Check for updates

Modelling the FluidFlower: Insights from Characterisation and Numerical Predictions

Christopher Green[1] · Samuel J. Jackson[1] · James Gunning[1] · Andy Wilkins[2] · Jonathan Ennis-King[1]

Abstract
The FluidFlower International Benchmark Study is a unique chance to assess the uncertainties introduced by numerical modelling through comparison to high-quality experimental data. A complete description of how the experiment was modelled by CSIRO is presented here, from the simplified model of the tracer tests and subsequent inversion of the permeability data using the spatial tracer concentration, the rapid development of an open-source finite volume simulator capability, through to the final numerical predictions and the reporting of key metrics.

List of Symbols

β Fluid phase (−)
c_i Tracer concentration (−)
c_t Total (rock plus fluid) compressibility (Pa^{-1})
d Molecular diffusion coefficient (m^2 s^{-1})
I_i Tracer intensity (−)
$\mathbf{g}$ Gravity (m s^{-2})
γ Surface tension (N m^{-1})
h Thickness (m)

✉ Christopher Green
 chris.green@csiro.au

 Samuel J. Jackson
 samuel.jackson@csiro.au

 James Gunning
 james.gunning@csiro.au

 Andy Wilkins
 andrew.wilkins@csiro.au

 Jonathan Ennis-King
 jonathan.ennis-king@csiro.au

[1]   CSIRO Energy, Clayton North 3169, VIC, Australia

[2]   CSIRO Mineral Resources, Queensland Centre for Advanced Technologies, Kenmore 4069, QLD, Australia

κ Fluid component (−)
k   Absolute permeability (m^2)
k_r Relative permeability (−)
μ Dynamic viscosity (Pa s)
m_l Permeability modifier (−)
P Pressure (Pa)
P_c Capillary pressure (Pa)
ϕ Porosity (−)
Q Volumetric injection rate (m^3 s^{-1})
ρ Density (kg m^{-3})
S Saturation (−)
t Time (s)
χ Mass fraction (−)
z Total mass fraction summed over all phases (−)

1 Introduction

Robust predictions derived from numerical modelling of underground CO_2 storage in porous media rely on suitable characterisation of the reservoir, ensuring that all of the important physical processes involved are captured, and finally, accurate numerical solutions of the governing physics. Each of these essential steps can be subject to uncertainties, and often require decisions to be made to simplify the analysis in order to make the modelling tractable. The basis of these simplifying assumptions is often a judgment of the user, based on experience, or simply due to a lack of reliable data for key properties. In practice, it is usually difficult to assess the reliability of numerical simulation predictions due to uncertainties in the geological models and the difficulty in obtaining high-quality experimental evidence.

Several previous studies have considered well defined models in order to compare the accuracy of numerical simulators (Pruess et al. 2004; Class et al. 2009). Even in these cases, where there is no uncertainty in the model, results between groups can differ due to choices made by each participant, even when using the same code. These uncertainties were considered in a subsequent benchmark study on CO_2 storage (Nordbotten et al. 2012), where variability between participants was categorised into several sources: the interpretation of the problem, the physical processes included, how properties are modelled through different length scales, and the numerical modelling details.

The FluidFlower international benchmark study (Nordbotten et al. 2022) provides a novel opportunity to compare the impact of these different modelling choices on the robustness of numerical predictions of a carefully controlled flow experiment that includes all of the main physical processes that may be present in underground CO_2 storage operations: buoyancy, structural capillary trapping, residual trapping, dissolution and convection. In particular, numerical results will be subject to several choices by each participant. These include how the experimental apparatus is modelled, how the tracer test data is used to inform the *ex-situ* petrophysical properties provided, and how the physical processes included in the numerical simulators are chosen. With the experimental results providing a well-controlled "ground truth", uncertainties in these modelling choices can be assessed on their impact on the final predictions.

In this manuscript, we detail how the modelling was approached in order to model the FluidFlower experiment. We simplify the tracer tests to a steady-state Darcy flow to allow rapid inversion of the data using the high-quality spatial concentration data. We then detail the rapid development of a finite volume simulator suitable for this modelling exercise, before presenting our "best case" numerical predictions of the experiment. Finally, we discuss how the sparse data statistics required by the organisers was calculated using a small set of ensemble models.

2 Model Characterisation

Participants were provided with high-resolution curvature-corrected images of the Fluid-Flower apparatus, along with average multiphase flow properties for each of the sand types, see (Nordbotten et al. 2022) for details. Images and pressure data from tracer tests were also provided, to enable further characterisation of the model properties.

The first step in this modelling study was to produce a digital mesh suitable for numerical computations from the images provided by the organisers. Initial attempts to segment the image into distinct sand types automatically through a thresholding process were unsatisfactory, due to the poor contrast between layers in some sections (particularly where the lighting in the photographs was dim), despite attempts to enhance contrast. A semi-automated picking routine was used to delineate sedimentary horizons and faults in the model, with the results passed to some code to create distinct zones from the horizons, see Fig. 1

It was decided at the start of this analysis to undertake all modelling in two dimensions only, due to the relative length scales of the experimental apparatus (Nordbotten et al. 2022). The FluidFlower thickness (nominal $\approx$ 19 mm) was provided at regular points in the vertical plane (Nordbotten et al. 2022). Bilinear interpolation was used to provide thicknesses at each point in the model, and this thickness was used to provide spatially-varying thickness-averaged petrophysical properties (porosity and permeability) of the form $u(x, y) = u_0 h(x, y)/h_0$ where u_0 is the property (porosity ϕ or permeability k), $h(x, y)$ is the linearly-interpolated thickness at any spatial point (x, y), and
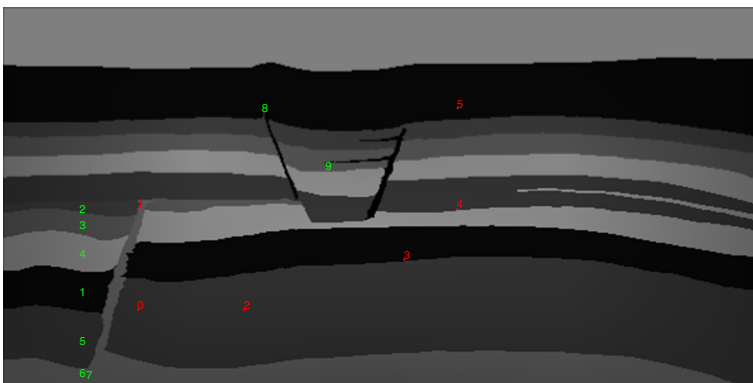

Fig. 1 Digitised and interpreted facies-label (green labels) image used as the basis for the parameterisation. The labels correspond to sand type {ESF (1); C (2); D (3); E (4); F (5); G (6); Fault 1 (7); Fault 2 (8); Fault 3 (9)}. The red labelled ports 0, 1,... 5 are used in the inversion, corresponding to ports {5_3(0), 5_7(1), 9_3(2), 15_5(3), 17_7(4), 17_11(5)}, see (Nordbotten et al. 2022) for full details

$h_0 = 0.019$ m is the nominal thickness, from the single porosity and permeability values measured for each sand facies by the experimental team (Nordbotten et al. 2022).

The initial characterisation of the model by facies was refined by inversion using data from the tracer tests (Nordbotten et al. 2022). A detailed description of the inversion methodology used is provided in Tian et al. (2023), so only a brief description of this characterisation is presented here.

