2020:01136- Unrestricted

Report

User Guide to Flow Diagnostics Postprocessing

Simulations in MRST and ECLIPSE Output Format

Author(s)
Knut-Andreas Lie
Stein Krogstad
Francesca Watson

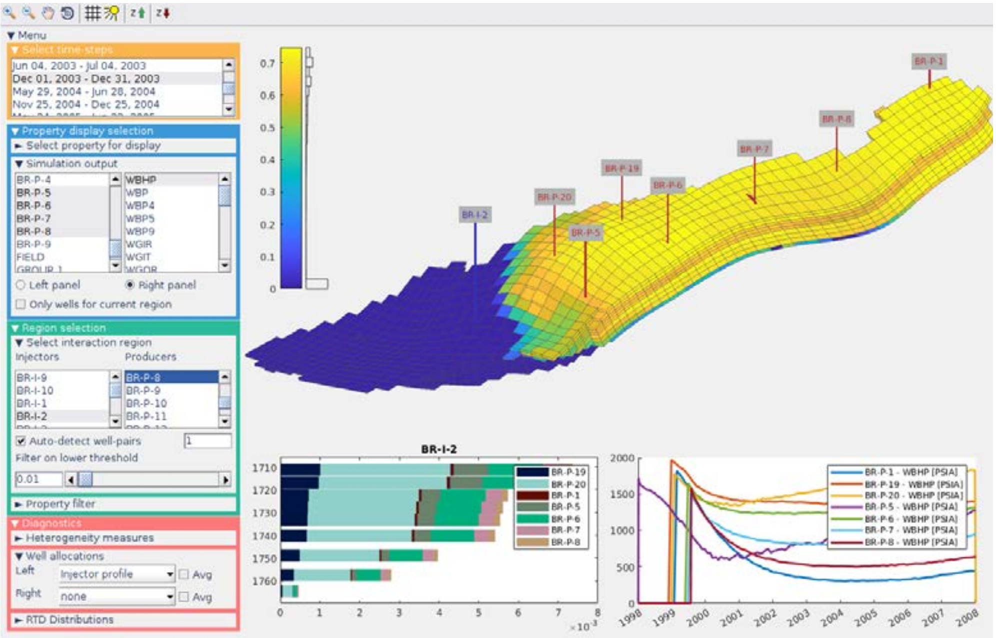

▼ Menu

▼ Select time-steps
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004

▼ Property display selection
► Select property for display
▼ Simulation output
BR-P-4
BR-P-5
BR-P-6
BR-P-7
BR-P-8
BR-P-9
FIELD
WBHP
WBP
WBP4
WBP5
WBP9
WGIR
WGIT
○ Left panel ● Right panel
□ Only wells for current region

▼ Region selection
▼ Select interaction region
Injectors: BR-I-9, BR-I-10, BR-I-1, BR-I-2
Producers: BR-P-8, BR-P-9, BR-P-10, BR-P-11
☑ Auto-detect well-pairs 1
Filter on lower threshold 0.01
► Property filter

▼ Diagnostics
► Heterogeneity measures
▼ Well allocations
Left Injector profile □ Avg
Right none □ Avg
► RTD Distributions

SINTEF Digital
Computational Geosciences
2020-11-03

SINTEF

SINTEF Digital
SINTEF Digital

Address:
Postboks 124 Blindern
NO-0314 Oslo
NORWAY
Switchboard: +47 40005100

info@sintef.no

Enterprise /VAT No:
NO 919 303 808 MVA

Report

User Guide to Flow Diagnostics Postprocessing

Simulations in MRST and ECLIPSE Output Format

KEYWORDS:
Flow Diagnostics;
MRST; ECLIPSE;
Reservoir Simulation;
Graphical User
Interface; MATLAB

VERSION 1.0	**DATE** 2020-11-03

AUTHOR(S)
Knut-Andreas Lie
Stein Krogstad
Francesca Watson

PROJECT NO. 102018629	**NUMBER OF PAGES/APPENDICES:** 3 + Appendices

ABSTRACT

Flow diagnostics are simple and controlled numerical flow experiments run to probe a reservoir model, establish connections and basic volume estimates, and measure heterogeneity in flow paths. All these quantities are quick to compute and can thus be used interactively to explore fluid communication in a geological model before or after more comprehensive multiphase flow simulations.

This user guide briefly outlines the various types of flow diagnostics and the corresponding graphical user interface (GUI) you can use to postprocess black-oil type multiphase flow simulations. The GUI can be used for multiphase simulations stored in either MRST format or ECLIPSE output format.

PREPARED BY Knut-Andreas Lie	K-A Lie Knut-Andreas Lie (Nov 9, 2020 13:18 GMT+1)	SIGNATURE
CHECKED BY Håvard Heitlo Holm	Håvard Heitlo Holm	SIGNATURE
APPROVED BY Trond Runar Hagen	Trond Runar Hagen	SIGNATURE

REPORT NO.	**ISBN**	**CLASSIFICATION**	**CLASSIFICATION THIS PAGE**
2020:01136	978-82-14-06438-4	Unrestricted	Unrestricted

Document history

VERSION	DATE	VERSION DESCRIPTION
1.0	2020-11-03	First Version

SINTEF

Technology for a better society

www.sintef.no

SINTEF Digital

User Guide to

FLOW DIAGNOSTICS POSTPROCESSING

Simulations in MRST and ECLIPSE Output Format

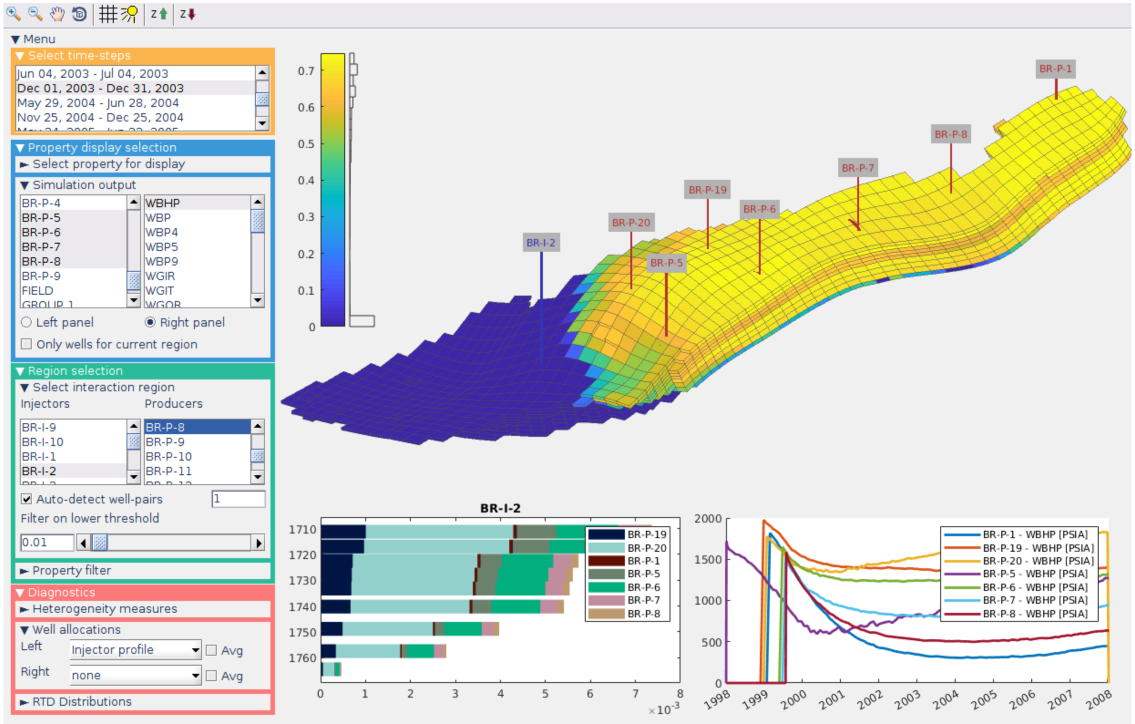

Version 1.0
October 2020

Contents

CHAPTER

1 INTRODUCTION

Flow diagnostics are simple and controlled numerical flow experiments run to probe a reservoir model, establish connections and basic volume estimates, and measure heterogeneity in flow paths. The result is a set of visually intuitive quantities that

- give the travel time for mass-less particles that passively follow the flow field from an injector into the reservoir and from a point in the reservoir to the nearest producer;
- delineate regions drained by given producers or swept (flooded) by given injectors;
- determine whether pairs of injectors and producers communicate or not and measure the relative strength of their connection;
- determine how flux is allocated between different injectors and producers;
- establish the volumetric region influenced by specific well-pairs;
- give the residence-time distribution of all flow paths connecting pairs of injectors and producers;
- measure the dynamic heterogeneity within drainage, sweep, or well-pair regions.

All these quantities are quick to compute and can thus be used interactively to explore fluid communication in a geological model before or after more comprehensive multiphase flow simulations. This

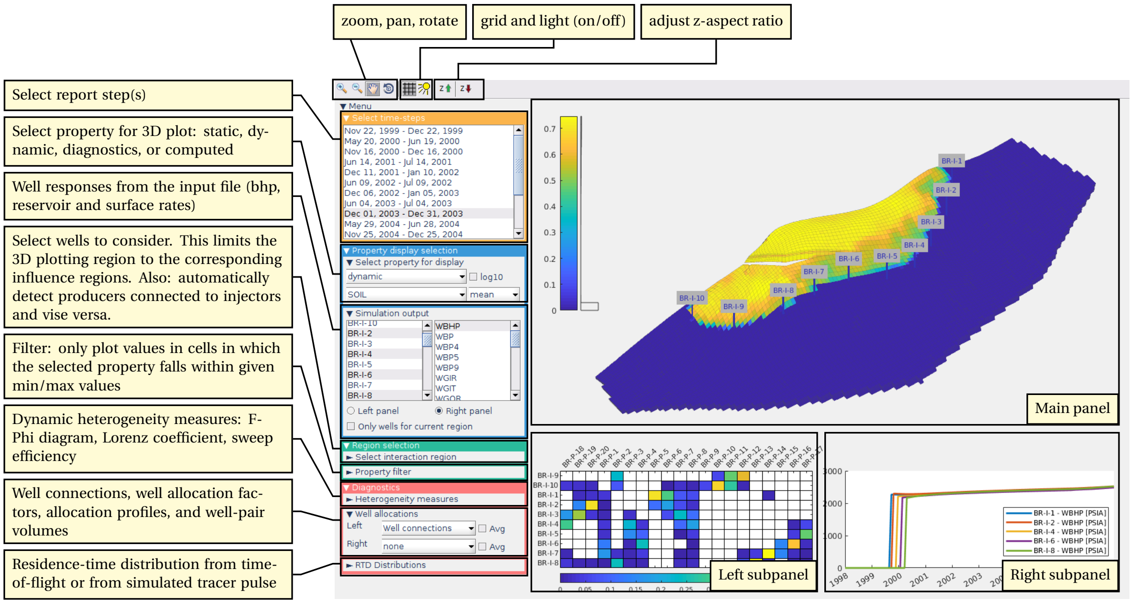
zoom, pan, rotate

grid and light (on/off)

adjust z-aspect ratio

Select report step(s)

Select property for 3D plot: static, dynamic, diagnostics, or computed

Well responses from the input file (bhp, reservoir and surface rates)

Select wells to consider. This limits the 3D plotting region to the corresponding influence regions. Also: automatically detect producers connected to injectors and vise versa.

Filter: only plot values in cells in which the selected property falls within given min/max values

Dynamic heterogeneity measures: F-Phi diagram, Lorenz coefficient, sweep efficiency

Well connections, well allocation factors, allocation profiles, and well-pair volumes

Residence-time distribution from time-of-flight or from simulated tracer pulse

Main panel

Left subpanel

Right subpanel

▼ Menu

▼ Select time-steps

Nov 22, 1999 - Dec 22, 1999
May 20, 2000 - Jun 19, 2000
Nov 16, 2000 - Dec 16, 2000
Jun 14, 2001 - Jul 14, 2001
Dec 11, 2001 - Jan 10, 2002
Jun 09, 2002 - Jul 09, 2002
Dec 06, 2002 - Jan 05, 2003
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004

▼ Property display selection

▼ Select property for display

dynamic ☐ log10

SOIL mean

▼ Simulation output

BR-I-10, BR-I-2, BR-I-3, BR-I-4, BR-I-5, BR-I-6, BR-I-7, BR-I-8

WBHP, WBP, WBP4, WBP5, WBP9, WGIR, WGIT, WGOR

○ Left panel ● Right panel

☐ Only wells for current region

▼ Region selection

► Select interaction region

► Property filter

▼ Diagnostics

► Heterogeneity measures

▼ Well allocations

Left Well connections ☐ Avg

Right none ☐ Avg

► RTD Distributions

FIGURE 1.1: The graphical user interface consists of three plotting panels: the main panel at the top shows a 3D view of the model, whereas the two subpanels along the bottom show various types of flow diagnostics. The menus to the left determine what is shown in these three panels. All loaded report steps are listed at start-up, but before you have selected a specific report step, the GUI can only visualize static properties from the geocellular model in the main panel and production data over the whole simulation period in each of the two subpanels.

user guide briefly outlines the various types of flow diagnostics and the corresponding graphical user interface (GUI) (see Figure 1.1) you can use to postprocess black-oil type multiphase flow simulations.

The GUI described herein can be used to postprocess multiphase flow simulations. The current version can postprocess (black-oil type) multiphase simulations stored in either MRST format or ECLIPSE output format, which is also used by other simulators like the open-source Flow simulator from the Open Porous Media (OPM) initiative [12]. For ECLIPSE output, the output file should contain phase fluxes at reservoir conditions, e.g., invoked by the keyword `FLORES` under `RPTSCHED` in the `SCHEDULE` section; see [13] for more details.

MRST also offers a similar GUI [17] to *preprocess* single simulation models or ensembles of model realizations, so that you can get a good picture of travel times, volumetric connections, and dynamic heterogeneity *before* conducting more comprehensive and expensive multiphase flow simulations.

In the following chapters, we first review basic flow diagnostics and then explain the content of the eight different sub-menus you can use to select report steps and plot various kinds of static, dynamic, flow diagnostic, and tracer data. To this end, we use an ECLIPSE simulation of the Brugge benchmark case [10] given in the `EGRID` format as an example. Overall, this file contains results from 129 report steps, spanning a simulation period starting from 1st January 1999 and ending on 10th January 2008.

Appendix A explains how to install the software and its prerequisites, whereas Appendix B gives a few more details about the implementation of the graphical user interface and how you can control and change its behavior without having to modify the source code and restart the program.

CHAPTER

2 WHAT IS FLOW DIAGNOSTICS?

To make the user guide more self-contained, we start with a brief summary of the basic flow-diagnostic quantities. You can find a more comprehensive introduction in the recent textbook by Lie [7, Chapter 13] or in one of the two original articles that introduce these methods [8, 14].

Many of the concepts that constitute what we herein refer to as "flow diagnostics" were originally developed within streamline simulation [2, 15, 19, 20]. Although MRST has functionality for computing streamlines for certain grid types, we have chosen to implement flow diagnostics using cell-wise quantities that are computed with standard finite-volume discretizations. This requires some small conceptual adjustments as to what the derived quantities mathematically represent. In particular, quantities that have a pointwise or line-specific interpretation in streamline simulation become volumetric quantities that represent the average over all streamlines passing through a cell or cell interface when calculated with a finite-volume method. More details are given in the following.

2.1 Time-of-flight

Flow diagnostics relies on two basic quantities: time-of-flight and influence regions. To define these, we assume a reservoir with porosity ϕ and a superficial Darcy velocity $\vec{v}$ defined as volumetric discharge per area. A streamline is defined as a curve that is tangential to the velocity field at every point; see Figure 2.1. In a steady (incompressible) flow field, there will be a unique streamline passing through each point $\vec{x}$ of the reservoir, starting at the nearest injector (or fluid source) and terminating at the nearest producer (or fluid sink). We can parameterize the path traced out by streamline Ψ by its arc length r, i.e., write $\vec{x}_\Psi(r)$. In streamline simulation [2], however, it is more common to use *time-of-flight* (TOF), which measures the travel time of a massless particle by the interstitial fluid velocity $\vec{v}/\phi$, as spatial coordinate along each streamline. We can define the TOF coordinate τ through two mathematically equivalent equations:

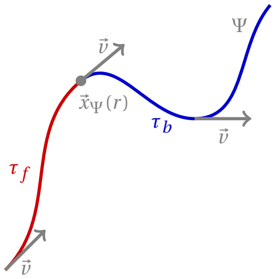

FIGURE 2.1: A streamline Ψ passing through a point $\vec{x}$ is tangential to the velocity field $\vec{v}$ at every point. The streamline can be parameterized by its arc length r, or the travel time defined relative to the interstitial fluid velocity $\vec{v}/\phi$.

$$\tau(r) = \int_0^r \frac{\phi(\vec{x}_\Psi(s))}{|\vec{v}(\vec{x}_\Psi(s))|}\,ds, \qquad \vec{v}\cdot\nabla\tau = \phi. \tag{2.1}$$

In flow diagnostics, we usually refer to two different TOF values: the *forward TOF* τ_f defined as the travel time from the nearest injector to a given point in the reservoir, and the *backward TOF* τ_b defined as the travel time from the given point to the nearest producer. The total travel time along a streamline from inflow to outflow is called *residence time* and equals the sum of forward and backward time-of-flight. Numerically, you can compute *point values* of τ by tracing streamlines and evaluating the integral in (2.1). Herein, we will follow [8, 14] and compute *volume averaged values* of τ from a standard finite-volume discretization of $\vec{v}\cdot\nabla\tau = \phi$, as first proposed by [3, 9]. The relationship among streamlines, time-of-flight, and residence time is illustrated in Figure 2.2 for a rectangular reservoir with two injectors and three producers.

