Report
R-15-17
June 2016

The BRIDGE project

Development, testing and application of a high performance computing framework for reactive transport modelling in crystalline rocks (iDP)

Jorge Molinero
Paolo Trinchero
Luis Manuel de Vries
Hedied Ebrahimi
Urban Svensson

SVENSK KÄRNBRÄNSLEHANTERING AB

SWEDISH NUCLEAR FUEL
AND WASTE MANAGEMENT CO

Box 250, SE-101 24 Stockholm
Phone +46 8 459 84 00
skb.se

SVENSK KÄRNBRÄNSLEHANTERING

ISSN 1402-3091
SKB R-15-17
ID 1512735
June 2016

The BRIDGE project

Development, testing and application of a high performance computing framework for reactive transport modelling in crystalline rocks (iDP)

Jorge Molinero, Paolo Trinchero,
Luis Manuel de Vries, Hedied Ebrahimi
Amphos 21

Urban Svensson
Computer-Aided Fluid Engineering

Executive summary

PFLOTRAN is an open source, reactive flow and transport simulator for modelling subsurface processes that has been developed with high performance computing and efficiency in mind. To this end it has been built on top of well-known frameworks like MPI, PETSc and HDF. It implements the finite volume method and has been written in Fortran 90, in a modular way. It has demonstrated peta-scale performance (i.e. more than 10^{15} floating point operations per second) in simulations of uranium migration in the superfund Hanford 300 Area as part of the SciDAC-2 programme (Hammond and Lichtner 2010, Hammond et al. 2011).

A powerful software for the simulation of groundwater flow in crystalline, fractured media is DarcyTools (Svensson and Ferry 2014). This software, owned by SKB, has been developed over the last 10 years and has been used in many different site investigation projects. Just like PFLOTRAN it is based on the finite volume method and uses MPI for its parallelization.

These two simulation codes are combined into a framework for reactive transport modelling in crystalline fractured media. This framework is denoted as DarcyTools-PFLOTRAN Interface (iDP). The interface is written in a modular, cross-platform and extensible way using Fortran and Python.

The accuracy of the reactive transport interface (iDP) has been verified by means of five benchmark problems, where the results computed using iDP are compared with those obtained with other popular geochemical codes (e.g. PHREEQC, Open GeoSys, PFLOTRAN standalone). In addition, iDP has been tested in 2 relevant case studies for SKB: (1) the degradation of cement grout applied in the Forsmark spent fuel repository and, (2) the penetration of glacial water and oxygen consumption in fractured crystalline rocks. These test cases of points 1 and 2 have been run successfully on 2 of the biggest supercomputers of Europe: the MareNostrum supercomputer at the Barcelona Supercomputing Center (BSC), and the Juqueen supercomputer at the Jülich Supercomputer Centre (JSC). The performance and scalability of iDP has been demonstrated up to 8000 cores in both supercomputers.

Sammanfattning

PFLOTRAN är ett beräkningsverktyg som har utvecklats, i öppen källkod, för simuleringar av reaktivt flöde och transport. Simulatorn är framtagen för modellering av hydrogeokemiska processer samt har utvecklats med högpresterande datorer och effektivitet i åtanke. För detta ändamål har den byggts ovanpå välkända ramverk som MPI, PETSc och HDF. Den använder sig av finita volymsmetoden och har skrivits i Fortran 90 på ett modulärt sätt. Den har uppvisat prestanda i peta-skalan (det vill säga mer än 1015 flyttalsoperationer per sekund) i simuleringar av uranmigration i området Hanford 300. Detta utfördes som en del av SciDAC-2-programmet (Hammond och Lichtner 2010, Hammond m fl. 2011).

DarcyTools (Svensson och Ferry 2014) är en kraftfull programvara för simulering av grundvattenflöde i kristallina sprickiga medier. Denna programvara, som ägs av SKB, har utvecklats under de senaste 10 åren och har använts i flera platsundersökningsprojekt. Precis som PFLOTRAN den är baserad på finita volymsmetoden och använder MPI i sin parallellisering.

Dessa två simuleringskoder kombineras till ett ramverk för reaktiv transportmodellering för kristallina sprickiga medier. Detta ramverk kallas Interface Darcytools -PFLOTRAN (iDP). Gränssnittet är skrivet på ett modulärt, plattformsoberoende och utbyggbart sätt i Fortran och Python.

Resultaten från användande av det reaktiva transportgränssnittet (iDP) har verifierats med hjälp av fem testproblem, där resultaten beräknade med hjälp av iDP jämförs med de som erhållits med andra populära geokemiska koder (t ex PHREEQC, Open GeoSys, PFLOTRAN fristående). Dessutom har iDP testats i två relevanta fallstudier för SKB: (1) degradering av cement som används för injektering i ett tänkt förvar för använt kärnbränsle i Forsmark och (2) penetration av glacialt vatten samt syreförbrukning i sprickigt kristallint berg. De två nämnda fallstudierna i punkt 1 och 2 har genomförts framgångsrikt på två av de största superdatorerna i Europa: superdatorn Marenostrum i Barcelona Supercomputing Center (BSC) samt superdatorn Juqueen vid Jülich Supercomputer Center (JSC). Beräkningsverktyget iDP har visat god effektivitet och skalbarhet för upp till 8000 kärnor i båda superdatorerna.

Contents

1 Introduction

1.1 Background and motivation

The integration of hydrogeology and geochemistry is crucial to reliably characterize and study candidate sites for the location of geological repositories for high level nuclear waste. Over the last decades, SKB has made large efforts to make possible this integration, especially during the site characterization and safety assessment of Laxemar and Forsmark. However, internal and external reviewers of the SDM and SR-Site projects (SKB 2008, 2011) have pointed out that, although this integration was achieved at a conceptual level, the numerical codes still lacked of the capabilities needed to adequately represent the mutual influence of the hydrogeological and geochemical processes.

Several reactive transport simulators have been developed during the last decades. Among the most frequently used by the scientific community are Phreeqc (Parkhurst and Appelo 2013), PHAST (Parkhurst et al. 2004), TOUGHREACT (Xu et al. 2006), CRUNCHFLOW (Steefel 2008), RETRASO (Saaltink et al. 2004), HYTEC (van der Lee et al. 2003), The Geochemist Workbench (Bethke 1994), RT3D (Clement 1997), etc. All these simulators, and others, have been extensively verified and applied and their results are documented into hundreds (probably thousands) of technical and scientific publications. However, the main limitation of all these simulators is that they were developed with a scientific perspective, and they are not optimized to be effectively used in 3D, site-specific problems. As an example, it would not be feasible to solve a typical hydro-geochemical Forsmark simulation, involving some tens thousands years, some millions of elements and hundreds of fracture zones, using one of these "conventional" scientific simulators.

A reactive transport simulator that is gaining increasing prominence among the scientific community is PFLOTRAN. A distinct feature of this code compared to the previously cited software is that it has been developed with high performance computing (HPC) in mind. PFLOTRAN is currently under active development and improvement by a group of researchers, most of them working at Sandia National Laboratory, Los Alamos National Laboratory, Pacific Northwest National Laboratory, Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory. Recently, a team of scientists (Hammond et al. 2011) has used PFLOTRAN to carry out stochastic simulations of uranium migration at the Hanford site. The numerical exercise, which included multiple stochastic reactive transport simulations, was run on the DOE's petascale Jaguar supercomputer using 40 960 processor cores and 450 000 computational hours.

PFLOTRAN is developed as open source code and is distributed under the GNU Lesser General Public License (LGPL) (FSF 2004). Amphos 21 has been working with PFLOTRAN for three years and is very active in the scientific community of developers that are continuously improving the code. The code comes with a set of standard scalability tests, which can be run to test the performance of the code on a given machine (Wylie and Geimer 2011).

Despite being a powerful code, PFLOTRAN is not specifically designed for reactive transport modelling in fractured media. In fact, it lacks of basic features such as the capability of generating and handling stochastic Discrete Fracture Networks (DFNs). On the contrary, this is one of the added values of the finite volume hydrogeological simulator DarcyTools (Svensson et al. 2010), which can generate complex DFNs, built upon hundred thousand discrete fractures, and reliably represent them onto an underlying continuum model, which is where the flow and conservative transport equations are solved. Moreover, DarcyTools can handle the complex layout of a repository and include it into the numerical calculations. A limitation of DarcyTools is that it only provides a simplified representation of geochemical reactions (e.g. using linear sorption or multi-rate mass exchange).

Given the aforementioned pros and cons (i.e. PFLOTRAN can efficiently solve complex multicomponent reactive transport problems but has limited capabilities when dealing with fractured media whereas, DarcyTools is specifically designed for the accurate simulation of flow and transport processes in fractured media but provides a simplified representation of geochemical reactions), it is evident that an interface between the two codes would benefit from their respective strengths.

It is with this motivation in mind that we have developed the interface between DarcyTools and PFLOTRAN (iDP), which provides a novel innovative numerical framework for the simulation of multicomponent reactive transport problems in fractured crystalline media.

The main added values of iDP are *efficiency* and *reliability*.

In iDP, the reactive transport problem at hand is solved using PFLOTRAN. This makes the tool particularly *efficient*, as the calculation can be carried out in powerful supercomputers taking profit of the good code scalability (Hammond et al. 2014). Moreover, iDP can be used to perform specific calculations in spatially limited volumes (e.g. the formation of a hyper-alkaline plume close to an access tunnel, the evolution of geochemical conditions close to the near-field, etc.), which means that the groundwater flow calculation need to be run only once and the results can be re-used in each of the local-scale reactive transport models.

The *reliability* of iDP is closely related to its efficiency. As the number of degrees of freedom of a model is no longer a strong limitation, the different processes and features can be represented in their full complexity, thus avoiding over-simplifications. This is the case for instance of a recent work (Trinchero et al. 2016), where iDP has been used to provide a detailed description of mineralogical heterogeneity and to show how coarse models are inadequate to reliably predict cesium transport and retention in these complex heterogeneous media.

1.2 Objective and scope

This report focuses on three main points, all related to iDP:

1. *Implementation details*. After discussing briefing the main processes and features of iDP (Chapter 2), a detailed account of the interface architecture is presented (Chapter 3). This chapter is intended to provide the user with enough details to understand the structure of the interface and the related modelling assumptions.
2. *Verification of the interface.* The correct implementation of the interface is verified using a number of benchmark problems, all taken from the literature (Chapter 4). The verification of iDP is crucial to build confidence in the reliability and robustness of the tool.
3. *Large-scale applications.* Chapter 5 and Chapter 6 demonstrate how iDP can be used to carry out large reactive transport calculations. The two simulation cases address two “typical” problems related to the construction and safety analysis of a repository for spent nuclear fuel: the formation of a hyper-alkaline plume as a consequence of the grout applied during tunnel construction and the migration and retention of cesium in a dual porosity system.

2 Theoretical background

The main tool available for performing quantitative integration of groundwater flow and geochemical processes is the so called “reactive transport modelling” (RTM). RTM is based on coupling the groundwater flow equation, the advection-dispersion transport equation and the set of algebraic equations resulting from a given hydro-bio-geochemical system. Groundwater flow and solute transport equations are continuous partial differential equations than can be solved numerically by several methods (being finite elements, finite differences and finite volumes the most used). Usually they are linear equations, except in the case of unsaturated flow and/or density-driven flow where nonlinearities appear. On the other hand, hydro-bio-geochemical systems can be solved assuming both, partial equilibrium approach (mass action law) or kinetic approach. Whatever the approach (or the combination of approaches) being used, the hydro-bio-geochemical system is formulated as a set of highly non-linear algebraic equations that must be solved by means of iterative numerical schemes.

In the iDP framework, the groundwater flow equation is solved by DarcyTools using the Finite Volume Method (FVM). A short description of DarcyTools can be found in the User’s Guide (Svensson et al. 2010):

“DarcyTools computes fracture network flows using a continuum model in which the mass conservation equation is associated to several mass fraction transport equations for the salinity and/or particle mass concentrations, and to a heat transport equation.”

For more information about the implementation details of DarcyTools, the reader is referred to the DarcyTools User’s Manual (Svensson et al. 2010).

In the iDP framework, reactive transport is solved by PFLOTRAN. Just as with DarcyTools, the Finite Volume Method is used as numerical scheme. A short description of PFLOTRAN can be found in the User Manual (www.pflotran.org):

“PFLOTRAN solves a system of generally nonlinear partial differential equations describing multiphase, multicomponent and multi-scale reactive flow and transport in porous materials. The code is designed to run on massively parallel computing architectures as well as workstations and laptops (e.g. Hammond et al. 2011). Parallelization is achieved through domain decomposition using the PETSc (Portable Extensible Toolkit for Scientific Computation) libraries for the parallelization framework (Balay et al. 1997).”

For more details on how PFLOTRAN implements reactive transport, the reader is referred to the PFLOTRAN User Manual, specifically the appendix with the governing equations (www.pflotran.org).

DarcyTools can use unstructured Cartesian grids (see Figure 2-1), where a side can be shared by 3 different volumes. So far, this type of unstructured mesh is not compatible with the unstructured mesh implemented in PFLOTRAN. For this reason, the current version of iDP assumes a structured grid topology. In DarcyTools, the unstructured grid can still be used as long as the sub-domain that is used in PFLOTRAN (for the reactive transport calculation) is structured. Within this sub-domain, the exact same mesh is used in DarcyTools and in PFLOTRAN. Since both codes are based on the Finite Volume Method, where the Darcy velocity is evaluated at the volume sides, the entire sub-domain velocity field, calculated in DarcyTools, can be imported in PFLOTRAN. In PFLOTRAN, this velocity field is then used to perform the reactive transport simulation.

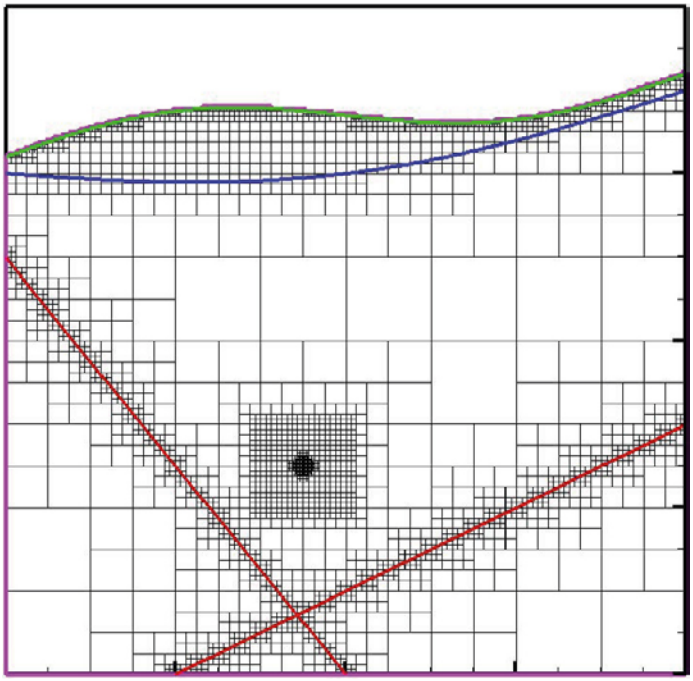

Figure 2-1. *Example of unstructured cartesian grid (from DarcyTools User's Guide).*

3 Implementation details

3.1 Application architecture

The full reactive transport framework consists of the following components (see Figure 3-1 and Figure 3-2):

- **DarcyTools**:
  Simulation of flow
- **iDP DarcyTools (IDP-DT) and iDP PFLOTRAN (iDP-PF)**:
  Reading a sub-domain[1] from the Darcy Tools output and creation of PFLOTRAN input.
- **PFLOTRAN**:
  Simulation of Reactive Transport.
- **iDP Post-processing (iDP-Post)**:
  Post-processing the output of PFLOTRAN for visualization purposes.
- **Visualization**:
  using ParaView to display the results.

The pre-processing done by iDP consists of 2 components:

- iDP-DT (Figure 3-3): This application is written in Fortran and links to the DarcyTools libraries to process the binary output format of DarcyTools. DarcyTools stores its output in so called ROF files and information about the mesh and markers in xyz files. The interface converts these files into a new set of binary files, which contain information of a sub-domain that can be read by iDP-PF. Note that iDP always works with sub-domains that must be at least one cell smaller in all the six directions (x^-,x^+,y^-,y^+,z^-,z^+) than the DarcyTools domain.
- iDP-PF (Figure 3-4): This application is written in Python 3 and reads the binary files produced by iDP-DT. It then generates HDF5 files (Folk et al. 1999) that can be read in parallel by PFLOTRAN. In addition, it generates the input file for the PFLOTRAN simulation. The user selects which geochemistry, initial conditions and boundary conditions to apply to the respective model. In addition, it has the following features:
  - The sub-domain can be divided into different horizontal layers. These can be assigned different initial and boundary conditions.
  - The cells that are marked in DarcyTools with a specific marker (or range of markers) can be imported as separate regions (with their own material and chemical conditions).
  - The subset of the previously mentioned marked cells can be imported as a separate region.
  - Marked cells below a certain depth can be ignored.
  - Random regions can be read in from a data file (produced by the free geostatistics program SGeMS (Remy et al. 2009)). This option is used to generate stochastic distributions of mineral abundance.

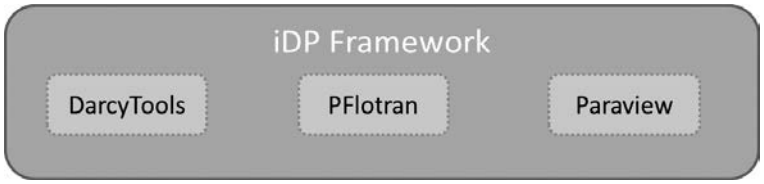

Figure 3-1. *Overall iDP Architecture.*

[1] We refer to "sub-domain" as a physical volume that is part of the whole DarcyTools domain whereas we use the term "region" in a way consistent with the PFLOTRAN notation. In PFLOTRAN, a region denotes a point, surface or volume within the problem domain.

Figure 3-2. *Individual iDP Components.*

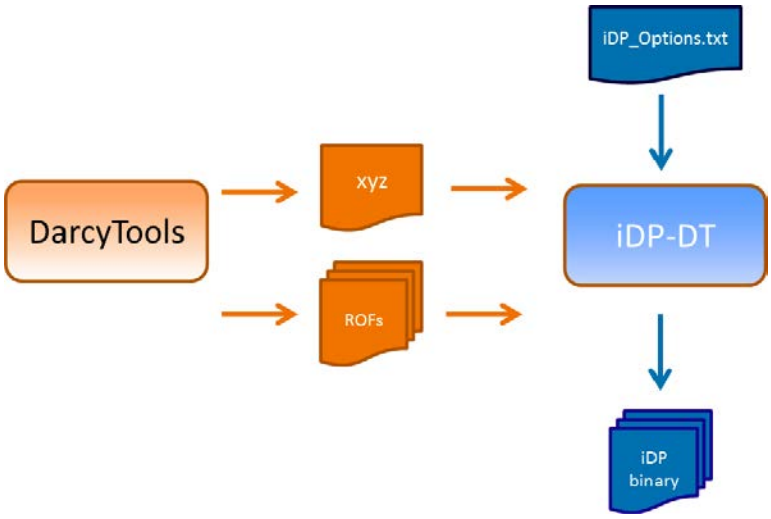

Figure 3-3. *iDP-DT Overview.*

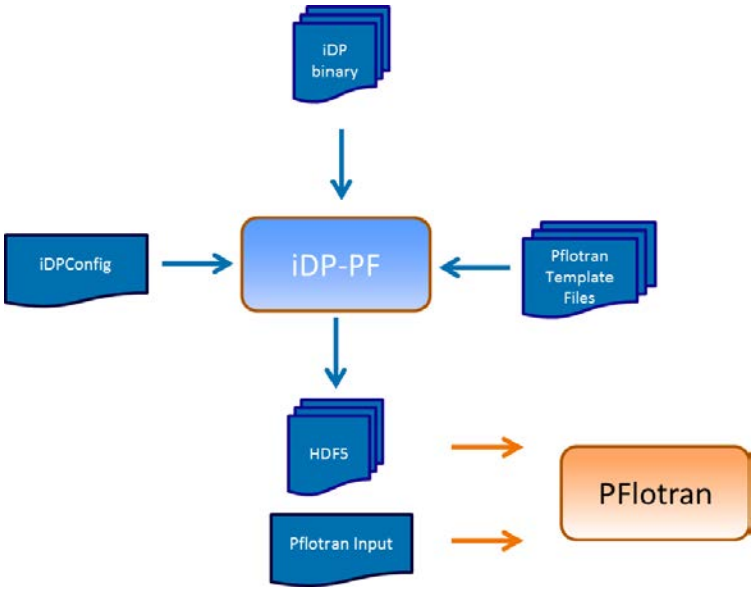

Figure 3-4. *iDP-PF Overview.*

The configuration file "iDP_Options.txt" specifies the sub-domain to be obtained from the DarcyTools output files. It also determines which markers should be used to extract regions. These regions can be assigned different materials, initial conditions or boundary conditions in the PFLOTRAN simulation. All output is written to binary files that can be easily read by PFLOTRAN.

The configuration file "iDPConfig.py" specifies the location of the input files, produced by iDP-DT. It also specifies all the details of the reactive transport simulation setup, like the exact domain, which geochemistry to apply and the simulation times.

The complete reactive transport framework is depicted in Figure 3-5.

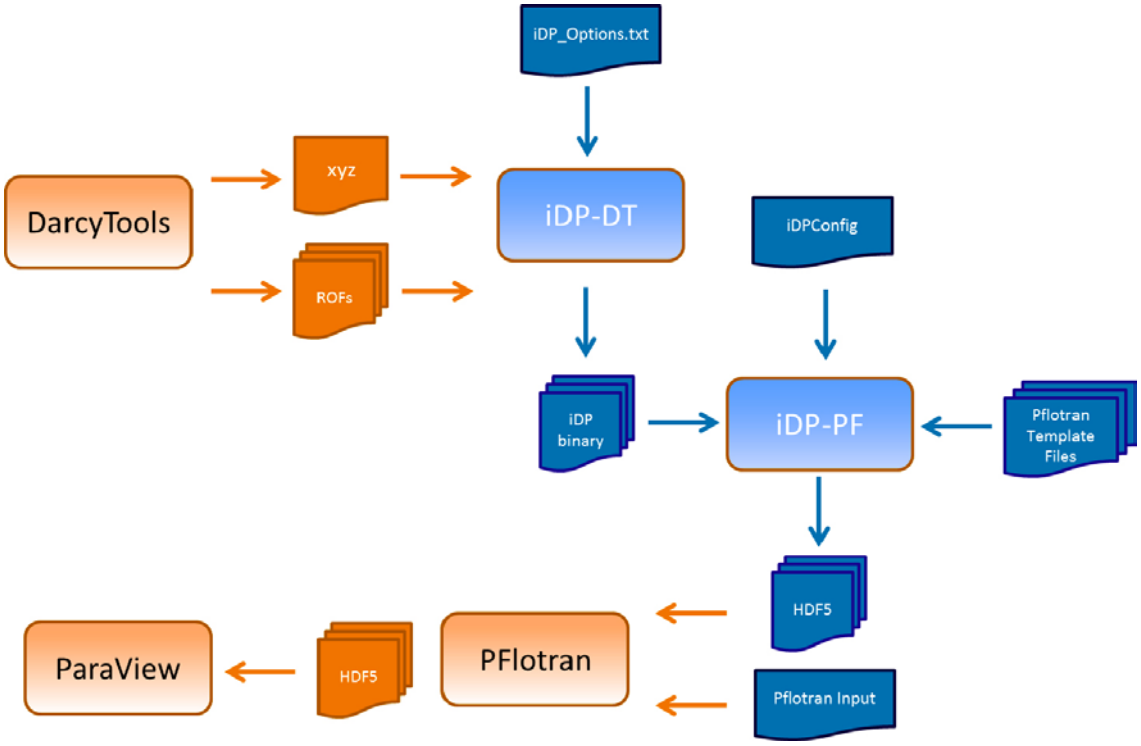

Figure 3-5. *Detailed overall iDP Architecture.*

3.2 iDP-DT functionality specifications

3.2.1 Overview

The iDP-DT application contains the following functionalities:

- Process configuration file.
- Generate PFLOTRAN sub-domain.
- Process DarcyTools binary files.
- Generate relative cell balance errors.
- Process DarcyTools markers.
- Process grouting cells.

3.2.2 Process configuration file

The configuration of iDP_DT is read from an input file. Note that the order of the lines is important. Lines that start with a '!' symbol are considered comment lines and are ignored. The variables are described in Table 3-1. Note also that the current version of iDP allows the flow velocity field to be updated (i.e. pseudo-transient flow conditions).

3.2.3 Generate PFLOTRAN sub-domain

The coordinate limits given in the configuration file are used to extract information about the new sub-domain from the Darcy Tools mesh file (normally named "xyz"). The assumption is that the volume is located in an area where Darcy Tools uses a structured grid. The same grid will be constructed for use in PFLOTRAN. In this way the velocity and permeability field values can be used directly without interpolation in PFLOTRAN.

