the Water in the Soil – Part 6

Geotechnical News, March 2012

William E. Hodge, PEng, M.ASCE

Now that I am writing the last article in this series I find myself wondering where on earth these ideas might have started out.

Maybe it was back in the 60's when Arrow Dam (now Keenelyside) was been built on the Columbia, and I was there as the junior engineer looking after earthworks and instrumentation. I remember at one stage the glacial till core earthfill was responding to roller compaction by making waves, as is inclined to happen when such material is placed too wet of optimum. As it happened, Arthur Casagrande was due to make one of his routine consulting visits just about then, so I installed a piezometer about 10 feet below grade and attached it to a pressure gauge. Then, as he watched, I had loaded dump trucks pass over the spot where the piezometer was buried. I wanted to see what he would say to the fact that the pressure on the gauge rose as the truck moved over the spot and then dropped back to zero as the truck moved away. Although he looked for a good while, sad to say, he went away without telling me what he thought about it. But now, half a century later, I think that observation might have done it for me.

Excess Pore Water Pressure

As a geotechnical engineer working in design and construction I was acutely conscious of being obliged to deal with soil behaviour at only one or other of two extremes: Fully drained, or no drainage at all. The real world was always somewhere in between - but inaccessible. This wasn't all that bad until earthquakes entered the scene. Then I felt our work was degraded to following some quasi-mystical beliefs set down by university diktat, and coming from the same place as the earthquakes - California. All strangely reminiscent of, and perhaps symptomatic of, times of on-campus student unrest. What forced us into that "soil-dynamics religion" was the absence of a clear understanding of the mechanics of pore pressure generation. And the devil in the mix was the undrained triaxial apparatus which while defending the "established truths" went about its business of mutilating entrapped sand in a manner reminiscent of what was done to nonconformists during the Inquisition. To get out of that mindset, and progress, it was necessary to become a geotechnical heretic.

The first step was to walk away from orthodox belief in the interpretation of what went on inside the membrane of the undrained triaxial machines. What actually happens within this sealed environment is that the vertically moving plunger results in either dilative or contractive deformation of the soil-structure, and that in turn results in either more or less solid area being pushed into physical contact with the membrane. Let's take the case of a contractive soil-structure. As specimen straining continues the volume of the soil-structure diminishes, and with it the proportion of the solid phase in contact with the membrane. Because of this, the load carried by the water inside the membrane will have to increase accordingly in order to maintain radial/horizontal force equilibrium as the solid phase retreats more and more from membrane contact. In consequence the water pressure in the specimen goes up. And at the same time the effective intergranular normal stresses imposed by the membrane on the soil-structure comes down. The resulting loss of shear strength is not because of the pore pressure increases, it is because of membrane interference.

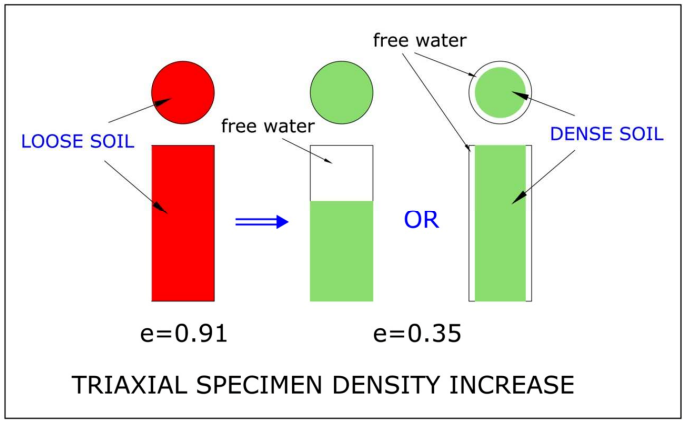

Figure 15: Contractive triaxial specimen

To clarify this important point I'll resort to a "*reductio ad absurdum*" style of reasoning. Figure

15 which shows to the left, to honest scale, the space required within the membrane to accommodate a mass of uniform spheres at their loosest packing (e=0.91). In the centre and to the right, the volume required by this same mass of spheres is shown for their densest packing (e=0.35). It is apparent that changing from the loosest to the densest packing (extreme contraction) must involve an increase in the proportion of the cell pressure conveyed to the water, with obvious consequences to the load bearing capacity of the soil column.

