
Discussion archives on groundwater modelling

Dec 14 2004, 7:20 am

I have been compiling selected extracts on groundwater modelling discussion from various mailing lists for past few months, as given below. I understand that in view of random and discontinuous placing, many parts may not be easily comprehensible. Still I think that it could be helpful to groundwater modellers to some extent.

Regards,
Kumar

=====
C. P. KUMAR
Scientist 'E1'
National Institute of Hydrology
Jal Vigyan Bhawan
Roorkee - 247667 (Uttaranchal)
INDIA
Web Page: <http://www.angelfire.com/nh/cpkumar/>
=====

* In my experience, PCG is faster than LMG for small models but LMG is faster for large models. Of course, for small models, speed generally is not so much of an issue.

* DAMP is used as a multiplier to scale the head change at each iteration, not RELAX. The DAMP parameter in the PCG solver package is analogous to the ACCL parameter in the SIP package. Although I cannot explain precisely what RELAX controls (since

I'm no mathematician), the typical range for RELAX is 0.97 to 1.0. In my experience, setting RELAX=0.97 can help to achieve convergence, and may speed up convergence on a "well-behaved" model, but not always.

* I'm sure you've heard "All models are wrong but some are useful." If you replaced all the wells in a model with drains, you would expect to get different results. How is this any different? You have replaced one package with a different one that is based on different assumptions so of course, they get different results. If you want to know which one is better, examine the assumptions behind the different packages and see which assumptions best match the conditions at the particular region to be modelled. If none of the available packages do what you need, write your own package to do what you need.

* The newer LPF package is the way that most other models work, i.e., you enter Kz and the inter-cell conductances are computed internally. The way that MODFLOW88/96 did it with a pre-computed vertical conductance always struck me as a bit odd; however, that has been our happy MODFLOW world for about 20 years. So now we have a new option that is mathematically more correct, but it is still disconcerting when the computed heads are different. I guess you'd have to ask a lawyer about who would win. My sense is that the differences could never be explained to a jury. It does present the MODFLOW modeling community with a problem, though. The way that you did it with MODFLOW-SURFACT is probably the way that the USGS should have done it as well - adding the input of Kz in the BCF Package.

* Input of vertical hydraulic conductivity has its advantages, however, McDonald and Harbaugh got it correct in keeping a fixed thickness of a cell for the vertical direction conductance - for gravity-segregated vertical equilibrium (GSVE), you have a transmissivity that is conductivity times saturated thickness which is valid only for the horizontal direction and not for the vertical direction (in actuality, the vertical direction physics changes from GSVE to unsaturated flow which would lower the conductivity further and that would be closer to the "fixed thickness" answers that do not reduce separation distance between 2 cells vertically even for consideration of a homogeneous case, but that aside). Nor does the document state where or under what conditions the use of a "variable thickness for vertical conductances" option performs better. That's why I was asking to see if anyone knew why this was done in the first place.

* I propose that the node location of a cell that contains the water table should be located at the water table instead of the center of the cell as is standard. Comparisons of this formulation to analytic calculations numerically implemented in the WATQ code indicate a much-improved solution at shallow depths compared to the calculations of the current LPF formulation. This formulation does not, however, take into account the dynamics of drainage from the unsaturated region above the water table, which can be a major influence. I want to repeat Richard Winston's point that "All models are wrong." A model is not intrinsically good or bad. That judgment depends on what the model is used for. What may be better in one situation may not be better in another. You have highlighted this concept in your discussion of GSVE. MODFLOW still exists for two contradictory reasons; 1) as the name implies it is modular and hence easily extensible to specialized situations. 2) It is widely known and accepted that the code has accurately solved groundwater flow problems. It is incorrect to extend the second point to expect MODFLOW to accurately solve all groundwater flow problems. Once again required accuracy like good and bad depends on the problem that needs to be solved. That MODFLOW produces two results for the same geometry and boundary conditions means that it approximates the physics of the flow situation differently. If proper care is taken, the ability to change the approximation of the physics of flow is a good thing because the code can be tailored to specific problems.

* Yes, having flexibility in averaging schemes is a good thing. The LPF package provides one more averaging scheme in the vertical direction. It does a harmonic average of the conductances, as opposed to the BCF package's harmonic average of vertical conductivity times area divided by separation distance as provided in the MODFLOW documentation for the various fully three-dimensional, and quasi-three-dimensional alternatives. However, I am guessing that people are encountering drastically different and surprising or unexpected results because of the vertical direction treatment. Specifically, the cell thickness varies as a function of its saturated thickness in the LPF package for vertical direction treatment. For a simple example, consider a homogeneous case with two cells where the cell below is getting drier. The vertical conductance quickly rises due to the cell below drying up, to such levels that drain and dry up the cell above (which also causes the conductance to increase in a cascading manner) which would be a very different result than with use of any of MODFLOW's BCF options. Like I stated earlier, in fact, when things dry up, it should slow down the conductance due to physics change to unsaturated flow, and even a gravity flow assumption will not speed it up.