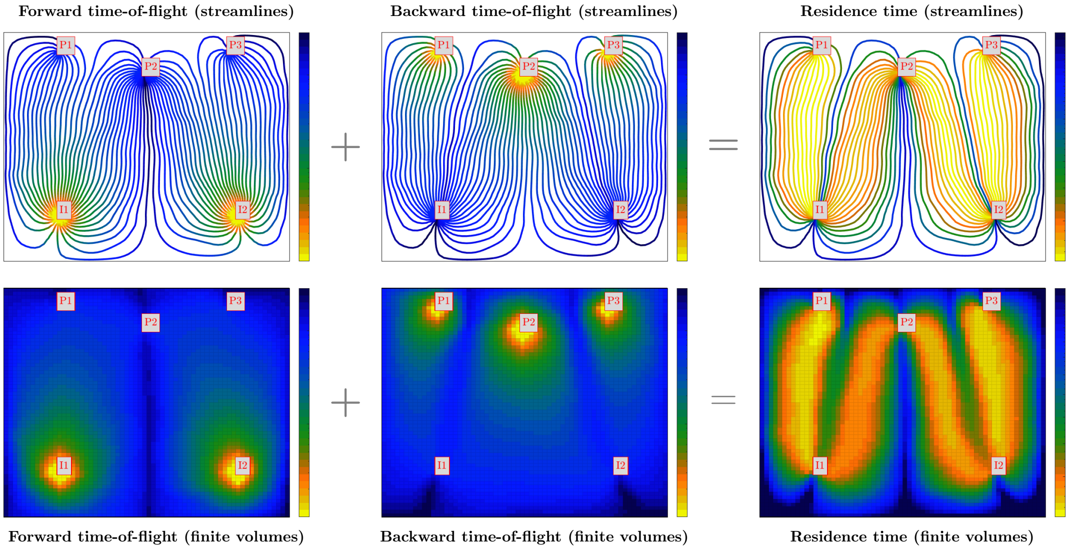

FIGURE 2.2: Relationship among streamlines, time-of-flight, and residence time for a rectangular reservoir with two injectors and three producers. Computing with a streamline method gives pointwise values, or to be more precise, increment values associated with line segments that cross individual cells, whereas a finite-volume method computes average values over all flow paths crossing each cell.

2.2 Influence regions

Influence regions are defined as the locus of the points that lie on all streamlines emanating from a given point, line segment, or volumetric object that represents a fluid source or sink. Since each point that is not a fluid source or sink in an flow field can only lie on a single streamline, these influence regions are distinct volumetric objects that introduce a natural volumetric partition of a fixed flow field. Assuming incompressible flow, we can define an influence region mathematically as

$$\vec{v}\cdot\nabla C = 0, \qquad C|_{\text{inflow}} = 1. \tag{2.2}$$

If this equation is solved with a monotone finite-volume method, the computed values will lie in the interval $[0,1]$ and represent the average over a finite volume. You may also see this quantity referred to as a *steady tracer*[1] *concentration*, which gives the portion of the total fluid volume passing through a point in the reservoir that can be attributed to a given fluid source or parts of the inflow boundary. Likewise, by reversing the sign of the flow field, we can compute influence regions associated with fluid sinks. In both cases, these C values form a partition of unity; see Figure 2.3. With a standard first-order discretization, the influence regions will contain significant numerical diffusion and will thus not have the sub-cell resolution and be as sharply defined as is possible if the regions are delineated by tracing a large number of streamlines. On the other hand, the advantage of using a finite-volume discretization is that this gives the regions that will be influenced by passively advected quantities in a standard multiphase simulation. Likewise, one avoids the problem of ensuring sufficient streamline coverage. (For improved accuracy and pointwise resolution, one can also use higher-order discontinuous Galerkin methods [3, 9, 11], which, like finite-volume methods, are applicable to general grid systems.)

Because each cell can contain several influence regions, it is sometimes advantageous to compute time-of-flight associated with each influence region by solving the following equation for τ

$$\vec{v}\cdot\nabla(\tau C) = \phi C, \qquad \tau|_{\text{inflow}} = 1 \tag{2.3}$$

for a fixed C from (2.2) to reduce undesirable averaging effects.

[1] We emphasize that the word "tracer" here refers to purely numerical quantities that should not be confused with the numerical modelling of *inter-well tracers*, i.e., substances that are either placed in well completions or injected along with displacing fluids to monitor the flow inside the reservoir.

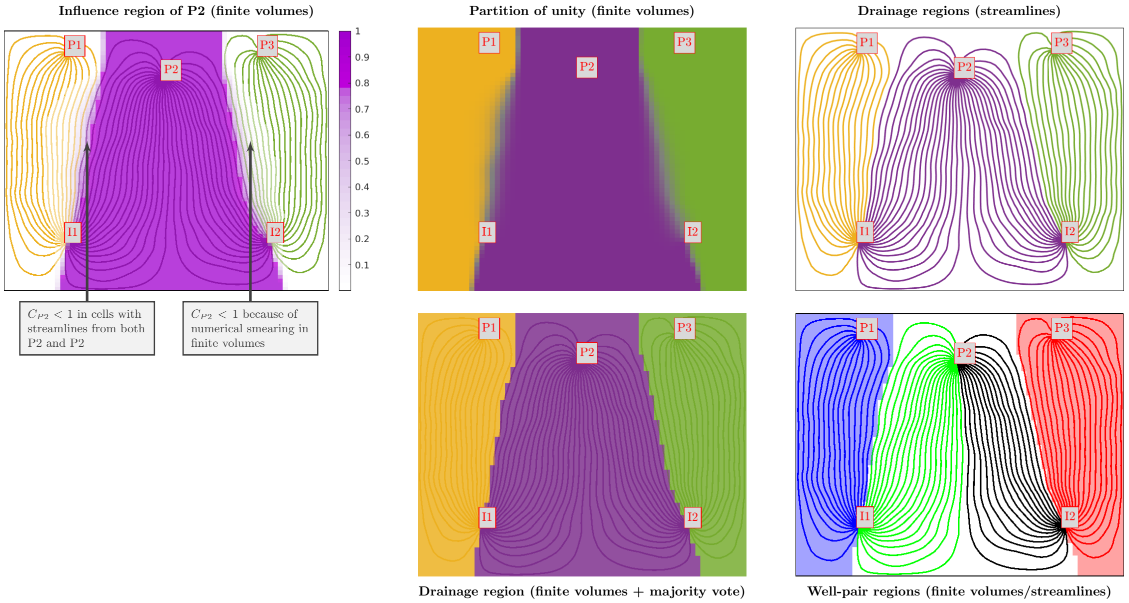
Influence region of P2 (finite volumes)
Partition of unity (finite volumes)
Drainage regions (streamlines)
P1
P2
P3
I1
I2
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
$C_{P2} < 1$ in cells with streamlines from both P2 and P2
$C_{P2} < 1$ because of numerical smearing in finite volumes
Drainage region (finite volumes + majority vote)
Well-pair regions (finite volumes/streamlines)

FIGURE 2.3: Volumetric partitions computed by labeling streamlines by the wells they connect to. Alternatively, one can solve one stationary transport equation of the type (2.2) for each well. The C-values for all injectors, or for all producers, define two partitions of unity of the reservoir volume. From these, we can define drainage and sweep regions and well-pair regions.

2.3 Volumetric partitions and well-allocation factors

From influence regions, we can define several quantities that represent the volumetric communication in the reservoir. Each influence region naturally delineates the drainage/sweep volume associated with a given producer/injector; for cells with multiple non-zero tracer values, we use a majority vote to assign a unique region to each cell. By computing the intersection of drainage and sweep regions, you can determine whether an injector has fluid communication with a producer and compute the corresponding flow volume (well-pair volume). This is illustrated in Figure 2.3.

You can also compute *well-allocation factors* that measure the communication strength between injectors and producers. The flux allocation from injector n to perforated cell c^m_k of producer m is defined as

$$\boldsymbol{a}^p_{mn}[c^m_k] = \boldsymbol{C}^i_n[c^m_k]\,\boldsymbol{q}[c^m_k], \tag{2.4}$$

where $\boldsymbol{C}^i_n$ denotes the vector of injector tracer values associated with well n and $\boldsymbol{q}$ the vector of volumetric source terms for all perforated cells. The flux allocation from producer m to injector n perforated in cell c^n_k is defined analogously as $\boldsymbol{a}^i_{nm}[c^n_k] = \boldsymbol{C}^p_m[c^n_k]\,\boldsymbol{q}[c^n_k]$. Collecting the allocation factors $\boldsymbol{a}^p_{mn}$ from all injectors n connected to producer m gives you the volumetric inflow rate that can be attributed to each of the connected injectors. Likewise, collecting the allocation factors $\boldsymbol{a}^i_{nm}$ gives the volumetric inflow rate into each of the connected producers that can be attributed to injector n. In the GUI, we plot these allocation factors as a set of stacked bars, one stack of bars for each perforation, where each bar represents one allocation factor, colored by a unique color for each connected well; see Figure 2.4.

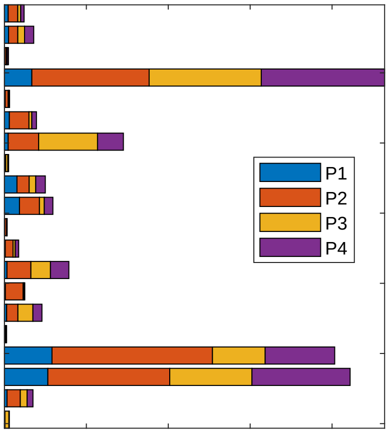

FIGURE 2.4: Allocation plot for an injector that communicates with four different producers. The stacked bars represent the flux out of each of the 20 perforations from heel (top) to toe (bottom).

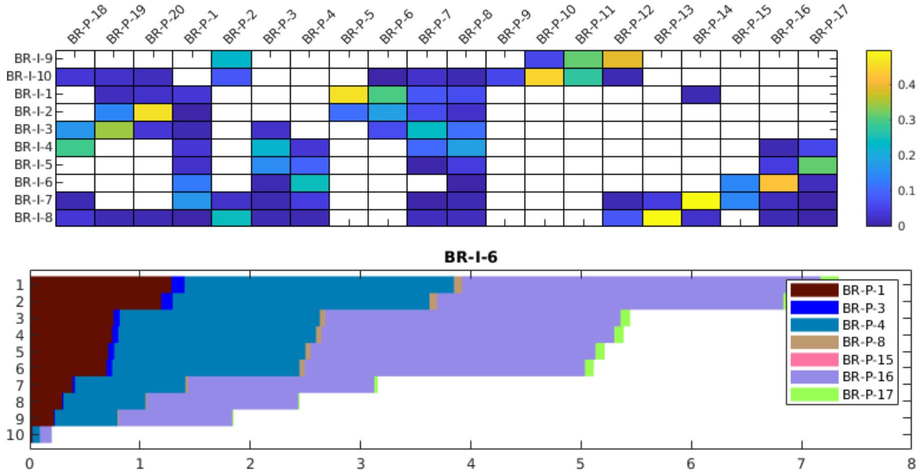

FIGURE 2.5: Visualization of well-allocation factors for a simulation of the Brugge model [10]. The upper plot shows the relative communication strength of each well pair over the whole simulation period, computed by summing the instantaneous injector allocation factors $\sum_k \boldsymbol{a}^i_{nm}[c^n_k]/\sum_k \boldsymbol{q}$ and integrating the result over all time steps in the simulation. The lower plot shows the instantaneous, cumulative allocation from toe to heel for a specific time step.

Example (Brugge field). Figure 2.5 shows two examples of how you can use this information to visualize the relative fluid communication in the Brugge simulation model [10]. The upper plot shows a matrix plot of the relative communication strength between the ten injectors (rows) and twenty producers (columns), averaged over all time steps in the simulation, Here, no color means no communication. If you look carefully, you can see that the upper plot shows significant communication between injector BR-I-6 and producer BR-P-16. The lower plot reports instantaneous outflow from well BR-I-6 at a specific time, shown as a cumulative plot from toe to heel (bottom to top of well). Here, you can see a large proportion of the outflow from BR-I-6 is allocated to producer BR-P-16.

2.4 Measures of dynamic heterogeneity

Secondary and tertiary recovery is usually strongly governed by the intrinsic variability (heterogeneity) in petrophysical properties. Classical sweep theory includes a number of static measures for characterizing heterogeneity, such as flow and storage capacity (Stiles' diagrams [18]), Lorenz coefficient, Koval factor, and Dykstra–Parson's permeability variation coefficient; see e.g., [6] for an overview. In flow diagnostics, some of these measures have been reinterpreted in a *dynamic* setting so that they measure the heterogeneity in flow paths (and connection structure) rather than measuring the heterogeneity in the spatial distribution of permeability and porosity. Large dynamic heterogeneity means large variations in the length and throughput of flow paths between injectors and producers, which in a water- or gas-flooding scenario typically manifests itself as early breakthrough of injected fluids.

F-Phi diagram: The first example of a dynamic heterogeneity measure is *flow and storage capacity*, which we compute from the total residence time (i.e., the sum of forward and backward time-of-flight) and the relationship $q_i\tau_i = V_i$ between pore volume V_i, flow rate q_i, and residence time τ_i of each cell c_i. To understand this measure, you can think of the reservoir as a bundle of streamtubes of infinitesimal width, sorted so that their residence times are ascending; see Figure 2.6. If we assume piston displacement (blue fluid displacing red fluid) inside each streamtube, the storage capacity Φ at time t is the volume of all streamtubes that have "broken through", i.e., the volume of all streamtubes that have a lower total residence time than t. With a slight abuse of notation, we write this as

$$\Phi(t) = \int_0^t \int_{\Psi(\tau)} \phi(\vec{x}_\Psi(s))\,ds\,d\tau. \tag{2.5}$$

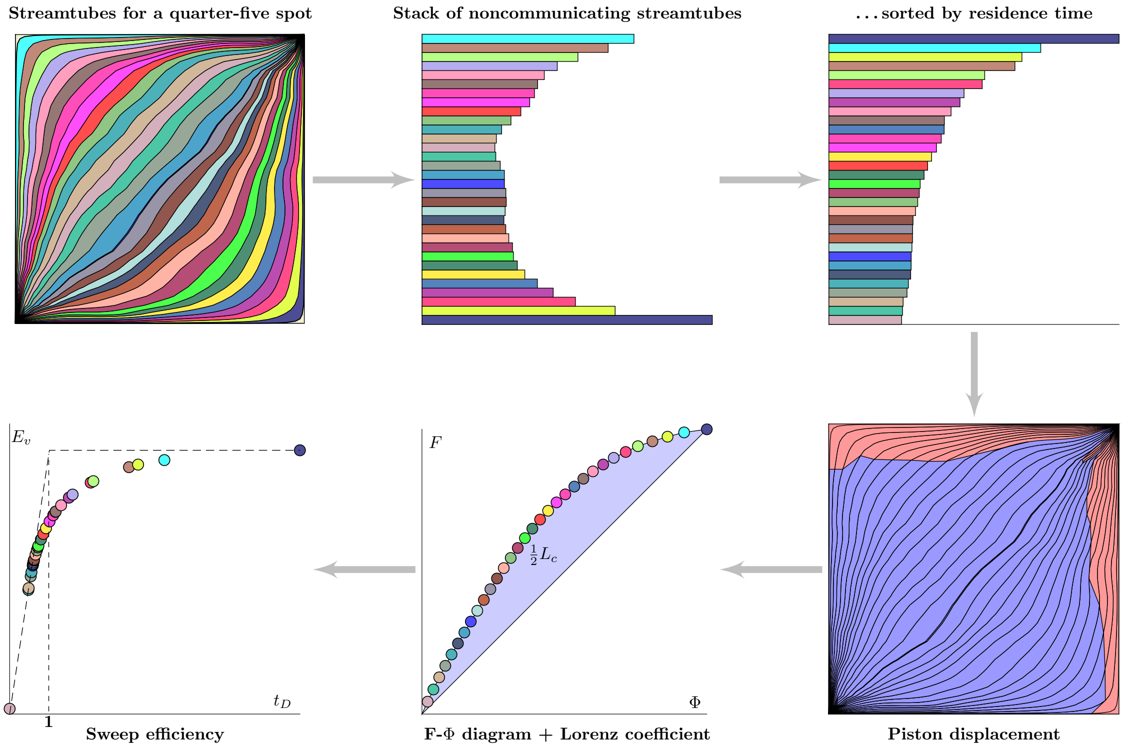

Streamtubes for a quarter-five spot

Stack of noncommunicating streamtubes

. . . sorted by residence time

Sweep efficiency

F-Φ diagram + Lorenz coefficient

Piston displacement

FIGURE 2.6: Illustration of how one can construct dynamic heterogeneity measures based on a bundle of streamtubes that divide the reservoir into a set of isolated flow channels that each has a given flow rate and a pore volume. The streamtubes are sorted according to residence time. If we now assume piston displacement with a blue fluid displacing a red fluid, we can record how the fractional flux F of blue fluid increases at the outlet as an increasing amount of streamtubes become completely filled by blue fluid. We also record the fractional volume Φ of the streamtubes that have been completely flooded. This gives the F-Φ diagram. The Lorenz coefficient is twice the area between the curves $y = F(\Phi)$ and $y = \Phi$. The sweep efficiency $E_v(t)$ is defined as the fraction of in-place fluid (red) that has been displaced by injected fluid (blue) by time t.