The idea that pore water pressure increases cause failure is simply wrong-headed. In fact, in terms of soil-structure stability, excess pore water pressure is not intrinsically a bad thing. But if it is changing in magnitude then it is a clear indication that the solid phase is trying to move through the liquid phase, and that things are not at rest. This is because deformation of the soil-structure results in the creation of pressures in the void water, and those responsive pressures act in a manner so as to oppose the movement of the soil particles. Essentially, the changes in pore water pressure are an effort of the system itself to rectify the situation; its own attempt to prevent movement and maintain the *status quo ante*.

In trying to visualize how the pore pressure generation mechanism works I found the analogy of a hydraulic piston helpful. I try to imagine what would be going on as a piston is being pushed into a rather leaky cylinder. Needless to say the piston is a particle and the leaky cylinder is the saturated soil-structure with drainage from a natural boundary some distance away. Leaving the unreal "undrained" condition behind us, and looking instead at an apparatus which does a fair job at representing soil behaviour in a natural setting and see if the "leaky piston" helps. What I have in mind is the laboratory consolidation machine, or oedometer.

For simplicity let's consider one-way drainage from an impervious solid base to an upper highly porous platen. When the consolidation force is applied to the platen, that force is transferred entirely to the topmost layer of particles, with the water continuum carrying virtually none of it. This is because, apart from having very little shear strength to provide bearing capacity, the water in physical contact with the porous platen can escape through it with very little resistance/effort. The soil-structure responds to the load by contracting into a more resistant intergranular arrangement. This involves all the particles moving towards the base, and this relative motion between the phases generates a pore water pressure field which grows in magnitude, particle after particle, until the solid base is encountered. At the base there can be no particle movement and therefore the pressure generation ends there. This generation of a hydraulic gradient within the specimen creates the required condition for seepage flow (leakage) from it. As consolidation progressed, and the soil-structure gets stronger, the rate of movement slows down, and with it, the generation of pore pressure. Eventually, the time comes when the soil-structure can carry the newly applied load without further movement, and consolidation leakage ends at this moment.

Calculating Pore Pressure Generation

Figure 16 illustrates the water forces generated by relative motion between the phases of a soil-structure immersed in water. As we are concerned here only with hydrodynamic forces, no effort is made to represent inter-particle forces on this schematic.

