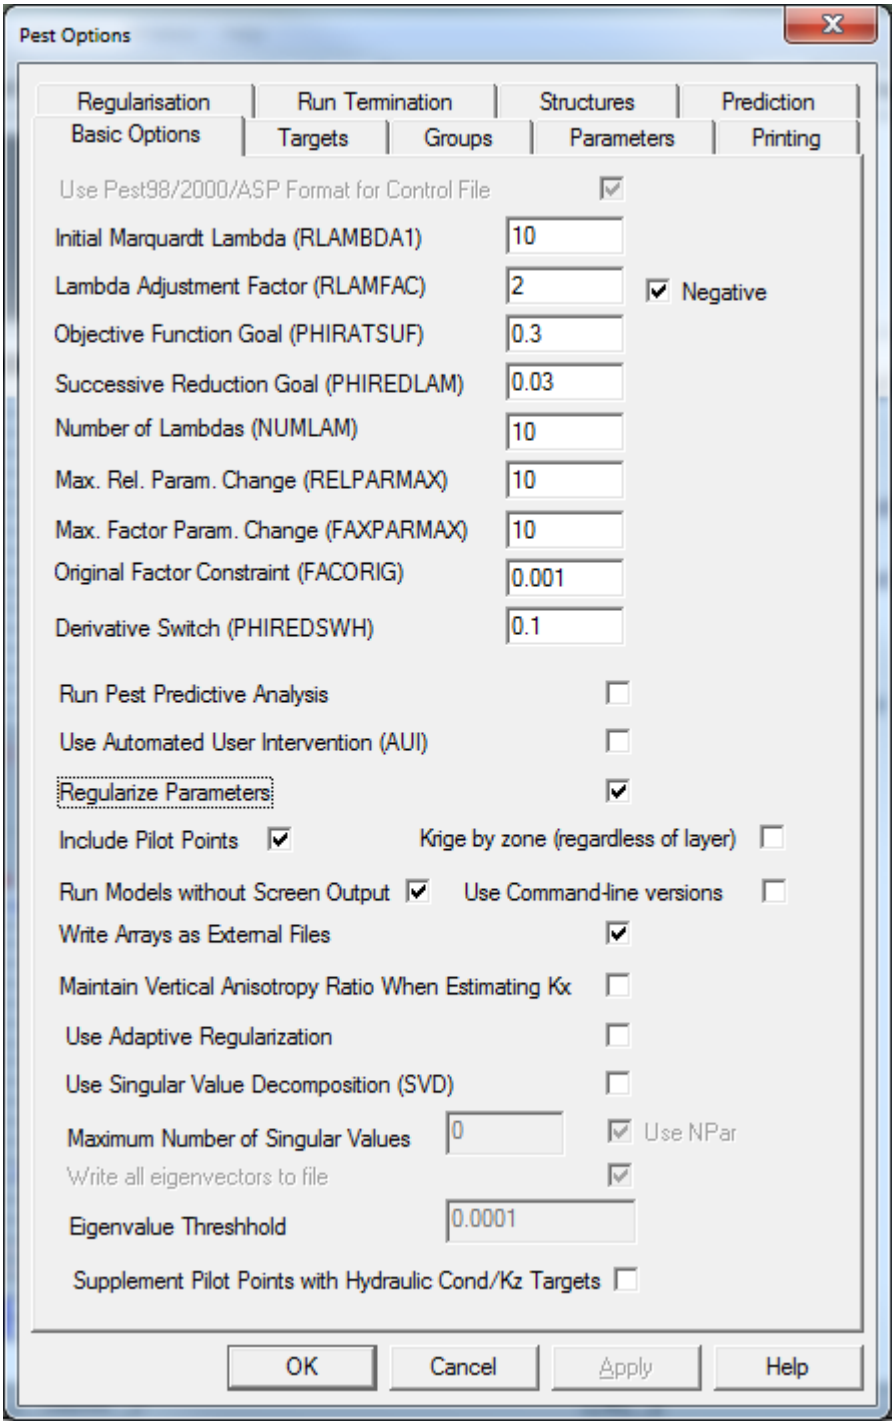

Now, you can select from the regularization options by clicking on the **Regularisation** tab. Fill out the dialog as below:

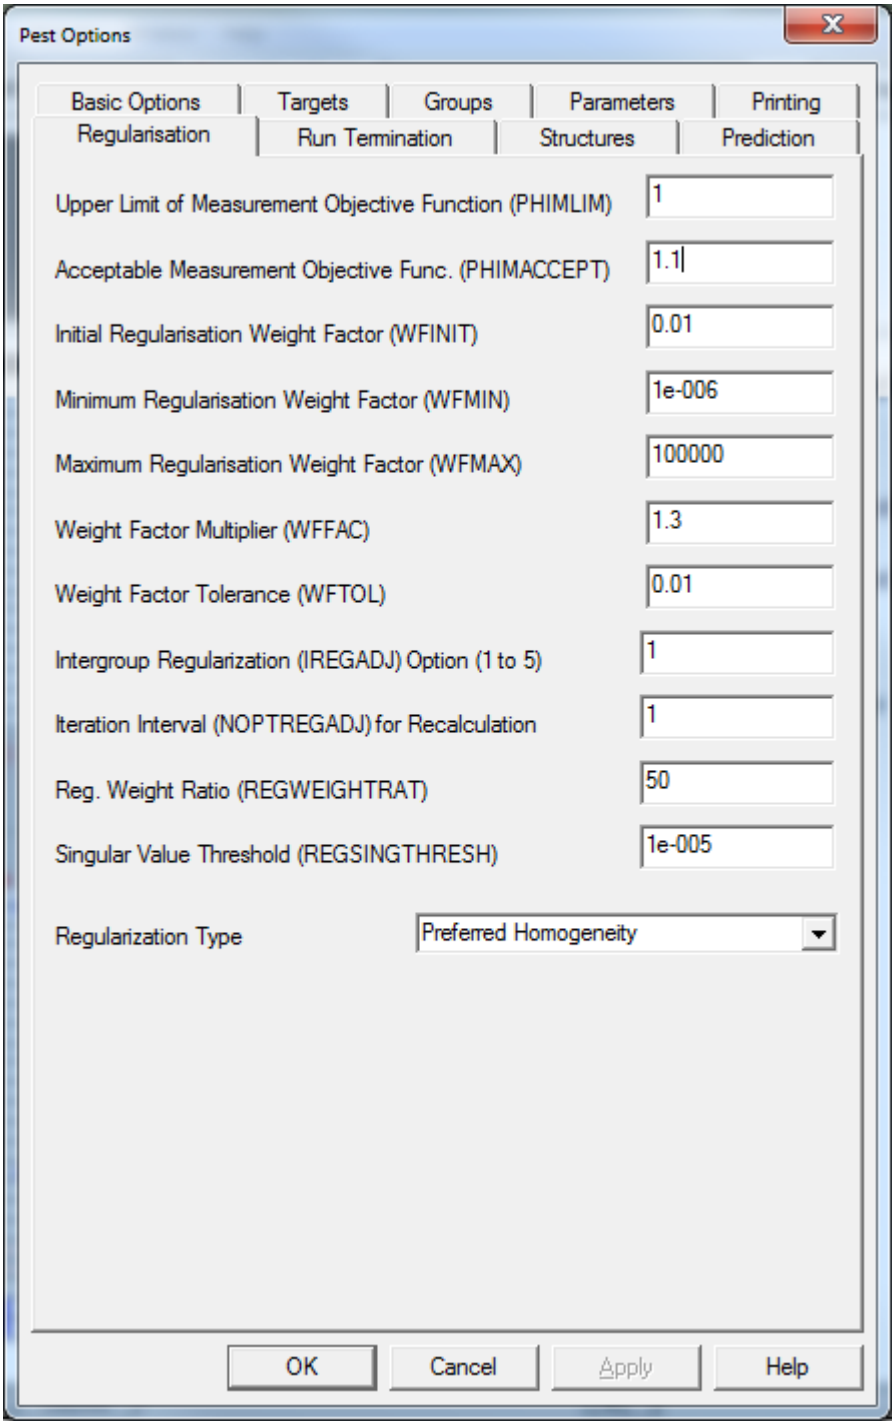

You should change PHIMACCEPT to be 1.1 (slightly higher than PHIMLIM) and WFMAX to 100,000.

Now create Pest datasets and run Pest as before.

When writing the new PEST control file, PPKREG sets the PESTMODE variable to “regularisation”, ie. it informs PEST that it must run in “regularisation” mode. When run in this mode, a number of control variables are required in the PEST control file, in addition to those required when PEST is run in

"parameter estimation" mode. One of these variables is PHIMLIM. This specifies the degree of model-to-measurement misfit that is allowed to occur in the present optimisation process. Because the attainment of a good model-to-measurement fit, and the simultaneous enforcement of homogeneity constraints, may place conflicting requirements on parameter values, a compromise between the two must be reached. The user determines the "compromise level" by setting a maximum model-to-measurement misfit that he|she will tolerate, this misfit being expressed in terms of the "measurement objective function". The maximum permissible value of the measurement objective function (ie. PHIMLIM) should be set a little higher than the objective function that it is possible to achieve without any regularisation constraints being enforced. From our previous run we know that we can obtain an objective function of almost zero without regularisation. In the present case we set PHIMLIM at 0.1, thus informing PEST that we will tolerate an rms model-to-measurement misfit of 7 cm. (Divide 0.1 by the number of observations, then take the square root to obtain this figure.)

Each prior information equation included in the parameter estimation process must be assigned a weight. As was discussed above, weights are calculated on the basis of geostatistical information available (or assumed) for the model area. If an observation or prior information equation is used for regularisation purposes, then it is assigned to the observation group "regul". As part of its regularisation functionality, PEST adjusts the weights assigned to all members of this group during each iteration of the optimisation process; however the relative weight values within this group remain the same. The "regularisation weight factor" by which the initial weights of all members of the group "regul" are multiplied during each optimisation iteration, is calculated in such a way as to respect the PHIMLIM value provided by the user as the maximum tolerable model-to-measurement misfit for the current case. An initial regularisation weight factor needs to be supplied by the user. In the present case this factor is supplied as 1.0. Similarly PEST needs to be informed of the upper and lower limits that it is allowed to use for the regularisation weight factor; default values offered by PPKREG are 10^{6} and 10^{-6} times the initial regularisation weight factor.

Inspect file *ex1ar.rec*. The indicator of model-to-measurement misfit when PEST runs in "regularisation" mode is the "measurement objective function". You will notice that PEST does not lower the measurement objective function to as low a value as it did during the previous PEST run when PEST was run in "parameter estimation" mode (where it was referred to simply as the "objective function" or "phi"). In fact it lowers it only to the value that you supplied for PHIMLIM. This gives PEST "room to move" in satisfying the other objective of the optimisation process, viz. the imposition of a homogeneity condition over the model domain.

Notice also that PEST runs for 15 optimisation iterations before it stops; however it had obtained a measurement objective function that was almost equal to PHIMLIM in only 7 iterations. Iterations after that were used for "fine tuning"; in fact if the default termination criteria for the regularisation process provided by PPKREG when it wrote the PEST control file *ex1ar.pst* weren't quite so tight, PEST would have ceased execution well before 15 optimisation iterations had elapsed.

At this stage you would normally import the MODFLOW-compatible real array *ex1a1._kx* into Groundwater Vistas for graphical display. If you did, you would find that the optimised parameter transmissivity distribution was as depicted in Figure 12.

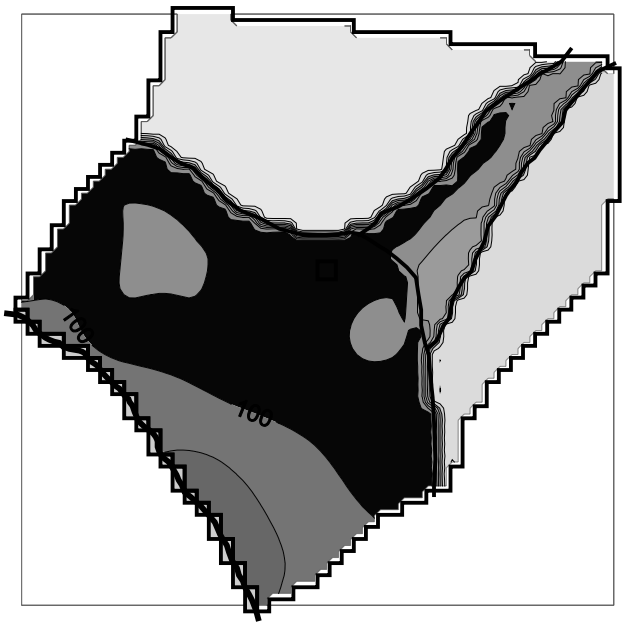

Figure 12. Transmissivity distribution within study area determined through calibration against heads with regularisation applied. Contour interval is 10m^2|day; darker shade indicates greater transmissivity.

A comparison between Figure 12 and Figure 2 indicates that the calibrated transmissivity distribution depicted in Figure 12 does indeed replicate the general rise in transmissivity to the south-east that is present in the "true" transmissivity distribution. However the transmissivity field of Figure 12 is a lot smoother than that of Figure 2. (It is of interest to note that if PHIMLIM had been set to 0.3 - which still indicates a very good fit between model outputs and field data, but gives PEST a little more "room to move" in terms of its capacity to apply regularisation constraints - PEST would have calculated a uniform transmissivity field in each of the model zones.)

When running MT3D on this version of the model, the cumulative mass budget shows 12,038 units of contaminant find their way to the river over the model simulation time.

You can probably see where pilot points can save a lot of time in coming up with a calibrated K field. Even if you do not agree with the philosophy behind the technique, you can use it to determine how K should vary to produce a good calibration. You can then take the continuous K field computed by PEST and use it to guide placement of zones in your model.

3D Pilot Point Example with SVD

One of the most powerful new features of Pest is called Singular Value Decomposition or SVD-Assist. SVD-Assist significantly reduces the number of runs you need to make when using pilot points. SVD-Assist also can better handle different parameter types that have very different sensitivities. A good example is horizontal and vertical K values. Prior to SVD, Pest often had a difficult time calibrating both Kx and Kz simultaneously. SVD seems to have solved that problem. The only down-side to SVD is that it is a more complex procedure to follow.

You use SVD-Assist with Pest by starting with a normal pilot point setup as in the first example with the one-layer model. Assign the calibration targets and pilot point locations. The SVD-Assist procedure then uses the following steps:

Create Sensitivity Run. This is a normal Pest run but with only 1 iteration. Pest uses this run to determine the sensitivities of each parameter (pilot point). To create this run, set the number of iterations to 1, initial lambda value to zero, and number of lambda values to 1.

Launch Sensitivity Run. After creating the Pest input files for the sensitivity run, launch Pest. It will run one model run for each parameter.

Launch SVDAPREP. SVDAPREP is one of the many utility programs that comes with Pest. The most important part of using SVDAPREP is to determine the number of "super" parameters for the next stage. Super parameters are combinations of your complete set of parameters. Usually the number of super parameters will be 10 to 20 percent of the total number. I have found that 15% generally works best. This means that if you have 400 pilot points, the number of super parameters should be 60. The practical implication of super parameters is that instead of making 400 runs per iteration, you only make 60 runs per iteration.

Launch SVD Pest Run. SVDAPREP creates a new Pest file based on your original Pest input files and the results of the sensitivity run. Run Pest using this new set of input files.

Create Final Pest Run. The results of the SVD Pest run are not useful in and of themselves. You need to run a program called PARREP to make sense of the SVD results. PARREP creates a new Pest file, which we will call the "final" Pest file. After running PARREP, change the number of iterations to zero and run Pest using this file. Pest will run the model only 1 time using the final parameters values.

Now, considering that most of the tasks in 1 through 5 above need to be done at the command line, this makes SVD a bit difficult to implement in practice (unless you are over 40 and still remember DOS). Recognizing this problem, we have created a special menu in Vistas that duplicates the steps above. All you have to do is execute the menu commands from top to bottom without skipping a step and it is relatively easy to use SVD. This new menu is under **Model|Pest|SVD Assist|…** as shown below:

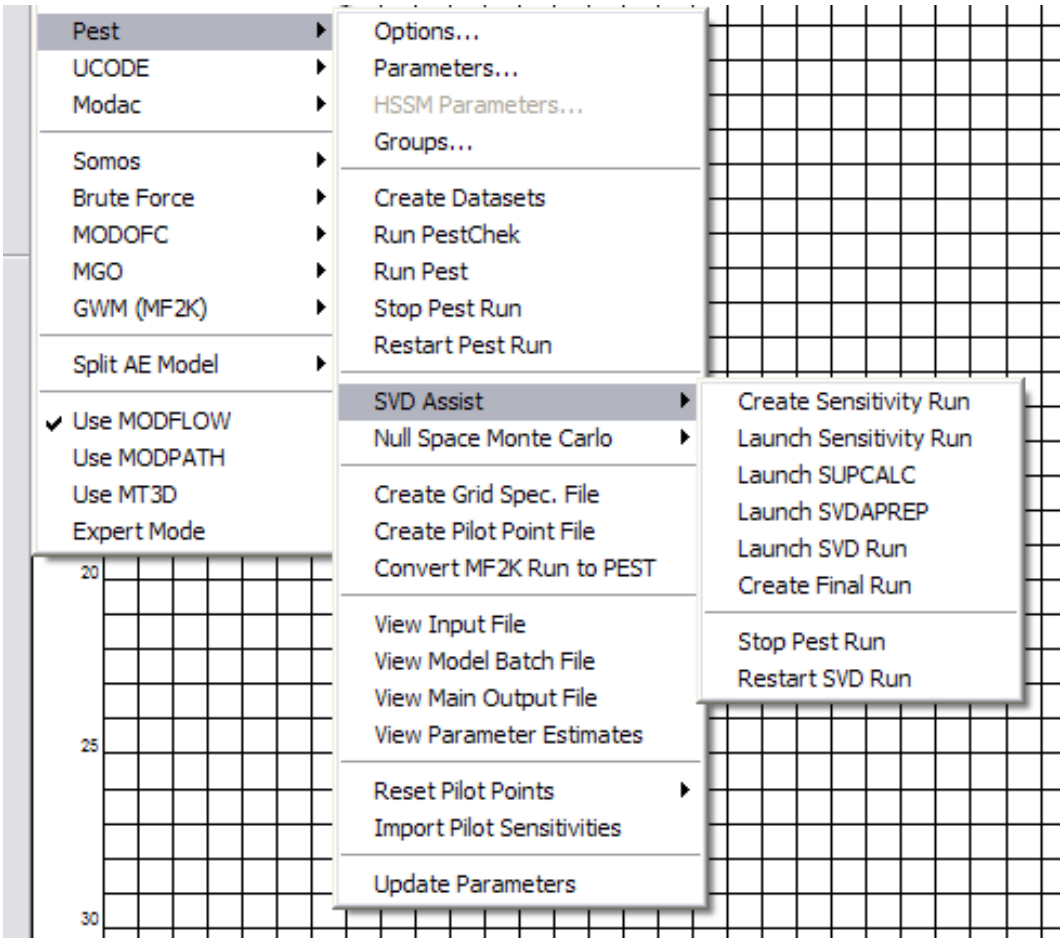

When using the **SVD Assist** menu in Groundwater Vistas, you do not need to make any changes to a normal Pest run. Just set up Pest as described in the previous examples. Once the Pest run is configured, simply start at the top of the SVD Assist menu and work your way down. Do not go to the next menu item until all of the steps above it have been accomplished. We will now work through an example using a simple 3D model to illustrate the SVD technique.

If your SVD Assist menu is grayed out and not accessible, it means you have the Standard version of Groundwater Vistas. Only the Advanced version and above can access this menu.

Setting Up Pest

This example is a variation on the basic Groundwater Vistas tutorial model. We created a heterogeneous hydraulic conductivity field, ran this hypothetical model, and extracted water level targets. You will start with a homogeneous model, add calibration targets, add pilot points, and then use SVD to calibrate the model. Start by opening the GV file called ***New_GV6_tutorial_pilotstart.gwv*** (NOTE: there are several files with similar names – be sure to get the right one). After opening this model, your screen should look like the one below:

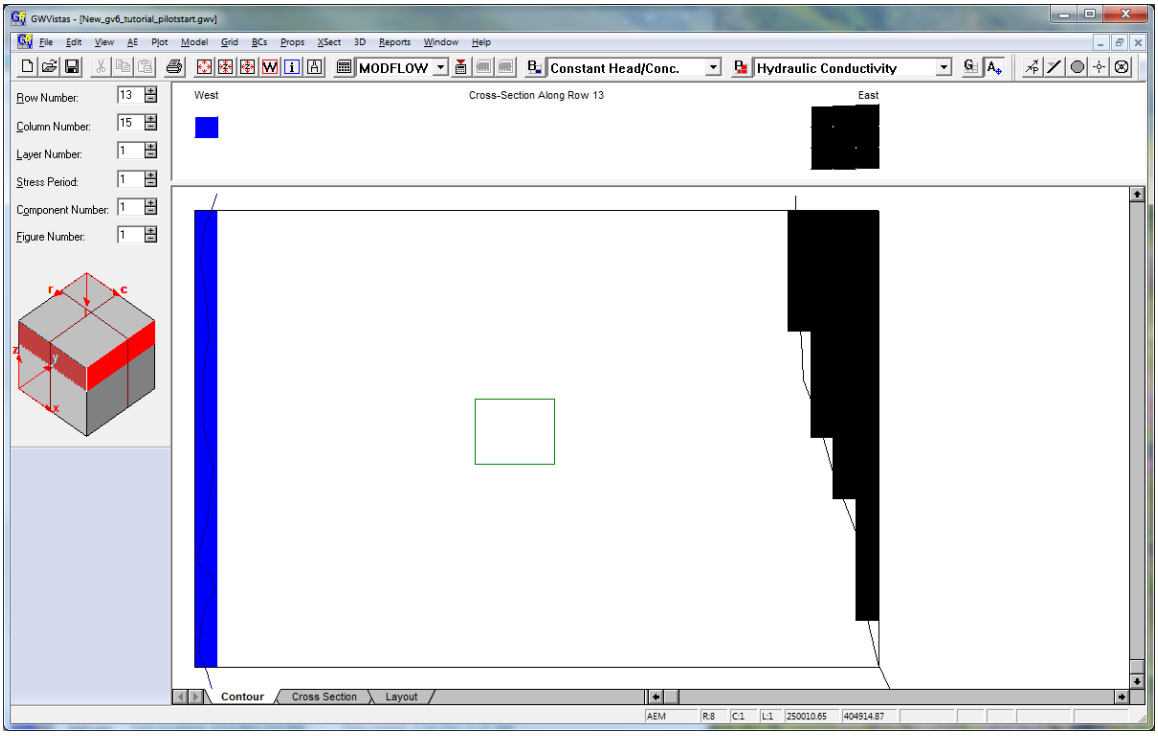

The model is complete with the exception of calibration targets and pilot points. The calibration targets will be imported from an external text file. Select **AE|Import|Target Text File** and browse to find the file ***c:\gwv6\tutorial\Targets_newpilot.dat***. Enter 1 line to skip, otherwise, the default settings on the import dialog are correct. GV should report 42 targets imported and your screen should look like the following:

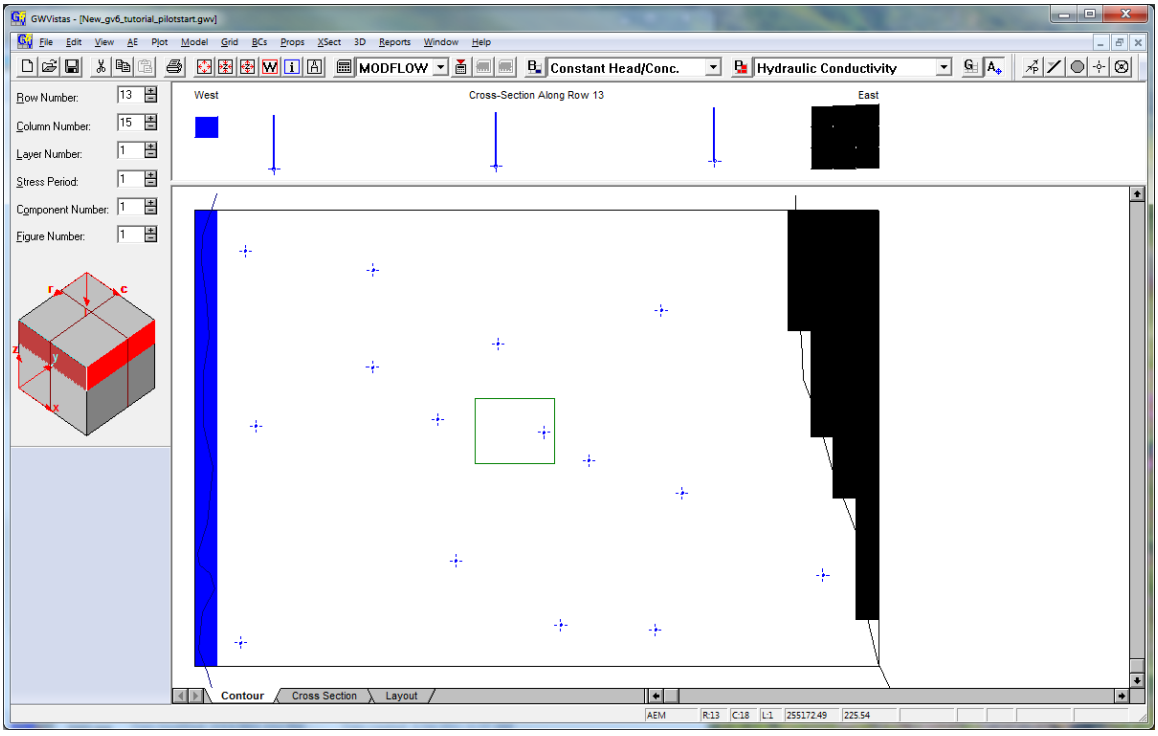

Now we are going to add pilot points. The biggest problem with pilot points is that you need a lot of them. That makes them tedious to add manually. We have created a simple way to add pilot points by placing them evenly spaced on a rectangular array. Start this process by selecting **AE|Pilot Points|Quick Pilots**. The following dialog is displayed:

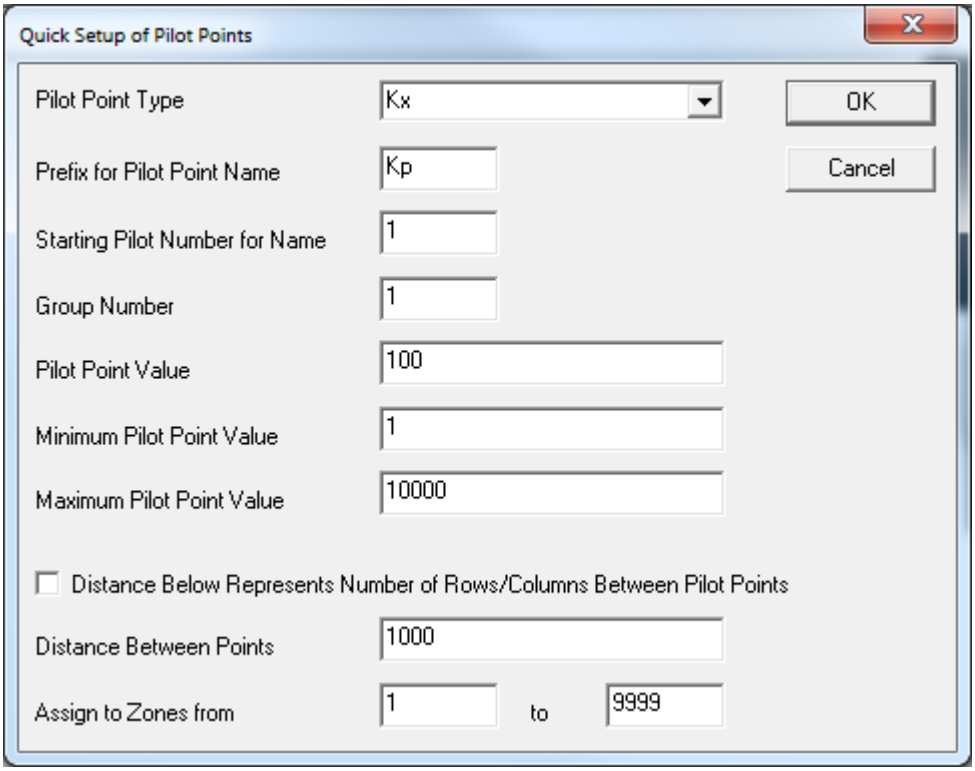

Change the minimum pilot point value to 1 and the maximum to 10000.

