



Strål
säkerhets
myndigheten

Swedish Radiation Safety Authority

Authors: Peter Robinson
Claire Watson

Research

2011:11

Handling Interfaces and Time-varying
Properties in Radionuclide
Transport Models

SSM perspective

Background

The application for a KBS-3 type spent nuclear fuel repository will be supported by a post-closure safety assessment (SR-Site) which includes modelling of radionuclide transport from an underground source term. In order to prepare for the review of the oncoming license application the Swedish Radiation Safety Authority (SSM) has performed research and development projects in the area of performance assessment modelling during recent years. Independent modelling capability has been established, both at SSM and with external consultants. This has included the development of approaches and models for consequence analysis (radionuclide transport) that can be used to support the review of a spent nuclear fuel repository.

Objectives

In this project the following three areas were investigated:

1. The Qeq concept is an important part of SKB's current methodology to model transport resistance at the buffer/fracture interface for radionuclide transport. The regulatory review of SKB's earlier safety analysis shows that the transport resistance offered by the buffer/fracture interface is a critical component for radionuclide transport. The Qeq concept is reviewed and calculations undertaken to explore whether this remains valid in situations where heterogeneity or spalling are present.
2. Some assessment calculations are undertaken to investigate the potential for changes in transport properties with time due to the effects of glacial episodes that could affect radionuclide transport to the surface.
3. SKB has developed a new code MARFA for handling spatially varying properties and time-varying flows relating to radionuclide transport problems for the far field. This code may be used in parallel with the older geosphere transport approach (FARF31) in SR-Site, therefore, the code is reviewed and analysed by Quintessa's independent model.

Results

Quintessa's QPAC code has been used to investigate the Qeq approach. The conclusions from this simulation study are the following. The basic approach to calculating Qeq values is sound, however, narrow channels could lead to the same release as larger fractures with the same pore velocity, so a channel enhancement factor of $\sqrt{10}$ should be considered. A spalling zone that increases the area of contact between flowing water and the buffer has the potential to increase the release significantly.

Quintessa's AMBER software has been used to explore the effects of glacial episodes on radionuclide transport with time-varying properties. The simulation results show that for both single and multiple glacial episodes the time-dependency of model parameters did not result in

much change to the calculated peak fluxes to the biosphere. These conclusions are preliminary and could be changed if different radionuclides are important in the SR-Site assessment.

A detailed review of MARFA code and the associated documentation has been undertaken. New semi-analytic methods have been developed in order to provide a means of checking the accuracy of MARFA calculations. The strengths and weaknesses of MARFA code have been explored in detail.

Project information

Project manager: Shulan Xu

Project reference: SSM 2010/829



Strål
säkerhets
myndigheten

Swedish Radiation Safety Authority

Authors: Peter Robinson and Claire Watson
Quintessa Ltd. UK

2011:11

Handling Interfaces and Time-varying
Properties in Radionuclide
Transport Models

Date: December 2010

Report number: 2011:11 ISSN: 2000-0456

Available at www.stralsakerhetsmyndigheten.se

This report concerns a study which has been conducted for the Swedish Radiation Safety Authority, SSM. The conclusions and viewpoints presented in the report are those of the author/authors and do not necessarily coincide with those of the SSM.

Summary

This report documents studies undertaken by Quintessa during 2010 in preparation for the SR-Site review that will be initiated by SSM in 2011. The studies relate to consequence analysis calculations, that is to the calculation of radionuclide release and transport if a canister is breached. A sister report documents modelling work undertaken to investigate the coupled processes relevant to copper corrosion and buffer erosion.

The Q_{eq} concept is an important part of SKB's current methodology for radionuclide transport using one-dimensional transport modelling; it is used in particular to model transport at the buffer/fracture interface. Quintessa's QPAC code has been used to investigate the Q_{eq} approach and to explore the importance of heterogeneity in the fracture and spalling on the deposition hole surface. The key conclusions are that:

- ▲ The basic approach to calculating Q_{eq} values is sound and can be reproduced in QPAC.
- ▲ The fracture resistance dominates over the diffusive resistance in the buffer except for the highest velocity cases.
- ▲ Heterogeneity in the fracture, in terms of uncorrelated random variations in the fracture aperture, tends to reduce releases, so the use of a constant average aperture approach is conservative.
- ▲ Narrow channels could lead to the same release as larger fractures with the same pore velocity, so a channel enhancement factor of $\sqrt{10}$ should be considered.
- ▲ A spalling zone that increases the area of contact between flowing water and the buffer has the potential to increase the release significantly and changes the functional dependence of Q_{eq}^{frac} on the flowing velocity.

Quintessa's AMBER software has previously been used to reproduce SKB's one-dimensional transport calculations and AMBER allows the use of time-varying properties. This capability has been used to investigate the effects of glacial episodes on radionuclide transport. The main parameters that could be affected are sorption coefficients and flow rates. For both single and multiple glacial episodes the time-dependency of model parameters did not result in much change to the calculated peak fluxes to the biosphere. This is because fluxes are calculated to be dominated by the poorly-sorbed radionuclides ^{129}I and ^{36}Cl . However, a small increase (less than an order of magnitude) in the overall flux contributed from the radium decay chain, which is important for long timescales, was calculated during the phase when the multiple glacial episodes are occurring. These conclusions are

preliminary and could be changed if different radionuclides are important in the SR-Site assessment.

The known shortcomings of its one-dimensional radionuclide transport modelling capability has led SKB to fund the development of the MARFA code, which it is understood will be used in addition to the existing methodology in SR-Site. A detailed review of this code and the associated documentation has been undertaken. New semi-analytic methods have been developed in order to provide a means of checking the accuracy of MARFA calculations. The key conclusions are that:

- ▲ MARFA can handle large networks in practicable run times, generally works well for single radionuclides or short chains and accurately handles advective systems where matrix diffusion effects are dominant.
- ▲ The code has, however, a number of important limitations. In particular, it is unable to handle long decay chains with short-lived radionuclides and calculations immediately after flow rate changes can be inaccurate.
- ▲ The documentation and Quality Assurance are poor. A large number of errors have been found in the User Guide and the associated test cases do not adequately test the use of the code for the anticipated applications.
- ▲ These limitations bring into question the code's suitability (in its present form) for performance assessment calculations for a deep radioactive waste repository.

Content

1	Introduction	1
2	The Fracture-buffer Interface	2
	2.1 Q_{eq} Formulae	2
	2.2 Modelling Approach	4
	2.3 Reference Calculations	9
	2.4 Fracture Heterogeneity	13
	2.5 Spalling	15
	2.6 Conclusions	16
3	Time-dependence in Radionuclide Transport Calculations	18
	3.1 Background	18
	3.2 Conceptual Model	18
	3.3 Calculations	26
	3.4 Conclusions	31
4	Review and Testing of MARFA	32
	4.1 General Overview	33
	4.2 Documentation	34
	4.3 The Code	35
	4.4 Algorithms	39
	4.5 Documented Tests	47
	4.6 Additional Tests	56
	4.7 MARFA 3.3	63
	4.8 Conclusions	64
5	Overall Conclusions	66
	References	68
	Appendix A: Nomenclature	71
	Appendix B: Semi-analytic Solutions for Radionuclide Transport with Piecewise-constant Properties	73

1 Introduction

This report documents studies undertaken during 2010 in preparation for the SR-Site review.

The studies described here relate to consequence analysis calculations, that is to the calculation of radionuclide release and transport if a canister is breached. A sister report [1] documents modelling work undertaken to look at the coupled issues of copper corrosion and buffer erosion.

The studies into some aspects of consequence analysis follow on from work in previous years ([2] and [3]). The objective of this work is to prepare for review of SR-Site by developing an understanding of the key issues and of SKB's assessment approach.

In this year's work, three aspects were examined.

The first area studied is the important interface between the buffer and a flowing fracture. SKB's use of the Q_{eq} concept is reviewed and calculations undertaken to explore whether this remains valid in situations where heterogeneity or spalling are present. This is documented in Section 2.

Secondly, some assessment calculations have been undertaken to investigate the potential for changes in transport properties with time to affect radionuclide transport to the surface. In particular, the importance of near-surface geochemical changes is considered and documented in Section 3.

Finally, the MARFA code is reviewed in some detail. This code may be used in parallel with the older geosphere transport approach (FARF31) and it is important to understand its strengths and weaknesses. The review is documented in Section 4 and additional mathematical details are given in Appendix B.

2 The Fracture-buffer Interface

Radionuclides released from a breached canister while the buffer remains intact will move through the buffer by diffusion. In order to enter the geosphere and ultimately reach the surface, they must enter flowing water. Several possible routes have been identified [4], and release to a fracture that intersects the deposition hole is an important route.

Release through this route is constrained by the nature of the interface between the buffer and the fracture. First, the fracture has a small aperture and diffusion to the small area that this implies is limited. Second, the entry of radionuclides into the flowing fracture water requires that they diffuse away from the interface.

SKB developed an approach to handling transport across this interface [5] which has been used in subsequent assessments and which can be expected to be used in SR-Site. This leads to the specification of an equivalent flow, Q_{eq} , at the interface. A review of the approach employed was reported in [6] where the derivations were found to be correct but concern was expressed that the application of the formulae in probabilistic cases needed to be handled carefully. In SR-Can, the potential for spalling was recognised and an adjustment was made to Q_{eq} to allow for this. This was reviewed in [2] and found to be provisional and speculative and not clearly documented. The approach is described as conservative in [7].

The SKB approach uses an idealised geometry. In particular, the fracture is open with a constant aperture. In reality, it is thought that fractures will have highly variable apertures and may be channelled. For the main geosphere transport calculations in SR-Can, SKB applied a channelling factor of 10 to the F-factors calculated in their flow models (implying that flow is largely confined to 10% of the fracture area), but no similar modification was discussed in relation to the fracture-buffer interface.

The purpose of the current study is to confirm that the basic Q_{eq} approach is sound and to explore the impacts of heterogeneity in the fracture and spalling on the deposition hole surface.

2.1 Q_{eq} Formulae

The Q_{eq} approach to transport across interfaces can be explained as follows.

Given a steady driving concentration (for example at the canister surface), C_0 (mol m^{-3}), if the steady release rate from the part of the system of interest

(e.g. the buffer and the fracture near the deposition hole) is Φ (mol y^{-1}) then this is equivalent to a flow rate Q_{eq} ($\text{m}^3 \text{y}^{-1}$) where

$$\Phi = Q_{eq} C_0. \quad (2.1.1)$$

Where there are two sections in sequence, an overall Q_{eq} can be calculated to be given by

$$\frac{1}{Q_{eq}} = \frac{1}{Q_{eq}^{(1)}} + \frac{1}{Q_{eq}^{(2)}}. \quad (2.1.2)$$

This is clear if Q_{eq} is thought of as a conductance with its reciprocal being a resistance.

For the specific case of interest, the two components are the resistance on the buffer side and the resistance on the fracture side. Writing the interface concentration as C_{int} , with the concentration in the fracture far from the deposition hole being zero, we write

$$\Phi = Q_{eq} C_0 = Q_{eq}^{buff} (C_0 - C_{int}) = Q_{eq}^{frac} C_{int}, \quad (2.1.3)$$

From which it is easy to derive the reciprocal relationship between the various Q_{eq} values:

$$\frac{1}{Q_{eq}} = \frac{1}{Q_{eq}^{buff}} + \frac{1}{Q_{eq}^{frac}}. \quad (2.1.4)$$

It is also clear from physical considerations that the smaller of the Q_{eq} component values determines the overall value and that the interface concentration will be closer to the concentration on the side with the larger Q_{eq} .

A detailed discussion was given in [6]; here we note the final results.

For the buffer

$$Q_{eq}^{buff} = \frac{4\pi R D_{eff}}{F/b} \quad (2.1.5)$$

where R is the distance from the canister to the fracture and F/b is calculated from a regression fit as

$$F/b = 0.9 - 1.466 \log_{10}(b/a) + 1.58 \log_{10}(R/a) \quad (2.1.6)$$

where a is half the relevant height of the canister where the concentration is fixed. The relevant height might be for a physical area where there is a breach in the canister, or simply half the distance between two fractures, where it represents the zone that would feed into a particular fracture.

For the fracture,

$$Q_{eq}^{frac} = 16b \sqrt{\frac{D_w L v}{\pi}} \quad (2.1.7)$$

where D_w is the pore-water diffusion coefficient for the fracture water (equal to the effective diffusion coefficient since the fracture is assumed to be open), L is the contact distance between the flowing water and the buffer, b is half the fracture aperture and v is pore-water velocity. Note that a factor of 2 has been added to this formula to account for the full fracture aperture.

Various geometrical approximations are made in deriving these results.

For the buffer, it is assumed that the same concentration applies all around the interface, which is clearly not the case as there is a particular flow direction that will lead to different upstream and downstream concentrations. Also, the radial geometry is ignored – in [7] it is stated that this leads to about a 5% error. The main effect is that the circumference of the canister is treated as if it were the same as the outer buffer circumference, so it is likely that this overstates the Q_{eq} and so is conservative. The analysis was also undertaken for a set of identical fractures, so that symmetry conditions could be imposed.

For the fracture, the key approximation is that the interface is taken to be linear in Cartesian coordinates with a constant flow velocity past it. In fact, of course, it is cylindrical and the flow will vary significantly along its length (the fastest flow being half way around between the upstream and downstream directions).

Our objective is to check that these approximations do not give significantly incorrect values for Q_{eq} .

2.2 Modelling Approach

The approach used here is to model the deposition holes and a fracture that intersects it numerically. The QPAC software [8] was used, modelling Darcy flow (in the fracture) and tracer transport (advection and diffusion of a marker species) in the buffer and fracture.

In order to reduce the size of the model, a symmetry plane was used along the direction of flow. A unit concentration is imposed on a band of the

canister, extending above and below the fracture plane. This might represent a region where a breach has occurred, and so matches one of the interpretations of the “a” parameter given above. All other concentrations are then relative to this.

The physical dimensions and layout are illustrated in Figure 1.

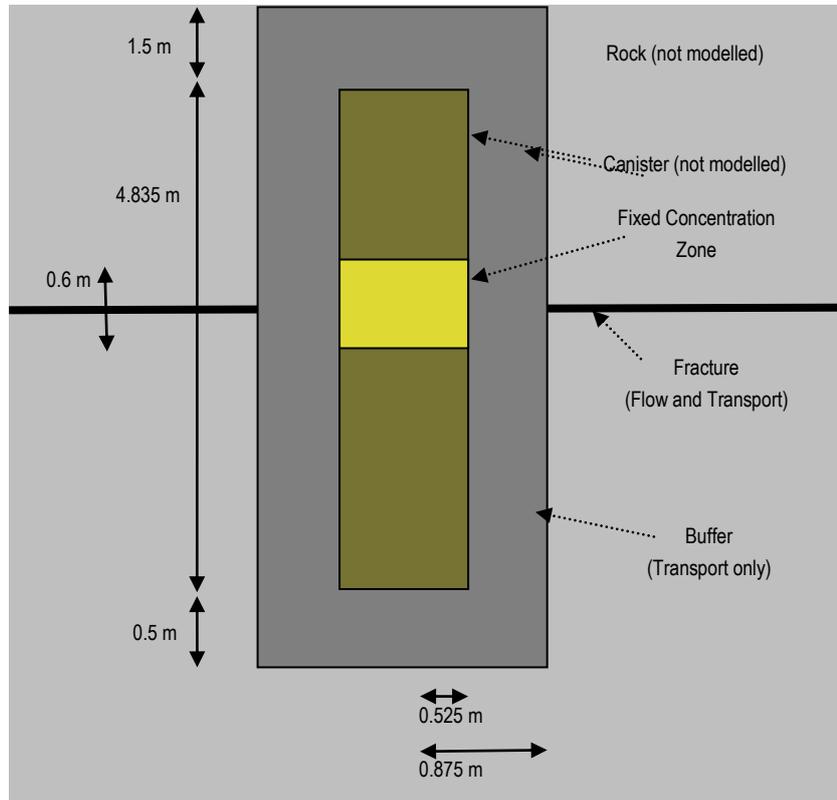


Figure 1: Schematic cross-section through the modelled region, showing the key features and dimensions

Because the interface is narrow, it is important to have sufficient grid cells near it. Convergence tests were run and a grid was selected which gave results that changed by less than 1% compared to a grid with half the number of cells.

The grid is defined in cylindrical coordinates as a series of uniform or geometric sections in each coordinate. Table 1 and Table 2 give the azimuthal and radial grid sizes respectively. In the angular direction, 18 cells are used each with an angle of 10 degrees.

The buffer and fracture were treated as separate subsystems, with a Joiner subsystem connecting them and providing the required continuity of concentrations and fluxes.

Table 1 Grid cell sizes in the z direction

From	To	Type	Number
Bottom of buffer	Bottom of Can	Uniform	3
Bottom of Can	Bottom of Fixed Concentration Zone	Uniform	4
Bottom of Fixed Concentration Zone	Bottom of Fracture Plane	Geometric	12, with last equal to fracture aperture
Bottom of Fracture Plane	Top of Fracture Plane		1
Top of Fracture Plane	Top of Fixed Concentration Zone	Geometric	12
Top of Fixed Concentration Zone	Top of Can	Uniform	4
Top of Can	Top of Buffer	Uniform	3

Table 2 Grid cell sizes in the r direction

From	To	Type	Number
Centre axis	Can radius	Uniform	3
Can radius	Buffer Radius	Geometric	15, with last equal to half fracture aperture
Buffer Radius	Near buffer radius (0.9 m)	Geometric	10, with first equal to fracture aperture
Near buffer radius	Edge of model (10 m)	Geometric	10, with first equal to 2 cm

Because of the small size of many of the cells, it is hard to visualise the grid in spatial coordinates. Instead, Figure 2 shows the grid in “index coordinates”, i.e. each cell is treated as having the same size. The zones are colour coded to show how they correspond to the real geometry. The view is from above, looking down at an angle. Figure 3 shows the buffer grid in spatial coordinates, noting that the cell boundaries are not all visible in this figure.

Figure 4 and Figure 5 show the fracture grid in the same ways. Note that the interfaces are shown in Figure 5 – these were hidden for the buffer grid. The near buffer zone is used for the heterogeneity tests.

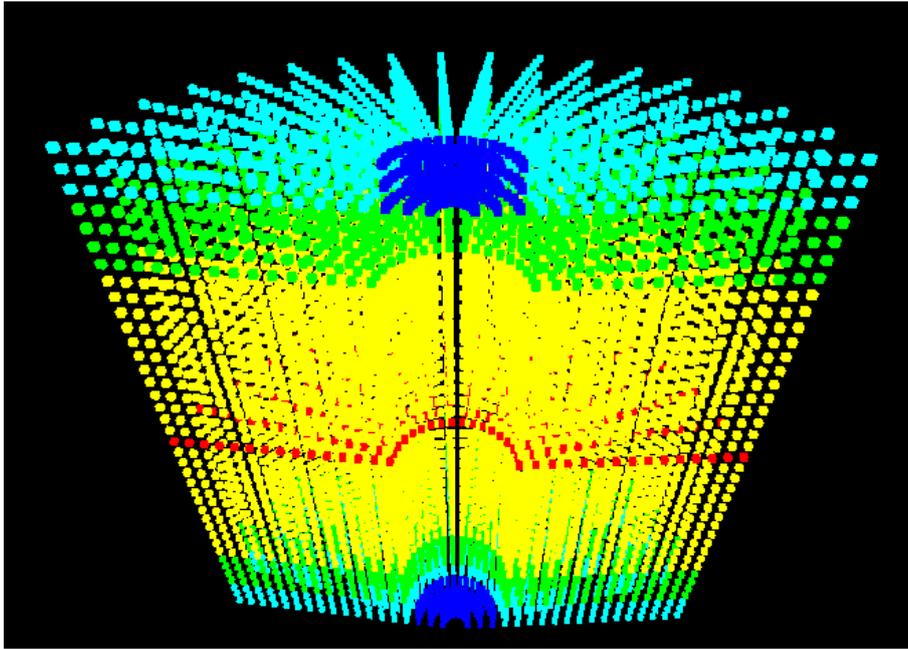


Figure 2: The buffer grid showing the different zones in index coordinates

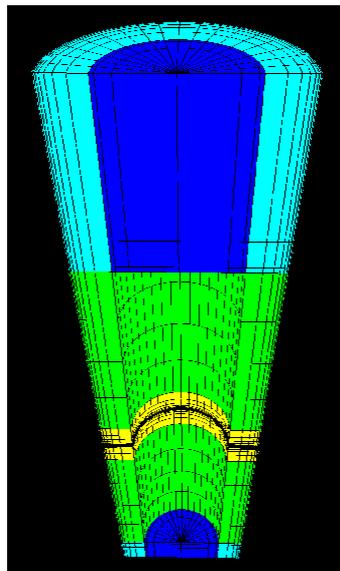


Figure 3: The buffer grid showing the different zones in spatial coordinates

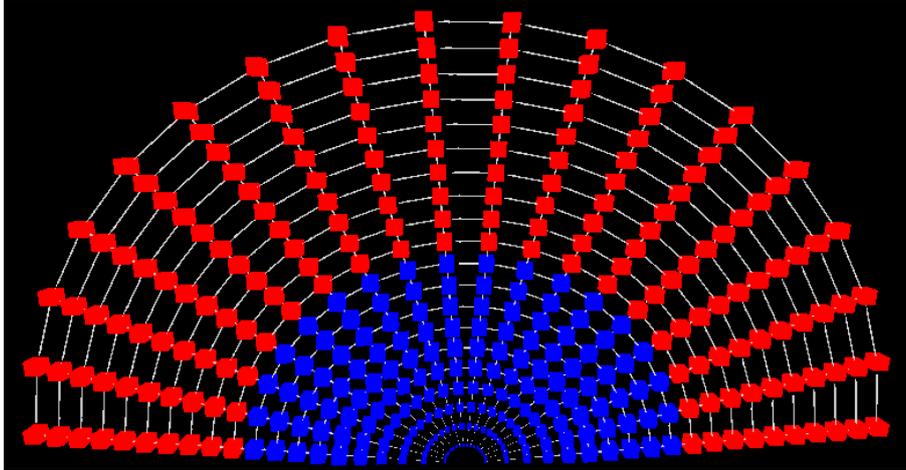


Figure 4: The fracture grid showing the different zones in index coordinates

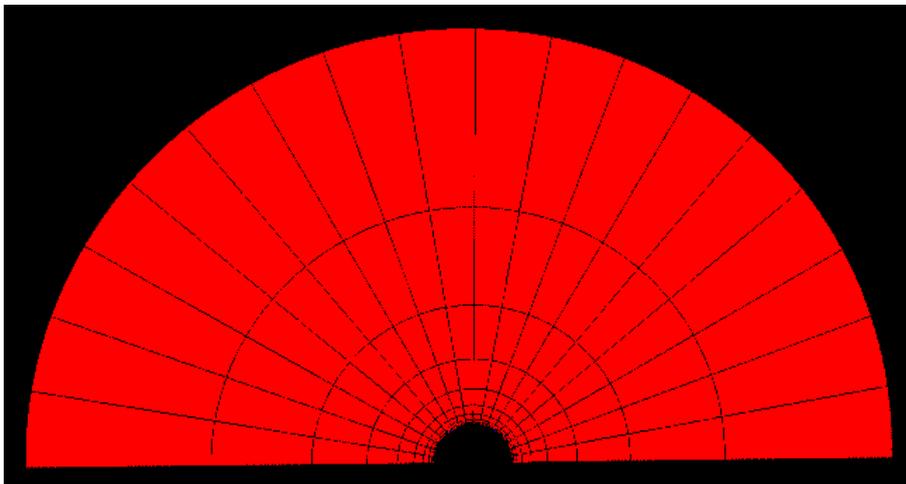


Figure 5: The fracture grid showing the different zones in spatial coordinates

The calculation of Q_{eq} only requires a steady state calculation. QPAC can calculate this directly. Given that results are required for a range of velocities, it is convenient to impose a time-dependent velocity and so calculate a sequence of results in a single QPAC run.

In order to calculate the individual Q_{eq}^{buff} and Q_{eq}^{frac} values, it is necessary to define a single interface concentration. The joiner subsystem calculates an interface concentration for each interface between a fracture and a buffer cell. We take a simple arithmetic average of these, on the basis that they all have similar areas and the concentration variation is small.

