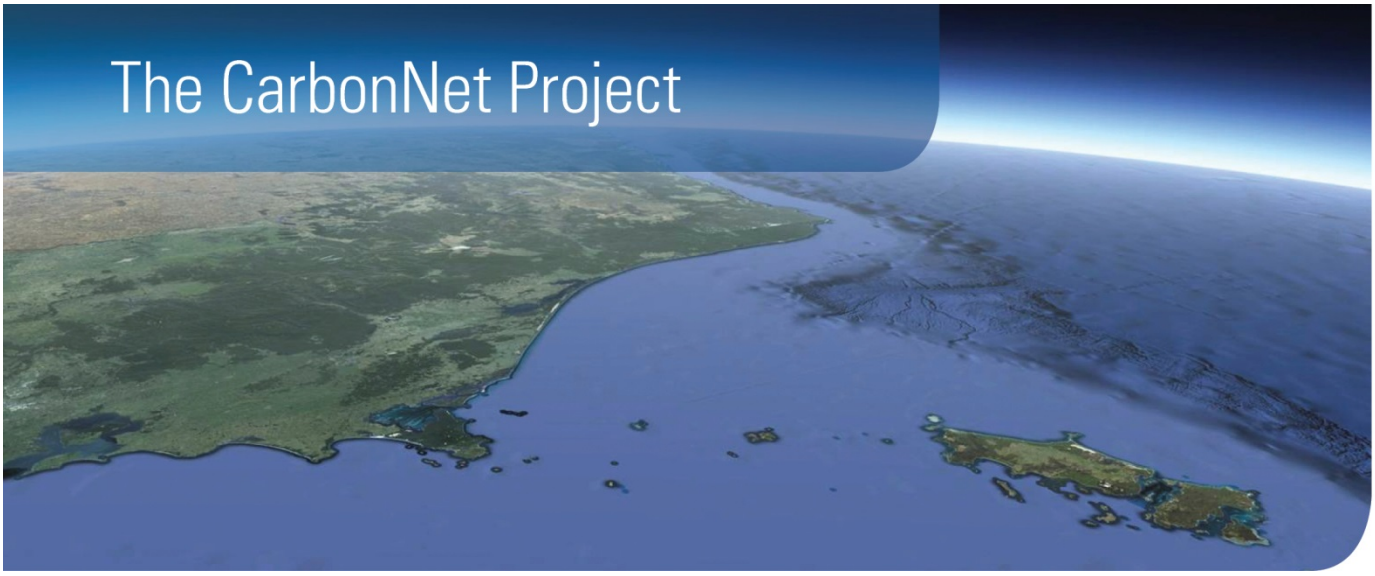

Probabilistic approach to CO_2 plume mapping for prospective storage sites: The CarbonNet experience.

GLOBAL CCS INSTITUTE

VICTORIA State Government

Australian Government

Disclaimer

This publication may be of assistance to you but the State of Victoria and its employees do not guarantee that the publication is without flaw of any kind or is wholly appropriate for your particular purposes and therefore disclaims all liability for any error, loss or other consequence which may arise from you relying on any information in this publication

The material and information contained in this publication is made available to further the Global CCS Institute's objective of accelerating the global adoption of safe, commercially and environmentally sustainable carbon capture and storage technologies in the public interest and is provided for convenience only. The Global CCS Institute, state, and any third parties who have contributed to the publication, do not give any representation or warranty as to the reliability, accuracy or completeness of the information, nor do they accept any responsibility arising in any way (including by negligence) for errors in, or omissions from, the information. No persons should act or fail to act on the basis of these materials

To the maximum extent permitted by law, the Global CCS Institute, state and any third parties who have contributed to this publication, disclaim all liability for any loss, damage, expense and costs incurred by any person arising out of using or relying on any material and information contained in this publication.

September 2016

Authorised by the Department of Economic Development, Jobs, Transport and Resources

Acknowledgement

The CarbonNet Project receives funding from the Australian and Victorian Governments as part of the Carbon Capture and Storage Flagships program.

The preparation of this paper is part of a funding agreement between the CarbonNet Project and the Global Carbon Capture and Storage Institute for Knowledge Share Activities.

Technical input and discussions with colleagues and external experts have significantly assisted with the preparation of this paper. In particular,

CarbonNet would like to thank Nick Hoffman and Mohammad Bagheri, and Laurent Alessio of LEAP Energy. Paul Tregoning and Anthony Purcell of ANU

provided important mathematical insight and advice

Executive Summary

In CO_2 storage, there is a requirement to predict the range of possible plume extents and travel paths and associate a probability with this range. This requirement is in the context that subsurface uncertainty is a given, and that no single plume prediction can be 100% precise. The probabilistic expectation of the plume at future times is used for project purposes and for regulatory assurance that the plume will remain within the defined storage boundaries (both geographical and stratigraphic) for the required period of time with an appropriate high level of confidence. In particular, Australian GHG storage regulations call for a prediction of all plume paths with more than 10% probability of occurrence (i.e. plume paths at P90 confidence level).

Here we outline a probabilistic approach based on reservoir modelling sensitivity and uncertainty analysis, adapted from the petroleum industry and suitable for high-mobility CO_2 plumes in thick and well-defined reservoirs. The method can also be extended to other basins and geological circumstances. In the petroleum industry, it is commonplace to evaluate resources in probabilistic terms with some objective parameter such as oil in place, recoverable reserves, or nett present value. This methodology can be adapted easily to objective measures such as vertical ascent of a plume relative to a caprock or lateral approach of the plume to a boundary or other geographic feature to be avoided (e.g. a mapped fault). What is novel in our approach is to analyse plume paths (extents) in a statistical manner to generate probabilistic maps and cross-sections of plume extents to inform on containment risks and areas with key monitoring requirements.

In our approach, the reservoir layering must first be analysed and the principal hydrodynamic flow units (HFU's) and the intervening seals identified. In the Gippsland Basin, multiple reservoir layers of 100-150m of multi-darcy, clean quartz-dominated sands form the main reservoir units and are proven by over 1,500 hydrocarbon exploration and development wells and are mappable on extensive 3D marine seismic data. The reservoirs are supported by an ideal, almost infinite aquifer which buffers pressure effectively and dissipates it regionally over short timescales (100 km in decades). In these reservoirs, CO_2 plumes are highly mobile and must be controlled by either structural trapping, or by careful mapping and use of non-structural (saline aquifer) storage.

The CarbonNet Project aims to store a nominal 125 million tonnes of CO_2 over 25 years in the same basin still in use for hydrocarbon extraction, and adjacent to an important onshore aquifer. Plume management and containment is therefore vital and high confidence must be placed on plume path modelling, including the analysis of rare statistical outliers.

The initial geologic static model contains a description of the HFU's and seals. A base case dynamic model with the most likely geological properties forms the basis of the sensitivity and uncertainty analysis. Subsurface uncertainty inevitably leads to potential variability of rock properties which will be expressed as different plume paths through time. A wide range of variability should be allowed in major properties such as vertical and horizontal permeability, porosity, residual saturation, relative permeability, etc.

A two-stage reservoir dynamic modelling process first examines each variable acting individually and identifies the key variable parameters compared to the base case - measured by key Objective Responses such as the vertical rise of the plume or lateral approach to a boundary. The second phase studies co-varying parameters using a formal Experimental Design, to identify the coupling of parameters and amplification of key responses, and define the weighting of a proxy model.

Proxy models represent the outcome (the Objective Response, measured from actual reservoir dynamic simulations) as a simple mathematical consequence of the weighted and scaled input variable parameters and the coupling between them. It is thus a rapid way to explore parameter space and predict outcomes of different combinations of inputs at different strengths, without running many computationally intensive full reservoir simulations. It is important to check that the predictions from the proxy agree sufficiently well with actual simulations in selected scenarios, especially when multiple parameter extremes are simultaneously invoked.

The proxy model is used to statistically explore parameter space and define probabilistic plume outcomes (Objective Responses such as approach to the lateral and vertical storage boundaries, to allow confidence to be placed in estimates of storage certainty.

A novel aspect of the current work is to define a 3D probability cloud for the plume extent, based on a Boolean analysis (a Boolean variable is a binary on/off – in this case plume present/absent). This allows maps and cross

sections of the 3D probability cloud to be displayed, illustrating the countours of probability in comparison to physical features such as vertical or lateral boundaries. An example of the approach is demonstrated for structural storage in the offshore Gippsland Basin, Australia. Maps of 90% probability of the plume path from the actual dynamic simulation runs are illustrated to meet Australian regulatory requirements.

MMV techniques and locations are significantly informed by the 3D probability forecasts. Expensive fixed assets can be concentrated in locations with high probability of detection, or where they offer protection against low probability but unwelcome outcomes. Volume based MMV is best deployed when the plume has a wide range of outcomes in 2D or 3D space.

In conclusion, the new approach offers a method to quantify plume path uncertainty in terms of lateral extent (i.e. map view) and vertical extent (cross-section), offering 3D understanding of plume containment with an appropriate high level of regulatory and public confidence.

Introduction

In CO_2 storage, there is a requirement to predict the range of possible future plume extents and travel paths and associate a probability with this range. This requirement is in the context that subsurface uncertainty is unavoidable, and that no single plume prediction from reservoir modelling can be 100% precise. However, multiple stakeholders need to know what is a reasonable expectation of future plume travel. The probabilistic expectation of the plume at future times is used for project purposes and for regulatory and public assurance that the plume will remain within the defined storage boundaries (both geographical and stratigraphic) for the required period of time with an appropriate high level of confidence. In particular, Australian GHG storage regulations call for a prediction of all plume paths with more than 10% probability of occurrence (i.e. plume paths at P90 confidence level).

There is a range of timescales that are interest to different parties. The timeframe of injection at a site is several years to decades, similar to the timeframe of oil and gas production operations, and this matches with a human and industrial timeframe for stakeholder interaction. Long-term storage requires centuries to millennia for dissolution of the plume to become significant and for mineralising reactions to commence significantly. Storage over geologic time is the ultimate goal, with dissolution, mineralisation, or physical entrapment of all the injected greenhouse gas.

Here we outline a probabilistic approach based on reservoir modelling sensitivity and uncertainty analysis, adapted from the petroleum industry (e.g. Ma and La Pointe, 2011), and suitable for high-mobility CO_2 plumes in thick and well-defined reservoirs. The method can also be extended to other basins and geological circumstances. In this work, Schlumberger Petrel and Eclipse software is used, but alternative static and dynamic modelling software can be used with no change in the method.

In resource extractive industries, it is commonplace to evaluate resources in probabilistic terms with some objective parameter such as oil in place, recoverable reserves, or nett present value. This methodology can be adapted easily to alternative physical objective measures such as vertical ascent of a plume relative to a caprock or lateral approach of the plume to a boundary or other geographic feature to be avoided (e.g. a mapped fault). What is novel in our approach is to analyse plume paths (extents) in a statistical manner in 3D space to generate probabilistic maps and cross-sections of plume extents to inform on containment risks and areas with key monitoring requirements.

In our approach, the reservoir layering must first be analysed and the principal hydrodynamic flow units (HFU's) and the intervening seals identified. In the Gippsland Basin, multiple reservoir layers of 100-150m of multi-darcy clean quartz-dominated sands form the main reservoir units and are proven by over 1,500 hydrocarbon exploration and development wells and are mappable on extensive 3D marine data (covering most of the offshore basin, beyond approximately 5 km from the coastline). The reservoirs are supported by an ideal, almost infinite aquifer (Kuttan et al., 1986, Varma and Michael 2012, Hofmann and Cartwright, 2013) which buffers pressure effectively and dissipates it regionally over short timescales (100 km in decades). In these reservoirs, CO_2 plumes are highly mobile and must be controlled by either structural trapping, or by careful mapping and use of non-structural (saline aquifer) storage.

The CarbonNet Project (CarbonNet 2015a, b, c) aims to store a nominal 125 million tonnes of CO_2 over 25 years in the Gippsland Basin (offshore, but within 20 Km of the coast – hereafter described as "nearshore"). Storage is planned in the same basin, some of the same reservoir and seal systems, and the same aquifer still in use for hydrocarbon extraction which is also adjacent to an important onshore aquifer. Plume management and containment is therefore vital and high confidence must be placed on plume path modelling, including the analysis of rare statistical outliers.

Probabilistic analysis in resource industries

In the petroleum and other resource industries, it is an established practice to evaluate hydrocarbon or mineral resources in probabilistic terms using a Monte-Carlo approach measuring some objective parameter such as oil in place, recoverable reserves, or nett present value. This methodology is well-established and is regularly used in sophisticated risk and reserves analysis and has an extensive literature (e.g. Newendorp 1976, Schuyler and Newendorp 2013). This modelling applies over the timescale of construction of the facilities and extraction of the resources – typically a few years to a few decades.

The probability distribution of the objective parameter is used to establish that an opportunity is significant enough to be worth the considerable expenditure of development, in the context of uncertain subsurface parameters, commodity prices, taxation, and operating and capital costs. Typically, key deciles of the probability distribution are derived and used to demonstrate that a positive commercial outcome is likely, at a degree of confidence such as P10, P50, or P90.

Sometimes in resource modelling, a range of geological and/or development parameters and scenarios (well spacing, drawdown etc.) are varied, and probabilistic outcomes are derived – again in terms of a 1-dimensional (1D) objective parameter. It is important to note that typically it is not appropriate in these circumstances to describe a particular scenario as ”the” P10, P50, or P90 case. The reason for this is that the particular combination of parameters, conditions, and geometries that results in the P90 outcome is unlikely to be representative of the set of variant cases that are locally close in probability space – for example a P90 outcome of total oil recoverable, based largely on low well productivity due to low permeability may sit adjacent to another case based on small structural volume or another case based on fault compartmentalisation. Therefore, one should be cautious in forming a mental image of “the P90 case”.

This objective-parameter probabilistic methodology can be adapted relatively easily to measures of plume movement such as vertical ascent of a plume relative to a caprock or lateral approach of the plume to a boundary, structural spill point, or other geographic feature to be avoided (e.g. a mapped fault). However, the timescale of modelling needs to be extended to a multi-century, or millennial timeframe, to show long-term containment of the CO_2.

If methods can be developed to view a range of spatial outcomes of plume movement, then P90 probabilistic boundaries for plume paths can be mapped, even if there is no physically associated “P90 case”. It is the purpose of this paper to derive practical methods to measure probabilistic outcomes in 2D (map) and 3D space. The method will be illustrated by worked examples, from the CarbonNet storage portfolio.

The spatial probability methods developed herein can also be re-applied to resource industries. In particular, reservoir models of the extraction of hydrocarbons from subsurface reservoirs often demonstrate bypassed or unswept oil or gas. Rather than targeting these bypassed volumes based on individual reservoir simulations, in which the volume and location of bypassed hydrocarbon is “certain”, but the model is known to be uncertain, statistical techniques such as those developed here can be used to map statistically significant volumes of bypassed hydrocarbons, and to statistically optimise the locations of infill wells.

Sensitivity and uncertainty analysis - methodology

The initial geologic static model contains a description of the HFU’s and seals as outlined above. A base case dynamic model with the most likely geological properties forms the basis of the sensitivity and uncertainty analysis. Subsurface uncertainty inevitably leads to potential variability of rock properties which will be expressed as different plume paths through time. A wide range of variability should be allowed in major properties such as Kv, Kh, porosity, residual saturation, relative permeability, etc.

As shown in Figure 1, a two-stage reservoir dynamic modelling process first examines each variable acting individually and identifies the key variable parameters compared to the base case - measured by key Objective Responses such as the vertical rise of the plume or lateral approach to a boundary. These objective measures are intrinsically one-dimensional, collapsing the full complexity of the situation to a simple observation parameter. There are strengths and weaknesses in such a simplification which will be addressed later in this report when 2D and 3D methods of measuring plume statistics are discussed. The second phase studies co-varying parameters using a formal Experimental Design, to identify the coupling of parameters and amplification of key responses, and define the weighting of a proxy model.

Sensitivity analysis in Reservoir Modelling is defined as a process of testing the robustness of a base case simulation to independent variation of many different parameters to determine which ones are most important and whether the base case is reliable. This analysis was undertaken on the base case simulation to understand how actual reservoirs might perform, given a wide range of subsurface variation, to allow for the fact that reservoir properties are never perfectly understood, except in the immediate vicinity of a wellbore. A sensitivity analysis tests the effect of potential variation in any uncertain parameters by varying one parameter at a time and observing the overall effect on the base case. The purpose is to determine the most influential parameters which can be displayed in a Tornado chart. Using this approach, the four key sensitivity parameters were identified.

CarbonNet adopted an approach of first identifying the principal HFU's, and then varying rock properties of the 3D static model within each HFU, and the intervening aquitards and seals. In this way, the overall flow architecture is preserved, but the rate of flow may vary, and hence the plume will take different pathways. The initial single-variable Sensitivity study was then extended to a multivariable Uncertainty study using experimental design (Alessio et al., 2005), where two, three, or even all four of the key parameters were varied simultaneously to analyse coupling between the effects of different parameters (do the effects of single variables add, multiply, or compound even more strongly when they operate at the same time as other variables, or do they offset each other and hence diminish when combined?).

Finally, a mathematical simplification known as Proxy Modelling was conducted to allow the rapid evaluation through Monte Carlo simulation of a very complex parameter space to define the range of responses with respect to two key 1D measures – approach towards key vertical or lateral boundaries of the GHG storage zone.

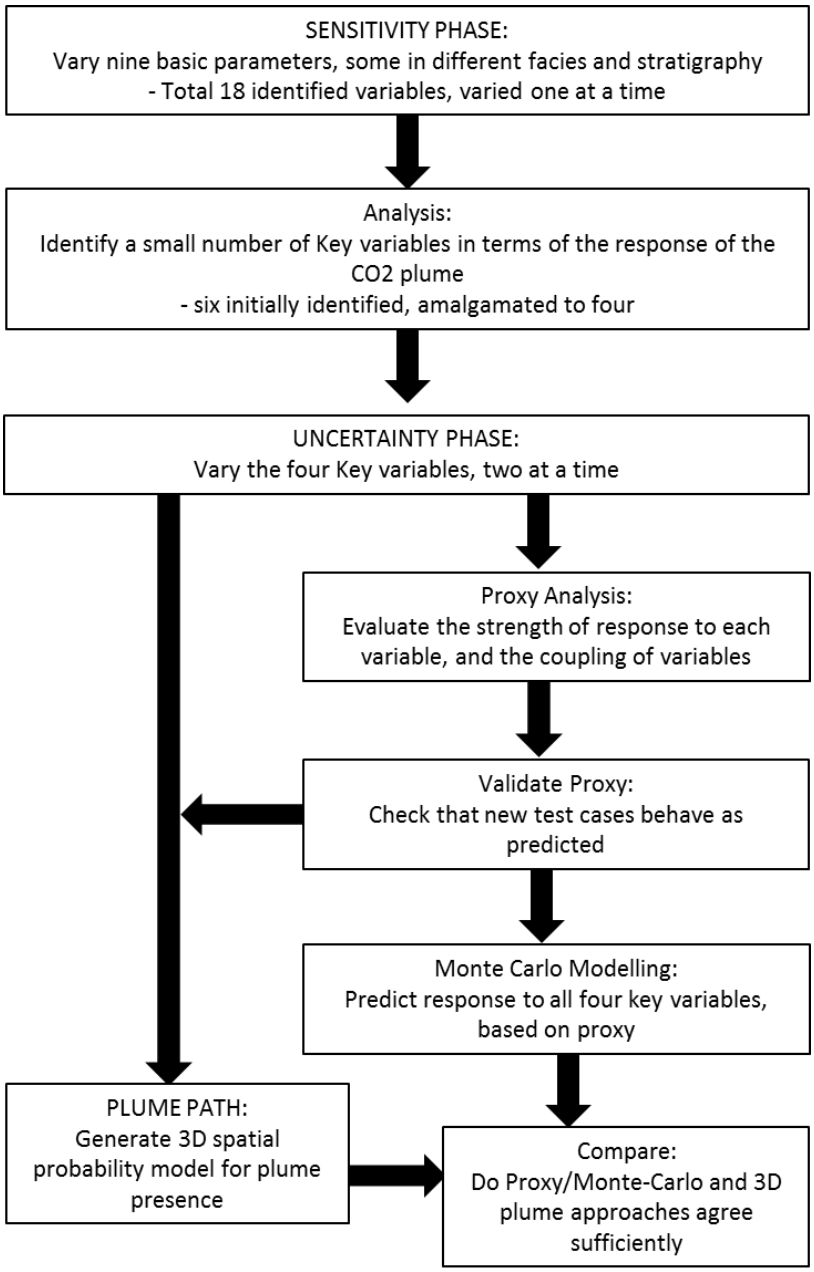

Figure 1: Concept of this study

As a separate stage of analysis, rather than simply analysing the 1-D responses, the suite of plume clouds output from Eclipse reservoir simulation as variable saturation was represented as either 2D maps or 3D spatial clouds of presence/absence of the plume. These binary maps and clouds were then added in 2D or 3D space to generate probability functions for presence/absence of the plume as a 2D map or 3D cube at key times.