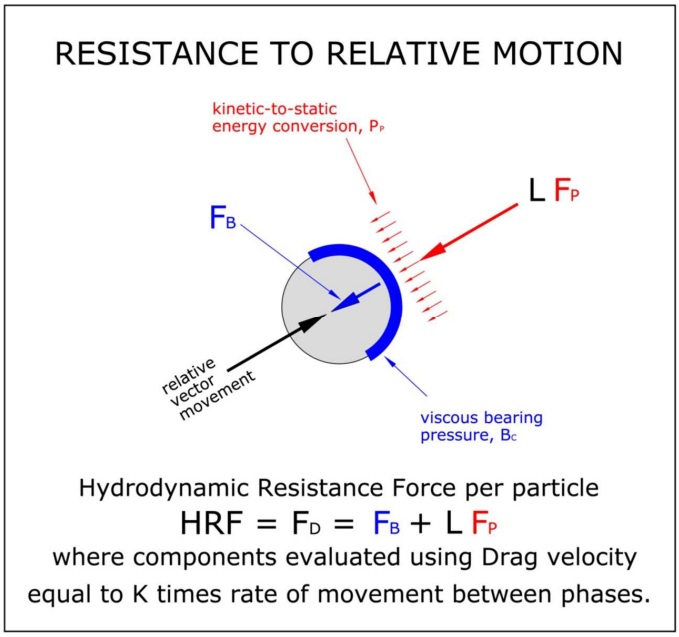

Figure 16: Resistance to relative motion

For a single particle, represented here by a sphere, the Crowding Factor, K is = 1. As discussed earlier, in the case of soil aggregations K > 1, where that value depends on particle sizes, packing density, and fluid velocity. Depicted here are the two component forces, viscosity and pressure, which together make up what I call the Hydrodynamic Resistance Force [HRF], and which I treat as the fundamental quantum of resistance offered by each soil particle to soil-structure deformation. It is here that the axiom "pore pressure is the response to movement, and not the cause of destabilization" is most clearly expressed.

In order to perform the tedious calculations required for determining the viscous drag and pore water pressures generated in saturated non-cohesive soil gradations, I wrote the computer program EPWPGRAD. This program is freely available from Geotechnical News as a Fortran compiled DOS executable file. Anyone who might want the source code can write me. The program works in the following manner:

A. The program requires the following input:

a) Soil gradation in terms of paired mesh size and percentage of soil passing that mesh for each of the soil fraction. In other words, the normal output determined during a sieve analysis.

b) Void ratio.

c) Water temperature.

d) Rate of relative motion between the phases.

e) Dimensions of a prismatic element of soil to be assessed.

f) Permeability of the soil if known; if not known a built-in subroutine PERMSOIL is used to estimate it.

B. The program then goes about the following routine:

g) For each soil fraction, an average size is used to determine the L-factor in this range. At the same time, the number of individual particles of this size within the soil element is found/calculated.

h) The permeability of the soil element for this particular rate of soil-structure movement is either taken as user input, or calculated by PERMSOIL.

i) The hydraulic gradient across the element is calculated from permeability, and then used to evaluate the element's Seepage Force in the direction of relative movement.

j) By a process of iterating on the void velocity, the unique overall value for the Crowding Factor, K, is found which would make the total Drag Force across the element numerically equal to the Seepage Force that would prevail for that same element of soil if it were subjected/exposed to the velocity of movement.

k) The magnitudes of the two components of HRF are calculated for each soil fraction using the L for that size, and the common K for the aggregation. By summing these components for each and every particle within the soil element the total force exerted against the upstream face of that rectangular prism is arrived at.

l) Energy and pressure gradients across the element in the direction of solid phase translation are then readily available as part of the program output.

As a point of interest, PERMSOIL goes about estimating soil permeability (hydraulic conductivity) in the following way. It takes as input the void ratio, particle size distribution, and water velocity being currently used in the parent/main program. It determines for itself the fluid (in this case, water) viscosity from the temperature given.

It uses the J.S. Kozeny (1931) inspired technique whereby an equivalent pipe diameter can be assigned to a particular soil aggregation. He

realized, quite brilliantly, that this could be justified by equating the Fluid Mechanics parameter, hydraulic radius, to the Soil Mechanics ratio of pore volume to surface area of all the grains. Once in the pipe analogy mode it is a simple matter to determine permeability from a combination of the Darcy-Weisback formula and the Colebrook equations for surface roughness (e/D = 0.05 adopted herein). Flows ranging from laminar to turbulent are assigned based on R_E, and where transient conditions are sometimes found to be appropriate for coarse sands.

C. The program provides the following output:

Figure 17 is a plot of three sets of data points produced by the computer program EPWPGRAD. The soil gradation is what I call fully proportionate, with a grain size ranges from 75mm down to 0.002mm. What I mean by fully proportionate is that each size is equally represented with respect to dry weight. In other words this is a perfectly well graded silt and sand and gravel.

