GEOMODELING: A TEAM EFFORT TO BETTER UNDERSTAND OUR RESERVOIRS

Part 7: Reservoir Engineers and Geomodeling

By Thomas Jerome, RPS, Samaneh Razzaghi, Encana Corporation and Martin Malek, Tamarack Valley Energy Ltd.

INTRODUCTION

After having discussed how geoscientists and geomodelers can efficiently collaborate, this series is now reviewing the point engineers can take in geomodeling projects. This paper focuses on the collaboration between reservoir engineers and geomodelers around the topic of flow simulation. The next two papers will look at geomodeling for reserve estimates and for production engineering respectively.

Many geomodels are built at the request of reservoir engineers. They are in need of a 3D grid, capturing the characteristics of the rocks (porosity, SW, permeability) to feed to their flow simulation software. In fact, this is largely what motivated the development of geomodeling in the first place and several decades ago. Geomodeling can be seen as a bridge between geoscientists and engineers (Figure 1). This places geomodeling at the intersection of two worlds which have always had difficulties communicating one with the other. The first part of this paper describes the common communication issues around geomodels. When not handled properly, these issues can be the cause of failure of our projects.

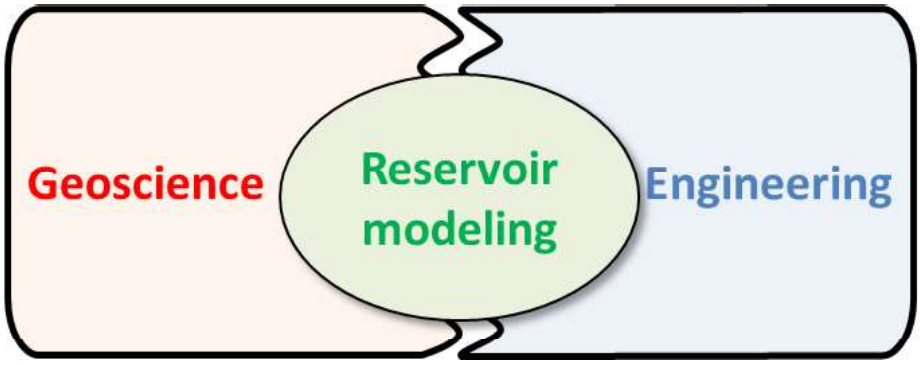

Figure 1. Reservoir modeling: at the intersection of geosciences and engineering.

A good 3D geological grid captures all the geological, petrophysical and geophysical information gathered by geoscientists about the reservoir. The geomodeler then transfers the information stores in the geological grid into a new 3D grid fit for the purpose of flow simulation. This new 3D grid is called a flow simulation grid hereafter. The reason why a specific grid is needed is also covered in the first part of this paper. The second part focuses on the technical processes involved (creation of the flow simulation grid as well as upscaling and downscaling of properties).

Transferring the uncertainties from the geomodel study into the flow simulation study is a major challenge. Which geomodels shall we use in flow simulation when potentially hundreds of 3D petrophysical distributions have been created by geostatistical algorithms? Historically, only one was used. Nowadays, several 3D geomodels are sent to flow simulation. In either case, it is up to the geomodeler to help his team to choose which 3D distributions should be used. It requires ranking the geomodels based on criteria relevant to flow simulation. This topic is covered in the last part of this paper.

COMMUNICATION, COMMUNICATION, COMMUNICATION…

What asset team doesn't joke about geoscientists not understanding what engineers need and/or about engineers not getting what geoscientists do? These jokes are as much a way to exorcise any possible communication issue to come as they are a way to vent out the frustration of on-going problems caused by miscommunication. And beyond that, these jokes are simply funny! Our teams are under a lot of pressure. A good joke is always a nice way to lift some of the tension we face and we should enjoy them for that!

Every geomodeler should be vigilant about this potential problem though. Too many geomodeling projects don't reach their full potential, because of miscommunication between geoscientists and engineers. It's unfortunate, but luckily it can be largely avoided. The remainder of this paper provides some ways to do so.

Nothing gives a reason for a good laugh (or a fair amount of frustration) more than a geomodeling project already in progress for a few weeks (months…) and everything has to be redone because the team suddenly realizes that the model doesn't take into account a few wells needed later for flow simulation. The question is not who shall have given the information to start with – the team, the geomodeler, the geoscientists or the engineers. The point is that it is a problem that a proactive geomodeler can easily fix, at the beginning of a project, by agreeing on the list of wells to be used.

Firstly, we must validate the list of wells with our geoscientists. On their side, it will be linked to which wells have geological/petrophysical/geophysical data that must be taken into account in the model. Secondly, we must crosscheck this initial list with the list of wells the engineers are looking at. Many wells will be on both lists. But engineers will also consider wells with some production history, even if these wells have no data useful for modeling the geology of the reservoir. Horizontal wells tend to fall in this category. They have some production attached to them, or they will in the future and so they must be taken into account for predicting future production. Yet, they might have no data usable for geomodeling per se. Too often geomodelers forget to make sure that these wells fall in the correct geological units. That is, wells known to have been drilled in a sand layer might end up crossing into some shale units located above or below the targeted sand. It happens when the geomodeler interpolated the horizons incorrectly between vertical wells, not realizing that it placed the horizontal wells in the wrong place. Geomodeling packages have options in their workflows to take into account the complete geometry of the horizontal wells.

At the same time, it is wise to check with the engineers what the lateral and the vertical extent of the volume of rock they need modeled is. Figure 3 in the second paper of this series (Jerome et al, 2015a) and its

(... Continued on page 16)

(... Continued from page 15)

associated paragraphs give an example of such problems.

Once the well list and the volume to model is approved by the whole team, the project can start. During the project, the geomodeler will communicate about his processes and his results to the geoscientists. Among other things, he will explain why he picked some specific geostatistical workflows and he will show that his model is indeed respecting the ideas the geoscientists have about the reservoir. It is wise to include engineers in these discussions. Firstly, it will give them more confidence in the project. Secondly, it will emulate discussions about the model inside the whole team. Geoscientists tend to focus their review on how the geomodel respects their ideas about 3D facies distribution. This is crucial, of course, but it can sometimes overshadow some mistakes a geomodel might have in term of respecting the laws of physics in general and the laws of flow dynamics in particular. Engineers will often spot such mistakes. During the presentation of his model to his team, the geomodeler should go as far as stating that he needs the geoscientists' feedback on the facies and the porosity as well as the engineers' feedback on the water saturation and the permeability models. In so doing, everyone knows what your expectations are for her/him.