The results of the 2D and 3D probability analysis were then compared with the results of the 1D objective parameter analysis (the Proxy model and Monte-Carlo simulation) to check whether there is sufficient agreement between the two alternative methods.

Applying this methodology and examining the variation of a range of input parameters allowed CarbonNet to statistically assess the resulting set of possible plume paths, and hence address the legislative requirement to understand and demonstrate all plume paths with a probability of occurrence greater than 10%.

Theory of sensitivity analysis

For any reservoir model, even though all available geoscience information is incorporated accurately and existing wells can be faithfully reproduced, there will be uncertainty about what the real-world reservoir parameters and layering would be between wells. Hence, a geological model is an approximation of the reservoir architecture that will be present around the injection site, based on petrophysical, geological and engineering judgement. Although seismic data offers a 3D picture of the subsurface, it has a limited spatial resolution and contains some artefacts of the data acquisition and processing methods. Most importantly, seismic energy does not respond directly and simply to rock porosity and permeability. Therefore there will always be uncertainty about the actual subsurface geology and this must be included in the resulting geologic model.

In order to manage this uncertainty, petroleum companies adopt a variety of approaches to simulating geologic variation between wells, tailored to the specific circumstances of their reservoir, and the geology of the basin in which it lies.

Given that the Gippsland Basin is known to contain laterally-consistent rock units, and the petrophysical properties of these units are relatively predictable (as demonstrated below in the analysis of flow units), and have geometries that can be mapped on good seismic data, and are not highly elongate, directional, or varying (for example, marine or terrestrial sequences with overall sheet-like sand reservoirs and shale seals, with aspect ratios ~ 5:1), a relatively straightforward approach has been adopted by CarbonNet.

CarbonNet Study Approach

Method: Starting from the best available geostatistical model, a sensitivity study was run to vary the global values of reservoir parameters such as porosity and permeability that ultimately determine fluid migration rates and flow paths.

Advantages: This method allows different flow paths and plume geometries to be readily simulated because of complex interactions of flow within directional reservoir units with variable porosity and permeability. So long as a wide enough range of parameter variation is allowed, the outcomes can be very different between runs. The improved speed of the runs allows more runs to be undertaken in practicable timeframes.

Disadvantages: This method does not specifically address changes in geometry of reservoir and seal units, therefore it can only be applied where these geometries are ***not*** highly directional (e.g. aspect ratios of >10:1), and the lateral facies variations are ***not*** sharp (e.g. shale-encased sand channels).

As noted, the data on Gippsland Basin reservoirs is extensive and high-quality (1500 exploration and development wells and abundant 3D seismic), and shows that reservoirs and seals have good lateral continuity, and that major channel sequences do not cut through the proposed storage sites. CarbonNet work on HFU analysis demonstrates that this approach is applicable to the Gippsland Basin, and is therefore sound and fit-for-purpose.

In order to understand fluid flow in the reservoirs, a sensitivity and uncertainty analysis as described above was undertaken to determine which reservoir parameters had the most impact on the storage of CO_2 in a structural trapping system, compared to the performance of the base case dynamic model. The overall plan of the study is illustrated in Figure 1.

Given the long execution times for reservoir simulation of CO_2 storage, it was important to keep parameter variations to a minimum, but most importantly, determine what the most influential parameters are with respect to CO_2 plume migration within the storage site. To assess these parameters, the Base Case was varied by altering a wide range of whole-model reservoir parameters including porosity, permeability (horizontal and vertical), fault transmissibility, residual CO_2 saturation and aquifer support.

The objectives of the study included:

- To understand the range of performance of the reservoir, to meet project objectives;
- To inform a regulatory application for a Declaration of Identified Storage Formation;
- To verify and optimise existing models;
- To identify independent reservoir parameters that are the key sensitivities for dynamic modelling (sensitivity analysis);
- To generate a proxy model for key observables such as approach to the western boundary and vertical migration and so model combined key parameters (uncertainty analysis/proxy modelling);
- Probabilistic analysis of vertical and lateral migration; and
- Provide P10, P50 and P90 plume migration pathways.

Interpretation of hydraulic flow units

In reservoir engineering, a hydrodynamic flow unit or HFU (alternatively known as a hydraulic flow unit) is a working description of how the geological architecture of reservoirs and seals stacks in a three-dimensional pattern to produce composite bodies which are connected together as flow units (also described as reservoirs or aquifers), and separated by aquitards or seals.

In an ideal geological situation, the flow units and intervening aquitards (seals) are sub-horizontal stratigraphic layers. The change from aquifer (reservoir) to aquitard (seal) is caused by wide-ranging changes in basin sedimentation, such as a change to shale deposition to create a semi-regional seal, or an influx of sand over time and its lateral transport over space to create a reservoir layer.

Some systems are more complex and the HFU's may cross stratigraphic boundaries, or exist as sub-sets of geobodies embedded within individual stratigraphic intervals. Because of these complexities, these systems are more difficult to model. An important step of reservoir simulation is to demonstrate that, in the basin context of the current model, the depositional systems and architectural stacking of geobodies leads to a rational and relatively simple set of HFU's.

HFU's should ideally be defined by post-hoc analysis (hindcasting) of flow performance within the local geology. This analysis should be based on history-matching of observed pressure depletion and flow connectivity during hydrocarbon production, aquifer abstraction, and fluid injection (if this has occurred). A priori prediction of flow units is more difficult, requiring the construction of multiple alternative geologic models to demonstrate that alternate stacking patterns of reservoir architecture do not lead to different flow unit outcomes in different scenarios.

However, if reservoir architecture is well-known and simple, the nett:gross ratios are high, and large amounts of well and seismic data exist (including production data from offset fields), then the potential variability in geologic models due to subsurface uncertainty is substantially reduced. In these circumstances, as is the case for the Gippsland Basin, the HFU architecture can be derived from the input well and seismic data and from the derived static model and dynamic reservoir simulation results.

Detailed performance within each HFU will still vary according to subsurface uncertainty, but the gross structure of HFU's will not change in these circumstances. CarbonNet's identification of relatively simple HFU's in the nearshore Gippsland Basin is an important component of describing the geology, supporting the dynamic modelling process and explaining why an eligible GHG storage formation exists.

As shown by multiple well data sets, there is a relatively high nett gross ratio with 60-80% reservoir-quality sand compared to non-nett units (shales and coals) in the Gippsland Basin. This high proportion of sand to shale indicates that thick amalgamated reservoir units will be the norm, and that non-layered HFU's will be rare.

The opportunities for shaly or sandy channels embedded in the opposite facies and giving rise to highly oriented permeability channels or permeability barriers is substantially reduced by the high nett:gross. While channels can be demonstrated from seismic observations within a few distinctive layers within the estuarine facies of the Lower Cobia Subgroup, these are not observed within the Halibut Subgroup – the zone of planned CO_2 injection, migration, and storage for the CarbonNet project.

Published data on the Latrobe Group reservoirs and the regional, local, and vertical pressure depletion from 50 years of production (e.g. Varma and Michael 2012), demonstrate that lateral pressure communication is exceptionally good, and vertical communication is good, with a few identified sub-regional aquitards (intraformational seals). These few functioning aquitards are very important for defining intraformational traps for oil and gas with storage over millions of years and CO_2 trapping over a millennial timescale (Hoffman and Preston, 2014).

Figure 2 is a cross section of the upscaled dynamic model used for CarbonNet reservoir simulation within a structural trap modelled in the nearshore Gippsland Basin. A complex pattern of good, intermediate, and poor reservoirs is depicted in rainbow colours, and low permeability units (shale and coal) in black. The corresponding CarbonNet HFU interpretation is summarised in Figure 3 (based on the results of numerous flow simulations from this study), where amalgamated reservoir units that function as a single HFU are coloured yellow, aquitards which are limited to ~50 years retention are coloured grey, and long-term effective seals with retention times >1000 years are coloured green. Note that on approach to the western boundary of the model, there is a dip reversal – which is why this boundary is important. If CO_2 were to reach it, then it would no longer be structurally trapped.

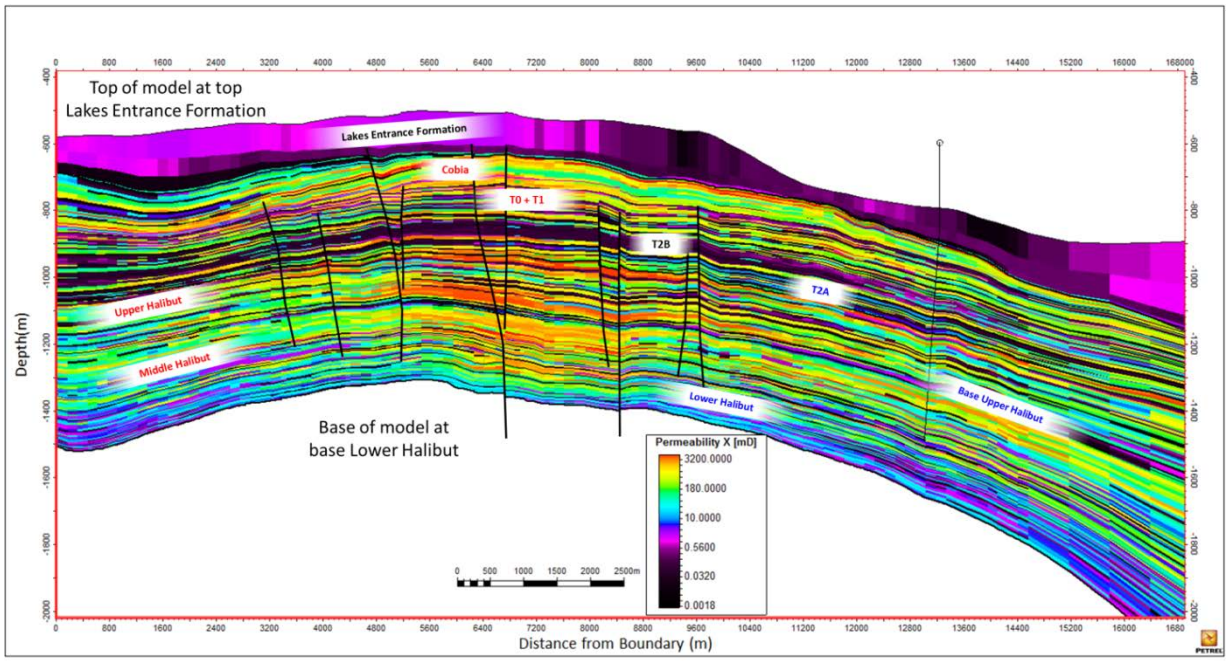

Figure 2: Reservoir permeability (kh) on key east-west cross-section of model.

A complex series of reservoir and seal units is contained in the static model. Dynamic simulation performance depends on the stacking of these reservoir units. Two major seal units with good continuity are the overall Petroleum Seal – the Lakes Entrance Formation, and the T2B intraformational seal.

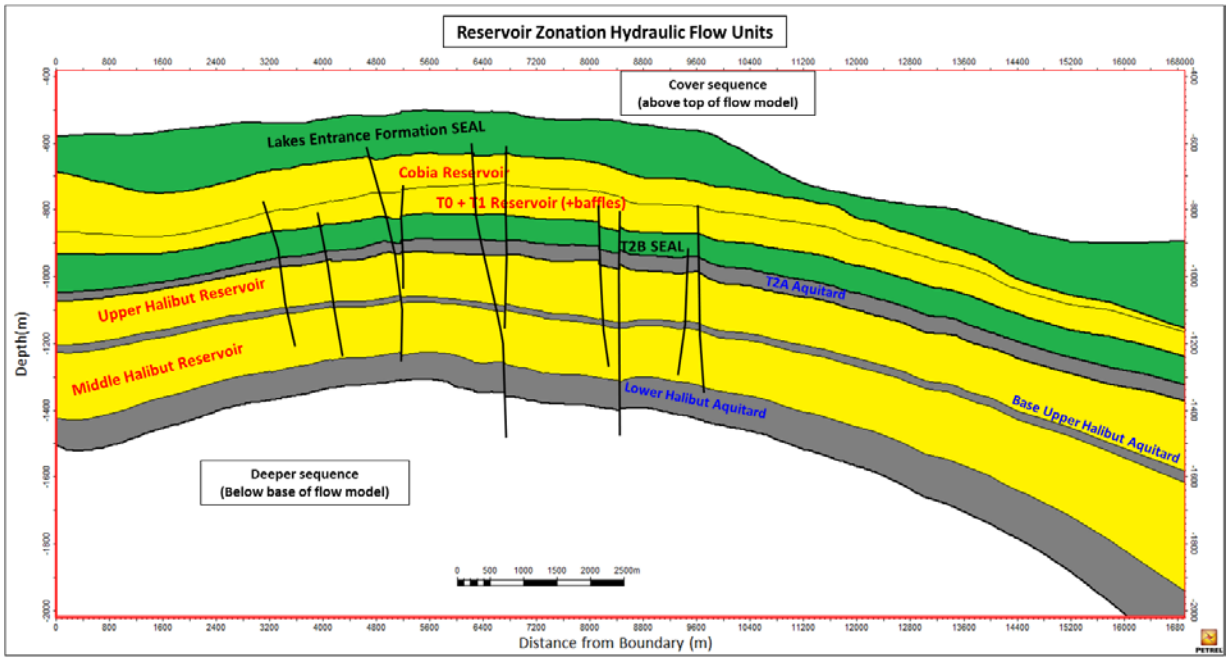

Figure 3: HFU interpretation of model stratigraphy

The complexity of Figure 2 can be simplified to a series of four distinct aquifers, two strong and effective seals, and two aquitards which act to delay, but not stop plume ascent.

Table 1 summarises the HFU's identified from CarbonNet dynamic modelling studies.

Table 1: HFU's identified from flow analysis of dynamic reservoir simulations

HFU	Type	Average Parameters over whole model		Comments
Lakes Entrance Formation[+]	Seal	Thickness kv md	150 m 0.008	Regional petroleum seal
Upper Cobia Subgroup[+]	Reservoir	Thickness kh	90 m 900 md	Excellent proven petroleum reservoir
Traralgon T0 + T1 [+]	Reservoir	Thickness kh	80 m 450 md	Reservoir with mapped coal baffles *
Traralgon T2B	Seal	Thickness kv	70 m 0.7 md	Continuous semi-regional pressure and salinity seal 99.9% effective in this study
Traralgon T2A	Aquitard	Thickness kh	60 m 500 md	Varies in effectiveness with scenario ~20% effective in this study
Upper Halibut Subgroup	Reservoir	Thickness kh	140 m 850 md	Effective reservoir
Base upper Halibut Subgroup	Aquitard	Thickness kv	25 m 0.6 md	Effective barrier for 50-200 years (scenario-dependent) 20-50% effective in this study
Middle Halibut Subgroup	Reservoir	Thickness	160 m	Effective reservoir

		kh 850 md	
Lower Halibut Subgroup	Aquitard	Thickness 100 m kh 20 to 200 md	May have storage capacity on-structure. Not a reservoir in downdip area near injection point

[+] CO_2 storage is not planned at this level.

* a baffle is a thin seal which cannot trap a significant column of CO_2 but can act to slow the ascent of a CO_2 plume by imposing a tortuous pathway.

Reservoir injectivity for each HFU

For each of the HFU's, the injectivity can be estimated based on nearby well core and log data and seismic mapping of key horizons. The simplest measure of injectivity is permeability times thickness. Studies by CarbonNet, published as Hoffman et al., 2015a, show that, for a world-class "Type One" storage site capable of safely injecting over 100 million tonnes of CO_2, reservoirs require to exceed a threshold of ~100 darcy-metres. Either of the upper or middle Halibut Subgroup zones achieves this level of injectivity independently and is therefore a world class reservoir in its own right. The planned injection zone of middle Halibut Subgroup reservoir is the preferred injection zone since it has better injectivity and also leads to better structural and stratigraphic utilisation – storage efficiency.

Table 2 summarises injectivity at a potential injector well location (values differ slightly from Table 1 which is an average over the whole model).

Table 2: Base Case injectivity of reservoir HFU's at a potential injector well location

Reservoir HFU	Thickness (m)	Kh (md)	Injectivity (darcy-metres)	Comments
Upper Cobia	62	745	46	Too shallow for supercritical CO_2 storage
T0 + T1	86	276	24	
Upper Halibut	140	689	97	Planned storage zone. Also available for future add-on injection
Middle Halibut	185	804	149	Main injection zone. Planned storage zone

If combined in future operations, the two Halibut Subgroup reservoirs/HFU's would offer a total of 250 darcy-metres. The assessed injectivity of these Upper and Middle Halibut Subgroup reservoirs/HFU's demonstrate that the CarbonNet site exceeds the threshold for commercial-scale injectivity.

Seal competency between HFU's

Seal competency consists of two components

- the complete blocking of a lithological unit to entry of a fluid (non-wetting and/or non-permeation)
- the rate at which a fluid could pass through it, if it were wetted / permeated (seal permeability or hydraulic resistance).

Fluid blocking

The first parameter is best assessed by MICP or other capillary entry pressure tests on real rock samples. This element has been assessed in Hoffman et al., 2012, where CarbonNet analyses demonstrate high capillary entry pressures for intraformational shales, including the Traralgon T2B. This high MICP value demonstrates the capacity of the storage formation to store a significant CO_2 column for long timescales (potentially millions of years).

Comparisons with oil and gas fields (many with an a significant CO_2 content) demonstrate that intraformational seals such as T2B are highly-qualified to retain oil and gas on million-year timescales and indeed are proven to do so in the immediate vicinity of the CarbonNet potential storage sites (Hoffman and Preston, 2014, CarbonNet 2015d).

Conventional reservoir modelling simulators such as Eclipse do not incorporate MICP as a parameter in fluid movement, since MICP is only important over very long timescales (centuries to millennia), which are beyond the time horizon of oil and gas production and relevant only to geological time scale storage projects such as CO_2 geosequestration or sour gas disposal. The 3D spatial distribution of oil and gas is assigned in the model when it is initialised and seal units are typically given zero hydrocarbon saturation. On production timescales it is not important whether a seal unit has a high MICP value and is therefore impenetrable, since the intrinsic low permeability of the seal unit precludes any significant fluid entry or permeation over the timescale of the simulation. A conservative view of storage containment can be derived by neglecting MICP, and considering only short- to medium-term storage through intrinsic seal permeability, which covers the timescale of 10^1 to 10^5 years.

Seal permeability (hydraulic resistance)

The second factor, of passive permeation of a seal, is fully modelled by conventional simulators such as Eclipse, and is best measured by ***hydraulic resistance*** **(R)**. Similar to electrical resistance, hydraulic resistance to vertical flow is defined as the thickness or height (H) of the material divided by the vertical permeability (k_v).

$R=H/k_v$

Hydraulic resistance is a characteristic of the properties and thickness of the seal unit, and is independent of fluid or pressure. High values are good, indicating low permeation rates, even if wetting does occur. The dimension of permeability is m^2, so the dimension of hydraulic resistance is m^{-1}, and 1 millidarcy has a magnitude of 10^{-16} m^2.

Table 3: Hydraulic resistance of identified intraformational seals and aquitards

Seal or aquitard HFU	Thickness (m)	Kv (md)	Hydraulic Resistance (10^{-16} m^{-1})	Comments
Lakes Entrance Formation	150	0.008	18750	Proven million-year petroleum seal for giant fields
T2B	70	0.7	100	Proven million-year petroleum seal for small fields. Effective (1000 year+) CO_2 seal from modelling
T2A	60	5	12	Only a baffle
Base upper Halibut Subgroup	25	0.6	42	Partially effective (up to 100 years) from modelling
Lower Halibut Subgroup	100	2	50	Not tested in this model

The proven million-year petroleum seal of the Lakes Entrance Formation is also an ideal seal for CO_2 and some of the CarbonNet storage sites can take advantage of this seal. In others, the burial depth of reservoirs under this seal is insufficient to maintain a dense (supercritical) phase. For this family of sites – the subject of the current study – the Traralgon T2B is required to be the effective topseal and this study is designed to investigate storage lifetimes below T2B seals.

By comparing the actual results of plume simulations and storage timeframes within this study (Table 3), it is evident that the T2B interval is the most effective intraformational aquitard (seal) within the Latrobe Group in the nearshore Gippsland Basin, with 1000+ year resistance to permeation and hydraulic resistance of 10 m^{-1}. This agrees with local well data (Hoffman et al., 2012, Hoffman and Preston, 2014, CarbonNet 2015d) that demonstrates several oilfields trapped for millions of years below T2 with water in the reservoirs above.

The base upper Halibut Subgroup shale is resistant to permeation, but is offset locally by faults in the crestal region and therefore its competence could be compromised. The base upper Halibut Subgroup shale is a

particularly useful retarding unit in many, but not all scenarios within this study, and even if it fails to be fully effective can still trap 50% to 75% of the vertically-ascending CO_2 for at least 300 years.

The T2A is a less resistive unit than T2B, as demonstrated by its worse (lower) value of hydraulic resistance and the outcome of simulations. For example, in ten CarbonNet scenarios within this sensitivity and uncertainty modelling, the CO_2 plume entered into the T2A, crossing thin subunits that had seal potential, but were offset by faults. In only two of these scenarios was the plume trapped within T2A for 300 years.

