

1. Introduction

Motivations for Geologic CO₂ sequestration

It is commonly accepted that rapidly rising atmospheric CO₂ levels caused by anthropogenic emissions will present one of the greatest environmental challenges of the 21st century. Most of these emissions are derived from the burning of fossil fuels. (globally 22x10⁹ tCO₂/yr). 80% of Britain's electricity is currently generated from fossil fuels, with 19% from nuclear and <1% from renewable sources. With most of Britain's nuclear capability due to be decommissioned within 20 years, even the most optimistic estimates for what can be achieved through energy efficient practices and the development of renewable energy sources predict that we will be dependent on fossil fuels for the foreseeable future. One tool that is available to reduce CO₂ emissions is for it captured upon emission and stored in other domains of our planet, such as the oceans, the terrestrial biosphere and the subsurface.

Potential Reservoirs

Table 1 shows estimates from *Johnston and Santillo* (2003) for the volumes of CO₂ that can be stored in different domains:

Domain	Sequestration potential (Gt C)
Oceans	1000s
Subsurface	1000s
Deep saline formations	100s-1000s
Depleted oil and gas reservoirs	100s
Coal seams	10-100s
Terrestrial	10s

Although it has a large storage potential, storage in the deep oceans does not appear satisfactory (*Johnston and Santillo 2003*). It is not clear what the effect this would have on fragile benthic ecosystems. Also, CO₂ is not removed from the global carbon cycle, but moved to another part of it. During oceanic circulation over 10²-10³yr timescales, much of this carbon will be returned to the surface, potentially causing problems for future generations.

Subsurface storage in geologic formations is viewed as a much more attractive option for a number of reasons (*Benson 2000*).

- Injection of CO₂ has long been used in the oil industry to re-pressurise hydrocarbon reservoirs and increase output. As a result, the expertise and technology required for this to be used as an emission reduction tool is already in place (eg. *Oldenburg et al 2001*)
- Collateral economic benefits are available: CO₂ enhanced hydrocarbon recovery, and enhance coal-bed methane recovery.
- Suitable target reservoirs are widespread, particularly in the North Sea, where it is estimated that there is potential for storage of 800,000Mt CO₂ in saline aquifers. (*Torvanger et al 2004*)
- Natural CO₂ reservoirs provide natural analogues for long-term CO₂ storage.

Carbon sequestration with enhanced gas/oil recovery (CSEG/OR) provides an attractive economic incentive for CO₂ sequestration. Also, these depleted fields will have been well mapped during extraction. However, there is concern that the integrity of such formations may be compromised by drilling of large numbers of wells and by structural changes occurring as a result of removal of hydrocarbons.

Injection of CO₂ into deep, un-mineable coal seams will lead to the adsorption onto the coal, allowing mineral sequestration of CO₂ along with release of methane as a commercial benefit. The commercial benefits inherent in these options mean that they may be advantageous for small-scale sequestration. However, they do not have the capacity to be available for sequestration for large volumes of CO₂.

By far the largest reservoirs available for CO₂ are deep saline aquifers. European CO₂ emissions are currently at 3000Mt CO₂/yr. Saline aquifers in the North Sea have the potential to store 800 000Mt CO₂, the total European emission of CO₂ for hundreds of years, so long as cap-rock integrity can be guaranteed (*Torvanger et al 2004*). If geologic sequestration is to become a serious, large-scale option for emission mitigation, it is likely that saline aquifers will become our principle targets.

If geologic sequestration is to have a positive environmental impact, then the injected CO₂ must be stored in the subsurface for as long as it takes for anthropogenic output

rates to drop to acceptable levels and for the carbon cycle to have recovered and stabilised (*Holloway 2001*). This constraint requires that deep aquifer sequestration must be capable of storing CO₂ for timescales of the order of 10⁴, or even 10⁵ years. In order to meet this requirement, we must have a good understanding of the processes that control CO₂ migration and reaction upon injection into subsurface aquifers, enabling predictive injection models to be developed, and we must be able to accurately monitor CO₂ migration in the subsurface through remote sensing methods. Consideration of fundamental research questions related to multiphase transport in porous media will strengthen the scientific foundations for CO₂ sequestration, allowing for enhancement of site-specific predictive modelling (*Benson 2000*).

Project Aims

At the expected pressure and temperature conditions for deep aquifers, CO₂ will exist as a supercritical fluid, with a density less than that of surrounding saline formation waters. It will therefore form a plume rising due to buoyancy through the aquifer until it reaches an impermeable aquitard, where it becomes trapped, flowing outward under gravity. The extent to which the plume spreads is vital for considering the storage potential of any aquifer – if it can spread too far, then it is possible that it will spread beyond the aquitard, or find discontinuities in the aquitard, which could allow rapid migration back to the surface.

During injection, CO₂ will also dissolve in the host formation brines, creating carboxylic acid, a highly corrosive fluid. This has the potential to react corrosively with the minerals of the aquifer, potentially leading to changes in rock mass strength, and importantly, porosity and permeability. These changes could in turn affect the dynamics of CO₂ sequestration. The possible interactions between CO₂ and the aquifer, and the effects that these may have on CO₂ migration in the subsurface, are not well understood. The purpose of this project is to investigate experimentally and theoretically the effect that changes in porosity and permeability due to matrix/fluid interactions will have on a gravity driven flow through a reactive porous media, and to consider our findings with respect to CO₂ sequestration.

2. A review of CO₂ sequestration literature

In order for geologic sequestration to become a viable option for the long term disposal of CO₂, containment within the targeted aquifers must be guaranteed on a timescale of 10⁴ years or greater (*Holloway 2001*). This guarantee requires an understanding of the behaviour of the CO₂-brine-aquifer system after the injection of large quantities of supercritical CO₂. Importantly, we must ensure that the supercritical CO₂ is not capable of migrating large distances, either vertically or horizontally, before becoming trapped. If CO₂ is able to migrate laterally over long distances through the aquifer, the chances of it encountering discontinuities in the overlying aquitard are increased. These discontinuities may provide a route through which the buoyant, supercritical CO₂ may migrate rapidly towards the surface, rendering sequestration both hazardous and ineffective.


Mechanisms for CO₂ sequestration in deep saline aquifers

There are 3 principal mechanisms that are capable of trapping CO₂ at depth for geological periods of time (*Johnson et al 2001*):

1. Structural trapping: At a pressure exceeding 7MPa (corresponding to injection depths exceeding ~800m), CO₂ exists as a supercritical fluid with a density of approximately 700kg/m³ (*Holloway 2001*). As such, it will be significantly less dense than any fluids already within the aquifer. Hence, upon injection, the supercritical fluid will rise through the aquifer due to gravity. Structural trapping describes the situation where low permeability layers such as shales or evaporites (known as aquitards) prevent further upward migration of the CO₂. Structural trapping can occur both in 'confined' or 'unconfined' aquifers (*Bentham & Kirby 2005*). Confined aquifers rely on a structural or stratigraphic geometry capable of trapping CO₂ in a manner analogous to trapping by cap rocks in oil and gas reservoirs. Unconfined aquifers have no specific traps or closures. Instead, upon injection, the CO₂ will migrate laterally below the aquitard, filling any small undulations in the aquitard boundary before spilling into the next. This will result in the CO₂ becoming spread

over a wider area, which in turn requires extensive reservoir modelling to verify aquitard integrity, but does allow for the possibility of increased CO₂ dissolution and mineral trapping due to the increased CO₂/host fluid contact area. It is currently believed that, over the short term at least, the majority of injected CO₂ will be sequestered in this manner (*Johnson et al 2001, Bentham & Kirby 2005*).