2.3 Reference Calculations

The first set of calculations is aimed at a direct comparison with the Q_{eq} formulae used by SKB. Geometric properties are as in Figure 1. An aperture of 0.1 mm was used initially, with pore velocities ranging from 0.1 to 10 000 m/y. Other transport properties are given in Table 3

Table 3 Reference transport properties

Property	Value Used	Note
Buffer porosity	0.43	from [9] Table A-11
Buffer effective diffusion	$1.2 \cdot 10^{-10} \text{ m}^2/\text{s}$	from [9] Table A-11
Fracture porosity	1	open fracture
Fracture effective diffusion	$0.0316 \text{ m}^2/\text{y}$	from [9] p158. D_w is the same as open fracture
Buffer and fracture retardation	1.0	non sorbed tracer

The Q_{eq}^{buff} value should be essentially independent of the flow velocity, as it depends only on the buffer properties. In practice there might be a small variation, because of the variation of concentrations at the interface for example.

In fact, the numerically calculated value does not change across the whole velocity range. A value for Q_{eq}^{buff} of $5.56 \cdot 10^{-3} \text{ m}^3/\text{y}$ is obtained. This compares to a formula (equation 2.1.5) value of $6.36 \cdot 10^{-3} \text{ m}^3/\text{y}$. Thus, there is a 15% difference on the conservative side, which is a little larger than claimed in [10].

We note that experiments while assessing grid convergence showed that the buffer beyond the immediate vicinity of the fracture has a rather small influence. Inappropriate gridding near the fracture can lead to significant errors (factors of two or more).

Figure 6 shows the concentration profile around the interface (at several velocities). There are two points to note here. First, the concentration is close to unity (recall that the canister concentration is unity) except for the very highest velocities, indicating that the buffer is not dominating the release. Second, the variation around the circumference is significant and so the fact that the calculated Q_{eq}^{buff} is constant suggests that the local flux scales with the local concentration all around the circumference.

Figure 7 shows the overall pattern of concentrations in the fracture for a central case. There is a clear upstream-downstream difference as expected. The lower velocity cases show less differences between upstream and

downstream but the concentrations are higher throughout. The highest flow case has a small plume downstream only.

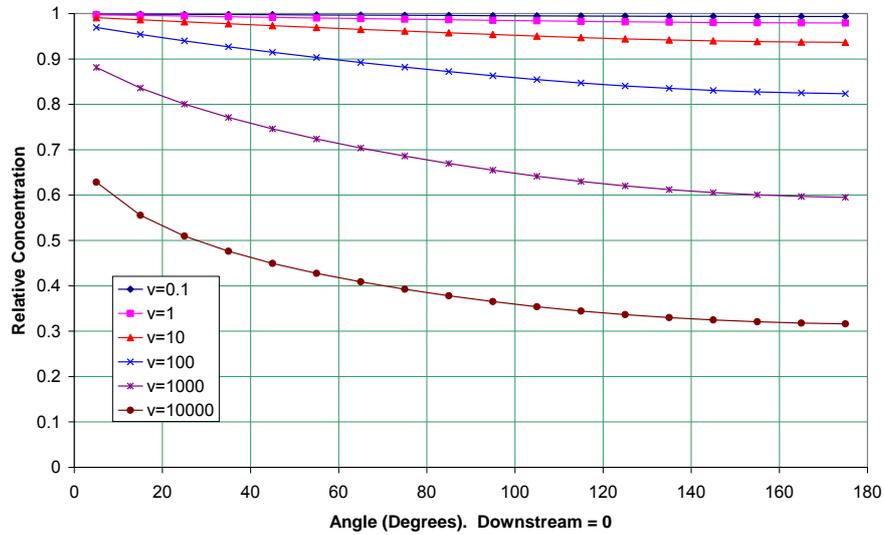


Figure 6: Concentrations around the fracture-buffer interface for a range of velocities

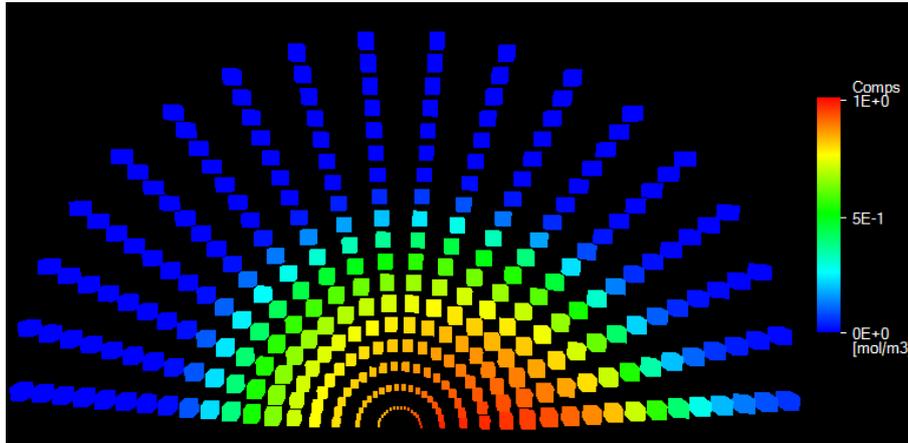


Figure 7: Concentrations in the fracture for the 100 m/y flow case, in index coordinates

The Q_{eq}^{frac} values should depend on the flow velocity. Figure 8 shows the numerically calculated and formula values (equation 2.1.7). It is clear that there only a small discrepancy (16% at low velocities and 24% at high velocities). The values are much lower than for the buffer except for the largest velocities. Figure 9 shows the average concentration on the interface. The buffer and fracture components become equal when the average concentration reaches 0.5.

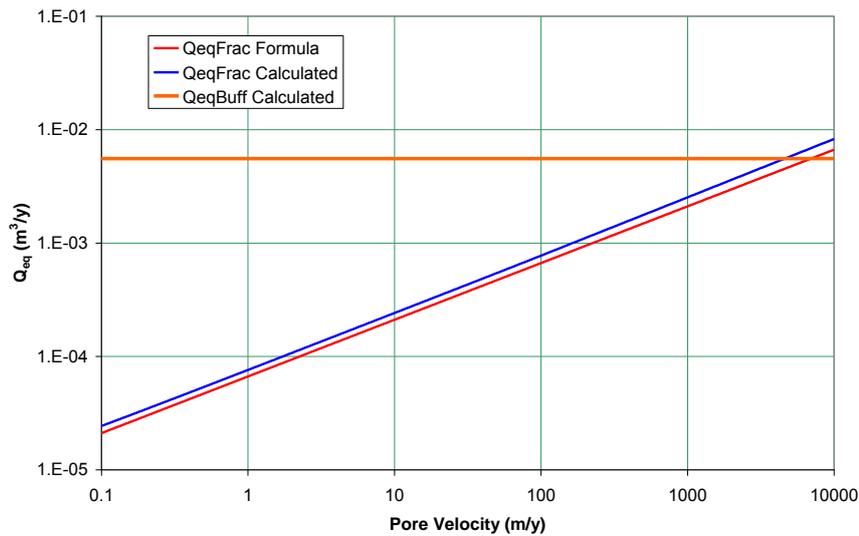


Figure 8: Comparison of calculated and formula values of Q_{eq} across a range of velocities

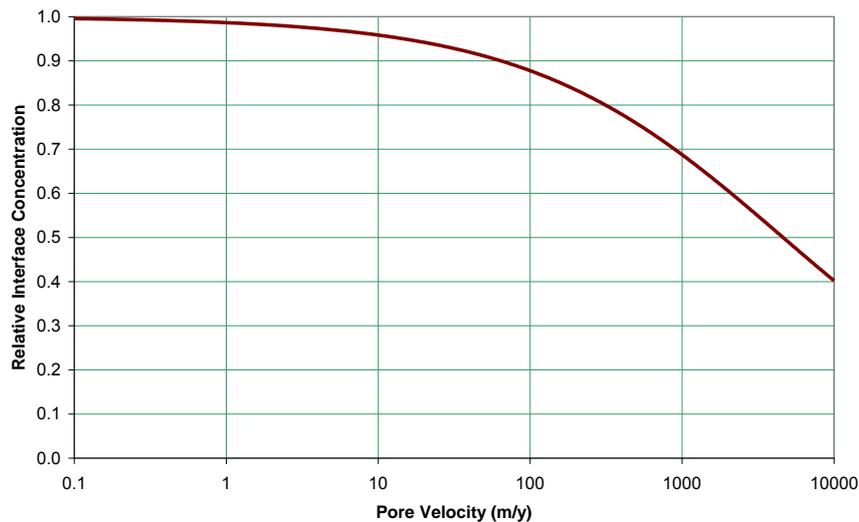


Figure 9: Average concentration on the fracture-buffer interface across a range of velocities

These calculations have been repeated for two other apertures: 0.3 mm and 1 mm. Figure 10 and Figure 11 show the results. The QPAC calculated Q_{eq}^{buff} values are 6.22e-3 m³/y and 7.13e-3 m³/y, compared to formula values of 7.12e-3 m³/y and 8.20e-3 m³/y. These are both about 15% different on the conservative side, very similar to the 0.1 mm case.

The Q_{eq}^{frac} values show that the buffer and fracture components become equal at much lower velocities for the larger apertures. It is clear that there is only a small discrepancy (16% at low velocities and 24% at high velocities) of the same size as in the case with a fracture aperture of 0.1 mm.

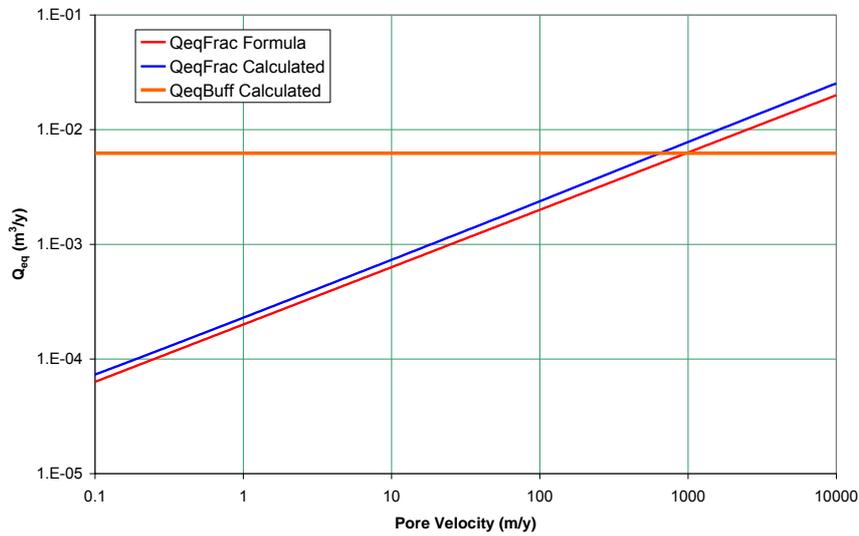


Figure 10: Comparison of calculated and formula values of Q_{eq} across a range of velocities for an aperture of 0.3 mm

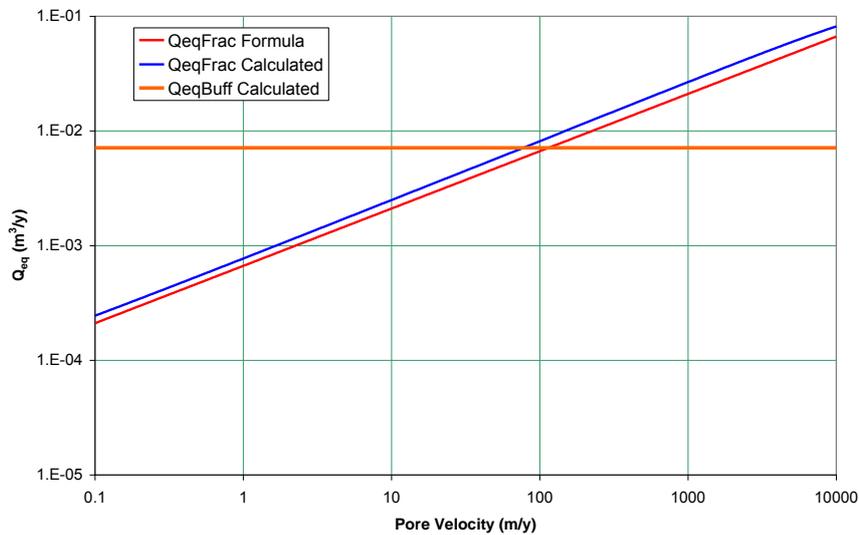


Figure 11: Comparison of calculated and formula values of Q_{eq} across a range of velocities for an aperture of 1 mm

These results indicate that the Q_{eq} formulae give a good approximation to the release in the simple planar geometry.

2.4 Fracture Heterogeneity

In order to treat heterogeneity in the fracture, we impose a varying porosity value. This is numerically simpler to handle than varying the aperture, where discontinuous cross-sectional areas would be hard to handle.

We start by taking a 0.2 mm aperture with a porosity of 0.5 throughout. This is expected to give similar results to the earlier calculations for a 0.1 mm open fracture and this is confirmed numerically – the Q_{eq}^{frac} values agree to 4 significant figures.

We now vary the porosity in a zone near to the buffer. The porosity further away is not varied because the grid cells are larger there and an average is more appropriate and because the impact is expected to come from the zone near to the fracture (as can be seen by the concentrations in Figure 7 for example). Conductivities are scaled with the square of the porosity, because it is intended that the porosity change should represent an aperture change and a Pouseille flow law can be assumed. The effective diffusion scales with porosity.

A randomly sampled porosity was used, with half the cells set to 0.01 and the other half to 0.99. Figure 12 shows the sampled porosity values. The concentrations for the 100 m/y velocity case are shown in Figure 13. The local impact of the porosity variations can clearly be seen.

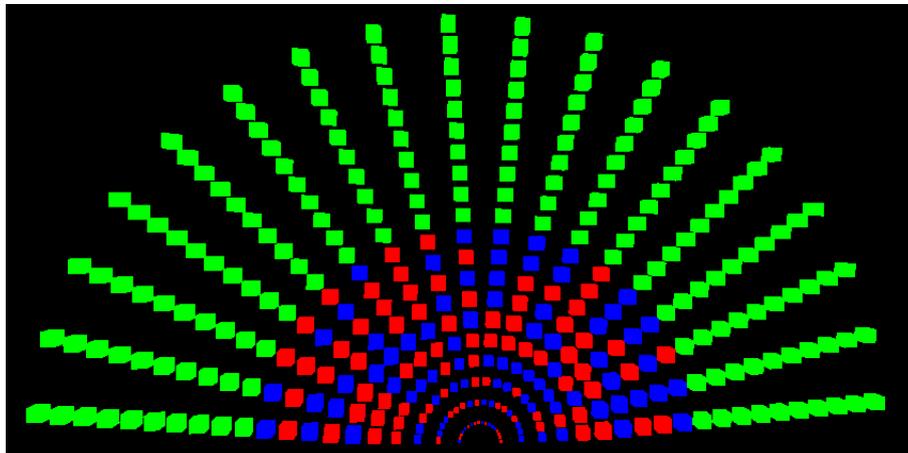


Figure 12: Porosity values for the heterogeneity test. Red = 0.01, Blue = 0.99, Green = 0.5

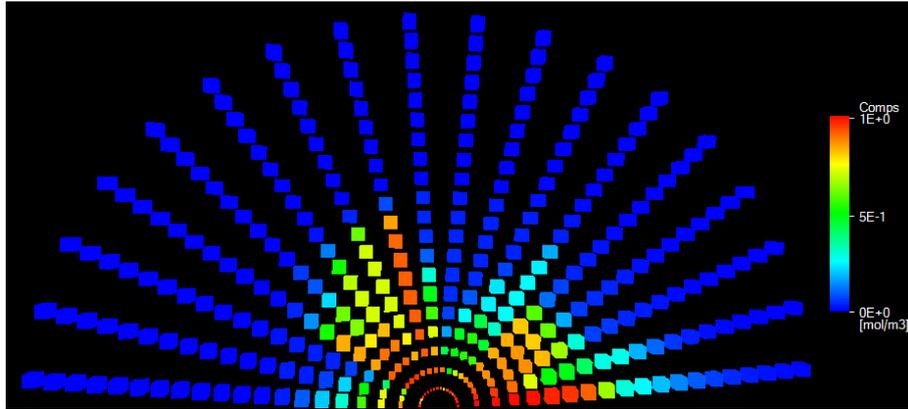


Figure 13: Concentrations for the heterogeneity test for a velocity of 100 m/y

The effect obviously depends on the porosity samples. A second concentration calculation for a different sample is shown in Figure 14.

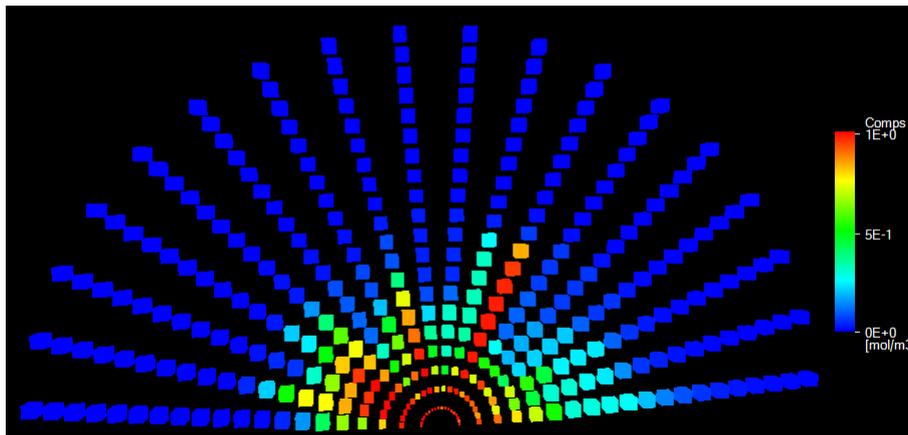


Figure 14: Concentrations for the heterogeneity test for a velocity of 100 m/y with a second sampled porosity field

The Q_{eq}^{frac} calculations differ for each realisation. Figure 15 shows results for 8 such realisations, compared to the uniform case. It is clear that the heterogeneity reduces Q_{eq}^{frac} and that the reduction depends both on the particular realisation and on the overall velocity. The velocity dependence depends on the details of the distribution of high and low porosities. It appears that at low velocities there is less sensitivity to the realisation than for the higher velocities. At low velocities there is up to a factor of two reduction from the uniform case; for intermediate velocities this increases to almost a factor 10 before reducing again. There is some indication that for very high velocities the factor again increases, but these are beyond the range of interest.

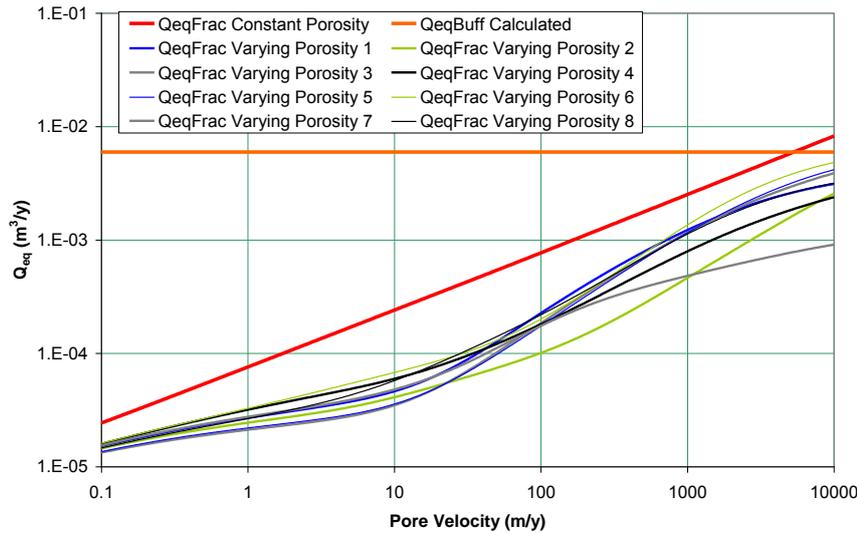


Figure 15: Comparison of calculated values of Q_{eq} across a range of velocities for constant and varying porosity fields

Overall, the effect of randomly varying porosity with no correlation structure (which can be interpreted as varying fracture aperture) is to reduce releases, so taking a constant fracture aperture is conservative.

The examples shown here did not have any correlation in the porosity choices and so did not include channels [11]. It is clear that the length scale over which the concentration falls away from the interface value is small (a few cm), so any channel of that size would act in the same way as a full fracture. Therefore, if there is a channelling factor of 10 (all the flow is in 10% of the fracture), it is clear that the effect can be the same as a 10 times larger velocity (giving a factor of $\sqrt{10}$ for Q_{eq}^{frac}).

2.5 Spalling

In order to look at the effect of a localised spalling zone at the edge of the buffer, we introduce a zone of high conductivity in a zone above and below the fracture. The size of the zone is chosen for gridding convenience to be the same as the unit concentration zone on the canister. It extends a distance of 2.5 cm into the rock. Thus there is an enhanced area of contact between the flowing water and the buffer.

Figure 16 shows results for the spalling case. The Q_{eq}^{buff} value in the spalling case is higher, as would be expected because of the much larger contact area, and varies with velocity to a small degree. More importantly, Q_{eq}^{frac} is much larger and appears to scale directly with velocity, rather than the square-root. This arises because the constriction at the interface has been removed and there is now advection from the spalling zone into the main part of the

fracture. It is the advection at this interface that controls the release – hence the linear dependence.

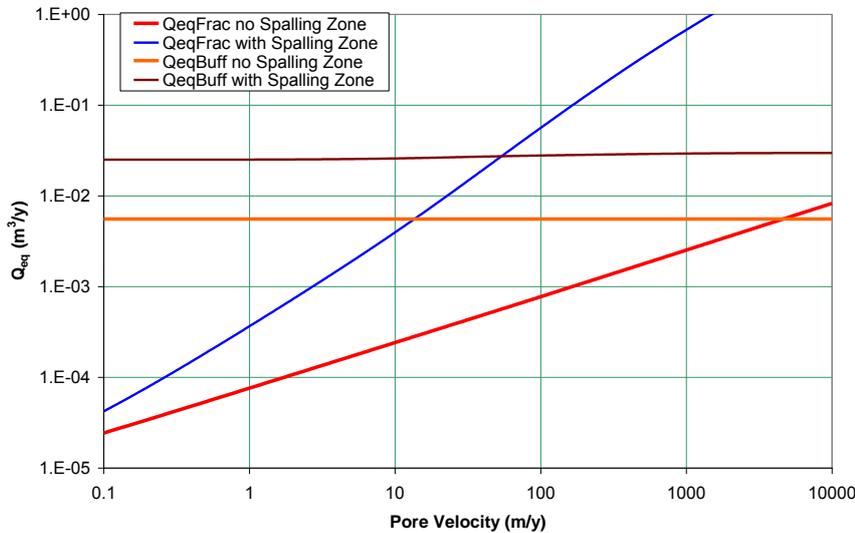


Figure 16: Comparison of calculated values of Q_{eq} across a range of velocities for the spalling case

Thus, the impact of spalling, even when it is localised, can be significant. This is an important area for review in SR-Site since the approach used in SR-Can was provisional.

There is a need to model the spalling zone in a more realistic way, but the calculations here show that simply increasing the contact area (without any significant extra flow) is significant.

2.6 Conclusions

The calculations presented here have shown that:

- ▲ the basic approach to calculating Q_{eq} values is sound and can be reproduced in a numerical model;
- ▲ the fracture resistance dominates over the diffusive resistance in the buffer except for the highest velocity cases;
- ▲ heterogeneity in the fracture, in terms of uncorrelated random variations in the fracture aperture, tends to reduce releases, so a constant average aperture approach is conservative;
- ▲ length scales for the release are such that narrow channels would lead to the same release as larger fractures with the same pore velocity, so a channel enhancement factor of $\sqrt{10}$ should be considered;

- ▲ a spalling zone that increases the area of contact between flowing water and the buffer has the potential to increase the release significantly and changes the functional dependence of Q_{eq}^{frac} on the flowing velocity.

3 Time-dependence in Radionuclide Transport Calculations

3.1 Background

To date, SKB have used a one-dimensional transport modelling approach using the FARF31 code [12], [13]. As stated in the main SR-Can report [4], one limitation of the migration path concept is that only steady-state velocity fields can be addressed whereas clearly the flow field will evolve in time (due to factors such as shoreline displacement and glacial advance and retreat). These and other limitations led to the development of the MARFA code, reviewed in Section 4.

