GEOMODELING: A TEAM EFFORT TO BETTER UNDERSTAND OUR RESERVOIRS

Part 5: Petrophysicists and Geomodeling

By Thomas Jerome, RPS, J. Kim McLean, independent petrophysicist and Laverne Chan, Enerplus

INTRODUCTION

Can you export for me the vshale (VSH), porosity, water saturation (SW) and permeability (PERM) logs? Thanks in advance for your help!

Such request is too often the only type of collaboration between petrophysicists and geomodelers. Of course, it's not true in every company nor every project, but it does happen more often than it should. Geologists, geophysicists and geomodelers can talk for days, if not weeks, about the structure, the stratigraphy and the facies distribution. But we tend to order logs the way we order take-away at a restaurant: we want our order fast, we want it in LAS format and we don't really care about how the logs were cooked in the kitchen.

That's unfortunate. All the properties that we need to model in 3D (porosity, SW…) are some expression of the rocks' and fluids' characteristics such as pore size, mineralogy and/or fluid composition. They are here to quantify how the rocks will behave when we attempt to produce and all these properties that we need to model are inter-connected one to the other to some degree. A good petrophysical analysis will properly capture the correlations between these different reservoir properties. As geomodelers, we must also respect these correlations when we model petrophysical logs in 3D. For example, if a rock's characteristics are such that low porosity values are necessarily correlated to low permeability and high water saturation, we must make sure that we respect this relationship in every cell of our 3D grid. We don't want to have cells with a mix of interpolated low porosity values, high permeability and low water saturation. It would contradict the nature of the reservoir and it might jeopardize the analysis engineers will derive from our models.

As geomodelers, we tend sometimes to consider that the challenges of the geomodeling process are building the grid and populating facies. Once this process is properly done, populating petrophysics is "easy". The aim of this paper is to show that an appropriate amount of time must in fact also be dedicated to distributing the petrophysical analysis in 3D correctly, and in a way that respects the nature of the rocks. This goal is easier to achieve if we keep our petrophysicists in the loop instead of trying to do it on our own.

Firstly, this paper reviews some basic definition of VSH, porosity, SW and permeability before looking at how the logs are being upscaled (blocked) into the 3D grid. It then looks at a question sometimes asked: will it really change something if I was just to model my logs in 3D without losing time modeling facies? At last, 3D petrophysical modeling itself is covered as well as some aspects of integrating petrophysical uncertainties into the geomodel.

SOME DEFINITIONS

Petrophysics is a vast domain which can't be summarized in a few paragraphs. Each rock requires specific petrophysical equations, or at least specific values for the parameters of these equations. Also, as for any other sciences, practitioners have developed multiple techniques to compute any given petrophysical log, different techniques requiring a different set of input logs or different techniques being based on a different conceptual understanding of the intrinsic characteristics of the rock. At last, the list of reservoir properties requested by engineers might be different from one study to the next (conventional vs unconventional, natural fractures to be modeled or not…). As it is impossible to cover everything in one paper, we have decided to focus on a simple set of petrophysical logs: VSH, porosity, SW and permeability. Discussion between your geomodeler and your petrophysicist will allow you to adapt what is proposed hereafter to your specific reservoir and your specific petrophysical analysis. In complement, we invite you to refer to one of the many excellent books available for more details about well logging (Ellis and Singer, 2007) and petrophysics (Doveton, 2014).

The purpose of well logging is to measure how the rocks and fluids in the vicinity of the wellbore react to different stimuli such as electricity (resistivity logs…), nuclear emission (density logs…) and acoustic waves (sonic…). Well logging provides an essential input to formation evaluation but also to completion evaluation. Among other things, once processed, cleaned and calibrated, well logs help to answer questions about the nature, location and quality of the hydrocarbons. Well logs are also used to quantify important rock and fluid properties such as VSH, porosity, permeability and SW, through the process of petrophysical analysis. These petrophysical logs are in turn needed as input to engineering studies, and so as input to geomodeling.

VSH quantifies the part of the rock that is "ineffective" in the sense that it will slow down or prevent the good flow of the hydrocarbons. GR is used as the base line for quantifying VSH as some clays such as illite contain radioactive minerals such as potassium which are detected by the GR log. Some clays such as kaolinite does not contain any of the radioactive minerals (potassium, uranium, and thorium), detected by the GR log, and are therefore not visible to this log. The presence and proportion of each type of clay, among which those not visible to GR, can be observed and measured in core samples and then used to correct the VSH logs as needed.

Two main types of porosity are defined: total porosity (PHIT) and effective porosity (PHIE). PHIE is the percentage of the rock volumes which represents the connected porosity. It is made of all the pores that are connected and form the pore network. Only fluids in the pore network can be moved through production (unless some recovery techniques are used to connect some of the non-connected pores to the pore network). The total porosity is the fraction of the rock made of all the pores, both connected and non-connected. PHIT includes PHIE and mathematically PHIT is always greater or equal to PHIE.

Total water saturation (SWT) and effective water saturation (SWE) are associated respectively to PHIT and PHIE. The water saturation represents the percentage of the pore volume (total or effective) filed with water.

Most engineers want to see PHIE and SWE modeled in 3D, as these are important input parameters to their own computations (flow simulation, volumes). That being said, some recovery techniques are influenced by the total pore volume, not just the connected pore volume. As such, geomodelers should always check with their engineers what type of porosity and SW are needed (if not both). It should also be noted that initial SW is needed for flow simulation and the current SW might be different to the initial SW. Initial means the state of SW in the reservoir before any production or enhanced recovery technique occurred. SW logs based on well logging done after EOR techniques, like waterflooding, were applied must be corrected or discarded. Waterflooding will increase SW in the vicinity of the water injectors and the SW logs will show higher values than were prevailing in the initial condition of the reservoir.

The discussion about which the type of porosity and SW needed (effective and/or total) must involve the petrophysicist as this request might change the way he/she will complete his/her petrophysical analysis.