The information extracted about the sub-domain is written as one record to a binary file using the format specified in Table 3-2. Note that real values of 8 bytes can be used by recompiling iDP_DT against the Darcy Tools R8 libraries, instead of the R4 libraries.

Table 3-1. Format of the iDP_DT configuration text file.

Variable Name	Description
Output Folder	Location where the binary output files will be written.
Number of velocity fields	Number of different velocity fields provided by Darcy ToolsEach component (x,y or z) can be read from a different DT binary output file (ROF).
Full path to x velocity File	ROF file for the x component of the darcy velocities (m/s).
Full path to y velocity File	ROF file for the y component of the darcy velocities (m/s).
Full path to z velocity File	ROF file for the z component of the darcy velocities (m/s).
Velocity ROF Time Step	The time step containing the velocities to be extracted.
PFLOTRAN Velocity ID	The ID that identifies the velocity field within PFLOTRAN.
Permeability-x ROF File	ROF file for the x component of the permeability field (m^2).
Permeability-y ROF File	ROF file for the x component of the permeability field (m^2).
Permeability-z ROF File	ROF file for the x component of the permeability field (m^2).
Permeability ROF Time Step	The time step from which to extract the permeabilities.
xyz mesh file	Darcy Tools mesh/grid file. This file contains information about the mesh, connectivity and the location of the markers.
PFLOTRAN sub-domain	These 6 lines define the sub-domain to be extracted from the Darcy Tools simulation domain.
minx	
minY	
minZ	
maxX	
maxY	
maxZ	
Number of markers	For each marker the following information is read:
Marker start number	Start of marker range to consider (inclusive).
Marker end number	End of marker range to consider (inclusive).
New Marker ID	All the cells marked in Darcy Tools that are within the range are assigned this new ID.
New marker ID for grout	All the cells that are within the range that have the grout filter applied, will get this new ID.
Grout filter value	if the tunnel side has a permeability larger than this threshold value, grouting will be applied, and the cell will be assigned the new marker ID for grouting.

Table 3-2. Binary format for grid information file.

Variable Name	Description	Format
nx	Number of cells in x direction	Real*4
ny	Number of cells in y direction	Real*4
nz	Number of cells in z direction	Real*4
minX	Lower x coordinate of the sub-domain	Real*4
minY	Lower y coordinate of the sub-domain	Real*4
minZ	Lower z coordinate of the sub-domain	Real*4
deltaX	Delta x for all cells	Real*4
deltaY	Delta y for all cells	Real*4
deltaZ	Delta z for all cells	Real*4

3.2.4 Process DarcyTools binary files

Process DarcyTools binary files (“ROF”) to obtain datasets (velocities, permeability and porosity) within the PFLOTRAN sub-domain. All the obtained datasets are written as binary files that can be read in effortlessly from the Python NumPy library.

3.2.5 Generate relative cell balance errors

To ensure that the DarcyTools flow simulation converged sufficiently to provide a velocity field that can be used to perform reactive transport simulations, the relative balance errors (cell mass balance residual/largest absolute side flow) for each cell are written to dataset.

3.2.6 Process DarcyTools markers

In the Darcy Tools input file (“cif.xml”) a unique marker ID can be assigned to each cell. A range of markers is often used to group a specific type of objects (e.g. markers 3 to 21 are access tunnels). iDP provides the possibility to select such a range and assign it to one PFLOTRAN region with a specific material type.

If one wants to apply a boundary condition to the interface of the marker cells and the other cells, they can be applied to the neighbouring cells. An example: Marker 2 defines open tunnel cells. The cells next to the open tunnel cells can be used to define the boundary conditions in PFLOTRAN (e.g. if the tunnel cells are deactivated and one wants to apply a boundary condition to the sides of the tunnels). To allow for this option all the neighbouring cells of the marker cells are exported with their respective neighbouring sides.

3.2.7 Process grouting cells

Since the marker IDs are often used to indicate tunnel walls, iDP implements an extra filter to apply an upper permeability threshold. If the permeability of a specific marker ID cell side is higher than the given threshold, the cell is marked with a new marker ID, a given grouting ID.

3.3 iDP-PF functionality specifications

3.3.1 Overview

The iDP-PF application contains the following functionalities:

- Process configuration file.
- Create layered regions.
- Convert random fields to regions.
- Marker functionality.
- Material assignment.
- Process velocities.
- Process porosity and permeability.
- Process mass balance errors.
- Generate PFLOTRAN input file.

3.3.2 Process configuration file

The configuration of iDP-PF is read from an input file. The input file contains a python data structure, called dictionary. A dictionary can contain several keys, where each key will have a value. In addition the values can be dictionaries as well. The definition of a key is case-sensitive. The main parts in the configuration file are described in Table 3-3.

Table 3-3. Main items defined in the configuration file.

Variable name	Description
logFile	The full path to and name of the log file. Note that instead of "\", the "/" needs to be used
Simulations	This dictionary defines a sequence of PFLOTRAN simulations to generate. For each simulation a PFLOTRAN input file will be written. • simulationID: the ID defined for a velocity field in the iDP-DT configuration • simulationTime: total time to simulation • wallclockStop: the simulation will be stopped after this time, even if it is not finished. • Restart (optional): A list with the filename of the checkpoint file and the time to restart from.
PflotranSettings	This dictionary defines the setup of the simulation in PFLOTRAN. Among the most important are: • Time steps • Solver settings • Chemistry Setup • Material Properties
RealType	Specifies if the binary files are written in Real*4 or Real*8 format
FileNames	The output file names for the HDF5 files and the PFLOTRAN configuration file
Regions	The sub-domain can be divided into several regions defined by depth. In this section the different regions and their properties are specified.

3.3.3 Handling of regions

The configuration for each region contains the values specified in Table 3-4.

Table 3-4. Region configuration.

Variable name	Description
ID	The name of the region.
markerID	If defined, the region contains all the cells marked by this ID.
Side	If the region is defined in terms of sides, these values define which side (1–6). If the region is defined by cells this is set to 0.
materialID	The ID of the material assigned to this region.
isExternalRegion	If the region includes at least one of the 6 sides of the whole PFLOTRAN domain, then the region is considered an external region and this value is set to "True". If it defines an internal region (like a marker ID region), this value is set to "False".
minDepth, maxDepth	If different geochemistry are assigned to different depths, these values define the depth range. If set to "None", the whole sub-domain is treated as a single domain.
Geochemical initial conditions	The geochemical initial conditions for the reactive transport simulation in PFLOTRAN.
Geochemical Boundary conditions	The geochemical boundary conditions for the reactive transport simulation in PFLOTRAN.
writeOrder	The order in which the regions and the initial conditions are written to the PFLOTRAN configuration file. This can be important if cells belong to two different regions (e.g. domain cells and marker cells).

3.3.4 Marker functionality

If a region is defined by a marker ID, the cell IDs are read from the binary file:

"markedCells_*markerID*_CellIDs.bin". If a boundary condition is defined for these cells, the cells and their side are read from "markedCells_*MarkerID*_BorderCellIDs.bin" and "markedCells_*MarkerID*_BorderCellSides.bin". Each side is read in as a separate region (region name with the extension: "_west", "_east", "_north", "_south", "_top", "_bottom"). In this way different boundary conditions can be assigned to each side.

3.3.5 Material assignment

The materials are defined by the variables described in Table 3-5.

Table 3-5. Material configuration.

Variable Name	Description
ID	The material ID. To assign a material to a region, this ID must be referenced in the region
Tortuosity	Tortuosity to be used in the simulation
PermX	Permeability to be used in the simulation
PermY	Permeability to be used in the simulation
PermZ	Permeability to be used in the simulation

3.3.6 Process velocities

For PFLOTRAN to be able to import velocities, the velocity HDF5 file should contain three arrays (for the three velocity components) that hold values for each cell (upwind direction). For each boundary condition another full array with velocity values is required. This means that most values will be zero, only the boundary cells will have a non-zero value.

3.3.7 Process porosity and permeability

The porosity HDF5 file contains a value for each cell.

The permeability file is only written for post-processing purposes, since the velocity field is already calculated.

3.3.8 Process mass balance errors

The field of relative mass balance errors is written in PFLOTRAN HDF5 output format so that it can be opened and analysed directly in ParaView (Henderson et al. 2004). Mass balance errors are computed by iDP-PF based on the information (i.e. Darcy fluxes and grid properties) previously written by iDP-DT.

3.3.9 Generate PFLOTRAN input file

iDP generates a PFLOTRAN input file that automatically links to all the HDF5 files that were generated. In addition it will apply the correct initial and boundary conditions to the regions specified in the configuration file. All PFLOTRAN settings can be modified through the iDP configuration file.

If the option for pseudo-transient flow conditions (i.e. multiple velocity fields) is used, multiple PFLOTRAN inputs, with the proper restarting options, are created.

3.4 Post-processing

3.4.1 Introduction

Due to large file sizes that result from “big” PFLOTRAN simulations (total size can be in the order of terabytes), inspecting the output files interactively with ParaView is not always feasible. For this reason and to ensure consistency in the graphics generated, a set of python scripts has been written to post-process the PFLOTRAN output. These scripts take care of the following:

- Merge different time output time step files into one HDF5 file.
- Split the full HDF5 into separate files, one for each type of data (i.e. specific species or minerals).
- Merge one or more species with the coordinates group.
- Calculate basic statistics (min/max/mean) for each time step for all the aqueous species and minerals.

3.4.2 PFLOTRAN HDF5 Data format

When referring to HDF5 in this document, what is meant is the data format that PFLOTRAN uses to store its output data. PFLOTRAN supports different types of output, but for running large models in parallel, HDF5 is the most flexible and efficient choice.

The structure of HDF5 files can be inspected using the application HDFView (www.hdfgroup.org). ParaView has a specific reader plugin for reading PFLOTRAN HDF5 files.

The structure of a PFLOTRAN HDF5 file is as shown in Figure 3-6 (left). The file should contain a Coordinates group that defines the structured grid and each output time step is stored in a separate group. Each time step group consists of many 3D datasets, one for each species in Figure 3-6 (right). PFLOTRAN outputs the concentration for the aqueous species, and the saturation index (SI), volume fraction (VF) and reaction rate for the minerals. In addition it can output pH, pe and Eh.

On large runs PFLOTRAN provides the (safer) option to write each time step to a separate HDF5 file (each with a coordinates group). This file can be viewed directly in ParaView, but to see the time evolution of a specific species, all time steps for that species need to be in the same file.

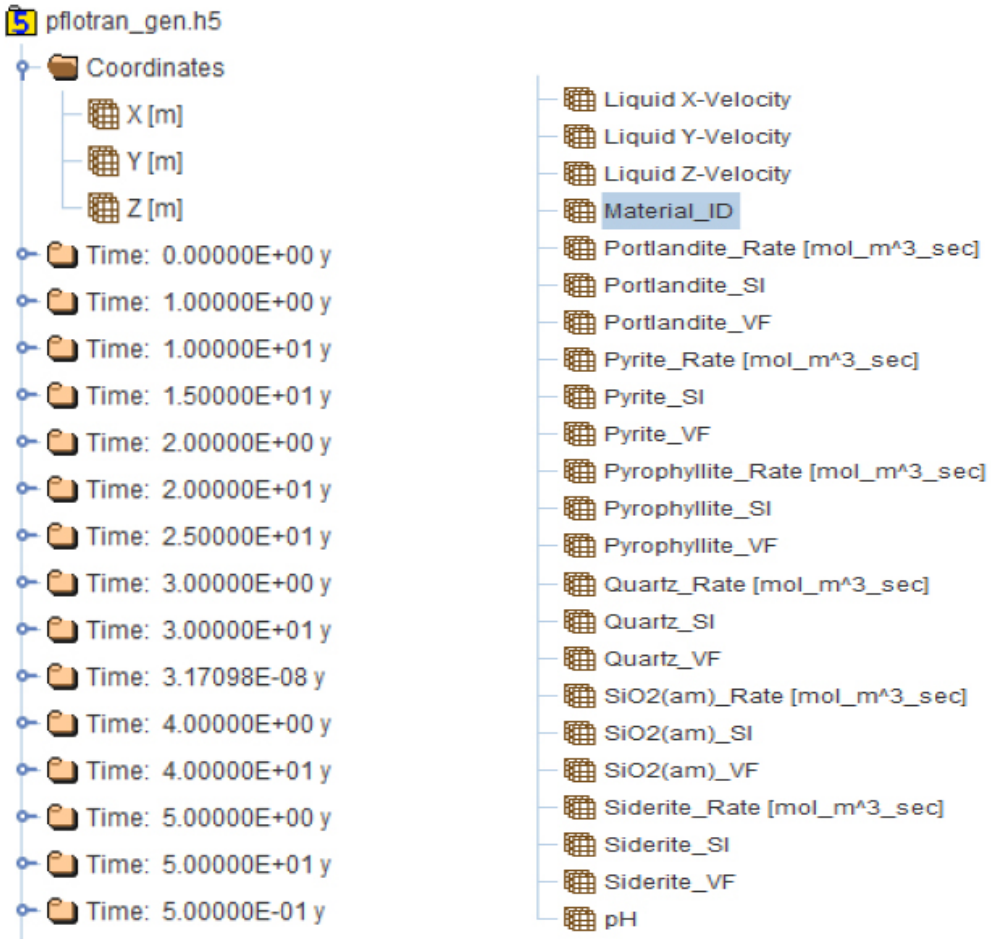

Figure 3-6. *Structure of PFLOTRAN HDF5 file. Left: The file contains one group for the coordinates. Each time step is written into a separate group. Right: Each time step contains the 3D datasets for all species & minerals*

3.4.3 Customizing ParaView Functionality

In order to speed up the post-processing with ParaView one can save a so called “State File”. The state file contains the information about all the filters to be applied, called “pipeline” in ParaView, and about which view(s) to use. When loading a state file, one can use a different data file as source. In this way it ensures a more consistent look for different ParaView results.

The state files can be saved as ParaView State File, with extension “.pvsm”, or as a python script. The python script gives the flexibility of customization. One could for example modify the python state file easily to save an image for each time step.

3.4.4 Commonly used ParaView filters

The filters most commonly used for looking at PFLOTRAN results are threshold and contour. These will allow the user to visualize for example the evolution of a specific plume.

3.4.5 Camera path and settings

A comprehensible way to present simulation results is to create animations from the post-processed 3D-Visualized results of a simulation. To do this, the following steps need to be performed in ParaView:

- Set up the visualization pipeline to combine and filter data.
- Set the position and orientation of the camera.
- Create the animation path.
- Set the animation timing, frame rate and resolution.

4 iDP Benchmarks

iDP has been tested against five different benchmark problems, which are well documented in the literature. In each of these exercises, the problem is solved using iDP, which means that (I) the groundwater flow calculation is carried out using DarcyTools, (II) iDP is used to extract the information, generate the input datasets (i.e. hdf5 files) and the PFLOTRAN input deck and (III) PFLOTRAN is used to run the calculation. The results obtained using iDP are then compared with independent calculations/solutions. A list of the different benchmark cases, their geometry (i.e. 1D, 2D or 3D) and the independent solutions used for the validation is listed in Table 4-1. The overall objective of these exercises is to verify that the whole iDP modelling procedure (i.e. extraction and creation of datasets, generation of PFLOTRAN input deck, etc.) is done correctly.

Table 4-1. List of benchmark cases with related independent solutions.

Name of the benchmark	Geometry	Independent solution
1D calcite	1D	PFLOTRAN standalone
Cation exchange	1D	PHREEQC
Calcite dissolution	1D	GeoSys-GEM
Copper leaching	3D	PFLOTRAN standalone
Matrix diffusion and oxygen penetration	2D	analytical solution

4.1 1D Calcite

This verification test is based on one of the reference examples included in the PFLOTRAN source code directory on the bitbucket PFLOTRAN repository. The results obtained with the default PFLOTRAN input file have been compared to those computed using iDP.

The test case consists of a one-dimensional domain that is initially filled with water in equilibrium with calcite. The mineral is homogeneously distributed with initial volume fraction of 10^{-5}. Acidic water (pH=5) enters the domain from the inlet boundary. The ingress of the acidic front is temporarily buffered by the dissolution of calcite. Once the mineral is completely depleted there are no additional pH buffers and thus the acidic front penetrates deep along the domain. A conservative tracer, with unitary concentration, is also added to the boundary water.

The parameters of the model are summarized in Table 4-2 and the initial and boundary waters are listed in Table 4-3. The PFLOTRAN simulations have been carried out:

1. Using the standard PFLOTRAN input file for the 1D calcite dissolution case. In this case, a constant uniform velocity is prescribed (see Table 4-2).
2. Using iDP. In this case, first a DarcyTools model, with the same physical parameters specified in Table 4-2, has been set-up and solved. iDP has then been used to generate the PFLOTRAN input file and the external velocity file, which is read by PFLOTRAN.

Breakthrough curves for the conservative tracer, pH and calcium at the outlet of the domain are shown in Figure 4-1 to Figure 4-3. A profile showing the distribution of calcite (i.e. calcite volume fraction) after 30 years from the beginning of the simulation is shown in Figure 4-4. For all the considered species, the match between the two solutions (i.e. PFLOTRAN vs. iDP) is very good and the verification exercise is considered successful.

Table 4-2. Parameters used in the problem.

Parameter	Value
Domain length (L)	100.0 m
Darcy Velocity (q_x)	1.0 m/y
Porosity (Φ)	0.25

Table 4-3. Initial and boundary waters used in the problem.

Parameter	Initial water		Boundary water	
pH	8		5	
Species	**Concentration (mol/L_w)**	**Note**	**Concentration (mol/L_w)**	**Note**
Ca^{2+}	5.0×10^{-4}	Equilibrated with calcite	5.0×10^{-6}	Charge balance
HCO_3^-	1.0×10^{-3} bar	Equilibrium with CO_2 gas at given partial pressure	1.0×10^{-3}	
Tracer	1.0×10^{-10}		1.0	

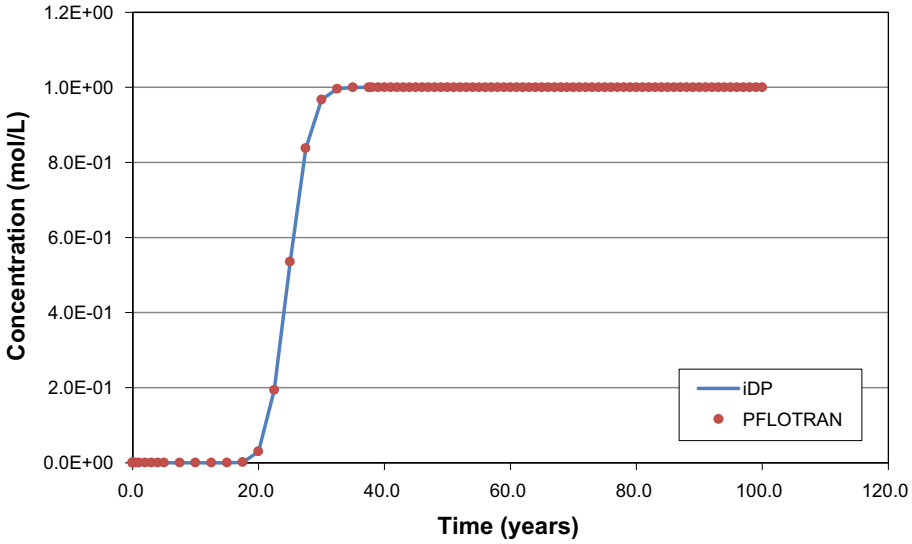

Figure 4-1. *Tracer breakthrough curve at the outlet of the domain computed with iDP (blue curve) and PFLOTRAN (red dots).*

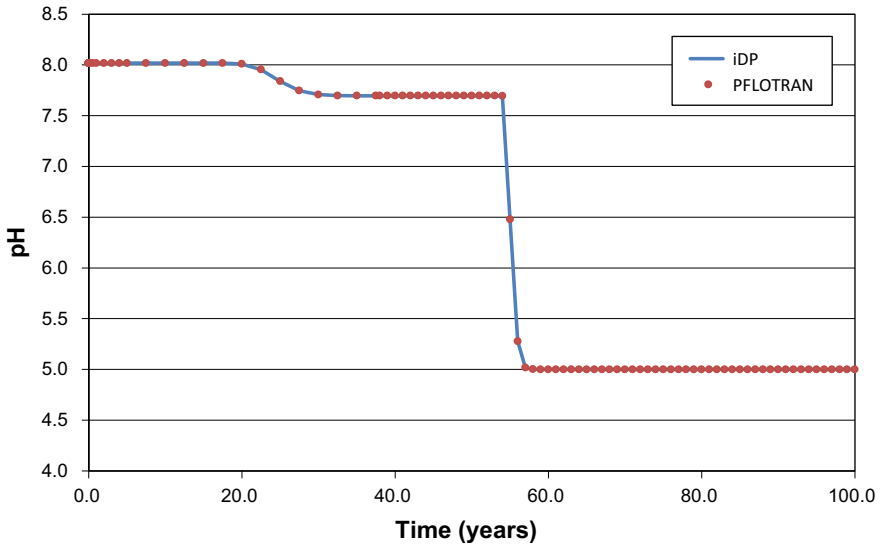

Figure 4-2. *pH evolution at the outlet of the domain computed with iDP (blue curve) and PFLOTRAN (red dots).*

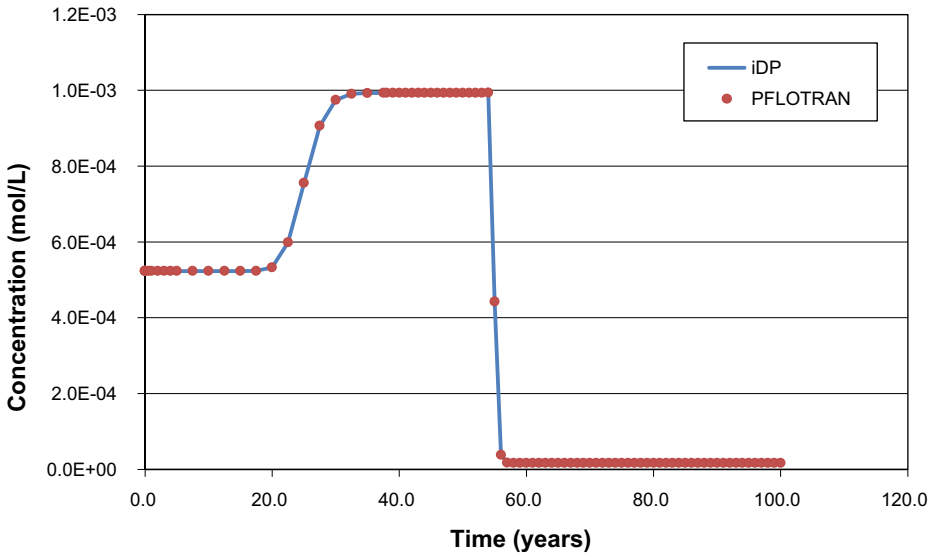

Figure 4-3. *Calcium breakthrough curve at the outlet of the domain computed with iDP (blue curve) and PFLOTRAN (red dots).*

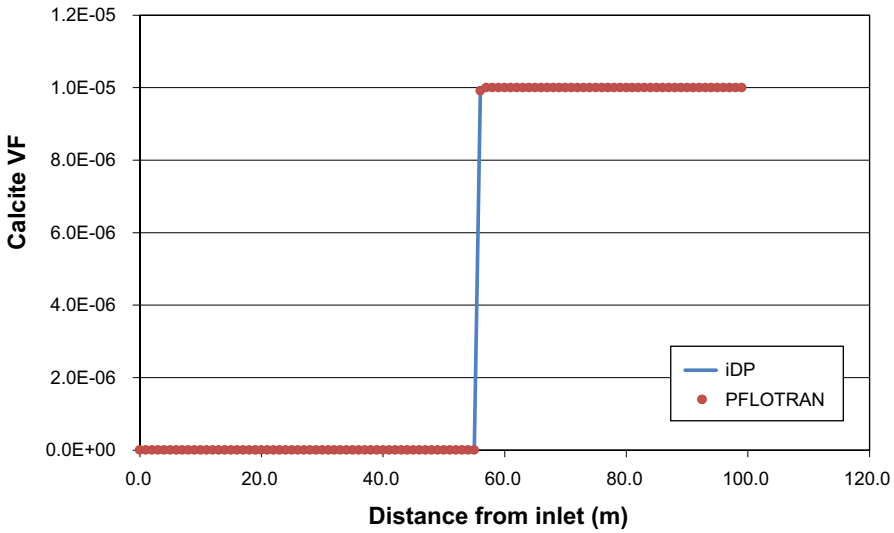

Figure 4-4. *Calcite volume fraction along the considered domain after 30 years from the beginning of the simulation.*

4.2 Cation exchange

Here, we use the example 11 of the PHREEQC User's Manual as benchmark for iDP. In the example 11 of PHREEQC, the chemical composition of the effluent from a column containing a cation exchanger is simulated. Initially, the column contains a sodium-potassium-nitrate solution in equilibrium with the cation exchanger. The column is then flushed with three pore volumes of calcium chloride solution. Calcium, potassium, and sodium react to equilibrium with the exchanger at all times.

