Milestone

M 2.1: A posterior ensemble of the dilation of the fracture network after assimilation of deformation data

Assessing the value of deformation measurments for diagnosing thermal breakthrough in enhanced geothermal sustems through the inversion of synthetic distributed strain sensing data.

Report information	
Work package	WP2: Digital Geothermal Reservoir Twin
Lead	ETH Zurich, Storra Dynamics GmbH
Authors	Michael Liem, Edoardo Pezzulli, Storra Dynamics GmbH
Approval	Vanille Ritz, ETH Zurich
Status	Final
Dissemination level	Public
Will the data supporting this document be made open access?	N.A.
Delivery deadline	15.06.2026
Submission date	27.07.2026
Intranet path	GENERAL/FILES/D.MILESTONES/Milestone_2_1.pdf

GEOTWINS Project Office | ETH Zurich, Dept. of Earth and Planetary Sciences
Sonneggstrasse 5, 8092 Zurich | geotwins-office@sed.ethz.ch | www.deepgeothermal.org

Co-funded by the European Union

This Project is funded by CETPartnership, the Clean Energy Transition Partnership under the 2023 joint call for research proposals, co-funded by the European Commission (GA N°101069750) and with the national funding organizations detailed on https://cetpartnership.eu/funding-agencies-and-call-modules.

Contents

Abstract . . . 1
1 Introduction . . . 3
2 Methodology Overview . . . 5
3 Geothermal Reservoir Model . . . 6
3.1 Domain and Grid . . . 6
3.2 Fracture Network . . . 6
3.3 Rock and Fluid Properties . . . 6
3.4 Initial Reservoir Conditions . . . 8
3.5 Boundary Conditions . . . 8
3.6 Well Configuration . . . 9
3.7 Simulation Setup . . . 10
3.8 Uncertain Parameters . . . 10
3.9 Reference Realisation . . . 11
3.10 Prior Ensemble . . . 11
4 Geomechanical Forward Model and Measurement Design . . . 14
4.1 Poro-Elastic Forward Model . . . 14
4.2 Reference Strain Response and Measurement Locations . . . 14
5 Scoping the Expected Value of Deformation Measurements . . . 17
5.1 Deformation Response Across the Prior Ensemble . . . 17
5.1.1 Strain Along the Wells . . . 17
5.1.2 Displacement Through the Reservoir and Overburden . . . 17
5.2 Parameter Controls on Deformation and Reservoir Performance . . . 20
5.3 Expected Variance Reduction from Deformation Measurements . . . 20
5.3.1 Expected Value of an Individual Displacement Measurement . . . 22
5.3.2 Expected Value of an Individual Strain Measurement . . . 22
5.3.3 Expected Variance Reduction of Distributed Strain Sensing . . . 23
6 Data Assimilation Methodology . . . 29
6.1 Overview . . . 29
6.2 State Vector . . . 29
6.3 Observation Operator . . . 30
6.3.1 Measurement Error Covariance . . . 30
6.3.2 Synthetic Observation Data . . . 30
6.3.3 Prior Simulations . . . 30
6.4 ES-MDA Update . . . 31
6.5 Data-Fit Diagnostics . . . 32
6.5.1 Misfit of the Ensemble Mean . . . 32
6.5.2 Misfit of the Ensemble Members . . . 32
6.5.3 Ensemble Spread . . . 33
6.5.4 Realized Noise Levels . . . 33
7 Geomechanical Inversion of Distributed Strain Sensing Data . . . 34

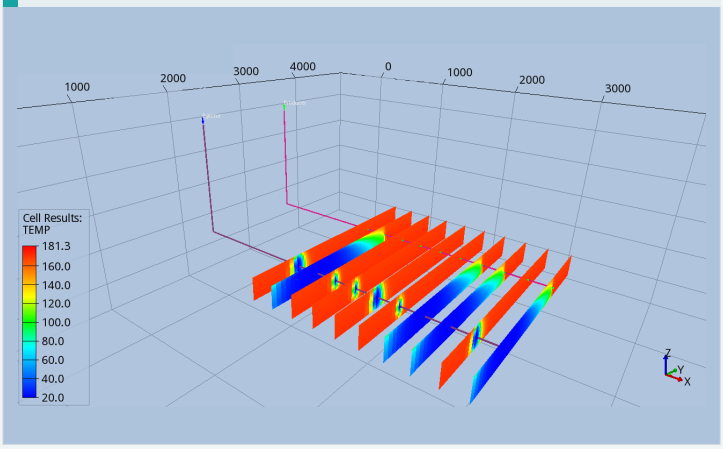
7.1 Reproduction of the Strain Observations . . . 34
7.2 Recovery of the Reservoir Properties . . . 36
7.2.1 Expected and Realised Variance Reduction . . . 37
7.3 Posterior Ensemble Forecast . . . 38
8 Conclusions and Outlook . . . 43
A Comparison of the Uniform Case with the Gringarten (1975) Analytical Solution . . . 47
A.1 Analytical Model of Gringarten (1975) . . . 47
A.2 Uniform Ten-Fracture OPM Setup . . . 48
A.3 Numerical validation study . . . 48
A.4 Conclusion . . . 51
A.5 Outlook . . . 54
B Reservoir Simulation Results: Uniform, Reference, Prior & Posterior . . . 56
B.1 Uniform Fracture Properties . . . 57
B.2 Reference Realisation . . . 63
B.3 Influence of Boundary Conditions and Matrix Permeability . . . 68
B.4 Prior and Posterior Ensemble Fields . . . 74
B.4.1 Overpressure . . . 74
B.4.2 Tracer Concentration . . . 77
B.4.3 Temperature . . . 77
C Ensemble Variance Diagnostics . . . 81
C.1 Correlation Analysis among the Parameters . . . 81
C.2 Correlation Analysis among the Measurements & Forecasts . . . 81
D ES-MDA Diagnostics . . . 83

Abstract

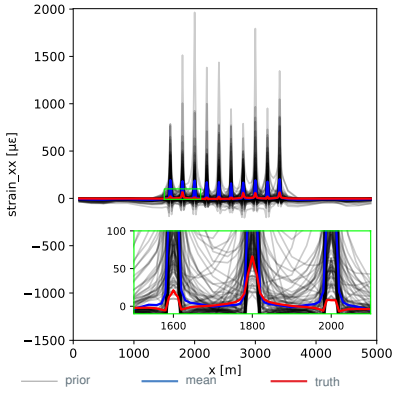
This report presents a proof-of-concept study of the value of deformation measurements in constraining the parameters of fracture networks that govern the performance of an enhanced geothermal system (EGS). The value of deformation measurements is assessed by performing a probabilistic inversion of synthetic distributed strain sensing data in order to obtain a posterior distribution of fracture parameters such as fracture permeability. The conceptual multi-stage enhanced geothermal system is represented by a fully three-dimensional reservoir with ten parallel fractures simplified as highly-localised permeable channels. To perform the probabilistic geomechanical inversion, an ensemble of simulations are performed with the open-source thermal-hydraulic reservoir simulator OPM Flow, that is one-way coupled to the fast, semi-analytical nucleus-of-strain approach for the simulation of drained poro-elastic deformation in homogeneous rock. Starting from a prior ensemble of 100 reservoir realisations spanning uncertainty in fracture and matrix permeability and porosity, 90 noisy strain measurements generated from a selected synthetic truth and recorded along the injection–production doublet are assimilated to produce a posterior ensemble of the uncertain reservoir properties. A conclusion of the preliminary study is that strain-sensing along the horizontal sections of the wellbore provides a sufficiently rich dataset to infer the unknown fracture permeabilities. When propagated forward, the obtained posterior of fracture permeabilities correctly reproduces the wellbore bottom hole pressures as well as the overpressure across the reservoir. Such a matched reservoir-pressure response leads to the significant reduction of uncertainties in the production rate and the thermal power output of the system. The matched reservoir-pressure response is also anticipated to improve assessments of fault stability and seismic risk. On the other-hand, strain responses are relatively insensitive to fracture porosity, and interestingly, a bias is noted where very large fracture permeabilities are underestimated by the inversion. Consequently, this bias translates to an underestimation of the transport of tracer and heat through such highly conductive fractures. As a result, forecasts of the thermal power output are estimated to be modestly overestimated when inverted from distributed strain data, since the arrival time of the cold thermal front is overestimated. The study demonstrates that a single early strain survey can characterise the evolving pressure architecture of a fractured geothermal reservoir long before thermal signals develop. Overall, this assessment provides a preliminary scoping of the value of deformation measurements, and particularly DSS, in characterising a fractured enhanced geothermal system and thereby inform preventative strategies for reservoir management. However, it remains to be seen if the value of geomechanical monitoring is preserved when the effects of thermal-elastic deformations, stress-dependent fracture permeability, and uncertain fracture geometries are incorporated in the analysis.

Visual Abstract

COLD-WATER TRANSPORT THROUGH CONDUCTIVE FRACTURES



DISTRIBUTED STRAIN SIGNAL ALONG THE INJECTOR



ES-MDA UPDATE OF FRACTURE PERMEABILITY



1 Introduction

Deep geothermal energy is one of the few renewable sources able to supply continuous, weather-independent baseload heat and power. In hydrothermal doublets and Enhanced Geothermal Systems (EGS), cold water is injected into a hot subsurface reservoir and produced from a second well, extracting thermal energy from the rock it passes through. The economic and operational viability of such a system depends on the hydraulic connectivity between the two wells, which is set by the geometry and properties of the fracture network that carries the flow.

Fractures in deep reservoirs span several orders of magnitude in permeability, from near-matrix values of a few millidarcy to highly conductive pathways exceeding 10^5 mD. Their spatial distribution is largely unknown before and during early production, yet it controls (i) the share of injected water that reaches the producer, (ii) its residence time in the hot rock, and hence the produced temperature, and (iii) the risk of thermal short-circuiting, in which rapid breakthrough along a dominant fracture cools the production well prematurely [1, 2]. Characterising the fracture network early is therefore central to forecasting and managing a geothermal doublet.

Conventional observations provide only a partial picture of the reservoir. Temperature and tracer measurements can reveal the pathways connecting the injector and producer, but their signals must first travel across the reservoir and may therefore take considerable time to develop. Well-pressure measurements respond much earlier but are available only at a small number of locations. Deformation measurements offer a potentially valuable additional source of information because changes in reservoir pressure deform the surrounding rock. If measured with sufficient spatial coverage, this deformation may reveal the evolving pressure distribution and, indirectly, the hydraulic properties of individual fractures.

Fibre-optic Distributed Strain Sensing (DSS) [3], records axial strain continuously along an optical fibre installed in or behind the well casing. A DSS survey can therefore provide spatially distributed observations along the injection and production wells soon after injection begins and well before a thermal signal reaches the producer. The usefulness of such a monitoring system depends on whether its measurements reduce uncertainty in the fracture-network parameters that govern reservoir performance. This reduction in uncertainty provides a measure of the information value of the measurements and can be used to assess whether a proposed monitoring system is worth deploying. It can also help determine which measurements, locations and observation times are most informative.

In this study, the value of deformation measurements is assessed using ensemble-based methods that estimate the reduction in uncertainty in the reservoir's fracture network and the thermal output of the geothermal system. Two complementary methods are used to assess the information provided by deformation measurements. First, expected variance reduction (EVR), calculated using ensemble-variance analysis (EVA) [4], estimates from the prior ensemble which reservoir properties and forecasts the proposed measurements are expected to constrain. Second, synthetic strain measurements are inverted using the Ensemble Smoother with Multiple Data Assimilation (ES-MDA) [5]. ES-MDA has a track record in geothermal characterisation, including permeability and tracer inversion at the Soultz EGS [6, 7]. ES-MDA updates the prior ensemble to produce a posterior ensemble consistent with the measurements. Comparing the prior and posterior ensembles then quantifies the reduction in uncertainty achieved through the inversion.

The report is structured as follows. Section 2 provides a brief summary of the methodology. Section 3 describes the geothermal reservoir model: the grid, fracture network, rock and fluid properties, well configuration, uncertain parameters, and prior ensemble. Section 4 presents the geomechanical forward model that converts the simulated pressure field to strain and defines the measurement points of the observation vector. Section 5.2 analyses the prior sensitivity of the observables, screening a priori what the strain data can and cannot constrain. Section 6 sets out the data assimilation methodology: the state vector, observation operator, prior ensemble construction, the ES-MDA update, and the data-fit diagnostics. Section 7 presents

and discusses the results, comparing prior and posterior ensembles in measurement, parameter, and forecast space. Section 8 summarises the findings and outlines directions for future work. Three appendices support the main text: Appendix A validates the OPM model against the analytical solution of [8] and demonstrates that a fully three-dimensional simulation is required to capture the dominant physics of the target configuration; and Appendix B documents the underlying reservoir-simulation results in detail, from the reference realisation to the prior and posterior ensemble fields. Appendix C and D diagnose in more detail the correlation structure of the ensemble and the inversion results.

2 Methodology Overview

This report documents a proof-of-concept study with the following workflow:

1. Construct a synthetic, fully three-dimensional geothermal doublet containing ten fractures with uncertain hydraulic properties simplified as highly-localised planar permeable channels.
2. Simulate the thermal-hydraulic reservoir response with the open-source simulator OPM Flow [9] and calculate the resulting undrained poroelastic deformation using a fast, semi-analytical nucleus-of-strain model [10].
3. Before any inversion is performed, estimate the expected variance reduction (EVR) from the prior ensemble using Ensemble Variance Analysis (EVA) [4] to identify which deformation measurements are expected to constrain the uncertain reservoir properties.
4. Apply ES-MDA ($N_A = 3$ iterations) to a synthetic strain observation to estimate the permeability and porosity of each fracture, the matrix permeability, and a boundary pore-volume multiplier, and thereby quantifying the resulting reduction in parameter uncertainty.
5. Propagate the posterior ensemble forward to assess how deformation measurements affect forecasts of reservoir pressure, production rate and thermal-power output.

The analysis begins with a prior ensemble of 100 reservoir realisations spanning uncertainty in fracture permeability and porosity, matrix permeability, and the boundary pore-volume multiplier. One realisation is selected as the synthetic truth. Ninety noisy strain measurements generated from this realisation and recorded along the injection–production doublet are then assimilated over three ES-MDA iterations. Because the true reservoir properties are known, this controlled synthetic experiment allows the ability of the strain measurements to recover those properties to be assessed directly.

The study focuses on the information contained in pressure-driven, drained poroelastic deformation. It does not yet account for thermoelastic deformation, stress-dependent fracture permeability, or uncertainty in fracture geometry. The results should therefore be interpreted as an initial assessment of whether deformation monitoring warrants further investigation under more realistic reservoir conditions.

3 Geothermal Reservoir Model

3.1 Domain and Grid

The synthetic reservoir occupies a domain of 5000 m × 4000 m × 800 m (x, y, z), representing a section of deep crystalline basement rock at depths of 3600–4400 m. The computational grid is a structured, non-uniform Cartesian mesh with $175 \times 81 \times 61 = 864{,}675$ cells. Cell sizes vary across the domain, with refined grid near fractures and coarse grid in the far field (see Figs. 1 and 2). The grid specification is summarised in Table 1.

Table 1: Grid specification summary.

Property	x	y	z	Total
Grid dimensions	175	81	61	864,675
Domain extent	5,000 m	4,000 m	800 m	16×10^9 m^3
Minimum cell size (fracture)	1 m	20 m	2 m	40 m^3
Maximum cell size (far field)	200 m	200 m	50 m	2×10^6 m^3

3.2 Fracture Network

The reservoir contains ten parallel vertical fractures oriented in the y–z plane and distributed uniformly in the x-direction with a spacing of 200 m. Each fracture has a length of 2000 m and a height of 100 m, resolved with 71 and 41 cells in y- and z-direction respectively. The fractures are modelled as a zone of 1 m thickness, resolved with a single grid cell in x-direction.

Each fracture zone is assigned to one of ten independent flow regions using the FLUXNUM keyword (FLUXNUM 1 through 10 for fractures 1–10). The surrounding matrix is assigned FLUXNUM 11.

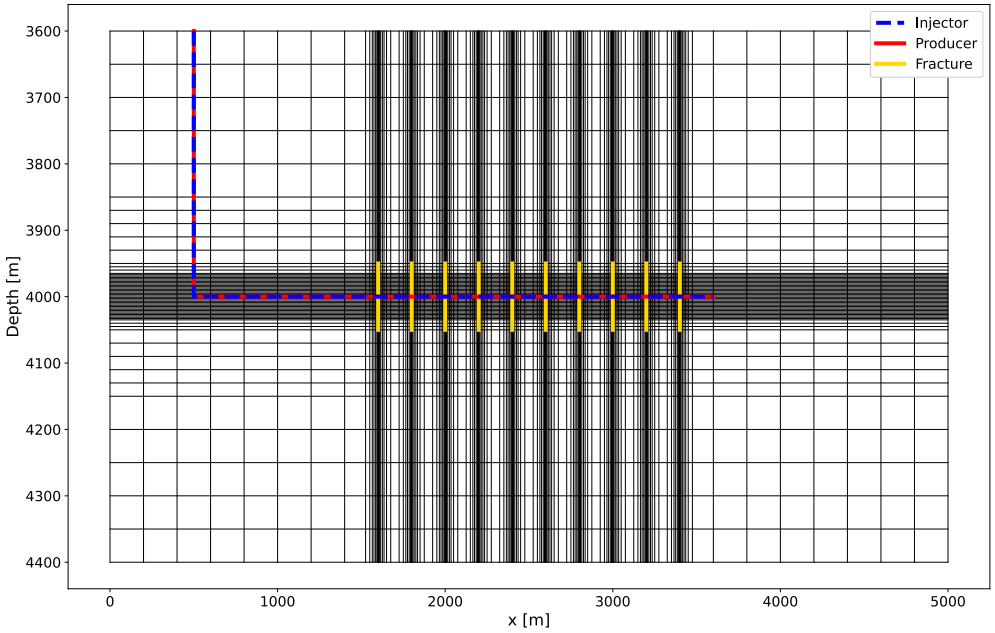
Each fracture zone is parameterised by assigning effective permeability and porosity values directly to the corresponding grid cells. These values represent the effective hydraulic properties of each fracture zone; they are applied as isotropic permeability (PERMX = PERMY = PERMZ) and porosity (PORO) in the simulator input. The fracture permeabilities and porosities are the primary uncertain parameters of this study (Section 3.8).

The fracture layout is illustrated in Figs. 1 and 2.

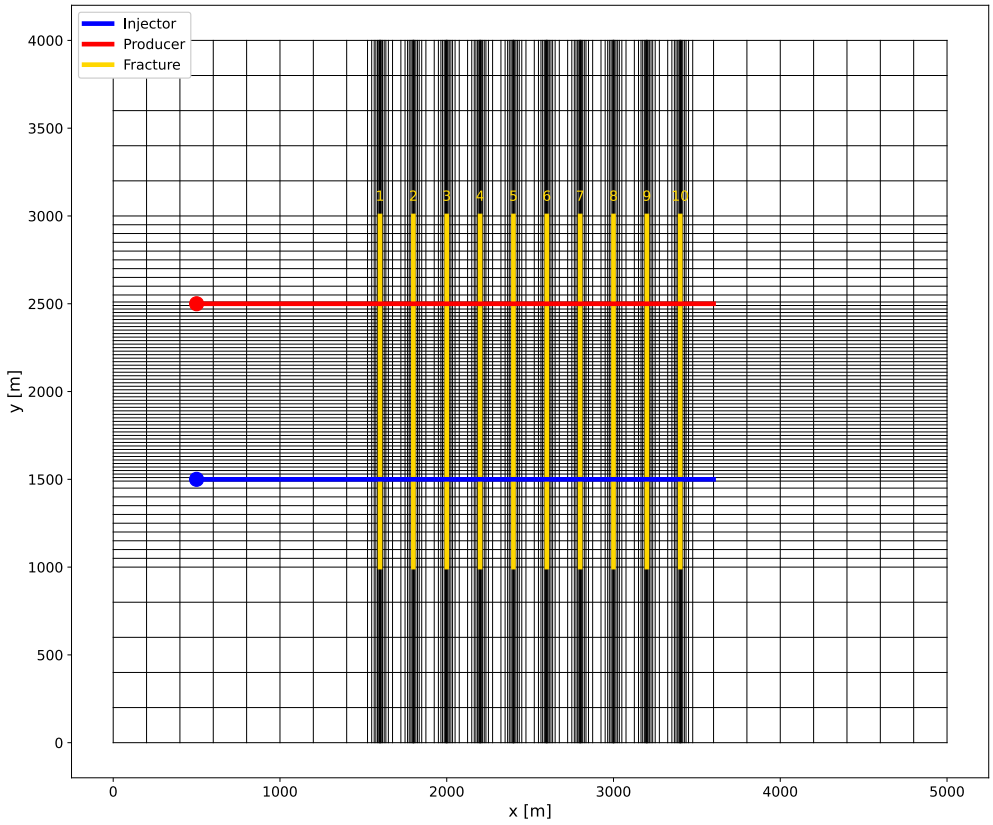
3.3 Rock and Fluid Properties

Matrix. The matrix permeability k_m is isotropic, with a reference value of 20 mD; it is treated as an uncertain parameter of the study (Section 3.8). The matrix porosity is fixed at $\phi_m = 0.15$. The rock heat capacity is 2200 kJ m^{-3} K^{-1} (2.2 MJ m^{-3} K^{-1}) and the thermal conductivity is 260 kJ m^{-1} d^{-1} K^{-1} (≈ 3 W m^{-1} K^{-1}), consistent with crystalline basement rock.

Fluid. The reservoir fluid is a single-phase sodium-chloride brine with a salinity of 1.4 mol kg^{-1} of water ($\approx$ 7.6 wt% NaCl, roughly twice seawater salinity). Its properties are evaluated as functions of pressure, temperature, and salinity from analytic correlations rather than from tabulated input (Section 3.7): the brine density combines the pure-water equation of state of Hu et al. [11] — a computationally cheap surrogate that reproduces the IAPWS-95 formulation (the Wagner–Pruß scientific standard for water) over the relevant range — with the salinity correction of the Batzle–Wang correlations [12]; the viscosity follows the temperature- and salinity-dependent Batzle–Wang correlation [12]; and the enthalpy combines a pure-water fit with the



(a) Cross-section (x–z plane) of the full model domain.



(b) Plan view (x–y plane) of the full domain.

Figure 1: Computational grid and doublet well configuration. Blue: injector well; red: producer well; gold lines: fractures. **(a)** Cross-section of the full model domain (x–z plane). The wells run vertically at $x = 500$ m and horizontally at 4000 m depth; injector and producer coincide in this projection. **(b)** Plan view (x–y plane) of the full domain, with the fractures numbered 1–10 (the numbering used throughout this report) and the well heads marked by dots.

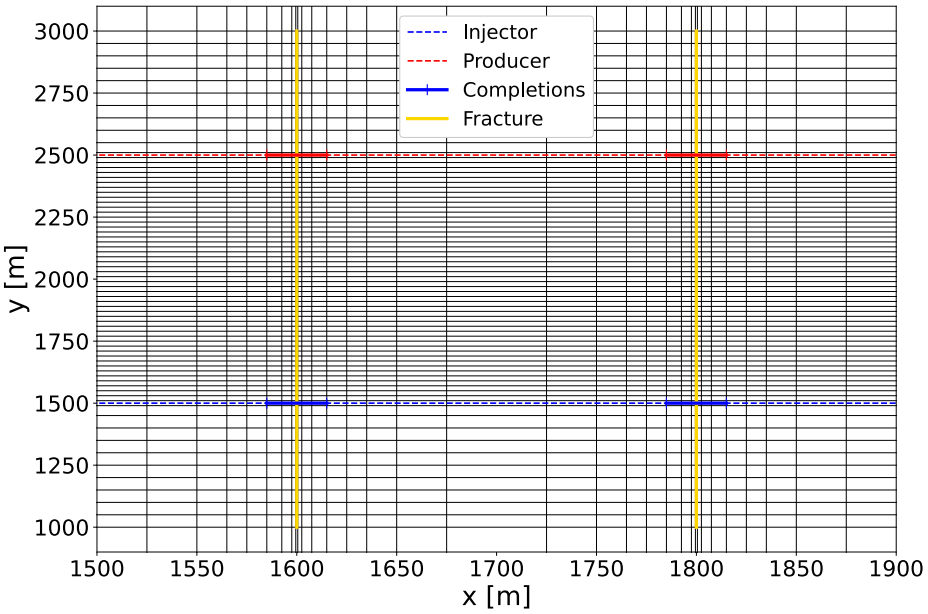


Figure 2: Zoomed plan view centred on fractures 1 and 2. The completion segments of the injector (blue) and producer (red) are indicated by thick line segments around each fracture (gold); the dashed lines mark the horizontal wellbores.

NaCl mixing-enthalpy model of Michaelides [13]. At standard surface conditions (15.56 °C, 1.013 bar) the resulting brine density is $\rho_w = 1051\,\mathrm{kg\,m^{-3}}$.

3.4 Initial Reservoir Conditions

The initial temperature follows a linear geothermal gradient of $30\,^{\circ}\mathrm{C\,km^{-1}}$ across the simulated depth range (Table 2).

Table 2: Initial temperature profile across the simulated depth range.

Location	Depth (m)	Temperature (°C)
Top of simulated domain	3600	158
Datum; horizontal well sections (Section 3.6)	4000	170
Bottom of simulated domain	4400	182

The initial pressure at the datum is 400 bar, corresponding to approximate hydrostatic conditions. The reservoir is initially fully water-saturated.

3.5 Boundary Conditions

The top and bottom faces of the grid are closed to fluid flow, with their temperature held fixed: they represent the confining formations above and below the reservoir interval, which seal the domain vertically but remain in thermal contact with it. The four lateral faces are open: a Dirichlet condition fixes their pressure and temperature to the initial equilibrium state (Section 3.4). The simulated block is only a finite window into a much larger geological formation, and the open lateral boundaries represent the pressure and thermal support of the surrounding rock and connected aquifer, allowing fluid to cross the boundary in response to the injection–production imbalance rather than imposing an artificial sealing wall at the edge of the modelled

domain. A fully closed (no-flow) variant, in which all six faces are sealed, is compared against this open configuration in Section B.3; the open boundaries are adopted as the base case throughout the main study.

Boundary-cell pore-volume multipliers. In addition to the open lateral faces, the storage capacity of the formation beyond the modelled window is represented by enlarging the pore volume of the outermost lateral cell rings. The multiplication is graded over the outermost two to four cell rings (depending on the multiplier magnitude), with factors decreasing inward such that no cell is scaled below its geometric pore volume; for the reference value *multpv* = 50, the four rings are scaled by factors of 28, 14, 6, and 2 (Fig. 3). Matching transmissibility multipliers are applied so that the flow resistance of the enlarged cells remains consistent with their increased volume. The graded pore-volume buffer damps the pressure response of the finite domain and, together with the Dirichlet faces, emulates the diffusive pressure support of the surrounding formation. Because the far-field storage is poorly known, the multiplier *multpv* is treated as an uncertain parameter.

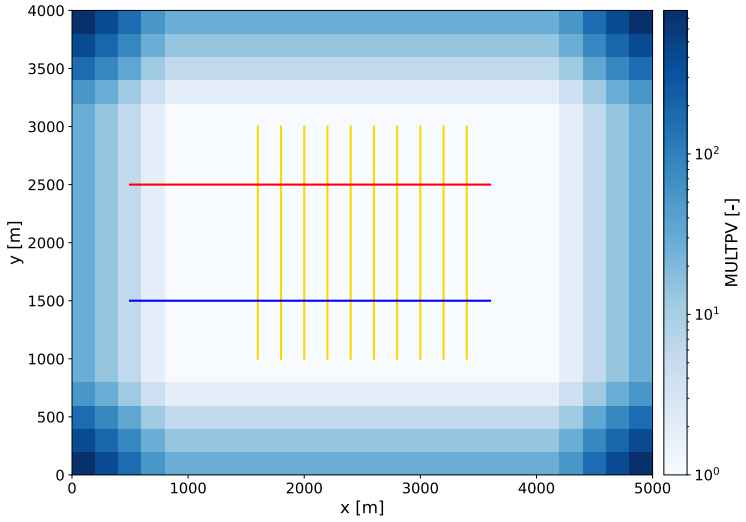


Figure 3: Plan view of the boundary-cell pore-volume multiplier field (logarithmic colour scale) for the reference value *multpv* = 50: the four outermost lateral cell rings are scaled by factors of 28, 14, 6, and 2, with corner cells carrying the product of both axes' contributions. The field is identical in every grid layer. Fractures (gold) and the horizontal well sections (blue/red) are shown for orientation.

3.6 Well Configuration

A geothermal doublet is formed by two deviated wells: an **injector** that delivers cold water into the fracture network and a **producer** that extracts the hot fluid. Both wells are initially drilled vertically and then deviated to horizontal at reservoir depth (4000 m, grid layer K = 31). The horizontal sections run parallel to the x-axis (I-direction), with the injector at J = 16 and the producer at J = 66, giving a centre-to-centre separation of 1000 m.

Each well is completed in the seven grid cells surrounding each fracture I-column, yielding 70 completion segments per well in total (7 segments × 10 fractures; see Fig. 2). The total per-fracture completion length is 30 m (cell widths 7.5 m, 5 m, 2 m, 1 m, 2 m, 5 m, 7.5 m), resulting in a total completed interval of 300 m per well.

Operating conditions are summarised in Table 3. Injection is at a fixed water rate with a maximum BHP constraint, while the producer operates under a fixed rate target with a minimum BHP constraint; both wells target the same volumetric rate at surface conditions, so the doublet is balanced. Cold water injection at 20 °C creates a thermal front that advances at different rates through each fracture depending on its permeability. A passive tracer (TR1) is co-injected to provide additional breakthrough diagnostics.

Table 3: Well operating conditions.

Parameter	Injector	Producer	Unit
Control mode	Rate	Rate	–
Target rate	10,000	10,000	$\mathrm{sm^3\,d^{-1}}$
Max BHP constraint	600	–	bar
Min BHP constraint	–	200	bar
Injection temperature	20	–	°C
Well diameter	0.2	0.2	m

3.7 Simulation Setup

Simulations are performed with OPM Flow, an open-source, industry-grade reservoir simulator that supports fully coupled thermal flow in porous media [9]. The runs use the `CO2STORE` model formulation with the `THERMAL` option. Although this formulation is designed for CO_2 storage, no CO_2 is present in this study: the gas saturation is zero at all times and the system remains single-phase brine. It is adopted nonetheless because it evaluates the brine properties described in Section 3.3 internally from analytic correlations, as functions of pressure, temperature, and salinity (specified as a molality), rather than requiring tabulated PVT input. An overview of the underlying models is given by Sandve et al. [14].

One limitation of repurposing the CO_2-storage formulation for a geothermal application is worth noting: the pure-water enthalpy correlation underlying the brine enthalpy is a polynomial fit calibrated for 0–99 °C and is extrapolated at reservoir temperatures, where it underestimates the water heat capacity by 3–5 % relative to IAPWS reference values; the resulting error in the advective heat balance and the thermal-front propagation speed is on the order of 1–2 %. Because the strain observations assimilated in this study are pressure-driven, and because the same model generates both the synthetic observations and the ensemble predictions, this does not affect the assimilation results. However, it would become relevant for thermal observations or absolute energy forecasts.

The simulation period spans 10 years (1 January 2000 to 1 January 2010). Output is written at 17 report dates, densely spaced around the onset of injection — at 5, 10, and 20 days and at the starts of February, March, April, and July 2000 — and annually thereafter. Each report date produces two kinds of output. Restart files contain the three-dimensional pressure, temperature, and density fields; the pressure fields are the input to the geomechanical forward model (Section 4). Summary vectors record the well quantities from which the scalar diagnostics of the sensitivity analysis (Section 5.2) are built — the temperature (BTEMP) and water rate (CWPR) at every completion segment, the bottom-hole pressure of both wells (WBHP), and the produced tracer concentration (WTPCTR1) — together with standard well-rate, cumulative-volume, and field-pressure diagnostics.

Numerically, each run uses 8 MPI ranks and a reduced initial time step of 0.01 d, which improves the non-linear convergence at the onset of injection and reduces the total runtime substantially.

3.8 Uncertain Parameters

Four groups of model parameters are treated as uncertain and form the state vector of the data assimilation study:

1. the effective permeability k_i of each of the ten fractures, spanning the range from near-matrix values (closed fracture) to highly conductive open pathways;
2. the effective porosity ϕ_i of each of the ten fractures;
3. the matrix permeability k_m; and

4. the boundary pore-volume multiplier *multpv* (Section 3.5), which controls the strength of the lateral pressure support.

This yields $2 \times 10 + 2 = 22$ uncertain parameters per realisation.

3.9 Reference Realisation

The synthetic reference ("truth") realisation assigns one of three discrete conductivity regimes to each of the ten fractures: a *closed* fracture with the hydraulic properties of the matrix, a *partially open* fracture, and a fully *open*, highly conductive fracture (Table 4). The class assignment (Table 5) was chosen to produce a representative heterogeneous system with a mixture of all three conductivity regimes. Note that these classes serve only to define the reference realisation; the prior ensemble samples the fracture properties continuously (Section 3.10).

Table 4: Discrete fracture conductivity regimes used to define the reference realisation. The closed regime (L) takes the matrix values.

Regime	Label	Permeability (mD)	Porosity
Closed	L	20	0.15
Partially open	M	1,600	0.18
Open	H	80,000	0.20

Table 5: Fracture regime assignment for the reference realisation. L = Closed, M = Partially open, H = Open (see Table 4).

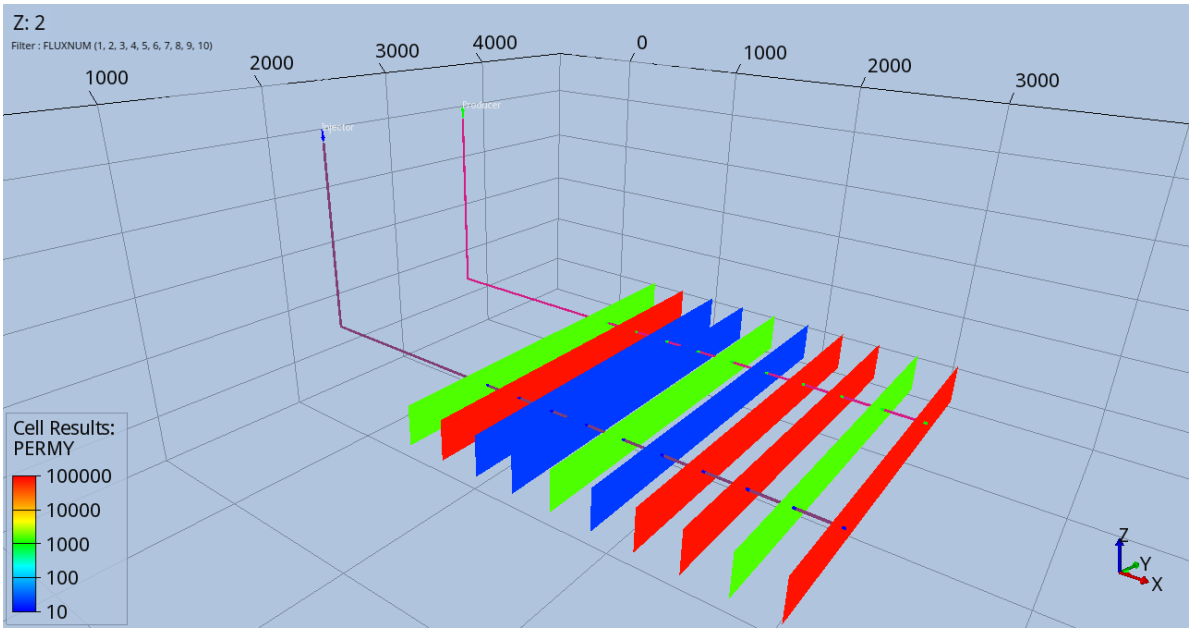
Fracture	1	2	3	4	5	6	7	8	9	10
Regime	M	H	L	L	M	L	H	H	M	H

The remaining uncertain parameters take their reference values $k_m = 20\,\text{mD}$ and $multpv = 50$, with open lateral boundaries. The reference simulation is run with OPM Flow and treated as the ground truth: the synthetic strain observations assimilated in this study are generated by applying the geomechanical forward model and the measurement operator (Section 4) to its output and perturbing the result with measurement noise (Section 6.3.1).

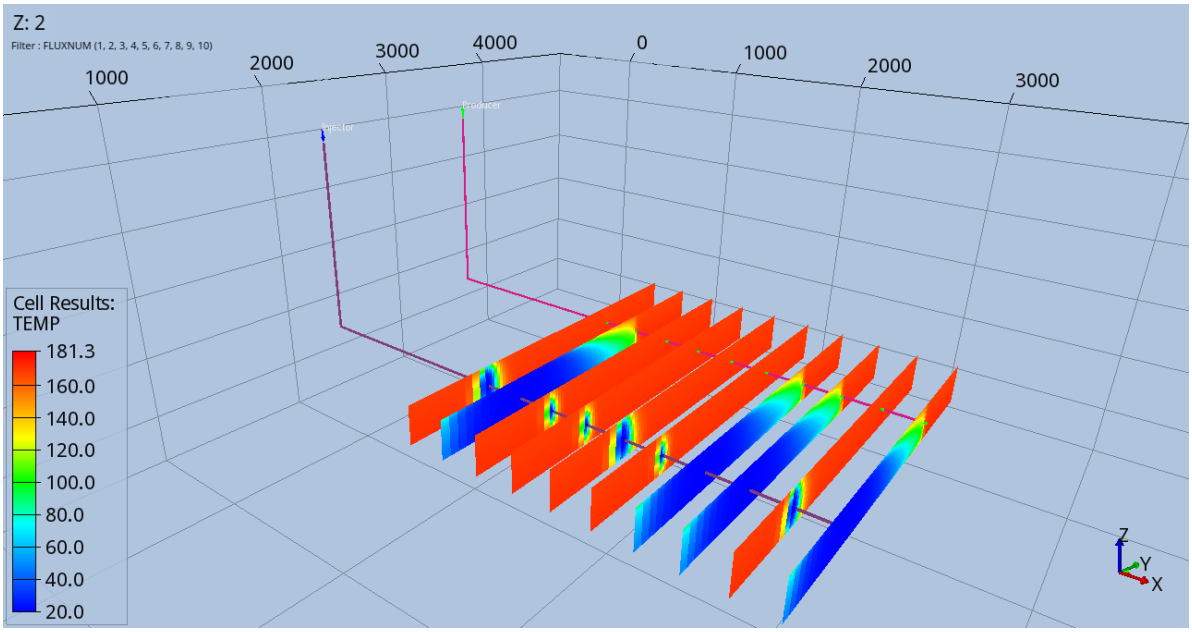
Fig. 4 visualises the reference realisation on the ten fracture planes: the permeability field of the class assignment, and the temperature at the end of the 10-year simulation. The open (H) fractures carry deep cold plumes around the injector, the partially open (M) fractures show a reduced cooled region, and the closed (L) fractures remain essentially at the undisturbed reservoir temperature — the fracture conductivity leaves a strong, spatially distinct footprint in the reservoir response. The flow behaviour of the reference realisation — its thermal front, pressure footprint, tracer transport, and per-fracture production response — is characterised in detail in Section B.2.

3.10 Prior Ensemble

The prior ensemble consists of $N_e = 100$ realisations drawn by Latin-hypercube sampling over the 22 uncertain parameters (Section 3.8), with the independent marginal distributions listed in Table 6: the permeabilities are log-uniform, spanning several orders of magnitude, the porosity is uniform, and *multpv* is drawn uniformly from its discrete set. Latin-hypercube sampling stratifies each marginal, which yields a space-filling design with negligible spurious correlation between the parameters (verified in Section C.1).



(a) Permeability (log colour scale).



(b) Temperature at the end of the simulation (1 January 2010).

Figure 4: Reference realisation displayed on the ten fracture planes (vertical exaggeration 2×), with the injector and producer trajectories. **(a)** Fracture permeability of the class assignment (Table 5): closed 20 mD (blue), partially open 1600 mD (green), open 80 000 mD (red). **(b)** Temperature after 10 years of operation: the cold-water plume (blue, down to the injection temperature of 20 °C) develops along the conductive fractures, while closed fractures stay near the initial reservoir temperature.

Table 6: Prior distributions of the 22 uncertain parameters.

Parameter	Distribution	Min	Max
Fracture permeability k_i (mD)	log-uniform	10	80,000
Fracture porosity ϕ_i (–)	uniform	0.10	0.25
Matrix permeability k_m (mD)	log-uniform	0.5	50
Boundary multiplier *multpv* (–)	uniform on $\{10, 20, \ldots, 80\}$	10	80

Fig. 5 shows the marginal prior distributions of all 22 parameters, together with the values of the reference realisation (Section 3.9). The spread of reservoir responses spanned by this prior — and its reduction by the data assimilation — is documented in Section B.4.



Figure 5: Marginal prior distributions of the 22 uncertain parameters ($N_e = 100$ realisations), one panel per parameter: fracture permeabilities k_i and matrix permeability k_m (both in $\log_{10}$ space), fracture porosities ϕ_i, and the pore-volume multiplier *multpv*. Red vertical lines mark the values of the reference realisation (Section 3.9).

4 Geomechanical Forward Model and Measurement Design

As injection and production change the pore pressure in the reservoir, the resulting expansion and contraction of the rock deforms the surrounding formation. This section describes the geomechanical forward model that maps the simulated pressure field to strain at arbitrary observation points (Section 4.1), the strain response of the reference realisation along the wells and the selection of the measurement points that define the observation vector of the data assimilation study (Section 4.2).

4.1 Poro-Elastic Forward Model

The mechanical response is computed with a nucleus-of-strain approach in a homogeneous, isotropic, linear-elastic half-space with Young's modulus $E = 3300\,\mathrm{MPa}$ and Poisson's ratio $\nu = 0.25$. The open-source implementation of Barbosa *et al.* [10] is used, which relies on exact analytical expressions for the poro-elastic forcing acting on a prism, which have been derived by [15, 16] for the analogous solution of the gravitational potential. The source of deformation is the pore-pressure change of each reservoir grid cell relative to the undisturbed initial state t_{ref},

$$\Delta p_c(t) = p_c(t) - p_c(t_{\mathrm{ref}}). \tag{1}$$

Each grid cell is represented as a rectangular prism carrying its uniform overpressure Δp_c; the displacement field at an arbitrary observation point is the superposition of the analytical prism solutions over all cells. Strain components are obtained from the displacement field by central finite differences with a step of 0.2 m, small compared to the minimum grid cell size of 1 m.

Two properties of this formulation are worth noting. First, the coupling is one-way: the pressure field drives the mechanical response, but the deformation does not feed back on the flow simulation. Consequently, the coupling assumes a drained poro-elastic response. Second, only pore-pressure changes act as sources; thermoelastic stresses caused by the cooled region around the injector are neglected at this stage.