Results

Sensitivity analysis

A 'sensitivity by variable' analysis was performed on the 125 Mt case to determine what the most influential uncertain parameters on CO_2 plume extension and CO_2 plume migration are. In addition to the global variability of parameters, key stratigraphic intervals were separately varied in permeability to see if they had important consequences for plume evolution. These include the T2 seal, the Halibut Subgroup reservoirs and seals, and the overlying Cobia Subgroup reservoirs.

The sensitivity analysis was therefore designed to study variation of the following static model parameters:

- Overall Kv/Kh (the ratio of vertical to horizontal permeability at cellular scale)
- Kv/Kh for sealing facies
- Kv/Kh for reservoir facies
- horizontal permeability multiplier for a range of subunits and facies
- Fault horizontal transmissibility (binary: open / closed)
- Porosity multiplier
- Diffusion coefficient
- Residual gas saturation Sgr
- Relative permeability tables
- Aquifer volume
- Average aquifer permeability

The parameter values range between a selected minimum and maximum as listed in Table 4. These ranges far exceed the range of variation observed in average reservoir properties within offset wells in the nearshore Gippsland Basin. For example, permeability is modelled to vary up or down by a factor of 10 in most cases, which represents the P01 and P99 percentiles of the individual well core measurements and are assessed to be at the P0.01 and P99.99 of HFU averages. For comparison with industry practice, an order of magnitude is also typical of the range used for petroleum sensitivity studies. To include the effect of potential unmapped faults and fractures below the limit of seismic resolution which could act as fluid conduits, seal permeability (both Kv/Kh and Kh) was allowed to reach 50 times the base case.

Table 4: Sensitivity variations applied in the study

Sensitivity Parameters	Low	Base	High	Distribution
K_v/K_h whole model	0.001	0.01	0.1	Triangular (log)
K_v/K_h for sealing facies only	0.001	0.01	0.5	Triangular (log)
K_v/K_h for reservoir facies only	0.001	0.01	0.1	Triangular (log)
K_v/K_h only for sealing facies within key T2B seals	0.001	0.01	0.5	Triangular (log)
T2B horizontal permeability multiplier (sealing facies only)	0.1	1	10	Triangular (log)
T2A horizontal permeability multiplier (sealing facies only)	0.1	1	50	Triangular (log)

T2 horizontal permeability multiplier (sealing facies only)	0.1	1	50	Triangular (log)
T2 interbedded sand horizontal permeability	0.1	1	10	Triangular (log)
Intra-formational shale (base Halibut) horizontal permeability	0.1	1	50	Triangular (log)
Halibut horizontal permeability (sand facies only)	0.1	1	10	Triangular (log)
Halibut horizontal permeability (sealing facies only)	0.1	1	50	Triangular (log)
Fault lateral transmissibility	closed	open	open	Rectangular
Porosity multiplier	0.6	1	1.4	Triangular (linear)
Diffusion coefficient	Off	Off	On	Rectangular
Sgr	0	0	0.15	Rectangular
Relative permeability table	Berea sandstone (1.2mL/min)	Cardium (IFT=56.2)	Cardium (IFT=19.8)	Triangular
Aquifer volume (Gm^3)	600	800	1000	Triangular (linear)
Average aquifer permeability (mD)	250	500	1000	Triangular (linear)

The 18 variable parameters in Table 4 define a base case and a set of 33 unique sensitivity runs - 15 triangular variables (low-base-high), and 3 rectangular (on-off). From these, we aim to identify the most influential uncertainty parameters (also called Key Uncertainty Parameters KUP), and verify plume conformance in all of them. It is important in this type of analysis to have an objective numerical measure of plume containment. In this study, two key lateral and vertical responses formed the objective numerical measures of the plume movement and position at the end of ~300 years of injection and migration. The principal response measures used to assess whether a parameter was important were a measure of lateral containment (approach of the plume towards the western boundary) and a measure of vertical containment (approach of the plume to top of the T2 seal). Both responses were given equal weight in assessing which influential parameters to evaluate further.

Sensitivity results

A “tornado” plot was created from all 33 sensitivities run against the base case, for the two key observable parameters – distance from western boundary (Figure 4), and vertical ascent relative to the top of the T2 seal (Figure 5). From this analysis, several key uncertain parameters which are most influential in the plume lateral extent and vertical migration are considered. The selection of key uncertain parameters are required for uncertainty analysis (and proxy modelling) in the next phase of response surface methodology.

In the tornado plots, red colours indicate the response when the parameter is increased, and blue the response when the parameter is decreased. It is clear that several parameters have negligible effect and therefore need not be considered further. Those parameters that do have effect, seem to be of similar importance in both lateral and vertical response, although some are of like polarity (blue for blue) and others of opposite polarity. This is because some parameters lead to overall faster or slower plume movement in both the horizontal and vertical directions, while other parameter variations such as seal permeability have a directional effect and tend to hold the plume down vertically and force it to flatten-out laterally, or conversely allow it to more freely ascend, with less lateral spread.

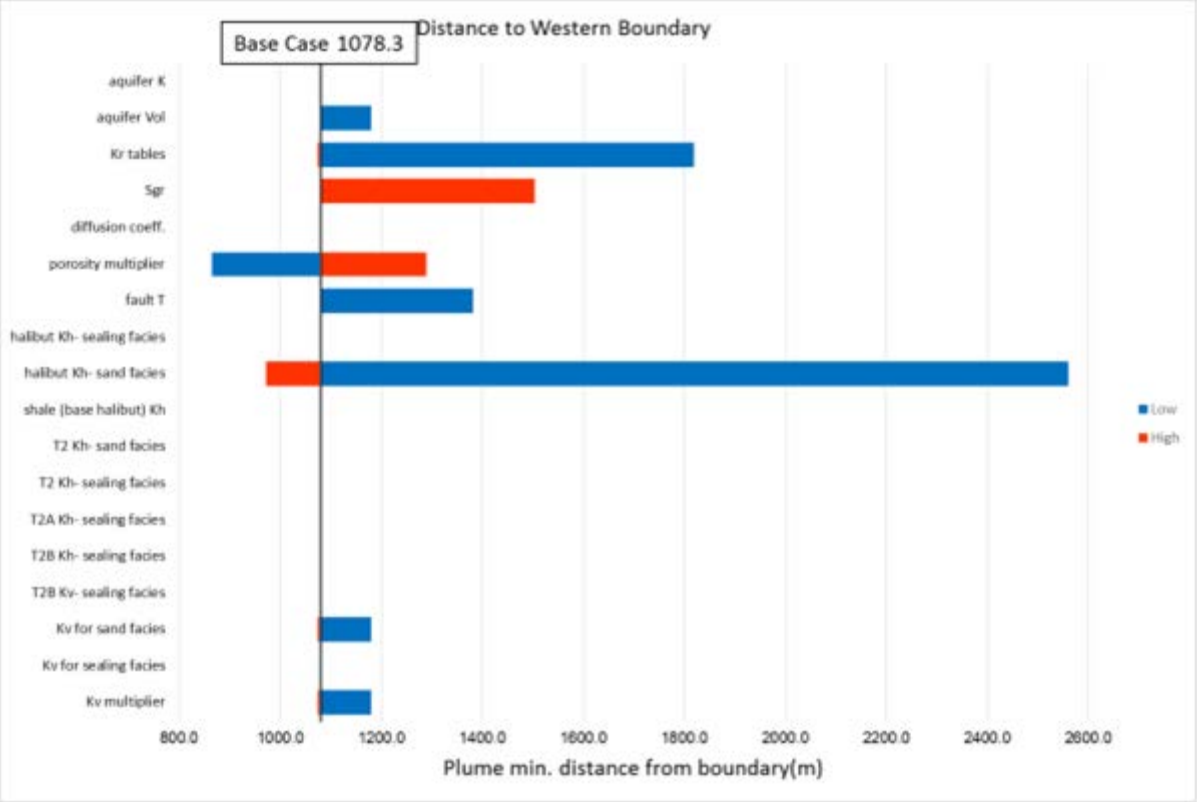

Figure 4: Tornado plot of CO_2 plume approach to the western boundary at I+300

Analysis of the many possible variables show that many are unimportant, or weak in terms of lateral response of the plume. Eight measurable responses are observed, some of which will be combined in later steps.

Sensitivity Result – approach to the western boundary

Figure 4 shows that in the base case, the plume approaches to a little over 1 km from the western boundary after 300 years. The closest approach noted in any sensitivity run is ~500 m and therefore the western boundary is never reached nor crossed. The conservative nature of the base case is evident in the results, with distance from the western boundary being biased to positive rather than negative changes (i.e. the most likely outcome is more distant from the boundary than the base case).

Sensitivity Result – vertical migration

Figure 5 shows that, in the base case, the plume rises to ~140 m below top T2 seal after 300 years. All the sensitivity runs are reported in terms of variation from this base case. The closest approach of the plume to top T2B is ~35m, but the top of T2B seal is never reached.

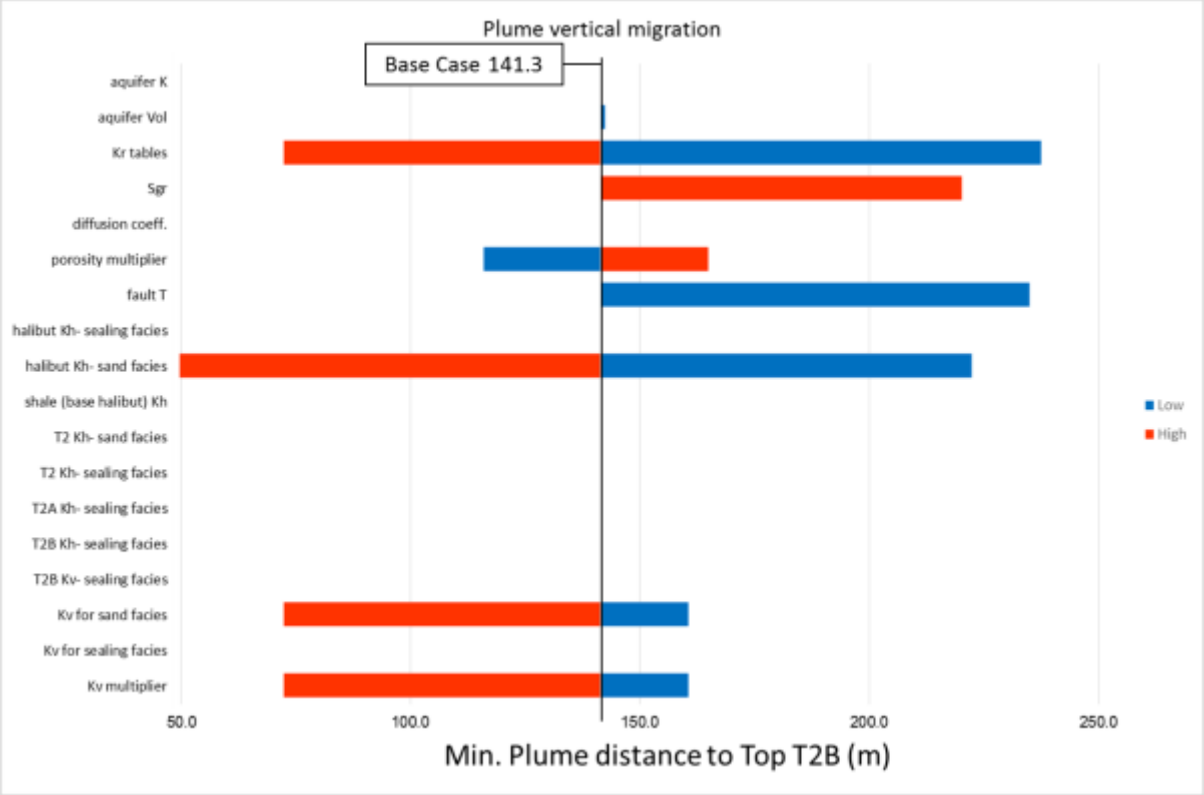

Figure 5: Tornado plot of CO_2 plume distance to top T2B at 300 years, relative to base case

Analysis of the many possible variables show that many are unimportant, or weak in terms of vertical response of the plume. Seven measurable responses are observed, some of which will be combined in later steps.

Analysis identified that six of the eighteen studied parameters were potentially significant. These are identified in bold blue text in Table 5, with some of the minor stratigraphic variations already consolidated to offer twelve items for consideration:

Table 5: Six potentially significant variables (blue) identified by the sensitivity study

Uncertain Parameters	Lateral response	Vertical response	Potentially significant?	Independent?
K_v/K_h	Medium, positive	Strong, negative	YES	YES
K_v/K_h for sealing facies only	Nil	Nil	No	Included in total kv/kh
K_v/K_h for reservoir facies only	Medium, positive	Strong, negative	YES	Included in total kv/kh
K_v/K_h only for sealing facies within key T2B seals	Almost Nil, negative	Nil	No	Included in total kv/kh
Horizontal permeability	Very Strong, positive for reservoirs only	Medium, positive for reservoirs only	YES	YES
Fault lateral transmissibility	Strong, positive	Very Strong, positive	YES	YES
Porosity multiplier	Strong, negative	Strong, negative	YES	YES
Diffusion coefficient	Nil	Nil	No	yes
Sgr	Strong, positive	Strong, positive	YES	YES
Relative permeability tables	Strong, negative	Very strong, negative	YES	NO (see discussion below)
Aquifer volume (Gm^3)	Weak, positive	Very weak negative	No	Both aquifer parameters have similar effect
Average aquifer permeability (mD)	Almost nil, positive	Almost Nil, negative	No	

Having allowed high values of Kv/Kh and high Kh (50 times base case) in seal units, to allow for the presence of potentially leaky fractures, it was surprising that the response to this high permeability multiplier was almost zero. The main response to Kv/Kh and Kh multipliers is through the permeability of reservoir units. We interpret that the main lateral and vertical migration routes are via tortuous connections of reservoir units and this flow path is enhanced when reservoir permeability is enhanced. When seal permeability is enhanced, it is still resistive to flow and offers no additional flow paths. Hence, when merging reservoir and seal permeability ranges, an upper limit of 10 times Base case was adopted, consistent with core permeability statistics.

It was noted that the complex distinctions of different stratigraphic subsets were not significant, and the parameter variations were therefore identified to be applied across the entire stratigraphy in the following stages of the study. The effect or consequence of each of the six parameters in terms of plume mobility or storage is also noted in Table 6.

Table 6: Selection of key uncertainty Parameters

Importance	Variable and alias	Effect on plume
1	Kv Vertical permeability (technically, Kv/Kh)	Vertical Mobility Has main effect within porous units by allowing plume to rise faster through stratigraphy along tortuous paths.
2	Kh Horizontal permeability	Horizontal Mobility Has main effect within reservoir units by allowing plume to spread wider and faster
3	Fault Transmissibility	Horizontal Mobility As above
4	Sgr	Storage Allows the plume be stored in a smaller or larger volume of rock in cases where migration occurs, by changing the amount of CO_2 left behind as residual saturation
5	Porosity	Storage Allows the plume be stored in a smaller or larger volume of rock by changing the pore space available
6	Relative permeability tables	Mobility AND Storage As above

A reduction of complexity was made by eliminating fault transmissibility as a key parameter for two reasons:

1) Lateral fault seal is both unlikely (from basin statistics) and actually beneficial to storage. The sensitivity analysis showed that the distance to both the western boundary and top T2 both increased if faults seal laterally. It is obvious that approach to the western boundary is reduced by lateral fault seal, since the movement of CO_2 is partially held back to the east flank. The reduced vertical rise is because less reaches the crestal, and therefore cannot rise so strongly In conclusion, no additional storage risk accrues if the faults do seal, and neglecting fault seal makes the modelling inherently conservative.

2) Vertical fault transmissibility of infrequent small faults and fractures is covered by allowing a wide range of dispersed (average) vertical transmissibility through seals (Kv/Kh).

An important element of identifying the key sensitivities for later multi-dimensional variation modelling (uncertainty modelling) is that the selected variables should be both ***independent*** and ***significant*** in their effect. In this way, a full range of parameter space can be explored, without repetition. For this reason, although relative permeability tables have a significant effect, they were assessed to not be independent, since they overlap with intrinsic permeability (K or Kh) in terms of a mobility effect and with porosity in terms of a ***storability*** effect.

It was thus determined that the four Key Uncertainty Parameters (Figure 6) were:

1. Overall Kv:Kh ratio
2. Overall Kh
3. Porosity
4. Residual Saturation (Sgr)

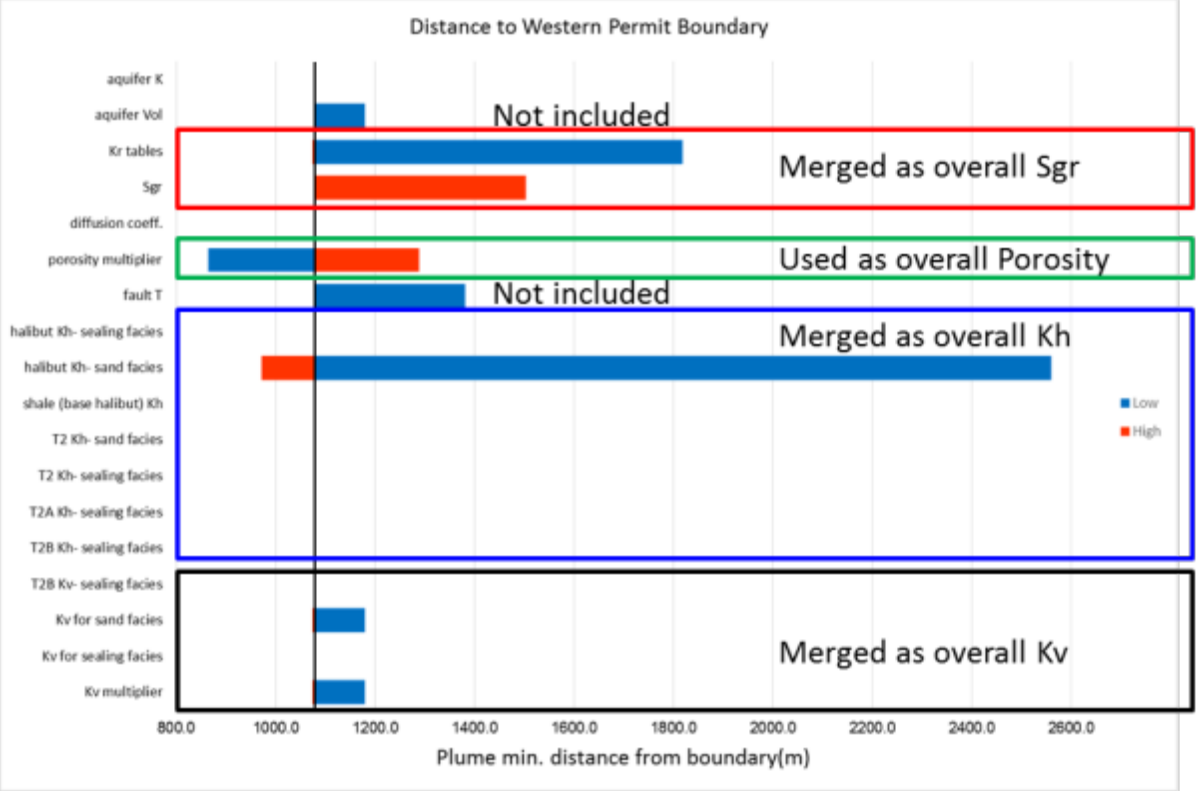

Figure 6: Merging of like parameters to the final four Key Uncertainty Parameters (KUP)

All Kh elements are merged into one. All Kv/Hh variables are merged into one. Kr tables and Sgr are merged into a single saturation term, and porosity stands as the fourth independent variable.

Table 7: Key Uncertainty Parameters and their ranges (Figure 7)

Variable	Distribution	Min	Base Case	Max
Kv:Kh	Triangular (log)	0.1 x base case	1% of Kh	10 x base case
Kh	Triangular (log)	0.1 x base case	Values from model	10 x base case
Porosity	Triangular (linear)	0.6 x base case	Values from model	1.4 x base case
Sgr	Rectangular (linear)	0.0		0.15

Variable	Distribution	Min	Base	Max
Kh	Triangular for Log10(Kh)	0.1 x base	Values from model	10x base
Kv	Asymmetric Triangular for Log10(Kv/Kh)	0.1 x base	0.01 x Kh	10x base
Porosity	Triangular	0.6 x base	Values from model	1.4x base
Sgr	Rectangular	-	0	0.15
Random	Gaussian			

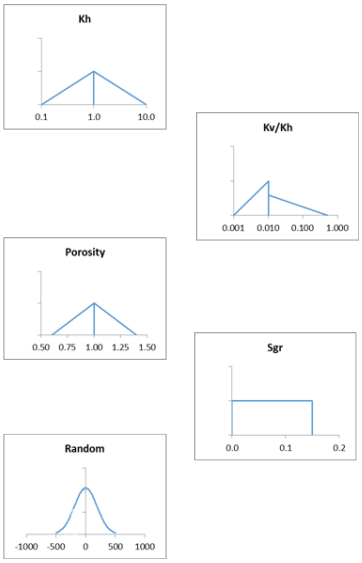

Figure 7: Key Uncertainty Parameters and their probability distributions

Three of the key uncertain variables are triangular – two of these are symmetric, and one asymmetric. One variable is a two-point (binary) distribution with rectangular form. In comparison, a random variable would have a Gaussian distribution.