Change the distance from 100 to 1000.

One key thing about pilot points is that the names all have to be unique. Using this quick pilot point technique, GV will add the pilot points using a root name of "Kp" followed by a sequence number.

The starting sequence number is listed at the top of the dialog called "Starting Pilot Point Number for Name". It is best to start with 1 for the first set, as shown above. However, it is very important that in subsequent layers (or any time you add more points), you not duplicate a previous number.

Each pilot point is given a default starting K value, in this case 100 ft|d. Before clicking the OK button, carefully compare your dialog to the one above. There are several fields that are all multiples of 10 so getting the number of zeroes correct is important. Click OK and your screen should look like the following:

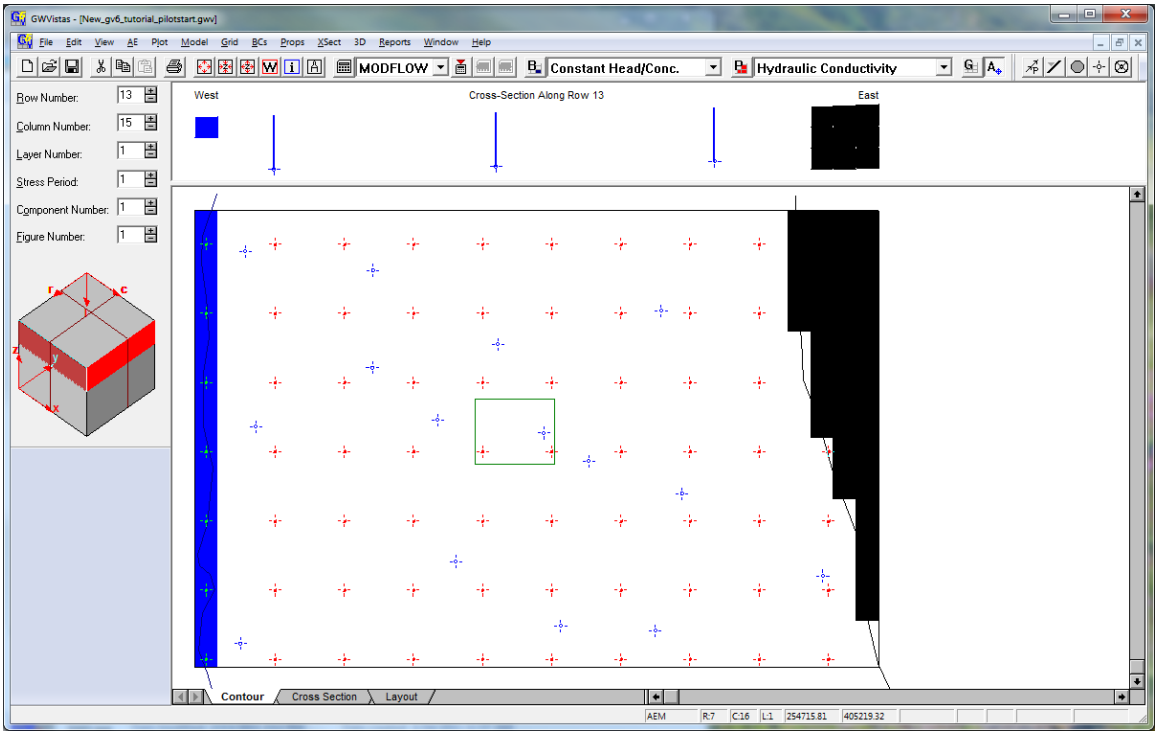

The red target symbols are the pilot points that you just added. **Now repeat this for layers 2 and 3**. Right now you are in layer 1. Press Ctrl-D to go down one layer or increment the layer number on the left side of your screen above the 3D cube. You know you are in the next layer because the red pilot points will disappear. After entering pilot points for layer 2, move to layer 3 and repeat.

You will note that the quick pilot points dialog automatically resets the starting sequence number ("Starting Pilot Point Number for Name") at the top of the dialog to track how many you have added. **Do not change this value.** That will assure that all pilot point names are unique.

After you have entered pilot points for layers 2 and 3, we will now add Kz pilot points using the same technique.

Go back to layer 1 and select **Add|Quick Pilot Points** again.

Change the sequence number ("Starting Pilot Point Number for Name") to 1.

Change the type of pilot point from Kx to Kz.

Change the pilot point prefix to Kz.

Change the default Kz value from 100 to 10.

Change the distance to 1000.

Change the minimum value to 0.1 and the maximum value to 1000. Kz values with thus range from 0.1 to 1,000 ft|d.

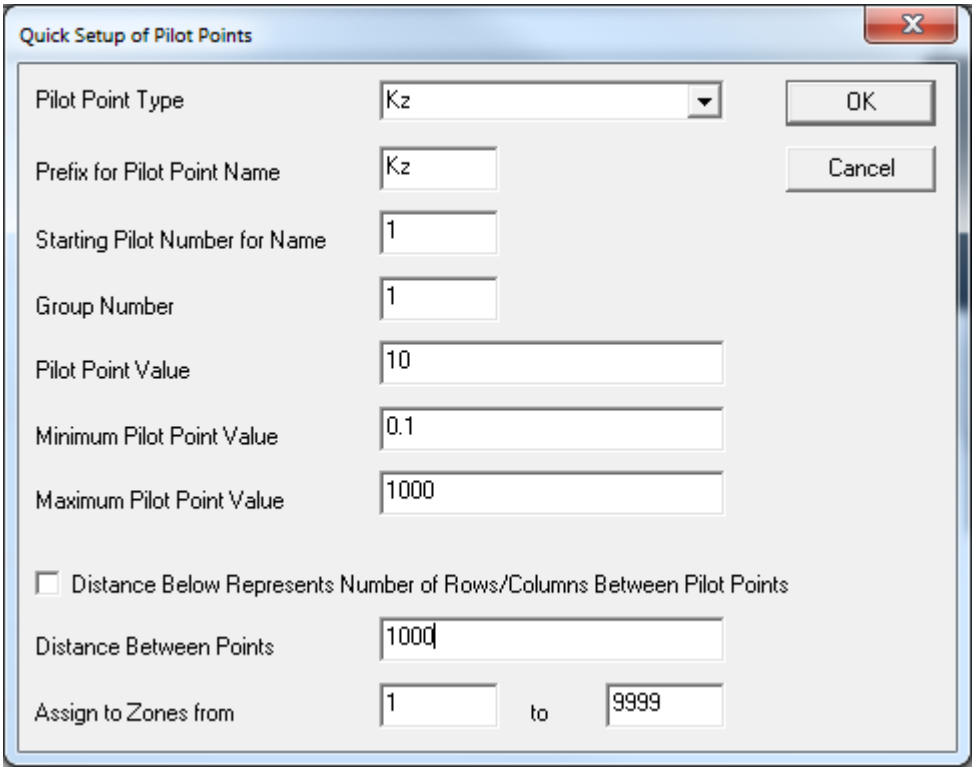

Repeat adding Kz pilot points to layers 2 and 3. You will note that after adding the Kz pilot points, your pilot points may disappear entirely! This is because they are drawn over top of each other and essentially "white out" each other. In order to avoid confusion, it is best to only display one type of pilot point at a time. Select **Plot|What to Display** and turn off the display of one of the pilot point types.

This is a good time to save your work! You might want to use **File|Save As** to rename the Vistas file.

The next step is to set the Pest options for using pilot points. Select **Model|Pest|Options**. On the Basic Options tab (the default one that is displayed), turn on regularization, use of pilot points, run the model without screen output, use command-line versions, and write BCF arrays to external files. When using pilot points, get in the habit of always checking these options, as shown below:

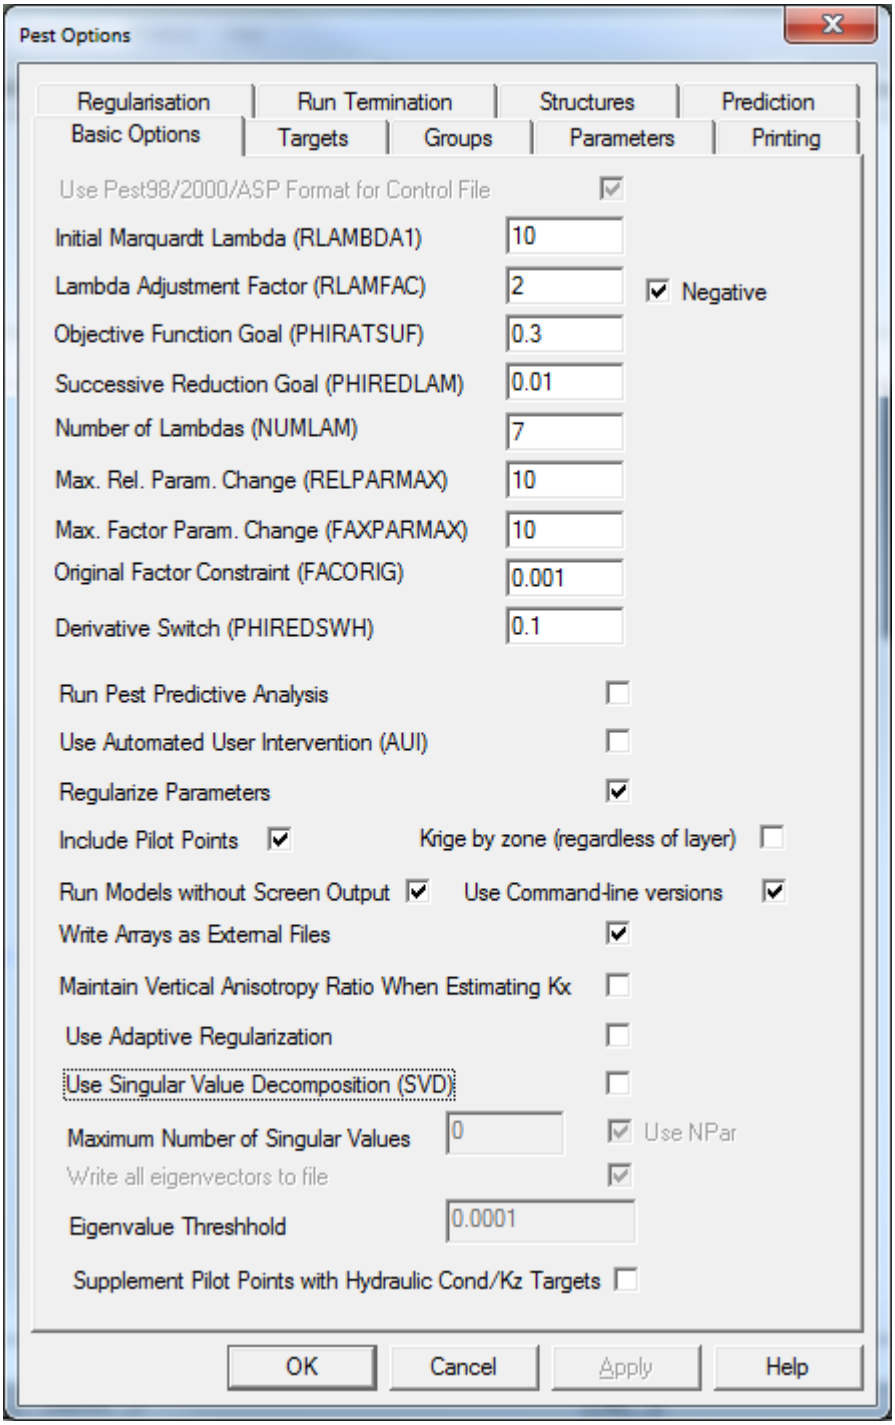

Now, click on the **Structures** tab. GV now does a pretty good job of setting reasonable default settings for these structures, which control the kriging of the pilot points. The only thing that we will change here is the maximum number of points to use in the kriging. Change that from 50 to 12. With gridded pilot points, 12 should be plenty for a smooth interpolation.

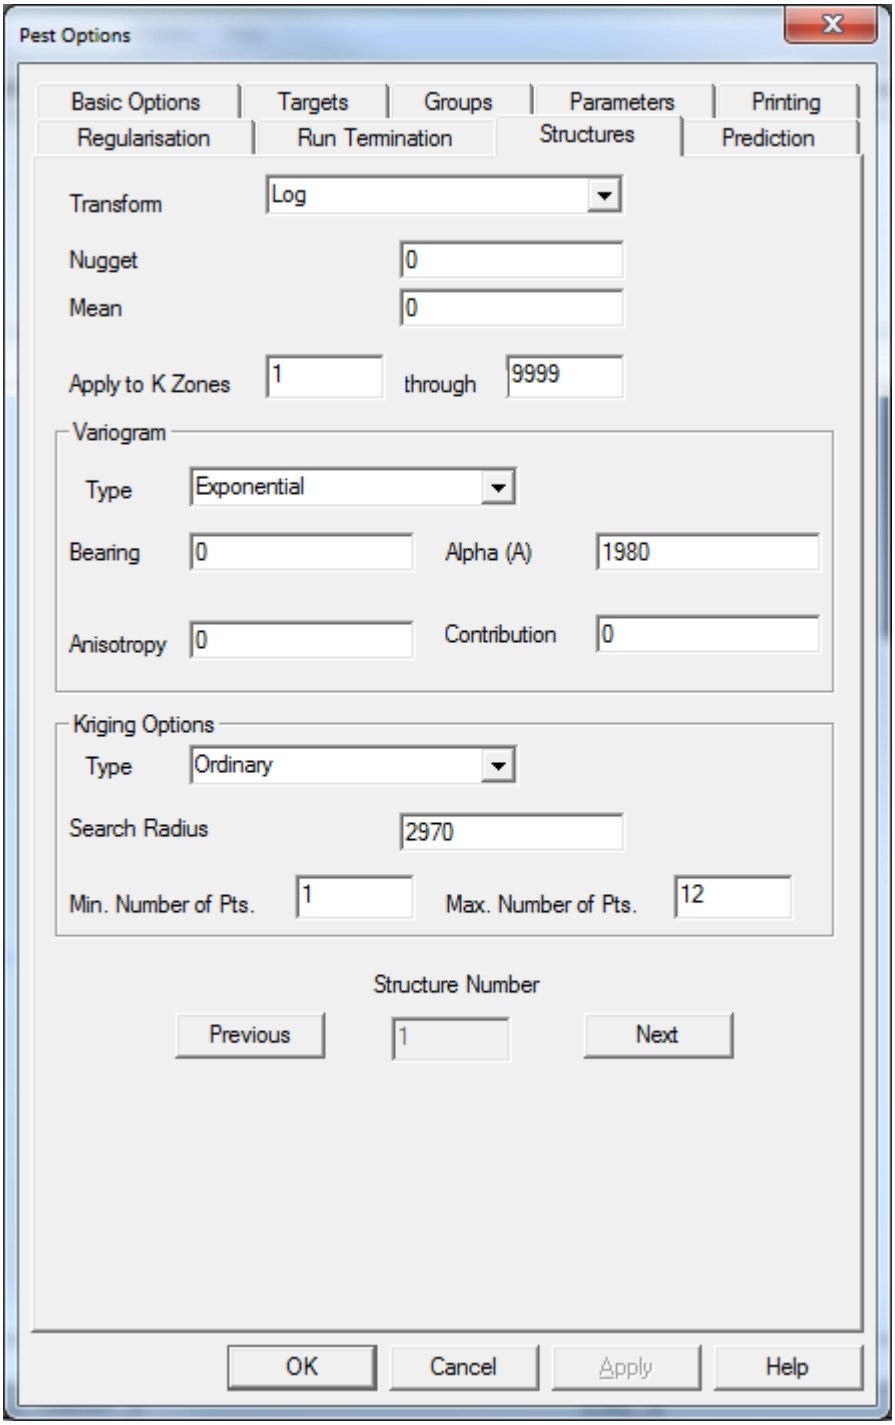

Now save your work again. The model is setup for pilot points. No other changes need to be made.

Using SVD Assist

To start the calibration, create the MODFLOW datasets using **Model|MODFLOW|Create Datasets**.

If this were a normal Pest run, you would select **Model|Pest|Create Datasets**. However, when using SVD with Pest, you use the **Model|Pest|SVD Assist|Create Sensitivity Run** instead. Use that command now.

You should see lots of DOS windows flashing on your screen as GV runs PPK2FAC and FAC2REAL for each layer and for each pilot point type. GV will then run PPKREG for each layer and pilot point type. Thus a total of 18 DOS windows will flash – much too fast to count. GV also automatically changed the number of iterations to -1, set the initial lambda to 0.0, and set the number of lambda values to 1. You should see a message that the Pest file was successfully created. GV will automatically run Pestchek.

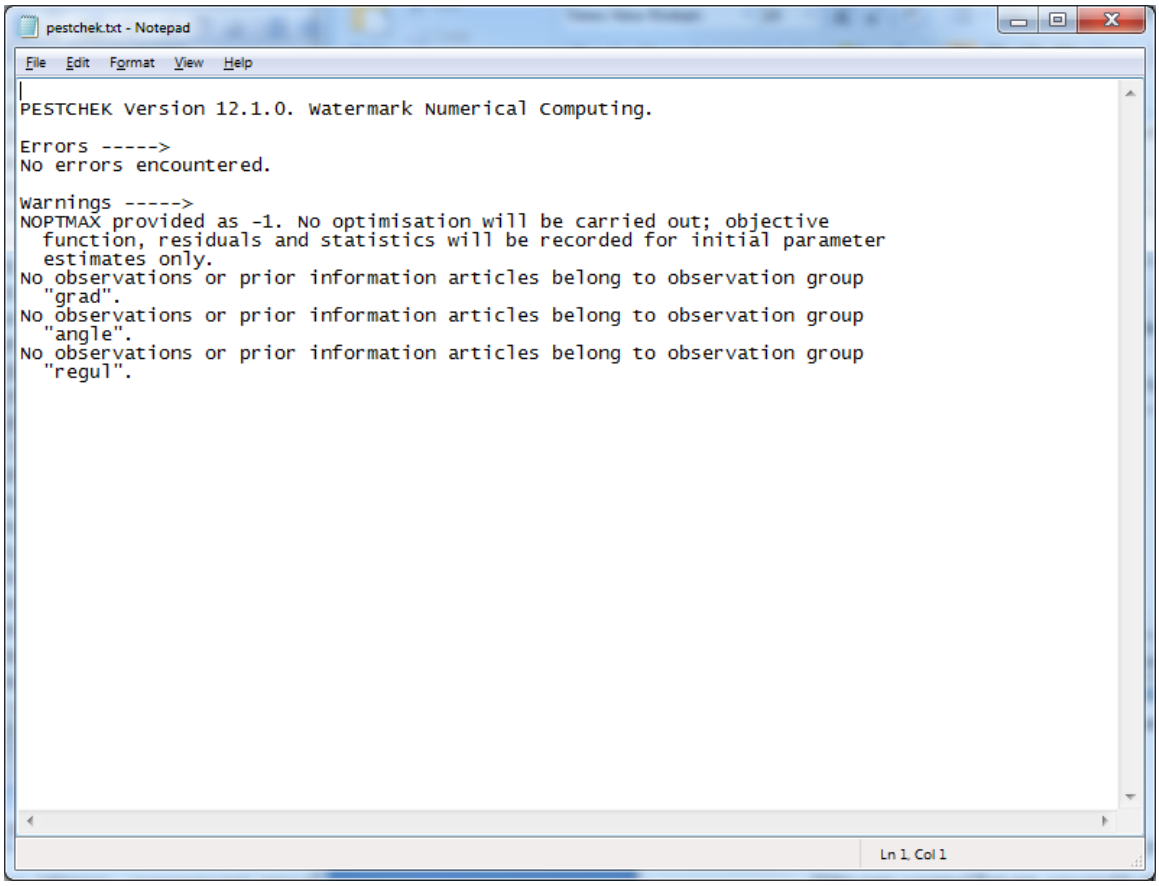

You should see some warnings but they are not serious. You can ignore them.

Now, run this initial sensitivity run. If you followed the steps for adding the pilot points, GV should run the model 402 times during the sensitivity run. This will take a couple of minutes. Select **Model|Pest|SVD Assist|Launch Sensitivity Run** and take a break.

You will note when Pest runs that the starting phi (sum of squared residuals) is 3,829.9 ft^2.

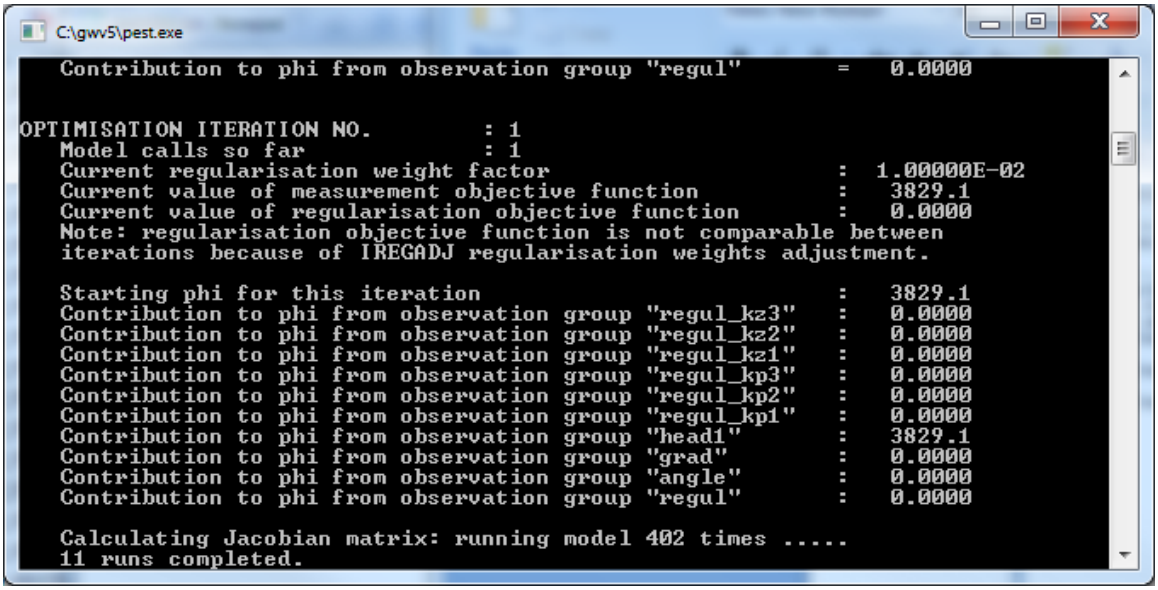

After the run finishes, we need to run SVDAPREP to reconfigure the Pest control file to use a set of super parameters. A super parameter is a linear combination of the base parameters (pilot points in our case). The question that is always asked is how many super parameters should you use? A new utility called SUPCALC estimates the minimum and maximum number of super parameters for a calibration. To run SUPCALC, first select **Model|Pest|Options** and select the **Regularization** tab. The value of PHIMLIM is the objective function (sum of squared residuals) you think is reasonable to attain. Since this is a synthetic case, leave the value at its default value of 1.0. In a real-world calibration, you would have to make a more realistic guess as to what phi is attainable. Now run SUPCALC using **Model|Pest|SVD Assist|Launch SUPCALC.** You should see the following in notepad:

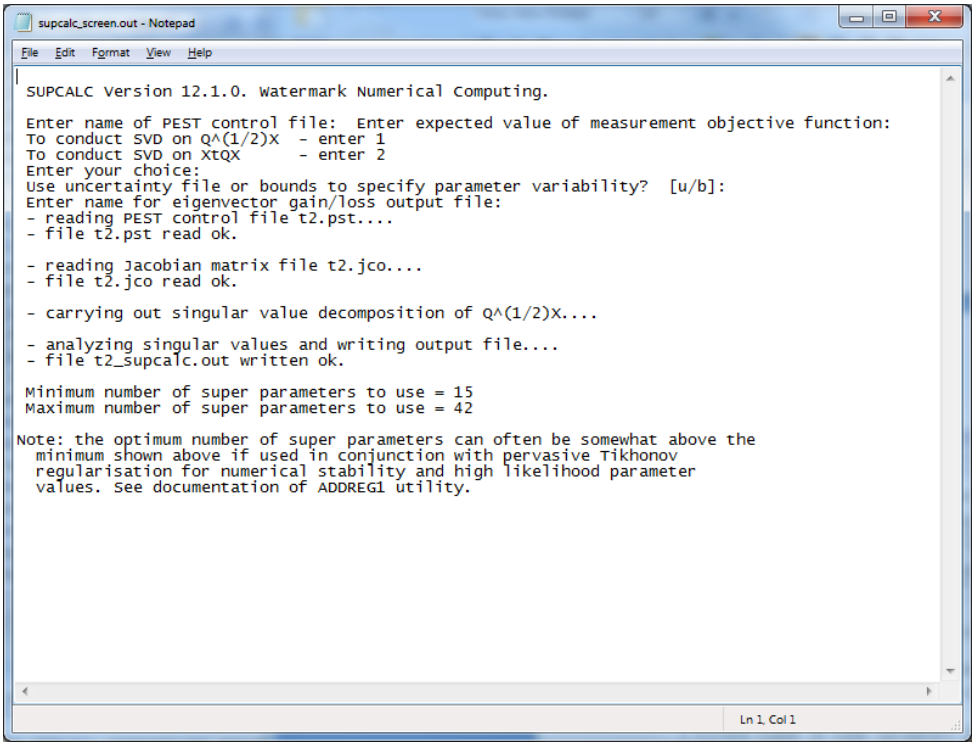

The important thing to note is the minimum and maximum number of super parameters which in this case should be around 15 to 42. We will use 40 but you might also try the lower end to see if it makes a difference.