Figure 2 gives an example of a project in which the engineer's feedback on water saturation was crucial. It is based on an anecdote that happened to one of the authors a few years ago. The reservoir was a simple sandy geological unit. There was no facies modeling per se as the whole unit was considered made of sand. The porosity modeling didn't cause any issue either. Water saturation proved more challenging (Figure 2A). The water saturation log showed low values everywhere with the exception of a zone close to the top of the unit, and only in the South-West corner of the reservoir. There, the water saturation was getting close to 100%. Due to a large number of input wells, this information was overlooked by the geomodeler. Water saturation was modeled using geostatistical techniques, in the same way it had been done on many other projects before. The 3D water saturation model was showing, locally, a zone of high values around the wells. Everything was consistent as far as geostatistics was concerned; the hard data were respected as well as the global saturation distribution and the global variogram. The geoscientists and the geomodeler reviewed the project. Satisfied by their model, they gave it to their engineers and they moved on to other tasks. Months later, the geomodeler and the geoscientists discovered that the engineers were struggling with the geomodel; water was literally "raining" in their model from the zone of high saturation. To them it was, in fact, impossible that such a zone of water saturation existed there. It did not make any sense in terms of flow dynamics. Gravity would have made this water drop to the bottom of the reservoir (water being denser than the oil in this reservoir). They decided to manually edit the saturation in the problematic area to get some good flow simulation results. Naturally, they were frustrated by this situation. Reviewing the geomodel and the input data, the geoscientists discovered the source of the problem - the water saturation logs were valid, but not the facies description. It had been missed that the reservoir was showing a local continuous shale in that zone. The water saturation model was correct, but the permeability model was not. High permeability values, believed to be in a pure sand unit, had been distributed in the whole sand. Instead, it should have been set to zero in the shale unit. In that case, the water would not have "rained" in the sand below. The geomodel was rebuilt. A zone of shale was added to the facies model. The water saturation was now modeled by facies – very low in the whole sand and close to 100% everywhere in the shale. At last, permeability was computed by facies as well - high in the sand, null in the shale.

This anecdote illustrates several important points. Firstly, the engineers might indeed spot issues with fluids and permeability models that geoscientists and the geomodeler himself might miss. Secondly, if the engineers are not involved in the review and they just received the geomodel as a package thrown over a fence, there is a greater chance that they will try to correct such problems themselves rather than reporting them to the team. The team might then end up with two geomodels; the original, and the one edited behind closed doors by the engineers. And who is to blame for such situations – the engineers for not communicating about what they saw or the geoscientists and the geomodeler for not properly involving them in the review process? To avoid having to argue about such questions at a later stage, we believe it is in any geomodeler's own interests to include engineers in their project at the same level than geoscientists are.

Involving engineers in their projects will also help the geomodelers to address two of their most common questions. Many engineers wonder why we are spending so much time building a facies model while they need only porosity, water saturation and permeability. Many also wonder why the 3D geological grid we work on has a complex mesh and millions of cells while they specifically asked for a "sugar box", simple 3D grid.

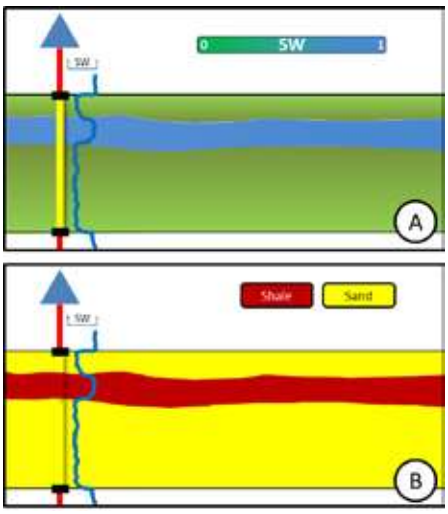

Figure 2. A) Initial, incorrect SW 3D model. B) Reinterpreted reservoir following engineer's feedbacks.

Flow simulation engineers need 3D grids which are aligned with the main direction of flows in the reservoir. In a simple, layer-cake reservoir with no fault and no folding, it means that the K axis of the 3D-grid should be indeed perfectly vertical. The horizontal mesh will be, for example, parallel and perpendicular to the horizontal wells around which the flow simulation is run. If the reservoir is fractured, the horizontal mesh will likely be built parallel and perpendicular to the main direction of the fractures. As a last example, if the reservoir is faulted, the horizontal mesh will likely be built parallel and perpendicular to the fault surfaces. In addition to this, the mesh of the flow simulation grid should be made of cells of constant size, with no truncated or eroded cells. These constraints ensure that the computations in the simulation software run faster and are more stable numerically.

The geological 3D grids are built to populate petrophysical properties in 3D. As these properties are primarily controlled by the facies distribution, we have to model facies in detail as well. Geostatistics are our main toolbox to do this. In (Jerome et al, 2015b), we explained that the orientation of the mesh of the geological 3D grid is the primary control on how the facies (and the petrophysics) are interpolated around the wells. Use a mesh that doesn't reflect the directions of sedimentations and you are likely to get an incorrect 3D facies distribution.

In fact, building a geological 3D grid and building a flow simulation 3D grid follows the same problem. In both cases, we need a 3D grid that is aligned with the main directions of the physical phenomena we are modeling. In flow simulation, it means solving the

equations of flow dynamics and the mesh must follow the directions of flow. In geomodeling, it means mimicking with geostatistics the results of physical phenomena like erosion and sedimentation and we need a 3D grid with a mesh parallel and perpendicular with the directions of deposition. Those directions are usually different from the directions of flow simulation. That's why we can't use the flow simulation grid to model facies, and in reverse that's why it is unwise to run flow simulation in a 3D grid fit for facies modeling. We need two 3D grids with specific cell size, with specific orientations for the mesh and with or without eroded cells. We need a geological 3D grid and a flow simulation 3D grid. It implies that we will have to transfer rock properties from the geological grid to the flow simulation grid (it will be covered in the next part).

Explaining the need for a specific geological grid by building an analogy with flow simulation constraints have proven efficient to the authors on several occasions.

Figure 3 and Figure 4 illustrate these points. Let's assume we have three vertical wells, each showing a succession of shale and sand. An engineer might convince a geomodeler to use a "sugar-box" 3D grid for modeling facies – that is a grid with constant cell size and horizontal and vertical mesh. The type of mesh flow simulation would be run into. If we do this, the facies model would look like a succession of horizontal sand and shale layers (Figure 3A). Let's assume that dipmeter data shows that the sands and the shales are in fact dipping. It makes sense to build a 3D geological grid with an inclined mesh (Figure 3B). The new 3D facies model is now very different from the original one. This second approach is better than the first one as it not only respects the facies at the wells, but also the information of the dipmeter data and, from there, the geological concept developed for the reservoir– the facies are dipping. Figure 4 shows how the water saturation model would look like. This is the type of property (with porosity and permeability) that reservoir engineers need. Shall we give them our geological, inclined 3D grid for their flow simulation? Maybe, if our engineers confirm such a 3D grid is good for their work; but very likely, they will ask for the 3D petrophysical models to be transferred into a sugar box grid.