Uncertainty analysis and proxy modelling

In the next step of the study, variables were allowed to co-vary, two at a time to explore whether there are non-linear interactions between the variables, and to define a probabilistic expectation of plume outcomes.

The objectives of the uncertainty analysis were to develop technically sound probability cases and observe plume pathway variability. Key sensitivity parameters were applied as scale factors, uniformly across the whole model, two at a time for the uncertainty analysis to identify any interactions.

Proxy models represent the outcome (the Objective Response, measured from actual reservoir dynamic simulations) as a simple mathematical consequence of the weighted and scaled input variable parameters and the coupling between them. It is thus a rapid way to explore parameter space and predict outcomes of different combinations of inputs, without running many computationally intensive full reservoir simulations. It is important to check that the predictions from the proxy agree sufficiently well with actual simulations in selected scenarios, especially when multiple parameter extremes are simultaneously invoked.

The proxy model is used in Monte-Carlo modelling to statistically explore parameter space and define probabilistic plume outcomes (as 1D Objective Responses), including the approach to the lateral and vertical storage boundaries and the chance of passing these boundaries.

Method of analysis

The task of how to set up a combination of simulation runs that can be used to train the proxy model is usually referred to as “experimental design” or Design of Experiment (DOE) and is discussed extensively in the petroleum literature (e.g. Allessio et al, 2005). An experimental design can be applied to physical or numerical experiments with or without additional random or external influences. The quality of a proxy model is obviously dependent on the simulation cases that were used to train it. Selecting the design of experiment, means selecting the sampling method.

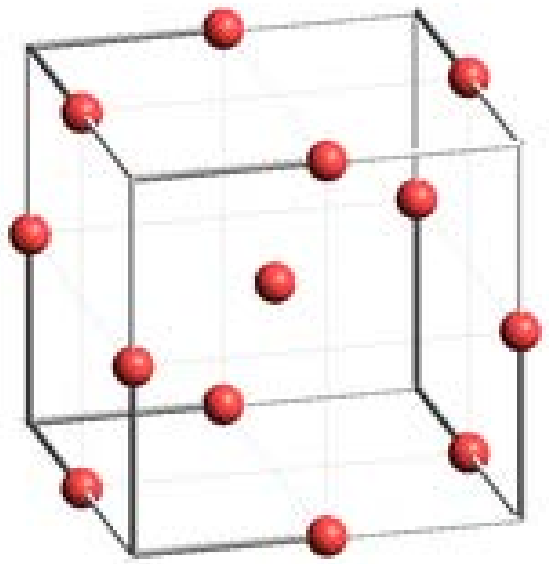

Figure 8: Box-Behnken sampler in three dimensions

Two variables at a time are adjusted to either high or low, defining the midpoints of the 3D wireframe elements. The corners of the wireframe would require three parameters to vary. Our case is 4-D, and the "corners" of the 4D hypercube are 4-variable cases.

The Box-Behnken sampling method was selected for this study (Figure 8). This sampling method is a deterministic sampler that is suitable for making quadratic proxy models because it uses three levels of each of the input parameters (fully LOW, neutral, fully HIGH). Box-Behnken does not sample the variables at intermediate strength and it is assumed that the three-point output (LOW-Neutral-HIGH) is sufficient to fit a quadratic equation. Experimental and random noise (if present) is addressed through repeat experiments or by multiple tests of the LOW and HIGH cases of one parameter, when other parameters are also independently varied. Figure 9 shows the setup of the Box-Behnken sampler for the four key uncertainty parameters in the CarbonNet study.

An optional alternative construction of the experimental design would be to sample probability space independently for the n parameters, and to conduct a Monte-Carlo style series of simulations that can cover the probability space adequately. There is considerable discussion in the petroleum literature as to whether a probabilistic approach or a regimented approach is best, and this depends on the nature of the "experiment" or simulation, on the number of variable parameters, and the preference to either minimise computing time or maximise information return from the modelling. We will not discuss these issues but note that with either approach, our proposed method of probabilistic 3D spatial analysis of the actual experimental outcomes is equally valid.

	Kv/Kh	halibut Kh	porosity	Sgr
Exp #	*F1*	*F2*	*F3*	*F4*
1	-1	-1	0	0
2	0	1	1	0
3	0	-1	0	1
4	1	0	1	0
5	0	0	-1	1
6	-1	0	1	0
7	1	0	0	1
8	0	0	0	0
9	0	-1	1	0
10	1	0	-1	0
11	0	0	0	0
12	0	0	1	1
13	1	-1	0	0
14	-1	0	-1	0
15	0	1	-1	0
16	0	0	-1	-1
17	0	0	1	-1
18	0	1	0	1
19	1	1	0	0
20	1	0	0	-1
21	-1	0	0	1
22	0	-1	0	-1
23	0	0	0	0
24	-1	1	0	0
25	0	1	0	-1
26	0	-1	-1	0
27	-1	0	0	-1

Figure 9: Setup of the Box-Behnken sampler for the four key uncertainty parameters

Note the apparent repeat of the base case 0-0-0-0 on three occasions. This repeat is not needed in our situation since there is no experimental drift or external uncertainty (simulation is an exact and repeatable "experiment"). In other situations, this repeat of an identical "experiment" allows random variability to be quantified.

Probabilities of combinations of variables

It is important to be able to estimate the resulting probability of a combination of independent variations of more than one parameter. To do this, we first assess the probability associated with a single extreme variation of one parameter. We have already assessed that the extreme upper and lower limits of the variables are at the 10^{-2} probability level.

However, in a 3-point sampling, the end member case is also representative of a range of nearby values. Figure 10 shows that for a Triangular distribution, the end members each represent about 12.5% of probability density space, while the central case represents about 75%. For a 3-point sampled rectangular distribution, the values are 25% and 50%, respectively.

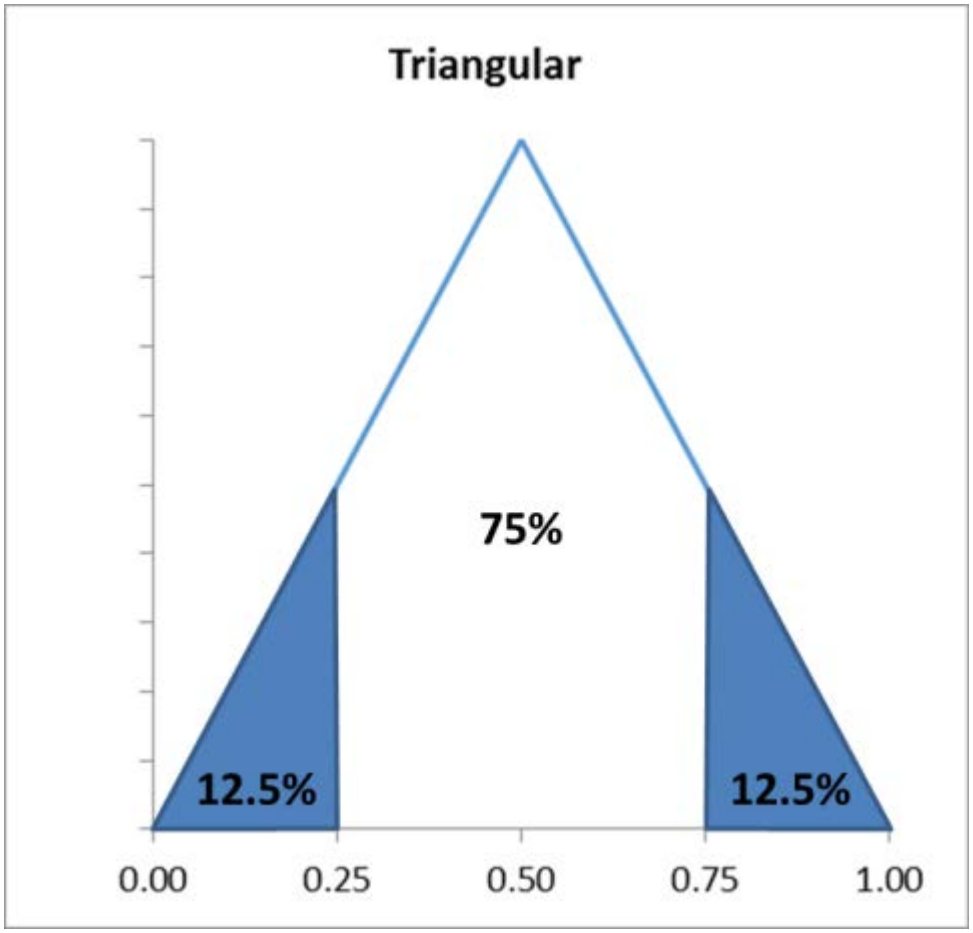

Figure 10: Representative probability of the extremes of a triangular distribution

For a three-point triangular distribution, the central point represents the majority (75%) of the whole range of outcomes from 0 to 1, while the high and low end points each represent 12.5% of outcomes.

In this study, three of the four key uncertainty parameters have triangular probability distributions, and one is rectangular and we can thus estimate the representative probability of a combination of two, three, or four extreme end-members of the four key uncertainty parameters (Table 8).

Table 8: Probability estimation of extreme multivariable cases

Number of variables	***Representative*** probability of one extreme case	Multiplied representative probabilities
One	16% (P84) or 10^{-1}	16% (P84) or 10^{-1}
Two		2.4% (P97.6) or 10^{-2}
Three		0.38% (P99.6) or 10^{-3}
Four		0.06% (P99.94) or 10^{-4}

An example of the analysis is presented in Figure 23. In this figure, all 24 possible combinations of two variables are included, plus the base case and the 8 single-variable runs. In addition, all downside cases of three and four variables are included to show that the cases where the plume approaches the western boundary are fully evaluated.

Proxy Model and Monte-Carlo simulation

In order to mathematically analyse the uncertainty of plume migration using the defined parameters, a surrogate model or proxy that mimics and simplifies the simulation model is required. This provides an

efficient tool to obtain results, without running the simulator, for a statistically large number (e.g. 10,000) of varying inputs and outcomes of interest, called 'responses'.

A Proxy equation is derived from the **actual** responses for each of the four key uncertainty parameters F 1 to F 4, to ***predict*** the response of partial variations and mixtures of the four key sensitivity parameters including one at a time, two at a time, and up to three or four at a time. For this study, second order additive and higher order polynomial proxies were tested.

[Eq 2] Second order Polynomial function (additive terms only)

$$Response=b0+b1*F1+b2*F2+b3*F3+b4*F4+b5*F1*F1+b6*F2*F2+b7*F3*F3+b8*F4*F4+b9*F1*F2+b10*F1*F3+b11*F1*F4+b12*F2*F3+b13*F2*F4+b14*F3*F4$$

This type of function is traditionally used to determine the response form with Box-Benhken designs. But it actually implies that the main effects are additive rather than multiplicative. For a number of physical processes involving fluid flow or volumetrics for instance, this can be shown to be not suitable (Alessio, 2005).

[Eq 3] Combined Product function (higher order polynomial)

$$Response=b0+F1^{b1}*F2^{b2}*F3^{b3}*F4^{b4}+b5*F1*F2+b6*F1*F3+b7*F1*F4+b8*F2*F3+b9*F2*F4+b10*F3*F4$$

This type of function can be very useful when the main effects are essentially multiplicative, rather than additive. The factors b5 to b10 are "cross-terms" i.e. factors for the main compounding effects of each parameter when combined with another.

In order to define the most appropriate form of the proxy, we tested the first order part of each equation, e.g where we only consider the single order sensitivities (no cross-parameter terms). That means, for Eq 1 and Eq 2: bi , i <5 <>0. We will call these two proxies additive and multiplicative respectively. We tested the two options on the different responses, and noted a slight superiority of the multiplicative fit, therefore Equation 2 was selected as the basis for detailed proxy fitting.

Using a minimising algorithm from Excel ('Solver'), combined with a simple macro, a best fit was generated. Note that a reduction of fitting parameters (Fi) was also conducted, which led to the results captured in Table XX. All cross-terms (*FiFj*) evaluated as zero, since an appropriate degree of cross-coupling is already included in the model because it is a multiplicative proxy, rather than additive, so the effects of different terms automatically amplify each other through multiplication.

Table 9: Final selected proxy fitting parameters

	Lateral response	Vertical response
F1 (Kv/Kh)	1.200	0.700
F2 (Kh)	0.951	0.830
F3 (porosity)	0.885	0.700
F4 (residual saturation)	0.923	0.700
FiFj	All cross-terms evaluated as zero	

The comparison of predicted result from the final proxy equations, and the input and test cases is displayed in Figures 11 and 12. An excellent fit is observed with only small residuals (R^2 of 0.97 and 0.93 respectively).

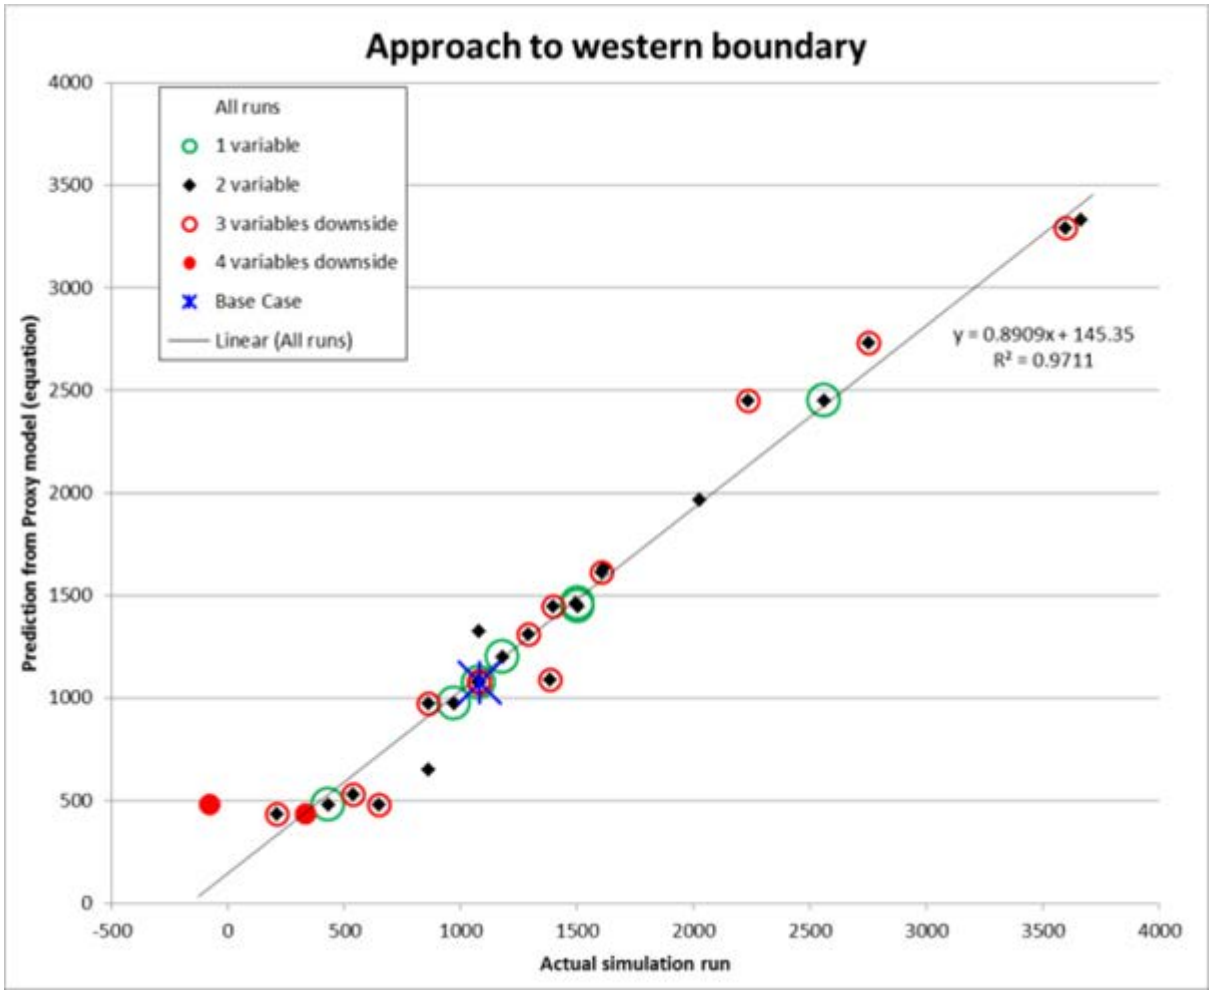

Figure 11: Proxy and actual simulation results: approach to the western boundary

Data points show the prediction of the proxy equation (vertical axis) compared to the actual simulation run outcome (horizontal axis). The proxy equation is trained on the 1- and 2- variable runs, and tested with the 3- and 4- variable cases. There is a good linear fit with a small departure for the 4-variable downside case where the actual simulation outcome is more extreme than the proxy prediction. The base case is conservative – the centre of gravity of the results is clearly more distant from the boundary than the base case.

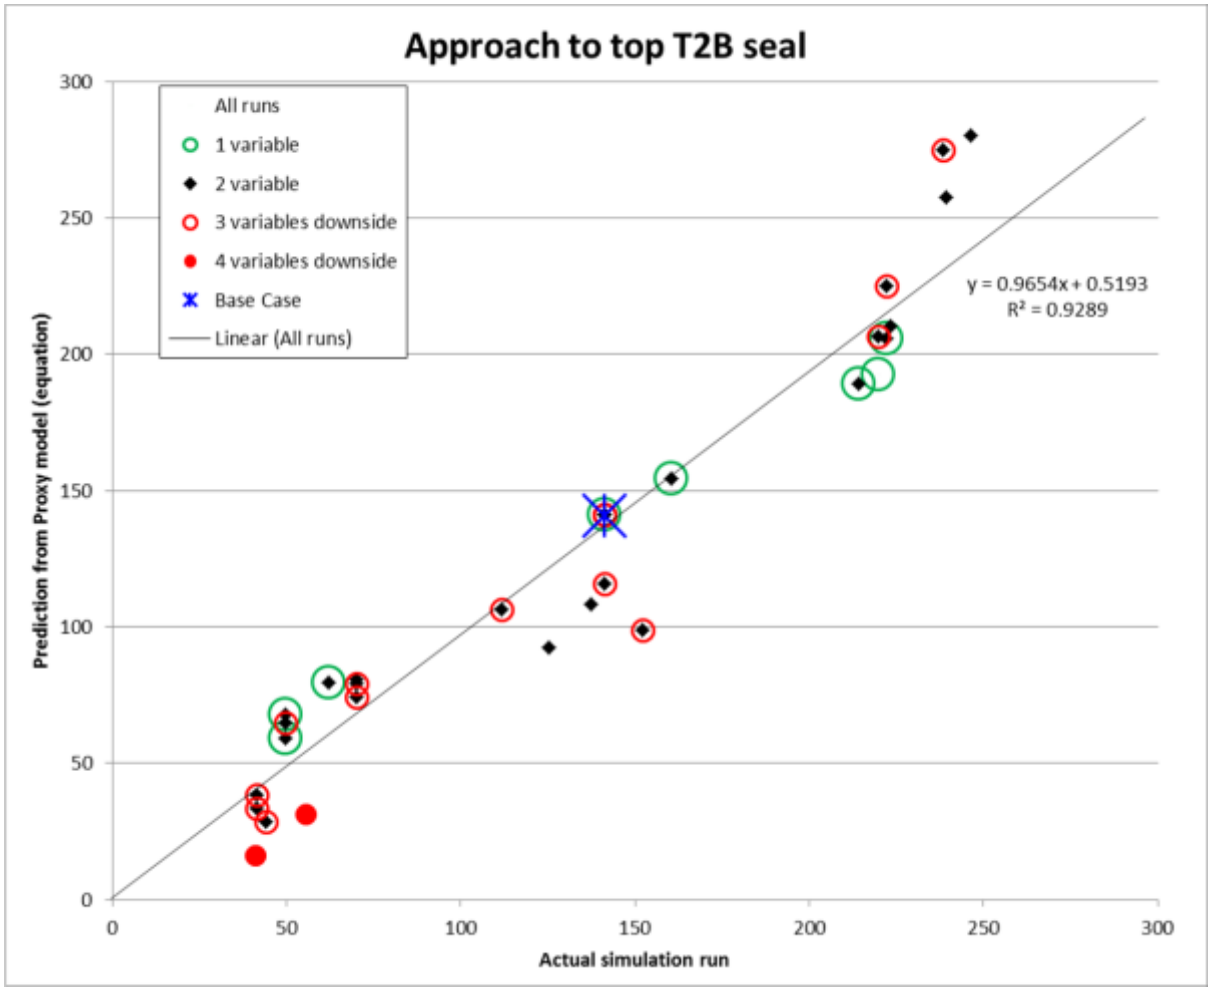

Figure 12: Proxy and actual simulation results: distance to top T2 seal

Data points show the prediction of the proxy equation (vertical axis) compared to the actual simulation run outcome (horizontal axis). The proxy equation is trained on the 1- and 2- variable runs, and tested with the 3- and 4- variable cases. There is a good linear fit with a small departure for the 4-variable downside case where the actual simulation outcome is less extreme than the proxy prediction. The base case is approximately at the mean of the distribution.