Quintessa's AMBER software [14] has previously been used to reproduce SKB's transport calculations (e.g., [2]). AMBER allows the use of time-varying properties although this has not previously been exploited in calculations undertaken by Quintessa for SKI or for SSM. Here this capability has been used to investigate the effects of glacial episodes on the radionuclide transport. The main parameters that could be affected are the sorption coefficients and flow rates. The capabilities of AMBER allow smooth transitions between parameter values for temperate, glacial advance/retreat and glacial completeness conditions, which is not possible with MARFA or the Quintessa semi-analytic approach described in Appendix B. We note that the QPAC code [8] could also have been used for these calculations, but that it was convenient to continue with the established AMBER model.

Here we focus on releases from the geosphere. The question of how to use these to calculate doses to humans is also important, but it is understood that the landscape dose factor (LDF) method used by SKB in SR-Can will be replaced in SR-Site. Therefore no dose calculations are presented in the current report, and results are shown in terms of fluxes.

3.2 Conceptual Model

3.2.1 Description

The calculations described here are based on the pinhole failure scenario considered by SKB in SR-Can [4]. Although new welding techniques may mean that the risk of pinhole failures in canister welds is negligible, this case is useful for examining the safety role that the buffer and host rock play in

the multi-barrier concept. It is also representative of other failure modes, e.g. canister corrosion.

A brief description of the pinhole failure scenario calculations is given here; full details can be found in the SR-Can reports [4], [9] and in [2]. As previously indicated, the SKB transport models employed for SR-Can did not include time-varying properties and therefore parameter values (such as those for sorption coefficients and flow rates) were employed that are suitable for temperate interglacial periods. In this study a base case calculation that adopts the same approach as SKB has been undertaken together with variant time-varying calculations that consider glacial cycles and allow parameters to vary accordingly.

In order to focus on the effects of time-varying properties, deterministic calculations have been employed here using best estimate parameter values. The calculations presented here consider only the Q1 pathway (a fracture intersecting the deposition hole) since, as discussed in [3], this helps to focus attention on the issues being addressed without the complicating factor of different transport pathways. A schematic illustration of the near field is shown in Figure.17.

The canister is assumed to have a pinhole defect from the start, but this does not become an open pathway due to water intrusion until 1000 years after repository closure. The defect eventually leads to failure of the canister after 10 000 years. The flux of radionuclides into the buffer is limited by their solubility in water.

Once radionuclides have entered the bentonite buffer, they are transported by diffusion. The buffer is assumed to remain intact throughout the simulation. Sorption onto the clay mineral surfaces is modelled using sorption coefficients (K_d).

As described above, the only pathway into the host rock considered in this study is Q1, a fracture intersecting the deposition hole. The flux from the bentonite buffer into this fracture is proportional to the equivalent flow rate (see Section 2).

Once in the fracture the radionuclides are transported to the surface by advection, the rate of which is calculated separately in flow models. Data obtained from SKB during earlier studies [3] in the form of average travel times has been used in the present model. Along the fracture pathway radionuclides diffuse into the surrounding rock matrix, where sorption may again occur; as before, this is modelled using sorption coefficients.

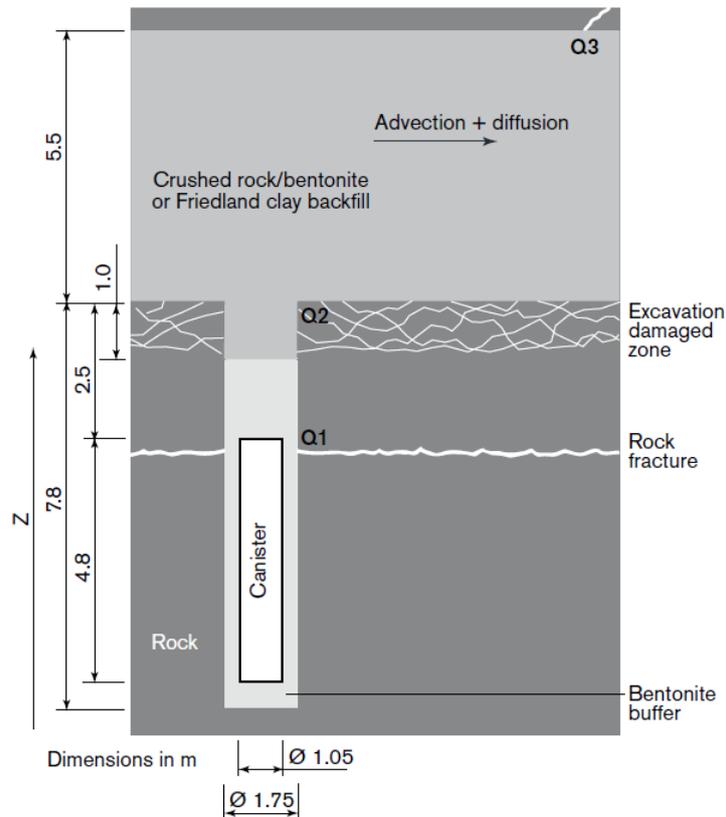


Figure.17: The near field, depicting the deposition hole, access tunnel and host rock (reproduced from [4]). Only path Q1 (to a fracture intersecting the deposition hole) is considered in the calculations reported here.

3.2.2 Glacial Phases

A glacial episode is likely to have a large impact on parameter values, as it will change flow regimes, groundwater compositions and alter the landscape above the repository dramatically. Glacial episodes are therefore used to investigate the effects of time-varying properties on the transport calculations.

Following Jaquet and Siegel [15] four phases of a glacial cycle are considered in the first time-varying variant case (all times are relative to the closure of the repository):

- I. 0 - 9 000 years: Temperate period (no ice sheet)
- II. 9 000 - 14 800 years: Glacial advancement (ice sheet progressively covers the area surrounding the repository)

- III. 14 800 - 55 800 years: Glacial completeness (ice sheet covers the area surrounding the repository).
- IV. 55 800 - 58 500 years: Glacial retreat (ice sheet progressively withdraws)

These timescales are purely illustrative; there is great uncertainty as to when the next glacial episode will occur, whether global warming will delay its onset etc.

Flow calculations conducted by Jaquet and Siegel and reported in the main SR-Can report [4] for a glacial episode at Simpevarp show large variations in salt concentrations and flow velocities (Figure 18). An initial increase in salinity during the advancement phase is a result of salt diffusing upwards from depth. High groundwater flows are seen when the ice sheet margin is located directly above the repository, during both advance and retreat. This results in salt being flushed out of the fractures and rock matrix, and concentrations remain low throughout the remainder of the simulation (which may be unrealistic).

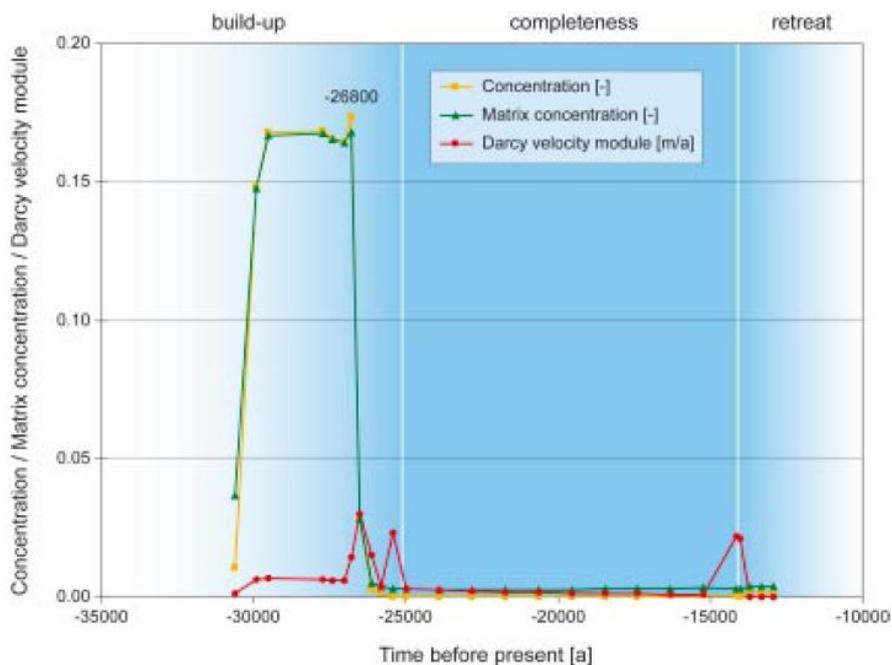


Figure 18: Calculated salt concentrations in the flowing water and rock matrix and the Darcy velocity during a glacial episode at Simpevarp (reproduced from [4])

3.2.3 Time-Dependent Parameterisation

This study focuses on the effects of the glacial cycle on the far-field parameters; there is much uncertainty surrounding the effects that glacial meltwaters will have at repository depth, but it is likely that they will have some effect on parts of the far-field at least. The issue of buffer erosion is considered in [1].

Far-field transport is taken to take place along a fracture that intersects the deposition hole. The flow velocity in the fracture is determined from the particle travel times calculated from the SKB discrete fracture network (DFN) models. The length of the path has been assumed to remain constant during a glacial episode (although glacial erosion and the thickness of ice at the surface may alter this in reality), but the travel time will change as flow rates increase and decrease during the glacial advancement and completeness phases.

An important safety function of the far-field is the retention of radionuclides by diffusion into the rock matrix and sorption onto mineral surfaces. The rate of diffusion is unlikely to be greatly affected by glaciations, but sorption could be as groundwaters become more dilute. Sorption is characterised by a sorption coefficient or K_d . In the initial calculations changes to sorption were only considered near the surface, but a variant calculation was undertaken to investigate the effect of sorption coefficients being altered throughout the far field.

Radionuclides enter the fracture at the deposition hole end of the transport pathway from the bentonite buffer. This transfer rate is characterised by the equivalent flow rate, Q_{eq} , (see Section 2) which depends on the flow velocity in the fracture at depth. As discussed above, a glacial episode will affect flow regimes so it is possible that this parameter may also be affected.

In summary, therefore, the far-field parameters that are assumed to evolve with time are: sorption coefficients (K_d); equivalent flow rate (Q_{eq}); and travel time.

3.2.4 Parameter Values

Sorption Coefficients

One of the prime functions of the host rock in the Swedish concept is that it provides retention of radionuclides that have escaped the near field by its sorption capability. The SR-Can data report [4] provides tables of sorption coefficients (K_d values) for two types of waters: saline and non-saline. The values for saline conditions are generally lower (i.e. fewer radionuclides are retarded due to sorption) than those for non-saline conditions. During glacial completeness and retreat, it is likely that dilute glacial meltwaters will pass through the host rock. Such dilute waters are likely to result in different K_d values but no consideration is given in either the SR-Can data report or the original source [16] to such conditions.

For this study data from the Finnish waste organisation, Posiva, has been considered [17]. Large uncertainties exist as to the composition of glacial meltwaters, particularly the composition that could penetrate to repository depth. Thus, for the majority of elements, the K_d values used by Posiva are taken to be the same as for brackish water. However for some elements, such as uranium, a much lower K_d is used (2 orders of magnitude lower in the case of U).

As the Posiva values were derived for groundwaters and rock types specific to the Finnish, not the Swedish, repository, the K_d values have not been used directly in this study. Instead, the SKB non-saline values were multiplied by the same factor used by Posiva to produce glacial values. Thus for the time-varying calculations the SKB non-saline values were used in the temperate period, SKB saline values were used for the glacial advance phase, and glacial values (the SKB non-saline data multiplied by the Posiva factor) were used during the glacial completeness and retreat phases. Smooth transitions rather than step changes were employed, with the transition taking 1 000 years. For the base case the non-saline data were used throughout. The data used in the calculations is given in Table 4.

Table 4: Sorption coefficients for non-saline, saline and glacial meltwater conditions. Data for non-saline and saline conditions taken from best estimate values in [4]. Data for glacial meltwater conditions reduced from non-saline values for Tc, Sm, U, Pu and Np in accordance with Posiva data [17]. Other elements adopt the non-saline values.

Element	K_d ($\text{m}^3 \text{kg}^{-1}$)		
	Non-saline	Saline	Glacial Meltwater
C	1e-3	1e-3	1e-3
Cl-	0	0	0
Ni	1.2e-1	1.0e-2	1.2e-1
Se	1e-3	1e-3	1e-3
Sr	1.3e-2	3.1e-4	1.3e-2
Zr	1	1	1
Nb	1	1	1
Tc	1	1	0
Pd	0.1	1e-2	0.1
Ag ⁺	0.5	5e-2	0.5
Sn	1e-3	1e-3	1e-3
I	0	0	0
Cs	1.8e-1	4.2e-2	1.8e-1
Sm	2	2	0
Ho ⁺	2	2	2
Pb ²⁺	1.3	2.1	1.3
Po ²⁺	1.3	2.1	1.3
Ra	1.3	2.1	1.3
Ac ⁺	3	3	3

Element	K_d ($\text{m}^3 \text{kg}^{-1}$)		
	Non-saline	Saline	Glacial Meltwater
Th	1.0	1.0	1.0
Pa	1	1	1
Np	9.6e-1	9.6e-1	9.6e-3
U	6.3	6.3	6.3e-2
Pu	5	5	2
Am	13	13	13
Cm	3	3	3

† Not reported by Posiva, non-saline value used for glacial.

‡ Not given by SKB, same value as Ra used.

It should also be noted that correlations between K_d values for different radionuclides are considered in the SKB transport calculations. Here no correlations were considered, and the best estimate values were adopted.

Particle Travel Time

DFN calculations have been carried out by Hartley et al. [18] to provide flow input to the transport models described in the main SR-Can [4]. However, this data is only valid for temperate climates. The work of Jaquet and Siegel [15] describes the modelling of groundwater flow in a glacial domain (see Section 3.2.2), but has not been applied to date to the Forsmark area.

One of the outputs of the DFN modelling is the particle travel time, the time taken for a particle released at depth to reach the surface. This is used in the AMBER calculations to give the velocity in the fracture. The mean travel time in the data provided by SKB for temperate conditions at Forsmark is 127 years.

The mean travel time at the Simpevarp subarea considered by Jaquet and Siegel in temperate conditions is ~459 years. During glacial advance/retreat phases, this is reduced to approximately 5 years, as the head gradient increases greatly and meltwaters flush through the system. During glacial completion the travel time increases greatly to ~1503 years; flows are effectively cut off and hydraulic gradients are reduced significantly. The values are included in Table 5.

As these values were calculated for Simpevarp and not Forsmark, they cannot be applied directly in the AMBER calculations but they can be used to provide an indication of the factor by which travel times may change. Travel times for the temperate period were reduced by a factor of 100 for the glacial advance/retreat phases and multiplied by 3 for completion.

Table 5: Mean travel times (the time taken for a particle released at depth to reach surface) calculated by Jaquet and Siegel [15] for different points in the glacial cycle. These values are relevant for the Simpevarp subarea.

Time (years before present)	Glacial Phase	Mean Travel Time (years)
-26 800	Build-up	7.4
-26 500	Build-up	5.2
-17 900	Completeness	1 503
-13 900	Retreat	5.7
-13 800	Retreat	6.9
0	Temperate	459

The travel times adopted for the AMBER calculations are shown in Table 6. Again, a smooth transition between values over a period of 1 000 years was employed. The temperate value was used for the base case calculation.

Table 6: Mean travel times (the time taken for a particle released at depth to reach the surface) adopted in the AMBER calculations for the Forsmark repository.

Glacial Phase	Mean Travel Time (years)
Temperate	127
Build-up/Retreat	1.3
Completeness	381

The values used here are simply for sensitivity analysis and should not be taken to be representative values for the Forsmark site. It should also be noted that Jaquet and Siegel state that the travel times for glacial completeness are likely to be overestimated due to the boundary conditions employed.

Equivalent Flow Rate

The equivalent flow rate determines the transfer of radionuclides from the bentonite buffer to the fracture in the host rock. In the base case calculation, a value of $1.16 \times 10^{-4} \text{ m}^3 \text{ y}^{-1}$ was used (averaged from the SKB DFN results). Similar data is not available for non-temperate flows, but the expectation is that this value will increase (at least during the advancement and retreat phases) due to the increased flow rates experienced. Therefore as a simple sensitivity calculation, this value was multiplied by a factor of 10 during the non-temperate phases of the glacial cycle. The equivalent flow rate is proportional to the square root of the flow velocity, as discussed in Section 2, therefore this value is reasonable for the advance/retreat phases, but is probably unrealistically large for the glacial completeness phase.

3.3 Calculations

3.3.1 Single Early Glacial Episode

Near Surface Effects

The first case considered was a single glacial episode which occurs relatively soon after the repository is closed, after a period of 9 000 years. It was assumed that only the near surface is affected by the glacier, thus K_d values were only modified in the final 100 m near the surface of the 500 m pathway. The total flux of activity to the biosphere is shown in Figure.19; for comparison, the flux calculated for the base case where the climate is assumed to remain temperate for the duration of the simulation is also shown.

During the glacial advance there is a very slight initial increase in flux in the glacial case compared to the base case, followed by a slight dip. The pinhole defect in the canister increases in size during this period, causing the increase in flux seen in both cases. The peak flux is approximately a factor of 1.2 larger in the glacial case, as the increased flow rates ahead of the glacier flush the contaminants out of the geosphere.

During the glacial completeness period, there is very little difference in the fluxes observed in the two cases despite the longer particle travel times. If doses were considered instead of fluxes, it is likely that a large decrease in the dose expected to humans in the biosphere would be observed due to the thickness of the ice on the surface and the reduction of flow rates there.

The peak in flux during the advancement phase is reflected in the glacial retreat phase but here the increase in the total flux is much larger (the flux in the glacial case is approximately 6 times larger than that in the base temperate case). The peak is caused by the increased flow rates, decreased particle travel times and lower K_d values during retreat. Following the peak there is a slight dip in the flux, as the temperate phase begins.

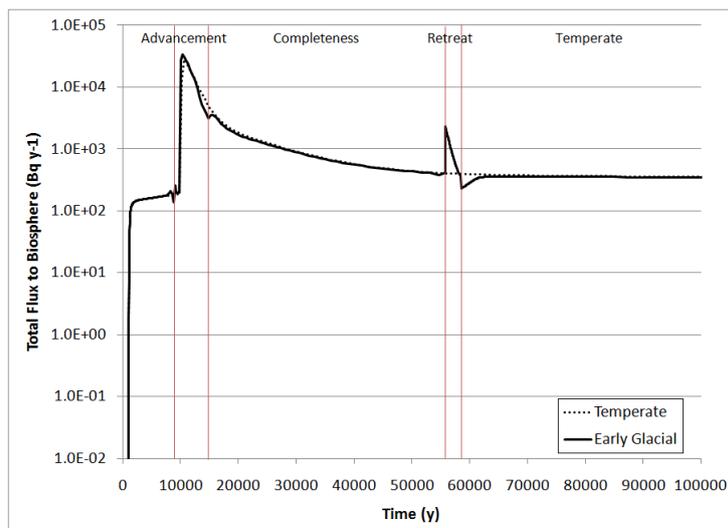


Figure.19: Total fluxes to the biosphere for the temperate (base) case and the single early glacial episode case

The lack of difference between the two cases is due to the fact that the main contributors to the total flux over this time period (100 000 years) are ^{129}I , ^{36}Cl (which are assumed not to sorb at all on the rock matrix or fracture surfaces) and ^{14}C (which has a constant K_d value throughout). The increased flow rates and decreased travel times during advancement and retreat have a small effect (the small peak seen in Figure.19), but these do not persist for long enough to have a big impact on the calculated fluxes.

The individual fluxes from each radionuclide are shown in Figure.20 for the temperate case and Figure.21 for the single early glacial episode. There is a small increase in the fluxes contributed by other radionuclides, for example ^{79}Se and ^{126}Sn , but the fluxes contributed by these radionuclides are several orders of magnitude lower than that of ^{129}I or ^{36}Cl , and thus the impact of glaciation is small. These increases are not attributed to changes in the sorption coefficients, which are also fixed for these radionuclides, but appear to be due to the increase in the equivalent flow rate during the non-temperate period. Fluxes of the other radionuclides are therefore limited by the rate of release from the source rather than the transfer rate from the buffer to the fracture.

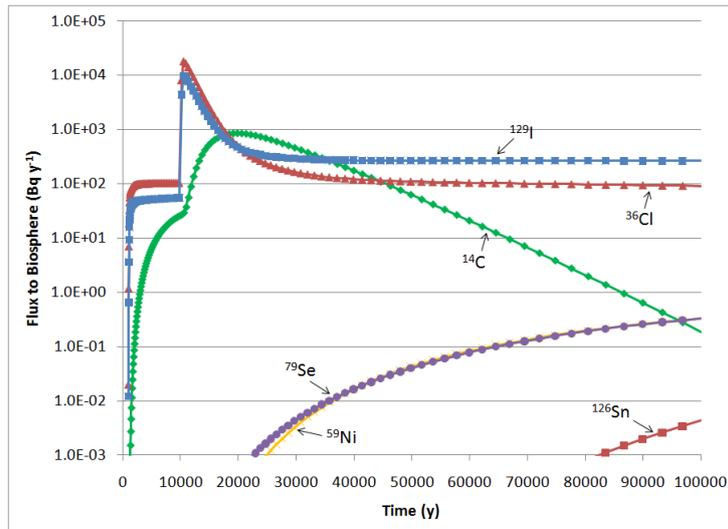


Figure.20: Fluxes of each radionuclide to the biosphere for the temperate case

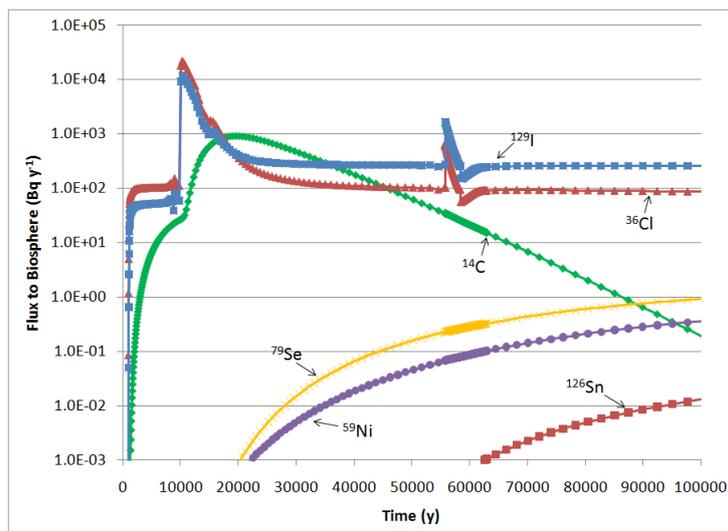


Figure.21: Fluxes of each radionuclide to the biosphere for the single early glacial episode case

Full Depth

The extent to which a glacial episode will affect system evolution at repository depths is uncertain; therefore a variant case was considered where K_d values were reduced in all fracture and matrix compartments, not just those near the surface. This was found to have no impact on the calculated total fluxes.

3.3.2 Multiple Late Glacial Episodes

Glacial advances and retreats are generally not single events; for example, in the last two million years it is estimated that there have been at least 60 such episodes. This variant case therefore considered 10 glacial phases, occurring every 100 000 years. The first episode now starts at 109 000 years. In this variant, the effects of glaciation were assumed to affect the whole of the geosphere not just the near surface.

As shown by the individual radionuclide fluxes in Figure.22 (compare to the temperate case, Figure.23), there is little impact on the main flux contributors; ^{129}I , ^{14}C and ^{36}Cl early on and ^{135}Cs and ^{107}Pd later in the simulation. However, there are significant differences for some of the radionuclides that provide smaller contributions to the total flux; for example ^{126}Sn , ^{59}Ni and ^{99}Tc , all of which have a higher peak flux in the glacial case.