Without proper understanding, by the geomodeler, of what is needed for facies modeling and later for flow simulation, either the facies model would be wrong (Figure 3A) or the 3D-grid sent to the engineers could potentially be inadequate for flow simulation (Figure 4).

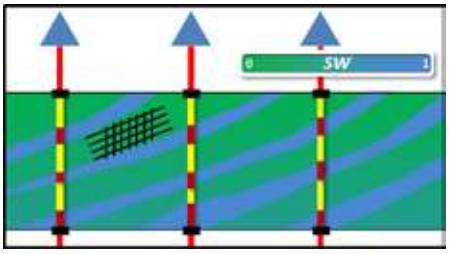

Figure 3. A) Incorrect, initial 3D facies model. B) Corrected 3D facies model, once taken into account the wells' dipmeter data.

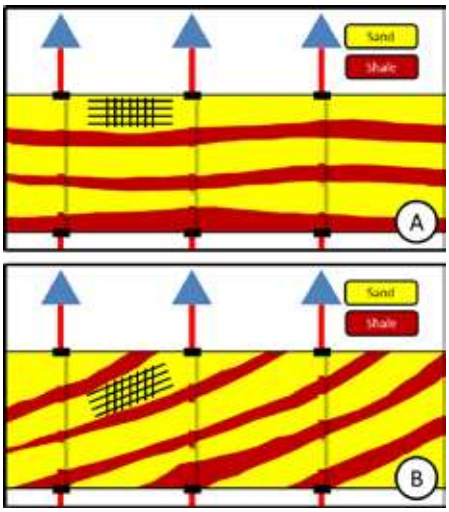

Figure 4. 3D SW model associated to the corrected 3D facies model (Figure 3B).

FROM 3D GEOLOGICAL GRIDS TO 3D FLOW SIMULATION GRIDS

Transferring information from the 3D geological grid into the 3D flow simulation grid is a two-step process. Firstly, the simulation grid is created. Secondly, the properties modeled in the geological grid are transferred into the flow simulation grid.

A flow simulation grid might be as simple as an upscaled version of the geological grid; the same dimension and same mesh orientation, but larger cells. For example, a geological grid of 25m*25m (Figure 5A) is upscaled to a flow simulation grid of cell size 50m*50m (Figure 5B, green mesh is the mesh of the flow simulation grid). In many projects, the flow simulation grid has larger cells than the geological grid. It's a way to limit the total number of cells, and so to limit the computation time. Because of this, engineers might ask why we build the geological 3D grid at a smaller scale if, ultimately, the grid will be upscaled to larger cells. Following the same approach than was developed in the end of the first part, the simplest is to explain that the physical phenomena we are modeling (facies, deposition, erosion…) require a high-resolution grid to be properly modeled; as shown by the average length of the facies along the input wells.

Knowing from the start of the project that the flow simulation grid will be simply an upscaled version of the geological grid has one benefit – it allows picking more appropriate dimensions for the area to model. Let's imagine a project in which the geomodeler knows he will create a geological grid of cell size 50m*50m. Let's assume the area to model is a square of 10,550m by 10,550m. This dimension can be split into 211 cells of exactly 50m. It works. But let's imagine now that the engineer explains, after months of geomodeling work, that he will need a simulation grid of cell size 100m. At this point, there is a problem – 10,500m can be split into 105 cells; 10,600m can be split into 106 cells; but 10,550m doesn't work. It would give 105 cells and a half. Had the geomodeler known the engineer's request from the start, he might have modeled a square of 10,600m instead of 10,550m.

Sometimes, the engineer knows he will use an upscaled version of the geological grid as flow simulation grid, but he doesn't know yet the exact cell size he will need. In such a case, the geomodeler should, at least, pick a number of cells that is a multiple of 2, 3, 4 and even 5, if possible. It will give some flexibility to the engineer at a later point.

In the previous example, 10,600m is split into 212 cells of 50m. As we have seen, 212 can be divided by 2, giving a simulation grid of 100m. But 10,600m can't be split into cells of 3*50m (150m, it gives 70 cells and 2/3 of a cell) nor into cells of 5*50m (250m, it gives 42.4 cells). It can be split into cells of 4*50m though (200m, it gives 53 cells). Not bad, but we can easily improve on it. Indeed, a square of 10,500m instead of 10,600m works well for each possible level of merge. If the 100m we remove were of no real importance, then it would make sense to decrease the modeled area slightly. At a negligible cost (losing 100m of modeled area), the geomodeler has now a grid that, very likely, will be good when the engineer decides what cell size he really needs.

The lateral extent of a model is never really well constrained. The modeled zone must at least cover the zone of interest to the team, but beyond that, it can be as large as needs be. For that reason, adjusting it based on the criteria discussed in the previous paragraph is often easy to do. Vertically though, the model is extremely well constrained by the geometry of the horizons delimiting the reservoir. We can't really "add" anything to allow for a perfect match when the grid is upscaled for flow simulation.

The flow simulation grid can be more complex than being just an upscaled version of the geological grid. Occasionally, it is

(... Continued on page 18)

(... Continued from page 17)

necessary to add some LGR (Local Grid Refinement, see Figure 6A). An LGR is a zone in which the reservoir engineer needs a resolution higher than that of the geological grid. A zone in a model of 50m*50m cell size might need to be split into cells of 1m by 1m. It is needed when the physical phenomena to be modeled in flow simulation is expected to change quickly over short distances. It is the case for the behavior of fluids around induced fractures for example. When this is needed, it is usually not just for one area. Often, multiple small zones will each need its LGR (for example to capture multiple zones of induced fractures along a multi-fracked horizontal well). If such small cells are needed, why not build the whole geological grid at this level of detail to start with? We could, but we would face two problems. Firstly, it would transform multi-million cell models into multi-billion cell models. Hardware and software might not be able to handle such large grids. Secondly, we have limited to no information on how the reservoir can change at a very fine scale. Even at a resolution of 50m*50m, a geological grid has a good deal of uncertainty in it; we don't know precisely what the facies distribution between the wells should be. With this in mind, what trust could we have in facies and petrophysical 3D distributions in a geological grid of 1m*1m? Probably little to none. We might as well stick to a resolution (50m*50m) at which we feel comfortable defending our results rather to adventure ourselves to a level of details (1m*1m) at which we know nothing about. LGR are examples of flow simulation grids in which the properties are not upscaled but downscaled. This process is described at the end of this part.