The sodium initially present in the water exchanges with the calcium of the boundary water and is released as long as the exchanger contains still sodium. Because potassium has a higher affinity with the exchanger, is released after sodium. Finally, when all of the potassium has been released, the concentration of calcium increases to a steady-state value equal to the concentration in the boundary water.

In this modelling case, first a Darcy Tools model, with the same physical parameters specified in Table 4-4, has been set-up and solved. iDP has then been used to generate the PFLOTRAN input files, containing porosity an velocity, which are input to PFLOTRAN. The initial and boundary waters used in the calculations are listed in Table 4-5.

Table 4-4. Parameters used in the problem.

Parameter	Value
Domain length (L)	0.08 m
Darcy Velocity (q_x)	2.78×10^{-7} m/s
Porosity (Φ)	0.1

Table 4-5. Initial and Boundary waters used in the problem.

Parameter	Initial water	Boundary water
pH	7	7
Species	Concentration (mol/L_w)	Concentration (mol/L_w)
Ca^{2+}	1.0×10^{-12}	6.0×10^{-4}
K^+	2.0×10^{-4}	1.0×10^{-12}
Na^+	1.0×10^{-3}	1.0×10^{-12}
Cl^-	1.0×10^{-12}	1.2×10^{-3}
Tracer	1.0×10^{-12}	1.2×10^{-3}

In Figure 4-5, the results computed using iDP are compared to the solution given by PHREEQC. The agreement between the two solutions is good; although the PHREEQC solution is affected by some numerical dispersion, which is amplified in the breakthrough curves of those species affected by cation exchange reactions. Reaction-dependent numerical dispersion in PHREEQC is a well-documented issue ("*The magnitude of numerical dispersion also depends on the nature of the modeled reactions; numerical dispersion may be large in the many cases--linear exchange, surface complexation, diffusion into stagnant zones, among others {...}*"; (Parkhurst and Appelo 2013)).

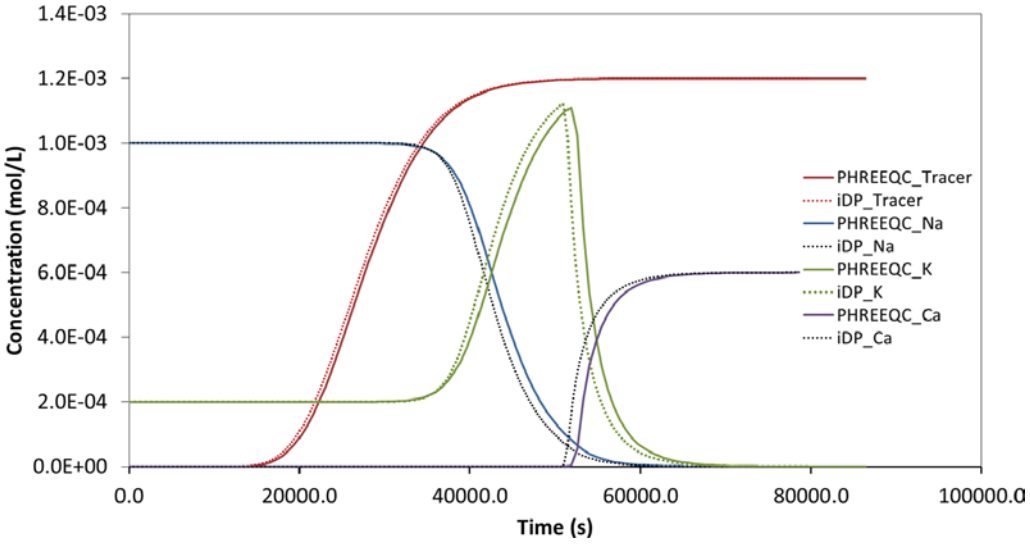

Figure 4-5. *Comparison of breakthrough curves computed using PHREEQC (solid line) and iDP (dots).*

4.3 1D Calcite dissolution and dolomite precipitation

This benchmark is presented in Section 15.1 of Kolditz et al. (2012) and is denoted as "*1D Reactive Transport: Calcite dissolution and precipitation"*. A one-dimensional column that initially contains calcite is continuously flushed with water that contains magnesium chlorine. With the movement of the water front, calcite starts to dissolve and dolomite is temporarily formed.

The parameters of the benchmark are summarized in Table 4-6. The initial and boundary waters used in this model are shown in Table 4-7.

To solve the benchmark in iDP, a DarcyTools model has been previously set-up and solved using the parameters listed in Table 4-6. iDP has then been used to generate the PFLOTRAN input files, containing porosity an velocity, which are input to PFLOTRAN.

In Figure 4-6, the results of iDP are compared to those reported in Kolditz et al. (2012). The agreement between the two solutions is good.

Table 4-6. Parameters of the benchmark problema.

Parameter	Value
Domain length (L)	0.5 m
Darcy Velocity (q_x)	3.0×10^{-6} m/s
Porosity (Φ)	0.32
Bulk Density	1830 kg/m^3
Longitudinal dispersivity	0.0067 m
Calcite abundance	2×10^{-4} mol/L

Table 4-7. Initial and boundary waters.

Parameter	Initial water	Boundary water
pH	9.91	7
Species	**Concentration (mol/L$_w$)**	**Concentration (mol/L$_w$)**
Ca^{2+}	1.227×10^{-4}	1.0×10^{-40}
HCO_3^-	1.227×10^{-4}	1.0×10^{-40}
Mg^{++}	1.0×10^{-40}	1.0×10^{-3}
Cl^-	1.0×10^{-40}	2.0×10^{-3}
Tracer	1.0×10^{-40}	2.0×10^{-3}

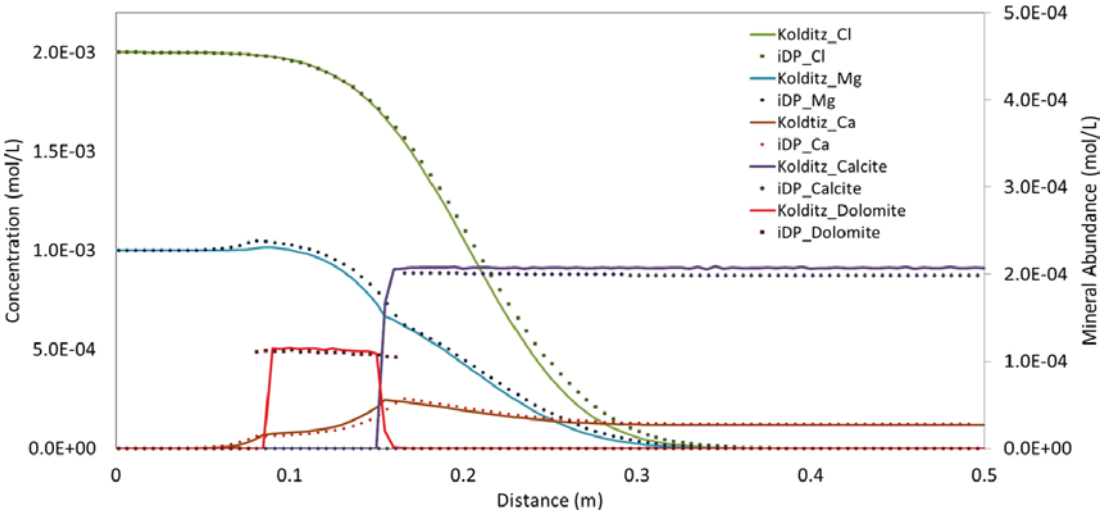

Figure 4-6. *Benchmark results from iDP compared vs. Kolditz (GEOSys-GEM). The profiles are taken at time t=2133.3 s.*

4.4 Copper leaching

The copper leaching problem simulates in-situ mining of copper by injection of an acidic solution (*pH=1*) into the subsurface. The acidic solution dissolves chrysocolla ($CuSiO_3 \times 2H_2O$) releasing copper into the solution. Copper is extracted above ground and secondary minerals precipitate in the Chrysocolla front's path (Lichtner 1996). A 5-spot well distribution (i.e. one injection well in the center and four extraction wells in the corner) is used in the simulation. Due to symmetry, only one quarter of the domain is modelled (Figure 4-7).

The parameters of the problem are summarized in Table 4-8.

The initial and inlet boundary conditions (i.e. initial and boundary waters) are described in Table 4-9. The minerals and related volume fractions are specified in Table 4-10. Minerals with volume fraction equal to zero are not initially present but can precipitate if over-saturation is reached.

Table 4-8. Copper leaching model parameters.

Parameter	Value
Domain length_X	16.0 m
Domain length_Y	16.0 m
Domain length_Z	128.0 m
Input rate	0.68 kg/s
Output rate	−0.68 kg/s
Porosity (Φ)	0.05
Permeability	1.5×10^{-13} m^2

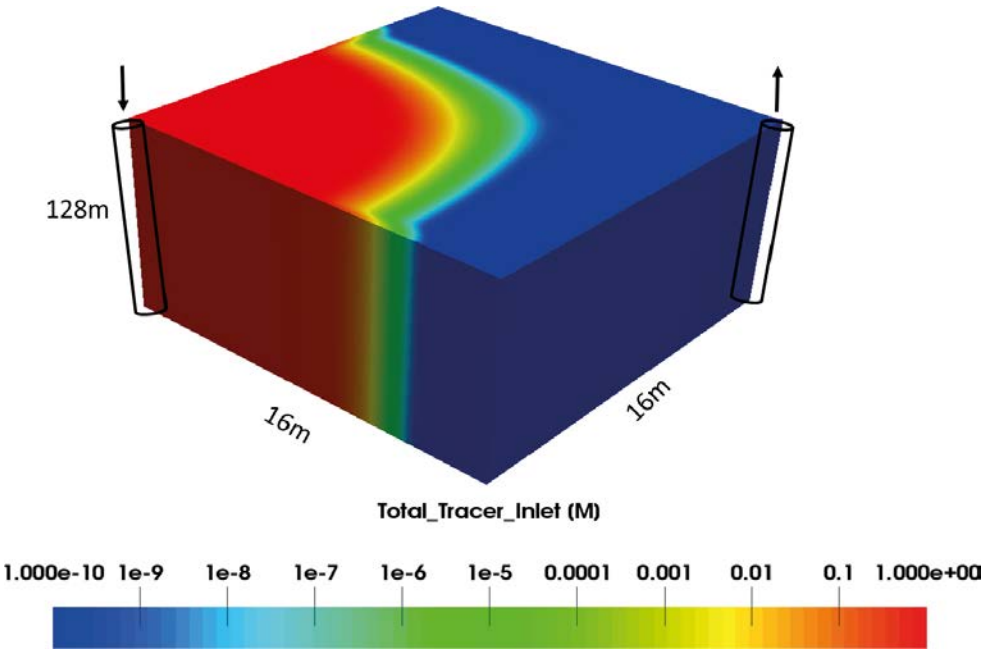

Figure 4-7. *Sketch of the copper leaching problem.*

Table 4-9. Copper leaching: initial and boundary waters.

Parameter	Initial water		Boundary water	
pH	8.0		1.0	
Species	**Concentration (mol/L_w)**	**Note**	**Concentration (mol/L_w)**	**Note**
Ca^{2+}	6.5×10^{-4}	Equilibrium with Calcite	1.1×10^{-2}	Equilibrium with Gypsum
Na^{+}	5.0×10^{-3}	Total aqueous component concentration	5.0×10^{-3}	Total aqueous component concentration
K^{+}	2.5×10^{-5}	Equilibrium with Muscovite	1.3×10^{-4}	Equilibrium with Jarosite
Al^{3+}	2.8×10^{-17}	Total aqueous component concentration	2.5×10^{-2}	Total aqueous component concentration
Cu^{2+}	6.4×10^{-9}	Equilibrium with Chrysocolla2	1.0×10^{-8}	Total aqueous component concentration
Cl^{-}	3.7×10^{-3}	Charge Balance	5.0×10^{-3}	Total aqueous component concentration
Fe^{2+}	1.2×10^{-23}	Equilibrium with Goethite	3.4×10^{-9}	Equilibrium with Goethite
$SiO_{2(aq)}$	1.8×10^{-4}	Equilibrium with Chalcedony	1.9×10^{-3}	Equilibrium with SiO_2(am)
HCO_3^{-}	−3.0	Equilibrium with CO_2(g)	−2.0	Equilibrium with CO_2(g)
SO_4^{2-}	5.0×10^{-4}	Total aqueous component concentration	6.1×10^{-2}	Charge Balance
$O_{2(aq)}$	−0.699	Equilibrium with O_2(g)	−0.699	Equilibrium with O_2(g)
Tracer	1.0×10^{-10}	Total aqueous component concentration	1.0	Total aqueous component concentration

Table 4-10. Minerals volume fractions present in initial water.

Parameter	Initial Mineral Volume Fraction	Specific Surface Area (cm^2 mineral/cm^3 bulk)
Alunite	0.0	1.0
Chrysocolla2	5.0×10^{-3}	1.0
Goethite	2.5×10^{-2}	1.0
Gypsum	0.0	1.0
Jarosite	0.0	1.0
Jurbanite	0.0	1.0
Kaolinite	5.0×10^{-2}	1.0
Muscovite	5.0×10^{-2}	1.0
SiO_2 (am)	0.0	1.0
Quartz	8.2×10^{-1}	1.0

Compared to the original PFLOTRAN problem, when running Copper Leaching in iDP two simplification assumptions are needed.

First, in the PFLOTRAN original problem the SOURCE_SINK card is used, which allows flow and transport boundary conditions (typically Neumann and Dirichlet, respectively) to be defined along 1D and 2D geometrical entities. In iDP, flow conditions are simulated in DarcyTools whereas the PFLOTRAN transport conditions always apply over cell faces.

The second difference is that, as specified in Chapter 3, iDP only works over sub-domains of DarcyTools. This means that the biggest DarcyTools sub-domain that can be extracted is one cell smaller in all the six directions (i.e. $x^{+}/x^{-};y^{+}/y^{-};z^{+}/z^{-}$), which in turns means that transport conditions, which are applied as Dirichlet conditions over the external faces of the first (bottom-left corner) and last cell (upper-right corner) of the sub-domain, are shifted one cell compared to the original problem.

The original problem set-up and the simplifications required by iDP are illustrated in Figure 4-8. To minimize the effect of these simplifications, an extremely fine mesh refinement has been used.

The results computed by PFLOTRAN standalone and iDP are used to produce maps of concentrations of Tracer, pH, Chrysocolla (volume fraction), Cu^{2+} and Jurbanite (volume fraction) (Figure 4-9 to Figure 4-13). The agreement between the two solutions (i.e. PFLOTRAN standalone vs. iDP) is good. Very small discrepancies can be observed close to the boundary as a result of the simplification assumptions discussed before.

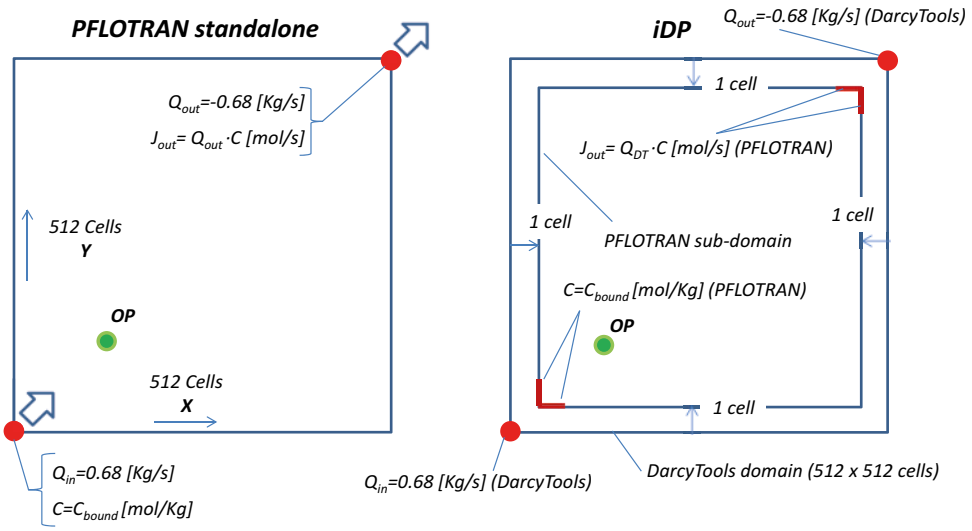

Figure 4-8. *Original set-up of PFLOTRAN (left) and set-up used in iDP (right). The green dot indicates observation point OP.*

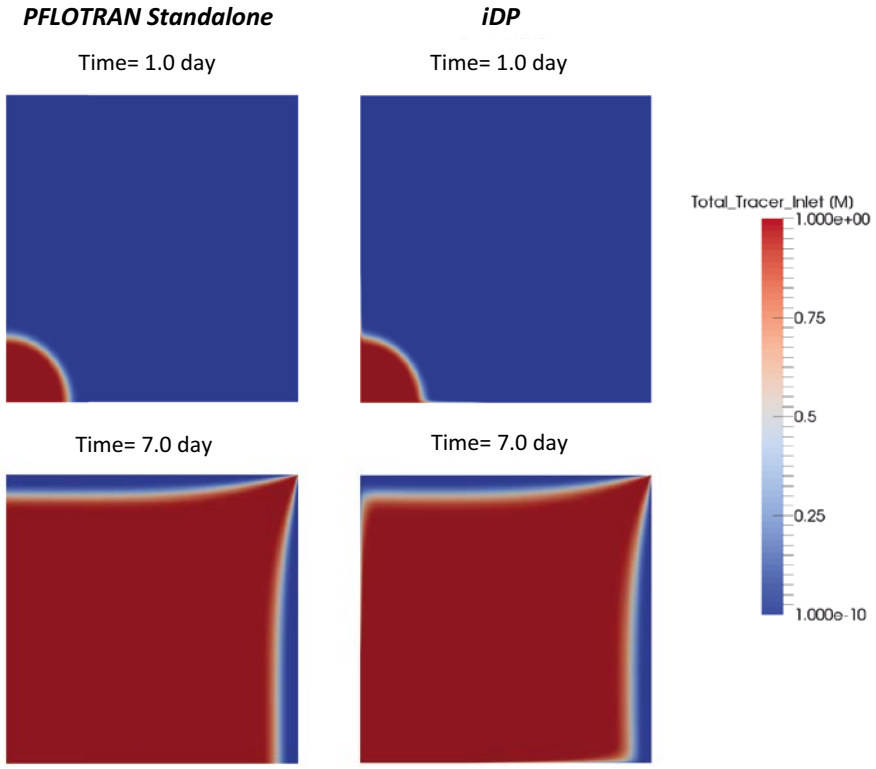

Figure 4-9. *Tracer concentration computed by PFLOTRAN (left) and iDP (right) at 1.0 day (top) and 7.0 day (bottom).*

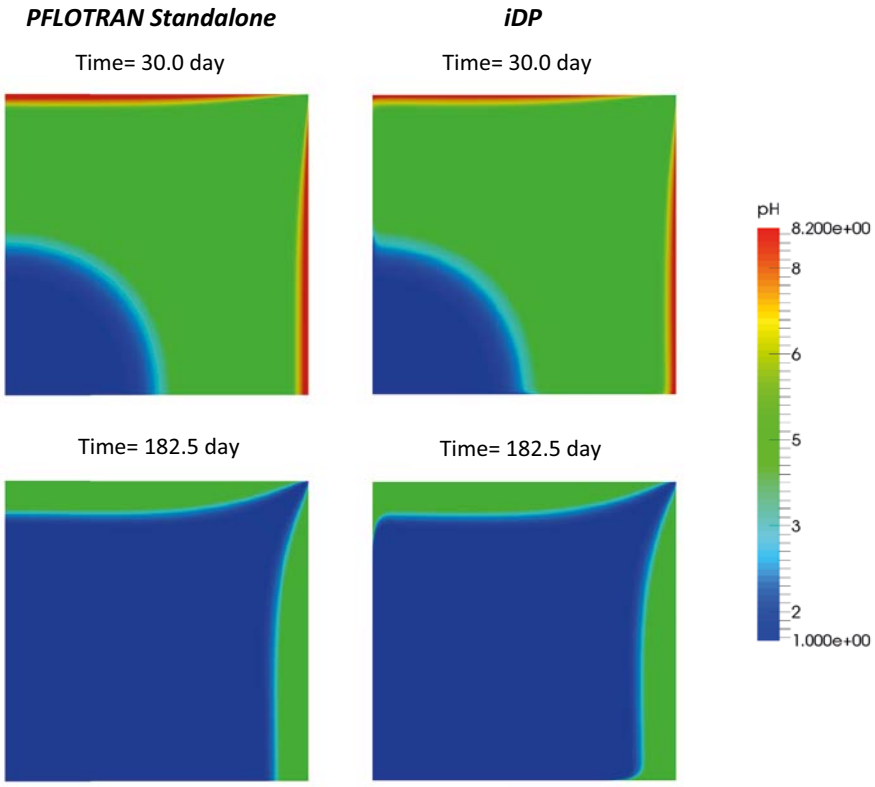

Figure 4-10. *pH computed by PFLOTRAN (left) and iDP (right) at 30.0 day (top) and 182.5 day (bottom).*

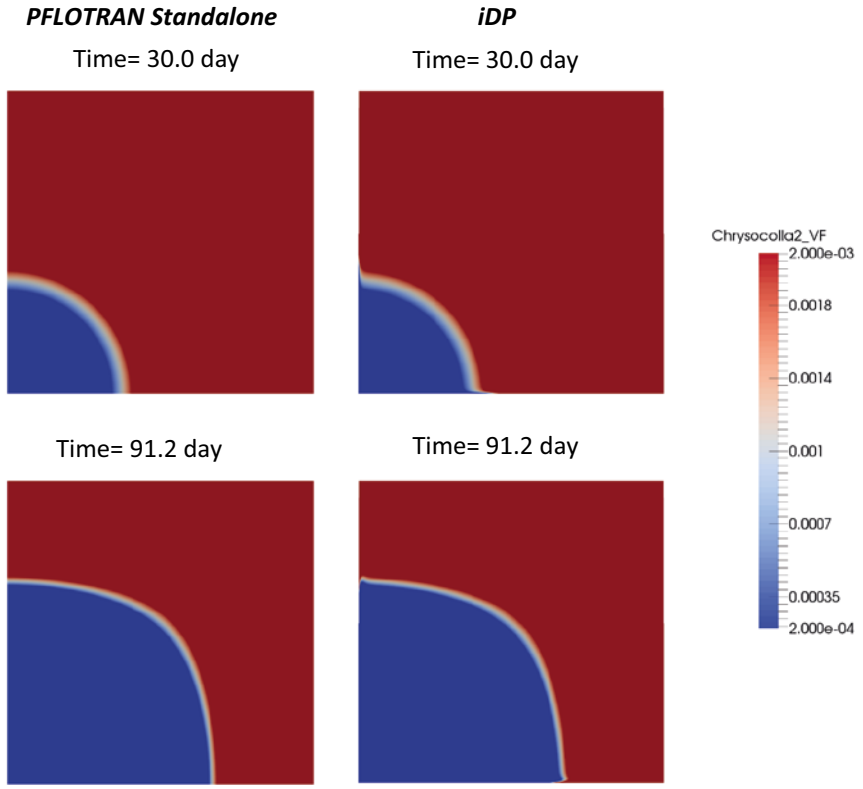

Figure 4-11. *Chrysocolla volume fraction computed by PFLOTRAN (left) and iDP (right) at 30.0 day (top) and 91.2 day (bottom).*

