REGIONAL MODELLING TO FORECAST THE TRANSIENT CHLORINATED SOLVENT RISK TO GROUNDWATER FOR WATER SUPPLY SECURITY

by

CHRISTOPHER JAMES BARRY

Department of Earth Sciences
College of Life and Environmental Sciences
University of Birmingham
August 2017

Abstract

The UK has an environmental legacy of chlorinated solvent contamination from historic industrial activity. Groundwater in Permo-Triassic Sandstone aquifers has been widely impacted. Future use of this resource requires methods to forecast regional-scale chlorinated solvent risk to groundwater assets. Current methods do not adequately represent chlorinated solvent dense non-aqueous phase liquid (DNAPL) source term behaviour and transient contaminant transport. This research investigates the hypothesis that improvements for regional contamination forecasting methods are possible and necessary. Experimentation with particle-based transport solutions shows that neglect of spatial dispersion is over-simplistic, especially where there are complex interactions between capture zones of different receptors. A fully dispersing Lagrangian solution, dynamic random walk (DRW), is developed that accurately models contaminant migration even with coarse-gridded flow models. For representation of the source term, a review of research into DNAPL behaviour is conducted. It is not feasible to derive a generic model for DNAPL dissolution into groundwater based on readily derived metrics at the regional scale. Therefore, a source term framework is developed that tests multiple conceptual models of DNAPL behaviour and therefore gives a holistic uncertainty analysis. Two case studies, based in historically industrial catchments in Birmingham and Stourbridge, validate the combined use of the DNAPL source term framework and the DRW transport solution. Chlorinated solvent sources are persistent, but careful asset planning informed by these regional risk modelling tools — DRW and the DNAPL source term framework — could allow continued use of solvent-impacted groundwater bodies.

Funding

This project is a NERC PhD CASE studentship funded by Severn Trent Water Ltd and the Environment Agency.

Acknowledgements

I would like to thank my PhD supervisors, Michael Rivett and Alan Herbert, for their guidance during the project which was always willing, kind and expert, independent of other pressures they faced. David Lerner, during the early stages of the project and John Tellam, later on, also provided useful comments and steer to the project.

Apart from the commitment of funding, personnel from Severn Trent Water Ltd and the Environment Agency have also greatly helped the project by providing data, giving personal time and clearly explaining the societal and commercial context on which the project was based. Particular thanks are due to Matilda Beatty, Amanda Coffey, Sid Priestley and Antonio Gennarini of Severn Trent Water Ltd and Tim Besien, Andrew Pearson, Arifur Rahman, John Davis and Alistair Brodie of the Environment Agency.

A PhD project can be a solitary experience and so I am grateful for the friendship and collaboration of other PhD and Post-doctoral researchers in hydrogeology at the University of Birmingham. I have particularly benefitted from regular Hydrogeology Research Group meetings with Ben Harvey, Simiao Sun, Ban To, Mahmoud Jaweesh, Omar Al-Azzo, Rohazaini Jamil, Lindsay MacMillan and Mark Cuthbert, in which we shared research updates and advice and campaigned for hydrogeology research. I would also like to thank Mark Hammerton for his collaboration during his MSc project

Finally, my wife Elizabeth deserves high recognition for her kind and patient practical and emotional support and for taking great interest in my work. Our parents also provided indispensable practical support particularly during the thesis writing period.

Reflection

"With joy you will draw water from the wells of salvation." — Isaiah 12:3 NIV

Those working in hydrogeology are often given cause to reflect sadly on past and ongoing loss of clean groundwater due to contamination and of course this is a great shame. But it is easy to take for granted that we have such an abundant drinking water source in groundwater at all. Humanity's capacity to spoil is more than matched by God's capacity to give. Michael Rivett (1989), in his work during the early stages of research into chlorinated solvents in groundwater, quoted the above promise from Scripture, and this promise still holds. The following work is presented with the hope that it will be of some use in providing clean water to future generations. But it is recognised that the complete solution for equitable and secure water supply, as well as all our other needs, does not depend on human effort.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

In the UK, as well as other historically industrial countries, a legacy of chlorinated hydrocarbon (CHC) solvents in the environment formed during the twentieth century. Groundwater surveys and monitoring at public water supply (PWS) wells has revealed that this legacy severely impacted groundwater quality under former industrial regions (section 1.1). For the UK, groundwater in the Permo-Triassic Sandstone, which has been impacted by CHC solvents in many areas, is an increasingly important resource for future water supply security (section 1.2). The benefit of remediating individual contaminated land sites for regional-scale groundwater quality is limited. Therefore, modelling solutions appropriate for multiple industrial spills at the regional scale have emerged to help water resource managers to forecast the chlorinated solvent risk for better informed asset management plans (section 1.3). This approach is much needed, but existing solutions have so far neglected the effects of transient groundwater flow and not incorporated representation of the dense non-aqueous phase liquid (DNAPL) behaviour of chlorinated solvents in the subsurface. This project aims to address these shortcomings by building on the concepts of existing tools and developing methods that can properly represent the chlorinated solvent source term to groundwater and transient dissolved phase transport (section 1.4).

1.1 Chlorinated Hydrocarbon Solvents in Groundwater

The widespread use of CHC solvents by UK industry started in the the 1920s but it was not until the 1980s that the impact of their disposal on groundwater became widely recognised (Rivett et al., 2006b). Rivett et al. (1990) were among the first to confirm the legacy of solvent contamination of groundwater, finding a strong association between solvent contamination and historic industrial land use in the Birmingham Permo-Triassic Sandstone aquifer. Subsequent surveys have found widespread groundwater contamination in other historically industrial cities in the UK (Burston et al., 1993; Tait et al., 2008) and abroad (Grischek et al., 1996; Graber et al., 2008). Tracing contamination back to individual companies, many of which have closed, is often infeasible. Litigation and remediation have their place, but do not solve the problem of the regional-scale legacy of chlorinated solvent contamination originating from a large number of forgotten spill sites under formerly heavily industrialised areas. Such areas include the Tame Valley region in central Birmingham and Stourbridge (chapters 4 and 5).

The chlorinated solvent contamination legacy arises not only from historically poor, or uninformed, handling and disposal practices, but also from their physical properties. Chlorinated solvents behave as DNAPLs in the environment, characterised by low aqueous solubility, a density higher than water and low viscosity (Kueper et al., 2003). Low viscosity and relatively high density causes rapid downward migration penetrating below the water table. The low solubility of chlorinated solvents means that they dissolve slowly, persisting as DNAPL in the environment for a long time after solvent

spillage ceases. This phenomenon is further exacerbated by the irregular geometries commonly adopted by chlorinated solvent source zones, termed source zone architectures, effectively resulting in large proportions of solvent mass being sequestered in slower-flowing portions of aquifer where it forms slowly dissolving sources that can persist for many decades (Brusseau and Guo, 2014). The complex and unpredictable nature of source zone architectures also creates difficulty in quantitatively representing the distribution and dissolution of chlorinated solvents into flowing groundwater. Upscaled modelling solutions for the chlorinated solvent DNAPL source term typically require a large number of parameters that are difficult to estimate and are difficult to transfer to contexts other than the case for which they are demonstrated (for example: Christ et al., 2010). The occurrence and behaviour of chlorinated solvents in the environment and attempts to model their behaviour are discussed more fully in sections 3.1 and 3.2.

1.2 The UK Permo-Triassic Sandstone Groundwater Resource

The UK Permo-Triassic Sandstone (or New Red Sandstone) was formed in terrestrial basins with the onset of the Variscan Orogeny at the end of the Carboniferous Period. Most of the sequence was deposited in fluvial environments, but some formations, such as the Bridgnorth Formation in the West Midlands, were formed in aeolian environments, and these tend to be more lithologically homogeneous (To, 2016). Coarse grain size and limited cementation give much of the Permo-Triassic Sandstone sequence moderate to high hydraulic conductivity and high storage (Allen et al., 1997). Groundwater flow is further enhanced in many (perhaps most) Permo-Triassic Sandstone aquifers by secondary porosity in fractures (Allen et al., 1998; Walker and Jameson, 2010; Brodie, 2012). Further discussion of the hydrogeology of West Midlands Permo-Triassic Sandstones is in sections 4.1.1 and 5.1.1.

Around a third of UK water supply is sourced from groundwater (UK Groundwater Forum, 2011). The Permo-Triassic Sandstone is second most important source of groundwater in the UK after the Cretaceous Chalk and its high storage properties relative to chalk mean it bears the greatest volume of groundwater (Daily and Buss, 2013; To, 2016). With continued population growth, increasing water demand *per capita* and surface water security challenges posed by the changing climate, the Permo-Triassic Sandstone groundwater resource is set to become more important for water supply security. Permo-Triassic Sandstone is widely distributed over the UK, outcropping in the South West, the West Midlands, the East Midlands, the North East and the North West (Allen et al., 1997). Some major bodies of Permo-Triassic Sandstone lie under large cities such as Birmingham, Coventry and Liverpool. This feature is advantageous, giving the opportunity of locally sourced water supply for large populations.

However, urban groundwater in Permo-Triassic Sandstone aquifers faces specific water quality pressures that hinder its use. It is now recognised that historic industrial activ-

ity which is generally concentrated in urban centres has left a legacy of contamination. Inorganic contamination from industry, such as heavy metal salts, are found in some areas, but the Permo-Triassic Sandstone's sorption capacity for cations has served to limit this problem (Ford and Tellam, 1994; Tellam and Barker, 2006). Organic contamination however, in particular chlorinated solvents, is a widespread problem for urban Permo-Triassic Sandstone aquifers (Rivett et al., 1990; Burston et al., 1993; Tait et al., 2008; Rivett et al., 2012). A more detailed review of organic groundwater quality in specific case studies is found in sections 4.1.2 and 5.1.2. Nitrate contamination impacts Permo-Triassic Sandstone groundwater in both rural and urban catchments as a result of agricultural activity (Rivett et al., 2007) and leaking sewers and landfill (Lerner and Harris, 2009). With regards to groundwater quantity, however, excessive quantities of groundwater due to water table recovery following a decline in industrial abstraction give rise to flooding problems for underground infrastructure (Greswell, 1992; Lerner and Barrett, 1996). This creates additional incentive, as well as future water supply security, to use the urban Permo-Triassic Sandstone resource. Such projects will require quantitative methods addressing these water quality risks to inform groundwater asset management and development.

1.3 Regional-Scale Modelling of Aqueous Contaminant Transport in Groundwater

Investigations at contaminated land sites often include detailed aqueous contaminant transport modelling studies (for example: Frind et al., 1999). However, it is recognised that the problem of chlorinated solvents in groundwater often extends beyond contaminated site boundaries and must be addressed at the regional scale (section 1.1; Rivett et al., 2012). Addressing the problem of contaminant transport at the regional scale requires different modelling solutions, because groundwater flow models are less detailed and source locations less well constrained.

Particle-tracking is commonly the basis of regional scale transport modelling tools because this method does not incur obvious numerical error on coarse grids (considered further in section 2.4). The simplest examples of this are the source protection zone (SPZ) and wellhead protection area (WHPA) methods (Environment Agency, 2015; Blandford and Huyakorn, 1991). Not specific to the problem of chlorinated solvents, these methods use capture zone and isochron delineation to identify particularly sensitive areas for regulation of potentially contaminating activities. Reverse numerical methods are also sometimes used to estimate sensitive areas for protecting wells (Frind et al., 2006; Black and Foley, 2013). Tait et al. (2004b) develop a more specialised method to identify well vulnerability from chlorinated solvents which, as well as using particle track-derived capture zones, takes account of land use and vadose zone characteristics.

Models based on the source-pathway-receptor paradigm with quantitative consideration of contaminant flux have also been developed, with Borehole Optimisation Sys-

tem (BOS) and CatchRisk applying a flux calculation to modelled pathlines (Tait et al., 2004a; Troldborg et al., 2008) and FRAC-WECO using a numerical transport solution (Jamin et al., 2012b). These tools use a land-use model giving the locations of historic industries associated with chlorinated solvent use to estimate potential point source locations. They are designed for water resource managers and environmental regulators as prioritisation tools, to identify the most problematic contaminant sources or the most at-risk wells. However, none of the existing methods for regional scale contaminant risk adequately address transient groundwater flow, especially relevant for urban Permo-Triassic aquifers with long-term rebound following the alleviation of industrial abstraction demand, or chlorinated solvent DNAPL source term behaviour. A detailed review of contaminant transport modelling methods with a focus on their application to the regional scale is in section 2.1 and the issue of representing chlorinated solvent source terms is considered in section 3.2.

1.4 A Project to Quantify and Forecast the Transient Chlorinated Solvent DNAPL Risk to Permo-Triassic Sandstone Groundwater Resources and Supply Wells

1.4.1 Project purpose

Some important public supply wells in the Permo-Triassic Sandstone are now abstracting water contaminated with dissolved chlorinated solvent. At the same time, large bodies of fresh groundwater, such as in the Birmingham aquifer, are mostly abandoned because of concerns about chlorinated solvents. These two situations embody the transient and the spatial uncertainties of the regional chlorinated solvent DNAPL risk. Both these uncertainties contribute to a considerable economic and resource risk to water utilities which face, as a result, uncertainty about the future viability of existing PWS wells and the feasibility of developing potentially abundant groundwater resources for future water security.

The transient aspect of this uncertainty arises from lack of quantitative constraint on the transient behaviour of chlorinated solvent DNAPL source terms. Although fresh release of chlorinated solvents into the environment has largely ceased, their low solubility means that chlorinated solvent source zones tend to persist and degrade slowly for decades after the cessation of spillage (section 3.2). The rate and depth-distribution of dissolution of chlorinated solvents into groundwater and the nature of degradation of the DNAPL source zone is a key uncertainty affecting the future water quality at receptors such as PWS wells.

The spatial uncertainty arises because accurate modelling of contaminant transport from point sources at the regional scale is difficult to accomplish with existing tools.

In the UK, an Environment Agency (EA)-led programme to develop regional resource models for major aquifers is well developed and there exist several regional groundwater models for Permo-Triassic Sandstone aquifers (for example: Cuthbert et al., 2001; Daily and Buss, 2013). However, these models have been developed to assess and forecast the distribution of the water budget to streams and other receptors, often with a focus on low-flow periods. They are ill-suited for use with conventional contaminant transport tools such as MT3D (Zheng and Wang, 1999). Poor constraint on the transport of dissolved chlorinated solvents makes it difficult to assess which are the most suitable portions of a groundwater resource to develop to avoid solvent contamination, even when the source locations are well constrained.

The purpose of this project is to address these uncertainties for more informed and confident decision-making about new and existing water supply assets. The chlorinated solvent DNAPL source term and aqueous transport may be studied intensively in the context of an individual contaminated land site (for example: Broholm et al., 1999; Brusseau et al., 2013; Wang et al., 2013; Rivett et al., 2014), but the potential for measurements and observations to be applied to other sites and to the regional scale is limited. The focus of this project, therefore, is the application of the science of DNAPL source terms and aqueous contaminant transport to the regional scale (kilometres to tens of kilometres).

1.4.2 Hypothesis and Objectives

The hypothesis of this project is that significant improvements are necessary and possible for the methods by which future chlorinated solvent risk to groundwater resources is assessed. The particular areas for improvement are a DNAPL-appropriate representation for the source term and transient contaminant transport modelling at the regional scale. The aim of this project is to demonstrate improved chlorinated solvent risk forecasting capability by focussing on these areas.

In order to evaluate this hypothesis, the following objectives were developed:

1. Develop a method for modelling transient aqueous contaminant transport from point sources at the regional scale that is accurate even with coarse-gridded groundwater flow models.
2. Develop a quantitative representation of the chlorinated solvent DNAPL source term. The source term model should fulfill the following criteria:
   (a) It should be consistent with existing research into chlorinated solvent behaviour in the subsurface.
   (b) It should not require a large number of parameters that are difficult to measure or estimate.
   (c) It should take account of historical usage and handling practices of chlorinated solvents by industry.

(d) It should produce a transient dissolved flux output that can easily couple with contaminant transport models, particularly that which is developed in this project, as the source term.

3. Substantiate the aqueous contaminant transport and DNAPL source term methods (as well as providing useful results for the future development and management of groundwater assets) by applying them to regional-scale case study settings in Permo-Triassic Sandstone Aquifers.

1.4.3 Project overview

The development of the methods is described in detail in chapters 2 and 3. Chapter 2 reviews the theory of and existing technologies for aqueous contaminant transport models in groundwater as well as methods specifically designed to predict the regional chlorinated solvent risk (section 2.1), before describing the development of new methods. The development of a contaminant transport method for use with regional models was not trivial and an early attempt, the pathline flux transfer (PFT) model (section 2.3), proved unreliable due to its neglect of transverse dispersion and its sensitivity to particle tracking errors (section 2.4 and 2.6). A superior method which better avoided these issues, the dynamic random walk (DRW) method, was subsequently developed (sections 2.2, 2.5 and 2.6).

Chapter 3 reviews the current understanding of the presence and behaviour of chlorinated solvents in the subsurface, where they behave as DNAPLs (section 3.1). There is also a review of numerical and analytical methods that have been developed for modelling DNAPL distribution and dissolution in groundwater (section 3.2). It is noted that existing methods for modelling the regional risk from chlorinated solvents have tended to represent the source term with a surface or soil-bound source paradigm, neglecting the fact that DNAPLs tend to migrate below the water table and are dissolved directly into flowing groundwater. Section 3.3 describes the development of a framework for modelling chlorinated solvent source terms, incorporating consideration of both historic industrial solvent usage and handling practices and the physical behaviour of DNAPLs. A key feature is that the upscaled equations which represent DNAPL dissolution into groundwater, which are empirically based and uncertain, may be chosen by the user. Section 3.4 outlines how the framework is implemented in code using the functional programming language `R` (R Core Team, 2014) and in particular how the program can use MODFLOW results to estimate the transient groundwater flow through a DNAPL source zone. Finally, section 3.5 demonstrates the source term framework in practice, using a contaminated land-scale case study to validate the method.

Chapters 4 and 5 are regional-scale experiments applying the aqueous transport and DNAPL source term methods in combination. The Tame Valley (chapter 4) is a historically industrial region which has been subjected to multiple organic groundwater quality surveys spanning three decades and using several wells (section 4.1). Thus it presents a rare opportunity to compare model results with historic data distributed

spatially and temporally. By selecting an appropriate groundwater flow model (section 4.2) and designing an uncertain DNAPL source term conceptual model within the DNAPL source term framework developed for this project (section 4.3), a calibration of the DNAPL source term and aqueous contaminant transport is attempted, with restrictions on degrees of freedom imposed in order to generate transferable results (section 4.4).

In the Stourbridge area a PWS well is currently experiencing the impact of tetrachloromethane (TeCM) contamination, probably from historic industrial usage and spillage (section 5.1). This case study area has a larger scale than the Tame Valley study and the aquifer layering is more complex, so that the depth dimension of the source term framework is tested. Additional uncertainty associated with the groundwater flow is tested in this study, by testing the effects of plausible boundary condition adaptations (section 5.2) and investigating different future abstraction scenarios (section 5.3). After a DNAPL source term model is coupled to the transport model (section 5.4 and 5.5), using stochastic parameter testing, it is possible to identify more confidently the approximate location of the source of contamination and to give recommendations for future asset management.

Finally, a summary of the results of the project and how effectively the project aims were achieved, together with remaining limitations and recommendations for future work, is presented in chapter 6.

Chapter 2

SOLVING POINT SOURCE AQUEOUS CONTAMINANT TRANSPORT WITH COARSE-GRIDDED REGIONAL FLOW MODELS: THE DYNAMIC RANDOM WALK METHOD

Urban groundwater offers a promising resource which could contribute greatly to future water security in the UK. However, the historical legacy of point source contamination originating from industrial spills and leaks over the last century presents an unknown risk to existing and prospective public water supply (PWS) wells (Rivett et al., 2012). Economical management and future development of urban groundwater resources therefore requires tools for forecasting transport from point sources to receptors. Contaminant transport modelling in groundwater is a well developed field of research. However, different contexts and scales of study have different requirements and are not suited by all contaminant transport solution methods due assumptions and numerical errors that become problems in different situations. At the scale of regional aquifers, there is a particular paucity of satisfactory contaminant transport solutions. In the UK, flow models for regional aquifers are designed to model water balances, a problem that does not require fine model grids, and as such regional models are constructed with very coarse grids. This feature all but precludes the accurate use of grid-based numerical transport solutions such as those within MT3D.

Section 2.1 discusses the approaches that have been taken for modelling risk to receptors at the regional scale. Both forward-tracking methods and reverse-tracking methods are reviewed. Typically, these methods choose between particle tracking and gridded numerical solutions, the MODPATH and MT3D paradigms respectively, to calculate the contaminant pathways and flux between sources and receptors. However, both these method classes have severe limitations when applied to coarse-gridded regional models. Gridded (Eulerian or part-Eulerian, technically) numerical models tend to incur artificial oscillations and excess dispersion (as demonstrated in section 2.6). Particle tracking is much easier to implement stably on coarse grids and the results can look realistic if it is felt that dispersion may be justifiably neglected. However, pathlines on coarse grids are susceptible to misdirection (section 2.4) and dispersion, especially transverse dispersion, can be an important process in situations with multiple potential receptors with adjacent capture zones (section 2.6). What is more, the treatment of pathlines that pass through partial sink cells is generally unquantified. For example, it may be unclear if a gaining river passing between a source and a receptor would protect the receptor. This issue is discussed, with a solution, in section 2.2. Additionally, existing regional risk estimation tools are generally restricted to steady-state flow fields, an assumption which is not easily applied to many urban regional aquifers.

Following sections 2.1 are a series of models and modelling experiments contributing solutions to these gaps in regional contaminant risk modelling. MassTrack (section 2.2) acts as a post-processor to MODPATH-generated pathlines that can track mass as well as location, calculating mass removed from particles that pass through discharging cells. The pathline flux transfer (PFT) model develops further the transport algorithms used in steady-state risk models such as Borehole Optimisation System (BOS) and CatchRisk (Tait et al., 2004a; Troldborg et al., 2008) to offer a pathline-based solution that can work with transient flow fields (section 2.3). However, the experiments of section 2.4 show that tracking single particles can be inaccurate on coarse grids because of the assumption of uniform flow distribution through cell faces, casting doubt on the validity of non-dispersing particle-based solutions. Section 2.5 sets

out the solution to this by developing a dispersing particle-based solution. As a Lagrangian solution, the dynamic random walk (DRW) method avoids incurring artificial oscillations and dispersion. By splitting and displacing particles according to a multivariate probability density function, and then grouping clusters with a fast coalescing algorithm (appendix A), full longitudinal and transverse dispersion may be calculated. This method is benchmarked against the Method of Characteristics (MOC) solution and the analytical Domenico (1987) solution (section 2.5.3, appendix B). The PFT and DRW methods and MT3D solutions are then applied to the case study of the Birmingham Aquifer, considering conceptual risk from historical industrial land use to multiple transient wells (section 2.6). This experiment shows that PFT's neglect of spatial dispersion results in missing several receptors, whereas the MT3D solutions incur numerical dispersion, such that risk to receptors is overly conservative and not well distinguished, and artificial oscillations, such that the results cannot be considered reliable. The DRW method, however, is able to model dispersion accurately, such that contaminant crosses capture zone boundaries but does not carry significant contaminant to regions of the model far removed from contaminant plume distributions, given the flow direction and dispersivity.

2.1 Review of Regional Transport Methodologies for Resource Management

The field of groundwater quality regional risk methodologies generally attracts simplified methods for contaminant transport. Often the flow models to which they are applied are at the regional scale and not optimised for transport simulations, so detailed numerical simulations may be unstable or give rise to artifacts. Risk simulations may require that multiple realisations are executed in order to account for model uncertainty and in these cases it is important that models can run quickly. Finally, the users of risk assessment tools are often not experts at implementing numerical simulations, so it is important that the software packages are easy to use.

Computational risk assessment tools for groundwater quality may be divided into two groups, broadly. Some methods, such as BOS (Tait et al., 2004a) and CatchRisk (Troldborg et al., 2008) simulate forwards in time (section 2.1.1). These methods tend to focus on identified contaminant sources and are designed to identify which receptors are put at risk. Other methods, such as the delineation of source protection zones (SPZs), simulate backwards in time and focus on the capture zones of receptors, delineating sensitive areas for environmental management (section 2.1.2).

Most regional risk assessment methods include a transport simulation module, as clearly the pathway from source to receptor is a crucial component in predicting the likely impacts that may be seen at a receptor. This presents a challenge, as the simulation of contaminant transport is a complicated task with potential pitfalls and regional groundwater models are usually not optimised for transport modelling. This is particularly true of transport from point sources which tends to produce low-dispersion narrow

plumes (Rivett et al., 2006a). Section 2.1.3 discusses these challenges and reviews how the available contaminant transport solutions perform in the context of regional scale transport simulation.

2.1.1 Forward tracking methods

Forward tracking risk methods focus on contaminant inputs to the groundwater system. A key component of this is considering the spatial distribution of activities at the surface and the impacts that these may have on the underlying groundwater. Vulnerability models assess aquifer vulnerability without further coupling to groundwater transport models. DRASTIC (Aller et al., 1987) attempts to map pollution potential (termed DRASTIC index therein) as a risk parameter which depends upon seven factors which include, broadly, the recharge, depth to water table and underlying geology. Each factor is weighted according to its importance and assigned a value that reflects how it may enhance the contaminant risk to groundwater. Aller et al. (1987) present eight distinct hydrogeological settings relevant to the USA and demonstrate the use of the tool, resulting in a DRASTIC index for each. Although the method is a useful exercise for prioritisation, the design is somewhat subjective and has been shown to be significantly dependent on the analyst's judgment (Tait et al., 2004b; Frind et al., 2006).

Tait et al. (2004b) builds on the concept of vulnerability assessment by including some representation of contaminant sources and receptors (including aqueous transport). Their method is specific to the problem of chlorinated solvent contaminant sources. They produced an approximate distribution of solvent sources by a postal address directory search for solvent-associated industrial activities, spatially referenced by post code. The receptor and transport component follows the SPZ methodology (Environment Agency, 2015), with travel time zones giving the transport factor to the risk assessment. The final risk assessment index combines the locations of supposed solvent sources, travel time to receptor zones and aquifer vulnerability distribution. The technique may be used to prioritise well quality inspections and treatment, as well as for planning new supply wells. The vulnerability assessment factors are more rigorously categorised than in the DRASTIC method, so there is greater consistency; this is helped by the method being developed for a more specific context (chlorinated solvents in sandstone). However, there is still subjectivity in how the categorisations are devised for the method.

Some risk tools simulate transport from source to receptor explicitly. BOS (Tait et al., 2004a; Chisala et al., 2007; Tait et al., 2008), CatchRisk (Troldborg et al., 2008) and FRAC-WECO (Jamin et al., 2012a,b) all employ three basic components: source distributions, a source term representation and contaminant pathway simulations. Each method requires that a database of potential source locations is constructed. Each source is then associated with a site area and linked with certain contaminants according to the type of industrial activity that occurred at the site. All three methods use

a surface leaching paradigm for the source term, whereby the strength of the source term J is given by:

$$J = fC_S RA \tag{2.1}$$

where C_S is solubility, R is recharge, A is site area and f summarises mass loss in the vadose zone. Subsurface dense non-aqueous phase liquid (DNAPL) dissolution is not considered, although the conceptual frameworks for these methodologies would not preclude such a development. For the groundwater flow field, the finite difference code MODFLOW is used (Harbaugh et al., 2000), with steady state models in each case. BOS and CatchRisk use the particle-tracking MODFLOW post-processor MODPATH (Pollock, 1994) to simulate pathways from sources to receptors, whereas FRAC-WECO uses MT3D (Zheng and Wang, 1999), which allows the inclusion of dispersion. Troldborg et al. (2008) correct for the neglect of dispersion, to an extent, by shifting sources which are just outside a receptor's capture zone across the divide, although it should be noted that around the capture zone boundary there is high variability in travel time to receptor as pathlines come close to the stagnation point at the capture zone tip. BOS uses stochastic MODFLOW, which has the effect of softening the capture zone boundaries due to uncertainty. BOS and CatchRisk both apply analytical transport equations to the simulated pathlines in order to estimate the concentration at the well. The travel time of the pathline is used to determine the contaminant degradation along the pathline and, in the case of CatchRisk, to estimate a first arrival time for the contaminant (no dispersion is applied). BOS lacks any accounting for delayed arrivals, so sources that are distant from the receptor will be equally important to those close by, apart from the effect of degradation if included.

2.1.2 Reverse tracking methods

The delineation of well capture zones and isochrons is perhaps the simplest example of a groundwater risk assessment methodology. In the UK, the Environment Agency uses the SPZ concept to manage land-use in the vicinity of public supply wells (Environment Agency, 2015). In the USA, the similar wellhead protection area (WHPA) concept is used (Blandford and Huyakorn, 1991). Although there is a variety of techniques used to delineate SPZs and WHPAs, depending on the context, most commonly reverse particle tracking is used, tracing back pathlines from the well in all directions, taking care to account for the stagnation points (Blandford and Huyakorn, 1991; Bakker and Strack, 1996). The WHPA methodology also allows for analytical solutions to be used in the absence of a groundwater flow model. SPZs are divided into three zones, with zone 1 being subjected to the tightest regulation as the zone which is closest to the well, extending to zone 3 which covers the entire capture zone of the well. These methods are relatively simple to implement, which gives the advantage that it is feasible to have a nationally consistent well protection scheme. The results are also easily plotted on a geographic information system (GIS) and are transparent to the public.

Other methods of capture zone delineation than analytical solutions and particle tracking exist. FlowSource (Black and Foley, 2013) is a finite difference solution to advective flux that can trace back water to its origin. FlowSource has the advantage of being less deterministic than particle tracking as the result does not give a hard capture zone boundary which would imply more certainty than is likely, especially considering the possible artifacts inherent in particle tracking (see section 2.4). However, FlowSource assumes complete intra-cell mixing of water, which introduces numerical dispersion. Additionally, FlowSource cannot track travel time.