2.Solubility trapping: Where the supercritical CO₂ comes into contact with formation brines, the CO₂ will react, dissolving to form an acidic, CO₂ enriched aqueous fluid:



CO₂ that becomes dissolved into the formation waters can only return to the surface if the formation waters do so. Given the depth of the formation, and the relatively high salinity, and hence density, of the fluids in question, this is unlikely to occur on geological timescales, meaning that the CO₂ is effectively sequestered.

3.Mineral sequestration: The injected CO₂ is capable of reacting with the aquifer itself, precipitating dawsonite and other carbonate minerals. This is the most effective method of sequestration in that the potential for migration is essentially eliminated, and, as carbonates are the densest of these 3 carbon-bearing phases, storage capacity is maximised. Unfortunately, reaction rates between the aquifer and supercritical CO₂, though relatively unknown, are believed to be slow, resulting in only small (<1%) amounts of mineral sequestration over the short term (*Johnson 2001*).

Modelling CO₂ injection

Numerical multiphase flow simulators have been available to the petroleum industry for some time. These can be adapted to consider the scenario of injection of supercritical CO₂ into deep aquifers. *Weir et al (1996)* use the TOUGH2 numerical simulator for this purpose. TOUGH2 is a general-purpose numerical model capable of simulating multiphase flow through porous media. It is commonly used in studies of hazardous and/or nuclear waste disposal sites. It is capable of modelling CO₂ migration within the aquifer, as well as dissolution of CO₂ into the formation brines, and has been applied to a number of scenarios to predict the evolution of CO₂/brine systems, and hence their storage capacity for CO₂. *Weir et al (1996)* use this model to

demonstrate how the lower density supercritical CO₂ might migrate upwards through the aquifer until it reaches a lower permeability aquitard, where it becomes structurally trapped. Over time, the CO₂ dissolves into formation waters, demonstrating the potential for structural and solubility trapping.

Saripelli et al (2001) use a semi-analytical model (PNLCARB) to demonstrate the growth of a CO₂ bubble around the point of injection, spreading radially outwards and migrating to the top of the formation, where it becomes trapped by a confining layer. Upon reaching this flat layer, the bubble spreads out laterally under buoyancy forces, continually spreading and thinning through time (fig 1). Assuming that this thin layer does not reach the edge of the formation, or a discontinuity in the aquitard, then it will be effectively structurally trapped, even without the presence of confining aquitard geometry. Furthermore, the thin layer has a large contact area with the formation brine, allowing for higher rates of CO₂ dissolution and hence solubility trapping.

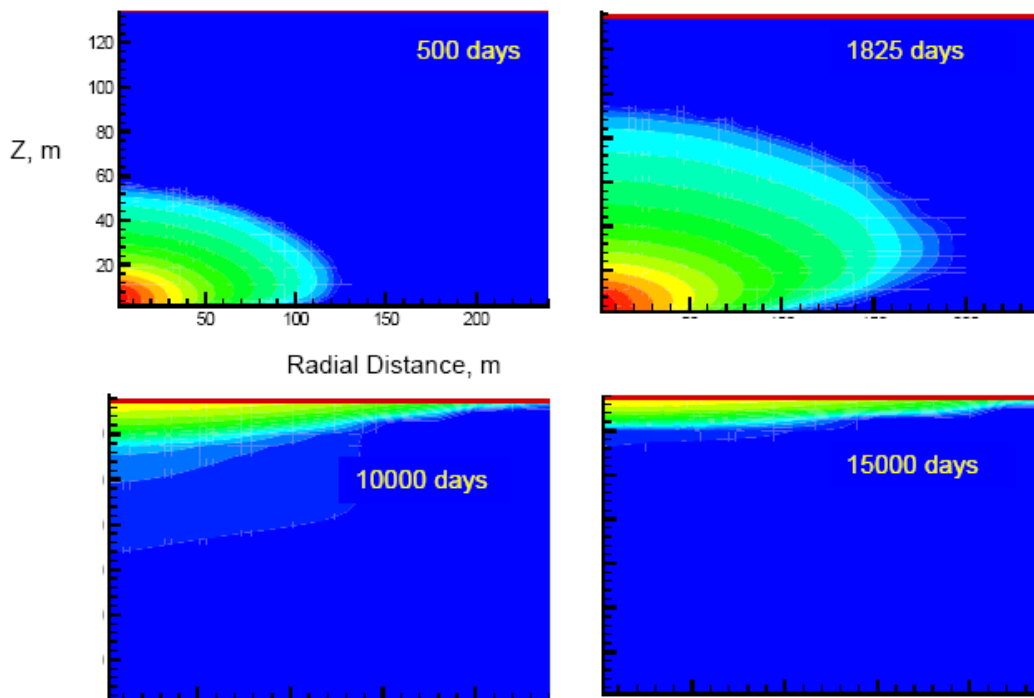


Fig 1. Numerical simulation from Saripelli et al demonstrating a CO₂ bubble growing radially outwards, rising to the top of the aquifer. Once trapped, it spreads, or slumps, under gravity, to form a thin layer.

These papers serve to highlight the importance of aquifer porosity and permeability to this process. A high porosity will give the reservoir a larger storage capacity. However, high porosity may well be accompanied by high permeability. This will increase the flow speed of the supercritical CO₂ through the aquifer, increasing the potential for it to escape before it can become sequestered. Evidently, site-specific evaluations of any potential sequestration site that consider these parameters are imperative.


Reactions between CO₂ and aquifer

The above models consider the aquifer itself as an inert media. This simplification is not justified given what we know about the reactivity of supercritical CO₂ and CO₂ enriched fluids. The acidification of the formation brines through the dissolution of CO₂ causes disequilibrium between aquifer and pore waters, resulting in both dissolution and precipitation reactions. As well as being capable of sequestering CO₂ as a mineral deposit, these reactions can also affect the porosity and permeability of the aquifer. As such, it is important that CO₂-aquifer reactions are considered when modelling CO₂ injection.

The reactions that we might expect to occur between CO₂ rich fluids and silicate aquifers are as follows:




Feldspar may react further to precipitate dawsonite:



These reactions are capable of sequestering CO₂ as a mineral. However, they are slow, at least an order of magnitude slower than carbonate reactions (*Grigg et al 2003*). This is why we do not expect to see significant mineral sequestration over short-term periods.

If carbonates are already present in the aquifer, it is likely that they will be dissolved:



This reaction does not sequester CO₂ directly, but it does buffer the pore waters, allowing solubility trapping to increase as more CO₂ becomes dissolved.

Reactive transport models combine multiphase flow simulators with kinetic and geochemical data for the CO₂-aquifer reactions. This allows us to predict both the migration of CO₂, and its interaction with the aquifer. Again, they are widely available owing to their usefulness when considering disposal of toxic and nuclear waste. Using the NUFT reactive transport simulator, *Johnson et al* (2001) demonstrate CO₂ injection into homogenous and stratified aquifers over a 20-year time-span. They predict that the majority of CO₂ remains structurally trapped, although 15% is dissolved into formation waters, and 1% is precipitated as minerals, principally dawsonite. They predict that this dawsonite precipitation decreases porosity by as much as 0.2% (fig 2).