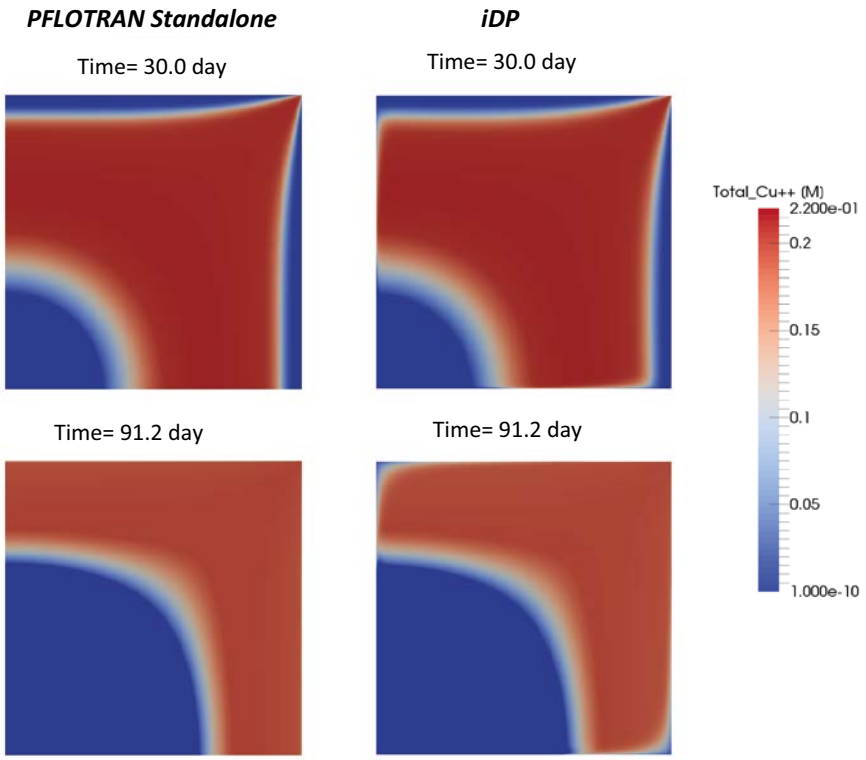

Figure 4-12. *Cu^{2+} concentration computed by PFLOTRAN (left) and iDP (right) at 30.0 day (top) and 91.2 day (bottom).*

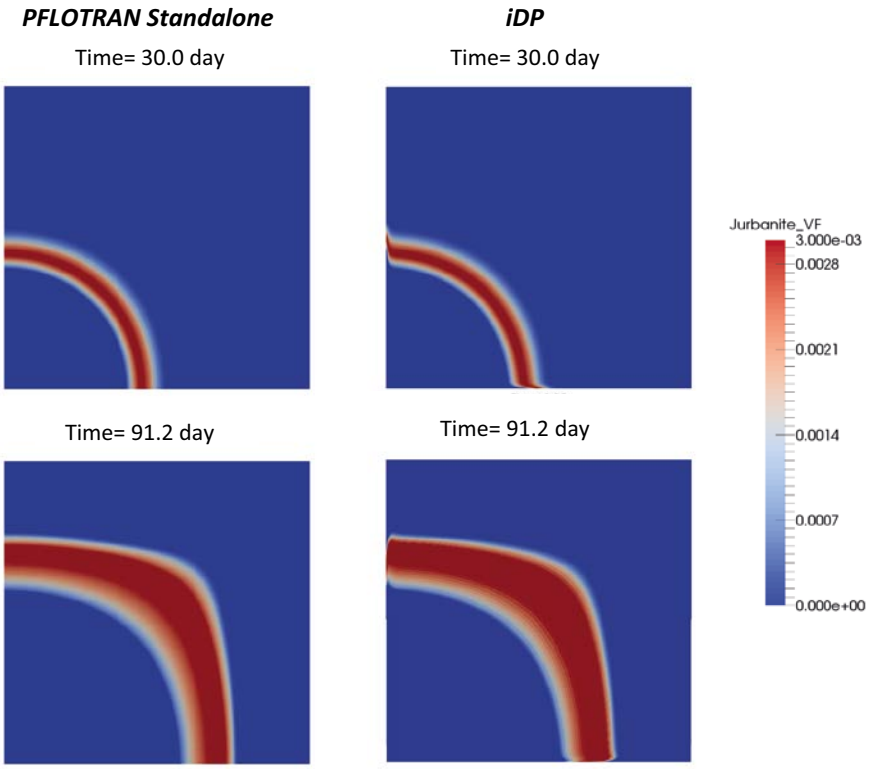

Figure 4-13. *Jurbanite Volume Fraction computed by PFLOTRAN (left) and iDP (right) at 30.0 day (top) and 91.2 day (bottom).*

Breakthrough curves computed using iDP and PFLOTRAN standalone at observation point OP (see Figure 4-8) for the conservative tracer, pH and copper are shown in Figure 4-14. The evolution of chrysocolla volume fraction computed at observation point OP using the two solutions is shown in Figure 4-15. For all the considered species and minerals, the match between the two solutions (i.e. PFLOTRAN standalone vs. iDP) is very good and the verification exercise is considered successful.

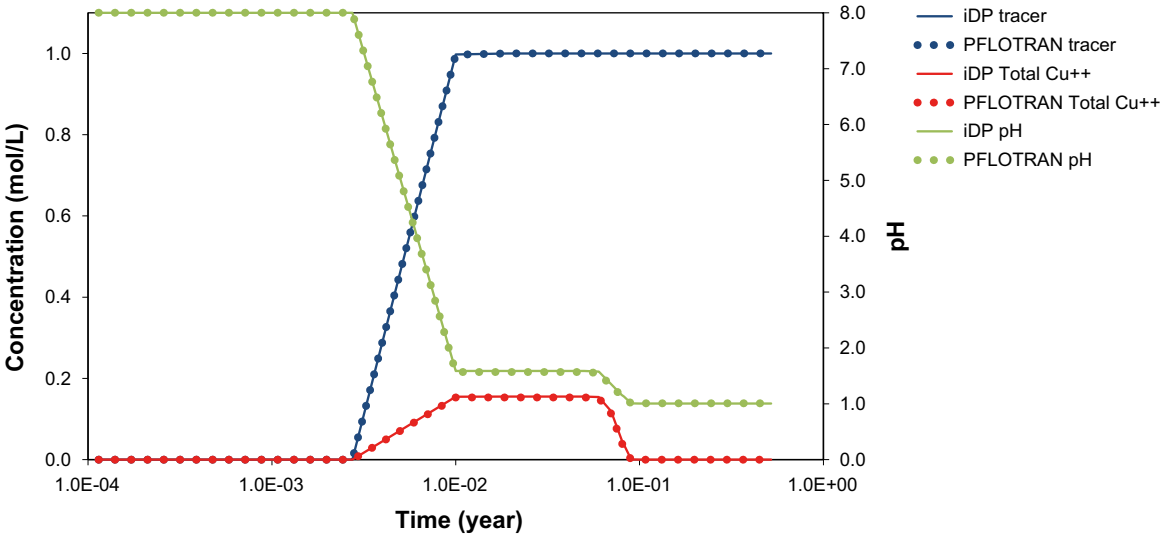

Figure 4-14. *Comparison of breakthrough curves computed using iDP (solid line) and PFLOTRAN standalone (dots). Breakthrough curves at computed at observation point OP (see Figure 4-8).*

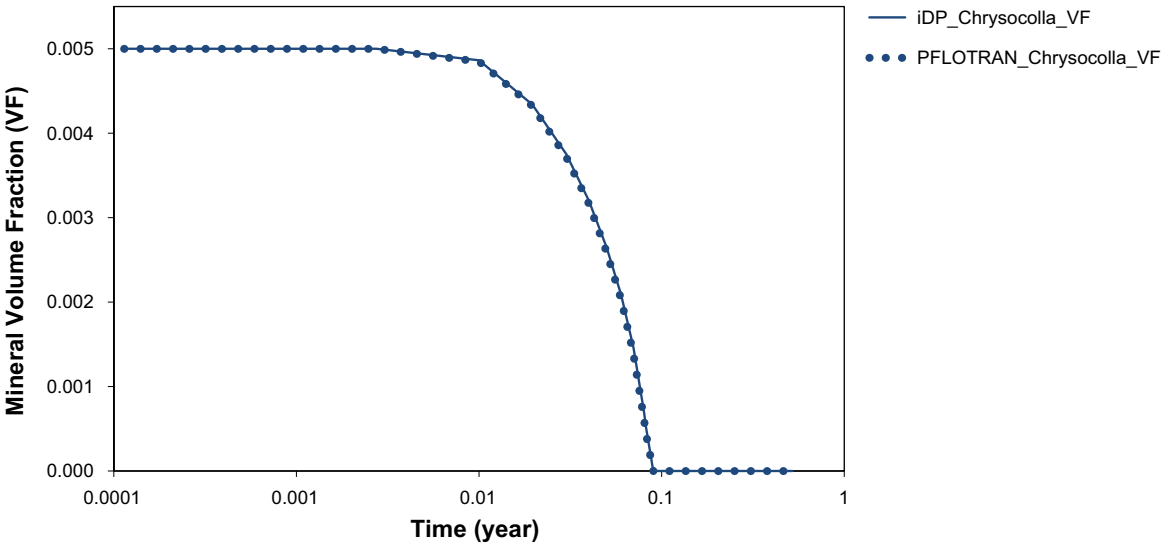

Figure 4-15. *Evolution of chrysocolla volume fraction computed at observation point OP (see Figure 4-8) using iDP (solid line) and PFLOTRAN standalone (dots).*

4.5 Matrix diffusion and oxygen consumption

Here, we simulate the ingress of oxygen in a fracture-matrix system and its consumption by abiotic processes. The results (i.e. steady-state concentration profiles of oxygen) are compared to the analytical solution developed by Sidborn et al. (2010).

A sketch of the model set-up is shown in Figure 4-16. Oxygenated water enters the model domain from the left boundary of the fracture and moves driven by advective fluxes along the fracture. Part of this oxygen diffuses into the matrix where is consumed by the homogeneous oxidation of ferrous iron. Ferrous iron is in turn released by the non-oxidative dissolution of biotite, which is present in the matrix (volume fraction, V_f=0.078) and is absent in the fracture. Given the relatively high biotite concentration, at some point in time the concentration profile of oxygen reaches a steady-state.

In this numerical exercise, the groundwater flow velocity field has been computed using DarcyTools. Darcy's velocities and material properties (i.e. porosity) have then been extracted using iDP and used as input to PFLOTRAN.

As mentioned before, biotite is assumed to be the source of dissolved ferrous iron. Its non-oxidative dissolution is represented using the following rate-law:

$$r_{diss}[mol_{Fe}m^{-2}s^{-1}] = k_{Fe,diss} \cdot \left(\frac{V_f^{biot}}{V_{f0}^{biot}}\right)^{2/3} \cdot \left(1 - \frac{c_{Fe}}{c_{Fe,sol}}\right) \quad \text{Equation 4-1}$$

where $k_{Fe,diss}$ is the maximum release rate of ferrous iron, c_{Fe} is the iron aqueous concentration, $c_{Fe,sol}$ is the iron solubility and V_{f0}^{biot} and V_f^{biot} are, respectively, the initial and the computed biotite volume fraction.

The dissolved ferrous ions are oxidized by oxygen as follows:

$$4Fe^{2+}+O_2+4H^+ \leftrightharpoons 4Fe^{3+}+2H_2O \quad \text{Equation 4-2}$$

A widely used rate law for this reaction can be written for pH > 4.5 (Stumm and Lee 1961):

$$R_{ox}[mol_{Fe}L^{-1}s^{-1}] = k_r c_{Fe} P_{O_2} c_{OH^-}{}^2 \quad \text{Equation 4-3}$$

where k_r is a homogeneous oxidation rate constant, P_{O_2} is the partial pressure of oxygen and c_{OH^-} is the concentration of OH^-.

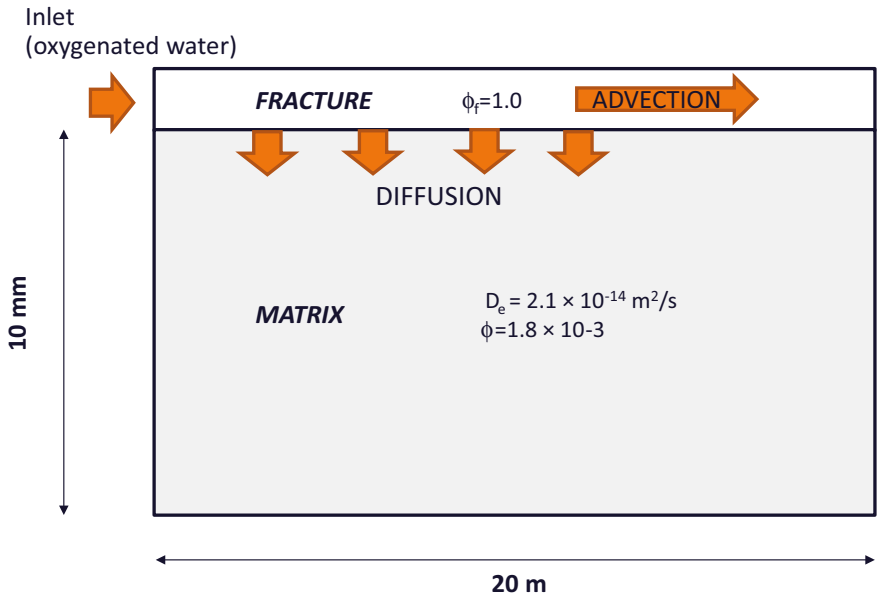

Figure 4-16. *Sketch of the oxygen intrusion problem. All the parameters are taken from Sidborn et al. (2010).*

In PFLOTRAN, custom kinetic reaction rates have to be defined using the "Reaction Sandbox" (see appendix of the PFLOTRAN User's Manual for further details). To implement a new reaction within the reaction sandbox, the user has to hardcode it in the related subroutines and classes of the code (i.e. "reaction_sandbox.F90" and "reaction._sandbox_base.F90"). Although this is a very powerful code capability, it is evident that the use of the reaction sandbox is not straightforward as it requires changes in the PFLOTRAN source code, which has then to be recompiled.

In order to keep this exercise at a simple level of complexity, Equation 4-3 has been simplified by assuming a constant OH^- concentration (as in Sidborn et al. 2010, who also fixed the hydroxide ion concentration, $[OH^-]$ to 1×10^{-6} mol/L_w). In this way, Equation 4-2 reduces to a second order kinetic reaction, which in PFLOTRAN can be easily defined using the "General_Reaction" card.

By default, dissolution/precipitation reactions are written in PFLOTRAN using transition state rate laws. The implementation of alternative formulations is not straightforward (they need to be hard-coded in the "Reaction Sandbox" subroutine). Thus, we have simplified the rate of biotite dissolution as follows:

$$r_{diss}[mol_{Fe}m^{-2}s^{-1}] = k_{Fe,diss} \cdot \left(1 - \frac{c_{Fe}}{c_{Fe,sol}}\right) \qquad \text{Equation 4-4}$$

It is worthwhile noting that in Equation 4-4, the rate of dissolution of biotite is no longer weighted by the relative mineral abundance.

The results of the iDP simulation show that a steady-state profile of oxygen is reached after 200 days of simulation time, approximately (Figure 4-17).

The oxygen concentration profiles in the rock matrix at different fracture locations are shown in Figure 4-18. The results are compared to those obtained using the analytical solution of Sidborn et al. (2010). When the concentration in the fracture is high (i.e. close to the inlet boundary), iDP underestimates the extent of oxygen penetration, whereas far from the inlet boundary (i.e. low oxygen concentrations in the fracture) the agreement between the two solutions is good.

The reason for the discrepancy observed at high fracture concentrations is to be found in the simplified approach used to represent the biotite dissolution rate (i.e. Equation 4-4). In fact, in PFLOTRAN the rate law is not weighted by the relative mineral abundance. This means that close to the injection boundary, that is where more biotite is depleted, the rate of mineral dissolution computed by PFLOTRAN is actually higher than that considered by the analytical solution resulting in a stronger buffering effect.

Figure 4-17. *Oxygen concentration computed using iDP after 200 days of simulation time. The domain is magnified by a factor of 400 in the direction perpendicular to the fracture.*

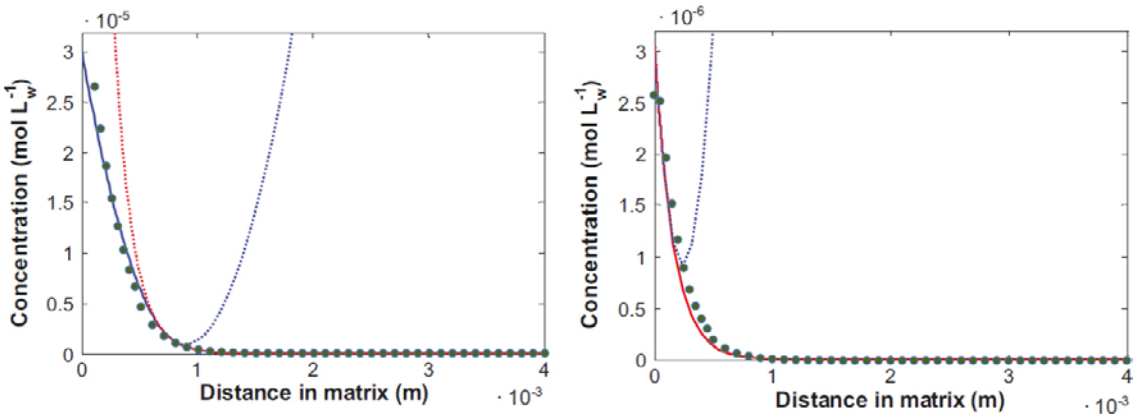

Figure 4-18. *Steady-state oxygen concentration profiles computed in the matrix at different fracture concentrations. iDP results (green dots) are plotted over the results of Sidborn et al. (2010) (solid line). Small dots indicate tha auxiliary solution used in the analytical calculation (for more details the reader is referred to Sidborn et al. (2010)).*

5 Test case #1: Grout degradation and hyper-alkaline plume development at the Forsmark spent fuel repository

5.1 Background

The iDP framework has been applied to simulate the cement grout degradation in the planned repository for nuclear spent fuel in the Forsmark area in Sweden (Figure 5-1). The spent fuel repository (Figure 5-2) is to be constructed at a depth of 500 (approximately) meters and designed to be isolated from people and surface biosphere for at least 100 000 years.

During the construction of the tunnels, water-conducting fractures will be sealed with cement-based grout. After closure of the tunnels, this could lead to the formation of hyper-alkaline (high pH) water (Soler 2010) that can be detrimental for the stability of bentonite, which is used in the waste barriers. The development of the hyper-alkaline water plume and the risk of the plume reaching repository depth have been evaluated. The formation and migration of the hyper-alkaline plume (high pH plume) during 20 000 years is simulated in a regional scale 3D model. The results show, in 3D, the general changes in geochemistry, the influence of the rock fractures on the movement of plume and the limited mobility of the high-pH plume.

It is worthwhile stressing that, though the parameterization of the model relies on specific features of the Forsmark site (e.g. hydraulic conductivity field, site topography, repository layout, etc.), this exercise should be considered as merely illustrative, as it is based on a number of simplifying assumptions (e.g. no land uplift, constant groundwater density, no matrix diffusion, etc.).

Figure 5-1. *Map of Sweden and location of the Forsmark area.*

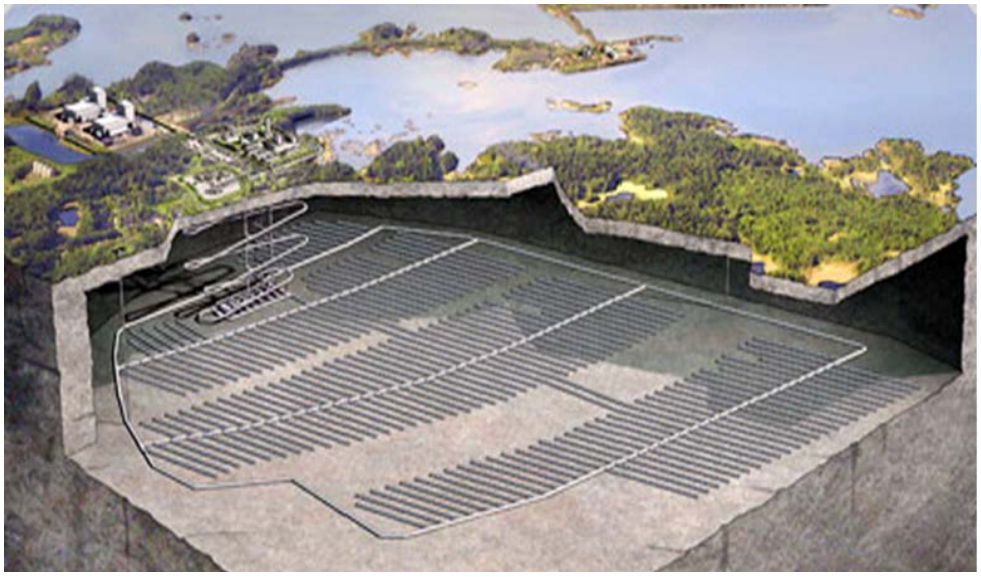

Figure 5-2. *Artistic view of The Forsmark Spent Fuel Repository (from SKB Website).*

5.2 Objective

The main objective of this test case is to demonstrate that the iDP framework is ready to use state-of-the-art HPC techniques to simulate complex, large-scale, highly nonlinear, long-term RTM problems.

5.3 Methodology

The groundwater flow in the Forsmark regional domain is solved in DarcyTools. The Darcy and mass conservation equations are solved for. Although certain density stratification can be expected, due to a variation of the salinity, the density-driven flow effect has been neglected in this simulation. A box-shaped sub-domain that covers the repository volume will be used for the reactive transport simulation.

When steady state has been reached, the DarcyTools output files are post-processed by iDP. The sub-domain is extracted into binary files and then the inputs are written to HDF5 files that can be read by PFLOTRAN. The actual input file for PFLOTRAN that contains references to these HDF5 files is also generated. The final input files are transferred to a large parallel system where PFLOTRAN will be run. To be able to process the output file, which is also written in HDF5, several post-process scripts to split the huge output file into smaller subsets are applied. The results are visualized with ParaView.

5.4 Technical platform

The DarcyTools simulation and iDP were run on a powerful workstation. The PFLOTRAN simulation run on the supercomputer called “MareNostrum”. MareNostrum is hosted by the Barcelona Super-computer Center (www.bsc.es). MareNostrum has 48 896 Intel Sandy Bridge processors in 3056 nodes, with more than 96.6 TB of main memory and 2 PB of disk storage. Currently, MareNostrum holds the 47th position in the TOP500 list of fastest supercomputers in the world. The actual simulation took place over 15 days using 2.77 million CPU hours. The calculation time on the supercomputer was offered by BSC in the framework of a “Grand Challenge” Program.

5.5 Conceptual model and numerical implementation

5.5.1 Hydrogeological model

The hydraulic conductivity field generated by DarcyTools is shown in Figure 5-3. The dimension of the domain is $2700 \times 2200 \times 548$ m^3.

At the top boundary the ground water table is fixed to a level of 0.5 m below ground level. All other boundaries are of zero flux type. A grouting condition is applied to seal the repository tunnel wall sections that contain fractures with a higher permeability (conducting water) than 1.99×10^{-15} m^2. The finite volume grid used to represent the repository is shown in Figure 5-4. Figure 5-5 shows a close up view of the same grid.