Two types of data from the tracer tests were provided; pressure measurements at each monitoring port for the duration of the tracer tests, as well as digital images of the injected tracer (Nordbotten et al. 2022). To assist in inversion, the tracer images were processed to produce difference images (using the background image before the tracer injections as the reference). Constant attenuation scaling was used to convert the image intensity to tracer concentration

$$c_i = \rho_t I_i, \tag{1}$$

where c_i is the tracer concentration (−) in the i^{th} pixel, $\rho_t = 1002$ kg.m^{-3} is the tracer density, and I_i is the (scaled) intensity (−) of the i^{th} pixel. Mass conservation was used to scale the concentrations to ensure a global mass balance using the (known) injection volume

$$V = Qt = \rho_t \sum_i \phi_i v_i I_i, \tag{2}$$

where V is the injected tracer volume (m^3), Q the (constant) volumetric injection rate (m^3. s^{-1}), t is time (s), ϕ_i the porosity of pixel i, and v_i its volume (area multiplied by interpolated thickness). This process was found to be reliable, with only small mass balance errors in a few individual pixels. The resulting concentration maps for each stage of the tracer tests are presented in Fig. 2.
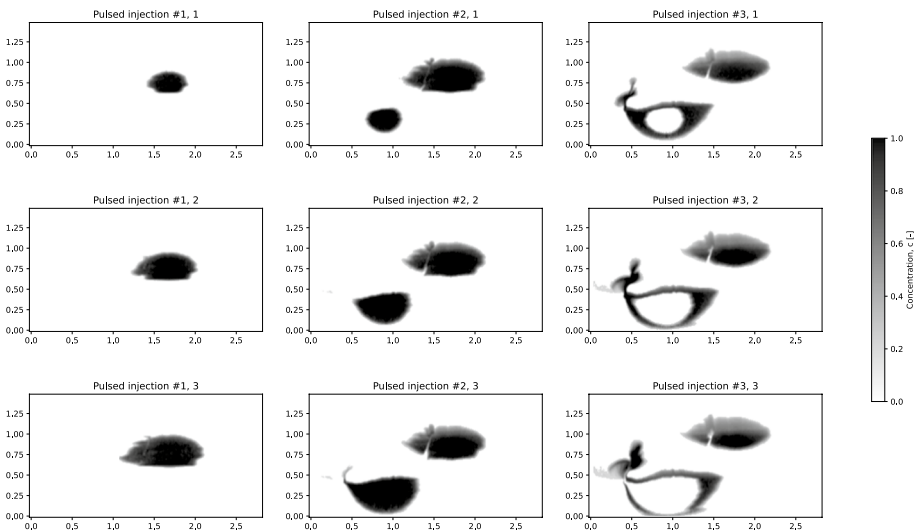

Fig. 2 Concentration of tracer computed from provided well test images. Each column represents one stage of the three-stage tracer test, and each row shows the evolution of the tracer plume during each stage, see (Nordbotten et al. 2022) for full details

An obvious feature of the pressure data is the instant increase from ambient pressure once each injection begins (and likewise, an instant drop back to ambient pressure once injection ceases), suggesting rapid pressure communication throughout the experimental apparatus. The characteristic time scale for pressure diffusion $t = \frac{\phi \mu c_t L^2}{k_{sc}}$, for the Fluid-Flower is very short, typically 5 ms using suitable values for $\phi = 0.45$, $\mu_w = 10^{-3}$ Pa.s, $c_t \approx 0.73 \times 10^{-9}$ Pa^{-1}, $L = 1$ m, $k_{sc} = 6.9874 \times 10^{-10}$ m^2. All pressures are rescaled to dimensionless pressure using a scaling $P_{sc} = Q_{sc} \mu_w / (k_{sc} h_0)$, with additional chosen constants $h_0 = 0.025$ m (average cell thickness), $Q_{sc} = 6.25 \times 10^{-7}$ m^3/s (the water injection rate) yielding $P_{sc} = 35.778$ Pa.

For this reason, the forward modelling was approximated as steady-state on the time scale of the experiment, with the pressure corresponding to the steady state solution of the thickness–averaged 2D pressure diffusion equation for single-phase flow

$$c_t \phi h \frac{\partial P}{\partial t} - \nabla.(k_w/\mu_w)h(x,y)\nabla P = Q(x,y,t) \tag{3}$$

where P is the pressure (Pa), h the FluidFlower thickness (m), ϕ the porosity, k the absolute permeability of the media (for water) (m^2), μ_w the water viscosity (Pa.s), and c_t the total (rock plus water) compressibility (Pa^{-1}).

The injection rates are held relatively constant over each 30 min interval, so the overwhelming majority of the data are collected with the system equilibrated in terms of pressure given the rapid pressure communication. We therefore proceed on the basis that time-stepping the pressure diffusion equation was unnecessary, and solve the steady-state Poisson equation with Q at two different injection ports (ports 9_3, 17_7), corresponding to the tracer experiment. The resulting fixed pressure fields and velocities are then constant over each injection period, which is sufficient to compute the advection of the tracers.

The numerical model was constructed on a structured 5 mm grid ($N_x = 568$, $N_y = 300$), with the upper boundary fixed at atmospheric pressure, and no–flow boundaries on the left, right, and base of the model.

Tracers were advected along streamlines using the fixed velocities computed from solving Eq. (3) and the upwinding scheme described in Koren (1993) for the duration of each tracer pulse. The time-stepping in the tracer computation was adjusted so the number of time steps divided the total tracer time exactly, which has the merit that any computed quantities from the tracer image are smooth differentiable functions of the parameters in the governing equation.

The cross-port pressure data provided during the tracer experiment was conspicuously noisy and clearly close to the noise floor of the instruments. Significant drift was evident in these measurements, and even the stable values showed the curious property of the amplitude not diminishing in a consistent way with the distance from the injection port. Further, the gauges were located some distance (20 cm or more) from the actual injection face and so measurements are subject to unknown frictional and other losses in the feed plumbing. This has implications in the inversion if this data is weighted very heavily, and thus limits their usefulness for inversion. By contrast, the tracer images were very clearly interpretable, rich in spatial content, and not obviously contaminated by experimental artefact of any significant kind. For these reasons, inversion proceeded using only the spatial concentration data. We consider the effect of including cross-well pressure data weighted accordingly in the inversion, see Appendix A for details.

The FluidFlower is initially filled with blue tracer. Each injection period was modelled using $Q = 2250$ml/h (6.25×10^{-7} m^3/s) for a 3×30 minute period at port 17_7 with clear

tracer, then 3×30 minutes at port 9_3 with clear tracer, then 3×30 minutes at port 9_3 with blue tracer (Nordbotten et al. 2022). The forward modelling operation, which implements standard tracer advection under an upwinding scheme, is expressed below as a function $f_t(\boldsymbol{m})$ which generates concentration profiles which are very close to unity inside the swept region, and fall rapidly to zero at the tracer front. The tracer advection is stepped forward for precisely the number of time steps needed for the injection, and numerical integrations of the total tracer mass over the modelling grid at the end of the simulation agree very closely with the mass known to be injected from Q in the tracer source.

Under the assumptions of the single–phase PDE and the fast equilibration time, the experimental 30 min wait times between injections does not need to be modelled, as nothing happens in the Poisson model if the sources is switched off, since the velocities are then instantly zero. The modelled tracer positions at the end of each 30 min injection period are compared to the experimental concentration data for inversion.

The inversion of this data was couched as a Bayesian inverse problem with a likelihood $P(\boldsymbol{y}|\boldsymbol{m})$ formed as a joint probability using (in general) pressure and tracer data $\boldsymbol{y} = \{\mathbf{P}_{\mathrm{obs}}, \mathbf{c}_{\mathrm{tracer}}\}$. The model was taken to be multiplier modifiers m_l ($l \in \mathcal{L} \equiv \{1, 2 \ldots 7, 9\}$) of the permeability parameters, per facies, and applied in a "paint-by-numbers" fashion over the labelled facies model. The thickness and porosity data were considered to be sufficiently precise and experimentally stable to be fixed for the purposes of model prediction and inversion. The Bayesian framework was completed with the provision of a weak prior problem of the modifier parameters, of Gaussian form $P(m_l) \sim N(1, \sigma^2)$ with $\sigma = 5$ for each parameter. The associated prior covariance is $C_p = \mathrm{diag}\{\sigma^2\}$. The model point estimate at the global maximum of the posterior probability is referred to as the MAP (maximum aposteriori) inversion.