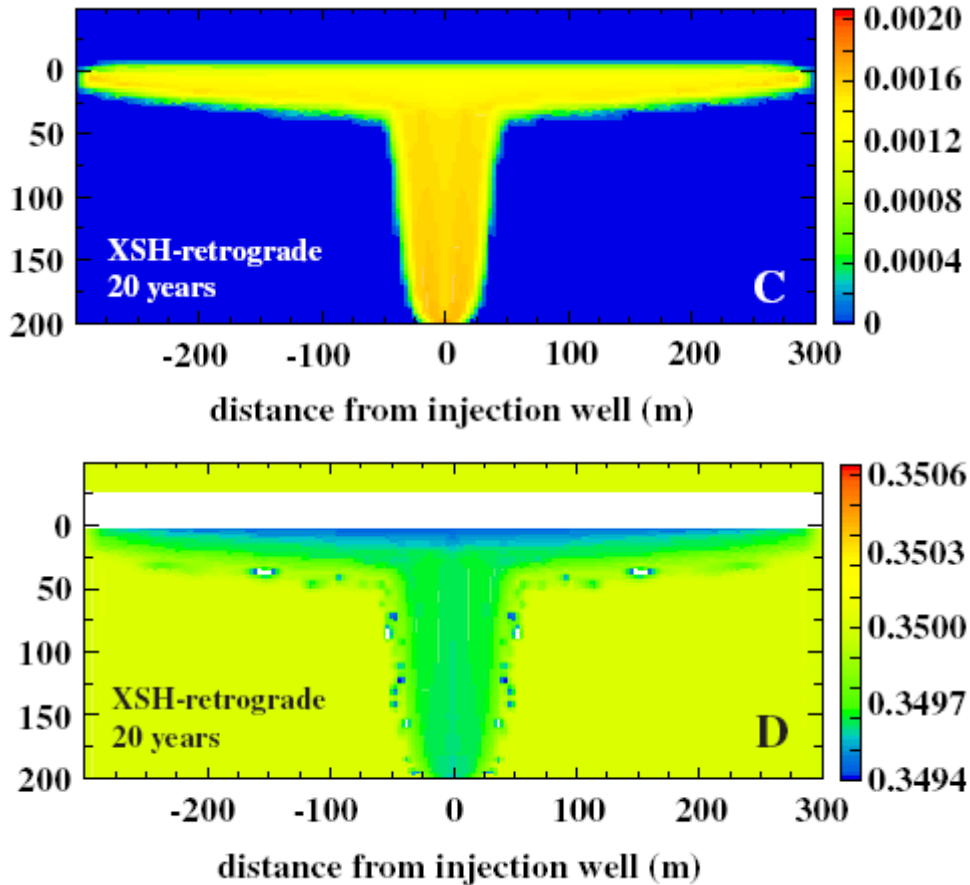


Fig 2. From Johnson et al (2001). Predictions from NUFT reactive transport simulations for dawsonite precipitation (C) and resulting porosity changes (D) during CO₂ injection into a Sleipner style aquifer

Lagneau et al (2005) use similar reactive transport codes to model CO₂ injection into both carbonate and non-carbonate aquifers. Again, effective mineral sequestration as dawsonite precipitate is predicted, preceded by an initial calcite dissolution event upon injection of CO₂ into the carbonate bearing aquifer.

The kinetic data required for these reactive transport models is difficult to obtain for aquifer pressure and temperature conditions. They are usually based upon either extrapolation or 'best guesses' (Bateman et al 2005), and as such are subject to error. Reactive through-flow experiments (Kaszuba et al 2003, Bateman et al 2005, Haszeldine et al 2005) demonstrate that the reactions occurring between the aquifer and the supercritical CO₂ may not be adequately described by the kinetic and geochemical data used by current reactive transport simulators. Kaszuba et al 2005 conducted experiments in which an arkose sandstone was allowed to react with CO₂

enriched brine for 80 days at temperatures and pressures close to that expected for sequestration scenarios. A suite of complex reactions between the silicates of the arkose and the fluid, including both dissolution and precipitation reactions, was observed. Clearly, the geochemical evolution of such systems is more complex than first imagined. *Bateman et al 2005* conducted a laboratory experiment pumping CO₂ enriched fluids through a porous medium consisting of quartz, feldspar and calcite. Again, calcite dissolution was observed to be the dominant reaction, whilst there was little or no evidence for silicate reactions and dawsonite precipitation as predicted by the various reactive transport models. Furthermore, observations of many analogous natural scenarios (*Haszeldine et al 2005*) indicate that long-term disequilibrium is possible between minerals and CO₂ rich fluids. This implies slower reaction rates than those predicted by previous reactive transport models, especially for silicate minerals.

Reactive transport models should be capable of producing reaction rates comparable to those observed in reactive through-flow experiments before they can be used with any degree of confidence. New models are currently under development that are capable of satisfying these criteria, as a result of being directly calibrated to laboratory observations (eg *Grigg et al 2003*). However, such models remain in their infancy. Clearly, the difficulties of predictive modelling - choosing the appropriate kinetic constants and pathways - means that nature and rates of reaction between supercritical CO₂ and the aquifers remain relatively unknown (*Lagneau et al 2005*). This is significant in that reactions may affect not only the degree of mineral sequestration, but also the physical properties of the aquifer, including both the porosity and the rock mass strength. Until the CO₂-aquifer reactions are better understood, we cannot hope to accurately predict the evolution of such systems, and hence the total capacity for CO₂ sequestration. A summary of estimates – from both numerical models and reactive through-flow experiments – for the expected CO₂/aquifer reactions, their timescales, and the resultant changes in porosity, are shown in table 2. Note the large variation between these estimates, caused principally by variations in aquifer mineralogy and errors in estimates of kinetic constants.

Effects of changes in porosity and permeability on a gravity-driven flow

A number of successful attempts have been made to model the effect of changes of permeability on reacting buoyancy currents (Phillips 1991, Raw & Woods 2003, Woods?). Fig 3 shows the predicted flow morphology when a negatively buoyant flow reacts to increase the permeability of the matrix. This predicts that dissolution reactions create a finger of high permeability at the base of the flow, behind a sharp reaction front. The flow then becomes focussed through this high permeability zone (fig 3). The first order controls on the evolution of the reaction front in time are the injection rate, the permeability on either side of the front, and the density difference between the injected and host fluids. The author is unaware of any attempt to model this effect experimentally.