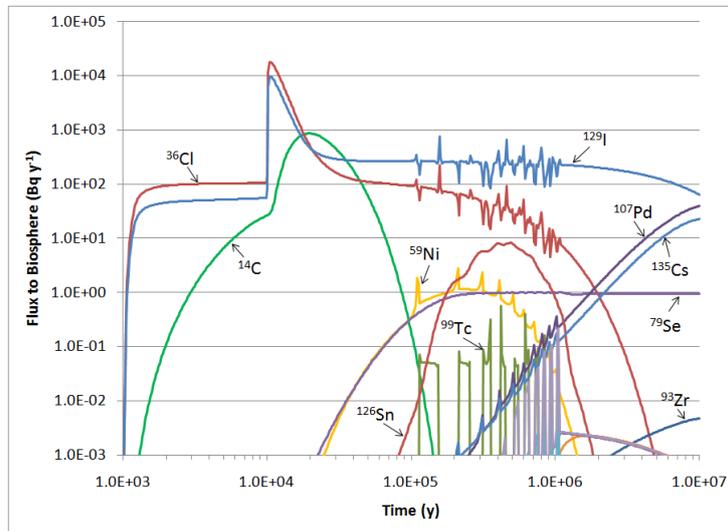


Figure.22: Fluxes of each individual radionuclide to the biosphere for the multiple late glacial episodes case

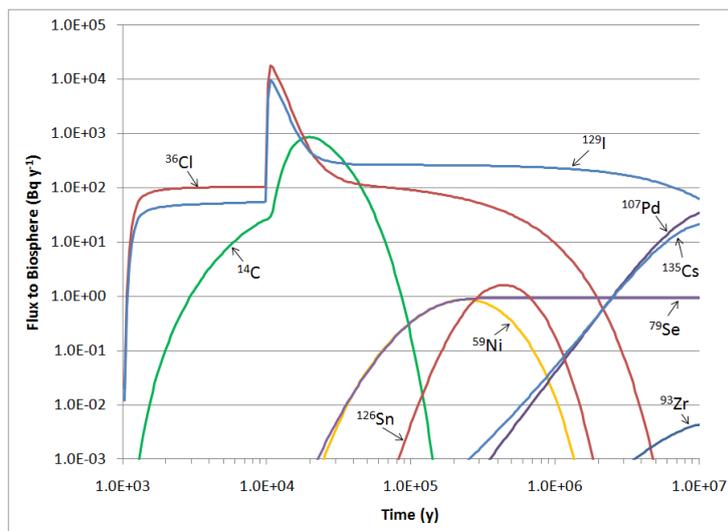


Figure.23: Fluxes of each individual radionuclide to the biosphere for the temperate case

Considering just the radium decay chain (Figure.24), which is important at long timescales, it is clear that the glaciation episodes do increase the flux seen at the surface whilst the glaciations episodes are occurring, by up to an order of magnitude, but that the effects are limited and the fluxes return to the levels seen in the base temperate case not long after the last glacier has retreated.

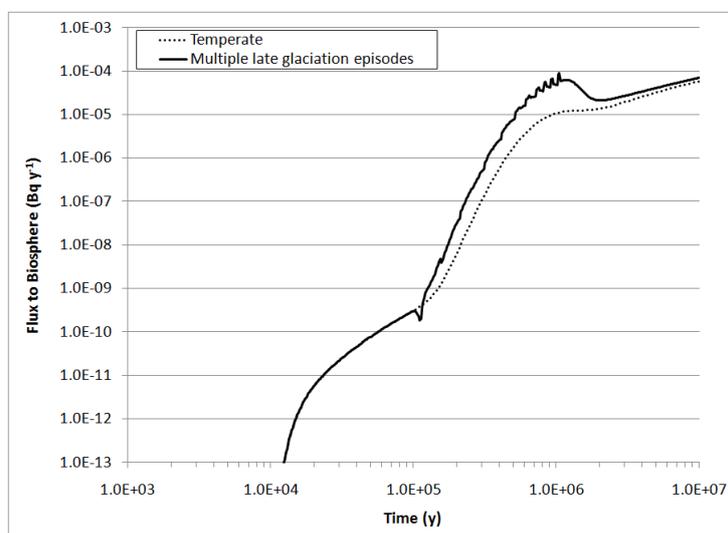


Figure.24: Total flux to the biosphere arising from the radium decay chain ($^{238}\text{U} \rightarrow ^{234}\text{U} \rightarrow ^{230}\text{Th} \rightarrow ^{226}\text{Ra} \rightarrow ^{210}\text{Pb} \rightarrow ^{210}\text{Po}$)

3.4 Conclusions

This study investigated the impacts of time-varying flow rates and sorption coefficients on transport calculations of the type used by SKB in the presence of glacial cycling. A single glacial episode occurring 9 000 years after repository closure resulted in short-lived, increases in fluxes of radionuclides to the biosphere during glacial advancement and (in particular) retreat when the flow rates increase due to penetration of glacial meltwaters. However, there is no impact on fluxes to the biosphere during the temperate periods. This is because the dominant radionuclides in the period of interest are ^{129}I , ^{36}Cl and ^{14}C , which are either assumed not to sorb or to have sorption coefficients that are not altered by the presence of the glacier, and so are largely unaffected by the glacial episode (except for a slight increase in flux during advancement and retreat).

For multiple glacial episodes starting later (100 000 years after closure), there was again little impact on the fluxes to the biosphere for the same reasons. However, a small increase (less than an order of magnitude) in the overall flux contributed from the radium decay chain, which is important for long timescales, was calculated during the phase when the glacial episodes are occurring. Once temperate conditions return and are maintained the calculated flux returned to the value seen when temperate conditions were assumed throughout.

These calculations highlight the fact that the timing of the onset of glaciation is important, as well as the number of episodes. If glacial episodes occur when members of the radium chain are the dominant radionuclides, it could lead to slightly elevated fluxes and possibly also therefore doses during temperate periods. A more detailed consideration of suitable parameter values, particularly K_d and LDF vales, would be required to reduce the uncertainties associated with these calculations. Consideration would also need to be given to the penetration depth of glacial waters, and whether the bentonite buffer could be eroded.

As mentioned in the introduction to this section, in order to calculate doses, the releases calculated here need to be used as input to a suitable biosphere model. This will change the relative importance of particular nuclides but it is not expected that the overall conclusions presented here would change.

4 Review and Testing of MARFA

The MARFA Code [19] has been developed for SKB in order to address some limitations in SKB's modelling capabilities for geosphere transport.

According to SKB (section 10.3.2 in [4]):

A limitation with the migration path concept is that only steady-state velocity fields can be addressed (adopting the snapshots in time approach for transport modelling), whereas clearly the flow field will evolve in time due to shoreline displacement. A second limitation with the current utilisation of the F-factor integrated over the migration path as an input parameter is that the solution is formally correct for single-member decay chains only. For longer decay chains, use of the integrated parameter F is strictly not correct if the channel width to flow ratio varies in space. An entirely new transport code under development, based on a Particles On Random Streamline Segments (PORSS) approach, will be able to handle both transient flow and variable conditions (including variable matrix parameters) for transport of single nuclides and decay chains. A first application of the new code, in parallel with FARF31, is planned within SR-Site.

MARFA was developed in response to this need. The MARFA user's guide [19], henceforth referred to as the User Guide, states that:

The physical processes represented in MARFA include advection, longitudinal dispersion, Fickian diffusion into an infinite or finite rock matrix, equilibrium sorption, decay, and in-growth. Multiple non-branching decay chains of arbitrary length are supported.

It also notes that:

Version 3.2 included new capabilities to accommodate transient flow velocities and sorption parameters, which are assumed to be piecewise constant in time.

No earlier versions of MARFA have been documented in SKB reports. The basis of the current review is the main part of the User Guide which relates to version 3.2.2. The User Guide also contains an appendix describing version 3.3 which extends the capabilities of MARFA to allow pathways through a fracture network to change with time. We comment on this only briefly in this review.

In addition to the User Guide and the associated scientific papers to which it refers [20, 21, 22], this review used a copy of the MARFA code. This was supplied by SKB with the test case inputs. The Fortran-90 source code was also supplied – this is not the subject of the current review, but was occasionally useful in clarifying various algorithmic details.

MARFA is designed to be linked to a discrete fracture network (DFN) code, such as CONNECTFLOW, but can also be run separately. The current review focuses on the basic capabilities of MARFA rather than the DFN aspects. In particular, the use of subgrid downscaling is not reviewed in detail.

Because the particle following methodology used in MARFA is novel, this review has looked in some detail at all aspects of the algorithms used as well as reviewing the code and its documentation. A key objective of the review is to identify the strengths and weaknesses of MARFA so that appropriate judgements can be made when reviewing its application in SR-Site. The current review is being undertaken before the SR-Site documentation is published, and hence the degree to which MARFA will be used is not known.

The review is structured as follows.

- ▲ Section 4.1 gives a general overview of the way MARFA works;
- ▲ Section 4.2 reviews the documentation;
- ▲ Section 4.3 looks at the way the code is used;
- ▲ Section 4.4 describes and reviews the algorithms behind the approach;
- ▲ Section 4.5 reviews the documented test cases;
- ▲ Section 4.6 describes new tests developed for this review, including the use of a new semi-analytic approach to handling piecewise-constant parameters;
- ▲ Section 4.7 discusses the MARFA 3.3 documentation; and
- ▲ Section 4.8 draws some conclusions.

More details of the new Quintessa semi-analytic approach are given in Appendix B.

4.1 General Overview

MARFA employs a “pathway stochastic simulation” approach. Although the User Guide does not discuss why this approach has been chosen, some of the background scientific papers emphasise the robustness and efficiency of the approach within the context of a network of 1D pathways arising from a discrete fracture network flow calculation.

The basic idea is that the distribution of arrival times for particles travelling through a single segment (a 1D pathway with constant properties) can be calculated through separate consideration of advective travel time and time interacting with the rock matrix. The distribution of arrival times is directly related to the cumulative breakthrough curve for a delta-function input. By

sampling arrival times for individual particles, and cumulating these through a complex network of segments, the breakthrough curve for the network can be reconstructed. The two key difficult aspects to handle are radioactive decay and switches in properties with time. In both cases, a particle may not arrive at the end of the segment before a change occurs. MARFA handles both situations by (stochastically) determining where the particle is at the time of the change and updating its arrival time to account for the new circumstances.

Once a large number of particles (at least millions) have been tracked, the overall cumulative breakthrough curve can be directly created and the rate of arrival can be derived from this, with appropriate smoothing.

Various approximations are introduced in each aspect of the algorithm, but the intention is that these remain insignificant in the overall calculation.

4.2 Documentation

This review is based on the single SKB report that has been produced – the User Guide [19]. This document provides some technical details but refers to published papers for some key results. It also contains a small set of test cases.

In this section, we discuss the overall documentation. More specific comments are made in the later sections.

The overall impression of the documentation is that it is rather inconsistent in the level of detail provided. In some cases, a series of detailed equations are provided with full explanations (e.g. when the handling of dispersion is discussed in section 3.1) while other aspects are discussed only briefly and without discussion (e.g. the retention functions used and how they are calculated).

The verification tests appear quite limited – no more than two nuclides are used for example. There is no discussion of many of the user controls – particular the number of particles, but also choices of importance factor (see Section 4.3.2), channelling, and type of source term sampling.

The description of the input files again offers little advice as to what values would be expected for some of the controlling parameters, or when it might be useful to use some of the options.

The preface of the document states that MARFA was developed under the CNWRA software QA procedures. It is not clear whether this implies that there is a suite of documentation covering user requirements, design and testing – but nothing of that sort is referenced in this User Guide.

4.3 The Code

For the purposes of the review, Quintessa was supplied with a copy of MARFA version 3.2.2 along with associated test input files and other files used by the code. The code was used for running test cases. Quintessa was also supplied with the Fortran-90 source code. This has not been reviewed, but was referred to during the review to confirm details of some of the algorithms. The retention functions are supplied with MARFA in small data files. These have been reviewed as part of the review of the algorithms in order to understand the impact of the various approximations that are made.

In this section, we set out our understanding of what MARFA does and comment on how this is presented and any potential approximations that are introduced. The aim is to understand the approach sufficiently well to discuss its strengths and weaknesses. In some cases this understanding is reinforced by later test cases.

4.3.1 Equations Solved

Rather surprisingly, the MARFA User Guide does not explicitly set out the equations that are solved. The equations given in the verification tests section are almost complete and so are repeated here:

$$\frac{\partial C_i}{\partial t} + v \frac{\partial C_i}{\partial x} - \alpha |v| \frac{\partial^2 C_i}{\partial x^2} = \frac{D_{eff}}{b} \frac{\partial M_i}{\partial z} \Big|_{z=0} - \lambda_i C_i + \lambda_{i-1} C_{i-1} \quad (4.3.1)$$

$$R_i \frac{\partial M_i}{\partial t} = \frac{D_{eff}}{\theta} \frac{\partial^2 M_i}{\partial z^2} - \lambda_i R_i M_i + \lambda_{i-1} R_{i-1} M_{i-1} \quad (4.3.2)$$

$$C_i(t,0) = \delta_{i1} f(t) \quad (4.3.3)$$

$$M(t, x, 0) = C(t, x) \quad (4.3.4)$$

$$M \xrightarrow{z \rightarrow \infty} 0 \text{ (unlimited diffusion)} \quad (4.3.5a)$$

$$\frac{\partial M_i}{\partial z} \Big|_{z=\Delta} = 0 \text{ (limited diffusion)} \quad (4.3.5b)$$

Appendix A gives the nomenclature used throughout this report.

The handling of dispersion in MARFA is consistent with a boundary condition “at infinity” in the fracture, that is to say that the system is treated as semi-infinite and the flux passing a particular point is reported as the flux at the end of the segment.

These equations assume no retardation within the fracture, but when limited diffusion in the matrix is considered, one of the inputs is k_a “a sorption coefficient for equilibrium sorption on fracture surfaces”. This is a nuclide-specific parameter with units given as [1/L]. This would seem to be incorrect – the expected units would be [L] and it would contribute to a fracture retardation value of $1 + \frac{k_{a,i}}{b}$ (here b is half the fracture aperture). A test case to check this is considered later. If the ES (equilibrium sorption) retention model is selected, a porous medium transport model is used – with retardation in equation (4.3.1) but no rock matrix – a test case for this is considered later.

The values that are allowed to be time-varying in MARFA are said to be the velocity and retention parameters. In fact, only the velocity and the matrix retention factors can be altered – the fracture surface sorption coefficient is assumed fixed.

4.3.2 Input files

The names of the MARFA input files are fixed, as are the names of the output files. The code is run by collecting the required input files in a directory and running MARFA in that directory.

The input files were found to be very user-unfriendly. There are no keywords in the files to act as a reminder of what the inputs are. Comments can be added to the end of the lines but this facility is not used in the test inputs supplied. The parameters that are required in the input seem to have been selected for the convenience of the code rather than the user. Rather than specify the basic physical properties, the user is required to calculate various combinations (e.g. $\kappa = \sqrt{\theta_m D_{eff} R_m}$). This is error-prone and hard to check. Evidence of this arises in Test 1, where the User Guide states that the matrix retention factors are 1000 and 100 but the supplied input file uses 500 and 100. Because none of these values appear explicitly in the input file, discovering this discrepancy took some time.

The *nuclides.dat* file lists the decay chains to be modelled. The decay constants are given and each nuclide is associated with an element. This association is by name, but in the remaining input files these names are never used – the order of the listed elements is crucial therefore – again, this is an error-prone approach. Each nuclide is also assigned an importance, which is used to scale the number of particles introduced – the User Guide warns that inappropriate use of this parameter can increase the variance in the results, by which it means that there may be a lack of convergence to the correct breakthrough curve.

The *source.dat* file starts by giving the number of particles. No advice is offered as to how many particles should be used. The test inputs use numbers ranging from 50 000 to 10 million. Next the units of the specified source are given. There is then an optional input line “UNIFORM IN TIME” which changes the generation of particles representing the source

term to be at equal time intervals rather than equal amount intervals. There is then a number of sources given, followed by a block for each source. The block gives the source a name and then associates it with a range of trajectories (as specified later in the *trajectories.dat* file). Again, rather than use the names of the trajectories their index numbers are used. This can be confusing – Test 1 has two trajectories named CAN1 and CAN2 and the name CAN1 is referred to in the source file. However, the trajectories line in the source file says “2 2” which indicates that the relevant trajectory is in fact CAN2! The source block ends with a line giving the number of times specified and then lines giving the time and source fluxes. The test files use uniformly spaced set of times, but presumably any monotonic time series is allowed. An example is given in the User Guide that suggests that the source becomes zero after the final specified time.

The *rocktypes.dat* file specifies retention properties and models for a set of rocktypes. The number of rocktypes is specified followed by a line that specifies the path to a data directory, which is only used when stochastic rocktypes are specified. The rocktype is then given a name (restricted to 5 characters) – these names are used in the trajectories file making it easier to see what is specified. One of three types of retention model is then selected (ES – equilibrium sorption, LD – limited matrix diffusion and MD – unlimited matrix diffusion). The parameters required depend on the model used. For ES, the overall porous medium retardation factor is required for each element. For MD, the $\kappa = \sqrt{\theta_m D_{eff} R_m}$ value for each element is required, while for LD κ , $\eta = \frac{\Delta}{D_{eff}}$ and k_a are required. As discussed

earlier, the units for k_a appear to be incorrectly stated – it should have units of length. When there are switches in properties (specified in *trajectories.dat*) the block of lines is repeated for each time period. Finally, a line specifies the dispersivity (which can be zero for a non-dispersive case), a flag indicating if this is a deterministic or stochastic rocktype (i.e. whether the subgrid downscaling is used). For stochastic rocktypes, two further parameters controlling the subgrid pathways are specified – a published paper is referenced to explain these.

The *trajectories.dat* file is written by CONNECTFLOW when MARFA is used for DFN networks. Here we focus on the self-contained version of this input that is used for the simpler test cases. The first line gives a path to the retention file data directory, where the pre-calculated retention functions are stored. This is followed by a number of trajectories. Each is given a name and a starting location followed by a set of segments. For each segment, the end coordinates are given, the rocktype specified (by name) and the τ and β parameters specified. τ is the water travel time for the segment and β is equivalent to the F-quotient. This is the quantity $1/(bv)$ integrated along the segment and so for a constant aperture segment we simply have $\beta = \tau/b$. In fact, these are the mean values before the dispersion model is applied. Finally, a hydraulic gradient is specified – this is only relevant for stochastic rocktypes. Some optional blocks can be specified. Most importantly a FLOW CHANGES block specifies the number of switches in

properties and gives the flow velocity scales after each time – these are scale factors compared to the initial period. The other blocks allow a subset of trajectories to be studied and specify a channelling factor. This latter option is only available in the CONNECTFLOW style input – it is used to scale the β values to reflect a reduction in contact area between the fracture and matrix if channelling is present.

When stochastic rocktypes are used a set of subgrid trajectories must be specified – we do not consider that option in the current review.

The final input file specifies the post-processing options. The file *postprocessor.dat* starts with the “gamma parameter” which controls the reconstruction of breakthrough curves, by controlling the range of times over which smoothing occurs. A value of 0.2 is recommended and it is suggested that slightly higher values may sometimes be useful. This is followed by a time of first release – it is not clear why the user has to specify this as it could be obtained from the results file. The next line allows all or only a subset of the trajectories to be analysed. When a partial set is selected the next few lines specify which trajectories to use. Finally, the times at which breakthroughs are to be reported are specified. This can be done as an explicit list of times (manual) or as a uniform or logarithmic (presumably geometrically spaced). The User Guide does not make it clear how the minimum and maximum output times are specified - looking at the source code suggests that they are automatically generated from the particle arrival results.

4.3.3 Running MARFA

Running MARFA is straightforward once the input files are set up. The run time obviously depends on the number of particles. A test with 40 million particles took around 2 minutes on a desktop PC to calculate the particle results and then a further 3 or 4 minutes to post-process these. While calculating, the memory requirement is small (less than 1 Mb) but during post-processing it rises significantly (to 770 Mb in the example) presumably because the results are all read back into memory. This memory requirement restricts the maximum number of particles to the order of 100 million. When too many particles are used, the code calculates the *results.dat* file but then crashes when available memory is exceeded. Run times with multiple segments are higher – a test with 10 segments in place of a single one took 5 times longer to calculate the particle results, with no change in the post-processing time.

Each time MARFA is run for the same input a different set of results is produced. This is because the stochastic sampling is initialised using the system clock. Running the same case several times gives an indication of the stochastic variability in the results. For testing and for sensitivity analysis it would be useful to be able to fix a particular random number seed so that the same results could be obtained repeatedly.

The only problem encountered in running the code was with the “time of first release” in the post-processor. For Test 1 this is set to 499, but when

more particles are used the dispersion can lead to particles arriving before that. The code spots this and writes an error message – unfortunately this message is followed by a report on every particle and so swamps the console with numbers and hides the message itself. It was necessary to redirect the output to a file in order to be able to read the error message.

4.3.4 Tabulated Response Function Files

The tabulated response function files, stored in the data directory, are crucial to the correct running of the code. They are only mentioned in passing in the User Guide and there is no discussion about how the level of discretisation (or time and η and the cumulative quantiles) was determined. Moreover, these files are supplied as text files, which makes them susceptible to accidental modification or to uncontrolled change. While it was convenient for the testing undertaken here to have these in text form, it would be better from a QA point of view to prevent modification – either by storing them as binary files or by incorporating a checksum in them that the code could cross check.

The numerical content of these files is examined in section 4.4.1.

4.4 Algorithms

The MARFA code uses many separate algorithms in order to calculate and report the final breakthrough curves. The level of documentation for these varies as does its location. Here we identify the key algorithms, record where they are documented and analyse them for accuracy.

In order to check aspects of the MARFA approach, we utilise the Laplace Transform approach. In particular, we use the Talbot inversion algorithm [23] which gives accurate results for the type of function of interest here. The inversion process can essentially be regarded as exact. This enables precise numerical checks to be made against the MARFA algorithms which use time-domain calculations of response function that can often involve infinite sums where convergence may be an issue.

4.4.1 Basic Particle Tracking

The basic algorithm in MARFA is the calculation of the distribution of travel times for a particle travelling through a single segment. In the first instance we can consider a non-decaying particle with constant properties and no dispersion. We start with the case of an infinite matrix, which is simpler to analyse than the finite matrix case (the equilibrium sorption case is not of particular interest here).

The CDF for the retention time t_{ret} for this case is given in [24] as $\text{erfc}\left[\frac{\kappa\beta}{2\sqrt{t_{ret}}}\right]$. This can be deduced directly from the basic Laplace transform solution for a continuous source input. In the MARFA notation this is $\frac{e^{-\tau s - \kappa\beta\sqrt{s}}}{s}$. The first term in the exponent represents a time shift of τ , the water travel time. The second term is the matrix response and inverts to the function given.

Rather than calculate the inverse error function on demand in order to get the sampled retention times, MARFA uses a stored calculation in the file *fable.dat*. By examining the source code, it is clear that the stored distribution is for $t^* = \frac{4t_{ret}}{\kappa^2\beta^2}$, that is the function is the inverse of

$$\text{erfc}\left[\frac{1}{\sqrt{t^*}}\right] \text{ corresponding to a Laplace transform of } \frac{e^{-2\sqrt{s}}}{s}.$$

In *fable.dat*, values are given for quantiles of 1e-6, 1e-5, 5e-5, 1e-4, 5e-4, 1e-3, 2e-3, 3e-3, 5e-3, 1e-2 and then at 1e-2 intervals until 0.94 after which 0.945 and 0.95 are specified. For values above 0.95 there is a comment in *fable.dat* that a power law proportionality should be used: $t^* \propto (1-r)^2$, where r is the required quantile. Inspecting the source code confirms that this is what was done. We can easily check that the values given are correct and that the extrapolation is valid for the large quantiles. This has been done and we confirm that the figures given are accurate to at least 5 significant figures.

Between the values given, a linear interpolation is used. In order to see how much of an approximation this is, we have calculated the interpolated result at the midpoint between each pair of results and checked this. The results are much less satisfactory. The largest relative errors are for the lowest quantiles, while the largest absolute differences occur for higher percentiles. The early relative errors might be important for cases where the early breakthrough is most important. Table 7 gives some examples.