To continue the calibration, select **Model|Pest|SVD Assist|Launch SVDAPREP**. A dialog is displayed with 3 options. Change the number of super parameters to 40 (between the range shown above from SUPCALC) and make sure the super parameter calculation option is set to 1. Click OK and GV will run the SVDAPREP utility to create a new Pest file (called *t2p_svd.pst*).

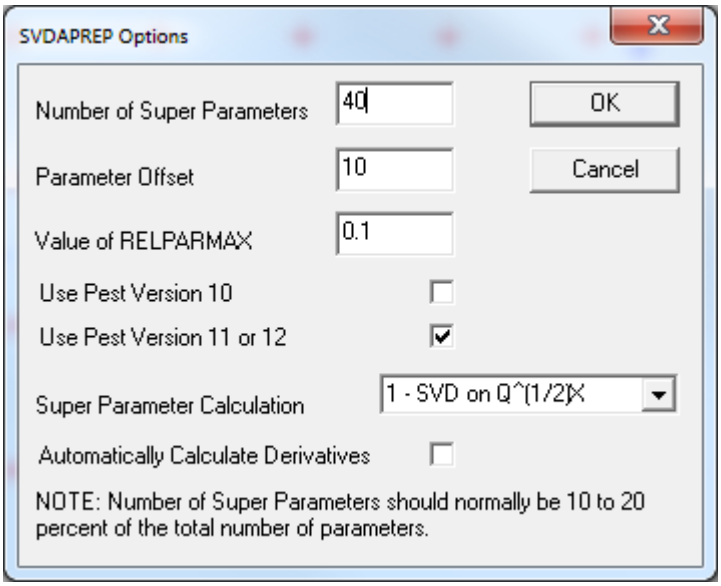

Next, select **Model|Pest|SVD Assist|Launch SVD Run**. You should now see that Pest is running the model 40 times per iteration (80 times after the first iteration when it switches to central differencing). Let it go at least 10 iterations. If you get tired of waiting, use **Model|Pest|Stop Pest Run** to kill the run. Pest always ends gracefully though so you can still use the results even if you end the run prematurely.

In order to have the right results for the next session on Null Space Monte Carlo, allow the run the go for 10 iterations. After it starts on the 11th iteration, select **Model|Pest|Stop Pest Run**. The last step is to use **Model|Pest|SVD Assist|Create Final Run**. This tells GV to run PARREP to assemble the final set of parameters, modify the Pest file to make only one model run, and run Pest with that final dataset.

To view the results of that final calibration run, use **Plot|Import Results** and click OK. Your screen should look like the following (the following is based on 10 iterations – you will also have to change the contouring parameters a bit):

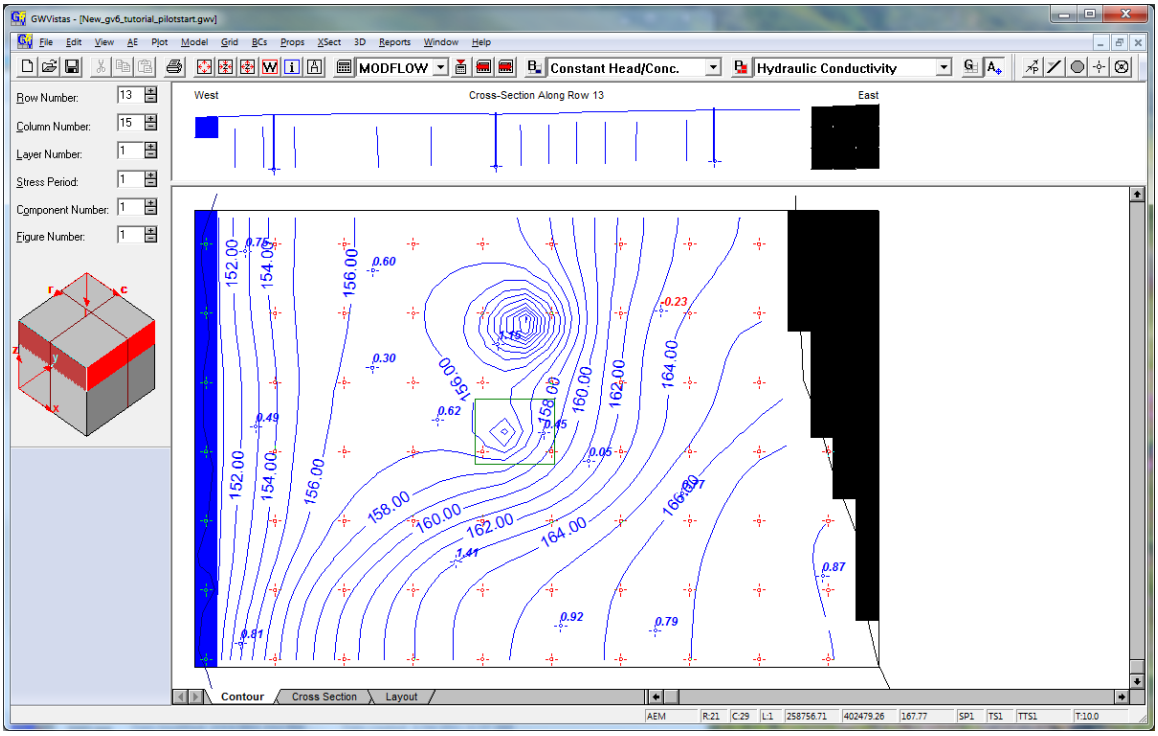

If you check the calibration statistics, you should see that Pest reduced the sum of squared residuals from the initial value of 3,830 down to about 31 (as long as you let it go at least 10 iterations). The plot of observed vs simulated head values is also quite good:

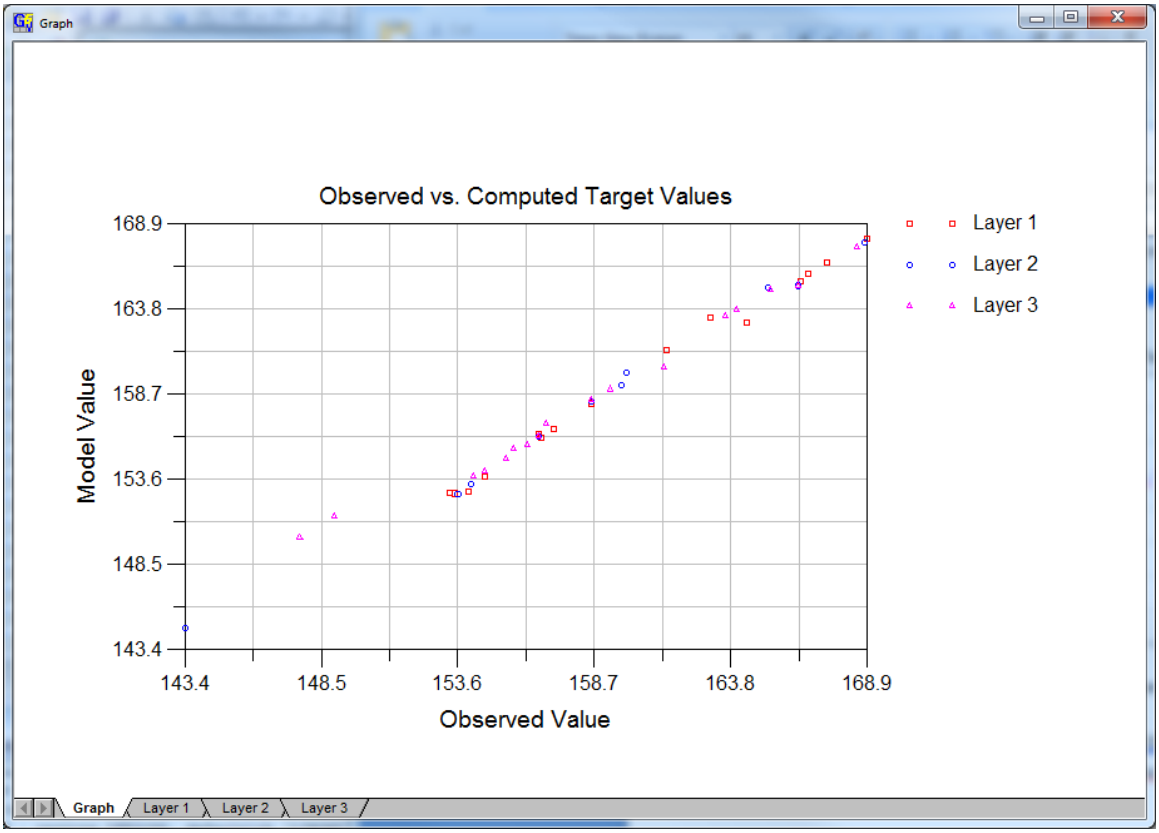

Please save your Groundwater Vistas file now! Use **File|Save As** and name the file *New_GV6_tutorial_pilotsetup.gwv.*

We are going to bring in the new parameter values from the Pest run. It is very important that you use **File|Save As** now to create a new Vistas file (name it ***New_GV6_tutorial_pilot_answer.gwv***) where these results will be stored. The reason we do this is that once you reset the hydraulic conductivities, you can no longer make a Pest pilot point calibration run with this version of the model. This model will be your final answer. However, if you want to go back, make changes to the calibration run, and recalibrate with Pest, then you need to use the one you saved previously.

To see the final hydraulic conductivity distribution for this run, select **Model|Pest|Update Parameters**. Answer Yes to reset the parameters and Yes to reset hydraulic conductivity values. The K values should look like the following:

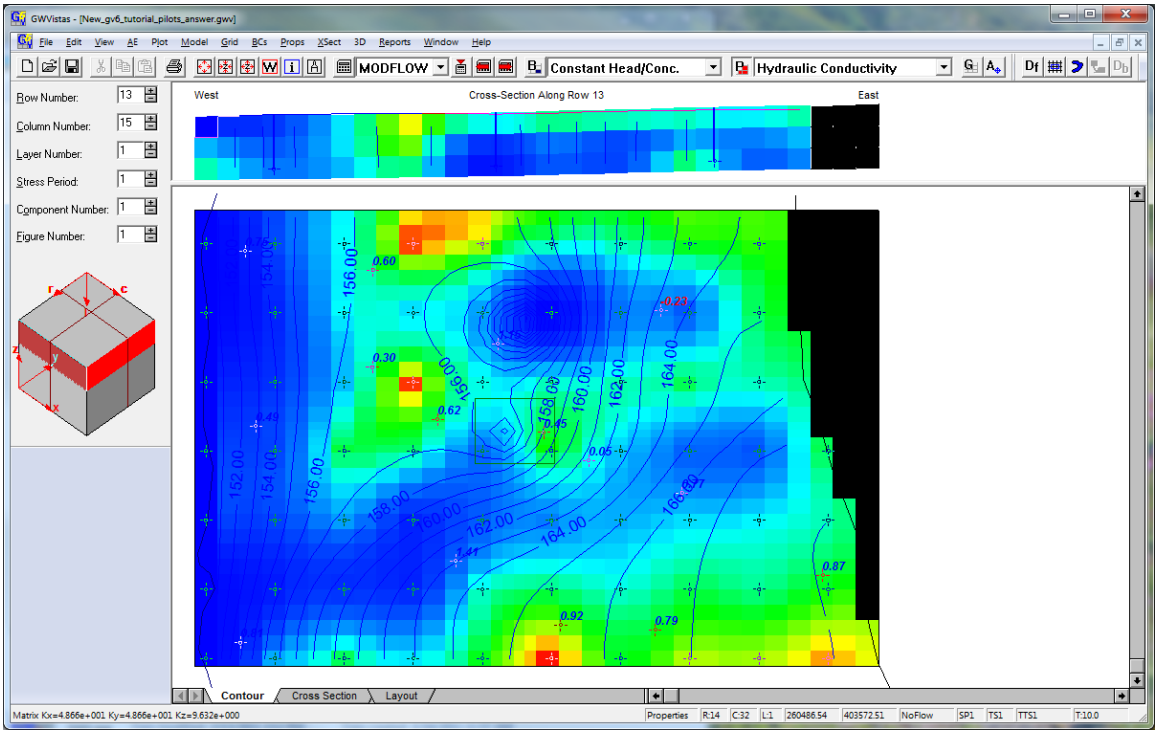

Save this as your final calibrated model using a new file name.

Now, open your previous Groundwater Vistas file (***New_GV6_tutorial_pilotsetup.gwv)*** that contained the setup for this calibration. Select **Model|Pest|Update Parameters**. Answer Yes to reset the parameters and **NO** to reset hydraulic conductivity values. This will update our pilot points to the final calibrated values but not import the final Kx and Kz matrices. This will serve as the starting point for the next session where we explore the new Null Space Monte Carlo procedure.

Select File|Save As and name the file *New_GV5_tutorial_NullSpace.gwv* which we will use in the next session.

PEST and Null Space Monte Carlo

Introduction

Another powerful new feature in Pest is called Null Space Monte Carlo. In this procedure, you can create any number of calibrated versions (called realizations) of your model. You saw in the last session that Singular Value Decomposition or SVD can dramatically reduce the time needed to calibrate very complex models. The Null Space Monte Carlo technique is designed to work on a model calibrated using SVD Assist with many parameters, as in a pilot point calibration. While pilot points are not required for Null Space Monte Carlo, a large number of parameters is a requirement.

Null Space Monte Carlo would be used after finalizing a calibration using SVD Assist. The procedure then uses the following steps:

Set Options. There are some options to set prior to running the Null Space Monte Carlo simulation. You need to reset the number of super parameters to be the same as you used in the previous SVD Assist run, for example.

Get Covariance Matrices. This step is only needed when you are using the random method for generating new pilot point values or if you have non-pilot point parameters in the calibration. If neither of these two situations are true, then GV will gray out this menu selection.

Update Jacobian Matrix. Remember that the Jacobian matrix is the matrix of sensitivities where the change in residual at each observation is divided by the change in the parameter value. When you run SVD Assist, the Jacobian Matrix is based on the initial parameter values. Pest needs to update this matrix using the final parameter values after the calibration is finished. Thus, this step will require one model run per parameter.

Generate Realizations. A realization is a set of parameters based on the original calibration but modified using either a random number generator or a geostatistical field generator. It is not clear which is better as this technique is so new.

Calibrate Realizations. The realizations generated in step 4 will generally result in a degradation in model calibration. Pest is now run through a large batch file to tweak each realization back into calibration. Only a couple of iterations are allowed to keep run time to a reasonable level so that some realizations might not be accepted in the end.

Now, considering that most of the tasks in 1 through 5 above need to be done at the command line, this makes Null Space Monte Carlo even more difficult to implement than SVD Assist. As with SVD Assist, however, Groundwater Vistas has been configured to make this as simple as possible. This new menu is under **Model|Pest|Null Space Monte Carlo|…** as shown below:

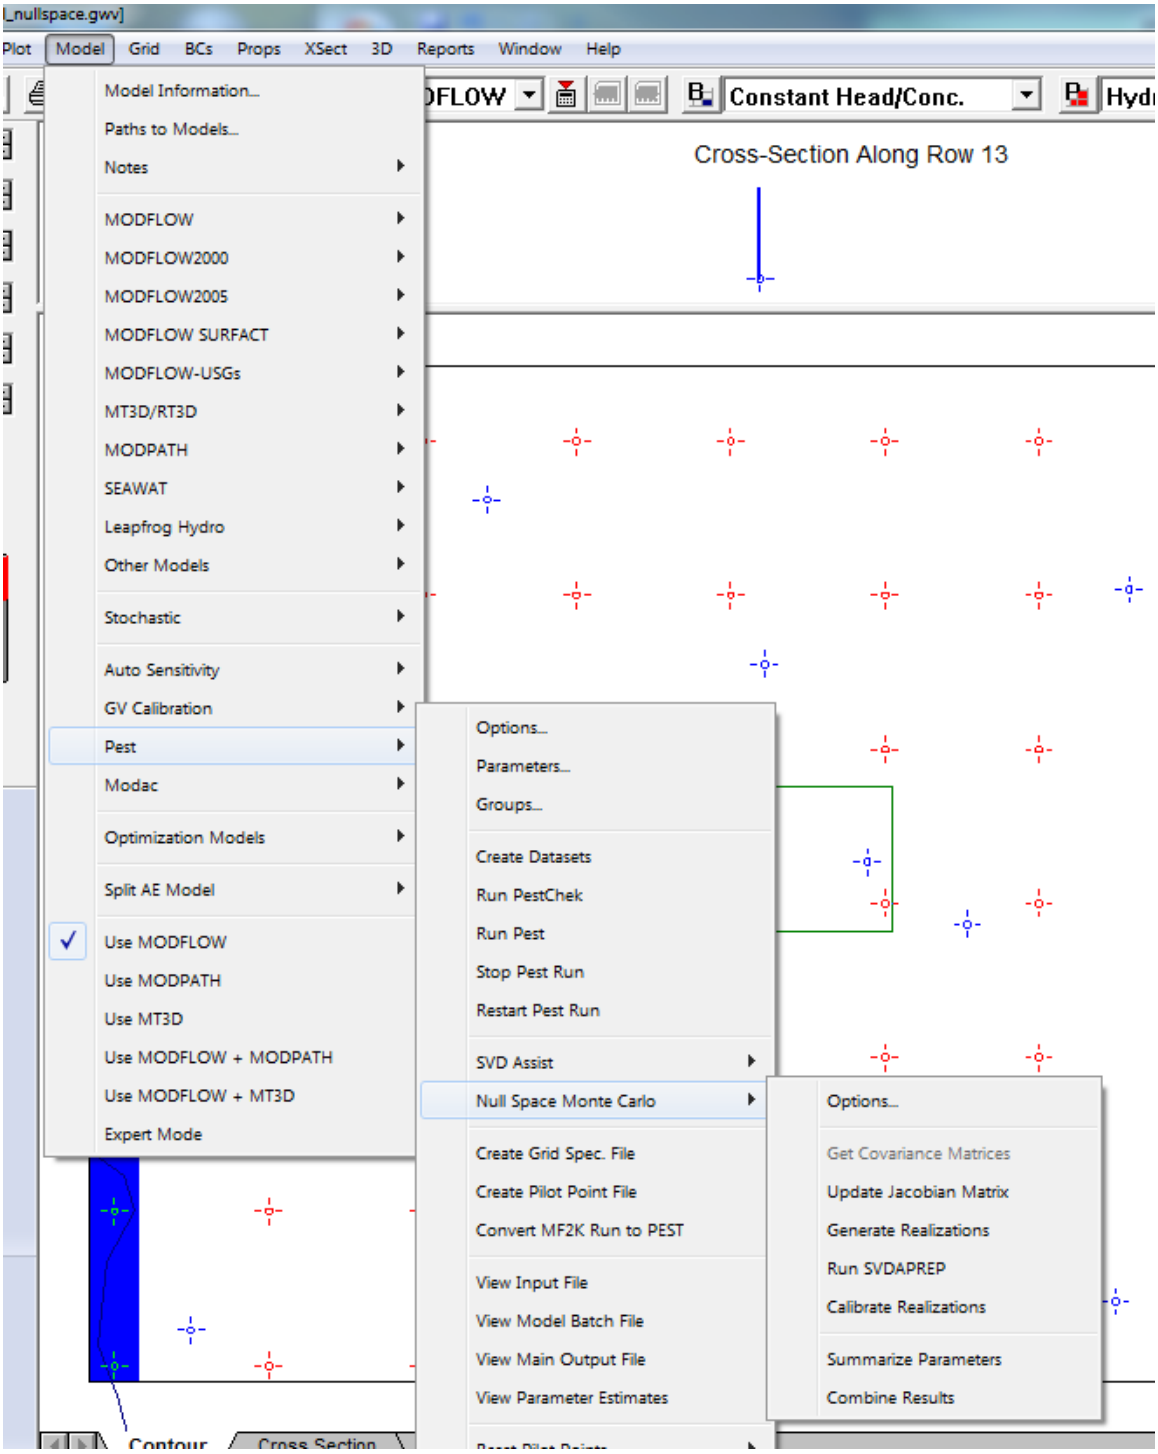

When using the **Null Space Monte Carlo,** you first calibrate your model using SVD Assist as in computer session 4. Then you update parameters so that the pilot points and|or other parameters are set to the calibrated values. Next, use the Null Space Monte Carlo menu and start at the top, executing each menu item in turn as you work your way down. We will now work through an example using the model calibration from Session 4.

Reset Pilot Point Bounds

Start by opening the model you created from the last session (***New_GV5_tutorial_NullSpace.gwv***). If you did not finish or your results were not good, there should be a file of the same name already in the c:\gwv6\tutorial directory that has already been prepared.

One important thing to keep in mind is that the bounds on the pilot point values are used to sample new pilot point values for each realization. That means that the wider the bounds, the more the initial realization will be out of calibration potentially and thus the more model runs to tweak the calibration back into an acceptable level. So, to start, we will make the bounds on our final pilot points +|- a factor of 5.

Select **AE|Modify|Pilot Points|Bounds** and enter 5 for maximum and 0.2 for minimum and then check the box at the bottom to indicate these are multipliers on current pilot point values

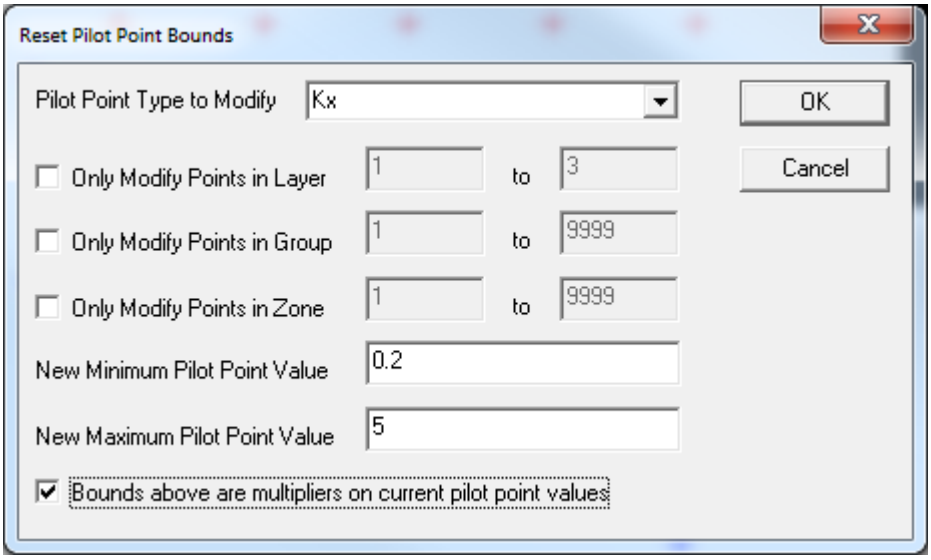

Repeat this for Kz pilot points.

Null Space Monte Carlo Options

Run the model to make sure that all files have been created for the MODFLOW simulation. Then select **Model|Pest|Null Space Monte Carlo|Options**.

Set the number of realizations to 10 (just to keep the run time down).

Change the number of super parameters at the top to 40.

Set the option for generating parameter sets to “Stochastic Fields Using Fieldgen”

Change the Objective Function Threshold to 50.

Change the Abandonment Threshold to 1000.

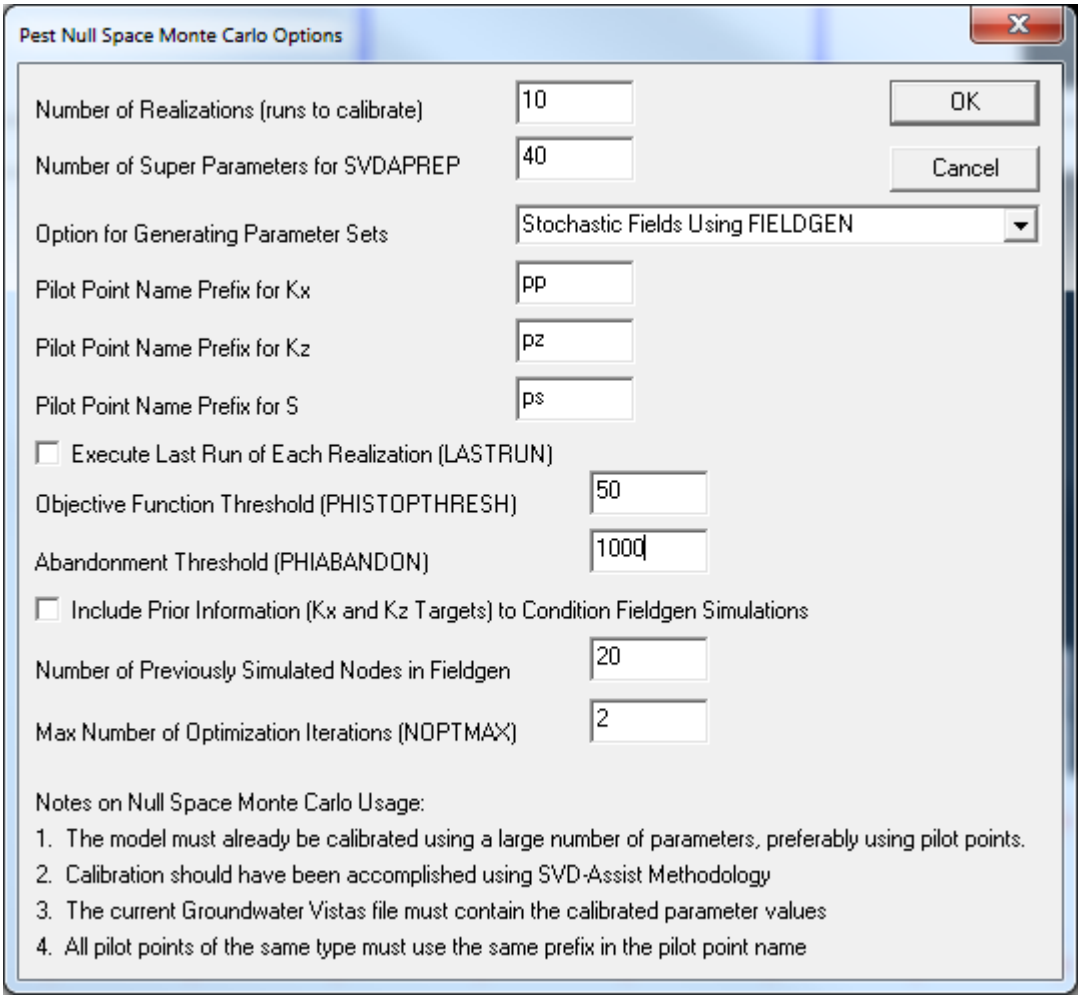

The Objective Function Threshold is the value of the sum of squared residuals that signifies an acceptable calibration. If the objective function falls to this level during the recalibration of a realization, Pest will stop and move on. On the other hand, if the objective function is above the abandonment threshold at any time, Pest will quit and move to the next realization. These values are designed to make the procedure more efficient by avoiding too much calibration or runs that are so poorly calibrated that time is not wasted on them.