Here,$\Psi(\tau)$ is interpreted as all streamtubes with residence time equal τ. The flow capacity F is the corresponding fractional flow, i.e., the fraction of injected fluid to the total amount of fluid produced,

$$F(t) = \int_0^t \int_{\Psi(\tau)} q(\vec{x}_\Psi(s))\, ds\, d\tau = \int_0^t \int_{\Psi(\tau)} \frac{\phi(\vec{x}_\Psi(s))}{s}\, ds\, d\tau. \tag{2.6}$$

Both quantities are normalized by their value at time infinity, giving relationships as in Figure 2.6. The $F(\Phi)$ diagram is generally a concave function, except in the special case of a perfectly homogeneous displacement, for which $F = \Phi$. The steep initial slope in a concave $F(\Phi)$ diagram corresponds to high-flow regions giving early breakthrough and the flatter tail corresponds to low-flow and stagnant regions. The more concave $F(\Phi)$ is, the larger is the spread in residence times for characteristic flow paths.

Lorenz coefficient: This scalar quantity, typically denoted L_c, measures the difference in flow capacity from that of a perfectly homogeneous displacement and is defined as twice the area under the concave $y = F(\Phi)$ curve and above the straight line $y = \Phi$; see Figure 2.6. This means that $L_c = 0$ for a homogeneous displacement and $L_c = 1$ for an infinitely heterogeneous displacement. Experience has shown that the dynamic Lorenz coefficient in many cases correlates (surprisingly) well with forecasts of hydrocarbon recovery predicted by more comprehensive flow simulations as long as the initial fluid distribution is relatively uniform. It can hence be used as an effective flow proxy in various reservoir management workflows; see [8, 14, 15, 21].

Sweep efficiency: The *volumetric sweep efficiency* $E_v(t)$ measures how efficiently injected fluids are used to displace in-place fluids. It is defined as the volume fraction of in-place fluid that has been displaced by injected fluid, or equivalently, the ratio between the volume contacted by the displacing fluid at time t and the volume contacted at time $t=\infty$. Using forward time-of-flight τ_f, we can express it as,

$$E_v(t) = \frac{1}{\Phi_0}\int_{\{\vec{x}:\tau_f(\vec{x})\le t\}} \phi(\vec{x})\,d\vec{x} = \frac{1}{\Phi_0}\int_\Omega\int_0^t \delta\big(s-\tau_f(\vec{x})\big)\phi(\vec{x})\,ds\,d\vec{x}, \tag{2.7}$$

where Φ_0 is the total pore volume. If we sort the indices of the cells according to ascending τ_f values, we can compute a cell-based estimate of sweep efficiency as

$$E_v(t) = \sum_{\{j|\tau_{f,j}\le t\}} V_j \Big/ \sum_{j=1}^{N} V_j. \tag{2.8}$$

Alternatively, one can show that E_v can be computed from the F-Φ diagram using the following formula

$$E_v = \Phi + (1-F)\,t_D, \qquad t_D = \frac{d\Phi}{dF}, \tag{2.9}$$

where t_D represents dimensionless time. In a homogeneous piston displacement, all the in-place volume will be displaced by time t_D, and thus t_D represents units of pore volume injected. Before first breakthrough ($t < \min\tau_f$), E_v equals the injected pore volume t_D. After breakthrough, Φ is the volume fraction of flow paths that have been fully swept, whereas $(1-F)\,t_D$ is the volume fraction contributed by the swept parts of the flow paths that have not yet broken through. This means that for a heterogeneous displacement, the *volumetric sweep-efficiency diagram* (t_D, E_v) is a concave curve bounded above by $y=\min(x,1)$, see Figure 2.6, which highlights fluid responses after first breakthrough.

Fractional recovery: The *fractional recovery diagram* $(t_D, 1-F)$ emphasizes breakthrough behavior and can have utility as a proxy for early-time fractional recovery. The name can be somewhat misleading and should not be confused with recovery factors computed based on multiphase simulations.

Multiphase extensions: The heterogeneity measures described thus far are all computed without taking any multiphase information into account, other than possibly through the stationary flow field. This may not always be sufficient to obtain good correlation with recovery factors predicted by multiphase flow simulation simulation, in particular if the fluid distribution inside the reservoir is largely nonuniform. relatively uniform. Approaches that compute, e.g., the Lorenz coefficient to the individual phases [4, 8] have the weakness that they only consider the heterogeneity within the phase volume and not how far, measured in τ_b, any movable volume is from a producer. The exception is sweep efficiency, which generalizes more naturally to individual phases. The oleic sweep efficiency, for instance, reads [21]:

$$E_{v,o}(t) = \frac{1}{V_o}\int_{\{\vec{x}:\tau_b(\vec{x})\le t\}} \phi(\vec{x})S_o(\vec{x})\,d\vec{x} = \frac{1}{V_o}\int_\Omega\int_0^t \delta\big(s-\tau_b(\vec{x})\big)\phi(\vec{x})S_o(\vec{x})\,ds\,d\vec{x}, \tag{2.10}$$

where S_o is the oil saturation and $V_o=\int_\Omega \phi S_o\,d\vec{x}$ is the total oil volume. Here we have used τ_b to approximate the oil volumes produced at time t rather than the oil volumes displaced (as would be the case using τ_f). For the total sweep (2.7), however, the forward and backward expressions are equivalent.

2.5 Residence-time distributions

We have already emphasized that time-of-flight computed from a finite-volume discretization of the second equation in (2.1) represents a volumetric average over all τ values inside each grid cell. In particular, one can show that the residence time represents the pore volume of the backward influence region from the outlet point divided by the outflux at this point; see [5, 7] for details. Hence, this residence time represents the average of a distribution that potentially can have (very) large variance, as illustrated in Figure 2.7. Using such residence times to compute dynamic heterogeneity measures may

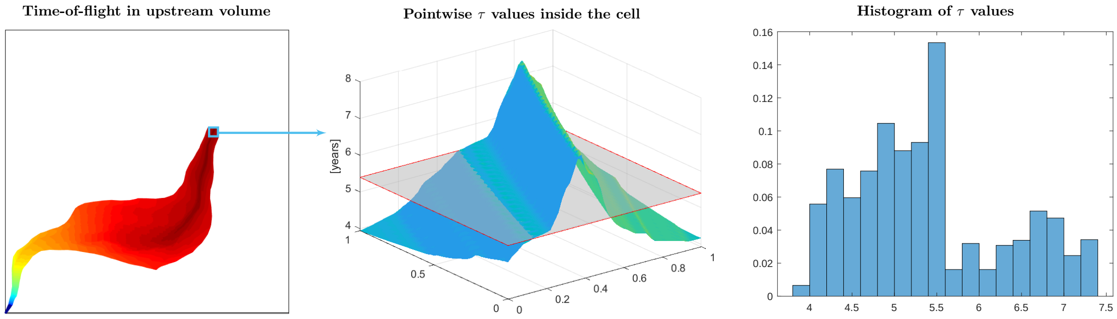

FIGURE 2.7: Illustration of subcell variation of time-of-flight. The left plot shows τ values inside the upstream volume of the cells, i.e., along all flow paths that pass through the cell. The colored surface in the middle plot exhibit the spatial variation of τ inside the cell, whereas the gray plane represents the average value. The right plot shows the histogram of τ values defined over a 51×51 mesh covering the cell.

thus potentially introduce a significant bias. Fortunately, this bias appears to be systematic [11], so that derived measures can still be used to, e.g., rank model ensembles or as simple reduced-order models to predict recovery of secondary oil recovery [8]. On the other hand, average TOF-values will in most cases *overestimate* the time to breakthrough in heterogeneous displacements (the arithmetic mean is significantly impacted by the large values that contribute little to flow). To compute times of first arrival more accurately, it is thus better to use a graph algorithm to compute the shortest path of the discrete flux graph. Likewise, one can estimate the variation σ_τ within a cell by computing second moment by solving $\vec{v}\cdot\nabla\sigma_\tau = 2\phi\tau$.

For a better description of the dynamic heterogeneity of a reservoir model, we can consider the *distribution* of time-of-flight for each cell. This is particularly interesting for cells perforated by production wells, since pointwise time-of-flight values in these cells describe the residence times, or time to breakthrough, for individual flow paths. To this end, we can solve the linear transport equation

$$\phi\frac{\partial\eta}{\partial t} + \vec{v}\cdot\nabla\eta = 0, \qquad \eta|_{\Gamma_i} = \delta(t), \qquad \eta(\vec{x},0) = 0, \tag{2.11}$$

and compute the evolution of a unit pulse from inflow Γ_i to outflow Γ_o, as illustrated in Figure 2.8. Herein, we use the backward Euler method for temporal discretization and the same upstream finite-volume method as for the steady-TOF equation for the spatial discretization. The resulting method is robust and can take large time steps, but is not very accurate. In each time step, we must solve a linear problem, but the coefficient matrix is triangular, possibly after a permutation, and hence inexpensive to invert. (This also holds for the discretization of (2.1) and (2.2).)

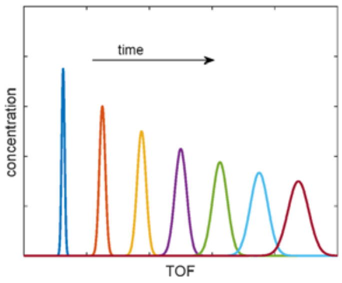

FIGURE 2.8: Numerically computed tracer pulse mapped onto a time-of-flight coordinate.

For each point $\vec{x}$, the normalized TOF-distribution $\mathcal{T}(\cdot;\vec{x})$ is now defined by the Dirac function

$$\mathcal{T}(t;\vec{x}) = \eta(\vec{x},t) = \delta\big(t-\tau(\vec{x})\big). \tag{2.12}$$

At the outflow, the normalized residence-time distribution (RTD) is given as

$$\mathcal{T}_o(t) = \frac{1}{F_o}\int_{\Gamma_o} \eta(\vec{x},t)\,\vec{v}\cdot\vec{n}\,ds, \qquad F_o = \int_{\Gamma_o} \vec{v}\cdot\vec{n}\,ds. \tag{2.13}$$

It follows from the definition of the Dirac distribution that $\int \mathcal{T}_o(t)\,dt = 1$. We can also use this distribution to compute flow and storage capacity [16]

$$F(t) = \int_0^t \mathcal{T}_o(s)\,ds, \qquad \Phi(t) = \frac{F_o}{\Phi_o}\int_0^t s\mathcal{T}_o(s)\,ds, \tag{2.14}$$

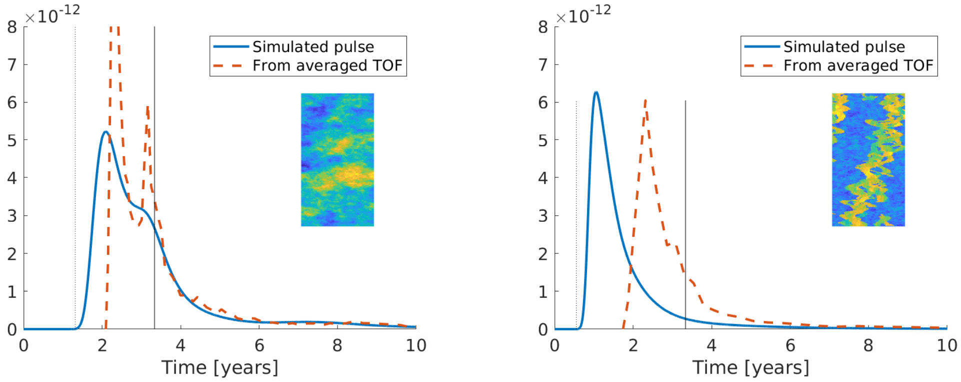

FIGURE 2.9: Residence-time distributions for two subsamples from the SPE 10 benchmark [1]: Tarbert (left) and Upper Ness (right). Thin solid and dotted lines represent the mean of the distribution, i.e., time to inject one pore volume and time to breakthrough for the fastest flow path.

where Φ_o is the total pore volume drained by the outflow boundary Γ_o. As before, both quantities are normalized so that $F(\infty) = \Phi(\infty) = 1$. From this definition, it also follows that the mean value of $\mathscr{T}_o(t)$ corresponds to the time $\bar{t} = \Phi_o/F_o$ it takes to inject one pore volume (1 PVI).

To illustrate the difference between averaged time-of-flights computed from (2.1) and the residence-time distributions, we include an example from [5, 7]. The solid lines in Figure 2.9 report $\mathscr{T}_o(t)$ without normalization as function of time for two rectangular reservoirs with an injector along the south boundary and a producer along the north boundary. (The integral of this curve equals the total allocation.) The leading pulse for the Tarbert layer is spread out and has a small secondary hump. For Upper Ness, the pulse breaks through earlier and is more focused because of high permeability channels connecting the south and north boundaries. The mean of each distribution equals 1 PVI by construction. This may not be apparent from the plots, since the distributions have very long tails, particularly in the channelized case. We can also estimate the same distribution from the average residence times τ_r; details are given in [7, §13.3]. Using averaged TOF-values introduces a significant delay in the breakthrough time, in particular for Upper Ness. The two types of distributions also suffer from different types of numerical errors: Tracing a pulse by a finite-volume method preserves flux allocation and not pore volume and can also contain significant temporal smearing. Backing out a distribution from averaged TOF values preserves total volume but not flux allocation. The corresponding F-Φ diagrams computed from the residence-time distribution are more concave and the Lorenz coefficients are somewhat larger than when computed from averaged TOF-values. However, the differences are not very large, and since the measures computed from averaged TOF-values carry a systematic underestimation bias and are much quicker to compute, they can thus still be robustly used to rank and discriminate different cases. For more accurate predictions of recovery and recovery factors, RTD-measures should be used [21].

2.6 Extension to compressible flow

So far, we have assumed that the flow field is steady, immiscible, and incompressible. Most multiphase simulations used in reservoir engineering consist of multiple time steps and assume *compressible* conditions. Typically, the flow paths and the inter-well communication will change over time because the fluid mobility changes when fluids displace each other. Flow in and out of wells varies with time because of pressure and mobility changes as a result of changes in well controls, in particular changes that turn wells on and off. The flow diagnostics discussed earlier in this chapter only depict time lines, volumetric connections, and heterogeneity in dynamic flow paths *at each instance in time*, assuming that the current flow field persists till infinity. This is the same as in streamline methods [2], in which the streamlines are instantaneous views of the flow field that are bound to change over time.