The forward model is validated by Barbosa *et al.* [10] and validated for a pressurised finite sphere in an elastic half-space [17] in a separate internal study [18], which also demonstrates the full inversion workflow on that conceptual geometry.

4.2 Reference Strain Response and Measurement Locations

In principle, DSS provides strain measurements distributed along the entire sections of both wells, with each measurement representing an average over a finite gauge length. However, the EVR and ES-MDA methodologies employed in the following sections require a measurement-error covariance matrix describing both the uncertainty of each measurement and the correlations between their errors. These correlations are difficult to constrain accurately for densely sampled DSS data. If closely spaced measurements are incorrectly treated as independent, the repeated information can cause the uncertainty reduction to be overestimated. To mitigate this effect, strain is sampled at selected locations along the injection and production wells that capture the principal deformation caused by the reservoir pressure response.

To select suitable measurement locations, the thermal-hydraulic response of the realisation chosen as the synthetic truth is first simulated with OPM Flow. The resulting reservoir pressure changes are then passed to the geomechanical forward model to predict displacement and strain along both wells. Fig. 7a shows the horizontal displacement (u_x) and axial strain (ε_{xx}) forecasted along the horizontal well sections, with a zoom view between two fractures shown in Fig. 7b. On the other-hand Fig. 6 shows the vertical displacement (u_z) and vertical strain (ε_{zz}) forecasted along the vertical sections of the well.

The open circles in these figures (Figs. 6 and 7) indicate the locations selected for the synthetic strain measurements. Along each horizontal well section, (ε_{xx}) is sampled at 36 locations, with three measurements

between successive fractures and a coarser spacing towards the ends of the well. Along each vertical section, (ε_{zz}) is sampled at nine depths, using a coarse spacing through the overburden with additional measurements near the reservoir. Across the injector and producer, this gives 72 horizontal and 18 vertical strain measurements, forming a 90-component observation vector. All measurements are evaluated two years after the start of injection and relative to the pre-injection strain state.

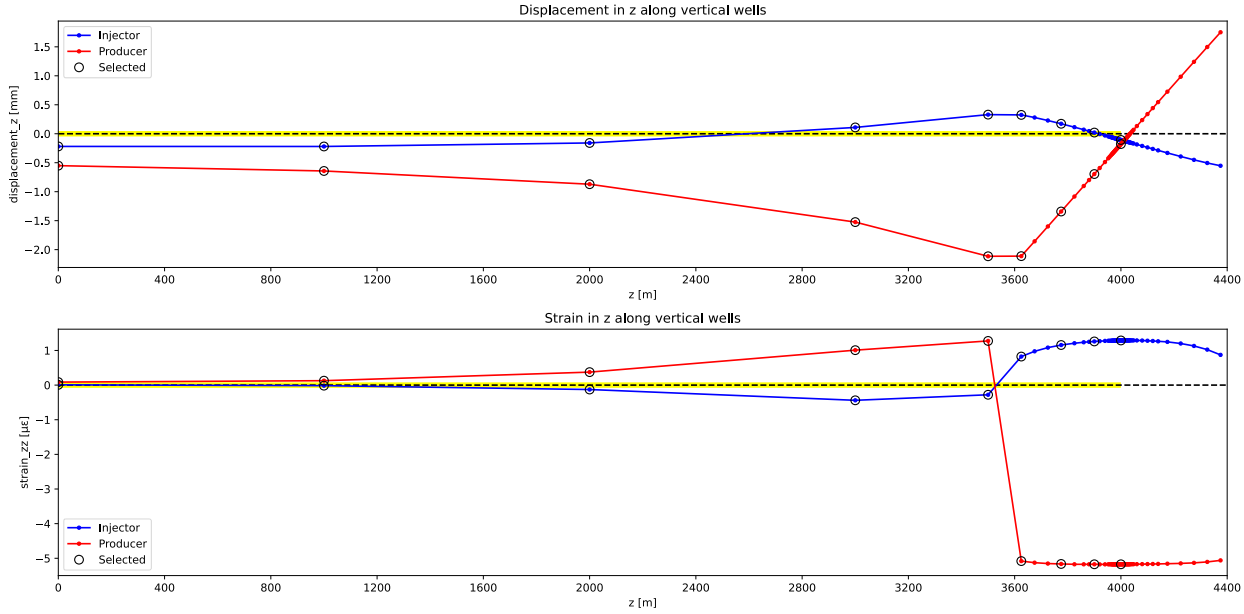
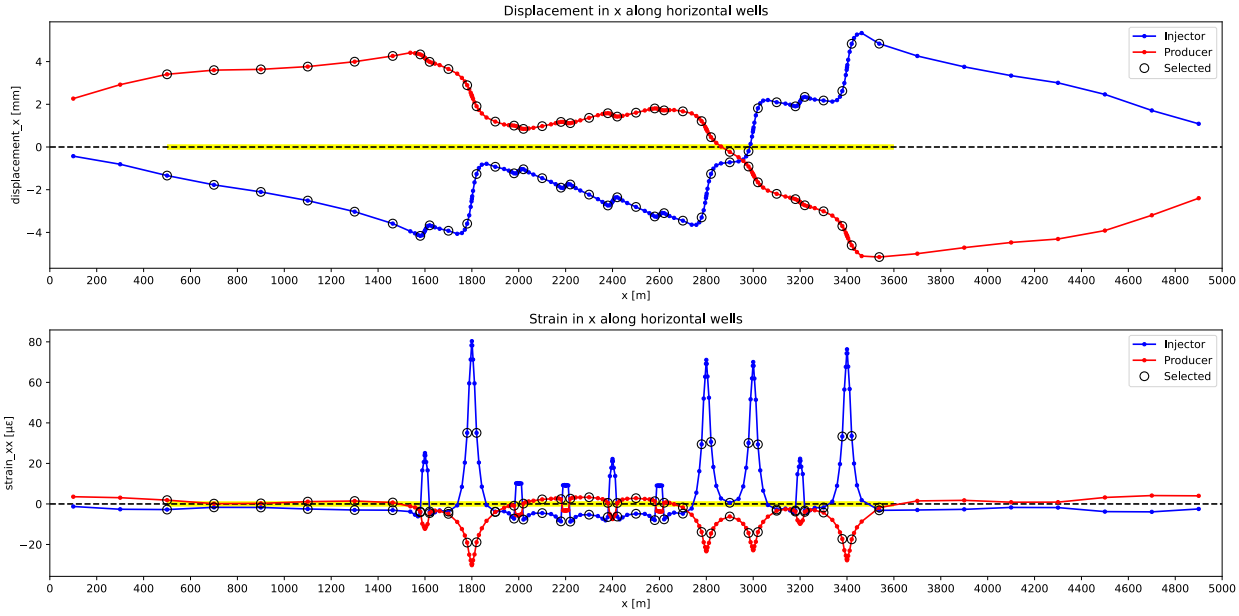
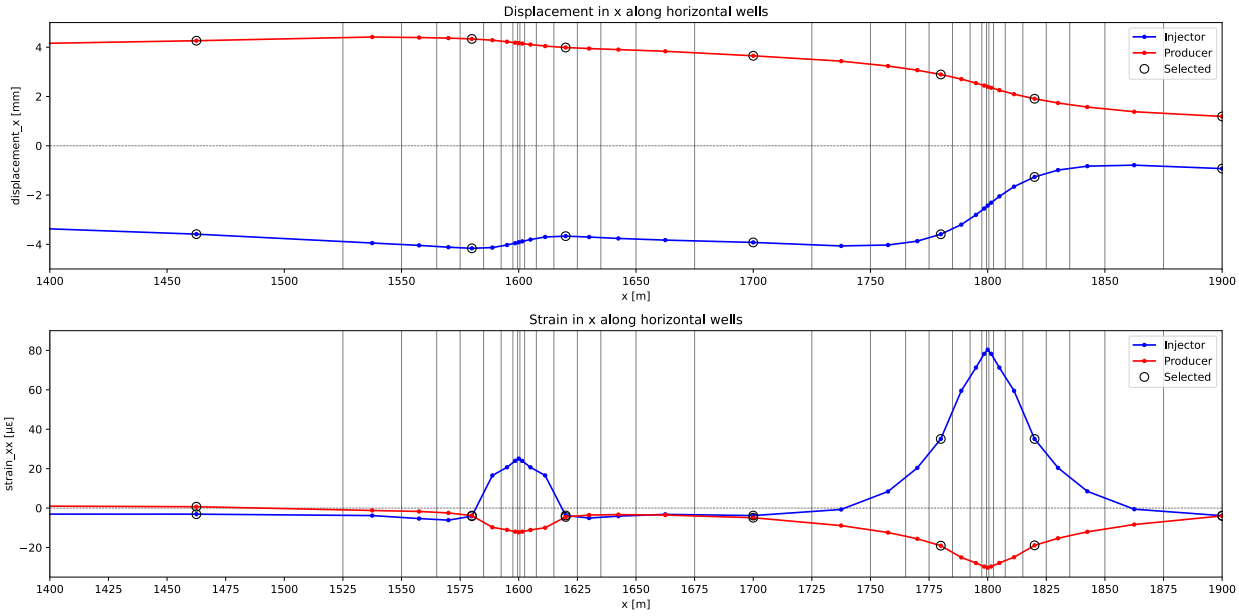


Figure 6: Forecasts of the vertical displacement u_z (top panel) and vertical strain ε_{zz} (bottom panel) along the vertical injector (blue) and producer (red) sections at $x = 500$ m for the reference realisation, after 10 years of operation. Dots mark the grid cell centres, black open circles the selected measurement points and the yellow bar marks the extent of the wellbore.



(a) Full well length.



(b) Zoom onto the first two fractures at $x = 1600$ and 1800 m.

Figure 7: Axial displacement u_x (top panels) and axial strain ε_{xx} (bottom panels) forecasted along the horizontal injector (blue) and producer (red) sections at 4000 m depth for the reference realisation, after 10 years of operation. Dots mark the grid cell centres, black open circles the selected measurement points; the yellow bar marks the extent of the wellbore, and the thin grey lines in (b) the grid cell boundaries.

5 Scoping the Expected Value of Deformation Measurements

Before assimilating any data, we first assess what deformation measurements are expected to reveal about the uncertain reservoir properties. By using the prior ensemble of $N = 100$ realisations simulated with the thermo-hydraulic and geomechanical models, two types of measurement responses are analysed: strain distributed along the injection and production wells, and displacement throughout the reservoir and overburden. Ensemble-variance analysis (EVA) is then used to estimate the uncertainty reduction that individual strain and displacement measurements could provide, given their assumed measurement errors. Finally, the same analysis is applied to the complete 90-component strain survey used in the subsequent inversion to estimate its expected value for constraining the reservoir parameters and improving pressure, production and thermal-power forecasts. These estimates are prospective: they describe what the measurements are expected to achieve before any observations are assimilated. In other-words, no ES-MDA inversion is performed, as all statements in this section are derived from the prior ensemble alone. Their predictions are later compared in Section 7 with the uncertainty reduction actually realised by the ES-MDA inversion.

Throughout the section, all observables are evaluated at 2 year forecasts after the start of injection/production of the geothermal site, with the exception of Section 5.3.3 varies the measurement time.

5.1 Deformation Response Across the Prior Ensemble

Before estimating their expected value, we first examine the deformation responses predicted across the prior ensemble. This establishes whether the strain and displacement signals are sufficiently large relative to their assumed measurement errors and whether they vary enough between realisations to distinguish different reservoir properties.

5.1.1 Strain Along the Wells

Fig. 8 shows the strain predicted along the vertical and horizontal sections of the injection and production wells for the 100 prior realisations, together with the prior mean and the selected synthetic truth. Along the horizontal sections, the strain response is concentrated around the fracture crossings and varies substantially across the ensemble. The ensemble spread at these locations is considerably larger than the assumed measurement error of 5 $\mu\varepsilon$, indicating that the measurements should be able to distinguish between reservoir realisations. The synthetic truth also remains within the prior range at the selected measurement locations, confirming that the prior ensemble adequately spans the strain response used in the subsequent inversion.

The strain distribution also indicates which parts of the well are sensitive to different aspects of the reservoir response. Measurements close to the fracture crossings respond strongly to the hydraulic behaviour of the neighbouring fractures, while the smoother strain response away from the fractures reflects the broader reservoir-pressure change. Along the vertical well sections, the largest strain variations occur near the reservoir, whereas the response decreases through the overburden and becomes small towards the surface. The horizontal and vertical measurements therefore provide complementary information about the local fracture response and the reservoir-scale pressure field.

5.1.2 Displacement Through the Reservoir and Overburden

The broader deformation response of the prior ensemble is examined using the predicted vertical displacement u_z, which is the displacement component most readily observed through surface geodetic measurements. Fig. 9 shows the reference response, prior mean and prior standard deviation at the ground surface and at a depth of 3500 m, immediately above the reservoir.

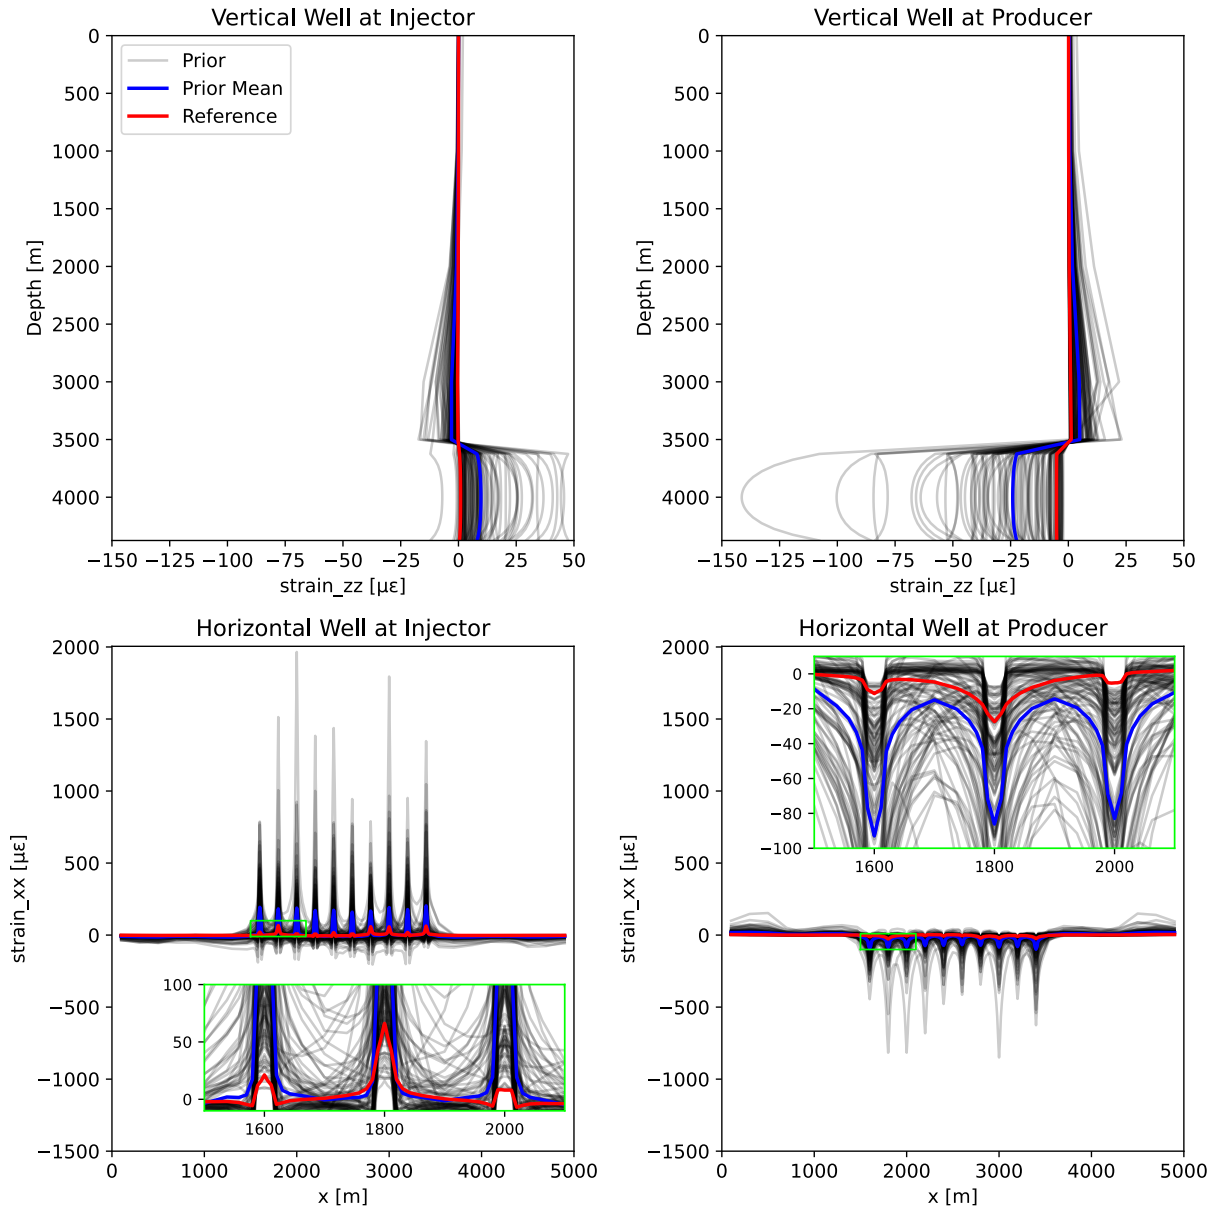
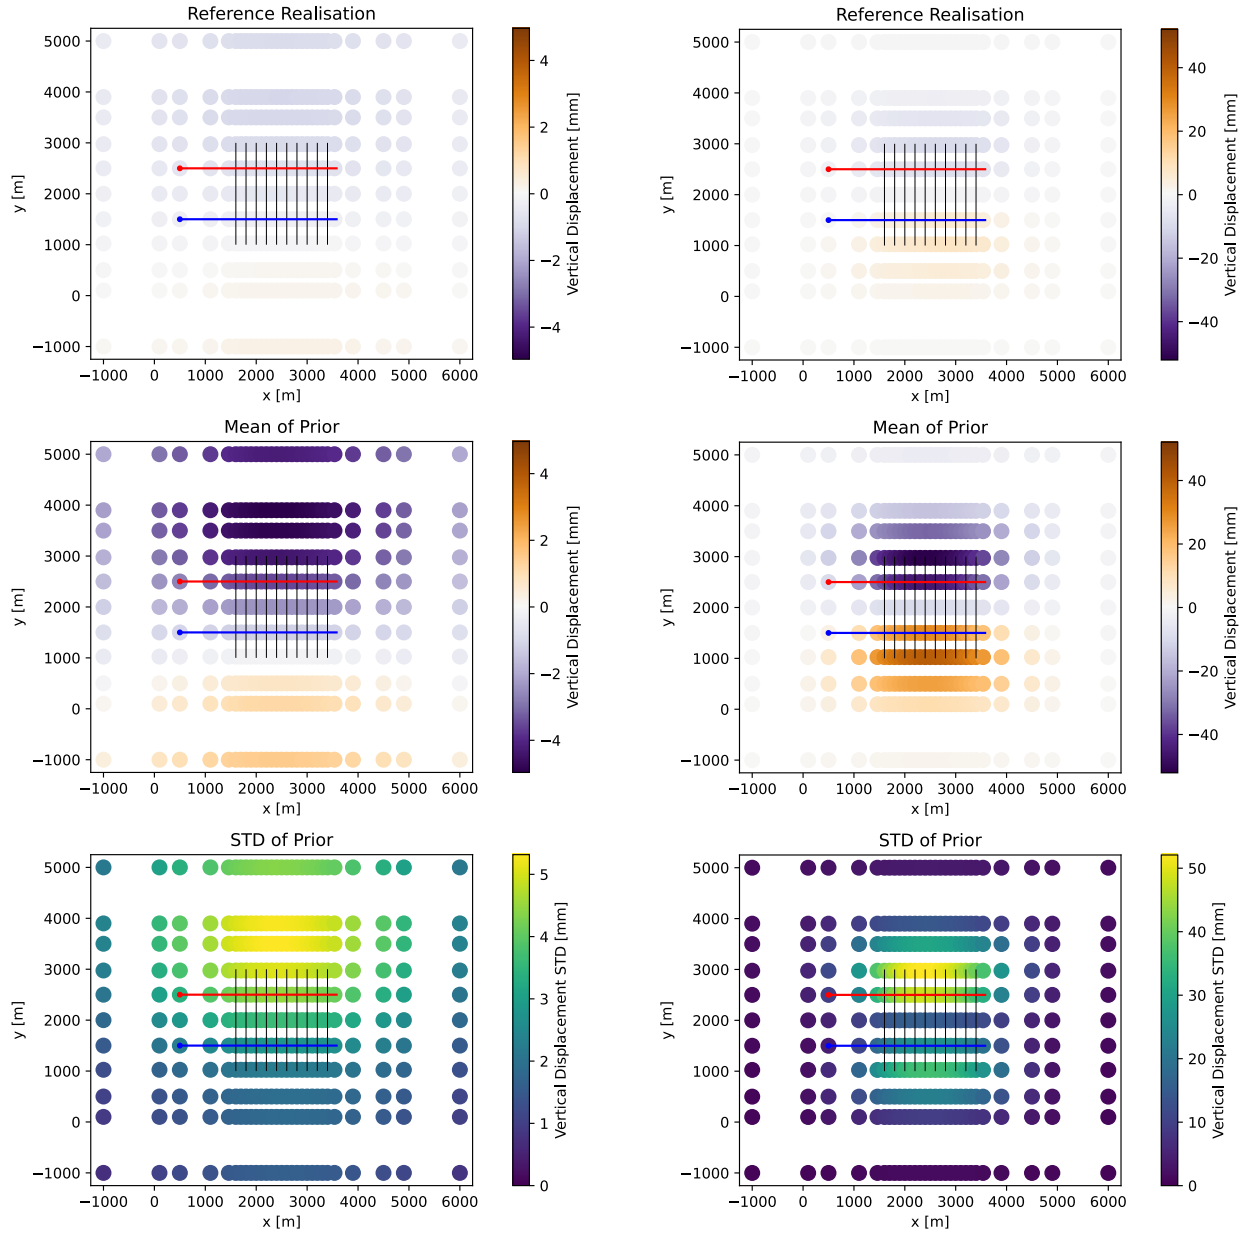


Figure 8: Prior ensemble spread of the strain observables at the measurement date (1 January 2002): vertical strain ε_{zz} over depth along the vertical well sections (top) and axial strain ε_{xx} along the horizontal sections (bottom), for injector (left) and producer (right). Grey: prior realisations; blue: prior mean; red: reference. The insets zoom onto the reference response around the first three fractures (x = 1600, 1800 and 2000 m), which the full-scale prior spread would otherwise hide.

At both depths, injection produces uplift around the injector, while pressure drawdown produces subsidence around the producer. Immediately above the reservoir, the predicted displacement and its variability reach several tens of millimetres and remain concentrated around the fractured region. As the deformation propagates through the overburden, the response becomes smoother and decreases substantially in magnitude. At the ground surface, the prior mean and standard deviation are reduced to a few millimetres, while the selected synthetic truth produces a displacement of approximately 1 mm. The standard deviation in surface displacement approaches the limit of detectability of geodetic measurements such as InSAR. This suggests the intuitive result that it is difficult to distinguish between reservoir realisations with different fracture parameters (at 4km depth) from the surface. Horizontal displacement and tilt at the surface are not analysed here.



(a) Ground surface.

(b) At 3500 m depth.

Figure 9: Ensemble forecast of vertical displacement u_z at the measurement date (after 2 years of injection/operation), relative to the pre-injection baseline, on the observation grid: reference realisation (top), prior ensemble mean (middle), and prior ensemble standard deviation (bottom). Blue/red: horizontal injector/producer sections; black: fracture positions. Note the ten-fold difference in colour scale between the two levels.

5.2 Parameter Controls on Deformation and Reservoir Performance

Before quantifying the expected uncertainty reduction, the prior ensemble is used to identify which reservoir properties control the deformation measurements and the principal indicators of reservoir performance. Fig. 10 shows the Pearson correlation between the 22 uncertain parameters and twelve representative observables. These comprise the bottom-hole pressures, produced tracer concentration, thermal-power output, surface displacement, and selected strain measurements along the injection and production wells. The correlations are ranked separately for each observable so that the parameters exerting the strongest influence can be identified.

Matrix permeability k_m exerts the strongest control on the observables, including bottom-hole pressure, surface displacement and much of the strain response. It also influences tracer breakthrough and thermal-power output because it controls the proportion of fluid transported through the slower matrix pathways rather than through the fractures. The shared dependence of deformation and reservoir performance on k_m indicates that deformation measurements can provide information about the pressure response and its consequences for production.

Individual fracture permeabilities are most clearly expressed in the horizontally distributed strain measurements. Strain measured at a fracture crossing is sensitive to the permeability of the neighbouring fracture, while measurements between fractures respond most strongly to the adjacent fractures. This local sensitivity is not apparent in the bottom-hole pressures or surface displacement, which primarily reflect the reservoir-scale pressure response. Fracture porosity and the boundary multiplier are only weakly correlated with the selected measurements, suggesting that they will be more difficult to constrain.

The tornado plots therefore indicate that deformation measurements may constrain both the overall pressure response and the hydraulic behaviour of individual fractures. However, correlation alone does not account for measurement accuracy or the combined information provided by multiple observations. The expected uncertainty reduction associated with the individual deformation measurements and the complete strain survey is quantified using EVA in the following subsection.

5.3 Expected Variance Reduction from Deformation Measurements

The correlation analysis identifies which reservoir properties influence each measurement, but it does not account for measurement error or quantify the uncertainty reduction that could be achieved. Ensemble-variance analysis (EVA) provides this estimate directly from the prior ensemble [4]. The key parameter which EVA estimates is the posterior variance $\sigma^2_{J|D}$ of a quantity J if a measurement D_{obs} is performed of the random observation vector D. The parameter J could be anything from the unknown fracture properties or the unknown thermal output of the well, while D could be any combination of measurements. The key requirement in this analysis is that the joint distribution between J and D follow the multi-variate normal distribution:

$$\begin{bmatrix} J \\ \mathbf{D} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_J \\ \boldsymbol{\mu}_D \end{bmatrix}, \begin{bmatrix} \sigma_J^2 & \mathbf{C}_{JD} \\ \mathbf{C}_{DJ} & \mathbf{C}_{DD} + \mathbf{R} \end{bmatrix} \right), \tag{2}$$

where μ_J and σ_J^2 are the prior mean and variance of the quantity of interest. The vector $\boldsymbol{\mu}_D$ and matrix $\mathbf{C}_{DD}$ are the prior mean and covariance of the noise-free modelled observations, estimated from the ensemble. The cross-covariance $\mathbf{C}_{JD}$ describes the linear statistical relationship between the quantity of interest and the modelled observations, with $\mathbf{C}_{DJ} = \mathbf{C}_{JD}^{\mathsf{T}}$. Finally, $\mathbf{R}$ is the measurement-error covariance matrix. The measurement errors are assumed to be zero-mean, Gaussian and independent of the model parameters and predictions.

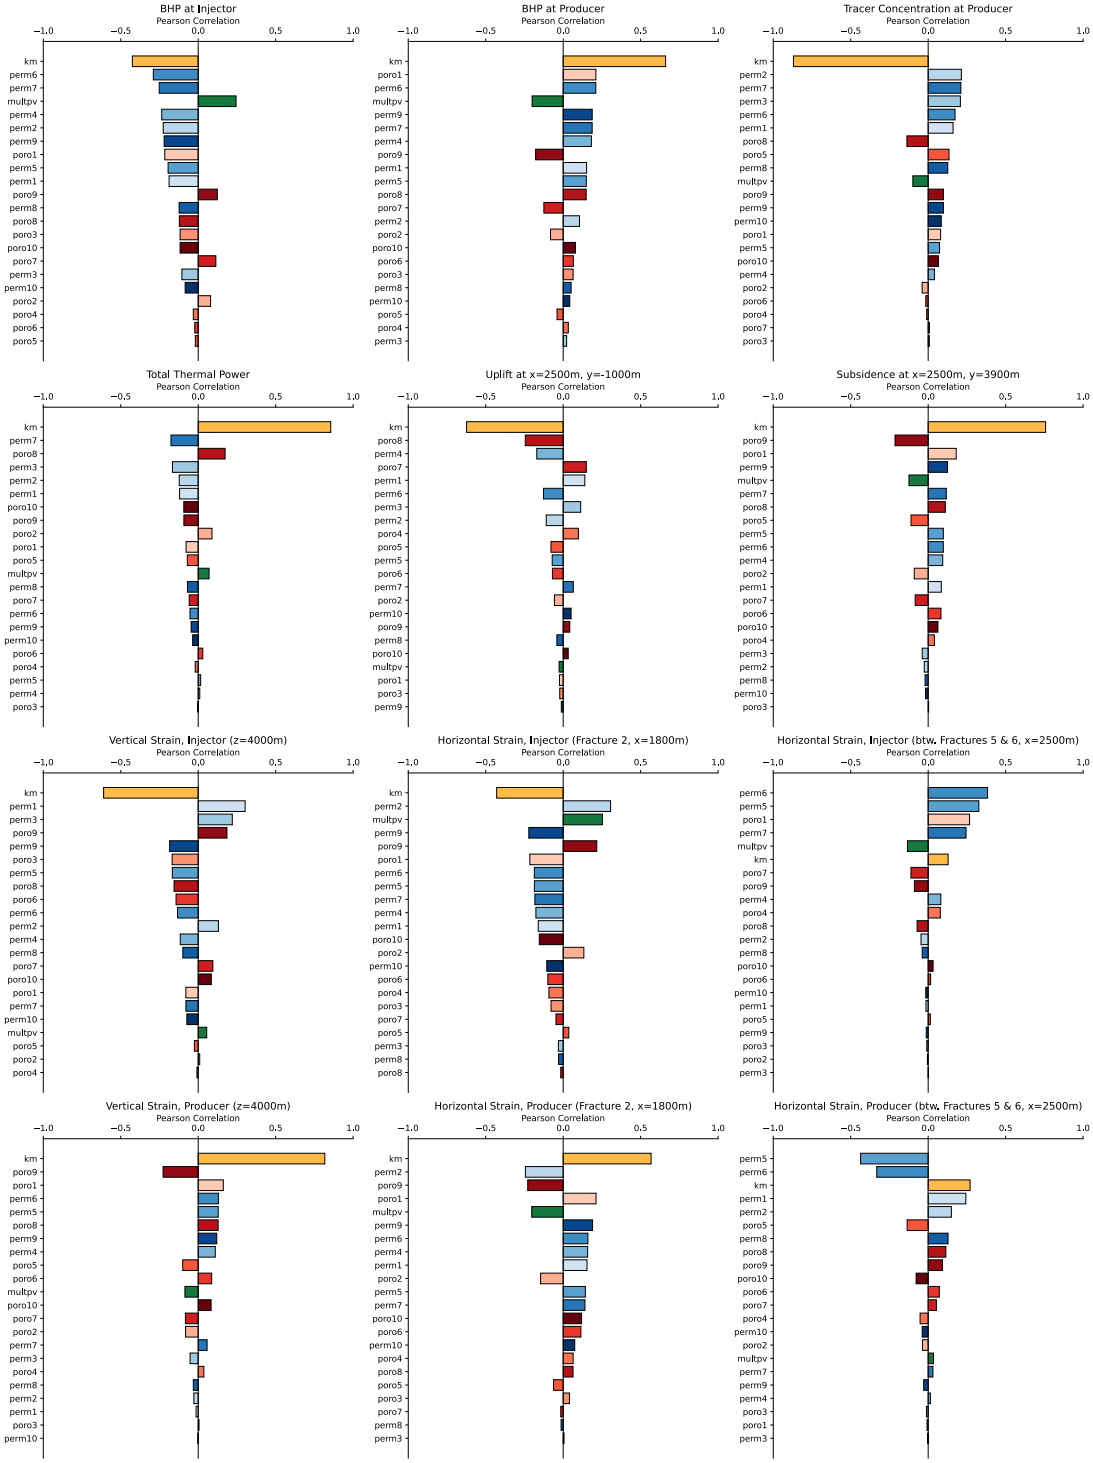


Figure 10: Correlation between the 22 uncertain parameters and twelve scalar measurements after 2 years of injection/production, evaluated across the prior ensemble and ranked by magnitude. Colours: fracture permeabilities (blues), fracture porosities (reds), matrix permeability k_m (orange), boundary multiplier *multpv* (green).

For a measurement $D = D_{obs}$, the posterior distribution of $J|D_{obs} \sim N\left(\mu_{J|D}, \sigma^2_{J|D}\right)$, where the posterior variance $\sigma^2_{J|D}$ can be analytically derived to be

$$\sigma^2_{J|D} = \sigma^2_J - \mathbf{C}_{JD}\left(\mathbf{C}_{DD} + \mathbf{R}\right)^{-1}\mathbf{C}_{DJ}\,. \tag{3}$$

Note, such an expression is independent to the value of D_{obs}, and therefore can be forecasted before the measurement is made.

Finally, the expected variance reduction (EVR) of the quantity J when observing D is simply the fraction of the remaining variance relative to the prior variance:

$$\mathrm{EVR}(J;D) = \frac{\sigma^2_J - \sigma^2_{J|D}}{\sigma^2_J}\,, \tag{4}$$

The EVA analysis first considers individual displacement and strain measurements to determine how their expected value varies with location. It then evaluates the complete 90-component strain survey used in the subsequent inversion. Measurement errors of 0.5 mm for displacement and 5 $\mu\varepsilon$ for strain are assumed. In the individual-measurement figures, the maximum EVR calculated for the ten fracture permeabilities and ten fracture porosities is used. The figures therefore show the fracture property that is best constrained at each measurement location, rather than the average reduction across all fractures.

5.3.1 Expected Value of an Individual Displacement Measurement

Fig. 11 shows the EVR obtained from a single displacement measurement immediately above the reservoir. The calculation is repeated for each displacement component and for the matrix permeability, fracture permeabilities, fracture porosities and boundary multiplier.

A single displacement measurement is principally informative about matrix permeability. Over much of the domain, one measurement is expected to remove approximately 60 – 75% of its prior variance. The information decreases around locations where the relevant displacement component changes sign and its amplitude approaches zero. Nevertheless, the broad spatial extent of the response indicates that the precise measurement location is relatively unimportant when the objective is to constrain the reservoir-scale pressure response.

Displacement above the reservoir is much less informative about individual fracture properties. Immediately above the reservoir, the best-positioned measurement removes at most approximately 20% of the variance of one fracture permeability, while the expected reduction at the ground surface is generally below 5%. Fig. 12 shows how the three components of the displacement field (u_x, u_y, u_z) reduce the posterior variance for different points in space (x, y, z). Fracture porosity and the boundary multiplier are also weakly constrained. Displacement measurements therefore provide useful information about the overall pressure response, particularly through matrix permeability, but offer limited ability to distinguish the hydraulic properties of individual fractures.

5.3.2 Expected Value of an Individual Strain Measurement

The expected value of a single strain measurement depends more strongly on its position along the wells. Fig. 13 shows the EVR of axial strain ε_{xx} along the horizontal well sections and vertical strain ε_{zz} along the vertical sections.

A single strain measurement can remove up to approximately 60% of the matrix-permeability variance. Along the horizontal wells, the most informative locations for matrix permeability occur outside the fractured interval, where the smoother strain response is controlled primarily by the broader reservoir-pressure change.

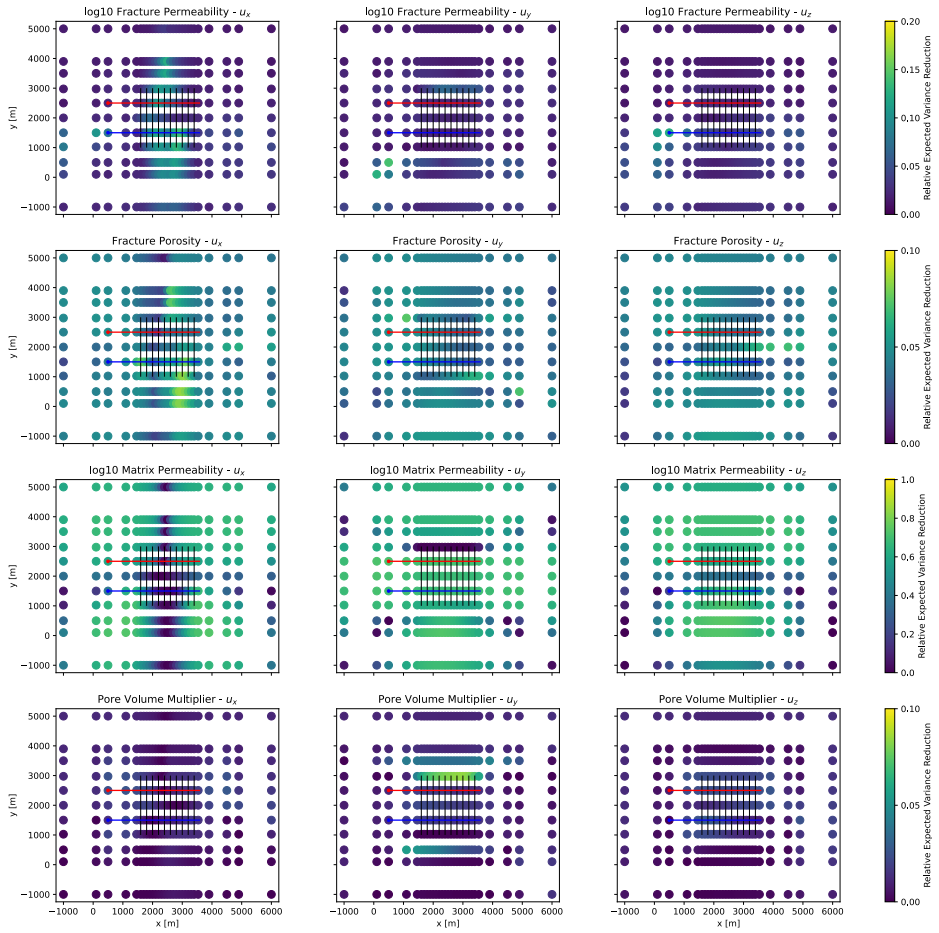


Figure 11: Relative expected variance reduction (EVR) from a single displacement measurement immediately above the reservoir, shown for each displacement component and parameter group. For the fracture permeabilities and porosities, the maximum EVR within each group is shown. Note that the colour scales differ between parameter groups.

Along the vertical sections, the information generally increases towards the reservoir and is greater at the producer than at the injector.

The expected value for fracture permeability shows the opposite spatial pattern. Measurements near the fracture crossings are expected to remove up to approximately 50% of the variance of the neighbouring fracture permeability at the injector and 40% at the producer. Away from the fractured interval, their value for individual fracture properties falls substantially. Fracture porosity remains weakly constrained at all locations, while the boundary multiplier is only modestly informed. Overall, the results are promising in that they suggest strain measurements near the fracture crossings have a high potential to resolve differences between individual fracture permeabilities. This local sensitivity motivates the distributed strain layout used in the inversion.

5.3.3 Expected Variance Reduction of Distributed Strain Sensing

The preceding results consider one measurement at a time. The proposed monitoring layout instead contains 90 strain measurements distributed along the vertical and horizontal sections of both wells. Fig. 14 shows the combined expected variance reduction for the 22 uncertain reservoir parameters and how this information develops with measurement time.

When the measurements are considered jointly, their expected value is substantially greater than that of any individual observation. The matrix permeability is expected to be almost fully constrained, with an EVR

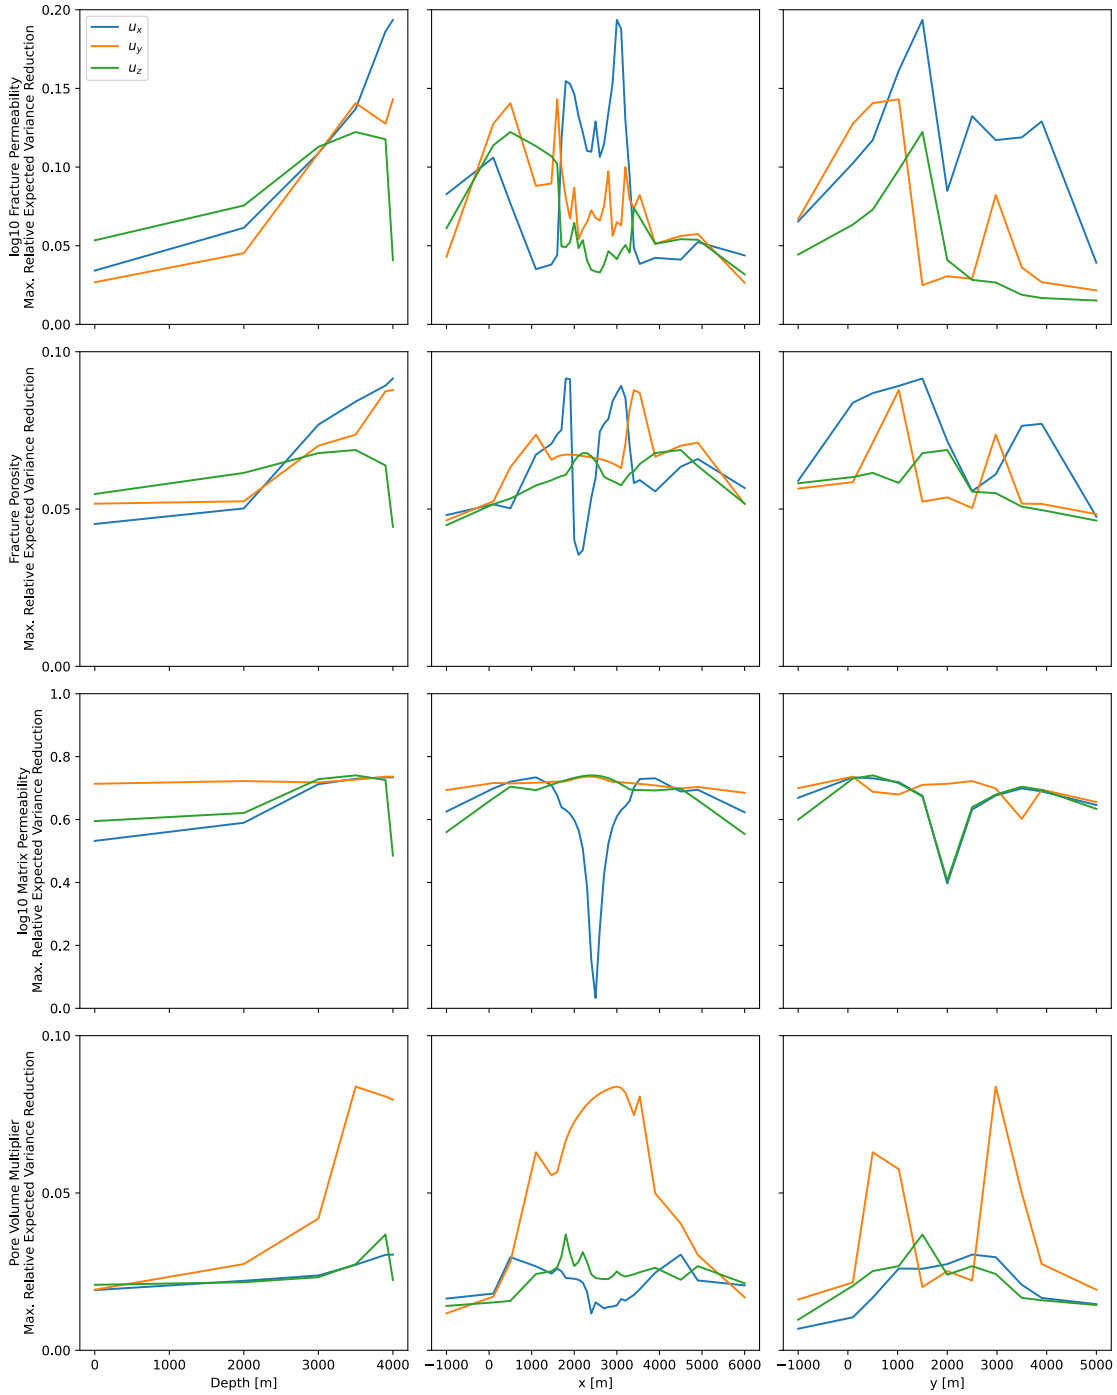


Figure 12: Maximum relative EVR of a single displacement measurement over the three-dimensional observation grid, collapsed onto profiles along depth (left), x (centre), and y (right), per parameter group (rows; for fracture permeabilities and porosities the maximum over the group) and displacement component (colours). Note the axis scales: $[0, 1]$ in the matrix-permeability row, $[0, 0.2]$ for the fracture permeabilities, and $[0, 0.1]$ for the porosities and *multpv*.

