The Oil Drum
DISCUSSIONS ABOUT ENERGY AND OUR FUTURE

Water in the Gas Tank

Posted by Stuart Staniford on March 26, 2007 - 10:24am
Topic: Supply/Production
Tags: ghawar, peak oil, saudi arabia [list all tags]

Further Forensics on Saudi Oil Supply

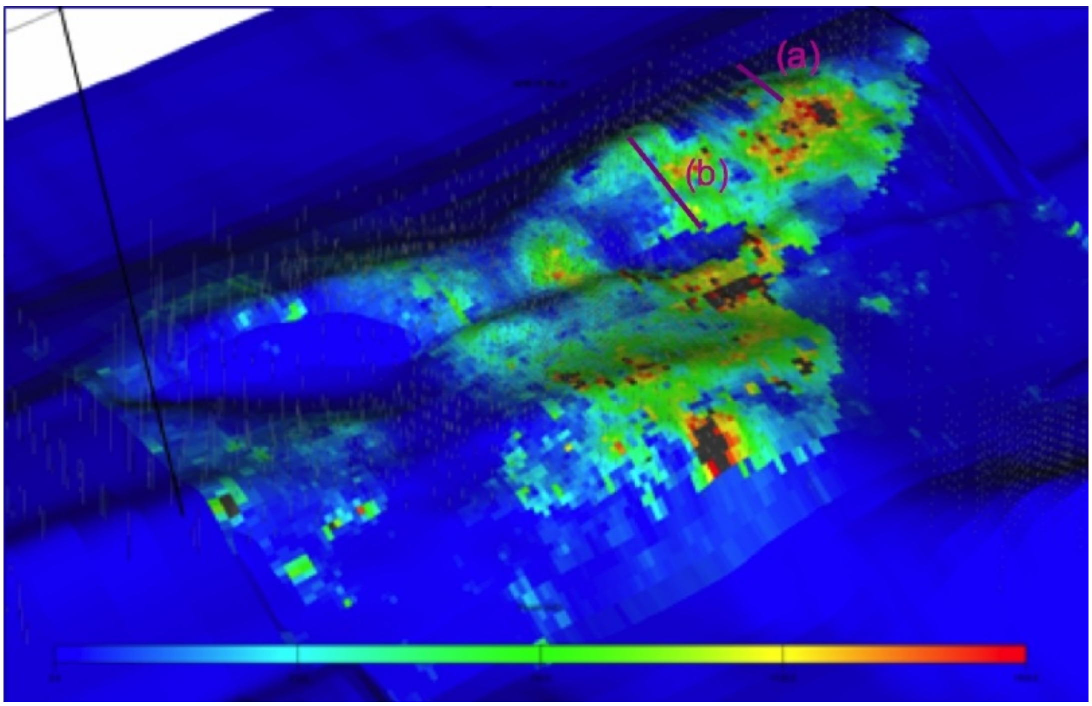

Reservoir simulation visualization of rock permeability in 'Ain Dar and Shedgum regions at northern end of Ghawar field. Scale runs from 0 (blue) to 1500 Millidarcies (red). Source: Figure 12 of Hussain et al, International Petroleum Technology Conference Paper #10395, November 2005.

Background

I have been arguing recently that, since about the middle of 2004, Saudi Arabian oil production has been supply constrained, and that in particular the production declines since the middle of 2005 are not voluntary. See Saudi Arabian Oil Declines 8% in 2006 and A Nosedive Toward the Desert. Also see, for a contrasting view, Euan Mearn's posts Saudi Arabia and that $1000 Bet, and Saudi Production Laid Bare. Behind the scenes, Euan and I have exchanged a lot of email trying to figure out the points of agreement between us, as well as the remaining areas of disagreement.

I see that Cambridge Energy Research Associates continues to have a radically different view of future (2015) oil supply:

> Saudi Arabia was ranked No. 1. Its output was forecast by Cambridge to grow to 14.3 million barrels per day from 2005 output of 12.7 million bpd.

I suppose if one thinks they produced 12.7mbpd in 2005, then one is in fantasy land about the past, never mind the future.

I'd like to start out this analysis by creating a lettering system for regions of interest in the Saudi production curve. The hope is to give us a common terminology for what we are referring to. I have lettered each separate feature (as I see it - we could always subdivide later if needed), so that we can debate what is going on at each phase. Here's my scheme, superimposed on a graph that shows Saudi production at the bottom, oil rig count in the middle, and sales weighted price of OPEC exports at the top:

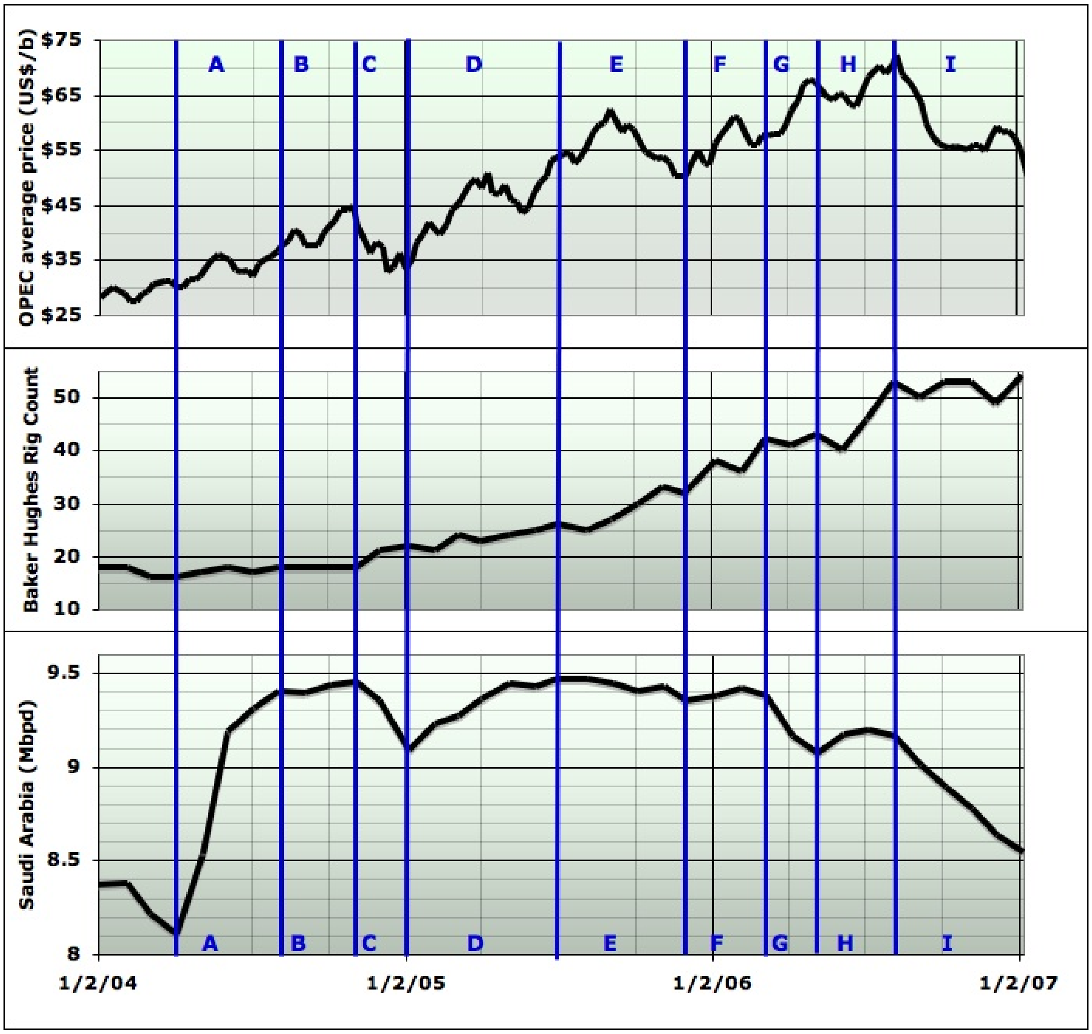

Top: Sales-weighted average crude price for OPEC exports from Jan 2004-Dec 2006 (from EIA). Center: Baker Hughes oil rig count in Saudi Arabia in same period. Bottom: Saudi Arabian oil production over the same period, average of four different sources (discussion here). Graphs are not zero-scaled to better show changes. Click to enlarge.

I am focussing now on the period since the beginning of 2004, since there is no controversy about what was happening prior to that. Everyone agrees that, before the second quarter of 2004, Saudi Arabia had at least some spare capacity and was acting as the swing producer. The regions I have labelled, and a brief description of discussion so far are as follows.

A:

This is the rapid rise in Saudi production from April 2004 to August 2004. I have suggested, based on analyzing the extensive discussion in OPEC Monthly Oil Market Reports, that this was a deliberate effort to moderate prices in the face of increasing demand and prices rising far above the agreed OPEC price band at the time. I'm not aware of any serious dissent on this explanation. Saudi Arabia has increased production to accomodate rising demand many times in the past, so this behavior was "in character".

B:

From August 2004 to November 2004, Saudi production goes flat despite the fact that prices have not been brought under control and are still rising. My interpretation is that this level of production was the maximum possible at the time. Others seem to believe that in general perhaps the Saudis changed their approach and decided higher prices were fine and they'd prefer to make more money by voluntarily reining in production. I'm not sure if anyone thinks this decision was taken as early as August 2004, but I think to explain the data in **B** as voluntary holding back, one would have to make that assumption (which is fairly problematic since the Saudis would have no basis for changing their view of the effect of these prices on the world economy yet). I believe Euan may agree with me that this is likely the maximum plateau production at the time. It's also potentially significant that the rise in rig counts begins at the end of this period **B**. Given the lead time to get a rig to Saudi Arabia, it's likely that the decision to start building up the rig fleet was taken sometime during period **B** (or possibly earlier).

C:

From November 2004 to January 2005, there is a sudden and fairly sharp decline in production by about 350kbd. This coincides with a sharp price break down from $45 to $35 (in this series which is generally a little lower the the oft quoted NYMEX front month, or West Texas Intermediate spot price) I think at this point Euan and I are in relatively close agreement that this is probably something along the lines of "Phew, prices are dropping, let's rest some of our more tired/problematic wells."