The inversion is performed using a Levenberg–Marquardt routine (Nocedal and Wright 1999; Madsen et al. 2004), which requires the Jacobian J of the forward response with respect to the unknown model parameters. Since the model dimensionality was very low and the forward model speed very fast (measured in seconds), this was computed using simple forward differences.

The negative log posterior $E(\mathbf{m}) \sim -\log(P(\boldsymbol{y}|\mathbf{m})P(\mathbf{m}))$ used in the optimisation step was written as a standard l_2 misfit energy

$$E(\mathbf{m}) = E_{\mathrm{pressure}}(\mathbf{m}) + E_{\mathrm{tracer}}(\mathbf{m}) + E_p(\mathbf{m})$$

where the cross–port pressure misfit, accumulated over only stable average measurements at ports p5.3_1, p5.7_1, p9.3_1, p15.5_1, p17.7_1, p17.11_1 is

$$E_{\mathrm{pressure}}(\mathbf{m}) = \tfrac{1}{2}\lambda_p \|\mathbf{P}_{\mathrm{obs}} - f_p(\mathbf{m})\|_2^2$$

and the tracer image mismatch is written as

$$E_{\mathrm{tracer}}(\mathbf{m}) = \tfrac{1}{2}\lambda_t \|\mathbf{c}_{\mathrm{tracer}} - \mathbf{f}_t(\mathbf{m})\|_2^2$$

The weights λ_p, λ_t were adjusted so the tracer data are dominant in the likelihood as this data is much more abundant and artefact–free. The prior Bayesian term amounts to

$$E_p(\mathbf{m}) = -\log(P(\mathbf{m})) \sim \tfrac{1}{2}\sum_{l \in \mathcal{L}} (m_l - 1)^2 / \sigma_p^2$$

and has a very benign influence on the inversion. In dimensionless units, the pressure data $\mathbf{P}_{\mathrm{obs}}$ are $O(1)$ numbers, but rather noisy, so setting $\lambda_p = 1$ seems appropriate. The tracer data $\mathbf{c}$ are processed from the digital images to have concentration values ranging over $0 < y < 1$. Since the associated l_2 norm has a very large number of voxels, λ_t is scaled such that the tracer misfit energy is $E_{\mathrm{tracer}}(\mathbf{m}) = 1000/2$ for a model that produces no tracer concentration ($\mathbf{f}_t(\mathbf{m}) = \mathbf{0}$), i.e. the information content is equivalent to 1000 measurements. In practice since the volume in which experimental and forward–modelled concentrations differ is only a small fraction of the image, the misfit energy from this term ends up being $O(10)$, perhaps equivalent to putting a 10-fold the emphasis on the tracer images as the cross-well pressure data.

Our preferred model choice was to omit the problematic pressure data, and run the inversion using the 3–injection tracer data alone. It was also considered reasonable to merge the parameters for regions 5 and 6, since region 6 is at the edge of the modelling region and will have a more fragile permeability inference. The corresponding parameter inferences are shown in Table 1, where dimensionless scaled model inference and corresponding actual unitised values are presented. The final column is the dimensionless uncertainty estimate $\sigma_l^2 = H_{ll}^{-1}$ for each parameter formed from the inverse Hessian matrix at the final optimum, where $H = J^T J + C_p^{-1}$ and C_p is the effective Bayesian prior covariance. The final inversion forward model images and associated data snapshots are depicted in Fig. 3, where good qualitative agreement between the model and experimental data is observed.

3 Numerical Model for Multiphase Predictions

In order to model the FluidFlower CO_2 prediction benchmarks, we solve the two-phase, two-component mass conservation equations

$$\frac{\partial}{\partial t}\left\{ \phi \sum_\beta S_\beta \rho_\beta \chi_\beta^\kappa \right\} + \nabla \cdot \boldsymbol{F}^\kappa - q^\kappa = 0 \tag{4}$$

where S_β is the saturation of phase β (either liquid or gas in this case), ρ_β is the density of phase β (kg m^{-3}), χ_β^κ is the mass fraction of fluid component κ (either water or CO_2 in this

Table 1 Inversion results from 3–tracer model. Permeabilities provided in benchmark description (k_l) (Nordbotten et al. 2022) are multiplied by m_l to give the permeabilities used in the modelling. Note: sands F and G merged to a common (averaged) permeability

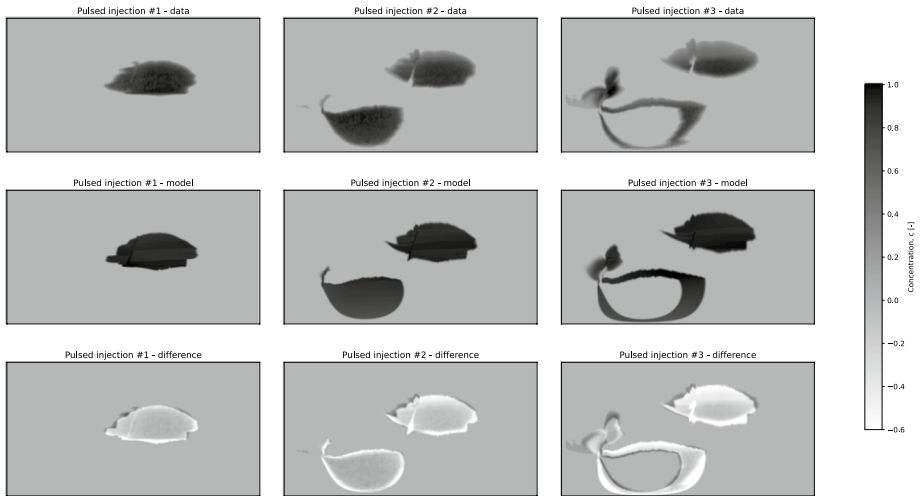| Facies	Index (l)	Modifier (m_l)	Permeability $m_l k_l$ (darcy)	Uncertainty σ_l
ESF	1	0.67	26.2	0.235
C	2	0.60	176	0.65
D	3	0.73	309	0.78
E	4	0.96	678	1
F	5	4.18	1080	0.61
G	6	4.18	1080	0.61
Fault1	7	2.74	1340	0.84
Fault2	8	–	0	–
Fault3	9	3.23	1575	0.84

Fig. 3 Tracer data image (top row), associated MAP forward models (middle row), and difference between data and model (bottom row). Each column represents one stage of the three-stage tracer test

case) in phase β, $\boldsymbol{F}^{\kappa}$ is the mass flux of fluid component κ (kg.m^{-2}.s^{-1}), and q^{κ} is an external source/sink (kg m^{-3} s^{-1}).

The fluid flux is a sum of advective and diffusive fluxes (no hydrodynamic dispersion is included in this model):

$$\boldsymbol{F}^{\kappa} = -\sum_{\beta} \left(\rho_{\beta} \frac{\boldsymbol{k}k_{\mathrm{r},\beta}}{\mu_{\beta}} \left(\nabla P_{\beta} - \rho_{\beta}\boldsymbol{g} \right) + \phi\rho_{\beta}d_{\beta}^{\kappa}\nabla\chi_{\beta}^{\kappa} \right), \tag{5}$$

where $\boldsymbol{k}$ is the absolute permeability tensor (m^2), $k_{\mathrm{r},\beta}$ is the relative permeability of phase β (-), μ_{β} is the dynamic viscosity of phase β (Pa.s), P_{β} is the pressure of each phase (Pa), $\boldsymbol{g}$ is gravity (m.s^{-2}), and d_{β}^{κ} is the molecular diffusion coefficient of component κ in phase β (m^2 .s^{-1}).