For some reservoirs, measuring total properties on core is more easily done (or simply feasible) than measuring effective porosity. For some reservoirs, it is the reverse. The key to a good petrophysical analysis is the capacity to calibrate it back to some reliable core measurements. If core studies measured total properties, the petrophysicist will compute PHIT and SWT from the logs, calibrate them to the core measurements, and then, if needed, the petrophysicist will derive PHIE and SWE properties from PHIT and SWT. If core studies measured effective properties, the petrophysicist will do it the other way around. VSH is one of the parameters linking PHIE and PHIT. Any uncertainty on VSH (because of clays not seen by the GR for example) will add uncertainty in the porosity computation.

SW is the most difficult log to quantify because the coring process doesn't preserve SW well. SW core measurement can be very unreliable. In conventional reservoirs, capillary pressure is extremely useful to quantify SW, but such technique can't always be applied in unconventional reservoirs. Also, SW being partly connected to the porosity, any porosity uncertainty (coming from VSH uncertainty among other things), will also create uncertainty in the estimation of SW.

Permeability values change depending on the direction considered. The vertical permeability quantifies how the fluids move in the vertical direction, while the horizontal permeability quantifies how the fluids move in the plane of deposition. For flat to near-flat reservoirs, "horizontal" does mean horizontal. In complex structural reservoirs (folds, faults, tilted blocks…), the plane of deposition might not be "horizontal" anymore. The same can be said for the "vertical" permeability. For reservoirs with natural and/or induced fractures, a more complex horizontal permeability field must be considered as fluids will move with more ease in the direction parallel to the fractures (if they are open) than perpendicular to them. Geomodels usually contain two permeabilities (vertical, horizontal). Some studies consider three permeabilities when two orthogonal horizontal permeabilities are needed. In this paper, we consider just one permeability as the message is about the general workflow, not about the specifics of any given type of reservoir. A discussion in your team will clarify what is needed in your case.

For consolidated reservoirs, permeability can be measured on core with some level of confidence. Cross-plots between core porosity and core permeability are used to define a mathematical relationship between the two properties. Once porosity logs are available (and calibrated to the core porosity measurements), the mathematical relationship porosity-permeability is used to compute permeability logs from the porosity logs. Permeability is very hard to measure, especially in non-consolidated rocks. As such, the porosity-permeability relationship contains a lot of uncertainty.

DEFINING THE RESOLUTION OF THE 3D GEOLOGICAL GRID BY UPSCALING THE WELL LOGS

One of the first tasks of a geomodel project is to decide at which resolution the model will be built. Shall we use cell size of 100m*100m horizontally? Finer maybe? 50m*50m? Coarser, such level of details being unnecessary? 250m*250m maybe? And what about the vertical cell size? 5m? 1m? 0.1m? Some geomodelers use the resolution that engineers will need for their 3D simulation grid. Simulation engineers might for example decide up-front for a cell size of 100m*100m by 1m vertically in the 3D simulation grid because it will limit the number of cells to a level manageable by the flow simulator. Experience shows though that it is wiser to select a cell size based on the expected spatial heterogeneity of the reservoir, and especially of the facies. If the reservoir contains geobodies of a few hundred meters width, then it is a good idea to have smaller cell size, maybe at 50m*50m or even less. There will always be time later to upscale the whole 3D geological grid to the resolution the engineers need for their own 3D grid. At that time, the then-built geomodel might even prove that the cell size asked by the engineers might oversimplify the complexity of the reservoir. The upscaling of a geological grid into a simulation grid will be covered in the paper on Flow Simulation and Geomodeling, in two issues from now. For the vertical cell size, we should use the upscaling the well logs as a means to decide what to do.

Facies data and petrophysical data are all physically stored as data along each of the object well of our project. In the meantime, we need to use these data to populate a different physical object: the 3D geological grid. While some geomodeling packages allow running geostatistics with wells directly as input, it is wiser to first upscale the facies and the petrophysical data into the 3D grid (step also known as blocking the well data in some packages – both terms are used hereafter). The 3D geological grid must be refined enough to capture the spatial characteristics of the reservoir, and this starts with respecting the characteristics of the reservoir along the wells. The logs are upscaled to the resolution of the 3D grid and the upscaled logs are compared to the original well logs. If they are close enough, the vertical cell size is good. If the upscaled logs have lost too much important detail shown on the original logs, the 3D grid must be refined vertically.

A typical workflow can go as follows. Firstly, the well data (facies and petrophysics) are analyzed to find the average thickness of the facies. This thickness gives an original vertical cell size. The well data are blocked into the 3D grid and compared to the original well data. At this stage, we can face two situations. If it is matching perfectly, we should go with a coarser cell size (maybe 2m instead of 1m originally) and we repeat the process until we reach a level where the upscaling doesn't respect the log data well enough anymore. The vertical cell size to use is the last one that worked well: it is the coarser cell size that respects the well data. On the contrary, if there is not a good-enough match, then the original vertical cell size was too coarse, and we need to refine progressively until no improvement can be seen. Once the vertical cell size captures the data resolution very well (ie going any finer doesn't improve the resolution, while it needlessly increases the number of cells in the model) we have found the cell size we need.

Comparing the original logs to the upscaled values can be done qualitatively or quantitatively. If the project contains only a few wells, displaying original and upscaled values side by side on a well display is a good way to analyze the results. A snapshot of such a display can help to explain this process in a report. When the project contains many wells, such qualitative approach becomes tedious if not impossible. An alternative is to compare the statistics of the original logs with the statistics of the upscaled values. If the upscaling worked well, then the original and the upscaled logs should have similar distributions and so similar percentiles, mean and standard deviation. Such analysis is enough to validate the vertical cell size. Nevertheless, it might be more interesting for your team that you analyze also the statistics of the meta-data of each well. Under the term meta-data, we are

(... Continued on page 16)

(... Continued from page 15)
grouping any type of computation your team has done on the well logs, independently from your work on the geomodel. Maybe a net-pay thickness has been computed at each well as well as a net porous thickness and an oil-column thickness. Those numbers computed at each well might have already been used by your team to make decisions about the next steps of the whole project. Proving that the blocked well data does respect these crucial results is a great way to have your team support your geomodeling project.

Many mathematical techniques exist to block the well data. Usually, the vertical cell size (25cm to 1m) is coarser than the log resolution (10cm if not less). So all these techniques correspond to some sort of averaging. For discrete properties like facies, the most common approach is to keep the facies which was "the most preponderant". For example, if in a given cell of 50cm height, the well shows 40cm of Sand and 10cm of Shale, it makes sense to assign the facies Sand to this cell. Is the "loss" of 10cm of Shale important? As long as it doesn't impact the statistics along the well, probably not.