D:

From January 2005 to July 2005, there is a gradual and slowing rise in production that roughly restores the November 2004 level. This occurs as prices are rising from $35 to $55. A mystery is why some of the 690kbd of additional production capacity from the Qatif megaproject does not allow for higher production during this period. I lean to the explanation that something else declined, but specific attribution is not possible at present (see discussion below). People who believe in voluntary restraint could argue that the Saudi's had by now decided the world could cope with $45 oil and wanted to see what $55 oil would do. However, the specific pattern of production rise seems very problematic for this explanation. As prices rise from $35, the Saudis increase production at the fastest rate we have seen since (suggesting an attempt to respond to demand and moderate prices), but then as prices continue to rise, they respond less and less. Why increase production at all if they just wanted prices to go up? And given that they decided to increase production, what are the implications of the gradual creep up, from a country that in the past could jump production by 1-2mbpd in a month or two? Rig count rises slowly but steadily through this period. (As an aside, the Oil Drum is founded in March during this interval as people everywhere start to wonder what is going on with oil markets. Twilight in the Desert is published in June).

E:

From July 2005 to Dec 2005, production is flat and then begins gently declining (continuing the deceleration trend of period **D**). This starts to occur in the face of sharp price increases to the 2005 peak of around $62 in September, and there is no noticeable break in the trend line as prices then fall back to to around $50 by December. I will argue below that the likeliest hypothesis for this decline is the beginning of the cresting of the flood front in North Ghawar. I am unclear how the voluntary restraint hypothesis can explain both the decision to push prices up to $62 by dropping production, and then allowing it to fall to $50 by not dropping production faster as price falls. Rig counts rise more rapidly through this period.

F:

There is a very slight production hike from Dec 2005 to February 2006. I now have ascertained that the Qatif redevelopment megaproject was not actually completed in all respects until December 2005. So it is possible that this represents a boost from some last portion of the Qatif redevelopment. Rig counts continue to rise.

G:

Production resumes declining , more rapidly now, from Feb to May 2006. This is occurring in the face of prices increasing sharply. I believe this is likely due to worsening problems in North Ghawar, though that can't be proven beyond a reasonable doubt. The voluntary restraint hypothesis would require that the Saudi's are deliberately dropping production to drive prices from the upper $50s to the upper $60s.

H:

Production increases again from May 2006 to July 2006, as prices first drop, and then crest in the $70s on fear of war with Iran or hurricanes in the Gulf. I interpret this as most likely the onset of the Haradh megaproject. The voluntary restraint hypothesis requires that the Saudis now change their minds and make a (very modest) production increase to try to ameliorate the price increases they have just been driving a few months ago. Rig counts continue to increase rapidly.

I:

Production and prices both continue to decrease. There may be scope for both explanations here, as problems continue in North Ghawar, but is also true that OPEC begins to talk of, and then carry out, small production cuts to prevent the price going too much lower than the levels they have now got used to (ie around $60). A few hundred kbd of Saudi production cuts might be accounted for this way. Rig counts have plateaued for the moment.

Overall, we are in need of about 2mbd of explanation - 1mbpd of net reduced production, 690kbd of additional C&C production from the Qatif/Abu Safar redevelopment that didn't show up in the top line, and 300kbpd of Haradh III redevelopment that also must have been offset. The possible classes of explanation are declines in existing production, failures of megaprojects to meet plans, inaccurate or misleading data, and voluntary production restraint.

Let me now turn to documenting some of the things I alluded to above.

Ghawar

In response to Euan's latest post, we received a flying visit from Fractional_Flow, an experienced reservoir engineer, who provided a link to a very interesting paper as well as a few key insights. This led me to dig further into the recent petroleum engineering literature on the Ghawar field since the publication of Matt Simmon's book Twilight in the Desert. This exercise has helped a lot in fleshing out what I think is happening and I'd like to discuss what I found. The papers themselves are copyrighted and behind a paywall, but I will link to the abstracts and make some fair-use excerpts.

It's also probably worth stressing here that I am **not** and never have been a reservoir engineer. I was academically trained in physics and computer science, and have mainly worked in the latter field. I am able to understand reservoir engineering textbooks and papers, but it's possible that lack of deep knowledge and operational experience may lead to errors. Unfortunately it seems that the world's reservoir engineers mostly work for someplace that doesn't want them discussing sensitive topics in public - but I encourage any of them to point out any errors in interpretation I might have made, and I'll do my best to fix them.

Let's start though by reminding ourselves of the basic facts about Ghawar (also see [Greg Croft's summary](#)).

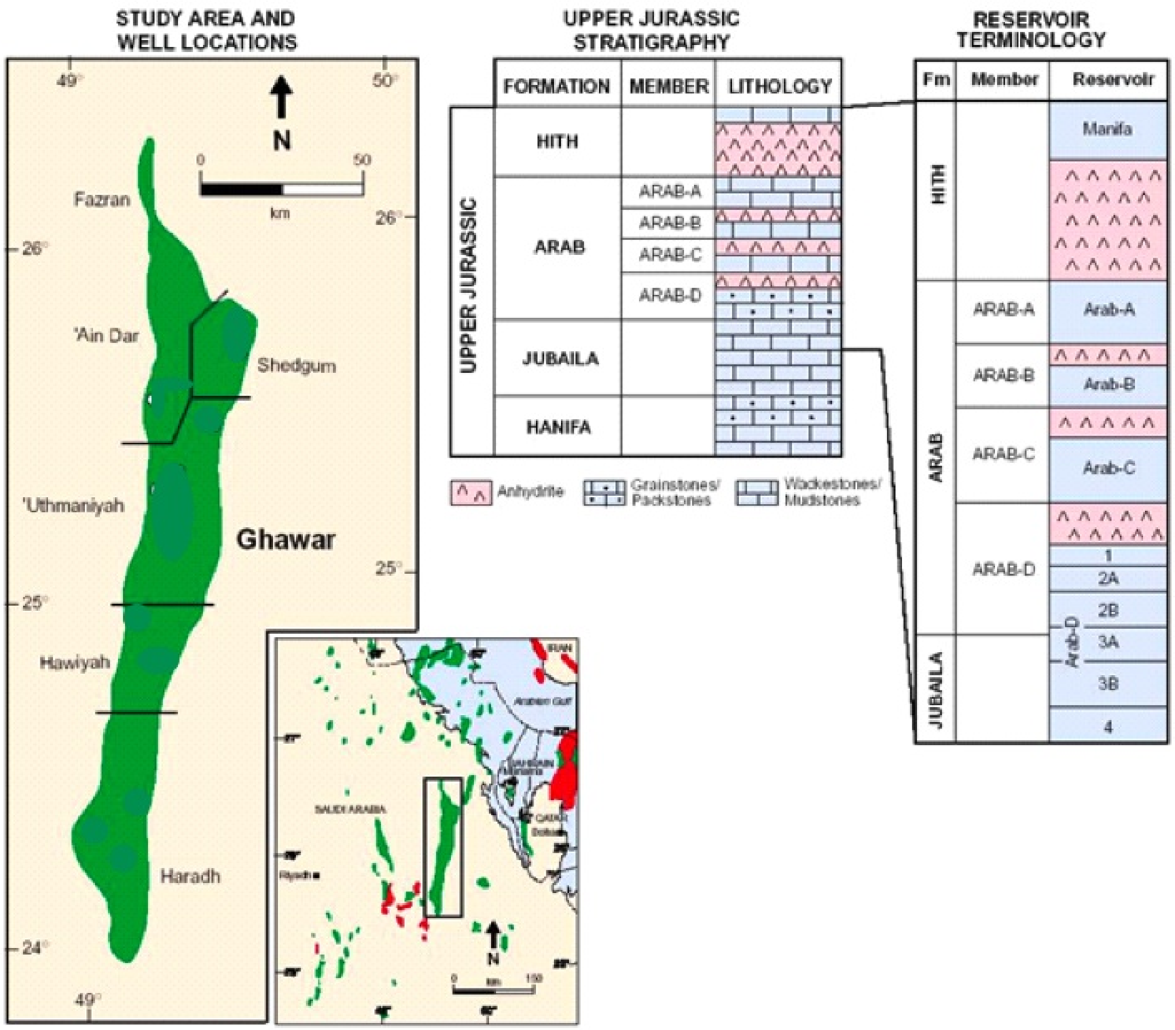

Summary of Ghawar operating areas and geology. Source: Figure 1 of Okasha et al, Society of Petroleum Engineers [Paper #105114](#), March 2007.

The whole structure is about 175 miles long, and around 20 miles wide at the widest - by far the largest field in the world. The oil in Ghawar is almost all in the pores of the Arab D carbonate rock shown in the Figure above (on which more in a moment). There are five major operating areas of Ghawar: 'Ain Dar, Shedgum, Uthmaniyah, Hawiyah, and Haradh. 'Ain Dar and Shedgum were the highest quality (highly permeable reservoir rock and sweeter crude), while the quality of both rock and crude tends to decrease from north to south. Accordingly, development of the field started at the north end, and each wave of new technology has tended to be applied at the top end first and then work south.

Let's look at 'Ain Dar and Shedgum first. This next picture is believed to show the permeability of the two regions. Ain Dar is the narrower longer ridge to the top (west) of the picture, while Shedgum is the broader shorter lump lower in the picture (ie to the east):

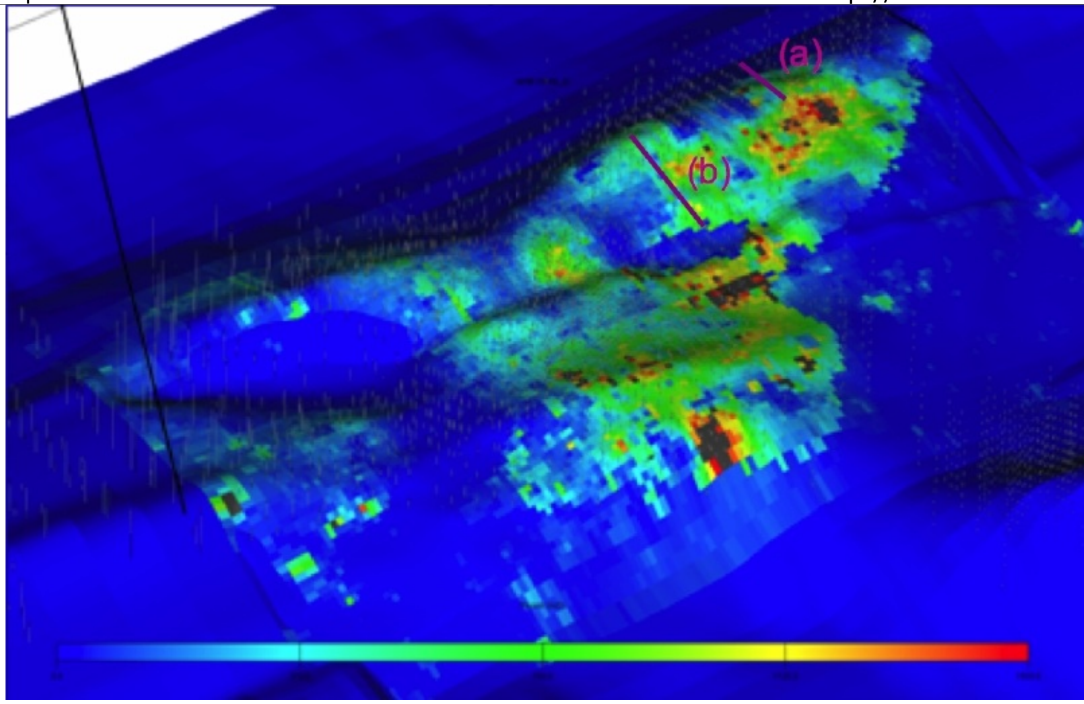

Reservoir simulation visualization of reservoir permeability in 'Ain Dar and Shedgum regions at northern end of Ghawar Field, probably as of mid 2005. Scale runs from 0 (blue) to 1500 Millidarcies (red).. Inferred approximate locations of reservoir simulation cross sections shown also. Source: Figure 12 of Hussain et al, International Petroleum Technology Conference [Paper #10395](), November 2005.