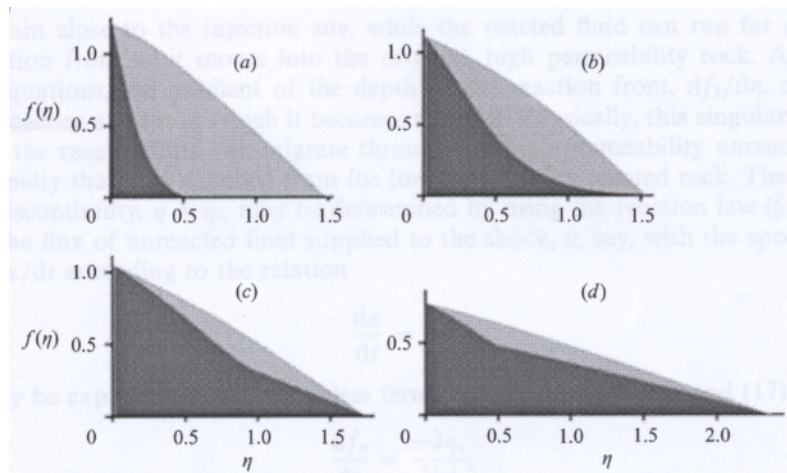


Fig 3. From Raw and Woods 2003. Calculations showing the shape of the reaction front and the distribution of the injected liquid in the case of a dissolution reaction. Calculations are shown for permeability ratios $k=2$, with (a) $\lambda=0.2$, (b) 0.5 , (c) 0.8 , and for (d) $k=5$ with $\lambda=0.8$. Note the formation of a higher permeability zone behind the reaction front through which the flow becomes focussed, affecting the morphology of the flow front.

Despite the difficulties in predicting CO₂-aquifer interactions, it is not unreasonable to expect that the dissolution of any calcite present will lead to an increase in porosity and permeability (fig 4)

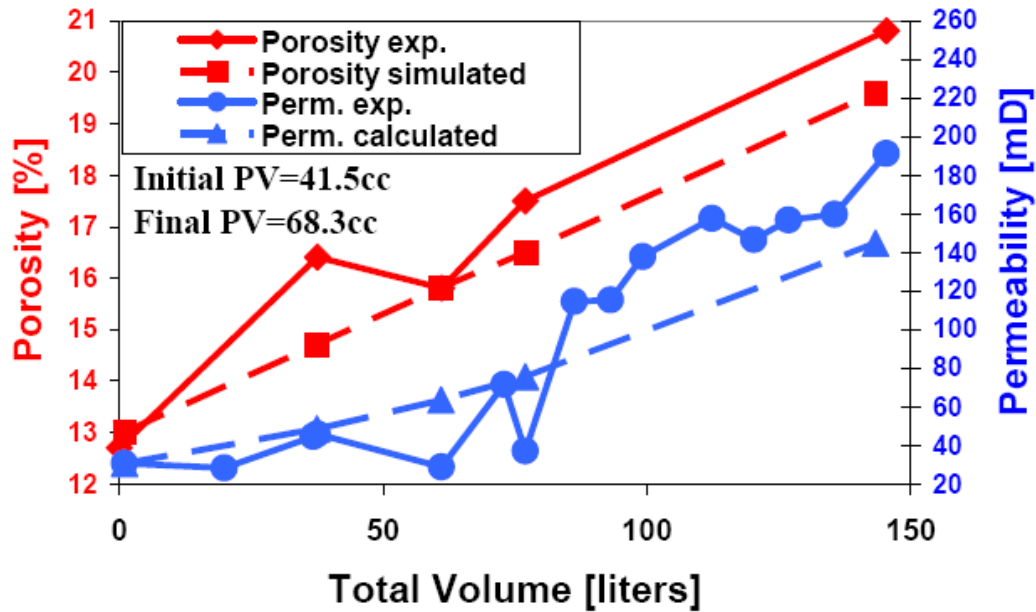


Fig 4. From Grigg et al (2003). Observed and predicted increases in porosity as a function of volume of supercritical CO₂ injected through a 50cm core containing dolomite, quartz and anhydrite

If this is indeed the case, then the dissolution models described by *Raw and Woods 2003* approximates the situation that may occur during CO₂ sequestration – where the dissolution of carbonates leads to the formation of higher permeability channels through which CO₂ becomes focussed. The aim of this project is to create laboratory models demonstrating reactive flow through porous media, and combine these with numerical models for analogue scenarios. These models can then be used to consider what effects the dissolution of carbonates from an aquifer may have on a buoyancy driven flow of supercritical CO₂.

3. Theoretical Model

(The following model has been developed in collaboration with the project supervisor, Prof AW Woods)

Modelling of the flow front:

One assumption used to model these experiments is that all the salt is dissolved by the fresh-water as it flows across the reaction front. This implies that the water becomes saline saturated below the front. Therefore its density will be the same as the saline solution already present in the aquifer. As a result, the flow front becomes passive, controlled only by flow across the reaction front. It cannot be modelled mathematically.

Modelling of the reaction front:

Consider the situation where a rapid dissolution reaction is occurring during injection of a buoyant fluid at a constant rate into a saturated aquifer, leading to an increase in permeability across a sharp reaction front:




Fig 5.

The speeds of fluid flow in front of and behind the reaction front, u_l and u_u respectively, will be controlled by the Darcy law equations:

$$u_l = -\frac{k_l}{\mu} \frac{\partial p_o}{\partial x}; u_u = -\frac{k_u}{\mu} \left[\frac{\partial p_o}{\partial x} + \Delta \rho g \frac{\partial b}{\partial x} \right] \quad (1)$$

Where, k_l , k_u , give the permeability in front of and behind the reaction front, p_0 is pressure, μ is the viscosity of the fluid, and $\Delta\rho$ is the change in density between the reacted and unreacted fluid.

Both fluids are incompressible, so we must assume conservation of volume:

$$Q = (H - b)u_l + bu_u \quad (2)$$

where Q is the rate of fluid injection, and H is the height of the tank.

The speed of the chemical reaction front will be controlled by the volume of fluid required to cause a reaction, and the volume of fluid passing through the aquifer:

$$\phi \frac{\partial b}{\partial t} = -\lambda \frac{\partial}{\partial x} [bu_u] \quad (3)$$

where ϕ is the porosity of the matrix, and λ is the volume of rock required to react with a unit volume of fluid

Combining equations (1)-(3) gives the governing equation, which describes the evolution in time of the reaction front.

$$\frac{\partial b}{\partial t} - \Gamma \frac{\partial}{\partial t} \left(\frac{b}{K + (1-K)b} \right)^{\frac{1}{2}} = \Gamma \Omega \frac{\partial}{\partial x} \left(\frac{b(1-b)}{K + (1-K)b} \frac{\partial b}{\partial x} \right)^{\frac{1}{2}} \quad (4)$$

where $\Gamma = \left(\frac{\lambda Q}{\phi H^2} \right)^{\frac{1}{2}}$, $\Omega = \left(\frac{\Delta\rho g H k_u}{Q \mu} \right)^{\frac{1}{2}}$, $K = \frac{k_l}{k_u}$

Table 2 shows the approximate values of the relevant variables for the experimental set-up. Fig 6 shows the full solution to this equation using the values in table 2. If this model is accurate, then we would expect to observe the reaction front evolving with a similar profile to that shown in fig 6.