It is also possible to apply the advective-dispersive transport equation in reverse, modifying the equation to form an adjoint equation (Neupauer and Wilson, 2004, 2005). This approach is demonstrated by Neupauer and Wilson (2005) and Frind et al. (2006). Whereas forward transport modelling shows the distribution of contaminant fate from a given source, reverse modelling shows the probabilistic source location for observed or hypothetical contaminant at a given receptor. Frind et al. (2006) are able to plot spatial distributions of intrinsic well vulnerability to potential contaminant sources at any given location in its capture zone, through a number of metrics such as mean travel time and expected attenuation of peak concentration. Such an approach is very useful for prioritising land-use management above used groundwater resources.

2.1.3 Performance of solution algorithms to contaminant transport with regional scale groundwater models

Sections 2.1.1 and 2.1.2 summarise regional risk assessment methods that make use of a variety of contaminant transport methods. The varying complexity of transport solution methods make them suitable to different tiers of risk assessment. Analytical solutions require a very simple conceptual model for the pathway between source and receptor and are generally only suitable for a single source-receptor pair. Analytical solutions offer several advantages where they are appropriate, however. They are relatively simple to understand and to implement, enabling the development of spreadsheet models that can be widely distributed to practitioners who have limited understanding of numerical modelling and may have limited access to advanced software. They are also fast, stable and do not produce artifacts. Examples include a tiered risk assessment spreadsheet model developed by Entec and the Environment Agency (EA) (Remedial Target Methodology: Carey et al., 2006) and ConSim, which also includes stochastic consideration of uncertainty (Environment Agency, 2003).

For regional scale transport simulations, modellers often use simple particle tracking for simulating migration of contaminants, as with the BOS and CatchRisk methodologies (section 2.1.1). Regional models are typically built with coarse grids in order to keep run times manageable and because the aims of the model are principally to forecast groundwater head and flux distributions. And particle tracking is generally much simpler to execute successfully than numerical advective-dispersive transport models. Another concern at the regional scale is model portability, as regional risk models must be usable and understandable by a diversity of resource managers across organisations,

as well as the public to an extent. This contrasts with studies of individual contaminated sites where confidentiality is more important. BOS uses particle-delineated steady-state capture zones to determine whether a receptor is at risk from each of a collection of sources, and then a one-dimensional steady-state transport equation to calculate the scale of that risk (Tait et al., 2004a; Chisala et al., 2007; Tait et al., 2008). CatchRisk uses a very similar approach, but uses pathline-derived travel times to determine the likely arrival time of a contaminant plume (Troldborg et al., 2008). However, simple advective particle tracking as used in BOS and CatchRisk neglects dispersion, especially transverse dispersion which would be significant even after plumes reach steady state.

There are different algorithms that can be used to find a particle's trajectory from a given location. Two options are implemented within the MOC transport solution for the advection process, which is represented by particle tracking (Zheng, 1993). The simplest is the 1st order Euler method, which explicitly projects the particle's velocity forward for a given time interval, and is only suitable for almost uniform flow fields and small time steps. The MOC algorithm in MT3D provides another method, the 4th order Runge-Kutta method, which uses the average of four different velocity estimates, both explicit and implicit, to estimate the average velocity along the trajectory during the time step. MODPATH (Pollock, 2012), the popular pathline tracer post-processor to MODFLOW, is able to be more flexible with its algorithm, because it does not use predetermined time step lengths. Each particle may be tracked independently, in contrast to the MOC method where all particles must be stopped at the end of time steps in order to simulate dispersion and other reactive transport processes. MODPATH therefore tracks particles between cell faces and the time steps from MODFLOW, saving particle locations at each of these points. This means that for each pathline segment, a particle is only inside one distinct cell. Therefore, MODPATH uses an analytical expression for the curved trajectory of the particle through the cell, based on linear interpolation of the flows across the cell faces. MODPATH 6 also considers the likely position of any boundary fluxes within the cell, with, for example, recharge applying to the top face and well abstraction drawing from the cell centre. MODPATH uses an analytical expression for trajectories within a cell, using linear interpolation to calculate the variable velocity components within the cell assuming uniform flow across cell interfaces. The implications of this are discussed in section 2.4.

Numerical transport solutions may, broadly, be divided into Eulerian, Lagrangian and mixed Eulerian-Lagrangian methods. Eulerian methods track the concentration field in finite blocks: the grid cells. Lagrangian methods use a set of infinitesimal points (particles) to describe a continuous concentration field. MT3D's finite difference and total variation diminishing (TVD) algorithms are pure Eulerian methods. The MOC algorithm (and its derivatives modified MOC (MMOC) and hybrid MOC (HMOC): Zheng, 1990) is mixed Eulerian-Lagrangian, using a set of particles to track the movement of contaminant mass due to advection but reverting to the finite difference method for all other processes within the transport equation including dispersion. MOC is designed to reduce numerical dispersion by decoupling the advection process from the grid (Zheng and Wang, 1999). However, it is still susceptible to numerical dispersion during the so-

lution of the dispersion step, which reverts to the finite difference method. The random walk particle method (RWPM) is a pure Lagrangian method, using particle movements to represent displacement due to both advection and dispersion (Salamon et al., 2006). The RWPM is often used for transport studies at very small scales, such that particles may even be considered to represent a mere handful of actual molecules (Jiang and Wu, 2013; Rahbaralam et al., 2015), though this level of detail is not feasible or necessary for aquifer-scale groundwater transport (Fernàndez-Garcia and Sanchez-Vila, 2011).

Eulerian numerical methods are susceptible to artificial oscillations and instability in the solution that commonly feature negative and extreme values. Explicit numerical solution schemes are particularly prone to this because fluxes due to large concentration gradients evaluated at the start of a time step, when projected to the end of the time step, can lead to unphysical (perhaps negative) values for concentration. MT3DMS uses the generalised conjugate gradient (GCG) solution scheme, an implicit method. The GCG scheme, although more complicated and in principle slower, is less prone to instability than explicit schemes and so it is easier to use and may use longer time steps (Zheng and Wang, 1999). The finite difference method within MT3D also has an option for upstream weighting, which, if used, takes the most upstream concentration gradient estimate for each cell when evaluating the advective flux. Because at the edge of a plume the concentration gradient will be greater upstream, this option has the effect of reducing the curvature of the concentration field. While this enhances the stability, it also means that some dispersion is artificially simulated while solving for advection, causing numerical dispersion. Nonetheless, it is still possible to have an unstable solution with upstream weighting and it is also possible to have numerical dispersion without it because intra-cell mixing is always assumed by Eulerian methods, as seen in section 2.6.

Pure Lagrangian methods, of which the RWPM is the most common in groundwater solute transport, are grid-decoupled. This means that, unlike for Eulerian methods, the state of the model is not an array of values for the centres of cells, but a table of particles each with real number (or floating point) co-ordinate references, such that the concentration field is described by any number of arbitrarily located points. This removes the need to assume any mixing and as such avoids numerical dispersion. Because Lagrangian methods are not locked to the grid, there is no need to consider dimensionless stability criteria such as the Péclet and Courant numbers, as the movement of mass may be represented over the most suitable distance automatically. Therefore, Lagrangian solutions are free from numerical instability and dispersion and have often been the choice for complex high-resolution transport simulations (Salamon et al., 2006).

Nonetheless, the RWPM has some complications and limitations. Because particles represent points rather than finite volumes, they cannot directly represent a concentration field. This makes it difficult to simulate non-linear reaction processes and to visualise the results spatially (Fernàndez-Garcia and Sanchez-Vila, 2011; Rahbaralam et al., 2015). Also, heterogeneous variation in dispersion coefficient can lead to an error termed the local mass-conservation problem, unless corrected for. This problem arises

at sharp interfaces of distinct dispersion coefficient, where dispersing particles would be expected to favour movement toward the region of higher dispersion coefficient. However, when dispersion length is estimated from the particle's starting location, this is not captured and the error often results in an artificial accumulation of solute mass in low dispersion coefficient (often corresponding to low permeability) regions. There are a number of solutions implemented for this problem (LaBolle et al., 1996, 2000; Salamon et al., 2006).

2.1.4 Dispersion in Groundwater

Section 2.1.3 discusses how well transport modelling methods can produce accurate solutions to given equations for advective-dispersive contaminant transport. However, the understanding of contaminant dispersion is currently debated and therefore the most appropriate equations for representing the process is uncertain. The conventional model of dispersion is the model of hydrodynamic dispersion, which is a combination of principally mechanical dispersion, whereby pathlines mix due to heterogeneous travel times and directions at multiple scales (pore scale and formation scale, for example) and molecular diffusion (Bear, 1979; Bredehoeft and Pinder, 1973; Fetter, 1999). While this model of dispersion is popular and probably the most common model of dispersion used in contaminant transport studies, both analytical and numerical, its physical basis is far from robust and was described as a "working hypothesis" by Bear (1979), one of the early proponents of this model (Konikow, 2011). The predictions of the hydrodynamic dispersion model are of a concentration field that evolves similar to anisotropic Fickian diffusion, which is symmetric (normally distributed) about the mean travel distance. However, there is a growing collection of field evidence that the hydrodynamic dispersion model is not able to reproduce observations of contaminant plumes, and that an asymmetric (log-normally distributed) evolution of concentration is in fact more commonly seen (Zheng and Gorelick, 2003; Bianchi et al., 2011; Hadley and Newell, 2014). Bianchi et al. (2011) suggest that rather than mechanical dispersion, inter-diffusion between immobile and mobile pore water is the dominant process by which contaminant mass is dispersed, especially in very heterogeneous situations with large permeability contrasts (Bianchi and Zheng, 2016). Their results are much better reproduced by a dual-domain single rate mass-transfer model, which has an additional process representing mass transfer between mobile and immobile pore water:

$$\phi_{\mathrm{I}} \frac{\partial C_{\mathrm{I}}}{\partial t} = \kappa \left(C - C_{\mathrm{I}} \right) \tag{2.2}$$

where ϕ_{I} is immobile porosity, C and C_{I} are mobile and immobile phase concentrations respectively, t is time and κ is the mass transfer coefficient.

The focus of this chapter is the development of solutions to contaminant transport that may be applied to the problem of regional contaminant transport, repurposing regional flow models without incurring solution-derived errors. The work presented aims to assess whether solution schemes produce reliable output in this context and

assessment of the understanding of the dispersion process is considered beyond its scope. Therefore the more common hydrodynamic dispersion model is adopted for the remainder of this chapter, though further development to allow the incorporation of dual-domain transport models is by no means precluded (section 6.5).

2.2 Tracking Particle Mass in Pathline Simulations

Particle-tracking as a solution to contaminant transport is generally viewed as a qualitative approach because it does not give a numerical estimate for concentration, merely the path taken by non-dispersing particles and which receptor captures them. This causes difficulties when particles pass through weak sinks. A model cell that has more water flowing in than flowing out, but still has some water flowing out is a weak sink. In these cases, an arbitrary decision must be made as to whether particles in weak sink cells terminate or continue (Pollock, 2012). Weak sinks may result from weakly abstracting boundary conditions or imperfect convergence of the flow solution.

This section describes an algorithm for tracking mass along pathlines. Conceptually, a mass value is attached to each particle. Particles are allowed to pass through all weak sinks, but when there is a net loss of water in a model cell, then mass is extracted from any particles that pass through. Because the mass of the particle should have no effect on the particle's trajectory (neglecting density effects), the package, called MassTrack (Barry, 2017d), can run as a post-processor to particle tracking. MassTrack is currently programmed to analyse pathlines produced by MODPATH version 5 (Pollock, 1994)[1]. By giving quantitative consideration to the fate of mass along pathlines, MassTrack helps to allow more accurate contaminant transport modelling based on pathlines. MassTrack is a component of contaminant transport solutions developed later in this chapter: the PFT and DRW models (sections 2.3 and 2.5). The package is included in the supplementary electronic material (appendix D.1).

The results of MassTrack show how much of a particle's mass is expected to reach the particle's destination. For example, if the particle passes through a cell in which a river boundary condition (Harbaugh et al., 2000) has been specified, if that boundary is gaining, then a portion of the particle's mass will be lost according to the strength of the abstraction, the total volume of water in the cell and the length of time the particle spends in the cell. The cell references at which mass is abstracted from particles is also saved and this can be used to calculate the flux to weak and strong receptors. Additionally, a first-order decay may be specified to represent degradation of mass.

2.2.1 Algorithm

MassTrack works as a post-processor to MODPATH, which is itself a post-processor to MODFLOW. The output from MODPATH includes a data table giving details of the

[1]MODPATH version 5 does not have its own formal documentation which can be cited.

pathlines traced by each of the released particles. The columns of this data table give the cell reference, co-ordinates and model time value for each saved point (generally the cell faces) along each pathline. Pathlines follow the transient groundwater flow field as it evolves with time and are distinguished from streamlines which represent steady-state snapshots of the flow field (Blandford and Huyakorn, 1991).

The MODFLOW model results are also used. The flows through each cell face are required and flows into and out of storage, as well as the saturated cell dimensions. If the system is unconfined, transient piezometric head values are required.

The basic conceptual elements of MassTrack are:

1. A particle travelling through a cell cannot gain mass, only lose or retain.

2. Each cell at each time point has a value for flow imbalance, which is the sum of the flows out of the cell through the cell faces and into storage minus the sum of the flows into the cell through the cell faces or from storage (and thus a positive imbalance in this case implies that a cell is losing water).

3. The proportional rate of water volume loss from a cell is the ratio of the flow imbalance to the volume of water in the cell. If this value is greater than zero, then this value is used for the proportional rate of mass loss from a particle within the cell.

4. The mass loss from a particle in a cell is the rate of mass loss times the amount of time spent in the cell.

5. First order reactive loss of mass may also be applied along a pathline, expressed as an exponential decay.

This is expressed mathematically as follows. A pathline has a calculated trajectory with P segments. A pathline could be divided into limitless segments, but it is sufficient and necessary that a segment is terminated at each cell face and in time by the end of each MODFLOW time step. At the beginning of segment p, at time t_{p-1}, the particle has mass m_{p-1}. The segment lasts for duration $\Delta t_p = t_p - t_{p-1}$ and the mass changes by $\Delta m_p = m_p - m_{p-1}$. The following paragraph explains how Δm_p is calculated. Figure 2.1 shows the meanings of the variables graphically.

A cell ($[i, j, k]$) has water volume

$$V_{\mathrm{w}}[i, j, k] = \Delta x[i]\Delta y[j]\Delta z[k]\phi \tag{2.3}$$

where Δz may refer to the cell height if fully saturated or the saturated thickness of the cell if it contains the water table (figure 2.1 shows the saturated case). ϕ is porosity. The net outflow of water from a cell may be expressed as:

$$\begin{aligned} Q_{\text{imbalance}}[i,j,k] =& Q_x[i,j,k] + Q_y[i,j,k] + Q_z[i,j,k] + Q_s[i,j,k] - \\ & Q_x[i-1,j,k] - Q_y[i,j-1,k] - Q_z[i,j,k-1] \end{aligned} \tag{2.4}$$

Therefore the proportional removal rate of water from a cell (presumed to be caused by boundary conditions) is:

$$\frac{Q_{\text{imbalance}}[i,j,k]}{V_{\text{w}}[i,j,k]} \tag{2.5}$$

The proportional removal rate of water may be thought of as the inverse of residence time.

With the assumptions that the abstraction is evenly distributed throughout the cell and that abstracted water has the same concentration as the source water (see section 2.2.4), the same proportionate reduction rate may be applied to the mass of particles in the cell. Conversely, a cell that is being recharged is assumed to be recharged with pure water, so if $Q_{\text{imbalance}} < 0$ then the particle mass is unchanged in that segment. Therefore:

$$\begin{aligned} \Delta m_p = & \text{ if } Q_{\text{imbalance}}[i_p,j_p,k_p] > 0\text{: } \frac{Q_{\text{imbalance}}[i_p,j_p,k_p]}{V_{\text{w}}[i_p,j_p,k_p]} \times m_{p-1} \\ & \text{else: } 0 \end{aligned} \tag{2.6}$$

If linear decay is specified, then the masses of the particle at the different locations along the pathline are reduced with time according to:

$$m_{p,\text{ decay}} = m_p e^{-\lambda(t_p - t_0)} \tag{2.7}$$

where λ is the first-order decay constant. Mass given to receptors during particle draining is also reduced in the same way, to reflect the fact that solute mass would have decayed before reaching a sink.

2.2.2 Implementation

This section gives an overview of the key features of MassTrack code and the packages it uses. MassTrack is written in the `R` programming language (R Core Team, 2014). MassTrack reads the pathline data file created by MODPATH as a data table and uses MODFLOW flow results converted to NetCDF format (Rew et al., 2016; Michna and Woods, 2016; Barry, 2017e). The `data.table` package for `R` provides memory-efficient tools for sorting, grouping and operating on data tables (Dowle et al., 2015). NetCDF data format allows fast and memory-efficient access to large arrays of data, of which only a small portion is generally required at one time.

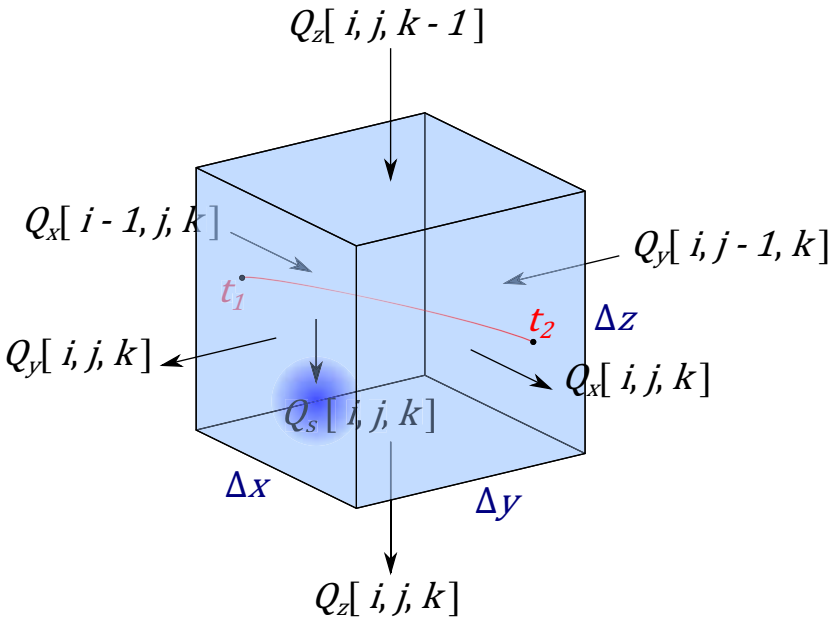

Figure 2.1: Visualisation of a pathline (red line) through a saturated cell, with water fluxes to and from the cell labelled. Q refers to a volumetric flux of water and Q_s is the flux to storage. The particle describing the pathline enters the cell at time t_1 with mass m_1 and exits the cell at time t_2 with mass m_2.

The flow budget data from MODFLOW includes flow through the right (x direction), front (y direction) and lower (z direction) faces, compensation flows to and from storage and flows to and from boundary conditions (which are not used by the current version of MassTrack). The flow and storage arrays are read for each cell which pathlines traverse, as well as for the appropriate adjacent cells which may be feeding or drawing from these cells. `R`'s index matrix feature, which allows irregular subsets of data arrays to be read very quickly, greatly helps the efficiency of this process. From this information, MassTrack calculates the net abstraction rate from each cell along each pathline. This data is divided by the cell volumes to give the proportional drainage rate of the cells which is used as the proportional mass reduction rate for particles in those cells, as described mathematically in section 2.2.1.

MassTrack returns a list of data, including the pathline data table with a mass column affixed and also columns for drained and degraded mass at the appropriate points in the path segments.

2.2.3 Example

A fictional model is used to demonstrate the functioning of MassTrack. The model (figure 2.2a) is built on a grid of 30×20 square cells. The model is unconfined, with water flowing between two fixed head boundaries in a direction oblique to the grid.

Between the boundaries are three wells operating transiently in cycles and a river boundary. The fixed heads are at 60 and 50 (units arbitrary) and the river, which is horizontally roughly halfway between the two fixed heads, has a stage of 52, and as such acts as a partial sink to groundwater except in places where the local effect of the wells causes drawdown of the surrounding groundwater. Fifty particles are released from a straight line roughly perpendicular to flow and close to the upstream fixed head boundary and each are assigned a starting mass of 1. Each particles is released at time 0 from the top of the water column (z-offset is 1).

The forward-tracked pathlines produced by MODPATH 5 from these particle starting locations are shown in figure 2.2b. Figure 2.2b also shows the evolving particle masses, represented by the size of the purple points, and figure 2.2c shows where the mass is removed from particles, mostly at the edge of the river boundary which abstracts most of the particles' mass. A few particles retain some mass beyond the river and deliver their remaining mass to a well. Figure 2.2d shows the interpreted solute flux (units of mass per time) to the river, revealing that the left part of the river receives most of the solute loading, probably because the drawdown from the rightmost well is greatest.

In this example, it would not be clear from particle tracking alone whether the river or the wells should be considered as the most significant receptor. MassTrack reveals where the solute mass represented by the particles is delivered. In this case the river is a sufficiently strong sink to abstract most of the passing mass. MassTrack does not, however, consider the height of particles within a cell, which could be important if the river should only take water from the top of the cells. This is discussed further in section 2.2.4.

2.2.4 Limitations

MassTrack only detects net drainage of a cell. Therefore, a well boundary that shares a cell with a recharge input will appear as a weaker sink to MassTrack. This is because MassTrack only considers the flows through the faces and into storage. In fact a cell with both sinks and sources should drain the mass from particles more quickly than the equivalent net sink in isolation. It would be possible, with further development, for MassTrack to analyse the boundary fluxes instead of the cell face flows and this would correct for this limitation, with the requirement that MODFLOW saves all boundary fluxes.

Some boundary conditions, such as rivers, are intended to represent a boundary close to the water table or the top of the model layer. MassTrack assumes (similar to numerical transport codes such as MT3D) that all sinks are evenly distributed within the cell that they occur. The modeller could divide the cells up so that such boundaries occupy more specific height ranges. Alternatively, MassTrack could be updated in the future to allow for specifications for the distribution of boundary sinks within a layer. This would be assisted by the fact that MODPATH saves z-offset information in the pathline data file. For example, if such an update were undertaken a user could set

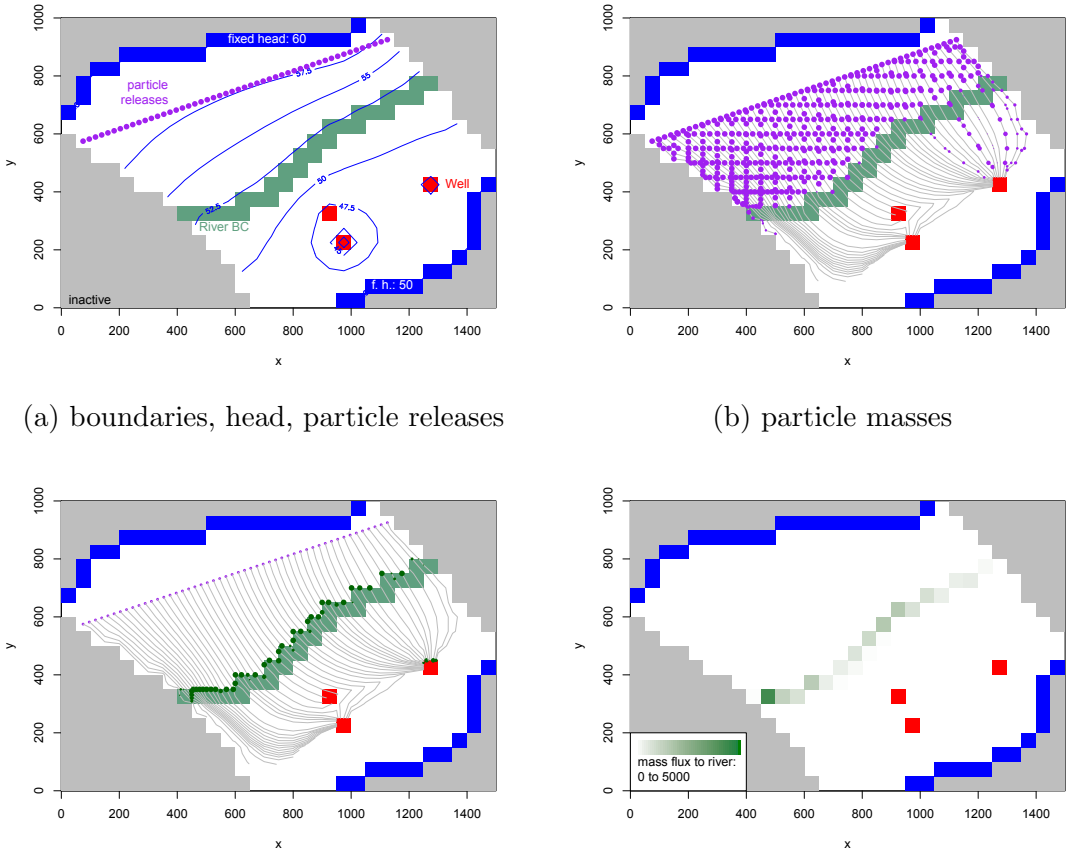

(a) boundaries, head, particle releases

(b) particle masses

(c) particle mass losses

(d) mass flux to river

Figure 2.2: Example results from MassTrack. (a) shows a summary of the set up of the model, including the particle release locations. All particles start with the same mass. (b) shows the particle masses along their pathlines, demonstrating the reduction of mass as the particles cross the river and arrive at wells. (c) shows the locations at which the particles lose their mass. For (b) and (c) mass is represented by point size. (d) shows the distributed mass flux to the river as the particles pass.

that a river boundary only drains particles that are in the top half of the cell (that is, z-offset is greater than 0.5). Alternatively, a calculation could be performed to determine the proportion of saturated thickness that should be drained by the river assuming no mixing of water, such that the river boundary acts as a strong sink in the upper portion of the cell and does not drain any solute mass from the lower portion, separated by a calculated threshold z-offset.

These limitations are not thought to be critical for the application of MassTrack in the methods featured later in this chapter (sections 2.3 and 2.5). In cases where weak receptors abstract from a cell that also receive recharge, the flux may be underestimated, but at strongly abstracting wells, all arriving particle mass will be drained quickly for typical recharge values, and if the sink is in a lower layer then this ceases to be an issue at all.

2.3 A Simple Particle-based Regional Risk Simulation

This method is a development of the BOS and CatchRisk methods for regional risk modelling (Tait et al., 2004a; Troldborg et al., 2008). An overview of these methods, their capabilities and limitations is given in section 2.1.1. Neither of these methods are able to consider transient flows. A single pathline is only able to represent a single transient flow path which does not persist. Therefore, in models that are transient and without a suitable steady-state summary model, multiple particle releases from each source location are required to capture the evolving set of pathlines.

This section describes the algorithm for the pathline flux transfer (PFT) method, developed for this project, which largely resembles the key components of BOS and CatchRisk but extended for transient flow fields. The incorporation of the MassTrack program (section 2.2, Barry, 2017d) to account for partitioning of flux to weak sinks along a pathline is a further development. The use of this method is demonstrated in section 2.6.

2.3.1 Algorithm

In the following section, a particle is a conceptual tool, simulating a small packet of water released from a specified location and used here to represent a mass of contaminant which in turn represents a contaminant flux over a certain time duration. A pathline is the trajectory followed by a single particle. A pathline follows the fully time-varying flow field, in which respect it is different from a streamline, which follows a time slice of the flow field. A pathway (as in source-pathway-receptor) is a full description of the linkage between source and receptor over time via a set of pathlines. Between a particular source and receptor pair, there may be a pathway during some periods and not at others.

The conceptual model is formed of four components. These are:

1. Pathline simulation.
2. Pathline analysis.
3. Source term.
4. Flux transfer algorithm.

Pathline Simulation

Any tool assessing the risk to a receptor must consider the pathway to the receptor, however simply. This requires some knowledge of the groundwater flow field.

The pathline simulation is done by particle tracking, performed using MODPATH (Pollock, 2012). The particles are released from the specified source locations. In a transient flow field, the pathway will also be transient in duration, length and existence (meaning that the transient capture region of the receptor will sometimes encompass sources and sometimes not). A single particle release will not capture this, so multiple particles are released from each source location, at specified points in time. This is why this tool is considered to be semi-analytical: it is spatially analytical and temporally discretised. Note that the retardation of the contaminant must be accounted for at this stage by adapting the true porosity value in the model to porosity multiplied by the retardation factor. MODPATH 5 does not have a parameter that represents the effects of chemical sorption, but if the assumption of equilibrium sorption (described by retardation factor) is valid, then the effects of retardation are equivalent to those of porosity — the average velocity of dissolved particles is inversely proportional to the retardation factor just as it is to the effective porosity.

DNAPL sources tend to be distributed vertically through the aquifer. In a multi-layered aquifer system, it may be appropriate to subdivide the source zone vertically. This can be achieved through multiple vertical release points at the source locations. This allows simulation of the different pathways in different aquifer layers (which may differ in hydraulic properties) and a depth-variant source term to account for the vertical distribution of the DNAPL material.

Pathline Analysis

MODPATH stores the pathline simulation information in a single data file. This file may be sorted and split by pathline so that each pathline may be operated on separately. The pathlines are input to MassTrack (section 2.2) and the output specifies which receptors the pathline outputs mass to by proportions and at what time. For each pathline-receptor pair, which may number more than the number of pathlines if there are weak sinks in the model, travel time and speed are calculated. Using Pythagoras' rule, the travel distance and speed as a function of travel time is calculated:

$$\begin{gathered} \delta s = \sqrt{\delta x^2 + \delta y^2 + \delta z^2} \\ v = \frac{\delta s}{\delta t} \end{gathered} \tag{2.8}$$

where s is path length, x, y and z are Cartesian spatial co-ordinates, v is speed and t is model time.

Source Term

The Source term defines the transient input flux to each pathline. The source term is defined for each source location as flux as a function of time. Each pathline is assigned a value for input flux according to the source term solved for the particle's release time

and this flux is redistributed along the path by advection, longitudinal dispersion and the other transport processes, using one of the flux algorithms described in section 2.3.1.