Compressibility and interphase mass transfer will generally cause fluid compression and expansion

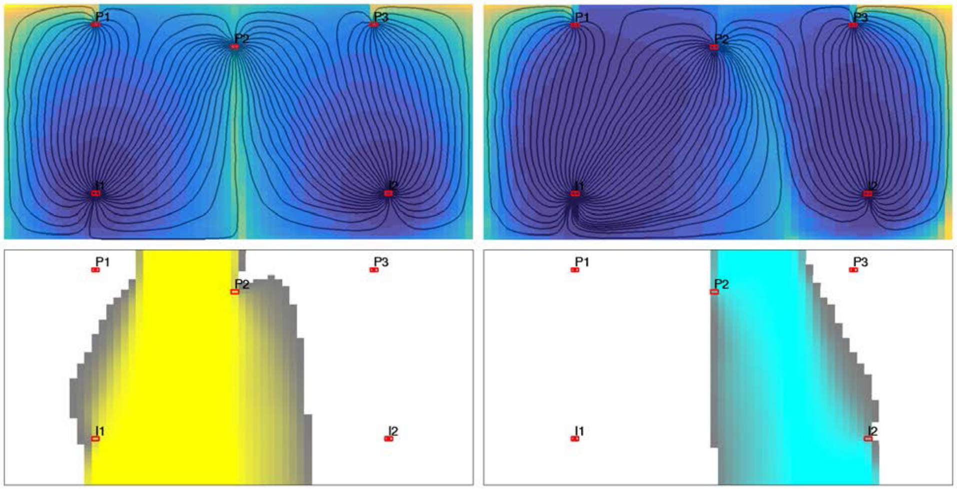

FIGURE 2.10: The flow field will typically change throughout a dynamic simulation, here exemplified by the streamlines and time-of-flight initially and at the end of a simulation in which the injection rate of the left injector increases relative to that of the right injector. Summing the tracer partitions for each time step gives a measure of how long each grid cell has been part of a given well-pair region. In the bottom row, bright colors signify that cells are part of the well-pair region for all time steps, whereas grayish colors signify that a cell is part of a region for a fraction of the simulation time; the more gray, the less time the cell has been in the region.

throughout the reservoir. However, these effects do not change the computation of time-of-flight and stationary tracer partitions. To see this, we start by writing the time-of-flight equation in conservative form, $\nabla\cdot(\vec{v}\tau) - \tau\nabla\cdot\vec{v} = \phi$, and let ϕ_i and τ_i denote the porosity and the unknown TOF value in cell i and v_{ij} be the flux from cell i to cell j. Moreover, let $U(i)$ be the set of upstream neighbors of cell i, i.e., the set of indices j such that v_{ji} is directed *into* cell i. Likewise, we let $D(i)$ be the set of downstream neighbors. For incompressible flow, the standard upwind finite-volume method reads,

$$\phi_i = \sum_{j\in U(i)} v_{ij}\tau_j + \sum_{j\in D(i)} v_{ij}\tau_i = -\sum_{j\in U(i)} v_{ji}\tau_j + \tau_i \sum_{j\in U(i)} v_{ji}.$$

because $v_{ji} = -v_{ij}$ and $\sum_{j\in D(i)} v_{ij} - \sum_{j\in U(i)} v_{ji} = 0$. For compressible flow, it follows by cancellation of down-wind fluxes that

$$\phi_i = \sum_{j\in U(i)} v_{ij}\tau_j + \sum_{j\in D(i)} v_{ij}\tau_i - \tau_i\Big(\sum_{j\in U(i)} v_{ij} + \sum_{j\in D(i)} v_{ij}\Big) = -\sum_{j\in U(i)} v_{ji}\tau_j + \tau_i \sum_{j\in U(i)} v_{ji}. \tag{2.15}$$

The main change lies in how we compute volumetric partitions and rate allocations in wells. For incompressible flow, it follows from mass conservation that outflow must equal inflow and hence we can connect all cells in the grid backward to fluid sources (injection wells or inflow boundaries) and forward to fluid sinks (producers and outflow boundaries). Compressible flow has additional fluid sources because of fluid expansion induced by decaying reservoir pressure, and not all flow through a given cell can thus be attributed to an injector or inflow boundary. Some cells may even not be connected to injectors at all. For simplicity, we have chosen to lump all such effects into an additional category called "reservoir" when reporting and plotting well-pair regions and well-allocation factors.

Given this choice, it is straightforward to extend existing tools from MRST to compute and display basic flow quantities such as well-pair flux allocations, well-pair volumes, and flooded and drainage regions as time-dependent variables. This enables you to quickly screen multiple output states from a reservoir simulator to get an idea of how the flow patterns vary throughout the course of time, e.g., as illustrated in Figures 2.10 and 2.11. You can also quantify changes by comparing differences in heterogeneity measures like Lorenz coefficient and sweep efficiency factor between individual time steps.

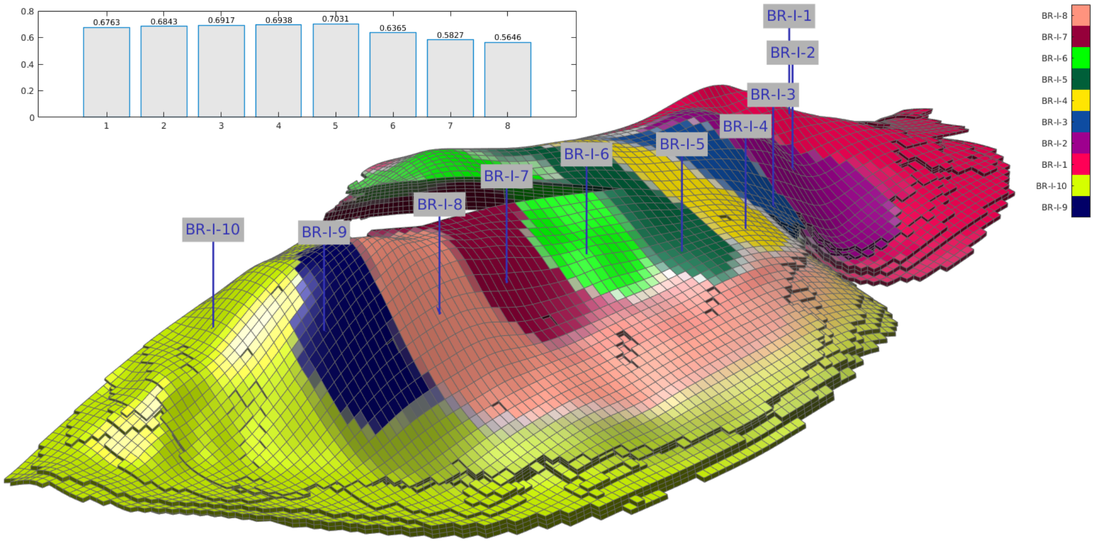

FIGURE 2.11: Time-dependent flow diagnostics for the Brugge field. The large plot shows how the sweep regions develop over a four-year period. Cells with bright colors are part of the same sweep region over the whole period, whereas grayish colors signify cells that are associated with different sweep regions over the time interval. The bar plot shows how the Lorenz coefficient for the whole reservoir develops over time.

CHAPTER

3 STARTING THE GUI AND SELECTING RESTART STEPS

To perform a flow diagnostic analysis of a multiphase simulation, you must do two things. First, you must select the output file and the range of restart steps (ECLIPSE) or timesteps (MRST) you want to load into the GUI. Multiphase simulations typically contain many report steps, which in turn amounts to a lot of data that need to be processed. It can therefore be advantageous to only load a limited number of steps, e.g., by specifying only a certain time period or every *n*th step. Second, you select the restart steps among those loaded for which you want to (compute and) display properties in the three plotting panels.

There are some minor differences when loading MRST simulations compared to loading ECLIPSE simulations. These are explained in the first section of this chapter.

3.1 Starting the GUI and loading a data set

To process ECLIPSE data you start the graphical user interface with the following command(s)

```
mrstModule add diagnostics
PostProcessDiagnosticsECLIPSE;
```

(The first line is only necessary if you have not yet loaded the `diagnostics` module into MRST.) This will bring up the dialog box shown in Figure 3.1 to the right, which asks you to navigate to the correct subfolder and select the file containing the case. Alternatively, you can specify the path of the file directly, e.g., something like (on Linux)

```
PostProcessDiagnosticsECLIPSE(...
      '/home/username/project/Brugge/data/HISTSW');
```

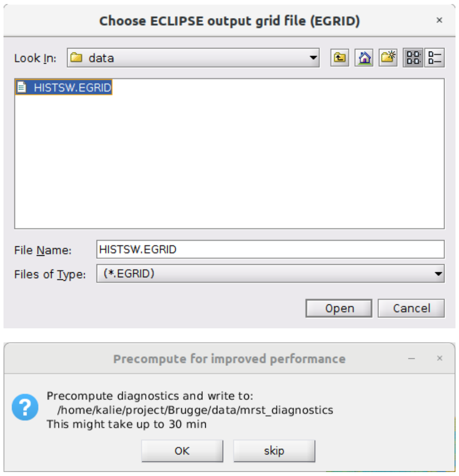
Choose ECLIPSE output grid file (EGRID)

Look In: data

HISTSW.EGRID

File Name: HISTSW.EGRID

Files of Type: (*.EGRID)

Precompute for improved performance

Precompute diagnostics and write to:
  /home/kalie/project/Brugge/data/mrst_diagnostics
This might take up to 30 min

FIGURE 3.1: Dialog boxes for loading ECLIPSE files and precomputing flow diagnostics.

If this is the first time you load this data set, you get another dialog box asking whether you wish to precompute flow diagnostics and save the result to disk. This is not necessary but will ensure that switching between different report steps is faster and that the whole user interface loads much faster the next time you wish to look at this particular data set. Here, we click OK on the precompute data dialogue box, which produces the following output in the MATLAB command window:

```
Creating MRST-grid ...done.
Creating simulation grid ...done
Reading restart: 100%,  done
Converting restart to mrst-states: 100%,  done
Computing diagnostics: 100%, done
```

The equivalent commands for loading output from an MRST simulation are

```
mrstModule add diagnostics
PostProcessDiagnosticsMRST(problem);
```

This assumes that you have set up the `problem` data structure using `packSimulationProblem` and run your simulation using the `simulatePackedProblem` function, both from the `ad-core` module of MRST.

Next, you get a menu (Figure 3.2) that asks you to select the report steps you want to load. The list box at the top gives all the report steps available. These can be marked individually, or you can mark the first report step and use the pop-up menu at the bottom to set a time interval between subsequent report steps. Using the two sliders, you can also limit the time interval in which report steps can be selected. In Figure 3.2, we have selected the first report step (Jan 01, 1998) and then loaded report steps with a time interval corresponding to approximately every 6th month onward. After you have finished your selection and have clicked the OK button, the menu disappears and MRST starts reading the input.

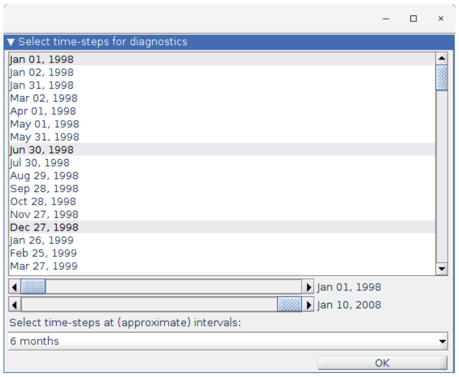
▼ Select time-steps for diagnostics

Jan 01, 1998
Jan 02, 1998
Jan 31, 1998
Mar 02, 1998
Apr 01, 1998
May 01, 1998
May 31, 1998
Jun 30, 1998
Jul 30, 1998
Aug 29, 1998
Sep 28, 1998
Oct 28, 1998
Nov 27, 1998
Dec 27, 1998
Jan 26, 1999
Feb 25, 1999
Mar 27, 1999

Jan 01, 1998
Jan 10, 2008

Select time-steps at (approximate) intervals: 6 months

FIGURE 3.2: Selecting report steps to be imported from the ECLIPSE output file. Here, we have selected (Jan 01, 1998) as the first step in the list box, and then used the pull-down menu at the bottom to pick steps for every 6th month.

For ECLIPSE input the software first creates a standard grid structure based on the grid information exported to the output file. This structure will be used to display 3D data. Next, it reads the exported transmissibilities and pore volumes and creates a simulation graph similar to the one that was used internally in the multiphase simulation. This graph will be used to compute flow diagnostics and residence-time distributions as discussed in Section 2.5. Finally, the software loads static and dynamic data from the restart file(s). The specific selection in Figure 3.2 gives the following output

```
Loading precomputed diagnostics 100%, done
Creating MRST-grid ...done.
Creating simulation grid ...done
Reading restart: 100%,  done
Warning: Current version requires at least one open producer and one open injector.
Skipping  4 of the 21 selected restart steps.
> In readAndPrepareForPostProcessor (line 43)
  In PostProcessDiagnostics (line 91)
Reading info from roughly 138 ministeps:  100%
Actual number of ministeps: 135
```

The warning informs you that no injector–producer pairs were active during the first four of the 21 selected restart steps and hence these states should not be loaded, since they cannot be used to compute any flow diagnostics.[1] Likewise, the output informs you that we have additionally loaded well responses from 135 time steps (called ministeps) used internally by the simulator over the selected time interval. The ECLIPSE output does not contain any cell properties for these ministeps, only the well responses that will be discussed in Section 4.2.

Then, the main GUI window shown in Figure 3.3 is launched. The GUI consists of a main plotting panel for 3D visualization of static and dynamic cell properties and two smaller panels for displaying various kinds of derived 2D (and 3D) plots (see also Figure 1.1).

3.2 Select report steps

The topmost menu shows a list of the report steps loaded into the GUI. To be able to inspect any data, except for static petrophysical data and production data, you must select at least one report step. You can select individual report steps using the left mouse button, multiple report steps by holding down Ctrl or Shift while clicking the left mouse button. Once you have selected a report step, you can also move up and down in the list using ↑ or ↓. Likewise, you can use the right mouse button to bring up a pop-up menu that lets you select all report steps or clear the current selection.

[1] Flow diagnostics *can* be defined for primary production by assuming quasi-steady flow conditions (i.e., constant pressure in each cell), but this is not implemented in the current version of the GUI.

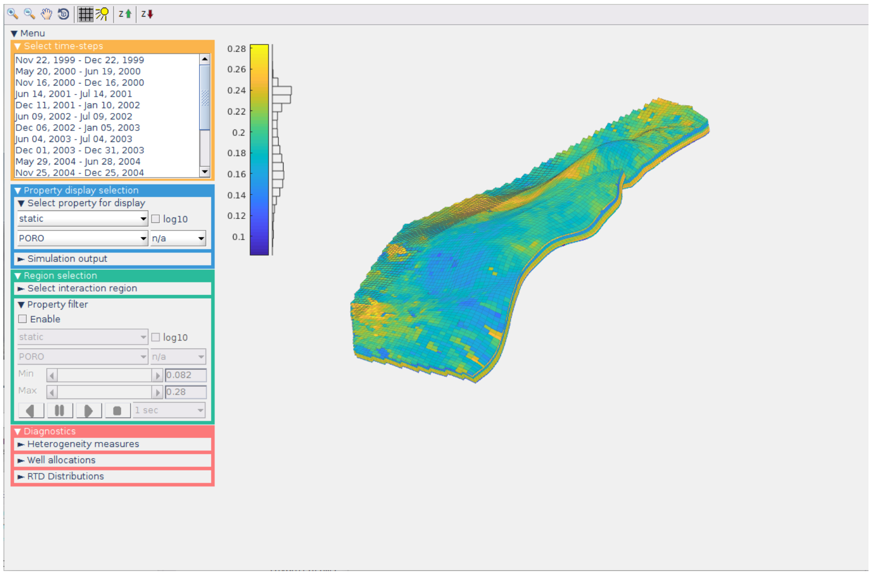

Nov 22, 1999 - Dec 22, 1999
May 20, 2000 - Jun 19, 2000
Nov 16, 2000 - Dec 16, 2000
Jun 14, 2001 - Jul 14, 2001
Dec 11, 2001 - Jan 10, 2002
Jun 09, 2002 - Jul 09, 2002
Dec 06, 2002 - Jan 05, 2003
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004

static □ log10
PORO n/a

static □ log10
PORO n/a
Min 0.082
Max 0.28

FIGURE 3.3: The GUI lists all loaded report steps at start-up, but before you have selected a specific report step, the GUI can only visualize static properties from the geocellular model in the main panel and production data over the whole simulation period in each of the two subpanels.

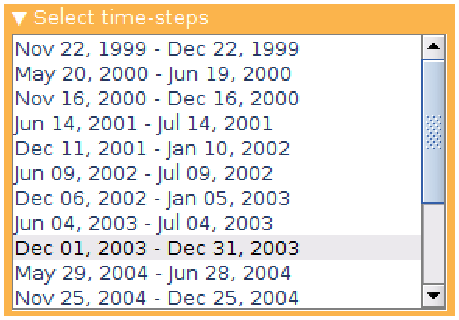

Nov 22, 1999 - Dec 22, 1999
May 20, 2000 - Jun 19, 2000
Nov 16, 2000 - Dec 16, 2000
Jun 14, 2001 - Jul 14, 2001
Dec 11, 2001 - Jan 10, 2002
Jun 09, 2002 - Jul 09, 2002
Dec 06, 2002 - Jan 05, 2003
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004

FIGURE 3.4: Menu for selecting report steps for display in the GUI.

If no report steps are selected, most of the menus are collapsed as shown in Figure 3.3. Once you select a report step, the GUI will immediately expand the menu for selection of interaction regions discussed in Section 5.2 and the menu for dynamic heterogeneity measures discussed in Section 6.1. Likewise, plotting of dynamic, diagnostic, and computed data is enabled (see Section 4.1).