Update Jacobian Matrix

The next step is to recomputed the Jacobian Matrix using calibrated parameter values. This will require one model simulation for each parameter. In this case, that would be 402 model runs.

Select **Model|Pest|Null Space Monte Carlo|Update Jacobian Matrix**

Groundwater Vistas will create a new Pest control file (t2.pst). Note that regularization is turned off for this Pest run. After the new Pest control file is created, GV will display the following message:

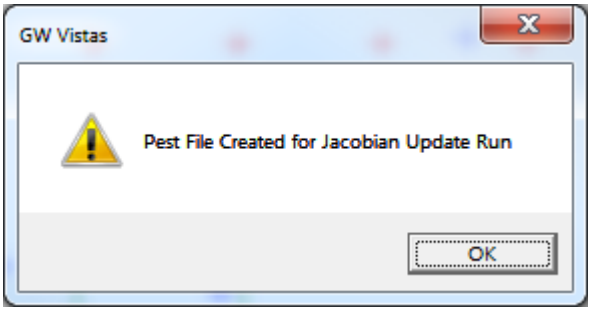

Click OK to continue. Groundwater Vistas will now run Pestchek on the new Pest control file and display the errors and warnings:

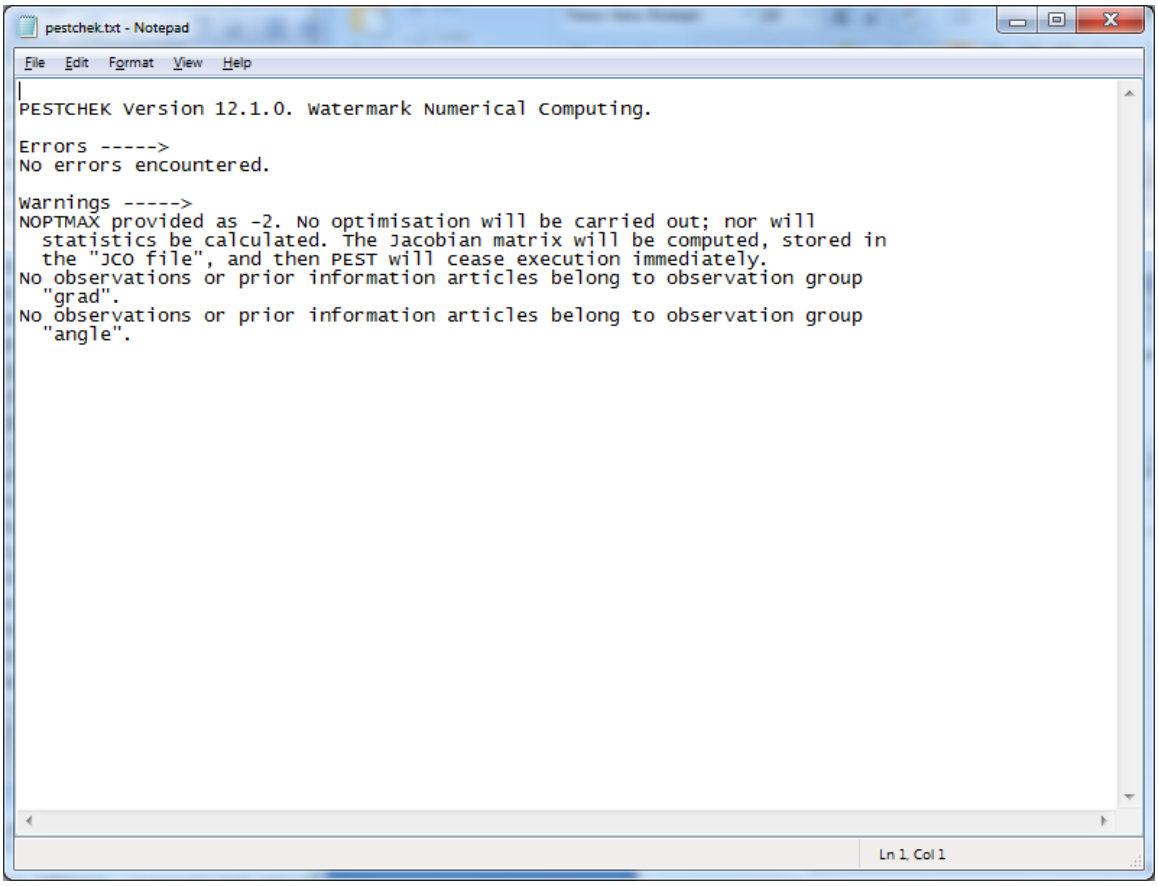

GV now asks if you want to continue. If there were errors, you should say NO. If there are only warnings, as shown above, then say YES.

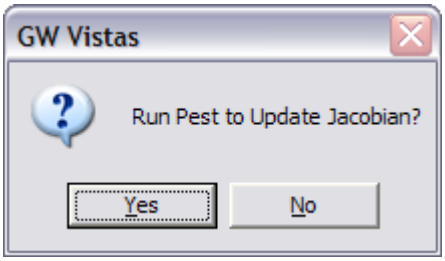

Pest will run. Note that the initial objective function (phi) should be the same as for your final calibration. If you ran 10 iterations of SVD, then it should be near 31:

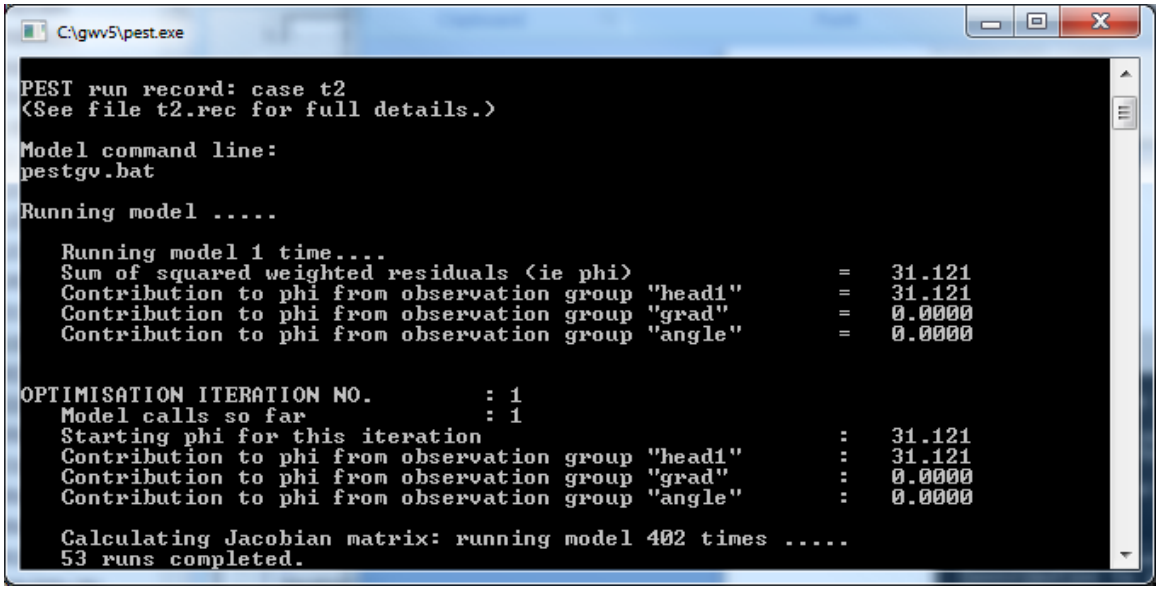

Generate Realizations

After the new Jacobian Matrix is created, it is time to create the parameter values for each realization. During this step, the program Fieldgen is run to create noisy Kx and Kz fields based on the variogram (structure) we used in the original calibration. Next, a program called PPSAMP samples pilot point values from these noisy fields and PNULPAR creates a final set of parameter files that Pest will use as starting guesses for parameters during recalibration of each realization.

Select **Model|Pest|Null Space Monte Carlo|Generate Realizations.**

You should see lots of DOS windows flashing as these utility programs generate the various files needed for the rest of the analysis.

Run SVDAPREP

Now, SVDAPREP must be run again to generate a revised SVD Pest control file that will be used to recalibrate each realization. GV will regenerate a main Pest control file, run SVDAPREP, and finally revise the SVDAPREP-generated Pest control file to reset some variables.

Select **Model|Pest|Null Space Monte Carlo|Run SVDAPREP**

Groundwater Vistas will first ask if you want to turn off regularization. It is probably best to answer YES. Without regularization, you get the first optimization iteration of each realization free (that is, no model runs). That can be quite a time savings. Probably the best approach is to try it the first time without regularization and then if it does not yield good results, run another set of runs with regularization.

Answer YES to this prompt:

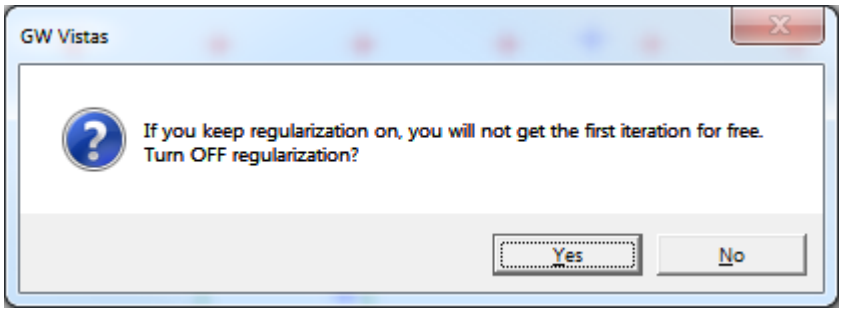

Groundwater Vistas now alerts you to the fact that it is going to regenerate a new Pest control file. This will contain no regularization and will be used by SVDAPREP in a subsequent step.

Click OK to continue.

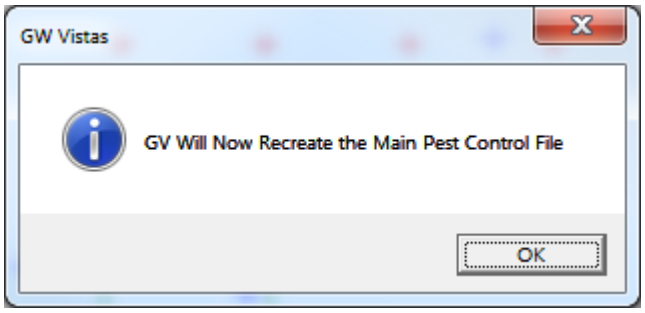

The new Pest control file was created.

Click OK to run Pestcheck:

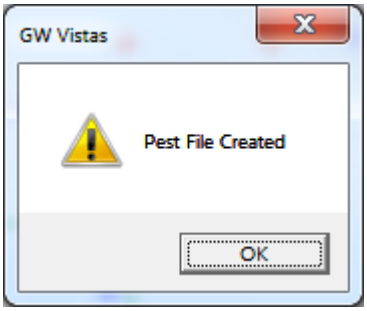

Hopefully Pestchek found no errors and only a few warnings as shown below. Close notepad after viewing the file.

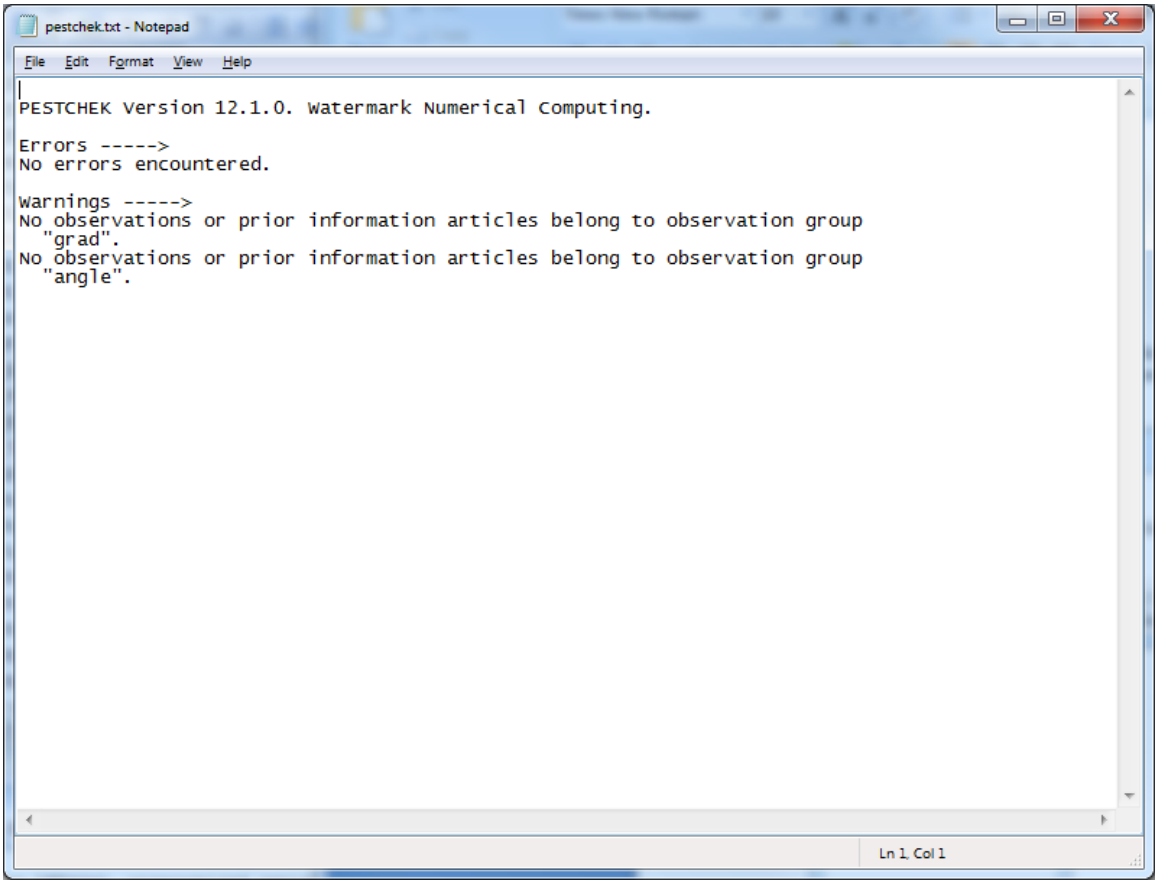

Groundwater Vistas is now ready to run SVDAPREP.

Click OK to continue:

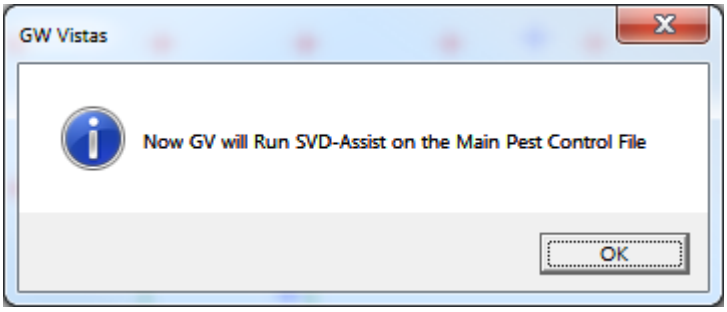

The SVDAPREP options should be filled out for you already. You should just confirm that the number of super parameters is correct (40 in our example). Note that the flag labeled "Automatically Calculate Derivatives" in ON. This tells Pest to run no model simulations in the first iteration.

Click OK to continue:

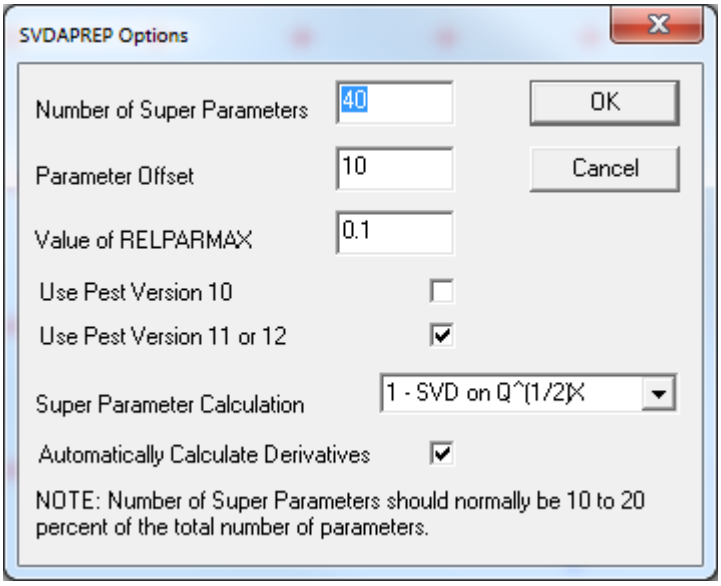

We should now have all of the many files we need to run Null Space Monte Carlo. If you go into Windows Explorer and browse to the c:\gwv6\tutorial\work directory and sort by date, you should see the most recently modified file is **t2_svd.pst** which is the final product of this whole sequence of commands.

Calibrate Each Realization

You are ready to calibrate each realization. To do this, select **Model|Pest|Null Space Monte Carlo|Calibrate Realizations.**

You will see a DOS window open like in a normal Pest run. Instead of running just one Pest simulation, however, it will be running 10 of them. The first iteration will be done without any model runs (except for Lambda testing). If it can reach a sum of squares of 50, it will then move to the next one. If not, it will try one more iteration before giving up. The file c:\gwv6\tutorial\work\Record.dat contains a summary of the results of each realization. You can open this file periodically to see how far it has progressed.

After each realization has been calibrated, GV will display the Record.dat file as shown below. You will see that there are a one or two of the simulations that did not achieve the desired sum of squared residuals of 50. In the example shown below, all 10 realizations met the requirement of 50 ft^2, however, your results may differ slightly. Do not be concerned if your results do not match this exactly.

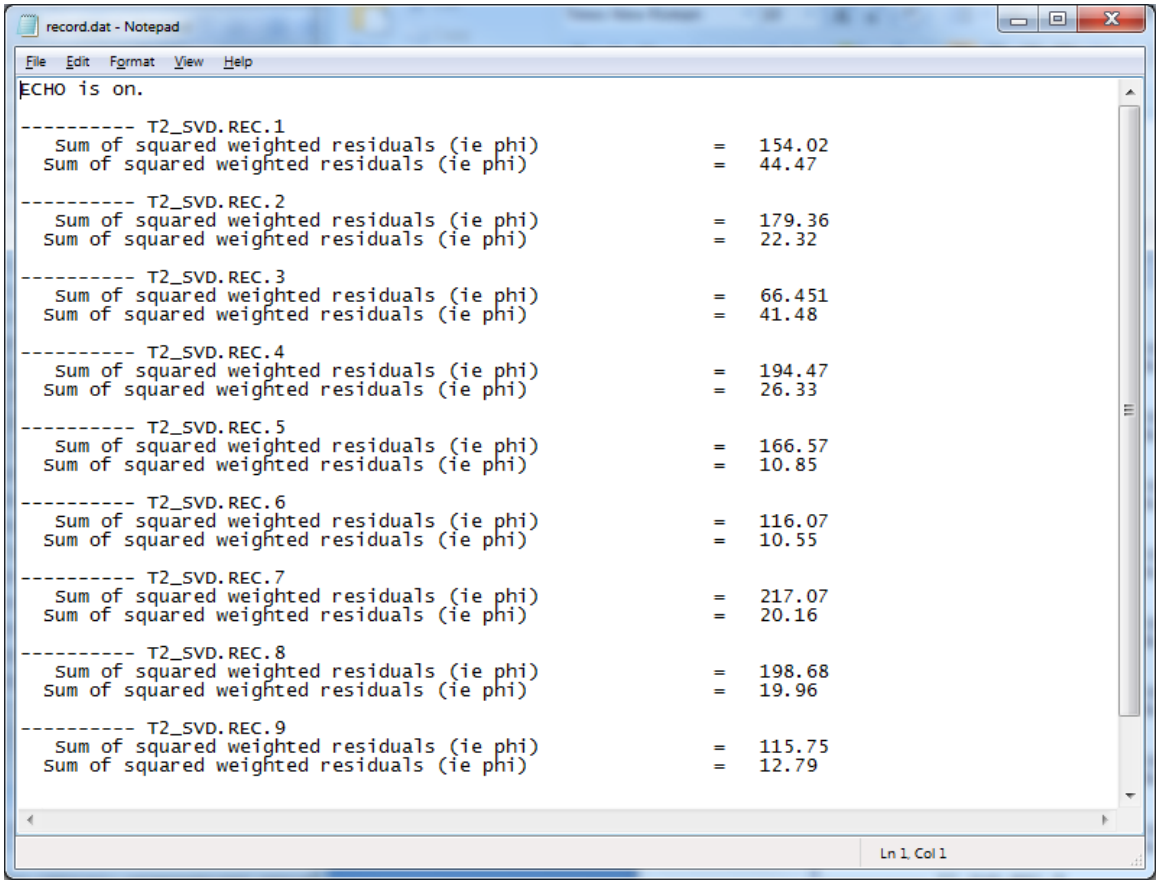

Post-Processing of Monte Carlo Results

So now you have 10 different model calibrations for this example. What do we do with this stuff? That is a good question. The final parameters for each realization can be found in the files called t2.bpa.# where # is the realization number. The Pest summary output file for each realization can be found in the files called t2_svd.rec.# where # is the realization number.

Post-processing of the realizations will be an ongoing development exercise in Groundwater Vistas. For now, there are a couple of things you can do. First, select **Model|Pest|Null Space Monte Carlo|Summarize Parameters**. Groundwater Vistas reads the file ***record.dat*** shown above along with the parameter files and summarizes the results for you. Two text files will be displayed. The first is called ***NullSpaceSummary.txt*** as shown below.

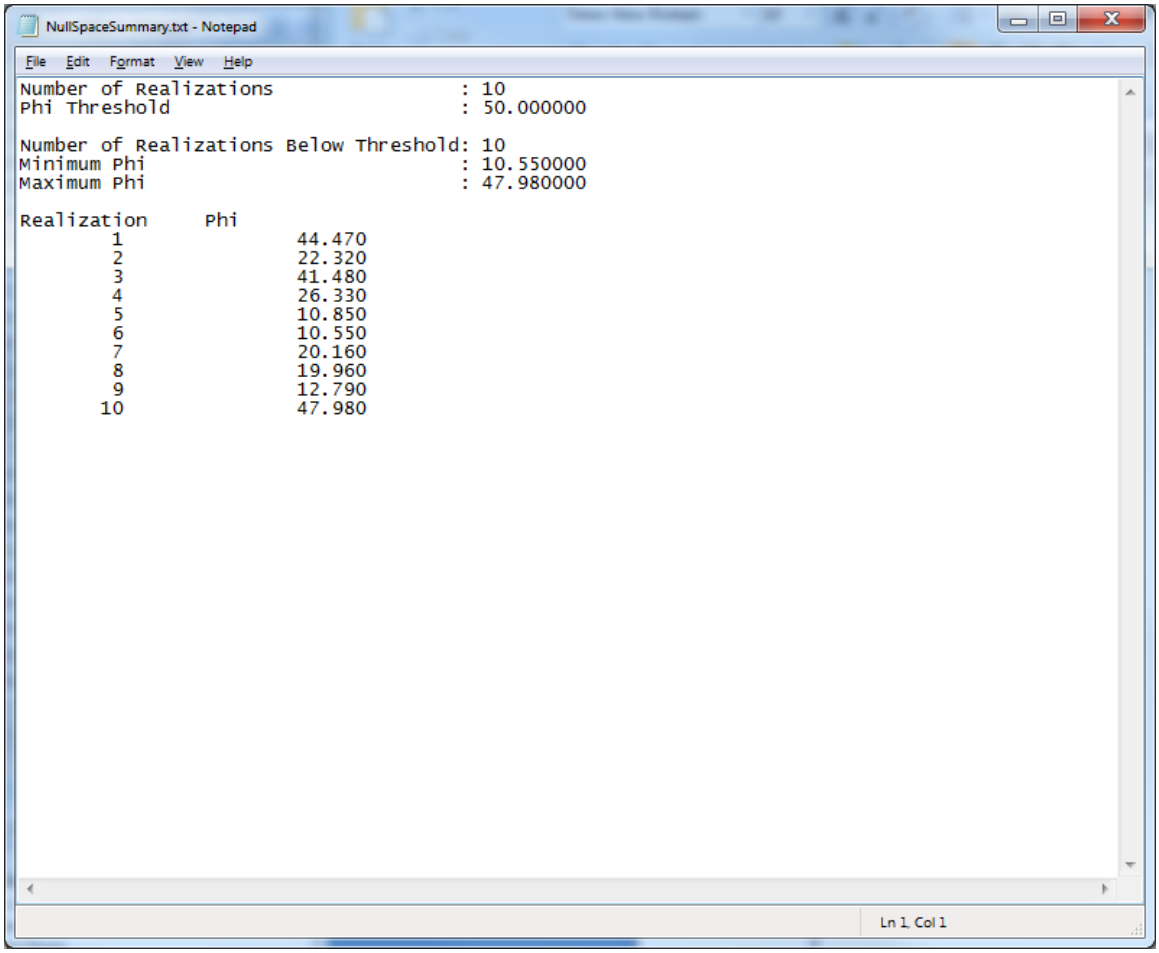

This file just summarizes the number of realizations, the number that met the phi threshold, and then the phi (sum of squared residuals) for each realization.