Monte Carlo Simulation of Plume extents

A Monte-Carlo simulation was then conducted, using the proxy equations to predict plume approach to the boundary and to top of T2B seal in a wide range of combinations of all KUP's, with effect between full negative and full positive, and statistically sampled values over the full range between these extremes. This ensures that any unusual combinations of variables will be modelled and a full range of outputs determined.

- Probability distributions were assigned as triangular, except for Sgr, which is less known and therefore given a rectangular distribution.
- The output function was adjusted to match mean, standard deviation, and P90 between input (actual runs) and output (proxy prediction + adjustment)
- As part of the adjustment process, and to ensure the full range of behaviour was simulated, an additional random variable was included (380m for lateral boundary, 30.5m for T2B).

The outcomes of Monte-Carlo modelling (Figures 13 and 14) show that there is a very high confidence that the plume will remain beyond the western boundary and below the T2B topseal, but that there is a small outlier representing failure cases at around 0.5% to 1.5% level of possibility.

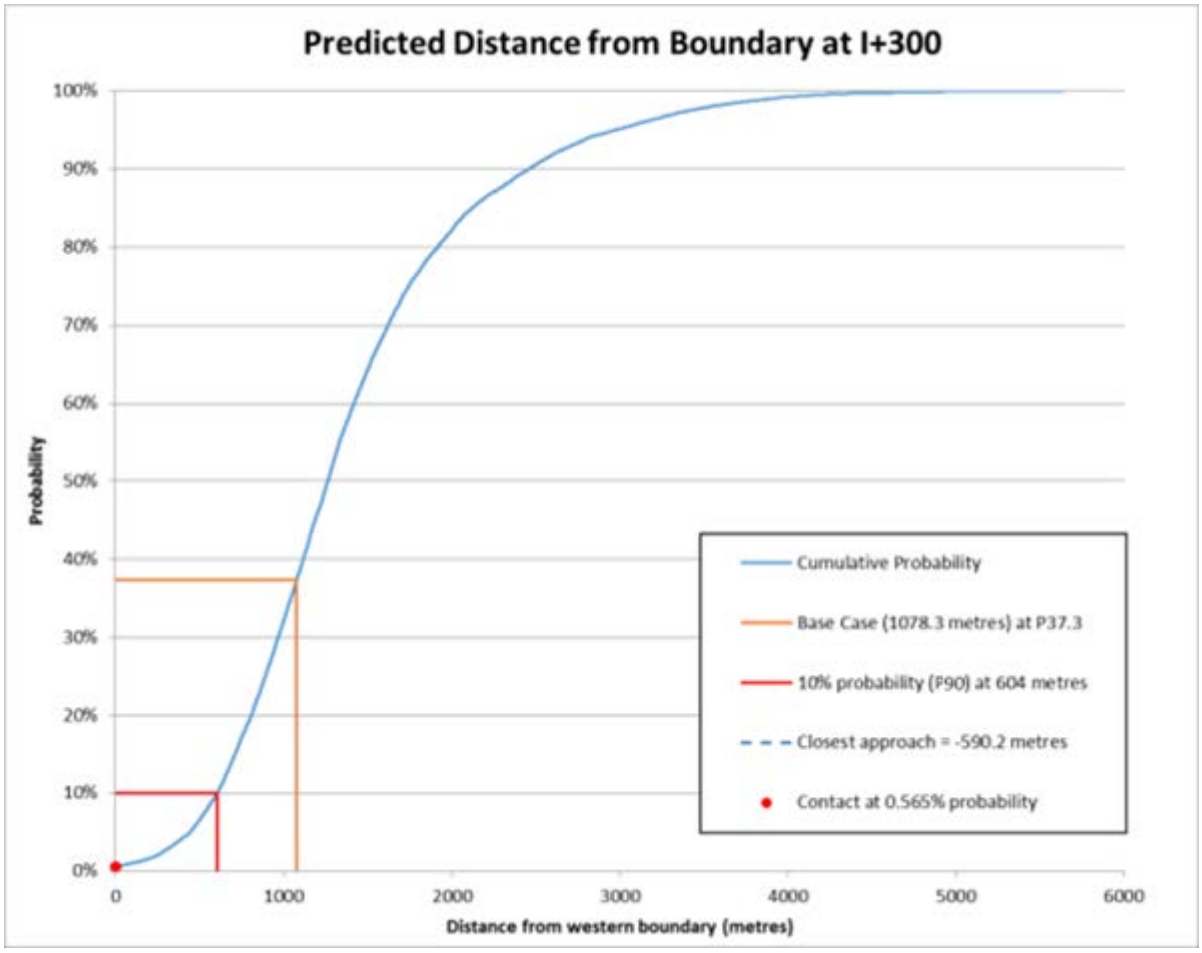

Figure 13: Monte-Carlo results for distance to western boundary

Simulation of a large number (10,000) of random combinations of parameter variations (each distributed continuously according to Figure 7) results in a prediction of the P90 approach and of the (un)likelihood of a failure case for lateral containment

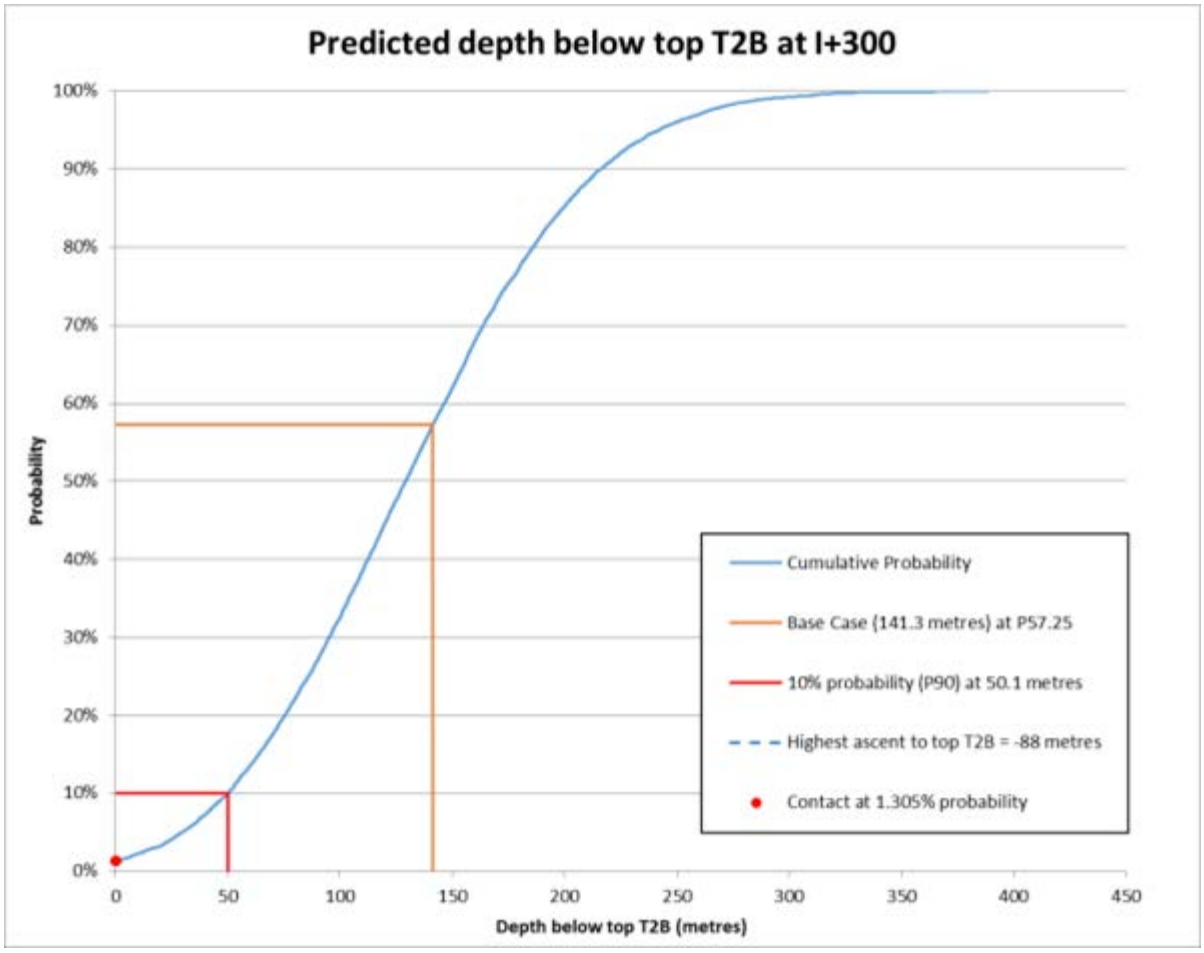

Figure 14: Monte-Carlo results for depth below T2B

Simulation of a large number (10,000) of random combinations of parameter variations (each distributed continuously according to Figure 7) results in a prediction of the P90 approach and of the (un)likelihood of a failure case for vertical containment

Probability of Plume Path in 2D and 3D

The novel aspect of the current work is to generate probability distributions in 2D and 3D space for the presence or absence of CO_2:

Statistical Prediction of Plume Paths

To generate a statistical prediction of plume paths space, CarbonNet adopted a Boolean statistical approach based on plume presence as described below:

For each sensitivity variation, and each key timestep, the full 3D volume of the plume saturation was converted to a binary Yes/No 3D volume for all grid cells, i.e., whether a plume enters a cell or not. This 3D binary grid was then either projected to a 2D map and used for 2D probabilistic summation or used directly in a 3D probabilistic summation. The reason for adopting this approach is to ensure that all possible plume paths are mapped, not just those that are taken by large volumes of fluid. This approach ensures that the fringes of the plume are fully-included in any plume path prediction, rather than a saturation-based measure, which would be biased to the plume core, and would tend to ignore the fringes if they represent less than 10% of the injected volume.

Key timesteps selected are:-

- 10 years after injection commences (to inform on early plume evolution)
- 25 years after injection commences (i.e., end of injection)
- 30 years after injection commences (nominal site closure, 5 years after end of injection)

- 50 years after injection commences (nominal issue of site closing certificate)
- 100 years after injection commences
- 200 years after injection commences
- At the end of the model runs, 300 years after injection commences

An example of the 2D method is illustrated in Figure 15. A 2D (map) projection of the 3D spatial distribution can be visualised easily. The plume for the base case, 10 years after the start of injection (I+10) is shown as a colourmap representing integrated CO_2 column density for mobile non-dissolved CO_2 (this is the effective total CO_2 column, but may represent several separate columns under different seals and baffles, rather than a single amalgamated column). Hotter colours have more CO_2 and, as expected, the CO_2 is concentrated in the crestal area of the structure by density segregation (buoyancy). The extreme limit of the plume is given by the zero edge of CO_2 density (white). Beyond this, there is absolutely no buoyant CO_2 saturation. The red line is the western boundary referred to in earlier sections.

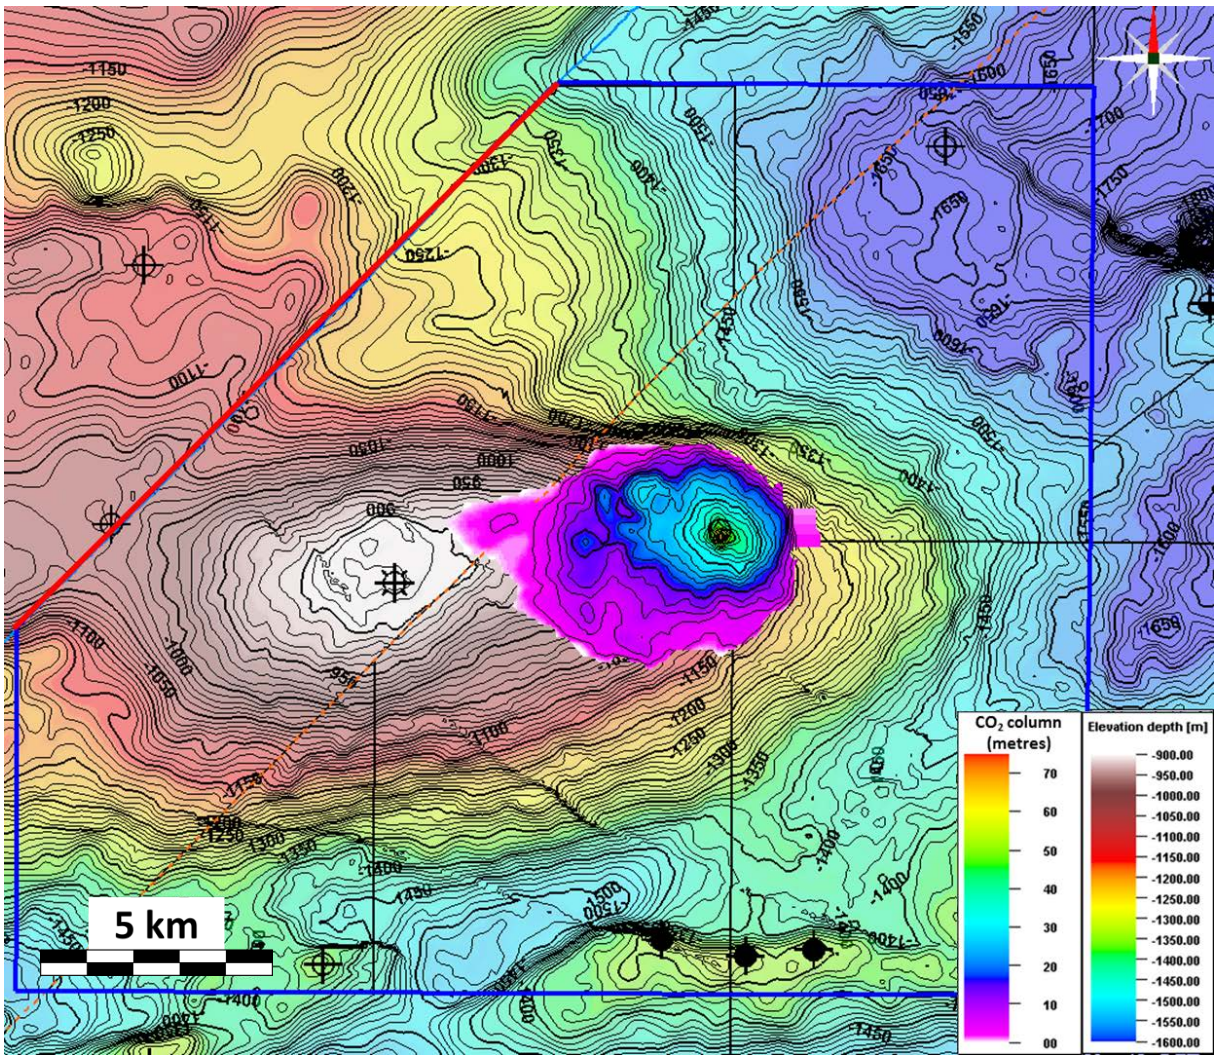

Figure 15: The buoyant CO_2 column for the base case, at (I+10)

CO_2 saturation in the base case dynamic model, ten years after injection commences (I+10) is displayed as a conventional colourmap, with contours, overlaid on a depth contour map of the structure. Injection occurs at a convenient location on the eastern flank of the structure, to keep the plume away from the western boundary.

In principle, one could add up similar column density maps for all variations, but that would tend to overemphasise the buoyant core of the plume and under-represent the thin edges of the plume, where the exact extent is of particular interest. Therefore, the Boolean map is produced (Figure 16, showing whether the

plume is present (red) or absent (no colour). This is exactly the same data as Figure 15, but all parts of the plume are given equal weighting.

With these Boolean maps, one can add up each variation that has been run and count the number of cases that enter each cell. This generates a probability map, with equal weight assigned to the fringes and the core of the plume. A map can then be produced showing the probability that any cell contains part of the plume, for all studied sensitivity cases.

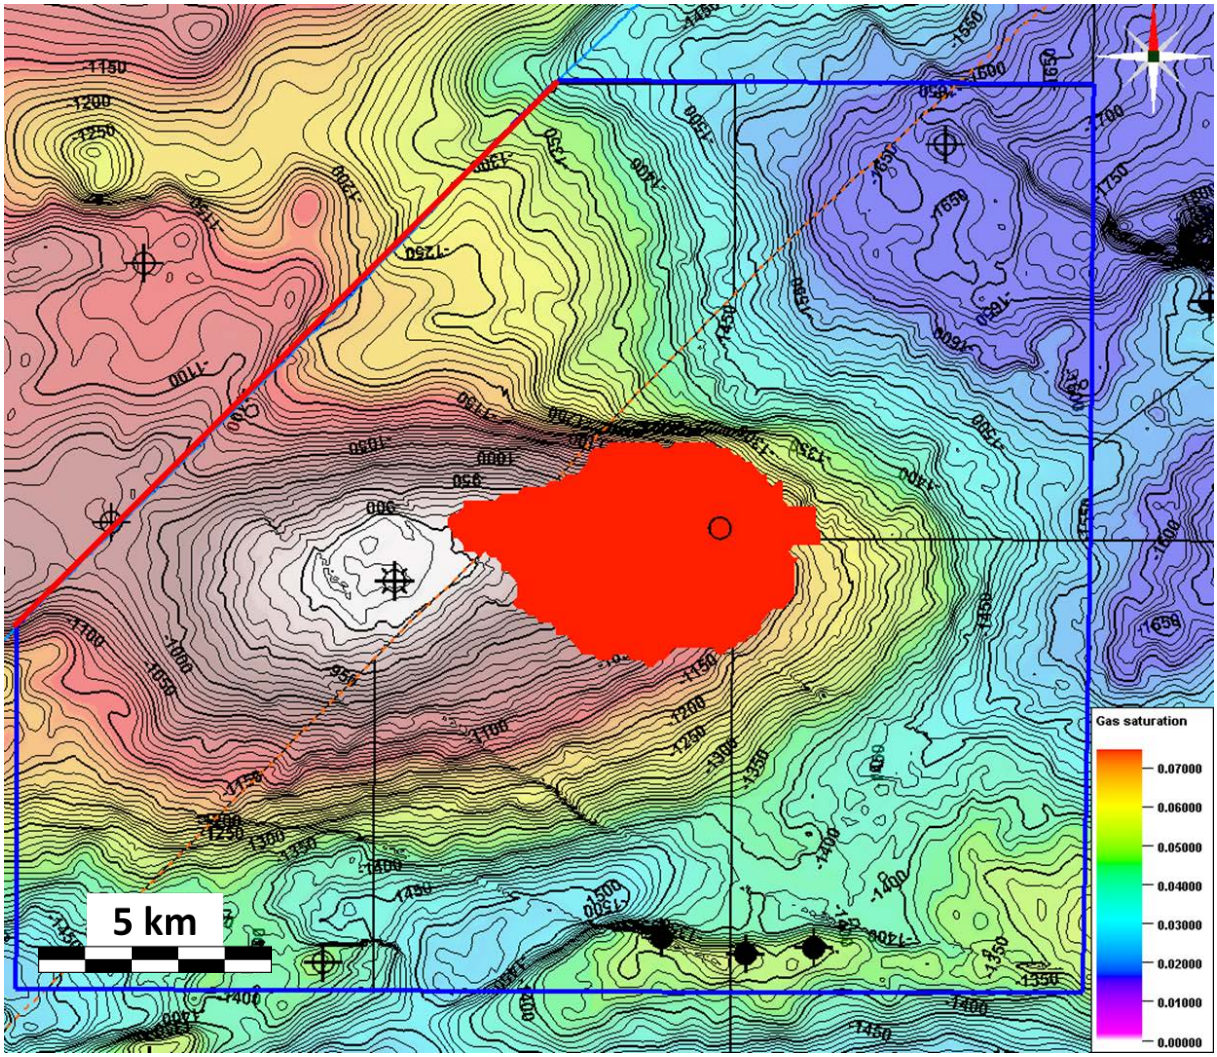

Figure 16: The buoyant CO_2 plume extent for the base case, at (I+10)

CO_2 presence at I+10 is displayed as a solid red shape, overlaid on a depth contour map of the structure. CO_2 is entirely absent outside the red shape.

2D Probability mapping

To generate a map of plume path probability, the set of all Boolean grids (e,g, Figure 15) is stacked, for a single timestep. Weighting is applied to the actual runs to reflect the combination of probabilities described in Table 8. Figure 17 shows this stack of grids at the timestep 10 years after injection commences (I+10). Contour interval is 5% with bold contours at P90, P50, and P10.