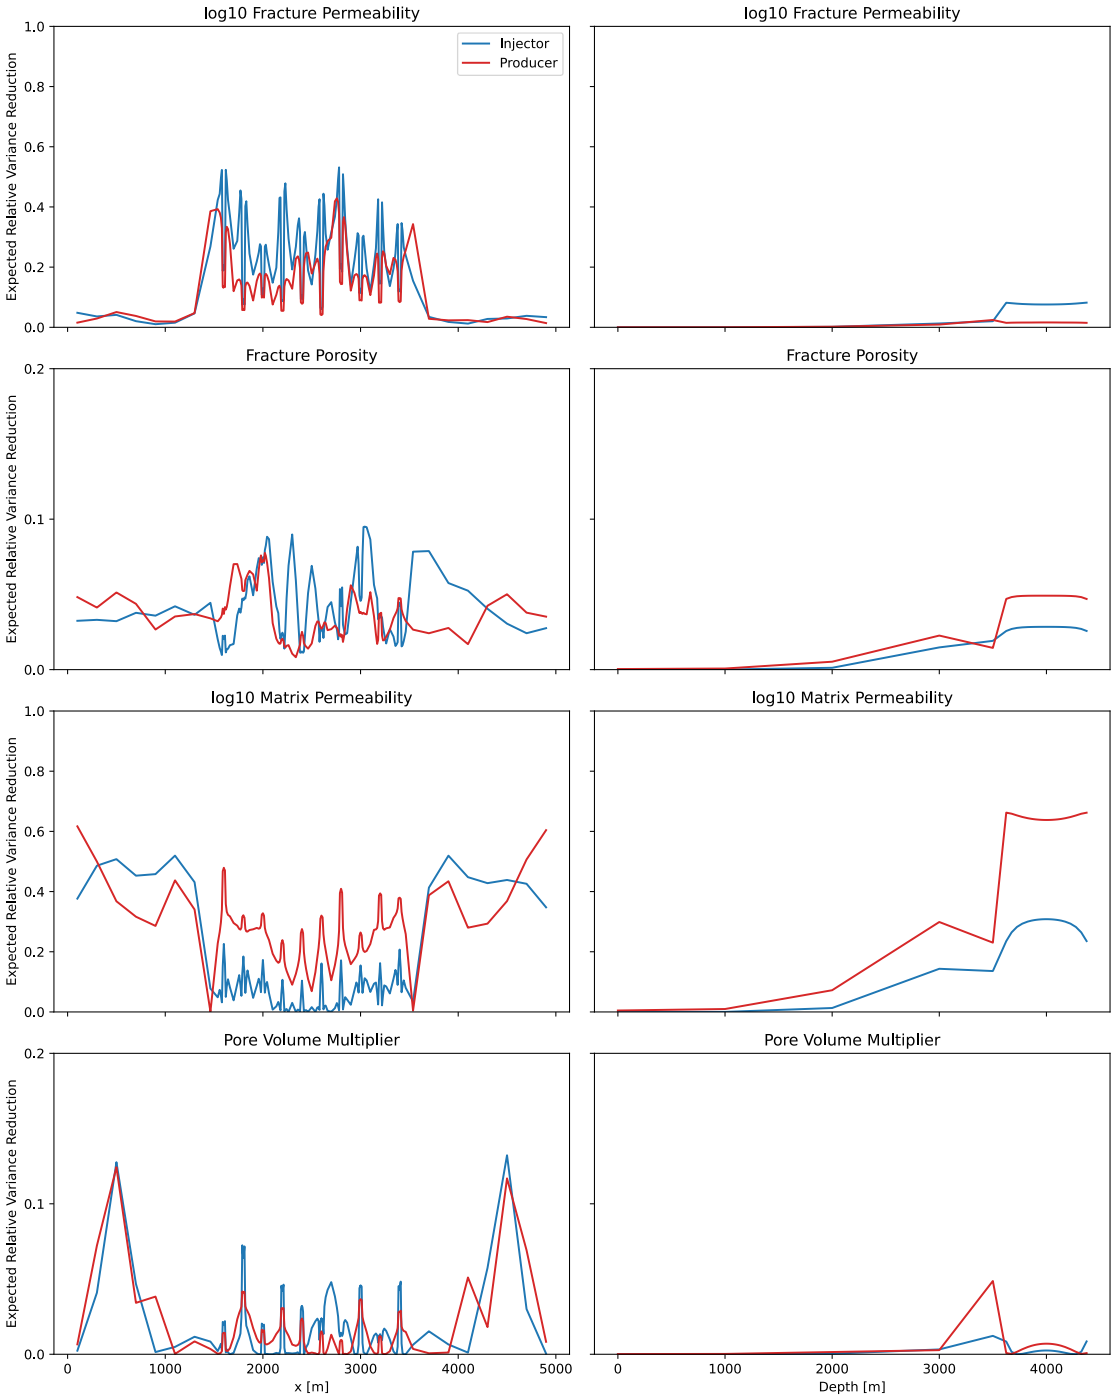


Figure 13: Relative EVR of a single strain measurement after 2 years of injection/production: ε_{xx} along the horizontal sections (left) and ε_{zz} along the vertical sections (right). For the fracture permeabilities and porosities, the maximum EVR within each group is shown. Note that the axis scales differ between parameter groups.

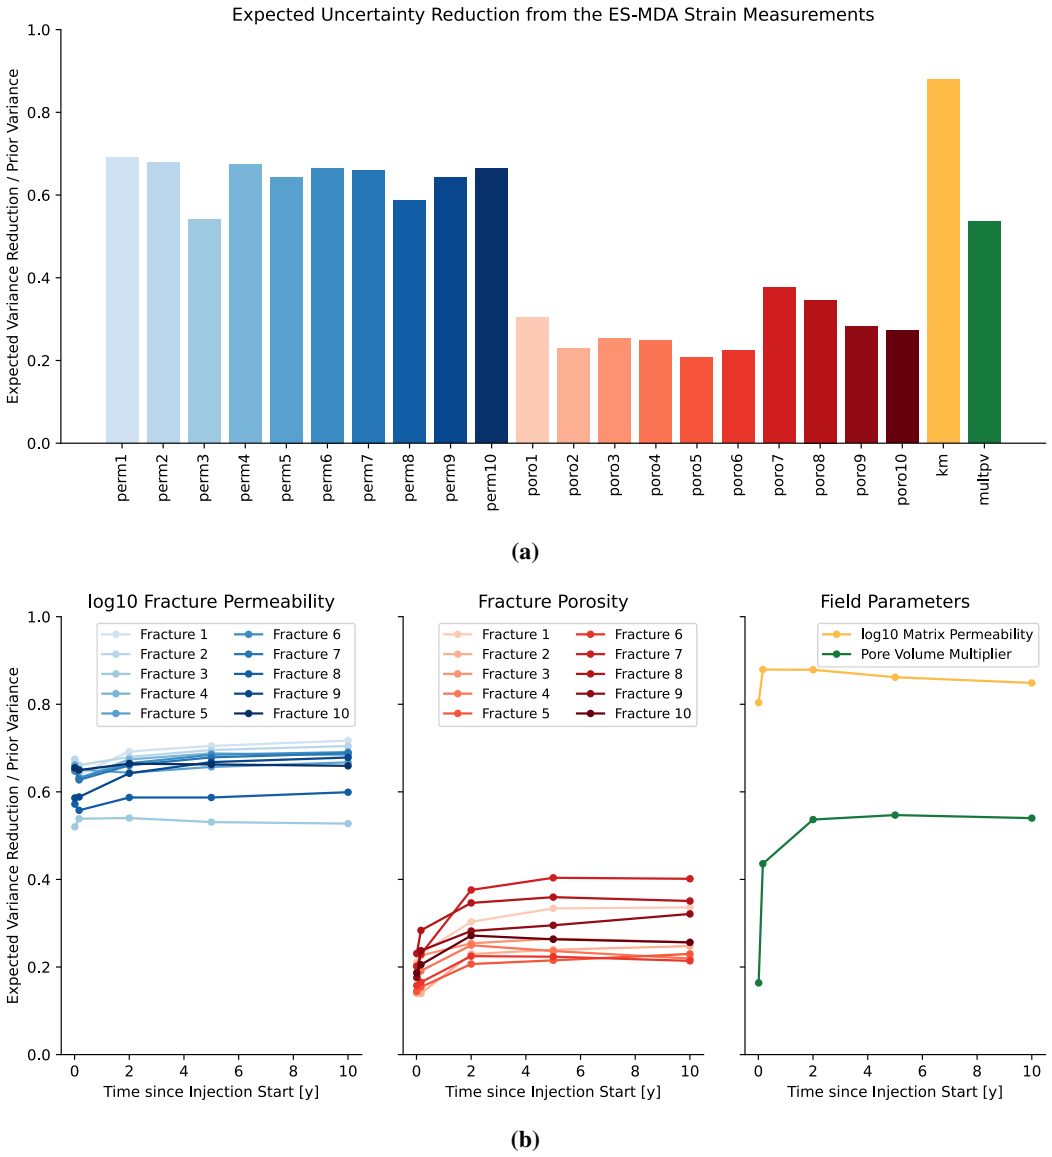


Figure 14: Expected variance reduction of the 22 uncertain parameters from the complete 90-component strain vector of the ES-MDA study with the 5 $\mu\varepsilon$ measurement error, relative to the prior variance: (a) after 2 years of injection/production, (b) as a function of the measurement time.

of 0.88. The ten fracture permeabilities have EVRs between 0.54 and 0.69, indicating that the distributed measurements can constrain all fractures rather than only the fracture nearest to a particular measurement. The boundary multiplier reaches an EVR of 0.54, whereas the fracture porosities remain the least constrained parameters, with EVRs between 0.21 and 0.37.

Interestingly, most of the pressure-related information develops early. Measurements obtained after only five days are already expected to remove approximately 80% of the matrix-permeability variance and 50 – 65% of the fracture-permeability variance. The information about fracture porosity and the boundary multiplier develops more gradually as the pressure perturbation propagates through the reservoir and interacts with the model boundaries. For all parameter groups, the expected uncertainty reduction changes little after approximately two years. The measurement time selected for the subsequent ES-MDA inversion therefore captures most of the information available from the pressure-driven strain response. However, effects such as temperature-induced deformations, and stress-dependent permeability are expect to effect such a conclusion.

Constraining the uncertain reservoir parameters should also reduce uncertainty in the predicted reservoir performance. Fig. 15 shows the expected variance reduction in the injector-producer pressure difference,

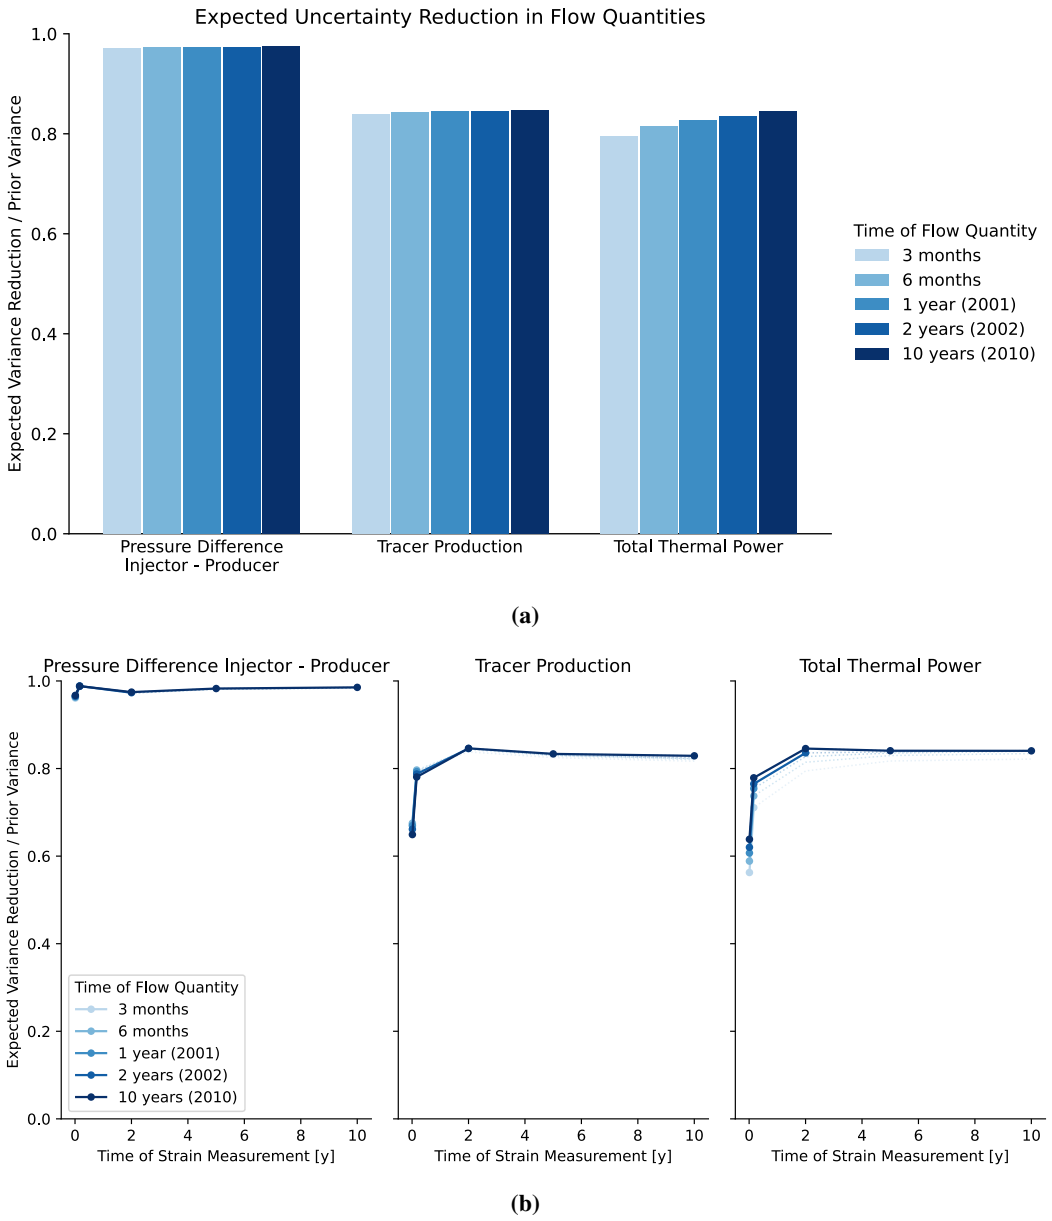


Figure 15: Expected variance reduction of the flow quantities from the 90-component strain vector, relative to the prior variance: (a) for strain measured at two years, with the flow quantities evaluated at five times between three months and ten years; (b) as a function of the strain measurement time, per flow-quantity time.

produced tracer concentration and total thermal-power output. The strain survey is expected to remove approximately 97% of the variance in the well's bottom hole pressure. The predicted reductions in tracer-concentration and thermal-power uncertainty are smaller, at approximately 80%, because these quantities also depend on transport properties that are less strongly expressed in the pressure-driven deformation. Pressure-forecast information is available almost immediately, whereas the expected value for tracer and thermal-power forecasts continues to increase during the first two years.

The EVA results therefore predict that deformation measurements should strongly constrain the reservoir-pressure response and provide useful, although less complete, information about transport and thermal performance. These estimates describe the expected reduction in ensemble variance before any data are assimilated; they do not demonstrate that the posterior will be unbiased or that the selected truth will be recovered. The predictions are tested in Section 7 by comparing them with the variance reduction and forecast accuracy actually obtained from the ES-MDA inversion.

Remark It should be noted that caution is needed in transferring the results to a real system. Within the model, strain measured after five days constrains even the ten-year thermal power because all forecast uncertainty resides in a few static parameters, which the early, pressure-driven strain response already senses. A real reservoir offers no such guarantee. Pressure baffles, permeability heterogeneity, and the background fracture network may obscure the pressure response, and postpone the variance reduction to later points in time if the pressure front does not reach such features. Dynamical processes like stress-dependent fracture permeability (aperture dilation), and the poro-elastic interactions between pressurised fractures will make the system behaviour more chaotic and time-dependent. As a result, the pressure and deformation response may detect and infer such behaviour only once it occurs in time, if at all. Moreover, the long-term thermal performance depends critically on the propagation of the thermal front. This transient thermal front will induce significant deformations that evolve over time, where it remains to be seen if they can be disambiguated from the pressure-induced deformations. Finally, interpreting the strain amplitudes presumes known elastic properties, which are themselves uncertain in the field. Overall, the high value of the earliest measurements is therefore partly an artefact of the compact parametrisation of the prior and the omission of strong transient effects in the simulated physics.

6 Data Assimilation Methodology

6.1 Overview

Data assimilation (DA) seeks a probabilistic estimate of the model parameters $\boldsymbol{x}$ that is consistent with both prior knowledge and observed data $\boldsymbol{y}_{\text{obs}}$. In the Bayesian framework, this is expressed via Bayes' theorem:

$$p(\boldsymbol{x} \mid \boldsymbol{y}_{\text{obs}}) \propto p(\boldsymbol{y}_{\text{obs}} \mid \boldsymbol{x})\, p(\boldsymbol{x}), \tag{5}$$

where $p(\boldsymbol{x})$ is the prior distribution, $p(\boldsymbol{y}_{\text{obs}} \mid \boldsymbol{x})$ is the likelihood, and $p(\boldsymbol{x} \mid \boldsymbol{y}_{\text{obs}})$ is the posterior distribution.

The Ensemble Smoother with Multiple Data Assimilation (ES-MDA) [5] provides an approximate, computationally efficient solution to Eq. (5) by representing distributions as finite ensembles and performing N_A successive update steps, each using an inflated covariance matrix. This iterative scheme better accommodates non-linear observation operators than a single linear update. Fig. 16 sketches the overall workflow.

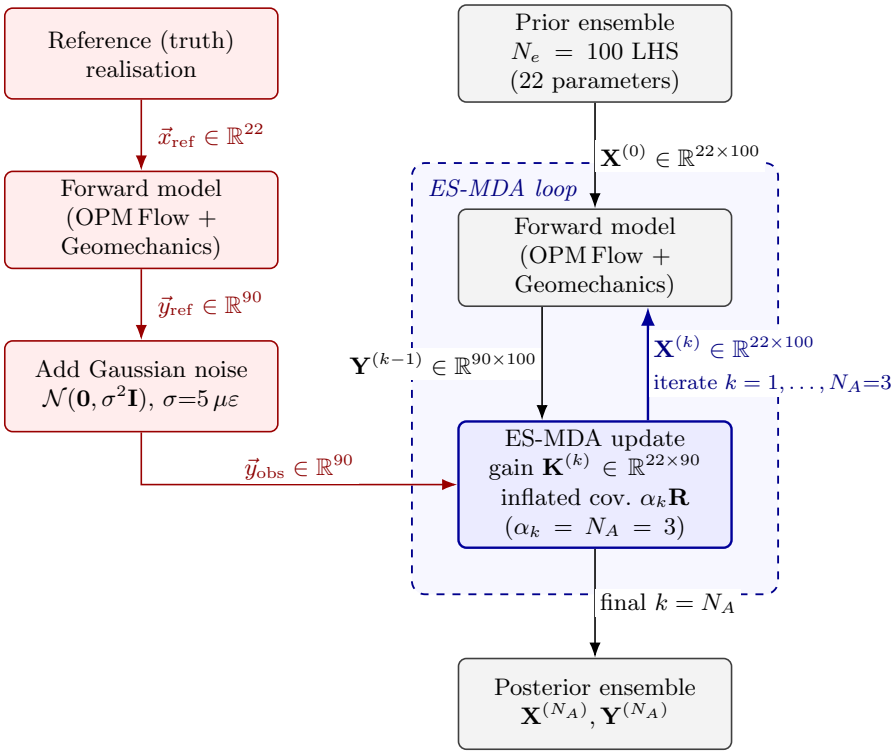


Figure 16: Data assimilation workflow. The reference (truth) realisation is run once through the forward chain and perturbed with Gaussian noise (σ = 5 $\mu\varepsilon$) to yield the synthetic data $\boldsymbol{y}_{\text{obs}}$ (red strand). The prior ensemble $\mathbf{X}^{(0)}$ enters the ES-MDA loop, in which each of the $N_A = 3$ cycles updates the ensemble with the inflated-covariance Kalman gain and re-simulates it through OPM Flow and the geomechanical forward model. Edge labels give the dimensions of the data objects (Sections 6.2 and 6.3).

6.2 State Vector

The 22 uncertain parameters of Section 3.8 are assimilated as the state vector

$$\boldsymbol{x} = \left[\log_{10} k_1,\ \ldots,\ \log_{10} k_{10},\ \phi_1,\ \ldots,\ \phi_{10},\ \log_{10} k_m,\ \mathit{multpv}\right]^T \in \mathbb{R}^{22}, \tag{6}$$

where k_i (mD) and ϕ_i are the permeability and porosity of fracture i, k_m (mD) is the matrix permeability, and *multpv* is the boundary pore-volume multiplier. Permeabilities are treated in $\log_{10}$ space to enforce

positivity and to improve the Gaussian approximation of their prior distribution. Although only the discrete values $multpv \in \{10, 20, \ldots, 80\}$ are simulated, the multiplier is carried as a continuous variable through the update and snapped back onto the discrete set afterwards. For an ensemble of N_e realisations the state matrix is $\mathbf{X} \in \mathbb{R}^{22\times N_e}$.

After each update, the parameters are clipped to physically admissible bounds,

$$k_i \in [5\,\text{mD},\ 100\,000\,\text{mD}], \qquad \phi_i \in [0.05,\ 0.30], \qquad k_m \in [0.25\,\text{mD},\ 100\,\text{mD}], \tag{7}$$

which are deliberately wider than the prior sampling ranges (Table 6), so that the update may move members somewhat beyond the prior support without producing unphysical values.

6.3 Observation Operator

The observation operator maps a parameter vector $\boldsymbol{x}$ to the strain observations defined in Section 4. It comprises the full forward chain of this study: the OPM Flow simulation of the realisation (Section 3.7), the extraction of the pressure fields at the baseline and measurement dates, and the geomechanical forward model (Section 4.1) evaluated at the measurement points. The resulting observation vector,

$$\boldsymbol{y} = \big[\underbrace{\varepsilon_{zz,1},\ \ldots,\ \varepsilon_{zz,18}}_{\text{vertical sections}},\ \underbrace{\varepsilon_{xx,1},\ \ldots,\ \varepsilon_{xx,72}}_{\text{horizontal sections}}\big]^T \in \mathbb{R}^{90}, \tag{8}$$

collects the vertical strains at the 18 measurement points on the vertical well sections and the axial strains at the 72 points on the horizontal sections, at the single measurement date $t_{\text{meas}} = 2 years$, relative to the pre-injection baseline (Eq. (1)). While the elastic model itself is linear in the pressure changes, the mapping from the parameters to the pressure field involves the full non-linear flow simulation, so the observation operator as a whole is non-linear.

6.3.1 Measurement Error Covariance

All strain observations are assigned the same accuracy of $\sigma = 5\,\mu\varepsilon$, a conservative figure for field-grade fibre-optic sensing. Measurement errors are assumed independent and Gaussian, so the measurement-error covariance matrix is

$$\mathbf{R} = \sigma^2 \mathbf{I}_{90}. \tag{9}$$

6.3.2 Synthetic Observation Data

The synthetic observed data are generated from the reference realisation by perturbing its simulated observation vector once with the same error model,

$$\boldsymbol{y}_{\text{obs}} = \boldsymbol{y}_{\text{ref}} + \varepsilon_{\text{obs}}, \qquad \varepsilon_{\text{obs}} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \tag{10}$$

so that the truth is not exactly reproducible from the data — as in a field application.

6.3.3 Prior Simulations

For each realisation of the prior ensemble (Section 3.10) the corresponding OPM input deck is assembled and the flow and geomechanical forward models are run. The prior observable matrix is

$$\mathbf{Y}_{\text{prior}} = \big[\boldsymbol{y}^{(1)}, \boldsymbol{y}^{(2)}, \ldots, \boldsymbol{y}^{(N_e)}\big] \in \mathbb{R}^{90\times N_e}, \tag{11}$$

where $N_e = 100$ (Section 3.10) and each column $\boldsymbol{y}^{(i)}$ is the observation vector (Eq. (8)) of the i-th prior realisation.

6.4 ES-MDA Update

ES-MDA [5] performs N_A successive ensemble smoother steps, each using an inflated measurement-error covariance $\alpha_k \mathbf{R}$. The inflation factors $\{\alpha_k\}$ are chosen to satisfy the *data-weighting condition*:

$$\sum_{k=1}^{N_A} \frac{1}{\alpha_k} = 1. \tag{12}$$

This study uses a uniform schedule $\alpha_k = N_A = 3$ for all iterations.

At each iteration $k = 1, \ldots, N_A$, a stochastic perturbation consistent with the inflated error model is drawn for each ensemble member i [19],

$$\boldsymbol{\varepsilon}_k^{(i)} \sim \mathcal{N}(\mathbf{0},\, \alpha_k \mathbf{R}), \tag{13}$$

and added to the member's simulated observations in the innovation term of the update (Eq. (17)); since the perturbations are zero-mean Gaussian, this is statistically equivalent to the classical scheme in which the observation vector itself is perturbed.

The ensemble-based cross-covariance and auto-covariance matrices at iteration k are computed from the current ensemble $\mathbf{X}^{(k-1)}$ and the corresponding simulated observations $\mathbf{Y}^{(k-1)}$:

$$\mathbf{C}_{xy}^{(k)} = \frac{1}{N_e - 1}\left(\mathbf{X}^{(k-1)} - \overline{\mathbf{X}}^{(k-1)}\right)\left(\mathbf{Y}^{(k-1)} - \overline{\mathbf{Y}}^{(k-1)}\right)^T \in \mathbb{R}^{22\times 90}, \tag{14}$$

$$\mathbf{C}_{yy}^{(k)} = \frac{1}{N_e - 1}\left(\mathbf{Y}^{(k-1)} - \overline{\mathbf{Y}}^{(k-1)}\right)\left(\mathbf{Y}^{(k-1)} - \overline{\mathbf{Y}}^{(k-1)}\right)^T \in \mathbb{R}^{90\times 90}, \tag{15}$$

where the overlines denote the corresponding ensemble means. The Kalman gain at iteration k is:

$$\mathbf{K}^{(k)} = \mathbf{C}_{xy}^{(k)} \left(\mathbf{C}_{yy}^{(k)} + \alpha_k \mathbf{R}\right)^{-1} \in \mathbb{R}^{22\times 90}. \tag{16}$$

The updated state for member i at iteration k is:

$$\boldsymbol{x}^{(k,i)} = \boldsymbol{x}^{(k-1,i)} + \mathbf{K}^{(k)} \left(\boldsymbol{y}_{\text{obs}} - \boldsymbol{y}^{(k-1,i)} - \boldsymbol{\varepsilon}_k^{(i)}\right), \tag{17}$$

with $\boldsymbol{x}^{(0,i)} = \boldsymbol{x}_{\text{prior}}^{(i)}$. No covariance localisation is applied: all 22 parameters are global rather than spatially distributed, so there is no natural distance between parameters and observations, and the ensemble size $N_e = 100$ is large relative to the dimension of the state.

After each update, the permeabilities are exponentiated from $\log_{10}$ space, the continuous parameters are clipped to the bounds (Eq. (7)), and *multpv* is snapped to the nearest member of its discrete set. The bounded parameters, not the raw update, are stored as the ensemble state $\mathbf{X}^{(k)}$, so that the state entering the next iteration is exactly the one whose forward simulations produce $\mathbf{Y}^{(k)}$. Each iteration therefore re-simulates the full forward chain for all $N_e = 100$ members, i.e the 10-year OPM Flow run followed by the geomechanical forward model. Once the observable matrix $\mathbf{Y}^{(k)}$ is obtained, that enters the next update. The ensemble after the final iteration, $\mathbf{Y}^{(N_A)}$, constitutes the posterior predictive distribution. In total, the study comprises $1 + (1 + N_A) \times 100 = 401$ forward runs (reference, prior, and three update iterations).

6.5 Data-Fit Diagnostics

The convergence of the assimilation is monitored with scalar diagnostics of the data fit, evaluated from the simulated observations of the ensemble after each iteration. Here, $y_j^{(k,i)}$ denotes observation j of member i at iteration k.

All χ^2 diagnostics are normalised per observation, so a value of one corresponds to a fit at the noise level.

6.5.1 Misfit of the Ensemble Mean

The first pair measures the misfit of the ensemble mean

$$\bar{y}_j^{(k)} = \frac{1}{N_e} \sum_{i=1}^{N_e} y_j^{(k,i)} \tag{18}$$

against the observed data, in physical units and normalised by the measurement error σ_j of observation j:

$$\mathrm{RMS}_{\mathrm{mean}}^{(k)} = \left(\frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \left(\bar{y}_j^{(k)} - y_{\mathrm{obs},j} \right)^2 \right)^{1/2}, \tag{19}$$

$$\chi^2_{\mathrm{mean},k} = \frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \frac{\left(\bar{y}_j^{(k)} - y_{\mathrm{obs},j} \right)^2}{\sigma_j^2}. \tag{20}$$

In this study all observations share the same accuracy, $\sigma_j = \sigma = 5\,\mu\varepsilon$ (Eq. (9)).

6.5.2 Misfit of the Ensemble Members

The second pair starts from the misfit of an individual ensemble member, defined by the same two measures applied to the member's simulated observations:

$$\mathrm{RMS}^{(k,i)} = \left(\frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \left(y_j^{(k,i)} - y_{\mathrm{obs},j} \right)^2 \right)^{1/2}, \tag{21}$$

$$\chi^2_{k,i} = \frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \frac{\left(y_j^{(k,i)} - y_{\mathrm{obs},j} \right)^2}{\sigma_j^2}. \tag{22}$$

The ensemble diagnostics are their averages over the ensemble:

$$\mathrm{RMSE}_{\mathrm{ens}}^{(k)} = \frac{1}{N_e} \sum_{i=1}^{N_e} \mathrm{RMS}^{(k,i)}, \tag{23}$$

$$\chi^2_{\mathrm{ens},k} = \frac{1}{N_e} \sum_{i=1}^{N_e} \chi^2_{k,i}. \tag{24}$$

These additionally penalise the ensemble spread, which Section 6.5.3 isolates.

6.5.3 Ensemble Spread

The gap between the member-based and the mean-based diagnostics is the spread of the simulated observations around their ensemble mean. Its basic quantity is the ensemble variance of a single observation,

$$s_j^{2\,(k)} = \frac{1}{N_e} \sum_{i=1}^{N_e} \left(y_j^{(k,i)} - \bar{y}_j^{(k)}\right)^2, \tag{25}$$

normalised by the measurement error to the spread contribution of measurement point j,

$$\chi^2_{\mathrm{spread},k,j} = \frac{s_j^{2\,(k)}}{\sigma_j^2}, \tag{26}$$

where a value of one means that the ensemble spread at that point matches the noise level. The spread diagnostic is their average over the observations:

$$\chi^2_{\mathrm{spread},k} = \frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \chi^2_{\mathrm{spread},k,j}. \tag{27}$$

Decomposing the squared residual around the ensemble mean gives the exact relation

$$\chi^2_{\mathrm{ens},k} = \chi^2_{\mathrm{mean},k} + \chi^2_{\mathrm{spread},k}, \tag{28}$$

so the ensemble diagnostics always lie above their mean-based counterparts and approach them only as the ensemble spread falls below the measurement error.

In physical units, the spread is summarised by the mean per-point ensemble standard deviation

$$\bar{s}^{(k)} = \frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} s_j^{(k)}, \tag{29}$$

a descriptive companion to χ^2_{spread} that, unlike the χ^2 diagnostics, admits no exact decomposition.

6.5.4 Realized Noise Levels

Even a perfect model — one that reproduces the reference response exactly — does not fit the noisy data perfectly: by Eq. (10), its residual is precisely the noise draw $\boldsymbol{\varepsilon}_{\mathrm{obs}} = \boldsymbol{y}_{\mathrm{obs}} - \boldsymbol{y}_{\mathrm{ref}}$. Inserting this residual into the definitions above yields the *realized-noise* levels

$$\mathrm{RMS}_{\mathrm{noise}} = \left(\frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \varepsilon^2_{\mathrm{obs},j} \right)^{1/2}, \qquad \chi^2_{\mathrm{noise}} = \frac{1}{n_{\mathrm{obs}}} \sum_{j=1}^{n_{\mathrm{obs}}} \frac{\varepsilon^2_{\mathrm{obs},j}}{\sigma_j^2}, \tag{30}$$

the misfit that the perfect model would attain for the specific noise draw of this study. They are the empirical counterparts of the nominal levels σ and one, around which they fluctuate. Diagnostics falling below these levels therefore do not indicate a better model, but overfitting to the specific noise realisation.

7 Geomechanical Inversion of Distributed Strain Sensing Data

The preceding EVA analysis estimated, from the prior ensemble, how much uncertainty the proposed deformation measurements were expected to remove. This section evaluates how much of that value is realised when the 90 synthetic strain observations are assimilated using ES-MDA. Because the observations are generated from a known synthetic truth, the posterior can be assessed in terms of both uncertainty reduction and accuracy. This distinction is important since a contraction of the ensemble does not necessarily imply that the true reservoir properties or forecasts have been recovered.

The results are presented in four stages. First, Section 7.1 examines whether the posterior ensemble reproduces the observed strain response and whether the assimilation reaches the assumed measurement-error level. Section 7.3 then assesses which reservoir properties are constrained, how accurately their true values are recovered, and how the realised variance reduction compares with the EVR predictions. The posterior ensemble is subsequently propagated forward in Section 7.3 to determine how the strain measurements affect forecasts of reservoir pressure, production, tracer transport and thermal-power output.

7.1 Reproduction of the Strain Observations

This subsection assesses if the updated ensemble reproduces the assimilated strain measurements. Fig. 17 compares the strain profiles predicted by the prior ensemble and the three ES-MDA iterations with the 90 noisy observations and the known noise-free response of the synthetic truth.

The prior ensemble spans a wide range of strain responses but contains the synthetic truth at all measurement locations. Its average standard deviation is 39 $\mu\varepsilon$, substantially larger than the assumed measurement error of 5 $\mu\varepsilon$, and the prior-mean response provides a poor fit to the observations. The ES-MDA updates rapidly reduce this spread. Most of the adjustment occurs during the first iteration, and after three iterations the posterior reproduces the vertical strain profiles and the individual fracture-related peaks along both horizontal wells. The posterior response remains close to the noise-free synthetic truth rather than matching every perturbation in the noisy observations, because the 90 measurements are fitted jointly through the common reservoir model.

Fig. 18 summarises the convergence using the measurement-space diagnostics defined in Section 6. The RMS mismatch and χ^2 of the ensemble mean measure agreement with the observations, while the ensemble standard deviation describes the remaining posterior spread. The RMS mismatch of the ensemble mean decreases from 17.5 $\mu\varepsilon$ in the prior to 4.5 $\mu\varepsilon$ in the posterior. Correspondingly, χ^2_{mean} decreases from 12.3 to 0.80, indicating that the final ensemble mean fits the observations within their assumed error. The average standard deviation of the simulated measurements decreases from 39 $\mu\varepsilon$ to 2.2 $\mu\varepsilon$, while the member-averaged χ^2 reaches approximately 1.2. The posterior ensemble is therefore both consistent with the observations and substantially narrower than the prior in measurement space.

This agreement establishes that ES-MDA can reproduce the complete observed strain response. It does not, by itself, demonstrate that every reservoir property has been recovered accurately, since different parameter combinations may produce similar strain profiles. The resulting parameter constraints and their agreement with the synthetic truth are examined next.

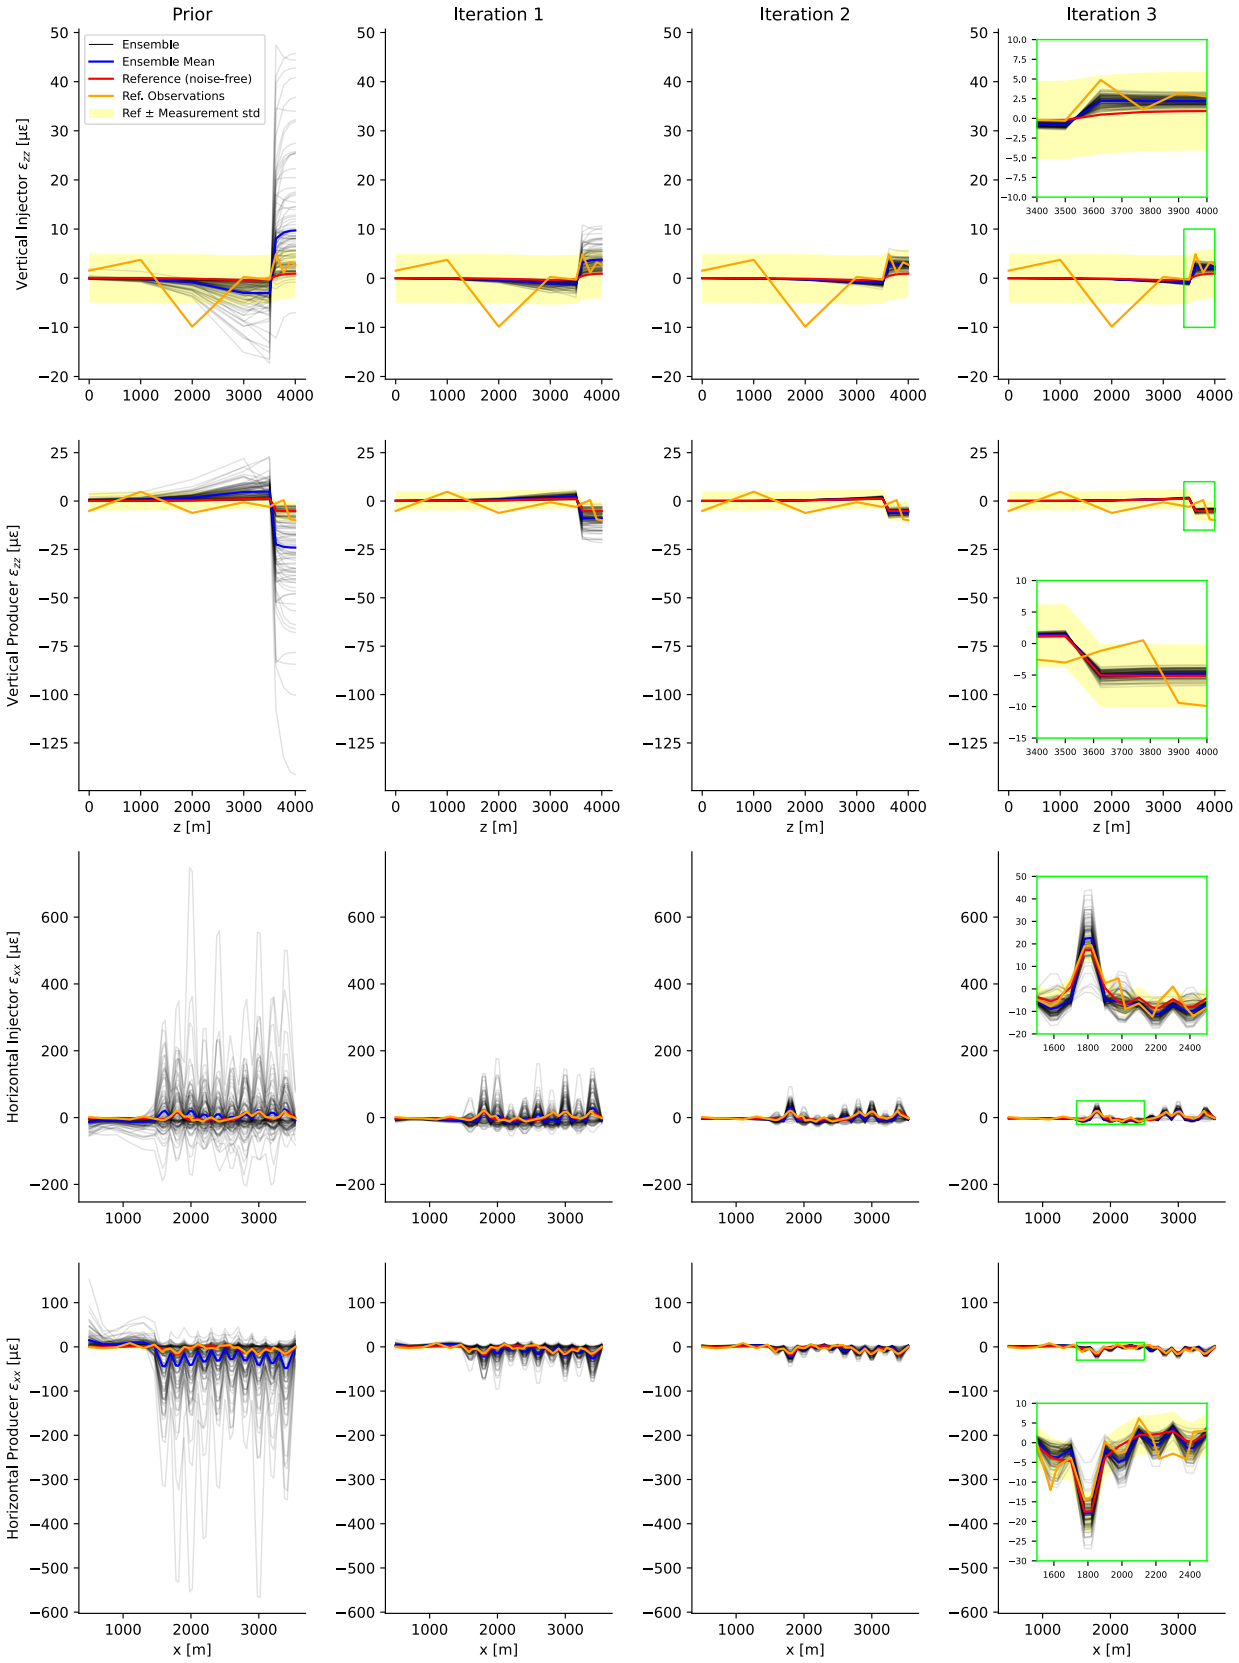


Figure 17: ES-MDA data fit: strain profiles along the four instrumented well sections (rows: vertical injector and producer, ε_{zz} vs. depth; horizontal injector and producer, ε_{xx} vs. x) for the prior and the three ES-MDA iterations (columns). Each panel shows the ensemble members (thin black), the ensemble mean (blue), the noise-free reference with the $\pm\sigma$ measurement band (red line, yellow band), and the assimilated observations $\boldsymbol{y}_{\text{obs}}$ (orange).