The mass balance is closed by noting

$$\sum_{\beta} S = 1, \qquad \sum_{\kappa} \chi_{\beta}^{\kappa} = 1 \ \forall\beta, \tag{6}$$

and that the phase pressures are related by the capillary pressure P_{c}

$$P_{\mathrm{c}} = P_{g} - P_{l}, \tag{7}$$

where P_{g} and P_{l} are the gas and liquid phase pressures, respectively.

Although hydrodynamic dispersion would be expected to influence the geometry of downwelling fingers during the anticipated convective mixing (amongst other factors that will also affect the size and position of the fingers, such as grid resolution), it is not expected to influence the mass flux of CO_2 from the gas phase to the liquid phase (Liang et al. 2018), which is the primary metric that we are concerned with.

The choice of capillary pressure and relative permeability models to implement is often informed by experimental values. In this case, in the absence of detailed experimental

curves for these quantities, we chose to use the common Brooks-Corey forms for both with only experience for justification (Brooks and Corey 1966)

$$P_c = P_e S_{\text{eff}}^{-1/\lambda}, \tag{8}$$

where P_e is the capillary threshold entry pressure, λ is the Brooks-Corey exponent, and

$$S_{\text{eff}} = \frac{S_l - S_{l,r}}{1 - S_{l,r}}, \tag{9}$$

where $S_{l,r}$ is the irreducible saturation of the liquid phase. The Brooks-Corey relative permeability model is (Brooks and Corey 1966)

$$k_{r,l} = k_{l0} S_{\text{eff}}^{c_l}, \tag{10}$$

$$k_{r,g} = k_{g0}(1 - S_{\text{eff}})^2 \left(1 - S_{\text{eff}}^{(2+c_g)/c_g}\right), \tag{11}$$

where k_{l0} and k_{g0} are the end point relative permeabilities, and c_l and c_g are coefficients. In all of the modelling, we used the capillary entry pressures and endpoint saturations provided for each sand type (Nordbotten et al. 2022), and constant exponents $\lambda = 2$, $c_l = c_g = 2$.

To solve the governing equations, we use the open-source multiphysics code MOOSE (Lindsay et al. 2022; Green et al. 2018; Wilkins et al. 2020, 2021). The governing equations are discretised in space using the finite volume method with a two-point flux approximation by default (although extended stencils are available in MOOSE if desired). A structured two-dimensional mesh is used, with each cell being 0.01m by 0.01m to match the discretisation required for reporting the spatial distribution of CO_2. Thickness-averaged porosity and permeability (where permeability is calculated through inversion described above) are used, such that the porosity and permeability in the centre of the FluidFlower are higher than at the edges to account for deformation of the perspex face, see Sect. 2. A constant molecular diffusion coefficient for CO_2 in the liquid phase, $d = 2 \times 10^{-9}$ m^2.s^{-1}, was used for all facies.

The finite volume discretisation is inherently conservative, and hence mass conservation is assured. Flux across the face of adjoining elements is calculated using linear interpolation of face-centred flux gradients, which makes the default implementation second-order accurate in space (Moukalled et al. 2016). Phase mobility $\rho_\beta k_{r,\beta} \mathbf{k}/\mu_\beta$ is upstream weighted (upwinded). Various temporal discretisations are available as part of the MOOSE framework, allowing higher-order time stepping if required. In this case, the governing equations are solved in a fully implicit manner using implicit Euler timestepping.

The finite volume method is a relatively new addition to MOOSE (Lindsay et al. 2022). Previous code for flow in porous media that has been implemented within MOOSE (Green et al. 2018; Wilkins et al. 2020, 2021) has used the finite element method. However, as the flow here was expected to be strongly dependent on capillary barriers, this study presented an opportunity to implement porous flow using a finite volume discretisation within MOOSE that would enable easy handling of sharp discontinuities in gas saturation at facies boundaries. Using the automatic differentiation capability within MOOSE (Lindsay et al. 2021), code development for this problem was rapid. Indeed, most of the coding was completed within only a few days. As with previous efforts for modelling flow in porous media using MOOSE, this capability has been open-sourced and is therefore available to all users.

Fluid properties are computed using high-accuracy equations of state. Water density is computed using the IAPWS Industrial Formulation 1997 (IAPWS 2007), while water viscosity is calculated using the IAPWS 2008 formulation (IAPWS 2008). The density of CO_2 is calculated using the Span and Wagner equation of state (Span and Wagner 1996), with viscosity calculated using (Fenghour et al. 1998). These choices are for both accuracy, with each equation of state widely used in practice for their demonstrated accuracy over a wide range of pressure and temperature ranges, as well as convenience, being already available in MOOSE.

The Span and Wagner formulation for CO_2 uses density and temperature as the primary variables with which to calculate properties such as density, enthalpy and internal energy. In order to use pressure and temperature as the primary variables, it is necessary to calculate density by iteration. As this can be a bottleneck in the computations, a tabulated version of this equation of state is used for speed. This is several of orders of magnitude faster than using the high-accuracy equations of state, and is comparable in time to simpler approximations of fluid properties.

The partitioning of CO_2 and water in each phase is computed using a high-precision brine-CO_2 equation of state that calculates the mutual solubility of CO_2 into the liquid brine and pure water into the CO_2-rich gas phase using an accurate fugacity formulation (Spycher et al. 2003; Spycher and Pruess 2005, 2010). As with many reservoir simulators, it is assumed that the phases are in instantaneous equilibrium, meaning that dissolution of CO_2 into the brine and evaporation of water into the CO_2 phase happens instantaneously.

The density of the aqueous phase with the contribution of dissolved CO_2 is calculated using a thermodynamically-consistent mixing rule for partial molar volumes

$$\frac{1}{\rho} = \frac{1 - \chi_{CO2}}{\rho_b} + \frac{\chi_{CO2}}{\rho_{CO2}}, \tag{12}$$

where ρ_b is the density of brine, χ_{CO2} is the mass fraction of CO_2 dissolved in the aqueous phase, and ρ_{CO2} is the partial density of dissolved CO_2 (Garcia 2001).

As water vapour is only ever a small component of the gas phase in the temperature and pressure ranges considered, the density of the gas phase is assumed to be simply the density of CO_2 at the given pressure and temperature.

No contribution to the viscosity of each phase due to the presence of CO_2 in the aqueous phase or water vapour in the gas phase is included. As a result, the viscosity of the aqueous phase is simply the viscosity of brine, while the viscosity of the gas phase is the viscosity of CO_2.

A persistent set of primary variables that remain independent in all phase states is used for this miscible multiphase flow problem. In this approach, the primary variables are pressure of a reference phase β, P_β, and total mass fraction of a fluid component summed over all phases

$$z^\kappa = \frac{\sum_\beta S_\beta \rho_\beta \chi_\beta^\kappa}{\sum_\beta S_\beta \rho_\beta}. \tag{13}$$

Using this set of persistent primary variables, saturation is calculated using a compositional flash, after which all fluid properties can be computed.