Each time you select a new time step, the content of the main panel is updated immediately (or as soon as the corresponding flow diagnostics computation is finished if you have chosen not to precompute flow diagnostics). This means that if you have selected to display one of the dynamic or diagnostics properties (see Section 4.1), you can use the ↑ or ↓ buttons to quickly see how these properties change with time. (Notice that this type of animation can be slow if you have not chosen to precompute flow-diagnostics properties when you loaded data.)

Note: By design, the content of each of the two subpanels for 2D plotting is not updated when you select new time step(s). Instead, each of these plotting panels is given a gray background color to signify that its content no longer refers to the current selection of report steps.

Chapter

4 Displaying static and dynamic properties

This chapter discusses how you can select cell properties to display in the main panel. The GUI can also display other types of simulation responses found in the output file(s) in the subpanels. Typical examples are production responses associated with individual wells, groups of wells, or the field as a whole.

4.1 Displaying cell properties

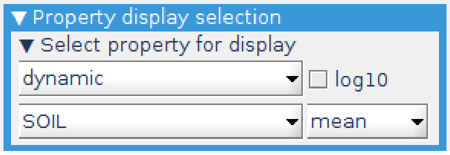

Figure 4.1: Menu for selecting the property displayed in the main plotting panel.

The 3D plotting panel occupies the largest portion of the GUI and can display a variety of cell properties on the full model or on a subregion selected as explained in Section 5.2. The quantities available for plotting vary with the data available in the output file and whether you have selected one or more restart steps. The 3D plot includes a color bar with an associated histogram giving the distribution of property values as shown in Figure 3.3. You can also use the eight buttons along the top of the GUI to zoom, pan, and rotate the view, toggle light and an extra wireframe grid on/off, and increase/decrease the z-aspect ratio.

You select which quantity to plot using the second menu box from the top (Figure 4.1). The menu box contains three drop-down menus and a check box:

- The first drop-down selects what type of property to display: static, dynamic, diagnostics, or computed from within this GUI.
- The second drop-down lets you select a specific property within the chosen property type.
- The third drop-down determines whether to plot the mean value, standard deviation, or maximum difference over all selected steps if you have selected multiple restart steps.
- The check box switches between displaying the property on a linear or logarithmic scale.

Before any restart step is selected you can only display `STATIC` properties describing the geology. Typical properties that are found in most ECLIPSE output files include (we stick to the naming conventions found in such files):

`PORO`	– porosity
`PERMX`	– permeability in the x-direction
`PERMY`	– permeability in the y-direction
`PERMZ`	– permeability in the z-direction
`NTG`	– net-to-gross
`DEPTH`	– depth to cell centroid

`DYNAMIC` properties consist of primary unknowns describing the reservoir state, including:

`PRESSURE`	– reservoir pressure (which in most cases is oil pressure)
`SOIL`	– oil saturation (e.g., as displayed in Figure 5.4)
`SWAT`	– water saturation
`SGAS`	– gas saturation

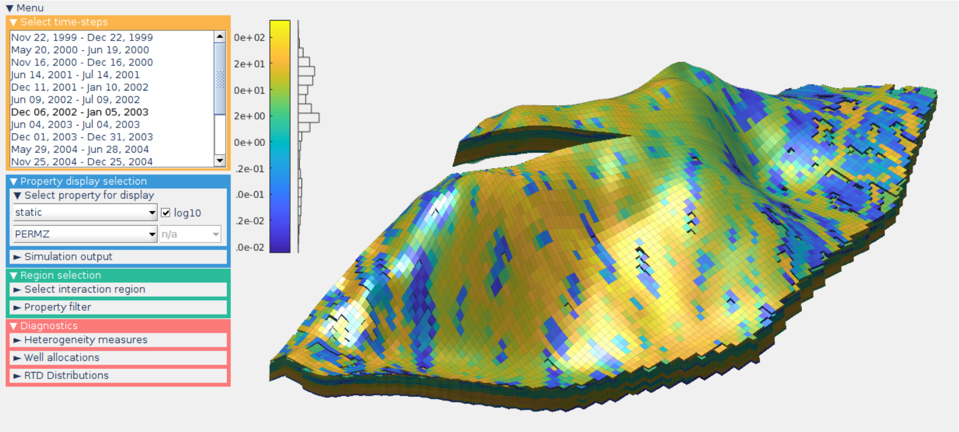
▼ Menu
▼ Select time-steps
Nov 22, 1999 - Dec 22, 1999
May 20, 2000 - Jun 19, 2000
Nov 16, 2000 - Dec 16, 2000
Jun 14, 2001 - Jul 14, 2001
Dec 11, 2001 - Jan 10, 2002
Jun 09, 2002 - Jul 09, 2002
Dec 06, 2002 - Jan 05, 2003
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004
▼ Property display selection
▼ Select property for display
static
☑ log10
PERMZ
n/a
► Simulation output
▼ Region selection
► Select interaction region
► Property filter
▼ Diagnostics
► Heterogeneity measures
► Well allocations
► RTD Distributions
0e+02
2e+01
0e+01
2e+00
0e+00
.2e-01
.0e-01
.2e-02
.0e-02

FIGURE 4.2: Vertical permeability for the Brugge model shown with a logarithmic color scale.

DIAGNOSTICS properties consist of a series of basic and derived quantities that give time lines and volumetric partition of the reservoir volume:

- Forward TOF – time-of-flight solving $\vec{v} \cdot \nabla \tau_b = \phi$ starting at injectors
- Backward TOF – time-of-flight solving $-\vec{v} \cdot \nabla \tau_b = \phi$ starting at producers
- Residence time – the sum of forward and backward time-of-flight, $\tau_r = \tau_f + \tau_b$
- Forward tracer – influence region solving $\vec{v} \cdot \nabla C^I = 0$ starting at each injector
- Backward tracer – influence region solving $-\vec{v} \cdot \nabla C^P = 0$ starting at each producer
- Tracer product – the product of forward and backward tracers, $C^I C^P$
- Sweep region – determined by majority vote over all C^I in each cell
- Drainage region – determined by majority vote over all C^P in each cell
- First arrival forward – time-of-flight for shortest forward flow path crossing cell
- First arrival backward – time-of-flight for shortest backward flow path crossing cell

The last category, COMPUTED, refers to cell properties computed inside the GUI. No such properties are computed in the current version, but functionality may be added later.

Example:

In Figure 4.2, we have used the GUI to display the vertical permeability of the Brugge field before selecting any restart step. The plot uses a logarithmic color scale (selected by the log10 check box). We have also used the buttons in the top menu to rotate the view, add lighting, and slightly increase the aspect ratio in the z-direction.

4.2 Simulation output (well responses)

The primary output from a reservoir simulation is in most cases well and field responses such as injection and production rates for individual phases, fractional fluid content in the inflow/outflow, bottom-hole pressures, field-average pressures, etc. In ECLIPSE output files, these are reported not only for each report step, but also at the end of any intermediate computational steps. For illustration herein, we selected to load 21 restart steps, approximately one half year apart, from a simulation of the Brugge benchmark case. The preprocessing summary on page 16 shows that the corresponding time interval consisted of 135 ministeps that each reports well responses.

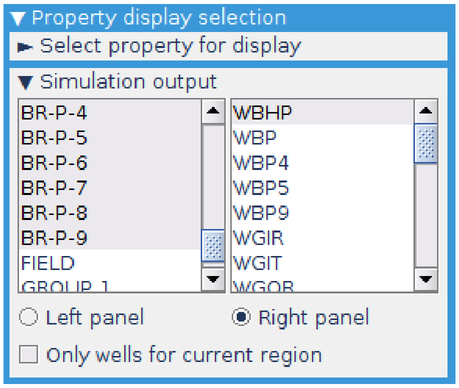

FIGURE 4.3: Menu for selecting simulation output to be displayed in the subpanels.

The list box to the left in the menu shown in Figure 4.3 lets you choose wells whose responses you want to plot. You can limit the list of wells to the current active region (set by the interaction-region and property-filter menus) by ticking the check box at the bottom of the menu. The list may also include groups of wells and `FIELD` that refers to the field as a whole. Once you have selected one or more entries from the list, the right list box is filled with the properties available for plotting. For simplicity, we have chosen to retain the names used by ECLIPSE for different output fields. In Figure 4.3, the properties available for injector BR-I-3 include the bottom-hole pressure (WBHP), 1-point to 9-point pressure averages (WBP to WBP9), gas-injection rate (WGIR), gas injection total (WGIT), and gas-oil ratio (WGOR), to name a few (see [13] for a full description of specific keywords). The two radio buttons let you select which of the two plotting panels to use. If you choose more than one property, they will be plotted together in the selected panel.

For MRST input, the output fields available from the well responses are bottom-hole pressure (bhp) and surface flow rates for water, oil, and gas (qWs, qOs, qGs respectively). When MRST input is used there is also a button labelled "plotWellSols" that launches the `plotWellSols` GUI. This provides more advanced functionality for plotting well data, including plotting more fields and using different units.

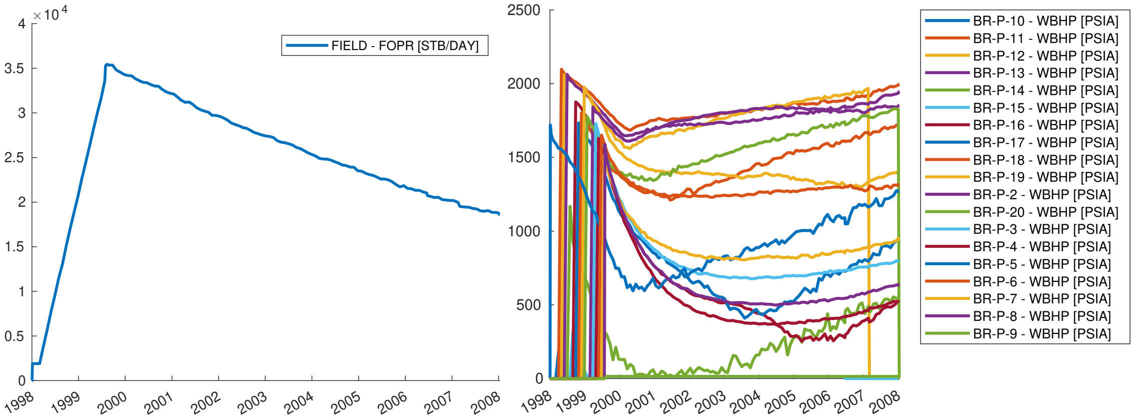

FIGURE 4.4: Field oil-production rate (left) and bottom-hole pressure for all producers.

Example:

Figure 4.4 shows two examples of simulation responses. The field-oil production rate in the left plot increases steeply during the primary-production period from January 1998 to August 1999. To produce this plot, we perform the following steps:

1. Select `FIELD` in the left list box, using the left mouse button.
2. Select `FOPR` in the right list box, using the left mouse button.
3. (Optionally, right-click inside the plot to export to a new figure window.)

The bottom-hole pressures in the right plot show that BR-P-10 is the first well to start producing and that the remaining producers are opened in the period up to August 1999. This plot is produced by the following steps:

1. Select `BR-P-1` to `BR-P-20` in the left list box, using the left mouse button with multi-selection.
2. Select `WBHP` in the right list box.
3. Right-click inside the plot to export to a new figure.
4. Manipulate the legend to get it outside of the plot.

A similar plot of bottom-hole pressure for injectors shows that these are turned on one-by-one from September 1999 to July 2000.

In this particular case, the output file also contains data for one group of wells and the time steps (found under the name `empty`).

CHAPTER

5 SELECTING A SUBSET OF THE RESERVOIR

Volumetric connections can be quite complex and difficult to understand in large reservoirs with many wells. To get a better understanding, you may therefore want to only view parts the reservoir and possibly analyze the fluid communication within individual drainage or flooding regions, or between pairs of injectors and producers. The GUI offers two different ways you can use to restrict your view to a subset of the reservoir. The first is a standard property filter that only affects the 3D display of the reservoir. The second lets you restrict both the 3D view and the computation of flow diagnostics shown in the two subpanels to a region of the reservoir.

5.1 Property filter

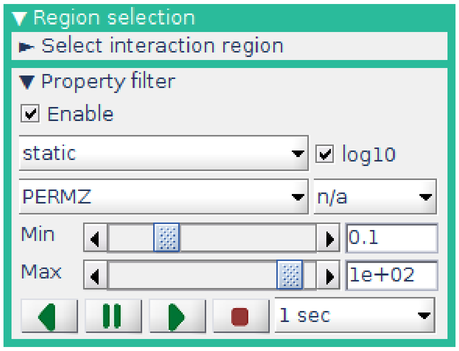
▼ Region selection

► Select interaction region

▼ Property filter

☑ Enable

static ☑ log10

PERMZ n/a

Min 0.1

Max 1e+02

1 sec

FIGURE 5.1: Property filter used to restrict the region displayed in the 3D panel. Can also play an animation with increasing filter values.

The property filter lets you display only those cells in the 3D plotting panel whose property falls within a certain range. You can for instance display cells with porosity value less than 0.1, horizontal permeability between 0.1 and 100 mD, cells with time-of-flight less than 5 years, all cells through which the average residence time of all flow paths is less than 10 years, and so on. The property used for filtering does not have to be the same as the property displayed. Notice that this selection only affects the 3D display and not the computation of flow diagnostics.

The main setup for the filter (see Figure 5.1) consists of a check box to turn the filter on/off, the same two selectors as in the property-display menu to select which property to filter with, and two sliders with input fields that let you determine the upper and lower threshold for the filter. With the ☐ log10 check box enabled, the sliders will run on a logarithmic scale rather than on a linear scale from minimum to maximum values. The drop-down menu below (n/a ▽) in Figure 5.1) tells whether the property filter should be applied to mean, standard deviation, or maximum difference. Settings for these last two need not be the same as in the property-selection menu.

Some properties, like time-of-flight and residence times, can in principle have values that span from zero to infinity, which makes visualization somewhat challenging. In the GUI, we cap all computed time-of-flights by an upper value defined by the `maxTOF` field in the `PostProcessDiagnostics` class (see Appendix B); the default is 500 years. This means that time-of-flight and residence times have already been filtered once when they appear in the property filter.

The menu also contains a “video control” that lets you create an animation effect. The animation starts by both thresholds set equal to the minimum property value and then gradually increases the upper threshold toward the maximum property value. The associated drop-down menu sets the approximate length of the animation in seconds. Once started, the animation will continue to loop until you push the pause (❙❙) or stop (■) buttons. Like with the rest of the filter, the property used to create the animation effect needs not be the same as the actual property displayed in the 3D plotting panel.

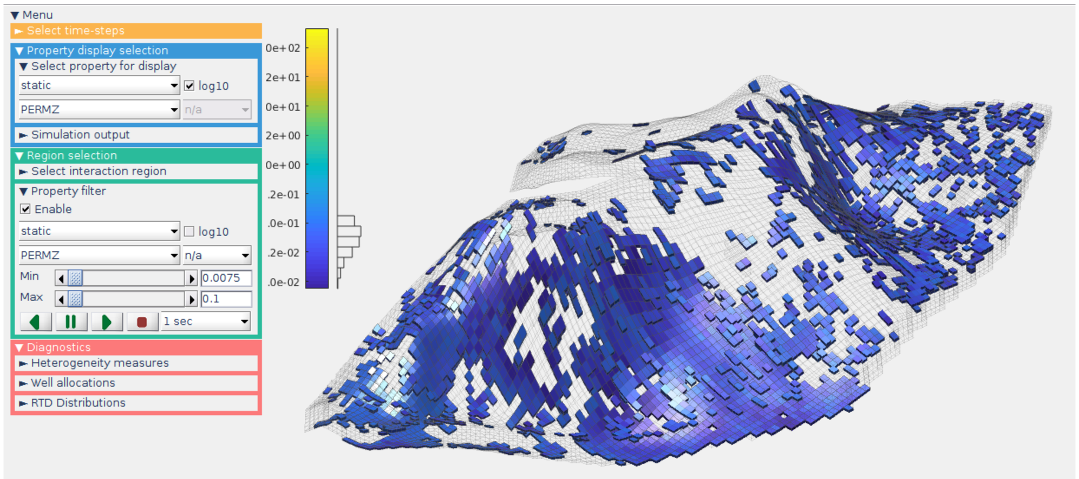

FIGURE 5.2: Vertical permeability for the Brugge model shown with a logarithmic color scale and filtered to only show values for which $K_z < 0.1$ md.

Example:

Figure 5.2 shows the same plot as in Figure 4.2 except that we now have filtered out all cells in which the vertical permeability exceeds 0.1 md. To show more clearly where such cells are found inside the reservoir, we recommend that you toggle the grid view (the fifth button with the mesh in the top row) so that this is "on".