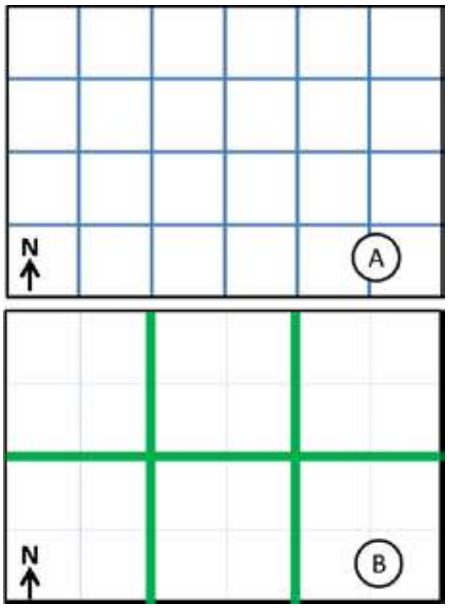

Figure 5.A) Original 3D geological grid. B) Upscaled 3D simulation grid (4 cells merged into 1 cell).

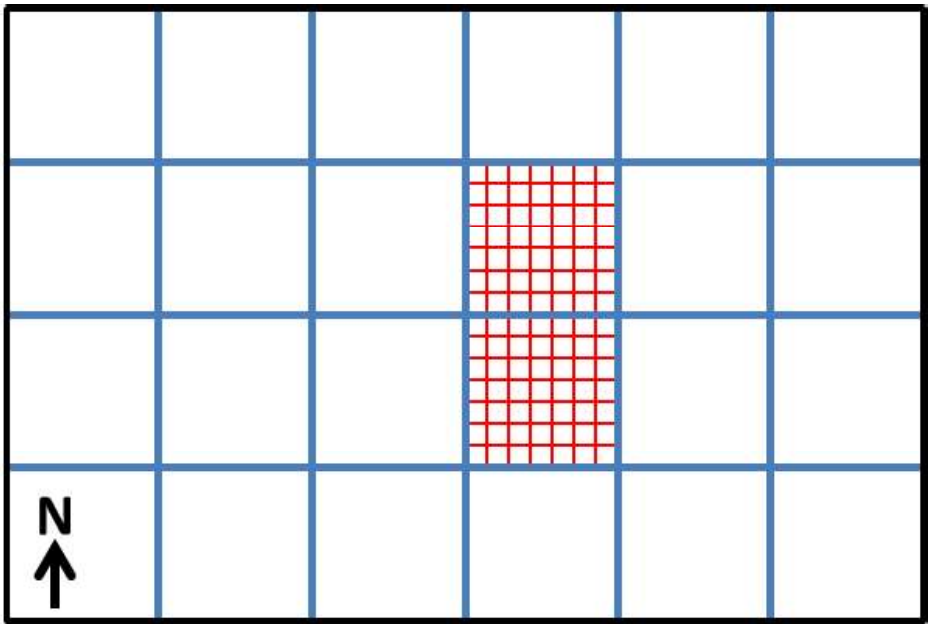

Figure 6. 3D simulation grid with an LGR.

At last, a flow simulation might be rotated compared to the orientation of the geological grid (Figure 7, left). If nothing really constrained the geological grid to be aligned North-South and East-West, while the simulation grid will need to be rotated, it is recommended to build the geological grid directly rotated. In so doing, we are back to the illustration of Figure 5. Such a situation happens, for example, in naturally fractured reservoirs in which the rotation will correspond to the average direction of general stress or to the average direction of the natural fractures.

Once the flow simulation grid(s) are created, the properties from the geological grid must be transferred into it (them). Two techniques are used. When the simulation grid has larger cells than the geological grid, property upscaling is used (Figure 8 and Figure 9); and secondly, when the simulation grid has smaller cells than the geological grid, property downscaling is used (Figure 10).

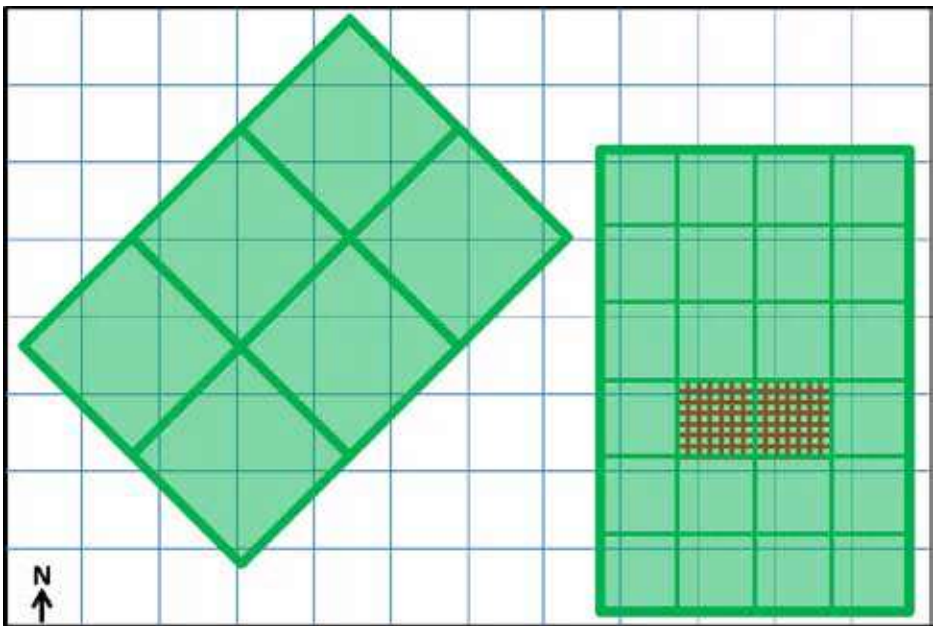

Figure 7. 3D geological grids versus two local 3D simulation grids.

It is also common that multiple flow simulation grids are generated from a unique geological grid (Figure 7, two flow simulation

Let's imagine a geological model of cell size 50m by 50m horizontally and 1m vertically. The engineer asks for a simulation grid of

cell size 100m by 100m horizontally and 3m vertically. We have now a single large cell where there were 12 small ones before. Each small, original cell has its own values of porosity and water saturation. What values should be stored in the corresponding large cells of the simulation grid? These are the questions that property upscaling has to answer. Fundamentally, it is the same problem as upscaling well logs into the cells of the geological grid (Jerome et al, 2015c). We apply some mathematical averaging technique that will compute a value in the simulation grid which will respect the characteristics of the 12 original cells.