Once the facies are blocked, the petrophysical logs are upscaled. It makes sense to do it in this order as each facies usually shows a specific range of values for each petrophysical property. But what values shall we block? Only those associated to the blocked facies or shall we average over all the values, no matter what facies they belonged to? In the previous facies example, shall we defined a blocked (averaged) porosity from the porosity values only from the 40cm of Sand, as Sand was defined as the blocked facies value? Or shall we do the averaging by including also the values from the 10cm of Shale? In the first case, we make sure that a blocked Sand has a value of blocked porosity that belongs to the expected range of porosity for this type of facies. But by doing so, we have increased the average porosity in those 50cm. In the second case, we use all the values, so the averaged porosity is a closer representation of the spread of porosity along these 50 cm. But we have now a "dirty" blocked Sand with a value of porosity belonging neither to a Sand nor to a Shale. Each approach has its pros and cons. To get past this problem, we suggest the following approach. Firstly, apply the idea that if a Sand is the blocked facies in a cell, then the average values of the petrophysical properties are defined from the portion of the well which was a Sand to start with. Then analyze your statistics. If everything is well and fine, you have made the right call. If the reservoir is laminated though and many blocked cells shows that the input facies log is close to 50/50 percent of Shale and Sand, the cell size might simply not be appropriate: you need a finer resolution. For an extremely laminated reservoir, this fine resolution might not be a valid option though, unless you are ready to work with a 3D geological grid made of tens if not hundreds of millions cells. An alternative is to revise the facies description on the wells and see if all the zones showing lamination of Sand/Shale can't be renamed as a third facies Laminated Sand or Shaly Sand (for example). This new facies is understood to be a mixed of Sand and Shale and as such it shows petrophysical values being averages of those found in pure Sand and pure Shale facies. In this example, moving from a 2-facies classification to a 3-facies classification allows going around our problem. The cells with 50/50 Sand/Shale are now 100% made of this third facies. Once the new classification is applied, we can go back to upscaling the petrophysical logs by facies.

With modern geomodeling packages, it is possible to run the process of well upscaling in a few minutes. Nevertheless, we recommend that you work on this step with great care as any error here will be felt everywhere for the remaining of the project. This step should be used to validate the vertical cell size as well as optimizing the blocking. It makes sense for the geomodeler and the petrophysicist to work together on this step, defining the objectives, which statistics to look at, which numbers to match, and ultimately proving to the team that the blocked data are a very good starting point to both 3D facies and petrophysical modeling in the geomodeling project.

DO WE REALLY NEED TO MODEL FACIES?

The workflow promoted in the whole series is to populate facies in 3D and then to populate petrophysics by facies. The present section gives us the occasion to illustrate what might go wrong if the petrophysics is interpolated without a facies framework in place.

Let's consider two wells with facies and VSH information (Figure 1). Well 1 shows 50m of Shale at the top followed by 30m of Sand, 20m of Shaly Sand and at last 10m of Sand. Well 2 is Sand all along except for 20m of Shaly Sand near the base of the reservoir. The Shaly Sand is interpreted as a continuous facies zone across the area while the Shale is believed to be of limited lateral extent. No data can tell where the Shale exactly stops between the two wells though. Each facies shows a specific range of VSH: the Sand is associated to VSH less than 30%, the Shaly Sand has VSH between 30% and 70% and the Shale is associated to VSH superior to 70%. We need to model the VSH and the facies in 3D while respecting our interpretation of the reservoir. Ideally, we need to be able to create multiple models showing different lateral extent of the Shale between the two wells. For an initial path, we would like a Shale ending approximately half-way between the two wells. To do the modeling, we have a 3D grid refined enough to capture these three facies' spatial distribution.

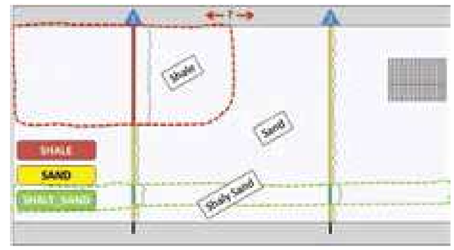

Figure 1. Schematic Shale, Shaly-Sand, Sand reservoir. Input data (Facies + VSH) and geological interpretation.

We decide to test two different modeling approaches. On one hand, we model VSH in 3D with a geostatistical algorithm and then we apply the VSH cut-offs to assign a facies code everywhere (Figure 2). On the other hand, we follow what we have advocated since the beginning: we model the facies in 3D with a geostatistical algorithm and then we model the VSH by facies, also with a geostatistical approach (Figure 3). Which approach gives the more reasonable results?

Both models show the proper relationship between facies and VSH (we don't have high values of VSH in the Sand for example). Both models also show the same lateral extent of the lower Shaly Sand zone. But the two models differ drastically in the upper part of the unit. When we model the facies first, we have what we expected: a sharp transition between Shale/Sand approximately half-way between the wells (Figure 3). We also know from experience that if we were to use a smaller variogram range for modeling the Shale, we will be able to "place" the limit Shale/Sand closer to well 1 or to well 2. We have an algorithm with which it will be easy to manage this uncertainty in multiple realizations. On the contrary, the other approach creates a lateral transition Shale/Shaly-Sand/Sand that we didn't interpret to start with (Figure 2). Why is that?

VSH is a continuous property which can takes a continuous range of values between 0% and 100%. All interpolators of continuous properties (as far as we know) create by definition models which vary continuously across the 3D grid. In the upper part of the unit, VSH will necessarily change progressively from 80% in average in the Shale on well 1, to 70%, 60%, 50%, 40% and 30% to finish around 20% in average in the Sand of well 2. It is unavoidable. When facies are calculated based on the cut-offs, it creates this halo of Shaly-Sand not observed at the wells. If the well data were showing that Shale and Sand are never in direct contact but instead are always separated by a zone of Shaly Sand, and VSH smoothly transition from high values to low values, then seeing Shaly Sand and a smooth VSH transition between the wells would be reasonable. But that is not what the wells are showing. Well 1 shows that VSH changes without transition between high and low values (limit Shale/Sand), but also between low and average values (limit Sand/ Shaly Sand). The smooth VSH model (Figure 2) and the forced succession Shale – Shaly Sand – Sand are not backed up by the well data.