5.2 Select wells and interaction regions

A primary purpose of the diagnostics GUI is to analyze the connections and communication between individual wells in the reservoir. The upper part of the region-select menu shown in Figure 5.3 lets you select individual wells or pair of wells and their corresponding influence/interaction regions. The menu box contains three parts:

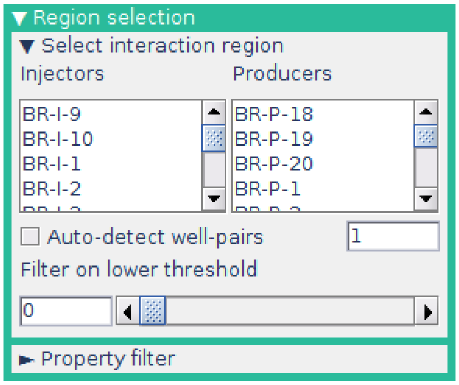

FIGURE 5.3: Menu for selecting wells and the subset of the reservoir used when presenting flow diagnostics and tracer analysis.

- Two list boxes giving the injectors and producers. Each of these lets you select individual wells or sets of wells using the mouse as described in Section 3.2. This will restrict the 3D plot of cell values discussed in Section 4.1 to the corresponding influence regions. Likewise, any type of flow diagnostics displayed will be restricted to these wells.

- A check box for enabling automatic selection of well pairs. If you select one or more injectors when auto-selecting is enabled, the software will automatically add all connected producers to your selection. Selecting one or more producers works the same way by adding all connected injectors. Whether an injector and a producer are connected is determined from the connection matrix $\boldsymbol{A}$ defined in (6.2) in Section 6.2. In the associated input area, you can set a lower threshold $c \in [0,100]$ on the connection strength required for wells to be auto-selected, so that injector number i is said to be connected to producer p if and only if

$$\boldsymbol{A}_{ip} \geq \frac{c}{100\max(m,n)} \sum_{j=1}^{n} \sum_{k=1}^{m} \boldsymbol{A}_{jk}. \tag{5.1}$$

- An input box and a slider to select a lower threshold ε that determines how much of the active influence/well-pair region is shown in the 3D plot panel. If you have selected injectors and pro-

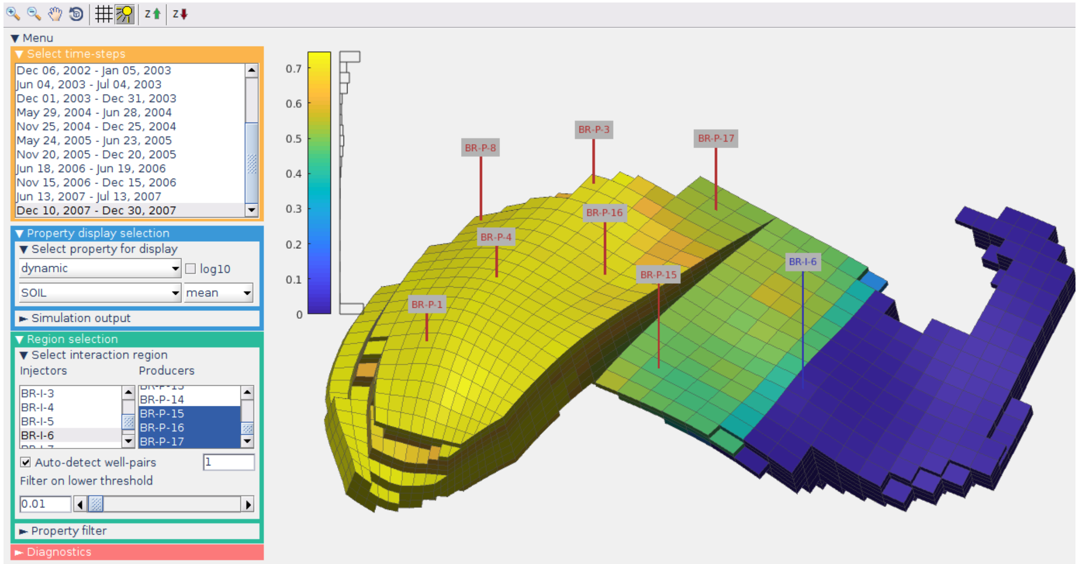

▼ Menu
▼ Select time-steps
Dec 06, 2002 - Jan 05, 2003
Jun 04, 2003 - Jul 04, 2003
Dec 01, 2003 - Dec 31, 2003
May 29, 2004 - Jun 28, 2004
Nov 25, 2004 - Dec 25, 2004
May 24, 2005 - Jun 23, 2005
Nov 20, 2005 - Dec 20, 2005
Jun 18, 2006 - Jun 19, 2006
Nov 15, 2006 - Dec 15, 2006
Jun 13, 2007 - Jul 13, 2007
Dec 10, 2007 - Dec 30, 2007
▼ Property display selection
▼ Select property for display
dynamic
□ log10
SOIL
mean
► Simulation output
▼ Region selection
▼ Select interaction region
Injectors
Producers
BR-I-3
BR-I-4
BR-I-5
BR-I-6
BR-P-14
BR-P-15
BR-P-16
BR-P-17
☑ Auto-detect well-pairs
1
Filter on lower threshold
0.01
► Property filter
► Diagnostics
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
BR-P-8
BR-P-3
BR-P-17
BR-P-16
BR-P-4
BR-P-15
BR-I-6
BR-P-1

FIGURE 5.4: Subregion of the Brugge model consisting of the communication volume between injector BR-I-6 and all connected producers. Color value shows water saturation.

ducers, the displayed region is given as all cells in which $C_i^I C_p^P > \varepsilon$, where C_j^I and C_j^P denote values of injector/producer influence region associated with well j and i, p vary over all the selected injectors and producers. If no producers are selected, p varies over all producers and likewise i varies over all injectors if no specific injector is selected.

The next three examples explain region selection in more detail:

Example:

Figure 5.4 shows one example from the Brugge model, where we have selected injector BR-I-6 and the software has automatically selected all connected producers. You can create the plot with the following steps:

1. Select a restart step (the figure shows Dec 10, 2003).
2. Toggle grid display off and light on (fifth and sixth button with the mesh/sun in the top row).
3. Select `dynamic` and `SOIL` in the "Property display" menu.
4. In the "Select interaction region" menu:
   a) Enable the "Auto-detect well-pairs" checkbox.
   b) Click on `BR-I-6` in the left list box. This will automatically also select all connected producers (shown in dark blue in the right list box).
   c) Set the lower threshold to 0.01, by manually specifying the value in the input box or by a single click on the right arrow of the slider.
5. Use the zoom, pan, and rotate buttons together with your mouse to manipulate the 3D view.

Example:

By filtering time-of-flight values upward, you can get a good idea of how far a displacement front will penetrate into the reservoir. To illustrate, we assume that water-flooding is described by quadratic Corey relative permeabilities and that the viscosity ratio between water and oil equals ten. This gives a displacement front moving approximately three times faster than a neutral fluid particle. To see approximately how far this displacement front would have travelled after 10 years of injection, we can thus apply an upper threshold of 30 years to the time-of-flight, as shown in Figure 5.5. This simple analysis is obviously not accurate, but gives you a rough estimate of how displacement fronts would move under prevailing drive mechanisms and flow conditions.

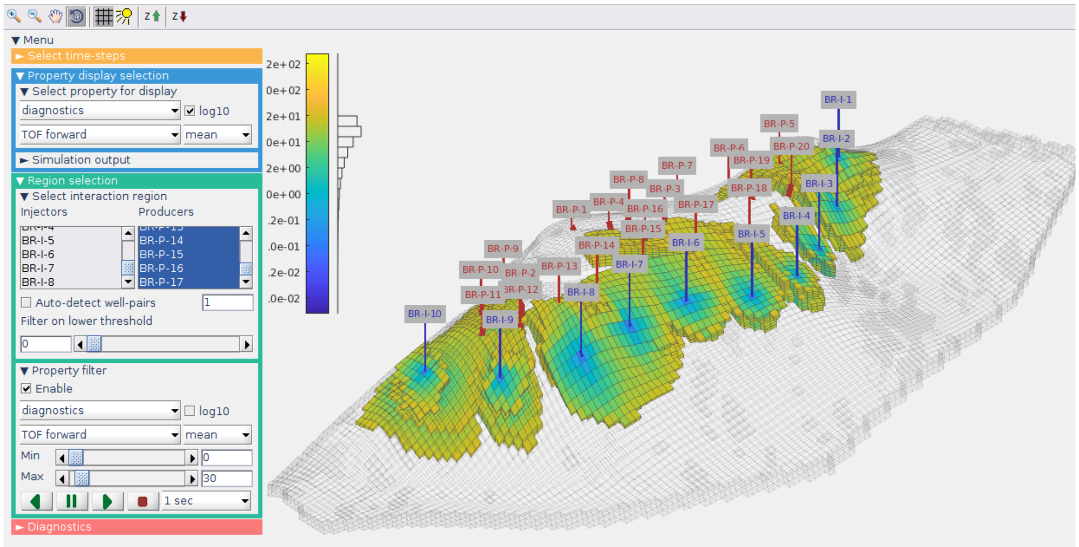

diagnostics ☑ log10

TOF forward mean

Injectors: BR-I-4, BR-I-5, BR-I-6, BR-I-7, BR-I-8

Producers: BR-P-13, BR-P-14, BR-P-15, BR-P-16, BR-P-17

☐ Auto-detect well-pairs 1

Filter on lower threshold 0

☑ Enable

diagnostics ☐ log10

TOF forward mean

Min 0

Max 30

1 sec

FIGURE 5.5: Time-of-flight bounded above by 30 years for the restart step starting on 4 June 2003. To create the plot, select all injectors and producers, and set the property selection and property filter as shown in the menus.

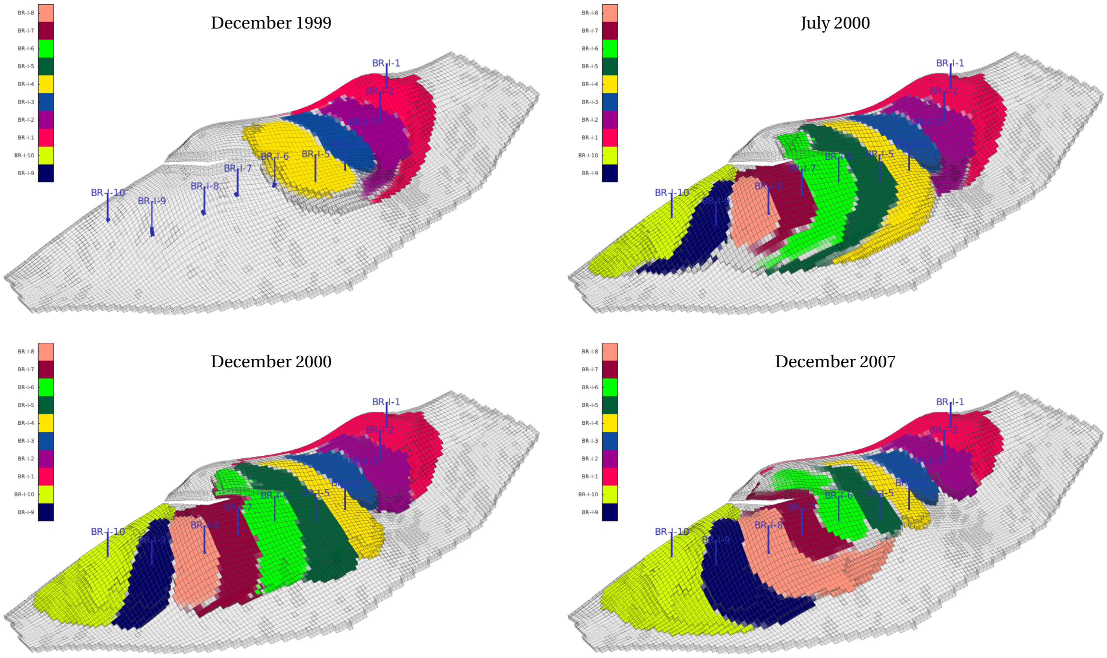

FIGURE 5.6: Sweep regions limited outward by $\tau \leq 50$ years for active injectors at four instances in time for the Brugge simulation. To create these plots, you can start from the same setting as in Figure 5.5, deselect all producers, switch the property from TOF to "sweep region", and then select the desired restart steps.

Example:

Figure 5.6 shows how the sweep region develops over time for the Brugge case. Only four injectors are active in the first report step. More injectors and producers come online in the second step (see Figure 4.4), whereas some producers have been shut down before we reach the last step.

CHAPTER

6 FLOW-DIAGNOSTICS ANALYSIS

Chapter 4.1 explains how you can display various basic cell-based flow-diagnostic quantities like time-of-flight, average residence time for all flow paths through a cell, influence regions (i.e., stationary tracer concentrations), and derived flooding/drainage volumes in the main panel. This chapter discusses all the other types of *derived* diagnostic quantities you can display in the subpanels, namely: (i) heterogeneity measures, (ii) volumetric partitions and well allocations, and (iii) residence time-distributions.

6.1 Heterogeneity measures

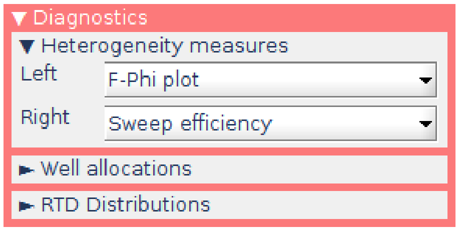

Left: F-Phi plot
Right: Sweep efficiency

FIGURE 6.1: Menu for selecting the type of heterogeneity measures to be shown in the left and right subpanels.

Section 2.4 explained how flow diagnostics uses residence time to reinterpret heterogeneity measures from classic sweep theory so that they characterize the *dynamic heterogeneity* of flow paths and their connection structure. The first entry of the diagnostics menu, shown in Figure 6.1, consists of two pull-down menus that each lets you choose between three standard measures of dynamic heterogeneity to plot in the lower-left or lower-right panel: **F-Phi diagram**, **Lorenz coefficient**, and **sweep efficiency**. (You can also choose the value `none` to clean the corresponding subpanel.) Clicking the right mouse button inside the main panel and each of the two subpanels brings up a clickable text box that lets you export the corresponding plot to a new figure window. This option is not available if you have activated the zoom, pan, or rotate buttons in the upper row.

The actual content shown in each of the subpanels depends on your selection of restart steps in the “Select time-steps” menu (Figure 3.4) and reservoir region in the “Select interaction region” menu (Figure 5.3):

- No wells selected: measures are reported for the reservoir as a whole.
- Only one well-pair: a measure is reported for the corresponding well-pair region.
- Multiple injectors/producers: a measure is reported for the whole active region.
- Only one injector: one measure for each well-pair region of all connected producers and one measure for the influence region as a whole.
- Only one producer: same as for one injector.
- Multiple restart steps: one measure or one set of measures for each step following the selection rules in the points above.

If you select many wells and/or many restart steps, the plots can be quite convoluted and difficult to interpret. We therefore recommend that you try to reduce your selection to a low number of wells and/or restart steps at a time.

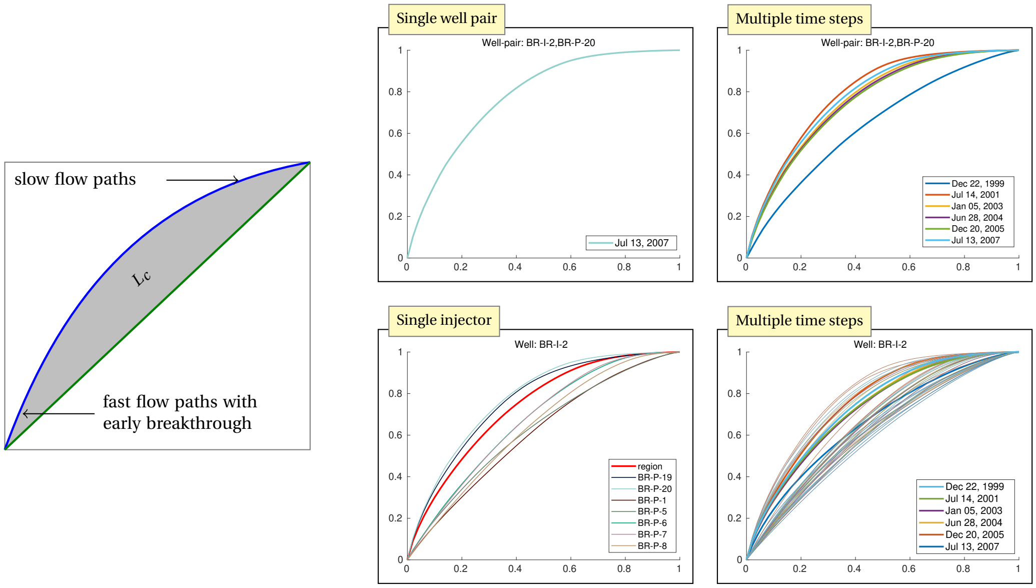

FIGURE 6.2: F-Phi diagrams for various selections of restart steps and reservoir regions.