Table 7 Check on the quantiles used in the *ftable.dat* file for infinite matrix diffusion

Quantile	Time calculated by interpolation	True quantile at that time	Discrepancy (True Quantile – Intended Quantile)
0.0000005	0.0417918	4.58633E-12	-5.00E-07
0.0000055	0.0930438	3.54706E-06	-1.95E-06
0.00003	0.11204925	2.39108E-05	-6.09E-06
0.0003	0.1486025	0.00024387	-5.61E-05
0.00075	0.174895	0.000720567	-2.94E-05
0.0015	0.197074	0.001444209	-5.58E-05
0.0025	0.218257	0.002468887	-3.11E-05
0.004	0.2404525	0.003926166	-7.38E-05
0.015	0.33549643	0.014623129	-3.77E-04
0.025	0.397124315	0.024822797	-1.77E-04
0.075	0.630872305	0.074992649	-7.35E-06
0.125	0.849885265	0.125021445	2.14E-05
0.275	1.6786956	0.275047323	4.73E-05
0.375	2.5418189	0.375057927	5.79E-05
0.475	3.92035405	0.475070467	7.05E-05
0.575	6.36449125	0.575087807	8.78E-05
0.675	11.3845455	0.675115201	1.15E-04
0.775	24.515592	0.775166569	1.67E-04
0.875	81.211197	0.875299712	3.00E-04
0.885	96.1548485	0.885325722	3.26E-04
0.895	115.607355	0.895356668	3.57E-04
0.905	141.589395	0.905394098	3.94E-04
0.915	177.399445	0.915440286	4.40E-04
0.925	228.72698	0.925498706	4.99E-04
0.935	306.094115	0.93557494	5.75E-04
0.9425	386.624805	0.942662877	1.63E-04
0.9475	464.433945	0.947678345	1.78E-04

We shall see whether these discrepancies are evident in the actual calculations when we look at specific test cases.

We also note that a piecewise-linear interpolation of the cumulative arrival time distribution is equivalent to a piecewise-constant interpolation of the arrival rate distribution. This would lead to a “staircase” effect in the output

in the absence of other smoothing processes such as dispersion. Again, this issue is addressed in the test cases.

4.4.2 Finite Matrix

For the finite matrix case, the Laplace transform of the retention function is reported in [24], but there is no discussion of how this is inverted. In this case, the Laplace transform solution for a continuous source input in the

MARFA notation is $\frac{e^{-\tau s - \kappa \beta \sqrt{s} \tanh(\kappa \eta \sqrt{s})}}{s}$. Again, the first term in the

exponent represents a time shift of τ , the water travel time and the second term is the matrix response. In this case, MARFA uses a stored calculation in the file *limdiff.dat*. This is a two-dimensional table, for the quantile and a scaled version of the η parameter. By examining the source code, it is clear

that the stored distribution is for $t^* = \frac{t_{ret}}{\kappa^2 \beta^2}$ (note the factor of 4 used in the infinite matrix case is not used here), and this has corresponding Laplace transform of $\frac{e^{-\sqrt{s} \tanh(\eta^* \sqrt{s})}}{s}$, where $\eta^* = \frac{\eta}{\beta}$. We note that the form given in [24] does not show the s dependence correctly.

Since we do not have an explicit time-domain representation, the checking is more difficult. We can use the Talbot inversion algorithm to evaluate the function at any time and η^* , so we start by evaluating it at the times in *limdiff.dat* for the various η^* that it contains. The range of η^* values in the table is 0.01 to 562.3 – we discuss what happens outside this range later. Within this range, the values in *limdiff.dat* appear to be rather inaccurate. For the larger values of η^* the accuracy is reasonable, but not nearly as good as for the unlimited diffusion case which it should be identical to (up to the factor of 4). For the smaller values of η^* the accuracy is poor. The result becomes more and more like a step function as η^* reduces, so this inaccuracy may not be significant (there will be little spread in the arrival times). As an example, the $\eta^* = 0.1$ values are compared in Table 8.

Outside the range of η^* in the file MARFA takes an equilibrium (step-function response) for small η^* and uses the largest η^* result for larger η^* . Both of these are sensible.

Table 8 Check on the quantiles used in the *limdiff.dat* file for limited matrix diffusion

Quantile	Time given in limdiff.dat	True quantile at that time	Discrepancy
1.00E-06	0.019582	5.22218E-07	4.78E-07
0.00001	0.02357	6.02462E-06	3.98E-06
0.001	0.037303	0.000892668	1.07E-04
0.01	0.048071	0.008386922	1.61E-03
0.05	0.060068	0.044063732	5.94E-03
0.1	0.067908	0.095779758	4.22E-03
0.2	0.077442	0.19524465	4.76E-03
0.3	0.084741	0.29484252	5.16E-03
0.4	0.091427	0.39694977	3.05E-03
0.5	0.098259	0.50414922	-4.15E-03
0.6	0.10531	0.60993702	-9.94E-03
0.7	0.11236	0.70451599	-4.52E-03
0.8	0.12304	0.81892248	-1.89E-02
0.9	0.1353	0.90624888	-6.25E-03
0.95	0.15038	0.96331262	-1.33E-02
0.99	0.17348	0.99315472	-3.15E-03
0.999	0.19752	0.99908508	-8.51E-05
0.9999	0.23376	0.99997105	-7.10E-05
0.99999	0.26238	0.99999859	-8.59E-06
0.999999	0.28403	0.99999988	-8.80E-07

The issue of interpolation is similar to the unlimited diffusion case, but is worse here because a much coarser set of quantiles is used (mainly 0.1 intervals as shown in the table above). Why 1e-4 is omitted is unclear. In any case the interpolation will be poor at this early arrival end of the retention function.

Overall, the limited diffusion response functions appear to be much less accurate than the infinite diffusion ones. The tabulated values are themselves less accurate and the interpolation errors are larger. This is manifest later in a test case.

4.4.3 Dispersion

MARFA handles dispersion by sampling a travel time from a distribution that depends on the dispersivity. For a system without matrix diffusion, a velocity of v , a path length of L and a dispersion length of $\alpha = L/Pe$, we can calculate the flux passing position L for a unit injected flux into a semi-

infinite system. The Laplace transform can be written $e^{-\frac{Pe}{2}\left(\sqrt{1+\frac{4s\bar{\tau}}{Pe}}-1\right)}$, where $\bar{\tau}$ is the mean water travel time. This can be inverted to give

$\sqrt{\frac{Pe}{4\pi(\tau')^3}}e^{-\frac{Pe(1-\tau')^2}{4\tau'}}$ where $\tau' = t/\bar{\tau}$, which is the result quoted in the

MARFA User Guide with Pe used in place of η .

MARFA samples from this distribution as described in [24] by solving the cumulative version for the time for each sampled cumulative position. It is unclear why this approach is used here while look-up tables are used for the more important matrix retention functions.

Once a value of the travel time has been sampled, a consistent value for β is calculated and used in the matrix retention calculations. This is clearly a sensible approach, although it is not precisely correct it should be sufficiently accurate for the advectively dominated cases of interest here. This is further investigated in the testing.

4.4.4 Multiple Segments

MARFA handles multiple segments by calculating the travel time for each segment separately as the particle passes through the system. There is an implicit assumption in this approach that there is no feedback between the segments which in turn is equivalent to an assumption of an advectively dominated system. The basic assumption for a single segment is that it can be treated as semi-infinite which is valid in the same circumstances.

Thus, the handling of multiple segments in MARFA is valid for the expected case of advectively dominated systems. It would not be appropriate to use the code when diffusion within the fractures was significant. For normal flow conditions this seems to offer no serious restriction, but there is a possibility that low flow conditions (e.g. when the surface is covered by an ice sheet) might be of interest. In such cases, checks should be made on the relative importance of the diffusive and advective transport processes.

4.4.5 Decay and decay chains

Radioactive decay is handled stochastically in MARFA. Given the specified decay rate a decay time is calculated and compared to the calculated travel time. If the decay time is less than the travel time then the particle has decayed in the segment. If the nuclide is at the end of a chain (so its daughter is not modelled) then the particle is simply eliminated from further

consideration. This approach is clearly correct, but has the effect of reducing the number of particles that arrive and so makes the variance in the final results greater.

When the daughter nuclide is of interest the particle changes to that daughter at the decay time and the rest of the segment travel is calculated with the daughter properties.

MARFA calculates the overall travel time in a simple fashion. The travel time for the daughter (for the whole segment) is calculated and a weighted average of this and the parent travel time is used. If t_A is the parent travel time, t_B is the daughter travel time, and t_d is the time at which the decay

occurs, then the overall travel time is taken to be $t_d + \left(1 - \frac{t_d}{t_A}\right)t_B$. It is noted

in the User Guide that the sampling for t_A and t_B must be fully correlated for this to work. This is clearly the case when the sorption properties are the same - t_A must equal t_B otherwise the correct retention distribution is not being respected.

An accurate algorithm for handling decay would need to calculate the position of the parent nuclide at the time of decay and then calculate the travel time for the daughter from that point. This is precisely the approach that is taken to handling time switches (see the next section). From the point of view of a particle, a decay that leads to a change of properties and a time switch that leads to a change of properties seem to be exactly the same, so it is unclear why different algorithms are used for handling the two situations.

The extent of the approximation introduced by taking the simple averaging approach is not immediately clear. This is explored further in the testing.

4.4.6 Time switches

MARFA allows the velocity and matrix retention properties to change at specified times. When such a change occurs, a series of calculations are undertaken to determine the overall travel time.

First, the depth of the particle in the matrix at the time of the switch is calculated. Then the time for the particle with the new properties to return to the fracture is determined and finally the retention time for the rest of the segment is calculated. Each of these is independently stochastically determined. The documentation does not make it clear how the fracture travel time is handled if the velocity changes and looking at the source code suggests that this aspect is being ignored. The fraction of the path that has been travelled is also not described in the documentation; the code contains a calculation that uses the square root of time to calculate the fraction travelled.

The distance into the matrix and time to return to the fracture are both sampled with the use of pre-calculated look-up functions. We look at the infinite matrix case first.

The User Guide gives an approximate calculation which ignores advection. It is noted that the full calculation is given in [25]. Comparing these formulae it is clear that the approximation is accurate if the matrix retention is dominant. We also note that the probability that the particle is in the fracture at the switch time appears to be ignored in the MARFA algorithm – again an assumption that the matrix dominates. This is also consistent with the apparent ignoring of the change in water travel time when velocity changes.

Thus, MARFA is likely to produce inaccurate calculations for cases where the effect of the matrix is weak (i.e., the time spent by radionuclides in the matrix is not the dominant contributor to overall transport times).

We note that the formulae given for G (a coefficient in the matrix depth probability density function) and t_0 (a dimensionless time) in this section of the MARFA User Guide are both missing porosity factors. The definition of the diffusion coefficient was changed from that used in the published papers. It is believed that the factor is correctly handled in the code.

The depth and return time functions are both normalised to remove the G dependence. They are then used together to give a return time. To do this the files *table1.dat* and *table2.dat* are used. The *table2.dat* file contains the inverse of the error function, evaluated at 0.01 intervals except at the top end where values for 0.999 and 0.9999 are also tabulated. These figures have been checked – they are accurate, so the only issue is interpolation error. The *table1.dat* file contains the required inverse of $\text{erfc}(1/\sqrt{t})$ which is needed for the return times. This again has been checked and found to be accurate. It is tabulated at generally 0.01 intervals with additional points at 0.005 intervals above 0.9 and at $1\text{e-}6$, $5\text{e-}6$, $1\text{e-}5$, (not $5\text{e-}5$), $1\text{e-}4$, $5\text{e-}4$, $1\text{e-}3$ and $5\text{e-}3$. As with the other tables, interpolation errors may be an issue.

The finite case is again more complicated. Both tables, *LD_table1.dat* and *LD_table2.dat* are two dimensional. They have not been checked in detail, but both use quite a coarse set of values so interpolation errors may be significant. The consequences of this are considered in the testing section.

4.4.7 Mass Discharge Calculation

Reconstruction of the rate of discharge from the particle arrival times is handled using an adaptive kernel method, which adjusts the degree of smoothing according to the local density of data values (less smoothing is used if there are more data points). This seems to be a sensible approach and is widely used. The details have not been reviewed, but it is expected that when there are “enough” particles arriving to get a sensible estimate then the methods used will do that. If there are not enough particles then no method can reconstruct the missing detail.

4.4.8 Stochastic Downscaling

The idea of filling in missing detail from a porous medium model of a fracture system by using small scale networks is a topic for a review of the whole approach to modelling fracture systems. In the current review, we have focussed on the deterministic aspects.

It is understood that the approach used in MARFA is to import trajectories calculated from the subgrids and to scale there parameters to match the travel time and F-quotient for the segment being replaced. As a basic approach this seems sensible.

If the subgrid approach is widely used in the SR-Site work then a more detailed review of that aspect will be necessary.

4.4.9 Other aspects

Various other aspects of the algorithms are mentioned in the User Guide.

The importance value for a nuclide allows the source strength for different nuclides to be weighted differently. It is possible to specify that uniform times are used for sources rather than uniform amount intervals.

When stochastic rocktypes are used, a channelling parameter can be specified to scale the effect of matrix diffusion.

None of these algorithms is documented in any detail and no tests are presented to verify their usage.

The source related algorithms have been checked in the testing described in Sections 4.5 and 4.6.

4.5 Documented Tests

Seven verification tests are documented. These tests include cases with:

- ▲ limited and unlimited diffusion
- ▲ multiple paths
- ▲ changes in flow rate; and
- ▲ changes in retention properties.

However, the tests do not look at:

- ▲ more than two nuclides in a chain;
- ▲ different choices for dispersivity;

- ▲ convergence with number of particles; or
- ▲ different decay rates, particularly shorter lived daughters.

All the tests are fairly minor variants about a central case. This gives no indication of how the code will perform in a real case. An example is given with a 4-member chain but only results for the 2nd member are reported.

In each case, results are compared to a benchmark result. The description for Test 1 suggests that this uses a discretisation for the matrix response and the method of lines for the fracture, but points out that there may be numerical dispersion in the results. In [24] comparisons against GoldSim calculations are presented, but no details of the discretisation scheme are given.

As discussed in Appendix B, we have developed a capability to solve the equations using a multi-Laplace transform approach. This enables accurate results to be calculated for the test cases with switches. For the cases without switches, the solutions developed in [3] can be used. The MARFA input files for each of the tests were provided, so we have rerun each of the tests.

4.5.1 Test 1

Test 1 is presented as having two nuclides each with 10 000 year half lives. In fact, the supplied test input contains a third nuclide with a half life of 1010 years. We include this third nuclide here. Other than the half life it has the same properties as the second nuclide.

The supplied *source.dat* file uses just 50 000 particles. It seems unlikely that this was the number used in the presented results; here we also present a case with 5 million particles. Figure 25 shows the comparison.

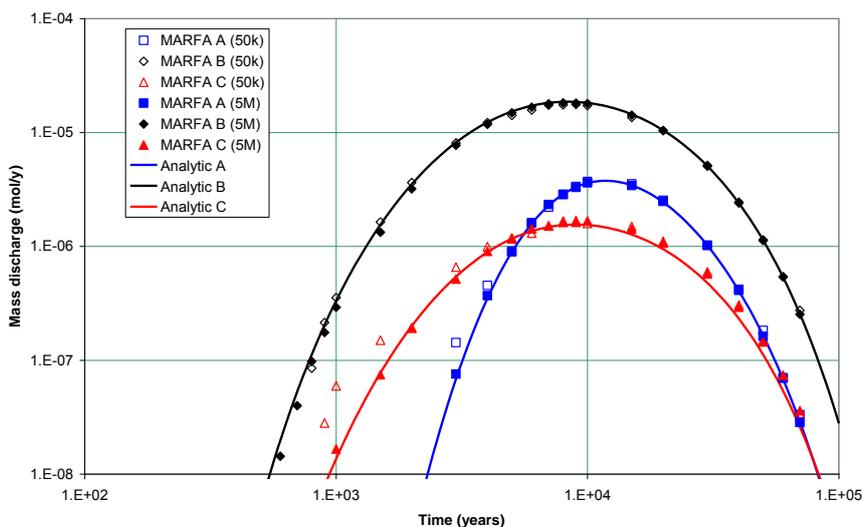


Figure 25: Comparison of MARFA calculations with the analytic calculation for Verification Test 1

The comparisons for the first nuclide are excellent for the larger number of particles. With fewer particles the early time calculations are less accurate as would be expected. The second nuclide shows similarly good comparisons, even for the smaller number of particles. For the third nuclide the early time effect is more severe for the small number of particles. There is also a definite tendency to give higher calculated fluxes at later times for this nuclide, possibly an indication of the errors introduced in handling decay using the weighted averaging approach although these would not be expected here because the properties are the same as its parent.

4.5.2 Test 2

Test 2 is also presented as having two nuclides but a third is included. In addition to limiting diffusion, the nuclides have different rock matrix retention factors (200, 500 and 100). The documentation suggests that the diffusion depth is set as 2 mm; however inspection of the input file reveals that 4.75 mm was in fact used.

The supplied *source.dat* file uses 5 million particles. Figure 26 shows the comparison.

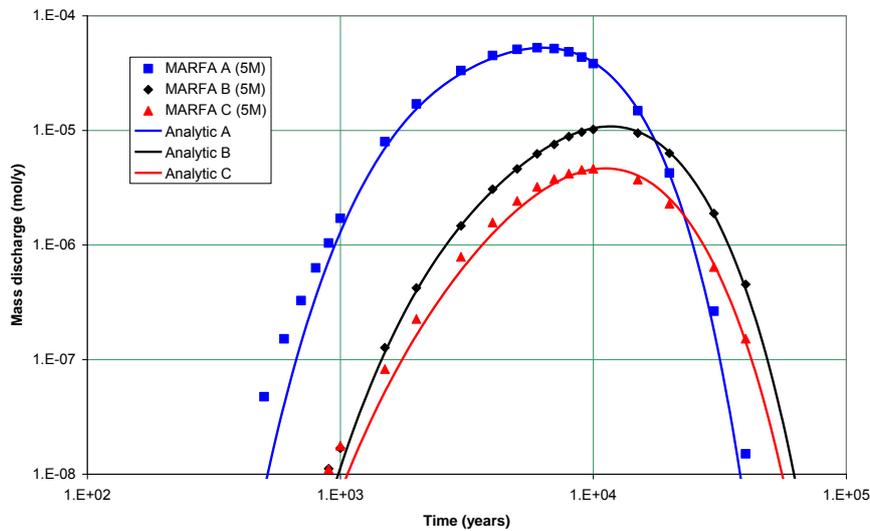


Figure 26: Comparison of MARFA calculations with the analytic calculation for Verification Test 2

The comparisons are not quite as good as for the infinite matrix case, with a systematic deviation at early times that may be due to the issues identified with the *limdiff.dat* file and interpolations.

For this case the value of η^* that would apply is 0.475 which means that interpolation errors should be small, but from the checking of *limdiff.dat* (section 4.4.2) is expected to overestimate early arrivals. Looking at the cumulative discharge shown in Figure 27, we can see that the early behaviour is poor. With 5 million particles, the number that has arrived for the $1e-5$ total discharge is only 50 on average – so the inaccuracy is possibly just statistical. By $1e-4$, 500 have arrived on average, so the factor of two discrepancy between the MARFA and analytic calculations is more significant and can be directly traced to the inaccurate quantiles in the file.

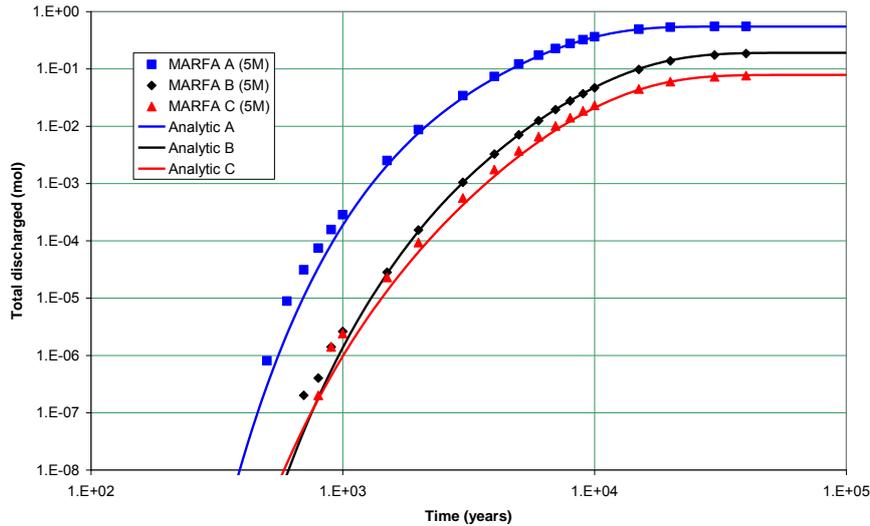


Figure 27: Comparison of cumulative MARFA calculations with the analytic calculation for Verification Test 2

A further possible reason for the difference is in the source term itself. The MARFA source term is rather coarsely discretised. This is easily tested by running a MARFA case with a finer resolution in the source. We used a time interval of 50 years up to 5000 years instead of the 250 year interval used previously, but the results are unchanged.

Another possible explanation of the early time discrepancy is the treatment of dispersion. To check this, a case was run with no dispersion. This shows exactly the same behaviour.

We are left with the inaccurate calculation of the quantiles in *limdiff.dat* as the most likely explanation for the early arrival times in MARFA for the unlimited diffusion case. We note that the comparison given in the User Guide shows the same behaviour, but this is not commented on in that document.

4.5.3 Test 3

Test 3 is for multiple sources. This test has not been repeated here. The calculations show the same behaviour as Test 2.

4.5.4 Test 4

Test 4 looks at changing the flow velocity at two switch times. As with the other cases, three nuclides are used but only the first two are reported in the User Guide. It is claimed that the other properties are for Test 1, but it is evident from the arrival curves before the first switch that this is not the case. Examining the input file reveals that the matrix retention for the first nuclide has been set to 500 (it was 1000 in Test 1). The test uses 10 million particles. The limitations of the particle approach are evident in the results obtained for the low flow period. Here the arrival rates are of the order of $1e-9$ mol/y, which corresponds to one particle every 100 years, or a total of 50 over the low flow period. With such small numbers, the results are bound to be subject to large statistical variability.

With the new capability described in Appendix B, we can obtain analytic calculations for this case; the comparison is shown in Figure 28.

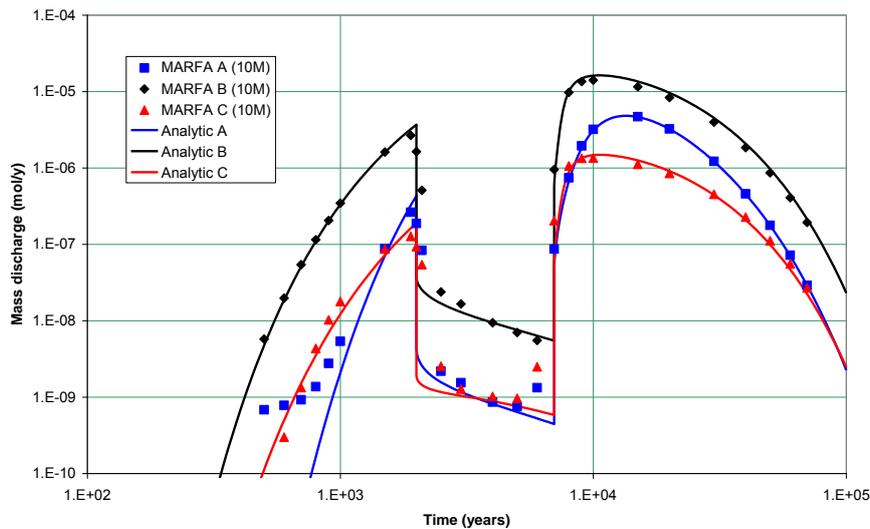


Figure 28: Comparison of MARFA calculations with the analytic calculation for Verification Test 4

The calculations match well except for the first nuclide at early times. We note that the benchmark solution reported in the User Guide appears to be inaccurate here (presumably due to numerical dispersion). There is a tendency for the MARFA calculations to smooth the transitions and to underestimate slightly the daughter nuclide peaks after the second switch.

4.5.5 Test 5

In addition to limiting diffusion, as claimed in the User Guide, the nuclides have different rock matrix retention factors (200, 500 and 100) compared to Test 4. The diffusion limit is 4.75 mm as in Test 2. The User Guide gives calculations from a single segment and then using 10 segments. We first compare against the single segment calculation in Figure 29.