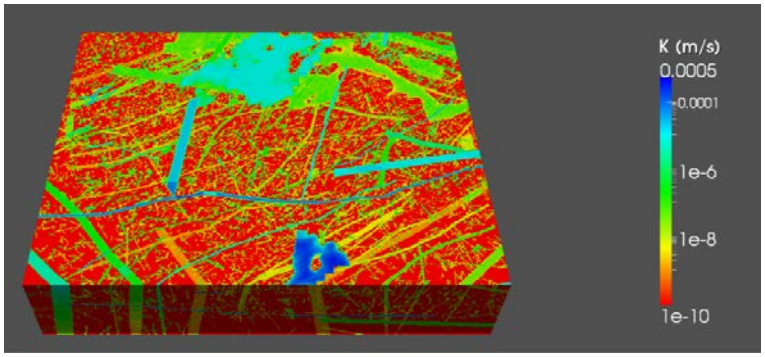

Figure 5-3. *Hydraulic Conductivity Field For the Forsmark Spent Fuel Repository area generated by DarcyTools.*

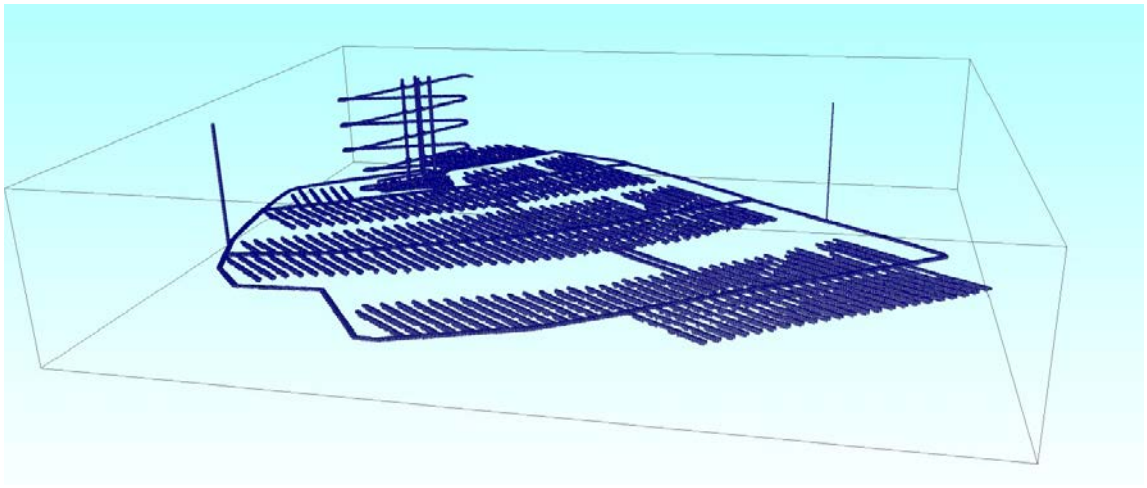

Figure 5-4. *Mesh used to discretize the repository.*

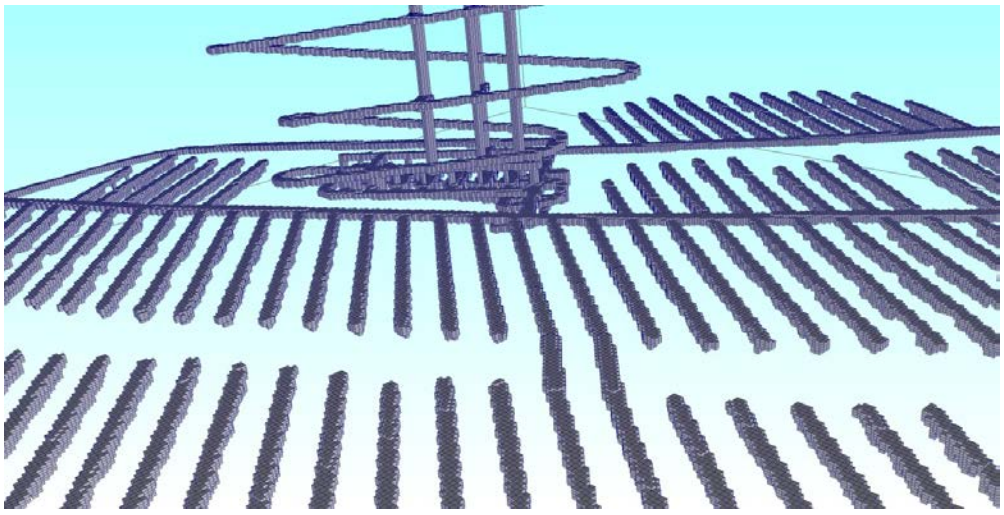

Figure 5-5. *Detailed view of the mesh used to discretize the model domain (deposition tunnels are shown).*

5.5.2 Geochemical model

The geochemical conditions in the Forsmark area are described in Salas et al. (2010). In this report the following groundwater types are distinguished at different depths:

- 0–20m: Near-surface water
- 20–200m: Shallow groundwater
- 200–600m: Intermediate depth groundwater
- > 600: Deep groundwater

In the conceptual model this is translated to the following water types:

- **0–36m: Near-surface water**
  This boundary water type is set to "Altered Meteoric".
- **36–200m: Shallow groundwater**
  In the model represented by a mix of the water types "Altered Meteoric" and "Littorina", with a Cl concentration of 0.03(M).
- **200–584: Intermediate depth groundwater**
  In the model represented by a mix of water types "Altered Meteoric" and "DeepSalineForsmark", with a Cl concentration of 0.15(M).
- **> 584: Deep groundwater**
  This boundary water type is represented in the model by a mix of water types "Altered Meteoric" and "DeepSalineForsmark", with a Cl concentration of 0.282(M).

The concentrations resulting from mixing (linear interpolation) the mentioned water types (definition of the water types can be found on page 144–145 of Salas et al. 2010) can be found in Table 5-1. These groundwater types have then been pre-equilibrated with calcite and quartz and charge balanced with chloride (Table 5-2). These pre-equilibrated groundwater types are used as initial and boundary conditions for the model, as shown in Figure 5-6.

Table 5-1. Chemistry of the different water types resulting from the mixing of original water types from Salas et al. (2010).

Parameter	Near surface water	Shallow grondwater	Intermediate groundwater	Deep groundwater
pH	7.31	7.40	7.39	7.45
pe	0.55	−0.06	−0.13	−0.76
Species	**Concentration (mol/kgw)**	**Concentration (mol/kgw)**	**Concentration (mol/kgw)**	**Concentration (mol/kg_w)***
Br	7.17×10^{-6}	4.50×10^{-5}	4.49×10^{-4}	8.51×10^{-4}
C	7.23×10^{-3}	6.46×10^{-3}	6.46×10^{-3}	5.77×10^{-3}
Ca	4.70×10^{-4}	9.38×10^{-4}	5.30×10^{-2}	1.01×10^{-1}
Cl	5.11×10^{-3}	3.00×10^{-2}	1.50×10^{-1}	2.82×10^{-1}
F	8.43×10^{-5}	7.63×10^{-5}	8.45×10^{-5}	8.47×10^{-5}
K	1.43×10^{-4}	6.02×10^{-4}	2.16×10^{-4}	2.82×10^{-4}
Li	2.02×10^{-6}	3.15×10^{-6}	7.49×10^{-5}	1.41×10^{-4}
Mg	3.09×10^{-4}	2.66×10^{-4}	2.86×10^{-4}	2.64×10^{-4}
Mn	0.0	2.57×10^{-3}	2.78×10^{-7}	5.32×10^{-7}
Na	1.19×10^{-2}	3.26×10^{-2}	5.11×10^{-2}	8.68×10^{-2}
S	8.86×10^{-4}	2.06×10^{-3}	8.02×10^{-4}	7.26×10^{-4}
Si	1.34×10^{-4}	1.33×10^{-4}	1.29×10^{-4}	1.25×10^{-4}
Sr	4.34×10^{-6}	8.01×10^{-6}	4.24×10^{-4}	8.07×10^{-4}

* kg_w denotes kilogram of pure water.

Table 5-2. Chemistry of the water types after initial equilibration in PFLOTRAN using the Hanford.dat thermodynamic data base (default data base).

Parameter	Near Surface Water		Shallow Grondwater		Intermediate Groundwater		Deep Groundwater	
pH	7.31		7.40		7.39		7.45	
Species	**Concentration (mol/kgw)**	**Equilibrated With**	**Concentration (mol/kgw)**	**Equilibrated With**	**Concentration (mol/kgw)**	**Equilibrated With**	**Concentration (mol/kgw)**	**Equilibrated With**
Br	7.17×10^{-6}		4.50×10^{-5}		4.49×10^{-4}		8.51×10^{-4}	
C	7.23×10^{-3}		6.46×10^{-3}		6.46×10^{-3}		5.77×10^{-3}	
Ca	1.06×10^{-3}	Calcite	1.33×10^{-3}	Calcite	1.42×10^{-3}	Calcite	1.62×10^{-3}	Calcite
Cl	6.33×10^{-3}	Charge Balanced	3.10×10^{-2}	Charge Balanced	4.75×10^{-2}	Charge Balanced	8.48×10^{-2}	Charge Balanced
F	8.43×10^{-5}		7.63×10^{-5}		8.45×10^{-5}		8.47×10^{-5}	
K	1.43×10^{-4}		6.02×10^{-4}		2.16×10^{-4}		2.82×10^{-4}	
Li	2.02×10^{-6}		3.15×10^{-6}		7.49×10^{-5}		1.41×10^{-4}	
Mg	3.09×10^{-4}		2.66×10^{-4}		2.86×10^{-4}		2.64×10^{-4}	
Mn	1.00×10^{-20}		2.57×10^{-3}		2.78×10^{-7}		5.32×10^{-7}	
Na	1.19×10^{-2}		3.26×10^{-2}		5.11×10^{-2}		8.68×10^{-2}	
S	8.86×10^{-4}		2.06×10^{-3}		8.02×10^{-4}		7.26×10^{-4}	
Si	1.01×10^{-4}	Quartz	1.01×10^{-4}	Quartz	1.01×10^{-4}	Quartz	1.02×10^{-4}	Quartz
Sr	4.34×10^{-6}		8.01×10^{-6}		4.24×10^{-4}		8.07×10^{-4}	

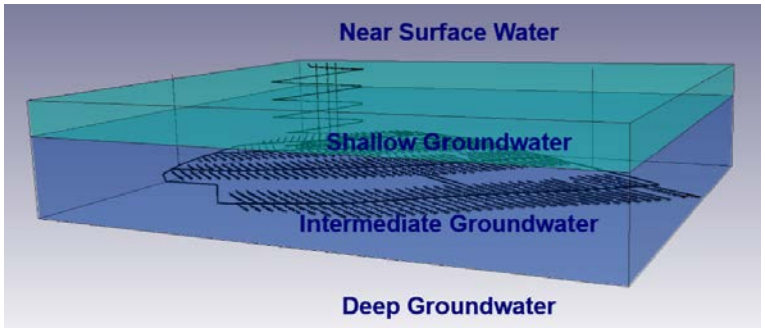

Figure 5-6. *Geochemistry setup. The Shallow and Intermediate groundwaters will be used as initial waters and as boundary waters for the sides. The Near Surface Water and Deep Groundwater will be used as top and bottom boundary waters.*

In PFLOTRAN terms, the initial condition type is set to "zero_gradient". It is assumed that the groundwater surrounding the domain is of the same type as the initial waters inside the domain. So at the sides of the box, the boundary condition "dirichlet_zero_gradient" is applied for the respective water type, meaning that on fluid-inflow boundaries a solution is specified (i.e. Dirichlet boundary condition), whereas zero diffusive gradient is considered on outflow (i.e. only advective transport is considered on outflow).

At the top the boundary condition "dirichlet_zero_gradient" will be applied with the "Near Surface Water" and at the bottom the "Deep Groundwater" will be applied.

Table 5-3 lists the minerals that are present at the start of the simulation (primary minerals). The whole domain is assumed to have the same distribution of minerals. All the mineral properties are taken from the thermodynamic database provided with PFLOTRAN ("*hanford.dat*"). Quartz and Calcite are considered to be present in abundance (volume fraction set equal to 0.1). A surface area of 1000 [m^2 mineral/m^3 mineral], corresponding to the surface area of a spherical grain with a radius of 3mm, is used. Quartz is treated kinetically (values for kinetic rates taken from Palandri and Kharaka (2004)) and the rest of the minerals, including the secondary minerals (see Table 5-4), are treated as equilibrium minerals. The minerals occupying the remaining volume fraction are considered as non-reactive.

The secondary aqueous species considered are presented in Table 5-5.

Despite of the nuclear waste repository being located in a low-permeability rock mass (with relatively low fracturing at repository depth), during the construction of the tunnels, the groundwater seepage into the repository must be minimized by sealing, with cement-based grouting, transmissive fractures and deformation zones. This grouting is represented by portlandite and CSH-1.8. After closure of the tunnels, the interaction of this grout with the surrounding water and rock will lead to the formation of a hyper-alkaline plume (high pH), as explained in Sidborn et al. (2014).

The minerals present at the cells where grout has been applied are listed in Table 5-6.

For the composition of the groundwater in the areas where grout has been applied, it is assumed that the "Altered Meteoric" water is equilibrated with Calcite and Portlandite and charge-balanced with pH (Salas et al. 2010). The resulting composition is listed in Table 5-7.

Table 5-3. Primary minerals; minerals present at the start of the simulation.

Mineral	Volume fraction (m^3 mineral/total volume m^3)	Surface area (m^2 mineral/m^3 mineral)
Quartz	0.1	1000
Calcite	0.1	1000

Table 5-4. Secondary minerals; the minerals that can precipitate (and dissolve again) during the simulation.

Brucite	CSH1.20	Gypsum
Chalcedony	CSH1.40	Portlandite
Calcite	CSH1.60	
CSH1.00	CSH1.80	

Table 5-5. Secondary species considered in the simulation.

OH^-	$MgSO_4(aq)$	$CaOH^+$
CO_3^{--}	MgF^+	$SrCl^+$
$CO_2(aq)$	$MgCl^+$	$SrCO_3(aq)$
$H_3SiO_4^-$	$MgOH^+$	SrF^+
$H_2SiO_4^{--}$	NaBr(aq)	$SrHCO_3^+$
$H_2F_2(aq)$	NaCl(aq)	$SrOH^+$
HF(aq)	$NaCO_3^-$	$SrSO_4(aq)$
HF_2^-	NaF(aq)	MnF^+
HSO_4^-	$NaHCO_3(aq)$	$MnCO_3(aq)$
KBr(aq)	$NaHSiO_3(aq)$	$MnSO_4(aq)$
KCl(aq)	NaOH(aq)	$MnHCO_3^+$
KOH(aq)	$NaSO_4^-$	$Mn(OH)_2(aq)$
KSO_4^-	$CaCO_3(aq)$	$Mn(OH)_3^-$
LiCl(aq)	$CaHCO_3^+$	$Mn(OH)_4^{--}$
LiOH(aq)	$CaSO_4(aq)$	$Mn_2(OH)_3^+$
$LiSO_4^-$	$CaCl_2(aq)$	Mn_2OH^{+++}
$MgCO_3(aq)$	$CaCl^+$	$MnOH^+$
$MgHCO_3^+$	CaF^+	

Table 5-6. Minerals present at the cells where grout is applied.

Mineral	Volume fraction (m^3 mineral/volume m^3)	Surface area (m^2 mineral/m^3 mineral)
Portlandite	0.0005	1000
CSH1.80	0.0005	1000

Table 5-7. Cement grout water composition after equilibration in PFLOTRAN using the Hanford thermodynamic database.

Parameter	Grout water	
pH	12.39	Charge balanced
Species	**Concentration (mol/kgw)**	**Equilibrated with**
Br^-	4.50E−05	
HCO_3^-	8.50E−06	Calcite
Ca^{++}	1.71E−02	Portlandite
Cl^-	3.00E−02	
F^-	7.63E−05	
K^+	6.02E−04	
Li^+	3.15E−06	
Mg^{++}	2.44E−08	Brucite
Mn^{++}	2.57E−03	
Na^+	3.26E−02	
SO_4^{--}	2.06E−03	
$SiO_2(aq)$	6.64E−06	CSH1.80
Sr^{++}	8.01E−06	

5.5.3 Simulation setup

The area outlined in Figure 5-7 represents the flow domain for the DarcyTools simulation.

A Discrete Fracture Network is first generated by DarcyTools and then converted into an equivalent continuous porous medium (ECPM), which is where the flow simulation is carried out. The flow simulation mesh consists of 105 044 948 cells. The porosity field and the resulting darcy velocity field for the sub-domain, containing the repository, were used as input for the PFLOTRAN simulation. This sub-domain is a regular mesh of 102 394 896 cells, each with a width and length of 4 meter and a 2 meter height (Figure 5-8).

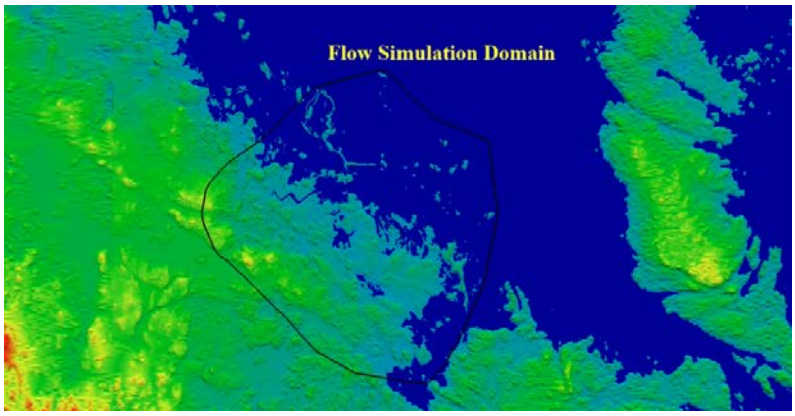

Figure 5-7. *DarcyTools flow simulation domain.*

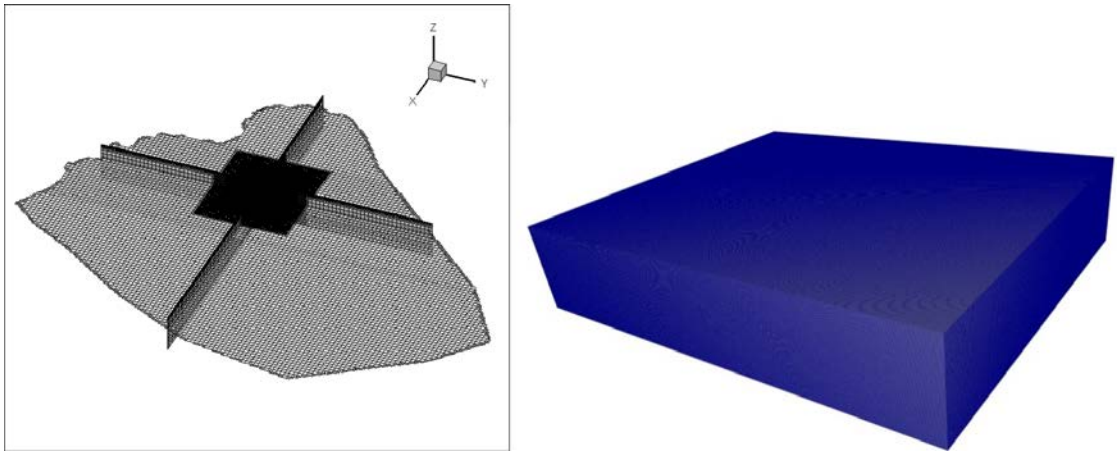

Figure 5-8. *(left) DarcyTools mesh (105 million cells). The refined cube (102 million cells) in the center represents the regular domain to be used for reactive transport (right).*

5.6 Results

5.6.1 Groundwater flow simulation

We have already presented and briefly discussed the hydraulic conductivity field used in DarcyTools (Figure 5-3). The discrete fractures can be clearly recognized and in the side of the domain one can see that the fracture density is larger closer to the surface.

The results from the DarcyTools steady state flow simulation, using this conductivity field, are shown in Figure 5-9. The shape of the flow lines is influenced by the underlying fracture distribution (see also Figure 5-10 with the combination of hydraulic conductivity and flow lines).

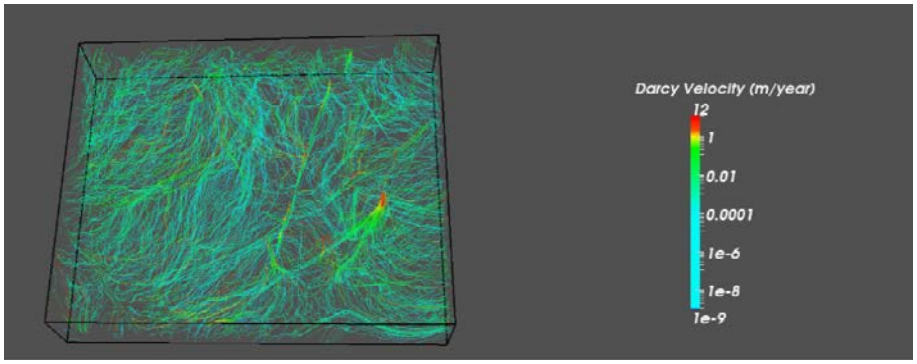

Figure 5-9. *Darcy velocity flow lines for the steady state solution.*

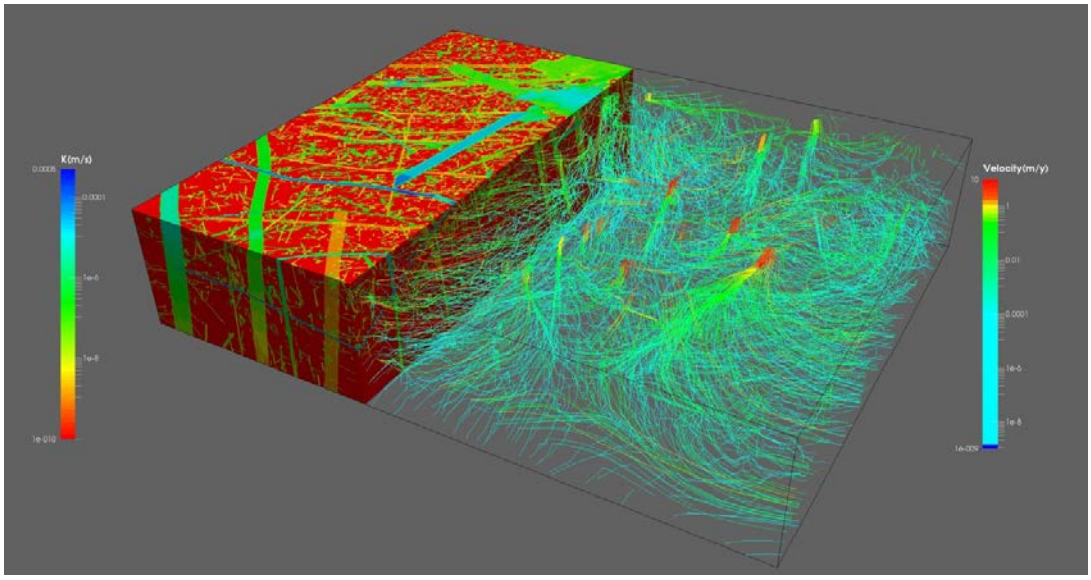

Figure 5-10. *Hydraulic conductivity and darcy velocity flow lines for the steady state solution.*

5.6.2 Reactive transport simulation

The dissolution of the Portlandite of the grout results in an increase of the pH in the groundwater. Figure 5-11 shows the growth of the plume of pH (only the cells with pH higher than 7.6 are plotted). The plume grows first in the shallow part of the domain that is where groundwater velocities are higher. Then, over the next 500 years it grows downward. However, the very high pH water (pH > 11) is restricted to a small area (see Figure 5-12) and does not affect the chemistry of the storage tunnels. After 20 000 years there are still some small areas with high pH.