The two regions together are almost 45 miles from north to south, and around 20 miles east to west. The variable we are looking at, permeability (often denoted K) is related to how easy it is to push fluid through the reservoir rock. The more pressure drop across a given volume of rock, the more fluid is going to flow through it (just as when you suck harder on a straw, you get more to drink). Part of the constant of proportionality between the pressure and the flow is called the permeability of the rock (the other part is the viscosity of the fluid). It is measured in Darcies (after the French scientist who developed the most widely used equation of motion for fluids in porous media). Rock with a permeability of one Darcy (equivalently 1000 Millidarcies) will allow for a flow of 1 cm/s of a fluid with viscosity 1 centiPoise under a pressure gradient of 1 atm/cm. To make sense of the "1 centiPoise part", note that water has a viscosity of 1 centiPoise at 20^{o}C. Ghawar oil is a little less viscous than water - varying from 0.62 centipoise in 'Ain Dar to 0.89 in Haradh. Highly permeable rock allows much more rapid production of oil (and of water).

As you can see, much of North Ain'Dar and Shedgum have very high permeabilities in the range of 750 mD. South 'Ain Dar has much lower permeability and is presumably far less productive in consequence. Also note the red and black regions in the picture. These are areas of very high permeability which likely allowed for the very high oil production early in the life of some Saudi wells, but also can channel large volumes of water later in the life of the field.

Data for the production history of 'Ain Dar and Shedgum together were given by Baqi and Saleri in their presentation to CSIS in February 2004 (the presentation that was essentially trying to refute Matt Simmons' concerns). Those data look as follows:

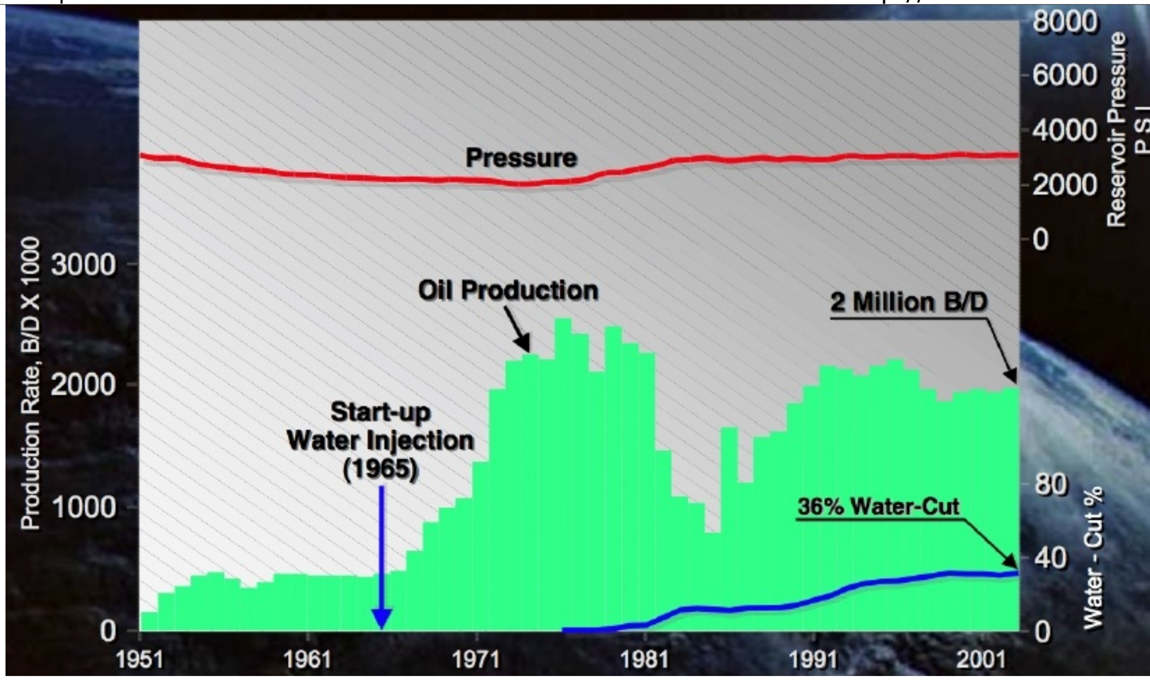

Saudi Aramco summary of production history for 'Ain Dar and Shedgum together. Source: [Fifty-Year Crude Oil Supply Scenarios: Saudi Aramco's Perspective](), Presentation to CSIS, Feb 24th, 2004.

As you can see, production has recently been running at around 2mbd. This graph shows annual data through 2002. If we look at the [Water Management in North 'Ain Dar]() paper, we get an annual production series just for North 'Ain Dar, which suggests that region was about one quarter of the total production of the whole 'Ain Dar/Shedgum north end:

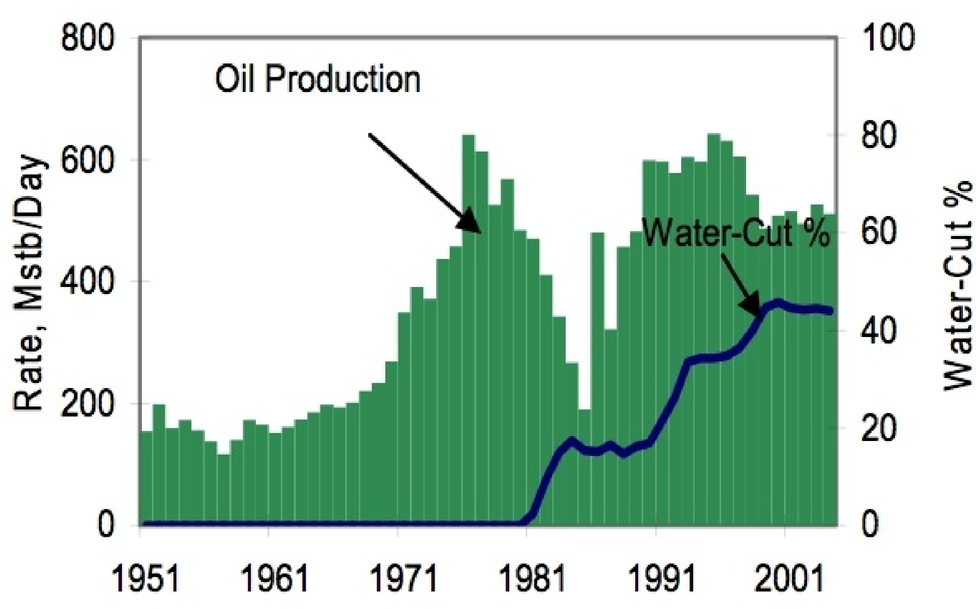

Fig. 1-Historical oil production and water-cut.

Production and water-cut history in the North 'Ain Dar region of Ghawar. Source: Figure 1 of Alhuthali et al, Society of Petroleum Engineers [Paper #93439](), March 2005.

Here the production data represent annual data through 2004. Assuming this data is accurate, production fluctuations in North 'Ain Dar cannot have contributed in a large way to changes in overall Saudi output through 2004 (though it is perhaps useful to note that 2004 is slightly down on 2003, whereas Saudi Arabia as a whole was up from 2003 to 2004 on an annual basis). At any rate, there is no evidence here of major trouble in North Ghawar production as of the end of 2004.

However, there are other indications in the paper suggesting that trouble cannot be far off at all. Let's look at this picture of the state of the reservoir in 'Ain Dar. This is taken from a paper titled *Water Management In North 'Ain Dar, Saudi Arabia* by Alhuthali *et al* and published by the Society of Petroleum Engineers as Paper #93439 in March 2005. Fractional_Flow pointed us to a [bootleg copy]() that somebody put up, but if that disappears, you can get the [official version]() behind the paywall.

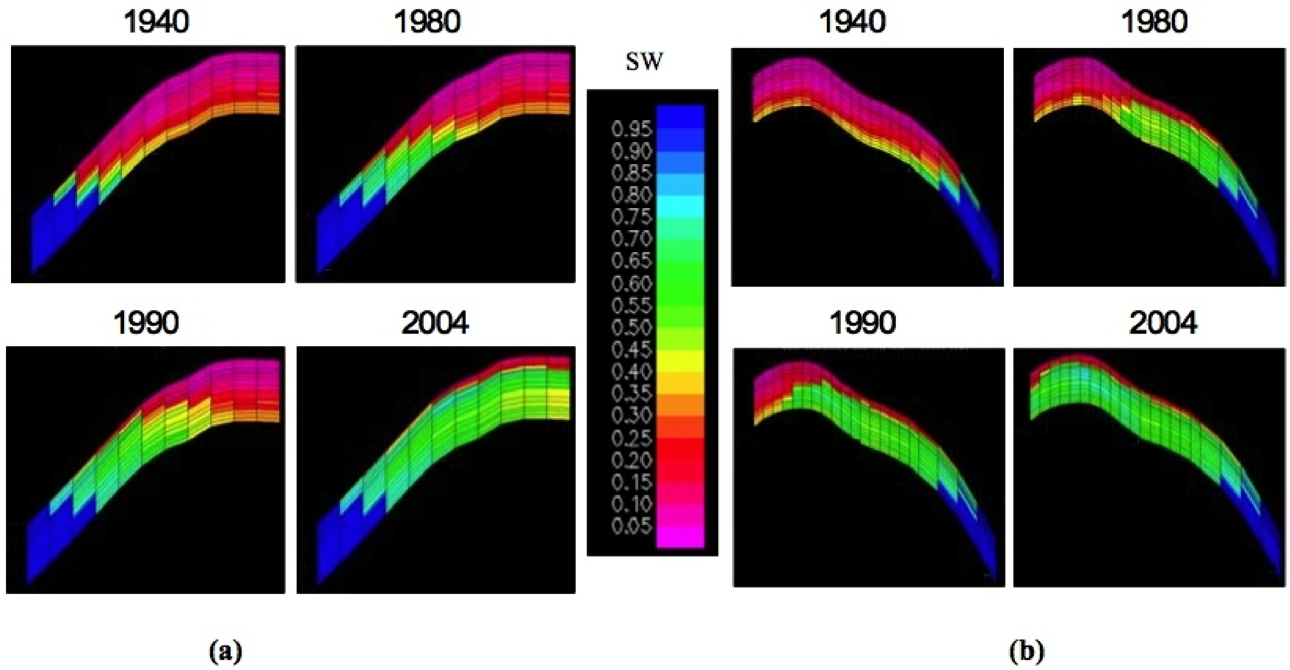

Two cross sections of a reservoir simulation of the northern portion of the 'Ain Dar region of Ghawar at various years. Color represents volumetric water saturation in the rock pores. Source: Figure 9 of Alhuthali et al, Society of Petroleum Engineers [Paper #93439](), March 2005.