4 Multiphase Flow Predictions

A dominating feature of this benchmark problem, independent of the visual appeal of the 'representative model' of typical geology at expected operating depths of > 900 m, is that the operating conditions are near–surface. One immediate consequence is that the density ratio between the gaseous CO_2 and the water was a factor of approximately 500, which is several orders of magnitude greater than that for depths commonly considered in practice, where the densities of supercritical CO_2 and formation water differ by less than a factor of two. Combined with the relatively high sand permeabilities and low injection rate which result in a small viscous pressure gradient during injection, it follows that migration of the gas phase CO_2 is primarily driven by buoyancy, while the spatial distribution of this phase is mainly controlled by the capillary entry pressures of each sand unit (and to a lesser extent the permeability contrast between sand units). The Capillary number $Ca = \mu u/\gamma$, where u is the Darcy velocity (m.s^{-1}) of the gas-phase and γ is the interfacial surface tension ($\gamma \approx 3 \times 10^{-2}$ N.m^{-1}), is estimated to be $< O(10^{-6})$. This problem is therefore amenable to modelling by invasion percolation (eg. Cavanagh and Haszeldine 2014; Trevisan et al. 2017). Indeed, one of the first attempts we made of modelling the plume of CO_2 was through a simple invasion percolation model which would run in less than a second. This provided a first-order sanity check for the problem, where we could see that CO_2 injected in the bottom port would fill the structure beneath the sealing unit before spilling up the lower fault, for example.

An unwelcome and perverse consequence of the density contrast was the numerical difficulty in modelling this experiment. As the density of CO_2 in the gas phase is so small, cells where a gas phase appears (when the amount of CO_2 present in the cell exceeds the equilibrium solubility in water) necessitates extremely small timesteps, as we now explain through a simple example. If we consider the cells in the immediate vicinity of an injection port, CO_2 saturates the liquid phase and a gas phase appears. However, as the density of CO_2 at the operating conditions is less than 2 kg.m^{-3}, any gas phase rapidly saturates the cell such that the gas saturation rises sharply in those cells, resulting in steep gradients in gas saturation between adjacent cells. All numerical solvers that we tried found it difficult to converge when this occurred, especially when multiple adjacent cells changed from single phase liquid to two phase liquid and gas. In order to reduce the numerical residual sufficiently, the timestep needed to be cut sharply, which slowed the overall simulation time. Through trial and error, we found that the fastest simulation time could be achieved by taking timesteps of only 2 s during the injection phase, which could be increased to 60 s during the post injection phase.

Due to the small timesteps required, the runtime of each model was greater than a day. As the models had only $\approx 40{,}000$ cells, using multiple processors could only slightly mitigate the time required due to the small timesteps, as interprocessor communication quickly negated any speedup gained by parallelisation. This meant that the computer wall time was greater than the timesteps within the simulation during the injection phase, which is a significant difference to the typical CO_2 modelling work that is performed, where wall time must be several orders of magnitude smaller than the simulation timesteps in order for modelling underground storage for hundreds of years to be feasible.

As the onset time for convective mixing was expected to be short given the high permeability of the sands (Emami-Meybodi et al. 2015), convection was expected to be significant, especially in the regions beneath the sealing units. Though instabilities due to numerical rounding will eventually lead to convection even in a homogenous sand, the onset of

convection can be delayed. In practice, the size and distributions of the sand grains within each sand unit will vary slightly, so that the sand unit is not entirely homogeneous. To avoid any delayed convection in the numerical simulations, a small random noise sampled from a uniform distribution of $\pm 0.01k$ was added to the permeability of each sand unit.

The results from the numerical model are presented in Fig. 4. This figure shows both CO_2 in the gas phase (white), and CO_2 dissolved in the liquid phase (grey). Carbon dioxide injected in the lower port rises due to density difference and spreads beneath the sealing unit. This structure fills until CO_2 spills into the heterogeneous fault in the lower left corner, whereby it rapidly rises, fills beneath a sand unit of higher threshold entry pressure, before the capillary pressure increase leads to breakthrough. The CO_2 then flows upwards, until it reaches the upper sealing unit. It then begins to fill beneath this seal until injection stops, see the upper left-hand part of Fig. 4. A second CO_2 plume is created during injection in the upper injection port. This CO_2 rises and spreads beneath the upper sealing unit. Some CO_2 reaches the high-permeability fault in this part, and some spreading along the two fingers of this fault is observed, as per the upper left-hand part of Fig. 4.

At the end of the five hour injection period, approximately 70% of the total injection CO_2 is in the gas phase, with the remaining CO_2 dissolved in the liquid phase. This can be seen in the upper-left part of Fig. 4, where most of the plume contains gas-phase CO_2. A small amount of dissolved CO_2 (grey) can be seen to diffuse beneath both the top and bottom plumes, and some dissolved CO_2 is observed to be diffusing into the sealing unit atop the bottom plume. Note that within the plume where gas phase CO_2 is present, the liquid in this region is fully saturated with CO_2, including the region in the vicinity of the injection
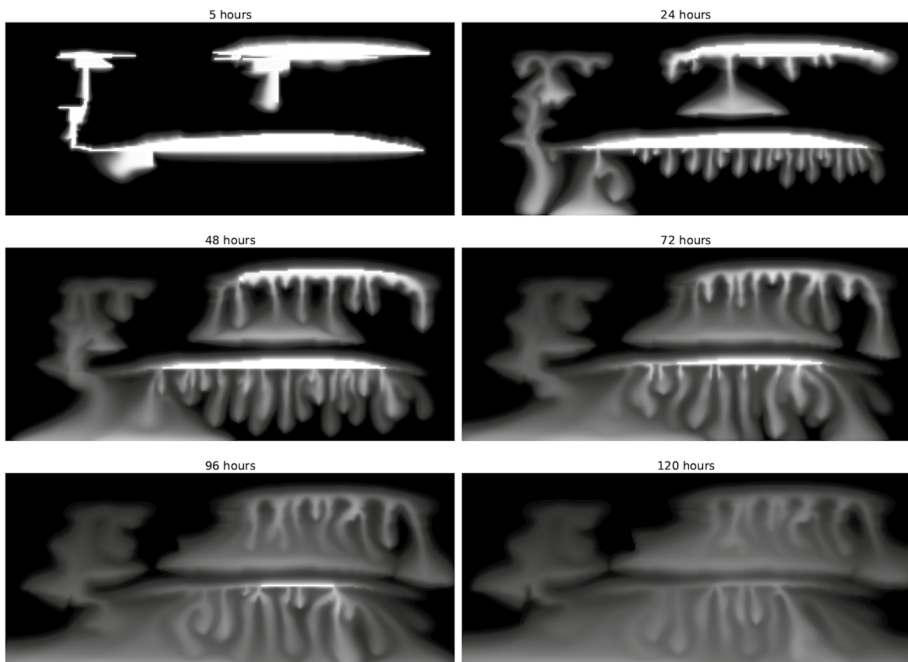

Fig. 4 Evolution of the CO_2 distribution over time. Black represents no CO_2, white is CO_2 in the gas phase, and grey is dissolved CO_2

ports (due to the instantaneous phase equilibrium assumption inherent in the equation of state used).

At the end of the first day, fingering is observed beneath the upper and lower plumes, see the upper-right part of Fig. 4, as the slightly denser CO_2 saturated water moves downwards. Significant slumping at both injection ports is also present, with all CO_2 near the ports having dissolved. The dissolved CO_2 near the lower injection port has reached the base of the FluidFlower and has begun spreading laterally, while the dissolved CO_2 near the upper injection port has reached the top of the lower sealing unit and begun to spread outwards, with only a small amount diffusing into the seal. We also note that all of the CO_2 that migrated up the lower fault into the upper-left part of the model has dissolved at the end of the first day, with the plume of dissolved CO_2 flowing down the fault and just reaching the base of the model after 24 h.

Over the remaining days of the simulation, CO_2 in the gas phase continues to dissolve into the water, as the downwelling fingers of CO_2-rich water grow and merge, and unsaturated water is drawn upwards towards the base of the two plumes. The interplay between the downwelling fingers and the unsaturated water drawn upwards results in lateral movement of the fingers, as observed in Fig. 4. The fingers formed beneath the lower plume reach the base of the model after three days, see Fig. 4. After four days, only a small amount of CO_2 in the gas phase is left in the model, with all CO_2 dissolved by the end of the fifth day.