Yes, there is a dilemma in GSVE models for vertical direction treatment because GSVE by its definition, applies only to horizontal direction flow - I have seen this in saltwater-freshwater GSVE models with multiple aquifers as well. And there too, when you give it specialized treatment, it is due to certain physically motivated considerations. Otherwise, the way MODFLOW BCF did it (or a use of grid-block cell thickness in the LPF package) is the best you can do with the VE assumptions being used, since conductance may be smaller than saturated levels due to unsaturated zone physics - but you don't know by how much, unless you simulate unsaturated zone physics. Using saturated thickness instead of cell thickness in the vertical direction leads to results that are more removed from this physics.

* This is one of the most challenging groundwater modelling projects you can find in real-world modelling... Ideally, you should use a saturated-unsaturated code, such as FEFLOW or Modflow Surface or similar. Such codes are able of simulating perched water conditions.

If you are using traditional MODFLOW, and assuming saturated conditions, it is a bit more complicated to simulate such a scenario. You will probably have some dry cell issues (and perhaps some convergence problems), but it is possible to simulate the open pit crossing multiple layers using drain boundary conditions.

For the upper layers, where seepage faces may occur, set the drain elevation slightly above the layer bottom elevation (say 0.1 m above cell bottom elevation). If you use Visual MODFLOW as the interface, that is very easy to do with formulas that will assign the correct elevation even in irregular layers. In that case, if the head in the pit wall is higher than the cell bottom elevation, Modflow's drain package will remove the outcroppng groundwater. The conductance value (another input to the drain package) is typically a calibration parameter, obtained through matching drain removal rates with real values observed at the mine site.

I suggest that you also use a solver that is not too aggressive, such as PCG2, and that you play a bit with re-wetting parameters and dumping factors until you obtain a stable solution. We've done several mine dewatering modelling projects using this approach and it definitely works for most cases, as long as you understand the code and it's limitations.

* Dry cells in a multi-layer model are not necessarily a bad thing unless they do not accurately reflect the conditions in the field. Perhaps I am stating the obvious, but just for the record, dry cells occur in MODFLOW models when the calculated water table is below the bottom elevation of the grid cell. Since MODFLOW only simulates saturated groundwater flow, it does not represent a head value in the unsaturated cell, so the cell becomes 'dry' and MODFLOW assigns a 'dummy' head value usually in the range of $-1e-30$ (or something else easily recognized as a dry cell). Although the presence of dry cells in a model is often a hassle because they create problems with the convergence of the numerical solution, dry cells do not necessarily indicate a 'problem' or error with the model. In your case, the dry cells don't appear to be causing a problem with the solution convergence (or at least, you haven't mentioned this problem) and you are getting a good mass balance, which is a good sign but doesn't necessarily indicate 'correctness' of the solution. If the water levels in the second layer of your model match closely with the observed field conditions, then the occurrence of dry cells in layer 1 appear to be reasonable.

However, if the dry cells are occurring where you would otherwise expect saturated conditions (based on water levels in the field) then you probably need to revisit your boundary conditions to make sure they are reasonable (conductance values are often the culprit), check your initial heads and make sure you are not starting the iterations with dry cells, and make sure you have enough recharge flux entering the system to accommodate flux leaving the system through wells, rivers, and other boundary conditions. If all this looks good, then Mark Wilsanac gave some very good suggestions on how to change the solver and rewetting package settings in order to deal with dry cells. - Patrick Delaney

* If the river boundary cells are the only place in your model where water may leave the system, then one of the most likely causes of the mass balance problem is the conductance values you are using for the river cells. Check your conductance values to make sure they are within reason. The conductance values can be reasonably estimated using the following formula:

$$C = (L \times W \times K_v) / B$$

where

L = length of the river in a cell W = width of the river in a cell K_v = vertical hydraulic conductivity of the riverbed in a cell B = thickness of the riverbed in a cell

Most of the popular graphical interfaces for MODFLOW allow you to enter these 'physical' parameters and will automatically calculate the conductance values for each river boundary cell.

Another possible cause for mass balance problems could be the solver you are using (SIP often produces larger mass balance than other solvers like the PCG or LMG) or perhaps the convergence criteria for the solver are not 'tight enough'.