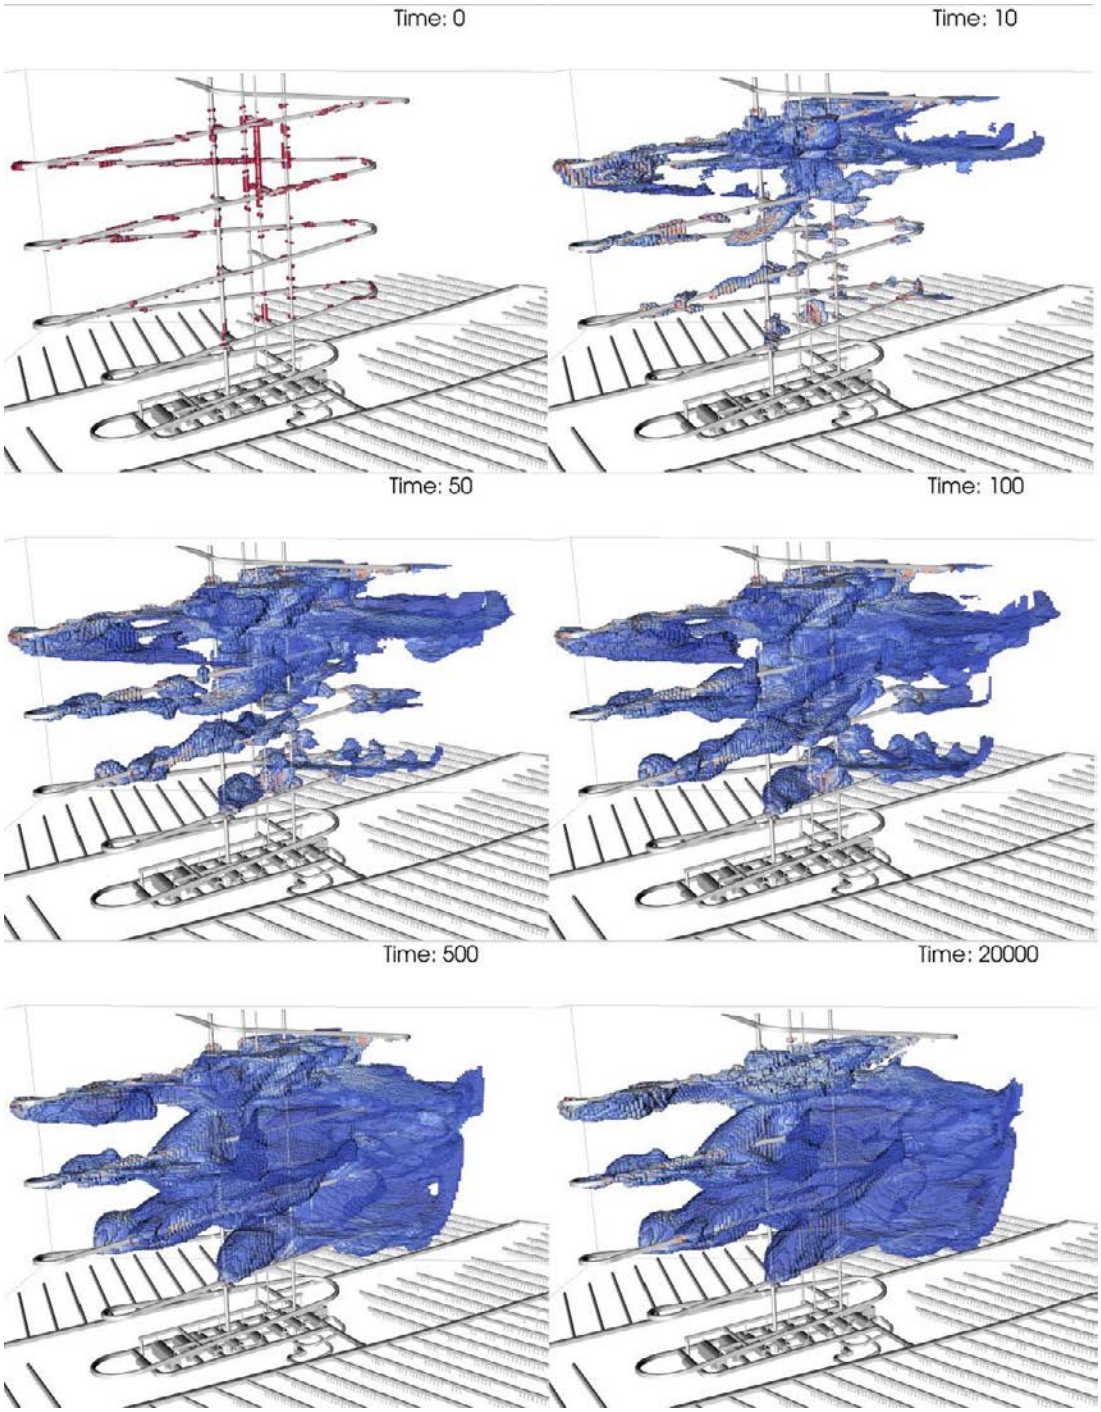

Figure 5-11. *Evolution of pH over time with pH > 7.6*

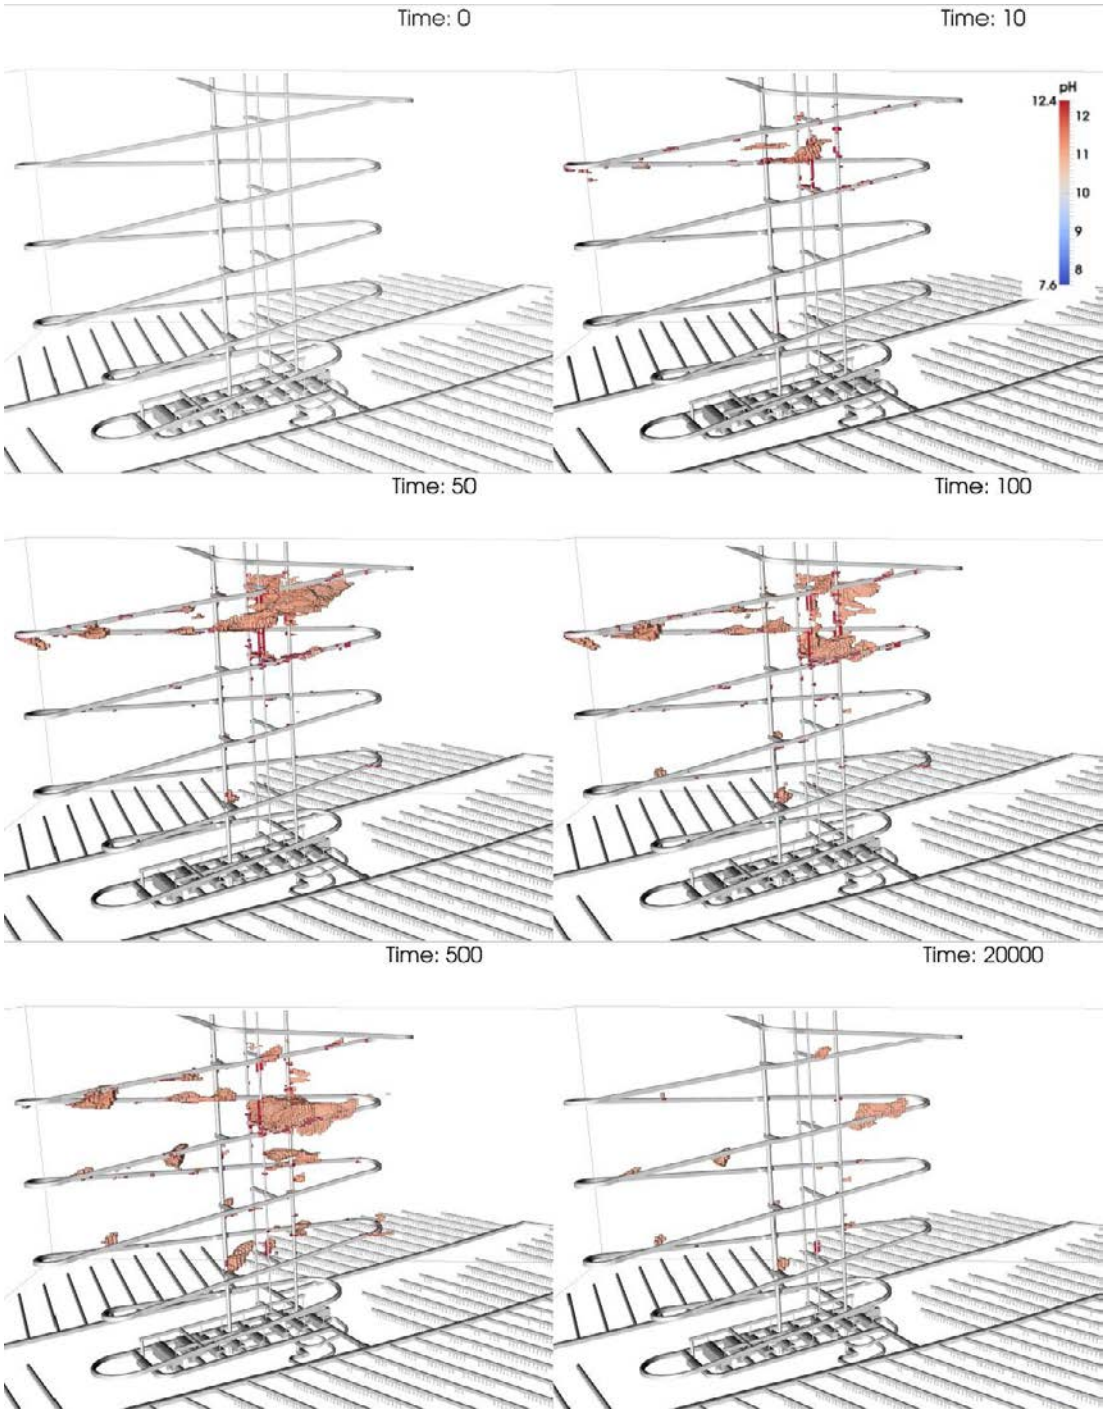

Figure 5-12. *Evolution of pH over time with pH > 11*

As a result of the higher pH, Brucite precipitates and quartz dissolves in the area of the pH plume (Figure 5-13 and Figure 5-14). Also, as a result of the release of Ca due to portlandite and CSH-1.8 dissolution, a clear precipitation of Calcite is observed close to the repository (Figure 5-15). Figure 5-16 displays the sequence of precipitation and dissolution of the different CSH types. After 10 years CSH 1.4 has precipitated (green cells), and the re-dissolved to precipitate CSH 1.2 (t=500y) and, after 10 000 years, the CSH 1.2 is re-dissolved to precipitate CSH 1.0.

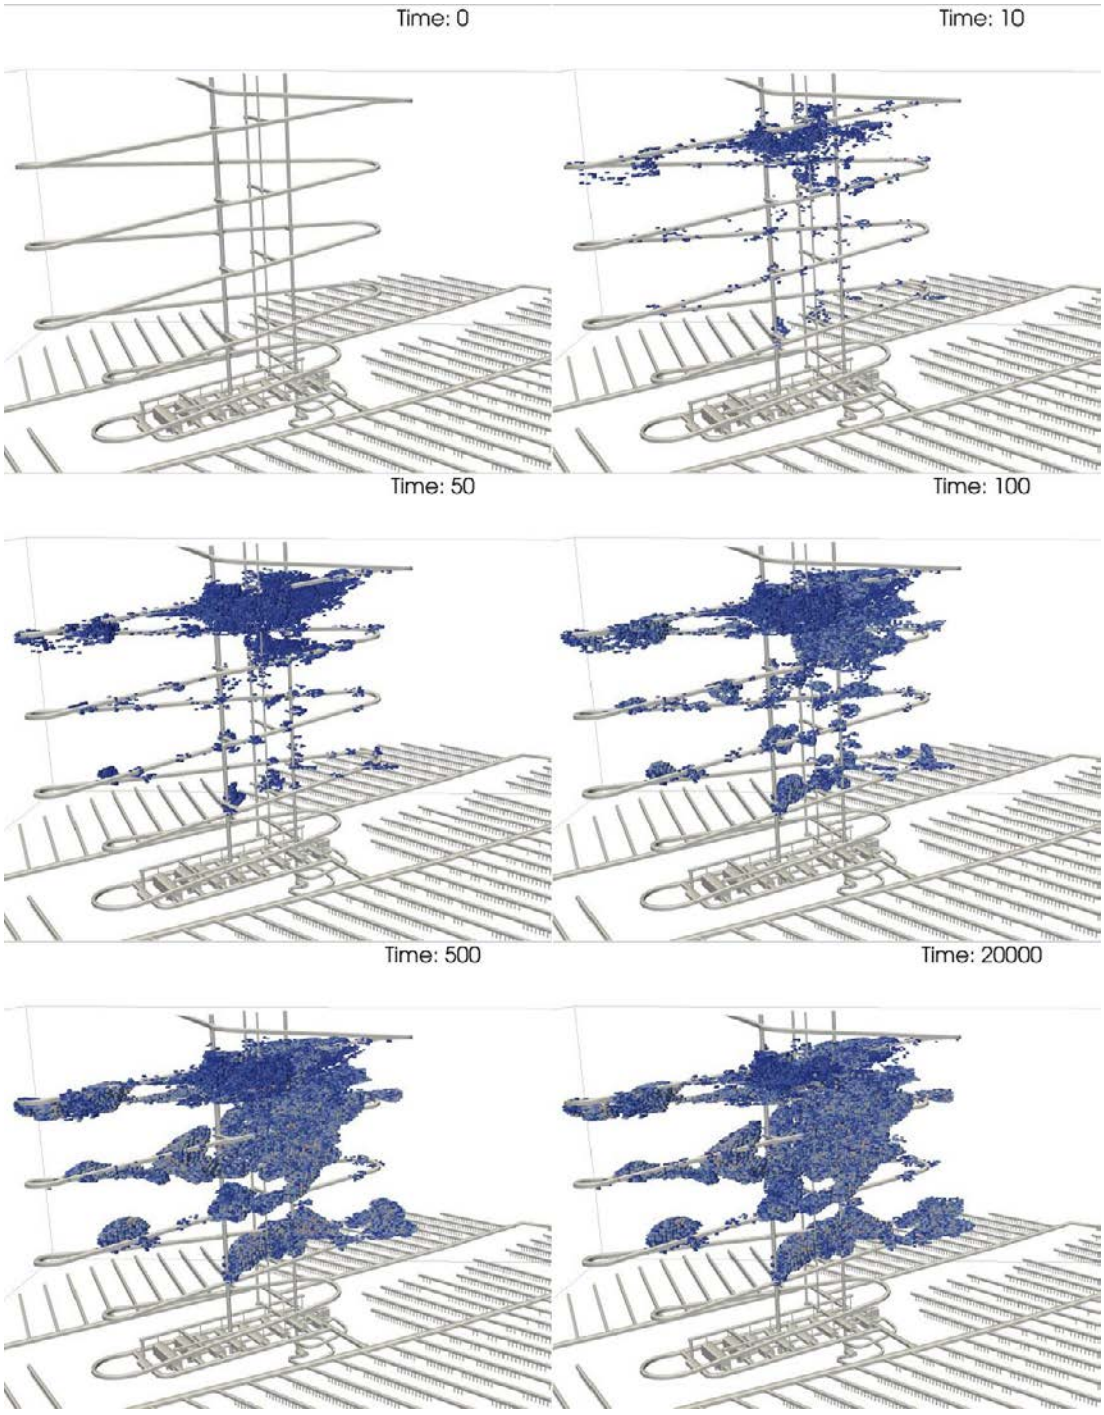

Figure 5-13. *Brucite precipitation (change in mineral volume fraction: ΔV_f) over time.*

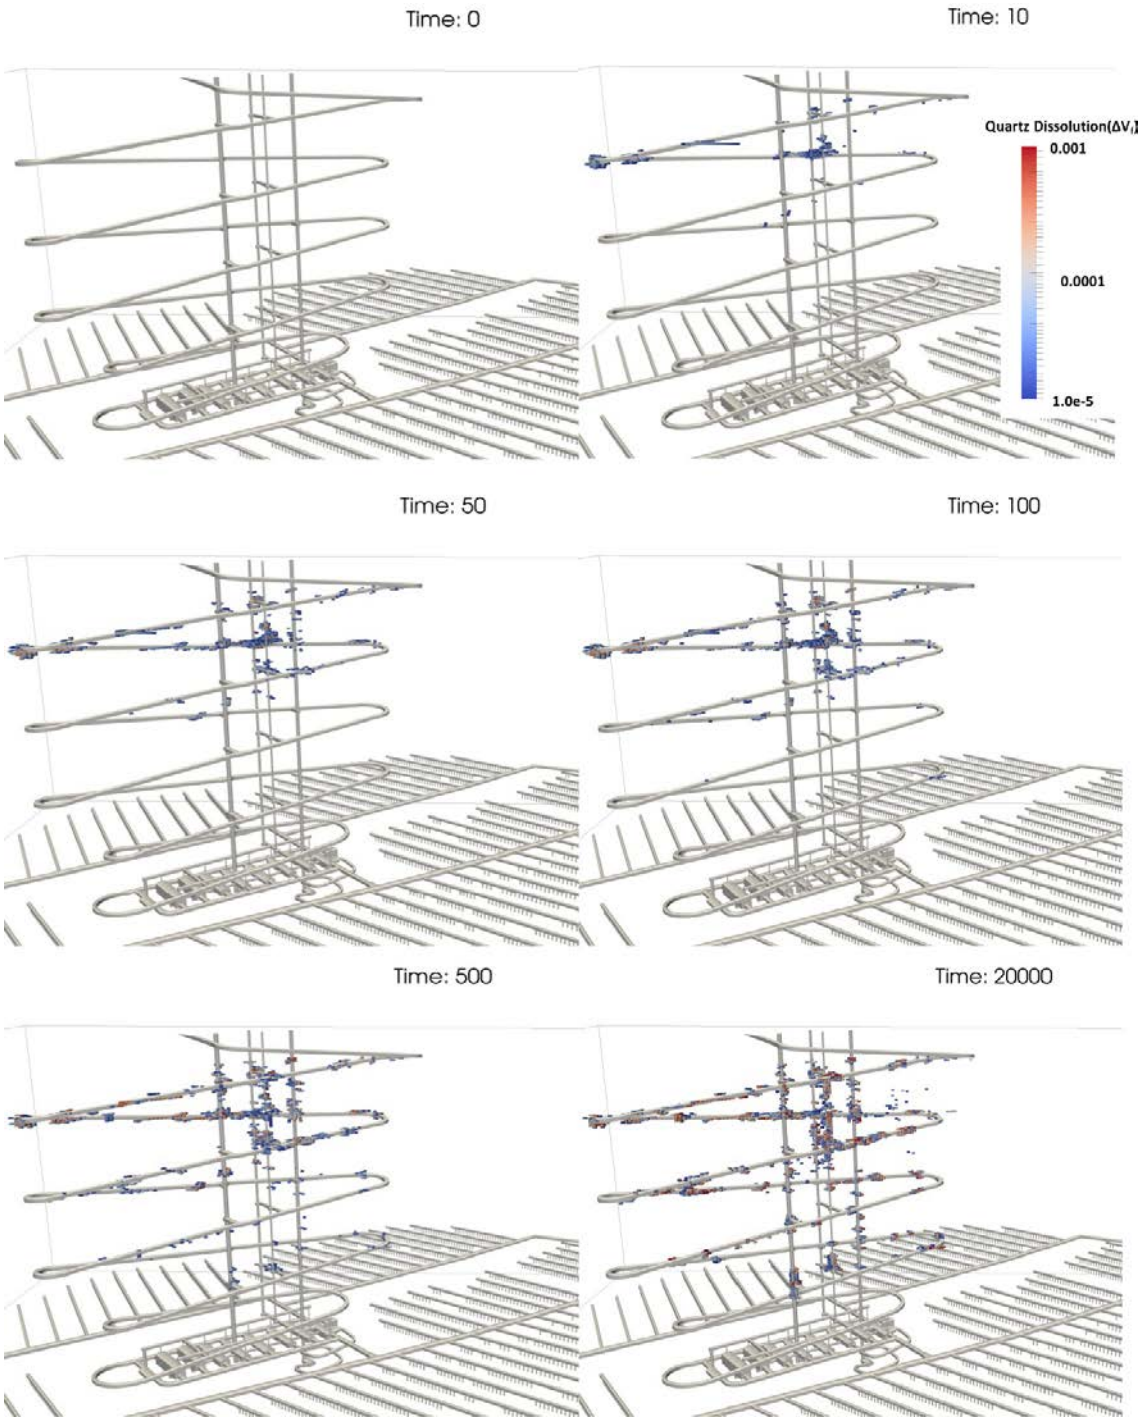

Figure 5-14. *Quartz dissolution* (*change in mineral volume fraction:* ΔVf) *over time.*

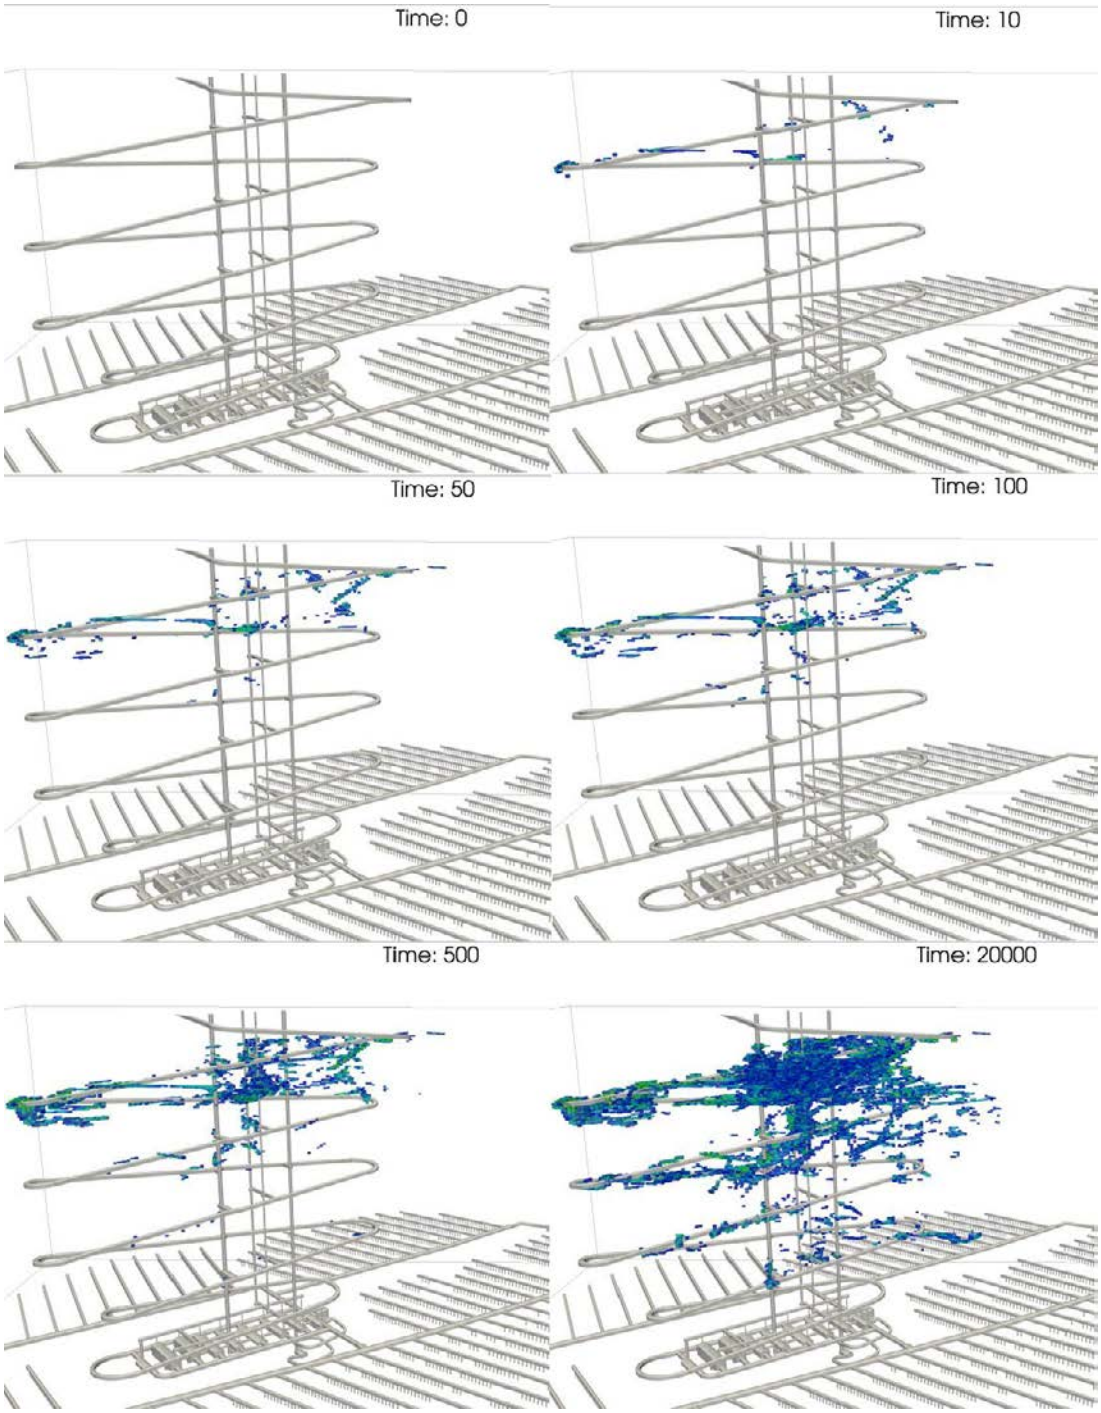

Figure 5-15. *Calcite precipitation (change in mineral volume fraction: ΔV_f) over time.*

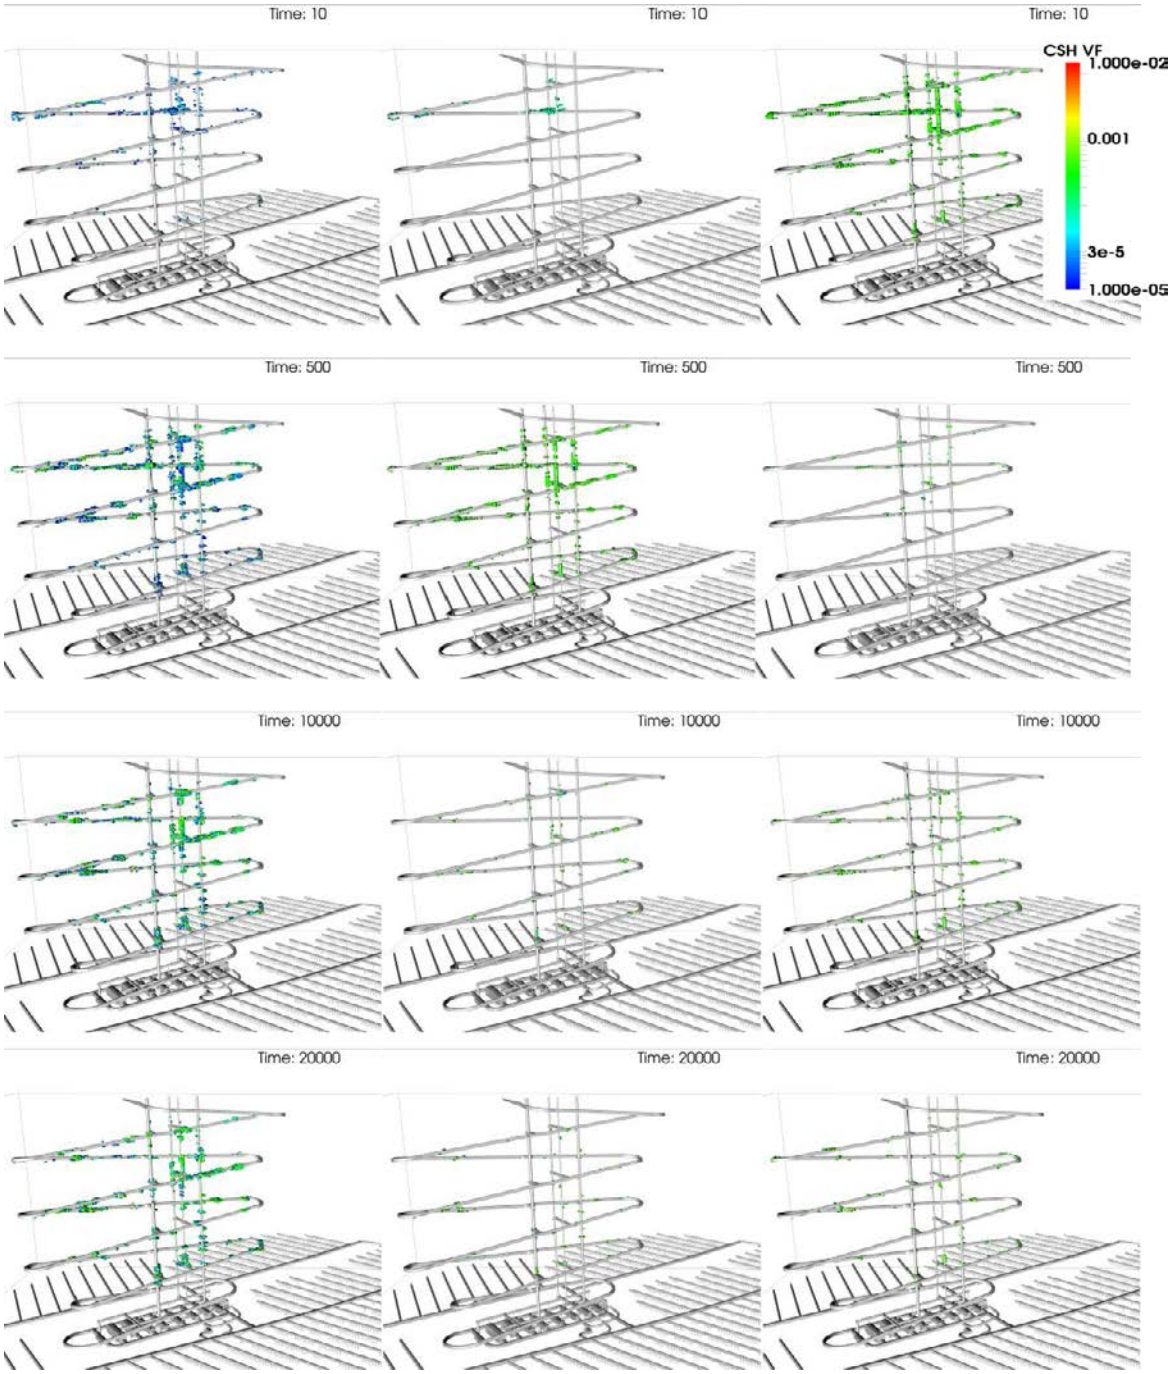

Figure 5-16. *CSH1.0 (Left), CSH1.2 (Middle) and CSH1.4 (Right) volume fraction evolution, showing precipitation and subsequent dissolution of the CSH phases.*

5.6.3 Model performance

Different scalability tests were carried out. These scalability tests are based on the same geochemical model presented before (Section 5.5.2). However, the domain size has been changed to span from a minimum of 125 000 cells to a maximum of 8 000 000 cells. The results of the tests, which have been carried out the MareNostrum supercomputer, demonstrate that iDP-PFLOTRAN scales well and given the problem size and HPC system, there can be found an optimal number of processes for which the problem scales optimally (Figure 5-17).