The second file is called ***NullSpaceParameters.txt*** and shows the mean, log mean, standard deviation, minimum value, and maximum value for each parameter. You can see from the example below that the range in parameter values can be quite extreme and yet each realization is well calibrated.

.

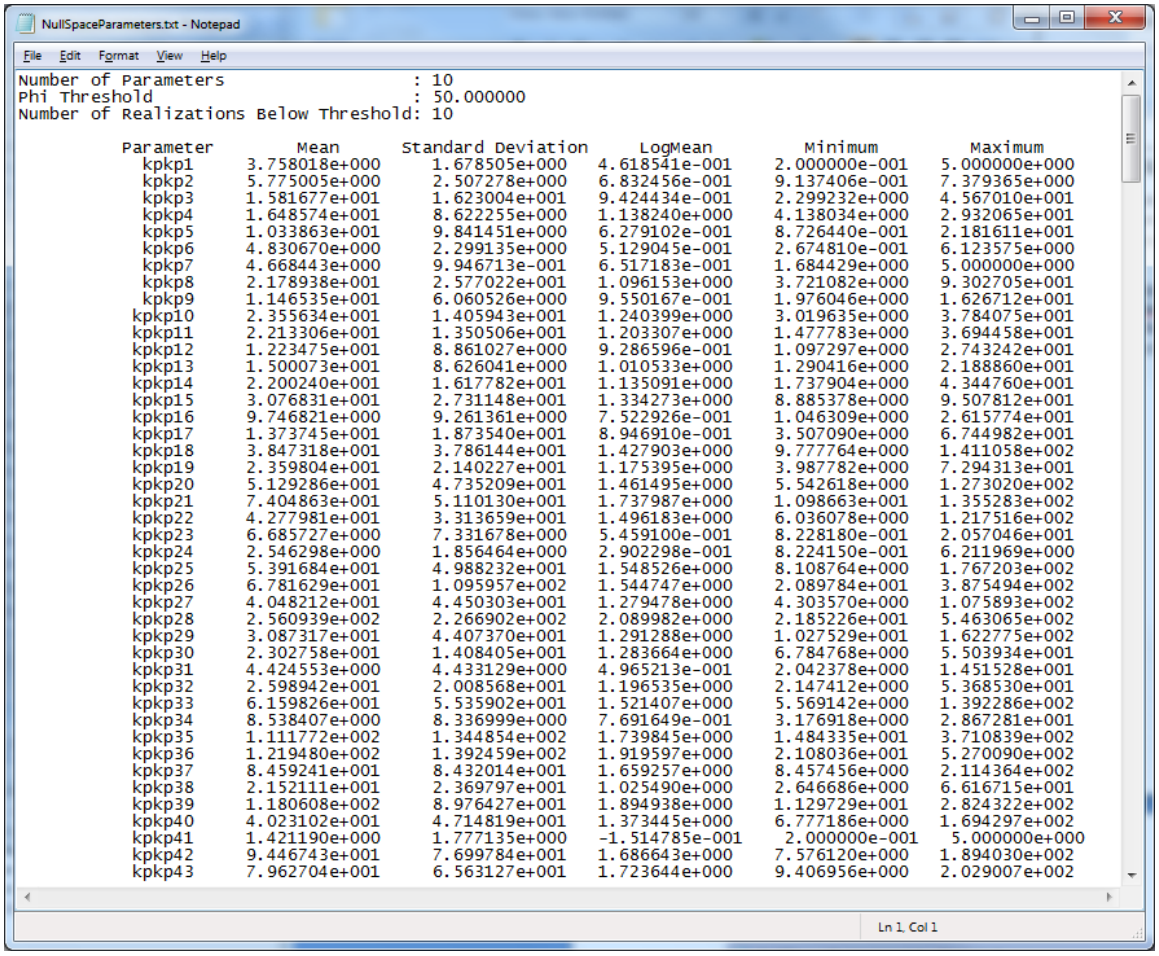
NullSpaceParameters.txt - Notepad

File Edit Format View Help

Number of Parameters : 10
Phi Threshold : 50.000000
Number of Realizations Below Threshold: 10

Parameter	Mean	Standard Deviation	LogMean	Minimum	Maximum
kpkp1	3.758018e+000	1.678505e+000	4.618541e-001	2.000000e-001	5.000000e+000
kpkp2	5.775005e+000	2.507278e+000	6.832456e-001	9.137406e-001	7.379365e+000
kpkp3	1.581677e+001	1.623004e+001	9.424434e-001	2.299232e+000	4.567010e+001
kpkp4	1.648574e+001	8.622255e+000	1.138240e+000	4.138034e+000	2.932065e+001
kpkp5	1.033863e+001	9.841451e+000	6.279102e-001	8.726440e-001	2.181611e+001
kpkp6	4.830670e+000	2.299135e+000	5.129045e-001	2.674810e-001	6.123575e+000
kpkp7	4.668443e+000	9.946713e-001	6.517183e-001	1.684429e+000	5.000000e+000
kpkp8	2.178938e+001	2.577022e+001	1.096153e+000	3.721082e+000	9.302705e+001
kpkp9	1.146535e+001	6.060526e+000	9.550167e-001	1.976046e+000	1.626712e+001
kpkp10	2.355634e+001	1.405943e+001	1.240399e+000	3.019635e+000	3.784075e+001
kpkp11	2.213306e+001	1.350506e+001	1.203307e+000	1.477783e+000	3.694458e+001
kpkp12	1.223475e+001	8.861027e+000	9.286596e-001	1.097297e+000	2.743242e+001
kpkp13	1.500073e+001	8.626041e+000	1.010533e+000	1.290416e+000	2.188860e+001
kpkp14	2.200240e+001	1.617782e+001	1.135091e+000	1.737904e+000	4.344760e+001
kpkp15	3.076831e+001	2.731148e+001	1.334273e+000	8.885378e+000	9.507812e+001
kpkp16	9.746821e+000	9.261361e+000	7.522926e-001	1.046309e+000	2.615774e+001
kpkp17	1.373745e+001	1.873540e+001	8.946910e-001	3.507090e+000	6.744982e+001
kpkp18	3.847318e+001	3.786144e+001	1.427903e+000	9.777764e+000	1.411058e+002
kpkp19	2.359804e+001	2.140227e+001	1.175395e+000	3.987782e+000	7.294313e+001
kpkp20	5.129286e+001	4.735209e+001	1.461495e+000	5.542618e+000	1.273020e+002
kpkp21	7.404863e+001	5.110130e+001	1.737987e+000	1.098663e+001	1.355283e+002
kpkp22	4.277981e+001	3.313659e+001	1.496183e+000	6.036078e+000	1.217516e+002
kpkp23	6.685727e+000	7.331678e+000	5.459100e-001	8.228180e-001	2.057046e+001
kpkp24	2.546298e+000	1.856464e+000	2.902298e-001	8.224150e-001	6.211969e+000
kpkp25	5.391684e+001	4.988232e+001	1.548526e+000	8.108764e+000	1.767203e+002
kpkp26	6.781629e+001	1.095957e+002	1.544747e+000	2.089784e+001	3.875494e+002
kpkp27	4.048212e+001	4.450303e+001	1.279478e+000	4.303570e+000	1.075893e+002
kpkp28	2.560939e+002	2.266902e+002	2.089982e+000	2.185226e+001	5.463065e+002
kpkp29	3.087317e+001	4.407370e+001	1.291288e+000	1.027529e+001	1.622775e+002
kpkp30	2.302758e+001	1.408405e+001	1.283664e+000	6.784768e+000	5.503934e+001
kpkp31	4.424553e+000	4.433129e+000	4.965213e-001	2.042378e+000	1.451528e+001
kpkp32	2.598942e+001	2.008568e+001	1.196535e+000	2.147412e+000	5.368530e+001
kpkp33	6.159826e+001	5.535902e+001	1.521407e+000	5.569142e+000	1.392286e+002
kpkp34	8.538407e+000	8.336999e+000	7.691649e-001	3.176918e+000	2.867281e+001
kpkp35	1.111772e+002	1.344854e+002	1.739845e+000	1.484335e+001	3.710839e+002
kpkp36	1.219480e+002	1.392459e+002	1.919597e+000	2.108036e+001	5.270090e+002
kpkp37	8.459241e+001	8.432014e+001	1.659257e+000	8.457456e+000	2.114364e+002
kpkp38	2.152111e+001	2.369797e+001	1.025490e+000	2.646686e+000	6.616715e+001
kpkp39	1.180608e+002	8.976427e+001	1.894938e+000	1.129729e+001	2.824322e+002
kpkp40	4.023102e+001	4.714819e+001	1.373445e+000	6.777186e+000	1.694297e+002
kpkp41	1.421190e+000	1.777135e+000	-1.514785e-001	2.000000e-001	5.000000e+000
kpkp42	9.446743e+001	7.699784e+001	1.686643e+000	7.576120e+000	1.894030e+002
kpkp43	7.962704e+001	6.563127e+001	1.723644e+000	9.406956e+000	2.029007e+002

Ln 1, Col 1

You can begin to see the degree of non-uniqueness in our models!

The next post-processing step is to assemble a file of parameters and heads from each realization so that we can start to analyze the results. Select **Model|Pest|Null Space Monte Carlo|Combine Results**. This creates another batch file with a loop that will run the model once per realization. The batch file finds the parameters from that realization, sets up a Pest run with zero iterations (so it runs the model once), and then takes the Kx arrays, Kz arrays, and heads and assembles three large binary files containing all of this information.

You can see the contents of the heads file using **Plot|File Operations|View Contents.** Browse to find the file called ***t2_nsmc.hds***. Then click the **View** button.

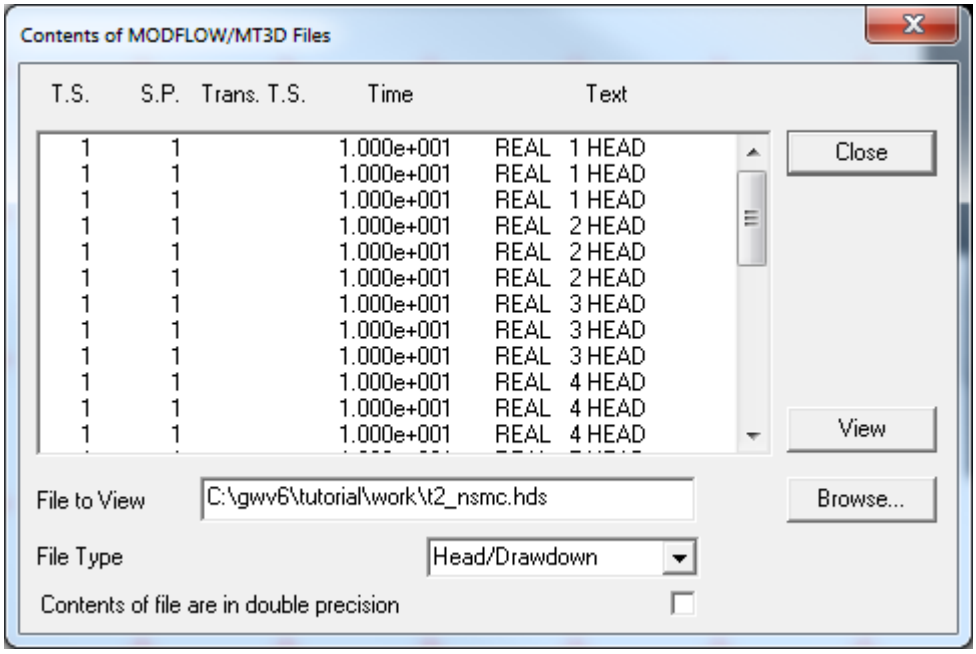

What you see is one set of heads and the name in the matrix is coded with the realization number "REAL 1". Similarly the Kx values are in a file called ***t2_kx.gst.*** If you view the contents of this file, you will see the following:

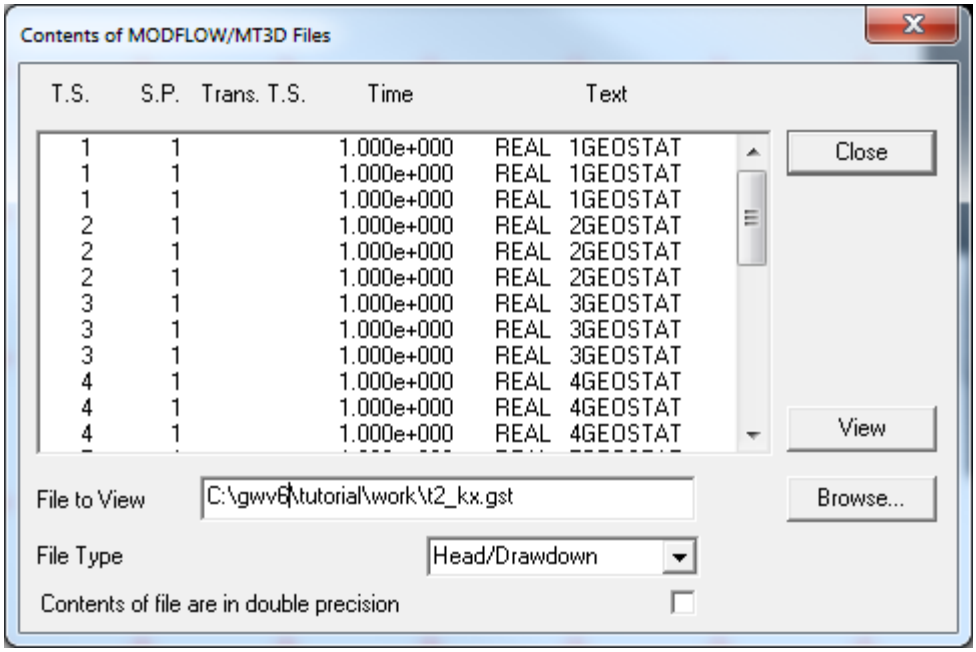

Note that this file also has the realization coded in the description of the matrix but also the time step number is the realization number. This allows you to import the results of a *.gst file to contour like it was a head file. We don't do this for the head file shown above because this information is processed in a different way.

We can now use some of the stochastic post-processing built into Groundwater Vistas to look at the 10 realizations we just combined into the one head file called ***t2_nsmc.hds***.

First, select **Plot|Stochastic|Import Target Data**. Browse to find the file ***t2_nsmc.hds*** and click OK. It should report 10 realizations were imported.

We can now create graphs for this information. One is a scatter plot of sum of squared residuals for each realization. Choose **Plot|Stochastic|Graph|Calibration Statistics**. The following dialog is displayed.

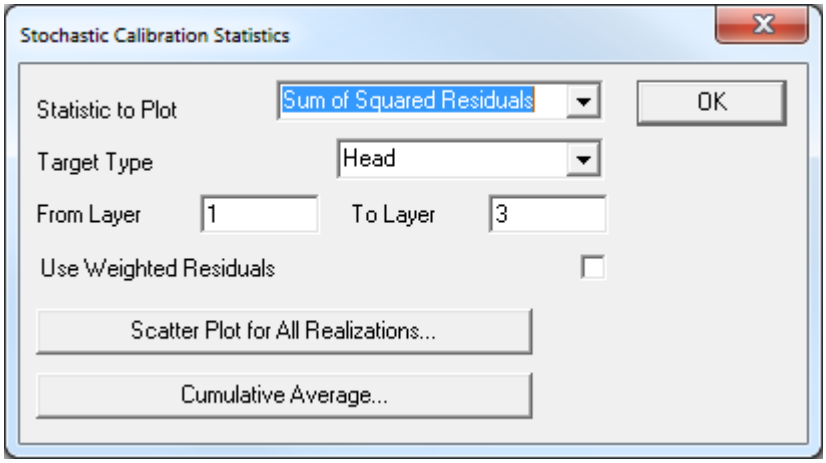

Click on the **Scatter Plot for All Realizations** button and the following graph should appear. This is just a visual summary of the text file you created earlier.

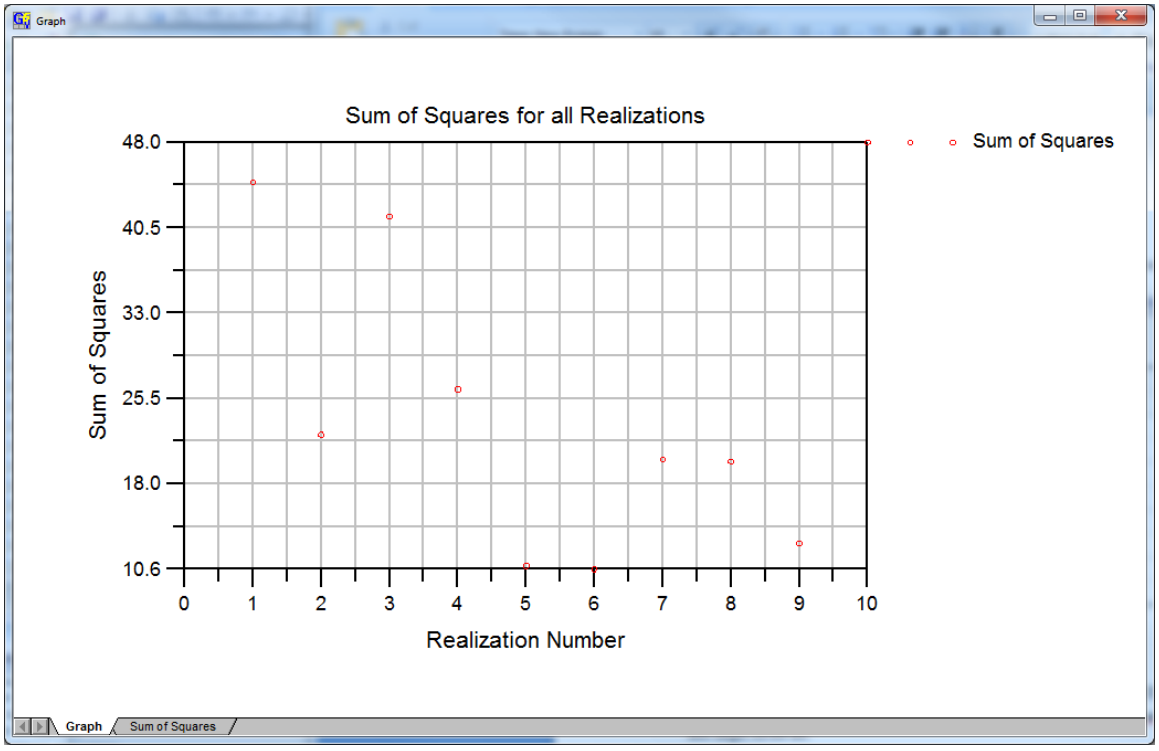

To see how heads vary at calibration targets, use **Plot|Stochastic|Graph|Target CDF** and just choose the first target on the list.

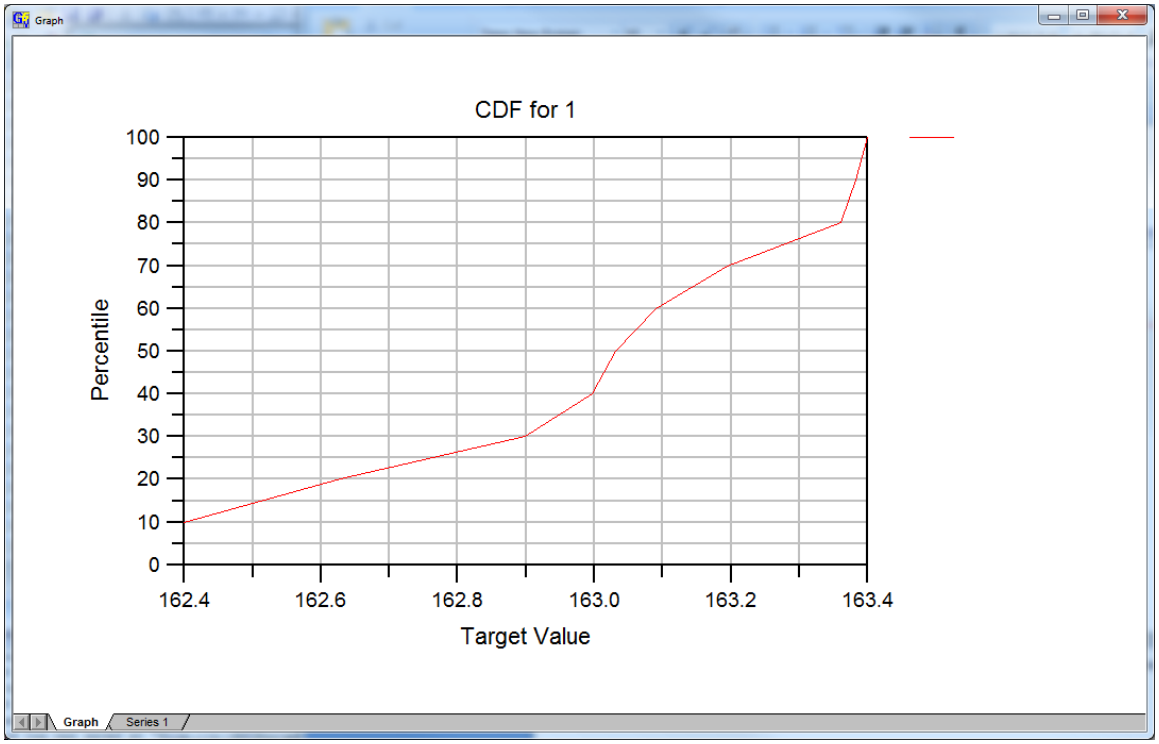

This graph shows that the range in head at this target is about a foot over all realizations.

Optimization

We will explore MODOFC and Brute Force in the following tutorial session. Both are included with GV Version 6.

MODOFC is probably the simplest of the classical optimizers to use. By *classical*, we mean that the optimization is done using linear programming techniques originally developed for management science. Classical optimizers include MODOFC, MODFLOW2000-GWM, SOMOS, MODMAN (developed for U.S. EPA by GeoTrans), AQMAN (old USGS model), etc. The nice thing about MODOFC (other than it's free!) is that it is totally self-contained. Many other optimizers, like MODMAN, require additional software. That adds cost and complexity to the whole process. With MODOFC, the optimizer is included within the MODFLOW software and runs as one application. MODFLOW2000-GWM is the successor to MODOFC and is also supported by Groundwater Vistas.

The MODOFC manual contains three simple examples to illustrate how to use the code for common applications. We have put those examples into Groundwater Vistas to get you started. Work through these examples to familiarize yourself with how MODOFC works.

Sample Problem I demonstrates a groundwater containment scenario in which head bounds, head difference constraints, and a capture zone are used. In this problem, wells are screened over multiple intervals and unconfined units. Sample Problem II demonstrates a construction dewatering project in which head bounds are used with construction costs in a confined aquifer. Binary variables and the branch and bound algorithm are used for this problem. Sample Problem III demonstrates a water supply problem in which drawdown and minimum total extraction is constrained. In this problem, the simulations use multiple stress periods.

MODOFC Problem I: CONTAIN

This sample problem represents a hazardous waste site in which contaminated groundwater must be contained. The sample problem has two layers, an upper unconfined unit above a lower confined unit with lengths of 600 feet and widths of 400 feet. Both layers have no flow boundary conditions along the top and bottom boundaries, fixed heads of 50 feet along the left boundary, and fixed heads of 45 feet along the right boundary. The layers are horizontally homogenous and isotropic with a vertical conductance between the two of 0.05 1/day. The upper unconfined unit has a hydraulic conductivity of 5 ft/day, a bottom elevation of 30 feet, and a uniform recharge of 0.002 ft/day. The lower confined unit has a transmissivity of 800 ft^2/day. Each layer is divided into 600 square block centered nodes each 20 ft by 20 ft. A production well is located in the lower unit at grid point (13,20,2) pumping at 10,000 ft^3/day. The grids for these layers are shown in the figures below.

It is assumed that an area in the upper layer of the model domain has been contaminated. From a design standpoint, at steady state, all contaminated water must be contained within a capture zone, while the production well continues to operate. Mounding above the ground surface cannot occur at injection wells (assumed at 65 feet), and the upper unit cannot be dewatered. Due to surface features at the site, there are only 6 possible well locations as shown in Figure 1 and well 5 is screened in both units. These design criteria can be converted into linear programming constraints, using head difference constraints and head constraints. First, a capture zone of three line segments has been defined around the contaminated area with a required gradient of 0.01 ft/ft. MODOFC will convert this into 23 head difference constraints. In addition, head difference constraints requiring a head difference of 0.01 ft have been defined in the source area to ensure upward vertical gradients and prevent contaminated water from leaking into the lower unit. A fixed head constraint has been used at the candidate injection well requiring that head remain below 65 feet. Finally, a minimum hydraulic thickness of 2 feet was allowed in the unconfined aquifer. It was assumed that extraction costs 5 times more than injection due to treatment costs.

This sample problem demonstrates several features of MODOFC. First, all possible head constraints are employed: head constraints, head difference constraints, and capture zones. Note the coordinates of the capture zone, which use the MODFLOW system of the origin at the upper left location of the grid, down is the positive Y-direction, and to the right is the positive X-direction. Because lower heads are required on the right side of the capture zone line, the first node is defined near Well 1. MODFLOW input file distances are used, for this case each cell is 20 ft by 20 ft. This problem also has an unconfined layer in it, so the solution will be determined through iteration, and care must be given to defining the first perturbation (PERTI), last perturbation (PERTF), perturbation scaling (PERTS), minimum aquifer thickness (BMIN), iteration convergence criteria (OCC), and whether iterations should be printed (ITPR). Also, repeated optimizations are not possible (ICOM = 1) and inclusion of binary variables is not recommended (INCST = 0) with unconfined problems. Because INCST = 0, values for installation costs (ICST), minimum pumping rates (PMN), and minimum (NWMIN) and maximum (NWMAX) number of wells to install are ignored. Wells 2 and 4 are screened over multiple units in this problem. Note the syntax for Wells 2 and 4, *i.e.* NL = 0 and extra rows are provided to set the pumping ratios for each screened layer.