For porosity, the upscaled value must respect the cumulated porous volumes of the 12 original cells. Arithmetic average weighted by cell size is usually applied. Upscaling water saturation follows a similar approach. We upscale it by arithmetic average weighted porosity to ensure that the hydrocarbon porous volume of the upscaled cell is equal to the sum of the hydrocarbon porous volumes of the original 12 cells.

One question remains– should we worry about upscaling facies?

Two approaches are used. In the first one (Figure 8, method A), the facies property is being upscaled first, and then the upscaled facies guides the upscaling of the petrophysical properties. The approach is again equivalent to the one used to upscaled facies description from the well into the cells of the geological 3D grid. The upscaled facies is the preponderant one in the original cells. In Figure 8, the upscaled upper cell is the combination of 4 cells of sand and 2 of shale; the upscaled facies takes the value sand. In a similar way, the 2 cells of sand and 4 cells of shale below are replaced by a single cell of shale. For the porosity, upscaling is done by arithmetic average, as explained previously. Only the cells with the same facies than the upscaled one are used. For the upper upscaled cell, it means that the averaging is done using the porosity of the 4 original sand cells. Similarly, the porosity in the lower cell is the average of the porosity in the associated 4 original shale cells. The same filtering technique is applied to compute the average water saturation values. In the second approach (Figure 8, method B), the facies is ignored. The 6 original values of porosity and water saturation are used to compute the average values in each upscaled cell.

The two methods give some significantly different results in situation, such as the one described here where the original cells belong to two different facies. When filtering by the facies, the resulting upscaled porosity and water saturation are characteristic to what is expected for the respective upscaled facies. The upper upscaled cell is sand and its associated upscaled porosity is around 30%, which is coherent with the values of porosity found on the logs and modeled in the sand facies in the geological grid. The same can be said for the lower upscaled shale cell. On the contrary, with the second method in which facies is ignored, the upscaled values are average of original numbers associated with shale and sand. The upscale cells are not really a sand, nor a shale, anymore. It is an "average facies" of some sort.

Which method is best? It will depend of the reservoirs and the decision should be made by the whole team. The key underlying question is to know which of the two methods creates an upscaled grid which behaves similarly to what would have been observed if flow simulation was run on the original, not-upscaled, geological grid. If time and resources allow, it might good to run flow simulation in a portion of the geological grid (so that flow simulation runs fast). Simulation is done in the original geological grid and it is also run in the flow simulation grid with properties upscaled with the first method, and then with the second method. The method associated to the flow simulation grid behaving the closest to the original geological grid is the one to keep.

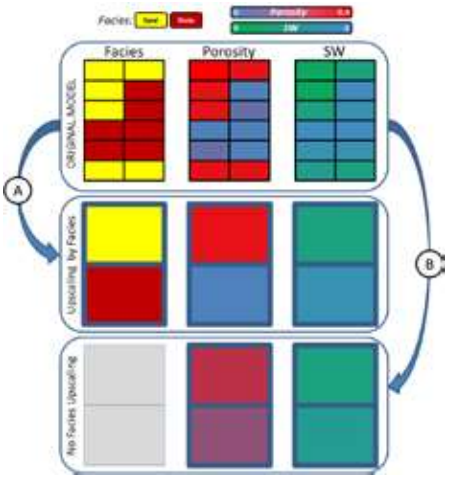

Figure 8. Upscaling porosity and SW by Facies versus upscaling the petrophysics without facies upscaling.

Property upscaling is also the time to validate the cell size for the flow simulation grid. We must ensure that the global geological characteristics stored in the geological grid are respected by the flow simulation grid.

For example, imagine an original geological grid showing a sandbar in the middle of a shale geological unit (Figure 9A). If we create a flow simulation grid with very large cells, the sandbar is split in two parts in the simulation grid (Figure 9B). The simulation grid is too coarse to respect the continuity of the sandbar. This would influence greatly the results of flow simulation. If a flow simulation grid with smaller cells is used (Figure 9C), the continuity is respected. This second simulation grid is more appropriate for this reservoir.

When asked about the cell size for the simulation grid, engineers often answer that they want a grid with as few cells as possible. The geomodeler should help his team to balance this constraint with another important one – the simulation grid must be detailed enough to respect the continuity of the geology of the reservoir.

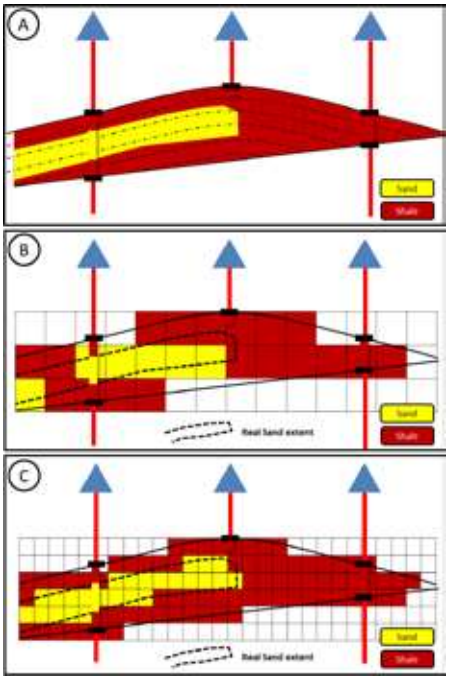

Figure 9. A) Facies model in the original 3D geological grid. Facies model in a coarse (B) and in a refined (C) 3D simulation grid.

Not every flow simulation grid is coarser than the geological grid. Sometimes, it is more refined and the properties must be transferred by downscaling instead of upscaling. Upscaling here means finding a single value to replace several original ones efficiently. Downscaling is the opposite – where the geological model is a single cell (50m*50m by 1m vertically for example), we now need values for multiple cells (for example, each of 10m*10m*1m). To do this, two approaches exist (Figure 10).

The first approach is probably the main one used in the industry and it is a very straightforward one. All the small cells of the simulation grid, which fall inside a large cell from the geological grid, are getting the value stored in that large cell. For example, with a sand-shale facies model (Figure 10A), the cells of the rotated flow simulation grid are assigned the facies sand if their center is falling in a sand cell of the original geological grid (Figure 10B, all the cells falling inside the

(... Continued on page 20)

(... Continued from page 19)

black shape, delimiting the sand zone of the geological grid).

This approach causes one problem though. Imagine a 50m*50m cell of the original geological grid with a porosity value of 22.4%. Does it really mean that we have a block of rock in the ground of 50m by 50m with a homogenous porosity of 22.4%? If the next cell of the geological grid has a stored porosity value of 30%, does it mean that we have a block of 50m by 50m of 30% porosity in direct contact with the block of 50m by 50m of porosity 22.4%? Obviously not. We believe instead that the porosity is gradually changing from 22.4% to 30%. And yet, downscaling with the approach described in the previous paragraph means that we believe our reservoir to be really made of "homogeneous blocks". Will the results of flow simulation in such a "blocky" simulation grid be close enough to what is really going to happen during production? If your team answers 'no' to this question, a second downscaling approach should be considered. It is a more time-consuming one, but it will allow seeing a different value inside each small cell of the flow simulation grid.