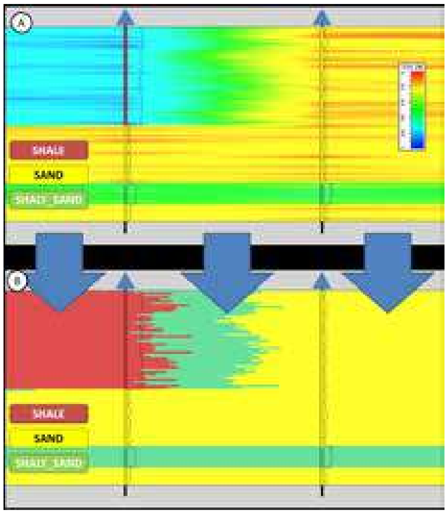

Figure 2. Initial model. Modeling first the VSH (A) and then the Facies from the VSH model (B).

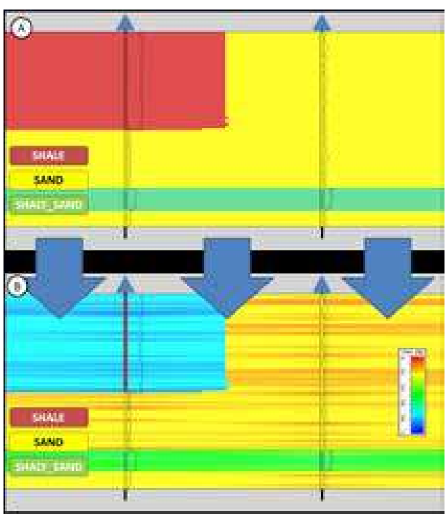

Figure 3. Improved model. Modeling the Facies (A) and then the VSH by Facies (B).

Even for models which must ensure a specific lateral succession of facies, modeling facies first makes more sense as we have algorithms in our geostatistical toolbox which gives us a lot of control on the process (plurigaussian simulations). It would still be much more preferable than modeling the VSH first.

In conclusion, only model continuous properties such as VSH, porosity, SW or permeability within domains in which they smoothly vary. Those domains will often be based on the subdivision of each geological unit into different facies zones. So, model your facies first and then model the continuous properties by facies.

MODELING PETROPHYSICS IN 3D: RESPECTING THE ROCKS AND FLUIDS CHARACTERISTICS

We have validated our choice of cell size. The well data are blocked in the 3D geological grid and we are even comforted in our workflow. It is time now to distribute the petrophysical properties… by facies of course .

The previous two papers talked about the geostatistical algorithms in general terms and about some specifics for modeling discrete properties. Much of what was covered previously also applies to modeling continuous properties such as VSH or porosity. The main algorithms are again kriging for interpolation and Gaussian simulation for generating multiple realizations. Both use variograms and some statistical parameters as input (mean value for the kriging, distributions for the simulation algorithms). Both can also take into account different types of secondary variables. Vertical Trend Curves (VTC) are the equivalent of the Vertical Proportion Curves (VPC) for discrete properties. They capture how the mean of a given continuous property varies with depth (see Figure 5 and Figure 11 for some examples). By depth, we mean by horizontal layer in the 3D geological grid. For example, VTCs would spot a decrease of porosity with depth due to compaction. Sometimes, continuous properties also show horizontal trends. Maybe the VSH increases from the North East corner to the South West for example. Such information can be captured as a map which can be taken into account by the geostatistical algorithms, in the way facies proportion maps are used to guide the modeling of facies. VTCs and trend maps capture how a property should vary spatially. It is also recommended to look at how properties are correlated one with the other. For example, if a cross-plot porosity versus VSH shows that the two properties are highly correlated, we can model first the VSH and then use the 3D distribution of VSH as a guide (a secondary variable) to model the porosity in 3D.

Ultimately, the key is to understand how each property varies spatially, how the different properties are connected one to the other and to convert all this knowledge into secondary data for the geostastistical algorithms. The geologist can help interpreting spatial trends while the correlations between properties are of course deeply rooted in the petrophysical analysis. While a geomodeler might guess them all properly, it is much more efficient to identify the meaningful correlations with the team's petrophysicist.

We illustrate the whole workflow hereafter with a real dataset. Carbonate or sandstone, conventional or unconventional, Canadian or international, none of this matters here as the general methodology can be easily adapted to any reservoir. For that reason, we won't detail any specifics. Similarly, it is nearly impossible to read the values on the axes of many of the associated pictures. This is done on purpose as the key is to focus on how one graph compares to another globally. Specific numbers won't add anything. You will be able to adapt it to your project by working with your team to identify what petrophysical properties should be modeled. In fact, we might have as well called our logs A, B, C and D instead of VSH, porosity, SW and PERM for this example.

Looking at a geomodel in 3D views, in cross-section and in map view is important of course. Such views must be used to validate the 3D petrophysical distributions of course. But sometimes, we conclude that because it looks good in 3D (or 2D), it is correct. This conclusion can be misleading when dealing with petrophysical modeling. In the example below, consider that every distribution we created did "look good" in 3D. Such displays did not help us improve our workflow. For that, we had to rely on three other types of views, which we believe geomodelers do not always use as much as we should: histograms (Figure 4), VTCs (Figure 5) and cross-plots between the different petrophysical properties (Figure 6, Figure 7 and Figure 8). On each of these three displays, the black colour is used for the original log data, while the blue and the red are used respectively for the blocked data and the 3D model. For each type of display, we should look at how the display from the original log data compares with the display of the blocked values and then how both compare with the display of the 3D distributed values. The first comparison relates to our discussion in the section about upscaling the well logs: a good well blocking will respect the original statistics of the different properties. Similarly, a good 3D petrophysical model should also respect the input statistics. By lack of space we did not include figures nor any discussion about trend maps, but those should be looked at as well, at the least to check if such trends exist or not, and if some do, to use it/them.