The BOS tool uses a surface-derived source term model, assuming that contaminant is leached from the surface by rainfall and transported through the vadose zone to the water table (Tait et al., 2004a). An alternative source term methodology more suitable for DNAPL source zones is presented in chapter 3. More simply, one can apply a constant unit source term in order to establish the attenuation of contaminant between source and receptor along the transient pathway.

Flux Algorithm

A flux algorithm is applied to the results of the pathline analysis, for each pathline-receptor pair, in terms of the travel time and speed. This specifies the input contaminant mass flux at the source and the dispersion of the aqueous contaminant. The total flux as a function of time will be sum of the transient fluxes from each pathline from which it abstracts some mass. Advective-dispersive transport along a one-dimensional pathline is expressed, in terms of flux by:

$$\frac{\partial J}{\partial t} = \alpha v \frac{\partial^2 J}{\partial s^2} - \frac{v}{R_f}\frac{\partial J}{\partial s} - \lambda J + S \tag{2.9}$$

where J is mass flux (units of mass per time), R_f is the retardation factor and λ is the first-order decay constant for reaction or biodegradation. The input flux is the source term and serves as a boundary condition to equation 2.9. R_f will always be 1 at this stage because retardation was accounted for in the MODPATH calculation by adjusting porosity.

The result of applying the flux algorithm to each captured pathline is the contaminant flux from source to receptor as a function of time, from each source, accounting for all the transient processes in the system. This method uses an analytical solution to calculate the transfer of flux from source to receptor, using the time-averaged velocity along each pathline as a constant velocity.

Analytical solutions for advective-dispersive contaminant transport only exist for constant-velocity scenarios together with a small number of other special situations (Gelhar and Collins, 1971). A pathline from source to receptor will tend to increase in velocity as it approaches the receptor if that receptor is a point abstraction. The velocity increase only occurs at a small distance from the abstraction and so the pathline velocity can be approximated by a constant velocity at total distance divided by total time. This assumption is questionable in aquifers with very variable or transient flow and the method could require further development in these situations.

This tool uses the one-dimensional Bear (1979) analytical solution to reactive contaminant transport in uniform flow from a constant (Heaviside unit function) source. Since each particle only represents a small amount of time, the source term for each pathline

should be a square pulse (constant source between two times and zero otherwise). This is equivalent to the difference between two Heaviside unit functions with equal magnitude and different start times. Since the Bear (1979) solution is a convolution of the source term and the differential equation (2.9), the analytical solution of contaminant transport from a square pulse source may be determined analogously, that is by the difference between two plumes with different release times which are the start and end times of the source pulse.

The flux J at the receptor at time t from a single captured pathline is accordingly:

$$J\{s,t\} = \frac{J_0}{2} e^{\frac{s}{2\alpha_s}\left(1-\sqrt{1+\frac{4\lambda\alpha_s}{v}}\right)} \operatorname{erfc} \frac{s - vt\left(1+\frac{4\lambda\alpha_s}{v}\right)}{2\sqrt{\alpha_s v t}} - \left(\text{if } t > T_{\text{pulse}}\text{: } \frac{J_0}{2} e^{\frac{s}{2\alpha_s}\left(1-\sqrt{1+\frac{4\lambda\alpha_s}{v}}\right)} \operatorname{erfc} \frac{s - v\left(t - T_{\text{pulse}}\right)\left(1+\frac{4\lambda\alpha_s}{v}\right)}{2\sqrt{\alpha_s v\left(t - T_{\text{pulse}}\right)}} \text{ else: } 0 \right) \tag{2.10}$$

where J_0 is the source term magnitude, α_s is longitudinal dispersivity (in units of length) and T_{pulse} is the amount of time (of the source term) represented by the particle. v is taken to be the mean velocity along the pathline (that is, $s/\Delta t$). The results of this method are demonstrated in section 2.6.

2.4 An Experiment for the Grid-Convergence of Particle Trace Solutions in Regions of Variable Flow

Section 2.3 introduced a semi-analytical solution for contaminant transport based on contaminant flux transferred along modelled pathlines. The dependence on pathlines to give the pathways, including duration, between sources and receptors means that the accuracy of pathline simulations is important. This section assesses the extent to which the trajectories of individual, non-dispersing particles may be relied upon and how that depends on the refinement of the flow model grid. If individual pathlines can be used with coarse-gridded flow models (such as regional groundwater resource models), then pathline calculations should be well grid-converged. The results also have implications for other risk models that non-dispersing pathlines to model source-to-receptor pathways such as Borehole Optimisation System (BOS) (Tait et al., 2004a), CatchRisk (Troldborg et al., 2008) and particle-based capture zone methodologies such as the UK SPZ methodology (Environment Agency, 2015).

Detailed models with fine grids are preferred for contaminant transport modelling, as these lead to greater stability and reduce numerical artifacts such as artificial dispersion. However, such models require large investments of time and money and it is useful

to have simpler models and tools to use for lower tiers of investigation in order to identify priorities for funding, as well as for simulations at larger scales, for which refined flow models may be impractical. Particle tracking offers a simpler transport algorithm and has been used as an integral part of a number of regional-scale risk assessment tools, such as BOS and CatchRisk. Particle trajectories are typically simpler to implement than numerical models and look realistic, if the neglect of dispersion can be justified. However, particle track solutions can be unduly deterministic, because a particle, representing a packet of water (perhaps originating at a contaminant source) can only either miss or be captured by a receptor; it cannot half reach a receptor. Therefore, it is important to know whether water really does go where computed particle trajectories would suggest. If not, risk assessment tools that rely on particle tracking ought to include some mechanism for showing the uncertainty. This uncertainty may be associated with the parameters given to the model. There may also be uncertainty due to numerical artifacts introduced by the particle-tracking technique, which is the subject of this experiment.

Particle tracking is a method of solving transport that considers only advection. Algorithms solve the locations of conceptual particles placed within groundwater by integrating the path through the velocity vector field. There are various methods implemented for this integration. The MOC transport solution available within MT3D (Zheng, 1990) uses either an explicit-style extrapolation or the fourth-order Runge-Kutta method which is more suitable for areas of variable flow because it captures more of the variation in the flow field. MODPATH (Pollock, 1994) uses an analytical equation within each finite difference cell (as the model is designed to work with MODFLOW results), taking a linear interpolation of the velocity vector between the cell faces. The difficulty arises in inferring a water velocity field from the MODFLOW results, which give the discharge rates through cell faces. MODPATH makes the heuristic assumption that the discharge is uniformly distributed across each cell face. This is probably the most practical assumption when interpolating for velocity and beyond this only refinement of the grid can give a more precise velocity field solution. This has repercussions for the application of particle tracking to coarse-gridded models.

This simple experiment is designed to test the accuracy of particle tracking in areas where the flow paths are converging or diverging. This is done by analysing a number of flow fields at different locations (for example, at flow divides or stagnation points, as well as more uniform regions) and testing the variation of the particle trajectories with changes in grid refinement. The hypothesis is that in regions where flow paths converge or diverge, coarse gridded flow models will lead to particle mis-direction errors with the particle-tracking algorithm implemented in MODPATH.

2.4.1 Method

Simple flow scenarios were conceptualised consisting of two components: uniform background water flow and a point abstraction. Flow equations for confined conditions were

used, so that superposition of the flow fields from the two components was exactly possible. Only steady-state horizontal flow is considered. The conclusions are generalisable to transient scenarios because the results are only dependant on the flow conditions at a point in time. With these two components, an analytical expression for Darcy flux at any point in x-y space may be deduced:

$$\begin{aligned} q_x &= q_w \cos\phi + q_{x,\mathrm{bg}} \\ q_y &= q_w \sin\phi + q_{y,\mathrm{bg}} \\ q_w &= \frac{Q}{2\pi b r} \\ r &= \sqrt{x^2 + y^2} \\ \phi &= \tan_2^{-1}(y, x) \end{aligned} \tag{2.11}$$

where q_x and q_y are the x and y components of Darcy velocity; $q_{x,\mathrm{bg}}$ and $q_{y,\mathrm{bg}}$ are the components of the uniform background flow; r and ϕ are polar co-ordinates; Q is the well's abstraction rate (technically injection rate, with negative indicating abstraction); b is the aquifer thickness. $\tan_2^{-1}$ is the generalised inverse tangent function which distinguishes the signs of both arguments so that the result is in the range $[-\pi, \pi]$.

With no background flow, all flow trajectories point directly to the well. Including a background flow causes the pathlines in some locations to diverge from each other, especially near the stagnation point. The models are designed to test the accuracy of MODPATH in flow systems with a spatially varying flow field.

The flow scenarios are investigated in sections by numerically simulating small square domains, extracted from the analytical solution by specifying fixed fluxes at the boundaries. Length scale is arbitrary, but for the sake of clarity, the square domains are set at 100 units length and 100 units thickness. The central abstraction is at $(x, y) = (0, 0)$ and abstracts at rate 1000 (although in fact the abstraction rate is also arbitrary because particle direction rather than velocity is of interest; what matters is the background flow rate relative to the abstraction rate).

The locations of the centre of the extracted sections are shown in figure 2.3.

Each square domain is tested with various grids: (columns×rows) 1×1 (that is, the grid cell is the size of the domain), 3×3, 3×9, 3×27, 9×3, 9×9, 27×3 and 27×27, excluding the edge cells in which the input and output fluxes are specified. Using odd numbers ensures that the well is equivalently positioned for each grid simulation with the central section — at the centre (in MODFLOW source and sink fluxes are applied to cell nodes).

The edge fluxes are calculated from the analytical equations, numerically integrated along the length of the cell to which they apply to give the analytically-derived mean flux along that boundary. The edge flux cells are separated using horizontal flow barriers of zero conductance using the HFB MODFLOW package (Harbaugh et al., 2000) so that the boundary fluxes are forced to apply to the central domain of the

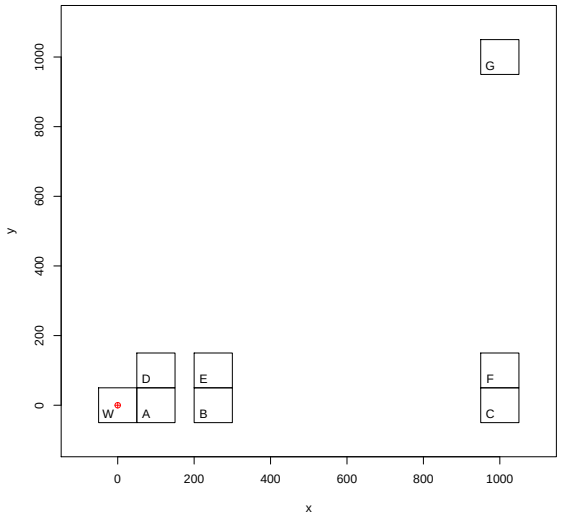

Figure 2.3: Locations of extracted submodels with labels used for identification. The well is within submodel "W".

model. One of the edge cells is assigned a fixed head rather than a fixed flux in order to constrain the simulation, although the value is arbitrary when using confined equations. Because the flows will balance, the flux resulting from the fixed head will be correct if the model is well converged. An arbitrary isotropic hydraulic conductivity of 10 is used. The set-up is demonstrated in figure 2.4. A fixed flux cell is also used for the cell containing the central abstraction in submodel W (see figure 2.3). Fixed flux conditions rather than fixed head conditions are used because fixed heads would apply at a finite difference from the cell faces (at the nodes of the boundary cells), whereas fixed fluxes are forced, with use of the HFB package, to apply at the cell faces.

The particles are released from the edges of the domain from 24 points — 7 on each side with shared corners. MODPATH does model each cell as a continuum, with linearly interpolated velocity, so particles will follow a different trajectory depending on the point at which they enter the cell. These trajectories are curved as the velocity varies linearly within the cell, although the particle locations are usually only saved at the cell boundaries, so when plotted they appear straight.

Eight sets of simulations were conducted, varying only the background velocity vector (because all parameter combinations could be reduced to a unique value of $\mathbf{q}_{\mathrm{bg}}/Q$). Q was set at -1000 for all simulations. The background flow vector is shown in table 2.1.

2.4.2 Results

Submodels C, F and G, the most distant from the well, are least sensitive to the grid refinement. This was expected, because far from the well the flow field is closest to

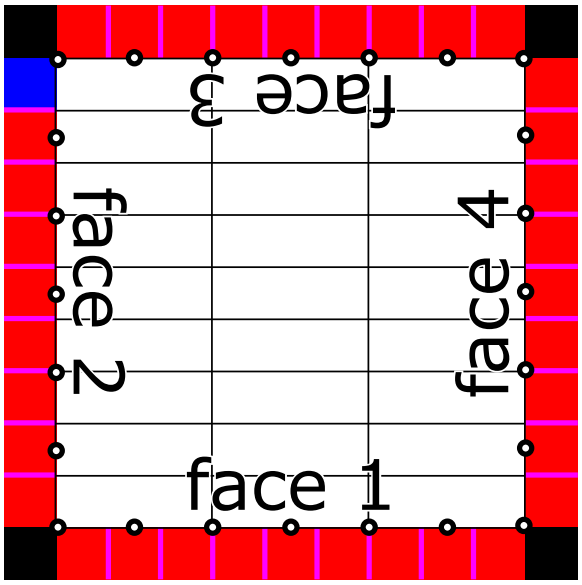

Figure 2.4: Generic grid-convergence submodel domain, in this case with 3 columns and 9 rows. Red cells represent fixed fluxes and the blue cell is a fixed head. Between each pair of edge cells is a zero-conductance flow barrier. The round white-filled points are the 24 particle release points.

Table 2.1: Background flow conditions for the different sets of simulations.

Simulation set	$q_{x,\mathrm{bg}}$	$q_{y,\mathrm{bg}}$
1	0	0
2	0.007	0
3	0.012	0
4	0.02	0.02

uniform. A single-celled domain calculates advective transport almost as effectively as a 27×27-celled grid, although some difference can be seen.

Closer in, performance is more grid-dependent, especially in areas with diverging or converging flow paths. In the purely convergent case, clearly all the particles from anywhere are captured, so the consequences of misdirection is not so severe, but it can be seen that in submodels W, A (figure 2.5) and D, trajectories are somewhat approximate with fewer cells, especially in the single-cell cases. Where there is a background flow and therefore a flow divide, imprecision becomes more important because particles may be tracked on the wrong side of the flow divide if the grid is insufficiently refined.

Plots are shown for submodels in which the discrepancy is seen most clearly. Figure 2.5 shows the poor representation of purely convergent flow. Figures 2.6, 2.7, 2.8 and 2.9 show inaccuracy representing a flow divide. In coarse-gridded flow models, the flow divides and stagnation points around wells may cross cells adjacent to the abstraction cells. The 1-by-1 cell models show that particle traces in these cases are in accurate and

will lead to incorrect capture zone delineations. Figure 2.10 also shows how particle tracking performs around a source or sink that is not dominating the flow around cells nearby. Figure 2.11 compares the performance of anisotropic refinement in different directions. This shows that pathline inaccuracy is greater for anisotropic cells, especially where elongate cells are oriented perpendicular to the direction of the well.

A key feature for judging the accuracy of a particle tracking simulation is whether particles exit the cell or submodel in the correct location. By extension, this will determine whether particles are reliably directed to the correct sink in the model. Travel times are also grid-dependent as a result of the grid-dependency of spatial routing, because inaccurate routing of pathlines brings particles too close or not close enough to stagnation points. Near stagnation points there are small areas with long residence times which could have significant bearing on travel times.

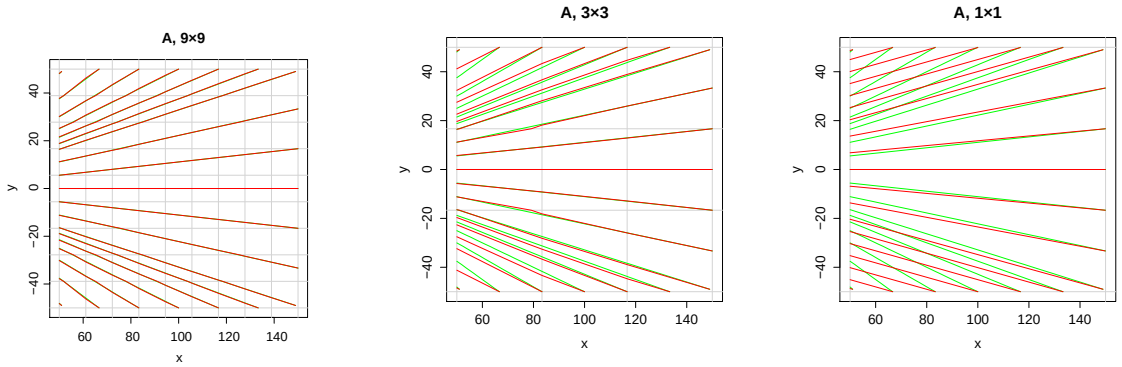

Figure 2.5: Simulation set 1 in block A, showing representation of purely convergent flow by various grids. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

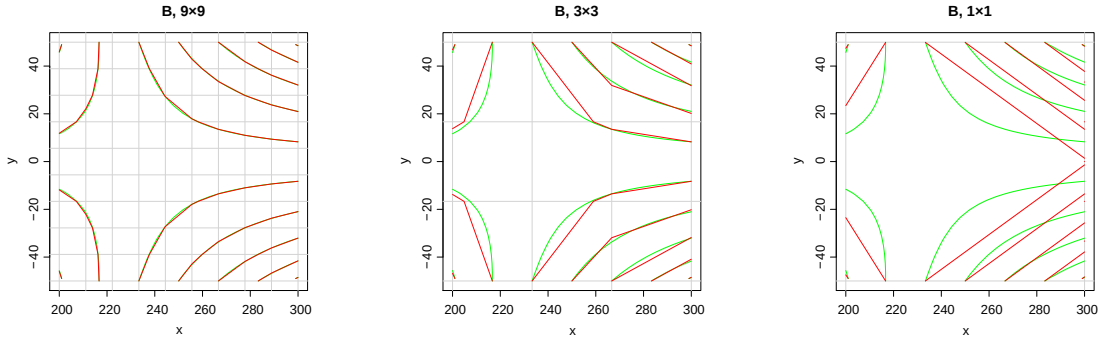

Figure 2.6: Simulation set 2 in block B, showing representation of a flow divide at the front of a capture zone by various grids. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

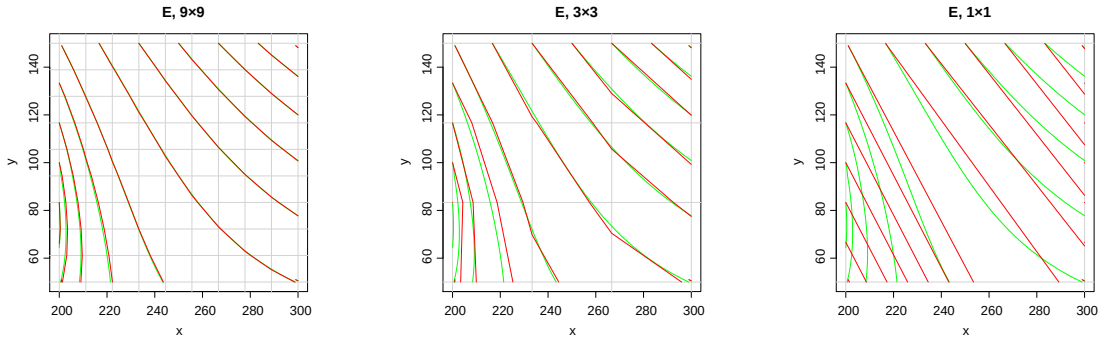

Figure 2.7: Simulation set 2 in block E, showing representation of flow near to a flow divide by various grids. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

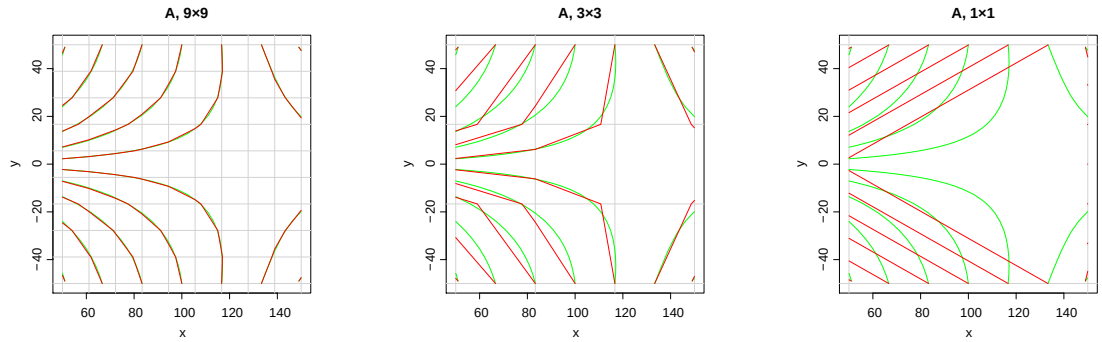

Figure 2.8: Simulation set 3 in block A, showing representation of flow near to a flow divide by various grids. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

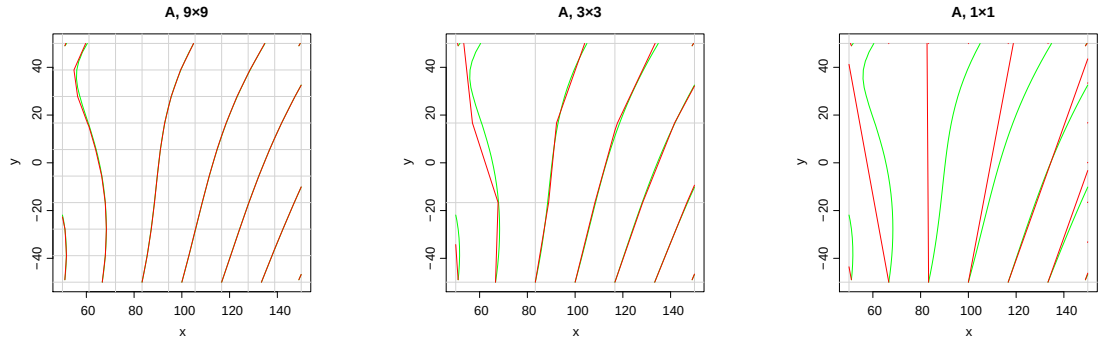

Figure 2.9: Simulation set 4 in block A, showing representation of a flow divide. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

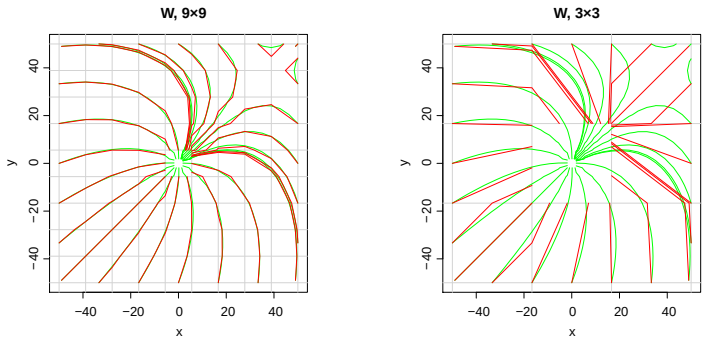

Figure 2.10: Simulation set 4 in block W (containing the central abstraction), showing representation of flow to the well by various grids. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

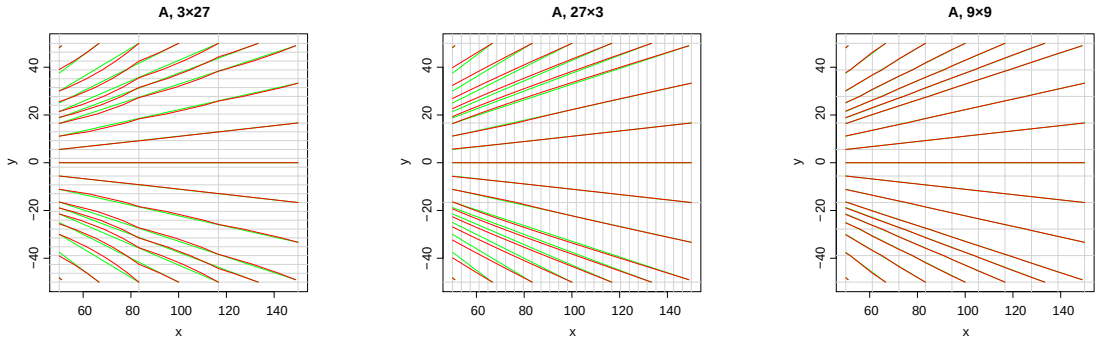

Figure 2.11: Simulation set 1 in block A, showing performance enhancement from anisotropic grid refinement in either direction compared to an equal isotropic refinement. Green lines are computed on a 27×27 grid and red lines are computed on the grid shown in each case.

2.4.3 Conclusions

The results from this experiment show that in regions of variability in groundwater flow direction, the particle trajectory results are dependent on the grid resolution used. A coarse grid in these regions is liable to misdirect the particles. This has implications for particle tracking-based risk tools, which may inaccurately delineate capture zones or wrongly identify priorities for remediation or treatment.

Although Eulerian methods such as the finite difference solution (Zheng, 1990; Black and Foley, 2013) may incur numerical dispersion because of complete intra-cell mixing, these methods do not suffer from misdirected advection. Dispersion, as well as being a physical process, has the advantage of effectively considering a continuum of trajectories, so that errors are not so severe. Trajectories that neglect dispersion (particularly transverse) allow for no exchange of contaminant between capture zones, which, asides from transverse dispersion being a real process, if the trajectories are misdirected then all contaminant flux from a source may be calculated to arrive at the wrong receptor. Numerical dispersion is also a problem and excessive dispersion can also direct contaminant mass down incorrect pathways. However, numerically dispersed results will tend to be too conservative with regards to spatial extent of contamination, rather than missing potential threats, although peak concentrations at the centres of plumes tend to be underestimated in numerically dispersed results[2]. Therefore neither Eulerian solutions nor particle-tracking solutions can give a reliably conservative estimate of contaminant risk with coarse-gridded flow models and it would always be possible that a risk has been underestimated or overlooked.

Section 2.5 introduces the DRW method, a Lagrangian solution to contaminant transport which avoids numerical dispersion but simulates both transverse and longitudinal dispersion. The DRW method is put forward in this chapter as the most suitable method for advective-dispersive contaminant transport with regional flow models. This method allows dispersion in all directions so uncertainty in capture zone boundaries is less of a problem. Although some particles in the DRW solution may be misdirected, the complete solution to the contaminant plume is spread over many particles, the majority of which will go to the correct receptor (to the extent that the flow model represents nature accurately).

2.5 The Dynamic Random Walk Method for Grid-Decoupled Simulation of Advective-Dispersive Contaminant Transport with Coarse-Gridded Flow Models

The process of dispersion is difficult to incorporate in models of contaminant transport. During model calculations, dispersion is often overestimated by solution methods

[2] These arguments assume that the solute is an unwanted contaminant.

due to the assumption of intra-cell mixing, which locally distorts the centre of mass of contaminant distributions, unless the cell size is very small, which is generally impractical at the regional scale. On the other hand, field observations suggest that the dispersivity of contaminants in groundwater is generally relatively small (compared to contaminant migration in free-standing water or air, for example), resulting in narrow plumes of contaminant (Rivett et al., 2006a). The technical difficulty of modelling contaminant dispersion in groundwater, as well as the uncertainty in estimating dispersivity values, means that some regional risk models use non-dispersing particle tracking to model the source-to-receptor pathway (Tait et al., 2004a; Troldborg et al., 2008).

However, the deterministic assignment of contaminant fate to the end of a one-dimensional pathline is often unjustifiable, because of uncertainties in the groundwater model, cumulative imprecision in the tracking of pathlines (section 2.4) as well as the physical spreading of contaminant mass by dispersion. As an example, the CatchRisk method is a simple pathline based method but acknowledges spatial uncertainty of particle trajectories and the possibility of dispersion by prescribing a five times uniform refinement of the flow model grid (which is not trivial) and the movement of sources just outside a well's capture zone to inside the capture zone (Troldborg et al., 2008). Therefore, especially for systems that have multiple receptors, it is necessary to simulate the dispersion process both to respect the fact that it physically occurs and as an acknowledgement of spatial uncertainty of contaminant distribution.

The dynamic random walk (DRW) method presented in this section simulates all contaminant processes, including dispersion, decoupled from the flow model grid. This avoids problems introduced by grid constraints, such as unstable oscillations and numerical dispersion. The MT3D solution methods (finite difference, MOC and TVD) allow decoupling from the flow model discretisation in the time dimension, but not the spatial dimensions, so contaminant mass is stored as an array with values at the cell centres only, which implicitly assumes complete intra-cell mixing. MOC uses particles to track contaminant advection but reverts to the finite difference method for the dispersion process; the particle concentrations after solving advection are homogeneously mixed into their destination cells before solving the other terms. Pure Lagrangian methods describe contaminant mass distribution as a distribution of massed points (particles) with real number co-ordinates that may be at any point in the model domain, not just cell centres. This removes the need to assume intra-cell mixing at any stage of calculation and therefore centre of mass can be preserved locally, which avoids numerical dispersion. Although Lagrangian methods are more commonly selected for very detailed, complex scenarios, it is argued in this chapter that a Lagrangian method may be a useful option for large scale coarse-gridded simulations.

The random walk particle method (RWPM) (Salamon et al., 2006) is a pure Lagrangian method, and the method described herein is an adaptation of this method designed for coarse gridded-regional models and also designed to require as few as possible solver-related (as opposed to hydrogeological) parameters as input. Many of the MT3D parameters (such as the Courant number and convergence criteria) are associated with Eulerian methods, so these are avoided. Further reduction of input options is achieved

by an automated dynamic organisation within the algorithm, so that the particle count can grow to describe extending plumes but is managed by coalescing of close-by particles to achieve an optimal particle count for the state of the solution at a given time. The solver-related options given by users and their effects are described in section 2.5.3. The method is called the dynamic random walk (DRW) method, as transport is simulated using a random walk and particles are managed by dynamic organisation. The implementation of DRW in the `DRW` package is included in the supplementary electronic material and the package documentation (appendix D.1; Barry, 2017c).