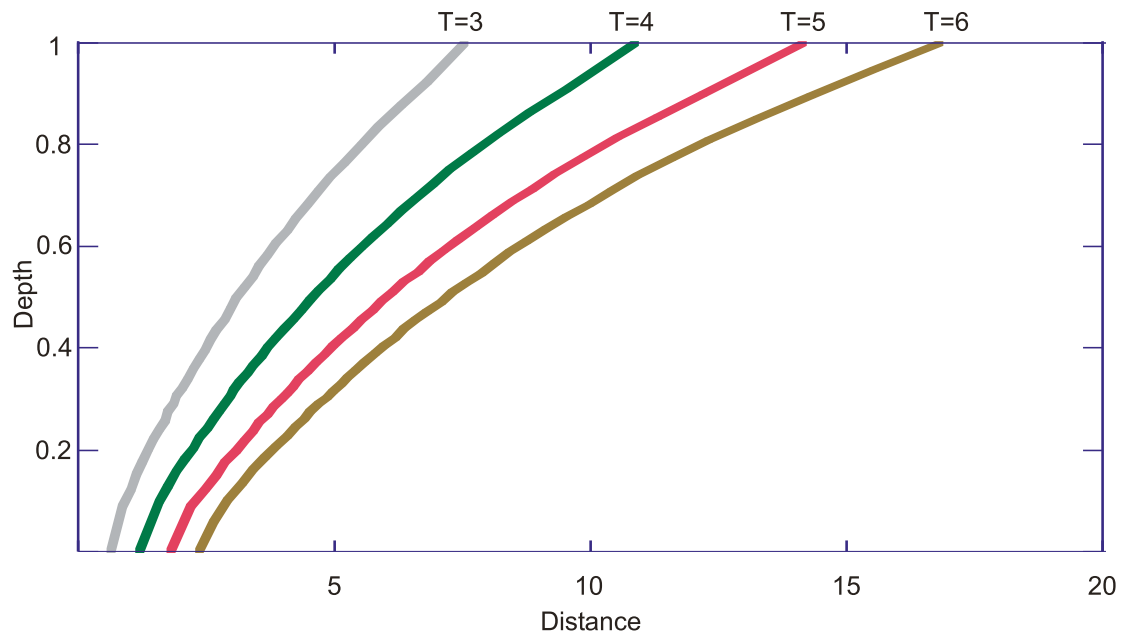


Fig 6. Full non-dimensionalised solution to eqⁿ 4, showing the predicted position of the reaction front through time

Parameter	Value
$\Delta\rho$	200kgm^{-3}
g	10ms^{-2}
H	0.1m
k_u	$4 \times 10^{-10}\text{D}$
ϕ	0.4
μ	10^{-3}Pas
Q	$1.5 \times 10^{-5}\text{m}^2\text{s}^{-1}$
K	2
λ	1.6

Table 4.

4. Experimental Work

Aims:

The aim of this experiment was to model the kind of reactive transport that we might expect during sequestration of CO₂. To achieve this, a number of criteria must be met:

1. In order for a gravity driven flow to form, the density of the injected fluid must be different to that of the host formation fluid.
2. The injected fluid must react with the materials making up the porous matrix, causing a dissolution reaction leading to a permeability increase. The host fluid must be inert with respect to both the injected fluid and the matrix.
3. The reaction must occur on a rapid enough timescale to affect the geometry of the gravity-driven flow.
4. This experiment aims to consider the situation where the aquifer is so large that the injection of CO₂ does not significantly affect the pressure within the reservoir. To model this ‘infinite aquifer’, fluids must be allowed to leave the tank without hindrance at the far end of the tank from the injection point.

Apparatus and materials

A very simple set of apparatus was used for these experiments. A Hele-Shaw cell with internal dimensions of 60x10x1cm was used as a 1-dimensional representation of the aquifer. This was later increased to 200x10x1cm to model flows with a higher aspect ratios. Initial experiments were affected by leakage of fluids below the lid of the tank. Attempts were made to seal the lid using window putty, silica gel, and then compressed rubber strips. However, all of these proved unsatisfactory. Eventually, an effective seal was created using strips of corking and increasing the number of screws used to hold the lid down. Holes of 7mm diameter were positioned at either end of the tank to allow fluids to enter and exit the tank. Ballotini being swept from the tank by the flow caused problems in initial experiments. To rectify this, perforated steel mesh filters were placed over both these holes to prevent any solids from leaving the tank. The dissolution of salt powder will decrease the total volume occupied by the matrix. This could lead to the development of a space between the lid of the tank and the

matrix. Left unblocked, fluids would flow along this gap, short-circuiting the experiment altogether. In order to prevent this, compressed sponge was placed along the top of the matrix, which could expand to fill any space created during dissolution. Fluids were pumped into the tank at a constant rate using a *Watson & Marlow* peristaltic pump, and allowed to leave the tank freely at the far end. A simple stopwatch was used to time the experiment.

The tank was filled with a mix of glass ballotini of diameter $0.45\text{mm} < d < 0.6\text{mm}$ and sieved salt powder ($d > 0.3\text{mm}$). The percentage by volume of salt powder was varied between 0 – 10% by volume. The porosities and permeabilities of these matrices were determined experimentally and are shown in table 3.

Salt content	Permeability (D)	Porosity
0%	3.80E-10	42.50%
2.5%	2.00E-10	42.10%
5%	1.99E-10	40.20%
10%	1.75E-10	38.60%
After 2.5% salt has reacted	3.74E-10	
After 5% salt has reacted	3.65E-10	
After 10% salt has reacted	3.54E-10	

Table 3. Experimentally determined porosities and permeabilities of the glass bead/salt powder mixes before and after reaction with fresh water.

The host fluid must be inert with respect to this glass bead and salt powder matrix. This was achieved by using fully saturated saline solution (25%wt NaCl), with a density of 1200kg/m^3 . The saturated solution will not be capable of dissolving the salt powder from the matrix. Fresh water, with a density of 1000kg/m^3 , was then injected into the tank. This was coloured with a small amount of food dye to aid visualisation. The fresh water would react with the salt powder, causing dissolution and an increase in porosity of the matrix.



Fig 7. Experimental Apparatus

Technique

The aim of this experiment is to model a homogenous aquifer. Hence, it is important that the salt powder and ballotini are well mixed to prevent any stratifications in salt content or in porosity. This is achieved by mixing the requisite amounts of salt and ballotini in small batches before putting them into the tank through a funnel. Once in the tank, the ballotini are mixed again using a stirrer to prevent any stratification developing. Once the tank is full, the top is lined with sponge, which is compressed so that it can expand when the matrix decreases in volume.

Once the tank is sealed, the matrix is saturated with the salt solution such that there are no air bubbles left in the tank. Dyed fresh water is then pumped into the tank at a constant rate controlled by the peristaltic pump, and allowed to flow through the matrix, creating a reaction front behind which the powdered salt has dissolved.

Visualisation of the reaction front

Where the salt is unreacted, there is a lower amount of fluid per unit volume within the bead pack, and a greater amount of the white salt powder. Hence, behind the reaction front there is a greater amount of dye per unit volume, as well as less salt. As a result, the reaction front can be observed due to the intensification of colour where the salt has dissolved. This intensification can sometimes be difficult to observe, but visualisation was improved by backlighting with a projector. The experiments were recorded with a digital video camera. The position of the reaction front at various points through time was highlighted on the outside of the tank with marker pen to further aid visualisation.

It is also desirable to observe the movement of fluid behind the reaction front, as this would allow us to detect any flow focussing. This was achieved by changing the dye colour within the injected fluid at certain points in time. This allows us to track the movement of the fluid injected at this point simply by observing the position of the colour change between the dyes.