This is from a numerical simulation of the development of the reservoir over time. However, Saudi Aramco history-matches their reservoir models with extensive amounts of well log data, so it is probably a reasonably accurate picture of the history of the field. The variable being plotted is water saturation - that is the percentage of the pore volume in the rock that is filled with water, rather than oil. In Ghawar, the reservoir pressure has been maintained such that there is no gas cap (the gas is dissolved in the oil), and so the pores are filled with oil, water, or a mixture. You can get the general idea of the figure - at the beginning (in 1940) the crest of the reservoir was almost all filled with pink (ie oil with less than 5% water) and red (less than 25%). As time has gone on and more water has been injected at the periphery, there is less and less dry oil like that, and instead there are green and pale blue areas which are 50%-80% water by volume.

Now, your first thought might be this: if there is 10% water and 90% oil in a particular volume of rock (pink areas in the figure above), then a well into that part of the rock would be receiving 10% water and 90% oil. Similarly, an area with 60% water and 40% oil might be producing at 60% water cut into a well into that area. However, this is not so: the difference is much more dramatic than that. The reason has to do with the physics of two phase flow in a permeable medium. If you want a mathematical treatment, [try this](), but let me try to illustrate the basic idea.

In a set of interconnected pores through which oil and water are being forced at pressure, the flow is too turbulent for large areas of the two fluids to separate out from one another. And yet, oil and water do not like to mix, and will tend to bead up in the presence of the other. If there is only a little water and a lot of oil, then the oil will form an interconnected network of fluid throughout the rock pores, whereas the water will tend to make small beads within the oil. Conversely, a little oil in a lot of water will result in a network of water throughout the rock, and small beads of oil within that network. Now, in either situation, the fluid that is interconnected can flow through the rock without making any change in the arrangement of beads and surfaces between oil and water. However, the fluid that is beaded up can only move by the beads physically moving around, and they are going to tend to get trapped by the rock pores.

So for this reason, in a mixture of almost all oil, the water cannot flow *at all*. Conversely, once there is almost all water, the oil cannot flow at all (which sets an upper limit on the amount of oil that can ever be recovered by a water flood). In between, there is a changeover in which the proportion of oil flowing to water flowing changes much more rapidly than the changeover of the actual mixing ratio. The curve that describes this is called the fractional flow curve.

For example, the tutorial I referenced earlier shows this picture for a typical fractional flow curve:

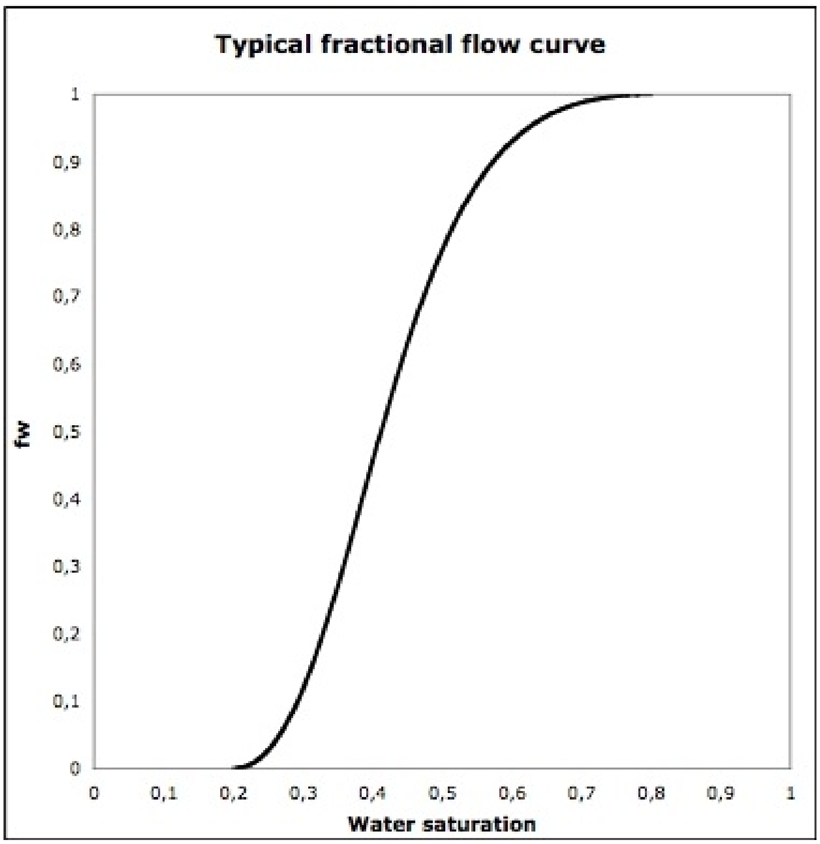

"Typical" fractional flow curve (from this tutorial).

So the way to read this is that when we are below 20% on the X-axis (less than 20% water in the oil), there is zero on the y-axis (the water will not flow through the rock at all). As we get above 20% water saturation, the flow of water increases rapidly, until above 80% water, there is no flow of oil at all. In the linear region at the center of the curve, the slope is about 3.6. That is, each 1 percentage point increase in water saturation results in a 3.6 percentage point increase in water flow in the rock.

So that implies that the difference between the red regions (20% water saturation) in the simulation picture above and the green regions (60% water saturation) could be very large. That might be the difference between hardly any water flowing, and almost all water flowing. But does this abstract "typical" fractional flow curve really apply to Ghawar? Or might things be very different there?

Well, I have managed to find in another SPE paper some actual data from which I was able to infer the fractional flow curve for three regions of Southern and Central Ghawar. In a different references, I found a carbonate fractional flow curve from the United Arab Emirates.

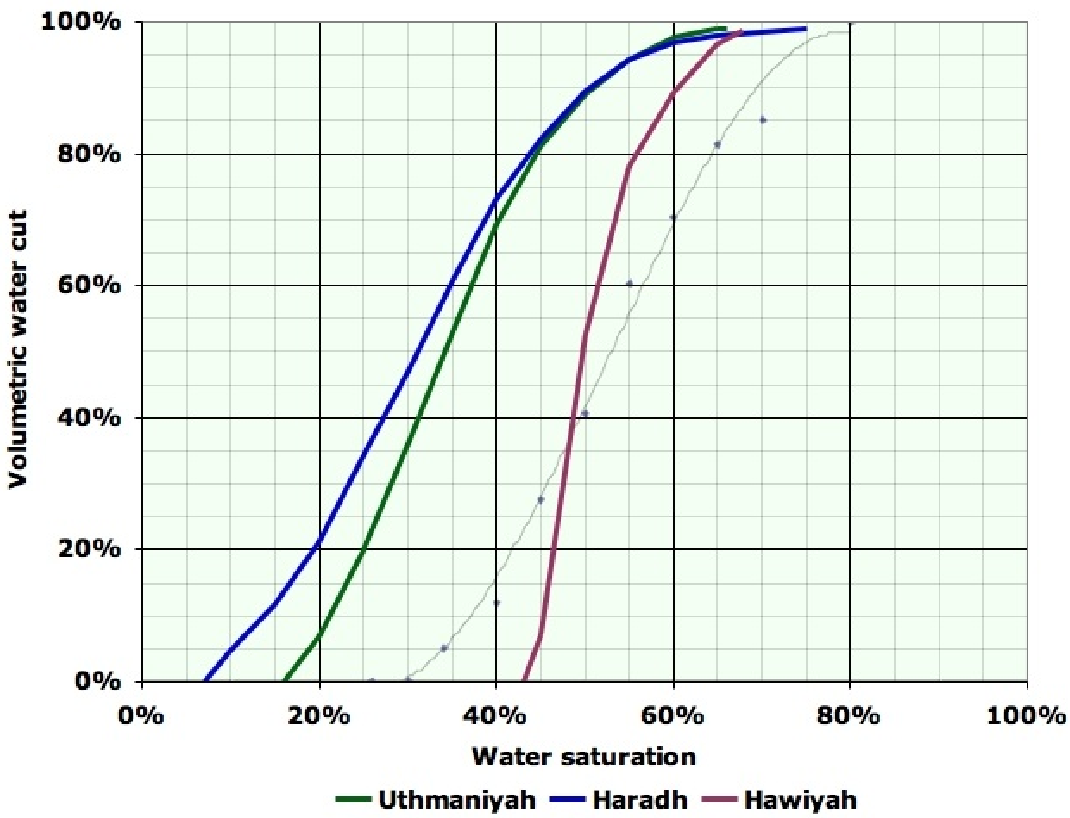

Fractional flow curve for three regions of Ghawar, with background of a fractional flow curve for carbonate rock in the United Arab Emirates. Source for Ghawar: Author's calculations[1] based on Figure 10 of Okasha et al, Fifty Years of Wettability Measurements in the Arab-D Carbonate Reservoir, Society of Petroleum Engineers Paper #105114, March 2007. UAE source, Slide 15 of this Powerpoint lecture.

Clearly, these have the same general character as the "typical curve" above, and again there is going to be a dramatic difference between red regions and green regions in the cross sections above.