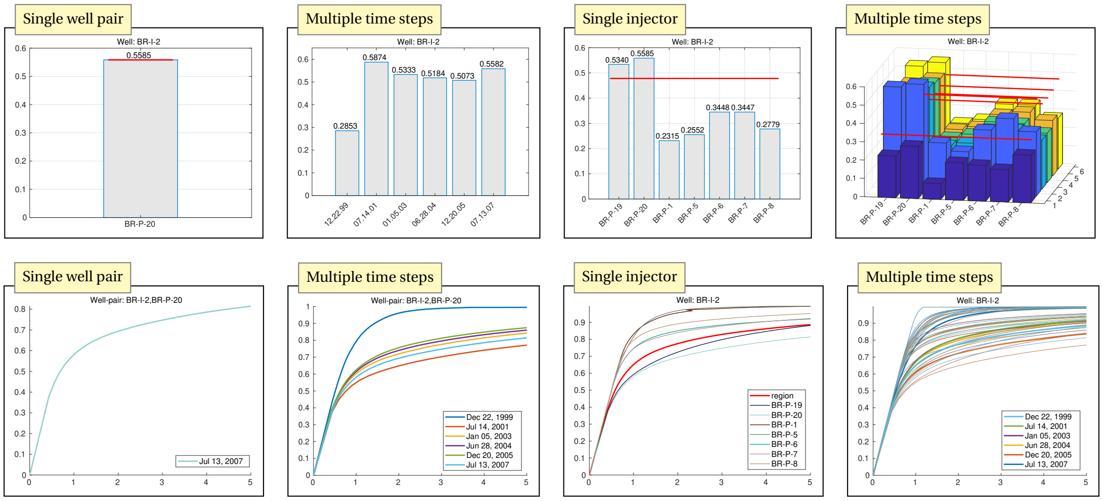

FIGURE 6.3: Lorenz coefficients and sweep efficiency diagrams corresponding to Figure 6.2.

Example:

The four boxes in Figure 6.2 give examples of how F-Phi diagrams are displayed in the two subpanels for various choices of reservoir regions and number of active restart steps. The upper-left box contains the diagram for a single well pair (BR-I-2,BR-P-20) in Jul 2007, whereas the lower-left shows the diagrams for the influence region of BR-I-2 and all the well-pairs that fall within this region. In the first restart step, only a few wells are turned on and the displacement between BR-I-2 and BR-P-20 is quite homogeneous (see Figure 5.6). As more wells become active, the well-pair region is affected and becomes more heterogeneous.

The upper-right box reports the F-Phi diagram for every third of the loaded restart steps, which corresponds to approximately every 1.5 years in the period from Dec 1999 to Jul 2007. The lower-right box contains the diagram for the full influence region of BR-I-2 (red curve) and for all the associated well pairs over the same period.

Figure 6.3 illustrates how the Lorenz coefficients and sweep efficiency diagrams corresponding to the different setups in Figure 6.2 are reported in the subpanels.

6.2 Volumetric partitions and well allocation

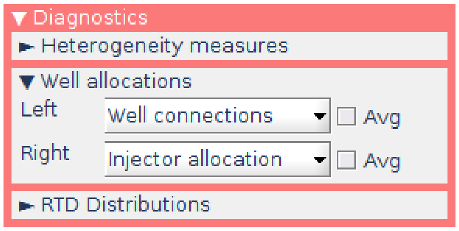

▼ Diagnostics

► Heterogeneity measures

▼ Well allocations

Left Well connections □ Avg

Right Injector allocation □ Avg

► RTD Distributions

FIGURE 6.4: Menu for selecting the type of well-allocation diagnostics to be shown in the left-/right subpanel.

The instantaneous influence regions naturally define volumetric partitions of the reservoir in the form of drainage/flooding volumes and well-pair regions. From the intersection of drainage and flooded regions, you can determine whether an injector has fluid communication with a producer and compute the corresponding flow volume. Likewise, you can measure well-allocation factors that describe how the injected flow volume of an injector will distribute to different producers, or how the produced volume from a producer can be attributed to push from different injectors.

The menu shown in Figure 6.4 consists of two drop-down menus, one for each of the two subpanels. With these, you can inspect various types of volumetric partitions and well-allocation factors. Volumetric partitions are usually associated with a single report step, but can also be reported in an averaged sense over multiple steps by activating the □ Avg check box. Well allocations can only be reported for a single producer or single injector at a time. The GUI issues a warning if the prerequisites of a particular plot are met. To produce a plot, you must then go back to the selector for wells and interaction (Figure 5.3) and update your well selection.

To explain the various quantities in this section, we start by defining some notation. Let C_i^I and C_p^P denote the stationary tracer concentration associated with injector i and producer p, respectively. Let c_k be a cell in the grid with associated porosity ϕ_i and bulk volume V_i. When a cell quantity like the volumetric flow rate q is evaluated in cell c_i, we write $q[c_i]$, like in Section 2.3.

Injector/producer volumes: The influence region of injector i is an instantaneous quantity defined as the portion of the whole reservoir in which $C_i^I > 0$. The total volume of this region is thus given as $V_i^I = \sum_k C_i^I[c_k]\,\phi_k V_k$. The menu choice "Injector volumes" subdivides this volume into flow volumes associated with individual injector–producer pairs, defined as follows

$$V_{i,p}^{IP} = \sum_k C_i^I[c_k]\,C_p^P[c_k]\,\phi_k V_k.$$

This is presented as a pie chart for all well pairs associated with a single injector or producer, see the upper row in Figure 6.5. Similar pie charts can report averaged volumes over multiple restart steps.

Injector/producer allocation: To understand the volumetric connection, it is often of interest to understand the fraction of the inflow to a given producer that can be attributed to each of the injectors. If c_k^p denotes perforated cell number k of producer p, the volumetric outflow $q(t_\ell)[c_k^p]$ during time step t_ℓ can be allocated to injector i by multiplying with the injector tracer concentrations C_i^I, that is,

$$a_{i,p}^k(t_\ell) = C_i^I[c_k^p]q(t_\ell)[c_k^p]. \tag{6.1}$$

The menu choice "Injector allocation" reports the sum over all perforations, $\sum_k a_{i,p}^k$, for all injectors connected to the selected producer. The result is presented as a pie chart if you have selected a single report step and as an area plot if you have selected multiple plots; see the second and third rows in Figure 6.5. Allocation factors are defined analogously for injectors and report how the injected volumes distribute their "push" to different producers.

Injector/producer profile: This menu choice gives cumulative sums of the injector allocation factors $a_{i,p}^k$ or producer allocation factors $a_{p,i}^k$ from toe to heel or bottom to top perforation; see the bottom two rows in Figure 6.5. This can be useful for viewing how fluid might preferentially move through different layers in the reservoir.

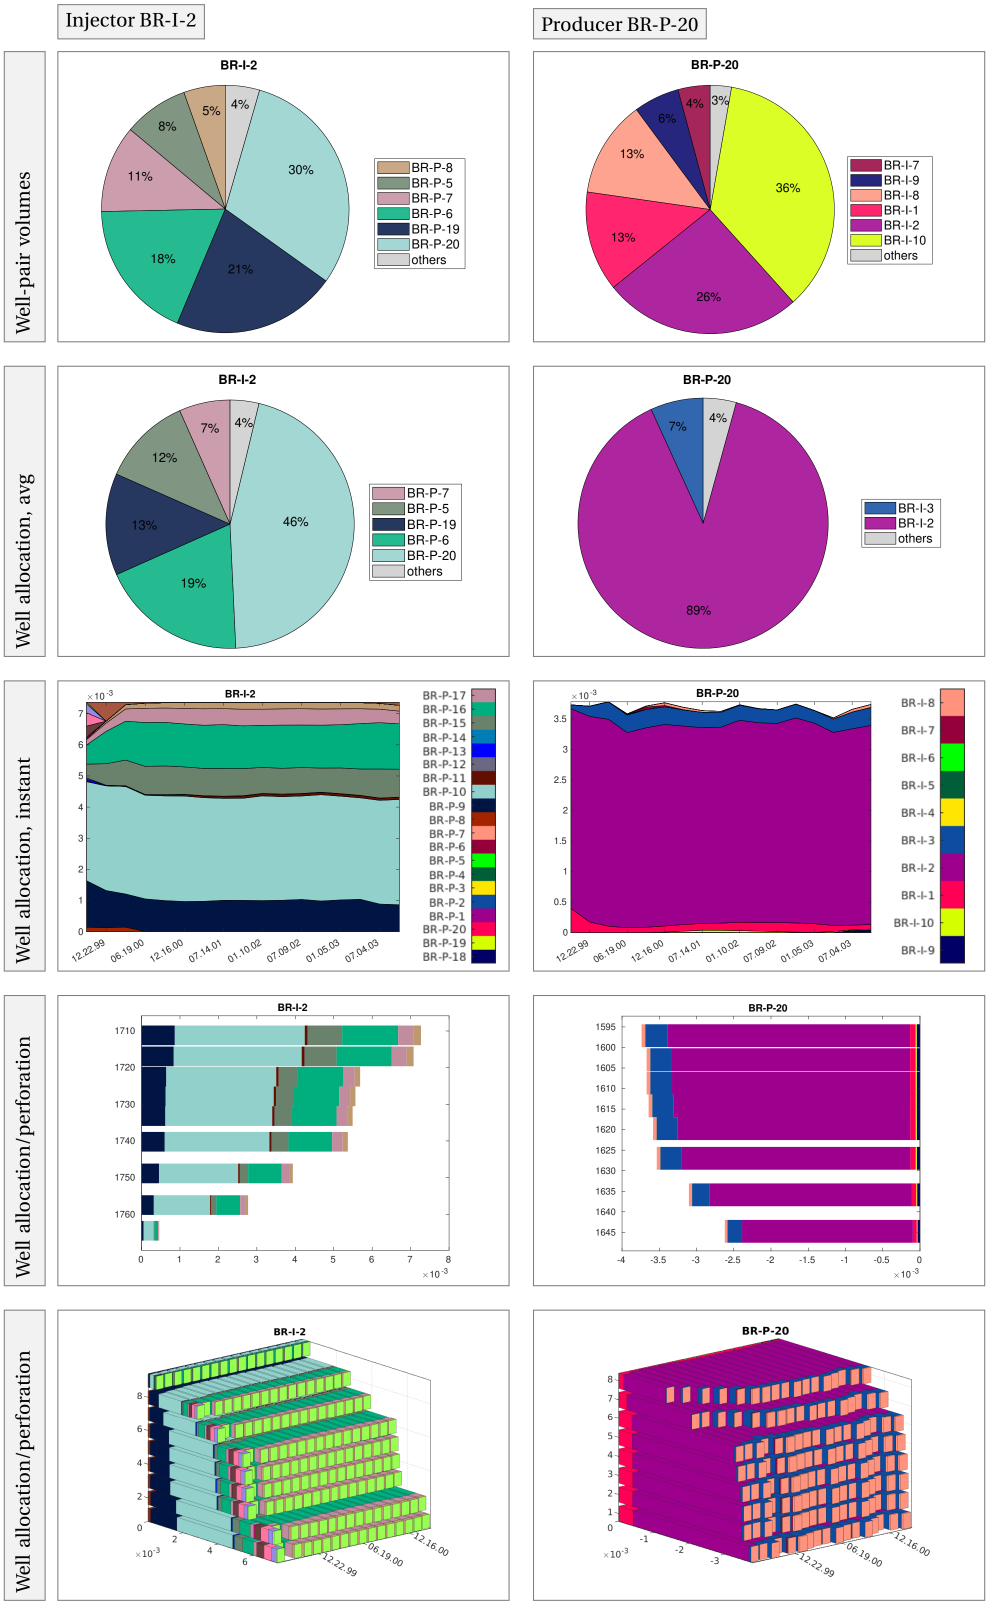

FIGURE 6.5: Volumetric partitions and well allocation for BR-I-2 and BR-P-20.

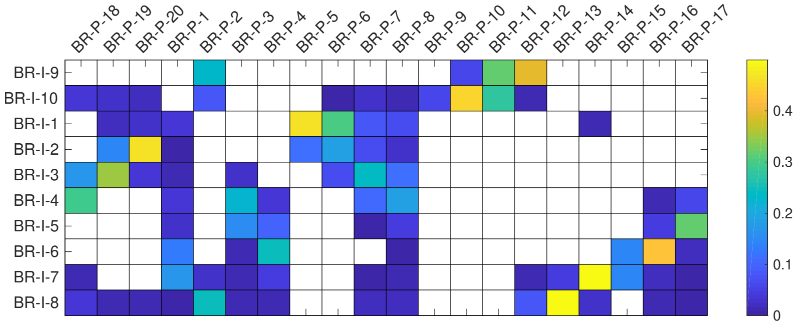

FIGURE 6.6: Plot of the well-connection strength. White color means no connection.

Well connection: To measure the communication strength between the n injectors and m producers, we introduce an $n \times m$ matrix $\boldsymbol{A}$ (see Figure 6.6) that is defined analogously to (2.4) and measures the average allocation factors over the whole simulation period. The matrix is defined as follows:

$$\boldsymbol{A}_{ip} = \frac{1}{T}\sum_{\ell=1}^{N}\sum_{k} C_p^P(t_\ell)[c_k^i]\,\big|q(t_\ell)[c_k^i]\big|\,\Delta t_\ell. \tag{6.2}$$

Here, T is the total time period covered by the report steps and t_ℓ is time at the end of step ℓ of length Δt_ℓ. The connection strength in each report step is measured by the product of the producer tracer value (influence value) C_p^P associated with producer p and the volumetric source term q, both evaluated in c_k^i, i.e., perforated cell number k of injector i. Or in words: we first sum the injector well-allocation factors (defined analogously to (6.1)) for all perforations of each well and then compute the weighted average over the loaded restart steps.

6.3 Residence-time distributions

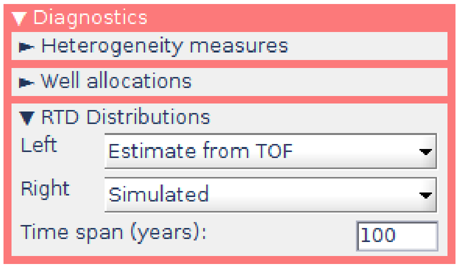
▼ Diagnostics
► Heterogeneity measures
► Well allocations
▼ RTD Distributions
Left: Estimate from TOF
Right: Simulated
Time span (years): 100

FIGURE 6.7: Menu for computing tracer distributions based on a single restart step.

The GUI can also display residence-time distributions (RTDs) for a simple and non-diffusive tracer, which we assumed to be passively advected by the total flux field. The distributions are either *estimated* from cell-averaged residence times as a crude approximation to the true RTD, or *computed* as a more accurate approximation by tracing a delta pulse through the domain as described in Section 2.5. This computation is simplified in the sense that it only considers *a single stationary velocity field*, assumes that one kilogram of tracer is injected instantaneously from each selected injector, and ignores diffusive effects. Hence, the resulting tracer should only be seen as a numerical quantity used to illuminate connections within the parts of the reservoir that contains movable fluids and should not be confused with RTDs of physical tracers, i.e., non-partitioning tracers that follow individual phases or partitioning tracers that distribute into both the aqueous and the oleic phase. (We have also developed more comprehensive solvers that account for multiple restart steps and simulate physical tracers, but these are not yet part of the public release.)

To display the RTD, you must first pick a single restart step from the "Select time-steps" menu and one injector from the "Region selection" menu. In addition, we advice to select all connected producers or a subset of these. (If no producers are selected, the GUI will try to display RTDs for all producers and the resulting figure legends can be quite messy.) Next, you set the time span for the residence distribution using the input field at the bottom of the menu shown in Figure 6.7 and then use one of the two drop-down menus to choose which quantity (estimated or computed) to display in the left and right subpanels, and an input field that gives the time-span of the simulation.