2.5.1 Algorithm

As the algorithm is described, a series of figures will illustrate how each step works. The flow model used is similar to the model used for section 2.2.3 and is a very simple fictitious model with a river acting as a weak sink. An intermediate time step is used for these figures, to show how the various processes operate on a collection of particles. The starting state for this time step is shown in figure 2.12

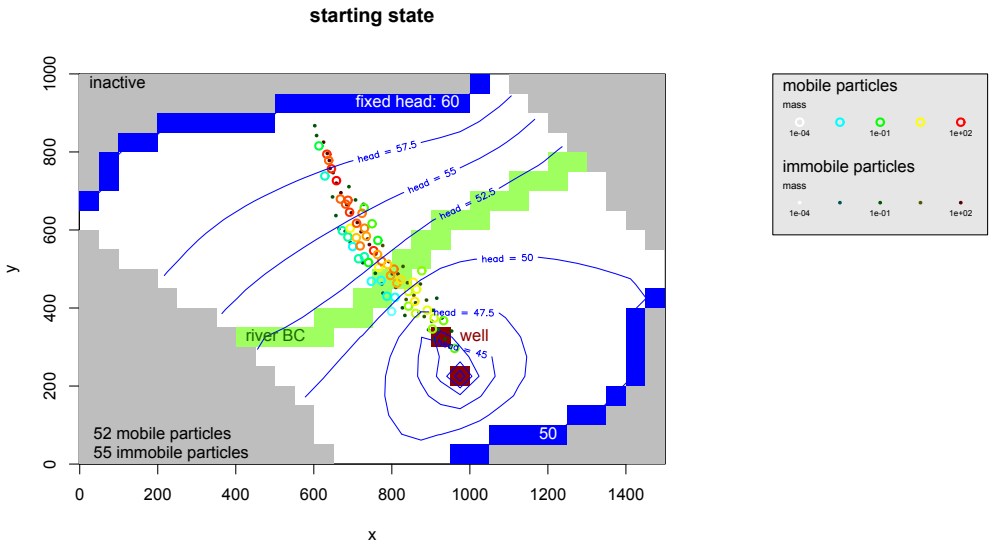

Figure 2.12: DRW algorithm demonstration: starting state (particle distribution and masses) for an intermediate time step.

Sources

The source process, which initiates particles, S_+, is defined as

$$\begin{aligned} \mathbf{r}_{p_0,n} &= S_{+,p}\left\{\mathbf{r}\right\} = \delta\left(\mathbf{r} - \mathbf{r}_{p_0}\right) \\ m_{p_0,n} &= S_{+,m}\left\{n\right\} = \int_{t_n}^{t_{n+1}} J_{p_0}\left\{t\right\}\,\mathrm{d}t \end{aligned} \tag{2.12}$$

$\delta\left(\mathbf{r} - \mathbf{r}_{p,0}\right)$ represents the delta function, reflecting that particles are infinitesimal points containing mass. $\mathbf{r}_{p_0}$ is the specified source locations and $J_{p_0}\left\{t\right\}$ is the list of corresponding transient source fluxes for each particle start location p_0. m_{p_0} is the mass of each particle p at source. n is used to denote time step number.

Source terms are specified by location (x, y, z) and transient flux. `R` (R Core Team, 2014) is a functional programming language and is capable of handling functions like other objects, allowing for intuitive and simple user input. A custom-made source term function, associated with a release location, may therefore be fed to the algorithm. Only mass loading-type sources are supported currently, but all source types can be converted to a mass loading source with prior calculations. The release locations are not constrained to be at the centres of grid cells.

For each time step, the source term functions are integrated across the time period of the time step, giving the mass that is injected into the system at each source location. These masses are then assigned to particles released at the appropriate locations, which are subsequently advected (section 2.5.1).

The collection of existing particles $1, ..., p, ..., N_p$ is then updated to include the collection of new particles for the time step. The implementation of the source term is shown in figure 2.13.

Sorption

Sorption, if included, is specified with a retardation factor, defined as:

$$R_f = 1 + \frac{K_d \rho_\mathrm{s}}{\phi} \tag{2.13}$$

where ρ_s is the dry bulk density of the sediment, ϕ is total porosity and K_d is the equilibrium distribution coefficient, describing the partitioning of contaminant between sediment and water (Fetter, 1999). K_d is a function of the organic carbon content of the aquifer material and the hydrophobicity of the dissolved contaminant (a more detailed discussion is in section 4.1.2. The validity of K_d (and by extension R_f) as parameters describing sorption behaviour depends on the validity of the assumption that sorption and desorption are in equilibrium and that the relationship between sorbed and aqueous concentration is linear. The DRW method can only simulate linear sorption because it tracks mass rather than concentration and the current implementation of DRW only supports equilibrium sorption, specified by R_f.

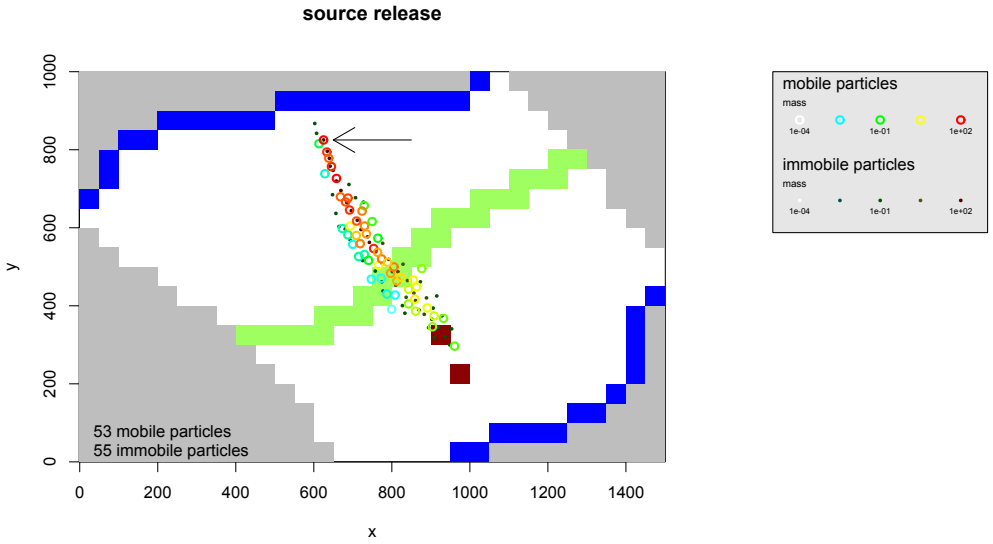

Figure 2.13: DRW algorithm demonstration: source release, adding particles (one in this case) at specified locations with calculated starting masses. The source location is indicated with an arrow.

DRW simulates sorption using a separate distribution of immobile particles which are excluded from the drift and dispersion processes. At the start of each time step, after new particles have been allocated at the source locations, the mobile and immobile particles exchange mass with each other. The equilibrium sorption process, R_{sorb}, is defined as:

$$\Delta m_{\text{sorb}} = R_{\text{sorb}}\left\{m, m_{\text{I}}\right\} = \left(1 - 1/R_f\right)\left(m - m_{\text{I}}\right) \tag{2.14}$$

The mass exchange between the particles is implemented as:

1. split each immobile particle p_{I} of mass $m_{\text{I},0}$ into a mobile particle with mass $m = \frac{m_{\text{I},0}}{R_f}$ and an immobile particle with mass $m_{\text{I}} = m_{\text{I},0}\left(1 - \frac{1}{R_f}\right)$
2. split each mobile particle p of mass m_0 into a mobile particle with mass $m = \frac{m_0}{R_f}$ and an immobile particle with mass $m_{\text{I}} = m_0\left(1 - \frac{1}{R_f}\right)$

(2.15)

This is shown in figure 2.14.

In this way new immobile particles act as sources creating new mobile particles out of their own mass and mobile particles similarly act as sources of immobile mass. The mobile to immobile ratio of mass both locally and globally is always $\frac{1}{R_f - 1}$ after the sorption calculation.

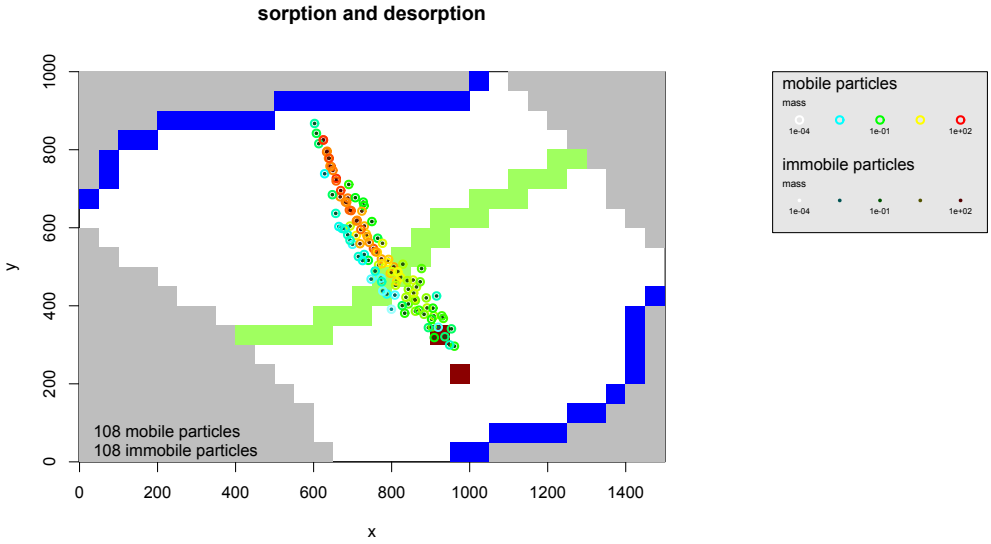

Figure 2.14: DRW algorithm demonstration: sorption, simulated by splitting particles into mobile and immobile mass according to the retardation factor.

Drift

Drift is the expected displacement of a particle. Intuitively it would seem to be the same as the advection, with the dispersion being a probabilistic displacement away from the expected advective transport. However where the dispersion coefficient is heterogeneous, this is only approximately true, because the mean movement of mass will be diverted towards regions of higher dispersion coefficient (LaBolle et al., 1996, 2000; Salamon et al., 2006). Therefore, the drift vector is represented as:

$$\left.\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}\right|_{\text{drift}} = \mathbf{v} + \nabla \cdot \mathbf{D} \tag{2.16}$$

For detailed investigation of groundwater flow in heterogeneous geology, the heterogeneity in dispersion coefficient ($\nabla \cdot \mathbf{D}$) is an important consideration, and neglecting it will result in too much model contaminant mass going into zones of low dispersion coefficient. However, for regional-scale studies where heterogeneity is not represented in detail, it is a reasonable approximation to take the advection term alone as the drift. Thus DRW represents drift as:

$$\text{Drift}\left\{\mathbf{r}_{p,n}, \mathbf{v}, t\right\} = \left(\Delta\mathbf{r}_{p,n}\right)_{\text{drift}} = \int_{t_n}^{t_{n+1}} \mathbf{v}\,\mathrm{d}t \tag{2.17}$$

The drift (assumed to be purely advection) is implemented by DRW using MODPATH 5 (Pollock, 1994). The mobile particle locations are used to specify release locations for

particles in MODPATH. The particles are given a start time according to the start time of the time step within DRW and MODPATH is instructed to run the time period of the DRW time step only. The pathline output from MODPATH is loaded and analysed, with the final position of the particles giving the new locations of the mobile particles, as shown in figure 2.15..

MODPATH integrates particle trajectories through the linearly interpolated velocity vector between each cell, which is the most accurate estimate possible given the grid, without supplying extra information that is not intrinsic to the MODFLOW results. MODPATH assumes that the specific discharge is uniform through each cell interface and for single particles and this can lead to trajectory misdirections in regions of velocity variability (section 2.4), but the effect is much less severe when tracking a distribution of a large number of particles.

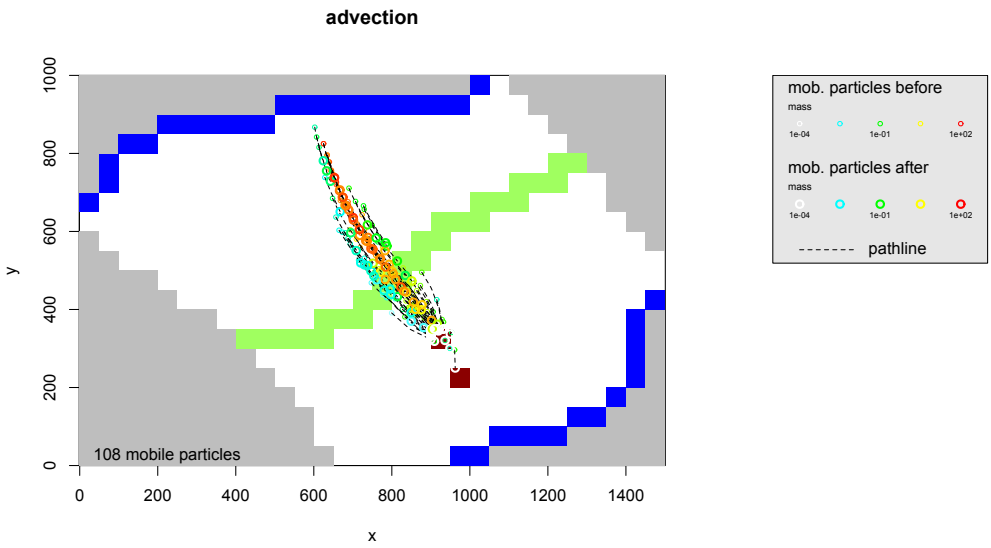

Figure 2.15: DRW algorithm demonstration: drift, approximated as advection and simulated by MODPATH 5.

Sinks

Mass loss from particles due to weak or strong groundwater abstractions are determined by comparing the MODPATH-generated particle trajectories with the flux outputs from MODFLOW using MassTrack (section 2.2; Barry, 2017d). The mass lost from particles as determined by the MassTrack algorithm is used for the sink term in DRW. The loss of mass from particles in cells with a net sink is shown in figure 2.16.

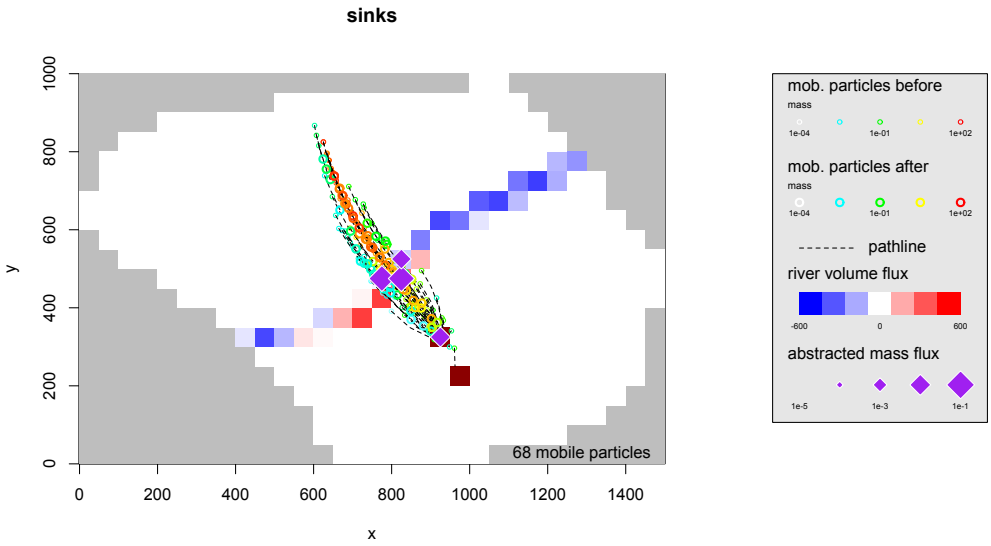

Figure 2.16: DRW algorithm demonstration: mass loss to sinks, calculated by MassTrack. The discharge rate to the river is plotted with a colour flood to indicate which cells are acting as sinks (in addition to the wells). Only wells with a negative flux from the river (that is, the flux is to the river) remove mass from the passing particles.

Reactions

This section discusses the degradation reaction process; the sorption process is discussed in section 2.5.1. DRW can only simulate first-order reaction processes because the model state is expressed in mass rather than concentration, so concentration is not known during calculation. First-order degradation is expressed as:

$$\left.\frac{\mathrm{d}m}{\mathrm{d}t}\right|_{\text{decay}} = -\lambda m \tag{2.18}$$

where λ is the decay constant. This is represented by:

$$\mathrm{R}\left\{m_{p,n}, t\right\} = -\left(\Delta m_{p,n}\right)_{\text{decay}} = m_{p,n}\left(e^{-\lambda \Delta t_{n:n+1}} - 1\right) \tag{2.19}$$

DRW implements first-order degradation within the MassTrack package (section 2.2, Barry, 2017d) because in this way the appropriate amount of degradation is calculated to occur along a pathline before it reaches a sink. Mass lost due to degradation is saved to the appropriate time step, but in the current version of DRW, is not further tracked as a reaction product as can be done in MT3DMS (Zheng and Wang, 1999).

There is not a figure to show the effect of the reaction calculation; the effect is simply a proportionate reduction in mass of each particle and mass lost in this way is saved in the DRW result by time step.

Dispersion

The dispersion process may be represented by a 2- or 3-dimensional probability density field, specifying the probalistic future location of a particle. For two dimensions, this is expressed as:

$$\Delta \mathbf{r}_{\text{disp}} = F\left\{\xi_1, \xi_2\right\} \tag{2.20}$$

where ξ_d is a random number in the range $[-1, 1]$ corresponding to the dimension d and F is the probability distribution function that governs the likelihood of the particle displacing to a given set of co-ordinates relative to its original position. The current version of DRW uses a Fickian dispersion model, in which dispersion occurs according to the diffusion equation with dispersion tensor $\mathbf{D}$.

In two dimensions, this derives as follows. The dimensions are used for the longitudinal and transverse dimensions x_L and x_T, relative to an individual particle's trajectory. Dimensionless displacements in the longitudinal and transverse directions are defined such that the problem becomes isotropic in dimensionless co-ordinates:

$$\begin{aligned} x'_L &= \frac{x_L}{2\sqrt{D_L \Delta t}} \\ x'_T &= \frac{x_T}{2\sqrt{D_T \Delta t}} \end{aligned} \tag{2.21}$$

The dimensionless ${r'}^2$ magnitude and direction relative to longitudinal θ_L of the displacement is then given as:

$$\begin{aligned} {r'}^2 &= {x'_L}^2 + {x'_T}^2 \\ \theta_L &= \tan_2^{-1}\left(x'_T, x'_L\right) \end{aligned} \tag{2.22}$$

The dispersion probability distribution function F is then defined as a function of two random numbers in the range $[-1, 1]$ to give r' and θ_L. The inverse error function of ξ_1 is taken to define a normally distributed displacement (such that extreme displacements, as ξ_1 approaches -1 or 1, are increasingly improbable) and ξ_2 is multiplied by π to give a uniformly distributed probability of selecting any angle:

$$\begin{pmatrix} r' \\ \theta_L \end{pmatrix} = F'\left\{\xi_1, \xi_2\right\} = \begin{pmatrix} \text{erf}^{-1}\,\xi_1 \\ \xi_2 \pi \end{pmatrix} \tag{2.23}$$

F' is the two-dimensional dispersion probability distribution function of equation 2.20, except in dimensionless space. Converting back to real co-ordinates, the displacement is:

$$
\begin{aligned}
x_L &= 2\sqrt{2D_L\Delta t}\, x'_L = 2\sqrt{2D_L\Delta t}\, r' \cos\theta_L \\
x_T &= 2\sqrt{2D_T\Delta t}\, x'_T = 2\sqrt{2D_T\Delta t}\, r' \sin\theta_L \\
\Delta x_{\text{disp}} &= 2\sqrt{2D_L\Delta t} r' \cos\left(\theta_L + \theta_{\text{traj}}\right) \\
\Delta y_{\text{disp}} &= 2\sqrt{2D_T\Delta t} r' \sin\left(\theta_L + \theta_{\text{traj}}\right)
\end{aligned}
\tag{2.24}
$$

where θ_{traj} is the trajectory of the particle's path during the advection step.

Dispersion may be specified as velocity-dependent, in which case longitudinal and transverse dispersivities in units of length are specified, which for each dispersing particle is multiplied by the average velocity in the time step to get dispersion coefficient in units of length squared per time.

Dispersed particles may move outside the active region of the model or outside the domain of the model. Contaminant mass lost in this way is saved as such by time step and may be analysed later to see whether a significant amount of mass has been lost in this way. In some contexts, this may be more physically appropriate than a no-flux boundary condition as assumed by MT3D.

The dispersion of particles is shown in figure 2.17.

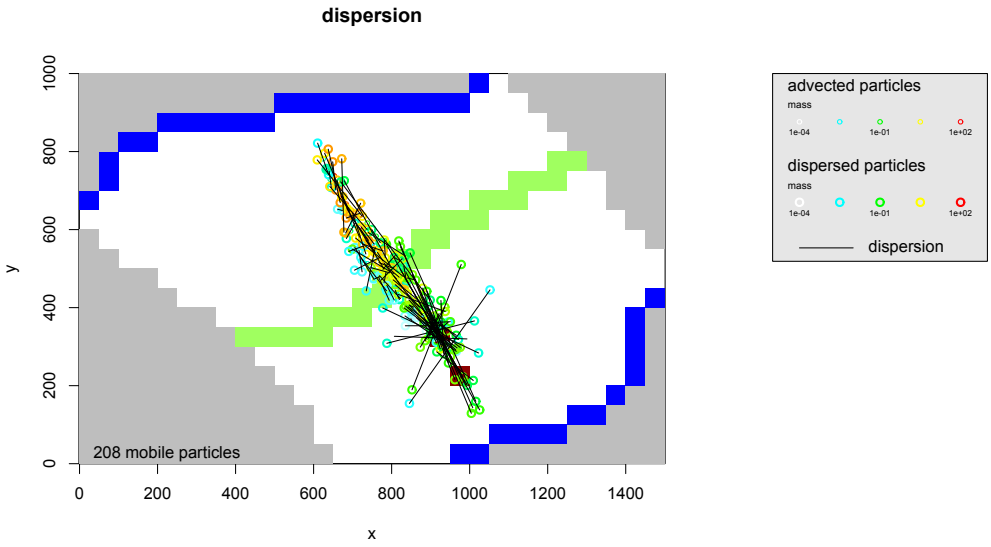

Figure 2.17: DRW algorithm demonstration: dispersion, represented by a splitting and probabilistic displacement of particles from their drifted location.

Dynamic organisation by coalescing

The dispersion algorithm for DRW spawns a fixed number of particles for each pre-existing particle. Unchecked, this would lead to an uncontrolled multiplication in

the particle count, which would clearly become unmanagable after a small number of time steps. However, many of these particles will be quite close to each other and therefore there is potential to aggregate particles without significant compromise on precision. On the other hand, particles that are dispersed to the peripheries of a contaminant plume describe new areas of contamination. Therefore the particle count ought to grow slowly with each time step to reflect the progressive pervasion of solute mass. DRW dynamically organises its active particles by coalescing particles that are very close together. With each grouping, the total mass and the centre of mass is preserved. In this way a roughly evenly spaced distribution of particles is used to describe the concentration field, with higher-mass particles in the centre of plumes and lighter particles at the fringes. Coalescing with a large search radius can lead to blotchy and under-dispersed results, but a reasonable search radius does not adversely effect the results. Sensitivity of results to the search radius is discussed in section 2.5.3, which gives guidance in the sensible choice of coalescing search radius.

The coalescing algorithm used by DRW (Barry, 2017a) is described in full in appendix A and its effects can be seen in figure 2.18.

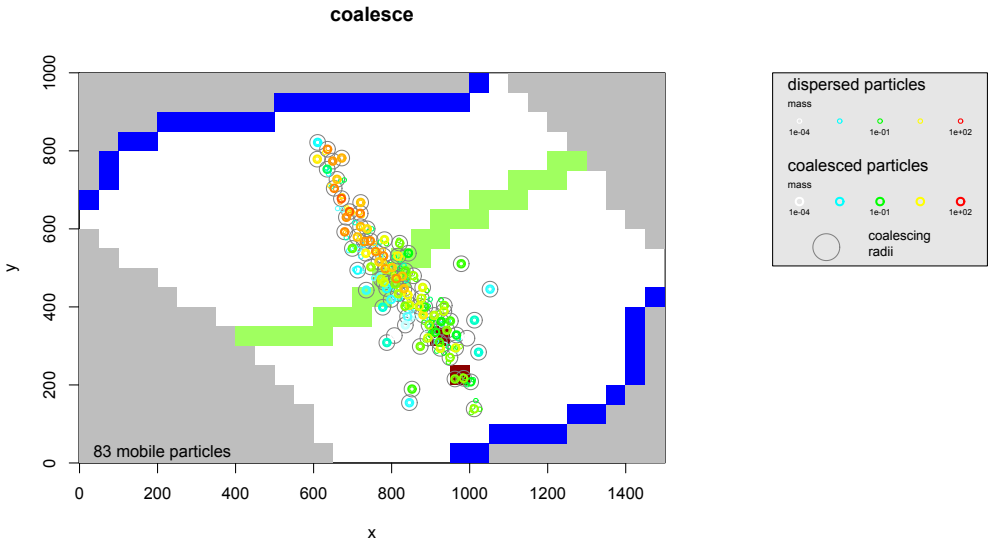

(a)

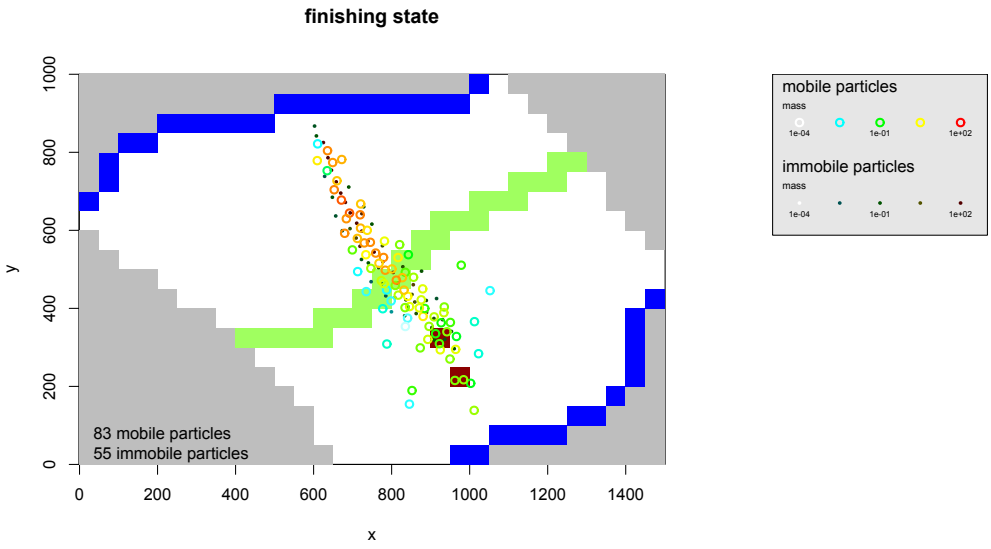

(b)

Figure 2.18: DRW algorithm demonstration: (a) particle coalescence, in which the particle distribution is simplified by grouping nearby particles in order to dynamically moderate particle count and (b) the final state of the model after this time step, which is then used as the starting state for the next time step.

Assumptions

In heterogeneous flow, the mean expected motion (sometimes termed the drift vector) of a dispersing particle is not exactly the same as advective water velocity. Rather, the true drift vector differs from the water velocity by a shift towards zones of higher dispersion coefficient (LaBolle et al., 1996, 2000; Salamon et al., 2006). For the application to the regional scale, however, where detailed heterogeneity is not represented in the groundwater flow model it is thought sufficient to approximate the drift vector as the advective velocity. However, where particles move close to the edge of the active domain of the model, this approximation will result in some overestimate of the amount of solute mass lost to the inactive portion of the model, if the inactive region represents aquitard rock. Hydrodynamic dispersion should be smaller in the aquitard on account of slower water velocity, so in fact less contaminant mass would disperse into these regions than DRW calculates. Away from model inactive model regions, this approximation is, largely, not an issue.

2.5.2 Implementation

DRW is scripted using the `R` programming language, making particular use of the packages `data.table`, `RNetCDF` and `Rflow`, (R Core Team, 2014; Dowle et al., 2015; Michna and Woods, 2016; Barry, 2017e,c). The model state is described as a table (or two tables if immobile particles are used), with a record for each particle giving its location and its mass. This is in contrast with Eulerian schemes that describe the model state as an array of concentration values representing each model node. The `data.table` package enables fast, memory-efficient operations on large data sets by reference.

The user defines the starting locations of the particles, which need not be the at the centres of cells. The source term functions are also given, each of which must be a function of time and depth-dependent source terms may be specified as multiple vertically separated releases. This format is well suited to the output of the DNAPL source term model framework developed in chapter 3, but it is general enough to describe any point source description. The start and end times of the model are given, as well as a time step. Only constant time steps are supported currently, but there is an option to use a previously simulated plume as the starting condition for a simulation, which would allow a period of model time to be simulated with a shorter time step for greater temporal refinement.

2.5.3 Validation

Model Description

A benchmarking model has been designed to allow a direct comparison between the results of DRW, MOC and an analytical solution. The benchmarking model has been

set up to be simple (solvable by an analytical equation) but refined, with small enough grid spacing to ensure that finite difference MT3D does not cause artificial dispersion. So asides from artifacts, MT3D, DRW and an analytical transport equation should all produce the same result. The groundwater model is a single-layer steady state model with a rectangular domain and is outlined in figure 2.19 and table 2.2, solved using MODFLOW 2000 (Harbaugh et al., 2000).