More specifically, for a problem of 100 Mcells several thousands of processors can be efficiently used.

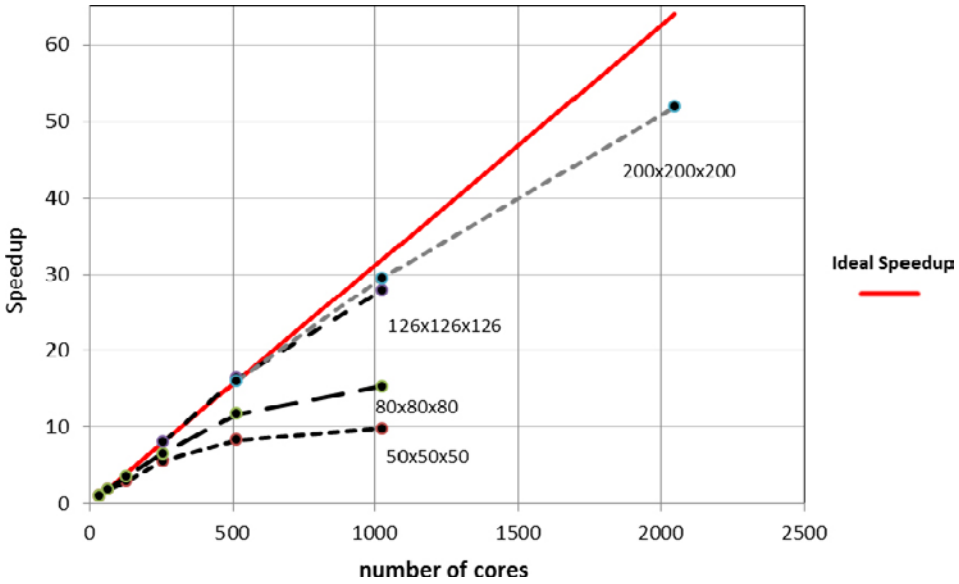

Figure 5-17. *Speedup curve results at MareNostrum (Barcelona Supercomputing Center). The numbers in the lines show the size of the simulation (grid dimensions). The speedup is computed as $S_p=T_{base}/T_p$ where T_p is the execution time on a base process configuration (in this case, T_p is the execution time using 32 cores) and T_p is the run time on an alternative configuration with p processes.*

6 Test case #2: Glacial water penetration and oxygen consumption – sub-mm modelling

6.1 Introduction

In the previous sections we have demonstrated the potential of iDP and PFLOTRAN to solve large-scale, complex and non-linear reactive transport problems. Here, we use iDP to study the potential ingress of oxygenated water in a fracture-matrix system. The numerical study relies on a very ambitious DarcyTools model, which provides a detailed description of the flow and advective transport patterns at the sub-mm scale. In iDP, the same scale is used to describe mineral-water interaction processes as well as homogeneous aqueous oxidation reactions. This section presents a concise discussion of the model and related results. For a more detailed discussion of this modelling case, the reader is referred to Trinchero et al. (2016).

6.2 DarcyTools model

6.2.1 Model domain

Our computational domain, (Figure 6-1), is a block with length 0.5 m, width 0.1 m and height 0.05 m. The block has a flowing fracture with a half aperture of 0.05 mm and the rest of the block constitutes the so called matrix. Any fluid and solute entering in the flowing fracture may penetrate the matrix by both advective and diffusive processes.

It is assumed that the matrix can be described by means of a discrete fracture network. The smallest fractures should represent the high porosity regions between grains. A power law describes the length distribution from this smallest scale up to the largest fractures, which are comparable to the domain in size. This matrix model is described in two unpublished notes, both by U. Svensson.

6.2.2 Computational grid

A Cartesian grid with uniform grid spacing, equal to 0.4 mm, in x and y directions is used. In the z-coordinates direction the grid spacing is 0.0125 mm in the flowing fracture and then expands up to 0.4 mm. Figure 6-2 illustrates the grid. In total about 54 million grid cells build this grid.

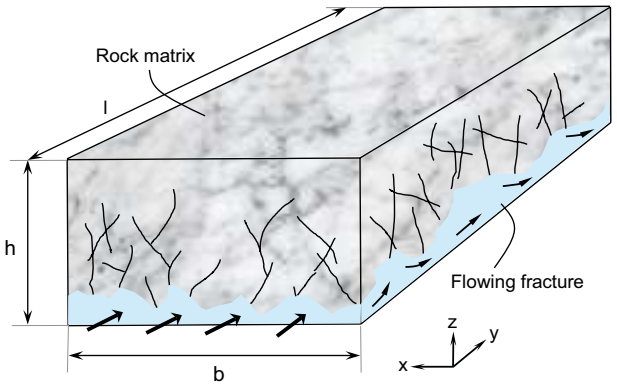

Figure 6-1. *Computational domain (b = 0.1 m, l = 0.5 m and h = 0.05 m).*

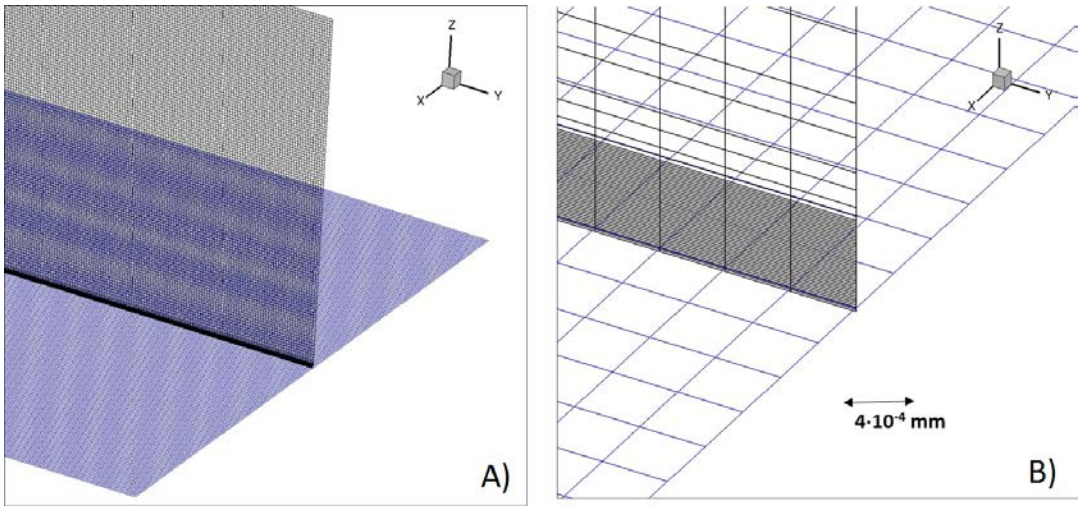

Figure 6-2. *Computational grid. (A) uniform grid is specified in the x and y directions, (B) while an expanding grid size is used in the z-direction.*

6.2.3 Fracture networks

The matrix is generated using two DFNs (Figure 6-3): a grain-scale DFN, with a length-scale interval of 4–5 mm and fracture aperture of 6×10^{-7} m, which represents the small micro-fractures between grains, and a larger scale DFN (5–50 mm with fracture aperture 8×10^{-7} m), which is used to reproduce the fewer larger scale micro-fractures. It is clear that the smallest fractures, representing intergranular porosity, dominate the rock matrix. However, the larger fractures are important for the interaction between the flowing fracture and the matrix. Fracture orientation and spatial centers are statistically independent and follow a Poisson process. For both DFNs, fracture size follows a power law distribution with an exponent equal to 2.6.

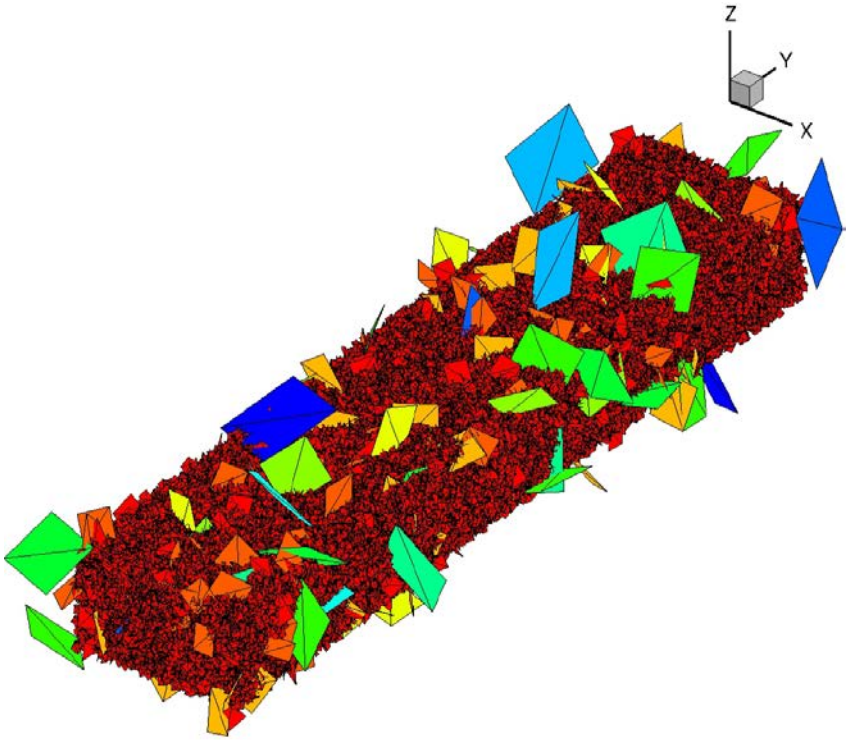

Figure 6-3. *Fracture network in the matrix, coloured by transmissivity (blue most transmissive).*

In DarcyTools, fracture properties are then represented as grid cell conductivities and porosities in the continuum model (Svensson 2001). Figure 6-4 shows the permeability field as represented on the computational grid. The blue parts (i.e. regions with very low permeability) should thus represent grains. It is also clear that larger fractures in contact with the flowing fracture may be important for the penetration of a fluid or solute into the matrix (see Figure 6-8).

The flowing fracture is assumed to be heterogeneous and is generated using a multivariate normal distribution. The mean permeability is set to $1 \times 10^{-10.7}$ m^2 whereas the correlation length and the log_{10} standard deviation are 5 mm and 0.6, respectively. Figure 6-5 illustrates the permeability of the flowing fracture.

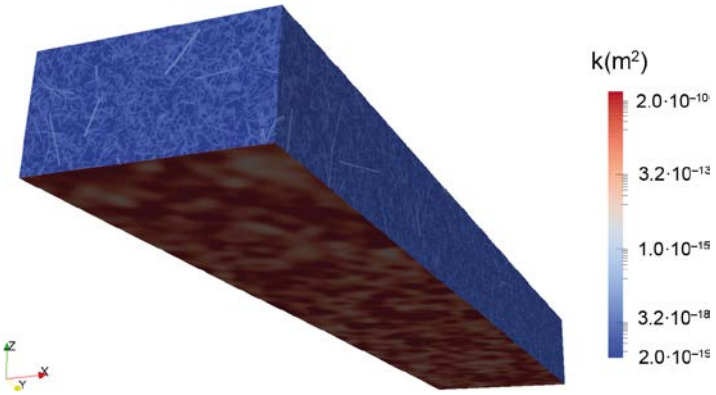

Figure 6-4. *x-component of permeability.*

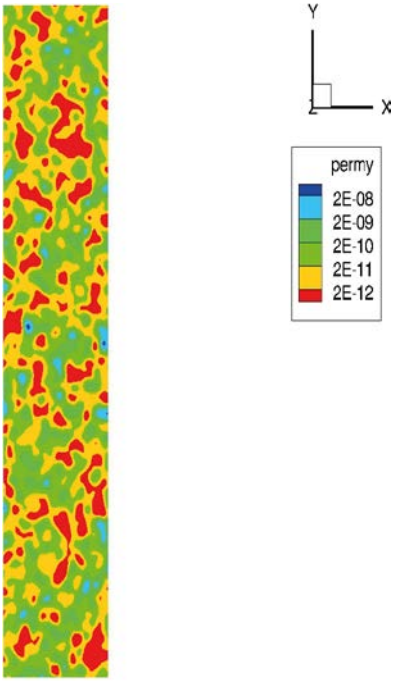

Figure 6-5. *Permeability distribution in the flowing fracture.*

6.2.4 Flow and pressure distributions

Darcy's law is used to solve the groundwater flow problem. Steady state conditions are assumed and pressure is prescribed at the low and high y-boundaries, with a gradient of 49.05 Pa. All other boundaries are of the zero flux type. The pressure distribution in the flowing channel is shown in Figure 6-6. The departure from a linear distribution is of course due to the heterogeneity introduced. The corresponding flow distribution is still more irregular (Figure 6-7). A strong channelling is generated as a consequence of the heterogeneity in the fracture.

In the matrix the flow is orders of magnitudes smaller, but still very irregular (Figure 6-8). In this figure the flow vectors in the flowing fracture have been removed in order to see the flow in and out of the matrix. The Darcy velocities are of the order of 10^{-12} m/s and hence extremely small.

6.2.5 Trajectories

In order to study the interaction between the flowing fracture and the matrix, particles were introduced at the inlet of the flowing fracture. The typical velocity in the flowing fracture was 10^{-6} m/s and the particles were tracked for 10^{5} s. A particle that does not enter the matrix should hence travel about 0.1 m. This is also seen in Figure 6-9.

However, the particles that enter the matrix are perhaps more interesting. Close-ups of some of these are shown in Figure 6-10. The maximum penetration depth, at this time, is about 1 mm.

Trajectories in the *xy*-plane are shown in Figure 6-11. The small irregularities of the trajectories are due to Brownian motion; this when the particles are still in the flowing fracture. The large departures from the mean path, seen in the close-up, are due to excursions into the matrix. For further details about the particle tracking scheme used in these calculations, the reader is referred to the PARTRACK algorithm described in the DarcyTools User's Manual (Svensson et al. 2010).

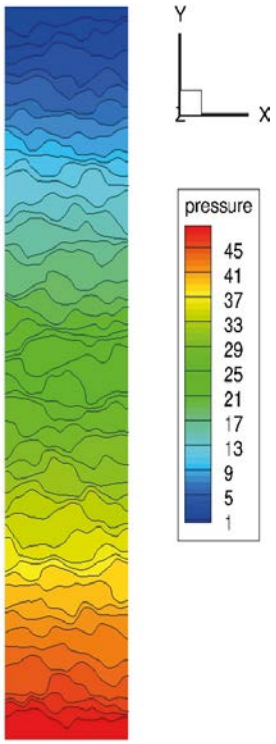

Figure 6-6. *Pressure distribution in the flowing fracture.*

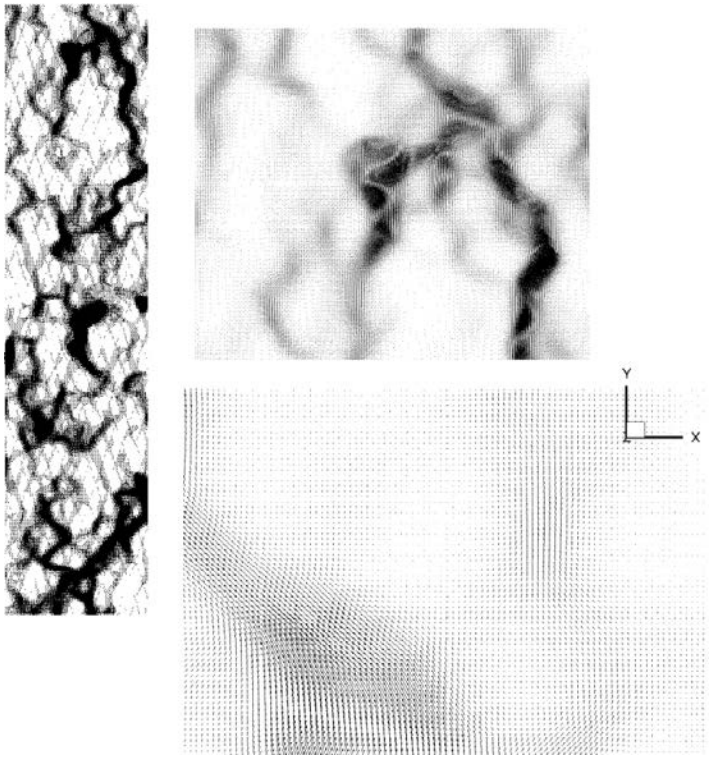

Figure 6-7. *Flow distribution in the flowing fracture. Left figure shows the whole xy-plane, top right figure a close-up at the high-y boundary and bottom –right an even larger magnification. The size of the arrows is proportional to the magnitude of the velocity vector.*

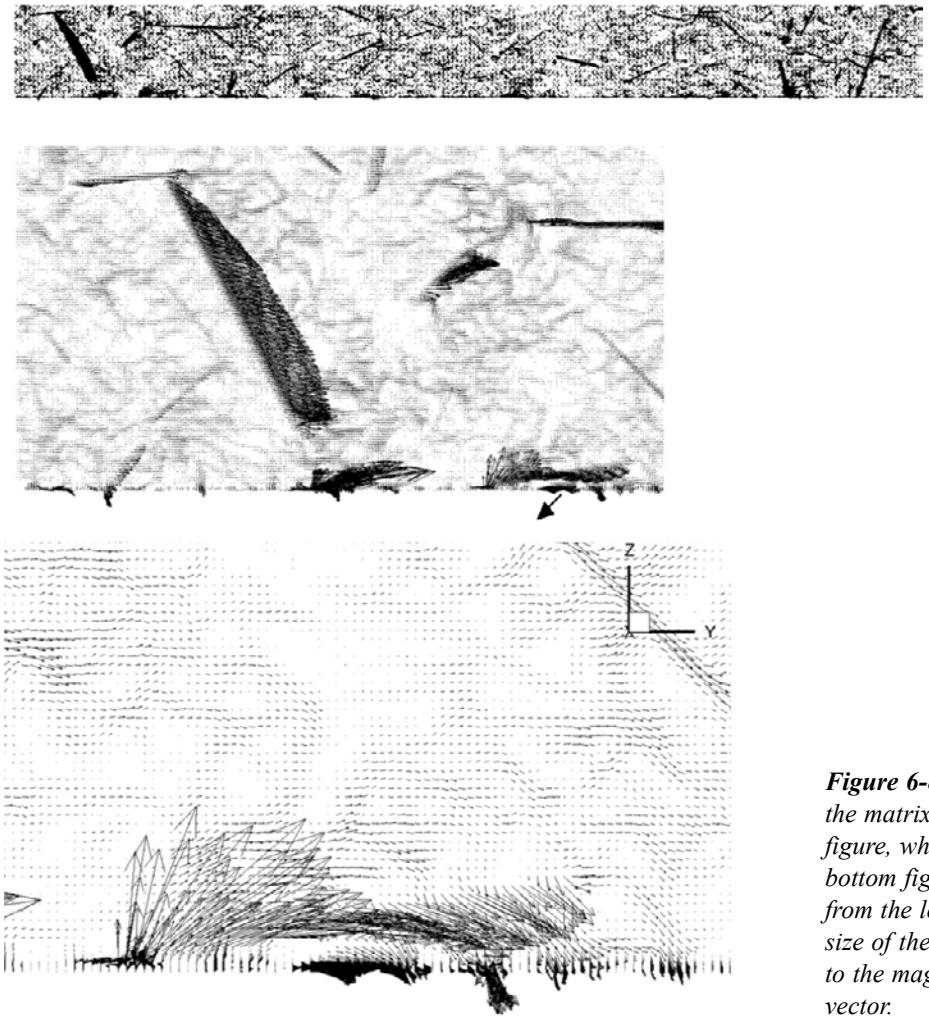

Figure 6-8. *Flow distribution in the matrix in an yz-plane. Top figure, whole plane. Middle and bottom figures are close-ups, from the low-y boundary. The size of the arrows is proportional to the magnitude of the velocity vector.*

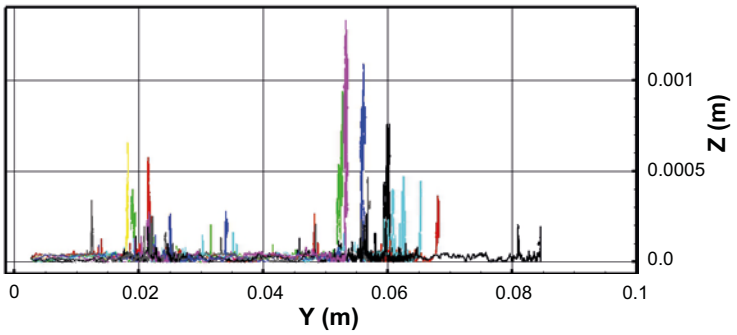

Figure 6-9. *Trajectories of a few particles released at the inlet of the flowing fracture (each particle trajectory is identified by a different colour). Trajectory segments moving perpendicular to the y axis indicate that the particles have penetrated into the rock matrix.*