Now, we do not have quantitative fractional flow curves for 'Ain Dar and Shedgum at the moment. However, we know one point on the curve for 'Ain Dar from Alhuthali *et al* (Paper #93439), which reports on measurements of rock right next to an abandoned water injector well. The pores of that rock had 21% oil, 79% water. Rock next to a water injection well is bound to contain only that oil which is never going to move under any further waterflooding, and so tells us that the fractional flow curve must reach the top of the graph (y-axis of 100%) when the x-axis is at 79% water saturation. This appears to be roughly consistent with the other curves on the graph.

Thus, we can conclude that as long as there is red/pink in the simulations above, it will be possible to produce oil at very low water cuts by drilling horizontally into those oil layers. The overall water cut of the field can be somewhat controlled by making decisions about how much to produce out of the red-pink zone, and how much to produce from the watery green areas. Once the red/pink zone is gone, then there will be no way to produce oil except out of those green areas, which probably correspond to water cuts in excess of 70% (perhaps much higher depending on the exact shape of the fractional flow curve in this part of Ghawar).

Now, how long before the red zone is gone? Let me remind you of the simulation picture again, to save on scrolling:

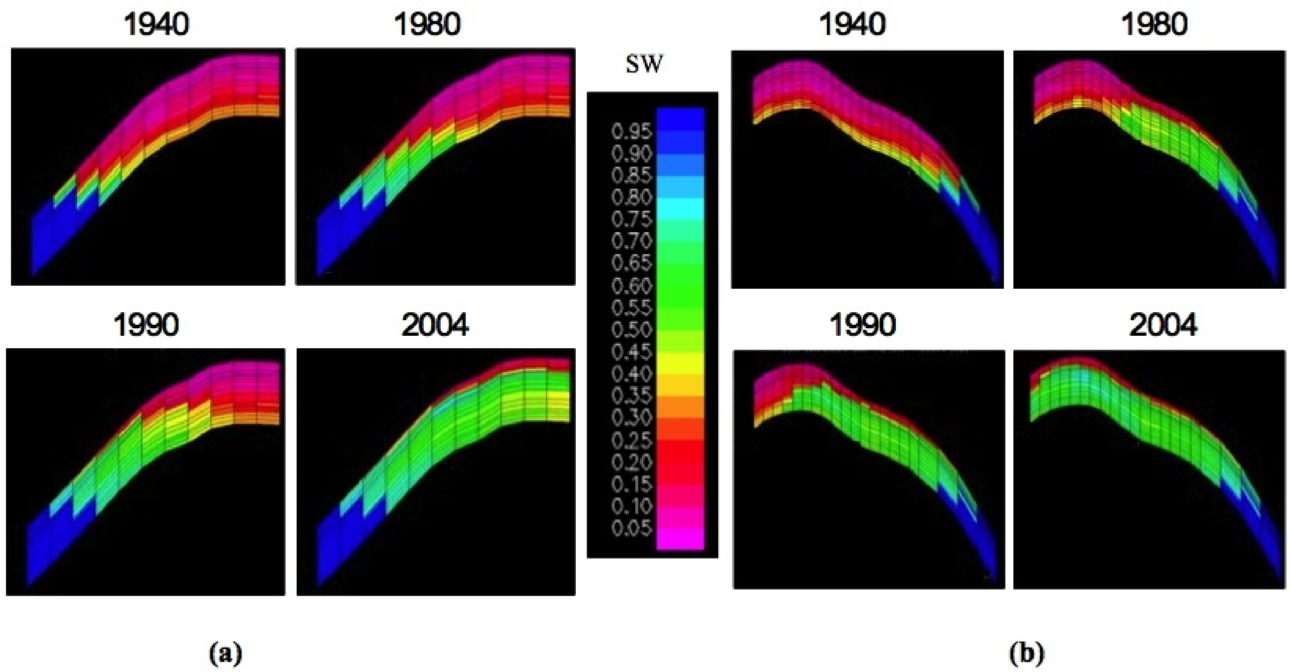

Two cross sections of a reservoir simulation of the northern portion of the 'Ain Dar region of Ghawar at various years. Source: Figure 9 of Alhuthali et al, Society of Petroleum Engineers [Paper #93439](), March 2005.

Now, after staring at the bumps in these figures, I think I have figured out roughly where in North 'Ain Dar these two cross sections -- a) and b) -- are. I have indicated them on the permeability visualization I showed before:

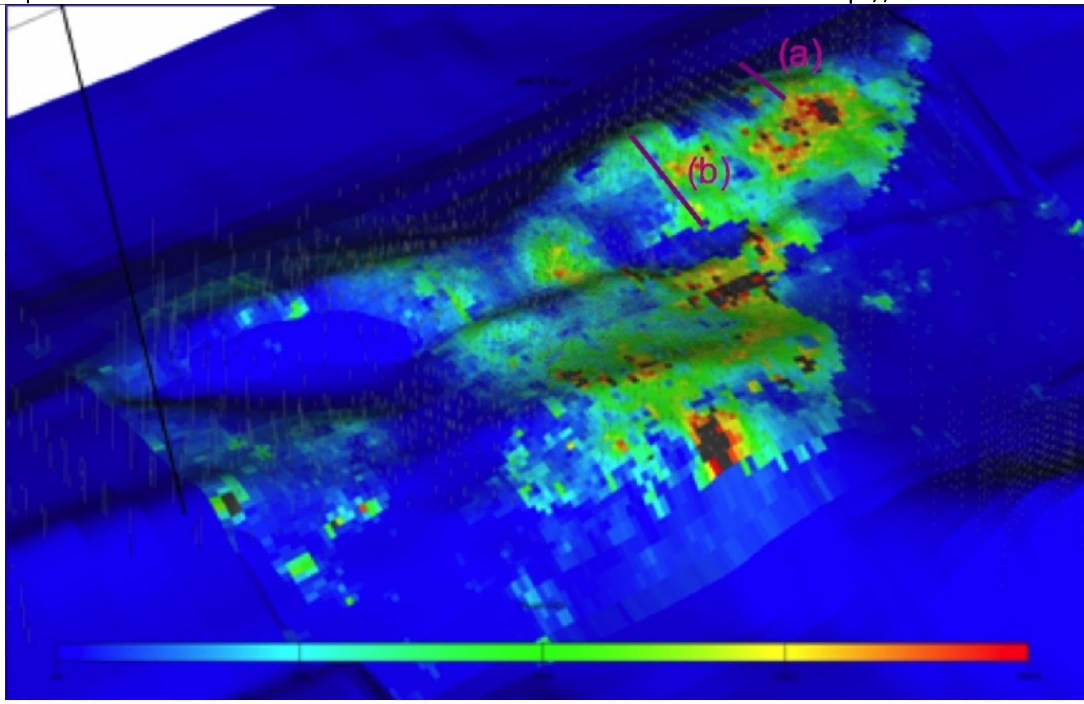

Reservoir simulation visualization of water saturation level in 'Ain Dar and Shedgum regions at northern end of Ghawar Field, probably as of mid 2005. Inferred approximate locations of reservoir simulation cross sections shown also. Source: Figure 12 of Hussain et al, International Petroleum Technology Conference [Paper #10395](), November 2005.

These placements are not exact - they could easily be a few miles north or south of where I have shown them, but I am pretty confident that I am in the right general region. Now, the conclusion I draw from this placement is as follows: the cross sections are reasonably representative of the crest of North 'Ain Dar. As the cross sections show only a reasonably thin layer of red/pink oil that, in two dimensions, is hugging the top of the reservoir, it seems very likely that, in three dimensions, we only have a layer of roughly that thickness throughout the entire crest of North 'Ain Dar.

So how long is that layer going to last?

Well, if we blow up just the 1990 and 2004 pictures for the (a) cross section, that is the western flank, we get this picture:

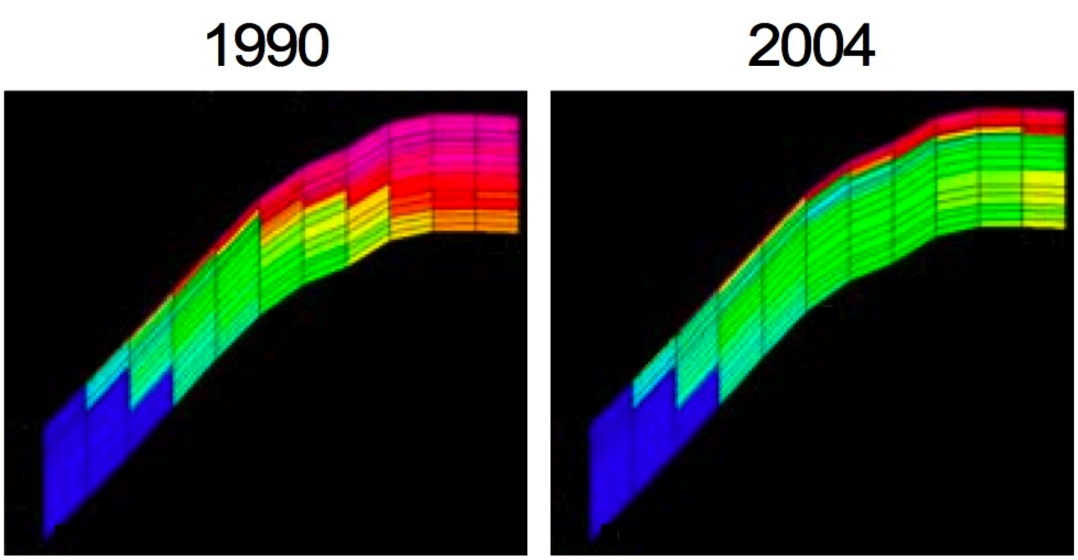

West flank cross section in reservoir simulation of the northern portion of the 'Ain Dar region of Ghawar. Blow-up of 1990 and 2004 only. Source: Figure 9 of Alhuthali et al, Society of Petroleum Engineers Paper #93439, March 2005.