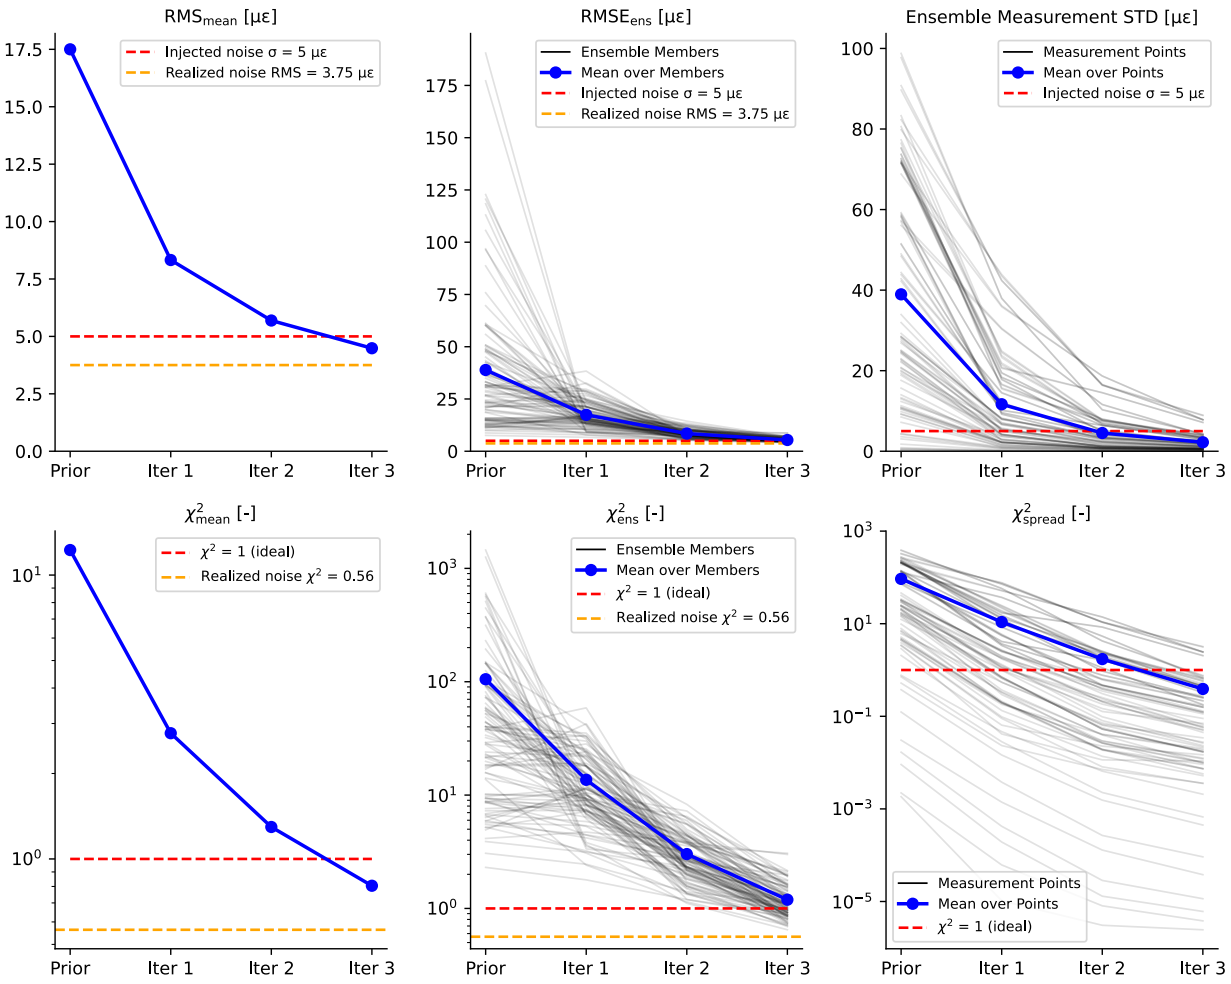


Figure 18: Data-fit diagnostics for each ES-MDA iteration, shown in physical units in the top row and normalised by the measurement error in the bottom row. The left column shows the misfit of the ensemble mean, the middle column shows the individual-member misfits (black) and their ensemble average (blue), and the right column shows the spread at each measurement point (black) and its average (blue). Red dashed lines denote the assumed noise level of $\sigma = 5\,\mu\varepsilon$ and $\chi^2 = 1$, while orange dashed lines denote the realised noise of the synthetic observations.

7.2 Recovery of the Reservoir Properties

The close reproduction of the strain observations does not imply that all 22 reservoir parameters are constrained equally. Because the synthetic truth is known, the posterior ensemble can be evaluated in terms of both its reduction in uncertainty and its proximity to the true parameter values. Fig. 19 compares the prior and posterior ensembles with the synthetic truth for the 10 fracture permeabilities and porosities, matrix permeability and boundary pore-volume multiplier.

Matrix permeability is the most accurately recovered parameter. Its posterior contracts to 23 ± 3 mD around a true value of 20 mD, consistent with the strong sensitivity of the reservoir-scale strain response to pressure diffusion through the matrix. The fracture-permeability results are more differentiated. For the partially open fractures the prior mean already lies close to the reference, so assimilation mainly tightens the spread. The open fractures show a clear gain, but the posterior does not fully recover their high permeabilities, which sit at the upper end of the prior range. Instead, the closed fractures are matched poorly, with the posterior ensemble mean *decreasing* in predictive accuracy compared to the prior mean. The fracture porosities show little improvement, assimilation reshapes their near-uniform prior into a roughly Gaussian posterior of slightly reduced spread, without moving it towards the reference values.

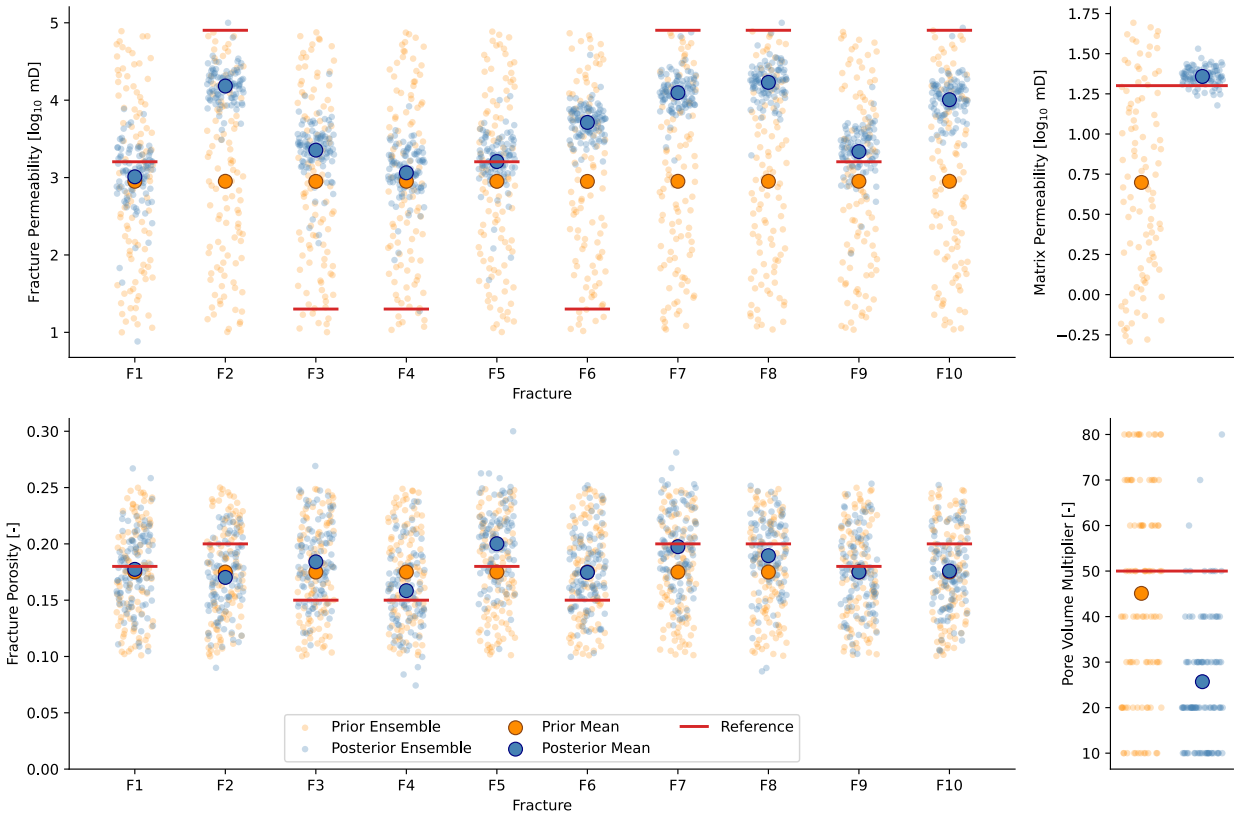


Figure 19: Prior and posterior estimates of the uncertain reservoir parameters. Small dots show the individual prior (orange) and posterior(blue) ensemble members, large dots their ensemble means, and red lines the synthetic truth. Fracture permeabilities and porosities are shown separately for fractures F1–F10.

This behaviour reflects the sensitivity of the pressure-driven strain response. Once a fracture is sufficiently conductive, a further increase in permeability produces little additional change in reservoir pressure or strain. Similarly, a sufficiently impermeable fracture carries little flow, making its exact permeability difficult to determine, and likely irrelevant. The strain measurements can therefore distinguish the fractures that contribute most strongly to the flow within the reservoir, but have limited ability to recover permeability values at the extremes of the prior range.

Fracture porosity remains weakly constrained. Its posterior means stay close to the prior means, and the reduction in ensemble spread is small compared with that obtained for permeability. This is consistent with the EVA results: a single pressure-driven strain survey is primarily sensitive to the permeability field controlling the pressure distribution and contains considerably less information about the fracture storage volume. The boundary multiplier also contracts but moves away from its true value, providing an example of posterior contraction without improved accuracy. It should therefore be treated as weakly identified by these observations.

Additional diagnostics of the convergence of the ES-MDA algorithm, and the achieved posterior distributions for each uncertain reservoir parameter are shown in Sections C and D.

7.2.1 Expected and Realised Variance Reduction

Fig. 20 compares the posterior contraction produced by ES-MDA with the variance reduction predicted by EVA. The bars show the posterior-to-prior variance ratio, while the black lines show the corresponding linear-Gaussian prediction, 1 − EVR. Lower values indicate a greater reduction in uncertainty.

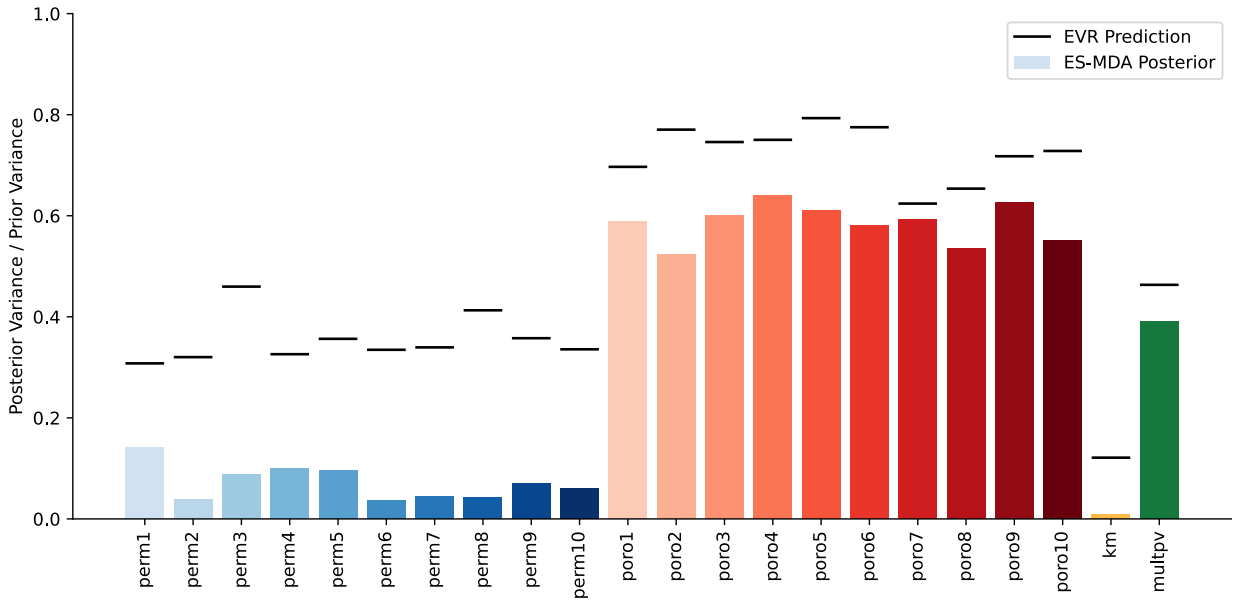


Figure 20: Comparison between the posterior variance realised by ES-MDA and that predicted by EVA. Bars show the posterior-to-prior variance ratio, $\sigma^2_{\text{post}}/\sigma^2_{\text{prior}}$, for each parameter. Black lines show the linear-Gaussian prediction $1 - \text{EVR}$ derived from the prior ensemble. A bar below its corresponding line indicates a stronger variance reduction than predicted.

EVA correctly predicts the broad ordering of the parameter groups. Both analyses identify matrix permeability as the most informative parameter, followed by fracture permeability, the boundary multiplier and fracture porosity. This indicates that EVA provides a useful prospective assessment of which properties a proposed strain survey is likely to constrain.

The magnitude of the realised variance reduction is more accurately predicted by ES-MDA compared to EVA. Fig. 20 shows ES-MDA reduces the posterior variance more strongly than EVA for every parameter. The difference arises because EVA represents the parameter–measurement relationship through a single linear-Gaussian approximation, whereas ES-MDA applies successive updates to a nonlinear, finite ensemble. For matrix permeability and the conductive fractures, this additional contraction is accompanied by movement towards the synthetic truth and therefore represents a genuine improvement. For the closed fractures and the boundary multiplier, however, the posterior variance predicted by ES-MDA can smaller without becoming more accurate. The comparison demonstrates why variance reduction must not be interpreted as accuracy. EVA is effective for identifying the parameters expected to benefit from the measurements, but the realised posterior must still be checked for bias. The underestimated permeabilities of the most conductive fractures are particularly important because they can affect the predicted transport of tracer and heat, as examined in the following section.

7.3 Posterior Ensemble Forecast

The practical value of the strain inversion lies in whether it improves forecasts that were not assimilated. The posterior ensemble is therefore propagated to the end of the ten-year simulation, eight years beyond the strain survey. The forecast is assessed progressively through the reservoir-pressure response, the distribution of production among the fractures, and the resulting tracer and thermal-power behaviour.

Fig. 21 compares the prior and posterior pressure forecasts. Assimilation substantially reduces the uncertainty in the injector and producer bottom-hole pressures and in the average reservoir pressure. The producer and reservoir-average pressures are reproduced closely, while the posterior injector pressure remains approximately 15 bar above the synthetic truth. The strain measurements therefore constrain the overall pressure

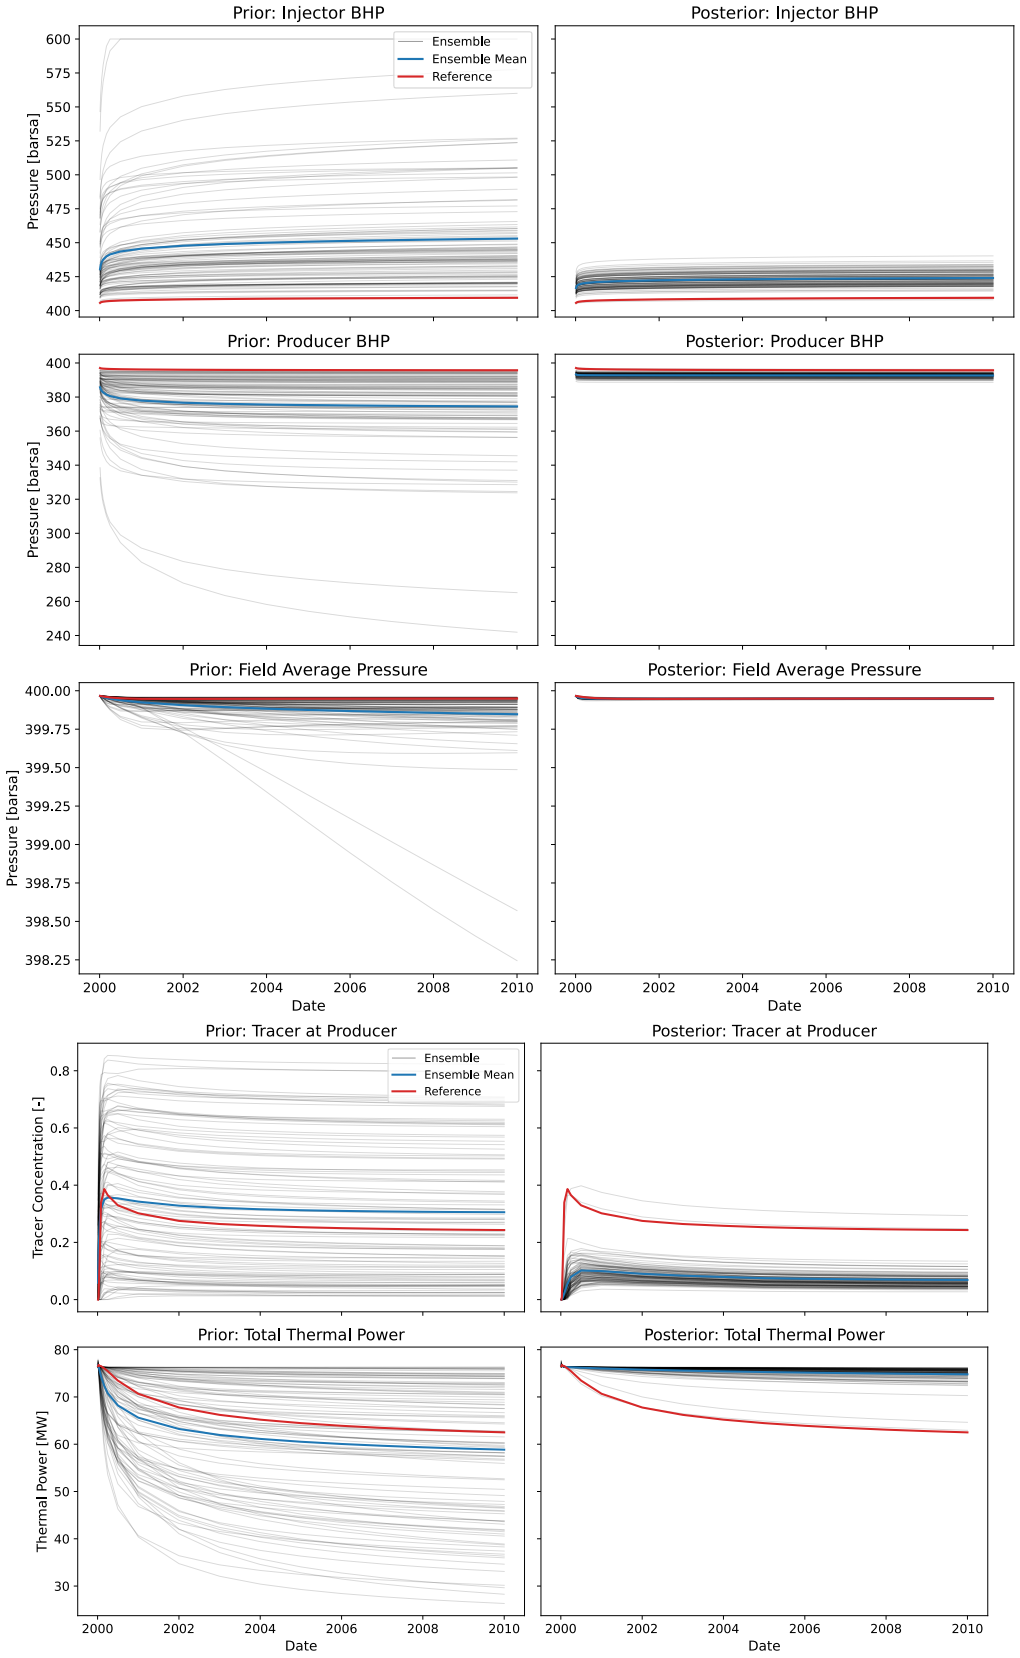


Figure 21: Prior and posterior forecasts of the reservoir pressure and production response over ten years. The upper panels show injector and producer bottom-hole pressures and field-average pressure; the lower panels show tracer concentration at the producer and total thermal-power output. The prior ensemble is shown on the left and the posterior ensemble on the right; individual realisations are shown in transparent black, the ensemble mean in blue, and the synthetic truth in red.

architecture very effectively, although the residual injector-pressure bias indicates that the transmissivity of the most conductive flow paths remains underestimated.

A consequence of the above results is that DSS is more effective at inferring the future flux rates, which are closely related to the thermal power of the system. On the other-hand, the method has more difficulty in assessing the parasitic loads the system may experience from injection, effecting the power-efficiency of the system. Furthermore, the extremely good forecast of the bulk reservoir pressure response indicates that strain sensing could provide an interesting solution for predicting fault slip and seismic risk evolving within the reservoir. This may be the most promising feature to explore next for DSS.

The ability of the inversion in predicting the future performance of the geothermal system can be further understood by looking at the performance forecasts for each fracture. Fig. 22 shows how the posterior distinguishes the four high-throughput fractures from the less conductive fractures and strongly reduces the prior uncertainty in their production rates and thermal-power contributions. However, production from the most conductive fractures is underestimated, while production from several low-throughput fractures is overestimated. As a result, the posterior also underestimates the cooling associated with the fractures, thereby forecasting thermal breakthrough *later* than the truth case. Such an outcome is a direct result of the bias the inversion has when inferring the permeability of the higher fractures. These results are confirmed through Fig. 23, where the thermal front of the high-permeable fractures of the posterior are retracted compared to the reference case.

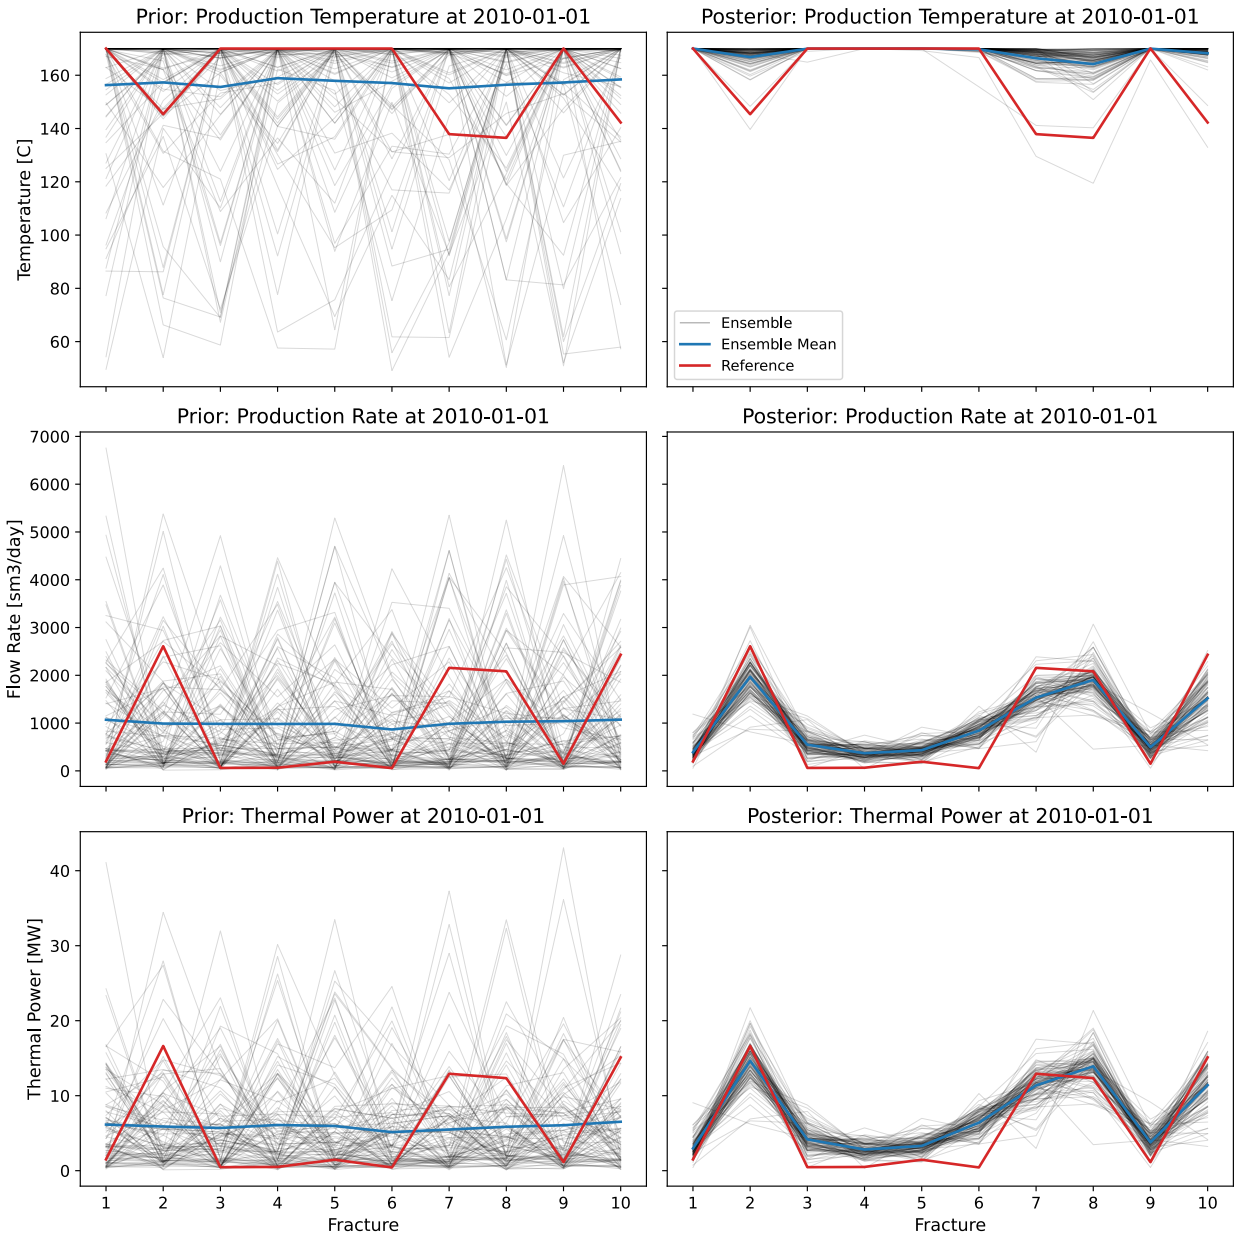


Figure 22: Prior (left) and posterior (right) per-fracture production temperature, flow rate, and thermal power at ten years (1 January 2010). Individual realisations in transparent black, ensemble mean in blue, reference realisation in red.

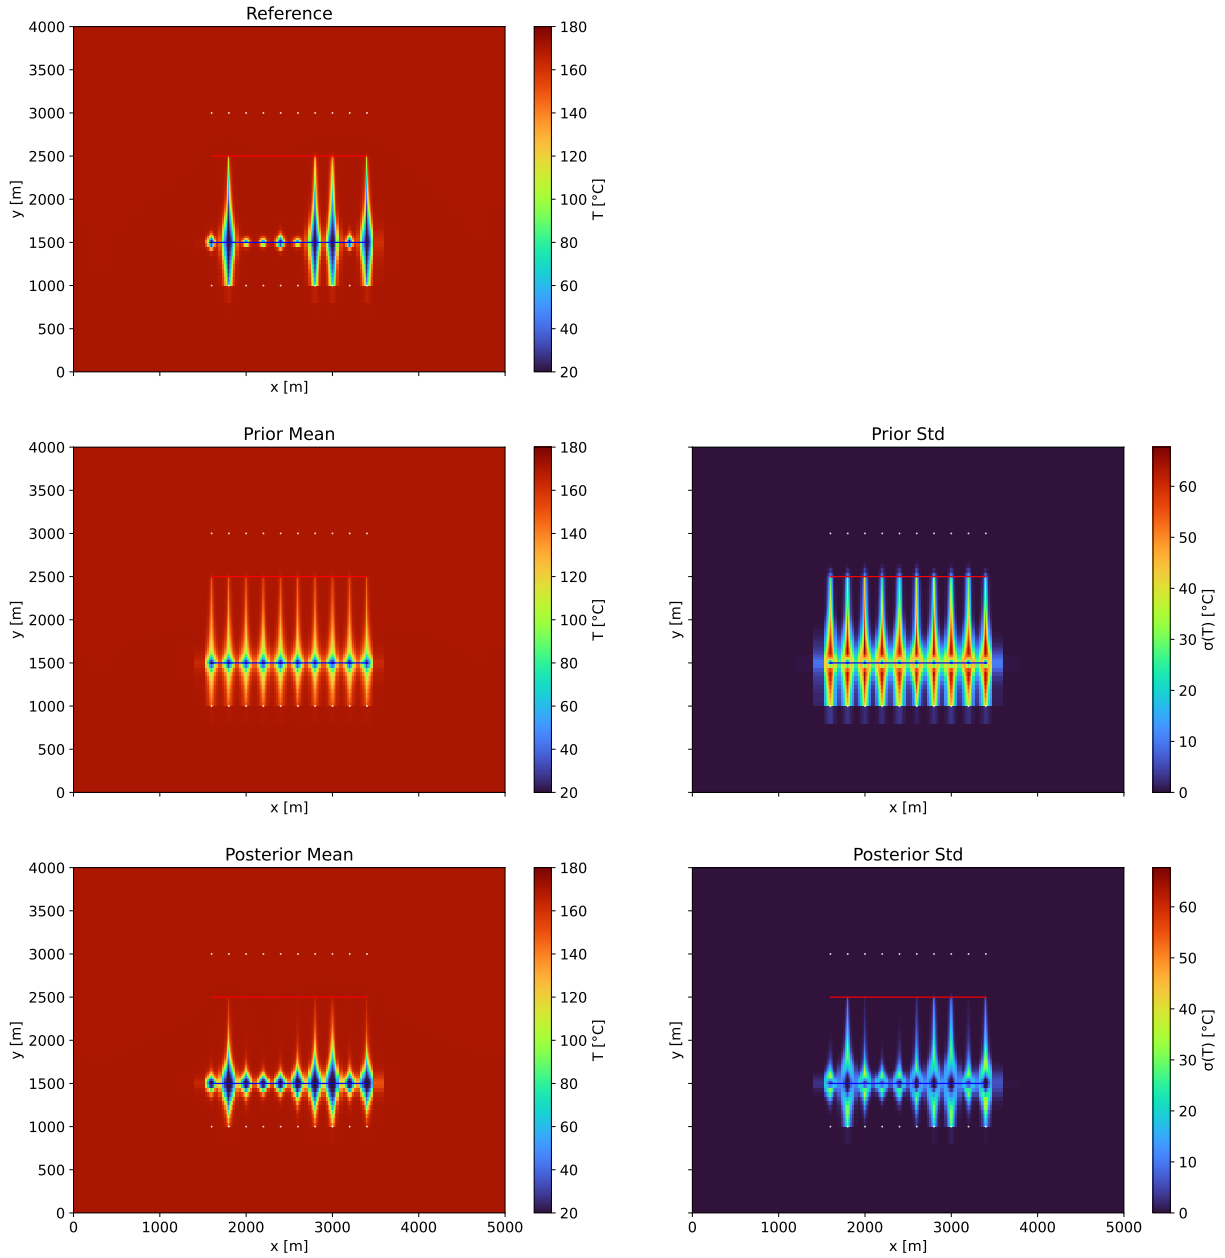


Figure 23: Temperature distribution at the well layer ten years after the start of injection. The top panel shows the synthetic truth, while the middle and bottom rows show the prior and posterior ensemble means (left) and standard deviations (right), respectively. The mean fields share the same temperature scale, and the standard-deviation fields share a separate zero-anchored scale. Assimilation identifies the fractures carrying the cold-water plumes and substantially reduces the prior uncertainty, although the posterior thermal fronts remain shorter than those of the synthetic truth.

8 Conclusions and Outlook

This proof-of-concept study assessed the value of pressure-induced deformation measurements for characterising a fractured enhanced geothermal system. This was achieved by leveraging an ensemble based analysis and inversion that performed 3D numerical thermal-hydraulic simulations of a conceptualised multi-stage EGS doublet with a one-way coupling to a poroelastic mechanical solver using the nucleus-of-strain approach for homogeneous drained elastic properties.

Prior to inversion, ensemble-variance analysis was used to scope the expected information provided by displacement and strain measurements. This allowed the fast comparison of the value of information of different deformation measurement locations and quantities. A synthetic distributed-strain survey was then assimilated using ES-MDA to determine how much uncertainty reduction was realised for one selected reservoir truth and how this information affected the subsequent reservoir forecast.

The principal result is that distributed strain measured along the wells contains sufficient spatial information to identify the hydraulic architecture of the fracture network. Although not proved herein, this result is particularly carried by the sensing of horizontal strain within the horizontal section of the well. The inversion accurately recovers the matrix permeability and distinguishes the fractures carrying most of the flow. It is less successful at recovering the absolute value of the highest and lowest permeabilities, with the most conductive fractures underestimated and the least conductive fractures overestimated. Fracture porosity remains weakly constrained because the early strain response is governed principally by the pressure distribution rather than by the storage properties of the fractures.

These parameter constraints achieved by the inversion have direct implications on the quality of the posterior reservoir forecasts. The posterior reproduces the bulk reservoir-pressure response and substantially reduces uncertainty in the bottom-hole pressures and the distribution of flow among the fractures. This provides a much better estimate of which fractures contribute to production and therefore of their contribution to the gross thermal output. However, the remaining bias in the highly conductive fractures delays the predicted tracer and cold-water fronts. The posterior consequently underestimates tracer breakthrough and remains optimistic about sustained thermal power, predicting approximately 75 MW after ten years compared with 63 MW for the synthetic truth.

The inversion of DSS measurements is also less reliable for forecasting the parasitic power required to operate the system. Although the bulk pressure field is recovered accurately, the injection bottom-hole pressure remains systematically overestimated. This pressure controls the pumping power required to maintain injection and therefore affects the net power and efficiency of the geothermal system. On the other-hand, the accurate recovery of the bulk pressure response may have an important application beyond production forecasting. Changes in pore pressure alter the effective stress acting on faults and are a principal control on injection-induced fault slip. Distributed strain measurements could therefore provide an early indication of how overpressure and pressure drawdown evolve within the reservoir and, when combined with information about fault geometry, the in-situ stress field and fault friction, could improve forecasts of fault stability and seismic risk. This possibility is not tested in the present model, which contains no explicitly slipping faults, but it may be the most promising extension of the DSS approach.

A further advantage is that much of the pressure-related information becomes available early. EVA predicts that a strain survey after only five days could remove approximately 80% of the matrix-permeability variance and $50 - -65\%$ of the fracture-permeability variance. Information about fracture porosity and the boundary multiplier develops more gradually, but the expected uncertainty reduction changes little after approximately two years. Within the present model, deformation monitoring can therefore characterise the pressure response well before tracer or thermal signals reach the producer.

These timing and uncertainty reductions should not be transferred directly to a field system. The experiment is a synthetic case, where the observations are generated and inverted using the same reservoir

and geomechanical models, so there is no structural model error. The prior also contains a compact set of static parameters within a known network of ten parallel fractures. A field reservoir may contain uncertain fracture geometry, background fractures, pressure baffles and spatially variable matrix properties that delay or obscure the propagation of the pressure signal. Under those conditions, early deformation measurements may constrain only the portion of the reservoir reached by the pressure perturbation, and their value may continue to evolve beyond the two-year period considered here.

The forward model also includes only pressure-induced, drained poroelastic deformation with prescribed elastic properties. Thermoelastic deformation will develop as the cold-water front advances and may become comparable with, or obscure, the pressure-induced strain. Stress-dependent fracture permeability (i.e aperture dilation), fracture interaction and irreversible deformation would introduce nonlinear, time-dependent and coupled feedbacks between pressure, deformation, flow and temperature. The pressure response may reveal these processes only after they occur, making repeated time-lapse surveys more important than the single strain snapshot considered here. Uncertain elastic properties and correlations between measurement errors would introduce further ambiguity when converting measured strain amplitudes into reservoir-pressure and permeability estimates.

The results should therefore be interpreted as a scoping assessment of the value of information provided by deformation monitoring, rather than a prediction of field performance. Within the simplified model, the outcome is nevertheless encouraging: distributed strain constrains the early pressure architecture, identifies the dominant flow paths and significantly improves reservoir forecasts before thermal observations become available. The next stage should determine whether this value is preserved when thermoelastic deformation, stress-dependent permeability, uncertain elastic properties, realistic fracture geometries and fault-stability calculations are incorporated into the analysis.

References

[1] C.-F. Tsang and I. Neretnieks. "Flow Channeling in Heterogeneous Fractured Rocks". In: *Reviews of Geophysics* 36.2 (1998), pp. 275–298. DOI: 10.1029/97RG03319.

[2] P. Fu et al. "Thermal Drawdown-Induced Flow Channeling in Fractured Geothermal Reservoirs". In: *Rock Mechanics and Rock Engineering* 49.4 (2016), pp. 1001–1024. DOI: 10.1007/s00603-015-0776-0.

[3] A. H. Hartog. *An Introduction to Distributed Optical Fibre Sensors*. Boca Raton, FL: CRC Press, 2017.

[4] Jincong He et al. "Quantifying Expected Uncertainty Reduction and Value of Information Using Ensemble-Variance Analysis". In: *SPE Journal* 23.02 (Jan. 2018), pp. 428–448. ISSN: 1086-055X. DOI: 10.2118/182609-PA. URL: https://doi.org/10.2118/182609-PA (visited on 07/29/2024).

[5] Alexandre A. Emerick and Albert C. Reynolds. "Ensemble smoother with multiple data assimilation". en. In: *Computers & Geosciences*. Ensemble Kalman filter for data assimilation 55 (June 2013), pp. 3–15. ISSN: 0098-3004. DOI: 10.1016/j.cageo.2012.03.011. URL: https://www.sciencedirect.com/science/article/pii/S0098300412000994 (visited on 05/23/2023).

[6] C. Vogt et al. "Estimating the Permeability Distribution and Its Uncertainty at the EGS Demonstration Reservoir Soultz-sous-Forêts Using the Ensemble Kalman Filter". In: *Water Resources Research* 48 (2012), W08517. DOI: 10.1029/2011WR011673.

[7] C. Vogt, C. Kosack, and G. Marquart. "Stochastic Inversion of the Tracer Experiment of the EGS Demonstration Reservoir in Soultz-sous-Forêts—Revealing Pathways and Estimating Permeability Distribution". In: *Geothermics* 42 (2012), pp. 1–12. DOI: 10.1016/j.geothermics.2011.11.001.

[8] A. C. Gringarten, P. A. Witherspoon, and Y. Ohnishi. "Theory of Heat Extraction from Fractured Hot Dry Rock". In: *Journal of Geophysical Research* 80.8 (1975), pp. 1120–1124. DOI: 10.1029/JB080i008p01120.

[9] A. F. Rasmussen et al. "The Open Porous Media Flow Reservoir Simulator". In: *Computers & Mathematics with Applications* 81 (2021), pp. 159–185. DOI: 10.1016/j.camwa.2020.05.014.

[10] V. C. F. Barbosa Jr. et al. "3D Displacement and Stress Fields of Compacting Reservoir: Alternative Solutions". In: *Brazilian Journal of Geophysics* 40.5 (2022), pp. 7–19.

[11] J. Hu et al. "PVTx Properties of the CO2–H2O and CO2–H2O–NaCl Systems below 647 K: Assessment of Experimental Data and Thermodynamic Models". In: *Chemical Geology* 238.3–4 (2007), pp. 249–267. DOI: 10.1016/j.chemgeo.2006.11.011.

[12] M. Batzle and Z. Wang. "Seismic Properties of Pore Fluids". In: *Geophysics* 57.11 (1992), pp. 1396–1408. DOI: 10.1190/1.1443207.

[13] E. E. Michaelides. "Thermodynamic Properties of Geothermal Fluids". In: *Geothermal Resources Council Transactions* 5 (1981), pp. 361–364.

[14] T. H. Sandve et al. "Convective Dissolution in Field Scale CO2 Storage Simulations Using the OPM Flow Simulator". In: *Proceedings of the 11th Trondheim Conference on CO2 Capture, Transport and Storage (TCCS-11)*. 2021.

[15] D. Nagy, G. Papp, and J. Benedek. "The gravitational potential and its derivatives for the prism". en. In: *Journal of Geodesy* 74.7 (Nov. 2000), pp. 552–560. ISSN: 1432-1394. DOI: 10.1007/s001900000116. URL: https://doi.org/10.1007/s001900000116 (visited on 07/10/2026).