5. Observations

Inert matrix model

When no salt powder is present, the injected fluid is inert with respect to the matrix. As a result, no reaction front forms, the flow simply advances as a linear front. The less dense injected fresh water rides over the denser saline fluid as shown in fig 8



Fig 8. With an inert matrix, the lighter fresh water rides over the saline solution, forming a linear flow front

Formation of a reaction front

When salt is present in the matrix, it begins to be dissolved by the injected fresh water. Because the kinetic rate constant for the dissolution of salt is short compared to the flow speeds in this experiment, the reaction zone is extremely localised, and for the purposes of this study, can be considered as a sharp transition or reaction front. Because a certain volume of injected fluid is required to react with the matrix, the reaction front always lags behind the flow front. In front of the flow front the matrix is unreacted and saturated with saline fluid. Between the reaction front and the flow front the fluid is the injected fresh water that has reacted, dissolving salt powder from the matrix and becoming saline. The area behind the reaction front has had the salt dissolved from it. As a result, the fluid behind the reaction front is unreacted fresh water. This is shown in fig 9




Fig 9. Morphology of the reaction front

We can verify that this is a reaction front, behind which all the salt powder has dissolved, and the injected fluid is unreacted, by measuring the salinity of fluids on either side of the flow and reaction fronts.

Fig 10. Salinity of the outlet fluid through time as the reaction front moves down across the outlet

Fig 10 shows the salinity of the fluid leaving the outlet through time. Before the flow front reaches the outlet, the output fluid is saline saturated – it is the initial host fluid. Even after the flow front reaches the outlet, the output fluid is still saline saturated. This fluid is dyed – it was initially injected as fresh water – but it has reacted with the salt powder in the matrix and become saline saturated. At the point where the reaction front reaches the outlet, there is a clear drop in salinity. This represents injected fresh water that has travelled entirely behind the reaction front before leaving the tank above the front. As a result, it remains unreacted, and therefore fresh. Unfortunately, the size of the outlet pipe is large with respect to the sharp reaction front. This means that it is impossible to observe the expected sharp change in salinity across the front. Instead we observe a hybrid signal as the reaction front moves down over the outlet, gradually decreasing the salinity of the output. However, it remains clear that there is a sharp drop in salinity behind the reaction front. Our modelling assumption that all the salt is dissolved at the reaction front, and the fluid behind it is unreacted, appears to hold.

Evolution of the flow front in time:

Initial experiments were carried out in a tank of 60cm length. However, for later experiments, this was increased to 200cm. This was to allow us to observe the flows for longer periods of time, as they evolve to form currents with higher aspect ratios. It is the observations from the longer tank that will form the basis of our experimental investigation and comparison with the numerical model. Fig 11 shows the evolution of the reaction front in time for the 3 principle experiments conducted, with the salt powder percentage of the matrix varied between 2.5, 5 and 10%.

Validation of modelling assumptions

Fig 12 shows how the position of the tip of the front, the position of the mid-point of the front, and the volume of fluid behind the front, evolve in time. Given that the injection rate is constant ($0.15\text{cm}^3/\text{s}$), conservation of volume requires that the position of the tip of the front (X), and the 2D volume behind the front (A) will be controlled by a power law relationship with T^1 . The observation that both X/T and A/T are constant for the three experiments confirms this assumption.

The theoretical model predicts that the overall gradient of the reaction front will be controlled by $T^{-0.5}$. Fig 13 confirms this prediction experimentally: $\text{gradient}/T^{-0.5}$ is constant. The observation of the predicted power law relationships between X , A , gradient and T is the first indication we have that the agreement between theoretical model and experimental observation is a good one. This agreement will be further discussed in later sections.

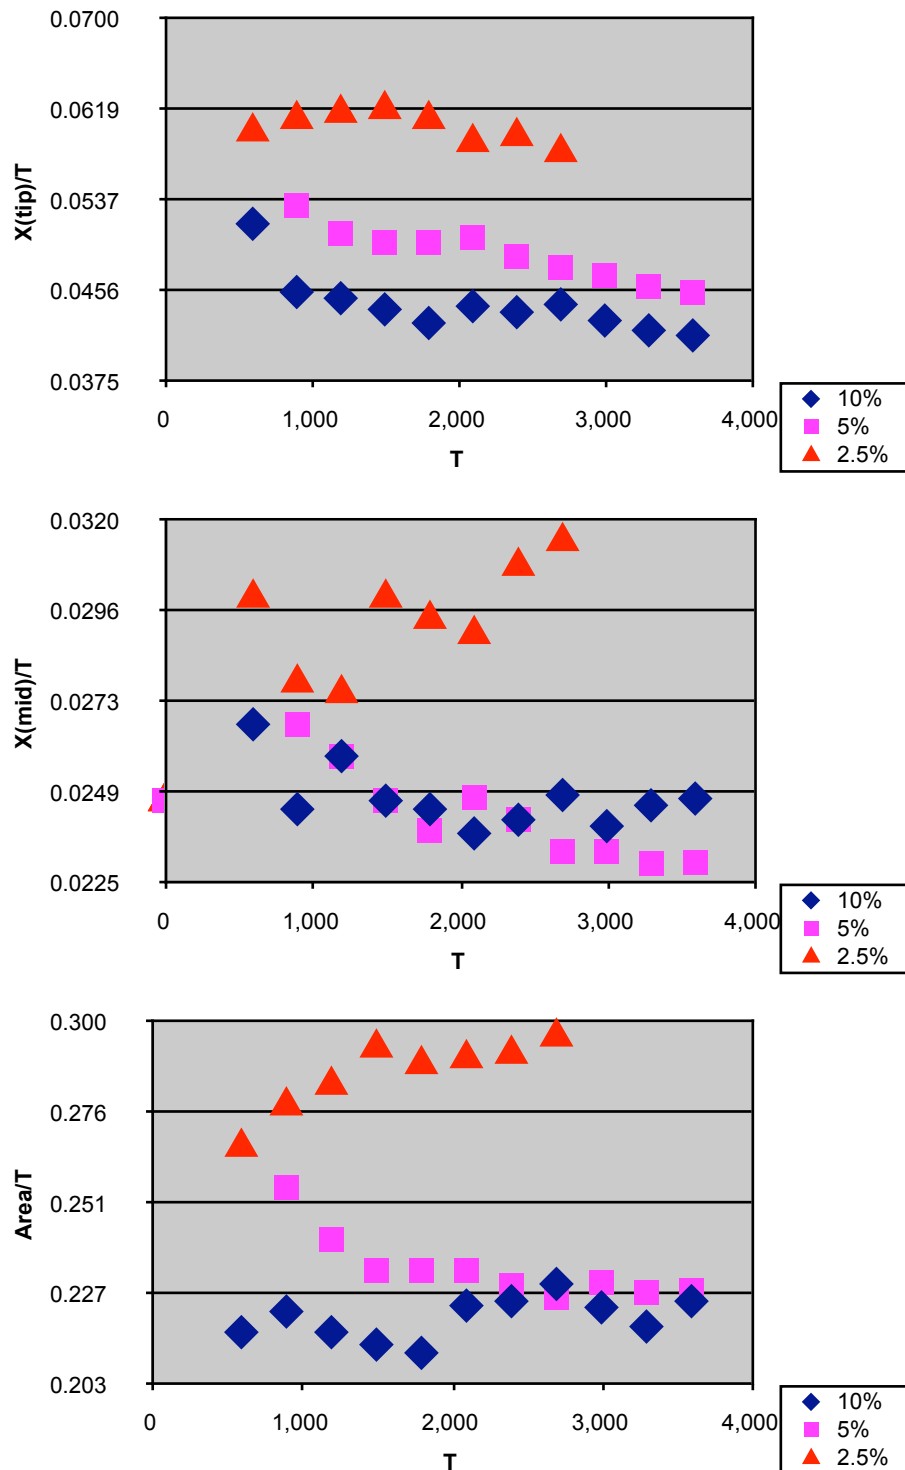


Fig 12. $X(\text{tip})$, $X(\text{mid})$ and A collapsed onto T^1 for the 3 experiments, demonstrating the linear dependence of X and A upon T for the 3 datasets




Fig 13. The gradient of the reaction front collapsed for the 3 experiments, demonstrating the dependence of Gradient on $T^{-0.5}$.