So my layer counting says that over 14 years we went from 10 layers of red/pink in the crest to 2 layers of red in the crest. And if we do the same for the (b) cross section on the eastern flank, we get this:

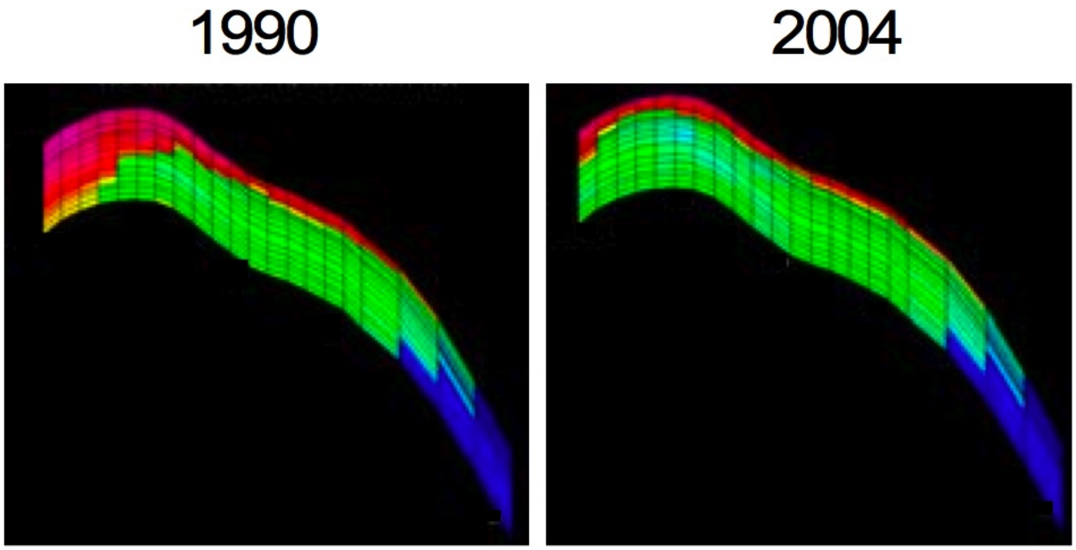

East flank cross section in reservoir simulation of the northern portion of the 'Ain Dar region of Ghawar. Blow-up of 1990 and 2004 only. Source: Figure 9 of Alhuthali et al, Society of Petroleum Engineers Paper #93439, March 2005.

Layer counting here says that 6-9 layers of red/pink have been reduced to 1 layer over the same 14 years, suggesting a similar pace. So, overall, we lose a layer about every 21 months, and there were 1-2 layers left at the time of the 2004 simulation. Now, we do not know for sure whether this simulation means by "2004", "2004.0" -- the beginning of the year, or the end of the year. However, labels on simulation visualizations in other papers shows annual data with a ".0", suggesting that is the convention. If so, then we would expect the green area to reach the top of the reservoir in some places in North 'Ain Dar in mid-late 2005, and in most places there by the middle of 2007. Thus in timing, this phenomenon would be consistent with the decline in output across production zones E-I (mid 2005 through the present).

It does not appear to me that, from a timing perspective, this phenomenon is well placed to explain any declines that offset whatever fraction of the 690kbd of Qatif megaproject production came on in late 2004 and early 2005 (zones B-D). We have production data for North 'Ain Dar which do not show more than a very small decline in 2004 on an annual basis, so any sizeable decline here could not have occurred before 2005.

Euan has questioned whether an increase in water cut like this could have a large enough impact on production to explain all the declines we see. Let me try to make an estimate on how large it could be. We cannot do this with any precision from the fractional flow curve because we don't have an exact fractional flow curve for this part of Ghawar. However, there is other data in Paper #93439 which allows us estimate this. In particular, the paper contains a discussion of what fraction of production comes from behind the flood front, versus in front of the flood front. They discuss the desirability of trying to produce from behind the flood front as much as possible in order to maximize overall recovery of the oil. However, it is clear from Figure 10, that they have been obliged to produce more and more oil from **above** the flood front in order to maintain production. Specifically, the fraction of production below/behind the flood front is shown as follows:

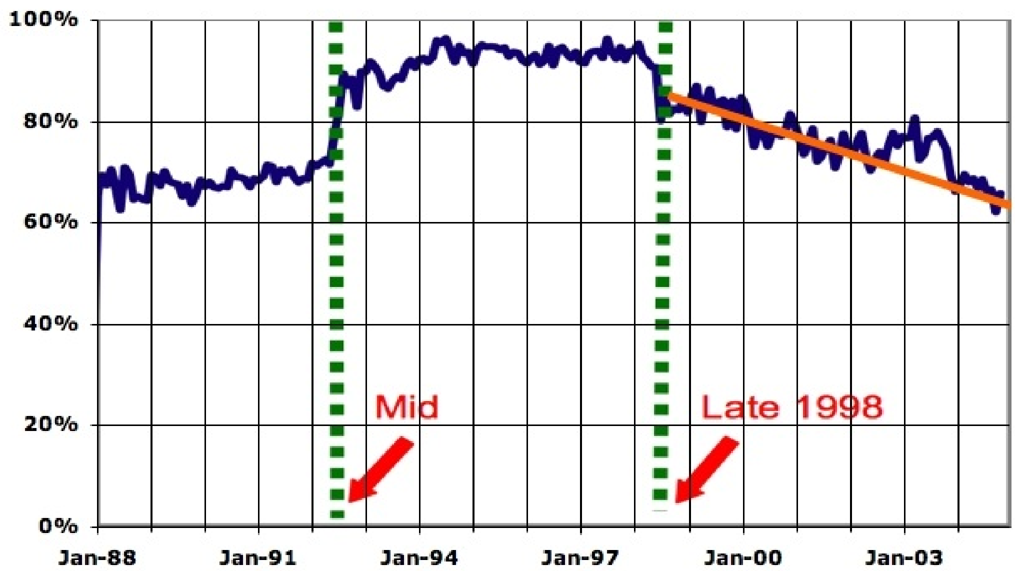

Fraction of oil produced from behind the flood front in North 'Ain Dar. Source: Figure 10 of Alhuthali et al, Society of Petroleum Engineers Paper #93439, March 2005. Orange line is my addition.

Now, this is dropping, presumably because the water behind the flood front has higher and higher water cut over time, and the water handling facilities are limited. If we extrapolate out the orange line to today, only 55% of production would be coming from behind the flood front. However, once the flood front has reached the crest (which I argue must already have happened already in most places), then there can be **no** production from above the flood front. This line of reasoning suggests that the flood front reaching the crest would produce around a 45% reduction in production over a period of a couple of years.

Now, we only have enough data to make a convincing case for when this will occur in North 'Ain Dar, which represented only 500kbpd of production in 2004. However, it is at least plausible, though not proven, that the same thing could be happening in approximately the same time frame in the whole 'Ain Dar/Shedgum region. 'Ain Dar and Shedgum have very similar reservoir properties, very similar oil qualities, almost identical thickness of original oil layer, and were brought onto production at the same time. So it is plausible that they have been produced in parallel, and will be reaching the end of their production plateaus at around the same time.

If this were so, a 45% reduction in 2mbpd of production (as of 2002) would represent about 900kbd (with a large, say 50%, uncertainty in that number for a SWAG at the error bar). This might be enough to explain most of the decline in production from June 2005 to the present (though we would still require some other field in decline earlier in order to offset Qatif and Haradh III). and make up the rest of the overall production decline.

Some further circumstantial evidence for this point of view comes from Hussain *et al*, International Petroleum Technology Conference Paper #10395, November 2005. This paper is titled *Optimizing Maximum-Reservoir-Contact Wells: Application to Saudi Arabian Reservoirs* and covers the process of selecting locations in 'Ain Dar and Shedgum in which infill Maximum Reservoir Contact wells (MRCs) would produce incremental recovery over what could be recovered via the existing wells. For each potential MRC (basically a horizontal multilateral), a detailed reservoir engineering simulation is done on a fine scale grid to assess the impact on recovery and the economics of the proposed well. What is interesting is that these wells (which presumably represent the best remaining opportunities in the field where there is still some localized piece of the oil layer left to go after) are anticipated to begin declining very quickly.

The new well studied in 'Ain Dar, which was "planned to be spudded early 2006" is simulated to begin declining sharply almost immediately:

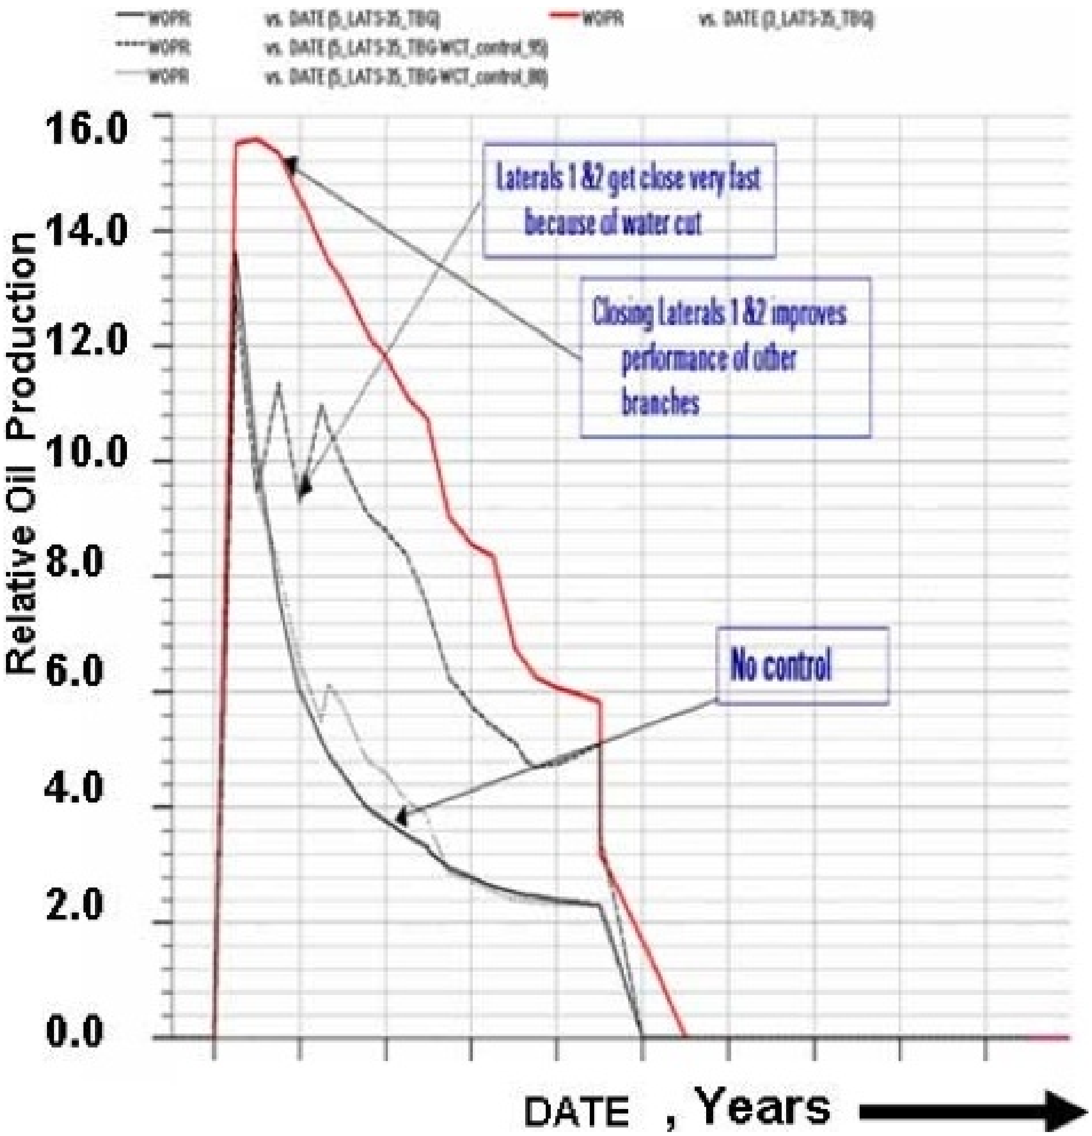

Simulated oil production profile for proposed horizontal multilateral well spudded early 2006 in 'Ain Dar. Source: Figure 22 of Hussain et al, International Petroleum Technology Conference [Paper #10395](), November 2005.