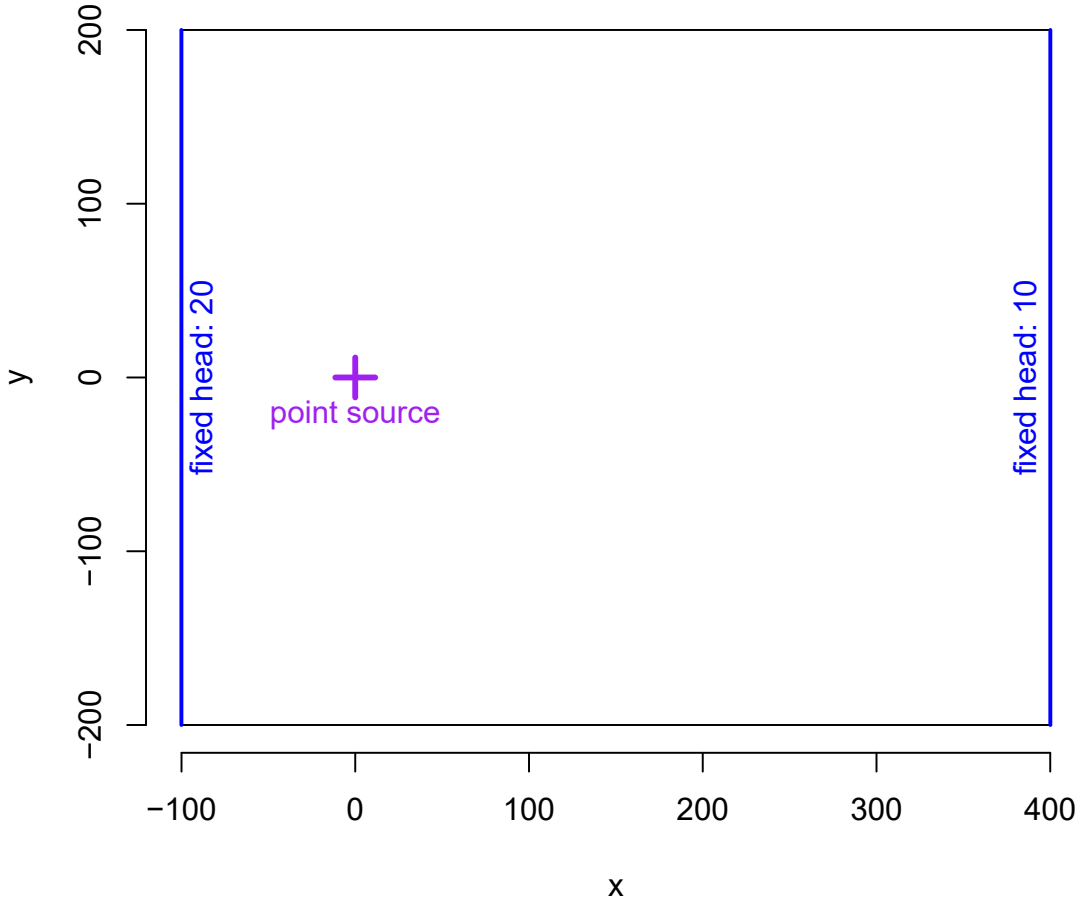

Figure 2.19: Set up of the benchmarking model, indicating the fixed head boundaries. The MODFLOW model is solved as a confined model and so the hydraulic gradient is uniform, resulting in a uniform Darcy velocity of 0.02. More details about the model set up are given in table 2.2.

A mass loading source is applied at the position $(0, 0)$, which is half-way between the two horizontal edges and closer to the upstream vertical boundary, allowing for transverse and backward dispersion (figure 2.19, table 2.2). For the MT3D solution, the Method of Characteristics (MOC) solution was used. The DRW solution was used with a constant time step length of 50 and a horizontal coalescing search radius of

5. Vertical dispersion was not modelled. The analytical solution used the Domenico (1987) solution for uniform flow fields for two dimensions:

$$C\{x,y,t\} = \frac{C_0}{4} e^{\frac{x}{2\alpha_x}\left(1-\sqrt{1+\frac{4\lambda\alpha_x}{v_x}}\right)} \operatorname{erfc} \frac{x - v_x t\left(1+\frac{4\lambda\alpha_x}{v_x}\right)}{2\sqrt{\alpha_x v t}} \times \left(\operatorname{erf} \frac{y+\frac{Y}{2}}{2\sqrt{\alpha_y x}} - \operatorname{erf} \frac{y-\frac{Y}{2}}{2\sqrt{\alpha_y x}}\right) \tag{2.25}$$

Parameters are as for equation 2.10 with, additionally, Y being the source width and α_y being the horizontal transverse dispersivity. Longitudinal and transverse dispersivities were 10 and 1 respectively for all simulations, and no reaction or sorption processes were included.

A table of parameters for the simulation set and plots of the results of this experiment are in appendix B, including plume time slices, transverse concentration sections, breakthrough curves and particle counts. Results for the DRW solution are compared directly with the analytical solution and the MOC solution.

Solution parameters

Two drawbacks of random walk-based methods are the computational time and the random variability in the result. These two features are in opposition to each other as one way to reduce the random variation in the result is to increase the particle count, which clearly requires more computational effort. With DRW, the user does not have direct control over particle count. Instead, the user chooses the value of certain solver parameters which act to manage the particle count: these are N_{dp}, which is the number of dispersed particle pairs produced for each prior particle during the dispersion process, and r_{co}, which is the coalescing search radius used for coalescing at the end of each time step. The user can also set the maximum number of particles, which is useful for ensuring that a simulation will complete in a reasonable amount of time, but has the danger of artificially reducing dispersion as it forces coalescing of the lightest particles that exceed the maximum, which are often the particles at the periphery of a plume.

Increasing N_{dp} has the effect of increasing the growth rate in the number of particles at the early stages of the simulation, but decreasing r_{co} has a much greater effect on the final number of particles used in the simulation (figures 2.20a and 2.20b). Coalescence should be regarded as a necessary evil: clearly the particle count cannot be left to grow exponentially with each time step, but an excessively large value of r_{co} could reduce the effect of dispersion by grouping particles at the fringe of the plume and exacerbate random variability (or blotchiness) by grouping many particles into single particles. The effect of increasing N_{dp} gives a moderate precision improvement for a moderate run time cost. Decreasing r_{co} has a much greater run time cost; if r_{co} is too large then the result will be inaccurate and blotchy, but once it is below a sensible value (which

can often by estimated intuitively), there is little benefit to further reductions. In this example model, reducing the coalescing radius from 5 to 2 doubles the run time, but with little discernible difference in the solutions. An increase to 10 introduces slightly more random variability. Increasing r_{co} further to 50 obliterates dispersion and gives extreme oscillations. However, this error is obvious and there is no danger that such results would be accepted. The effects of varying N_{dp} are comparatively minor, in terms of both result variability and particle count. The results for varying r_{co} and N_{dp} are shown in figure 2.20 and more detailed results are in figures B.4 and B.16–B.20 in appendix B.

The effects of the grid cell size may be seen in comparing the plots in figures B.4 (most cells 2×2), B.21 (most cells 20×20) and B.26 (most cells 100×100). The DRW solution is completely grid-independent in this experiment, with no differences between these three results apart from random variations. In more complex situations where there is variability in the flow field, flow models have some imprecision with coarse grids which will impact upon the DRW result. But there is no grid-dependence for the DRW solution inherent to the method itself. MOC performs well with moderately coarse grids, but the effects of the grid are seen with very large cells as oscillations. Additionally, because the DRW result is not a cell-by-cell array of values, the fine detail of the plume may be captured even if the grid is coarse, as many particles may be located in a single cell.

Physical processes

The results show that the DRW is very capable of simulating linear reactive transport processes, provided that good solver parameters are chosen, as discussed above. For example, smaller values of dispersivity require a smaller coalescing radius. The results from the DRW method are compared with equivalent analytical solutions, using equation 2.25. The mean arrival time is correctly matched to the analytical solution for different water flow speeds (modified by altering the porosity value). Results for the simplest case, with no reaction and moderate dispersivity, are shown in figure B.4, in comparison with the Domenico (1987) analytical solution and MT3D (MOC).

Dispersion can be correctly reproduced in both the longitudinal and transverse directions, as long as r_{co} is reasonable. Back-dispersion occurs for larger values of dispersivity, when the dispersion step is likely to send particles further back than the advective displacement. In this example, back-dispersion occurs when the longitudinal dispersivity is 10, but not 1, with porosity 0.1. Back-dispersion is shown in both the MOC and DRW solutions, but more so in the MOC solution.

Equilibrium sorption (retardation) and first-order degration results and variation with porosity are well represented by DRW and match the analytical and MOC breakthrough curves closely (figures B.6c–B.11c).

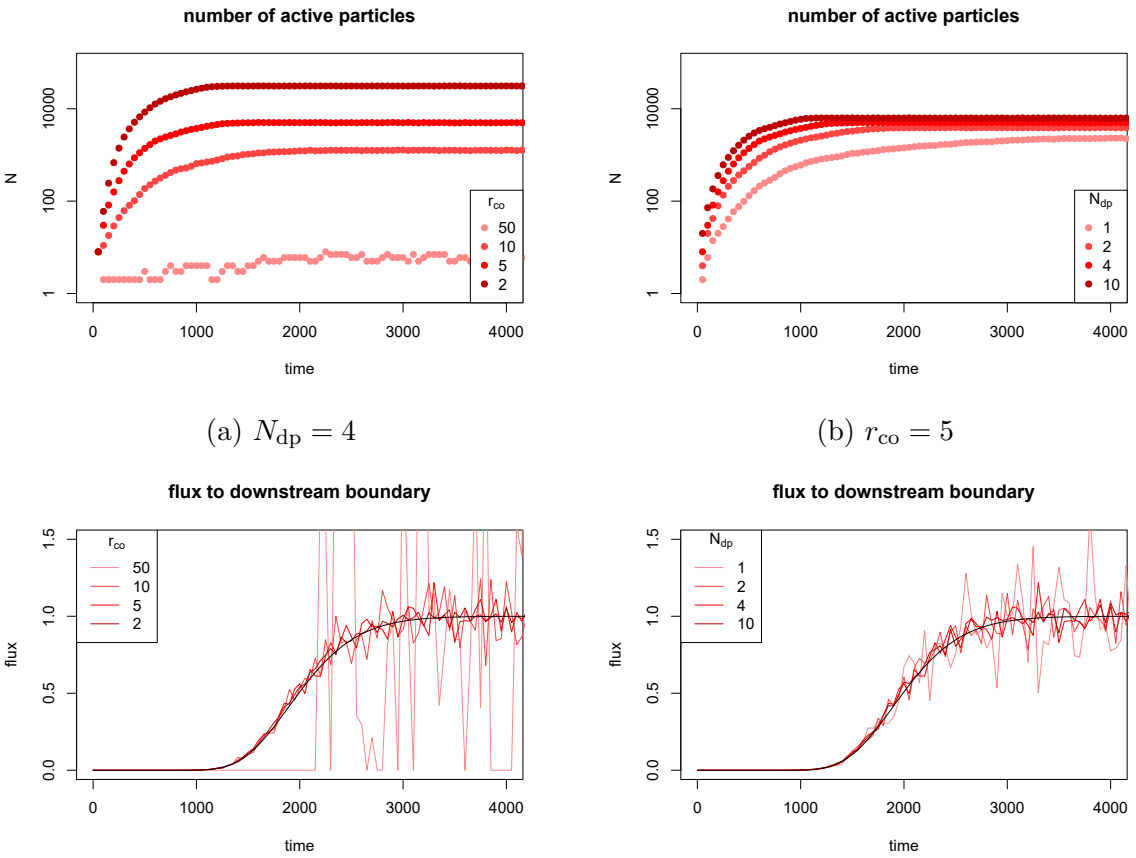

(a) $N_{\mathrm{dp}} = 4$

(b) $r_{\mathrm{co}} = 5$

(c) $N_{\mathrm{dp}} = 4$

(d) $r_{\mathrm{co}} = 5$

Figure 2.20: Effects of (a,c) varying r_{co} and (b,d) varying N_{dp} on (a,b) the particle count and (c,d) the random variability of the DRW solution result. The black lines on the breakthrough show the analytical solution to integrated downstream flux (equation 2.25).

Conclusion

The validation of the DRW method shows that it is able to produce accurate results for advective-dispersive transport and is not influenced by grid refinement. Advection, dispersion and reaction are all reproduced with similar results to the analytical and MOC solutions and accuracy is not compromised by a coarser grid. Random oscillations from DRW solutions are characteristic of all random walk methods, but are not unstable and may be reduced by adjusting solver parameters r_{co} and N_{dp} or smoothed as a post-processing step.

The two main solver-related options, r_{co} and N_{dp} are fairly intuitive in their meaning, but it can be difficult to judge *a priori* what the optimum values should be. Comparison of results in figure 2.20 suggests that the value of N_{dp} is not critical, with a slight increase in random variability for lower N_{dp}. An excessively large value for r_{co} will obliterate dispersion and give an inaccurate and blotchy result. However, r_{co} need not be tiny and the accuracy of the solution is converged for small and moderate values of

r_{co}.

The results from the benchmarking model show that DRW may be reliably used as a contaminant transport model. It is able to accurately model advection, dispersion and first-order reaction processes without overly conservative solver parameters which would be excessively costly of computer memory and run time. Section 2.6 demonstrates how the DRW method may be effectively applied to the problem of regional contaminant transport from point sources.

Table 2.2: Set up and properties for the benchmarking model. Default values (def.) are those used for most tests, but some tests will vary parameters from the defaults. Note that the units of the model are arbitrary, so no units are given.

parameter	value
Darcy velocity	uniform 0.02, to the right
dimensions	$500.1 \times 400 \times 10$
model origin co-ordinates	$(-100.05, -200)$
time	1 steady-state stress period of duration $10,000$
finite difference grid	column spacings (def.): $0.1, 0.95, 2, 2, ..., 2, 2, 0.95, 0.1$ row spacings (def.): $1, 2, 2, ..., 2, 2, 1$ layer elevations: 0 (base); 10 (top) Note that the edge columns are very thin and contain the fixed head boundaries and that, asides from these, the edge rows and columns are half thickness so that the source release $(0, 0)$ is in the centre of a cell. Other finite difference grids tested have 20 and 100 as the main cell dimension instead of 2.
effective porosity, ϕ_e	$0.03, 0.1$ (def.), 0.3
longitudinal dispersivity, α_x	$1, 3, 10$ (def.), 30
transverse dispersivity, α_y	$\alpha_x/10$
retardation factor, R_f	1 (def.), $2, 5$
first-order degradation half life	∞ (def.), $2000, 500$
flow boundaries	left: fixed head 20 right: fixed head 10 front, back, top and bottom: no flow
solute flux boundaries	point source release with constant flux 1 at $(0, 0)$ fixed head at right acts as a mass sink other three edges are no flux for MOC, but not a boundary for the analytical solution or DRW; this discrepancy only becomes significant in this case with larger transverse dispersivity

2.6 Transient Contaminant Risk Model from Industrial Point Sources to Abstraction Wells in the Birmingham Aquifer

The regional transport methodologies discussed in this section are demonstrated with a regional model: the EA's regional model of the Birmingham aquifer (Daily and Buss, 2013). The model represents the period from 1850 to 2012. Spatially the model uses 200×200 m grid cells (uniform) covering a 20×29 km domain, with three layers of variable thickness representing the Mercia Mudstone, the Bromsgrove Sandstone and collectively the Wildmoor and Kidderminster Sandstones. Further discussion and evaluation of this groundwater model is section 4.2.

These simulations are carried out to test the efficacy of the transport methodologies described in this chapter (sections 2.5 and 2.3) and to compare them with existing numerical methods, as implemented in MT3D. The Birmingham Aquifer regional model is a fully transient model, so only methods allowing for transient flow fields are applicable.

2.6.1 Historical contamination of the Birmingham Aquifer

The city of Birmingham has a long industrial history, reaching a peak in the mid-twentieth century. Many industrial activities are likely to have been associated with the use of chlorinated solvents and observed solvent contamination in groundwater is widespread, although particularly focussed in a small number of regions (Rivett et al., 2012). A fuller description of the solvent-related industrial history of Birmingham is in section 4.1.2.

2.6.2 Model parameters

For this study, mapped historical industries (figure 2.21) associated with metals and engineering works are considered to be potential point sources of chlorinated solvents, even if there is no site-specific evidence to suggest that there was solvent usage and spillage at the site (which, in fact, is the case for most historic industrial sites). This is a model of relative contamination at different locations and through time. It is not an attempt to reproduce actual concentration values of a particular substance, which is studied in chapter 4 and requires detailed quantitative consideration of the source term (chapter 3). The source fluxes, therefore, are conceptual, not associated with any real substance, and are therefore assigned units of d^{-1}, indicating that they represent a rate of arrival of an undefined substance. Flux from the point sources is constantly released, starting in 1925, in layers 2 and 3 (in the vertical midpoints) at a rate that is proportional to the site's area, normalised by the horizontal area of grid cells ($40,000$ m^2), reflecting the coarse assumption that a larger industrial site is likely to produce more waste. The conceptual concentration to receptor, which is

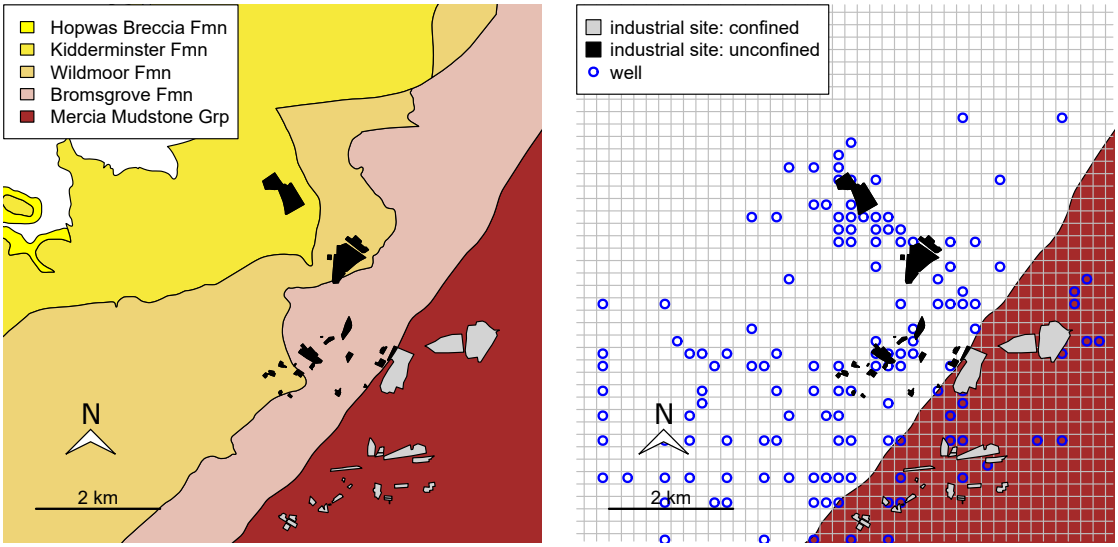

(a) (b)

Figure 2.21: Map of historic industries associated with metals and engineering works over the Birmingham aquifer. (a) shows the geological base map and (b) shows the model grid and well boundaries for the regional model, with grey areas representing inactive areas of the model (Daily and Buss, 2013). The well locations shown in (b) are snapped to the centres of model cells and no not represent precise locations. They include historic as well as present-day abstractions. The historic industries spatial dataset was compiled using the EDINA Historic Digimap service, as detailed in appendix C.

the conceptual flux diluted by the receptor's abstraction rate (in units of $m^3\ d^{-1}$), has units of m^{-3}. The scale of conceptual concentration is arbitrary, with the relative values between wells being significant. The releases are directly below the centroids of the industrial site footprints and for MT3D these releases are snapped to the nearest cell centres, aggregated where multiple releases are assigned to the same cell.

Moderate hydrodynamic dispersion is specified, with 200 m and 20 m for the longitudinal and transverse dispersion coefficients respectively. 0.264 is used for the effective porosity, uniformly applied (Allen et al., 1997; Daly, 2005).

The transport calculation is attempted using MT3D (Zheng and Wang, 1999), the PFT method (section 2.3) and the DRW method (section 2.5). The DRW solution uses a time step of 200 days and coalescing search radii of 50 m horizontally and 20 m vertically to moderate particle count (section 2.5.1). This value for r_{co} does not reduce the effects of dispersion. Section 2.5.3, in particular figure 2.20c, shows that apart from extreme values, the solution is well converged with respect to r_{co}.

2.6.3 Results and Discussion

This discussion considers the spatial distribution and the transient result. The spatial discussion looks at the distribution of the result, in particular whether all wells in

a probably highly contaminated area have similar values and whether wells outside the probable contaminated areas show significant modelled conceptual concentration results. The transient discussion considers the transient concentration response for a set of individual wells that represent wells in different positions relative to the potential contaminant sources and the flow direction.

Spatial analysis

Figures 2.23, 2.24 and 2.25 show the distributions of conceptual concentration predicted by the different transport methods. All models predict, unsurprisingly, that active abstractions in the vicinity of historic industrial sites are abstract greater concentration than those far removed from industrial areas. Contamination risk to wells in the confined part of the aquifer, even near the fault, is minimal, especially for the non-MT3D solutions.

The three MT3D methods (finite difference, MOC and TVD) all show significant oscillations, demonstrated by negative results (figure 2.23), even with upstream weighting, an option that has the effect of increasing dispersion in order to stabilise a simulation. Clearly, these results are unsatisfactory and with the large cell size it is unsurprising. In some cases, the wells that are of interest have a negative value for conceptual concentration. The MT3D solutions also exhibit greater dispersion, as evidenced by a low contrast in conceptual concentration across wells covering a moderate distance. The wells north-west(upstream) of the historic industrial areas show evidence of back-dispersion (figure 2.24c). DRW also exhibits some back-dispersion here, though less than MOC (figure 2.24b) and the PFT method is completely free from back-dispersion (figure 2.24a).

The PFT model neglects dispersion of the path of contaminants. Although the flux is dispersed longitudinally in the transient flux calculation, the dispersion is only in the time dimension and not spatial. Therefore, there is very little distribution of conceptual concentration between wells that are in close vicinity: the stronger or most upstream of the wells will generally capture the particles and the result for conceptual concentration at nearby wells will be zero (compare figures 2.24 a and b).

Transient

Figure 2.22 shows the locations of seven wells, labelled A to G, that are focussed on in this discussion. These wells are chosen for the features of the numerical methods that they highlight, not for any significance from a management perspective. They are mostly historical wells.w

Wells A and E (figures 2.26a and 2.26e) are close downstream from a largest industrial site, with A more directly downstream (figure 2.22). As would be expected, A abstracts greater concentration than E according to all simulations. However, there could be finite concentration at E if dispersion, especially transverse, is accounted for. The PFT

method calculates almost no concentration at well E because most pathlines starting at the nearby source are captured by well A. The DRW method disperses particles outside A's capture zone and so, while it simulates a lower concentration for E, there is a significant contaminant threat for E even with the protection of A. Figure 2.23 shows the different abstracted concentrations for A and E in 1988 clearly. The MT3D methods show artificial oscillations in that area.

Well B (figure 2.26b) is in the confined aquifer. As the Birmingham Fault acts as an imperfect hydraulic barrier (as it is represented in the flow model), there should be minimal contaminant reaching wells in the confined aquifer, except due to dispersion across the Birmingham Fault. None of these wells are shown as being at risk by the PFT method, which neglects spatial dispersion, but well B registers a result from the DRW method. Although it is possible that dispersion actually could cause contaminant mass to cross the fault, it is likely that the DRW result for well B is partly due to the approximation that the drift vector is identical to the water velocity (section 2.5.1). Dispersion across the fault is much more significant for MT3D, as can be seen in figure 2.23.

Wells C and D (figures 2.26c and 2.26d) are away from the centres of the contaminant plumes (especially C) coming from the currently mapped distribution of industries and as such the simulated conceptual concentrations at these wells are relatively low. Nonetheless, MT3D calculates a finite concentration for both due to numerical dispersion. The Finite Difference method with upstream weighting calculates greater concentration for both of these than without upstream weighting, indicating that it is numerical dispersion that is taking conceptual flux to these wells. Similarly well G (figure 2.26g) is upstream of the industrial area and only captures conceptual flux when methods produce artificial back-dispersion. DRW can produce back-dispersion (section 2.5.3) and shows a small response in wells C and G (compared to MT3D results) but it is impossible for the PFT method. In this study, back-dispersion is shown by all the MT3D methods and in particular by the Finite Difference method with upstream weighting.

Well F (figure 2.26f) operates in the heart of the historical industrial area and has maintained a relatively consistent and strong abstraction rate through the most recent time periods of the Birmingham groundwater model (about 1990–2012). The PFT and DRW methods have good agreement for conceptual concentration results at well F. The MT3D methods are similar to each other at well F, starting from a higher conceptual concentration value initially, gradually approaching the value of the other methods, and all appear to settle on an approximate steady-state concentration value. By the end of the model, well F is the only active well in its vicinity and so will be the only major receptor for the surrounding sources. Because all methods use the same input flux, the fact that all transport methods tend to a similar conceptual concentration result at well F indicates that all the methods are correctly balancing contaminant mass.

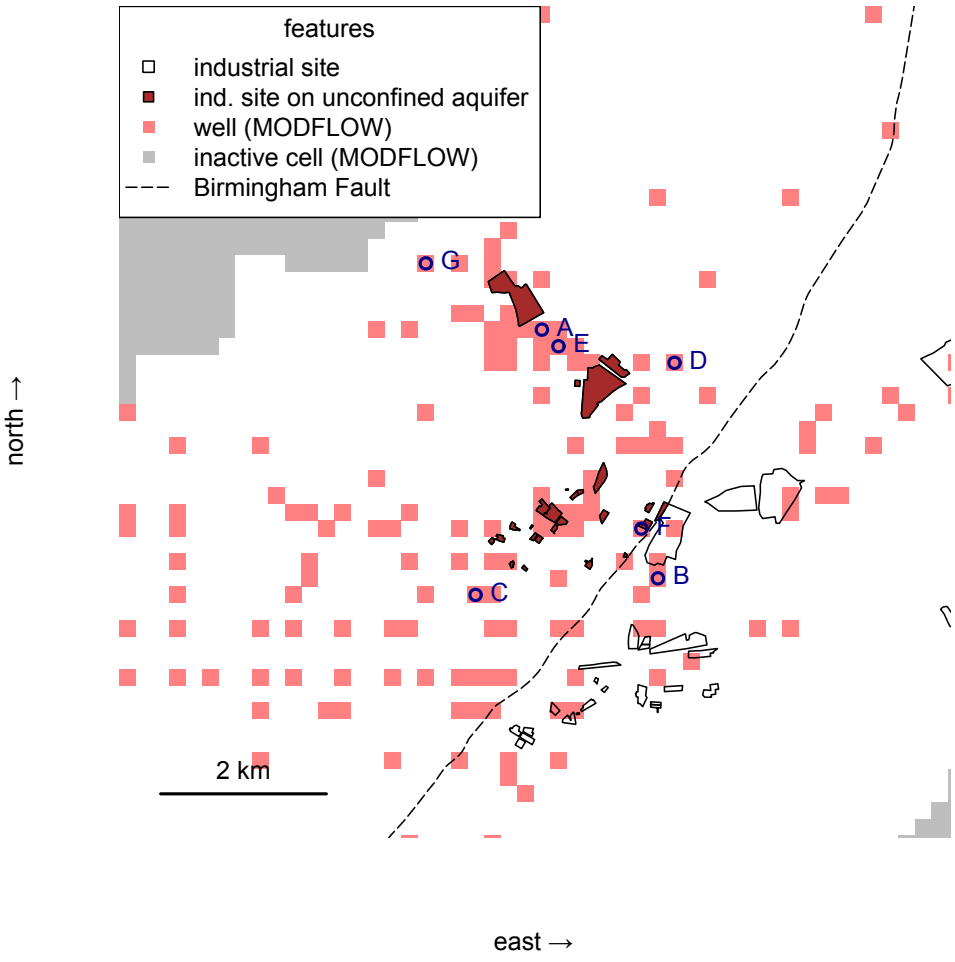

Figure 2.22: Overview map of the Birmingham Aquifer groundwater model, showing the model domain and well boundaries, the Birmingham Fault (British Geological Survey, 2013), the footprint of historical industry and seven wells (A–G) that are focussed on in the discussion of section 2.6.3. These wells are chosen as they distinguish model capabilities, rather than for any significant management concerns. The Birmingham Fault separates the confined aquifer (to the south-east) from the unconfined aquifer.