The second approach is based on one observation. When we use geostatistics to populate facies in the geological grid, we are not populating the volume of each cell with facies values. What geostatistical algorithms are really doing is computing a facies value at the XYZ point located at the center of each cell (Figure 10C). We don't say anything about what should be the facies values in the remaining volume of the 50m by 50m cells.

This fact can be used to our advantage when doing property downscaling. The facies values from the geological grid are painted only in the cells of the flow simulation which contain the center of each original cell (Figure 10D). It leaves a lot of cells with no facies value assigned to them. In a second step, geostatistics are used to populate these cells, using each of the painted cells as a hard data. With this process, facies, and then petrophysical properties such as porosity and water saturation can be downscaled and still show variation from one small cell of the simulation grid to the next. It gives a more realistic tone to the model than the blocky aspect that the first downscaling approach creates.

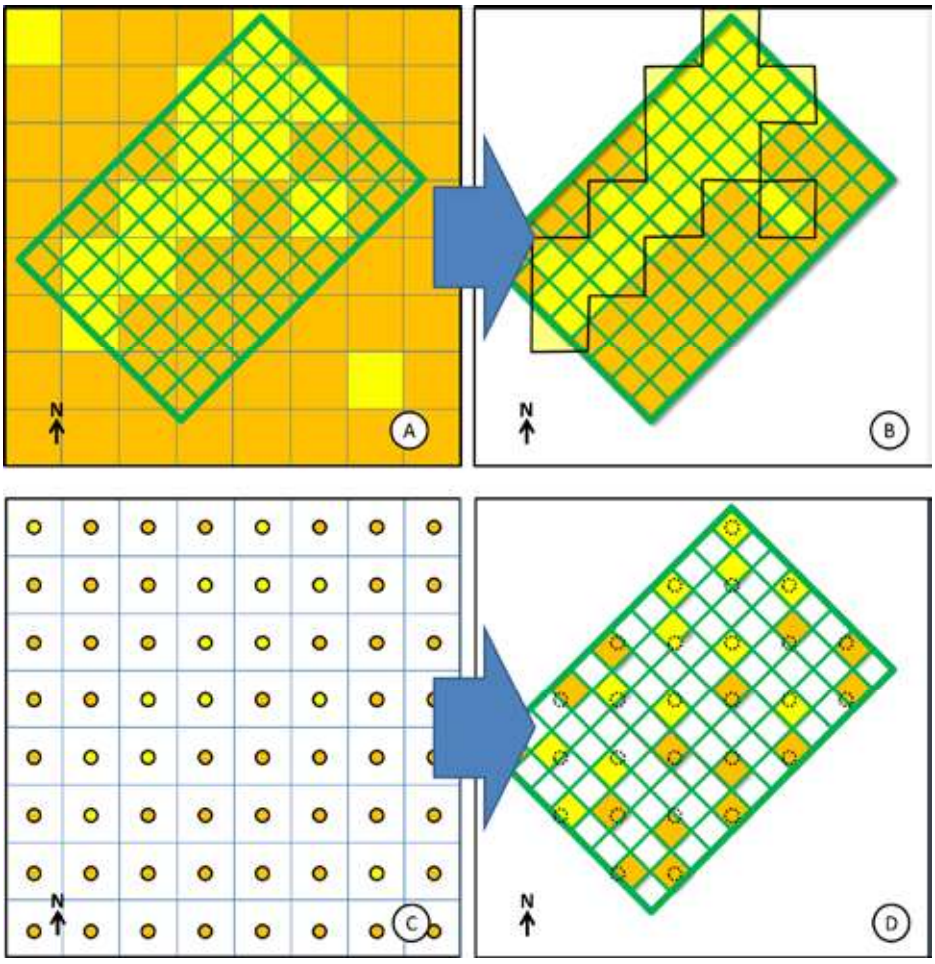

Figure 10. A) Facies model in the original 3D geological grid. B) "block" downscaling in the 3D simulation grid. C) Facies model in the original 3D geological grid in cell-center view. D) "center-to-center" downscaling in the 3D simulation grid.

RANKING GEOMODELS FOR THE PURPOSE OF FLOW SIMULATION

The previous section describes the fundamentals of creating the geometry of a flow simulation grid and how to transfer facies and petrophysical properties. One big issue remains. What do we do when we have created hundreds of geological models, each one made of a 3D distribution of facies and of 3D distributions of petrophysics. Do we give only one to the simulation engineer? Do we give several realizations? And, depending upon the approach, which of the hundreds shall we give? Do we transfer all the realizations? Answering the last question is the easiest one. No, we do not transfer all the realizations (much to the relief of our engineers!). If each realization takes hours, days or even weeks to run, one can only imagine how long it would take to run all of our realizations.

In that case, should we send just one then? This approach is still used a lot. It tends to be part of some discussion between the geomodeler and the engineer which goes like this: "I have hundreds of realizations, which ones do you want?"... "Hundreds!?!?! I can't run flow simulation on hundreds of realizations! Give me just one, the one you want!" And it is up to the geomodeler to pick one. Of course, it is possible to pick just one. We can use the concept of ranking that we'll cover in the next few paragraphs. But, is this acceptable? Providing only one to the engineer means that we suddenly ignore all the uncertainty we have identified with the geoscientists. The model is now completely deterministic, giving a false sense of certainty to the engineers, to the team and to the decision-makers of our companies.

How tempting it might be to send only one model to flow simulation; the team and the company will make more grounded decisions about the reservoir if the geological uncertainties are taken into account in flow simulation. We can't just send one realization. We need to send several of them to sample adequately the geological uncertainty space. The question is to know which realizations.

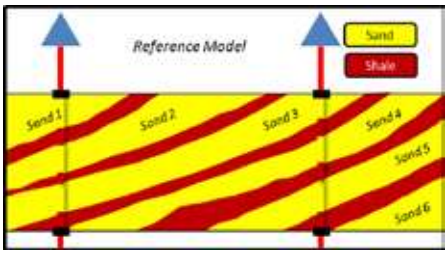

Figure 11. Facies distribution. Base case scenario with no connectivity between the sands.

The most common approach is to use the in-place volume, specific to each realization as a guide. The use of geomodels in reserve computations will be detailed in the next paper of this series. Hereafter, the topic is covered in very simple terms.