Figure 1
Layer 1 of Sample Problem I

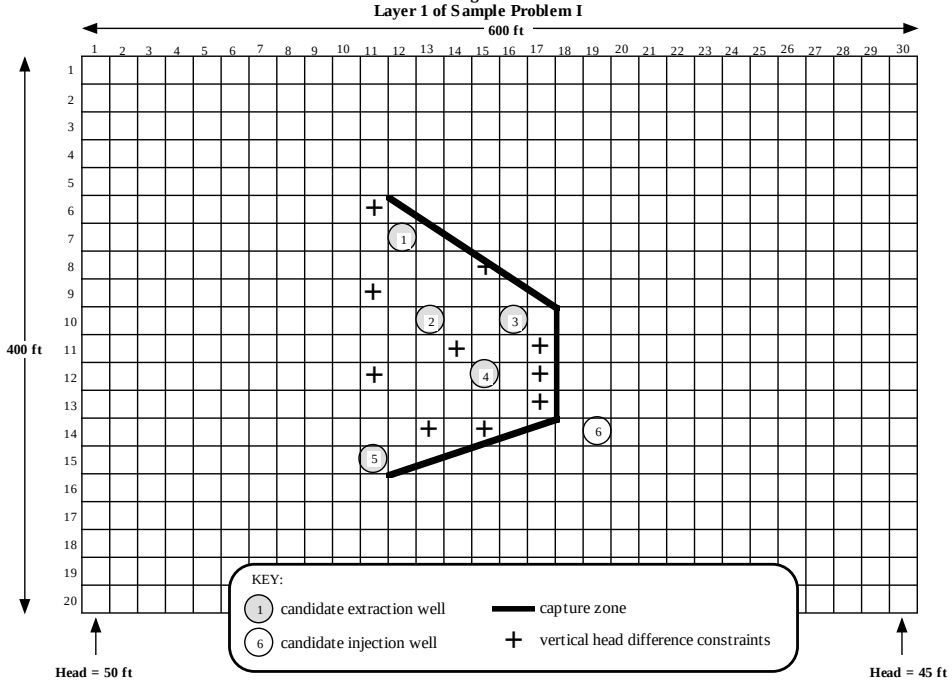

Figure 2
Layer 2 of Sample Problem I

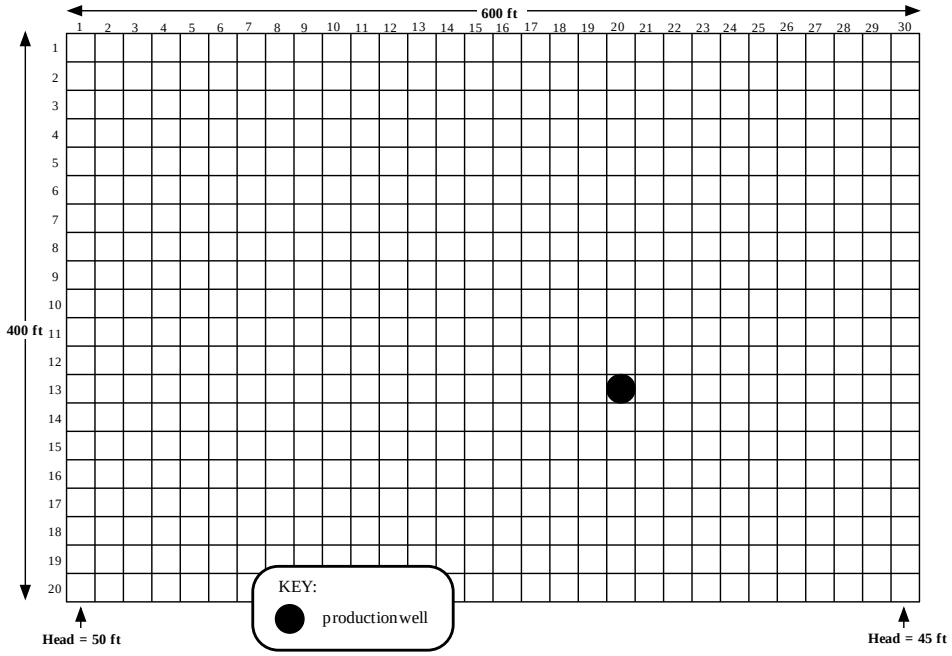

Run Groundwater Vistas Version 6 and open the file *modofc_contain.gwv* in the c:\gwv6\tutorial directory. Your screen should look like the one below:

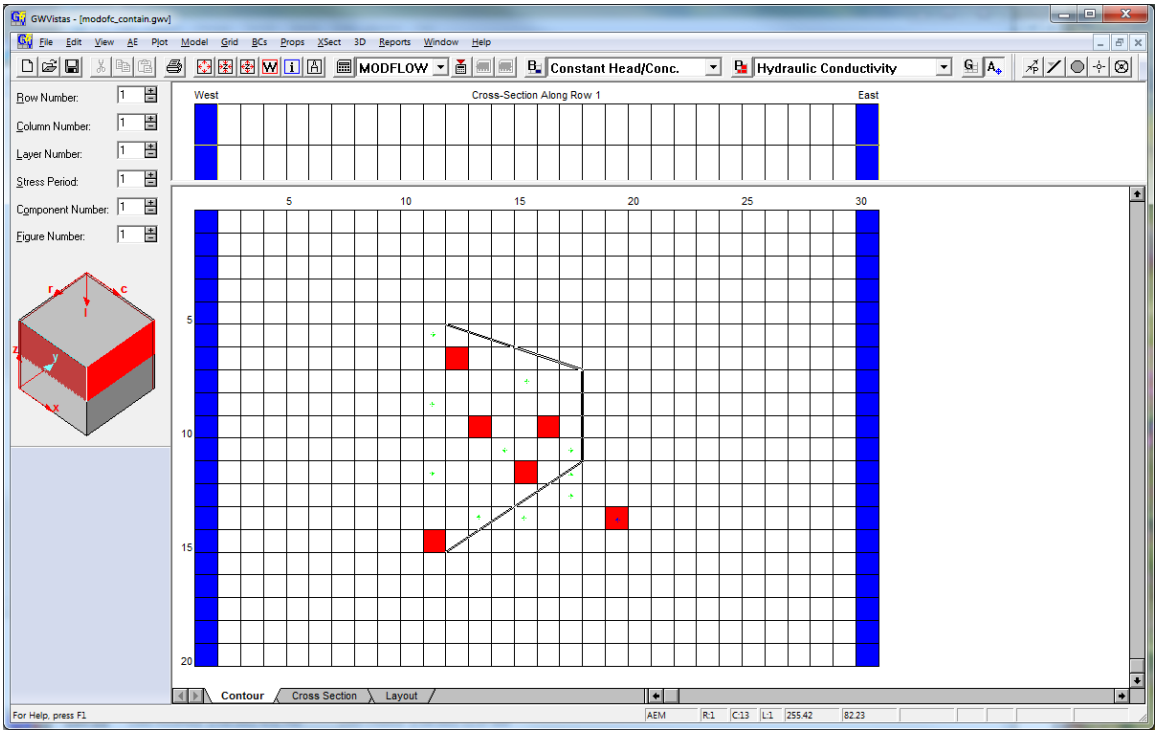

You will probably notice that the capture zone lines do not match the figure from the MODOFC manual. This is not an optical illusion, the data sets that come with MODOFC do not match the picture in the manual. The problem will only run with the capture zone as defined in the data files and as shown in your GV model.

The first thing you need to do when running MODOFC is to shut off the output control file. For some reason, MODOFC does not like the saving of binary head, drawdown, and cell-by-cell flow files. Select **Model|MODFLOW|Packages** and confirm that the output control file is off:

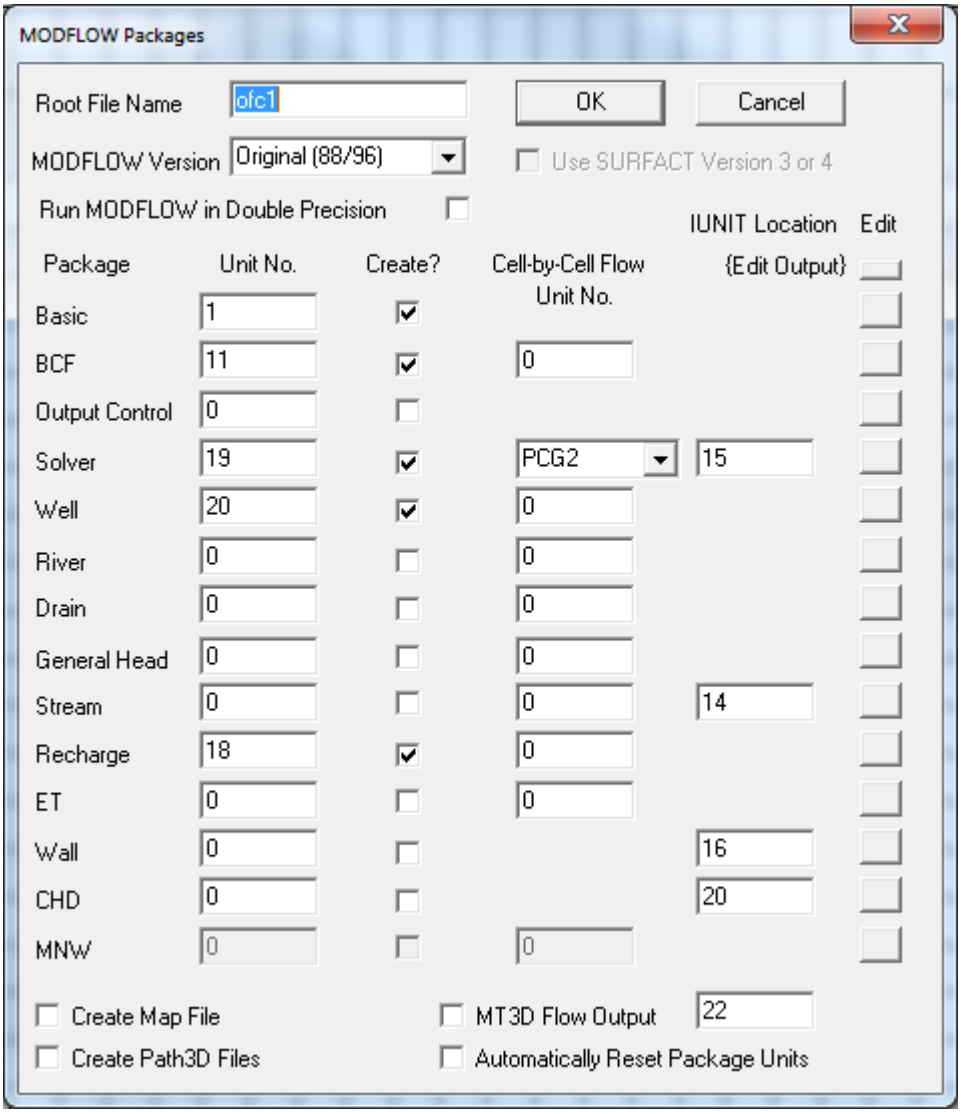

Also, make sure that all cell-by-cell flow unit numbers are zero, as shown above. MODOFC is a MODFLOW96 version, so also make sure that the MODFLOW version is set to *Original (88|96)*.

Now, double-click on one of the *targets* on the screen. Groundwater Vistas uses calibration targets as optimization constraints. The definition of the northern-most target is shown in the following dialog:

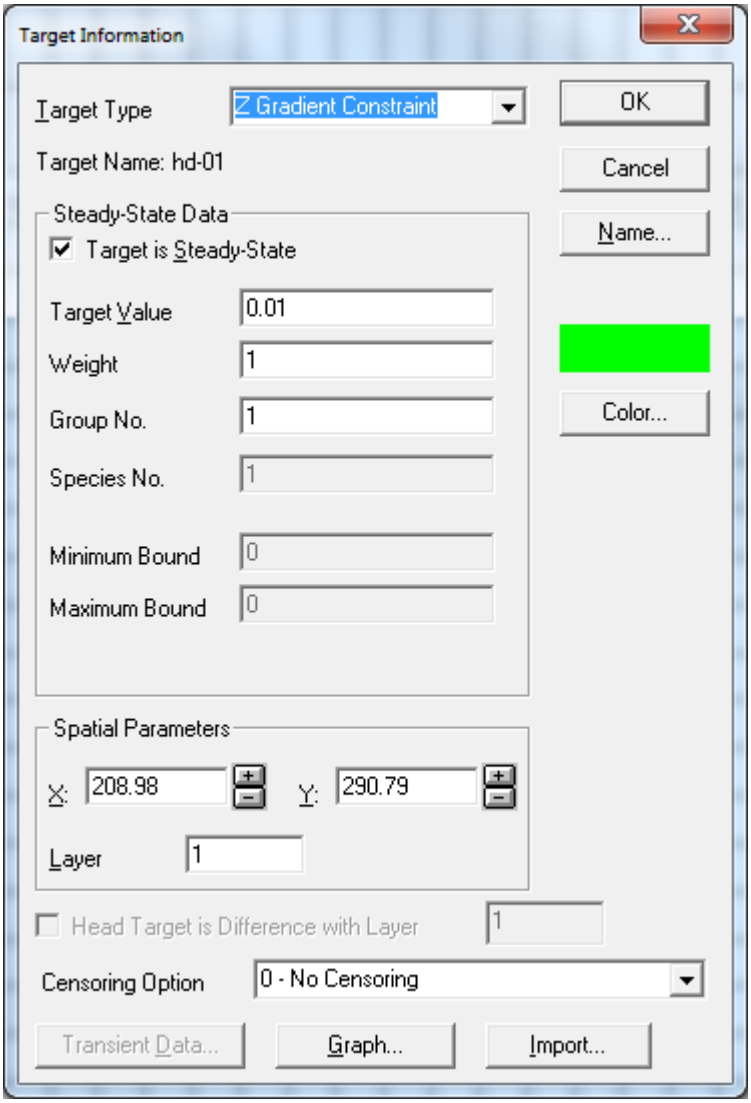

Notice that the target type is not head or drawdown but is changed to *Z Gradient Constraint*. If you pull down the target type list you will see there are several other types of constraints as well. The target value in this case is the constraint value for MODOFC. The weight is not used. One important thing to remember is that you MUST give the target a name.

The other type of constraint in this problem is the *capture zone*. MODOFC allows you to simplify the definition of constraints through the use of capture zone lines. Groundwater Vistas supports this option through the use of analytic line boundaries. Double-click on the southern-most line boundary and your dialog should look like the one below:

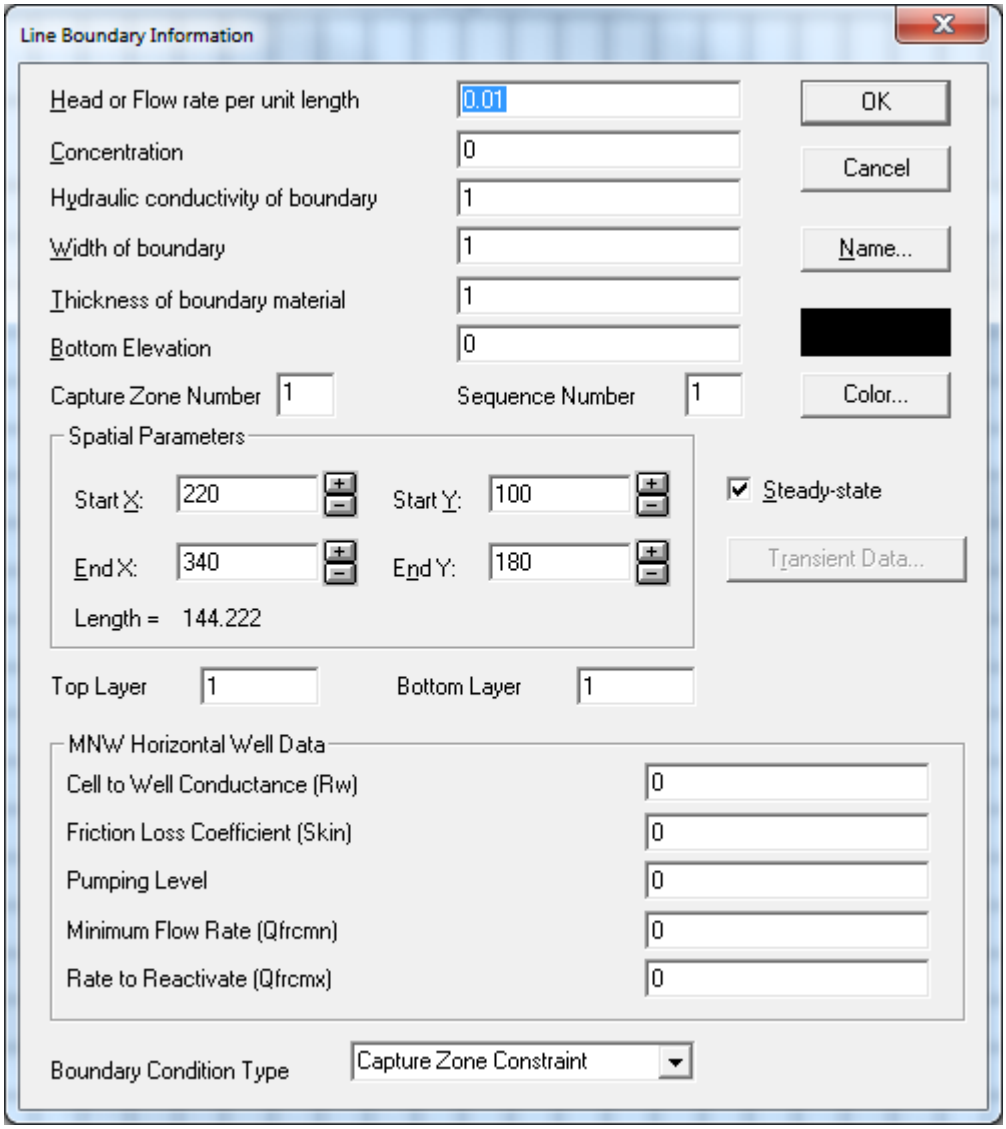

Note that the boundary condition type is *Capture Zone Constraint*. This type of line boundary is only used with MODOFC and is not supported by any of the other optimizing models in GV. There are several important things to remember when defining a capture zone for MODOFC. First, the area that you want to capture is on the left as you drag the line from its beginning point to its end point. You must also give the capture zone a number, called *Capture Zone Number*. Since a capture zone may have more than one line, you also give it a sequence number starting with 1. The one above is the first line for Capture Zone 1. You must start numbering capture zones at 1. The only other thing to remember is that the *Head or Flow rate* here represents a change in head used to construct head difference constraints by MODOFC.

Now look at the wells defined in GV. Select BCs|Well and then double-click on the well in the southwest corner of the capture zone. Your dialog should look like the following:

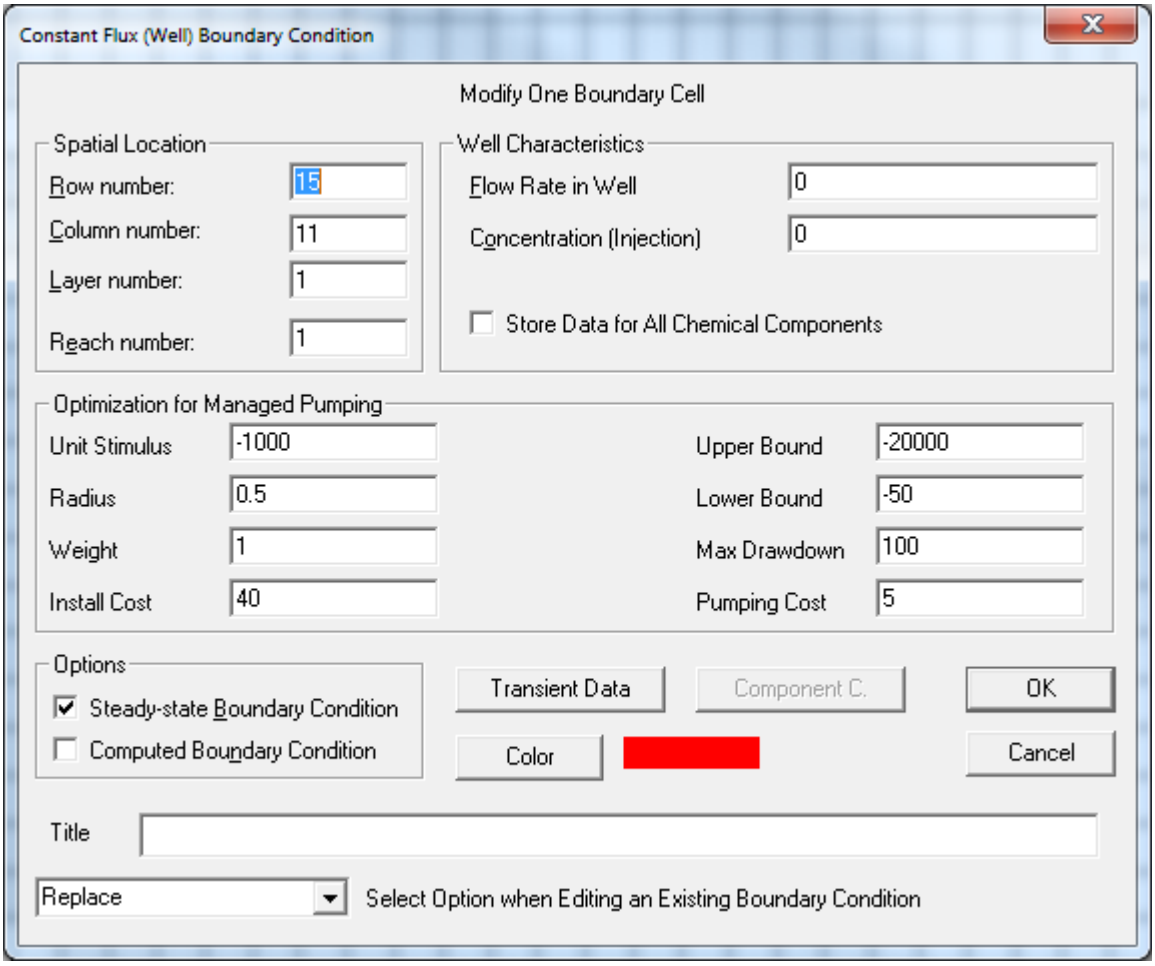

You may have wells in your model that are always on and those that are used in the optimization process. These are distinguished from one another by the reach number. You should use special reach numbers for those wells that are part of the optimization. In this example, all wells with reach 1 are part of the optimization. You must also enter a zero pumping rate for these wells. You will see a special section of the well dialog that is labeled *Optimization for Managed Pumping*. Not all of this information is used by MODOFC. The things that are used include the Upper Bound, Lower Bound, Install Cost, and Pumping Cost.

Now you have defined the problem except for MODOFC options. The model design has constraints and candidate wells. Fortunately, MODOFC has relatively few options. While this somewhat limits power users, it is nice for those of us who are not optimization experts! First, look at the general MODOFC options by selecting **Model|Optimization Models|MODOFC|Options - Main Options tab**:

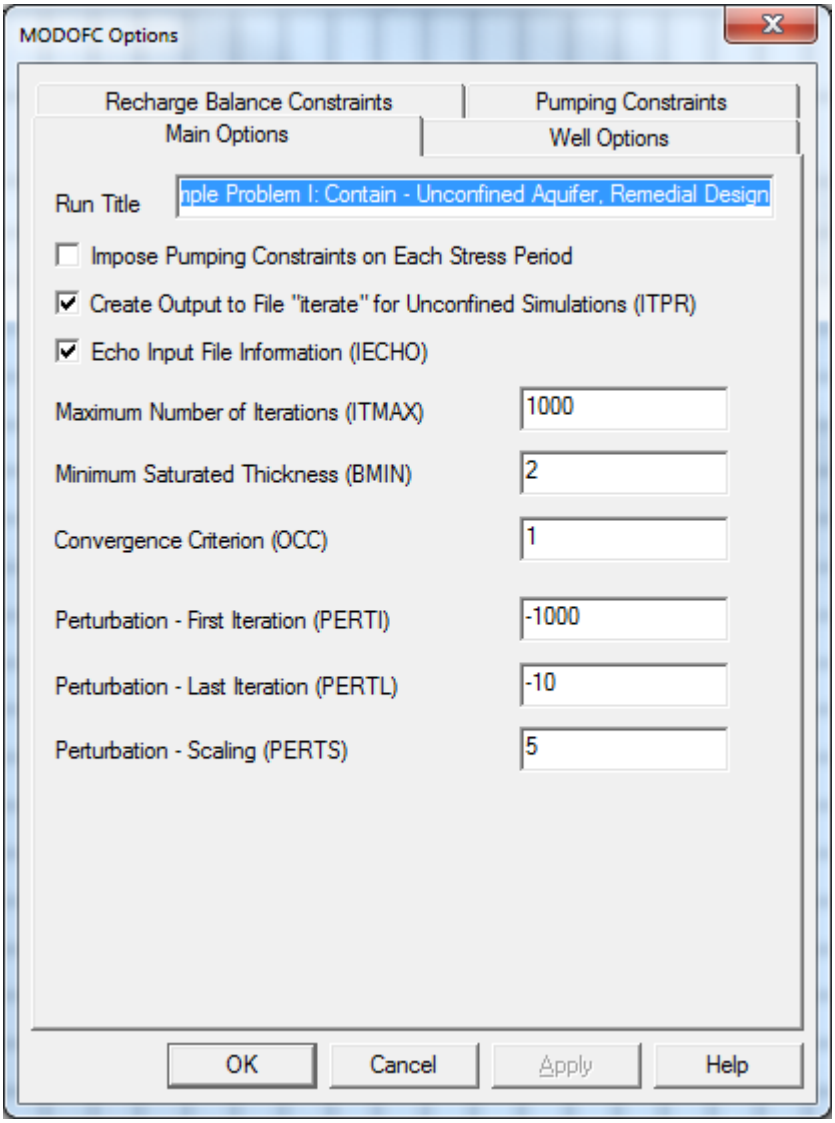

Probably the most important information on this dialog are the perturbation values. The first perturbation (PERTI) in unconfined aquifers is the pumping rate used in the first iteration. It should be about what you expect the pumping rate to be in the final solution. If you specify a value that is too large, the aquifer can be dewatered. If the value is too low, the aquifer response is not large enough for MODOFC to determine optimal rates. In confined systems, the last perturbation value (PERTL) and scaling factor (PERTS) are not used. In unconfined aquifers, iteration is required and these parameters determine how the perturbation factor is reduced during iteration. The author, David Ahlfeld, recommends that PERTL be 0.5 percent of PERTI and PERTS be set to 5. I have found that some experimentation if usually required for these values.

The other value to consider is the convergence criterion (OCC). For unconfined systems, OCC is the change in pumping rate that must occur for the iterations to stop. In this synthetic example, it is quite low. In practice, you may need to make this a much bigger number.