Just as an afterthought, the fact that you have more water entering the system from the rivers than leaving the system through the rivers does not, in itself, indicate a problem with the water balance. This indicates you have 'losing rivers' because the water table adjacent to the river is lower than the river stage. The excess water entering the system from the rivers may be leaving through other avenues such as evapotranspiration, pumping wells, drains, specified head boundaries, etc. - Patrick Delaney

* Check out if you don't have any dry well cells, or any other dry boundary cell. If a cell goes dry during simulation, it becomes inactive for MODFLOW, and any boundary assigned to that dry cell will be ignored. That means that your wells may be inactive (even though you assigned them to be pumping all the time). That typically produces a noticeable change in the budget, but sometimes the heads may not change too much, unless you look at specific well cells very closely.

My suggestion is to first look into the budget for changes in well pumping/injection rates. If the total rate is different from those you entered into the model, then you have a dry well cell. Zoom in all well cells to identify which ones have become dry during that stress period (look in all model layers).

* With a model set up with cell-to-cell size changes below 50%, there can be instabilities or failure to converge. This can be due to other reasons, or cell size change coupled to highly variable conditions in a transient solution. If a geologic unit pinches out, that should be represented primarily by changes to cell properties, and secondarily by changes in cell dimension (aspect-ratio rules apply vertically as well as horizontally). Remember to preserve aspect ratios of each cell below 10:1, and below 5:1 to be safe. It is acceptable to have a layer with variable properties from cell-to-cell. It is rarely acceptable to use MODFLOW cell dimensions to closely simulate physical variations in layers when they pinch out.

* I have not used the WHS solver sold by Waterloo hydrogeologic. while I therefore cannot comment on this solver and its claimed superiority to SIP, I feel quite sure that SIP has been the best performing solver for almost all models I have constructed over the past few years, and never had problems with mass balance which could not be successfully overcome by using SIP. I do recall a paper published in Ground Water years ago when WHS proponents at software developer Waterloo Hydrogeologic compared their WHS to SIP and PCG (among other public domain solvers). I felt that sip had been shortchanged and misrepresented in their analysis in that they had claimed that sip yielded a much greater mass balance error than that obtained when using their WHS solver. What they had failed to do in their analysis, was to accordingly decrease the HCLOSE with their corresponding decrease in the acceleration parameter. This is the equivalent of not closing tightly enough on heads, and therefore, ending up with an excess mass balance error. This was an improper use of SIP.

I have found that lowering of the acceleration parameter, and corresponding decrease in the HCLOSE will yield a very low (acceptable) to zero mass balance error, but one must be careful in that the acceleration parameter is not set so low that the model converges prematurely with insufficient head change.

* Check all inputs in the flow model first. Make sure your flows are realistic. Make sure your K's and recharge are realistic. Refine your grid and use dispersivity values that will minimize numerical dispersion. Make sure your grid cell size complies with appropriate Peclet criteria (as a rule of thumb, for finite differences, Pe should be ~ 2 ; that means your grid size should be roughly 2 times your dispersivity value, in the area where, for MOC, this can be higher). I suggest you read some transport modelling books that explain this in detail. Be careful with using constant concentrations, they can act as a source, but also as sinks (if the concentration in neighboring cells is higher than the ones you specified in the constant concentration node).

* When she is referring to inflow, in plane, and outflow, I think what she is doing is looking at the color-coding for the velocity vectors or pathlines. When she sees that there are very few arrows or pathlines with a green color (indicating 'inplane' or flat horizontal flow in the layer) she gets concerned for some reason. Most likely the lack of 'in plane' flow has very little to do with anything she is concerned about, it just reflects the fact that she has slight vertical gradients in the model.

WRT to the vanishing concentrations, you could also suggest her to make sure she didn't input a default degradation rate of 0.5 or something very high and then forgot about it, or the other possibility is she has assigned a sorption coefficient of 1 (mistaking it for a retardation coefficient), or perhaps just a very large sorption coefficient of 0.01 L/mg which would cause the plume not to move at all, thus giving the impression that it is disappearing from the model, when in fact it hasn't moved at all. Another possibility is that she initially assigned non-conservative values of sorption and decay coefficients when she set up the transport model scenario, and she is now trying to modify these values in the Setup/Transport Engine dialog instead of in the Input mode (I've seen this problem a few times) or Visual MODFLOW.

Finally, the fact that the solution solves in two iterations seems to indicate that the total run time of the MT3DMS solution may have been left as the default of 1 day (perhaps she assumed it automatically picks up the end time of the flow simulation). She can change this by selecting Run/MT3DMS/Output Times and entering the desired Simulation Time.

* The simplest way of defining a starting concentration is to use the Recharge Coverage and specify a polygon shape that denotes the contaminant concentration to be applied to the model.

* What to do when your contaminant plume does not migrate as expected.

When using MT3D or RT3D for a contaminant transport simulation, if the plume fails to move as expected (or does not appear at all in your Visual MODFLOW output), here are a few common causes and their solutions.