One interesting feature of the results presented in Fig. 4 is that it appears that most of the CO_2 that invades the lower sealing unit is CO_2 that enters from above, that is, CO_2 that is injected in the upper port above the sealing unit, rather than CO_2 from the lower injection port. As the CO_2 above the sealing unit dissolves, the density of the saturated water increases slightly. This saturated water sinks and spreads along the sealing unit, whereby dissolved CO_2 diffuses into the sealing unit. This is especially evident in the final snapshot at the end of five days shown in the lower-right part of Fig. 4.

These numerical results show good qualitative agreement with the experimental results (Fernø et al. 2023). A detailed quantitative comparison between the metrics reported in the sparse data specified in the benchmark description (Nordbotten et al. 2022) is presented in Flemisch et al. (2023), so no comparison is provided in this manuscript for brevity.

5 Conclusion

This paper describes our approach to modelling the novel FluidFlower International Benchmark experiment. For model calibration or inference, sand permeabilities scale factors were optimised to match modelled tracer data to digitised tracer concentration images. The forward model in the fitting phase employed a simplified steady-state Poisson model to minimise numerical cost. The multiphase flow experiment was then modelled using a rapidly developed finite volume capability in the open-source simulation framework MOOSE. A small set of ensemble simulations was undertaken to explore key sensitivities and to compute the statistical metrics required for the benchmarking exercise. The spatial distribution of CO_2 predicted in the numerical models showed good qualitative agreement with the suite of experimental results, and key metrics show good quantitive agreement with experiments for many of the metrics.

Appendix 1: Inversion with Cross-well Pressure Data

In the inversion described in Sect. 2, cross-well pressure data was excluded due to noise, drift, and inconsistencies with the expected dependence on distance from the injection ports. An alternative approach to the inversion is to include the cross-well data, but downweight this data as described in Sect. 2, and to allow independence in sands F and G (regions 5 and 6 in Fig. 1). The result of this enriched inversion is shown in Table 2: it is apparent that the parameters differ from the previous inversion, but fall within the estimated standard deviation associated with the estimated "statistical power" of the data embedded in the likelihood weightings.

Appendix 2: Sensitivity Studies

In order to assess the robustness of the numerical predictions, sensitivity to key parameters was assessed through multiple scenarios of the base model used with modified properties. A total of 12 scenarios were modelled:

- All permeabilities multiplied by 0.25, 0.5, 0.75, 1.25, 1.5, 2;
- Permeabilities calculated by the alternative inversion described in Appendix A;
- All liquid residual saturations $S_{l,r}$ multiplied by 0.5, 2;
- All gas residual saturations $S_{g,r}$ multiplied by 2;
- All gas relative permeability end points k_{g0} multiplied by 2;
- Temperature increased/decreased by 2°C;
- Threshold entry pressure of sealing unit P_e halved.

The parameters for these scenarios were chosen based on experience and estimates of the possible uncertainty in value provided, and also their expected influence on the results. For example, the onset and rate of convective mixing is strongly dependent on the absolute permeability of the porous media, so this is a key sensitivity that can be explored. Relative permeability endpoints can affect the distribution and saturation of CO_2 in the plume, so may influence the amount of CO_2 that migrates through the heterogeneous fault into the upper-left part of the model.

Table 2 Inversion results from 3–tracer inversion, with cross-well data, and facies F and G independent

Facies	index l	m_l	permeability (darcy) $(m_l k_{w,l})$	σ_l
ESF	1	0.89	34.6	0.24
C	2	1.08	317	0.65
D	3	0.78	332	0.78
E	4	1.2	840	1.0
F	5	5.7	1477	0.61
G	6	2.43	1186	0.84
Fault1	7	4.27	2082	0.84
Fault2	8	–	0	–
Fault3	9	2.41	1176	0.84

An example of the results for several different scenarios is shown in Fig. 5, where the distribution of CO_2 in the gas and liquid phases is compared to the base model presented in Sect. 4 at the end of three days. Using the alternative inversion strategy where cross-well pressure is included (Appendix A), we observe only minor differences in the CO_2 distribution, mainly in the location of the fingers. This is probably expected, as only minor differences between the permeability scaling was found when cross-well pressure data was included in the inversion.

More significant differences are observed when decreasing and increasing the permeability by a factor of two, see Fig. 5. As expected, halving the permeability in each facies slows the rate of dissolution due to convection. Fewer, smaller fingers are present in this case when compared to the base case. As a result, more CO_2 remains in the gas phase after three days in this scenario, and only a small amount of the CO_2 injected in the lower port has reached the base of the model, and similarly for the upper injection port. Interestingly, halving the permeability also increases the amount of CO_2 that migrates up the lower fault into the upper-left part of the model. Increasing the permeability of each facies by a factor of two has the opposite affect on the distribution of CO_2. Convection has obviously begun earlier, and lead to a faster rate of dissolution (as expected). In this scenario, nearly all of the CO_2 in the model has dissolved, and much more dissolved CO_2 has reached the base (or the top of the lower sealing unit in the case of CO_2 injected through the upper port). Doubling the permeability in each sand also reduces the amount of CO_2 that reaches the upper-left part of the model through the heterogeneous fault.
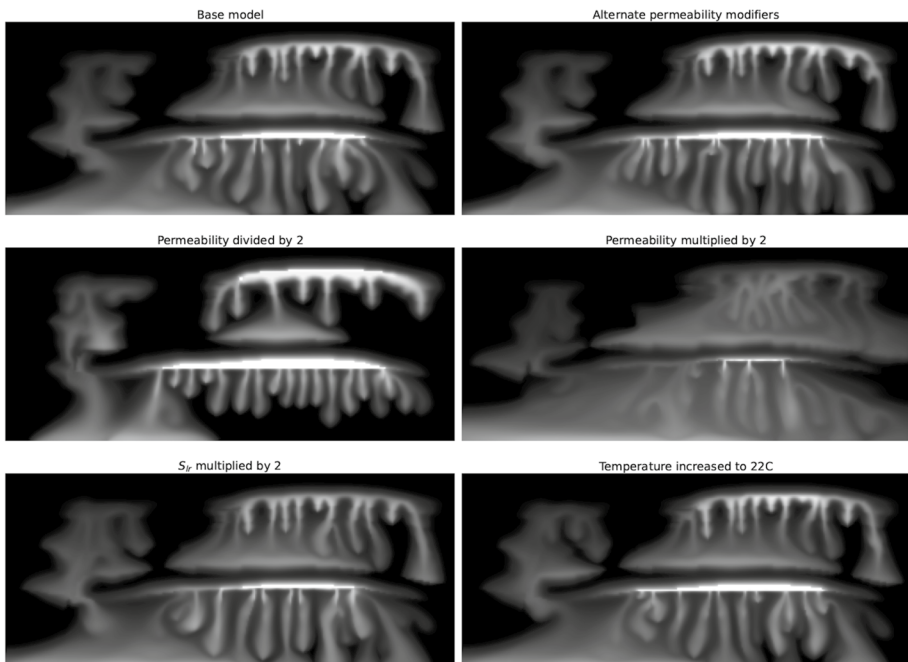

Fig. 5 Comparison of the CO_2 distribution for various scenarios at 72 h. Black represents no CO_2, white is CO_2 in the gas phase, and grey is dissolved CO_2

The maximum saturation in each grid block was reduced by doubling the irreducible water saturation, which promoted more up-fault migration from the CO_2 injected in the lower port, see Fig. 5. A consequence of increased migration through the fault into the upper-left region of the model is a smaller amount of CO_2 present beneath the lower sealing unit. Though the fingers of dissolved CO_2 look similar to those in the base case, less gas phase CO_2 remains in this lower structural closure in this case compared to the base case.