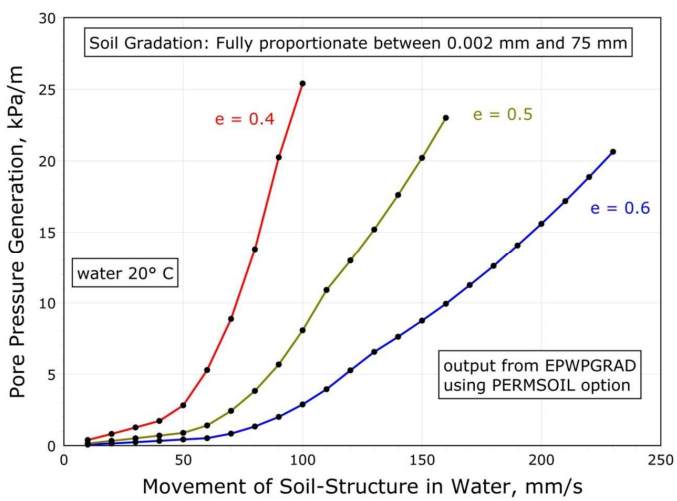

Figure 17: Pore pressure generation v. soil-water relative motion

A range of packing densities (e = 0.4, 0.5 and 0.6) was evaluated for the purpose of illustrating the strong influence of this parameter. The permeability for each of these "specimens" at 20°C was calculated for the appropriate flow type (laminar to turbulent) using the built-in PERMSOIL subroutine. The Crowding Factor found for these void ratios (respectively) were in the ranges: 6.4 to 10.9; 3.6 to 6.7; and, 2.3 to 4.5.

The plot in Figure 17 shows the theoretical relationship linking rates of movement between the phases with the pressure generated in the water phase as it opposes motion. The magnitude of pore pressure generation is shown in terms of gradient, and this is because it is built up, one particle after another in sequence, increasing progressively in the direction of the relative motion of the solid phase. It is only the pressure component [F_P] which is involved here, since the viscous component [F_B] cannot be seen by pressure sensitive devices.

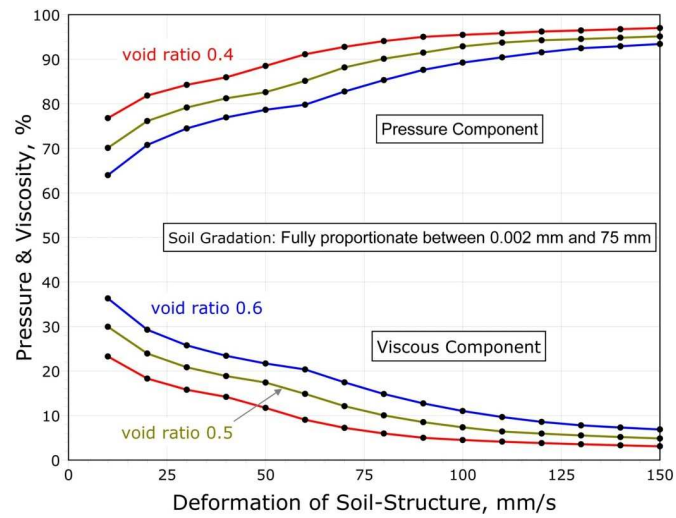

Figure 18: Pressure & viscosity % v. soil-water relative motion

Figure 18 shows how each of the two separate hydrodynamic components, that is, viscosity and pressure, contribute to the overall resistance. Here it may be seen that for the range of conditions depicted, pressure is the dominant component, and the contribution of viscosity becomes less as velocities increase and void ratios decrease.

Deformation

Up to this point the computations have been dealing with the type of motion that is best described as translational – the case of an intact, and unchanging, arrangement of separate particles which make up a stable soil-structure moving as an undisturbed fixture through water.

As seen in Figure 17 void ratio is a sensitive parameter in this context. And so, to put a number on the additional contribution made by soil-structure deformation to pressure generation, the procedure involves looking at void ratio changes, which are an accompaniment of deformation, as the key to the solution.