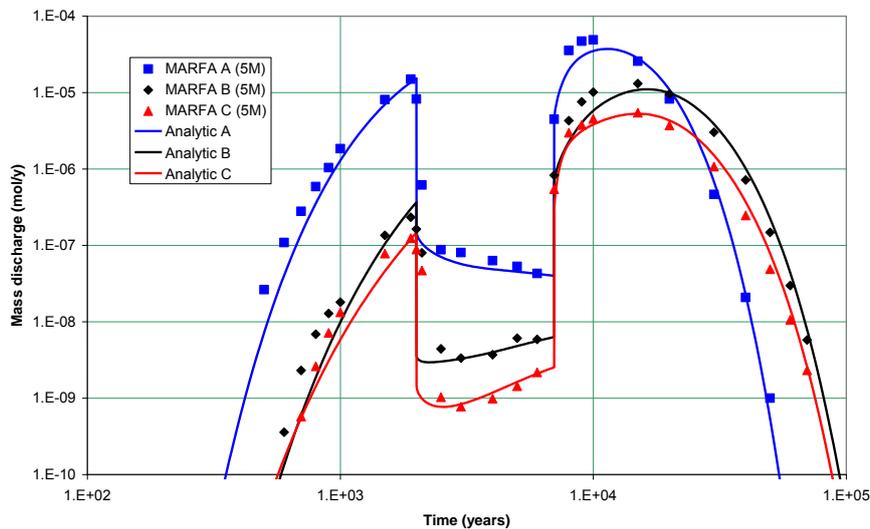


Figure 29: Comparison of MARFA calculations with the analytic calculation for Verification Test 5

The agreement is generally poorer than for the infinite matrix case. The User Guide suggests that the inaccuracy is due to the approximations involved when decay and a time switch occur in the same segment. A test is given with 10 segments which we compare against in Figure 30. The early calculations and the peak after the second switch are clearly much better handled but calculations for the daughter during the low flow period appear to be worse (the numbers of particles are small so this may be a statistical artefact).

The explanation put forward in the MARFA documentation does not seem plausible – the same effect ought to be seen in the infinite diffusion case. We believe that the true explanation relates to the coarse discretisation in *limdiff.dat* (the peak for the first nuclide corresponds to around a 10% overall arrival which is the point at which the discretisation becomes coarse). Using more segments effectively reconstructs the overall breakthrough as a convolution – with contributions from across the range of arrival times. This smoothes the breakthrough times and counteracts the coarse discretisation.

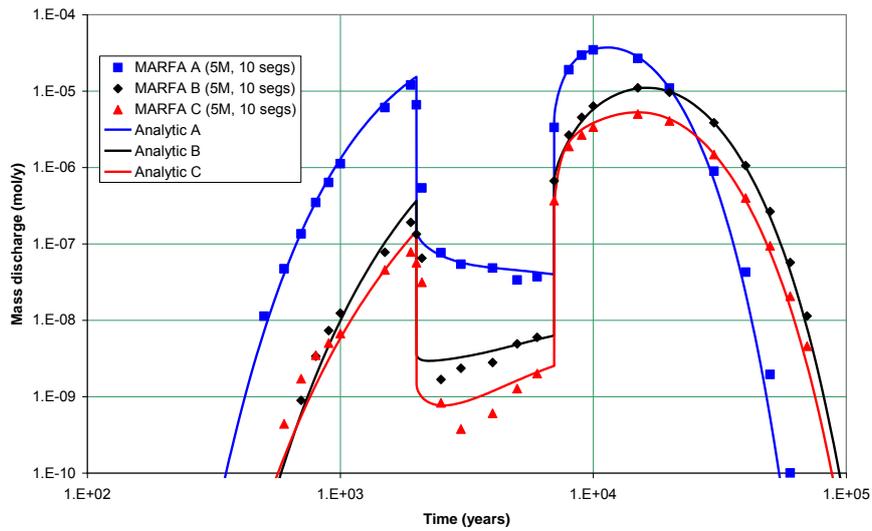


Figure 30: Comparison of MARFA calculations with the analytic calculation for Verification Test 5 (with 10 segments)

4.5.6 Test 6

In Test 6 a high flow period is used. It is stated that the test is otherwise the same as Test 4. In fact it is quite different. The decay constants for the first two nuclides are 0.001 and 0.0001 y^{-1} ; the matrix retardations are both set to 200 and the dispersivity is doubled to 1.0 m. The test also uses the 10 segment approach.

We compare against the analytic solution in Figure 31.

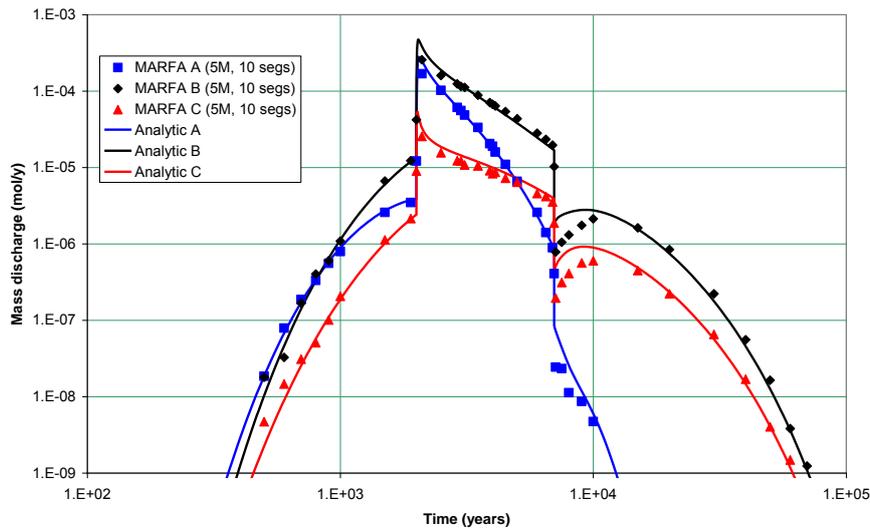


Figure 31: Comparison of MARFA calculations with the analytic calculation for Verification Test 6

The agreement is generally good up to the second switch. It then becomes rather inaccurate until the falling portion when the accuracy is restored. This effect is seen in the comparison presented in the User Guide.

4.5.7 Test 7

Test 7 is said to be identical to Test 5 except that the matrix retardation factors are reduced by a factor of 10 during the low flow period. In fact it is based on Test 4 with the retardation reduced for the first two nuclides.

Figure 32 shows the comparison. The calculations agree reasonably well except at early times – as they did in Test 4. We note that the analytic calculation shows a short spike release immediately after the second switch – this arises because the concentration in the fracture is higher in the middle phase (because the matrix retardation is lower) and this water is flushed when the flow rate increases. The total amount released in this spike is small as can be seen in the MARFA calculations.

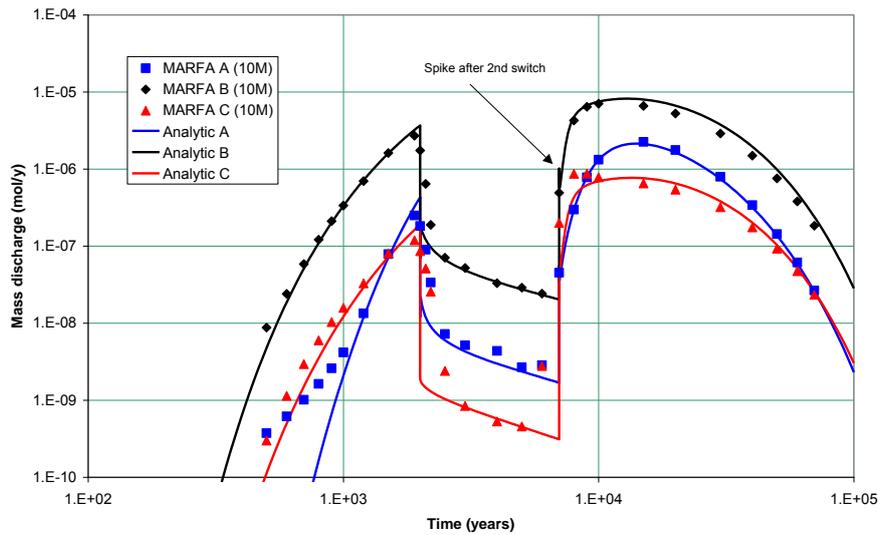


Figure 32: Comparison of MARFA calculations with the analytic calculation for Verification Test 7

4.5.8 Example Simulation

The example simulation uses a 4-member chain with a short lived first member. The User Guide shows calculations only for the second nuclide (Np237). The figure has a mislabelled time axis (powers are shown negative). It is difficult to resolve the spikes at 10 000 and 11 000 years.

The non-glaciation case calculations are much smoother than the glaciation case, but this appears to be due to using different settings in the post-processing rather than anything fundamental (the glaciation case uses a value of gamma of 0.6 – whereas the non-glaciation case uses 0.2). This was presumably necessary to resolve the early spikes, but is not discussed in the User Guide.

This appears to have a detrimental effect on the daughter nuclides, as can be seen in Figure 33.

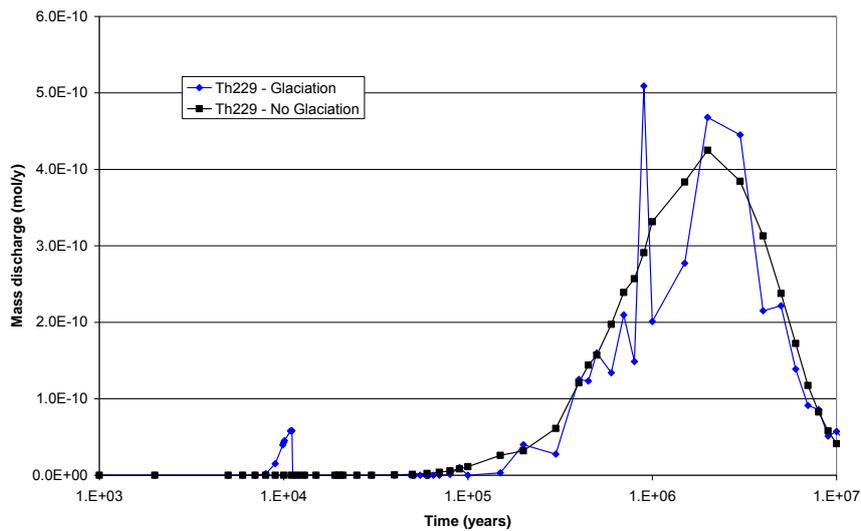


Figure 33: MARFA calculations for Th229 for Acceptance Test 1

4.5.9 Conclusions on Documented Testing

Two main features of the documented testing stand out:

- ▲ the tests are limited in scope and too similar to each other; and
- ▲ the descriptions given are inaccurate in many places.

The tests that are reported show that the code functions correctly in the broadest sense – it gives correctly shaped arrival distributions. There is some evidence that the limited diffusion calculations are poorer than the infinite diffusion cases, reflecting the poor accuracy and coarse discretisation in the supplied data tables.

The acceptance test suggests that handling short spikes is problematic – setting the post-processor gamma parameter high to capture the spikes tends to introduce instabilities at other times.

4.6 Additional Tests

In this section, additional test cases are presented to investigate aspects of MARFA's behaviour not covered by the documented tests. Clearly there are many tests that could be added, but here we focus on some key areas that may be relevant to the usage of MARFA in a safety case. These are:

- ▲ staircase calculations for limited diffusion;
- ▲ a wider range of dispersivity;
- ▲ cases where matrix diffusion is less important;

- ▲ the use of fracture surface sorption;
- ▲ handling of realistic decay chains; and
- ▲ short-term changes in properties.

These are discussed in the following sections.

4.6.1 Staircase Calculations

We ran a very simple MARFA case based on Test 2. The dispersivity was set to zero, the source term was set to a top-hat function lasting one year and the decay was set to zero. The output should therefore reproduce the basic limited diffusion breakthrough curve. The output times were set to 100 logarithmically spaced.

Figure 34 shows the results. The coarse discretisation together with the linear interpolation lead to a “staircase” effect in the MARFA calculations – the discharge rate is constant between times when the overall breakthrough reaches the tabulated values (e.g. $1e-3$ is reached around 1000 years and $1e-2$ is reached near 1600 years).

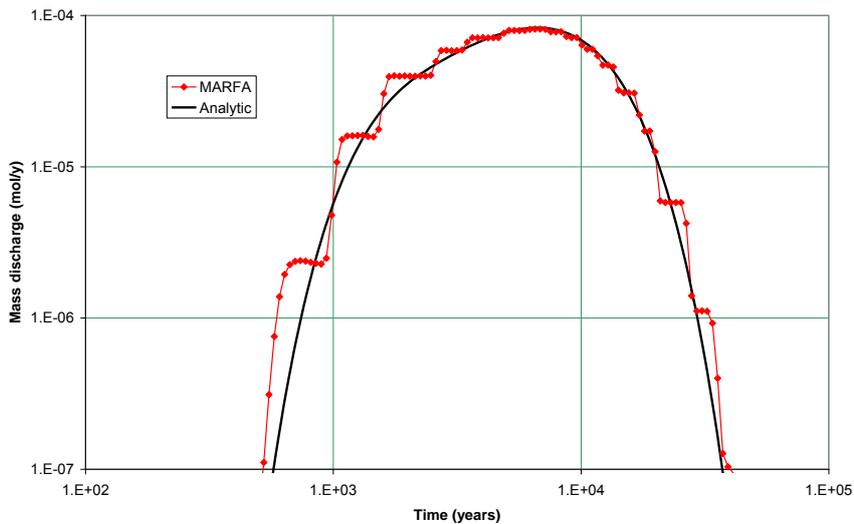


Figure 34: Comparison of MARFA calculations with the analytic calculation for a simple limited diffusion case

Clearly, this effect is generally hidden by dispersion or by a long-term source. It will also be hidden when there are a series of segments. This is therefore unlikely to be an issue in the practical use of MARFA unless it is used to calculate such simple breakthrough curves.

4.6.2 Wider range of dispersivity

The dispersivity used in all the test cases is low – generally 0.5 m with one case using 1 m. In some cases, especially if MARFA were to be used to represent the whole geosphere pathway, much higher values may be appropriate, corresponding to Peclet numbers of around 10.

We therefore created a modified version of Test 4 with a dispersivity of 10 m. The comparison is shown in Figure 35.

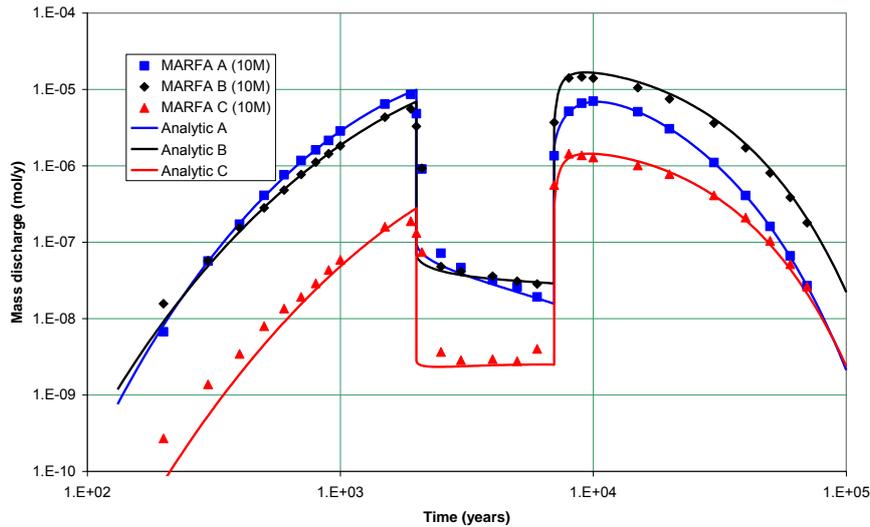


Figure 35: Comparison of MARFA calculations with the analytic calculation for a higher dispersivity

The calculations match well, even for the first nuclide at early times. The early breakthrough of the third nuclide is slightly too early, but the comparison is much the same as in Test 4. So, higher dispersivity cases appear to be handled correctly.

4.6.3 Weak matrix diffusion

Here we take a case where the matrix effect is much less dominant. It is based on Test 5 but the matrix depth and matrix retardation were reduced by a factor of 10 (giving 100 times less matrix capacity). The dispersivity was increased by a factor of 10 (numerical difficulties were experienced in the analytic calculations with the smaller dispersivity).

Figure 36 shows the results. The agreement is significantly poorer than for the dominant matrix cases, confirming that there is an implicit assumption in the MARFA approach that radionuclide transport in the matrix dominates.

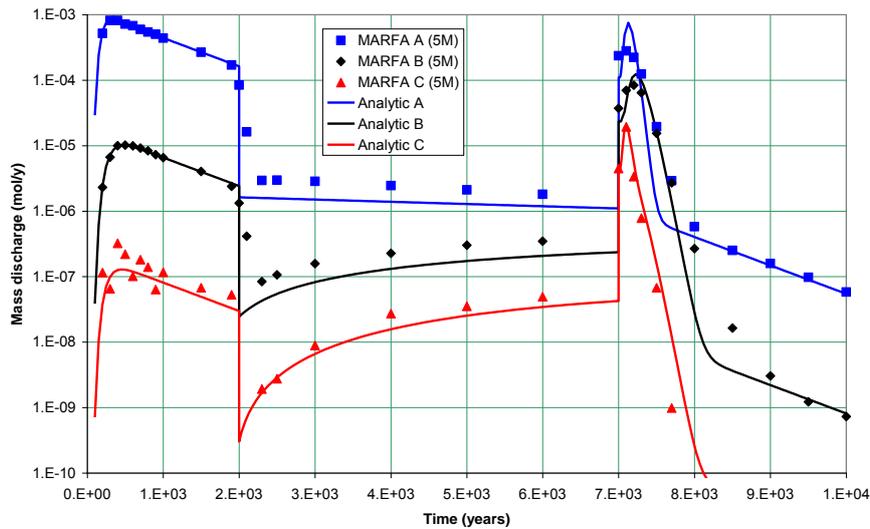


Figure 36: Comparison of MARFA calculations with the analytic calculation for a less dominant matrix

It is therefore concluded that MARFA cannot be used for cases where the matrix has only a small effect.

4.6.4 Fracture surface sorption

Fracture surface sorption is only available in MARFA for the limited diffusion case. We therefore take a variant of the previous test with the k_a set to $1e-3$, $2e-3$ and $3e-3$ for the three nuclides. As discussed earlier, we believe that this has units of metres and that it corresponds to a fracture retardation factor of $1 + k_a / b$, that is 11, 21 and 31 in this case.

Figure 37 shows the comparison. There is a clear difference of about a factor two in the results after the first time switch. The MARFA rate falls by a factor of 50 across the switch, whereas the analytic calculation gives a factor of 100, reflecting the change in velocity. Increasing the retardation in the fracture make the matrix even less dominant, and this therefore confirms the previous results showing that MARFA is inaccurate in these cases.

The good match before the switch confirms the interpretation of how the k_a value is used.

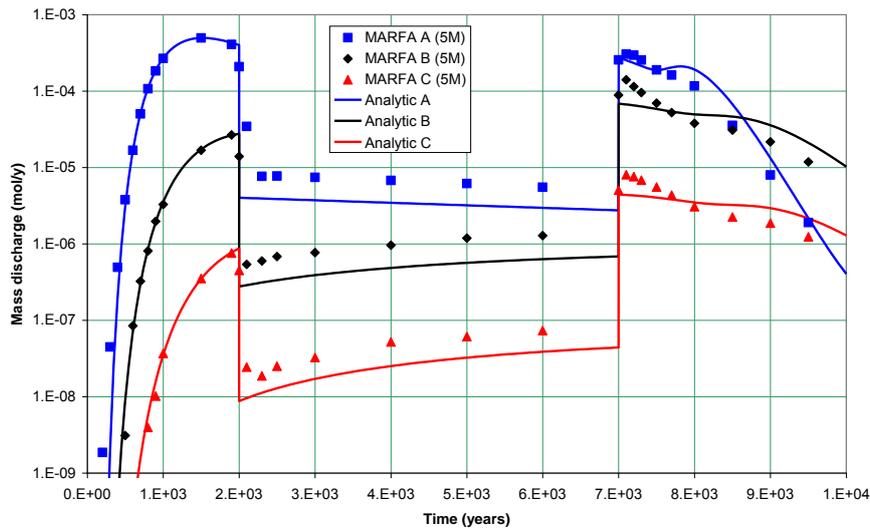


Figure 37: Comparison of MARFA calculations with the analytic calculation for a less dominant matrix and fracture surface sorption

4.6.5 Realistic decay chains

Here we use a real decay chain with data that broadly matches that use in the SR-Can study. The input file for the analytic test is shown below. This gives all the input parameters in a readable format.

```

<CHAIN_MEMBERS>
  U238 WITH DECAY RATE 1.55e-10
  U234 WITH DECAY RATE 2.83e-6
  Th230 WITH DECAY RATE 9.00e-6
  Ra226 WITH DECAY RATE 0.000433
  Pb210 WITH DECAY RATE 0.0311
</CHAIN_MEMBERS>
<SOURCE>
  TYPE = EXPONENTIAL
  DECAY RATE = 0.001
  INITIAL SOURCE:  U238=0.001  U234=0  Th230=0  Ra226=0
  Pb210=0
</SOURCE>
<SEGMENT>
  LENGTH = 500.0
  APERTURE = 0.002
  VELOCITY = 12.5
  FRACTURE RETARDATION:  U238=1  U234=1  Th230=1  Ra226=1
  Pb210=1
  MATRIX PENETRATION DEPTH = 0.03
  MATRIX POROSITY = 0.001
  MATRIX RETARDATION:  U238=2e6  U234=2e6  Th230=3e6
  Ra226=6e6  Pb210=6e6
  MATRIX DIFFUSION:  U238=6e-7  U234=6e-7  Th230=6e-7
  Ra226=6e-7  Pb210=6e-7
</SEGMENT>
PECLET = 25.0

```

```

BOUNDARY AT INFINITY
<TIMES>
  201 GEOMETRIC FROM 1e3 TO 1e7
</TIMES>

```

The MARFA files were set up to match this – using 5 million particles initially, but this gave very poor results with zero Pb210 flux so a 100 million particle case was run. The results are presented in terms of a biosphere dose, using typical dose conversion factors to emphasise which nuclides are significant. Figure 38 shows the comparison.

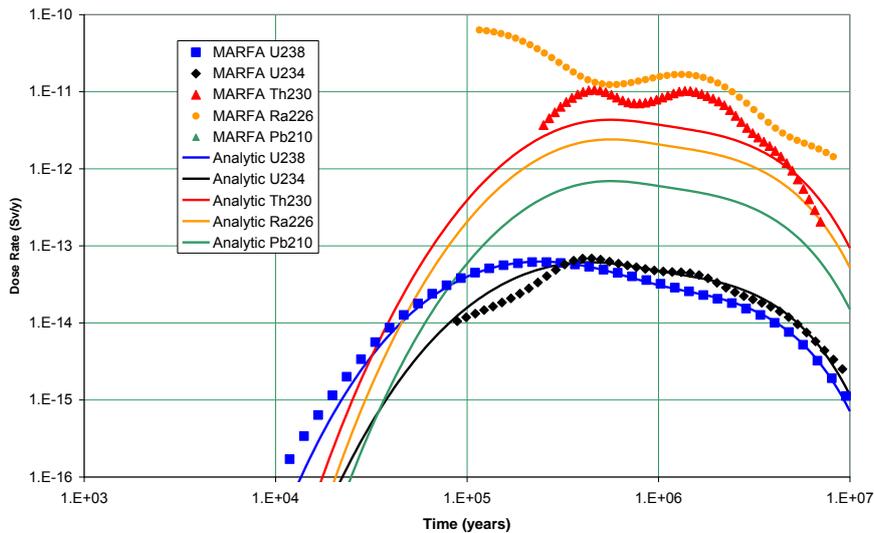


Figure 38: Comparison of MARFA calculations with the analytic calculation for a realistic decay chain

The calculations for the first nuclide (U238) match well, and the MARFA calculations for the second nuclide (U234) are broadly correct, although there are some oscillations that are not present in the analytic calculation. However, the shorter-lived daughters that dominate the dose are very poorly modelled. This is perhaps unsurprising, since the numbers of particles of the daughter that arrive are small. Note that the limited ranges indicate the times when some MARFA particles arrived.

It is therefore concluded that MARFA is not capable of handling the short-lived daughters that grow in for the U-238 chain.

4.6.6 Short-term changes

The examples suggest that short duration changes in properties lead to calculations that are hard for the MARFA post-processing to handle. To investigate this further, we modify Test 6 so that the increased velocity occurs over a much shorter period of 100 years (from 2000 to 2100 years) and in more dramatic way – an increase of a factor of 100 instead of 10. We

use the default setting of the post-processing gamma parameter (0.2). Note that the MARFA calculations split the trajectories into 10 equal segments as was done in Test 6.

Figure 39 shows the comparison, focussing on the time around the change of properties. The MARFA calculations during the high flow regime match the analytic calculation well except for the very sharp initial spike. After the properties revert to their original values, the agreement is much poorer.