[16] D. Nagy, G. Papp, and J. Benedek. “Corrections to “The gravitational potential and its derivatives for the prism””. en. In: *Journal of Geodesy* 76.8 (Nov. 2002), pp. 475–475. ISSN: 1432-1394. DOI: 10.1007/s00190-002-0264-7. URL: https://doi.org/10.1007/s00190-002-0264-7 (visited on 07/10/2026).

[17] D. F. McTigue. “Elastic Stress and Deformation Near a Finite Spherical Magma Body: Resolution of the Point Source Paradox”. In: *Journal of Geophysical Research* 92.B12 (1987), pp. 12931–12940. DOI: 10.1029/JB092iB12p12931.

[18] E. Pezzulli. *Validation of Geomechanical Inversion via ES-MDA: Conceptual Spherical Source Study*. Internal Technical Report. Storra Dynamics GmbH, 2026.

[19] G. Burgers, P. J. van Leeuwen, and G. Evensen. “Analysis Scheme in the Ensemble Kalman Filter”. In: *Monthly Weather Review* 126.6 (1998), pp. 1719–1724. DOI: 10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2.

[20] S. J. Zarrouk and H. Moon. “Efficiency of Geothermal Power Plants: A Worldwide Review”. In: *Geothermics* 51 (2014), pp. 142–153. DOI: 10.1016/j.geothermics.2013.11.001.

[21] R. DiPippo. *Geothermal Power Plants: Principles, Applications, Case Studies and Environmental Impact*. 4th ed. Oxford: Butterworth-Heinemann, 2016.

A Comparison of the Uniform Case with the Gringarten (1975) Analytical Solution

This appendix presents a validation and sensitivity study of the OPM geothermal model against the analytical solution of [8] for heat extraction from parallel vertical fractures. The study serves four purposes. First, a simplified single-fracture OPM model that closely matches the Gringarten assumptions reproduces the analytical outlet-temperature curve with excellent agreement, confirming that OPM Flow is a valid simulator for geothermal heat extraction (Setup 1; Section A.3). Second, the study establishes that a 17-cell cross-fracture grid with a 1 m fracture cell is sufficient for the present configuration; further grid refinement does not improve accuracy. Third, a systematic parameter variation across fifteen setups quantifies the sensitivity of the thermal breakthrough curve to matrix permeability, overburden heat supply, geothermal gradient, and flow-domain geometry, informing which reservoir parameters most critically affect simulation results. Fourth, the comparison shows that the idealised Gringarten model cannot capture the dominant physical effects present in a realistic 3D setup — principally the heat supply from overburden and underburden formations — so that numerical simulation is required for representative forecasting.

A.1 Analytical Model of Gringarten (1975)

Gringarten et al. present an analytical theory for heat extraction from a hot dry rock (HDR) geothermal system consisting of an infinite series of parallel, equidistant, vertical fractures of uniform aperture $2b$, separating blocks of homogeneous, isotropic, impermeable rock. Cold water enters the bottom of each fracture at $z = 0$ at constant injection temperature T_{W_0} and exits at the top at $z = H$; heat transfer occurs by horizontal conduction from the rock matrix into the fracture and by forced convection along the fracture. The model is inherently one-dimensional in the cross-fracture (x) and along-fracture (z) directions, and is strictly valid only for an *infinite* array of fractures (so that the no-flow symmetry boundary condition holds at the midplane between adjacent fractures).

The governing equations are the energy balance for the water temperature $T_W(z,t)$,

$$b\rho_W c_W \left[\frac{\partial T_W}{\partial t} + v\frac{\partial T_W}{\partial z}\right] = K_R \left.\frac{\partial T_R}{\partial x}\right|_{x=b},$$

and the heat conduction equation for the rock temperature $T_R(x,z,t)$. The solution is expressed in terms of three dimensionless groups:

$$X_{ED} = \frac{\rho_W c_W Q\, x_E}{K_R H}, \tag{31}$$

$$t_D = \frac{(\rho_W c_W)^2}{K_R \rho_R c_R}\left(\frac{Q}{H}\right)^2 t, \tag{32}$$

$$\beta = \frac{\omega H}{T_{R_0} - T_{W_0}}, \tag{33}$$

where Q is the volumetric flow rate per fracture per unit height (m^2/s), x_E is the fracture half-spacing, K_R the rock thermal conductivity, $\rho_R c_R$ the volumetric rock heat capacity, and ω the geothermal gradient. The parameter X_{ED} compares advective heat transport to the conductive capacity of the rock blocks; t_D is the dimensionless time; and β measures the importance of the geothermal gradient (set to zero when the initial rock temperature is uniform along the flow path). The outlet dimensionless temperature

$$T_D = \frac{T_{R_0} - T_W(H,t)}{T_{R_0} - T_{W_0}} \tag{34}$$

is obtained via Laplace transform and numerical inversion, and is tabulated as $T_D(t_D)$ curves for selected values of X_{ED}.

A.2 Uniform Ten-Fracture OPM Setup

The uniform case is defined by ten identical, parallel, vertical fractures with equal permeability and aperture: each fracture is a parallel-plate fracture of 1 mm aperture, upscaled into its 1 m-wide grid cell to an effective permeability of 83 353 mD and porosity of $\phi_f = 0.151$ (see Section B.1, where the flow behaviour of this configuration is characterised in detail). The horizontal injector and producer well sections cross all ten fractures perpendicularly and are separated by a distance $H = 1000\,\mathrm{m}$, so that the flow from injector to producer runs horizontally *within* each fracture plane. Each fracture has half-spacing $x_E = 100\,\mathrm{m}$, giving a total reservoir width of $2x_E n_{\mathrm{frac}} = 2000\,\mathrm{m}$. The fracture height, perpendicular to both the flow direction and the fracture normal, is $h = 100\,\mathrm{m}$. Note that Gringarten's flow coordinate z thus maps onto the horizontal injector–producer direction of the OPM model, and the fracture height h onto the vertical model axis. Note that these dimensions describe the active fracture zone only; the full OPM simulation domain is larger, including overburden, underburden, and buffer cells along the flow path (see Section 3.1). The total volumetric injection rate is $q = 1 \times 10^4\,\mathrm{m^3/day}$ ($\approx 0.116\,\mathrm{m^3/s}$) at surface conditions, distributed equally across all fractures. Cold water is injected at $T_i = 20\,^\circ\mathrm{C}$ into a reservoir initially at $T_{\mathrm{rock}} = 170\,^\circ\mathrm{C}$.

The thermal properties used in the dimensionless scaling are: rock thermal conductivity $K_r = 3.0\,\mathrm{J/(m\,s\,K)}$, volumetric rock heat capacity $\rho_r c_r = 2.2 \times 10^6\,\mathrm{J/(m^3\,K)}$, water density (surface conditions) $\rho_w = 1051\,\mathrm{kg/m^3}$, and water heat capacity $c_w = 4180\,\mathrm{J/(kg\,K)}$.

Substituting these values into Eqs. (31) to (33), the flow rate per fracture per unit height is

$$Q = \frac{q}{n_{\mathrm{frac}}\,h} \approx 1.16 \times 10^{-4}\,\mathrm{m^2/s}, \tag{35}$$

and the key dimensionless parameters are

$$X_{ED} = \frac{\rho_w c_w}{K_r} \cdot \frac{Q}{H} \cdot x_E \approx 16.95, \qquad \beta = 0, \tag{36}$$

since both wells are completed at the same depth (4000 m) and the flow is horizontal, so there is no systematic temperature variation along the flow path. Substituting all values, the dimensionless time and temperature evaluate to

$$t_D = \frac{(\rho_w c_w)^2}{K_r\,(\rho c)_{\mathrm{eff}}} \left(\frac{Q}{H}\right)^2 t = 3.01 \times 10^{-8}\,\mathrm{s^{-1}} \cdot t \approx 0.951\,\mathrm{yr^{-1}} \cdot t, \tag{37}$$

$$T_D = \frac{T_{\mathrm{rock}} - T_{\mathrm{prod}}}{T_{\mathrm{rock}} - T_i} = \frac{170\,^\circ\mathrm{C} - T_{\mathrm{prod}}}{150\,^\circ\mathrm{C}}, \tag{38}$$

where t is the elapsed time and T_{prod} is the produced water temperature. The value $X_{ED} \approx 17$ places the uniform case in the thick-block regime, between the reference curves for $X_{ED} = 8$ and $X_{ED} \to \infty$ published by [8].

A.3 Numerical validation study

Motivation. The OPM forward model extends the Gringarten idealisation in several respects: a three-dimensional finite-volume grid, realistic temperature-dependent fluid properties, a geothermal gradient, overburden and underburden formations, and a finite-permeability rock matrix. A systematic study across fifteen single-fracture setups serves to (i) validate OPM against the analytical benchmark by progressively recovering the Gringarten assumptions, (ii) confirm that the chosen grid resolution is sufficient, (iii) quantify

the sensitivity of the breakthrough curve to each physical parameter, and (iv) identify which 3D extensions are responsible for the departure of the full OPM model from the analytical solution. Parameters are varied one at a time; in each case the outlet temperature is plotted against the dimensionless time t_D and compared with the Gringarten curves at $X_{ED} = 0.5$, 1, 2, 4, 8, and ∞.

Single-fracture simplified setups. The base geometry uses a single fracture with the same permeability, fracture half-spacing, and well geometry as the uniform case. The grid has NX = 17 cells in the I-direction (perpendicular to the fracture surface), with the fracture occupying the middle cell. The J-direction (flow direction) uses either 50 uniform 20 m cells covering a 1000 m inter-well domain (setups 1–12), with injection and production wells placed at the grid boundaries, or 81 cells covering a 4000 m domain (setups 13–15), consisting of 1000 m of unfractured matrix, a 2000 m fracture zone, and again 1000 m of unfractured matrix; the wells are placed 1000 m apart, centred within the fracture zone. The K-direction is either 41 layers (fracture zone only, no overburden) or 61 layers (41 fracture layers plus, on each side, 10 graded layers — 5 of 20 m nearest the fracture zone and 5 of 50 m outermost, 350 m in total — increasing the vertical extent from 100 m to 800 m). Table 7 summarises all fifteen setups.

Key results. The main observations from each group of setups are as follows; results are shown in Fig. 24.

Grid refinement and 3D effects (Setups 1–3). All three setups show very good to excellent agreement with the analytical solution, confirming that the 3D fracture geometry and 20 m grid spacing in the flow direction are acceptable and do not introduce significant additional discrepancy (Fig. 24a).

Matrix permeability (Setups 3–5). Matrix permeability has a strong effect: at $k_m = 200$ mD (Setup 5) the horizontal heat supply from the matrix is so large that breakthrough is drastically delayed; at $k_m = 20$ mD (Setup 4) agreement is good overall, but the temperature front is delayed at intermediate t_D; at $k_m = 2$ mD (Setup 3) agreement is very good. The Gringarten model assumes zero matrix permeability, so the closest match is obtained at the lowest matrix permeability (Fig. 24b).

Overburden and underburden (Setups 3 and 7). Adding 10 overburden and 10 underburden layers (350 m per side; Setup 7), without changing any other parameter, shifts the OPM curve far to the right, beyond the $X_{ED} \to \infty$ reference. This is the single largest contributor to the discrepancy between the Uniform 10-fracture OPM case and the analytical solution; note that increasing k_m to 200 mD has a yet larger absolute effect, but is not representative of the target reservoir (Fig. 24c).

Temperature gradient and boundary conditions (Setups 3, 6–9). Adding a geothermal gradient has negligible effect on the comparison (Setup 6 ≈ Setup 3). Adding overburden with or without a gradient (Setups 7 and 8) gives equally bad agreement. Imposing fixed-temperature boundary conditions at the top and bottom of the domain (Setup 9) does not improve agreement relative to Setup 8, confirming that the overburden heat supply operates through conduction and convection, not through the boundary condition (Fig. 24d).

Overburden properties (Setups 6, 8, 10–12). Setting `THCONR` = 0 outside the fracture zone (Setup 10) gives only partial recovery; setting `PERMZ` = 0.02 mD (Setup 11) barely helps. Setting both together (Setup 12) restores very good agreement, confirming that overburden heat supply has both a conductive and a convective component (Fig. 24e).

Domain extension in J (Setups 6, 9, 13–14). Extending the J-domain to 81 cells with buffer cells on each side (Setup 13), while keeping no overburden, already shifts the curve to bad agreement. This is analogous to adding overburden and underburden in K: both increase the effective thermal inertia of the system by surrounding the fracture zone with additional stagnant matrix volume, thereby delaying thermal breakthrough. Combining the extended J-domain with overburden (Setup 14) cumulates both effects, producing even worse agreement (Fig. 24f).

Table 7: Single-fracture validation setups. Grid dimensions are NX×NY×NZ. k_m is the matrix permeability; "OB" indicates overburden/underburden layers present (10 graded layers above and below the fracture zone — 5×20 m plus 5×50 m, 350 m per side — NZ increased from 41 to 61). $\texttt{THCONR}_{\text{OB}}$ and $\texttt{PERMZ}_{\text{OB}}$ give the thermal conductivity and vertical permeability assigned *outside* the fracture zone (K < 11 and K > 51); "std" denotes the same values as the matrix. "Ext. J" indicates an extended flow domain with 15 buffer cells on each J-side of the fracture zone (NY increased from 50 to 81). "Fixed-T BC" indicates fixed-temperature boundary conditions at the top and bottom faces. The result column gives qualitative agreement with the Gringarten (1975) analytical curve.

Setup	Grid	k_m [mD]	∇T [°C/100m]	OB	$\texttt{THCONR}_{\text{OB}}$	$\texttt{PERMZ}_{\text{OB}}$ [mD]	Ext. J	Fixed-T BC	Result
1	$83 \times 200 \times 1$	2	0	—	—	—	—	—	Excellent
2	$17 \times 50 \times 1$	2	0	—	—	—	—	—	Very good
3	$17 \times 50 \times 41$	2	0	—	—	—	—	—	Very good
4	$17 \times 50 \times 41$	20	0	—	—	—	—	—	Good
5	$17 \times 50 \times 41$	200	0	—	—	—	—	—	Very bad
6	$17 \times 50 \times 41$	2	3	—	—	—	—	—	Very good
7	$17 \times 50 \times 61$	2	0	✓	std	std	—	—	Bad
8	$17 \times 50 \times 61$	2	3	✓	std	std	—	—	Bad
9	$17 \times 50 \times 61$	2	3	✓	std	std	—	✓	Bad
10	$17 \times 50 \times 61$	2	3	✓	0	std	—	—	OK (partial recovery)
11	$17 \times 50 \times 61$	2	3	✓	std	0.02	—	—	Bad (minor improvement)
12	$17 \times 50 \times 61$	2	3	✓	0	0.02	—	—	Very good
13	$17 \times 81 \times 41$	2	3	—	—	—	✓	—	Bad
14	$17 \times 81 \times 61$	2	3	✓	std	std	✓	✓	Bad
15	$17 \times 81 \times 61$	20	3	✓	std	std	✓	✓	Very bad

Comparison with the Uniform case (Setups 1, 6, 4, 14, 15 + Uniform). Setup 15 uses the full original parameters of the Uniform case ($k_m = 20\,\text{mD}$, extended J-domain, overburden) and shows very bad agreement, closely matching the Uniform 10-fracture result. Inspecting the individual fracture curves of the Uniform case reveals a spatial pattern: the central fractures agree well with Setup 15, while breakthrough is progressively more delayed towards the outer fractures at either end of the fractured zone. This is consistent with the outer fractures having access to additional unfractured matrix in the I-direction (perpendicular to flow) beyond the half-spacing boundary, analogous to the domain-extension effect seen in J and K; the inner fractures, flanked by other fractures on both sides, behave more like the idealised single-fracture setup. The remaining offset between Setup 15 and the central Uniform fractures is attributed to minor numerical differences in the larger multi-fracture grid (Fig. 24g). Note that this single-fracture equivalence holds only under *uniform* fracture properties; in the heterogeneous case studied in the main report, where permeability and porosity differ from fracture to fracture, each fracture must be modelled individually and a single-fracture analogue is no longer representative.

A.4 Conclusion

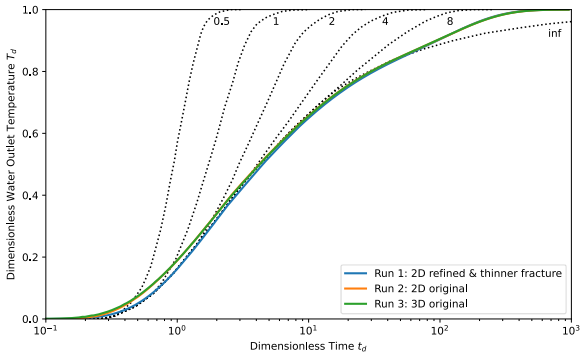
A significant mismatch is observed between the Uniform 10-fracture OPM simulation and the Gringarten reference curve at $X_{ED} \approx 17$. The following paragraphs identify the main physical and numerical factors that contribute to this discrepancy, together with the expected direction of each effect on the OPM outlet temperature curve relative to the analytical solution. Where relevant, the numerical study of Section A.3 provides direct confirmation.

Heat supply from overburden and underburden. Gringarten's model supplies heat exclusively from the rock perpendicular to each fracture (the x-direction), within a half-spacing $x_E = 100\,\text{m}$. No heat is conducted vertically from outside the fracture height. The OPM grid, however, includes formation cells above ($\text{K} < 11$) and below ($\text{K} > 51$) the 100 m fracture zone, all assigned the same thermal conductivity $\texttt{THCONR} = 260\,\text{kJ/(m day K)}$. As the fracture zone cools, heat conducts in from the full overburden and underburden column. The cells immediately adjacent to the fracture zone boundary begin supplying heat within days; at the 1000-year simulation horizon the vertical diffusion depth $\sqrt{2\alpha t}$ reaches roughly 250 m, so a significant column of rock outside the fracture zone is actively contributing heat throughout the simulation. This extra heat supply continuously retards the outlet temperature drop, shifting the OPM curve far to the *right* — beyond even the $X_{ED} \to \infty$ reference curve, which assumes only lateral (x-direction) heat supply. This is confirmed as the dominant cause of the large observed discrepancy by Setups 7–12 in Section A.3: adding overburden alone (Setup 7) causes bad agreement; setting $\texttt{THCONR} = 0$ outside the fracture zone (Setup 10) gives only partial recovery; only setting both $\texttt{THCONR} = 0$ *and* `PERMZ` reduced (Setup 12) restores very good agreement.

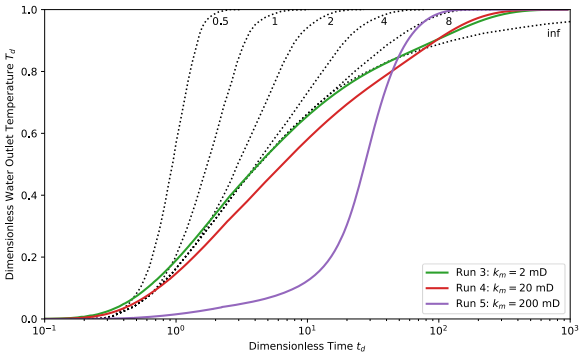
Effective rock heat capacity and matrix porosity. OPM's `HEATCR` keyword (Reference Manual, Section 6.3.84) gives the rock heat capacity per unit *bulk* volume in $\text{kJ/m}^3\text{/K}$; OPM adds the pore brine contribution $\phi_m \rho_w c_w$ separately. The correct effective heat capacity for the Gringarten dimensionless scaling is therefore

$$(\rho c)_{\text{eff}} = \text{HEATCR} + \phi_m \rho_w c_w = 2200 + 0.15 \times 1051 \times 4.18 \approx 2860\,\text{kJ/(m}^3\,\text{K)}, \tag{39}$$

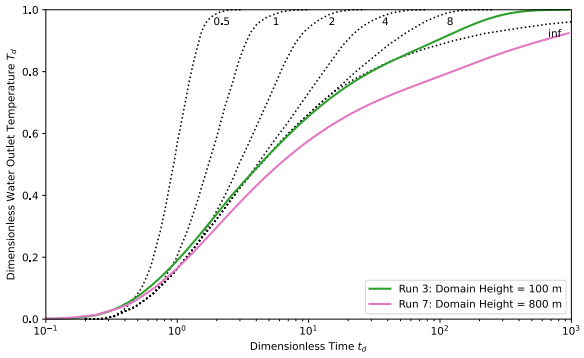
30% larger than `HEATCR` alone. Using only $\text{HEATCR} = 2200\,\text{kJ/(m}^3\,\text{K)}$ would overestimate t_D by a factor of $2.86/2.2 = 1.30$; this study uses the correct effective heat capacity $(\rho c)_{\text{eff}}$ throughout. The keyword `THCONR` represents combined rock and fluid conductivity ($260\,\text{kJ/(m day K)} = 3.01\,\text{J/(m s K)}$); since the



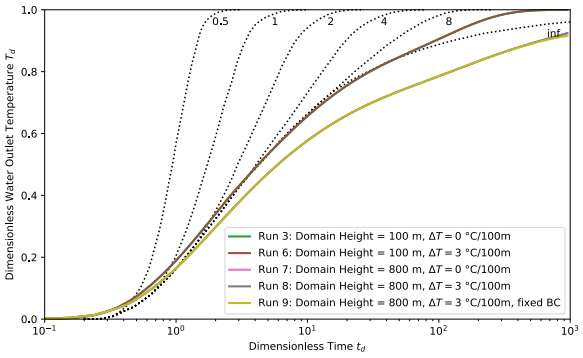
(a) Influence of grid refinement.



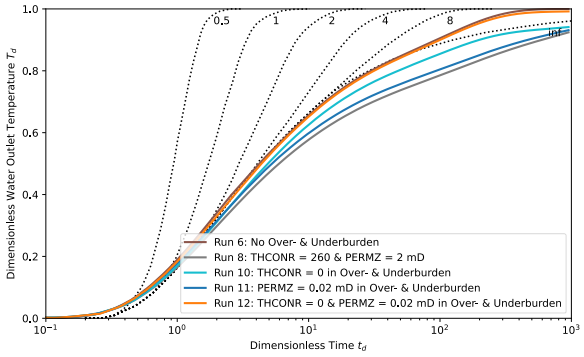
(b) Influence of matrix permeability.



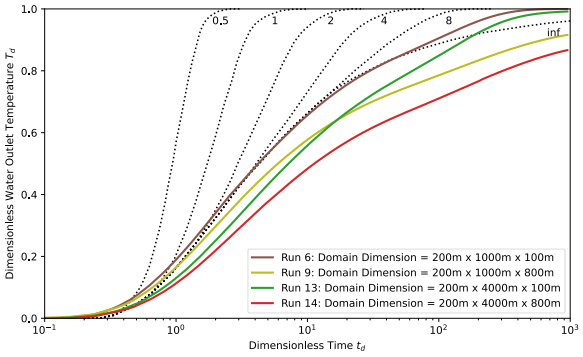
(c) Influence of over-/underburden.



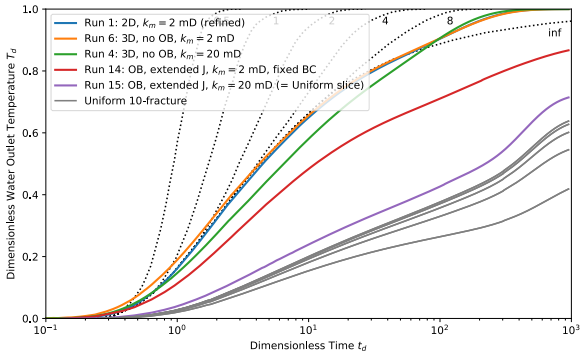
(d) Influence of temperature distribution.



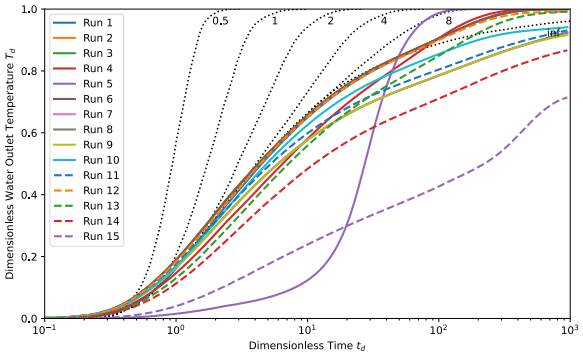
(e) Influence of over-/ underburden properties.



(f) Influence of extension in y-direction.



(g) Comparison to uniform 10-fracture setup.



(h) All runs.

Figure 24: Numerical validation study: dimensionless outlet temperature T_D vs. dimensionless time t_D for the 15 single-fracture setups compared with the Gringarten (1975) analytical curves at $X_{ED} = 0.5, 1, 2, 4, 8$, and ∞.

fluid contributes only $\approx 3\%$, no correction to K_r is needed. `HEATCRT` = 0 confirms temperature-independent heat capacity, consistent with Gringarten.

Three-dimensional versus one-dimensional heat conduction. Gringarten's model restricts heat conduction to the direction perpendicular to the fractures (the x-direction only). OPM solves the full three-dimensional heat equation, so additional heat conducts along the flow direction and in the vertical direction. This extra heat supply slows the cooling of the produced water, shifting the OPM curve to *later* dimensionless times (rightward shift).

Temperature-dependent fluid properties. Gringarten assumes constant ρ_w and c_w. OPM uses fluid property tables: at reservoir conditions (170 °C, 400 bar) the brine density falls to approximately 974 kg/m^3 (versus 1051 kg/m^3 at surface conditions) while c_w rises only slightly, to approximately 4210 J/(kg K). The volumetric thermal capacity $\rho_w c_w$ is therefore only about 7% lower at reservoir conditions than the surface values used in the dimensionless scaling. A lower thermal capacity means the fluid carries less heat per unit volume, so the thermal front advances faster than the surface-property mapping implies. This shifts the OPM curve *earlier* in dimensionless time (leftward shift).

Well completions confined to a single depth layer. Gringarten's multi-fracture solution assumes that fluid is injected and produced uniformly across the full fracture height $y = 100$ m. In the OPM model both wells are completed only at the centre of the fracture zone (`COMPDAT` K = 31–31), corresponding to a single layer of approximately 2.4 m out of the 41-layer, 100 m fracture zone. Comparing Setups 2 (NZ = 1, trivially full-height completion) and 3 (NZ = 41, single completion layer) both yield very good agreement, suggesting this effect is small for the present configuration. Some rightward shift may nonetheless arise from fluid having to spread vertically from the completion layer and from the BTEMP at K = 31 not fully representing the flow-weighted fracture average; this can be tested by extending completions to K = 11–51.

Buoyancy-driven flow redistribution. Gringarten assumes steady, vertically uniform plug flow. In the OPM model, the fractures are vertical and flow is horizontal, so the gravitational body force acts along the fracture height (y-direction, perpendicular to flow). Cold injected water ($\rho \approx 1050\,\mathrm{kg/m^3}$ at 20 °C) is denser than hot reservoir water ($\rho \approx 974\,\mathrm{kg/m^3}$ at 170 °C); it therefore sinks towards the bottom of the 100 m-high fracture while hot water rises towards the top. This gravitational tonguing causes the cold thermal front to advance faster near the fracture base, leading to slightly earlier average breakthrough than the vertically uniform Gringarten model predicts (leftward shift). The effect is limited by the small height-to-length ratio of the fracture ($y/H = 100/1000 = 0.1$), which constrains the vertical segregation distance relative to the horizontal flow path.

Surface versus reservoir volumetric flow rate. The injection rate q is specified in surface conditions (sm^3/s). At elevated reservoir temperatures the in-situ volumetric flux is larger due to the lower water density, so heat is transported slightly faster than the dimensionless time based on surface conditions suggests. This is a minor leftward shift.

Numerical dispersion. The finite-volume grid in OPM introduces artificial diffusion proportional to the grid spacing along the flow direction. This smears the thermal front, causing the outlet temperature to decrease earlier than the true physics warrants (leftward shift). The magnitude scales with cell size; coarser grids produce stronger dispersion.

Geothermal gradient. Gringarten's multi-fracture solution assumes a uniform initial rock temperature ($\beta = 0$ in the notation of [8]), where β is the dimensionless geothermal gradient *along the flow path*. In Gringarten's geometry flow is directed upward from injector to producer; in the present OPM model both wells are completed at the same depth (4000 m) so the flow is horizontal. There is therefore no systematic temperature variation along the flow path, and the $\beta = 0$ assumption is consistent with the OPM setup. The geothermal gradient does impose a temperature variation across the fracture *height* (y-direction): with a gradient of approximately 30 °C/km and a fracture height of $y = 100$ m, the initial rock temperature varies by only ±1.5 °C about the mean of 170 °C. This is less than 1 % of the total temperature difference $T_{\text{rock}} - T_i = 150$ °C and has a negligible effect on the comparison, as confirmed by Setups 3 and 6 in Section A.3.

Fracture thermal mass. Gringarten assumes infinitely thin fractures that contribute no thermal mass; the surrounding rock is treated as fully impermeable (zero matrix porosity). In the OPM model the fracture zone has porosity $\phi_f = 0.151$, while the surrounding matrix rock has porosity $\phi_m = 0.150$. Because $\phi_f \approx \phi_m$, the fracture zone is thermally nearly indistinguishable from matrix rock: the extra fluid volume relative to matrix is only $(\phi_f - \phi_m) \times d_{\text{cell}} = 0.001 \times 1\,\text{m} = 1\,\text{mm}$ of effective aperture per fracture face. The associated additional thermal mass is negligible, and this source of discrepancy is effectively absent in this model.

Summary. Table 8 summarises all sources of discrepancy. The dominant effect is heat supply from the overburden and underburden, which is absent in Gringarten's model and shifts the OPM curve far to the right. The rock heat capacity scaling (30% larger than `HEATCR` alone once pore-fluid storage is included) and three-dimensional heat conduction are moderate contributors. Temperature-dependent fluid properties act in the opposite direction (leftward shift), but are only a minor contributor. The single-layer well completions, buoyancy-driven flow, and surface-to-reservoir flow rate correction are all minor. Geothermal gradient and fracture thermal mass are negligible.

Table 8: Summary of factors causing discrepancy between OPM and the Gringarten (1975) analytical solution for the uniform case. "Left" denotes a shift of the OPM curve to earlier dimensionless time relative to the analytical curve; "right" denotes a shift to later dimensionless time.

Source	Direction	Magnitude
Overburden/underburden heat supply (vertical conduction from K<11, K>51)	Right (beyond $X_{ED} \to \infty$)	Large
Effective rock heat capacity (`HEATCR` + $\phi_m \rho_w c_w$; accounted for in this study)	Right (t_D overest. 30%)	Moderate
3D vs. 1D heat conduction	Right (slower cooling)	Moderate
Fluid properties: surface vs. reservoir $\rho_w c_w$	Left (faster cooling)	Minor
Single-K well completions (K= 31 only; fracture zone spans K= 11–51)	Right (slower apparent cooling)	Minor
Buoyancy-driven flow redistribution (gravitational tonguing)	Left (earlier breakthrough)	Minor
Surface vs. reservoir volumetric flow rate	Left	Minor
Numerical dispersion	Left (faster apparent cooling)	Grid-dependent
Geothermal gradient (±1.5 °C across height y, perpendicular to flow)	Negligible	Negligible
Fracture thermal mass ($\phi_f \approx \phi_m$, effective aperture 1 mm)	Negligible	Negligible

A.5 Outlook

The following modifications to the OPM setup or to the dimensionless scaling would reduce the discrepancy between the simulation and the Gringarten analytical solution, ordered roughly by expected impact.

OPM: restrict thermal conduction to the fracture zone. Set `THCONR` = 0 in all cells outside the fracture zone (K < 11 and K > 51) using an `EQUALS` block. This removes the overburden/underburden heat supply and brings the OPM geometry closer to Gringarten's assumption of an isolated, laterally bounded rock block. Of all possible modifications this is expected to produce the largest leftward shift, as confirmed by Setup 12 in Section A.3.

Analytical: use reservoir-condition fluid properties in the t_D scaling. The dimensionless time formula uses surface-condition values $\rho_w = 1051\,\mathrm{kg/m^3}$ and $c_w = 4180\,\mathrm{J/(kg\,K)}$. At reservoir conditions (170 °C, 400 bar) the brine density falls to approximately $974\,\mathrm{kg/m^3}$ and c_w rises slightly to approximately $4210\,\mathrm{J/(kg\,K)}$, giving $(\rho_w c_w)_{\mathrm{res}} \approx 4.10 \times 10^6\,\mathrm{J/(m^3\,K)}$ versus $(\rho_w c_w)_{\mathrm{surf}} \approx 4.39 \times 10^6\,\mathrm{J/(m^3\,K)}$. Replacing the surface values with reservoir values reduces t_D by approximately $(4.10/4.39)^2 \approx 0.87$, shifting the OPM curve roughly 13 % to the left.

OPM: make the matrix effectively impermeable. Reduce `PERMX`, `PERMY`, and `PERMZ` of all matrix cells to 0.001 mD. This forces all flow through the fractures, as Gringarten assumes. Because the matrix already carries only ≈ 5 % of the total flow, this modification is expected to have a minor effect, but it removes one source of non-Gringarten behaviour cleanly.

Analytical: account for vertical heat supply in the effective X_{ED}. The extra heat from overburden and underburden can be approximately absorbed into an effective block half-thickness by noting that the vertical diffusion depth at time t is $\delta(t) = \sqrt{2\alpha t}$ (with $\alpha = K_r/(\rho c)_{\mathrm{eff}}$). Adding δ to x_E gives a time-dependent effective dimensionless spacing:

$$X_{ED}^{\mathrm{eff}}(t) = \frac{\rho_w c_w}{K_r} \cdot \frac{Q}{H}\left(x_E + \sqrt{2\alpha t}\right). \tag{40}$$

At $t = 20\,\mathrm{yr}$ this gives $\delta \approx 36\,\mathrm{m}$ ($X_{ED}^{\mathrm{eff}} \approx 23$), and at $t = 1000\,\mathrm{yr}$ $\delta \approx 257\,\mathrm{m}$ ($X_{ED}^{\mathrm{eff}} \approx 61$); in both cases X_{ED}^{eff} is firmly in the $X_{ED} \to \infty$ regime. This is consistent with the OPM curve lying to the right of the $X_{ED} \to \infty$ reference, and confirms that the $X_{ED} = \infty$ Gringarten curve is the most appropriate comparison for the current setup.

B Reservoir Simulation Results: Uniform, Reference, Prior & Posterior

This appendix documents the flow response of the reservoir model underlying the data assimilation experiments of the main study. Sections B.1 and B.2 characterise the reservoir-flow behaviour (thermal front, pressure footprint, and tracer transport) that underlies the geomechanical observables — first for an idealised *uniform* configuration in which all ten fractures are identical, then for the *non-uniform* fracture-property field of the reference realisation (Section 3.9); both use the base-case settings of the main study, i.e. open lateral boundaries (Section 3.5), $k_m = 20\,\mathrm{mD}$, and $multpv = 50$. Section B.3 then varies the lateral boundary condition and the matrix permeability in a deterministic sensitivity study and justifies the open lateral boundaries adopted as the base case. Finally, Section B.4 documents the effect of the data assimilation itself, contrasting the prior and posterior ensembles in terms of the simulated reservoir fields and the production response.

Gridded fields are shown throughout in three section views: a y–z section along fracture 1 (the fracture closest to the $x = 0$ model edge), a y–z section along a second fracture in the interior of the array, and an x–y plan view at the well layer (K = 31, 4000 m depth). The quantities shown are the temperature, the overpressure $p - p_0$ relative to the initial pressure field, and the injected tracer concentration.

In addition, the production response is resolved per fracture. Both wells are completed only in the fracture cells (seven completions per fracture, see Fig. 2), so each fracture i is associated with the set C_i of producer completions lying in its plane. The per-fracture production rate is the sum of the connection water rates q_c (reported at surface conditions),

$$Q_i = \sum_{c \in C_i} q_c, \tag{41}$$

the production temperature is the flow-weighted average of the completion block temperatures T_c,

$$T_i = \frac{\sum_{c \in C_i} q_c\, T_c}{\sum_{c \in C_i} q_c}, \tag{42}$$

and the delivered thermal power per fracture and in total follows from an energy balance over the doublet, assuming isothermal flow in the wells and evaluating all properties at surface conditions,

$$P_i = \rho_w\, c_{p,w}\, Q_i \left(T_i - T_{\mathrm{inj}}\right), \qquad P = \sum_{i=1}^{10} P_i, \tag{43}$$

with the surface water density $\rho_w = 1051\,\mathrm{kg\,sm^{3-1}}$, the pure-water specific heat capacity $c_{p,w} = 4.18\,\mathrm{kJ\,kg^{-1}\,K^{-1}}$, and the injection temperature $T_{\mathrm{inj}} = 20\,^{\circ}\mathrm{C}$. Equation (43) is an approximation: it neglects heat exchange between the wellbores and the surrounding rock and replaces the brine enthalpy difference by a constant pure-water specific heat. Both simplifications overpredict the delivered power — by 5–7 % through the salinity-reduced specific heat of the brine, plus the heat lost by the hot production stream to the colder shallow rock on its 4 km ascent. Accurate wellbore modelling and thermal-power prediction are beyond the scope of this study, however; Eq. (43) serves only as a consistent measure for comparing configurations and realisations.

The extracted thermal power translates into net electric power through a conversion efficiency η,

$$P_{\mathrm{el}} = \eta\, P. \tag{44}$$

At production temperatures of 140–170 °C the conversion is performed by binary (organic Rankine or Kalina cycle) plants, whose net efficiencies typically lie between 6 % and 12 %; a worldwide review of operating geothermal plants reports an average of 12 % across all technologies, with the highest values reserved for flash and dry-steam resources well above 200 °C [20, 21]. Deep geothermal doublets in the Upper Rhine

Graben operate at temperatures comparable to this study: the Soultz-sous-Forêts EGS plant generates about 1.5 MW of electric power from a resource reported between 155 °C [20] and 175 °C at the plant [21], and the neighbouring Insheim plant about 5.0 MW [21] from about 165 °C. With $\eta \approx 0.1$, the 64–77 MW of thermal power extracted by the doublet studied here corresponds to roughly 6–8 MW of electric power.

B.1 Uniform Fracture Properties

In the uniform configuration every fracture is assigned the same effective properties, obtained by upscaling a parallel-plate fracture of aperture $a_f = 1\,\text{mm}$ (cubic law, $k_f = a_f^2/12$) into the 1 m-wide fracture cell by arithmetic averaging with the surrounding matrix: $k = 83\,353\,\text{mD}$ and $\phi = 0.151$ for $k_m = 20\,\text{mD}$. All fractures therefore conduct like the open (H) regime of Table 4, and any difference between the ten fracture responses is caused solely by the position of the fracture within the array — most visibly the outermost fractures, which exchange heat with a larger undisturbed rock volume. This idealised symmetric configuration is the three-dimensional counterpart of the Gringarten setup studied in Section A.

In the section views (Figs. 25 to 27), the cold plume nucleates at the injector and advances along the fracture towards the producer, remaining vertically confined to a narrow band around the well layer and sagging slightly downwards at late times as the cold, denser water sinks. The plan view confirms that all ten fractures respond near-identically, while the matrix between them stays essentially at the initial temperature. The overpressure forms an antisymmetric injector–producer dipole of only a few bar that is fully established within the first month and barely evolves thereafter, decaying smoothly towards the open boundaries. The tracer, by contrast, traverses the fracture within weeks and subsequently widens into a wedge-shaped tagged region between the wells — illustrating the separation of time scales between pressure, tracer, and thermal transport.

Fig. 28 summarises the per-fracture production response. Although all fractures are identical, their position in the array leaves a clear, symmetric imprint: after ten years the outermost fractures produce the warmest water (149 °C, against 142 °C at the centre of the array) and carry the highest rate. The temperature ordering follows from heat supply: an interior fracture competes with its neighbours for the heat stored in the shared matrix blocks, whereas the outermost fractures border a laterally undisturbed rock volume that keeps replenishing heat. The rate ordering has two contributions. Already at the first report step, before any thermal effect, the outer fractures carry about 15 % more flow than the central ones (1100 versus $960\,\text{sm}^3\,\text{d}^{-1}$) — pure pressure interference within the parallel array, in analogy to a row of equal wells in which the outermost wells capture the largest share, aided by the matrix inflow available on the outward-facing side. Over time this spread then widens (the outer fractures grow to $1175\,\text{sm}^3\,\text{d}^{-1}$ while the central ones drop to $924\,\text{sm}^3\,\text{d}^{-1}$; the total is fixed by the well control), which is a thermal effect: the viscosity of water increases from about 0.16 mPa s at the initial 170 °C to about 1.0 mPa s at the injection temperature, and the cold plumes of the interior fractures merge into one cold, viscous region, so the interior flow paths become progressively more resistive and flow shifts towards the warmer outer fractures. The feedback is self-limiting — the additional flow also cools the outer fractures faster — which keeps the redistribution at the level of a few percent. Since the delivered thermal power is the product of rate and temperature, both effects act in the same direction and the outer fractures contribute the largest thermal power.

Fig. 29 shows the corresponding well response. The doublet is balanced (both wells operate at a fixed rate of $10\,000\,\text{sm}^3\,\text{d}^{-1}$), and with open lateral boundaries the field-average pressure remains pinned at its initial 400 bar throughout. The differential pressure between the wells, however, grows from 3.8 bar in the first weeks to 6.0 bar after ten years (the injector rising by 1.3 bar, the producer falling by 0.9 bar). With fixed rates and unchanged permeability, the inter-well pressure drop is proportional to the viscosity integrated along the flow paths, and the expanding cold region is up to six times more viscous than the native 170 °C brine; the growth of the differential pressure is therefore essentially a viscosity effect (the higher density

of the cold water enters only weakly, through buoyancy, since injector and producer are completed at the same depth). The tracer, injected continuously with the cold water, breaks through along the fracture planes within weeks and peaks at a produced concentration of 0.55 after about three months; the remainder of the produced stream is resident, tracer-free matrix water drawn in laterally near the producer. The subsequent decline to 0.36 — despite continuous tracer injection — is the same viscosity-driven flow redistribution seen in the differential pressure: as the cold plume grows, the direct fracture paths from the injector become more resistive and carry a diminishing share of the produced stream, while the producer draws an increasing share of untagged matrix and far-field water. Only on time scales well beyond the simulated ten years, once the matrix water in the capture region of the producer is displaced, would the produced concentration rise again. The total thermal power declines from 77 MW to 64 MW, a reduction of approximately 17 % over ten years of operation.