My interpretation of this is as follows: there is nowhere left that one could put an MRC into a reasonably dry oil layer anywhere in 'Ain Dar. If there was, Saudi Aramco would have illustrated their paper with that well, instead of this one, and that hypothetical well would have had a longer plateau.

Similarly, the new MRC studied in Shedgum was "spudded in late 2004", and was anticipated to be on plateau through the end of 2008, before beginning to decline rapidly:

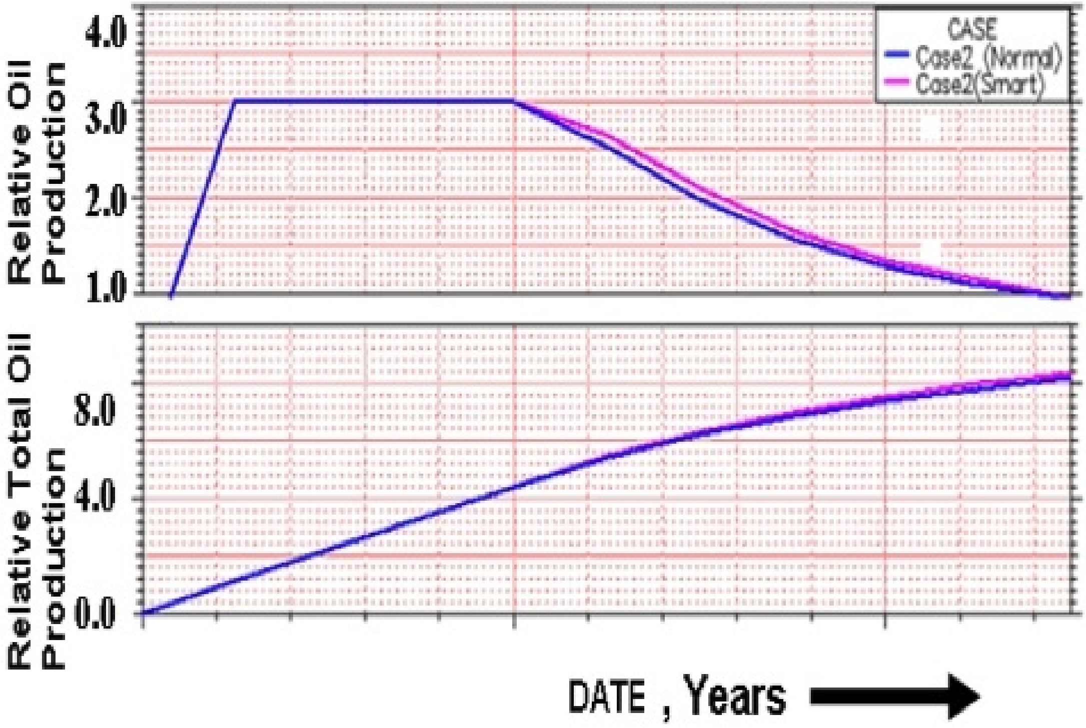

Simulated oil production profile for proposed horizontal multilateral well spudded late 2004 in Shedgum. Source: Figure 22 of Hussain et al, International Petroleum Technology Conference [Paper #10395](), November 2005.

My interpretation of this is that the last remaining spots of dry oil layer in Shedgum will be gone by the end of 2008, implying that there are probably some places in Shedgum where the oil layer is gone already.

Uthmaniyah

The situation in Uthmaniyah is less clear. There is a paper Water Production Management Strategies in North Uthmaniyah Area, Saudi Arabia, which covers some of the same ground as the North 'Ain Dar paper. However, what it particularly lacks is any production data, so we cannot estimate either how much Uthmaniyah has produced, or any trend in that production.

The paper does have some interesting hints of trouble. For example, "In August 2004, additional twenty-seven high water cut wells were put on cyclic production." Cyclic production means taking wells in wetter zones of the field off line for six months at a time to allow water to settle away from the well bores, and then running them again. It's interesting that this was happening right at the start of zone B, when Saudi production stopped rising even as OPEC was losing control of prices. However, 27 wells, which were presumably below the Saudi average at the time of 6kbd each, can only represent a number not much larger than 100kbd, and perhaps significantly smaller. So this is not evidence of much unless it were a harbinger of larger problems.

The paper does show cross-sections of a reservoir simulation again:

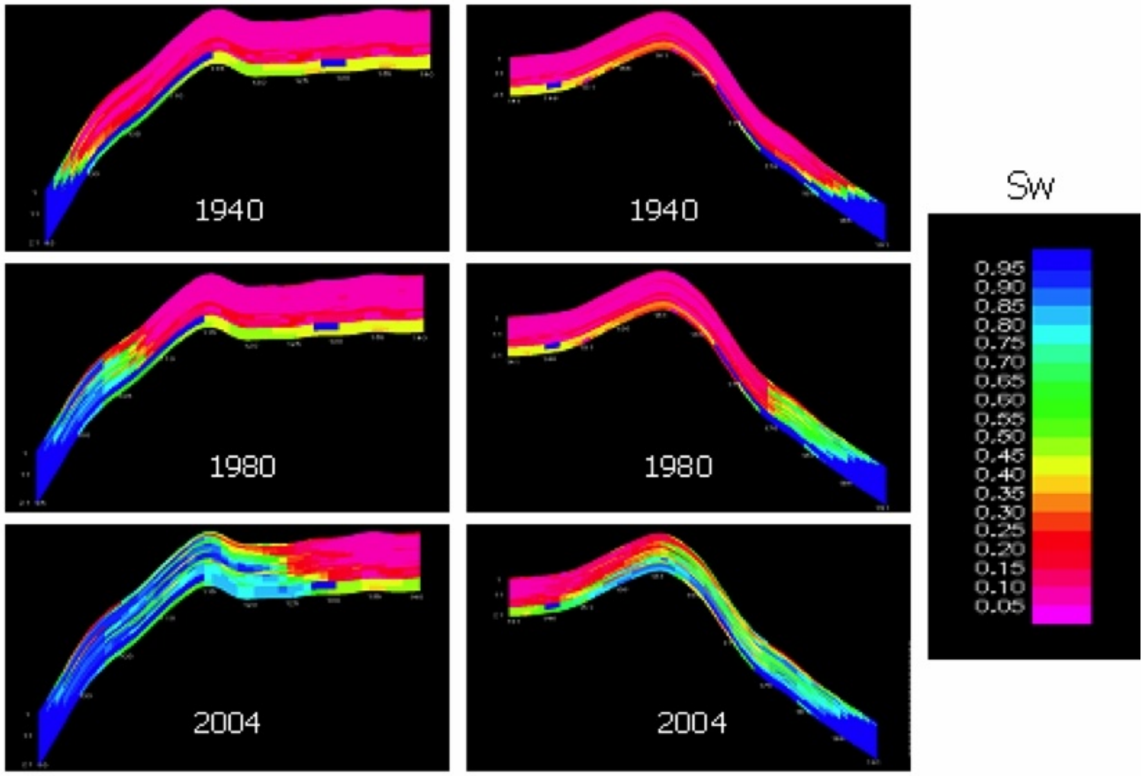

Two cross sections of a reservoir simulation of the northern portion of the Uthmaniyah region of Ghawar at various years. Source: Figure 12 of Al-Mutairi et al, Water Production Management Strategies in North Uthmaniyah Area, Saudi Arabia, SPE 98847, June 2006.

Several features are interesting. It looks as though there is more dry oil left in Uthmaniyah than in 'Ain Dar. However, it also looks like it is harder to flood uniformly: the simulation shows some streaking of water breaking through the oil layers that was not evident in the Ain Dar simulation. Some explanation for this is found in the fractional flow curve, which we actually have for Uthmaniyah:

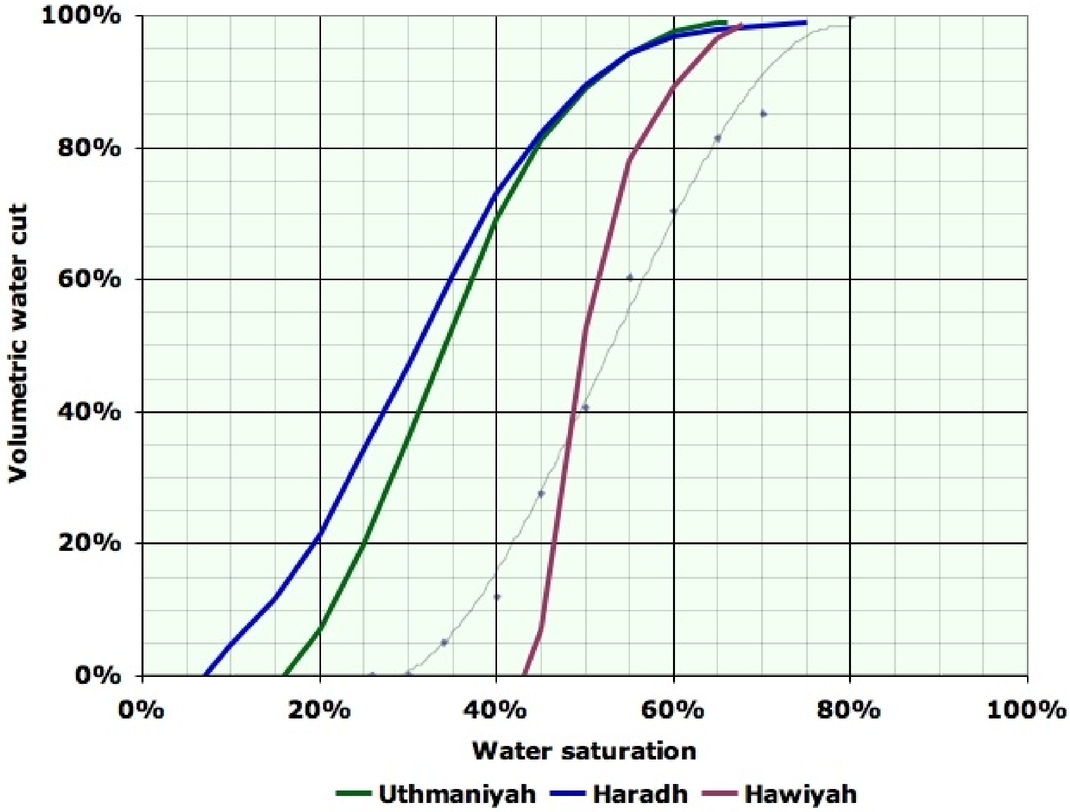

Fractional flow curve for three regions of Ghawar, with background of a fractional flow curve for carbonate rock in the United Arab Emirates. Source for Ghawar: Author's calculations[1] based on Figure 10 of Okasha et al, Fifty Years of Wettability Measurements in the Arab-D Carbonate Reservoir, Society of Petroleum Engineers [Paper #105114](), March 2007. UAE source, Slide 15 of [this Powerpoint lecture]().