Now, select **Model|Optimization Models|MODOFC|Options** - Well Options tab and look at the following dialog:

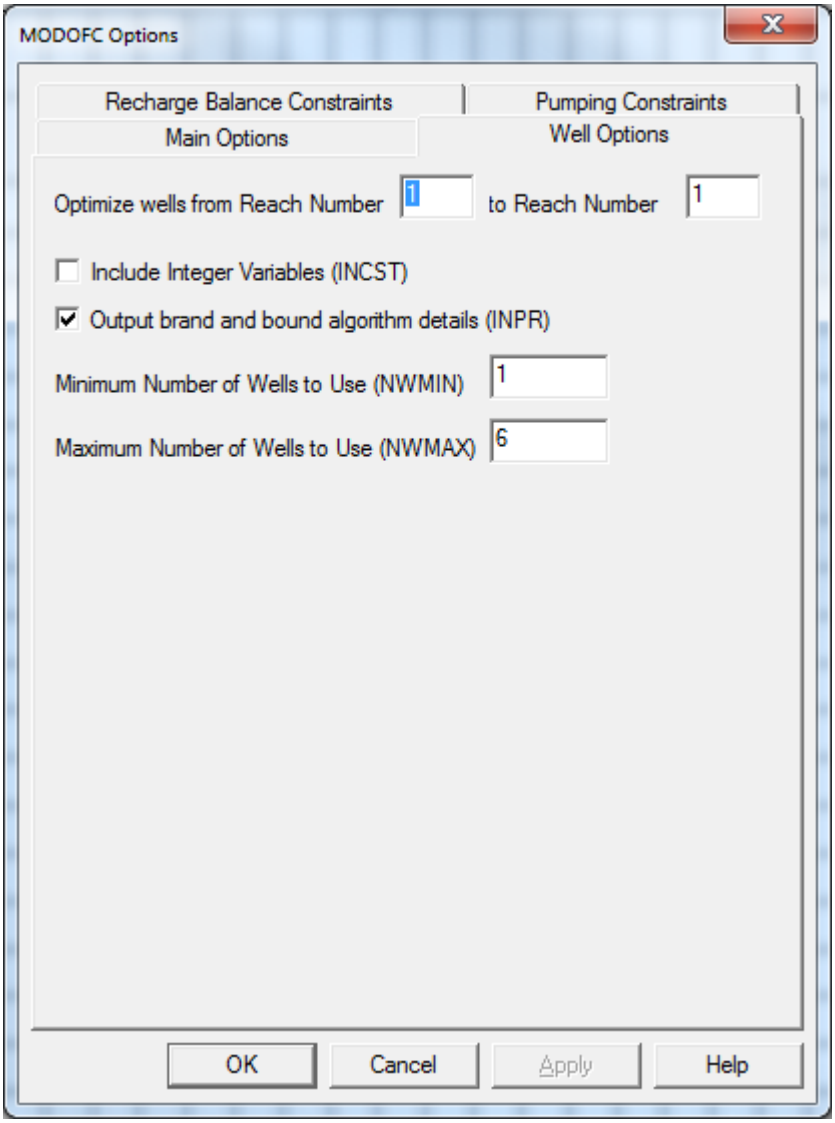

The first option on the dialog defines which wells in your model are going to be part of the optimization. In this case, only wells with a reach number of 1 are going to be used. The other important option is INCST, which includes or excludes integer variables. If you include integer variables, then MODOFC will constrain the solution using specified minimum pumping rates, well pumping|installation costs, and minimum|maximum number of wells. Without integer variables, these factors are not considered. Given that choice, why not always include integer variables? Unfortunately, using integer variables in unconfined aquifers often does not work. In this example, we will not include integer variables.

Now let's run MODOFC. The first step once you have defined constraints, candidate wells, and options, is to create MODFLOW files. Select **Model|MODFLOW|Create datasets**. Next, select **Model|Optimization Models|MODOFC|Create Input File**. MODOFC always looks for a file called *opt.in*, which is created when you select this menu item. If everything was set correctly, you should be able to choose **Model|Optimization Models|MODOFC|Run MODOFC**.

When it is finished, select **Model|Optimization Models|MODOFC|View Solution** to see the results. The first part of the file is the important stuff:

MODOFC VERSION 2.1 - SOLUTION OUTPUT FILE

MODOFC 2.1 Sample Problem I: Contain - Unconfined Aquifer, Remedial Desi

++

MODOFC Version 2.1 Optimization Results

++

Optimal Solution Found

++

PROBLEM SOLUTION

++

Objective Function Value 2.2435E+04

Pumping Rates Listed For Each Well

Name	Stress periods	Extraction	Injection
well1	1	1.5250E+03	
well2	1	1.2747E+03	
well3	1	2.7335E+02	
well4	1	8.0558E+02	
well5	1		5.3112E+02
well6	1	7.7213E+01	
Total Rates		3.9558E+03	5.3112E+02

++

The results are almost identical to the example output that is included with MODOFC. Note that all wells are needed, even though wells 1 and 2 pump most of the water. This happens because we could not include integer variables to throw out wells unless they have a certain minimum pumping rate.

Try turning on the integer variables and see what happens (Model|Optimization Models|MODOFC|Well Options). You will probably get the same answer. That is because our minimum flow rate for the wells was 50 ft^3|d. Edit all of the wells in layer 1 except the far eastern one which is the injection well. Change

the minimum flow rate from –50 to –500. Now run it. In this case, you probably get the common error *Problem infeasible*.

MODOFC Problem II: DEWATER

This sample problem represents a dewatering problem for a construction site. The sample problem has one confined layer with a length of 3,000 feet and a width of 2,000 feet. The layer has boundary conditions of no flow along the top and bottom boundaries, fixed head of 80 feet along the left boundary, and fixed head of 60 feet along the right boundary. The layer is homogenous and isotropic with a transmissivity of 50 ft^2|day. The aquifer is divided into 600 square block centered nodes each 100 ft by 100 ft. The grid is shown in Figure 3.

In this problem, the groundwater needs to be lowered to a depth of 50 feet so that footings can be installed in the region shown on Figure 3. This design criteria can be imposed by defining 10 upper bound head constraints as shown in Figure 3. Seven candidate well locations are selected with minimum pumping rates of 100 ft^3|day, maximum pumping rates of 10,000 ft^3|day, pumping costs for the construction period of 10 $ day|ft^3, and construction costs of $2000 per well. To improve reliability, at least three wells must be installed.

This sample problem primarily demonstrates the use of binary variables. To include the cost for well installation, binary variables must be used, thus INCST = 1. This enables several other constraints, so meaningful values must be provided for installation costs (ICST), minimum pumping rates (PMN), and minimum (NWMIN) and maximum (NWMAX) number of wells to install. Also, if the user wants to record branch and bound algorithm solution steps, INPR must be set to 1.

Figure 3
Sample Problem II

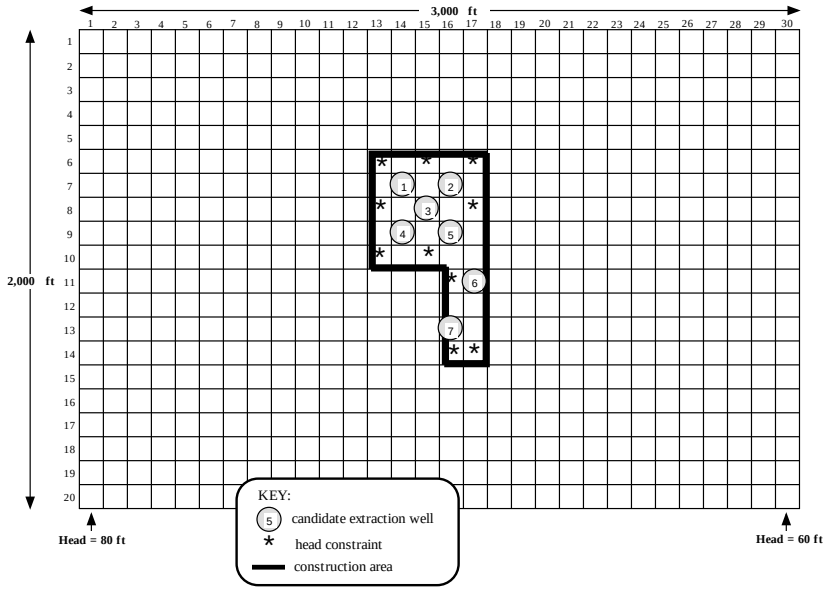

We have constructed this example as GV file *modofc_dewater.gwv*. Open up this design and it should look like the following:

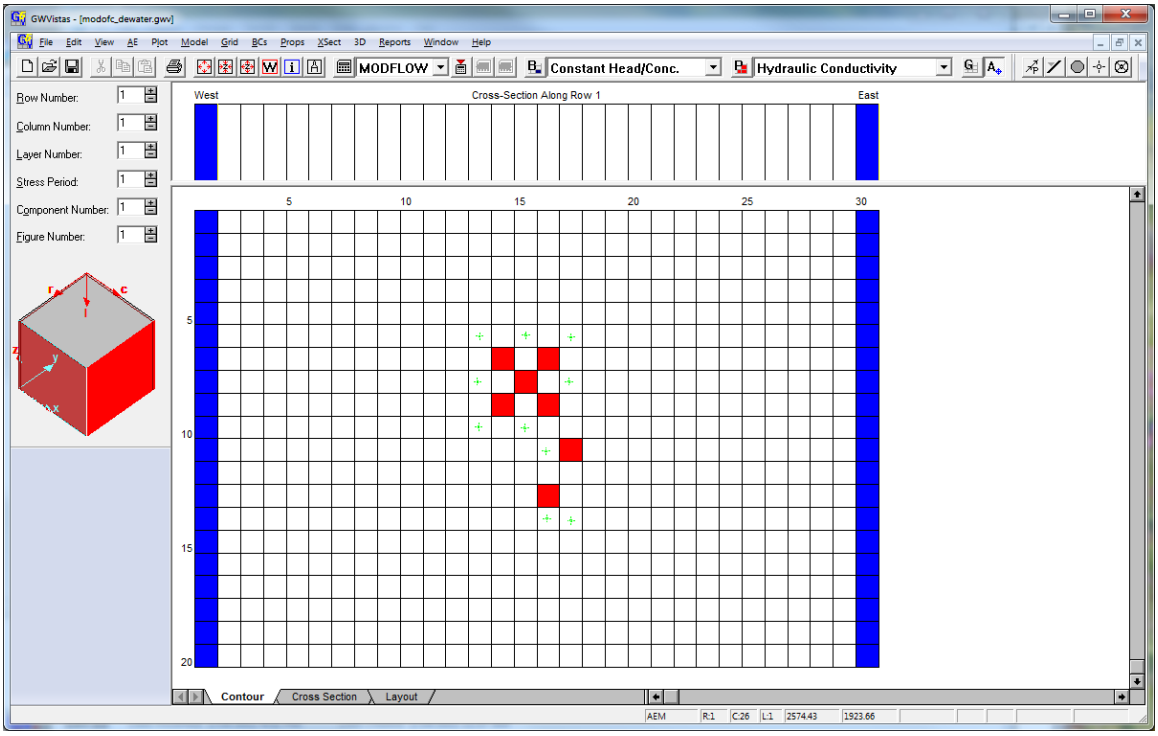

The main difference in this example is in the constraints. Instead of requiring an inward and upward gradient, we want a certain amount of drawdown for dewatering. Edit one of the targets (constraints) and you will see that the target type is *Max Head Constraint.*

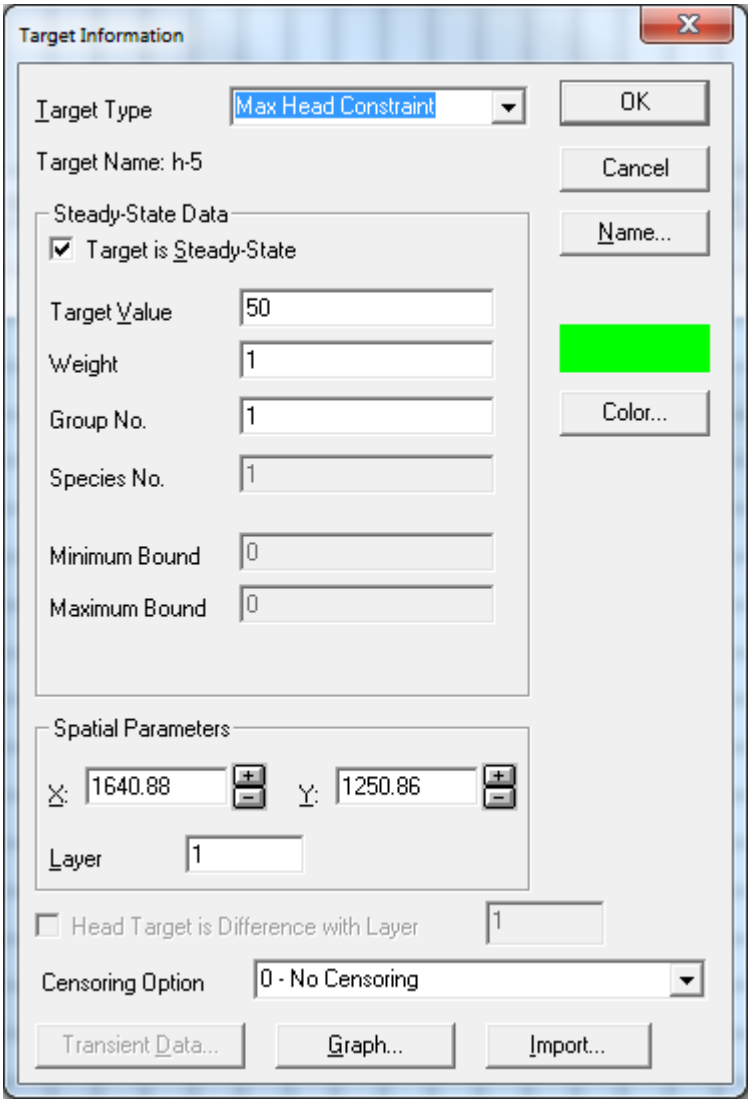

This means that the wells will be optimized so that the water level is at or below this value. The head value is defined as the target value, 50 in this case.

The wells for this example are similar to the previous one except that we are now entering cost information.

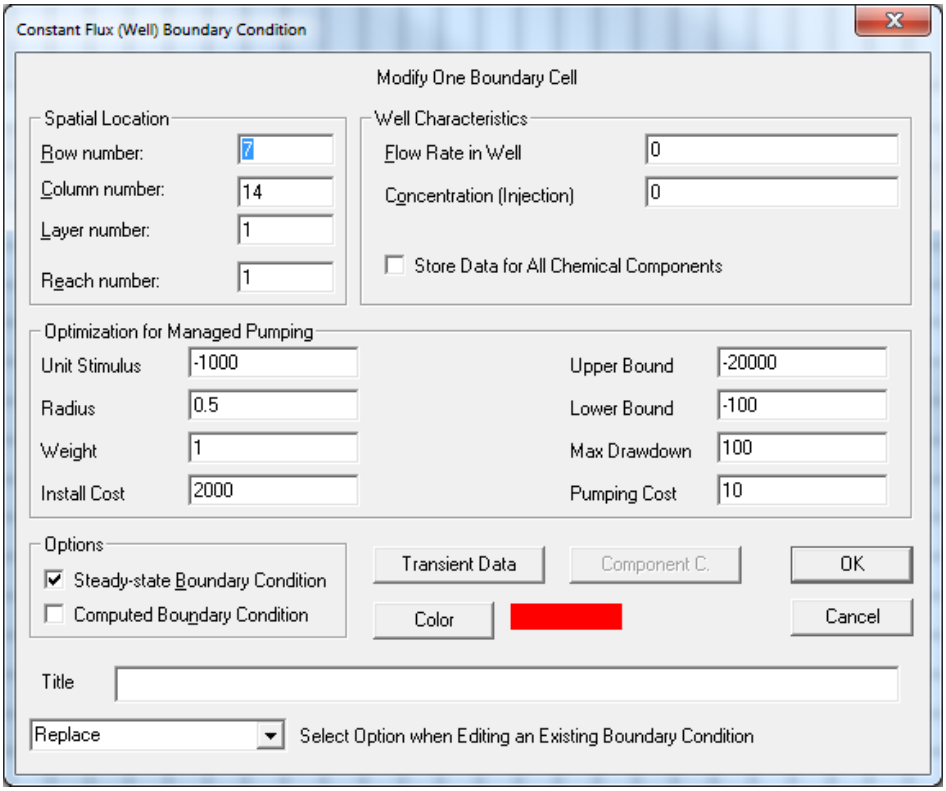

Go ahead and run this example. Your answer should be the same as the following:

MODOFC VERSION 2.1 - SOLUTION OUTPUT FILE

MODOFC Problem

++

MODOFC Version 2.1 Optimization Results

++

Optimal Solution Found

++

PROBLEM SOLUTION

++

Objective Function Value 3.4799E+04

Pumping Rates Listed For Each Well

Name	Stress periods	Extraction	Injection
well1	1	1.2425E+03	
well2	1	0.0000E+00	
well3	1	0.0000E+00	
well4	1	6.9412E+02	
well5	1	0.0000E+00	
well6	1	0.0000E+00	
well7	1	9.4327E+02	
Total Rates		2.8799E+03	0.0000E+00

++

Try turning off the use of integer variables (Model|MODOFC|Well Options) and see how that effects the answer.

MODOFC Problem III: SUPPLY

This sample problem represents a water supply problem in which limited drawdown is allowed near a river. The sample problem has one confined layer, 6,000 feet by 5,000 feet, with a river running through the center of it. The layer has boundary conditions of no flow along the left and right boundaries, fixed heads along the top boundary which vary from 100 feet to 86 feet, and fixed heads along the bottom boundary which vary from 80 feet to 66 feet. The layer is homogenous and isotropic with a transmissivity of 5000 ft^2|day and a storativity of 0.05. The aquifer is divided into 750 square block centered nodes each 200 ft by 200 ft. The stream flows from left to right with summer heads which vary from 50 feet to 32 feet along its length. The fall and winter heads are 1 foot higher and the spring heads are 2 feet higher. During summer the average stream depth is 3 feet. The stream is approximately 10 feet wide with a bed conductance of 500 ft^2|day. Recharge also varies throughout the year: 0.0005 ft|yr in winter, 0.002 ft|yr in spring, 0 ft|yr in summer, and 0.001 ft|yr in fall. The grid and stream are shown in Figure 4.

In this problem, the objective is to maximum groundwater extraction over a three-year period (12 seasons) while maintaining a certain level of water in the stream. This design criteria can be imposed by defining upper head constraints in the groundwater at locations under the stream, as shown in Figure 4. In addition, there are limits on the maximum (80,000 ft^3|day) and minimum (30,000 ft^3|day in the first year and 25,000 ft^3|day in the second and third year) total rates of water extracted in a three-month period. These values define the total pumping constraints. Four possible well locations are shown in Figure 4 with maximum pumping rates at each well of 50,000 ft^3|day and pumping costs of -0.001 $ |ft^3|day (i.e., each 1000 cubic foot of water pumped is worth $1). The wells and constraints are only active for various stress periods as shown in the opt.in file.

This sample problem primarily demonstrates optimization of groundwater problems with multiple stress periods. Because the recharge and streamflows vary with the season, a transient groundwater model is required to properly depict the site. It may be advantageous to have wells pump at different rates for each season. This is exhibited in Well 2 which is required to have one pump rate for winter (stress periods 1, 5, and 9), another for spring (stress periods 2, 6, and 10), another for summer (stress periods 3, 7, and 11) and another for fall (stress periods 4, 8, and 12). Well 1 on the other hand is required to have the same pump rate for the entire 3 years (stress periods 1-12). Well 3 will only be pumping in year 2 (stress periods 5-8) and Well 4 can only pump in spring (stress periods 2, 6, and 10) and fall (stress periods 4, 8, and 12). Similarly, constraints are only active for certain stress periods: b-01 and b-03 for the entire 3 years (stress periods 1-12) and b-03 and b-04 for summer (stress periods 3, 7, and 11) and fall (stress periods 4, 8, and 12).

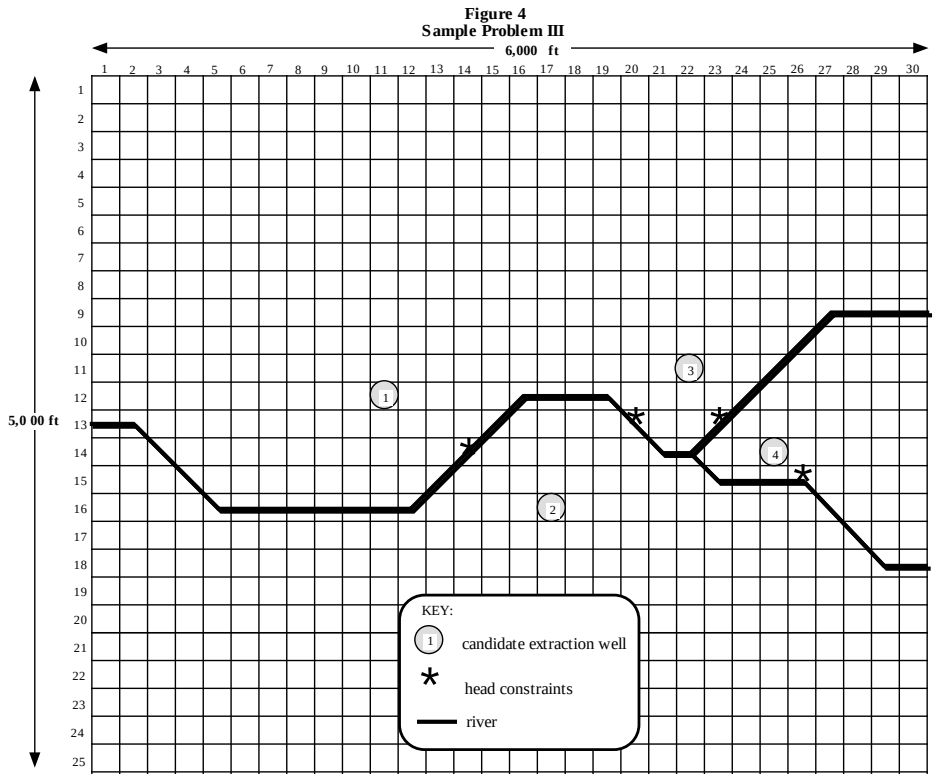

This example illustrates some very complex use of when wells can pump during a transient run. Groundwater Vistas right now cannot handle this degree of complexity, so you will need to edit the *opt.in* file in a text editor to complete it. First, load *modofc_supply.gwv* into Groundwater Vistas. Create the MODFLOW files and then the MODOFC file. Find the file c:\gwv6\tutorial\work*opt.in* and edit it in Notepad, Wordpad, or something else. Look for the following lines in the file:

well1 t 11 22 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '1,12'

well2 t 12 11 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '1,12'

well3 t 14 25 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '1,12'

well4 t 16 17 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '1,12'

And change them to:

well1 t 11 22 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '5-8'

well2 t 12 11 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '1-12'

well3 t 14 25 1 1 e 100.000000 50000.000000 0.000000 2000.000000 '2,6,10:4,8,12'

well4 t 16 17 1 1 e 100.000000 50000.000000 0.000000 2000.00 '1,5,9:2,6,10:3,7,11:4,8,12'

Likewise, find the *Bounds on Head* section and change the following:

h-1 t 14 14 1 > 67.000000 '1-12'

h-2 t 13 20 1 > 64.000000 '1-12'

h-3 t 13 23 1 > 62.000000 '1-12'

h-4 t 15 26 1 > 60.000000 '1-12'

To:

h-1 t 14 14 1 > 67.000000 '1-12'

h-2 t 13 20 1 > 64.000000 '3,4,7,8,11,12'

h-3 t 13 23 1 > 62.000000 '1-12'

h-4 t 15 26 1 > 60.000000 '3,4,7,8,11,12'

Save these changes and run MODOFC.

Using Brute Force Optimization

The Brute Force optimization technique is a new way of optimizing pump & treat systems. Unlike more MODOFC, Brute Force uses particle tracking to determine whether a system is successful. Other optimizing software uses head, drawdown, or gradient constraints to design the system as we saw with our MODOFC examples. The problem with these more classical techniques is that defining the gradient constraints is difficult and often the optimization is not successful. The benefit of using particle tracking is that the particles are either captured or they are not. The Brute Force routine continues until all particles, or a specified percentage of particles, are captured in the plume area. This makes interpreting the results very straightforward. The disadvantage of Brute Force is that it requires more simulations than many of the other optimizers.