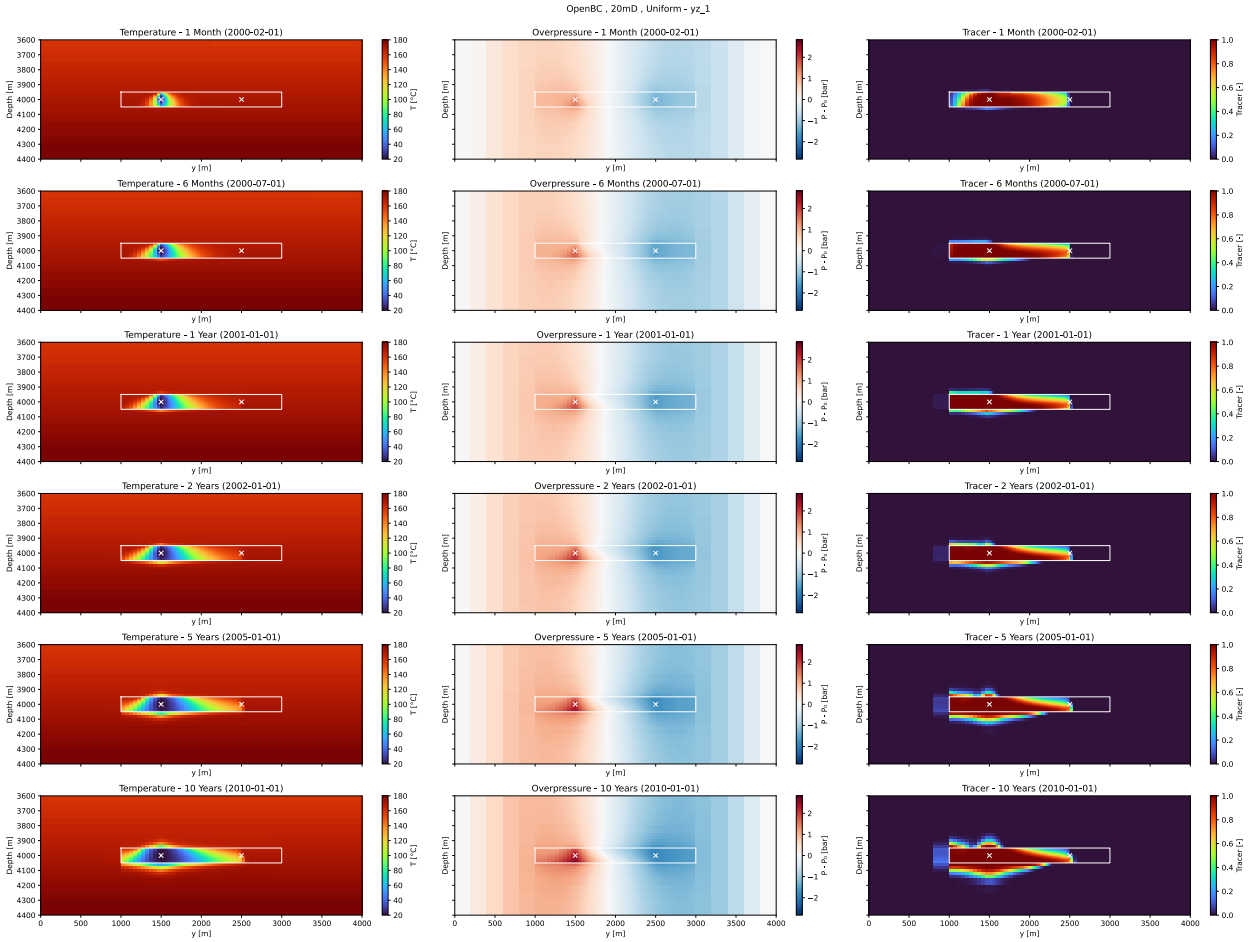


Figure 25: Uniform configuration (open boundaries, $k_m = 20\,\text{mD}$): y–z section along fracture 1. Columns: temperature, overpressure $p - p_0$, and tracer concentration; rows: 1 month, 6 months, 1, 2, 5, and 10 years after the start of injection. The horizontal injector and producer sections lie at $y = 1500\,\text{m}$ and $y = 2500\,\text{m}$ at 4000 m depth.

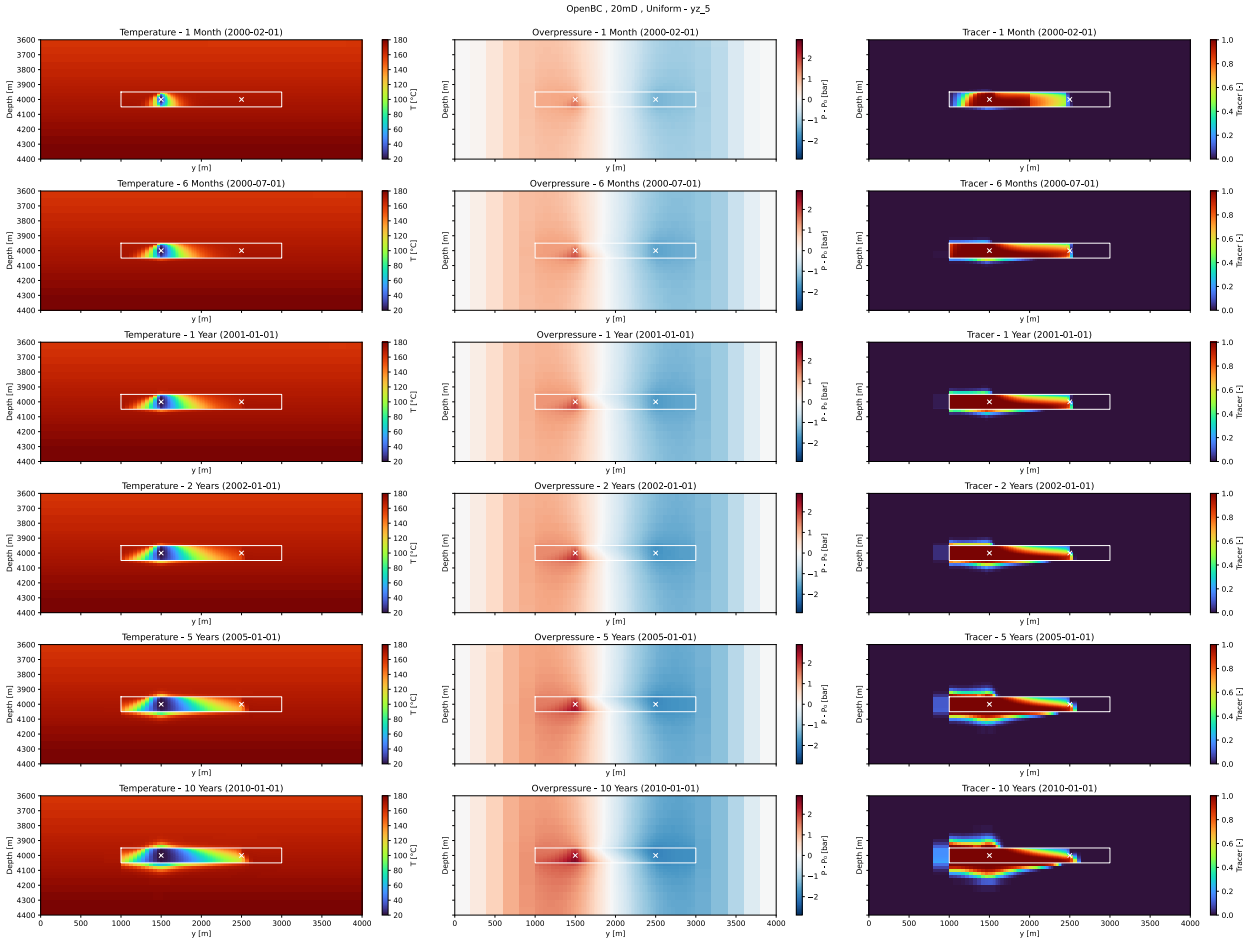


Figure 26: Uniform configuration (open boundaries, $k_m = 20\,\text{mD}$): y–z section along fracture 5, in the interior of the fracture array. Layout as in Fig. 25.

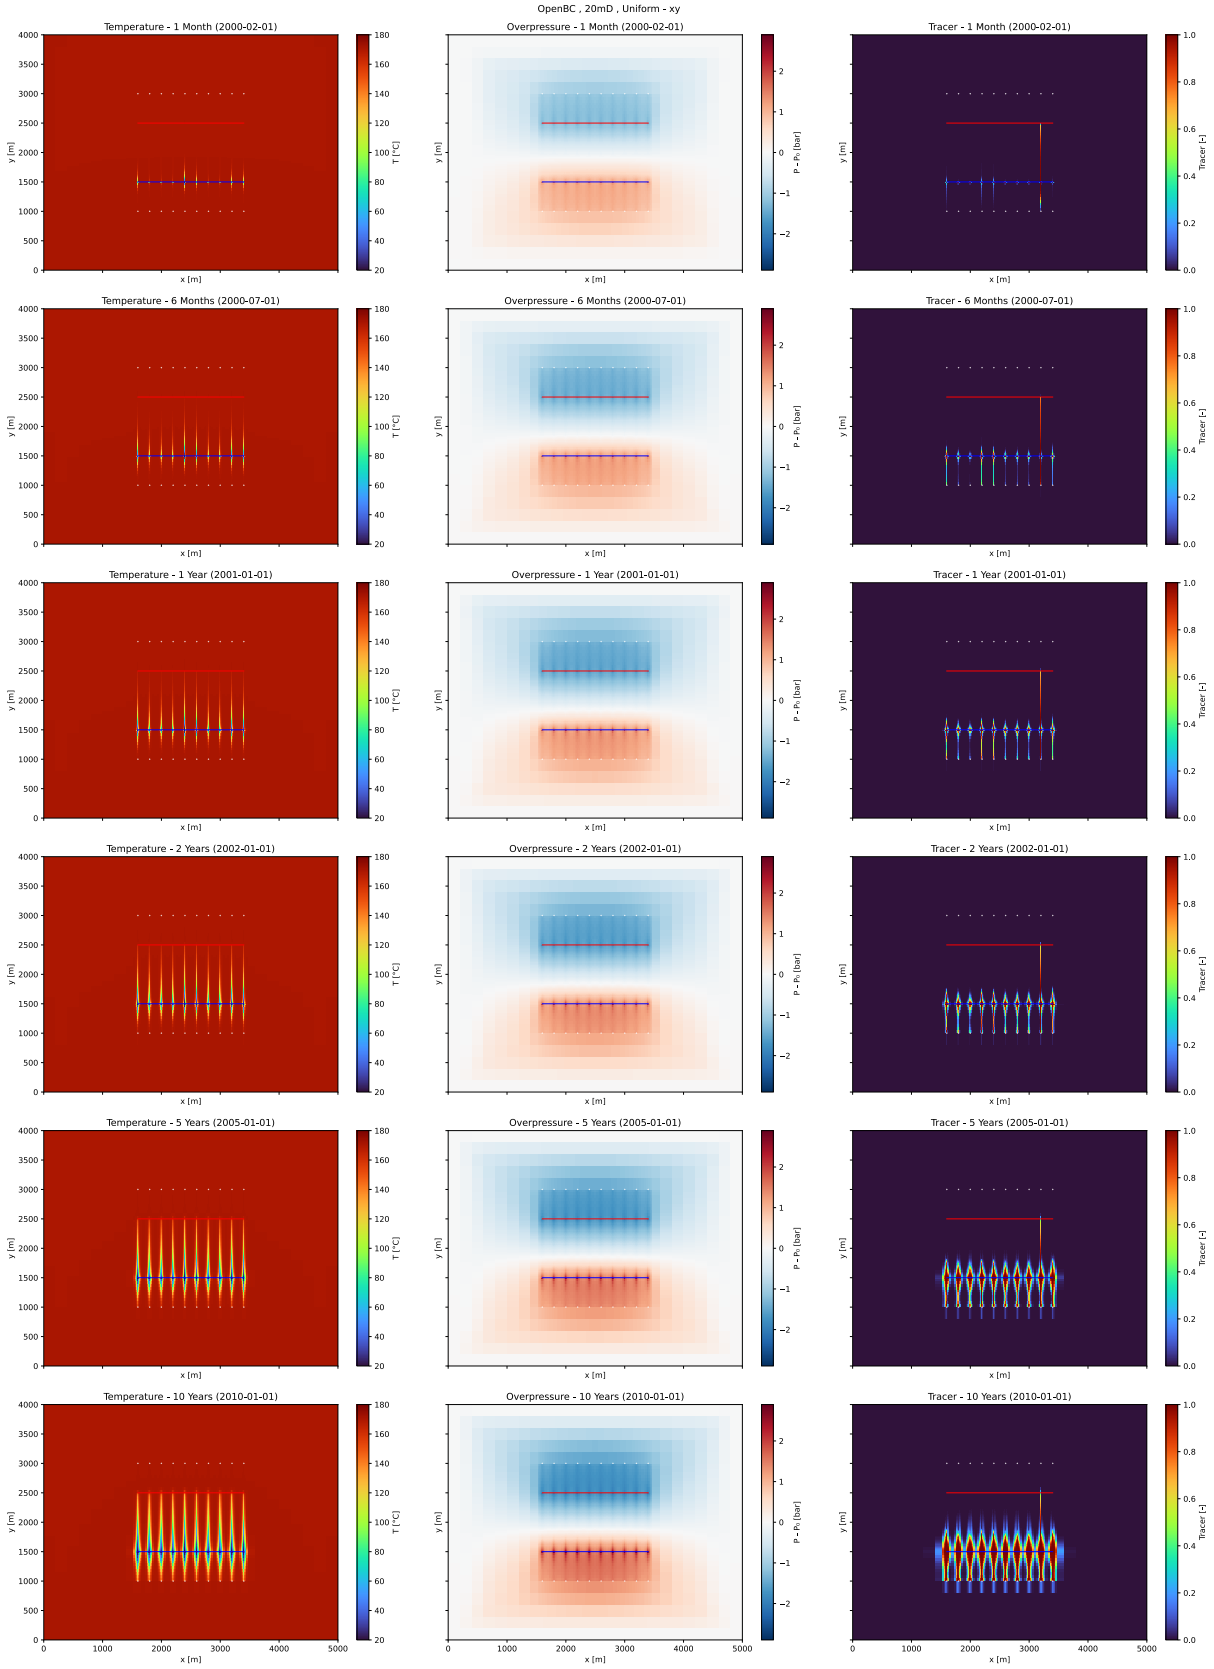


Figure 27: Uniform configuration (open boundaries, $k_m = 20$ mD): x–y plan view at the well layer (K = 31, 4000 m depth). Columns: temperature, overpressure $p - p_0$, and tracer concentration; rows: 1 month to 10 years. All ten fracture columns respond near-identically, up to edge effects at the outermost fractures.

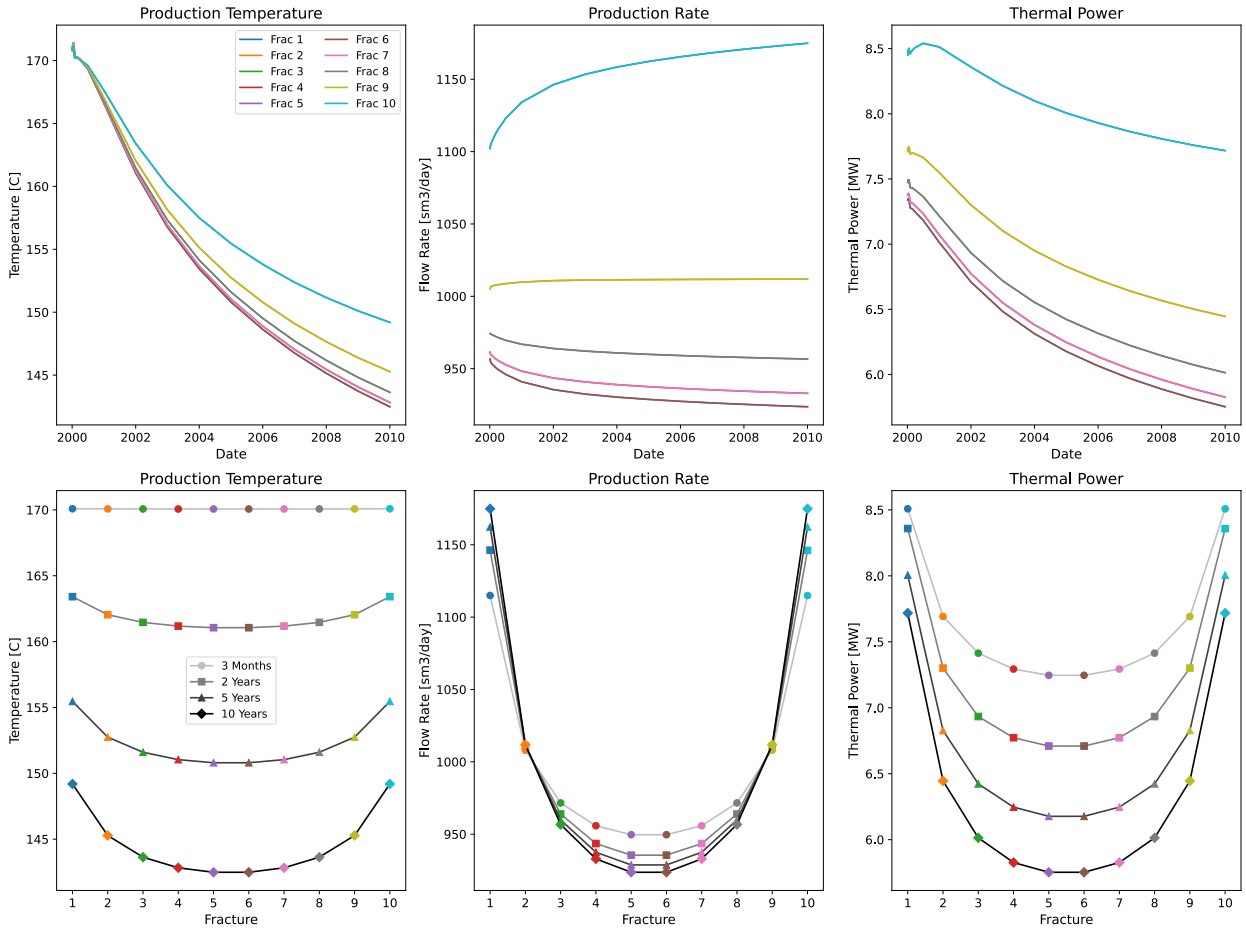


Figure 28: Uniform configuration (open boundaries, $k_m = 20\,\text{mD}$): per-fracture production response. Top row: production temperature, production rate, and thermal power over time (colour = fracture). Bottom row: the same quantities per fracture at 3 months, 2, 5, and 10 years after the start of injection (light to dark).

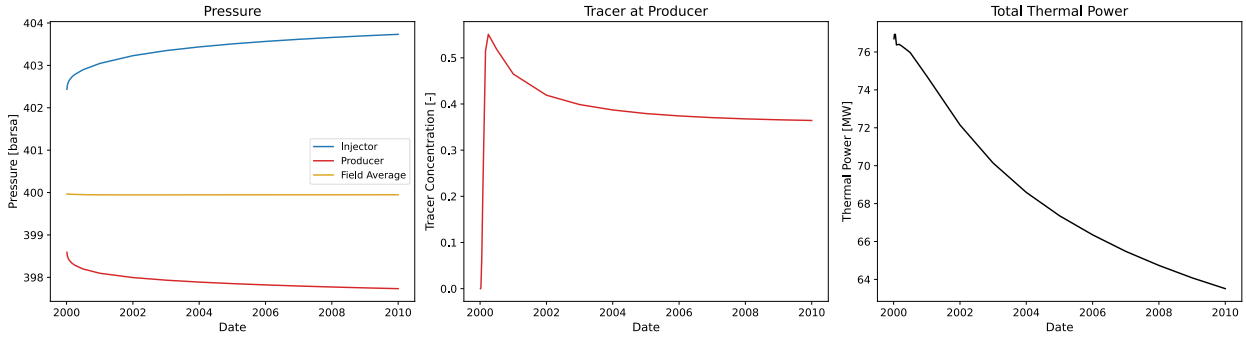


Figure 29: Uniform configuration (open boundaries, $k_m = 20\,\text{mD}$): bottom-hole pressure at the injector and producer with the field-average pressure (left), tracer concentration at the producer (centre), and total thermal power summed over all fractures (right).

B.2 Reference Realisation

The non-uniform configuration with open boundaries and $k_m = 20\,\mathrm{mD}$ is the reference ("truth") realisation of the main study: the fracture regime assignment M H L L M L H H M H of Table 5 with the class properties of Table 4. Its input deck is identical to that of the reference simulation from which the synthetic strain observations are generated (Section 3.9). The same views as in Section B.1 are shown, with the interior y–z section now taken along fracture 7: fracture 1 belongs to the partially open (M) regime (1600 mD) and fracture 7 to the open (H) regime (80 000 mD), so the two sections contrast the conductive regimes directly, while the plan view adds the closed (L) fractures.

The section views make the flow concentration directly visible: along the partially open fracture 1 (Fig. 30), the cold plume and the tracer barely leave the injector even after ten years, in stark contrast to the uniform case, whereas along the open fracture 7 (Fig. 31) the tracer breaks through within the first month and the thermal front advances most of the way to the producer, whose completions on this fracture cool by more than 30 °C. The plan view (Fig. 32) shows the same contrast across the whole array: cold and tracer fingers reach the producer only along the four open (H) fractures, with the closed (L) fractures indistinguishable from the surrounding matrix. The overpressure dipole is no longer uniform across the array but concentrates on the conductive fracture columns, and its amplitude is more than twice that of the uniform configuration — consistent with the larger differential pressure discussed below.

In the per-fracture production response (Fig. 33) the fracture class now dominates over the position in the array. The four open (H) fractures (2, 7, 8, and 10) carry about 94 % of the produced water (2100–2600 $\mathrm{sm^3\,d^{-1}}$ each), the three partially open (M) fractures contribute about 5 % (150–200 $\mathrm{sm^3\,d^{-1}}$), and the closed (L) fractures are hydraulically irrelevant (roughly 60 $\mathrm{sm^3\,d^{-1}}$, i.e. matrix-like inflow). Thermal breakthrough is accordingly confined to the H fractures, whose production temperature falls to 136–145 °C within ten years, while the M and L fractures still produce at the initial 170 °C — but at rates too small to matter, so the thermal power is delivered almost exclusively by the H fractures. Within the H class, the positional effect of Section B.1 reappears: the interior pair (7, 8) cools fastest and loses flow over time, whereas fractures 2 and 10, which border larger undisturbed rock volumes, remain 5–9 °C warmer and gain flow.

The well response (Fig. 34) is qualitatively that of the uniform case, amplified by the flow concentration: the differential pressure grows from 8.6 bar to 13.7 bar — more than twice the uniform values — because the same total rate is forced through four instead of ten conductive fractures, which deepens the cold plumes and strengthens the viscosity feedback; the relative growth over the ten years ($\approx 60\,\%$) is the same. The tracer breaks through faster (its peak is reached already after two months, reflecting the higher flow velocities in the H fractures) but at a lower level, 0.39 declining to 0.24: with only four fast paths the swept region is narrower, and a larger share of the produced stream is drawn from the untagged matrix. The total thermal power declines from 77 MW to 63 MW (18 %), barely distinguishable from the uniform configuration: with rate-controlled wells, the field-integrated quantities are almost insensitive to the fracture-property field, whereas the per-fracture temperatures and rates differ strongly — which is precisely what motivates the per-fracture observations assimilated in the main study.

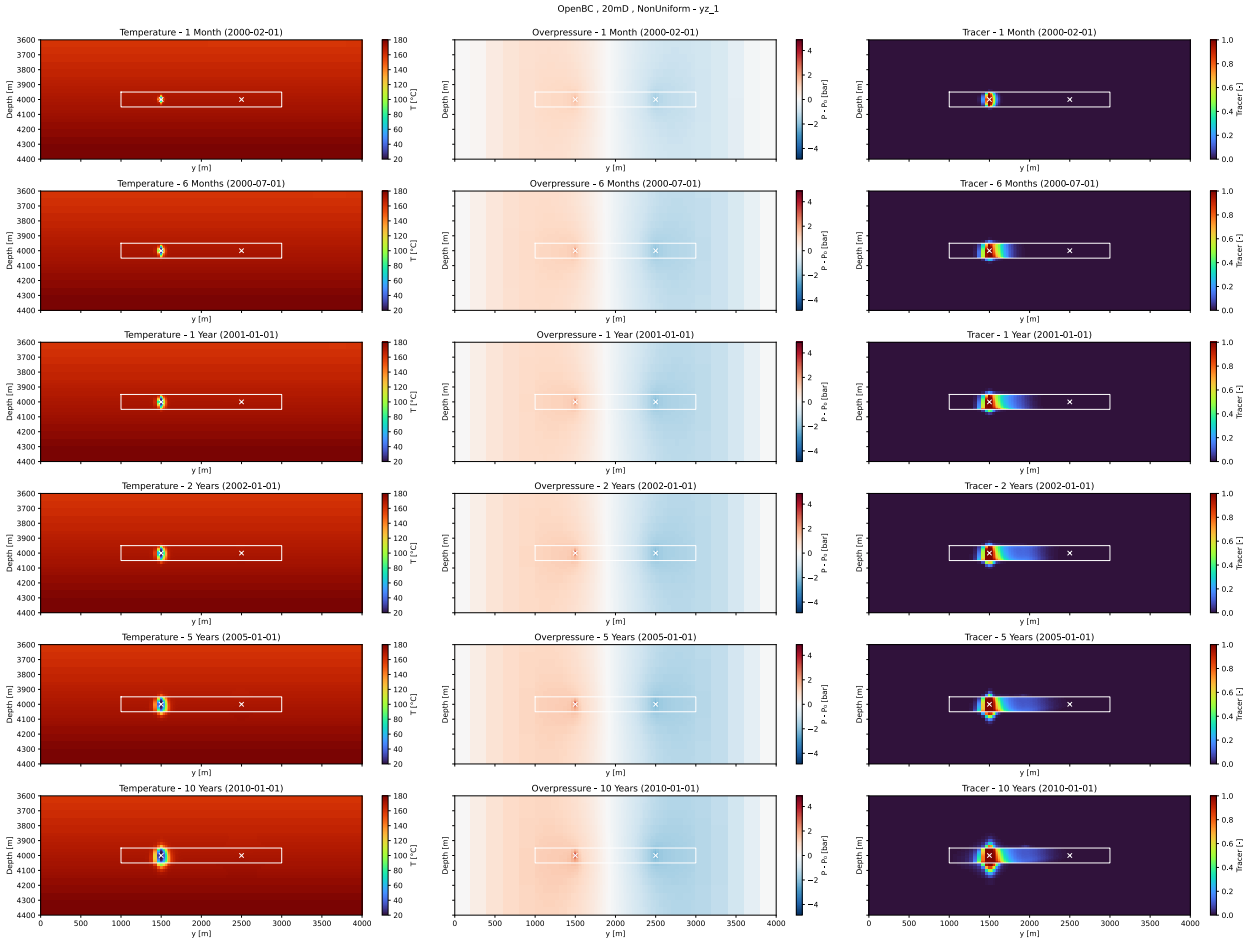


Figure 30: Reference realisation (non-uniform, open boundaries, $k_m = 20\,\text{mD}$): y–z section along fracture 1 (partially open, 1600 mD). Columns: temperature, overpressure $p - p_0$, and tracer concentration; rows: 1 month, 6 months, 1, 2, 5, and 10 years after the start of injection.

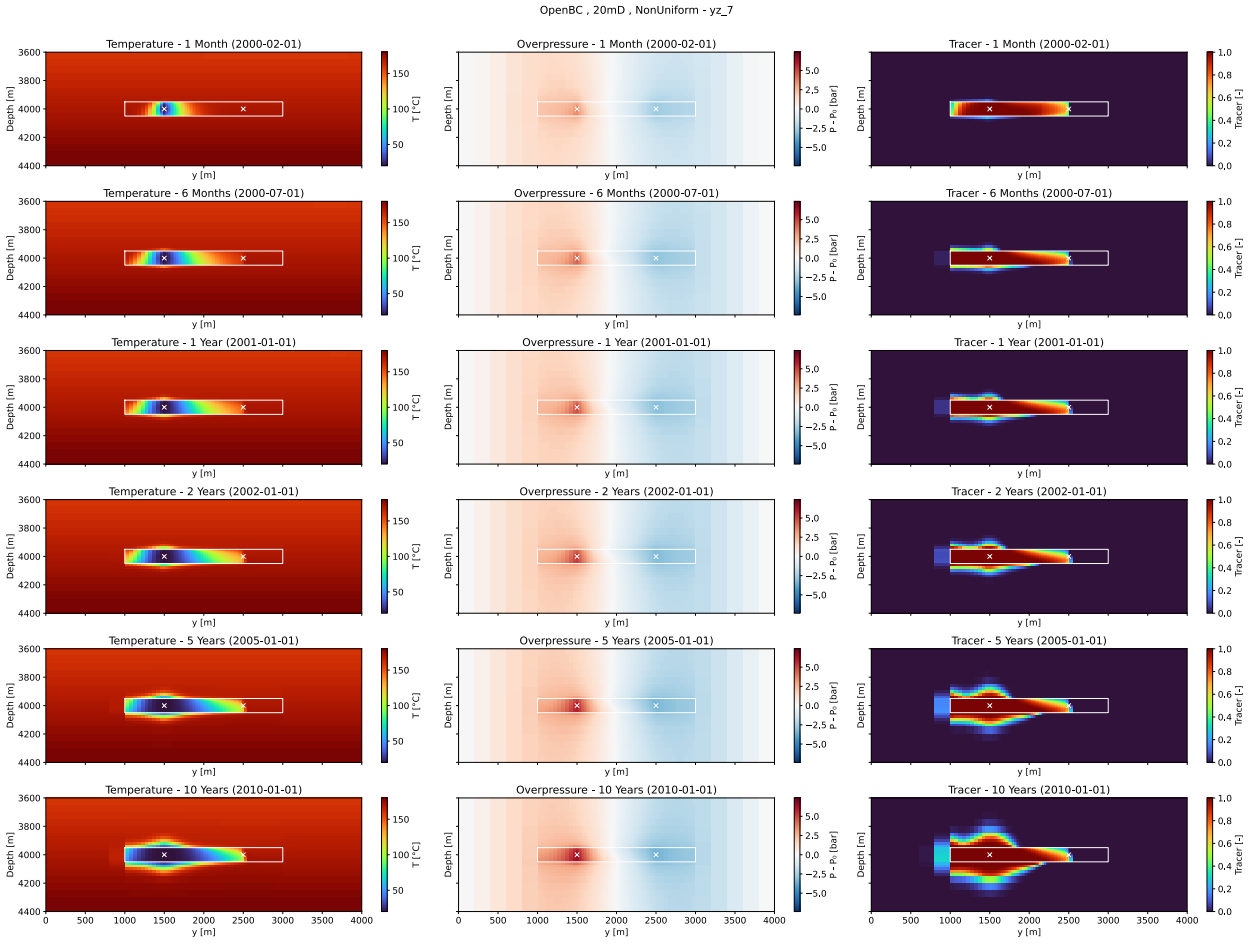


Figure 31: Reference realisation (non-uniform, open boundaries, $k_m = 20$ mD): y–z section along fracture 7 (open, 80 000 mD). Layout as in Fig. 30.

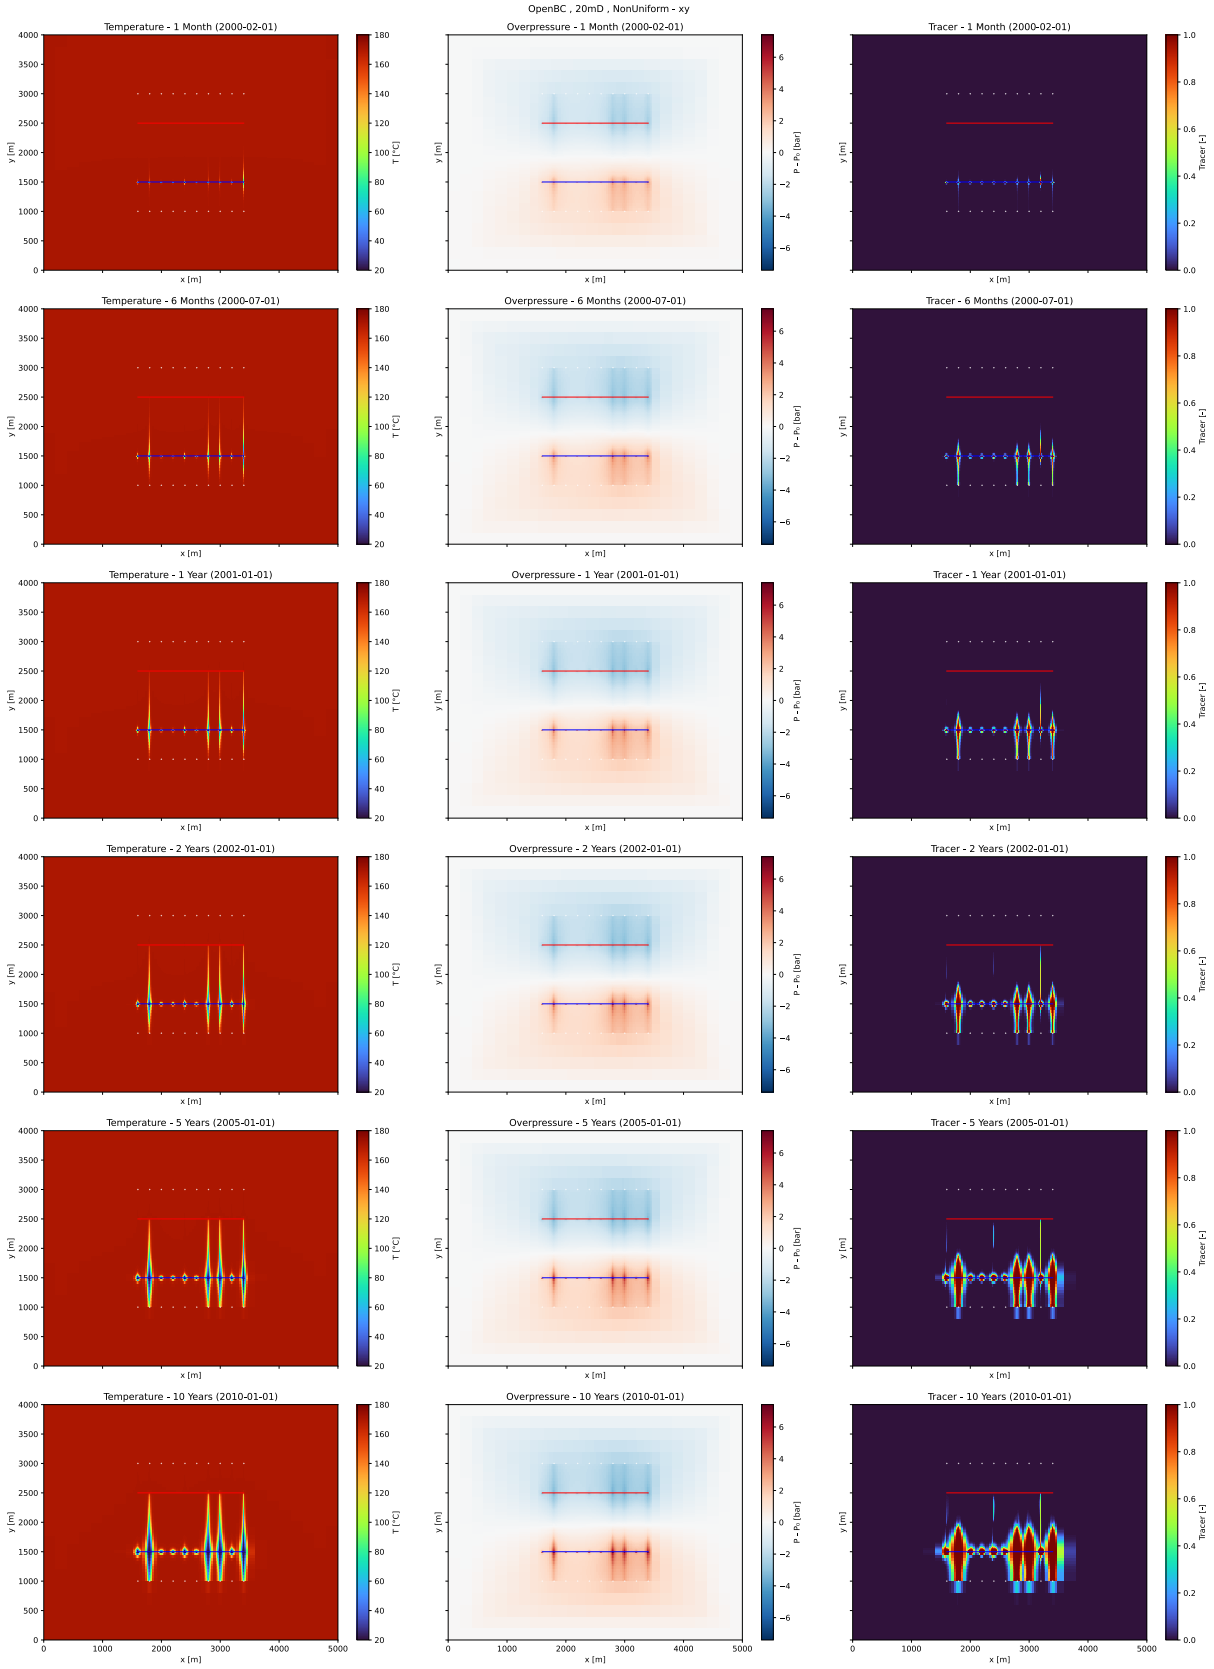


Figure 32: Reference realisation (non-uniform, open boundaries, $k_m = 20\,\text{mD}$): x–y plan view at the well layer (K = 31, 4000 m depth). Columns: temperature, overpressure $p - p_0$, and tracer concentration; rows: 1 month to 10 years. The open (H) fractures develop the deepest cold plumes and the strongest tracer breakthrough, while the closed (L) fractures remain essentially undisturbed.

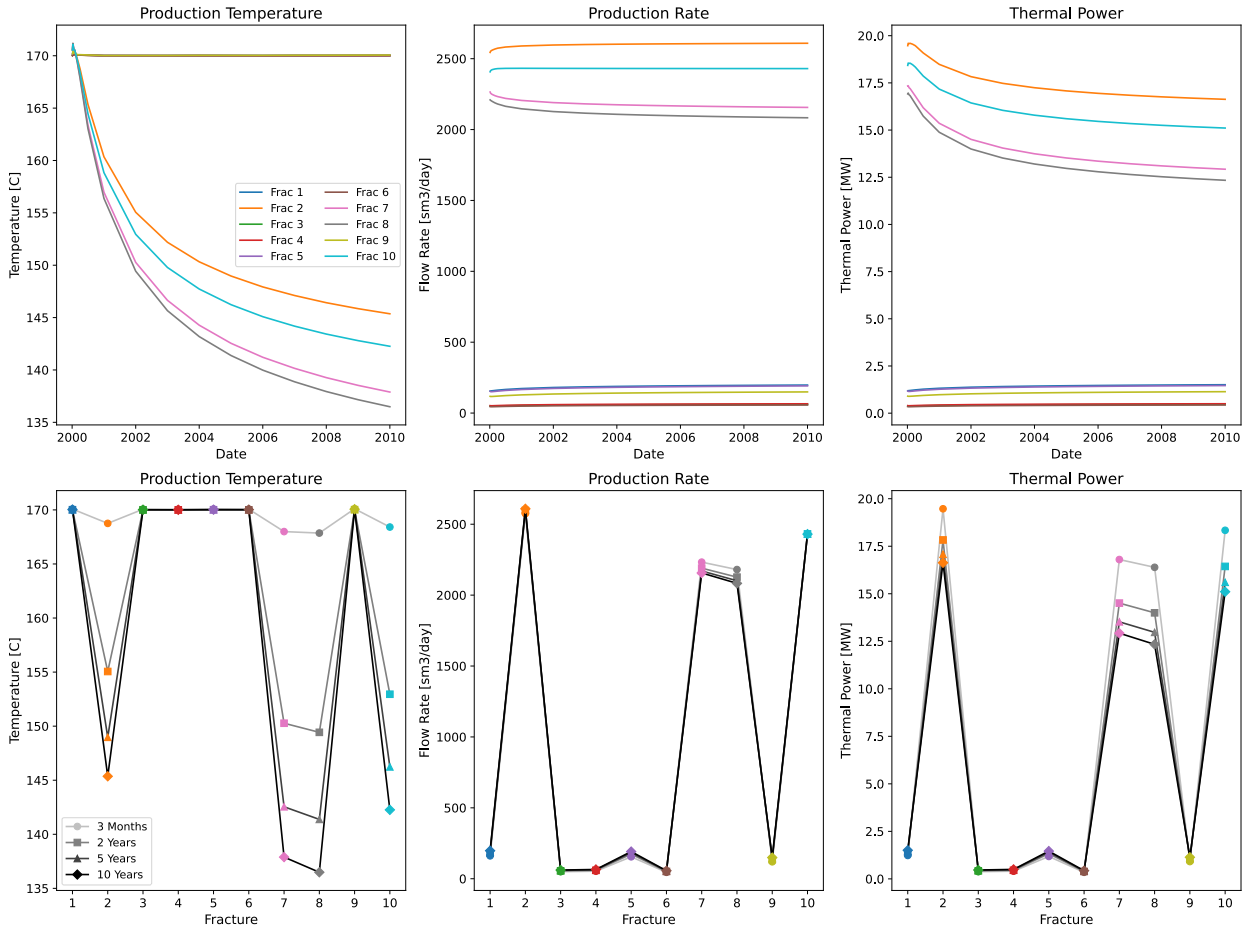


Figure 33: Reference realisation (non-uniform, open boundaries, $k_m = 20\,\text{mD}$): per-fracture production response. Top row: production temperature, production rate, and thermal power over time (colour = fracture). Bottom row: the same quantities per fracture at 3 months, 2, 5, and 10 years after the start of injection (light to dark).

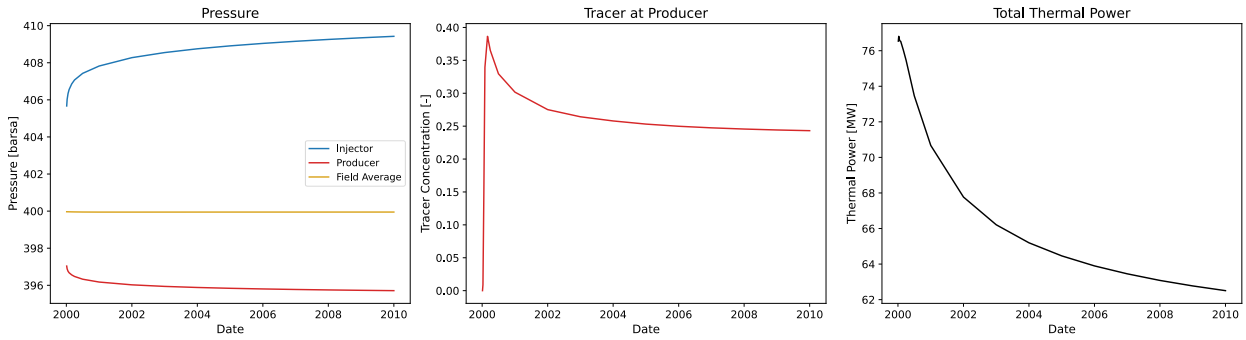


Figure 34: Reference realisation (non-uniform, open boundaries, $k_m = 20\,\text{mD}$): bottom-hole pressure at the injector and producer with the field-average pressure (left), tracer concentration at the producer (centre), and total thermal power summed over all fractures (right).