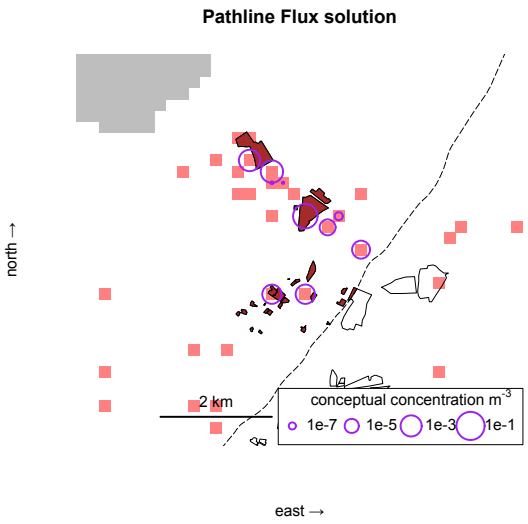

(a)

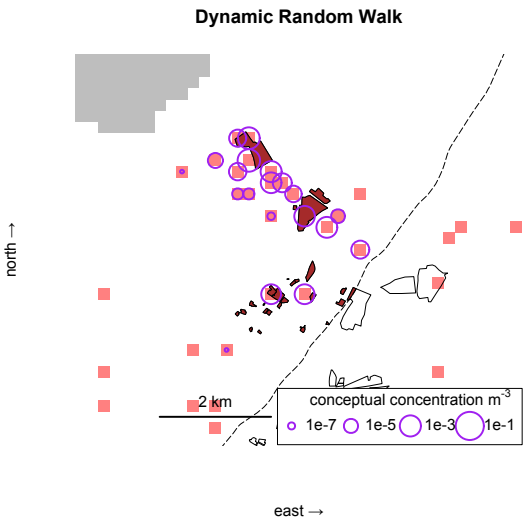

(b)

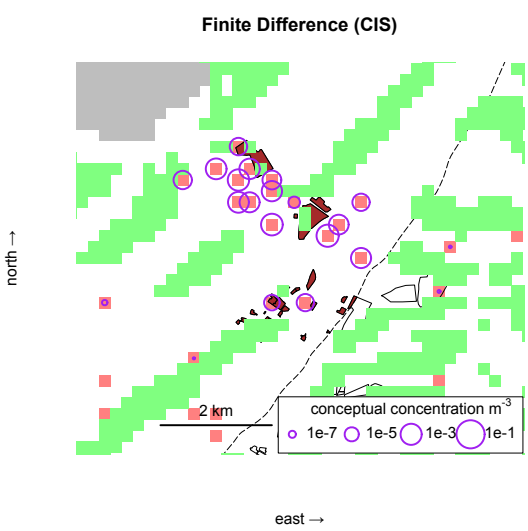

(c)

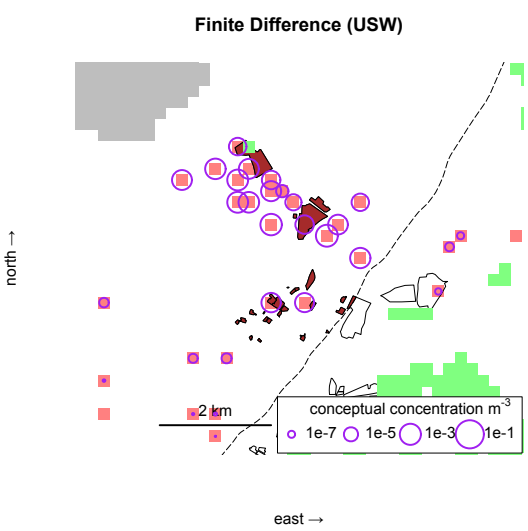

(continued on next page)

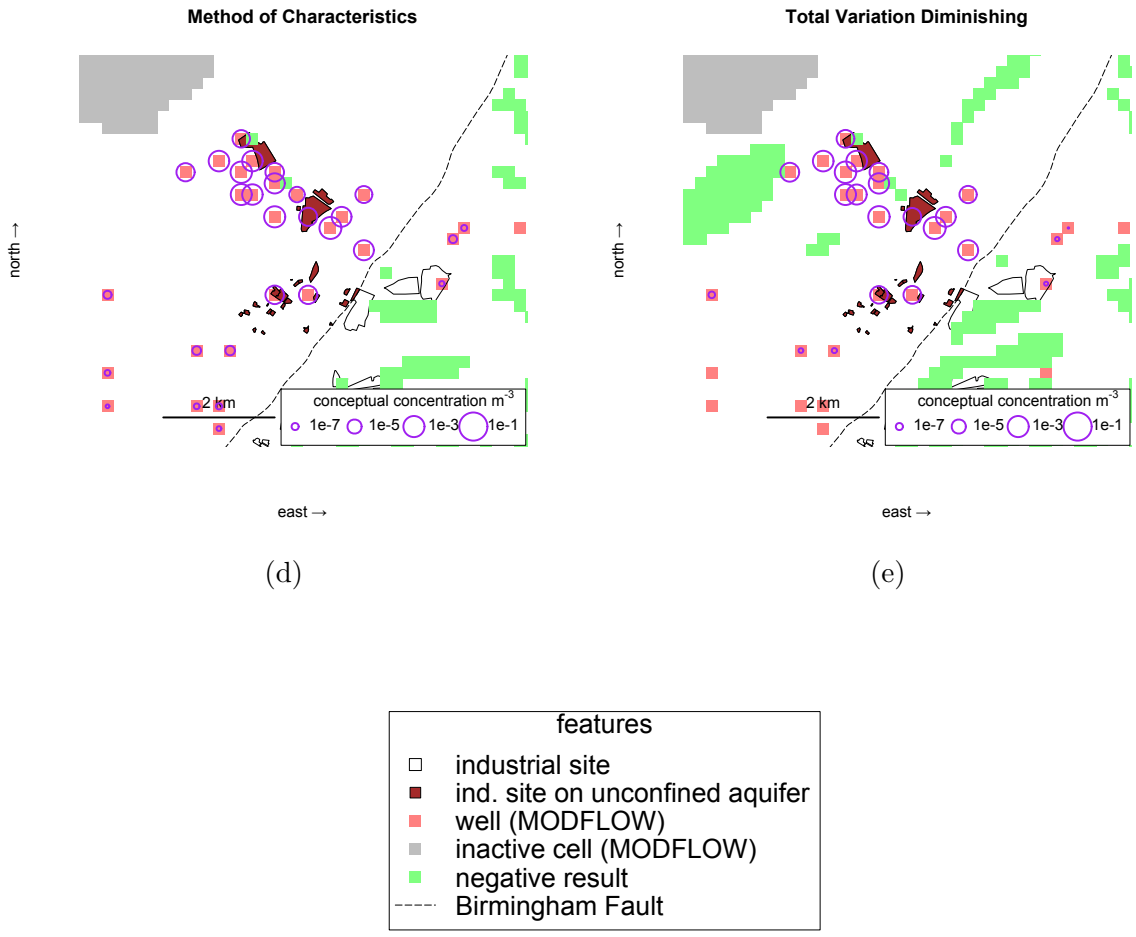

(d) (e)

Figure 2.23: Comparison of methods for transport to active abstractions in the Birmingham aquifer; a time slice of 1988. Conceptual concentration is a relative metric in this example, with particles carrying conceptual flux being released under each industrial site on the unconfined aquifer, weighted proportional to the site area. Resulting conceptual flux to receptor wells, diluted by abstraction rate, is plotted as purple circles. CIS and USW mean central in space and upstream weighted, referring to options in the Finite Difference advection package. Negative results are plotted for MT3D packages to indicate regions of artificial oscillations.

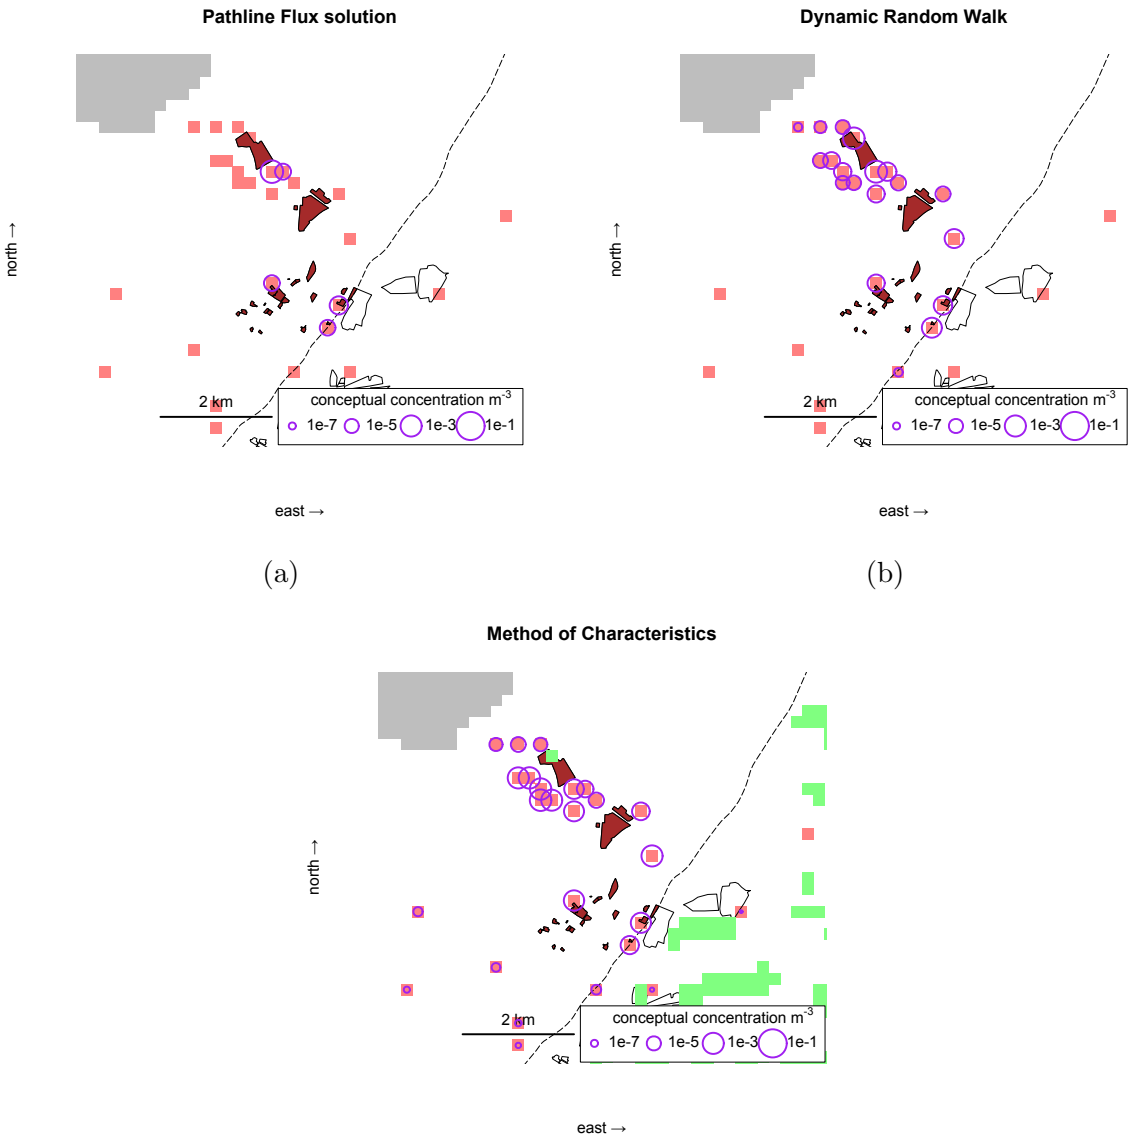

Figure 2.24: Comparison of methods for transport to active abstractions in the Birmingham aquifer; a time slice of 1998. See figure 2.23 for further description and legend.

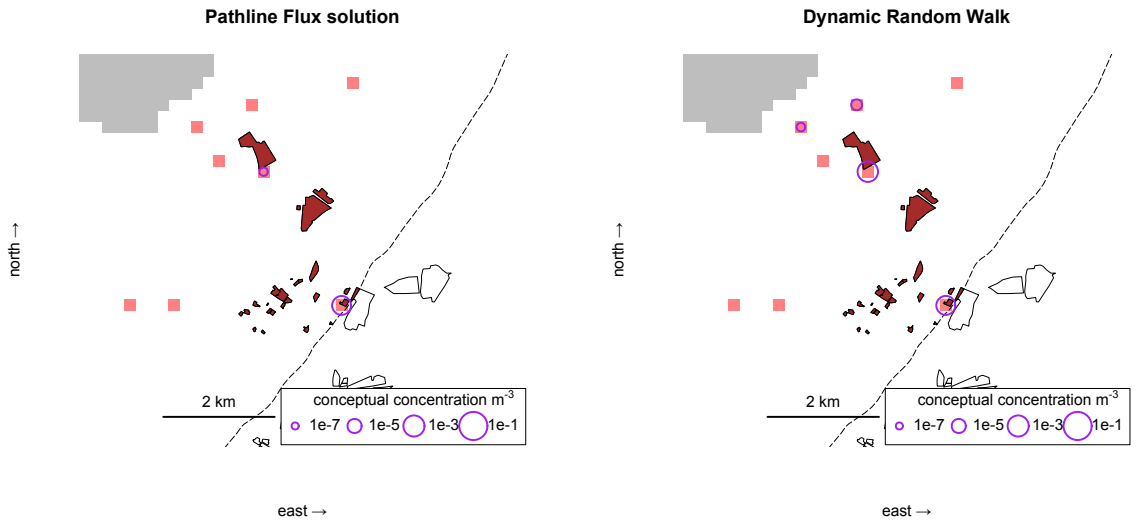

(a) (b)

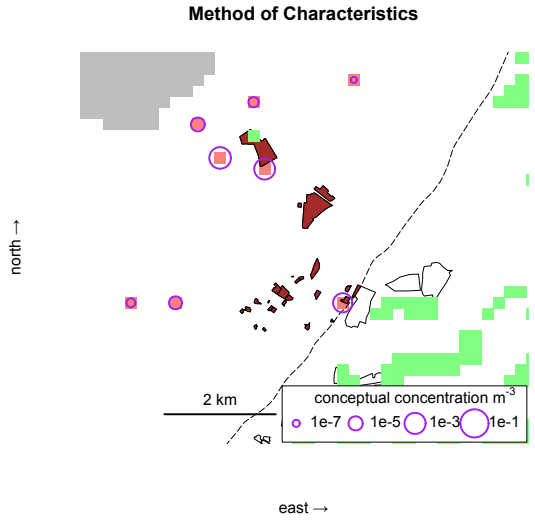

(c)

Figure 2.25: Comparison of methods for transport to active abstractions in the Birmingham aquifer; a time slice of 2008. See figure 2.23 for further description and legend.

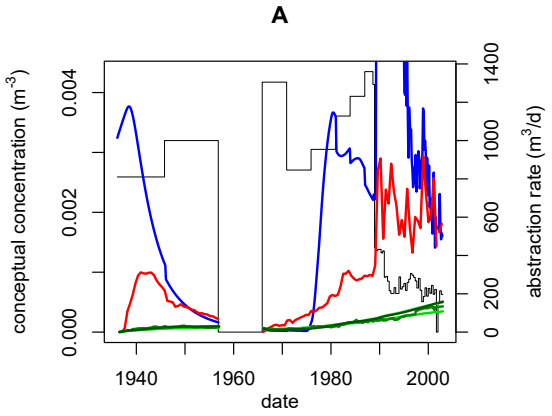
A
conceptual concentration (m-3)
0.000
0.002
0.004
abstraction rate (m3/d)
0
200
600
1000
1400
1940
1960
1980
2000
date

(a)

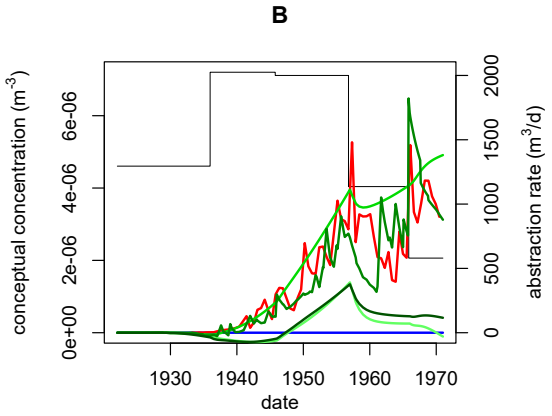
B
conceptual concentration (m-3)
0e+00
2e-06
4e-06
6e-06
abstraction rate (m3/d)
0
500
1000
1500
2000
1930
1940
1950
1960
1970
date

(b)

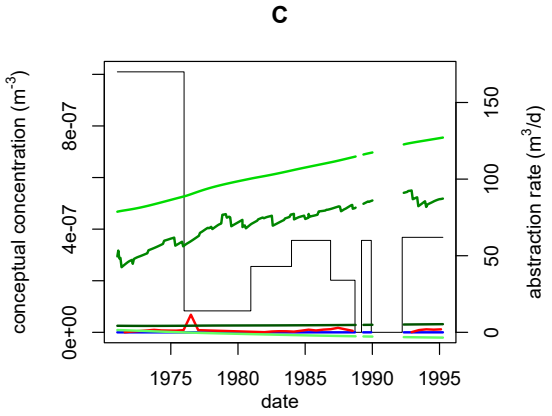
C
conceptual concentration (m-3)
0e+00
4e-07
8e-07
abstraction rate (m3/d)
0
50
100
150
1975
1980
1985
1990
1995
date

(c)

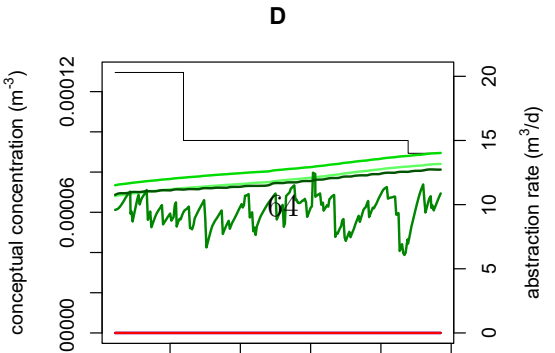
D
conceptual concentration (m-3)
00000
0.00006
0.00012
abstraction rate (m3/d)
0
5
10
15
20

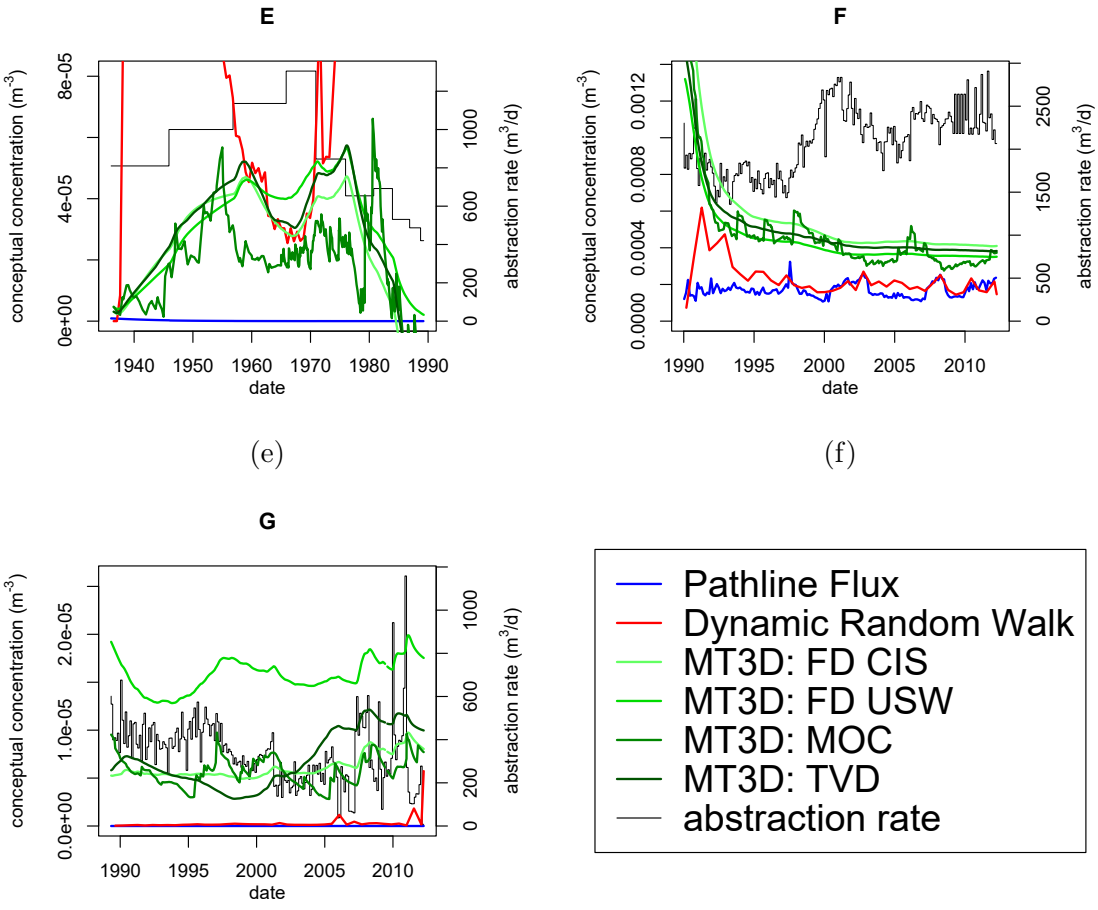

(e) (f)

(g) (h) legend

Figure 2.26: Transient conceptual concentration results at selected wells (see figure 2.22). Plotted are the estimates from the six different methods (PFT, DRW and the four methods within MT3D), using the left axes and the modelled abstraction rates from the well using the right axes (black lines). Note that all axes may be different for each plot. Negative values for the MT3D solutions are indicative of non-converged results.

2.6.4 Conclusion

There is no analytical solution to a problem of this complexity, with multiple sources and receptors (contrast section 2.5.3), so it is impossible to compare the results with an exact solution. Nonetheless, some important observations are possible about the representation of dispersion by the different methods. The PFT model simulates temporal longitudinal dispersion but not spatial dispersion. As such, contaminant flux to wells very close to each other is not partitioned, but given to the dominant well. Without modelling spatial dispersion, particularly transverse dispersion, it would not in fact be clear whether the dominant abstraction in each case should receive all the contaminant flux or whether there is really dispersion across capture zone boundaries. Given also the concerns about particle trajectory errors (section 2.4), this method should be used with caution. The MT3D methods, by contrast, tend to over-estimate dispersion in this context of coarse-gridded regional models and contaminant is spread over an area far too broad, due to numerical dispersion. Additionally, the solutions are oscillatory. Clearly, a 200×200 m cell size is unsuitable for a low-dispersion point source transport simulation with Eulerian or part-Eulerian methods.

The DRW method avoids both these pitfalls, with a spatial representation of dispersion which is not influenced by the model grid. It is therefore concluded that the DRW method is the most suitable method for simulating contaminant transport from point sources on regional scale models. The user will need to understand the significance of the coalescing search radius (appendix A) and be aware of the assumptions inherent in the method (section 2.5.1). However, the DRW method is a flexible and comprehensively documented package (Barry, 2017c). With demonstrated accuracy, an accessible user interface and self-describing outputs, the DRW is ideally suited to the problem of regional contaminant transport from point sources.

In order to produce real concentration outputs and for validation against concentrations observed in the field, the DRW method must be coupled to a physically-based source term model. A chlorinated solvent-appropriate source term methodology is developed in chapter 3 and the DRW method is validated against field data in chapters 4 and 5.

Chapter 3

A TRANSIENT MODEL FRAMEWORK FOR THE GENERIC REPRESENTATION OF CHLORINATED SOLVENT DNAPLS

Chlorinated hydrocarbon (CHC) solvents were used extensively by certain UK industries for most of the 20th century. The potential chronic toxicity of CHC solvents was not apparent until the 1970s by which time a legacy of subsurface pollution had formed (section 3.1). CHCs form dense non-aqueous phase liquids (DNAPLs) in the environment, which take on complex subsurface architectures and dissolve slowly into groundwater. Representation of the DNAPL source term is beset by two major sources of uncertainty. Historically, there is normally very little constraint on the usage and leakage rates at industrial sites over time. And physically, the unpredictable DNAPL source zone architecture makes it hard to represent the source zone contaminant flux with a transferable upscaled model.

The model described and presented within this chapter aims to gather both historical information and findings from hydrogeological studies of DNAPL sources in order to build a generic model. The goals of this model are (1) to isolate issues of uncertainty within the conceptual understanding of DNAPL contaminant sources so they may be addressed transparently, (2) to allow consistent comparison of the risk posed by various CHC user sites, (3) to provide a depth-varying transient source term representation for input into regional transport models and (4) to provide a model framework which allows users to trial different parameter values as well as different models for DNAPL behaviour. The model can use MODFLOW output for the transient groundwater flow through the source zone as part of the expression for the rate of mass transfer by dissolution and its output may be used as the source term for transport simulations such as MT3D or those described in chapter 2. Used in isolation, the source term model can allow water resource managers to assess source zone metrics such as depth penetration and source longevity within the bounds of uncertainty of the inputs.

A detailed numerical model of the dissolution of a DNAPL source zone is not feasible at the regional scale because of the influence of very small-scale variations in aquifer properties on DNAPL behaviour, so empirical models are be used. An empirical model must be firmly based on the results of scientific investigation and section 3.1 reviews the development of scientific understanding of the subsurface distribution and dissolution of DNAPL source zones. Section 3.2 overviews computational models of DNAPL distribution and dissolution, both detailed numerical experiments and attempts to develop upscaled models which are based on empirical parameters. These reviews are summarised by table 3.2 which indicates the increasing difficulty of predicting DNAPL behaviour with increasing scale of investigation. The legacy of historic spills at the regional scale (that is, generally, spills distributed over the area of a town or city) is the largest scale of interest for describing the DNAPL source term and the purpose of this work is to fill this gap. The solution is described in the following sections, with section 3.3 giving the conceptual model for the source term framework and section 3.4 outlining its implementation as a package in `R` (R Core Team, 2014; Barry, 2017b). Finally, in section 3.5, the source term model framework is demonstrated with a published pump and treat experiment on a contaminated land site in the USA, the Tucson International Airport Area (TIAA). Because the site has more available information than most regional-scale studies would, the study allows a more detailed calibration to data. Chapters 4 and 5 present case studies applying the DNAPL source term

framework in regional-scale aquifer contexts.

3.1 Review of Chlorinated Solvent DNAPL Presence, Subsurface Distribution and Dissolution Behaviour in the Environment

3.1.1 Groundwater contamination by chlorinated solvents related to historical industrial activities

Chlorinated hydrocarbon solvents were extensively used by various industrial activities for the majority of the twentieth century. An early study in Reading, UK, found that spilt chlorinated solvents had entered groundwater and could contaminate drinking water wells, but this study went largely unnoticed and the work was not further publicised by its authors (Lyne and McLachlan, 1949; Rivett et al., 2006b). Since then, the problem of chlorinated solvents in groundwater only came to public attention in the late 1970s and into prominence in the 1980s (Westrick et al., 1984; Schwille, 1988; Rivett et al., 1990, 2006b). This coincided with a general increase in public awareness and legislation of environmental issues (COPA, 1974; EPA, 1990; Franklin et al., 1995) as well as some high-profile public health disasters involving organic contaminants (Gibbs and Levine, 1981; Harr, 1995). In the UK, awareness of the chlorinated solvent problem picked up slowly, with early surveys conducted by Fielding et al. (1981) and Folkard (1986) among others.

Rivett et al. (1990) undertook the first UK city-wide survey of CHCs in groundwater in the Birmingham aquifer, finding a high incidence of aqueous chlorinated solvents in groundwater, especially trichloroethene (TCE), around areas associated with twentieth century industry. The Birmingham sampling survey was repeated in subsequent decades (Taylor, 1998; Rivett et al., 2005; Botha, 2006; Murcott, 2008; Rivett et al., 2012) and although the number of available monitoring points declined with time due to site closures, the results have revealed the evolution of water quality with time in the Birmingham Aquifer. Chlorinated solvent contamination is persistent even beyond the period of usage of the chemicals by industry, with some samples even showing greater concentration than in the earliest surveys, especially tetrachloroethene (PCE), whose peak usage by UK industry came later than TCE (Rivett et al., 1990, 2012). It has been hypothesised that the gradual closure of industrial sites which operated abstraction wells along with the UK's industrial decline may have released solvent plumes that were previously contained by these wells (Rivett et al., 2005; Murcott, 2008).

Subsequent surveys have found widespread chlorinated solvent contamination under many cities with industrial histories, elsewhere in the UK and globally (Burston et al., 1993; Grischek et al., 1996; Rowe et al., 2007; Tait et al., 2008; Graber et al., 2008). This has caused particular problems for water companies seeking to exploit the urban groundwater resource and for the Environment Agency which is responsible for

returning water bodies, including aquifers, to "good" status according to the European Union Water Framework Directive. Many of the sources of this contamination are unknown, at least not on the level of individual sites, and many former industrial areas in western countries are now redeveloped. Probably, therefore, the majority of chlorinated solvent source zones have not been actively remediated or contained and the chlorinated solvent legacy is now a public problem, not easily attributable to individuals or organisations. Therefore, water resource management strategies cannot rely on remediation efforts alone and must look to regional risk assessment methodologies (Tait et al., 2008; Troldborg et al., 2008; Rivett et al., 2012; Jamin et al., 2012b).

3.1.2 The DNAPL behaviour of chlorinated solvents in the subsurface environment

The classical paradigm for a contaminant source term is of an area of contaminated land from which contaminant is leached into the aqueous phase by infiltrating rainfall (Hadley and Newell, 2014) and this conceptual understanding still underpins some chlorinated solvent risk models (Tait et al., 2004a; Troldborg et al., 2008). Whilst this understanding is appropriate for many contaminants, particularly those of agricultural origin, chlorinated solvent spills generally behave differently to this (Pankow and Cherry, 1996). Because of their low solubility (generally 0.1–10 mg l^{-1}), chlorinated solvents tend to remain as a non-aqueous phase liquid (NAPL) in the environment for a long time, often many decades, dissolving slowly. Geological and the resulting source architectural considerations discussed in section 3.1.3 show that, additionally, chlorinated solvent source terms tend not to be dissolved efficiently, so that downstream concentrations are well below solubility. This increases the persistence of chlorinated solvents in the subsurface and means that pump and treat remediation operations usually must run for a long time (Blum and Annable, 2008; Akladiss et al., 2011).