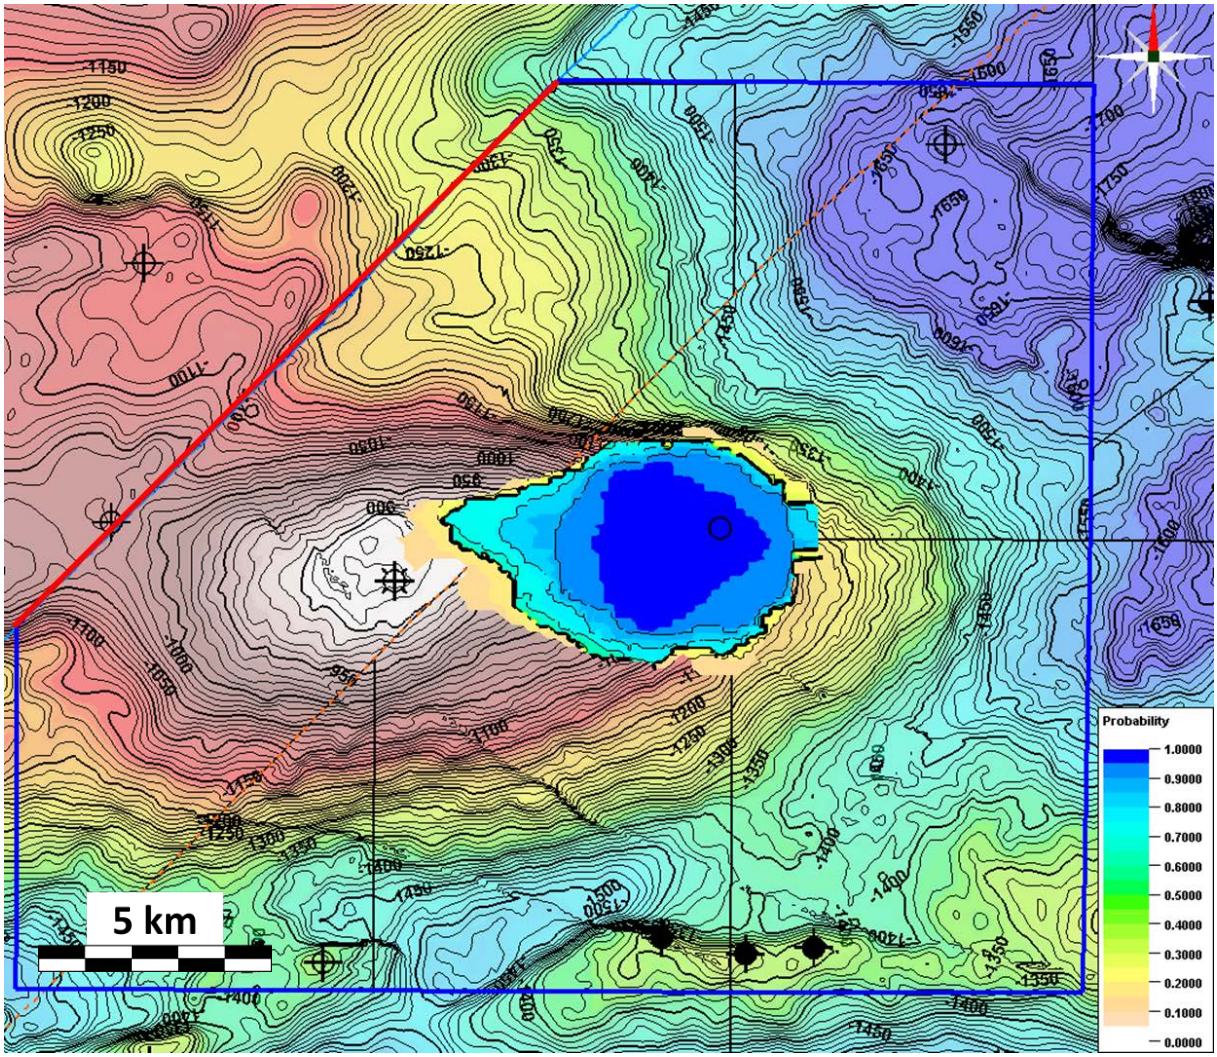

Figure 17: Boolean probability map at (I+10)

The entire set of CO_2 presence/absence shapes is stacked for all sensitivity and uncertainty runs, with appropriate weighting to derive a spatial probability map of plume presence at I+10.

For clarity, and to emphasise that there is a degree of spatial uncertainty to the discrete permeability pathways that determine the exact limits of a plume, this probability map is then smoothed (Figure 18). Since the goal is to define the plume extent at more than 10% probability level (P90), the probability is clipped at the P90 contour. This represents the best available prediction of the probability of plume presence for I+10 at P90 probability level – i.e. the set of all lateral plume extents with more than 10% probability of occurrence. Contour interval is 10% with bold contours at P90, P50, and P10.

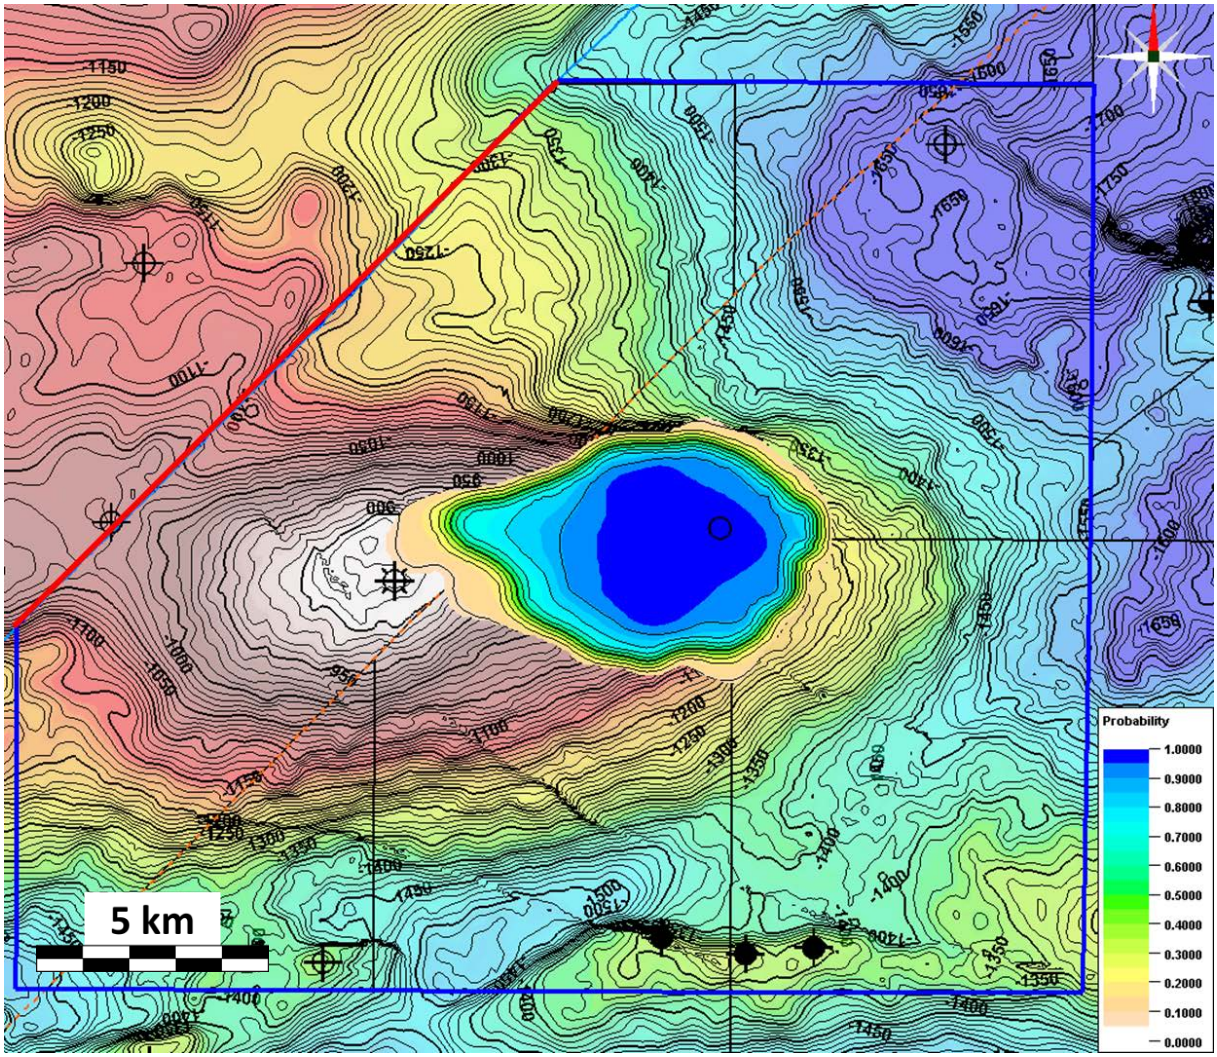

Figure 18: Smoothed P90 plume probability map for I+10

The probability map is smoothed and clipped at P90.

Time evolution of probability maps

The process of generating smoothed maps was repeated for all the key timesteps (Figure 19). The small inset at lower right is a stack of all the P90 contours from (I+10) to (I+300). This inset is enlarged in Figure 20.

Figure 20 is an overlay of all the P90 contours through time – i.e. the area swept out by the plume from commencement of injection to I+300 years. This illustrates that at no time do plume paths reach or exceed the structural limits of the storage formation, up to and including I+300.

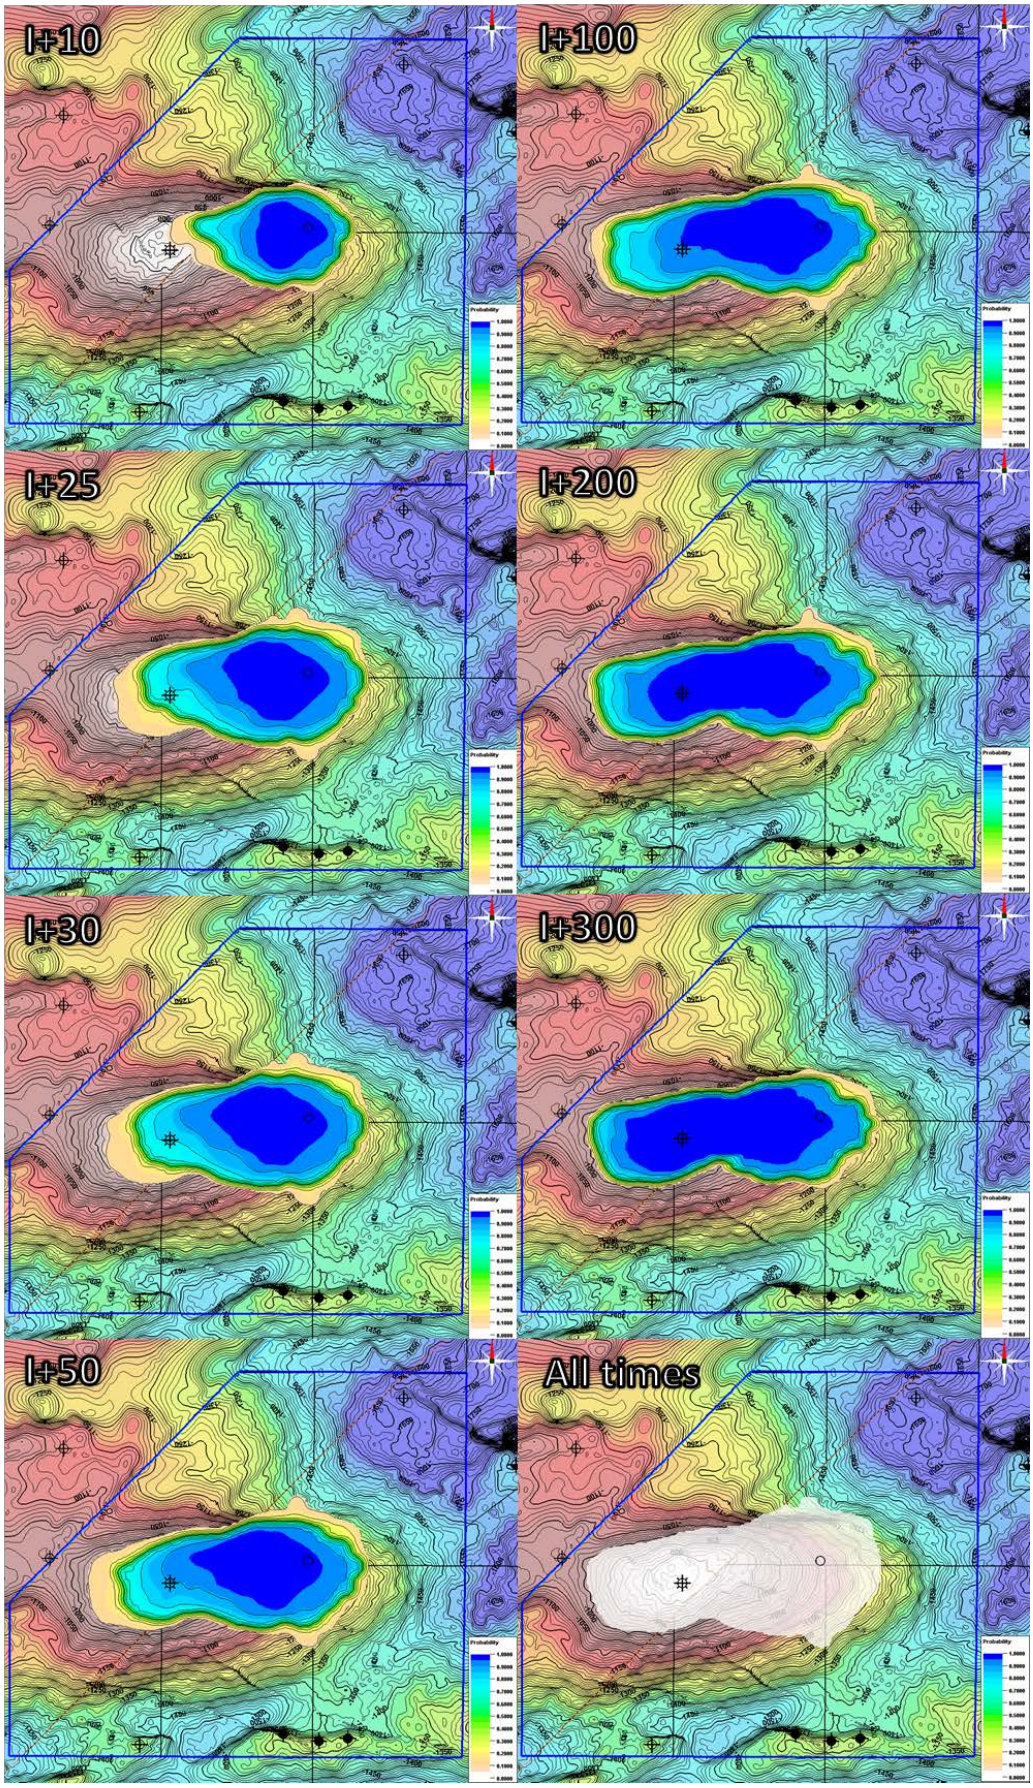

Figure 19: Smoothed P90 plume probability maps for timesteps I+10 to I+300

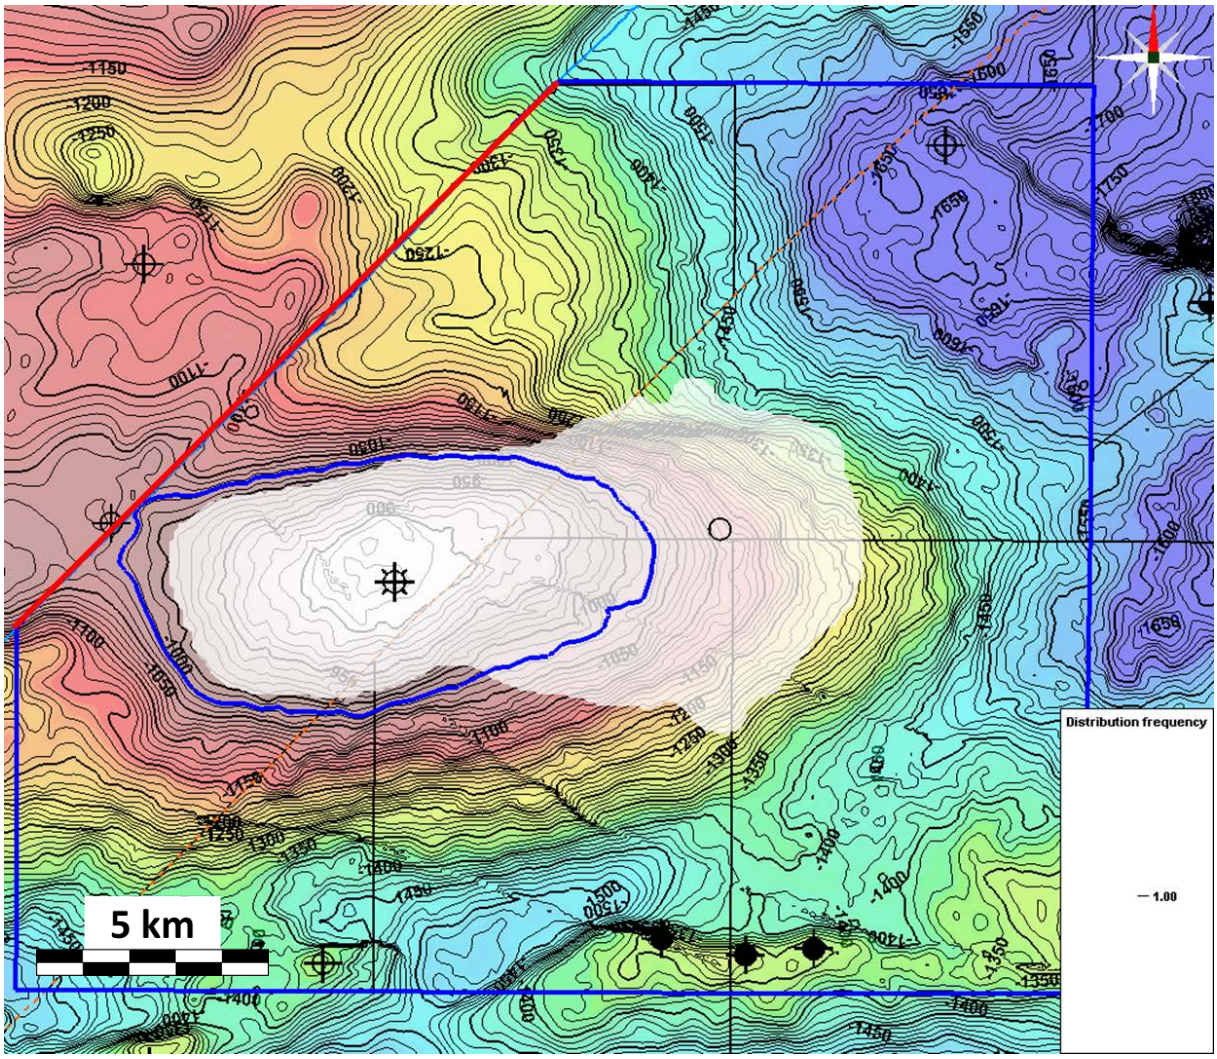

Figure 20: Set of all plume paths at P90 probability level for all timesteps

The white area is the compilation of all P90 maps from Figure 19. Structure contours are overlain, showing how the plume remains within the controlling structural closure as it moves updip.

3D extent of plume

A 3D probabilistic visualisation of the plume and its evolution through time is a little more complex to visualise, but is mathematically identical to the process described above. The 3D result is a Boolean measure of plume presence/absence for each cell of the 3D reservoir model, and summing those Boolean occurrences across all sensitivity and uncertainty cases determines a 3D plume probability volume.

The evolving plume path can then be visualised in arbitrary cross-sections of the 3D model chosen to illustrate key aspects of plume behaviour. Figures 21 to 24 show a conventional saturation plot of the evolving CO_2 plume evolving through key timesteps.