In this example, an initial modeling workflow is defined, run and the computed model is analyzed (Figure 4 to Figure 8). A few problems are found which lead us to modify the workflow. Figure 9 to Figure 12 illustrate the improvements we gained in doing so. This loop should be

(... Continued on page 18)

(... Continued from page 17)

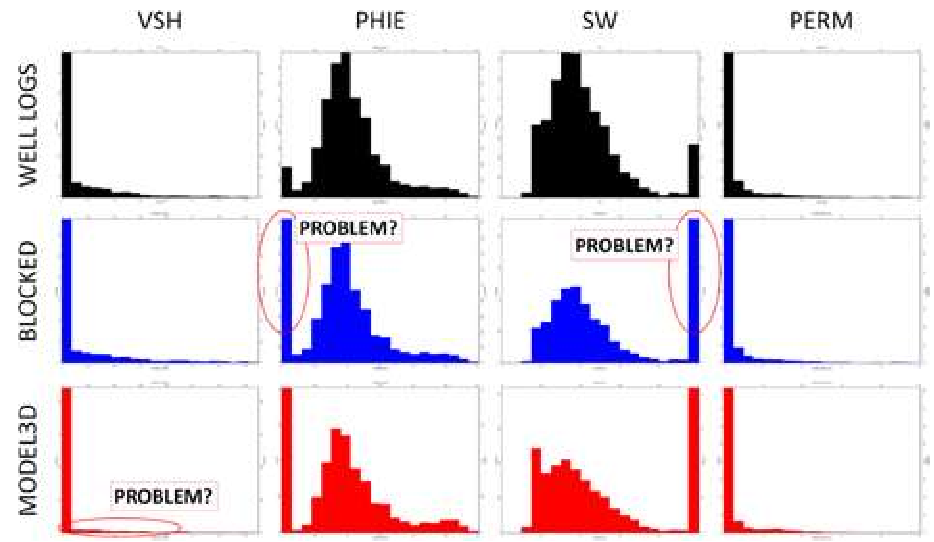

Figure 4. Comparing the distribution of each petrophysical properties (VSH, Porosity, SW and Perm) as logs, as blocked values and in the 3D petrophysical model.

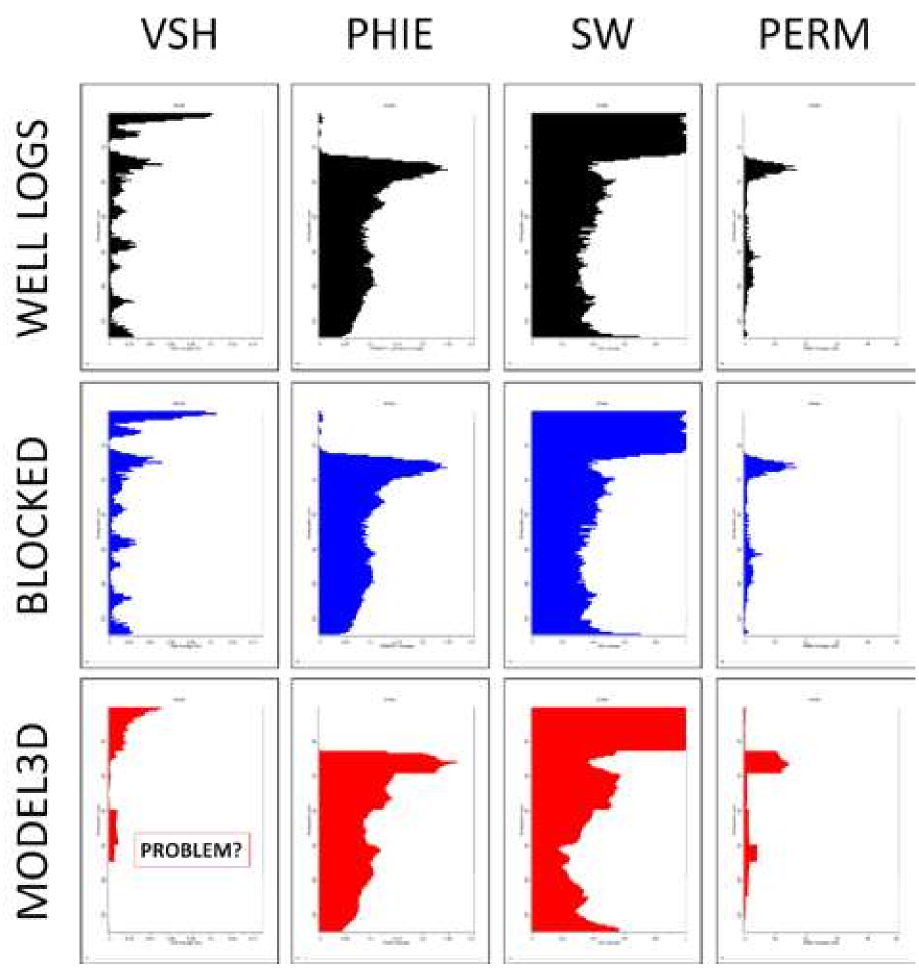

Figure 5. Comparing the VTC of each petrophysical properties (VSH, Porosity, SW and Perm) as logs, as blocked values and in the 3D petrophysical model.

repeated until no more fine-tuning is needed and the team considers that the geomodel does properly capture the characteristics of the petrophysics in this reservoir.

Firstly, the well data were blocked. The distributions of the blocked VSH and PERM are very similar to the distributions of the original well logs (Figure 4, first two rows). The blocked porosity and SW are close as well, but they show some suspicious bars at (porosity=0%) and at (SW=100%). This is the kind of thing to look for; discrepancies between original and blocked statistics. In this case, it is normal in fact: a post-process was applied to assign these values in some specific K layers. As it relates to something specific to this dataset and it does not add anything to the overall workflow, no more details will be given here. The VTCs (Figure 5) and the cross-plots (compare Figure 6 and Figure 7) confirm that the upscaling was done correctly as the plots on the blocked cells match the plots of the original logs. The cross-plots of porosity-permeability show that the permeability is obviously a mathematical function of porosity. At this time though we will ignore that fact and we treat the permeability like the other three properties. It can be spotted that something is a little suspicious already though: the cross-plot of the blocked values show that the mathematical relationship is not respected at 100% anymore (Figure 7, zoomed in area). But as it concerns only a few points, we decide to ignore this at first.