We have created a simple BF example that you can compare with MODOFC. Open the file *modofc_contain_noInject2_brute.gwv*. This is a modified version of the first MODOFC example problem with one layer and no injection. The model looks like the following:

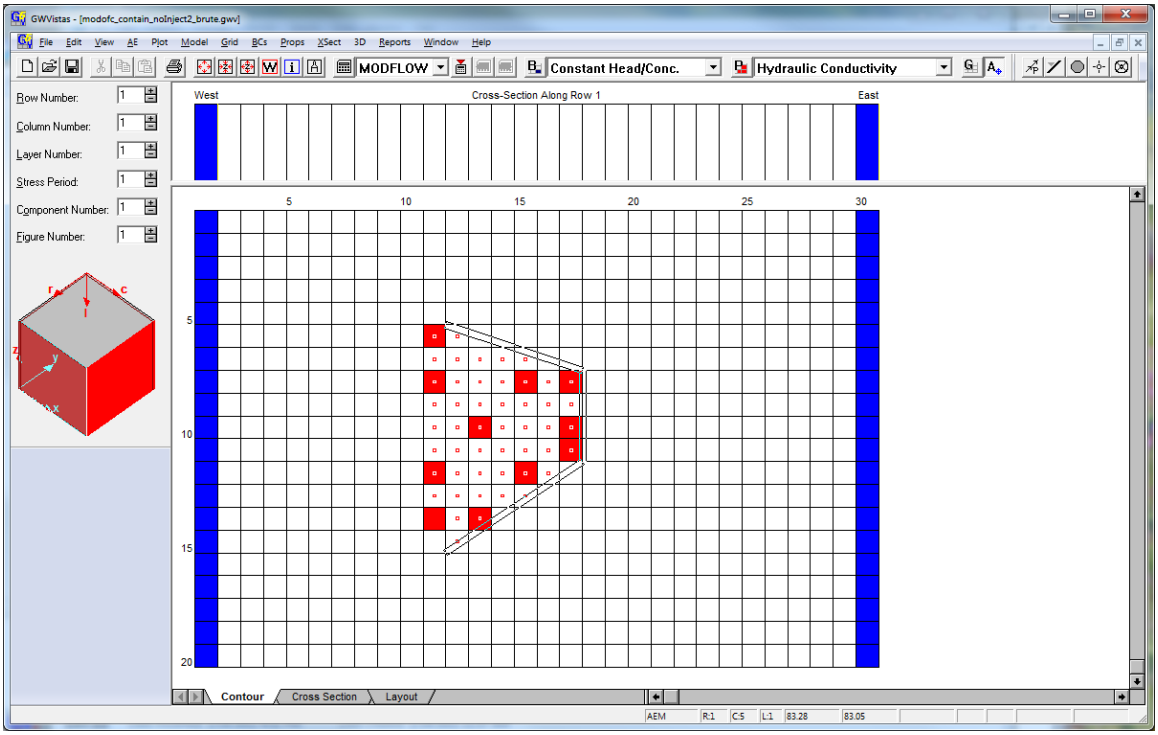

You can see the lines defining the capture zone and potential wells that are similar to the first MODOFC example. You also see a bunch of particles within the capture zone. These are for use with BF. The first thing to do is run one MODFLOW and one MODPATH simulation. These files form the basis for the next BF run. After running MODFLOW and MODPATH, your screen should look like the following:

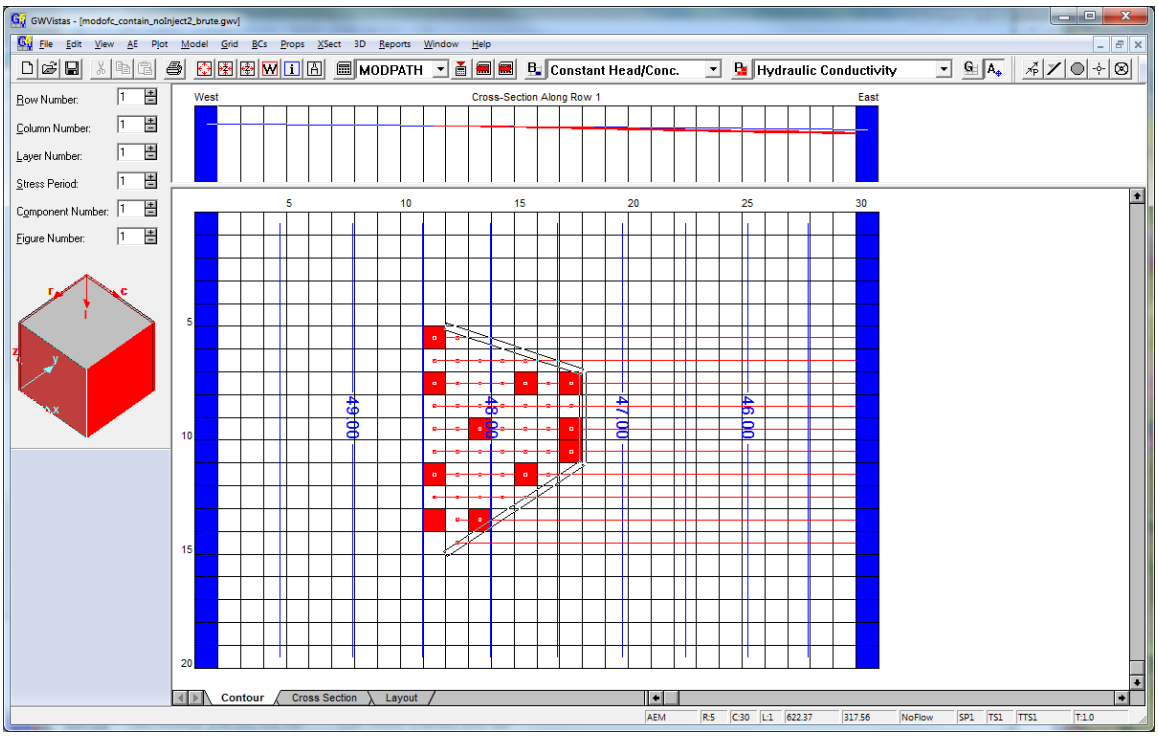

We are starting with no pumping so the particles just go east to the constant head boundary. Now, let's look at the Brute Force options. Select **Model|Optimization Models|Brute Force|Options**. The dialog is shown below:

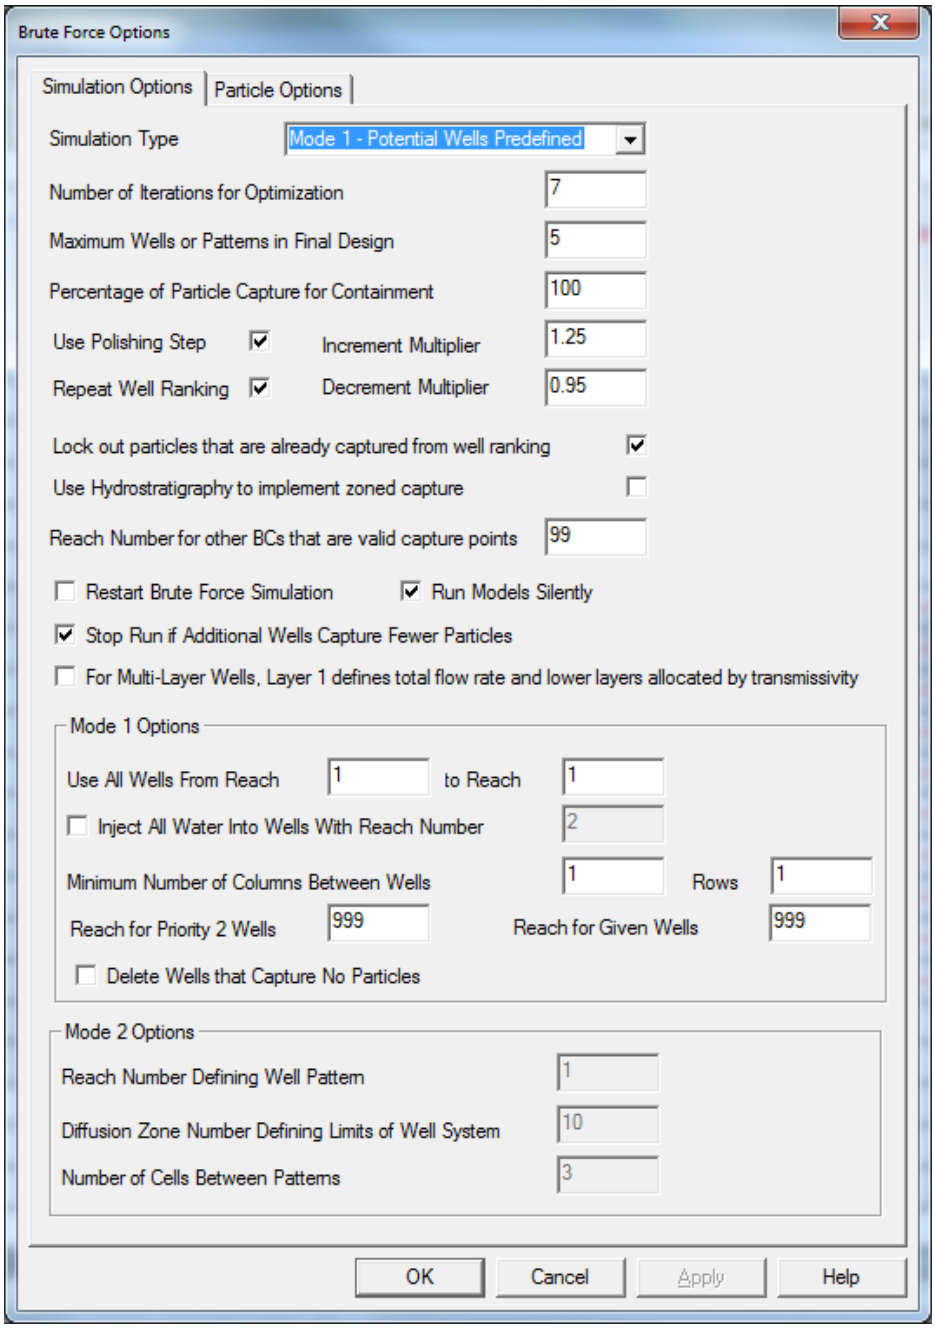

Probably the most critical option for BF is the number of iterations for optimization (7 in this case) and the increment multiplier (1.25). Also important is the *unit stimulus* for the pumping wells, which is –1,000 $ft^3|d$ in this case. BF will use the –1,000 pumping rate to rank the wells so it is important that this rate be large enough to actually capture particles. BF will then select the best well and run 7 simulations multiplying the pumping rate (-1,000 initially) by 1.25 until it hits either 7 simulations or the maximum pumping rate defined for that well (-20,000 in this example). You need to make sure that BF can test the whole range from unit stimulus to max. pumping rate during the number of iterations specified.

Now, select **Model|Optimization Models|Brute Force|Create Input Files**. GV will create the main BF input file (brutefrc.dat) and a new well file and particle file. To run BF, select **Model|Optimization Models|Brute Force|Run Brute Force**.

This run goes pretty fast. Select **Model|Optimization Models|Brute Force|View Output File** and you should see at the end of the file that BF came up with 1 well (row 10, column 17) pumping –2,441 ft^3|d to contain this plume. Plug this pumping rate in and rerun MODFLOW|MODPATH to confirm it. Your results should be similar to those below:

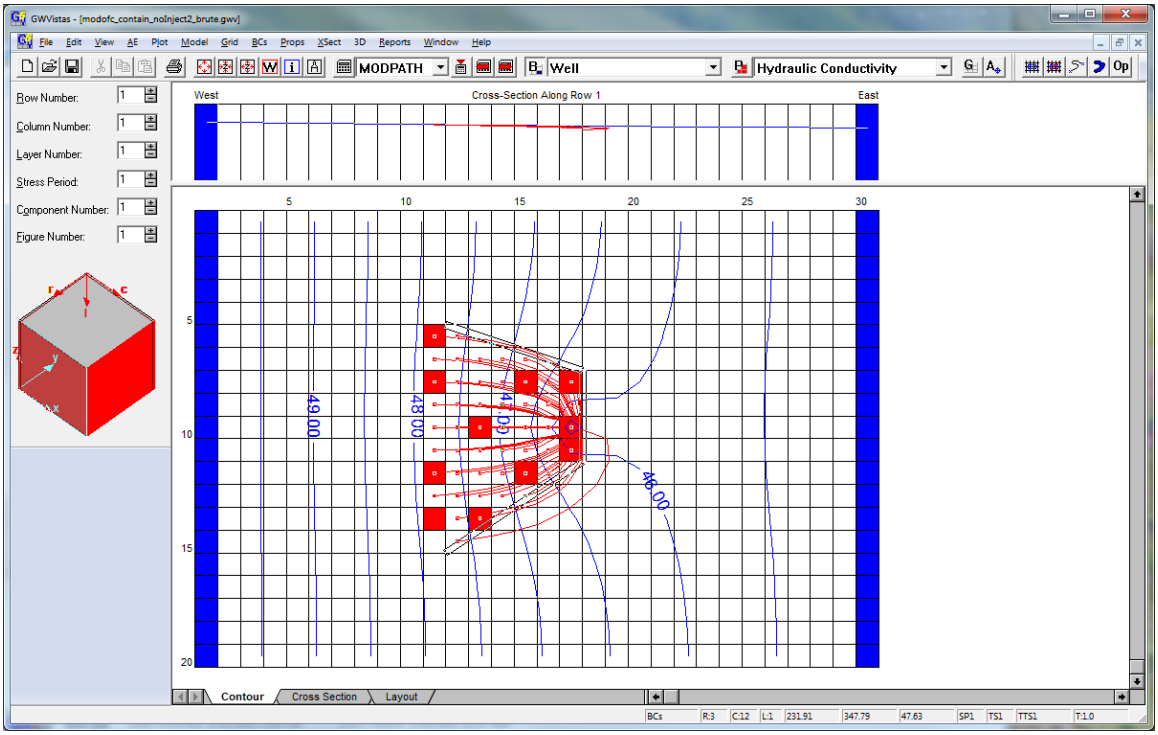

Try running the MODOFC version of this model – *modofc_contain_noInject2.gwv*. How do the results compare?

Using MGO with Groundwater Vistas

MGO is set up just like MODOFC. You start by defining the locations of pumping|injection wells that will be optimized and the constraints for the management problem. An example is presented from the MGO manual that is similar to the ones presented in the preceeding sections for MODOFC and Brute Force. Launch Groundwater Vistas and load the example called *mgo_ex1.gwv*.

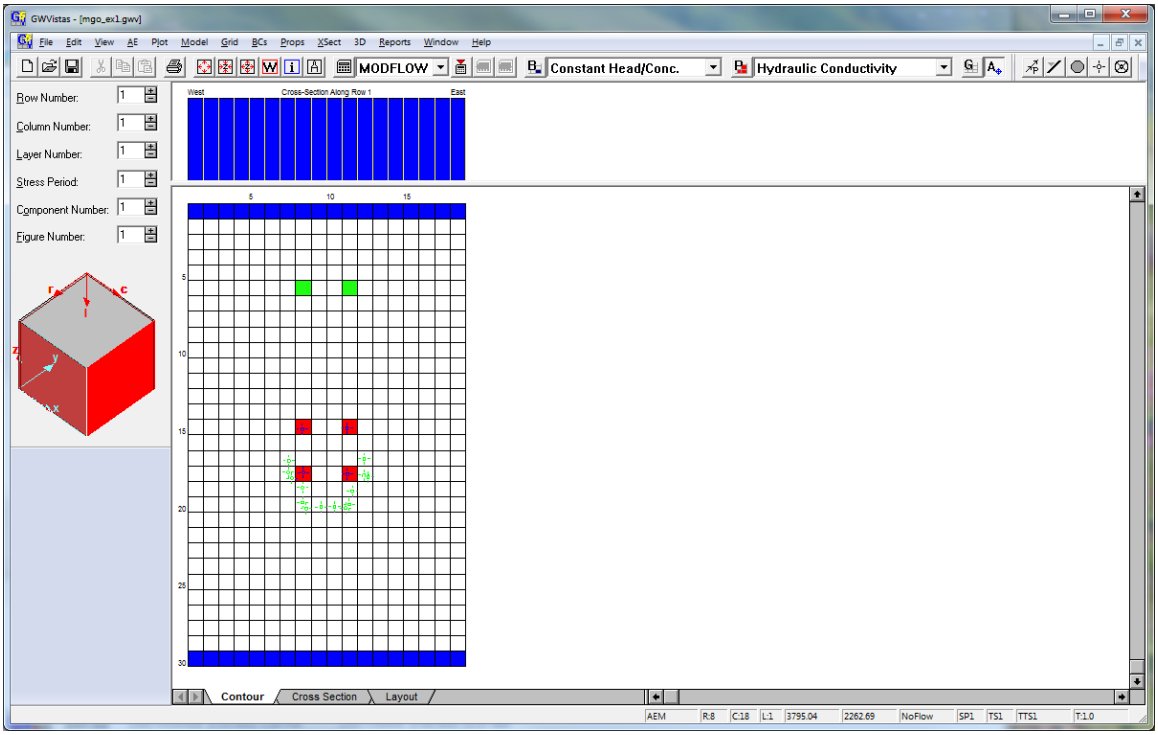

The constraints are lines up in a “U” shape south of 4 red pumping wells. Two green injection well sites are to the north of the pumping wells. The goal of this simulation is to determine how much water to pump from the red wells in order to contain the area outlined by the X and Y gradient constraints.

As with MODOFC, there is an MGO selection on the Model menu in Groundwater Vistas. Select **Model|Optimization Models|MGO|Options** now. The first tab is for Basic Options and primarily selects the type of optimization technique. The default method is called genetic algorithm.

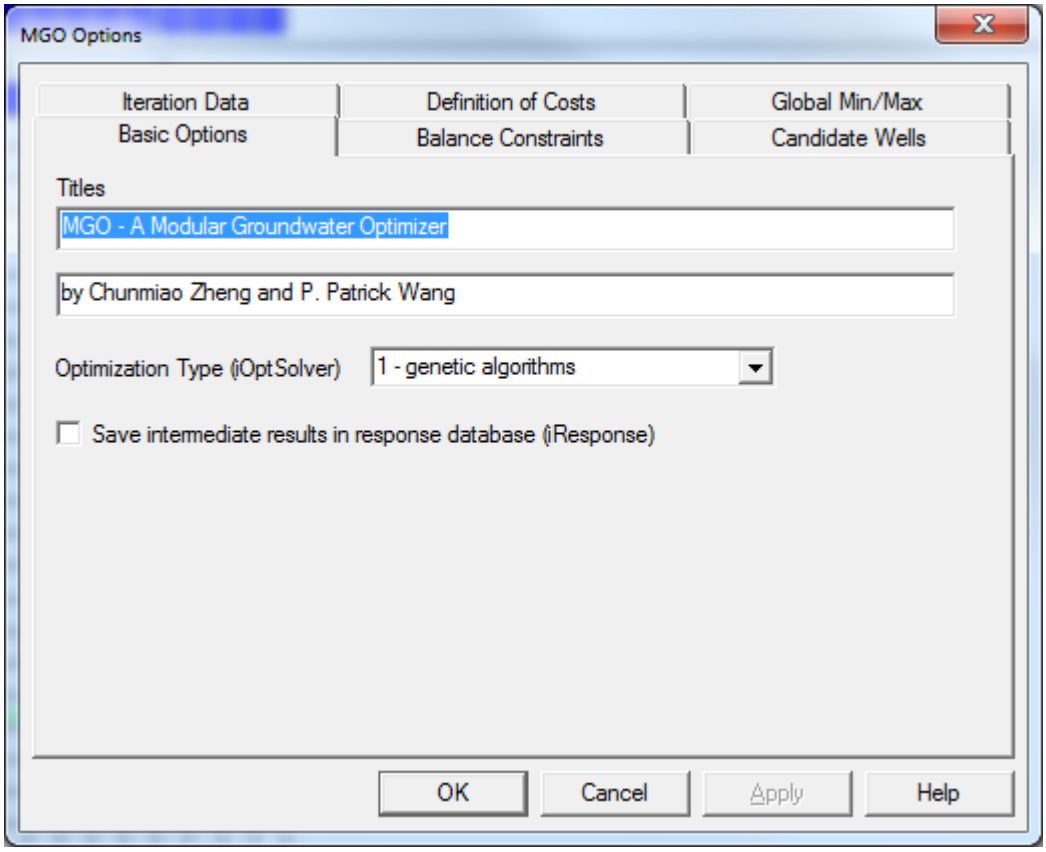

Click on the Balance Constraints tab next. This tells MGO that all of the water pumped from the 4 red wells must be injected into the two green wells sites. These sites are defined by a reach number of 5 and 6.

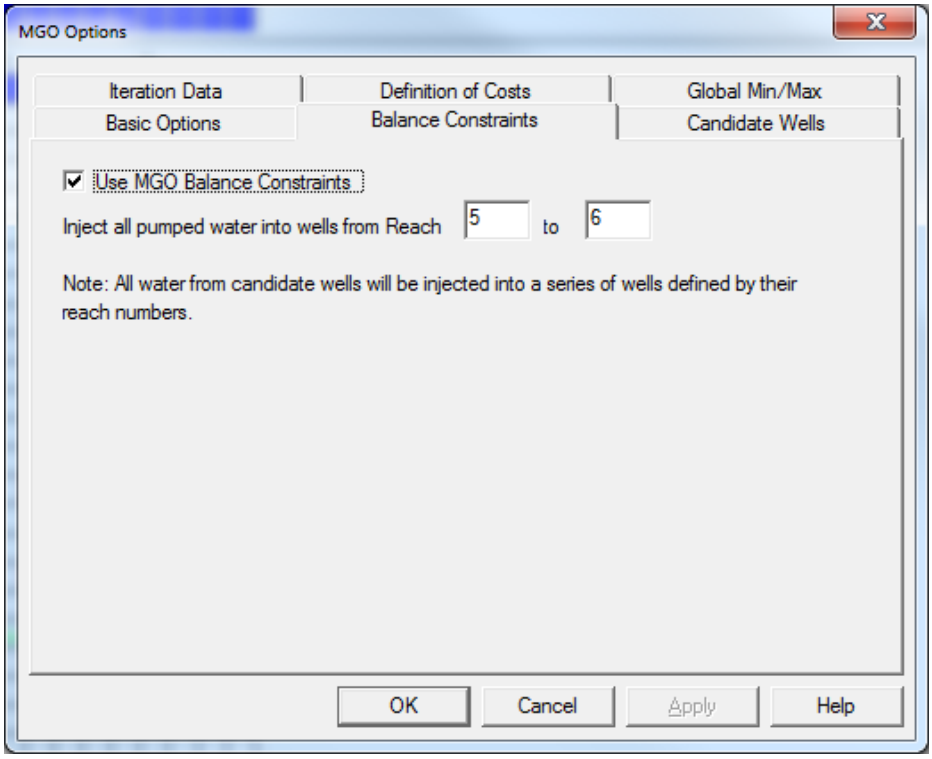

Now click on the Candidate Wells tab. This determines which wells are going to have their rates optimized. In this case, wells with reach numbers between 1 and 4 will be the candidate wells.

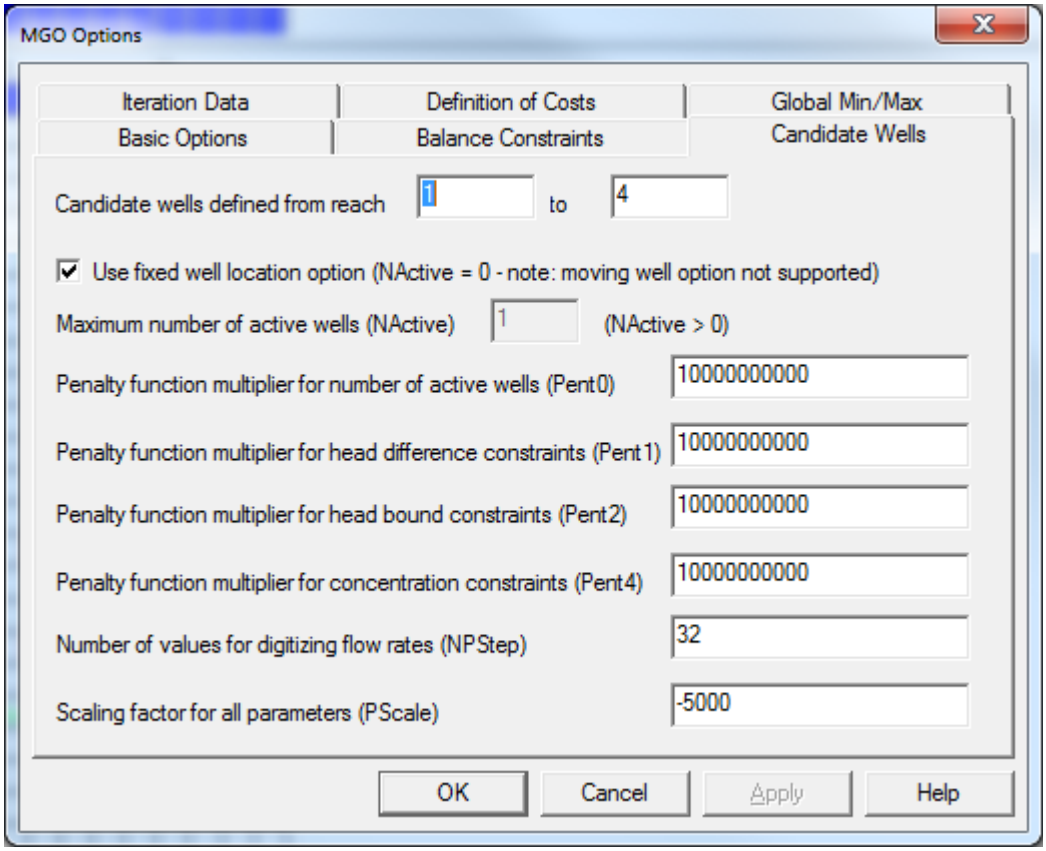

Next, click on the Definition of Costs tab. In this simple problem, the only cost is for pumping and the type of optimization is minimization. Since the cost per unit flow is 1.0, only pumping rates will be minimized.

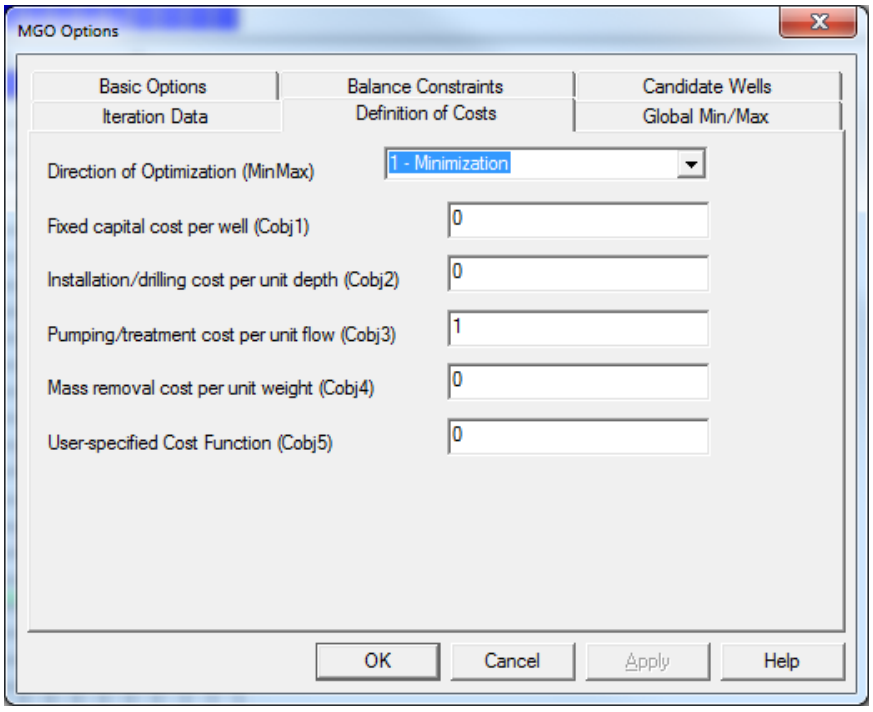

To run this problem, select Model|MODFLOW|Create Datasets and then select Model|Optimization Models|MGO|Create MGO Input File. Now, select Model|Optimization Models|MGO|Run MGO.

After the simulation is finished, you can view the final well configuration chosen by MGO by selecting **Model|MGO|View Final Well File**. In this case, two of the four pumping wells were chosen.