Problem #1: Overlay is not active, or masked by another overlay

Solution #1: As with Head Equipotentials, Particles, and even Basemaps and other overlays, you must make the Concentration overlay active in order to see it. Click the F9-Overlay button at the bottom of your screen, and ensure the overlay has been checked off (to make it active). It is also possible that other overlays are masking your concentrations. To resolve this, change the Overlay Order setting to User Defined, then highlight the Concentration overlay, and use the arrow buttons to move it up the list.

Problem #2: Simulation Output time

Solution #2: In Visual MODFLOW, the Simulation Output time can be different for your flow and transport simulations. Check to see that you have defined the correct simulation time(s) in the Run Menu, by selecting MT3D (or RT3D) / Output - Time Steps. Additionally, in the Output Menu of Visual MODFLOW, ensure that you have selected the Concentration overlay as the active overlay, then select the Time button in the left toolbar, to change output times as desired.

Problem #3: No concentration input assigned, or not assigned properly

Solution #3: Ensure that you have correctly defined at least one cell with an Initial Concentration, or a transport boundary condition of one of the following types: Constant Concentration, Recharge Concentration, Evapotranspiration Concentration, or Point Source. Please NOTE that Concentration values entered for Concentration Observation Wells are used for calibration purposes only; they do not contribute mass to a transport simulation.

Problem #4: Inadequate flow gradient

Solution #4: Before running your transport simulation, run just the flow simulation (MODFLOW). Using Pathlines or Velocity Vectors, ensure that there is a sufficient flow gradient to cause plume migration.

Problem #5: Incorrect reaction parameters

Solution #5: Check the reaction options in the Main Menu, by clicking on Setup / Numeric Engines / Transport. Check that the correct Sorption and Reaction options have been selected.

First order decay rates (dissolved and sorbed phases) may have a tremendous impact on the plume mass, since during the simulation, some contaminant mass may be removed from the model. If the decay rate is too high, your plume will show very small concentrations, and/or will not be visible at all at later simulation times.

Decay rates (due to biodegradation or other mass-removal processes) can be taken from literature values, but this should be done with caution, as this parameter is highly site-specific for most compounds. A good reference on decay rates used in natural attenuation studies can be found in:

Wiedemeier et al., 1999: Technical Protocol for Implementing the Intrinsic Remediation with Long-Term Monitoring Option for Natural Attenuation of Dissolved Phase Fuel Contamination in Ground Water. Air Force Center for Environmental Excellence, Brooks AFB.

This document can be downloaded from:

<http://www.afcee.brooks.af.mil/products/techtrans/monitorednaturalatt...>

The decay rate (λ , or sometimes called K_d , with units of 1/day), is typically obtained from half-life values, converted into appropriate units using the following relationship:

$$\lambda = \ln(2)/t_{1/2}$$

Where $t_{1/2}$ = half-life of the compound

In addition, check that your K_d value is correctly defined in the Species Parameters tab. A very large K_d value will result in an extremely high retardation factor, which can result in lack of contaminant movement through your model.

For typical organic compounds (such as TCE, DCE, PCE, BTEX, etc.) where linear, reversible sorption can be assumed, retardation will be calculated using the following formula:

$$\text{Retardation} = 1 + (\text{Bulk Density}/\text{porosity}) \times (K_d)$$

Where K_d = Partition coefficient

For hydrophobic organics, Kd can be determined using the relationship below:

$$K_d = K_{oc} \times f_{oc}$$

Koc - octanol-carbon coefficient f_{oc} - organic carbon fraction in the aquifer

A more detailed overview on Kd's can be found in:

US-EPA, 1999. Understanding variation in Partition Coefficient, K_d, values. EPA 402-R-99-004A&B. US EPA Office of Air and Radiation.

This document can be downloaded from:

<http://www.epa.gov/radiation/cleanup/partition.htm>

* From the Modflow standpoint, water that exits the GW system through of a drain cell is accounted for in the budget, but is otherwise not simulated. The situation for river cells is similar--Modflow does not simulate any flow in the river, it only simulates exchange between the GW system and each individual river cell. So there is no way to route water from a drain cell to a river. The Streamflow Routing (STR) Package, on the other hand, does simulate flow in surface-water streams, accounting for flow rates in the stream and calculating stream stage, which is then used as the external head in the GW/SW exchange calculation. If you want to simulate exchange between drains and rivers, you could use the STR Package for both types of features.

* How to import modflow model files to GW Vistas 4?

Select File/New and click the MODFLOW button. On the next dialog, click browse to go find the name file (*.nam). If this is an older MODFLOW88 dataset, then browse to find the *.bas file. Some datasets (especially those that were created without the use of a preprocessor) can be troublesome.

* I want to input separately 2 sets of recharge into a modflow model: (1) Effective recharge from rainfall, and (2) urban leakage (same format as the (1)), which will overlap. Since both components are fairly complex and over a long time (1200 Stress Periods), I would like to keep them as independent as possible (avoiding the data compilation option, if you see what I mean).