B.3 Influence of Boundary Conditions and Matrix Permeability

The influence of the lateral boundary condition and of the matrix permeability is assessed with eight deterministic configurations, spanning open versus fully closed lateral boundaries (Section 3.5), matrix permeabilities k_m of 20 mD and 2 mD, and the two fracture-property fields of Sections B.1 and B.2; all configurations use the reference boundary pore-volume multiplier *multpv* = 50. Figs. 35 and 36 compare the temperature, overpressure, and tracer fields of all eight configurations ten years after the start of injection. The boundary condition leaves no visible imprint on the temperature and tracer fields and changes the overpressure only marginally. The matrix permeability, in contrast, is a first-order parameter: with $k_m = 2$ mD the flow is confined almost entirely to the fracture planes, so the cold plume and the tracer form a thin band along the well layer with little leak-off into the matrix, the thermal front travels farther along the fractures and reaches the producer within the simulated period, and the overpressure dipole is several times stronger because the same rate must be driven through a much less permeable system.

The production response (Fig. 37) confirms this picture. The open- and closed-boundary curves overlap to within line width for every fracture and every quantity, so the boundary condition is irrelevant for the production behaviour over the ten-year horizon. This insensitivity is by construction: the lateral boundaries lie at least 1.5 km from the doublet, and the graded *multpv* = 50 pore-volume buffer (Section 3.5) endows the outermost cell rings with the storage of a much larger formation volume, so the balanced doublet never probes the difference between the sealed wall and the Dirichlet faces within the simulated period — apart from the slow field-wide depressurisation quantified below. Reducing the matrix permeability by a factor of ten, however, changes the thermal performance qualitatively: while the 20 mD configurations end at production temperatures of 142–149 °C, the 2 mD configurations drop to 100–105 °C (uniform field), because the fracture flow is fed by a much smaller matrix volume and sweeps correspondingly less rock heat. The per-fracture spread between outer and interior fractures also shrinks, reflecting the reduced matrix participation. The total delivered power (Fig. 38) accordingly declines by 46 % (uniform) and 58 % (non-uniform) over ten years at 2 mD, compared to 17–18 % at 20 mD.

The reservoir-scale balance is summarised in Fig. 39 and Table 9. The mass-weighted average temperature declines by only about 0.02 °C in all configurations — the extracted heat is negligible compared to the heat content of the model domain. The cooled water, however, contracts, and how this contraction is accommodated depends on the boundary condition: with open boundaries the deficit is replenished by lateral inflow at essentially constant pressure (a net mass gain of 0.93 Mt and 0.43 Mt for the uniform field at 20 mD and 2 mD, respectively), whereas in the sealed reservoir the mass stays constant and the contraction depressurises the system instead, the average pressure falling by 0.2–0.4 bar over the ten years. The ratio of pressure to temperature change makes the transition explicit: it grows from about 0.9 bar °C^{-1} for the open, 20 mD configurations, through about 8 bar °C^{-1} for the open, 2 mD configurations — where the tight matrix throttles the boundary inflow — to 17–18 bar °C^{-1} for the closed reservoirs.

Both ratios in Table 9 are quantitatively consistent with the equation of state of water at the mean reservoir conditions of 170 °C and 400 bar, where $(\partial\rho/\partial T)_p \approx -0.92\,\mathrm{kg\,m^{-3}\,°C^{-1}}$ and $(\partial\rho/\partial p)_T \approx 0.053\,\mathrm{kg\,m^{-3}\,bar^{-1}}$, corresponding to an isothermal compressibility $\kappa_T = \rho^{-1}(\partial\rho/\partial p)_T \approx 5.8\times10^{-5}\,\mathrm{bar^{-1}}$ (IAPWS values for pure water). In the open configurations the pressure stays essentially constant, so the tabulated $\Delta\rho/\Delta T$ is an average of the isobaric derivative over the cooled region; its magnitude of 0.4–0.7 kg m^{-3} °C^{-1} lies below the in-situ value because the expansivity of water decreases strongly towards the injection temperature (the mean slope between 170 °C and 20 °C is $-0.64\,\mathrm{kg\,m^{-3}\,°C^{-1}}$). In the sealed reservoir, by contrast, both the water mass and the pore volume are fixed, so the average

Table 9: Start-to-end (ten-year) changes of the reservoir-scale balance for the eight configurations: total water mass, average density, mass-weighted average temperature and pressure, and the density-to-temperature and pressure-to-temperature ratios.

		Open		Closed	
Property	Unit	20 mD	2 mD	20 mD	2 mD
Uniform					
Mass change	Mt	0.927	0.433	−0.001	−0.001
Density change	$\mathrm{kg\,m^{-3}}$	0.0148	0.0071	0.0006	0.0005
Temperature change	°C	−0.021	−0.016	−0.023	−0.017
Pressure change	bar	−0.019	−0.129	−0.393	−0.305
$\Delta\rho/\Delta T$	$\mathrm{kg\,m^{-3}\,°C^{-1}}$	−0.713	−0.434	−0.026	−0.027
$\Delta p/\Delta T$	$\mathrm{bar\,°C^{-1}}$	0.9	7.9	17.3	17.7
Non-uniform					
Mass change	Mt	0.860	0.313	−0.001	−0.001
Density change	$\mathrm{kg\,m^{-3}}$	0.0137	0.0051	0.0006	0.0003
Temperature change	°C	−0.020	−0.012	−0.022	−0.013
Pressure change	bar	−0.017	−0.095	−0.365	−0.222
$\Delta\rho/\Delta T$	$\mathrm{kg\,m^{-3}\,°C^{-1}}$	−0.681	−0.419	−0.025	−0.026
$\Delta p/\Delta T$	$\mathrm{bar\,°C^{-1}}$	0.9	7.8	16.6	17.2

density cannot change — the tabulated $\Delta\rho/\Delta T$ collapses to nearly zero — and the density balance $\mathrm{d}\rho = (\partial\rho/\partial T)_p\,\mathrm{d}T + (\partial\rho/\partial p)_T\,\mathrm{d}p = 0$ instead dictates the depressurisation,

$$\left.\frac{\Delta p}{\Delta T}\right|_{\rho} = -\frac{(\partial\rho/\partial T)_p}{(\partial\rho/\partial p)_T} \approx 17.2\,\mathrm{bar\,°C^{-1}}, \tag{45}$$

in close agreement with the 16.6–17.7 $\mathrm{bar\,°C^{-1}}$ observed for the closed configurations: the sealed response is simply the isochoric limit of the water equation of state.

In summary, the boundary condition affects neither the production response nor the observables assimilated in the main study over the ten-year horizon; it merely sets the level of the field-average pressure, which drifts by less than half a bar even in the fully sealed case. The open lateral boundaries are therefore adopted as the base case (Section 3.5). The matrix permeability, by contrast, is a first-order parameter for the long-term thermal performance, roughly halving the delivered power over ten years when reduced from 20 mD to 2 mD.

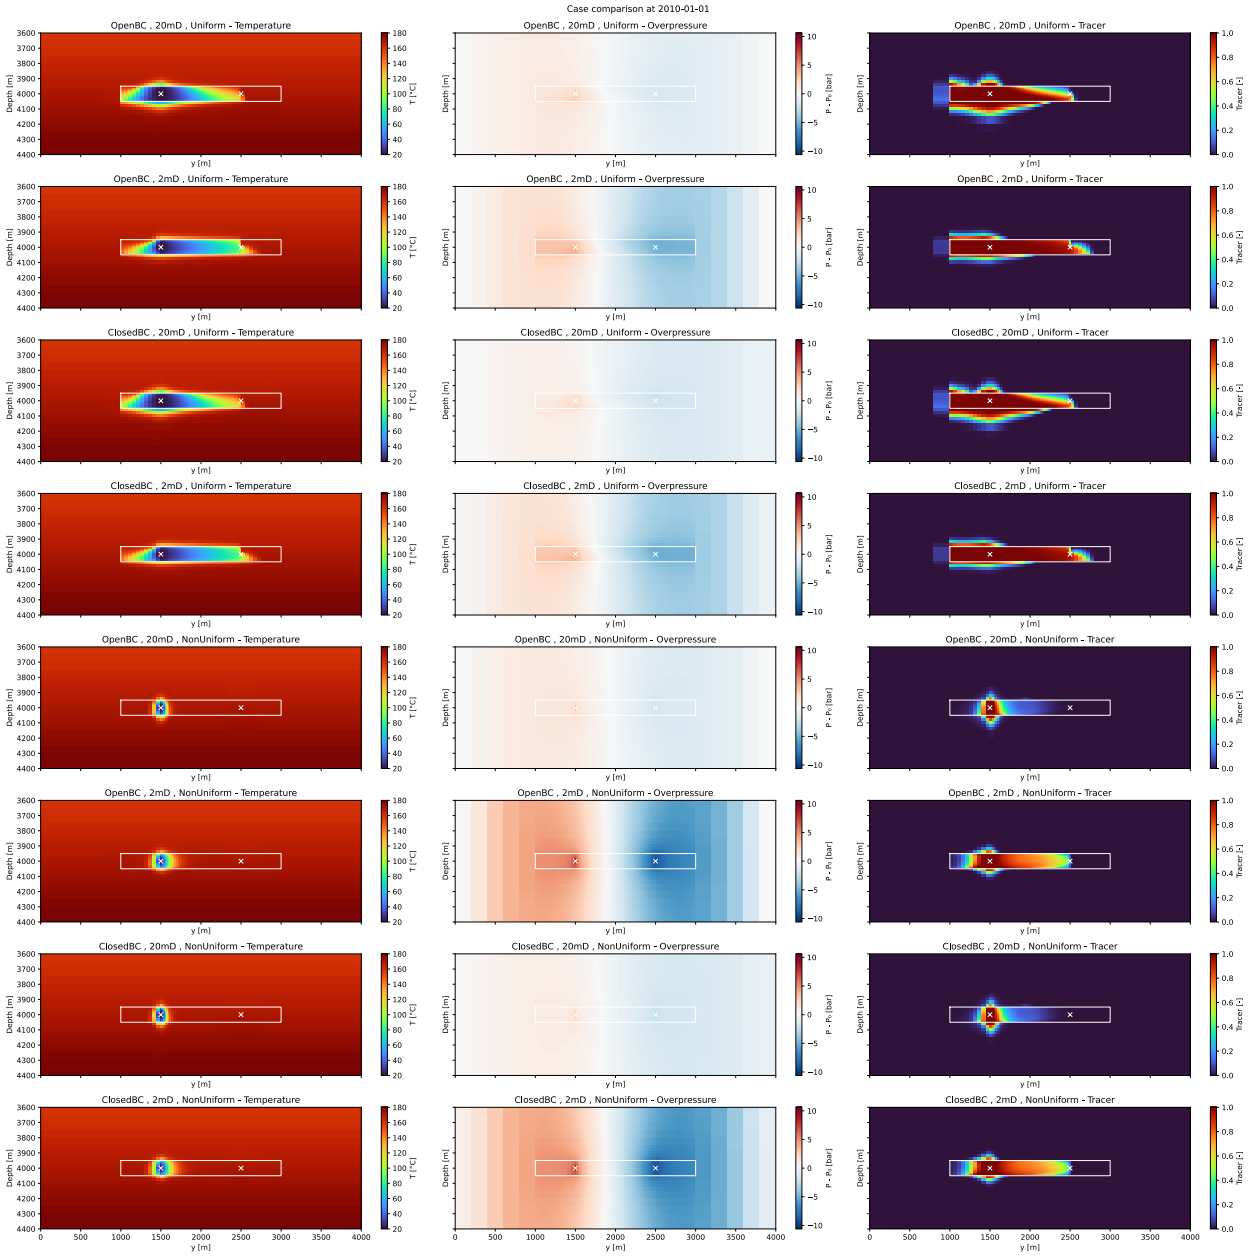


Figure 35: Comparison of the eight configurations ten years after the start of injection: y–z section along fracture 1. Columns: temperature, overpressure $p - p_0$, and tracer concentration; rows: the {open, closed} × {20 mD, 2 mD} configurations of the uniform (top four) and non-uniform (bottom four) fields.

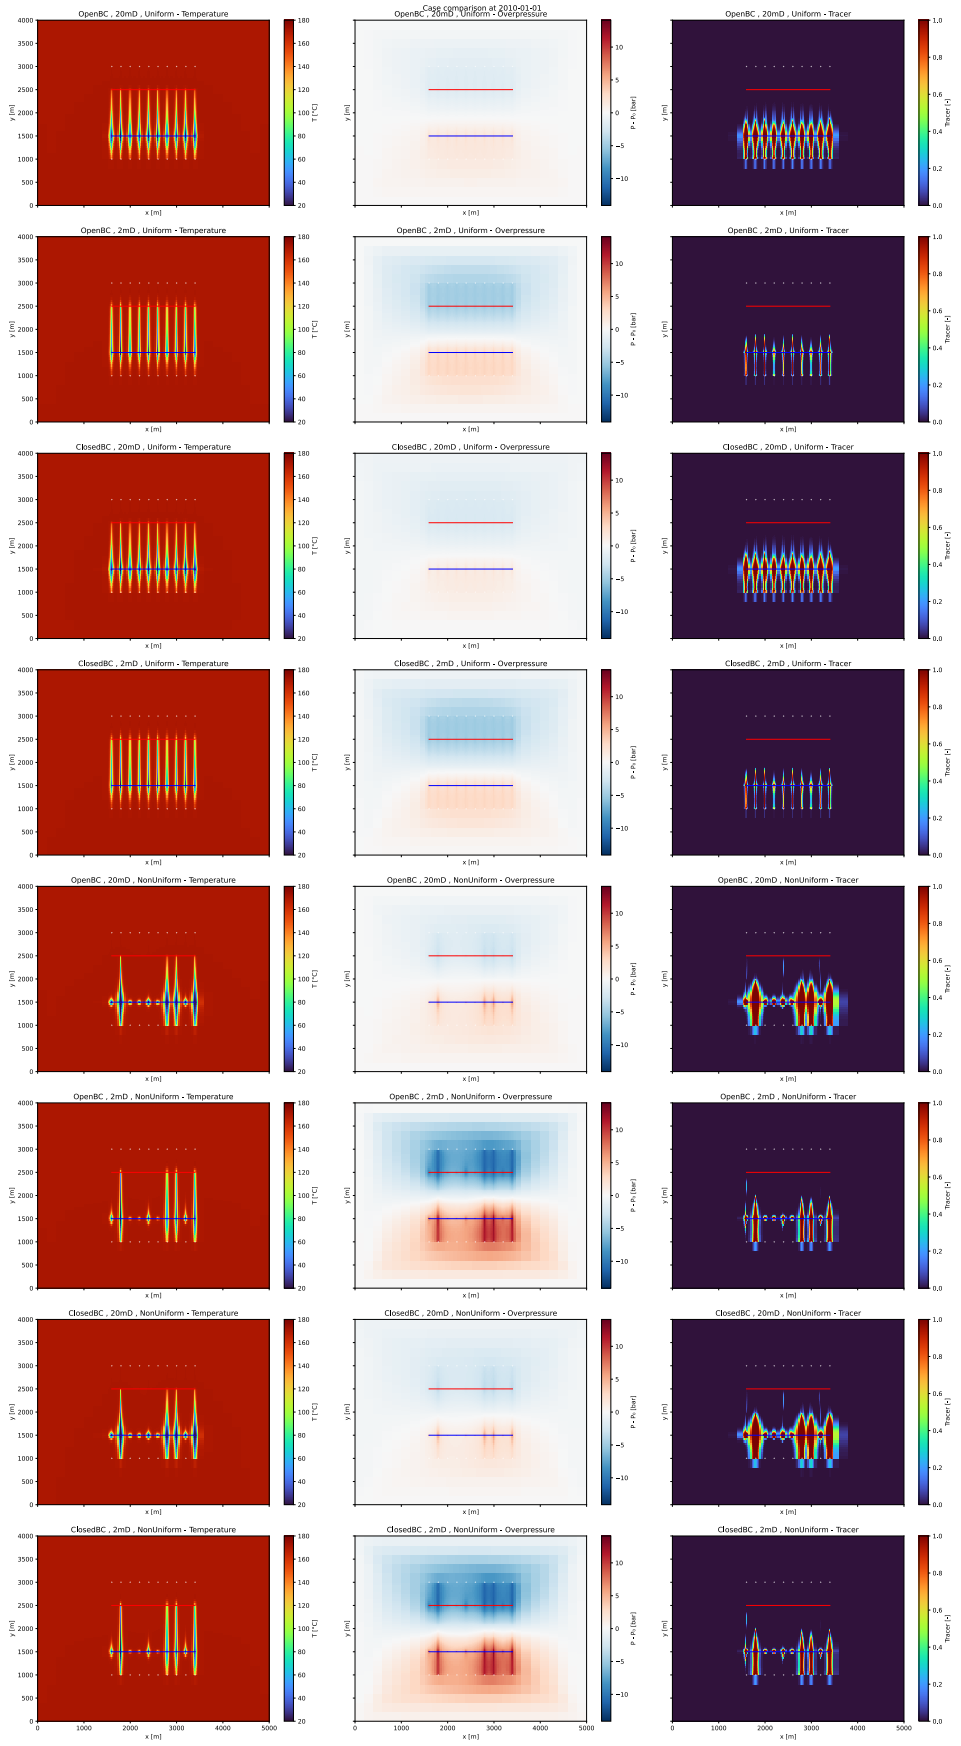
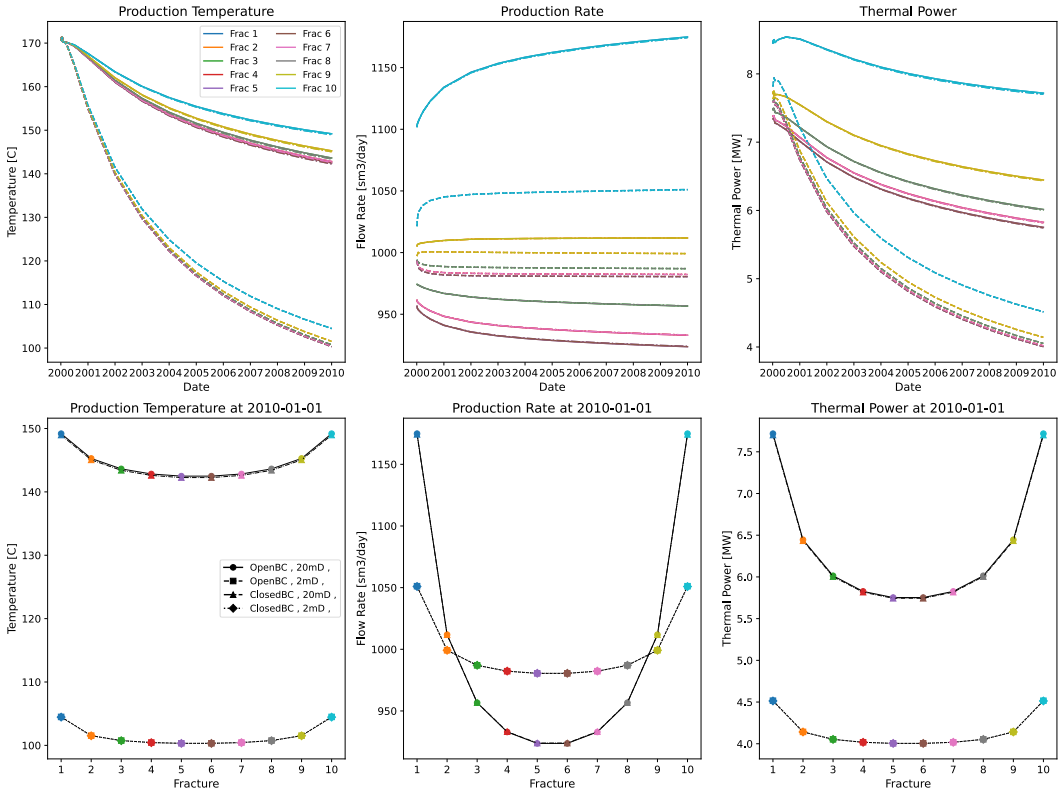
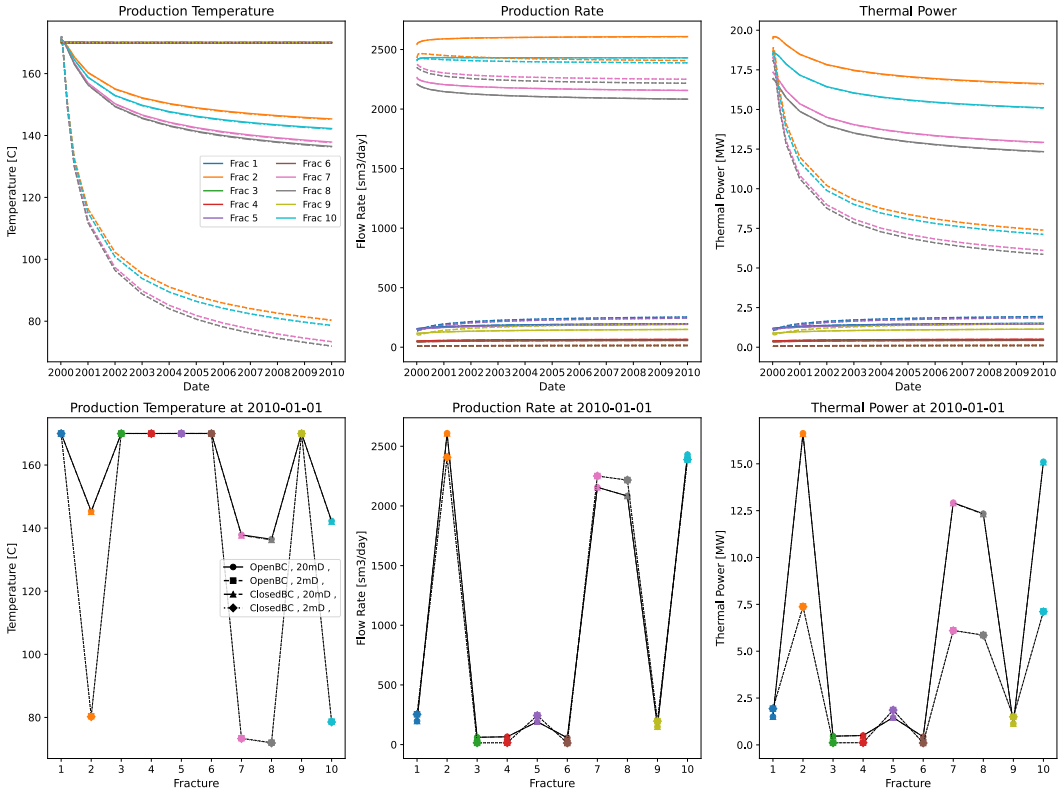


Figure 36: Comparison of the eight configurations ten years after the start of injection: x–y plan view at the well layer. Layout as in Fig. 35.



(a) Uniform



(b) Non-uniform

Figure 37: Per-fracture production temperature, rate, and thermal power for the four {open, closed} × {20 mD, 2 mD} configurations (line style/marker = configuration, colour = fracture). Top rows: time series; bottom rows: per-fracture values after ten years.

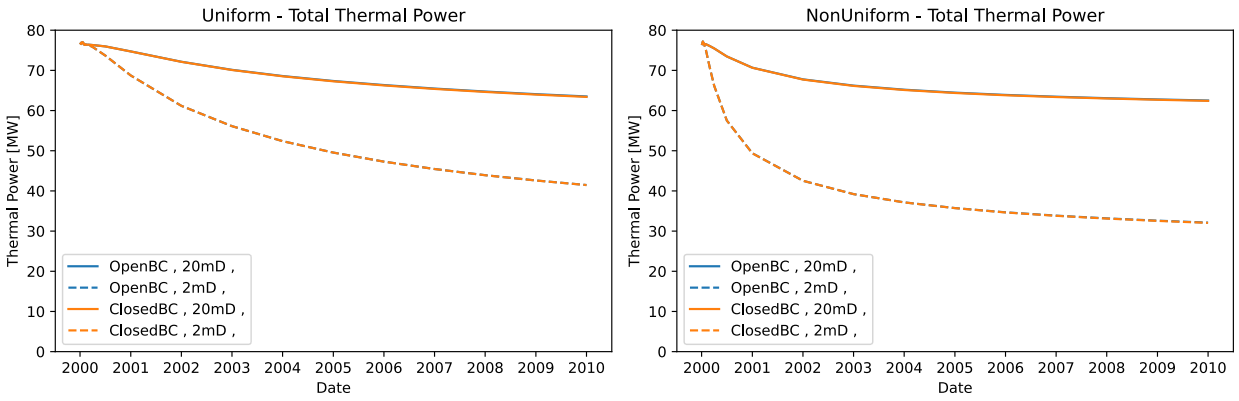
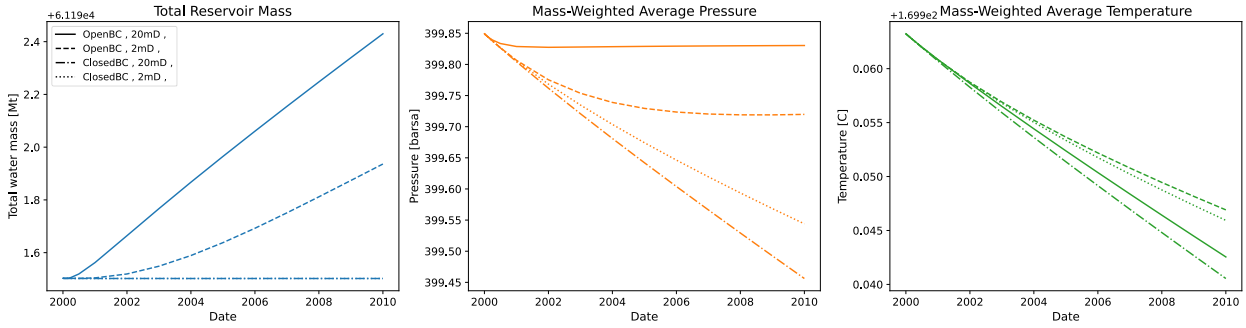
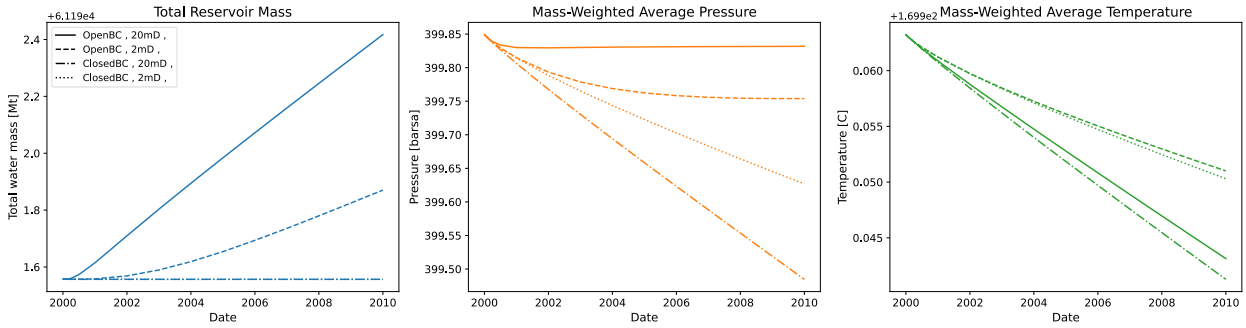


Figure 38: Total thermal power over time for the uniform (left) and non-uniform (right) fields, comparing the four {open, closed} × {20 mD, 2 mD} configurations. The open- and closed-boundary curves are indistinguishable; the matrix permeability controls the decline.



(a) Uniform



(b) Non-uniform

Figure 39: Total reservoir water mass and mass-weighted average pressure and temperature over time for the four {open, closed} × {20 mD, 2 mD} configurations. Open boundaries admit a net inflow at nearly constant pressure; closed boundaries conserve mass and depressurise instead.

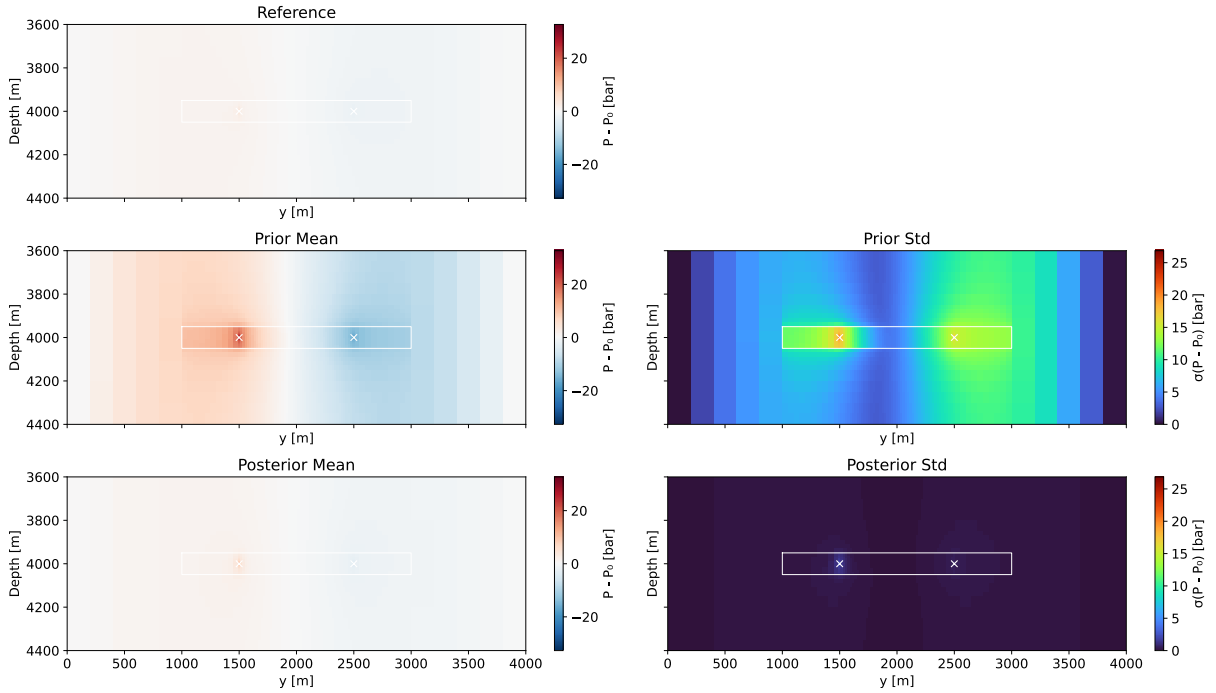
B.4 Prior and Posterior Ensemble Fields

This section provides the detailed forecast summarised in Section 7.3, documenting the effect of the data assimilation on the simulated reservoir fields themselves and thereby complementing the parameter- and observable-space results of Section 7. For the prior ensemble and the posterior ensemble (the $N_e = 100$ realisations after the last of the $N_A = 3$ ES-MDA iterations), the ensemble mean and standard deviation of the overpressure, tracer, and temperature fields are compared with the reference realisation ten years after the start of injection (1 January 2010) — eight years beyond the measurement date t_{meas}, so the posterior fields constitute a genuine forecast. The gridded fields are shown in the usual three section views, with the interior y–z section taken along fracture 7. Each figure shows the reference field above the prior and posterior ensemble means (left column) and the corresponding ensemble standard deviations (right column). The overpressure is evaluated per realisation as $p - p_0$ before averaging, and the sample convention $(N_e - 1)$ is used.

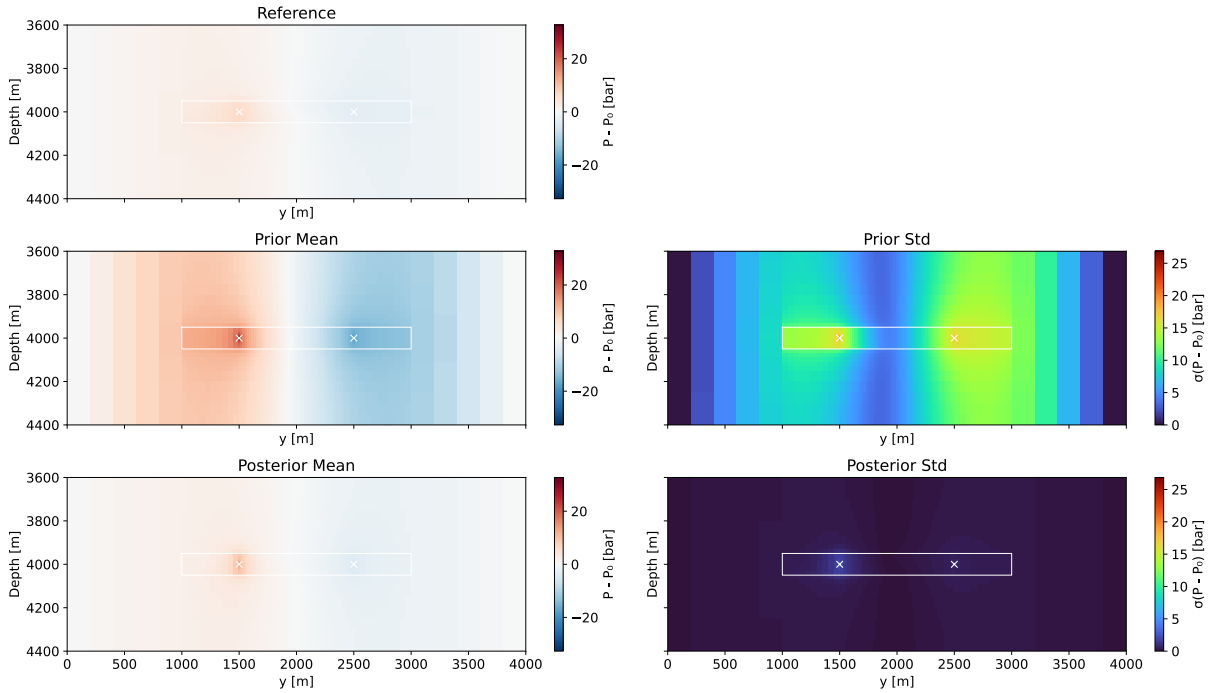
The three quantities are ordered along the causal chain of the assimilation: the strain observations (Eq. (8)) respond directly to the overpressure field, which therefore carries the immediate data constraint (Section B.4.1); the tracer shows how this constraint propagates to the advective transport between the wells (Section B.4.2); and the temperature — the quantity of interest for the geothermal forecast, and the slowest field to respond — inherits the constraint through the heat transport (Section B.4.3). Neither the tracer nor the temperature is observed; any uncertainty reduction in these fields is obtained solely through the strain-constrained fracture properties.

B.4.1 Overpressure

Figs. 40 and 41 show the overpressure statistics. In the reference realisation the overpressure field at ten years is weak: with four open (H) fractures carrying the flow and open lateral boundaries, the injection–production dipole amplitude remains at a few bar. The prior mean instead shows a pronounced dipole — positive around the injector row, negative around the producer row, of the order of 10 bar — because prior realisations with low overall transmissivity — predominantly low fracture permeabilities, low matrix permeability k_m, or both — must sustain far larger driving pressure differences to deliver the imposed well rates, and these members dominate the average. Low-k_m members additionally exhibit a field-wide pressure decline of up to about 30 bar: the thermally contracting reservoir is then poorly connected to the lateral boundary storage. The boundary multiplier *multpv* itself has almost no influence on the overpressure spread. The prior standard deviation exceeds 25 bar next to the wells and decays smoothly outward across the whole model, reflecting the diffusive character of the pressure field. After assimilation, mean and spread collapse onto the reference: the posterior mean is visually indistinguishable from the reference field, and the posterior standard deviation retains only a small residual at the wells. This is the expected behaviour — the strains measure the pressure-induced deformation directly, so the overpressure is the best-constrained field.



(a) Fracture 1



(b) Fracture 7

Figure 40: Ensemble statistics of the overpressure $p - p_0$ ten years after the start of injection (1 January 2010): y–z sections along fractures 1 and 7. Left column: reference realisation, prior ensemble mean, and posterior ensemble mean (shared colour scale); right column: prior and posterior ensemble standard deviation (shared zero-anchored scale).

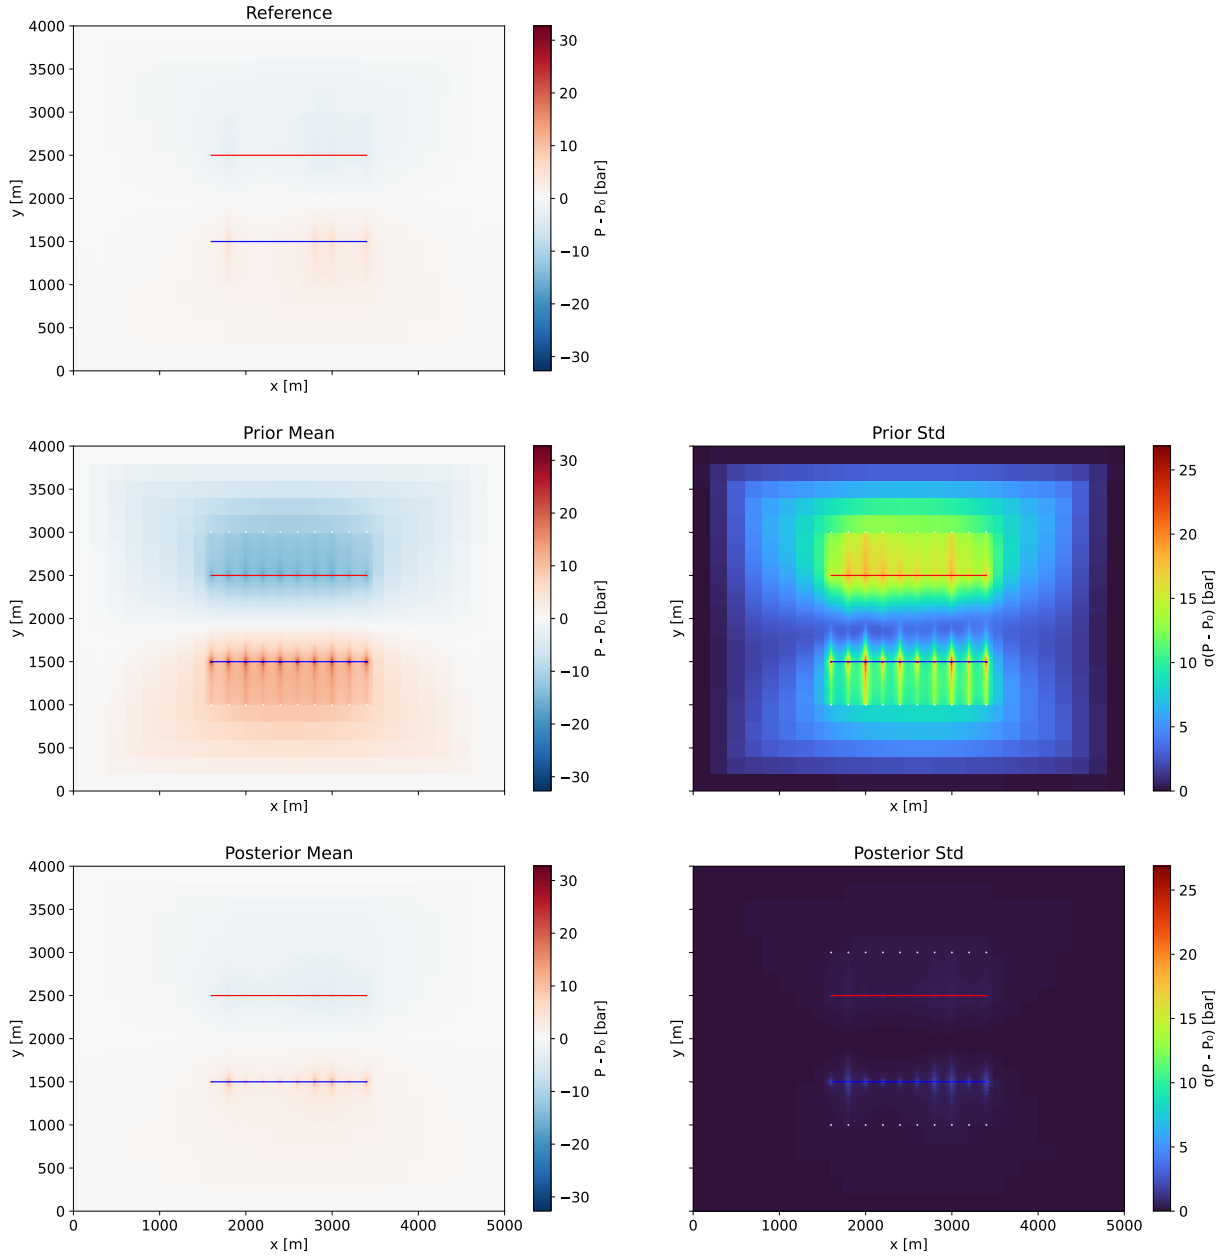


Figure 41: Ensemble statistics of the overpressure $p - p_0$ at ten years: x–y plan view at the well layer (K = 31, 4000 m depth). Layout as in Fig. 40.