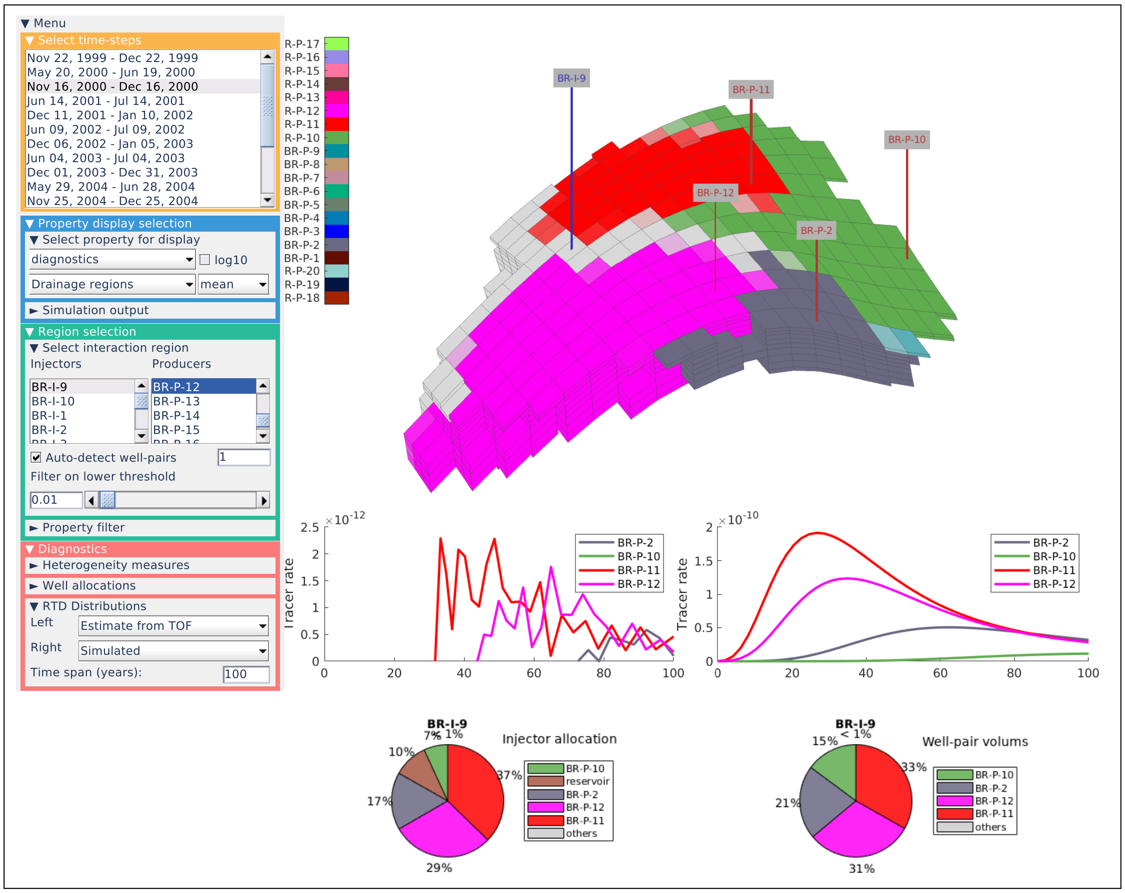

Select time-steps: Nov 16, 2000 - Dec 16, 2000

Select property for display: diagnostics, Drainage regions, mean, □ log10

Select interaction region: Injectors BR-I-9; Producers BR-P-12

☑ Auto-detect well-pairs: 1

Filter on lower threshold: 0.01

RTD Distributions: Left: Estimate from TOF; Right: Simulated

Time span (years): 100

FIGURE 6.8: Computed and estimated tracers, drainage regions, and injector allocation and volumes for all producers connected to injector BR-I-9.

Example:

Figure 6.8 shows estimated and computed residence-time distributions for all producers connected to injector BR-I-9 as given by the flux field from the restart step from Nov 16 – December 16, 2000. The fluids are moving relatively slow in this model and to show increase and decay in the residence-time distribution for at least two of the tracers, we computed tracer response over a period of 100 years (which obviously is much longer than any *physical* tracer experiment would be conducted). The computed residence time correlates well with the injector allocation: the largest fraction of the injected tracer goes to BR-P-11, which also has the largest flux allocation. The *estimated* RTDs appear quite jagged and show a much later breakthrough of tracer than the computed RTD. This should be expected since they are computed from time-of-flight values that represent the *average* over many flow paths with a wide span in residence times.

Bibliography

[1] M. A. Christie and M. J. Blunt. Tenth SPE comparative solution project: A comparison of upscaling techniques. *SPE Reservoir Eval. Eng.*, 4:308–317, 2001. doi:10.2118/72469-PA. URL `http://www.spe.org/csp/`.

[2] A. Datta-Gupta and M. J. King. *Streamline Simulation: Theory and Practice*, volume 11 of *SPE Textbook Series*. Society of Petroleum Engineers, 2007. ISBN 978-1-55563-111-6.

[3] B. Eikemo, K.-A. Lie, H. K. Dahle, and G. T. Eigestad. Discontinuous Galerkin methods for transport in advective transport in single-continuum models of fractured media. *Adv. Water Resour.*, 32(4): 493–506, 2009. doi:10.1016/j.advwatres.2008.12.010.

[4] O. Izgec, M. Sayarpour, and G. M. Shook. Maximizing volumetric sweep efficiency in waterfloods with hydrocarbon f-ϕ curves. *J. Petrol. Sci. Eng.*, 78(1):54–64, 2011. doi:10.1016/j.petrol.2011.05.003.

[5] S. Krogstad, K.-A. Lie, H. M. Nilsen, C. F. Berg, and V. Kippe. Efficient flow diagnostics proxies for polymer flooding. *Comput. Geosci.*, 21(5-6):1203–1218, 2017. doi:10.1007/s10596-017-9681-9.

[6] L. W. Lake. *Enhanced Oil Recovery*. Prentice-Hall, 1989.

[7] K.-A. Lie. *An Introduction to Reservoir Simulation Using MATLAB/GNU Octave: User Guide for the MATLAB Reservoir Simulation Toolbox (MRST)*. Cambridge University Press, 2019. ISBN 9781108492430. doi:10.1017/9781108591416.

[8] O. Møyner, S. Krogstad, and K.-A. Lie. The application of flow diagnostics for reservoir management. *SPE J.*, 20(2):306–323, 2014. doi:10.2118/171557-PA.

[9] J. R. Natvig, K.-A. Lie, B. Eikemo, and I. Berre. An efficient discontinuous Galerkin method for advective transport in porous media. *Adv. Water Resour.*, 30(12):2424–2438, 2007. doi:10.1016/j.advwatres.2007.05.015.

[10] L. Peters, R. Arts, G. Brouwer, C. Geel, S. Cullick, R. J. Lorentzen, Y. Chen, N. Dunlop, F. C. Vossepoel, and R. Xu. Results of the Brugge benchmark study for flooding optimization and history matching. *SPE Reser. Eval. Eng.*, 13(03):391–405, 2010. doi:10.2118/119094-PA.

[11] A. F. Rasmussen and K.-A. Lie. Discretization of flow diagnostics on stratigraphic and unstructured grids. In *ECMOR XIV – 14th European Conference on the Mathematics of Oil Recovery*, 2014. doi:10.3997/2214-4609.20141844.

[12] A. F. Rasmussen, T. H. Sandve, K. Bao, A. Lauser, J. Hove, B. Skaflestad, R. Klöfkorn, M. Blatt, A. B. Rustad, O. Sævareid, K.-A. Lie, and A. Thune. The Open Porous Media flow reservoir simulator. *Computers & Mathematics with Applications*, 2020. doi:10.1016/j.camwa.2020.05.014.

[13] Schlumberger. *ECLIPSE: Reference Manual*. Schlumberger, 2014.1 edition, 2014.

[14] M. Shahvali, B. Mallison, K. Wei, and H. Gross. An alternative to streamlines for flow diagnostics on structured and unstructured grids. *SPE J.*, 17(3):768–778, 2012. doi:10.2118/146446-PA.

[15] G. Shook and K. Mitchell. A robust measure of heterogeneity for ranking earth models: The F-Phi curve and dynamic Lorenz coefficient. In *SPE Annual Technical Conference and Exhibition*, 2009. doi:10.2118/124625-MS.

[16] G. M. Shook and J. H. Forsmann. Tracer interpretation using temporal moments on a spreadsheet. Technical Report INL report 05-00400, Idaho National Laboratory, 2005. doi:10.2172/910998.

[17] *Flow Diagnostics in MRST: A flow diagnostics preprocessor for model ensembles.* SINTEF Digital, Oct. 2020.

[18] W. E. Stiles. Use of permeability distribution in water flood calculations. *Journal of Petroleum Technology*, 1(01):9–13, 1949. doi:10.2118/949009-G.

[19] M. R. Thiele. Streamline simulation. In *6th International Forum on Reservoir Simulation*, Schloss Fuschl, Austria, 3–7 September 2001. URL https://www.streamsim.com/papers/rsf2001.pdf.

[20] M. R. Thiele. Streamline simulation. In *8th International Forum on Reservoir Simulation*, Stresa / Lago Maggiore, Italy, 20–24 June 2005. URL https://www.streamsim.com/papers/rsf2005.pdf.

[21] F. Watson, S. Krogstad, and K. Lie. Flow diagnostics for model ensembles. In *ECMOR XVII – 17th European Conference on the Mathematics of Oil Recovery*. EAGE, 2020. doi:10.3997/2214-4609.202035133.

Appendix

A Setting up the software

A.1 Prerequisites

The flow diagnostics GUI requires the following installed software:

- MATLAB version 2015b or newer.
- The MATLAB Reservoir Simulation Toolbox version 2017b or newer.

That said, we generally recommend that you use the most up-to-date version of MRST.

A.2 Installing MRST

MRST consists of a core part and a set of add-on modules. These are hosted as a collection of software repositories on bitbucket.org/mrst. Official releases are provided as self-contained archive files that can be downloaded from the website (mrst.no).

Assume now that you have downloaded the tarball of one of the recent releases; here, we use release 2017b as an example. Issuing the following command

```
untar mrst-2017b.tar.gz
```

in MATLAB creates a new folder `mrst-2017b` in you current working director that contains all parts of the software. Once all code has been extracted to what we henceforth refer to as the *MRST root* folder (`ROOTDIR`), you must navigate MATLAB there, either using the built-in file browser, or by using the `cd` command. Assuming that the files were extracted to the home folder, this would amount to the following on Linux/Mac OS,

```
cd /home/username/mrst-2017b/             % on Linux/Mac OS
cd C:\Users\username\mrst-2017b\          % on Windows
```

MRST is activated through a startup script, which also scans your installation and determines which modules you have installed. In the MRST root folder, the script is called by the command:

```
startup;
```

The startup script will display a welcome message showing that the software is initialized and ready to use. If you start MATLAB from within the MRST root folder, the startup script is run automatically. Alternatively, if you do not want to navigate to the root folder, you can call `startup` directly

```
run /home/username/mrst-2017b/startup    % or C:\MyPath\mrst-2017b\startup
```

For more information about the installation process and how MRST organizes its various components, you should consult [7, Appendix A].

Appendix

B Brief overview of the GUI class

The diagnostics GUI is implemented as a MATLAB class object. This object contains all the loaded and computed data as well as various containers, data, and structures used to control the visualization and describe its current state. By modifying this class, you can manipulate and control the behavior of the GUI without having to edit the source code and restart the program. Invoking the GUI with the call

```
d = PostProcessDiagnosticsECLIPSE()
```

will create and return the following MATLAB class object

```
  PostProcessDiagnostics with properties:

                      Figure: [1×1 Figure]
                      Axes3D: [1×1 Axes]
                     Axes2DL: [1×1 Axes]
                     Axes2DR: [1×1 Axes]
                    colorBar: [1×1 ColorBar]
                    colorHAx: [1×1 Axes]
                  colormap3D: ’default’
                           G: [1×1 struct]
                          Gs: [1×1 struct]
                    WellPlot: [1×1 WellPlotHandle]
                        Data: [1×1 struct]
                        Menu: [1×1 UIMenu]
                       Props: [1×1 struct]
                    Measures:
                  Allocation:
               Distributions: []
                       Patch: [1×1 CellDataPatch]
                      maxTOF: 1.5768e+10
    currentInteractionRegion: [43396×1 logical]
         currentFilterRegion: [43396×1 logical]
          currentDiagnostics: [1×10 struct]
            interactiveModes: [1×1 struct]
                    camlight: [1×1 Light]
                 outlineGrid: [1×1 Patch]
                        info: []
                layoutParams: [1×1 struct]
                   simOrigin: ’ECLIPSE’
```

Starting from the top, `Figure` is a handle to a standard MATLAB figure object; `Axes3D`, `Axes2DL`, and `Axes2DR` are handles to standard MATLAB axes objects representing the main panel and the two subpanels, resepctively; whereas `colorBar` and `colorHAx` are handles to the colorbar and the axes of the associated histogram. You can manipulate all these as you would do with any other graphical object in MATLAB, e.g., to change their position, background color, etc. As an example, you can directly control the size and position of the GUI by a call like the following

```
d=PostProcessDiagnosticsECLIPSE(); d.Figure.Position = [1324 572 1230 773];
```

Likewise, you can set the y-axis of the plot in the lower-left panel to logarithmic scale by

```
d.Axes2DL = 'logarithmic';
```

The text string `colormap3D` contains information about the current color map in the 3D plot: `'injectors'` for plots of sweep regions, `'producers'` for drainage regions, and `'default'` for all other quantities.

Continuing down the list, G is a standard MRST grid structure described in Chapter 3 of [7] (but with some extra fields), whereas the structure array Gs describes the simulation grid as a graph. The WellPlot field holds an instance of the WellPlotHandle class from the mrst-gui module. This class contains information about all the wells and how they currently are presented in a 3D axes:

```
WellPlotHandle with properties:

        injectors: [10×1 struct]
        producers: [20×1 struct]
          Visible: 'mixed'
  visibleInjectors: [10×1 logical]
  visibleProducers: [20×1 logical]
```

Here, the each element in the injectors and producers arrays contains a struct with three elements: body is a standard Line object in MATLAB that represents the well path from heel (top) to toe (bottom), label is a standard Text object representing the well label, and connector is a Line object that represents the line that connects the heel of the well and its label. As an example, the following call changes the font size for the labels and the width and color of the connector lines for all producers:

```
for i=1:numel(d.WellPlot.producers)
  d.WellPlot.producers(i).label.FontSize = 8;
  d.WellPlot.producers(i).connector.LineWidth = 3;
  d.WellPlot.producers(i).connector.Color = [.6 .6 .6];
end
```

The last two arrays of logical numbers tell whether each well should be visible in the 3D plot or not. By manipulating these arrays, you can hide a well that is selected as active in the “Region selection” menu or show wells that are not selected as active, e.g.,

```
d.WellPlot.visibleInjectors(4:7) = false;
d.WellPlot.visibleProducers(11:15) = true;
```

The Data structure array contains of a number of fields that represent static, dynamic, and diagnostics properties, the reservoir states from the restart steps, information about well communication, fluid and PVT data, and the unique colors assigned to each well. As an example, we can look at the static data field. For the Brugge field, Data.static is an array with seven element representing pore volume, permeability in the three axial directions, net-to-gross, depth, and pore volume. Looking at the first element gives us

```
>> disp(d.Data.static(1))
      name: 'PORO'
    values: [43396×1 double]
    limits: [0.0818 0.2839]
```

The Data.states field contains an array of structure arrays, one for each restart step, that contain the corresponding reservoir states: pressure, saturations, fluxes, time, mobility, shrinkage (reciprocal formation-volume) factors, etc.

```
>> disp(d.Data.states{1})
    pressure: [43396×1 double]
           s: [43396×2 double]
        flux: [121863×2 double]
     wellSol: [30×1 struct]
        time: 62208000
        date: [22 12 1999]
         mob: [43396×2 double]
           b: [43396×2 double]
```

Similarly, Data.diagnostics contains the corresponding stationary flow diagnostics, represented as two substructures. The first structure D is an extended version of the basic flow-diagnostics data object discussed in [7, §13.1.1] and contains forward and backward time-of-flight, concentrations for stationary injection/production tracers, injector/producer partitions, and time to first arrival:

```
>> disp(d.Data.diagnostics(1).D)
        inj: [10×1 double]             % Index of injectors in list of wells
       prod: [20×1 double]             % Index of producers
        tof: [43396×2 double]          % Forward and backward time-of-flight
    itracer: [43396×10 double]         % Influence regions for injectors
       itof: [43396×10 double]         % Time-of-flight for each influence region
        ifa: [43396×10 double]         % First arrival time from injectors
      ipart: [43396×1 double]          % Injector partition
    ptracer: [43396×20 double]         % Influence regions for producers
       ptof: [43396×20 double]         % Time-of-flight for each influence region
        pfa: [43396×20 double]         % First arrival time from producers
      ppart: [43396×1 double]          % Producer partition
```

The second structure WP describes all well pairs

```
>> disp(d.Data.diagnostics(1).WP)
    struct with fields:
     pairs: {1×200 cell}               % Names for well pairs, e.g., {'BR-I-1','BR-P-2'}
      vols: [1×200 double]             % Well-pair volumes
    pairIx: [200×2 double]             % Indices for each pair (injector,producer)
       inj: [1×10 struct]              % Well allocations, depth to perforation, injectors
      prod: [1×20 struct]              % Same as above, but for producers
```

We can also mention the field Data.wellComunication, which contains the matrix (6.2) measuring well connection strength.

The Menu field holds an instance of the UImenu class and contains all data structures necessary to represent the input menus to the left in Figure 3.3. Each of these is realized as a set of class objects. The Props structure contains the names of all the properties that can be displayed in the 3D panel along with limits on their range, whereas Measures and Allocations contain the names of dynamic heterogeneity measures and quantities representing volumetric and well allocation. The Distributions field is not used herein. The rest of the fields are primarily internal data structures used to represent graphical states and accelerate visualization by storing data that would otherwise have to be recomputed.