Gaussian Simulation was used for modeling VSH, porosity and SW, both in the original workflow and the modified workflow. What changed is which secondary variables we used. In the original workflow, VSH is modeled without any secondary variable, while their respective VTCs were used as input for modeling both the porosity and SW. The initial workflow did not use any of the cross-plots as input. In the initial workflow, permeability is also modeled by Gaussian Simulation, without any secondary variable (one might have argued that using porosity could have made sense).

The statistics of the resulting model of VSH, porosity, SW and PERM are then compared to the statistics from the blocked values. The histograms are well preserved (Figure 4) expect for the VSH. The VTCs for the porosity, SW and the PERM are also well respected (Figure 5), which is interesting for the permeability, considering that this VTC was not used as input. This illustrates something important about geostatistics: always run your workflows without secondary variables first; then add the secondary variables and see their impact on the model. In some instances, the well data are all you need to get it right. In others, the statistics of the model without secondary data will contradict the secondary data you intended to use. Which one is correct then? It depends on the reservoir. At the least, you should look into it instead of just adding the secondary variable without giving it any thought. In our example, the VSH model looks problematic again as its VTC really does not match the vertical trend from the data. At last, we look at the different cross-plots (Figure 8, the blue dots are the points from the blocked data; they are displayed here to help visualize the differences with the cross-plot from the

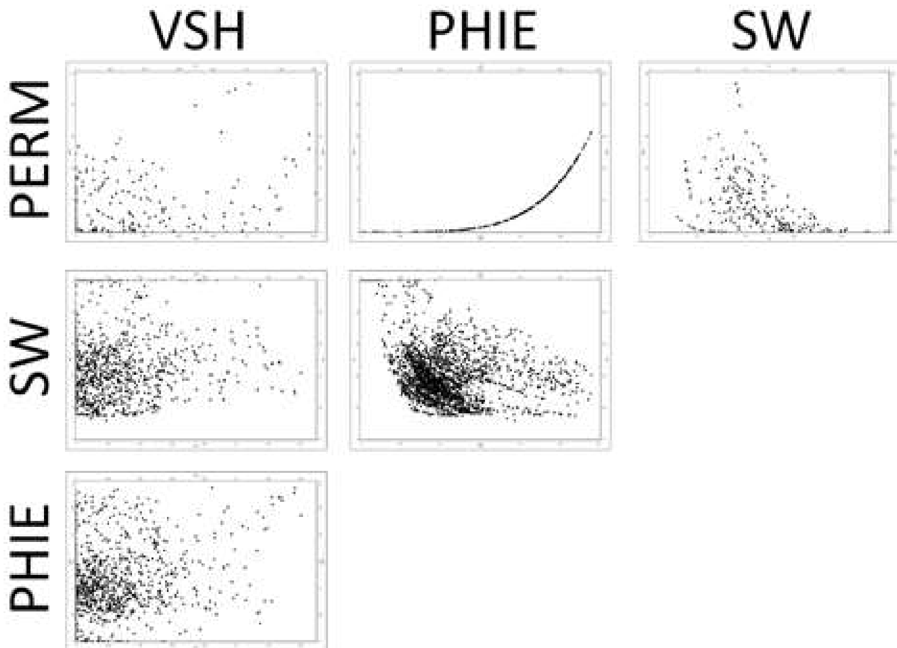

Figure 6. Cross-plots between the different petrophysical properties (VSH, Porosity, SW and Perm) as logs on the wells (reservoir zone).

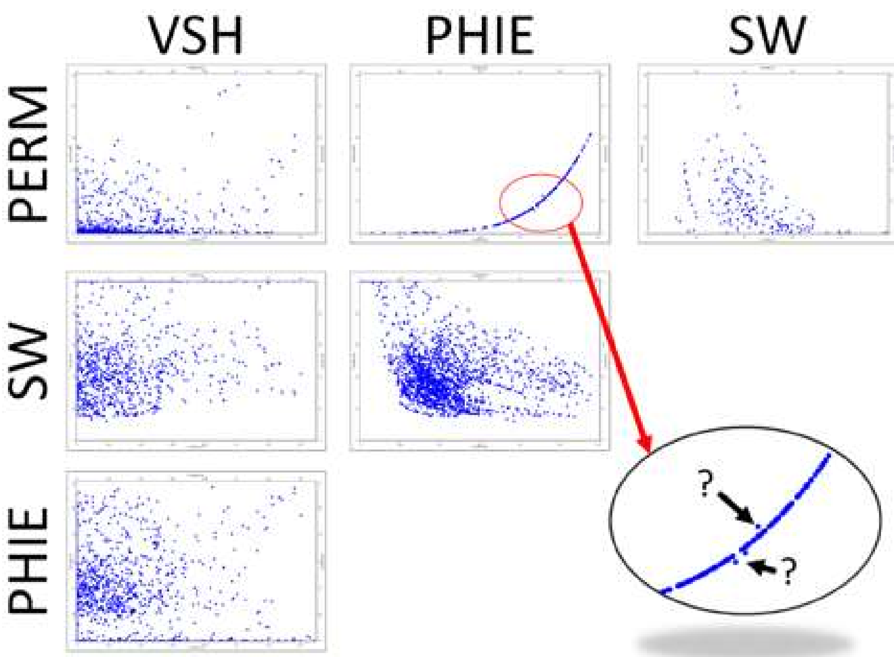

Figure 7. Cross-plots between the different petrophysical properties (VSH, Porosity, SW and Perm) as upscaled values in the 3D grid.

3D modeled distributions). Most of the cross-plots do not match the relationship seen on the input blocked data. This seems normal as we did not use any cross-plot correlation as secondary variable. In the first path at editing our workflow, we will focus on the problem with two cross-plots: porosity versus SW and porosity versus PERM.

For the first cross-plot, the blocked data showed a clear negative correlation while the modeled properties show a positive correlation. This means that we have cells with values of porosity and SW that are not supposed to be associated together in this reservoir. We need to fix this incorrect association. We decide to ignore the SW VTC and instead to use the porosity model as a secondary variable. The porosity is modeled as before, with the porosity VTC as a secondary variable.