The key properties of chlorinated solvents which affect their behaviour in the subsurface are their density, which is greater than water, their viscosity, which is generally low and their interfacial tension with water, resulting from their low solubility (Pankow and Cherry, 1996). This set of properties (especially low solubility and high density relative to water) classes chlorinated solvents as DNAPLs. The density and water solubility of some key chlorinated solvents are given in table 3.1. Chlorinated solvent DNAPLs are able to migrate downward beyond the capillary fringe and water table (Schwille, 1988), often quite rapidly, stabilising within hours or days (Kueper et al., 1993; Gerhard et al., 2007; Sweijen et al., 2014), although some other DNAPLs take much longer to stabilise due to higher viscosity and density closer to water, (Gerhard et al., 2007). The migration of DNAPL is stopped by two mechanisms. DNAPL may encounter a physical barrier such as a relatively impermeable geological layer which presents a high entry pressure to the migrating DNAPL, in which case it will pond and accumulate, migrating laterally if the accumulating NAPL head is sufficient to overcome the lateral capillary barriers and perhaps penetrating through the barrier if the NAPL head becomes sufficient for it to exploit a pathway through the layer

(Johnson and Pankow, 1992; Pankow and Cherry, 1996; Kueper et al., 2003). NAPL migrating through a porous medium has a residual saturation (Mercer and Cohen, 1990) which gives rise to the second mechanism by which DNAPL may be immobilised (figure 3.2). When DNAPL recedes (wetting) after fully invading a medium (draining), a residual trail of low NAPL saturation remains — this saturation is the residual saturation and is reached when NAPL blobs in the pore spaces are no longer connected through the pore throats and as such become isolated and immobile (Pankow and Cherry, 1996; Kueper et al., 2003). These two mechanisms are commonly termed pools and ganglia respectively and may be thought of as two distinct domains of a source zone (figure 3.1), which have distinct transient dissolution characteristics (Pankow and Cherry, 1996; Kueper et al., 2003; Christ et al., 2006).

NAPL migration displacing water can occur in different ways depending on the relative properties of the different liquids (particularly the relative viscosities and interfacial tension: Liu et al., 2015). Whereas in the vadose zone, where the displaced fluid is air (very low viscosity and low interfacial tension with chlorinated hydrocarbon NAPL), the downward migration of DNAPL occurs without great capillary resistance, in the phreatic zone the NAPL-water interfacial tension is higher, such that the two fluids resist juxtaposition and as such the migration of NAPL through saturated water is in narrow fingers through preferential pathways (Anderson et al., 1992; Held and Illangasekare, 1995; Parker and Park, 2004). If it were the case that DNAPL migrated as a stable front, it may give rise to a source zone that could be approximated as a cylinder, at least for relatively homogeneous media, but in fact DNAPL tends to migrate as fingers exploiting pore-scale variations in pore aperture sizes as it penetrates to deeper layers of the aquifer. As such, it is highly unpredictable, defying simplistic representation. Key features of a the DNAPL distribution in a source zone (termed the source zone architecture) from a practical perspective are the depth of penetration (Sweijen et al., 2014) and the proportion of NAPL in pool zones (Christ et al., 2006, 2010; Alexandra et al., 2012), which has a strong influence upon the rate of depletion of the source zone (section 3.1.3). Depth of penetration through porous media may depend upon DNAPL properties, especially density (Sweijen et al., 2014), and geological heterogeneity, especially the lateral continuity of less permeable layers (Kueper et al., 1993; Lawrence et al., 2006). However, although general observations may be discerned about source zone architectures and their dependence on the geology, the sensitivity of DNAPL movement to the subtlest variations in aquifer properties means that it is very hard to predict the lateral migration pathways and source zone architecture accurately (Brewster et al., 1995).

3.1.3 The dissolution of chlorinated solvent DNAPLs into groundwater

The dissolution of chlorinated solvents into groundwater is an important subject for predicting the strength of contaminant flux emanating from source zones, and as such it has been the subject of a number of laboratory and field experiments. Imhoff et al.

Table 3.1: Density and water solubility of selected chlorinated solvents at, or close to 20°C. Sources: National Center for Biotechnology Information, 2017d,a,c,b. There is significant variation in the experimental results for the solubility of TeCM.

solvent	density (kg m^{-3})	water solubility (g l^{-1})
trichloroethene (TCE)	1464	1.28
tetrachloroethene (PCE)	1623	0.206
1,1,1–trichloroethane (1,1,1–TCA)	1338	1.29
tetrachloromethane (TeCM)	1594	0.8

(1994) conducted one of the first column experiments, featuring a thick disc of sand invaded with TCE at an immobile saturation. Water was flushed through while both one-dimensional DNAPL saturation distribution and effluent concentration were monitored and a broad dissolution front was observed at the leading end of the contaminated disc. Effluent concentration remained at TCE solubility (1280 mg l^{-1}) for most of the duration of the experiment until the dissolution front reached the tail end of the contaminated disc. Imhoff et al. (1994) interpreted the broad dissolution front as indicating non-equilibrium dissolution of TCE but an alternative explanation could be that faster streamlines of water (due to small-scale heterogeneities) flush TCE more quickly and then become free of NAPL, but that each individual streamline in fact has a sharp dissolution front. Suchomel and Pennell (2006) also suggest that there is not a significant decrease in source strength until most of the ganglia zones are depleted, using findings from their sand tank DNAPL flushing experiment; in a ganglia-dominated source zone ganglia produce the most significant source flux and are flushed at a roughly constant rate until depleted. On the other hand, source zones with lower proportions of DNAPL mass contained in ganglia will show a greater reduction in source strength as DNAPL mass is depleted (Suchomel and Pennell, 2006; Chen and Jawitz, 2009).

There have been multiple attempts to derive a function for DNAPL dissolution based on dispersivity and DNAPL dimensions (Oostrom et al., 2006). For pools, a steady state relationship was proposed by Hunt et al. (1988):

$$C = C_\mathrm{s} \operatorname{erfc}\left(\frac{z}{2\sqrt{\frac{D_z x}{v}}}\right) \tag{3.1}$$

(C_s: aqueous solubility; z: height above pool; D_z: vertical dispersivity; x: distance along the pool in direction of flow, up to the total length of the pool; v: pore water velocity). However, sand tank experiments have consistently found that the downstream concentrations just a small distance above the pools are much lower than the solubility of the NAPL (Pearce et al., 1994; Chrysikopoulos et al., 2000), which may suggest that pool dissolution flux is limited by NAPL-water contact time rather than by diffusion rate. Further, researchers designing experiments have tried to ensure that the top of the DNAPL pool is flat, but this has proved to be very difficult and the complex pool-top topography may enhance mass transfer rates due the increased interfacial area (Whelan et al., 1994; Dela Barre et al., 2002). This observation, while

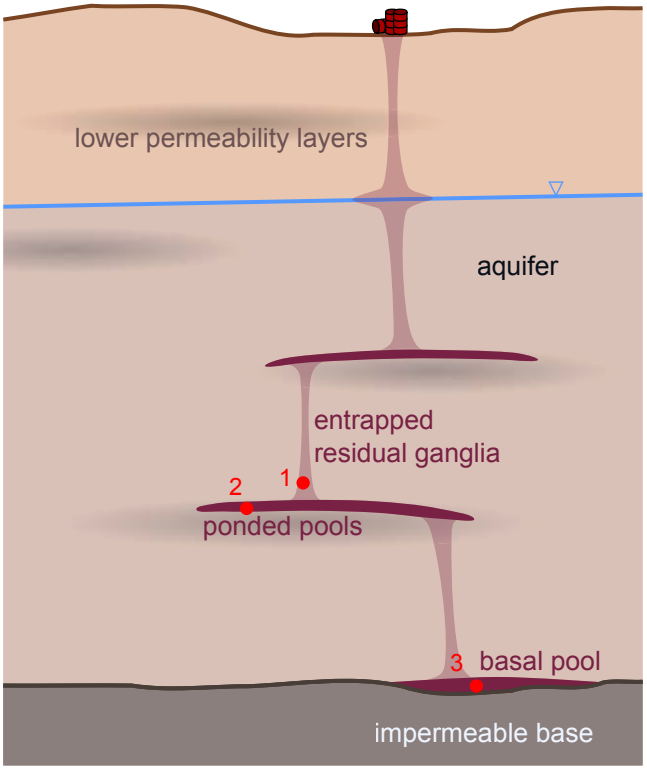

Figure 3.1: Illustration of a DNAPL source zone in a mildly heterogeneous porous medium over a perfect aquiclude, highlighting the different domains of DNAPL mass. The red points are referred to by figure 3.2.

complicating the interpretation of sand tank experiments, is probably more reflective of field conditions, where it is likely that the top surface of DNAPL pools would develop a topography that is controlled by small-scale pore size variations.

Other tank experiments have demonstrated the effects of aquitard material in the vicinity of a DNAPL source. Matthieu et al. (2013) used X-ray diffraction to show that TCE can become intercalated in the crystal structure of clay. This results in aqueous solvent being locked away into a persistent low flux source term. Yang et al. (2015) studied the effects of diffusion and subsequent back-diffusion in prolonging a contaminant source. Although their experiment did not use DNAPL material, they simulated the aqueous transient aqueous concentration that might be expected around a DNAPL source. Their set of experiments did not include a control with no aquitard, but they measured concentration profiles of the clay (kaolinite) layer at various stages during the experiment and did demonstrate that aqueous contaminant mass accumulates in aquitards quickly and then diffuses back out over a much longer time period.

DNAPL release field experiments have enabled the process of dissolution and plume formation to be studied in more realistic settings, with the sacrifice that monitoring is much more difficult. Controlled releases also allow researchers to know much more about the initial conditions of a source zone than at a normal contaminated site. It has been understandably difficult to find suitable locations for controlled re-

lease experiments and, as such, the controlled release experiments at the Canadian Forces Base, Borden, Ontario have given uniquely valuable insights into the dissolution of chlorinated solvent DNAPL sources. Early experiments (Kueper et al., 1993; Brewster et al., 1995) were focussed on imaging the DNAPL architecture and later experiments focussed on downstream concentrations in the plume (Broholm et al., 1999; Rivett and Feenstra, 2005; Rivett et al., 2006a). Broholm et al. (1999) released a multi-component (TCE, PCE and trichloromethane (TrCM)) source through a tube and allowed the DNAPL to redistribute itself into a heterogeneous source zone. They found that the plume generally showed concentrations up to 50% effective solubility, consistent with the presence of DNAPL (Pankow and Cherry, 1996) and had transverse dimensions similar to the source zone except penetrating deeper. The deeper plume was attributed to vertical flows induced by natural water table fluctuations, but may have been exacerbated by vertical pathways formed by the many boreholes drilled for the experiment. The results of Rivett et al. (2006a) also suggest only minimal transverse and vertical dispersion, and that for a longer plume. Rivett and Feenstra (2005) conducted an experiment with an emplaced source, by contrast, designed to be as regular as possible, with a cuboidal homogeneous block of sand with 5% NAPL saturation emplaced below the water table. Although equilibrium concentrations were achieved 1 m downstream of the source zone, saturation and component fraction profiles from within the source zone revealed that dissolution was in fact very heterogeneous, surprising given the homogeneity of the source zone aquifer material (which was mixed before it was emplaced) and the minimal calculated effect of the NAPL at 5% saturation on the hydraulic conductivity. This observation was the first demonstration of dissolution fingering in a field setting and was attributed to the positive feedback loop caused by subtly increased permeability in sections of the source zone where saturation starts to decrease. Dissolution fingering may reduce the efficiency with which ganglia zones are flushed. Based on remaining DNAPL masses within the source zone after the experiment, combined with numerical simulations (Frind et al., 1999), it was estimated that the source zone would persist for 23–32 years, with PCE by far the most recalcitrant solvent on account of its lower solubility. The three-year span of data suggest that equilibrium dissolution will hold for most of the source zone duration (Frind et al., 1999; Rivett and Feenstra, 2005), although it is hard to extrapolate conclusions to the end period of the source term.

Drawing conclusions from accidental source zones at contaminated land sites is more difficult because of lack of early-stage data. Conversely however, studying old spills can give more information about aged source zones (Rivett et al., 2014), which is the condition of most chlorinated solvent source zones today. Over the last decade, source zone characterisation has focussed on finding time-continuous relationships between fractional mass reduction and fractional flux (or concentration) reduction (Brusseau et al., 2007; DiFilippo and Brusseau, 2008; Brusseau and Guo, 2014; Zhu and Sun, 2016). Sources with comparatively little flux reduction resulting from partial source zone removal may not appear to respond favourably to remediation in terms of downstream concentrations, but their longevity will be reduced by such an operation. Conversely, sources whose flux tails off quickly with a small reduction in mass will show good reduc-

tion in downstream concentrations from partial source zone removal (Brusseau et al., 2007), but the source zone is still likely to persist for a long time (Brusseau and Guo, 2014). Commonly, though, a source zone will not display a simple mass-flux relationship and may evolve in multiple distinct steps, complicating predictions (DiFilippo and Brusseau, 2008, 2011). Johnston et al. (2014), on the other hand, provide anecdotal evidence that old source zone depletion functions may eventually resolve to an exponential decay, in line with the conclusions of Chen and Jawitz (2009). Mathematical functions to describe these relationships are discussed further in section 3.2.

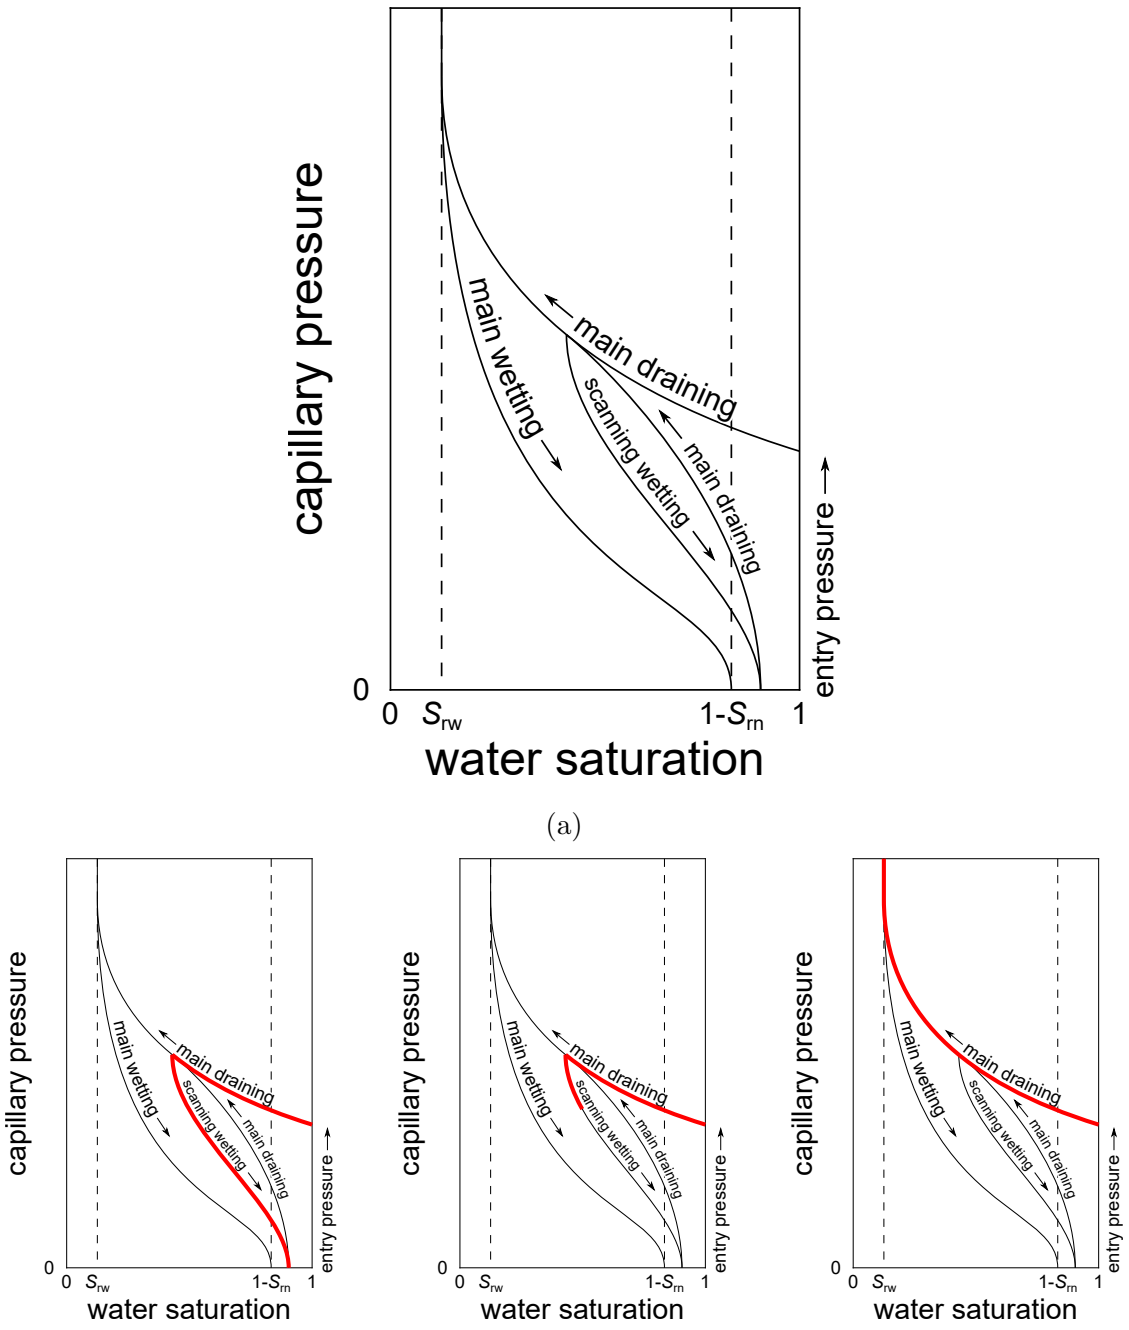

Figure 3.2: (a) A typical saturation-capillary pressure relationship for DNAPL invasion (drainage) and recession (wetting) (Pankow and Cherry, 1996; Guarnaccia et al., 1997) and typical saturation histories for points indicated in figure 3.1, where the source zone has resulted from a continuous slow leak: (b) ganglion, (c) intermediate pool and (d) basal pool. S_{rw} and S_{rn} are residual saturations for water and NAPL respectively. Main curves are spawned from a state in which only one phase is mobile and scanning curves from a state in which both NAPL and water are mobile.

3.2 Review of attempts at Generic and Quantitative Representation of Chlorinated Solvent DNAPLs in models

3.2.1 Numerical models of DNAPL behaviour

A number of numerical schemes have been developed for the simulation of NAPL migration and dissolution. Each model tracks the evolution of multi-phase saturations and more advanced schemes allow the simulation of multiple components within each phase, which may be exchanged between phases (such as with the dissolution process) (Christ et al., 2005). The generic simulated equations may be expressed as (Guarnaccia et al., 1997; Barry et al., 2002; Christ et al., 2005):

$$\frac{\partial \phi S_\alpha C_{i,\alpha}}{\partial t} = \overbrace{-\nabla\cdot\left(\phi S_\alpha C_{i,\alpha}\mathbf{v}_\alpha\right)}^{\text{advection}} + \underbrace{\nabla\cdot\left(\phi S_\alpha C_{i,\alpha}\mathbf{D}_\alpha\cdot\nabla\frac{C_{i,\alpha}}{\rho_\alpha}\right)}_{\text{dispersion}} - \overbrace{\phi S_\alpha R_{i,\alpha}}^{\text{intraphase reactions}} + \underbrace{Q_\alpha\left(\text{if } Q_\alpha < 0\text{: } C_{i,\alpha} \text{ else: } C_{i,\text{source}}\right)}_{\text{mass source or sink}} - \overbrace{\sum_\beta j_{i,\alpha\to\beta}}^{\text{interphase mass transfer}} \tag{3.2}$$

for component i in phase α, subject to a number of boundary conditions that assert mass conservation of each component (Guarnaccia et al., 1997; Barry et al., 2002). ϕ is porosity, S saturation, C concentration, $\mathbf{v}$ phase velocity, $\mathbf{D}$ dispersion coefficient tensor, ρ density, R is a reaction rate, Q a source (positive) or sink (negative), j is an interphase flux and β represents the set of phases not including α. Additional constraining boundary conditions are applied, sometimes implicitly, when simplifying assumptions are used or certain processes neglected. Equation 3.2 couples the transport of multiple components with the flow of multiple phases and as such flow and transport cannot be decoupled, in contrast to most aqueous phase transport models (chapter 2). The velocity is given, according to an expanded form of Darcy's law, by (Guarnaccia et al., 1997):

$$\mathbf{v}_\alpha\phi = \mathbf{q}_\alpha = \frac{\kappa k_\alpha}{\eta_\alpha}\left(\nabla P_\alpha - \rho_\alpha\mathbf{g}\cdot\nabla z\right) \tag{3.3}$$

κ is the intrinsic permeability of the porous medium, k is relative permeability, η viscosity, P pressure, $\mathbf{g}$ gravitational acceleration and z elevation.

The relative permeability is related to the saturation of the phase, according to a k–S–P model which relates relative permeability to saturation and saturation to capillary pressure. The k–S–P relationship is generally hysteretic, with, for example, successive draining events by NAPL invasion following different saturation paths and different relative permeability paths (Guarnaccia et al., 1997; Gerhard and Kueper, 2003b;

figure 3.2). The field of k–S–P models started in petroleum science. The Brooks-Corey model has often been imported from this subject area to water resources studies (Brooks and Corey, 1964; Christ et al., 2005; Basu et al., 2008), but efforts have been made to develop k–S–P models more suitable for the water resources context (Gerhard and Kueper, 2003a,b). The DNAPL3D model is a finite difference code for DNAPL migration that is based on these latter k–S–P relationships (Gerhard et al., 2007).

The results of numerical modelling studies of DNAPL migration has given several insights into the general patterns of DNAPL behaviour in the subsurface and have been useful in guiding the development of upscaled models, not least in highlighting the high complexity of DNAPL behaviour and cautioning against attaching too great a weight to any upscaled model (Christ et al., 2006). Kueper and Frind (1991b) used one of the earliest numerical DNAPL models (Kueper and Frind, 1991a) with a randomly heterogeneous permeability field, demonstrating the strong effect of even mild heterogeneity, with very fine variations in aquifer capillary properties controlling the apparently erratic pathway used by downward DNAPL migration. Studies such as this show that mechanistically modelling DNAPL in perfectly homogeneous media would bear no relation to field site reality or even perhaps to laboratory conditions. Nonetheless, more heterogeneous aquifers do appear to hold up more DNAPL mass in higher layers, with more uniform aquifers leading to deeper, more vertical descents (Basu et al., 2008). Gerhard et al. (2007) used DNAPL3D to simulate the effects of NAPL phase properties on migration, discerning that low viscosity DNAPLs such as PCE tended to stop migrating earlier, penetrate deeper and distribute more broadly than more viscous DNAPLs such as coal tar. Lower density DNAPLs also tended to spread further horizontally (Gerhard et al., 2007; Sweijen et al., 2014).

Numerical DNAPL models have also studied the problem of DNAPL dissolution into flowing groundwater, obviously a key question for addressing the DNAPL source term. Christ et al. (2006), using UTCHEM (Delshad et al., 1996), showed that dissolution rate is reasonably constant until most of the entrapped residual (ganglia) mass is depleted, whereupon the source zone declines more slowly. Christ et al. (2010) developed this further, showing that the presence of a basal pool tends to result in a two-stage mass depletion profile, with the basal pool causing a persistent steady flux until it is almost depleted. However, their simulated heterogeneous permeability field is isotropic, which would have diminished the importance of intermediate pools of DNAPL. By contrast, Parker and Park (2004) used an anisotropic variable permeability field to generate a spill zone using their own DNAPL percolation model. Like Christ et al. (2006), they aimed to use their simulated source zone to develop an upscaled model of DNAPL dissolution. Their numerical simulation conceptualises DNAPL migration in two modes — vertical movement in fingers, between layers and horizontal spreading within layers. Vertical percolation occurs if the entry pressure of the underlying model cell is overcome by the DNAPL head and if not, lateral spreading may occur if the entry pressure of laterally adjacent cells are overcome. If not, then the DNAPL lens will mound until sufficient head is built up to overcome the entry pressure somewhere. It was assumed that DNAPL head is constant over vertically continuous flow paths, probably justified given the low viscosity of TCE, the specific chemical they modelled. One incidental

result from the percolation simulation is that the average saturation of the pool zones are hardly more than the ganglia zones, with the final saturation distribution not much exceeding 0.3. This observation is consistent with field observations that DNAPL saturation at field sites rarely exceeds 0.25 (Stewart and North, 2006; Wang et al., 2013).

3.2.2 Analytical models of DNAPL behaviour

Given the dependence of each DNAPL source zone on fine-scale heterogeneity, it is impractical to suppose that a generic source term representation could depend on a fully mechanistic numerical model (Parker and Park, 2004). However, representation of the dissolution of DNAPL is hard to upscale, so that a generic analytical representation of DNAPL source terms is difficult to develop. Perhaps the multiple analytical source superposition technique (MASST) comes closest to this; this method superposes the analytical flux from many sub-sources (distinct pools and ganglia, conceptualised as oblate and elongate cuboids respectively: Sale and McWhorter, 2001). However, this method presupposes a detailed knowledge of the distribution of DNAPL and a regionally-applicable generic source term model cannot assume *a priori* knowledge of source zone distributions.

Anderson et al. (1992) and Johnson and Pankow (1992), in their linked studies, were amongst the earliest to conceptualise the DNAPL source zone with two domains, ganglia and pools (section 3.1). Using their own and Schwille's (1988) laboratory experiments, they develop analytical expressions for the eluting concentration from each domain. There is more that can be measured about the NAPL source from laboratory set-ups than field sites however, and as such these methods are probably too dependent on source zone dimensions to be useful for generic quantitative risk assessment from real spill sites.

Parker and Park (2004) conducted numerical DNAPL distribution and dissolution simulations (section 3.2.1) in an attempt to devise an upscaled analytical model for DNAPL source terms. They define an effective mass transfer coefficient κ_{eff}[1] which takes account of source zone architecture and matches numerically simulated concentrations to analytically derived results. Assuming a low Damköhler number (that is, reaction rate is low compared to advection), their expression approximates to:

$$\frac{C}{C_s} = \frac{\kappa_{\text{eff}} L}{q} \tag{3.4}$$

where L is the source zone length and q is the Darcy velocity through the source zone. Their method is able to match numerical results for both pool-dominated and ganglia-dominated portions of the source zone. κ_{eff} captures the state of the source zone at a given time and if this method was applied to the long-term evolution of a source zone

[1] Not related to κ in section 3.2.1.

then κ_{eff} would be regarded as a transient parameter that would reduce as a source became more pool-dominated.

More recent focus on analytical DNAPL source term representations has attempted to develop empirical functions that relate the eluting concentration from the source zone to the source zone mass, both in terms of proportional depletion from initial values. Linear and non-linear power relationships between proportionate mass reduction and proportionate decrease in source term concentration have been suggested (Zhu and Sykes, 2004; Falta et al., 2005b; Jawitz et al., 2005), expressed as:

$$1 - \frac{C}{C_0} = \left(1 - \frac{m}{m_0}\right)^{\Gamma} \tag{3.5}$$

where C is downstream averaged concentration, m is remaining source zone mass, the subscript $_0$ denotes initial values and Γ is an empirical constant. The dependence of the model on the initial values of concentration and mass implies the importance of source zone history. These initial values are very hard to estimate because they cannot be measured today. Chen and Jawitz (2009) investigated this model further with a series of sand tank experiments and found that the estimate of Γ changed throughout the experiments, but did so predictably, always converging on 1, so that various complicated DNAPL source terms would all tend towards a linear relationship between mass and concentration, such that:

$$1 - \frac{C}{C_0} = \left(1 - \frac{m}{m_0}\right)^{1-\frac{m}{m_0}(1-\Gamma_0)} \tag{3.6}$$

The equilibrium streamtube model (ESM) attempts to characterise the DNAPL source zone in terms of a small number of statistical parameters, removing the need for empirical parameters that are hard to estimate (Jawitz et al., 2005). The source zone is conceptualised as a collection of non-interacting streamtubes, some of which are partially occupied by NAPL, with the assumptions that there is no transverse dispersion between streamtubes, that NAPL saturation does not significantly modify water relative permeability and that dissolution is not rate-limited. Each streamtube has a trajectory-integrated NAPL saturation and a non-reactive travel time (simply the water travel time along the length of the streamtube), which may be combined to give a reactive travel time, the duration of flushing required to clean the streamtube. The number of contaminated streamtubes at time t is then the number of streamtubes which originally had a reactive travel time greater than t. Average effluent concentration is expressed as the solubility concentration times the proportion of contaminated streamtubes, and decreases as streamtubes are purged of NAPL. The ESM has been validated using a numerical study (Basu et al., 2008) and a sand tank experiment (Chen and Jawitz, 2009) and the characterisation of reactive travel time distribution has been demonstrated at a field site (Wang et al., 2013). Zhang et al. (2008) have shown that

neglecting the changes in water relative permeability can lead to errors when saturation, although this should not be an issue with saturations more commonly observed in the field (Wang et al., 2013).

3.2.3 Major uncertainties for DNAPL models

The challenge of designing a representation of the DNAPL source term that may be transferred to cases beyond the test case is great. It requires that the metrics used to describe the source are possible to measure or estimate, ideally with minimal expense, and are known to be reasonably consistent between sources in similar circumstances. Table 3.2 demonstrates the increasing difficulty of applying models in contexts closer to the field scale and ultimately the regional scale. Methods and models developed for one tier of DNAPL research do not usually lend themselves to application at the next tier. With computational and laboratory studies it is possible to know the detailed NAPL saturation distribution. Metrics such as the source zone length (Parker and Park, 2004), the ganglia-to-pool ratio (Christ et al., 2006, 2010) and the reactive travel time distribution (Basu et al., 2008; Chen and Jawitz, 2009) depend on this detailed characterisation that is generally not available at field sites. It is very hard to determine the saturation distribution in the field, as a large proportion of DNAPL mass tends to occupy a very small volume of aquifer (Rivett et al., 2014) and as such very finely-spaced sampling is required to find the most significant pools of DNAPL (Parker et al., 2003; Lawrence et al., 2006). Geophysical methods may give a more detailed picture (Kueper et al., 1993; Power et al., 2014) but because of the complexity of the geological signal in most cases these methods are useful for time-lapse studies, but less so for assessing static source zone architecture. Wang et al. (2013) have demonstrated the use of partitioning tracer tests to estimate the reactive travel time distribution, giving a practicable and relatively non-invasive technique for generating usable source zone metrics for the ESM.