The final scenario depicted in Fig. 5 shows the effect of a small temperature increase. Increasing temperature by two degrees Celsius slightly reduces the densities of each fluid and also the solubility of CO_2 (by approximately 2%), which results in only minor changes to the CO_2 distribution.

The difference in the required metrics for all scenarios are shown in Figs. 6 and 7 for boxes A and B, respectively. For box A, all scenarios show a maximum for the mobile gas-phase CO_2 at the end of the injection period, while a spread in predictions for the amount of mobile gas-phase CO_2 after 72 h in box A is observed, with this metric ranging from approximately 1.75 g to 0 g. In nearly all cases, there is a negligible amount of CO_2 in the immobile gas phase after 72 h (that is, CO_2 at the residual saturation). The one exception is the case where the entry pressure of the sealing unit was reduced in an attempt to enable gaseous CO_2 to breach the sealing unit.

The plots of mobile and dissolved CO_2 exhibit fairly regular bumps throughout the simulations. These bumps corresponds to the sharp peaks in immobile CO_2 also shown in Fig. 6. This phenomena is due to the finite grid blocks aligned with the base of the plume of gas-phase CO_2 beneath the lower sealing unit and the way we define mobile and immobile CO_2. If we consider the plume at the end of the injection period, the bottom of the trapped CO_2 is horizontal, and aligns with the structured mesh. As time increases, the dissolved CO_2 in the lower-most row of grid blocks containing CO_2 diffuses into the row of grid blocks below the bottom of the plume. This reduces the saturation in the lower-most
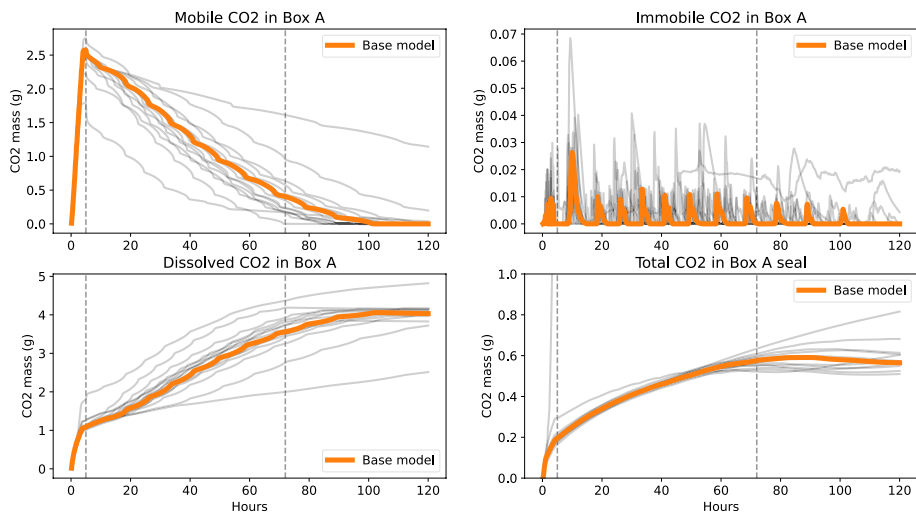

Fig. 6 CO_2 in various states in Box A for all scenarios. Grey lines represent each scenario. Dashed vertical line at 5 h represents the end of injection. Dashed vertical line at 72 h represents the time sparse data was reported (Nordbotten et al. 2022)
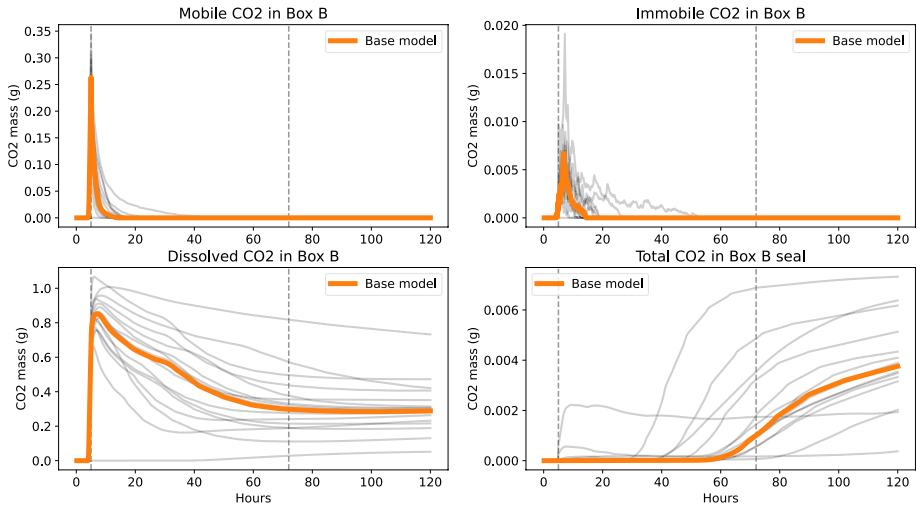
Fig. 7 CO_2 in various states in Box B for all scenarios. Grey lines represent each scenario. Dashed vertical line at 5 h represents the end of injection. Dashed vertical line at 72 h represents the time sparse data was reported (Nordbotten et al. 2022)

row of grid blocks containing CO_2 in the gas phase. This continues until the saturation in this row reaches its residual value, at which point there is now immobile CO_2, and we see a spike in immobile CO_2 in Fig. 6. The immobile CO_2 in this lowest row of blocks continues to dissolve as the denser CO_2-saturated water sinks (a process enhanced by convection). This continues until all immobile CO_2 in this row is dissolved, in which case the process is repeated in the next row above the previous lowest (as the plume of mobile CO_2 has shrunk vertically). As there is now mobile CO_2 in this row, dissolution increases slightly, and the amount of mobile CO_2 decreases accordingly, hence the bumps. This continues until all of the CO_2 in the gas phase dissolves. This artefact could be mitigated by refining the mesh in this location, albeit at the expense of computational time.

Similar trends are observed in the amount of dissolved CO_2 in box A, with a range of values between 1.8 g and 4.5 g observed. The onset and rate of convection was observed to be sensitive to absolute permeability in the scenarios shown in Fig. 5 through the difference in size and number of fingers. In most cases, we can see an increase in dissolution within the first 20 h, which is due to the onset of convection and the subsequent accelerated dissolution due to this process.

The amount of CO_2 in the lower sealing unit within Box A is observed to be fairly insensitive to these scenarios (with only the case where breakthrough into the sealing unit was enabled, which is the rapidly increasing line shortly after injection beings that is far outside the scale of this figure showing any significant difference).

The sensitivity of the reported metrics for box B in the upper-left region of the model to these scenarios is shown in Fig. 7. In all cases, there is no CO_2 in the gas phase in this box after the first two days in any scenario (and generally none after the first day), hence there is no mobile or immobile CO_2 after 72 h in any scenario. There is some spread in the amount of dissolved CO_2 in box B and also the amount of CO_2 in the sealing unit within box B, see Fig. 7. This is expected from the results shown in Fig. 5, where the amount of

CO_2 that migrates into this region shown to be sensitive to properties such as absolute permeability and relative permeability endpoints.

One interesting thing to note for both box A and box B is that most of the CO_2 in the sealing unit is dissolved CO_2 injected in the upper port that has sunk and spread along the sealing unit, and not CO_2 that has diffused into the seal from below, as might be expected.

If we treat the results from these scenarios, very cavalierly, as samples from a distribution, means and standard deviations for a given prediction can be estimated using a classical frequentist prediction interval $\bar{X}_n \pm T_a s_n \sqrt{1 + 1/n}$, using the sample mean $\bar{X}_n$ and standard deviation s_n. Here n is the number of samples (the scenarios in this case), and T_a is drawn from the appropriate percentile of the Student's t-distribution. Using this ad–hoc methodology, estimates of the mean and standard deviations of the sparse data were computed, the results of which have been reported in the the comparison paper (Flemisch et al. 2023).