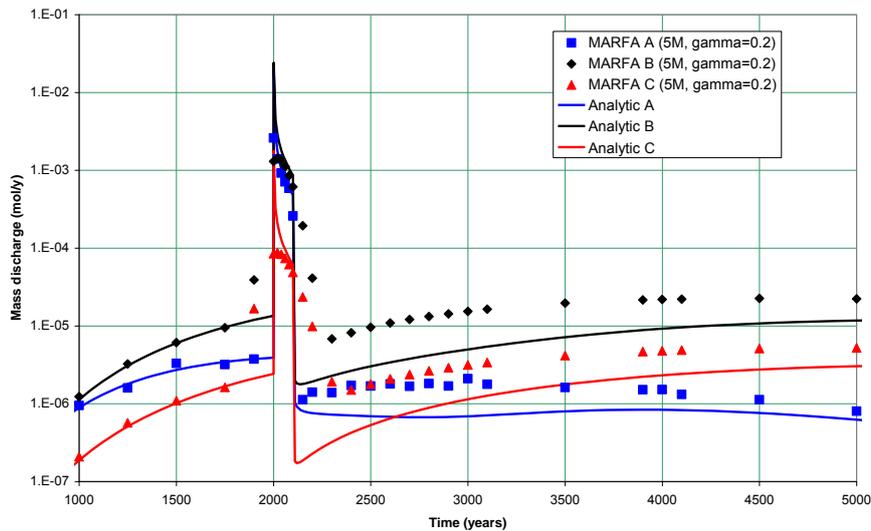


Figure 39: Comparison of MARFA calculations with the analytic calculation for a case with a short-duration velocity increase

Figure 40 shows the full time range, and it is clear that the agreement is restored at longer times. The discrepancy is quite marked, so a further check was undertaken using QPAC [8]. This clearly agrees with the analytic calculation, confirming that the problem is with MARFA.

We believe that the behaviour is related to the treatment of the position along the path at a switch time – this has the potential to give the wrong distribution of mass that leads to an enhanced release after the spike.

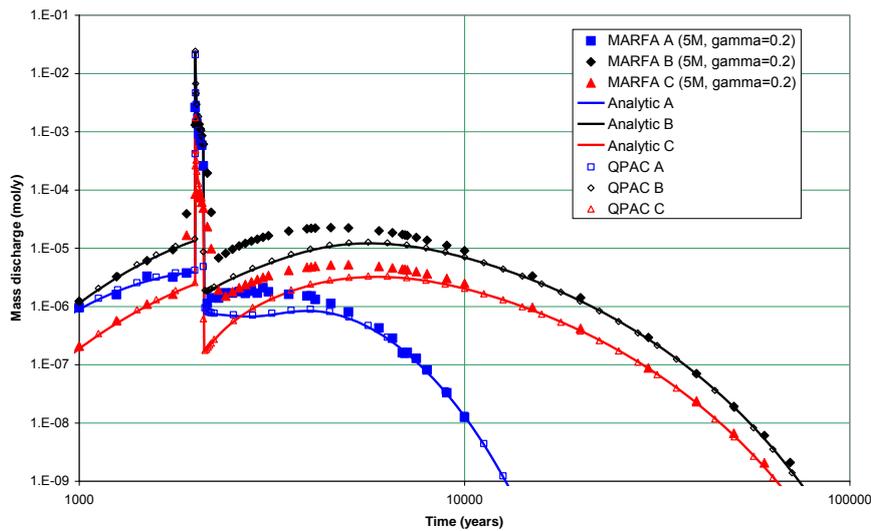


Figure 40: Comparison of MARFA calculations with the analytic calculation for a case with a short-duration velocity increase for the full time range

4.7 MARFA 3.3

The MARFA User Guide contains an appendix which describes MARFA version 3.3.

This version of MARFA does not use fixed transport paths. In order to follow an evolving flow field, each particle uses flow rates relevant to the time that it starts a new segment.

This algorithm appears to be sensible, and is making approximations at a similar level to the fixed flow direction version. The documentation notes that some additional issues arise in using this approach: stuck particles, stranded particles and node renumbering. These problems appear to arise in the way that CONNECTFLOW creates the files used by MARFA rather than in MARFA itself.

Loops in the flows reported by CONNECTFLOW can arise because of inaccuracy in the finite-element solution to the flow field. Finite-element flow solutions are not locally mass-conserving (they are mass conserving in a global sense) and this can lead to local sources and sinks. The same effect can lead to nodes where all the flows are inwards, meaning that the particle becomes stuck.

The approach used in MARFA is to monitor such occurrences and to exclude the particles from the final breakthrough calculations. This introduces a bias and so it is important that this is a rare occurrence.

Node renumbering is handled by mapping node numbers to a master list. This issue presumably arises if different grids are used for different time periods (changes in boundary locations may lead to this perhaps).

The test cases used for MARFA 3.3 are simply cross checks with MARFA 3.2.2. An example application to a 100 m cube with a single flow direction change is presented. It is reported that 40% of particles encounter at least one loop but only 4% are terminated (as stuck or stranded) and only 0.1% of the residence time is spent in loops. Presumably, the stochastic selection of the next segment to follow is generally sufficient to get a particle out of a loop. It is not possible to judge whether the presented results are correct from this example. Some smaller node network examples that could be checked would be a useful addition.

4.8 Conclusions

The conclusions from the review of the MARFA code and documentation are as follows.

1. The basic algorithms employed are sound, but some aspects appear to have been implemented using overly coarse look-up tables which lead to poor approximations for simple test cases.
2. Decay chain handling has not been demonstrated for realistic chains. The code is not capable in its current form of handling short-lived daughter nuclides correctly.
3. The descriptions of the verification tests in the User Guide are full of errors in reporting what was actually calculated. It is necessary to refer to the input files to find out what was actually done.
4. The scope of the test cases is rather limited.
5. There are some problems with handling changes of flow rate, leading to inaccurate releases after a switch.
6. The new Quintessa semi-analytic solutions developed here enable simple cases (with just a few segments) to be analysed accurately.
7. The input file syntax is prone to errors and hard to check because of the use of derived parameters as input and the way aspects are cross-linked by number rather than name.

The main strengths of MARFA are:

- ▲ it can handle large networks in practicable run times;
- ▲ it works well for single nuclides or short chains without extreme decay rates;

- ▲ it accurately handles advective systems where matrix diffusion effects are dominant; and
- ▲ dispersion is handled well (for high Peclet numbers).

The weaknesses identified here are:

- ▲ it is unable to handle long decay chains with short-lived nuclides;
- ▲ memory usage prevents very large numbers of samples being used;
- ▲ diffusive systems cannot be handled;
- ▲ systems where transport in the rock matrix is not dominant are not handled well;
- ▲ calculations immediately after flow rate changes can be inaccurate;
- ▲ handling short spikes is difficult; and
- ▲ the technical aspects of the User Guide are not self-contained, with key results only presented in referenced publications.

5 Overall Conclusions

The studies that have been undertaken provide valuable material in the area of consequence analysis calculations ahead of the SR-Site review. Detailed conclusions drawn for each of the three areas studied are given in the previous sections, but a summary of the most important issues is given here.

The key conclusions for the study of the Q_{eq} concept are that:

- ▲ the basic approach to calculating Q_{eq} values is sound and can be reproduced in QPAC;
- ▲ the fracture resistance dominates over the diffusive resistance in the buffer except for the highest velocity cases;
- ▲ heterogeneity in the fracture, in terms of uncorrelated random variations in the fracture aperture, tends to reduce releases, so the use of a constant average aperture approach is conservative;
- ▲ narrow channels could lead to the same release as larger fractures with the same pore velocity, so a channel enhancement factor of $\sqrt{10}$ should be considered;
- ▲ a spalling zone that increases the area of contact between flowing water and the buffer has the potential to increase the release significantly and changes the functional dependence of Q_{eq}^{frac} on the flowing velocity.

Quintessa's AMBER software has previously been used to reproduce SKB's one-dimensional transport calculations and AMBER allows the use of time-varying properties. This capability has been used to investigate the effects of glacial episodes on radionuclide transport. The main parameters that could be affected are sorption coefficients and flow rates. For both single and multiple glacial episodes the time-dependency of model parameters did not result in much change to the calculated peak fluxes to the biosphere. This is because fluxes are calculated to be dominated by the poorly-sorbed radionuclides ^{129}I and ^{36}Cl . However, a small increase (less than an order of magnitude) in the overall flux contributed from the radium decay chain, which is important for long timescales, was calculated during the phase when the multiple glacial episodes are occurring. These conclusions are preliminary and could be changed if different radionuclides are important in the SR-Site assessment.

The key conclusions for the study on the MARFA code are that:

- ▲ MARFA can handle large networks in practicable run times, generally works well for single radionuclides or short chains and accurately handles advective systems where matrix diffusion effects are dominant.

- ▲ The code has, however, a number of important limitations. In particular, it is unable to handle long decay chains with short-lived radionuclides and calculations immediately after flow rate changes can be inaccurate.
- ▲ The documentation and Quality Assurance are poor. A large number of errors have been found in the User Guide and the associated test cases do not adequately test the use of the code for the anticipated applications.
- ▲ These limitations bring into question the code's suitability (in its present form) for performance assessment calculations for a deep radioactive waste repository.

References

- [1] Benbow, S. J., Robinson, P. C. and Watson, S. P. (2010) Analysis of Barrier Performance: Quintessa's Modelling of Copper Corrosion Scenarios with an Eroding Buffer ahead of SR-Site. Quintessa Report to SSM QRS-1503A-1.
- [2] Maul, P., Robinson, P., Bond, A. and Benbow, S. (2008) Independent Calculations for the SR-Can Assessment: External review contribution in support of SKI's and SSI's review of SR-Can, SKI Report 2008:12.
- [3] Robinson, P., Maul, P. and Watson, C. (2010) Radionuclide Transport: Preparation During 2009 for the SR-Site Review. SSM report 2010:18.
- [4] SKB (2006) Long-term safety for KBS-3 repositories at Forsmark and Laxemar - a first evaluation. Main Report of the SR-Can project. SKB report TR-06-09.
- [5] Neretnieks, I. (1986) Stationary Transport of Dissolved Species in the Backfill Surrounding a Waste Canister in Fissured Rock: Some Simple Analytical Solutions. Nuclear Technology 72, 196.
- [6] Maul, P. and Robinson, P. (2005) Modelling Radionuclide Transport at Interfaces in the SKB Deep Repository Design, Quintessa Report to SKI, QRS-1097G-2 Version 1.0.
- [7] Neretnieks, I. (2006) Flow and Solute Transport in a Damaged Zone due to Spalling. SKB Report R-06-91.
- [8] Quintessa (2010) QPAC: Quintessa's General Purpose Modelling Software. Quintessa Report QRS-QPAC-11.
- [9] SKB (2006) Data report for the safety assessment SR-Can. SKB report TR-06-25.
- [10] Liu, J. and Neretnieks, I. (2004) Coupled Transport/Reaction Modelling of Copper canister Corrosion aided by Microbial Processes. Radiochim. Acta **92**, 849-854.
- [11] Hakami E (1995) Aperture Distribution of Rock Fractures. Doctoral Thesis, Department of Civil and Environmental Engineering, Royal Institute of Technology, Stockholm.
- [12] Norman, S. and Kjellbert, N. (1990) FARF31 - A far field radionuclide migration code for use with the PROPER package.

SKB report TR-90-01.

- [13] Elert, M., Gylling, B. and Lindgren, M. (2004) Assessment model validity document FARF31. SKB report R-04-51.
- [14] Quintessa (2009) AMBER 5.3 Reference Guide. Quintessa report QE-AMBER-1, v5.3.
- [15] Jaquet, O. and Siegel, P. (2006) Regional groundwater flow model for a glaciation scenario: Simpevarp subarea - version 1.2. SKB report R-06-100.
- [16] Crawford, J., Neretnieks, I. and Malmström, M. (2006) Data and uncertainty assessment for radionuclide Kd partitioning coefficients in granitic rock for use in SR-Can calculations. SKB report R-06-75.
- [17] Pastina, B. and Hellä, P. (2010) Models and Data Report 2010. Posiva report 2010-01.
- [18] Hartley, L., Hoch, A., Jackson, P., Joyce, S., McCarthy, R., Rodwell, W., Swift, B. and Marsic, N. (2006) Groundwater flow and transport modelling during the temperate period for the SR-Can assessment: Forsmark area - version 1.2. SKB report R-06-98.
- [19] Painter, S. and Mancillas, J. (2009) MARFA version 3.2.2 user's manual: migration analysis of radionuclides in the far field. SKB Report R-09-56.
- [20] Painter S., Cvetkovic V and Pensado O (2006). Time-Domain Random Walk Method for Simulating Radionuclide Transport in Fractured Porous Rock. Proceedings of the 2006 International High-Level Radioactive Waste Management Conference, Las Vegas, Nevada, April 30–May 4, La Grange Park, Illinois: American Nuclear Society, 2006. 28.
- [21] Painter S (2008). Accommodating Transient Velocities in Time-Domain Particle Tracking Algorithms of Radionuclide Transport. Proceedings of the 12th International High-Level Radioactive Waste Management Conference (IHLRWM); 7–11 September 2008; Las Vegas, Nevada. LaGrange Park, IL: American Nuclear Society (CD ROM). 2008.
- [22] Painter S and Cvetkovic V (2005) Upscaling Discrete Fracture Network Simulations: An Alternative to Continuum Transport Models. Water Resources Research 41, doi:10.1029/2004WR003682, 2005.
- [23] Talbot A (1979) The Accurate Numerical Inversion of Laplace Transforms. J. Inst. Math. Appl. 23.

- [24] Painter, S., Mancillas, J., Cvetkovic, V. and Pensado, O. (2008) Time-Domain Particle Tracking Method for Simulating Transport with Retention and First-Order Transformation. *Water Resources Research*, 44, doi:10.1029/2007WR005944, 2008.
- [25] Sudicky, E. A. and Frind, E. O. (1984) Contaminant Transport in Fractured Porous Media: Analytical Solution for a Two-Member Decay Chain in a Single Fracture. *Water Resources Research* 20(7), 1021–1029, 1984.

Appendix A: Nomenclature

Note that the notation used in MARFA differs from that used in Appendix B. Here, both notations are recorded. In some cases a symbol is used for different purposes – the context should make it clear which version is being used.

Subscripts and superscripts are used to distinguish values for various nuclides, for fracture and matrix and for different stages in time-dependent calculations. These are not recorded here.

Symbol	Description
a	limited matrix diffusion depth [m]
a	half the height of the canister zone where the concentration is fixed [m]
b	half the fracture aperture [m]
C	concentration in the water [moles/m ³]
D_{eff}, D_m	effective diffusion coefficient in the rock matrix [m ² /year]
d_L	dispersion length [m]
F	F factor = $L\delta_f / v$ [y m ⁻¹]
G	A coefficient in the matrix depth probability density function, stated in the MARFA documentation to be $\sqrt{R/D_{eff}}$, but believed to be $\sqrt{\theta R/D_{eff}}$.
G_{in}	specified flux at $x=0$ [mol/y]
G_{out}	output flux [mol/y]
h	fracture aperture [m]
k_a	area-based surface sorption coefficient [m ³ /m ²]
K_d	matrix sorption coefficient of nuclide n [m ³ /kg]
L	path length [m]
M	concentration in the matrix pore water [moles/m ³]
M	total mass in part of the system [mol]
Pe	Peclet number (for dispersion) [-]
Q_{eq}	equivalent flow rate [m ³ /y]
r	quantile in a probability distribution
R	retardation (due to surface sorption if any) of nuclide n [-]

Symbol	Description
R	The distance from the canister to the fracture [m]
s	Laplace variable [y^{-1}]
t, T	time [y]
t^*	dimensionless time [-]
t_0	A dimensionless time, stated in the MARFA guide incorrectly to be $\frac{R\Delta^2}{D_{eff}}$ and, correctly to be $(\kappa\eta)^2$ which is $\frac{\theta R\Delta^2}{D_{eff}}$.
U, V	coefficients in the semi-analytic solution
w	flow path width [m]
x	distance along the flow path [m]
v	water velocity [m/year]
z	distance into the rock matrix [m]
α	coefficient in the semi-analytic solution [m^{-1}]
α	dispersion length [m]
β	transport parameter (same as F) [$y m^{-1}$]
κ	derived matrix parameter, $= \sqrt{\theta_m R_m D_{eff}}$ [$m/y^{1/2}$]
δ_f	surface area of matrix per unit volume of flowing water [m^2/m^3]
Δ	matrix penetration depth [m]
η	derived matrix parameter, $= \Delta / D_{eff}$ [$y m^{-1}$]
η^*	dimensionless parameter, $= \eta / \beta$ [$y m^{-1}$]
ε_m, θ_m	matrix porosity [-]
λ	decay constant [per year]
ϕ	factor in solution of matrix transport equation, $= \sqrt{\frac{\varepsilon_m R_m (\lambda + s)}{D_m}}$ [m^{-1}]
Φ	release rate in Q_{eq} calculations [$mol y^{-1}$]
τ	water travel time [y]

An overbar denotes the Laplace transform throughout the report. Multiple overbars and a tilde are used to denote multiple Laplace transforms in Appendix B.

Appendix B: Semi-analytic Solutions for Radionuclide Transport with Piecewise-constant Properties

The general solution for radionuclide transport in a fracture with adjacent rock matrix when all the properties are constant in time is set out in [3]. This was presented in terms of the Laplace-transformed solution.

Here we aim to generalise the solution to handle piecewise-constant properties. We use the same notation as that employed in [3]. In order to make progress, we start with the simplified case of a single nuclide in an advection-only segment with an infinite rock matrix.

Our objective is to allow the non-geometric properties to change values at two switch times and to derive a generalised transform of the solution in analytic form which can be inverted to retrieve the time-domain solution.

Given an overall time, t , we shall denote the time spent in each of the stages as t_0 , t_1 and t_2 , with the properties given the same subscripts. The maximum possible times in the first two stages are T_0 and T_1 and the switch times are therefore T_0 and T_0+T_1 . All three stages are only relevant for times greater than the second switch time. Up to the first switch time the properties have been constant so the conventional Laplace transform solution applies. After the first switch, we will use a double-transform approach – applying a Laplace transform with respect to both t_0 , and t_1 . After the second switch we shall use a triple transform with respect to t_0 , t_1 and t_2 . The time-domain solution will be recovered by inverting the relevant multi-transform by using a product set of Talbot inversion contours.

This approach is believed to be novel and so it is worth setting down the basic approach in some detail.

The solutions described here have been coded as an extension to the code used in [3]. Self-consistency checks have been made to verify the correct implementation of the algorithms. These include mass conservation checks and checks with switch times when no property values actually change. The comparisons with MARFA in the main text also give confidence that the code has been correctly implemented.

B.1 Multiple-transforms and Inversion

In the following we denote the Laplace variables as s_0 , s_1 and s_2 rather than the simple s that is conventionally used for the single transform.

B.1.1 Single (standard) transform

The Laplace transform for a function $f(t_0)$ is defined as

$$\bar{f}(s_0) = \int_0^{\infty} e^{-s_0 t_0} f(t_0) dt_0 \quad (\text{B. 1.1})$$

where s_0 is (in general) complex.

The Talbot inversion algorithm [23] enables the time-domain function to be recovered numerically but accurately. We usually write:

$$f(t_0) \approx \frac{1}{n} \sum'_{k=0}^{n-1} \text{Re} \left[w_k e^{s_k t_0} \bar{f}(s_k) \right], \quad (\text{B. 1.2})$$

where Σ' implies that the $k=0$ term is given a weight of one half, and

$$\begin{aligned} s_k &= \lambda(\alpha_k + i\nu\theta_k) + \sigma \\ \theta_k &= \frac{k\pi}{n} \\ \alpha_k &= \theta_k \cot \theta_k \\ w_k &= \lambda(\nu + i\beta_k) \\ \beta_k &= \theta_k + \frac{\alpha_k(\alpha_k - 1)}{\theta_k}. \end{aligned} \quad (\text{B. 1.3})$$

The standard choice of parameters suggested by Talbot are $\sigma = 0$, $\nu = 1$, $\lambda = 6/t_0$, with n chosen to give sufficient accuracy, but typically $n = 32$.

We note here that the use of the real part of the term in the sum is actually a short cut allowing only half the contour to be used. This is possible because, for a real function, the terms will occur in conjugate pairs. The full formula would be:

$$f(t_0) \approx \frac{1}{2n} \sum_{k=-(n-1)}^{n-1} w_k e^{s_k t_0} \bar{f}(s_k). \quad (\text{B. 1.4})$$

When nesting the inversion formulae it is necessary to use this full form for the inner transforms as they are not real-valued functions.

B.1.2 Double transform

We define the double Laplace transform for a function $f(t_0, t_1)$ as

$$\bar{\bar{f}}(s_0, s_1) = \int_0^\infty \int_0^\infty e^{-s_0 t_0 - s_1 t_1} f(t_0, t_1) dt_0 dt_1 \quad (\text{B. 1.5})$$

where both s_0 and s_1 are (in general) complex.

The double inversion formula is then

$$f(t_0, t_1) \approx \frac{1}{4n_0 n_1} \sum_{k_0=-(n_0-1)}^{n_0-1} \sum_{k_1=-(n_1-1)}^{n_1-1} w_{0k_0} w_{1k_1} e^{s_{0k_0} t_0 + s_{1k_1} t_1} \bar{\bar{f}}(s_{0k_0}, s_{1k_1}), \quad (\text{B. 1.6})$$

where the two contours can in general have different numbers of points and different parameter values.

B.1.3 Triple transform

We now define the triple Laplace transform for a function $f(t_0, t_1, t_2)$ as

$$\bar{\bar{\bar{f}}}(s_0, s_1, s_2) = \int_0^\infty \int_0^\infty \int_0^\infty e^{-s_0 t_0 - s_1 t_1 - s_2 t_2} f(t_0, t_1, t_2) dt_0 dt_1 dt_2, \quad (\text{B. 1.7})$$

and the triple inversion formula is

$$f(t_0, t_1, t_2) \approx \frac{1}{8n_0 n_1 n_2} \sum_{k_0=-(n_0-1)}^{n_0-1} \sum_{k_1=-(n_1-1)}^{n_1-1} \sum_{k_2=-(n_2-1)}^{n_2-1} w_{0k_0} w_{1k_1} w_{2k_2} e^{s_{0k_0} t_0 + s_{1k_1} t_1 + s_{2k_2} t_2} \bar{\bar{\bar{f}}}(s_{0k_0}, s_{1k_1}, s_{2k_2}). \quad (\text{B. 1.8})$$

B.1.4 Example – exponentially decaying function

As an example, which is used in the test cases, consider an exponentially decaying function with an initial value of one and decay rates μ_0 , μ_1 and μ_2 in the three stages. We denote the solution as three different functions, one for each stage, but these are used to construct the full solution:

$$f(t) = \begin{cases} f_0(t) & t \leq T_0 \\ f_1(T_0, t - T_0) & T_0 < t \leq T_0 + T_1 \\ f_2(T_0, T_1, t - T_0 - T_1) & T_0 + T_1 < t \end{cases} \quad (\text{B. 1.9})$$

For times up to T_0 only a single transform is needed:

$$f_0(t_0) = e^{-\mu_0 t_0} \text{ with } \bar{f}_0(s_0) = \frac{1}{\mu_0 + s_0}. \quad (\text{B. 1.10})$$

For times after that and up to T_0+T_1 a double transform is needed:

$$f_1(T_0, t_1) = e^{-\mu_0 T_0 - \mu_1 t_1} \text{ with } \bar{\bar{f}}_1(s_0, s_1) = \frac{1}{(\mu_0 + s_0)(\mu_1 + s_1)}. \quad (\text{B. 1.11})$$

For later times a triple transform is needed:

$$f_2(T_0, T_1, t_2) = e^{-\mu_0 T_0 - \mu_1 T_1 - \mu_2 t_2}$$

with $\bar{\bar{\bar{f}}}_2(s_0, s_1, s_2) = \frac{1}{(\mu_0 + s_0)(\mu_1 + s_1)(\mu_2 + s_2)}. \quad (\text{B. 1.12})$

Of course, in this case the time-domain solution is simple and known, but the multiple-Laplace transforms could also have been derived directly from the differential equations.