For the second cross-plot, the blocked data were showing that the permeability was defined as a function of porosity (as confirmed by our petrophysicist). We can still recognize the relationship in the cross-plot from the 3D model. But many cells do not respect it at all. Of course, this might be fine as the relationship itself has some uncertainty in it. So the current cross-plot could be the sign that we have included some uncertainty in the model. This is true, but not what we intended, and clearly we did not control it. If uncertainty might be needed, maybe it is better to apply a range of mathematical relationships. As stated earlier, using the porosity model as a secondary variable would improve the permeability model, but not as well than if we simply apply directly the mathematical relationship PERM=function(porosity) defined by our petrophysicist. Once we have modeled the porosity in 3D, we apply the function to get permeability everywhere in the 3D grid. This approach is applied in the modified workflow.

In this path, we do not edit how the VSH is modeled. It would still need to be taken care of in another loop.

Analyzing the edited SW and PERM models shows that the problematic cross-plots are now cleaned (Figure 9 for porosity-SW, Figure 10 for porosity-PERM). Interestingly, the new SW model still respects well, and in fact slightly better, the vertical trend captured by the VTC (Figure 11) than the original SW model did, even while this VTC is not used anymore to model the SW. When looking into it, it does make sense. Porosity and SW are negatively correlated: in average, the higher the water saturation, the lower the porosity. The porosity and the SW VTCs show similar relationships (Figure 5): if a K layer has a high average water saturation value, it tends to have a low average porosity value. Our modified workflow still respects the porosity VTC and it now uses the porosity-SW correlation. This modification turned out to be enough to see the vertical trend from the SW VTC respected too.

Similarly, it is interesting to see that the cross-plot of SW-permeability has improved too, even if it was not used as input. SW and PERM have now better distributions and a side-effect was to fix that cross-plot as well.

As mentioned earlier, the work on this specific workflow should be continued and maybe looped a few more times to see if certain displays can be improved even more. Fixing the VSH would be a priority. This work is beyond the scope of this paper as it would mean applying the same approach already described here.

(... Continued on page 20)

(... Continued from page 19)

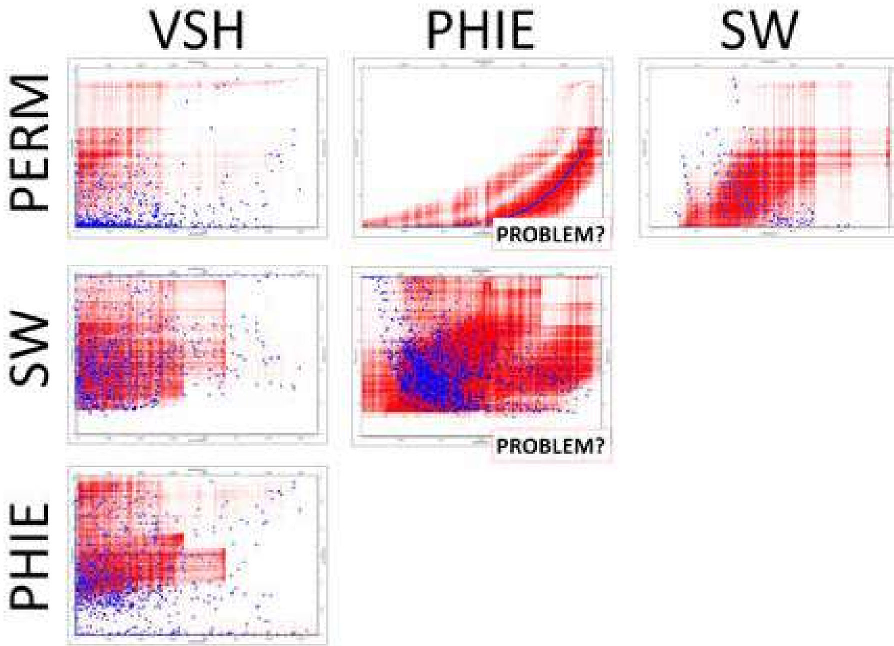

Figure 8. Cross-plots between the different petrophysical properties (VSH, Porosity, SW and Perm) as 3D petrophysical model

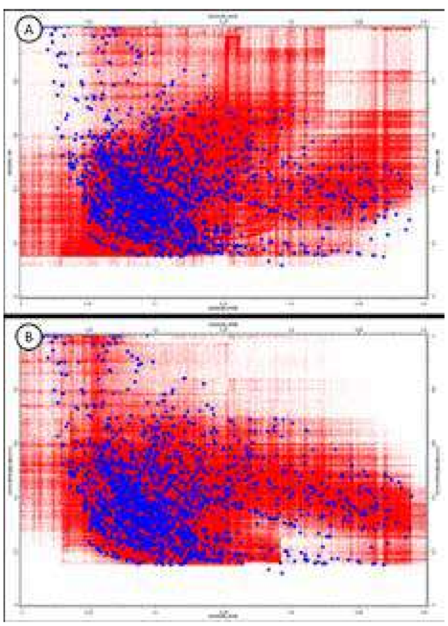

Figure 9. Cross-plots Porosity vs SW. (A) Original incorrect 3D SW model. (B) Corrected 3D SW model.

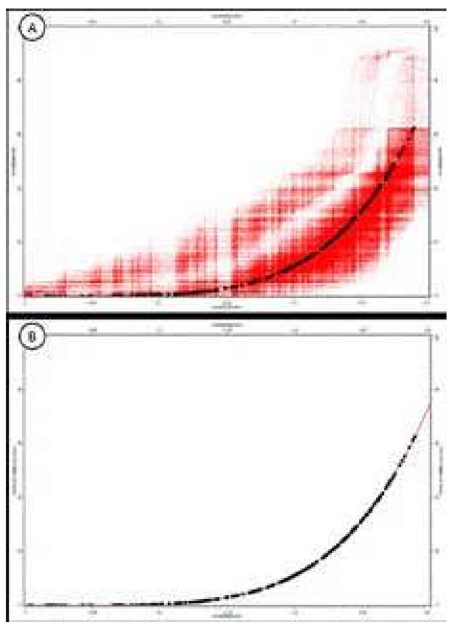

Figure 10. Cross-plots Porosity vs Perm. (A) Original incorrect 3D Perm model. (B) Corrected 3D Perm model.

INTEGRATING PETROPHYSICAL UNCERTAINTY

Like all of us, petrophysicists are dealing daily with uncertainty. Well logging is not error free. Interpreting the logs can be challenging. The core measurements might be inconclusive. And above all remains the fact that we work with a very limited dataset to understand a very complex environment – our reservoirs.