It is possible to do what you want using just the Recharge package. Here is how you do it.

Define each recharge distribution for each stress period using parameters. You can use one set of parameters for the rainfall and another for the urban leakage. With each parameter, you can use multiplier arrays and zone arrays to define the spatial distribution of the recharge. Then, for each stress period, use the parameters for that stress period to apply recharge from both rainfall and urban leakage. MODFLOW-2000 will add the contributions from all the parameters you use in a stress period to determine the total recharge.

You should be able to do this with any reasonably up-to-date GUI for MODFLOW-2000. If you are using a GUI and you can't figure out how to do this, contact the GUI developer for assistance. If you aren't using a GUI, you can check the input format in the Online Guide for MODFLOW-2000 at <http://water.usgs.gov/nrp/gwsoftware/modflow2000/MFDOC/guide.html> or in the original MODFLOW documentation as modified in "Time-varying-parameters.pdf" (distributed with MODFLOW-2000).

* Since I know you are using Modflow VKD (based on MF96, not MF2K) and that the wells package is almost certainly being used for many abstraction wells, I suggest 2 routes:

1. A FORTRAN utility to combine two recharge files - a days work? 2. To use a MF package that you aren't already using - but with some manipulation e.g., the river package. By setting the stage of the river as above any elevation in the model (1000m?) and the conductance equal to the quantity you want to recharge for that SP - with a 1 m difference between stage and bottom elevation you will get the appropriate amount leaking (recharge) to each cell.

You can achieve the same result with the drains package by specifying 2 drains in each cell. The trick is to have one with a positive conductance and the other negative, but the same value. Make the difference in elevation 1m and the magnitude of the two conductance values the quantity you want to leak/abstract. If the elevation is outside the range of variation in the model, you end up with a constant recharge/abstraction from the model.

For composite/trick boundary conditions of this type I suggest drawing figures for them like those in the MF88 book, it will help to understand what I'm trying to say.

* Design the shoreline?

Depending if you are using FEMWATER or MODFLOW, it is handled differently. With FEMWATER, you want the boundary of your 3D mesh to follow the shoreline. With MODFLOW, you are working with a 3D grid. So, you need to mark the cells that contain the shoreline with a fixed head boundary condition, and the cells outside this area will be disabled.

If you are attempting to model the interaction with the seabed, then you will need to set the top elevations of the mesh or grid to match that of the seabed. However, generally you end your model at the shoreline.

* MT3D Performance

What you need most depends on where your bottleneck occurs. You need to monitor CPU and RAM usage on the computer during the MT3D run. If you run Windows 2000 or XP, you can right click (for a right-handed pointing device) on the Taskbar (grey bar at bottom of screen), then select Task Manager in the menu pop-up. Once Windows Task Manager opens, select the tab labeled "Performance". "Performance" shows CPU usage graphically and memory usage in numeric form. If the "CPU Usage History" (top graph) stays maxed at 100%, then a CPU speed increase should help. In parallel, observe the "Physical Memory" values. If the "Available" value becomes a small fraction of the "Total," AND the "Commit Charge, Total" exceeds the "Physical Memory, Total" then you are likely to have a RAM-limited condition. In this case, more RAM would help. With 300,000 cells, consider 1 Gbyte or more. If both things are happening (unlikely), then upgrading CPU and RAM should help. Be aware that upgrading one component sometimes results in the other becoming the limiting condition. That is why it is "unlikely" both limits are happening simultaneously. Usually one tops out before the other. If you run Win98... consider upgrading the OS.

* There are various things that you have to do to get Surfact running in Vista 3:

1. You can only have a limited number of packages. For example, if you are doing contaminant transport modelling then you have to turn some other packages off, otherwise Surfact will not run e.g. you can turn the cell by cell flow packages off (just type 0 into the packages number in Model - Modflow - Packages). Another place you can turn packages off is in the output control: Turn off the drawdown unit.

2. You have to modify some things in the Model - Modflow options and some in the Modflow - Surfact options. These are:

Model - Modflow - Packages: change modflow version to Surfact; change solver to PCG4; and untick automatically reset package units

Model - Modflow - BCF package: tick the box to use the BCF4 package

Model - Surfact - Packages: add a unit number (one that you have not yet used) and tick the box for BCF4; add the unit number to PCG4 and tick the box (use same unit number as shown in Model - Modflow- Packages); add a unit number and tick box to RSF4 is you want the model to simulate surface seepages (and remember to turn this unit on in Model - Surfact - RSF4 package).

Model - Paths to models: set modflowwin32 option to DO NOT USE; put in correct path for Surfact in the modflow box.

I can't think of anything else right now - I always forget one thing and spend a couple of hours wondering why Surfact won't run. If you have any difficulties it would be helpful to know if the Surfact Dos window displays and if so, what it says in the window.