Data Availibility Statement The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

References

Brooks, R.H., Corey, A.T.: Properties of porous media affecting fluid flow. J. Irrig. Drain. Div. **92**, 61–88 (1966)

Cavanagh, A.J., Haszeldine, R.S.: The Sleipner storage site: Capillary flow modeling of a layered CO_2 plume requires fractured shale barriers within the Utsira Formation. Int. J. Greenhouse Gas Control **21**, 101–112 (2014)

Class, H., Ebigbo, A., Helmig, R., Dahle, H.K., Nordbotten, J.M., Celia, M.A., Audigane, P., Darcis, M., Ennis-King, J., Fan, Y., et al.: A benchmark study on problems related to CO_2 storage in geologic formations. Comp. Geosci. **13**(4), 409 (2009)

Emami-Meybodi, H., Hassanzadeh, H., Green, C.P., Ennis-King, J.: Convective dissolution of CO_2 in saline aquifers: progress in modeling and experiments. Int. J. Greenhouse Gas Control **40**, 238–266 (2015)

Fenghour, A., Wakeham, W.A., Vesovic, V.: The viscosity of carbon dioxide. J. Phys. Chem. Ref. Data **27**, 31–44 (1998)

Fernø, M.A., Haugen, M., Eikehaug, K., Folkvord, O., Benali, B., Nordbotten, J.M.: Meter-scale CO_2 injection experiments and data set. Transport in Porous Media (2023). this SI

Flemisch, B., Nordbotten, J.M., Fernø, M., Juanes, R., Class, H., Delshad, M., Doster, F., Ennis-King, J., Franc, J., Geiger, S., Gläser, D., Green, C., Gunning, J., Hajibeygi, H., Jackson, S.J., Jammoul, M., Karra, S., Li, J., Matthäi, S.K., Miller, T., Shao, Q., Spurin, C., Stauffer, P., Tchelepi, H., Tian, X., Viswanathan, H., Voskov, D., Wang, Y., Wapperom, M., Wheeler, M.F., Wilkins, A., Youssef, A.A., Zhang, Z.: The FluidFlower international benchmark study: Process, modeling results, and comparison to experimental data. Transport in Porous Media (2023). this SI

Garcia, J.E.: Density of aqueous solutions of CO_2. Technical Report LBNL-49023, LBNL (2001)

Green, C., Wilkins, A., La Force, T., Ennis-King, J.: Community code for simulating CO_2 storage: Modelling multiphase flow with coupled geomechanics and geochemistry using the open-source multiphysics framework MOOSE. In: 14th Greenhouse Gas Control Technologies Conference Melbourne, pp 21–26 (2018)

IAPWS: Release on the IAPWS Formulation 2008 for the Viscosity of Ordinary Water Substance. Technical report, IAPWS (2008). www.iapws.org/relguide/visc.pdf

IAPWS: Revised Release on the IAPWS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam. Technical report, IAPWS (2007). www.iapws.org/relguide/IF97-Rev.pdf

Koren, B.: A Robust Upwind Discretization Method for Advection, Diffusion and Source Terms, vol. 45. Centrum voor Wiskunde en Informatica, Amsterdam (1993)

Liang, Y., Wen, B., Hesse, M.A., DiCarlo, D.: Effect of dispersion on solutal convection in porous media. Geophys. Res. Lett. **45**(18), 9690–9698 (2018)

Lindsay, A., Stogner, R., Gaston, D., Schwen, D., Matthews, C., Jiang, W., Aagesen, L.K., Carlsen, R., Kong, F., Slaughter, A., et al.: Automatic Differentiation in MetaPhysicL and Its Applications in MOOSE. Nuclear Technology, 1–18 (2021)

Lindsay, A.D., Gaston, D.R., Permann, C.J., Miller, J.M., Andrš, D., Slaughter, A.E., Kong, F., Hansel, J., Carlsen, R.W., Icenhour, C., Harbour, L., Giudicelli, G.L., Stogner, R.H., German, P., Badger, J., Biswas, S., Chapuis, L., Green, C., Hales, J., Hu, T., Jiang, W., Jung, Y.S., Matthews, C., Miao, Y., Novak, A., Peterson, J.W., Prince, Z.M., Rovinelli, A., Schunert, S., Schwen, D., Spencer, B.W., Veeraraghavan, S., Recuero, A., Yushu, D., Wang, Y., Wilkins, A., Wong, C.: 2.0 - MOOSE: Enabling massively parallel multiphysics simulation. SoftwareX **20**, 101202 (2022). https://doi.org/10.1016/j.softx.2022.101202

Madsen, K., Nielsen, H.B., Tingleff, O.: Methods for non-linear least squares problems (2004)

Moukalled, F., Mangani, L., Darwish, M., et al.: The Finite Volume Method in Computational Fluid Dynamics, vol. 6. Springer (2016)

Nocedal, J., Wright, S.J.: Numerical Optimization. Springer, New York (1999)

Nordbotten, J.M., Fernø, M., Flemisch, B., Juanes, R., Jørgensen, M.: Final Benchmark Description: FluidFlower International Benchmark Study. (2022). https://doi.org/10.5281/zenodo.6807102

Nordbotten, J.M., Flemisch, B., Gasda, S., Nilsen, H., Fan, Y., Pickup, G.E., Wiese, B., Celia, M.A., Dahle, H., Eigestad, G., et al.: Uncertainties in practical simulation of CO_2 storage. Int. J. Greenhouse Gas Control **9**, 234–242 (2012)

Pruess, K., García, J., Kovscek, T., Oldenburg, C., Rutqvist, J., Steefel, C., Xu, T.: Code intercomparison builds confidence in numerical simulation models for geologic disposal of CO_2. Energy **29**(9–10), 1431–1444 (2004)

Span, R., Wagner, W.: A new equation of state for carbon dioxide covering the fluid region from the triple-point temperature to 1100 K at pressures up to 800 MPa. J. Phys. Chem. Ref. Data **25**, 1509–1596 (1996)

Spycher, N., Pruess, K.: CO_2-H_2O mixtures in the geological sequestration of CO_2. II. Partitioning in chloride brine at 12–100C and up to 600 bar. Geochim. Cosmochim. Acta **69**, 3309–3320 (2005)

Spycher, N., Pruess, K.: A phase-partitioning model for CO_2-brine mixtures at elevated temperatures and pressures: application to CO_2-enhanced geothermal systems. Transp. Porous Med. **82**(1), 173–196 (2010)

Spycher, N., Pruess, K., Ennis-King, J.: CO_2-H_2O mixtures in the geological sequestration of CO_2. I. Assessment and calculation of mutual solubilities from 12 to 100C and up to 600 bar. Geochim. Cosmochim. Acta **67**, 3015–3031 (2003)

Tian, X., Wapperom, M., Gunning, J., Jackson, S., Wilkins, A., Green, C., Ennis-King, J., Voskov, D.: The study of the history matching for FluidFlower experimental results. Transport in Porous Media (2023). this SI

Trevisan, L., Illangasekare, T.H., Meckel, T.A.: Modelling plume behavior through a heterogeneous sand pack using a commercial invasion percolation model. Geomech. Geophys. Geo-Energy Geo-Resour. **3**, 327–337 (2017)

Wilkins, A., Green, C.P., Ennis-King, J.: PorousFlow: a multiphysics simulation code for coupled problems in porous media. J. Open Source Softw. **5**(55), 2176 (2020). https://doi.org/10.21105/joss.02176

Wilkins, A., Green, C.P., Ennis-King, J.: An open-source multiphysics simulation code for coupled problems in porous media. Comput. Geosci. **154**, 104820 (2021). https://doi.org/10.1016/j.cageo.2021.104820