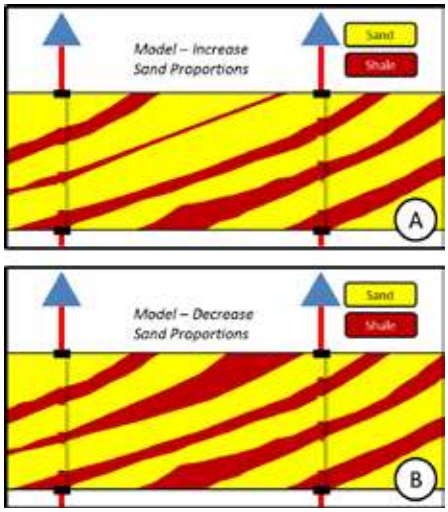

Figure 12. Facies distribution. Alternative scenarios with more (A) or less (B) Sand than in the base case (Figure 11).

Figure 11 represents the base case realization of a sand-shale reservoir (a variation of the example used in Figure 3 and Figure 4). The thickness of the shale separating the sands 2 and 3 is uncertain. It could be thinner (Figure 12A) or thicker (Figure 12B). In the first case, the in-place volume is higher than in the reference model, while it is lower in the second case (a thicker shale means less sand in sands 2 and 3). The in-place of the reference model is in the middle of the range. If a single model is sent to flow simulation, it might be tempting to send the reference case; if the in-place volume is an average of the in-place volumes of all the realizations, it is tempting to assume that this specific realization is also going to behave in an "average" way in terms of flow simulation. If we have hundreds of realizations and the two examples shown in Figure 12A and Figure 12B are two extreme in-places, then, following the same logic, we might send the three realizations to flow simulation. We might expect that the model showing a low in-place volume will show less production than the reference case which itself will be less performant than the model showing a high in-place volume.

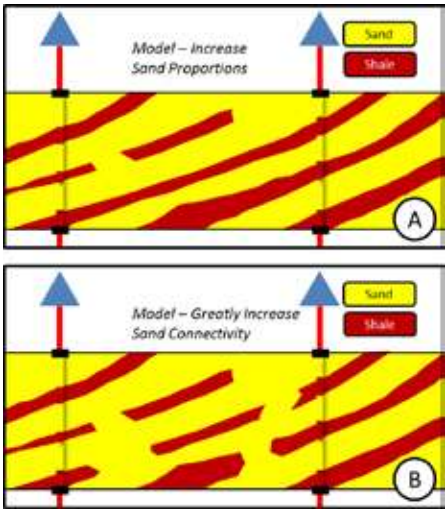

Figure 13. Facies distribution. Alternative scenarios with some (A) and a lot (B) of connectivity between the sands.

Unfortunately, ranking realizations based on in-place volumes can be misleading because in-place volumes tell nothing about the connectivity in the reservoir. Connectivity is a key controller in flow simulation.

Let's consider some new variations around the reference case. This time, the uncertainty is in terms of the continuity of the different shales. Do we really have continuous shales isolating completely the different sands, as in the reference case (Figure 11)? Or, are the shale discontinued and the sands connected? Figure 13 shows two scenarios where the sands are more (Figure 13A) and more (Figure 13B) connected. These two new realizations might have in-place volumes very similar to the in-place volumes of the reference case. Nevertheless, in terms of flow simulation, they will behave completely differently from the reference case. Without digging more about the consequence in term of flow simulation, we can easily imagine that the sands in the reference case will need to be produced independently one from the other, while connected sands might be produced as one block.

Now, we have five realizations. The reference case, (Figure 11 as well as upper central model in Figure 14), a realization of isolated sands with a lot of shale (Figure 12B as well as upper left model in Figure 14), a realization of isolated sands with less shale (Figure 12A as well as upper right model in Figure 14) and lastly, two models of connected sands, where the level of connection keeps increasing from model to model (Figure 13 as well as the middle and lower central models in Figure 14). Which ones shall we send to flow simulation? Changes in connection have a larger impact than changes in in-place volumes. It is preferable to send the reference case and the two models of connected sands. The reference case is expected to be a pessimist case in terms of flow simulation, because of the low level of connectivity, while the highly-connected model is expected to be the more optimistic scenario.

Figure 14 is the opportunity to see how information, provided by the geoscientists, is translated into geomodeling constraints which will lead to different results of flow simulation. On one hand, uncertainty in the sand proportions will lead to adding uncertainty in the input facies proportions used in geostatistical algorithms such as SIS (Sequential Indicator Simulation). This uncertainty will impact mostly the output range of in-place volumes. On the other hand, uncertainty in how the sands are connected will lead to adding uncertainty in the dimensions of the variograms used in geomodeling. The space of uncertainty in this simple example is the two-dimension space with sand proportion on one axis and connectivity on the second.

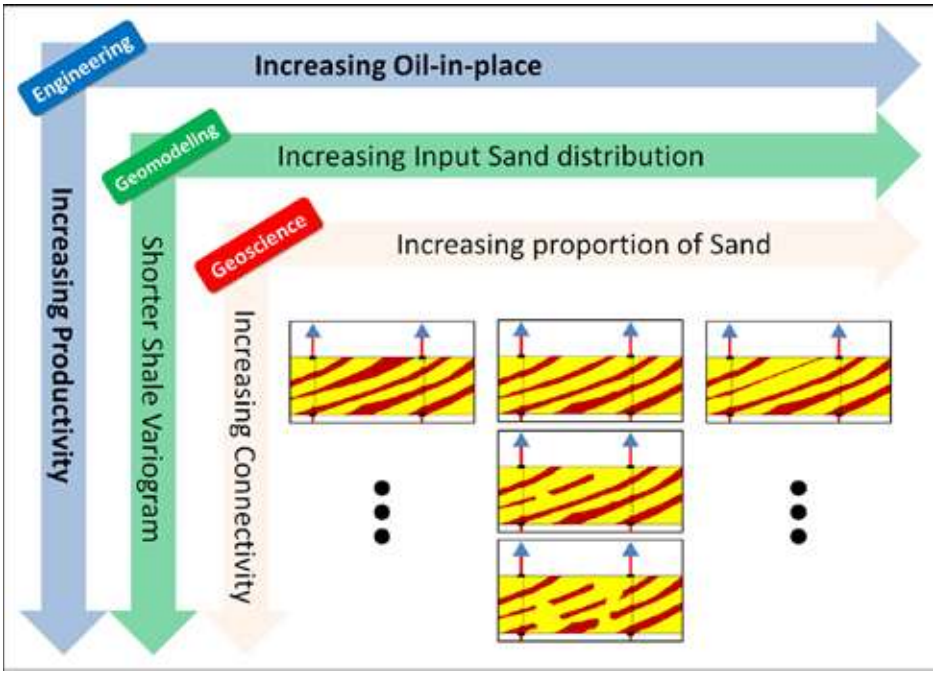

Figure 14. Translating the geosciences scenarios into geomodeling constraints and their impact in engineering terms.