* The most common problem when using SURFACT is setting up the program file. Select Model/Paths to Models. Change the MODFLOWwin32 Option at the top to "do not use". That tells Vistas that you are not using one of our model DLLs. Then change the MODFLOW program from MFWIN32.DLL to whatever your SURFACT program is (usually it is msft.exe).

* I believe that if you are using the latest version of SURFACT (V2.2) that problem 1. Below is no longer an issue. It used to be that SURFACT could only open a limited number of files but that is no longer the case. Also, if you upgrade to Vistas Version 4 and SURFACT V2.2 the link between the two programs is easier to establish. The same holds true for all of the other versions of MODFLOW that GW Vistas supports (MODFLOW88, MODFLOW96 in double precision, MODFLOW2000, MODFLOWT, MODFLOW-SURFACT, SEAWAT, and SEAWAT2000).

* VS2DI or VS2DT, being Richards's equation-based models, have certain PROs and CONs when used for groundwater recharge estimation:

Pros: the Richards equation models are the most theoretically based and allow representation of the flow processes in porous mediums more realistically than the water-balance models. They have been well tested against field and laboratory experiments and proved to provide good correlation with observed results.

Cons: the major complicity of the Richards equation, comparing to the saturated flow models, is that its coefficients are non-linear functions of the pressure potential. To approximate these functions, three different algebraic equations are usually used (named after their authors): Brooks and Corey's, Gardner's, and van Genuchten's. Richards's equation can be solved with a very limited number of boundary conditions. The condition for water at the boundary can be specified either as the flux of

water across the boundary, the head at the boundary or as a combination of specified head and specified flux. The Richards equation with these boundary conditions is a nonlinear partial difference equation that has no close-form or analytical solution. To solve the Richards equation, a finite-difference method of numerical approximation is applied in VS2DI. Searching for a best numerical method for the Richards equation became a special field in hydrologic science, particularly in 1980's, however, almost all implemented approaches cannot assure that the model will unconditionally converge and simulation results will be obtained. It is an art to get the model to work correctly (produce results for the whole simulation period with a good water balance) when you have varying flow boundary conditions in your profile. Our experience also proved that applications of Richards's equation-based models to highly heterogeneous soils with variable hydraulic properties can be extremely difficult to execute and time consuming.

These days there more and more examples when less sophisticated in flow equation but much more advanced in terms of weather/plant effects water-balance model HELP is used to estimate groundwater recharge.

* I will suggest you keep the north and south boundary as variable head boundary. This I am suggesting you treating there is no river, lake etc. If these features are present on north or south side of model domain, then you have look for another type of boundary. Regarding limited observation well data and that too is limited to layer1, then it will not be possible to do any realistic modelling. If there is no scope of generating additional data, I would like to suggest that you confine your modelling activity to first layer only.

* What makes a good correlation depends on what your expectations are. If you are conducting exploratory research, maybe 0.2 isn't bad. If you are testing a known process having some natural variability, 0.6 might be acceptable. But if you are calibrating instruments, maybe 0.9 isn't good enough.

If you don't have any clear expectations, how do you tell if a correlation is good? The answer comes in three parts.

1. Value - Square the correlation coefficient (called the coefficient of determination or just R-square). R-square is the proportion of variation in the dependent variable (y) that can be accounted for by the independent variable (x). You might be able to decide how good your correlation is from a gut feel for how much of the variability you wanted your model to account for.
2. Significance - Every calculated correlation coefficient is an estimate. The "real" value may be somewhat more or somewhat less. You can conduct a statistical test to determine if the correlation you calculated is different from zero (or some other number if it's relevant). The larger the calculated correlation and the greater the number of samples, the more likely the correlation will be significantly different from zero. For example, a correlation of 0.59 (R-square of 0.35) would be significantly greater than zero based on about 25 samples, but a correlation of 0.01 wouldn't be significant with 250 samples.
3. Plots - You should always plot the data used to calculate a correlation to ensure that the coefficient adequately represents the relationship. Specifically, the data should be linearly related and free of outliers. There are a few other things to look for (e.g., hidden trends, autocorrelation) that I won't go into here.

So, the reviewer who said that "the "R square" value of 0.35 has no significance and is just as good or as bad as say 0.01" would only be correct if she looked at a statistical test of whether the value was significantly different from zero. R square does not need to be more than any particular value to draw a comparison. Remember, too, that "no relationship" may also be an important finding.

* If you are considering GUI's for your project, you may want to look at Visual MODFLOW Pro. In particular, if cell splitting, or grid refinement is an issue for you. Visual MODFLOW includes two distinct features that help users with this.