As you can see, the fractional flow curve for Uthmaniyah is not too nice. Once we get out of the pink (ie above 15% water saturation), water movement starts, and by the time we are into green (above 50% water saturation) water cut is up to 90%. Thus production outside of the dry oil will be a slow business in Uthmaniyah, and ultimate recoveries are bound to be lower. If we blow up just the 1980 and 2004 pictures, we can try to estimate how much oil is left:

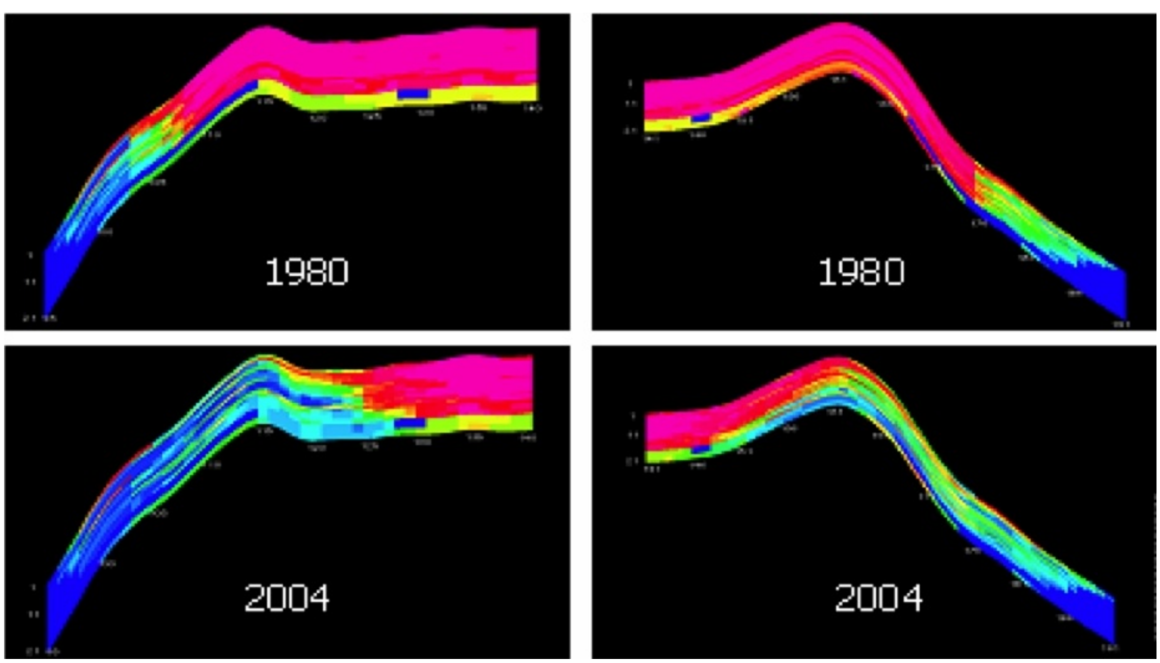

Blow up of two cross sections of a reservoir simulation of the northern portion of the Uthmaniyah region of Ghawar for 1980 and 2004. Source: Figure 12 of Al-Mutairi et al, [Water Production Management Strategies in North Uthmaniyah Area, Saudi Arabia](), SPE 98847, June 2006.

Eyeballing this, and assuming these profiles are representative, it looks like about a quarter of the pink oil there in 1980 was still there in 2004, suggesting plateau might be maintained for another 8 years give or take a year or two (and modulo periodic breakthroughs from those streaks). After that, it appears that water cuts will start to rise rapidly and within another five years production will inevitably decline greatly as there are only yellow and green regions to try and produce from with very high water cuts.

A Note on Well Productivity

Finally, I wanted to comment briefly on Euan's piece on well productivity, which I think might have been overly persuasive to some people. The first thing to note is that the main reason for the high well productivity of Saudi wells is that the rock is exceptionally permable by global standards. This allows a relatively small number of wells to drain a large volume of rock. It does not mean that once that rock has no dry oil in, production cannot decline rapidly, or that lots of further drilling into the depleted reservoir will necessarily do a great deal of good (it is noteworthy that the Saudis are putting in a fairly small number of very carefully planned wells in North Ghawar to extract the last oil they consider worth going after, rather than putting in large numbers of wells which would do little good).

The second thing to note is that the OPEC data Euan used is annual. Not only that, it ends right as the alleged problem is just getting going. This conceals any signal of the problem in the data. I did an exercise to infer what the monthly well productivity would look like through the end of monthly data. I did this by assuming that well retirements occur at the same rate in the monthly data as they did on average in the OPEC annual data from 2001-2005. I further assume that the rigs Baker Hughes counted in Saudi Arabia drill wells each month at the same average rig productivity as the OPEC annual data show from 2001-2005. I linearly extrapolate the ratio of oil to NGLs to continue as it has been going. I applied this procedure from the beginning of 2004 on, and its reasonableness is attested to by the fact that it predicts the OPEC reported number of producing wells at the end of 2004 and 2005 to better than 1%. (Further interpolation into the next year is based on the OPEC actual, rather than the monthly interpolated estimate from the prior year).

With those assumptions, the well productivity graph looks like this:

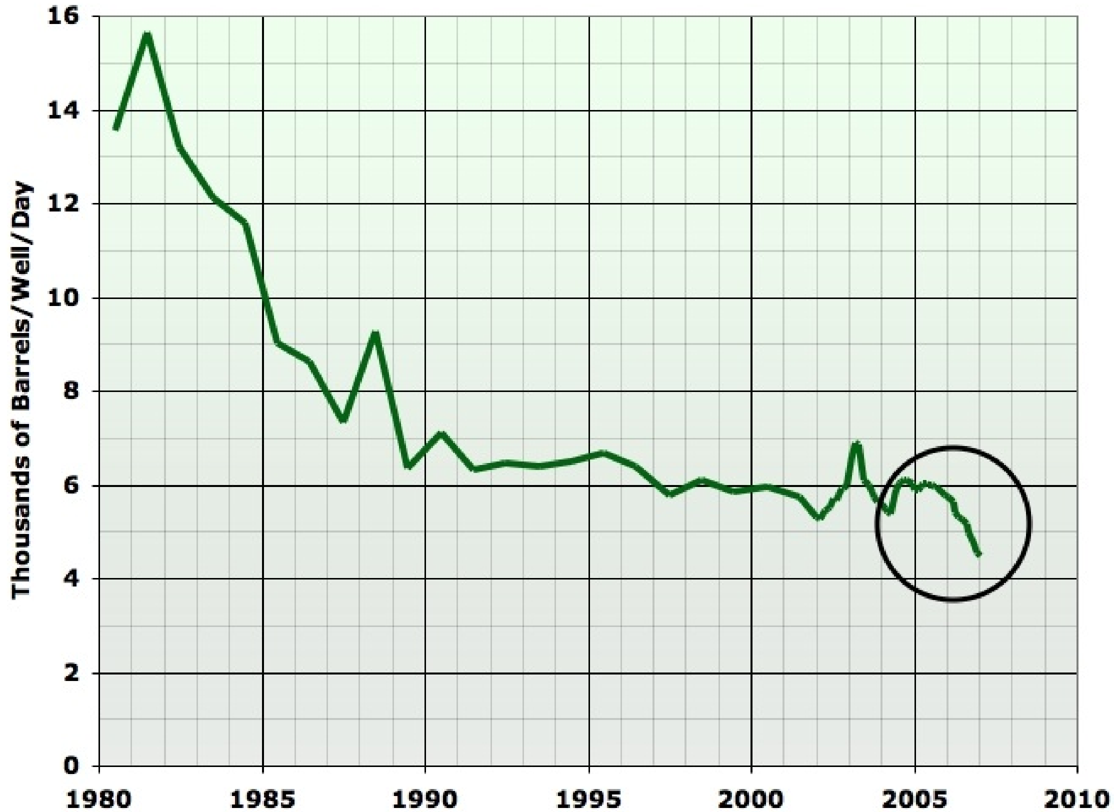

Estimate of well productivity. Annual through 2002 and monthly afterwards, according to the procedure described in the text.

As you can see, the well productivity shows evidence of an alarming decline starting in mid 2005, and accelerating through 2006. This is consistent with my overall hypothesis that declines in Saudi production are in significant measure forced, rather than voluntary.

The Future

Does this allow us any greater confidence in predicting the future? Probably not too much. On the one hand, it appears that Uthmaniyah may not imminently follow 'Ain Dar and Shedgum, and that the worst of the rapid damage in north Ghawar **may** already have occurred. On the other hand, the data appear to me to suggest that 'Ain Dar and Shedgum declines are not likely to be sufficient to fully explain the declines (net of megaprojects) that we have observed in total production. One possibility is that some other fields or regions of Ghawar are also declining (eg we know next to nothing about Hawiyah). Another possible explanation is that the redevelopment of Qatif and Abu Safah was largely a failure. A third possibility is a hybrid explanation in which some of the late 2006/early 2007 decline is voluntary as price has sagged below $60, over and above the forced decline. That might allow for a partial recovery of production this summer. However, this last would be somewhat inconsistent with the fact that Saudi Arabia's Asian customers have been put on allocation.

Time and further investigation are required.

Footnotes

1. The water cut was calculated as Krw/(Krw+Kro/Fv), where Krw is the relative

permeability for water from the figure, Kro is the relative permeability of oil also from the figure, and Fv is the oil formation volume factor from Greg Croft.