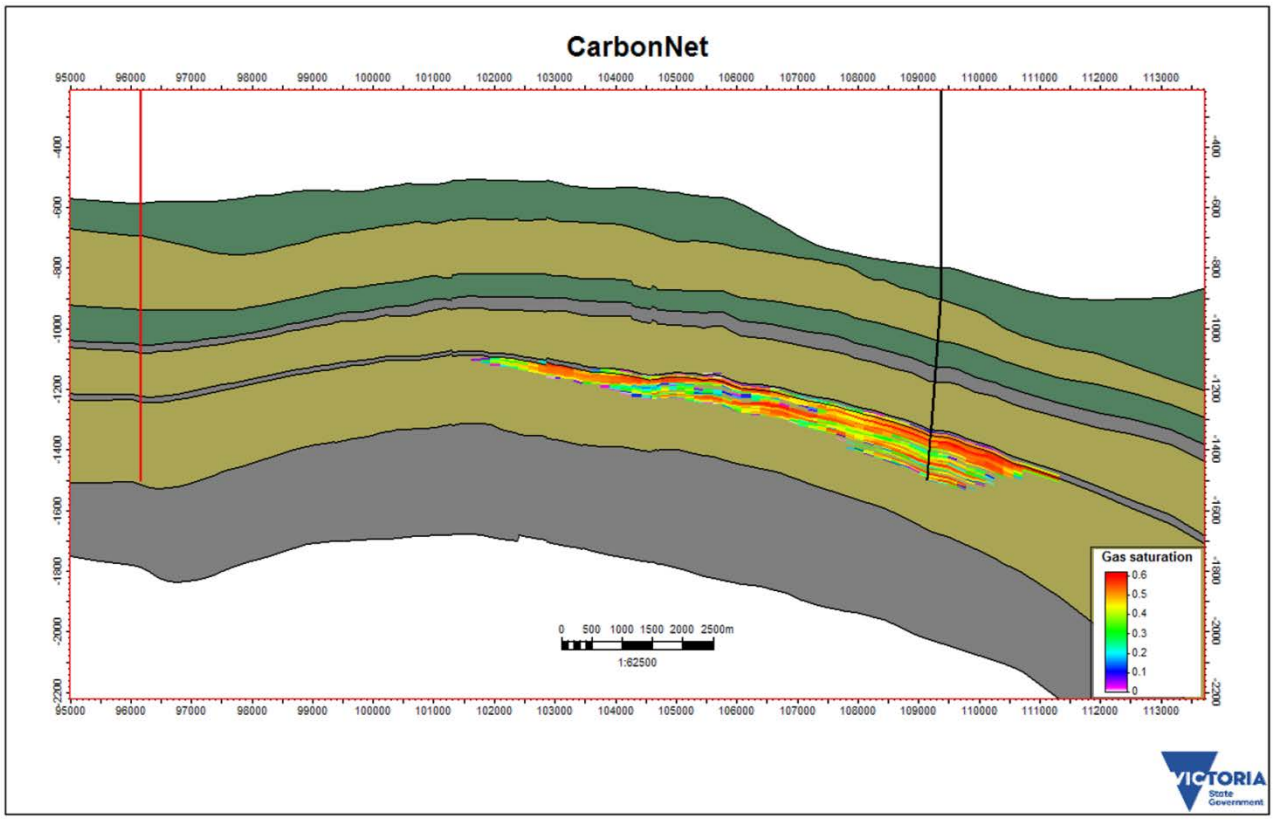

Figure 21: Plume evolution 25 years after start of injection in Base Case (I+25 = end injection)

Plume concentration as a colourmap 5 years after the end of injection, showing the plume moving away from the injection point and then migrating buoyantly updip towards the structural crest.

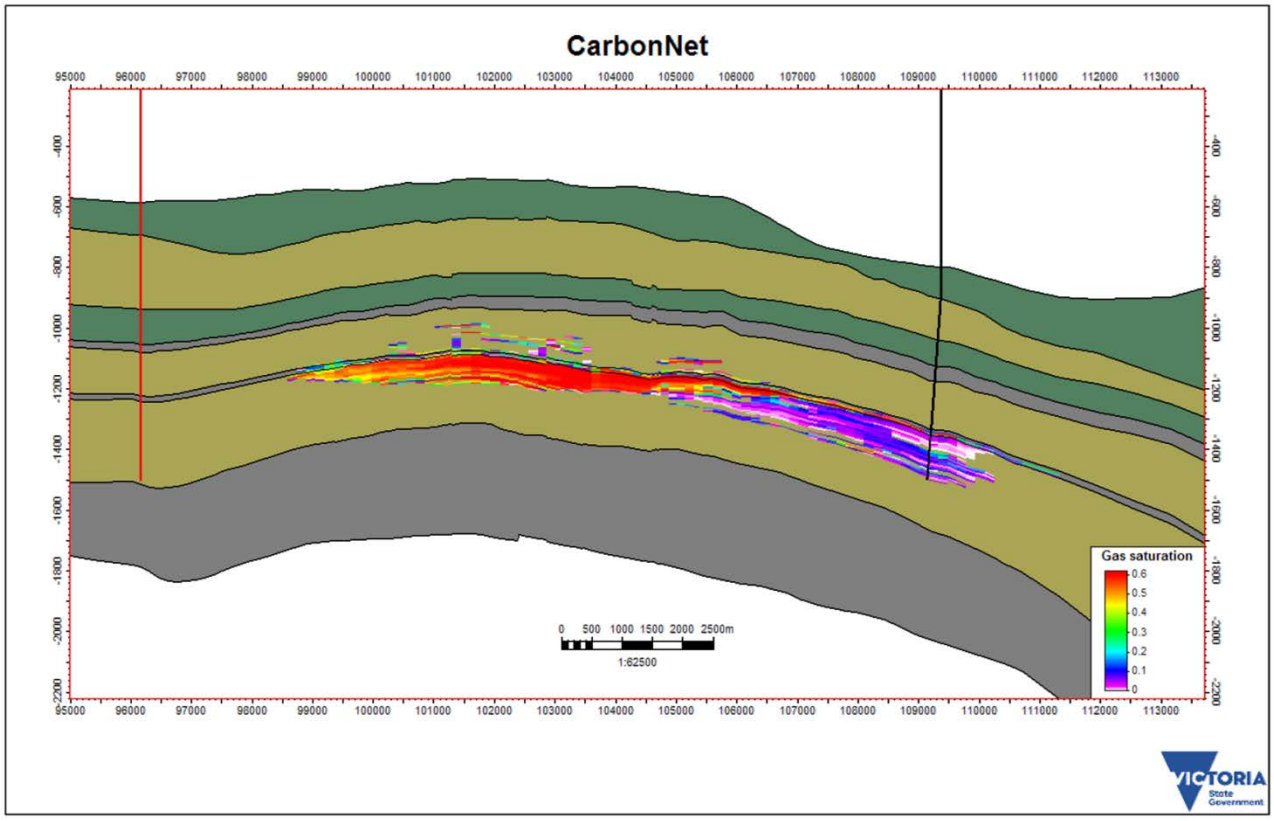

Figure 22: Plume evolution 100 years after start of injection in Base Case (I+100)

After 100 years the plume has migrated updip leaving residual saturation behind. The majority of the plume is concentrated in the structural closure, beneath the base Upper Halibut aquitard.

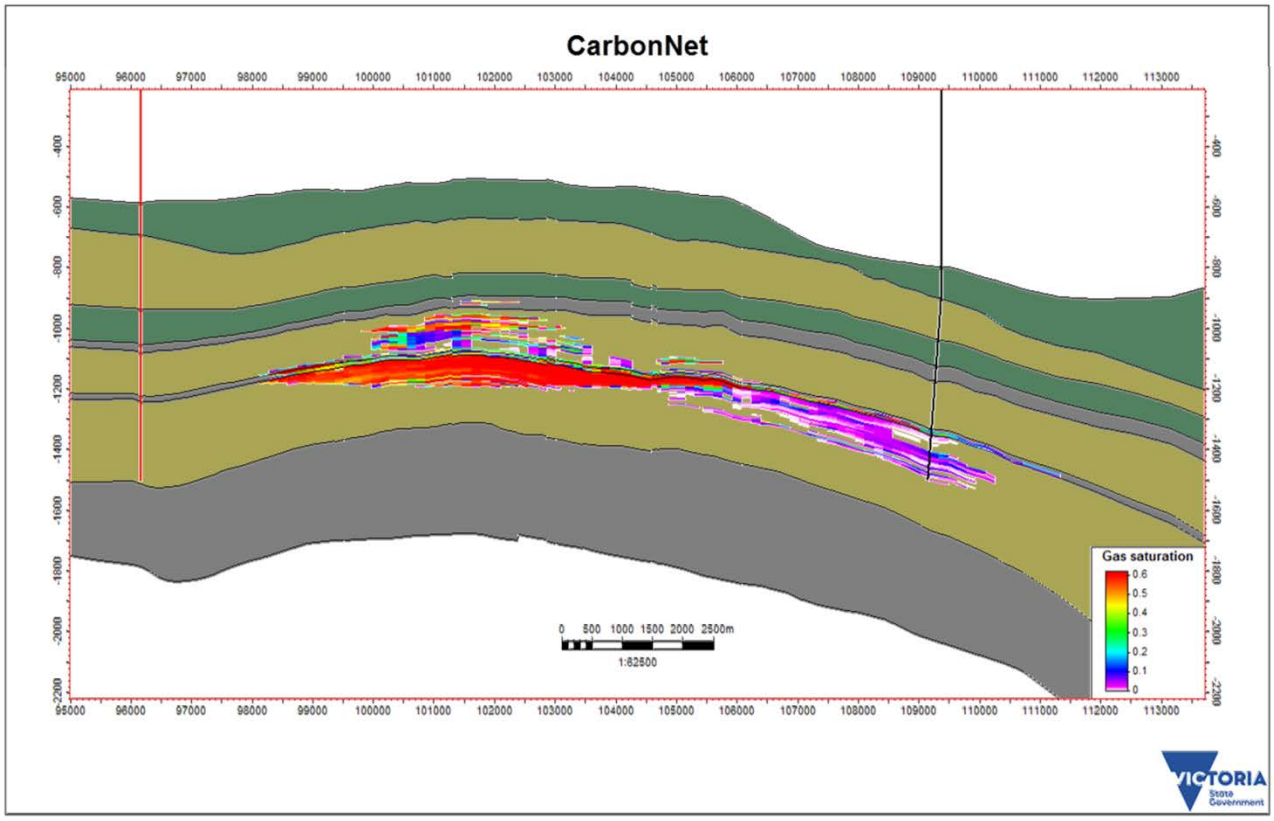

Figure 23: Plume evolution 200 years after start of injection in Base Case (I+200)

After 200 years, modest volumes of CO_2 have penetrated the base upper Halibut aquitard and are trapped at higher structural levels below T2B

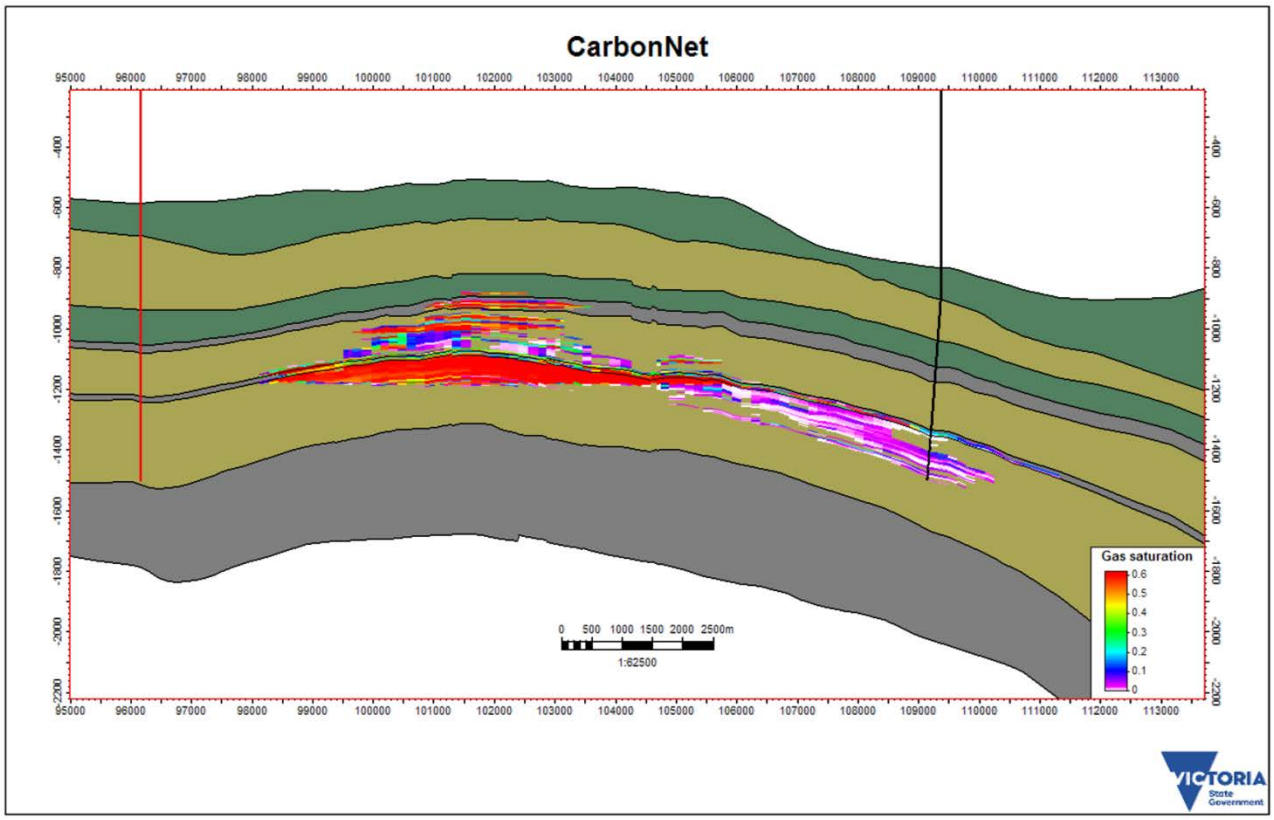

Figure 24: Plume evolution in Base Case at end simulation (, I+300)

After 300 years, modest volumes of CO_2 have penetrated the base upper Halibut aquitard and are trapped at higher structural levels below T2B and within basal T2B sand interbeds.

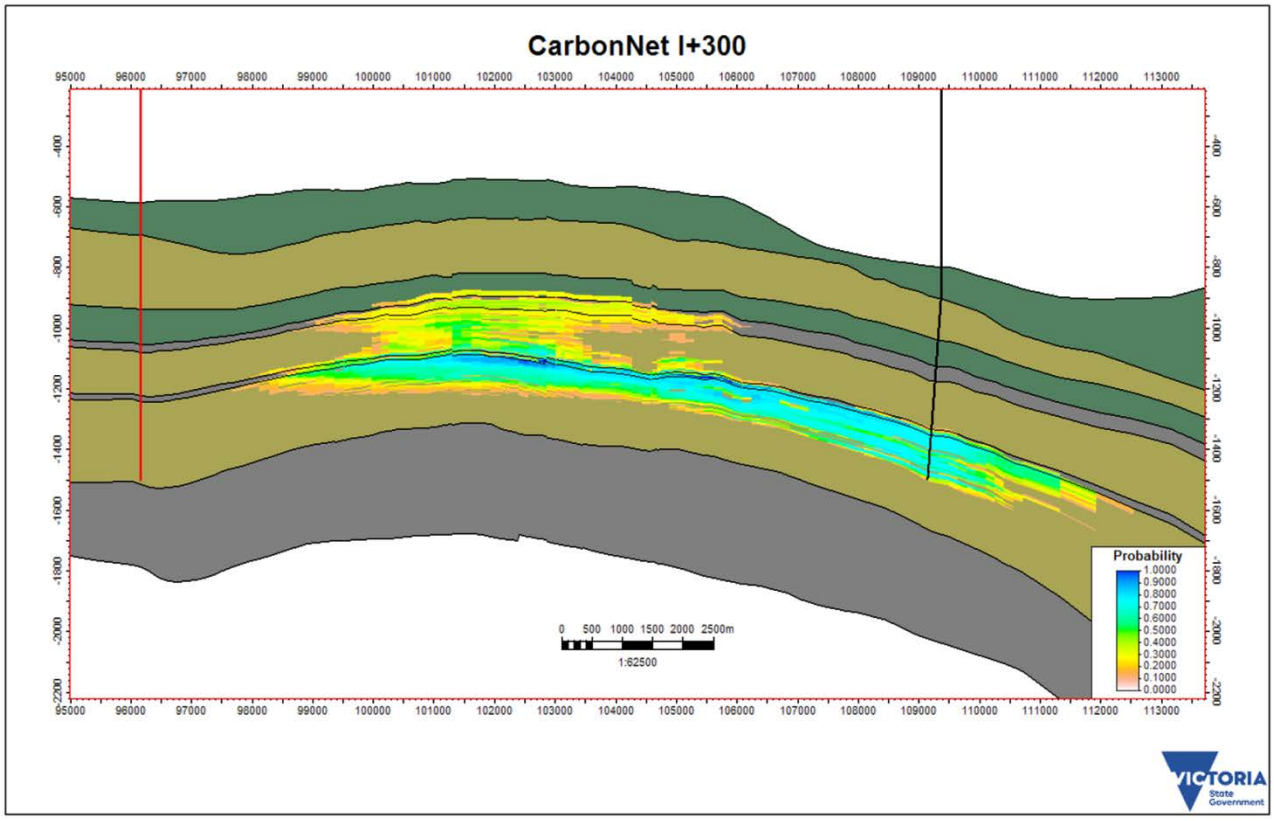

Figure 25: 3D plume probability cross section at I+300

A colourmap of plume probability from the set of all modelling runs showing high probability pathways in hotter colours and low probability paths around the fringes of the distribution.

The end-product of the 3D probability analysis is illustrated in Figure 25. This shows the probability of plume presence on the same cross section, at I+300. The high-probability pathways of plume movement are emphasised by the "colder" colours – typically above fault pathways through the base Upper Halibut aquitard. These zones are candidate sites for MMV – whether it be 3D seismic, other remote sensing, or well-based.

Comparison of 2D and 3D methods

An interesting subtlety arises when we compare the results of the 2D map and 3D voxel methods for spatial probability analysis. Figure 26 shows the probability map derived from 2D summation of projected maps of the plume extent and Figure 27 shows the result of a full 3D volume summation, and then projection to a map of the maximum value along any vertical section.

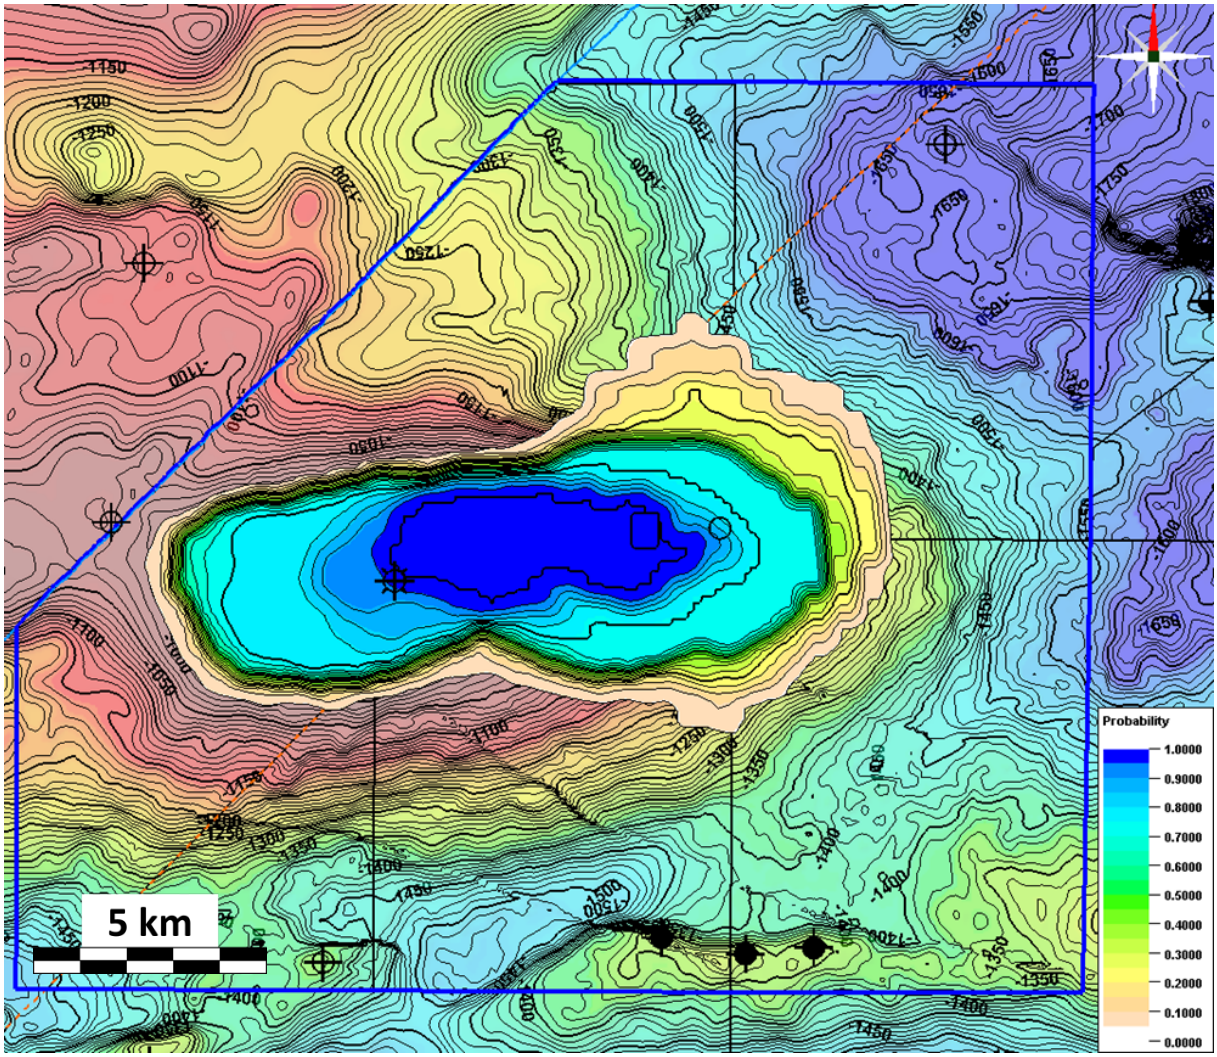

Figure 26: Plume probability map at I+300 from 2D map stacking

A colourmap of plume probability from the set of all modelling runs at I+300, produced by the 2D method