Focussing of flow above the reaction front

The zone behind the reaction front will have a higher permeability than the unreacted matrix. Thus, we would expect the flow to become focussed through this upper channel, avoiding the lower permeability zone below. We can determine whether this is the case by changing the dye colour of the injected fresh water at various points through time. The resulting observations are shown in fig 14



Fig 14 Flow focussing

Fig 14 shows how the fresh fluid injected into the tank (red) flows above the reaction front (marked in yellow), and does not at first penetrate into the unreacted, less permeable matrix (C). The flow is focussed above the reaction front. Gradually the injected fluid does move across the reaction front (E), though very slowly compared with the horizontal speed of the flow. At a later time, the colour of the injected fluid is changed again to black (E,F). Again, this fluid is trapped in the higher permeability upper layer. The majority of the injected red and black fluid has short-circuited the matrix along the upper, more permeable channel to exit the tank through the outlet, without interacting with the unreacted matrix in the lower channel. If a similar phenomenon were to occur in a CO₂ sequestration scenario, the formation of a high permeability channel might allow rapid and extensive lateral transport of CO₂ away from the site of injection.

Discussion

Comparison between experimental and theoretical models

Fig 15 shows a comparison of model and experimental data.

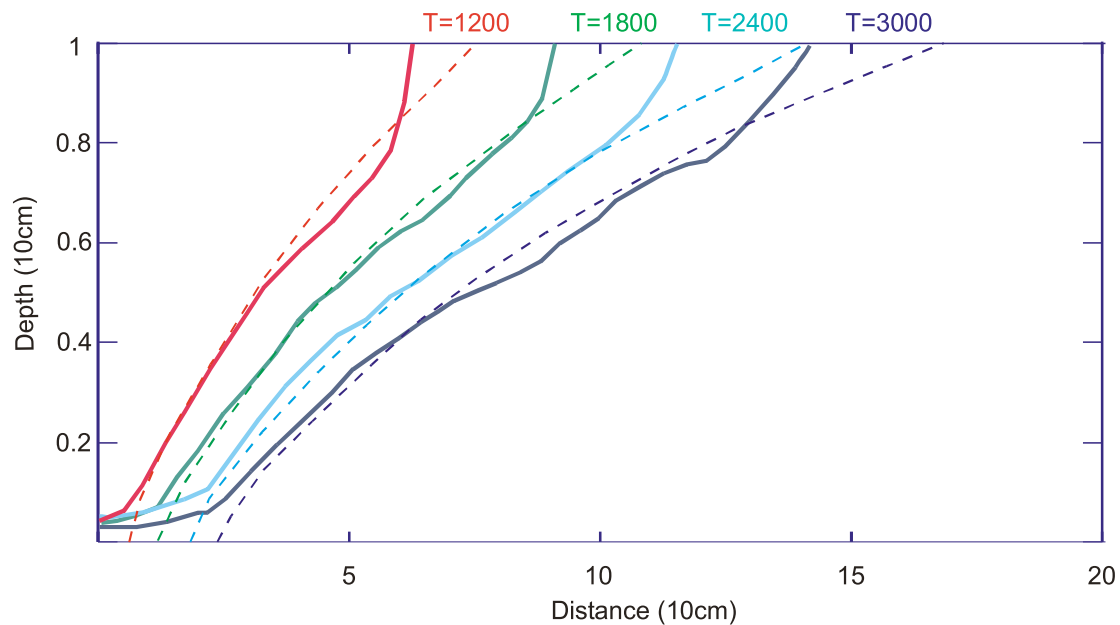


Fig 15. Dashed lines = model predictions ($\Omega = 5$). Solid lines = experimental observations (salt powder = 5%). Given the approximations used in creating the model, the agreement with experimental data is very good.

It is evident from by inspection of fig 15 that the model predictions and experimental observations are in good agreement. When combined with the observations made in the previous section on the power law relationships between A , X and T , and gradient $T^{-0.5}$, this agreement is a satisfactory confirmation of the ability of our theoretical model to predict the evolution in time of a dissolution front in a confined reactive porous media.

Applying this model to reservoir scale processes

In order to develop this laboratory model for the formation of a dissolution front, simplifications were made that mean that the model is not directly analogous to CO_2 sequestration scenarios. However, it is still possible to use these findings to consider the effect that changes in permeability may have on reservoir scale flows of supercritical CO_2

There are two principle factors that contribute to the morphology of the reaction front: the relative magnitude of the permeability change across the front ($K=k_l/k_u$), and the effect of gravity acting on the density difference between the fluids, represented by the non-dimensional parameter B .

Assigning specific values for the above variables is not trivial for reservoir situations. Some are poorly constrained; most will vary, possibly by orders of magnitude, depending on the specific injection conditions both within and between particular aquifers. Nonetheless, attempts can be made to approximate a more generic injection scenario. Table 4 gives typical values for the relevant variables under CO₂ injection conditions. Using these values we predict $B \approx 0.1-1$. Given that in our experiments, $B=10$, we might then expect that this flow regime is not that dissimilar.

Parameter	Value
$\Delta\rho$	200kgm ⁻³
g	10ms ⁻²
H	50m
k_u	10 ⁻¹³ D
ϕ	0.3
μ	5 x 10 ⁻⁴ Pas
Q	5 x 10 ⁻⁴ m ² s ⁻¹

Table 4.

Effects of changes to controlling variables, B and K

It is instructive to consider what this model predicts will happen both in the situation where buoyancy varies without change to K and in the limit that $B \ll 1$. This will allow us to consider the relative importance of the two dimensionless parameters to the evolution of the front. Fig 16 shows the variation in morphology of the reaction front caused by variations in B , where K is held constant. The increased gravity effect does cause a decrease in gradient of the reaction front, leading to the evolution of a narrower, more elongate reacted zone, as expected. However, these changes, though observable, are slight. It would seem that gravity does not play a significant part in this system.

By considering the situation where $B \ll 1$, we can observe the effect that changes in porosity alone have on the evolution of the front.




Fig 17. Generation of a dissolution finger in the limit of negligible gravity drive ($B \ll 1$) for a permeability ratio $K = 0.5$.

Fig 17 shows the evolution of the reaction front in the limit that $B \ll 1$, i.e. that, after a nudge from buoyancy to start dissolution at the top of the aquifer, the evolution of the front is controlled entirely by dissolution. This fig demonstrates that, even in the limit of a negligible gravity drive, the reaction front fingers out to the tip of the flow along the top of the matrix. This fingering is created entirely by focussing of the flow due to dissolution. It would seem that the evolution of the front is principally controlled by the magnitude of the permeability contrast across the front, K . The magnitude of this control, in the limit $B \ll 1$, is demonstrated in fig 18.