From the data plotted in Figure 17 it can be seen, for any chosen rate of relative motion, the pore pressure increases with decreasing void ratio. This is what we know as contractive behaviour. Similarly, it can be seen that dilation would cause pore pressure reduction; again, something which fits well with accepted and rational ideas. And finally, needless to say, if there are no void ratio changes then there is no reason to expect other than translational pressure changes in the pore water.

What this suggests therefore is that in order to evaluate the response of pore water to deformation we need to superimpose the effects of void ratio changes on those associated with translation of the intact structure as it moves relative to the fluid phase. As a consequence of this reasoning, the methodology I propose for the evaluation of pore pressure changes (either positive or negative depending on whether the soil-structure responds to the deformation in a contractive or dilative way) is to first create a translational plot, and from this, determine the magnitude and rate of additional pore pressures contributed by how the void ratio alters with time.

Hydraulic Gradients

The most important fact to bear in mind about pore water pressure is that a hydrostatic pressure distribution has no influence whatever on the behaviour of the soil-structure, however loose and unstable. I've been down to 135 feet in open water protected by no more than swim trunks and experienced no distress whatever because of the added 58 psi (400 kPa) of water pressure. And I'm sure enough that mineral grains of quartz aren't any more sensitive.

In still open water there is no hydraulic gradient. A deeply submerged slope will have very high pore water pressures, but they are irrelevant if there is no pressure differential across the bottom/ground, such as might be brought about by the passage of storm waves. Otherwise it is like standing in still air under atmospheric pressure. You are not aware of the air pressure - and, it isn't until a wind picks up that you know the air is there at all. So, having the pore water pressure at a single point in the soil is of little use to us. We need data from at least three separate points to know the hydraulic gradient magnitude and direction (vector) before we can have some idea about what might be going on.

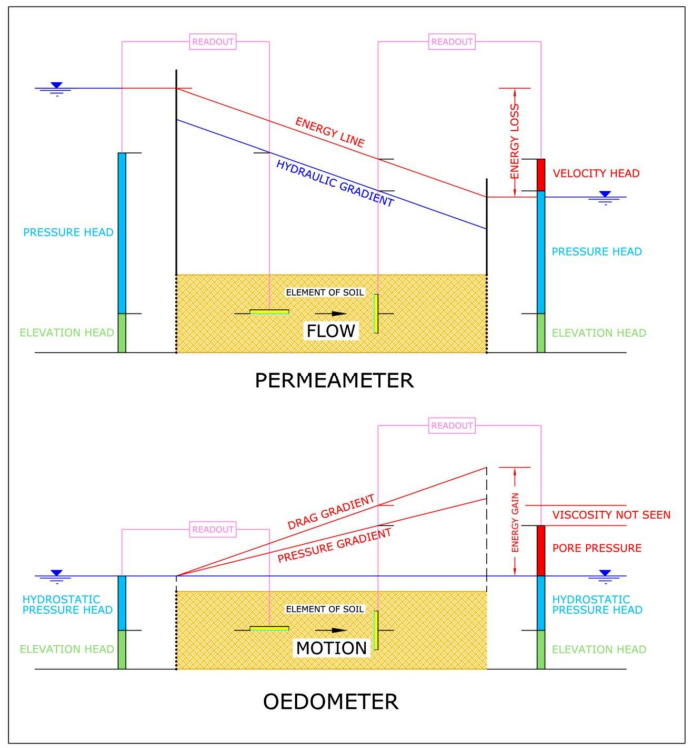

Figure 19: Pressure gradients in seepage and relative motion

Figure 19 was prepared to illustrate the similarities and differences between seepage forces and forces generated during two-phase relative motion. The top sketch shows a permeameter, which is the laboratory equivalent of steady state seepage, as for instance under a dam. The hydraulic gradient is the locus of the pressure head along the flow direction; it is parallel to the "energy line" and lies below it by an amount equal to the velocity head. There's nothing new here except that it shows the difference the orientation of a piezometer's sensor can make to the reading. If the sensor confronts

the flow the velocity head will register, otherwise it will not be seen.