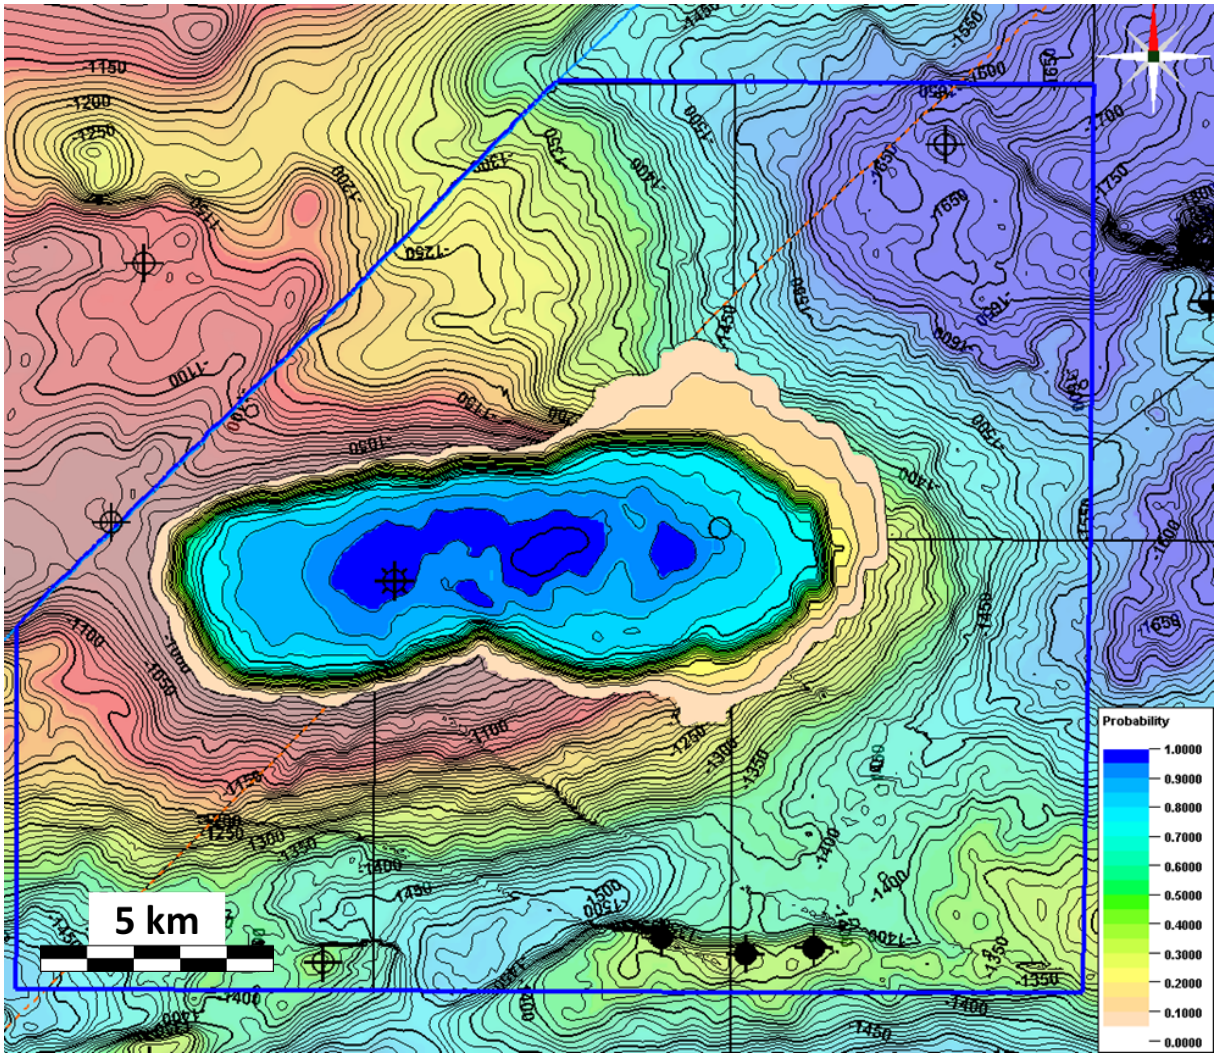

Figure 27: Plume probability map at I+300 from 3D voxel stacking

A colourmap of plume probability from the set of all modelling runs at I+300, produced by the 3D method

It is immediately obvious that the two maps differ. Let us consider why this is, and what implications it may have for MMV. The reason for the difference is that the 2D method identifies areas where the plume may exist, at any depth, and adds them up. The 3D method only adds up plume extents if they overlap in 3D space – i.e. they exist in the same Voxel. Figure 28 shows a cross-section of the model with the 3D and 2D extents of 100% certainty of encountering the plume.

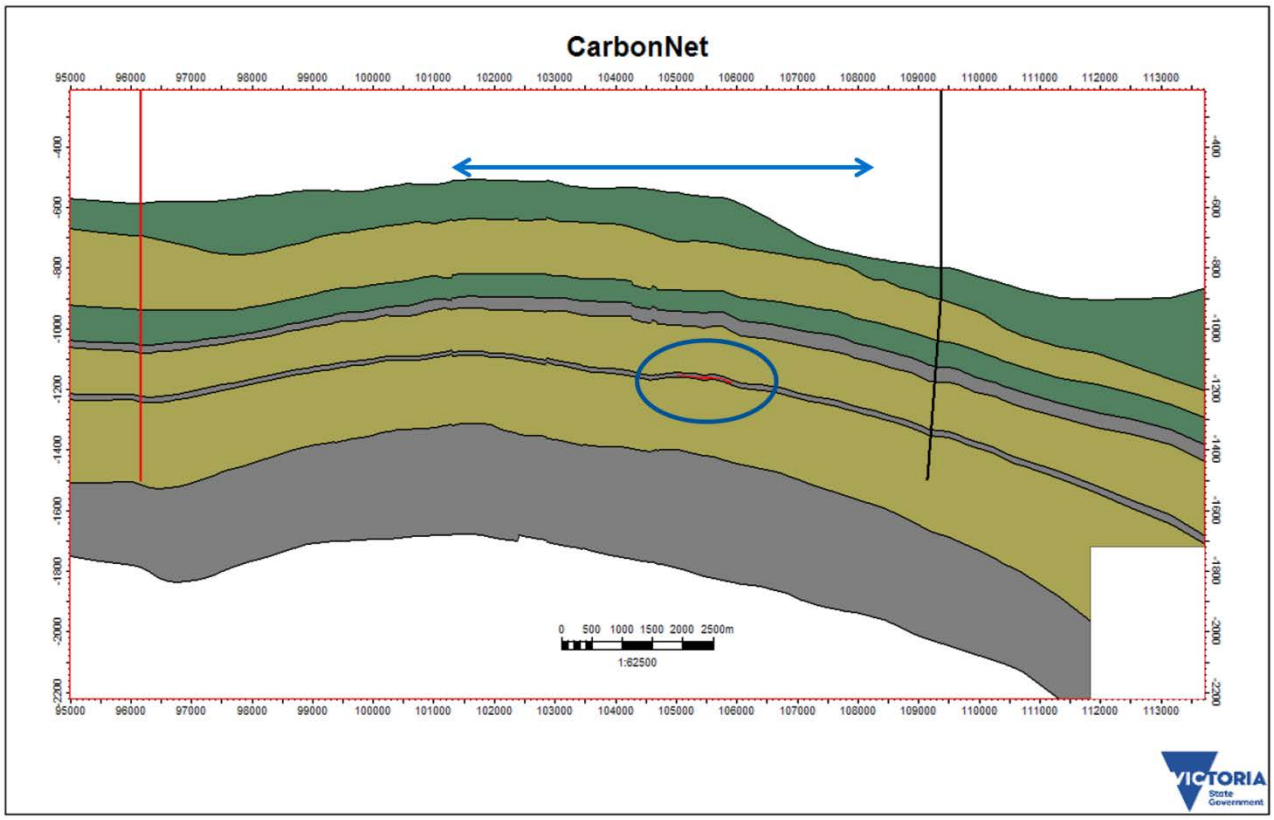

Figure 28: Zones of 100% plume probability at I+300 from 2D vs 3D method

In 2D map terms, the plume is certain to extend over the area between the blue arrows, although it may occur at different depths in different model runs. In 3D terms, the plume is only certain in a very limited spatial volume (red) highlighted by the blue ellipse.

This distinction between CO_2 at any depth and CO_2 at a chosen depth informs on the choice and deployment of MMV methods. Although Figure 26 indicates a wide area within which well-based MMV would be potentially capable of detecting the CO_2 plume, Figure 28 shows that unless other information is available about which reservoir level will contain the majority of the CO_2 plume, then choice of measurement depth is uncertain.

A volume-based MMV technique such as 4D timelapse seismic automatically covers all depths of interest. Providing that a seismic contrast is produced by CO_2 saturation, then a signal will be observed. In highly permeable reservoirs where the plume has high mobility, such as the CarbonNet sites, or Sleipner, then 3D seismic is likely to be the preferred monitoring tool due to its ability to sense at a wide range of depths and inform on the plume path in 3D, rather than having to de deployed redundantly at multiple levels in case of the plume taking certain paths. Volume-based MMV is more appropriate in cases where there is appreciable 3D uncertainty about flow paths. If only a single reservoir-seal pair exists, then well methods would be more reliable, but when multiple potential paths exist, then volume methods are preferred.

Comparison of 3D method to 1D Proxy

The same two objective measures - lateral approach to the western boundary and the vertical approach to top of the T2 seal can be measured from the 3D plume, at the P90 level. These can easily be compared to the Monte-Carlo Proxy prediction (Figures 13 and 14).

Table 10: Comparison of 3D cloud and Proxy prediction

Objective measure	3D cloud P90	1D Proxy / Monte-Carlo P90
Approach to Western Boundary	~900 m	~600 m
Approach to T2B Key Seal	~50 m	

In terms of vertical approach to the top of the T2B seal, the two methods produce a very similar answer, with the Proxy prediction suggesting a slightly closer approach. The difference is unlikely to be statistically significant. In terms of lateral approach, the proxy model is more conservative in terms of risk and predicts a closer approach. Any operational procedures and MMV plans based on the proxy prediction would be compatible with the 3D plume analysis.

We therefore conclude that although the Proxy formulation may be imperfect and limited to one-dimensional objective parameters, it is fit-for-purpose in the CarbonNet reservoirs. The full 3D plume method developed in this paper has been found to confirm Monte-Carlo forecast based on 1-D objective measures and Proxy equations, in this case. The same may not be true for other projects in other reservoirs, especially where non-structural storage is involved in more open saline aquifers.

Testing to failure

It is important to extend the modelling ranges far enough to include failure cases, since if failure does not occur, the limits are unknown. Only when failure does occur can the limits be quantified. By including a wide enough variation of each parameter, and allowing parameters to vary in a mutually reinforcing manner, the trap failure can be simulated and its consequences understood. For example, in our modelling a four-variable case with high Kh, High Kv/Kh, low porosity and low Sgr leads to vertical escape of CO_2 beyond the T2 topseal. This case has a nominal associated probability of 10^{-4}, so a first-pass assessment of the trap suggests 99.99% confidence in vertical containment.

However, in real geology, these factors are not entirely independent. There is a strong positive relationship between porosity and (log) permeability, so the combination of high permeability and low porosity is highly unlikely. In addition, the combination of high Kh and high Kv/Kh leads to unrealistic vertical permeabilities of 10 darcies for sands. The highest Kv actually measured in the basin is only 5 darcies (for a sand of 30% porosity).

An objective assessment of the probability associated with this failure case, given the coupling of real-world parameters, is more like 10^{-6}, giving a confidence in vertical containment of 99.9999%. This is a much higher confidence than the Monte Carlo simulation gives, showing the difference between “real” experimental runs and the approximation of the Proxy equation.

Discussion

The 2D and 3D Boolean probabilistic methods described here are a powerful way to visualise and quantify the expectation of future plume paths at given levels of probability. Note that the outcome of this probabilistic analysis will be quite different for different traps:

- Some traps will be better understood through local and regional data and hence have less uncertainty, leading to tighter probability contours
- Some geological contexts will offer more consistent reservoir and seal units and hence more consistent HFU's, leading to tighter probability contours
- Some trap types will inherently have a relatively tight distribution of probability contours while other types will inherently have looser contours

All factors being equal, a site with tighter probability contours offers more confidence in the forecast future distribution of CO_2 and is to be preferred for storage confidence, unless there are demonstrably no potential resource conflicts or environmental receptors within the extended probability envelope. This comparison must take account of the volume of CO_2 stored and the physical scale of the trap, as key elements of the "equal factors".

Even if no potential conflicts exist, for regional storage efficiency a tightly-clustered probability envelope is preferred so that multiple storage sites can coexist without interference. This requirement may be relaxed if the overall constraints on storage are non-spatial – for example, the mutual pressure effect of multiple storages may define the ultimate limit of site spacing, rather than physical interference of their plumes in a pressure-balanced aquifer.

MMV methods need to take account of the 3D probability forecasts and ensure that point-based measurements are placed in high probability zones, or in zones where they would inform on potentially high risk outcomes. Volume-based measurements need to cover a wide enough area of the probability cloud, again, to inform on lower probability outcomes such as lateral or vertical movement exceeding the P90 forecast.

Structural vs Saline Aquifer traps

For a structural trap, it is relatively easy to demonstrate, through well, seismic, or other remote monitoring, that the CO_2 is contained within an overall structural trap, and is accumulating at the crest in a buoyant fashion. The structural trap can be shown to be secure and effective so long as:

- The topseal has sufficient sealing capacity and is geographically extensive.
- No leaks are observed during operation.
- No adverse geochemical reactions are predicted or observed.

Even if slow movement of CO_2 is still taking place as the plume adjusts to its structural confinement, the plume will settle into the crestal area with slowly increasing CO_2 saturation and storage efficiency, expelling excess formation water downward.

In an aquifer trap, it is more difficult to demonstrate the "90% of plume paths" of Australian legislation, since many variables can have a distinct effect, especially at long times, on the actual plume path. A reference here is the behaviour of the Sleipner plume (Chadwick and Noy, 2010, 2015) which moved in directions and to distances that were initially unexpected, while remaining confined below the effective topseal. Only after many timelapse snapshots could the permeability model be calibrated sufficiently to reproduce the essentials of plume shape in forward models.

Non-structural traps are subject to greater uncertainty of plume movement, since the plume is not focussed into a strong structural closure. In saline aquifer traps, the plume is predicted to migrate steadily up-dip on a gentle slope before it slows down and stops due to a diminution of buoyancy forces as residual trapping and dissolution reduce the volume of available mobile CO_2. In practice, the plume will interact with small undulations in the top reservoir/base seal, with cross-cutting faults, and with stratigraphic variations in porosity and permeability. Key determinants of plume direction and distance include:

- Variations in residual CO_2 saturation (Sgr) will strongly affect the distance the plume travels before it is consumed by a combination of residual trapping and dissolution, precursors to geologic time scale mineralisation.
- Details of depth conversion, especially lateral velocity variation, will affect whether the plume swings slightly left or right and, in time, may influence it to enter an entirely different migration chain of connected swells and small culminations.
- The presence of faults may deflect or channel the plume.
- Stratigraphic variations (seal and reservoir fairways or alignments of permeability) can strongly steer a plume in the absence of a strong updip direction.

Summary

In conclusion, the novel 3D probabilistic approach developed in this paper extends existing dynamic modelling sensitivity and uncertainty methodology and the use of estimates of objective parameters based on proxy models and offers a new method to quantify plume path uncertainty in 2D terms of lateral extent (i.e. map view) and vertical extent (cross-section), derived from a full 3D understanding of plume containment with an appropriate high level of regulatory and public confidence.

The modelled storage can thus be tested against variation of a wide range of subsurface parameters and two types of result can be derived:

1) 2D maps and cross-sections, and 3D volumes of plume probability at future times, demonstrating containment at the required confidence level.
2) Theoretical tests of the “failure” of the trap – i.e. how far do constraints have to be relaxed before the plume does exceed the chosen boundary – and how (un)likely is this given the extent of geological knowledge?

This second aspect is addressed by identifying any “failure” cases, and assessing their nominal probability as calculated in Table 8, then reviewing the model input data to confirm whether this is statistically valid and that dependencies between parameters have been properly assessed.

Probability modelling in 3D space allows confirmation that less precise proxy methods do or do not provide an adequate forecast of plume behaviour at the desired probability level. In the CarbonNet case, at P90 level, the Proxy method is sufficient.

However, the insights provided from the 3D plume analysis method significantly inform the selection of practical MMV methods and locations in 3D(X, Y, and Z) space that are flexible enough to detect the full range of potential plume variation and will avoid stranded assets being placed in locations where the plume cannot be sensed.

Note that in storage concepts with a lower structural focusing effect (less buoyancy, lower permeability, less structural dip) and especially in saline aquifer traps, and in storage sites with multiple potential reservoir intervals (as seen in this study), the 3D plume probability method is recommended over a simpler Proxy approach since the plume shape may be highly variable in all three spatial dimensions and selection of a useful set of 1-dimensional objective measures may be difficult.

References

Alessio, L.D., Coca, S., Bourdon, L.M., 2005. Experimental Design as a Framework for Multiple Realisation History Matching: F6 Further Development Studies. APOGCE, Society of Petroleum Engineers paper 93164.

CarbonNet, 2015a. The CarbonNet Project: a historical perspective http://hub.globalccsinstitute.com/sites/default/files/publications/155928/carbonnet-project-historical-perspective.pdf accessed 26/04/2016

CarbonNet, 2015b. CarbonNet storage site selection & certification: challenges and successes. http://hub.globalccsinstitute.com/sites/default/files/publications/196158/carbonnet-storage-site-selection-certification-challenges-successes.pdf accessed 26/04/2016

CarbonNet, 2015c. Site characterisation for carbon storage in the near shore Gippsland Basin. http://hub.globalccsinstitute.com/sites/default/files/publications/196108/carbonnet-project-site-characterisation-carbon-storage-near-shore-gippsland-basin.pdf accessed 26/04/2016

CarbonNet, 2015d. 3D mapping and correlation of intraformational seals within the Latrobe Group in the nearshore Gippsland Basin. http://hub.globalccsinstitute.com/sites/default/files/publications/196708/carbonnet-project-3d-mapping-correlation-intraformational-seals-latrobe-group-nearshore-gippsland-basin.pdf accessed 26/04/2016

Chadwick, R.A., and Noy, D.J., 2010. History – matching flow simulations and time-lapse seismic data from the Sleipner CO_2 plume. In: Vining, B.A., Pickering, S.C. (Eds.), Petroleum Geology: From Mature Basins to New Frontiers– Proceedings of the 7th Petroleum Geology Conference. Geological Society, London, pp. 1171–1182.

Chadwick, R.A., and Noy, D.J., 2015. Underground CO_2 storage: demonstrating regulatory conformance by convergence of history-matched modelled and observed CO_2 plume behaviour using the Sleipner time-lapse seismics. Greenhouse Gases: Science and Technology 5, 1–17.

Hoffman, N and Preston, J.C., 2014 Geochemical interpretation of partially filled hydrocarbon traps in the nearshore Gippsland Basin, APPEA Journal 2014.

Hoffman, N., Arian, N., and Carman, G. 2012. Detailed seal studies for CO_2 storage in the Gippsland Basin, September 2012. EABS IV Conference

Hoffman, N., Carman, G., Bagheri, M., and Goebel, T. 2015a Site characterisation for carbon sequestration in the near shore Gippsland Basin. Proceedings, 14th Annual CCUS Conference, Pittsburgh

Hofmann, H., and Cartwright, I., 2013 Using hydrogeochemistry to understand inter-aquifer mixing in the on-shore part of the Gippsland Basin, southeast Australia. Applied Geochemistry 33 (2013) 84–103

Kuttan, K., Kulla, J.B. and Newman, R.G., 1986—Freshwater influx in the Gippsland Basin: impact on formation evaluation, hydrocarbon volumes and hydrocarbon migration. APPEA Journal, 26 (1), 242–249.

Ma and La Pointe, 2011. Uncertainty Analysis and Reservoir Modelling. 314pp AAPG Memoir 96. ISBN-13: 978-0-89181-378-1

Newendorp, 1976. Decision Analysis for Petroleum Exploration. 668pp Penwell Corp. ISBN-13: 978-0878140640

Schuyler and Newendorp, 2013. Decision Analysis for Petroleum Exploration. 588pp Planning Press. ISBN-13: 978-0966440140

Varma, S., and Michael, K., 2012 - Impact of multi-purpose aquifer utilisation on a variable-density groundwater flow system in the Gippsland Basin, Australia. Hydrogeology Journal (2012) 20:119-134