1) Cell refinement tool 2) Grid smoothing tool

1) Cell refinement/splitting tool: You'll need no more than a few clicks to edit your grid in a very flexible manner. To split your cells, just click on 'Refine by ...', and specify how many times you want the cell to be split into (2, 3, 4...). Then, select the area you need to refine with your mouse. You can apply this to one row or column, or all cells in the domain - whatever you wish. The same applies for splitting layers and coarsening the grid.

2) Grid smoothing tool: After grid refinement, you may want to use the grid smoothing routine to produce a smooth transition between coarser areas to highly discretized ones. Again, just a few clicks. This is actually a perfect tool to help you better understand the effects grid size and refinement will have on your model results.

Another feature that may be useful for you is Visual MODFLOW Pro's batch capabilities. It will allow you to prepare all data input files in different model projects and run them in 'batch mode'.

* The interpretation method is different according to the geology. You have methods for porous media, others for fissured aquifer, and others for double porosity media. Methods also vary according to the type of aquifer (confined or unconfined; semi-confined with a leaky aquifer...) and corrections can be made according to the type of well (partially penetrating well; pumping well with head loss; skin effect and so on). Finally, boundary effects (barrier or recharge) may affect significantly your interpretation.

In order to remain pragmatic I would suggest you to start with the simplest method: the Theis formula, (or even the Jacob approximation if your pumping time is big enough) and check whether you get a good fit or not with realistic parameters. Even a fissured unconfined aquifer can be interpreted with Theis when your radius of influence is wide so that the fissure network can be assimilated to an homogeneous equivalent porous media and the depletion of the water level is moderate compared to the thickness of the aquifer. This should give you a rough estimate of your transmissivity, which is in many cases, sufficient.

If you need more accuracy you can use specific software. You will find on the net many software proposing many different methods. You don't need Pest as they have their own fitting engine. AQTESOLV for instance proposes methods with well storage effect. However, as soon as a software proposes its own adjustment, it's becoming very difficult to know what's you're doing and you end up playing a video game. So be careful, especially if you lack experience in this domain, not to use complex methods with many parameters, which will end in a perfect fit but with unrealistic results. In my company (French geological survey) we use ISAPE, a software which is no more commercialized, but whose philosophy is to let the user propose realistic values and where you can add or remove well effects or boundary influence. I wish this approach were more widespread.

* I think your scenario will represent two different geologic materials with peculiar intrinsic properties. Hence, I suppose this will require two separate simulations. To the best of my knowledge, MODFLOW 2000 only allows one K matrix values per layer per simulation and this is incorporated within the flow package. However, if the introduction of the permeable barrier is controlled along defined specific paths or unique relatively smaller area(s), then you can use one simulation with two stress periods and incorporate Horizontal Flow Barrier Package in the second period to account for the permeable barrier distribution. This is not supposed to be of much problem especially if you incorporate your modflow with GIS package like GRASS GIS.

* Sophisticated modeling programs like Visual Modflow Pro, or any of a number of similar packages, will run properly with a variety of inputs that may not be physically realistic - and thus provide garbage for output. With that in mind, here are some thoughts on your questions:

1) You may divide a single aquifer into more than one layer if there is some reason to look at vertical stratification within the aquifer. For example, a thick aquifer may have different hydraulic or transport properties in different vertical zones. You may also divide a single aquifer into several layers if you want to examine the vertical flow patterns due to pumping from different zones within the aquifer.

2) The top of Layer 1 could be the ground surface, the water table (not as common because it can move vertically), or the top of the uppermost aquifer or aquitard of interest depending on the physical system being modeled.

3) You'll need to assign boundary conditions no matter what the size of the physical system is. If you don't have natural physical boundaries, one strategy is to assign boundary conditions at a distance far enough away from the area of interest to minimize the effects of boundary assumptions on the solution.

* It is appropriate to use a single layer if you feel that it is reasonable to ignore vertical flow in your model. Assuming that you would be using the Layer Property Flow (LPF) Package of Modflow-2000, and if the aquifer represented by layer 1 is not overlain by a confining unit, then specifying the top of layer 1 as ground surface is reasonable. The reason that you would be better off this way than specifying the water table elevation as the layer top is that the LPF Package only provides two options with respect to the confined/unconfined status of a layer: Confined or Convertible. If the head in a cell goes above the specified top of a convertible layer, the cell is treated as confined. Every model needs to have at least one fixed head, either as a constant-head cell or as a fixed external head (such as the stage at a river cell), or the solver will not be able to reach a solution. You may need to expand the domain of your problem to encompass a fixed head.

* 1- Model design depends greatly on your project goals. For example, you could have one aquifer, but an overlying layer (and/or underlying layer) as well. Or, you could be modelling just 1 aquifer, but split it up into several layers in VMOD to provide more detailed flow information (i.e., if the aquifer is 100 feet thick, having just 1 layer will not provide very detailed output. You can subdivide the 1 "real-world" layer into 10 "model" layers, all with the same hydrological properties, in order to obtain more detailed output). Or, you can simply design a 1-layer model, if that is all that your project requires.