In the first stage

$$\frac{df_0}{dt_0} = -\mu_0 f_0 \quad (\text{B. 1.13})$$

Transforming this and using the initial condition gives

$$\bar{f}_0 = \frac{1}{\mu_0 + s_0}. \quad (\text{B. 1.14})$$

How, for the second stage

$$\frac{df_1}{dt_1} = -\mu_1 f_1 \quad (\text{B. 1.15})$$

with initial condition (at $t_1=0$) being $f_1(T_0, t_1=0) = f_0(T_0)$. Then the (single) transform with respect to t_1 is

$$\bar{f}_1 = \frac{f_0(T_0)}{\mu_1 + s_1}. \quad (\text{B. 1.16})$$

Now we apply a second transform with respect to T_0 and obtain

$$\bar{f}_1 = \frac{\bar{f}_0}{\mu_1 + s_1} = \frac{1}{(\mu_0 + s_0)(\mu_1 + s_1)}. \quad (\text{B.1.17})$$

The result for the final stage follows in the same way.

Note that for this simple case it would have been possible to replace the initial condition for stage 2 in (B.1.16) by the known solution because the transform can be inverted directly. For the general case this is not possible; indeed, the case of interest will involve a spatially varying function which we do not want to evaluate at discrete positions since this would introduce discretisation errors into the result.

B.2 Advection and Infinite Diffusion for a Single Nuclide

The equations for the fracture and matrix concentrations are as follows. These are the same equations as used in the main text for MARFA, but with different notation and with the input specified as a flux rather than a concentration.

$$R_f \frac{\partial C_f}{\partial t} = -v \frac{\partial C_f}{\partial x} - \lambda R_f C_f + \delta_f D_m \left. \frac{dC_m}{dz} \right|_{z=0} \quad (\text{B.2.1})$$

$$\varepsilon_m R_m \frac{\partial C_m}{\partial t} = D_m \frac{\partial^2 C_m}{\partial z^2} - \lambda \varepsilon_m R_m C_m \quad (\text{B.2.2})$$

$$hw \left[v C_f \right]_{x=0} = G_{in}(t) \quad (\text{B.2.3})$$

$$\left[C_m \right]_{z=0} = C_f \quad (\text{B.2.4})$$

$$C_m \xrightarrow{z \rightarrow \infty} 0 \quad (\text{B.2.5})$$

The result of interest is the flux at the far end:

$$G_{out}(t) = hw \left[v C_f \right]_{x=L}. \quad (\text{B.2.6})$$

We will also calculate the total amount in the fracture and in the matrix. These are useful checks on the solutions.

$$M_f(t) = hw R_f \int_0^L C_f dx. \quad (\text{B.2.7})$$

$$M_m(t) = hw\delta_f\varepsilon_m R_m \int_0^L \int_0^\infty C_m dz dx. \quad (\text{B.2.8})$$

The properties that can vary with time are R_f , ν , R_m and D_m ; their values are denoted with additional subscripts 0, 1 or 2, e.g. $R_{f,0}$, $R_{f,1}$ and $R_{f,2}$.

B.2.1 Stage 0

Before the first switch time, the single transform solutions are straightforward.

$$\bar{C}_{f,0} = \frac{1}{hw\nu_0} \bar{G}_{in,0} e^{-\alpha_0 x_0} \quad (\text{B.2.9})$$

$$\bar{C}_{m,0} = \bar{C}_{f,0} e^{-\phi_0 z_0} = \frac{1}{hw\nu_0} \bar{G}_{in,0} e^{-\alpha_0 x_0} e^{-\phi_0 z_0} \quad (\text{B.2.10})$$

$$\bar{G}_{out,0} = \bar{G}_{in,0} e^{-\alpha_0 L} \quad (\text{B.2.11})$$

$$\bar{M}_{f,0} = \frac{R_{f,0}}{\nu_0} \bar{G}_{in,0} \frac{1 - e^{-\alpha_0 L}}{\alpha_0} \quad (\text{B.2.12})$$

$$\bar{M}_{m,0} = \frac{\delta_f \varepsilon_m R_{m,0}}{\nu_0} \bar{G}_{in,0} \frac{1 - e^{-\alpha_0 L}}{\alpha_0} \frac{1}{\phi_0} \quad (\text{B.2.13})$$

where

$$\phi_0 = \sqrt{\frac{\varepsilon_m R_{m,0} (\lambda + s_0)}{D_{m,0}}}, \quad (\text{B.2.14})$$

$$\alpha_0 = \frac{R_{f,0} (\lambda + s_0) + \delta_f D_{m,0} \phi_0}{\nu_0}. \quad (\text{B.2.15})$$

Note that we have written the concentration solutions here in terms of coordinates denoted x_0 and z_0 because these will become the coordinates for the initial condition in the second stage.

B.2.2 Later Stages

After each switch the solutions for the concentrations and flux out each have three components: from the ongoing injection; from material that was in the fracture at the switch time; and from material that was in the matrix at the switch time.

The solutions for the later stages are best calculated through a set of recurrence formulae for the coefficients in general expressions for the fracture and matrix concentrations. We will see that it is possible to write these concentrations as sums over the basic solution forms for each stage.

So, we write

$$\phi_N = \sqrt{\frac{\varepsilon_m R_{m,N} (\lambda + s_N)}{D_{m,N}}}, \quad (\text{B.2.16})$$

$$\alpha_N = \frac{R_{f,N} (\lambda + s_N) + \delta_f D_{m,N} \phi_N}{v_N}. \quad (\text{B.2.17})$$

and then write

$$\tilde{C}_{f,N} = \frac{1}{hw} \sum_{i=0}^N U_N^i e^{-\alpha_i x_N} \quad (\text{B.2.18})$$

$$\tilde{C}_{m,N} = \frac{1}{hw} \sum_{i=0}^N \sum_{j=0}^N V_N^{ij} e^{-\alpha_i x_N} e^{-\phi_j z_N} \quad (\text{B.2.19})$$

and derive the U_N and V_N from the previous stage U_{N-1} and V_{N-1} . These coefficients are independent of the coordinates, but will depend on all the parameters and Laplace variables. The area factor has been written outside the sum as it appears in all the terms but cancels out in the calculated outputs.

The tilde overbars here are used to denote that the number of transforms depends on the stage – it is always one more than the N subscript.

Given these solution forms, the flux out and masses can be calculated as follows:

$$\tilde{G}_{out,N} = v_N \sum_{i=0}^N U_N^i e^{-\alpha_i L} \quad (\text{B.2.20})$$

$$\tilde{M}_{f,N} = R_{f,N} \sum_{i=0}^N U_N^i \left(\frac{1 - e^{-\alpha_i L}}{\alpha_i} \right) \quad (\text{B.2.21})$$

$$\tilde{M}_{m,N} = \varepsilon_m R_{m,N} \delta_f \sum_{i=0}^N \sum_{j=0}^N V_N^{ij} \left(\frac{1 - e^{-\alpha_i L}}{\alpha_i} \right) \frac{1}{\phi_j} \quad (\text{B.2.22})$$

The solution obtained for Stage 0 can be written in this framework as

$$U_0^0 = V_0^{00} = \frac{\tilde{G}_{in,0}}{v_0}. \quad (\text{B.2.23})$$

We now develop the recurrence relations for the later stages.

The ongoing injection is the simplest to handle. The solution here has the same form as the solution in stage 0, with the injected multiply-transformed (along the lines described in B.1.4). Thus, the first contributions to the solution in stage N are as simply:

$$\left. \begin{array}{l} U_N^N + \\ V_N^{NN} + \end{array} \right\} \leftarrow \frac{\tilde{G}_{in,N}}{v_N}. \quad (\text{B.2.24})$$

The notation here indicates that the term on the right is added to each of the coefficients on the left.

Next we deal with the initial material in the fracture.

The concentration in the fracture at the start of a stage is directly related to the concentration at the end of the previous stage with a ratio of retardations being needed because it will be total mass that is conserved across the switch rather than concentration. Thus

$$C_{f,N}(x, T_{N-1}) = \frac{R_{f,N-1}}{R_{f,N}} C_{f,N-1}(x, T_{N-1}). \quad (\text{B.2.25})$$

The transformed equations can then be solved to give

$$\tilde{C}_{f,N,F} = \frac{e^{-\alpha_N x_N}}{v_N} \int_0^{x_N} e^{\alpha_N x} R_{f,N-1} \tilde{C}_{f,N-1}(x) dx \quad (\text{B.2.26})$$

$$\tilde{C}_{m,N,F} = \tilde{C}_{f,N,F} e^{-\phi z_1}. \quad (\text{B.2.27})$$

Here an additional subscript F is used to indicate that this is the concentration arising from material in the fracture at the start of the stage.

Now, we can substitute the general form for the concentration into the integral to obtain

$$\begin{aligned}\tilde{C}_{f,N,F} &= \frac{1}{hw} \sum_{i=0}^{N-1} U_{N-1}^i \frac{R_{f,N-1}}{v_N} e^{-\alpha_N x_N} \int_0^{x_N} e^{\alpha_N x} e^{-\alpha_i x} dx \\ &= \frac{1}{hw} \sum_{i=0}^{N-1} U_{N-1}^i \frac{R_{f,N-1}}{v_N} \frac{e^{-\alpha_i x_N} - e^{-\alpha_N x_N}}{\alpha_N - \alpha_i}\end{aligned}\quad (\text{B.2.28})$$

This demonstrates that the sum of exponentials form is preserved and provides contributions to the coefficients at the next stage:

$$\left. \begin{array}{l} U_N^i + \\ U_N^N - \\ V_N^{iN} + \\ V_N^{NN} - \end{array} \right\} \leftarrow \frac{R_{f,N-1}}{v_N (\alpha_N - \alpha_i)} U_{N-1}^i. \quad (\text{B.2.29})$$

Here the contribution is relevant to four coefficients, adding to two of them and subtracting from the other two.

Finally, we deal with material in the matrix at the switch time, using a subscript M to denote this.

The concentration in the matrix at the start of a stage is directly related to the concentration at the end of the previous stage with a ratio of retardations being needed because it will be total mass that is conserved across the switch rather than concentration. Thus

$$C_{m,N}(x, z, T_{N-1}) = \frac{R_{m,N-1}}{R_{m,N}} C_{m,N-1}(x, z, T_{N-1}). \quad (\text{B.2.30})$$

We directly derive the result for a single component in the sum: $V_{ij}^{(N-1)} e^{-\alpha_i x} e^{-\phi_j z}$.

With this as an initial condition, the matrix concentration is

$$\tilde{C}_{m,N,M} = \tilde{C}_{f,N,M} e^{-\phi_N z_N} + V_{N-1}^{ij} e^{-\alpha_i x_N} \frac{\varepsilon_m R_{m,N-1}}{D_{m,N}} \frac{e^{-\phi_j z_N} - e^{-\phi_N z_N}}{\phi_N^2 - \phi_j^2}, \quad (\text{B.2.31})$$

which can then be used to calculate the solution for the fracture:

$$\tilde{C}_{f,N,M} = V_{N-1}^{ij} \frac{\delta_f \varepsilon_m R_{m,N-1}}{v_N} \frac{e^{-\alpha_i x_N} - e^{-\alpha_N x_N}}{\alpha_N - \alpha_i} \frac{\phi_N - \phi_j}{\phi_N^2 - \phi_j^2}. \quad (\text{B.2.32})$$

From this we can see that there are two types of contribution to the coefficients at the next stage. The first arises from material that moves in the matrix but does not enter the fracture between stages:

$$\left. \begin{array}{l} V_N^{ij+} \\ V_N^{iN-} \end{array} \right\} \leftarrow \frac{\varepsilon_m R_{m,N-1}}{D_{m,N}(\phi_N^2 - \phi_j^2)} V_{N-1}^{ij}, \quad (\text{B.2.33})$$

while the second arises from material that does re-enter the fracture during the stage:

$$\left. \begin{array}{l} U_N^i + \\ U_N^N - \\ V_N^{iN+} \\ V_N^{NN-} \end{array} \right\} \leftarrow \frac{\delta_f \varepsilon_m R_{m,N-1} (\phi_N - \phi_j)}{v_N (\alpha_N - \alpha_i) (\phi_N^2 - \phi_j^2)} V_{N-1}^{ij}. \quad (\text{B.2.34})$$

B.3 The Finite Matrix Cases

The solution developed above was for the case of an infinite matrix. The finite matrix case requires only very minor changes. For a penetration depth a , the equation for α must be modified:

$$\alpha_N = \frac{R_{f,N}(\lambda + s_N) + \delta_f D_{m,N} \phi_N \tanh a \phi_N}{v_N}. \quad (\text{B.3.1})$$

The general expansion for the matrix contribution is now written:

$$\tilde{C}_{m,N} = \frac{1}{hw} \sum_{i=0}^N \sum_{j=0}^N V_N^{ij} e^{-\alpha_i x_N} \frac{\cosh \phi_j (a - z_N)}{\cosh \phi_j a}, \quad (\text{B.3.2})$$

and the mass in the matrix is now

$$\tilde{M}_{m,N} = \varepsilon_m R_{m,N} \delta_f \sum_{i=0}^N \sum_{j=0}^N V_N^{ij} \left(\frac{1 - e^{-\alpha_i L}}{\alpha_i} \right) \frac{\tanh a \phi_j}{\phi_j}. \quad (\text{B.3.3})$$

The contribution from mass in the matrix is now

$$\left. \begin{array}{l} U_N^i + \\ U_N^N - \\ V_N^{iN+} \\ V_N^{NN-} \end{array} \right\} \leftarrow \frac{\delta_f \varepsilon_m R_{m,N-1} (\phi_N \tanh a \phi_N - \phi_j \tanh a \phi_j)}{v_N (\alpha_N - \alpha_i) (\phi_N^2 - \phi_j^2)} V_{N-1}^{ij}. \quad (\text{B.3.4})$$

The other results are unchanged. The new forms tend to the simpler infinite case if the penetration depth tends to infinity.

B.4 Dispersion

Dispersion can be added quite simply if the assumption is made that each segment can be treated as semi-infinite, with the flux passing $x=L$ being the required output. This is the assumption made in MARFA and is appropriate for an advectively-dominated system.

Then, the formula for α is adjusted and the dispersive contribution to the flux is added. Denoting the dispersivity as d_L (corresponding to a Peclet number of L/d_L), we write

$$g_N = R_{f,N}(\lambda + s_N) + \delta_f D_{m,N} \phi_N \tanh a \phi_N. \quad (\text{B.4.1})$$

and then

$$\alpha_N = \frac{g_N}{v_N} \frac{2}{1 + \sqrt{1 + \frac{4g_N d_L}{v_N}}}. \quad (\text{B.4.2})$$

Various minor adjustments are needed to the other equations. The flux out is now

$$\tilde{G}_{out,N} = v_N \sum_{i=0}^N U_N^i (1 + \alpha_i d_L) e^{-\alpha_i L}. \quad (\text{B.4.3})$$

The contribution from ongoing injection is

$$\left. \begin{array}{l} U_N^N + \\ V_N^{NN} + \end{array} \right\} \leftarrow \frac{\tilde{G}_{in,N}}{v_N (1 + \alpha_N d_L)}. \quad (\text{B.4.4})$$

The contributions from the previous stages become

$$\begin{aligned}
\left. \begin{array}{l} U_N^i + \\ V_N^{iN} + \end{array} \right\} &\leftarrow \Psi_{N-1}^i \\
\left. \begin{array}{l} U_N^N - \\ V_N^{NN} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i d_L)}{(1 + \alpha_N d_L)} \Psi_{N-1}^i \\
\Psi_{N-1}^i &= \frac{R_{f,N-1}}{v_N (\alpha_N - \alpha_i) (1 + \alpha_i d_L + \alpha_N d_L)} U_{N-1}^i
\end{aligned} \tag{B.4.5}$$

and

$$\begin{aligned}
\left. \begin{array}{l} U_N^i + \\ V_N^{iN} + \end{array} \right\} &\leftarrow \Psi_{N-1}^{ij} \\
\left. \begin{array}{l} U_N^N - \\ V_N^{NN} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i d_L)}{(1 + \alpha_N d_L)} \Psi_{N-1}^{ij} \\
\Psi_{N-1}^{ij} &= \frac{\delta_f \varepsilon_m R_{m,N-1} (\phi_N \tanh a \phi_N - \phi_j \tanh a \phi_j)}{v_N (\alpha_N - \alpha_i) (1 + \alpha_i d_L + \alpha_N d_L) (\phi_N^2 - \phi_j^2)} V_{N-1}^{ij}
\end{aligned} \tag{B.3.4}$$

B.5 Decay Chains

When decay chains are considered, there are additional terms for daughter nuclides arising from decays that occur during each stage. The general form of the solution for the concentrations will contain terms from all ancestors to the particular nuclide including from previous stages.

Introducing a superscript to denote the nuclide (with zero being the head of the chain), we will write for the K -th nuclide:

$$\tilde{C}_{f,N}^K = \frac{1}{hw} \sum_{k=0}^K \sum_{i=0}^N U_{N,i}^{K,k} e^{-\alpha_i^k x_N} \tag{B.5.1}$$

$$\tilde{C}_{m,N}^K = \frac{1}{hw} \sum_{k=0}^K \sum_{r=k}^K \sum_{i=0}^N \sum_{j=0}^N V_{N,ij}^{K,kr} e^{-\alpha_i^k x_N} e^{-\phi_j^r z_N} \tag{B.5.2}$$

The solution for the first nuclide in a chain is as previously derived. Later nuclides include terms from the decay during the stage. In the equations this arises as a source term, which has similar form to the initial conditions terms already analysed, making it straightforward to calculate the extra relationships.

Here we write the full set of equations, including dispersion and finite matrix diffusion.

$$\tilde{G}_{out,N}^K = v_N \sum_{k=0}^K \sum_{i=0}^N U_{N,i}^{K,k} (1 + \alpha_i^k d_L) e^{-\alpha_i^k L}. \quad (B.5.3)$$

$$\tilde{M}_{f,N} = R_{f,N} \sum_{k=0}^K \sum_{i=0}^N U_{N,i}^{K,k} \left(\frac{1 - e^{-\alpha_i^k L}}{\alpha_i^k} \right) \quad (B.5.4)$$

$$\tilde{M}_{m,N} = \varepsilon_m R_{m,N} \delta_f \sum_{k=0}^K \sum_{r=k}^K \sum_{i=0}^N \sum_{j=0}^N V_{N,ij}^{K,kr} \left(\frac{1 - e^{-\alpha_i^k L}}{\alpha_i^k} \right) \frac{\tanh a \phi_j^r}{\phi_j^r}. \quad (B.5.5)$$

The contribution from ongoing injection is

$$\left. \begin{array}{l} U_{N,N}^{K,K} + \\ V_{N,NN}^{K,KK} + \end{array} \right\} \leftarrow \frac{\tilde{G}_{in,N}^K}{v_N (1 + \alpha_N^K d_L)}. \quad (B.5.6)$$

The contributions from the previous stages become

$$\begin{aligned} \left. \begin{array}{l} U_{N,i}^{K,k} + \\ V_{N,iN}^{K,kN} + \end{array} \right\} &\leftarrow \Psi_{N-1,i}^{K,k} \\ \left. \begin{array}{l} U_{N,N}^{K,K} - \\ V_{N,NN}^{K,KK} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i^K d_L)}{(1 + \alpha_N^K d_L)} \Psi_{N-1,i}^{K,k} \\ \Psi_{N-1,i}^{K,k} &= \frac{R_{f,N-1}^K}{v_N (\alpha_N^K - \alpha_i^K) (1 + \alpha_i^K d_L + \alpha_N^K d_L)} U_{N-1,i}^{K,k} \end{aligned} \quad (B.5.7)$$

and

$$\left. \begin{array}{l} V_{N,ij}^{K,kr} + \\ V_{N,iN}^{K,kK} - \end{array} \right\} \leftarrow \frac{\varepsilon_m R_{m,N-1}^K}{D_{m,N}^K \left((\phi_N^K)^2 - (\phi_j^r)^2 \right)} V_{N-1,ij}^{K,kr}, \quad (B.5.8)$$

$$\begin{aligned} \left. \begin{array}{l} U_{N,i}^{K,k} + \\ V_{N,iN}^{K,kK} + \end{array} \right\} &\leftarrow \Psi_{N-1,ij}^{K,kr} \\ \left. \begin{array}{l} U_{N,N}^{K,K} - \\ V_{N,NN}^{K,KK} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i^k d_L)}{(1 + \alpha_N^K d_L)} \Psi_{N-1,ij}^{K,kr} \\ \Psi_{N-1,ij}^{K,kr} &= \frac{\delta_f \varepsilon_m R_{m,N-1}^K (\phi_N^K \tanh a \phi_N^K - \phi_j^r \tanh a \phi_j^r)}{v_N (\alpha_N^K - \alpha_i^k) (1 + \alpha_i^k d_L + \alpha_N^K d_L) \left((\phi_N^K)^2 - (\phi_j^r)^2 \right)} V_{N-1,ij}^{K,kr} \end{aligned} \quad (B.5.9)$$

From decay in the fracture

$$\begin{aligned}
 \left. \begin{array}{l} U_{N,i}^{K,k} + \\ V_{N,iN}^{K,kN} + \end{array} \right\} &\leftarrow \Psi_{N,i}^{K-1,k} \\
 \left. \begin{array}{l} U_{N,N}^{K,K} - \\ V_{N,NN}^{K,KK} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i^K d_L)}{(1 + \alpha_N^K d_L)} \Psi_{N,i}^{K-1,k} \quad . \quad (B.5.10) \\
 \Psi_{N,i}^{K-1,k} &= \frac{\lambda^{K-1} R_{f,N}^{K-1}}{v_N (\alpha_N^K - \alpha_i^K) (1 + \alpha_i^K d_L + \alpha_N^K d_L)} U_{N,i}^{K-1,k}
 \end{aligned}$$

From the matrix, two components as before:

$$\left. \begin{array}{l} V_{N,ij}^{K,kr} + \\ V_{N,iN}^{K,kK} - \end{array} \right\} \leftarrow \frac{\varepsilon_m \lambda^{K-1} R_{m,N}^{K-1}}{D_{m,N}^K ((\phi_N^K)^2 - (\phi_j^r)^2)} V_{N,ij}^{K-1,kr} \quad , \quad (B.5.11)$$

$$\begin{aligned}
 \left. \begin{array}{l} U_{N,i}^{K,k} + \\ V_{N,iN}^{K,kK} + \end{array} \right\} &\leftarrow \Psi_{N,ij}^{K-1,kr} \\
 \left. \begin{array}{l} U_{N,N}^{K,K} - \\ V_{N,NN}^{K,KK} - \end{array} \right\} &\leftarrow \frac{(1 + \alpha_i^k d_L)}{(1 + \alpha_N^K d_L)} \Psi_{N,ij}^{K-1,kr} \quad . \quad (B.5.12) \\
 \Psi_{N,ij}^{K-1,kr} &= \frac{\delta_f \varepsilon_m \lambda^{K-1} R_{m,N}^{K-1} (\phi_N^K \tanh a \phi_N^K - \phi_j^r \tanh a \phi_j^r)}{v_N (\alpha_N^K - \alpha_i^k) (1 + \alpha_i^k d_L + \alpha_N^K d_L) ((\phi_N^K)^2 - (\phi_j^r)^2)} V_{N,ij}^{K-1,kr}
 \end{aligned}$$

2011:11

The Swedish Radiation Safety Authority has a comprehensive responsibility to ensure that society is safe from the effects of radiation. The Authority works to achieve radiation safety in a number of areas: nuclear power, medical care as well as commercial products and services. The Authority also works to achieve protection from natural radiation and to increase the level of radiation safety internationally.

The Swedish Radiation Safety Authority works proactively and preventively to protect people and the environment from the harmful effects of radiation, now and in the future. The Authority issues regulations and supervises compliance, while also supporting research, providing training and information, and issuing advice. Often, activities involving radiation require licences issued by the Authority. The Swedish Radiation Safety Authority maintains emergency preparedness around the clock with the aim of limiting the aftermath of radiation accidents and the unintentional spreading of radioactive substances. The Authority participates in international co-operation in order to promote radiation safety and finances projects aiming to raise the level of radiation safety in certain Eastern European countries.

The Authority reports to the Ministry of the Environment and has around 270 employees with competencies in the fields of engineering, natural and behavioural sciences, law, economics and communications. We have received quality, environmental and working environment certification.

Strålsäkerhetsmyndigheten
Swedish Radiation Safety Authority

SE-171 16 Stockholm
Solna strandväg 96

Tel: +46 8 799 40 00
Fax: +46 8 799 40 10

E-mail: registrator@ssm.se
Web: stralsakerhetsmyndigheten.se