The bottom sketch was constructed to make as clear a comparison as I can between the pressure losses of seepage and the pressure gains of motion. In this case I've used the oedometer as the laboratory equivalent of relative motion between the phases. Referring back to the free-body diagram in Figure 16 we see that relative movement between the phases results in hydrodynamic pressures being generated across the particle diameter: This constitutes an elemental hydraulic gradient. If we now consider the sequence of adjoining particles within a soil-structure we can appreciate that what results is in effect a continuous potential gradient. This, like the seepage gradient, is a vector, and the "motion head" will register on a sensor only to the degree that this vector is orthogonal to the sensing face.

The basic difference between the two sketches is that instead of the motion gradient being externally imposed, as is the case for steady state seepage, this gradient is built up from within, by virtue of the forces imposed on the water by the moving solids, one particle after another. In both cases the gradient increases upstream.

The real value of this comparison is to justify the assumption made in EPWPGRAD, and that is that the relative motion velocity used to compute the viscous and pressure component forces is equal to the approach velocity used in computing the Seepage Force. This is apparent so once the velocities at the left hand side boundary of both top and bottom soil elements are compared.

Practical Implications

If this new way of accounting for pore water pressure has any value then it should be able to give us some practical help in the various aspects of geotechnical engineering practice.

Laboratory

My views on the undrained triaxial compression test have been expressed already. As I see it, the intrinsic mechanical problem with this device, and the impression it gives that increasing pore water pressure leads to failure, is that within a sealed membrane, with no place for the water to flow, there is no possibility of a hydraulic gradient existing – the water potential is short-circuited. Here I am referring to things at the specimen and sensor scale. Of course at the microscopic scale, water must flow around individual particles as they are shuffled around, but the pressure sensor itself shorts these out.

In both the permeameter and the oedometer water is allowed to flow, and so there is no conceptual difficulty with regard to hydraulic gradients. Here it would be of interest to know the pressure distribution within the specimens. That would be a check on the notion that for smaller particles undergoing slow deformation the viscous component should predominate, thus reducing the measurable pressure component accordingly. And in consolidation testing, what if the particles are small enough that the only resisting force to motion is the viscous component? Would the motion be controlled by viscous creep rather than by seepage flow?

Site Investigation

The CPT probe penetrates the ground at a rate of 20 mm/s, recording the pore pressures caused by the cavity expansion straining at the tip. A typical trace of these on-the-run dynamic pore pressure responses [Bq] shows large swings from positive to negative Bq as the cone passes through contractive and dilative strata. A particular type of trace, which I have heard people call "hydrostatic", and is at first sight puzzling, presents a challenge to the hypothesis being advocated here to find an explanation.

These "hydrostatic" traces are apparently quite common in the sands of the Fraser River channel/delta. The name comes from the fact that the dynamic pore pressures follow a straight line coincident with the open water pressure line. In other words, the cavity expansion results in no pore pressure change. My interpretation, using the ideas presented in these articles, is as follows:

Bearing in mind that it is the cone tip which is moving and not the soil, the fact that there is no change in pore pressure in response to soil-structure deformation can only be because there is no change in void ratio during deformation. This is consistent with the classic case of constant volume straining. My opinion as to how sand could end up in this rather unique state of packing is that these sands were placed by bedload transport: a deposit moved so far and so often that it now exists at the constant volume void ratio. This suggests to me that these sands, although soft, are not liquefiable.