It does make sense to include the petrophysical uncertainty in our geomodeling workflows, as we do with geological or geophysical uncertainty. In practice though, few studies seem to consider this uncertainty. If you have

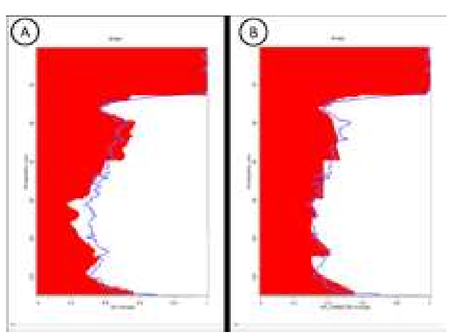

Figure 11. VTC of the 3D SW model. (A) Original incorrect 3D model. (B) Corrected 3D model.

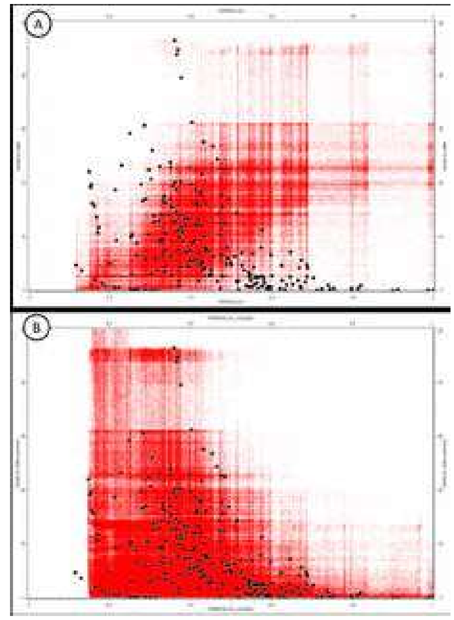

Figure 12. Cross-plots SW vs Perm. (A) Original incorrect 3D SW and Perm models. (B) Corrected 3D SW and Perm models.

to prioritize, we recommend that you first check if your geomodeling workflow properly takes into account the complex relationship between the properties as described in the previous section. Then, once checked, and fixed if needed, look at the petrophysical uncertainty.

As described in this paper, we should ask for three types of input from our petrophysicist: well logs of course, a description of the complex relationships between the different properties (what log was modeled from what log) and the mathematical equations linking some of these properties (if any). Each of these data can carry some uncertainty.

Instead of one version of the log curves, you could ask for several versions them. SW is a good candidate for this request. It is a complex property to model and your petrophysicist might want to see the impact of using slightly different values for the parameters of the Archie's equation. He could generate for you several logs of SW. The overall relationship between the properties (what log was computed from what log) is the least likely to change in a project. But it could be good to ask about, just in case. At last, your petrophysicist might decide to give several versions of the mathematical equations he is using (different porosity-permeability functions maybe).

Ultimately, these three types of uncertainty can be used, either through simple editing of your current geomodeling workflows, or in extreme cases, in defining several workflows, each one capturing a different petrophysical analysis.

CONCLUSION

While we easily involve the geologist and the geophysicist in our work, we tend to do it less with our petrophysicist. Many geomodelers are geologists or geophysicists by background. Maybe it leads us to spend more time on building the horizons, the faults and then facies distribution in 3D, and less on the challenges of populating the petrophysics in 3D.

It would be a mistake to not consider the complex relationship existing between the different properties (cross-plots). Our petrophysicist can help us a lot to understand these relationships. Many would be happy to open their kitchen and answer to all the questions we might have about their work.

The next paper will focus on the relationship between geophysics and geomodeling.

TO GO BEYOND

The CSPG and the CWLS are great sources of information about petrophysics, with their courses, their technical presentations and the papers they published.

The Reservoir magazine houses several papers or series of papers about petrophysics through the year. Among them we have a series on FMI data published in 2015 and the always popular papers written by our colleague Ross Crain for years now.

Ross' website is also considered a very good source of information by the many in our industry. It is definitely worth a look (www.spec2000.net).

REFERENCES

Crain, R. Crain's Petrophysical Handbook. www.spec2000.net

Doveton, J.H., 2014. Principles of Mathematical Petrophysics. Internal Association for Mathematical Geosciences (IAMG). Studies in Mathematical Geosciences No. 9. Oxford University Press. 248 pages.

Ellis, D. V. and Singer, J.M.Well, 2007. Logging for Earth Scientists. 2nd edition. Springer. 692 pages.

THE AUTHORS

Feel free to contact the authors about this paper or about the series.

Thomas Jerome is the Calgary manager of the Reservoir Modeling team in RPS (Thomas.Jerome@rpsgroup.com). Geologist by background, Thomas has spent the last 15 years working on geomodeling projects, first in Middle-East and in Calgary, Canada, since 2008. He has worked on conventional and unconventional reservoirs, both Canadian and international. Thomas is leading this CSPG's series on Geomodeling.

J. Kim McLean is an independent petrophysicist (kim.7.mclean@gmail.com). Kim attended the University of New Orleans for her Bachelor's of Science degree in geology, and decided she liked geology so much that she attended Central Washington University to get a Master's of Science in Geology with a focus on structural geology. After completing her graduate studies, she eventually landed at Paradigm, where she learned a great deal about petrophysics while supporting and mentoring clients on a formation evaluation software application. Kim then transitioned to the operator side, working as a petrophysicist for an oil sands asset in Canada.

Laverne Chan is a petrophysicist for a North American E&P Company, Enerplus, (lchan@enerplus.com), and has spent the last 10 years of her career in the oil and gas industry in Western Canada. After receiving her B.Sc. in Mechanical Engineering from the University of the West Indies in 2005, Laverne spent the first 4 years of her career as a Wireline Logging Engineer working in the fields of Western Canada. During that time, Laverne returned to the University of the West Indies where she completed her M.Sc. in Petroleum Engineering after which she continued her career in the role of Formation Evaluation Engineer and Petrophysicist in Calgary. While in the role of Petrophysiscist, Laverne went on to complete her M.Eng. in Reservoir Characterization at the University of Calgary and continues to work today in both conventional and unconventional reservoirs in North America.