2- In Visual MODFLOW, Layer 1 represents the uppermost layer of your profile, and the top of Layer 1 is therefore the ground surface. Layers in Visual MODFLOW represent soil layers (or conceptual soil layers)...the location of groundwater in these layers is determined by your model inputs, and the resulting calculations of the flow model.

3- The assignment of boundary conditions is a question that often goes beyond the scope of Technical Support, and enters the realm of Extended Modeling Support. However, to provide you with some guidelines:

-If you have no "obvious" water inputs (streams, rivers, lakes, injection wells, recharge from rainfall, etc.), you can try assigning an upgradient equipotential line, and downgradient equipotential line, using Constant Head Cells at the upgradient and downgradient edges of your model, to represent the regional water table (other boundaries may be more appropriate, depending on your specific model domain and objectives). You can then add additional inputs (extraction wells, contaminant sources, etc.) to the model region, which will be influenced by your regional flow gradient.

4- I would recommend taking some time to run through the Tutorials included with your Visual MODFLOW software (the PDF instruction files are located on your installation CD-Rom, and the files are automatically copied to the Tutorial subfolder in your Visual MODFLOW program folder). These tutorials are designed to teach you how to work with your software, and how to design a model, import data for model inputs, run the various packages, calibrate your model, examine your output in 2D and 3D, and export data.

* You can use the field interpolator program in the PMWIN package. Using the field interpolator you can interpolate the field data in txt file to the grid point of your model. On the other hand, you can import any DXF file, which illustrate the changes in hydraulic properties to help you in the graphical interface.

* Upwind discretization always has an artificial diffusion term no matter how refined your grid is (even though it diminishes with grid refinement). Central discretization always produces an artificial dispersion term that causes "wiggles" however you can't control it with grid refinement. For almost any Pellet number the upwind scheme will produce positive coefficients for your coefficient matrix, while that isn't true for central differences. However central differences has "second order accuracy" while upwind only has first order. You can visualize numerical diffusion and dispersion mathematically by deriving the "equivalent or modified equation" see <http://widget.ecn.purdue.edu/~jmurthy/me608/main.pdf> for some cool notes on this (as good as any book I've read).

Physically you can visualize numerical diffusion by imagining a 1D discretization of a domain $| : | : | : |$ where $|$ are faces and $:$ is the center of the cells.

If your initial property is 1 on a given cell say i and 0 on all others it would look like this:

$| 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0$

So using upwind with a $courant = v dt / dx$ of say 0.5 and a velocity from left to right after the first time step the accurate solution would be something like (zooming on cell i):

$| 0 \{ 0.5 | 0.5 \} 0 |$

Note that I have created a new cell delimited by the $\{ \}$ brackets this "new cell" that doesn't exist in our domain is exactly half of cell i and half of cell $i+1$. The property's value will be 0.5 inside this new cell and 0 outside it (meaning that cell i and $i+1$ would have half with 0.5 and half with 0). This would be the accurate solution because a $courant$ of 0.5 makes the property shift half a cell every time step. However we are bounded to our initial discretization so the solution we will obtain will be:

$| 0 | 0.5 | 0.5 |$

so you can see than instead of having a 0.5 of property in a volume the size of a cell we have 0.5 in a volume the size of 2 cell thus lowering the concentration. This happened because the definition of concentration is $C = \lim_{vol \rightarrow 0} DM/DVOL$ and we implemented it as $DM/DVOL$. If we would diminish the size of our grid so that $courant = 1$ we could obtain the exact solution. So the smaller your grid is the more accurate your solution will be. This is the physically correct way of solving the diffusion problem another way is to use higher order schemes witch usually maintain the sharpness of the gradients but usually remove mass from all the wrong places (even though they conserve it). Looking at the same example a central differences discretization would derive something like:

$| -0.25 | 0.75 | 0.25 |$

Cell $i-1$ will produce a negative property because the concentration at the face between $i-1$ and i is 0.5 cell $i+1$ will obtain the accurate concentration of 0.25 (0.5 in half the volume of a cell). But we are producing the famous "wiggles" close to the properties gradient. There are smarter second order schemes like *qwick [sic?]* or TVD schemes.

* In visual MODFLOW there is provision for importing MODFLOW files (*.bas files). In PMWin there is provision for conversion of MODFLOW88/96 (*.nam). I have tried to import *.bas file in Visual MODFLOW and found that model data is not correctly coming to model. These options may be available in GMS but I am unable to locate the exact sub-menu. Number of options is available for opening different format files In the FILE menu. But It does not permit opening of *.bas or *.nam files. This option must be in different menu.