Quantifying the level of connectivity and understanding its correlation to flow

simulation is still a topic of active research. A bibliographic review would be in order to understand how it is applied in the type of reservoirs studied by your team.

Another approach to ranking is to run streamline simulation (Figure 15) on as many of the geomodel realizations as possible. Streamline is a type of simplified simulation which runs much faster than full, true flow simulation. Streamline simulation could be used to evaluate how the different realizations are going to behave. From there, a few realizations are picked and sent to true, intensive flow simulation computations. More approaches exist to rank based on expected behavior in flow simulation. Geomodelers should discuss the topic with their reservoir engineers to see which one(s) seem(s) more appropriate to the project at hands.

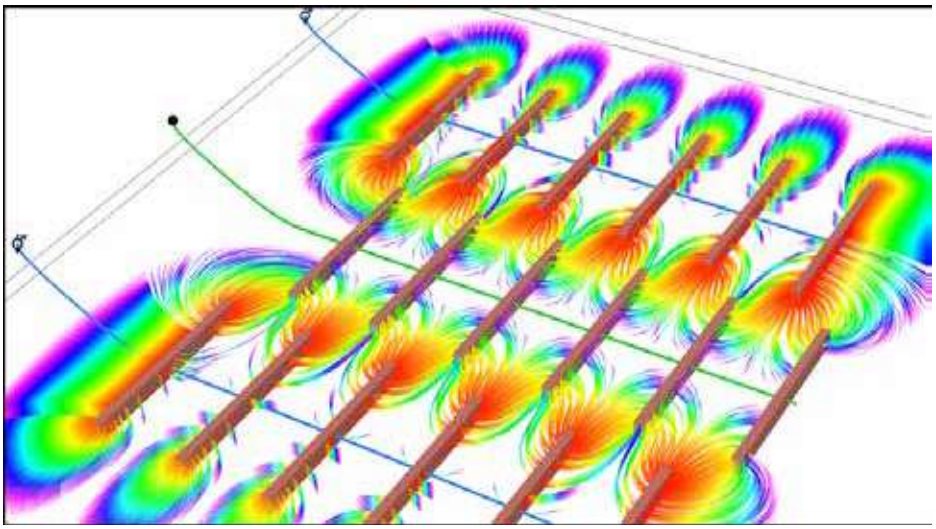

Figure 15. Example of streamline simulation (Lavoie and Thenin, 2013). Image courtesy of RPS.

CONCLUSION

Improving communication between the geomodeler and his team is a topic discussed in all the papers of this series. However, if the geomodeler has to pick one place where improving the communication might have the biggest impact, he should focus on the one with the reservoir engineers. Firstly, the communication between geoscientists and engineers are often challenging, and the communication around a geomodeling project will be no exception, unless attention is given to this challenge. Secondly, transferring the information from the geological model into the flow simulation grid is the moment where considerable knowledge about the reservoir can be corrupted or even lost. If the upscaling is not done properly, the upscaled grid might not reproduce well enough the behavior expected in the geological grid. Also, it is the moment where all the uncertainty identified by the geoscientists might get lost if the wrong realizations, or if not enough realizations, are sent to flow simulation.

For all these reasons, if reservoir engineers are involved in the geomodeling project, it is essential to communicate with them as often as possible.

The next paper will focus on the use of geomodels as inputs for in-place volume computations.

TO GO BEYOND

Many geomodeling packages have tools to create simulation decks, run simulation in the background by calling the flow simulation software and then display the time-dependant results. It might be interesting to discuss about these tools with the engineers.

Similarly, flow simulation software can compute statistics on the properties, as well as run some geostatistics. The geomodeler should get familiar with these. It will give him a better idea of how the engineer will review his model and how he might edit it.

When a field has had some production, running flow simulation starts with history matching. The goal is to validate and edit, if need be, the geomodel and the input dynamic parameters so that the engineers can reproduce the past production in their simulation. Once done, they can move to forecasting how the field's production will evolve moving forward. For more details about this important topic, the reader can refer to (Gilman and Ozgen, 2013) or (Oliver et al., 2008).

For more details about some techniques of ranking, the reader can refer to (Datta-Gupta and King, 2007) and (Jamshidnezhad, 2015).

REFERENCES

Datta-Gupta, A. and King, M. J., 2007. Streamline Simulation: Theory and Practice. SPE Textbook Series, Vol. 11. SPE. 394 pages.

Gilman, J. R. and Ozgen, C., 2013. Reservoir Simulation: History Matching and Forecasting. SPE. 120 pages.

Jamshidnezhad, M., 2015. Experimental Design in Petroleum Reservoir Studies. Gulf Professional Publishing. 176 pages.

Jerome T., Yang, J. and Gentile, S., 2015a. Geomodeling: a Team Effort. Part 2: the Geomodeling Workflow. The Reservoir Magazine, CSPG, March 2015.

Jerome T., Gentile, S. and Yang J, 2015b. Geomodeling: a Team Effort. Part 3: Geostatistics. The Reservoir Magazine, CSPG, April 2015.

Jerome T., McLean, J.K. and Chan, L., 2015c. Geomodeling: a Team Effort. Part 5: Petrophysicists and Geomodeling. The Reservoir Magazine, CSPG, April 2015.

Lavoie, R. and Thenin, D., 2013. Waterflooding Tight Oil Fields: The Static and Dynamic Modeling Aspects. RPS L&L series, Sep 2013.

Oliver, D. S., Reynolds, A. C. and Liu, Ning, 2008. Inverse Theory for Petroleum Reservoir Characterization and History Matching. Cambridge University Press. 380 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.

Samaneh Razzaghi is a Senior Reservoir Engineer with Encana Service Company Ltd. (samaneh.razzaghi@encana.com). She holds a Ph.D in Reservoir Engineering and has more than 10 years of experience in Numerical Reservoir Modeling. Samaneh has worked on both conventional and unconventional reservoirs such as Shale gas and Oil Sands in the WCSB, and Light Oil Carbonate Reservoirs in the Middle East.

Martin Malek is a Senior Development Engineer for Tamarack Valley Energy Ltd (Martin.Malek@TamarackValley.ca). Since graduating from the University of Calgary, Martin has worked in both subsurface and production capacities in the WCSB and Permian Basin in Texas. The majority of his expertise has been in development of unconventional reservoirs such as the Wolfcamp in Texas, and Cardium in Alberta.