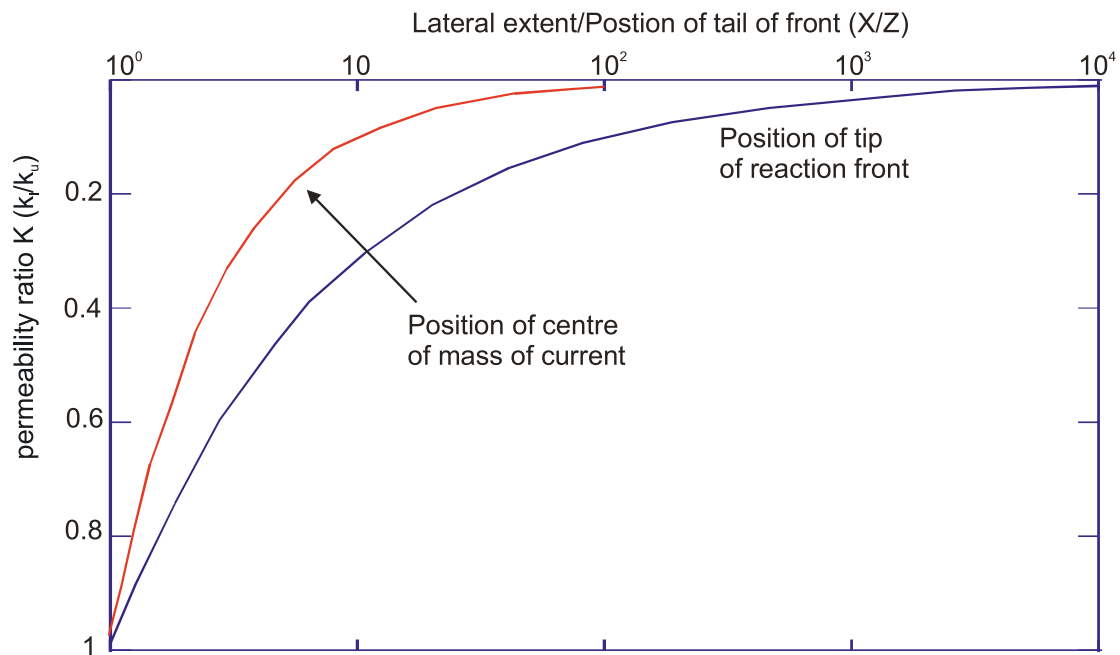


Fig 17. Gradient of the reaction front as a function of K in the limit $B \ll 1$.

Where there is no change in permeability across the front, there is no flow focussing, and, if there is no gravity influence, the front remains vertical ($X/Z = 1$). Where there is a large permeability change, flow focussing and fingering along the top of the aquifer creates a shallow gradient front ($X/Z \gg 1$). The magnitude that this effect can have is enormous. An order of magnitude change in permeability ($K=0.1$) can lead to the tip of the reaction front moving over 100 units ahead of where we would expect it to be if we did not include dissolution. An order of magnitude permeability change would not be unexpected if 2-3% calcite of the rock mass was dissolved, unblocking pore throats and increasing connectivity.

Discrepancies between model and CO₂ sequestration scenarios

The experimental and theoretical models presented here are by necessity somewhat simpler than might be envisaged for actual CO₂ sequestration scenarios. A number of assumptions have to be made in order for the model to be used to make predictions about reservoir scale processes.

Importantly, this model considers the situation where the viscosity of the injected fluid is similar to that of the host fluid. In CO₂ sequestration scenarios, this is unlikely to be the case: the viscosity of supercritical CO₂ is $\sim 10^{-4}$, an order of magnitude

smaller than the viscosity of host brines ($\sim 10^{-3}$). Such viscosity variations between phases flowing through a porous media can lead to the formation of instabilities in the flow front, and hence the formation of fingering patterns where the less viscous fluid tends to penetrate and finger through the more viscous one (Homsy 1987). Such a phenomenon might be expected to occur in CO₂ sequestration scenarios, but has not been considered here.

This model considers a situation where a dissolution reaction occurs on a rapid timescale compared to the flow speed, creating a zone of higher permeability behind a sharp reaction front. If this is to be applied to a CO₂ sequestration scenario, two important assumptions must be made: that potentially reactive minerals are present in the aquifer, and that these minerals will dissolve rapidly compared to the expected timescales of plume migration, which are on the order of 20-100yrs.

Most estimates for the speed of silicate reactions suggest that the rates of reaction are too slow to have significant effects on the timescales of plume migration. The most likely mechanism to cause such rapid dissolution is the reaction of carbonates with supercritical CO₂ and brine, forming bicarbonates and cations.



It is therefore unlikely that an aquifer that does not contain carbonate will be subject to any rapid dissolution reactions. This model cannot be applied to aquifers that do not contain carbonates.

However, even arkose sandstones, such as the Utsira (Sleipner Field) formation, may contain as much as 5% calcite as interstitial cement. Dissolution of 2-3% of the rock mass is enough to cause order of magnitude changes in permeability (Grigg et al 2003). Such an order of magnitude change is large to significantly affect the morphology of a buoyant plume of supercritical CO₂ (fig 18). Hence, the assumption that there will usually be enough potentially reactive minerals present that significant dissolution can occur to cause permeability variation does not appear unreasonable, simply because many aquifers, even those identified as 'arkose sandstones' may contain enough calcite for dissolution to be significant.

The reaction kinetics of CO₂-brine-aquifer systems are still very poorly understood. Estimates for the time taken to dissolve significant amounts of calcite vary by orders of magnitude, from a few days through to 10s of years. If the dissolution can occur over a few days or months, then this is a short timescale compared to the 20-100 years over which a plume may migrate, and the reaction front can be considered to be a sharp transition. If however, the reaction is slower, then dissolution will occur over a much broader reaction zone behind the flow front. In this situation, the model discussed here would not be appropriate.

Improvements in estimates of the reaction kinetics are unlikely to be made without study of full-scale CO₂ sequestration experiments, involving drilling out of the aquifer around the injection point for detailed study. Until these improvements are made, we cannot be sure if this model is appropriate for CO₂ sequestration scenarios.

Conclusions

Despite a number of simplifications, the model presented here successfully demonstrates the effect that dissolution reactions can have on the morphology of a buoyancy driven flow through a porous media. The morphology of the front is controlled almost entirely by the changes in permeability across it. A large change in permeability will create a shallow gradient front, the tip of which may extend laterally far further than we would expect if dissolution were not taken into account.

It is possible (though still uncertain) that this may be an analogue for processes occurring in CO₂ sequestration scenarios as carbonates within the aquifer are dissolved by reaction with injected CO₂. This is significant in that it may allow CO₂ to be transported laterally beyond the known limits of an aquitard far more rapidly than expected.

The author is unaware of any attempts to consider the effects that such a reaction could have on a migrating plume of buoyancy driven CO₂. However, the findings of this project demonstrate the very significant magnitude of the effect that such a permeability change will have. If we are to accurately model CO₂ injection,

permeability changes due to reaction should be considered when attempting to produce predictive coupled reactive transport models for CO₂ migration in deep aquifers.