Slope Stability

Where steady state seepage conditions prevail within a natural slope, or on the downstream side of an embankment dam, we expect to see a loss of hydraulic potential as we go downslope. Since this reflects the water energy lost to particle drag forces as the water moves through a stable soil-structure, this is as it should be. This downhill drag is the destabilizing influence trying to flatten the slope. After allowing for seasonal alterations in differential head across the system, we don't expect to see that hydraulic flux change over time. And that is what we hope to see from any piezometers we have installed for monitoring the slope. A pressure distribution other than the established pattern would indicate either a change in the permeability of the section (soil erosion), or be a warning sign of movement within the slope.

The upper part of Figure 19 is an instance of steady state seepage through a stable soil-structure. It is the lower part of this sketch, based on the ideas introduced here, that provides some additional insight into what might be going on within the slope. This applies more to natural slopes which are often composed of fine grained soil, and thereby prone to a larger viscous component of energy, than to the coarser soils used in earthfill embankments. A piezometer will not see the viscous drag forces pulling down the slope, it will only show the pressure component. But in either slope, a change in piezometric head, not attributable to changing potential difference across the slope, is a definite warning sign: And this holds true whether the head increases or decreases.

Ground Improvement

Once we acknowledge the fact that escalating pore pressures are a result, and not the cause, of soil-structure contraction or collapse, then it comes time to look again at what we think we are doing when we install vertical drainage devices in the ground to enhance the groundwater's natural drainage. Certainly, in the case of non-granular compressible strata, we hasten ground settlement by such means as wick-drains. And that is a good thing. We are venting the pore pressures which are resisting and retarding downward movement. And when we install similar vertical drainage elements in what is feared to be liquefiable sands, it is exactly the same thing we accomplish: We speed up post-failure settlement. And that's about all.

Soil Grains are not Spheres

The numbers of particles of gravel, of sand, and of silt required to make up one cubic centimetre of soil are: One single piece of gravel would do it; 40,000 sand grains would be needed; and, for silt, the number is a staggering one billion. I used spheres of 10mm, 0.3mm, and 0.01mm diameter, at a relative density of 50% to calculate these ballpark figures. Now, since it takes a million ccs to make a cubic metre, then, whether they be silt or gravel, there is obviously no way of dealing with real soils in the field other than statistically. Even in the laboratory, determining the size and shape of each particle in a small specimen of sand is utterly impractical. And once it came to dealing with soil grains theoretically I knew I had no option but to simplify the shape to the extreme.

Spheres have the great benefit that their shape, no matter what way you approach them, is exactly the same. And their geometry is entirely defined by one dimension – diameter. The next simplest shape is a cube, again definable by a single length, but a cube looks different depending on the view point. Also, a cube has kinematic characteristics which are absent in a sphere and so difficult that my theoretical work is constrained to spheres.

The need to have a simple geometry which is amenable to mathematical treatment is more of a scientific necessity than an engineering one. It is a fundamental tenet of scientific advancement that propositions describing natural phenomena be expressed in mathematical terms. That way the reasoning can be fully and continuously traced through the mathematical formulation: It makes the proposition more amenable to falsification, and allows it to be either dismissed or subsequently built upon by others. As this is a new idea I'm presenting here, it needs to leave an uninterrupted mathematical trail behind.

Before ending this series I would like to draw attention to something I've being watching myself with some amusement while writing these articles. And that is the difference between the pore water pressure equation I suggested in Part 1 and what I am offering now. There's a substantial difference. Even before Part 1 went to print I knew it could be improved on. But I decided to leave it alone, and let it stand. I wanted it to be a benchmark for myself, to see how much my ideas would change over the year and a half it took to get to where I find myself at now: The end of this series of articles.

Acknowledgements

It is quite necessary for me at this juncture to mention my good old friend Nigel Skermer. He has willingly read the drafts of each of these articles and I am most grateful to him for guiding me towards some measure of logical continuity, and we hope, avoiding serious lapses in reason. Thank you, Nigel.

Thanks also to John Gadsby for agreeing to publish this series, and of course, to Lynn Pugh for keeping me in line throughout the effort to put the ideas into print.