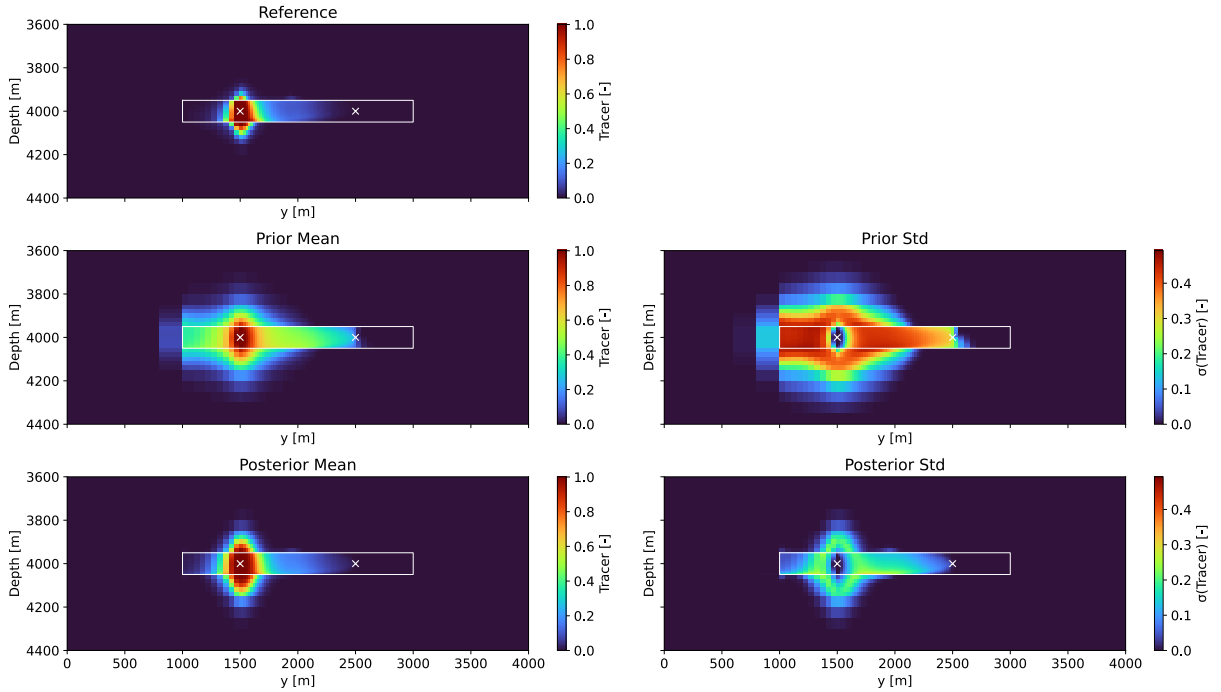
B.4.2 Tracer Concentration

Figs. 42 and 43 show the corresponding tracer statistics. In the reference realisation the plan view exhibits the strongly differentiated per-fracture transport of Section B.2: the tagged water sweeps far along the open (H) fractures while the closed (L) fractures remain essentially untouched. The prior mean erases exactly this contrast: since every fracture draws its permeability independently, averaging over the ensemble assigns the same intermediate concentration plume to all ten fractures. The prior standard deviation reaches $\sigma \approx 0.45$ in a band covering the entire fracture array — close to the maximum spread a quantity bounded in $[0, 1]$ can attain, reflecting the nearly binary breakthrough/no-breakthrough alternative between realisations. The posterior mean recovers the per-fracture breakthrough pattern of the reference, and the residual standard deviation is confined to narrow margins around the plumes: the throughput of each fracture is essentially resolved, and what remains is a moderate uncertainty about the exact front position.

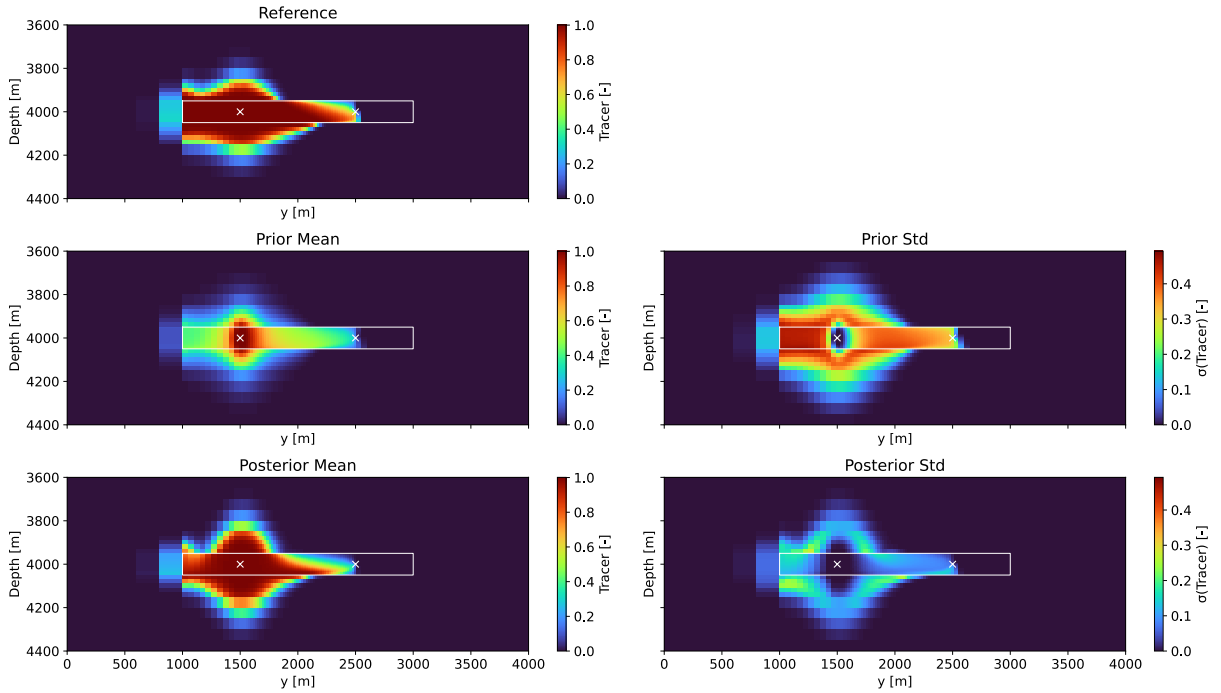
B.4.3 Temperature

The temperature (Fig. 44) shows the same transition, now for the quantity that determines the geothermal forecast. The prior mean smears an identical elongated cold tongue along every fracture, reaching from the injector far towards the producer, and the prior standard deviation spans the full fracture box with values up to about 65 °C — largest between the wells, where the realisations disagree about whether thermal breakthrough has occurred after ten years.

The posterior mean recovers the qualitative plume pattern of the reference: a compact cooled region around the injector on the partially open fracture 1, an elongated tongue towards the producer on the open fracture 7, and plan-view plumes developed only along the open (H) fractures. The quantitative match is more modest: the posterior plumes along the open fractures remain visibly shorter and paler than their reference counterparts, which reach the producer row; faint cooling appears along fractures that are closed in the reference; and the front region is additionally smeared because the mean averages over members with differing front positions.



(a) Fracture 1



(b) Fracture 7

Figure 42: Ensemble statistics of the tracer concentration ten years after the start of injection (1 January 2010): y–z sections along fractures 1 and 7. Layout as in Fig. 40.

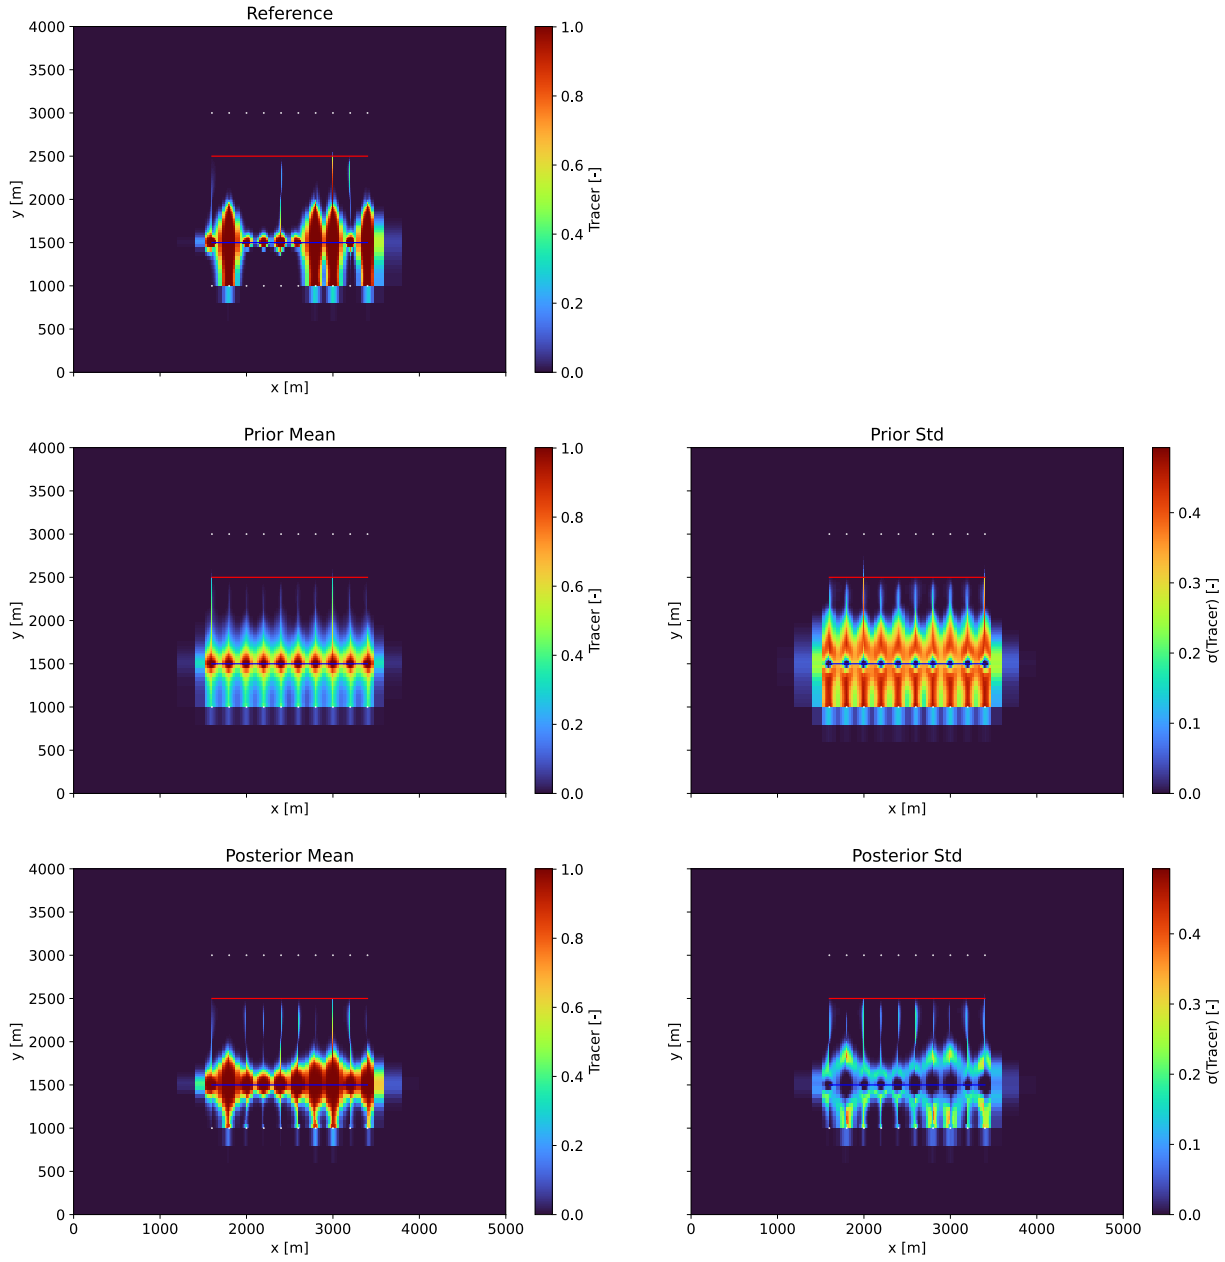
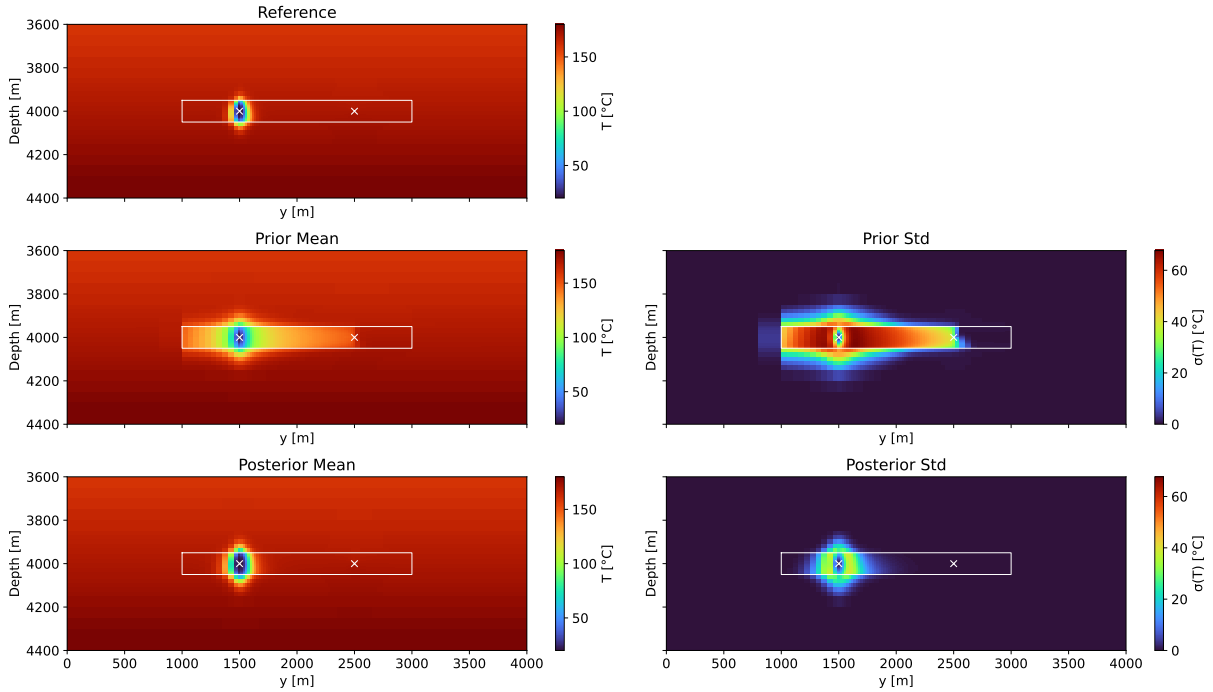
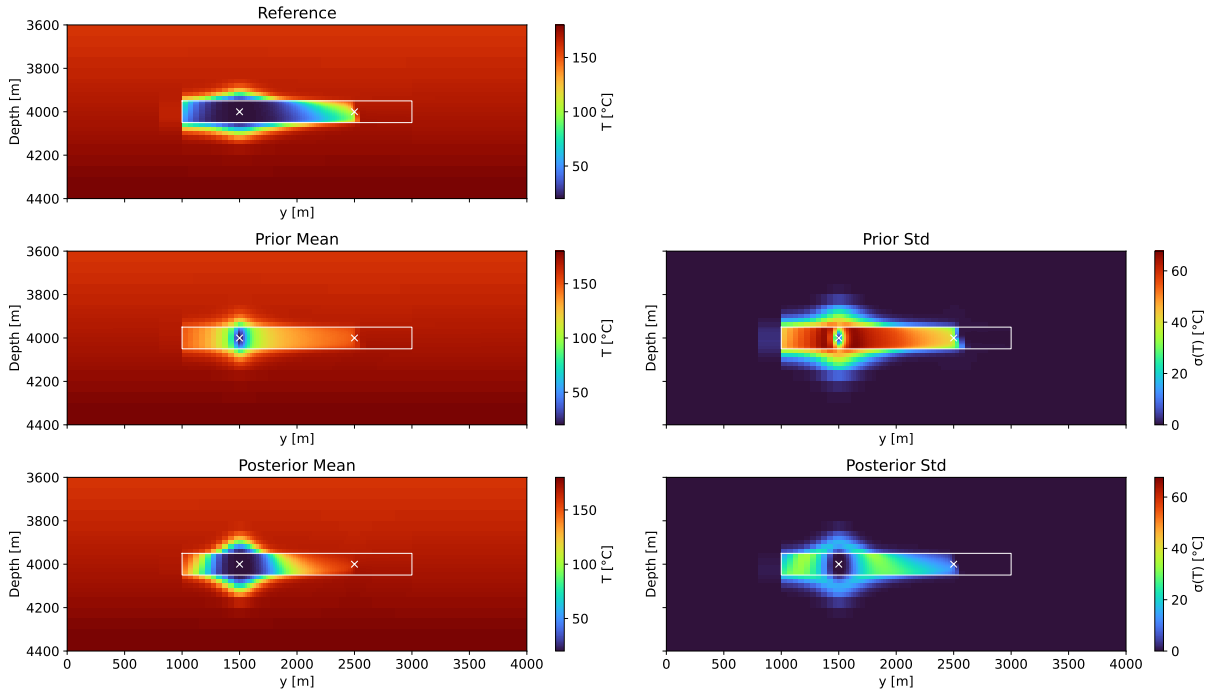


Figure 43: Ensemble statistics of the tracer concentration at ten years: x–y plan view at the well layer. Layout as in Fig. 40.



(a) Fracture 1



(b) Fracture 7

Figure 44: Ensemble statistics of the temperature ten years after the start of injection (1 January 2010): y–z sections along fractures 1 and 7. Layout as in Fig. 40.

C Ensemble Variance Diagnostics

This subsection analyses the linear (Pearson) correlation structure of the prior ensemble at two levels: among the uncertain parameters (a check of the sampling design), and among the measurements (how much of the measurement set is redundant).

C.1 Correlation Analysis among the Parameters

As a sanity check of the Latin-hypercube sampling, Fig. 45 shows the correlation matrix of the 22 uncertain parameters (Section 3.8) across the prior ensemble. The matrix is close to the identity: the largest off-diagonal magnitude is 0.30, between $\log_{10} k_6$ and $\log_{10} k_7$. This is consistent with pure sampling noise: for $N_e = 100$ the correlation estimate of an independent pair has a standard deviation of $1/\sqrt{N_e} = 0.1$, and the extreme over 231 parameter pairs is expected at 2.5–3 times that value. The prior therefore carries no spurious parameter correlations that could masquerade as data-induced structure in the update.

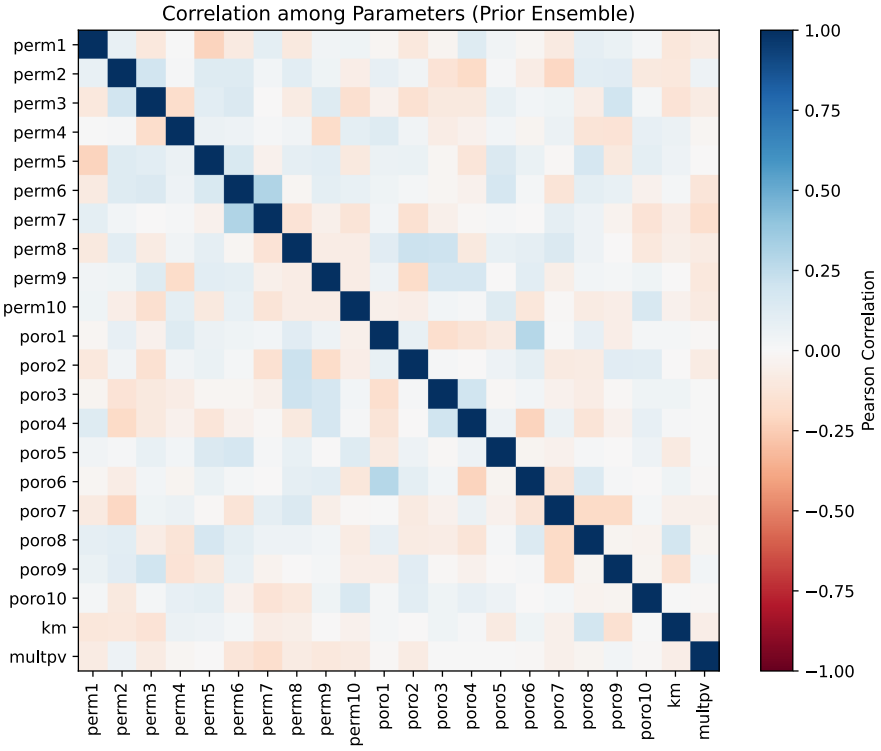


Figure 45: Correlation among the 22 uncertain parameters across the prior ensemble.

C.2 Correlation Analysis among the Measurements & Forecasts

The measurement set comprises twelve scalar diagnostics: the bottom-hole pressures of both wells, the produced tracer concentration, the total thermal power delivered by the doublet (computed from the produced rates and temperatures via Eq. (43) in Section B), the surface displacement at the points of extremal prior-mean uplift and subsidence, and six strain values: the vertical strain ε_{zz} at reservoir depth and the axial strain ε_{xx} at fracture 2 ($x = 1800$ m) and midway between fractures 5 and 6 ($x = 2500$ m), on each well.

Fig. 46 shows the correlation matrix of these twelve measurements across the prior ensemble (mean absolute off-diagonal correlation: 0.45). The pressure-supported quantities — the bottom-hole pressures, the surface displacements, and the strain at reservoir depth and at the fracture crossings — form a single strongly correlated block. The two transport quantities are almost perfectly anti-correlated with each other,

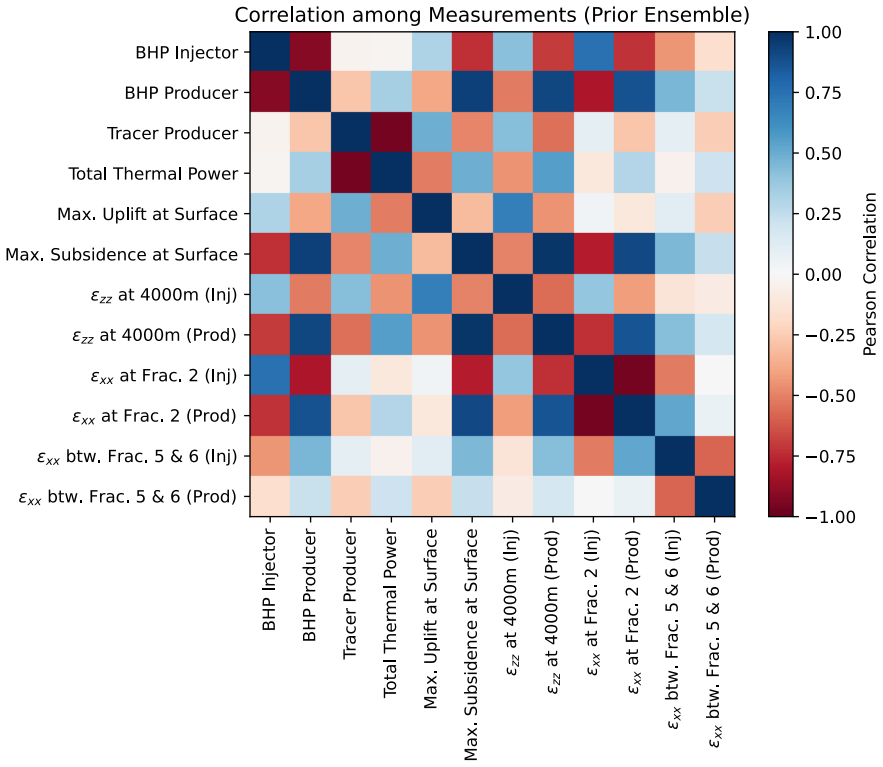


Figure 46: Correlation among the scalar measurements and forecast objectives across the prior ensemble.

but couple more weakly to this block. The three most correlated pairs illustrate the redundancy: the surface subsidence extreme and the vertical strain at reservoir depth on the producer (+0.98), the produced tracer concentration and the thermal power (−0.96), and the axial strain at the fracture-2 crossing on the two wells (−0.96). Within the block the measurements are largely redundant: assimilating several of them adds little over assimilating one. The between-fracture strain values stand apart: they correlate only weakly with the block, and with each other across the two wells, indicating that they carry complementary information.

D ES-MDA Diagnostics

Fig. 47 shows how the ensemble estimate of each uncertain parameter evolves over the three ES-MDA iterations, including changes in both its mean and spread. Fig. 48 compares the resulting posterior distributions with the prior distributions and the known synthetic truth, showing which parameters are constrained and whether their estimates become more accurate.

Because the reference is known in this twin experiment, the error with respect to the *truth* can be monitored directly (Fig. 49). The figure decomposes the per-parameter error into the bias of the ensemble mean (top row), the RMS error over the individual members (middle row), and the ensemble standard deviation (bottom row), the three being linked by $\mathrm{RMS}^2 = \mathrm{bias}^2 + \sigma^2$ for every parameter.

The matrix permeability error collapses almost completely: its member RMS error falls from 0.83 to 0.08 decades, with bias and standard deviation contracting in step (from 0.60 and 0.58 to 0.06 each) onto the reference. The fracture permeabilities behave very unevenly, and the split follows the reference regime. The partially open fractures are recovered best, but they also start closest to the truth case, since their reference of 1600 mD lies near the centre of the log-uniform prior, so their bias is already small before assimilation. The update therefore mainly collapses the spread, taking their member RMS error from about 1.15 to 0.3–0.5 decades while the bias stays small (top row). The open fractures improve strongly too, but incompletely. Their error of their ensemble mean falls by about a decade (from 1.95 to 0.7–0.9), yet almost a decade of bias remains. They are correctly identified as the most conductive but systematically under-predicted, as the update is unable to reach the extreme upper tail of the prior. The closed fractures essentially fail to improve, as their RMS error barely changes (from 2.0 to 1.8–2.4 decades) with the bias even worsenening in some cases, placing them about two decades above their true 20 mD permeability. The open and partially open fractures are thus drawn towards the reference while the closed ones are not.

The fracture porosities barely move, as their member RMS error and spread decrease only marginally, and the error of their mean not at all. The pore-volume multiplier is the most revealing case. Its member RMS error is *non-monotone*, such that is improves in the first iteration (from 23 to 20) but then degrading to end *above* its prior value, at 28. Its standard deviation nonetheless contracts throughout (from 23 to 14) while its bias grows monotonically (from 5 to 24): the ensemble tightens around an increasingly wrong value, the signature of an overconfident misidentification rather than a genuine gain.

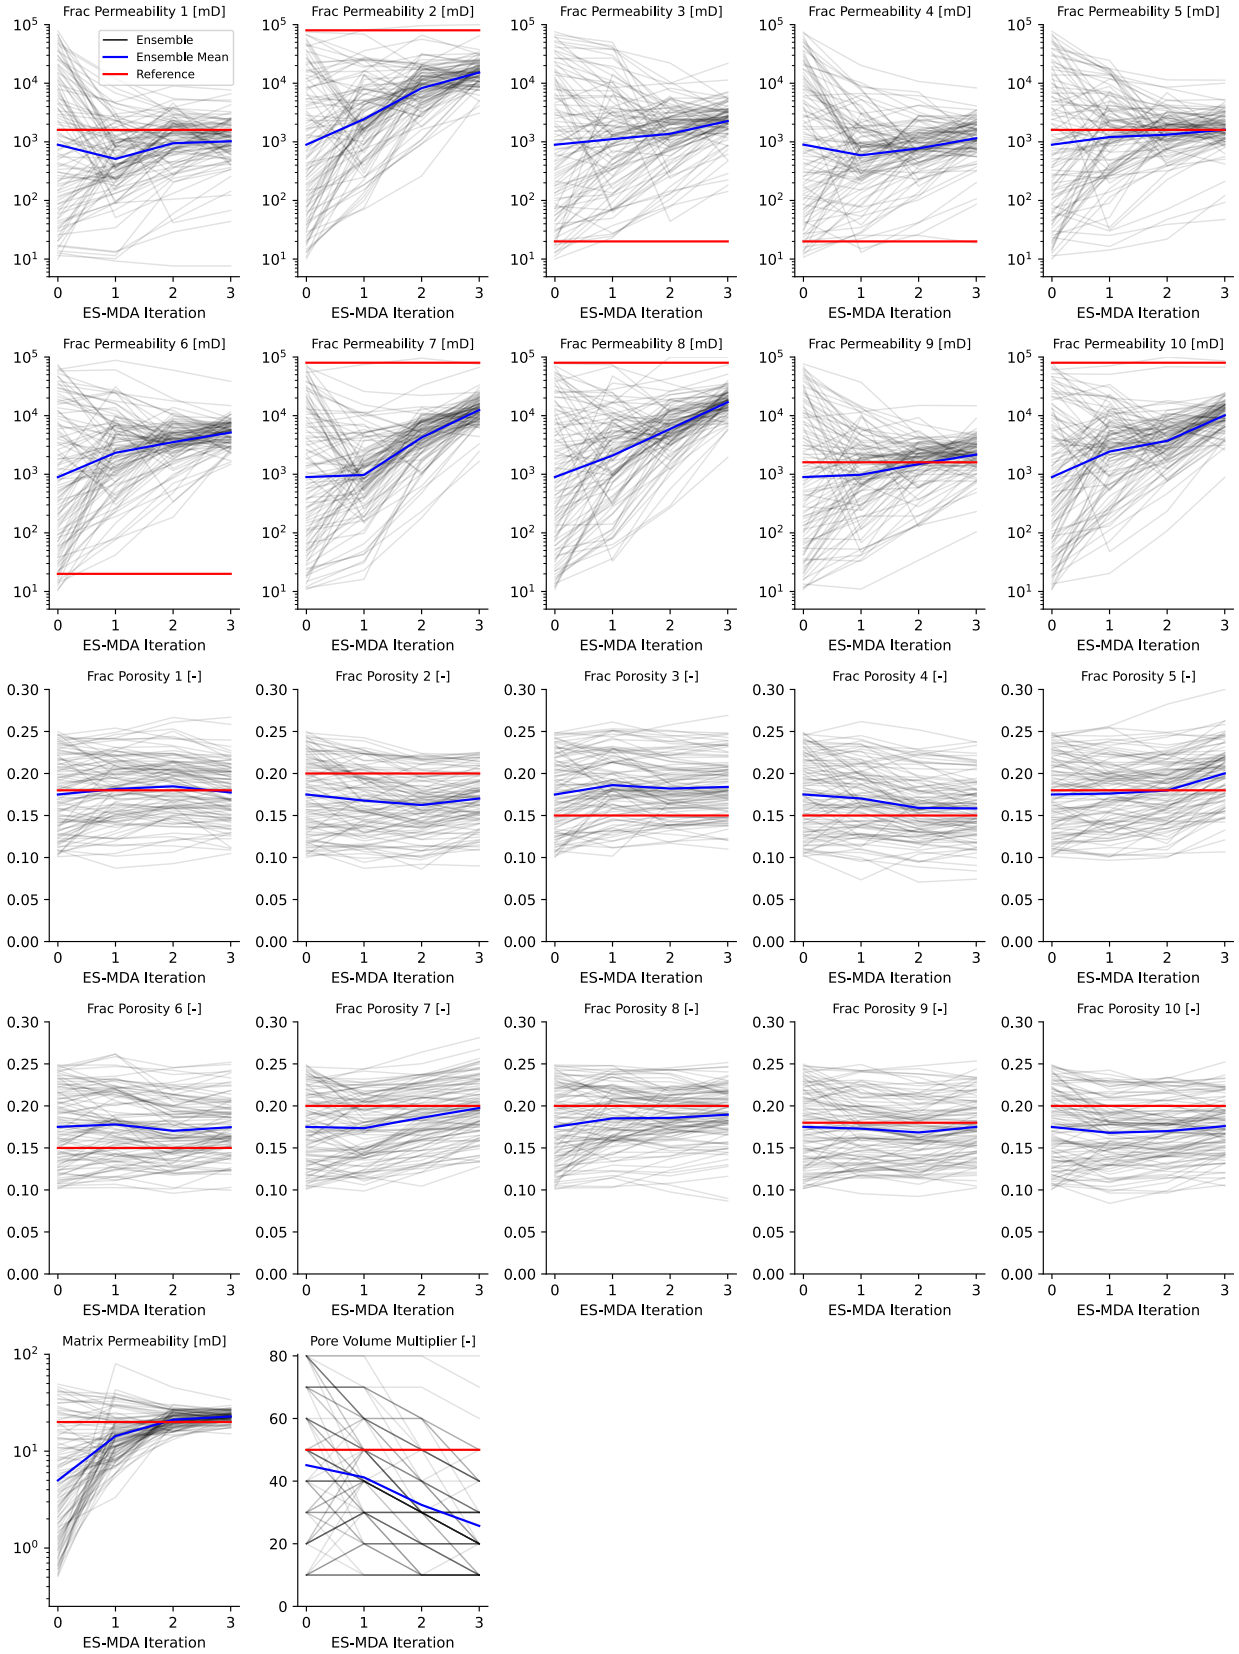


Figure 47: ES-MDA state evolution. One panel per parameter, showing every ensemble member (thin black), the ensemble mean (blue), and the reference value (red) over the iterations (0 = prior). Permeabilities on logarithmic axes.

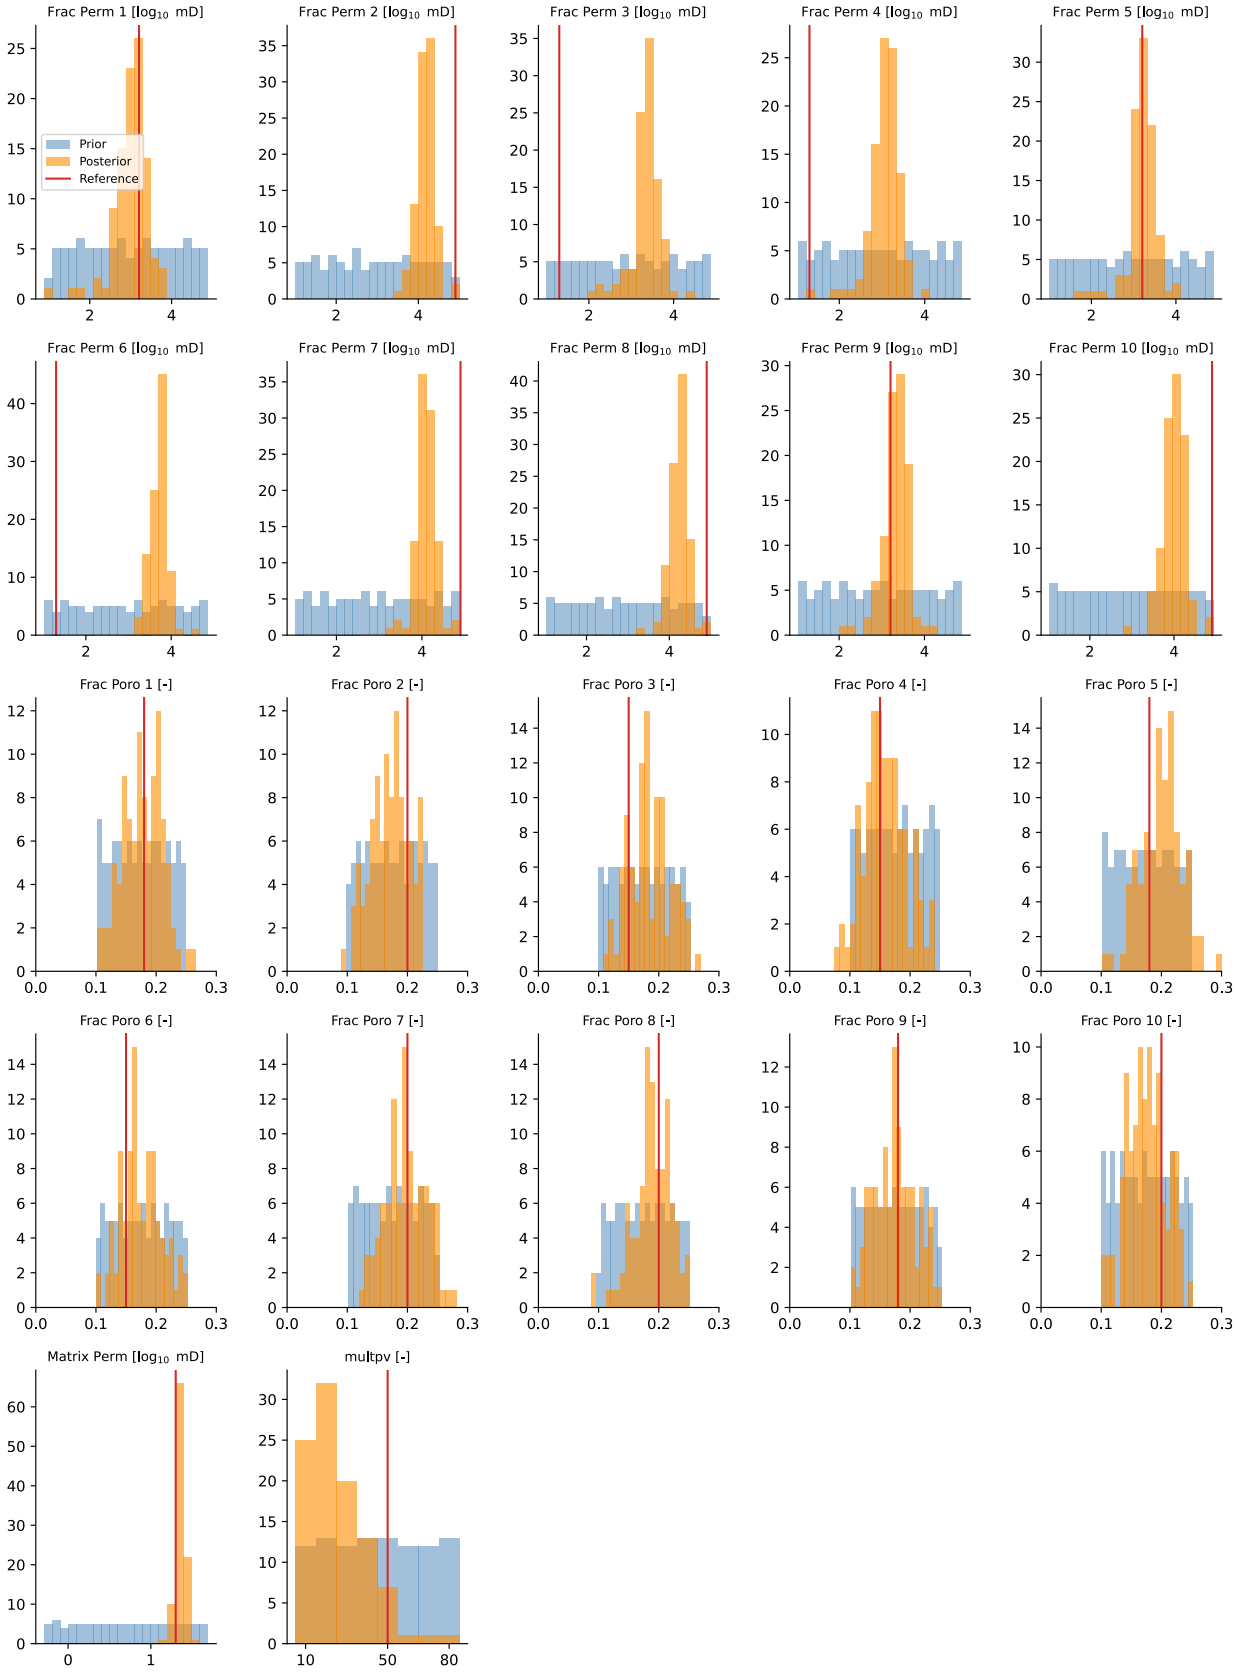


Figure 48: Prior (blue) vs. posterior (orange) parameter histograms in state space ($\log_{10}$ for the permeabilities), with the reference value (red line).

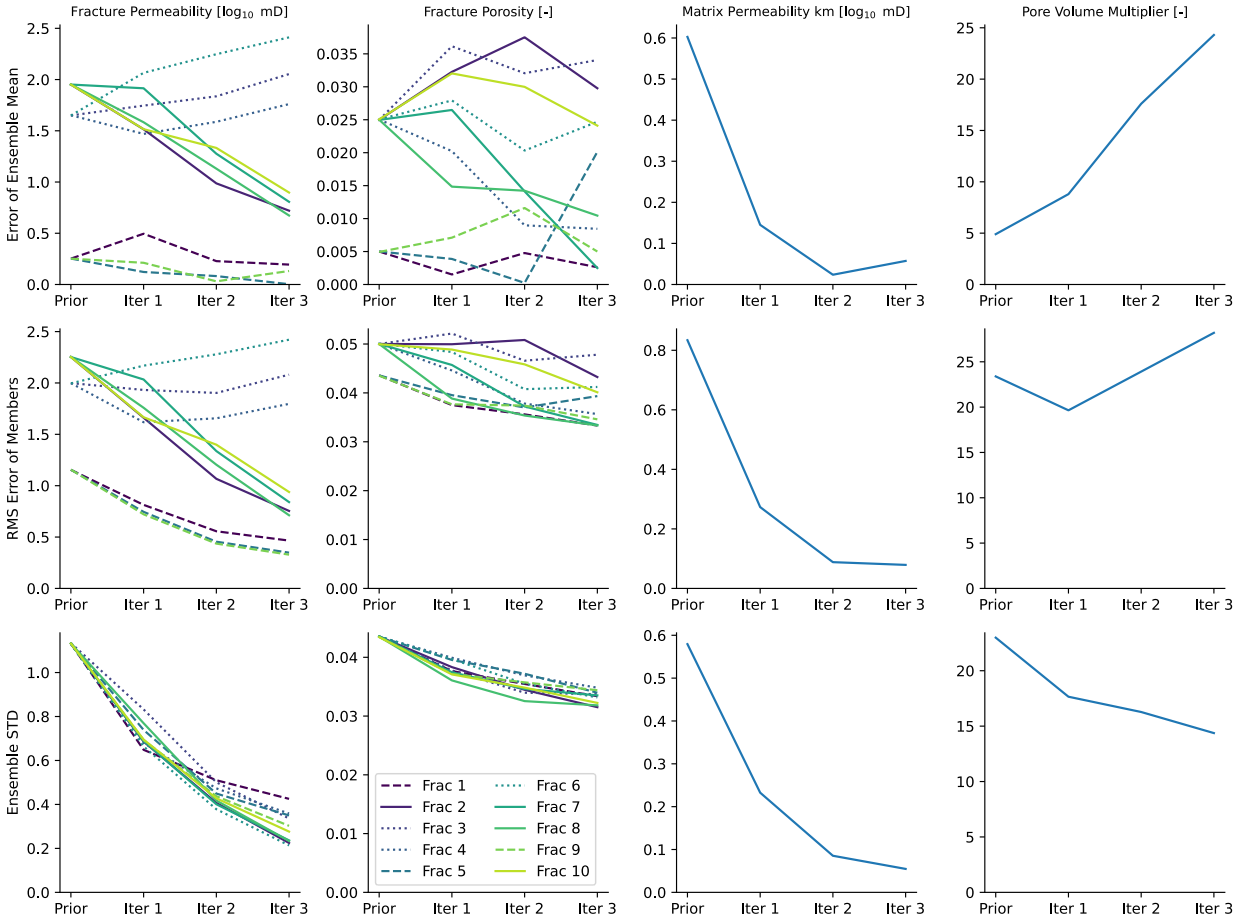


Figure 49: Per-parameter error with respect to the reference over the ES-MDA iterations, in state space. Top row: absolute error of the ensemble mean; middle row: RMS error over the ensemble members; bottom row: ensemble standard deviation. For the fracture permeabilities and porosities the line style marks the reference regime of each fracture: solid for open (H), dashed for partially open (M), and dotted for closed (L).