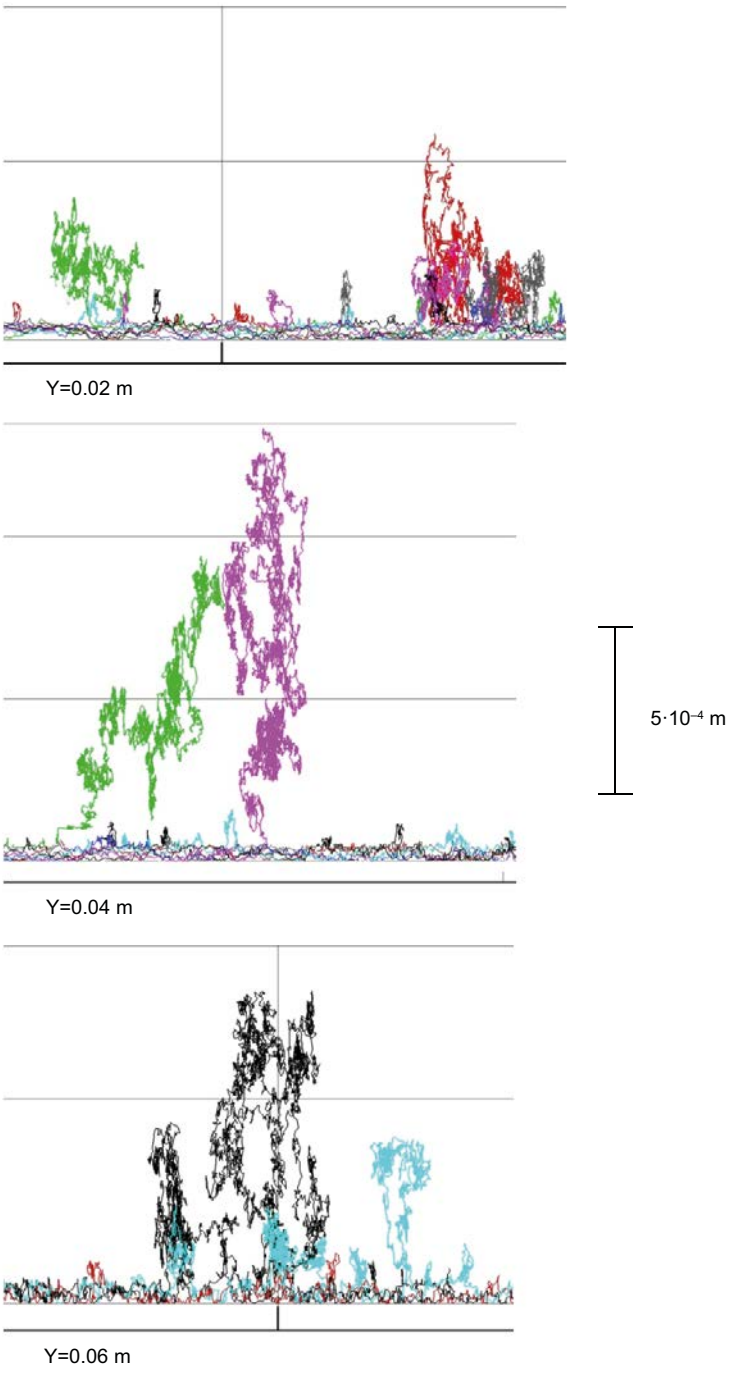

Figure 6-10. *Close-ups of the trajectories shown in Figure 6-9.*

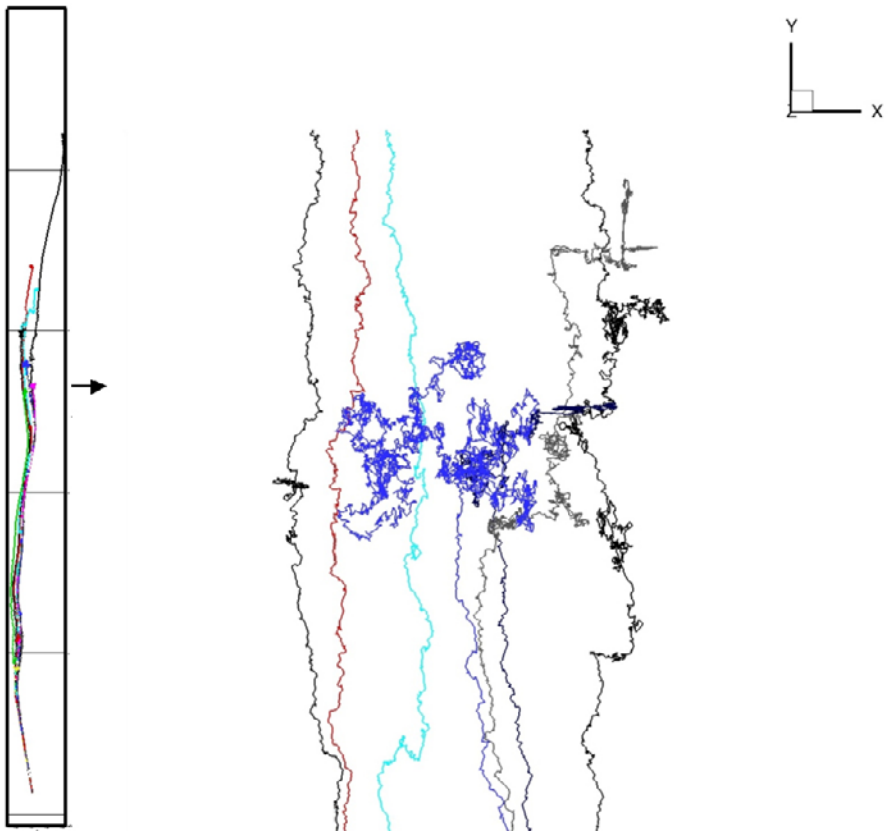

Figure 6-11. *Trajectories shown in an yx-plane. Left figure shows whole plane, right figure is a close-up.*

6.3 iDP: Reactive transport calculations

6.3.1 Model set-up

The model set-up is analogous to that used in the benchmark on "matrix diffusion and oxygen consumption" (Section 4.5); the main difference being that here both the fracture and the matrix are heterogeneous. The chemical composition of the initial and the boundary water is the same specified as in the benchmark test mentioned above. The oxidative boundary water, which has a redox potential pe=12.5, enters the domain from the inlet boundary of the fracture. As specified in Section 6.2.1, solute species entering in the flowing fracture are moved along the fracture by advection and may penetrate the matrix by both advective and diffusive processes (although, in the matrix, transport processes are clearly dominated by diffusion). Biotite and calcite are present in both the fracture and the matrix with a volume fraction of 0.1. The biotite specific surface, its dissolution rate and the parameters used to simulate the homogeneous oxidation of ferrous iron are the same used in Sidborn et al. (2010) and briefly outlined in Section 4.5.

iDP was used to extract Darcy's fluxes and porosity values in a sub-domain of the DarcyTools Model, which consists of 2 975 544 grid cells (102 × 204 × 143). The dimensions of this sub-domain are indicated in Figure 6-12. The average hydraulic gradient in the fracture is aligned along the y axis.

6.3.2 Results

The reactive transport calculation was carried out in Juqueen using 1040 processor cores. Each processor has around 1 GB of memory, meaning that total memory used by the simulation was 1 TB, approximately. The maximum wall-clock time was set to 11 h. Thus, approximately 12 000 h of supercomputing allocation time were consumed. With this allowed wall-clock time, the simulation was run out to 15 days in time, approximately. This simulation time might appear to be rather modest. However it should be considered that this model has singularities (e.g. reaction processes described at a pore scale with a computational grid of millions of element) that result in formidable computational challenges. Moreover, the coupled kinetic reactions and the high groundwater velocities along the fracture (Figure 6-13) introduce additional stiffness to the problem, making necessary the use of a very tight time stepping.

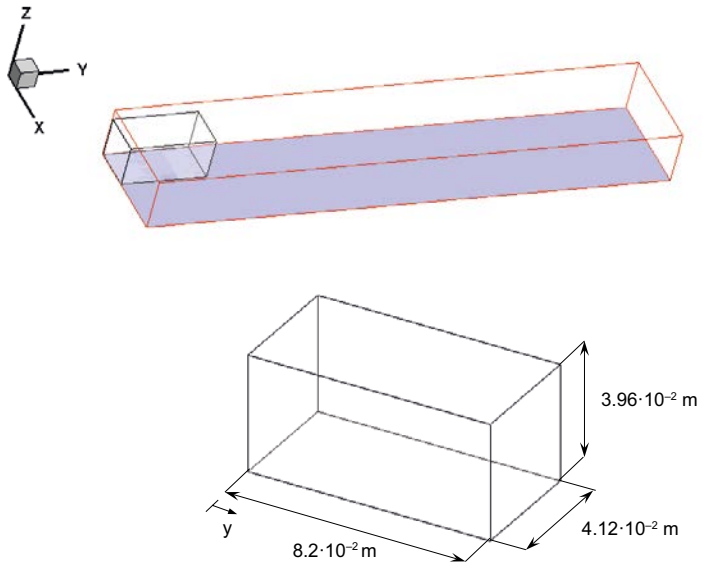

Figure 6-12. *DarcyTools domain (red outlines) and sub-domain used in the iDP calculations (black outlines). The shaded volume indicates the extension of the fracture.*

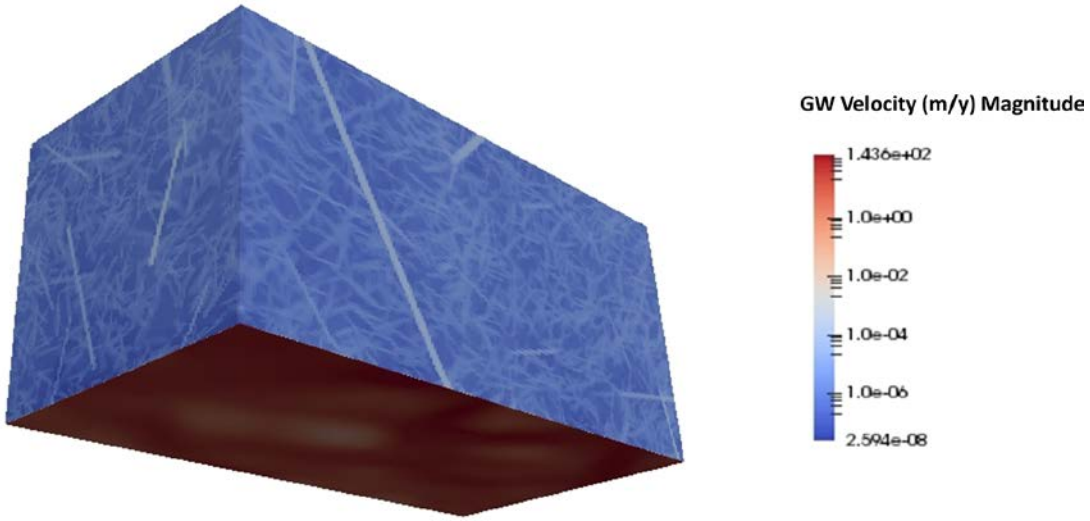

Figure 6-13. *Groundwater velocity (magnitude) in the extracted sub-domain.*

The displacement of the conservative tracer along the fracture and its progressive penetration into the matrix is shown in Figure 6-14 (natural scale) and Figure 6-15 (log scale). From a careful visual inspection of Figure 6-15 one can see the influence of the more conductive regions of the matrix (see also the distribution of velocity in the matrix – Figure 6-13), which constitute a complex network of preferential patterns.

The evolution of the redox potential is shown in Figure 6-16. Consistently with what already observed by Sidborn et al. (Sidborn et al. 2010), after a quick (and modest) penetration of the redox perturbation front, a steady-state profile is reached. In this steady state, advective processes, which "carry" oxygen from the inlet of the fracture, are perfectly counterbalanced by geochemical reactions, which oxidize iron and consume oxygen. By comparing Figure 6-15 and Figure 6-16 it might appear that redox front penetrates into the matrix deeper than the conservative tracer. This is indeed not true and this false perspective is related to the fact that the redox potential, *pe*, is a logarithmic measure (i.e. the logarithm of the electron concentration in the solution). In Figure 6-16, the colour scale of pe spans from a minimum of −4 to a maximum of 12.5, which means more than sixteen orders of magnitudes in a natural scale of electron concentration. This range is much higher than that used to show the tracer distribution.

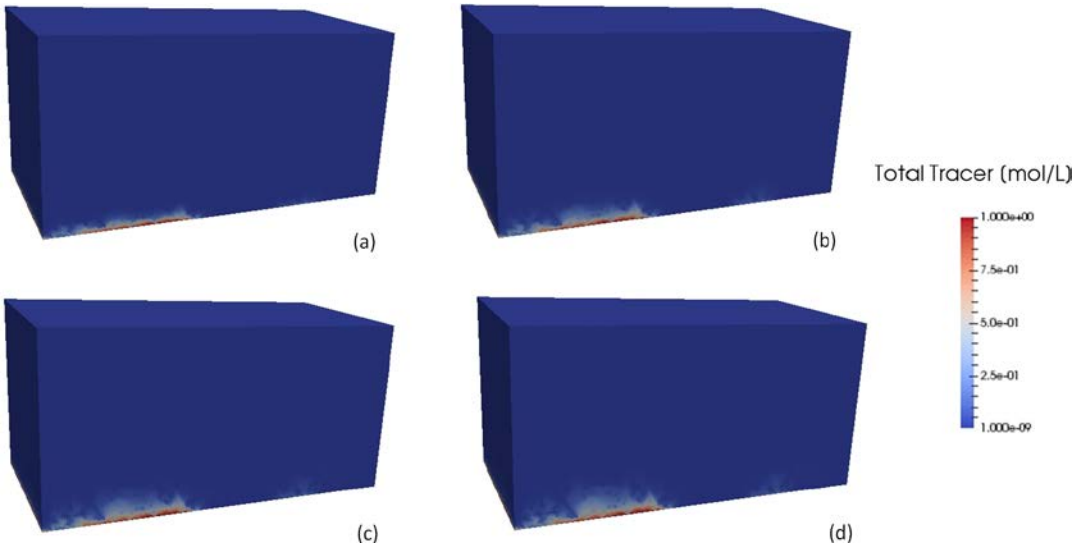

Figure 6-14. *Tracer concentration at (a) 2.3 day, (b) 6.6 day, (c) 10.7 day and (d) 14.6 day. Tracer distribution is plotted in natural scale.*

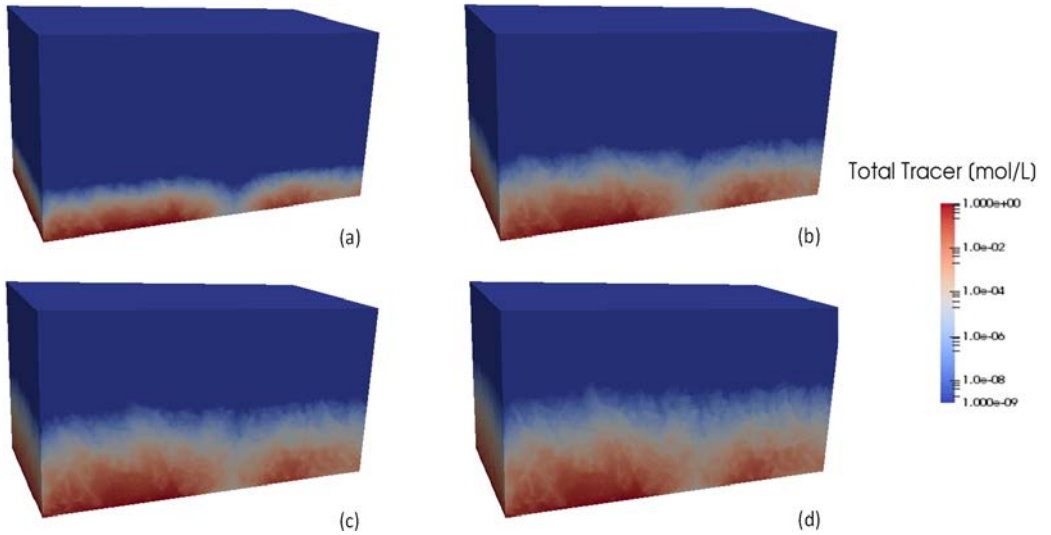

Figure 6-15. *Tracer concentration at (a) 2.3 day, (b) 6.6 day, (c) 10.7 day and (d) 14.6 day. Tracer distribution is plotted in log-scale.*

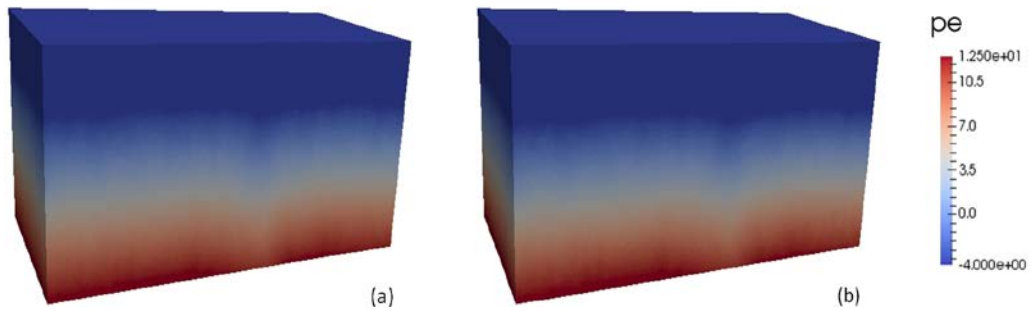

Figure 6-16. *pe distribution at (a) 2.3 day and (d) 14.6 day.*

7 Conclusions

DarcyTools and PFLOTRAN are two powerful last-generation simulators of groundwater flow in fractured rocks and reactive transport processes, respectively. These two simulation frameworks have been combined to form a very powerful integrated simulation framework for reactive-transport in crystalline rocks framework, denoted as iDP (standing for interface between DarcyTools and PFLOTRAN). The interface is written in a modular, cross-platform and extensible way using Fortran and Python.

The combined reactive transport framework has been verified, by means of five different benchmark cases. In all these benchmarks, a good agreement is observed between the results computed by iDP and the independent numerical and analytical solutions.

iDP has been used to solve three case studies, which are of interest for SKB: (1) the degradation of cement grout applied in the Forsmark spent fuel repository, (2) the macroscopic penetration of glacial water and oxygen consumption in fractured crystalline rocks, and (3) the microscopic penetration of glacial water and oxygen consumption in rock matrix.

The test cases have been run successfully on two world-class supercomputers: the MareNostrum supercomputer at the Barcelona Supercomputing Center, in Spain, and the JUQUEEN supercomputer at the Jülich Supercomputer Centre, in Germany. The performance and scalability of iDP has been demonstrated up to 8000 cores in both supercomputers.

The BRIDGE project has demonstrated that High Performance Reactive Transport Modelling is a feasible technological option. This represents a step forward in the integration of hydrogeological and geochemical processes. In particular, a broad range of reactive transport problems of interest for the safety analysis of a deep geological repository can be now studied in an integrated fashion, without making use of problem oversimplifications.

References

SKB's (Svensk Kärnbränslehantering AB) publications can be found at www.skb.se/publications.

Balay S, Gropp W D, McInnes L C, Smith B F, 1997. Efficient management of parallelism in object-oriented numerical software libraries. In Arge E, Bruaset A M, Langtangen H P (eds). Modern software tools for scientific computing. Boston: Birkhäuser, 163–202.

Bethke C, 1994. The geochemist's workbench: a user's guide to Rxn, Act2, Tact, React, and Gtplot. University of Illinois.

Clement T P, 1997. RT3D: a modular computer code for simulating reactive multi-species transport in 3-dimensional groundwater aquifers. PNNL-SA-28967, Pacific Northwest National Laboratory, Richland, Washington.

Folk M, Cheng A, Yates K, 1999. HDF5: a file format and I/O library for high performance computing applications. Proceedings of Supercomputing 99, 5–33.

FSF, 2004. General Public License. Retrieved from http://www.fsf.org

Hammond G E, Lichtner P C, 2010. Field-scale model for the natural attenuation of uranium at the Hanford 300 Area using high-performance computing. Water Resources Research 46. doi:10.1029/2009WR008819

Hammond G E, Lichtner P C, Rockhold M L, 2011. Stochastic simulation of uranium migration at the Hanford 300 Area. Journal of Contaminant Hydrology 120, 115–128.

Hammond G E, Lichtner P C, Mills R T, 2014. Evaluating the performance of parallel subsurface simulators: an illustrative example with PFLOTRAN. Water Resources Research 50, 208–228.

Henderson A, Ahrens J, Law C, 2004. The ParaView guide. Clifton Park, NY: Kitware.

Kolditz O, Görke U-J, Shao H, Wang W (eds), 2012. Thermo-hydro-mechanical-chemical processes in porous media: benchmarks and examples. Berlin: Springer.

Lichtner P C, 1996. Modeling reactive flow and transport in natural systems. In Ottonello G, Marini L (eds). Proceedings of the Rome Seminar on Environmental Geochemistry, Castelnuovo di Porto, Italy, 22–26 May 1996. Pisa: Pacini Editore, 5–72.

Palandri J L, Kharaka Y K, 2004. A compilation of rate parameters of water-mineral interaction kinetics for application to geochemical modeling. Open File Report 2004-1068, U.S. Geological Survey, Denver, Colorado.

Parkhurst D L, Appelo C A J, 2013. Description of input and examples for PHREEQC version 3: a computer program for speciation, batch-reaction, one-dimensional transport, and inverse geochemical calculations. Techniques and Methods 6–A43, U.S. Geological Survey, Denver, Colorado.

Parkhurst D L, Kipp K L, Engesgaard P, Charlton S R, 2004. PHAST: a program for simulating ground-water flow, solute transport, and multicomponent geochemical reactions. Techniques and Methods 6-A8, U.S. Geological Survey, Denver, Colorado.

Remy N, Boucher A, Wu J, 2009. Applied geostatistics with SGeMS: a user's guide. Cambridge: Cambridge University Press.

Saaltink M W, Battle F, Ayora C, Carrera J, Olivella S, 2004. RETRASO, a code for modeling reactive transport in saturated and unsaturated porous media. Geologica Acta 2, 235–251.

Salas J, Gimeno M J, Auqué L, Molinero J, Gómez J, Juárez I, 2010. SR-Site –hydrogeochemical evolution of the Forsmark Site. SKB TR-10-58, Svensk Kärnbränslehantering AB.

Sidborn M, Sandström B, Tullborg E-L, Salas J, Maia F, Delos A Molinero J, Hallbeck L, Pedersen K, 2010. SR-Site: Oxygen ingress in the rock at Forsmark during a glacial cycle. SKB TR-10-57, Svensk Kärnbränslehantering AB.

Sidborn M, Marsic N, Crawford J, Joyce S, Hartley Lidiart A, de Vries L M, Maia F, Molinero J, Svensson U, Vidstrand P, Alexander R, 2014. Potential alkaline conditions for deposition holes of a repository in Forsmark as a consequence of OPC grouting. Revised final report after review. SKB R-12-17, Svensk Kärnbränslehantering AB.

SKB, 2008. Site description of Forsmark at completion of the site investigation phase. SDM-Site Forsmark. SKB TR-08-05, Svensk Kärnbränslehantering AB, Stockholm.

SKB, 2011. Long-term safety for the final repository for spent nuclear fuel at Forsmark. Main report of the SR-Site Project. SKB TR-11-01, Svensk Kärnbränslehantering AB.

Soler J M, 2010. Reactive transport modeling of grout–rock interaction at the ONKALO site. Posiva Working Report 2010-73, Posiva Oy, Finland.

Steefel C I, 2008. CrunchFlow: software for modelling multicomponent reactive flow and transport. Berkeley, CA: Earth Sciences Division, Lawrence Berkeley National Laboratory.

Stumm W, Lee G F, 1961. Oxygenation of ferrous iron. Industrial & Engineering Chemistry 53, 143–146.

Svensson U, 2001. A continuum representation of fracture networks. Part I: Method and basic test cases. Journal of Hydrology 250, 170–186.

Svensson U, Ferry M, 2014. DarcyTools: a computer code for hydrogeological analysis of nuclear waste repositories in fractured rock. Journal of Applied Mathematics and Physics 2, 365–383.

Svensson U, Ferry M, Kuylenstierna H-O, 2010. DarcyTools version 3.4 – concepts, methods and equations. SKB R-07-38, Svensk Kärnbränslehantering AB.

Trinchero P, Molinero J, Deissmann G, Svenson U, Gylling B, Ebrahimi H, Hammond G, Bosbach D, Puigdomenech I, 2016. Implications of grain-scale mineralogical heterogeneity for radionuclide transport in fractured media. Transport in Porous Media. doi:10.1007/s11242-016-0765-0

van der Lee J, De Windt L, Lagneau V, Goblet P, 2003. Module-oriented modeling of reactive transport with HYTEC. Computers & Geosciences 29, 265–275.

Wylie B J N, Geimer M, 2011. Large-scale performance analysis of PFLOTRAN with Scalasca. In Proceedings of 53rd CUG Meeting, Fairbanks, Alaska, 23–26 May 2011. Cray User Group, Inc.

Xu T, Sonnenthal E, Spycher N, Pruess K, 2006. TOUGHREACT – a simulation program for non-isothermal multiphase reactive geochemical transport in variably saturated geologic media: applications to geothermal injectivity and CO_2 geological sequestration. Computers & Geosciences 32, 145–165.

Arkitektkopia AB, Bromma, 2016

SKB is tasked with managing Swedish nuclear and radioactive waste in a safe way.

skb.se