Nonetheless, the existing upscaled models for the DNAPL source term are appropriate for, at best, evaluation of individual source zones and are designed for application to a single contaminated land site. To develop a source term representation that can be applied regionally to spills which are little known and for which subsurface investigation is infeasible is a different problem, largely untried. Given the complexity of DNAPL distribution and dissolution it would be very difficult for such a model to deliver a reliable deterministic solution but, through scenario analysis and comparison with historic data, it may show useful predictive results about the regionally distributed chlorinated solvent source to an aquifer and its dependent wells. The following sections describe a methodology for the regional DNAPL source term, drawing upon historical trends of usage and spillage as well as synthesising the scientific understanding of physical DNAPL processes, attempting to account for the redistribution to the environment all chlorinated solvent mass imported to an industrial area.

Table 3.2: Summary of DNAPL source zone investigation contexts, and available methods for investigating source zone architecture and modelling dissolution.

context	source zone investigation methods	source dissolution models
numerically simulated	source zone characteristics are known from the model results	numerical simulation, against which upscaled methods may be tested (Parker and Park, 2004; Christ et al., 2006; Basu et al., 2008)
sand tank	detailed imaging (Zhang et al., 2008; Chen and Jawitz, 2009; Christ et al., 2012)	MASST, ESM and other semi-analytical methods (Zhang et al., 2008)
controlled release	intensive multi-level sampling (Rivett and Feenstra, 2005) and geophysical surveys (Kueper et al., 1993)	empirical mass removal-flux reduction relationship (DiFilippo and Brusseau, 2008)
contaminated site	intensive multi-level sampling (Lawrence et al., 2006; Rivett et al., 2014) and partitioning tracer tests (Wang et al., 2013)	ESM (Wang et al., 2013) and empirical mass removal-flux reduction relationships (DiFilippo and Brusseau, 2008)
regional legacy of spills	conceptual representation of source zone history, distribution and behaviour informed by literature survey	this work

3.3 A Generic Modelling Framework for Historical Chlorinated Solvent Releases in Groundwater

Table 3.2 shows that representing the DNAPL source term becomes increasingly difficult with the progression from the computer to the laboratory to the field and finally the regional scale. Transferable metrics for DNAPL dissolution are just starting to be applied to individual contaminated land sites, using a partitioning tracer test to characterise a source zone (Wang et al., 2013). At the regional scale, where, for an aquifer under a historically industrial area, there may be many individual spills, such characterisation is infeasible and there is a need for some meaningful generic characterisation of the DNAPL sources in this context beyond a simple recharge-leached concept of a source zone (Hadley and Newell, 2014). The method presented in this chapter serves as a framework within which a quantitative physically- and historically-based DNAPL source term representation may be built. Although, given the lack of site data avail-

able in the regional context, it is not likely that a source term representation within this framework will constitute a full and detailed description of DNAPL source zone processes, having a model that is specified in terms of DNAPL-relevant processes is advantageous. It means that the source term will respond to the correct system stresses (for example, the source term is sensitive to groundwater flow rather than recharge) and that uncertain features of the DNAPL source term are explicitly represented in terms of DNAPL-related parameters, even if they are empirical (such as Γ).

3.3.1 Conceptual Model

The conceptual model covers both the historic usage rate and handling of chemicals on site and the behaviour in the underlying aquifer. It is split into four stages:

1. site import and usage;
2. export and spillage;
3. infiltration into and through the aquifer as DNAPL;
4. mass transfer including dissolution.

The movement of chemical mass between these stages is demonstrated in figure 3.3.

The site import is, in a manner of speaking, the source term to the source term model. It gives the rate at which CHC has been brought onto site over time. This mass is then routed down various avenues (export, safe disposal, vapour or infiltrating DNAPL), according to the practices of the time (figure 3.3). The import rate may be known, but for most sites there will be sparse if any records. Therefore, import rate will be based on an estimated demand for the site, using the historic production rate for the UK (figure 3.4) and scaled to an estimate of the site's likely peak usage. The start and closure years of site operations may be approximately determined from studying historical maps and internet searches. These dates can be used to clip the usage rate history.

Export from site may occur when chemicals are used as raw material for a product or if they are sent away to be disposed of safely. CHC solvents were used primarily for degreasing agents, implying that very little would have been exported as part of products. With improved legislation and awareness (COPA, 1974; EPA, 1990; Franklin et al., 1995), more waste, on average, would have been taken off site for safe disposal. Although this transition would have occurred at a certain time for each individual site, the generic source term model seeks to capture average behaviour across all sites apart from any exceptions where detailed information is held. Therefore the recommended default for the sludge factor, the proportion of chemical which is not disposed of safely, is a linear decrease from 1 to 0 between 1974 and 1990 (table 3.3).

Infiltration into the aquifer uses one lumped parameter in this model: infiltration fraction, the fraction of DNAPL mass that is spilt or disposed of on site that reaches

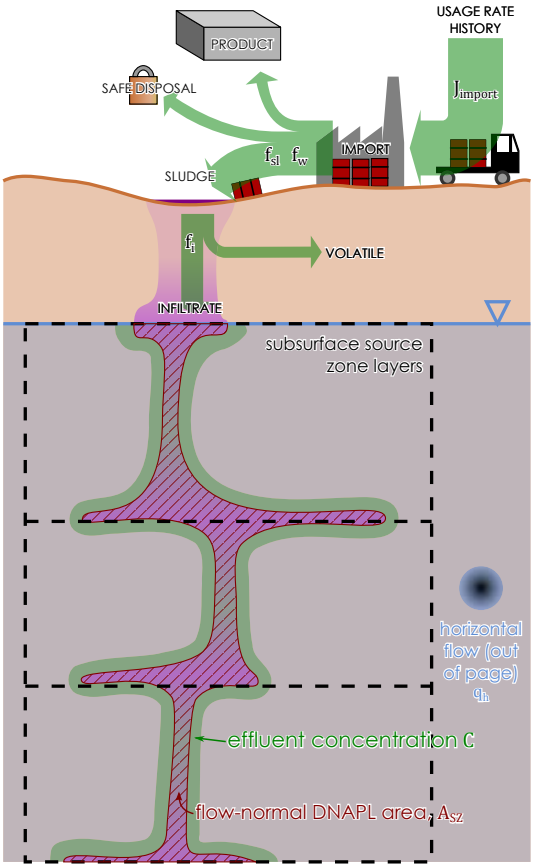

Figure 3.3: Conceptual representation of the DNAPL source term framework.

the water table, rather than vaporising. In fact, this part of the process is very hard to estimate and will depend strongly on local soil conditions, vadose zone thickness, air temperature, specific features of the spill and even perhaps the average wind speed. A conservative approach would be to assume that all spilt DNAPL reaches the water table; it would take intensive site investigation to refine further than this. However, the user may want to distinguish between parts of the aquifer that are more vulnerable on account of depth to groundwater or soil and this would be the parameter to vary in order to show these distinctions. The parameters relating to the import, export and infiltration of DNAPL are listed in table 3.3.

The phreatic portion of the source term model may be divided into multiple depth sections and is run as an explicit one-dimensional numerical calculation. The infiltrating DNAPL acts as the input to the model into the uppermost depth section. Each depth section has a maximum mass of NAPL that it can contain based on the ganglia width, the maximum pool width, the residual saturations, porosity and NAPL density (or, if a custom DNAPL distribution and dissolution model is used, possibly a different set of parameters). Once a section is full, further DNAPL migration into that cell spills into the cell below (this step is called cascading within the `DNAPL` package: section 3.4; figure 3.5). Thus the DNAPL source zone forms from the top down. Each of these sub-source zones will then cause an aqueous flux as a function of the DNAPL mass

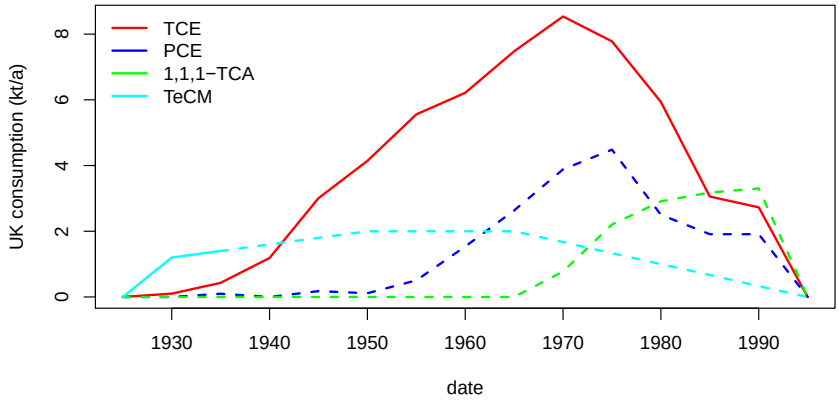

Figure 3.4: Estimated consumption rates of chlorinated solvents in the UK. Dashed lines indicate poor constraint. PCE, TCE and 1,1,1–TCA histories are taken from Rivett et al. (1990) and TeCM is estimated from a value for national consumption in a 1930s report (I.C.I., 1931) and general statements about historic usage patterns in a more recent report (Walker and Jameson, 2010).

present and the Darcy velocity. There are three built-in functions, but a custom-made function can be given. Two of these are based simply on equations 3.5 and 3.6, which give effluent concentration from the source zone, converted to flux by:

$$J = CA_{\mathrm{SZ}}q_{\mathrm{h}} \tag{3.7}$$

q_{h} is the horizontal magnitude of Darcy velocity and A_{SZ} is flow-normal source zone area (figure 3.3):

$$A_{\mathrm{SZ}} = w_{\mathrm{g}}h_{\mathrm{L}} + w_{\mathrm{p}}h_{\mathrm{p}} \tag{3.8}$$

w_p is the current average width of pools in a layer. Pools are allowed to spread laterally until they reach $w_{\mathrm{p,max}}$. Other parameters are given in table 3.4. The average width of pools within a layer is calculated assuming that pools are circular and, when newly formed, that saturation exponentially declines from $1 - S_{\mathrm{r,w}}$ at their base to $S_{\mathrm{r,n}}$ at their top:

$$\begin{aligned} \bar{S}_{\mathrm{n},0} &= \frac{1}{\log \frac{1-S_{\mathrm{r,w}}}{S_{\mathrm{r,n}}}} \left(1 - S_{\mathrm{r,w}} - S_{\mathrm{r,n}}\right) \\ w_{\mathrm{p}} &= \sqrt{\frac{4m_0}{\phi\rho_{\mathrm{n}}\bar{S}_{\mathrm{n},0}h_{\mathrm{p}}\pi}} \end{aligned} \tag{3.9}$$

A third model, termed the dual domain DNAPL model extends the constant Γ model (equation 3.5) to apply separately to ganglia and pools, applying, in effect, $\Gamma = \infty$ to

the ganglia domain, which means that there is no reduction in effluent concentration from ganglia until they are completely depleted, and a flexible Γ value to the pool domain, meaning that the pools are the persistent portion of the source term. The user chooses one of these models to apply or, alternatively, can supply a custom-made model which may represent a different empirical DNAPL dissolution model or set of mass transfer processes.

The parameters relating to the phreatic portion of the source zone model are listed in table 3.4.

Table 3.3: Description of parameters describing DNAPL spillage and infiltration rates.

parameter	symbol	explanation	suggested default or guidelines
import (usage) rate	J_{import}	Function of time: rate at which site imports solvent.	A transient value based on national consumption trends (figure 3.4) and an estimated likely peak usage rate.
waste fraction	f_{w}	Proportion of import ending up as waste (as opposed to being used as raw material).	1
sludge fraction	f_{sl}	Proportion of waste disposed of or leaked unsafely on site.	1 before 1974 (COPA, 1974), 0 after 1990 (EPA, 1990), linearly decreasing in between.
infiltration fraction	f_{i}	Proportion of sludge infiltrating rather than volatilising.	Conservatively 1.

Table 3.4: Parameters used in the three provided distribution and dissolution models (section 3.3.1).

parameter	symbol	explanation	suggested default or guidelines
ganglia width	w_{g}	Characteristic diameter of ganglia. Wider ganglia are flushed more efficiently.	
maximum pool width	$w_{\mathrm{p,max}}$	Average width that pools reach before they are likely to find a new vertical migration pathway and stop expanding.	
residual NAPL saturation	$S_{\mathrm{r,n}}$	Saturation of residual NAPL in rewetted aquifer after maximum NAPL invasion.	0.1
residual water saturation	$S_{\mathrm{r,w}}$	Saturation of water upon maximum NAPL invasion.	0.1
total porosity	ϕ		Allen et al. (1997) summarise data for UK aquifers.
NAPL density	ρ_{n}		table 3.1
mass-flux power	Γ or Γ_0	Power used in mass-flux power law relationship, either initially or constantly for the constant and converging power models (equations 3.5 and 3.6).	< 1 implies ganglia-dominated, > 1 means that the source will be more persistent, dominated by recalcitrant pools.
solubility	C_{s}		table 3.1
layer thicknesses	h_{L}	For the DNAPL source term model.	Evenly dividing the saturated thicknesses of each MODFLOW model layer.
number of layers	N_{L}	Number of layers for DNAPL source term model.	1 or more for each MODFLOW model layer.

3.4 Program Structure

This section is outlines the structure of the DNAPL source term program, showing how it works. The program is written as a package in `R`, called `DNAPL` (R Core Team, 2014; Barry, 2017b), which is included in the supplementary electronic material (appendix D.1). More detailed documentation of individual functions may be found within the package.

Figure 3.5 gives an illustrated overview of the program flow for the DNAPL source term model framework and figure 3.6 shows how different elements of the program structure relate to each other and the inputs. The program is divided into the spill model, which encompasses items 1–3 of the conceptual model (section 3.3.1) and the subsurface model, which is item 4 of the conceptual model. For the spill model, it is necessary to specify the rate at which solvent was imported to the site and the proportion of solvent through time that is spilt rather than exported and then the proportion of spillage that infiltrates to the subsurface (table 3.3). For the subsurface model, the DNAPL distribution and dissolution model must be specified (section 3.4.2) as well as the transient groundwater flow through the source zone, which may be given as MODFLOW results to be interpreted (section 3.4.3). This model is then implemented by the source term framework as an explicit scheme, calculating mass transfers between mass domains (including dissolution) and directing mass that exceeds the capacities of domains (section 3.4.4).

3.4.1 Surface spill model

The surface spill model (figure 3.5) specifies how much of the imported contaminant to a site ends up infiltrating to the water table. This component is relatively simple. It multiplies the import rate through time with values for f_{i}, f_{w} and f_{sl} to give the eventual flux to the water table. This time series is then used as the input to layer one in the solution.

The estimated national usage rates (figure 3.4) are included within the `DNAPL` package as data tables. Historic national consumption of solvents reflects both the number of solvent-using sites and the industrial practices of the times. Solvent-associated industries (such as metal and engineering works and tanneries) existed before the uptake of solvents by industry and continued beyond the cessation of solvent usage. However, solvent usage over time does, in part, follow the industrial productivity over time in the UK. National trends can be scaled to usage rates of individual sites using a utility function within the package. Alternatively, the usage rate at a site may be assumed constant (or described by any other function of time).

The sludge fraction f_{sl} is input as a two-column data table of values against time, allowing this value to vary with changing practices as suggested in table 3.3.

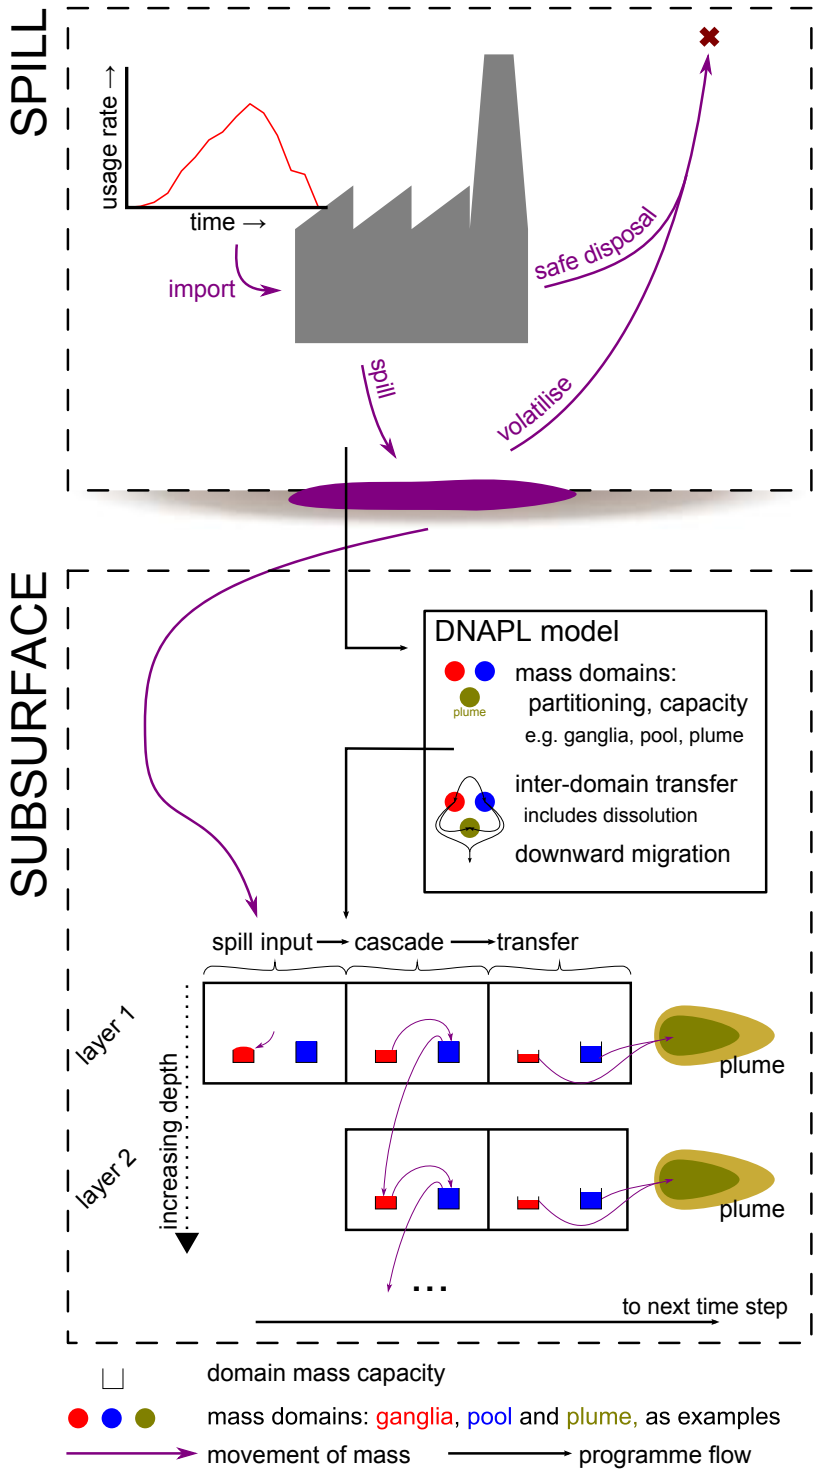

Figure 3.5: Schematic representation of the components and logic of the DNAPL source term model framework. The illustration uses the dual-domain ganglia and pool model (section 3.3.1) as an example; other distribution and dissolution models will have different domains.

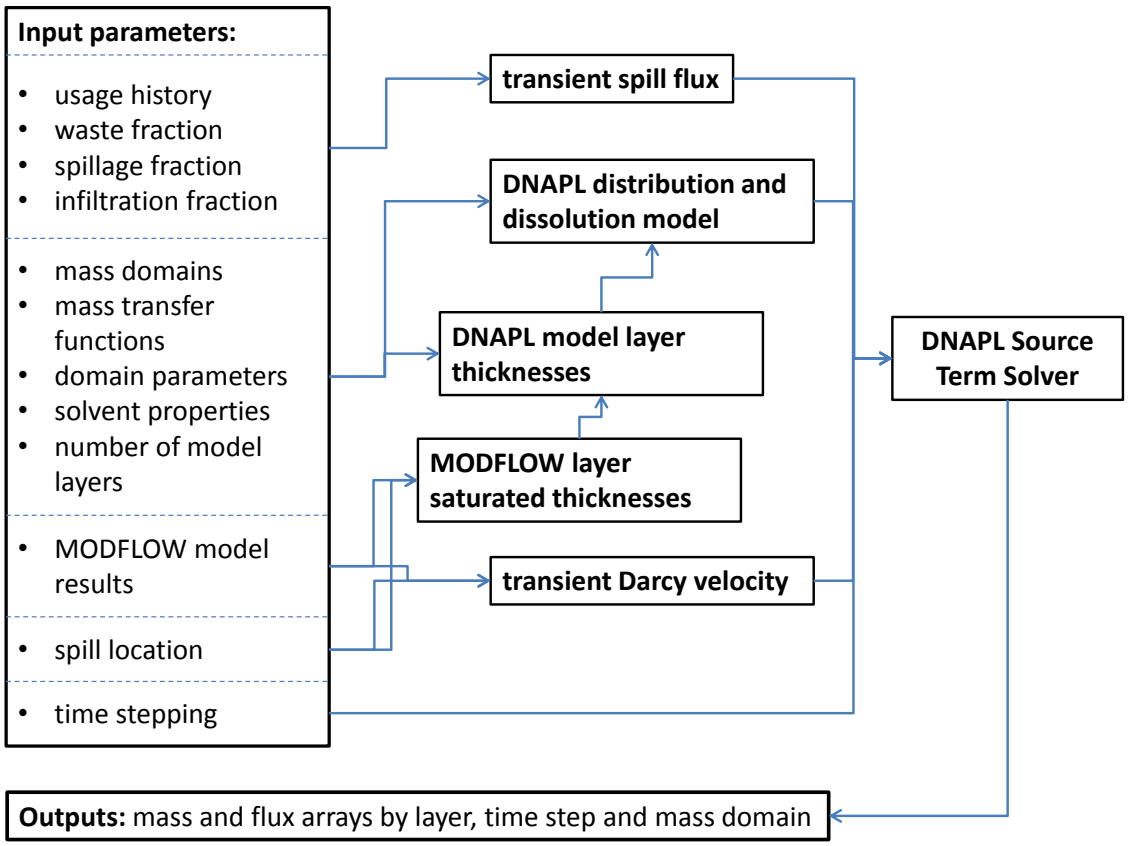

Figure 3.6: Relationships between inputs and the components of the `DNAPL` package (Barry, 2017b). This figure gives a summary of the basic components of the package and more detailed options can be found within the package documentation. MODFLOW data is given in NetCDF format using the `Rflow` package, although layer thicknesses and the Darcy velocity may alternatively be given directly, as in section 3.5. The spill location instructs the solver from which location to read MODFLOW model information.

3.4.2 The NAPL distribution and dissolution model

The DNAPL dissolution and distribution model (figure 3.5) is specified by the user with a number of complex inputs. The model conceptualises the contaminant as existing in a number of mass domains, which could simply be the different phases present (for example, NAPL and aqueous) or could subdivide further (for example, NAPL mass can be divided into ganglia mass and pool mass). These domains interact with each other according to a set of mass transfer functions that operate between pairs of domains. Each domain has a maximum capacity for each layer that governs when a layer is full and when mass penetrates to deeper layers or is redistributed between domains within a layer. The contaminant plume is always given a separate domain which can contain limitless mass. The source term flux from a layer is calculated as the rate of change of plume mass. Within the `DNAPL` package, distribution and dissolution models are declared as `DNAPLmodel` objects, which is a formal object class defined

within the package. The package is deliberately set up to enable users to define custom `DNAPLmodel` objects, but inevitably this is fairly involved and will require a good knowledge of the `R` programming language. Much simpler wrapper functions are provided for a simple interface to create the distribution and dissolution models described in section 3.3.1: `cstG.DNmodel`, `cnvG.DNmodel` and `DDpg.DNmodel` for the constant Γ, converging Γ and dual-domain pool and ganglia models respectively. `DNAPLmodel` objects contain the following components (or slots) which are described in the following paragraphs:

1. list of domains;
2. mass transfer processes;
3. layer-by-layer domain mass capacities;
4. cascade destinations for domains which overfill;
5. domain to which infiltrating mass is initially given.

The list of domain names is given. The built-in models have a NAPL domain and a plume domain for the constant Γ and converging Γ models and pool, ganglia and plume domains for the dual-domain pool and ganglia model.

The mass transfers between domains are a core part of DNAPL distribution and dissolution models and typically describe phase changes such as dissolution or vaporisation. Each transfer is specified as a formal `Mredistribution` object with three components (slots):

1. domain the flux is from;
2. domain the flux is to;
3. a function, giving flux as a function of mass in the "from" domain and the "to" domain.

A mass transfer function may give a negative result which would imply a flux from the "to" domain to the "from" domain and this is useful for defining transfers that maintain an equilibrium. Because the mass of a domain has an obvious minimum value (0), the solver method automatically limits mass transfer calculations to avoid the mass of any domain decreasing below 0 (section 3.4.4). In this way, the model does not suffer from instability.

The dissolution and distribution model must contain values for the maximum mass of each domain in each layer, which is specified as a matrix with rows representing the layers and columns representing the domains. Each value will normally depend on parameters such as the layer thickness, source zone characteristic width, porosity, residual saturations and solubility. The plume domain has limitless capacity, as it

represents mass removed from the source zone in flowing groundwater. When the mass of a domain in a layer exceeds its capacity, it cascades to a destination which is specified in the cascade destinations.

Each item in the table of cascade destinations corresponds to a mass domain with limited mass and specifies the name of the domain to which any overflowing mass is directed as well as whether that mass should go into the same layer, the layer below or (unusually) above. Finally, the name of the domain in layer 1 to which the spill input should be directed must be specified.

The three wrapper functions for the built-in models internally determine all these components with the given parameters.

3.4.3 Groundwater flow model

Although not a mandatory feature, the `DNAPL` package allows easy linking with MODFLOW model results (Harbaugh et al., 2000) using the `Rflow` package, which converts MODFLOW model discretisation and results to NetCDF format (Rew et al., 2016; Michna and Woods, 2016; Barry, 2017e). The DNAPL source term solver can read saturated layer thicknesses and transient flow rates through time at a given location. Using utility functions within the package, DNAPL model layer thicknesses are determined with one or more layer per MODFLOW layer from the MODFLOW layer elevations and the Darcy velocity is interpreted from horizontal cell face water flows. The Darcy velocity is a key component of determining the dissolution flux from NAPL domains and the layer thicknesses may be essential for determining the mass capacity of certain domains within a layer. Users can, alternatively, specify layer thicknesses and Darcy flux explicitly, so a MODFLOW model is not required, and this approach is taken in section 3.5. The DNAPL source term framework is linked to MODFLOW model results in chapters 4 and 5.

3.4.4 Solution

The DNAPL source term model is executed as an explicit 1-dimensional solution. As a 1-dimensional model, run times are normally very quick and therefore it is practical to use a small explicit time step scheme to increase the accuracy of the solution. Additionally, mass transfers are all automatically limited to avoid the mass of any domain in any layer decreasing below 0, which avoids any potential instability. The program sequence is demonstrated in figure 3.5.

With each time step, input mass from the spill is added to the specified initial domain in layer 1. Any mass cascades are then calculated. Mass cascading is iterated until no domain in any layer has a mass greater than its limit. The cascades may form a chain, whereby an overflow may cause the destination layer-domain to overflow to a further layer-domain and so on. If fresh spill is added to a source zone for which the NAPL

domain is full in the top three layers, for example, then NAPL mass will be cascaded into the fourth layer. Mass transfers are then calculated according to the specified mass transfer functions (section 3.4.2). After this, the cascading step is again calculated in case of any new overfull domains, though this is not shown in figure 3.5 for simplicity.

The DNAPL source term solver is the `DNST` function and returns objects with formal class `DNAPLSourceTerm`. This object class contains several pieces of information, including meta-data for future reference, but the key result is the mass array, which has three dimensions representing the layers, time steps and mass domains respectively. Additionally, the source term effluent flux is interpreted from the plume domain as the rate of change of plume mass with change. The effluent flux array may be used to specify the source term for an aqueous contaminant transport model and the dynamic random walk (DRW) model is able to read `DNAPLSourceTerm` results directly (Barry, 2017c).

3.5 Demonstration of the Modelling Framework using Field Results from a Pump and Treat Operation

The generic DNAPL source term model has been developed for the purpose of representing a regional distribution of uncertain and speculative source zones and their dissolution into flowing groundwater. This study applies the source term methodology to a well-studied DNAPL source zone that has been subjected to a pump and treat remediation operation in order to validate the methodology. The site chosen is the Tucson International Airport Area (TIAA), about which detailed information has been put into the published literature (Brusseau et al., 2013).

3.5.1 The Tucson International Airport Area site

The Tucson International Airport Area (TIAA) in southern Arizona, which is a Superfund site, was historically engaged in activities using chlorinated solvents, especially TCE. The central section of the site has a hangar where aircraft parts were cleaned using the substance. Today there is a 600 m long dissolved TCE plume extending west from this source.

The TIAA is located in the Tucson Basin, consisting of several hundreds of metres thickness of alluvial sediments, interbedded locally with volcanic-origin rocks and tuffaceous sediments. At the TIAA, the subsurface is composed of clay and silty clay units with lenses of silt, sand and gravel. There is also a prominent sand and gravel layer, about 3 m thick and 34 m deep, across the site. This sand and gravel layer is phreatic and confined, with a piezometric head about 29 m below the ground surface, forming a shallow aquifer. The shallow aquifer is separated on the site from the deeper regional

aquifer by clay, but merges with the regional aquifer at the western edge of the site, resulting in a contaminant pathway from the shallow aquifer to the regional aquifer, which is an important drinking water source (Brusseau et al., 2013).

There have been multiple remediation projects across the TIAA site. A long-duration pump and treat operation was conducted in the southern division of the site on a different spill zone between 1987 and 2000 (Brusseau et al., 2007). Source term decline was marked at first, but showed diminishing returns after a few years, with recalcitrant contaminant mass forming a persistent source zone. In the central division of the TIAA site, which is the focus of this study, Brusseau et al. (2013) oversaw a pump and treat operation from 2007 to 2012, again showing a declining source term with diminishing returns (figure 3.7).

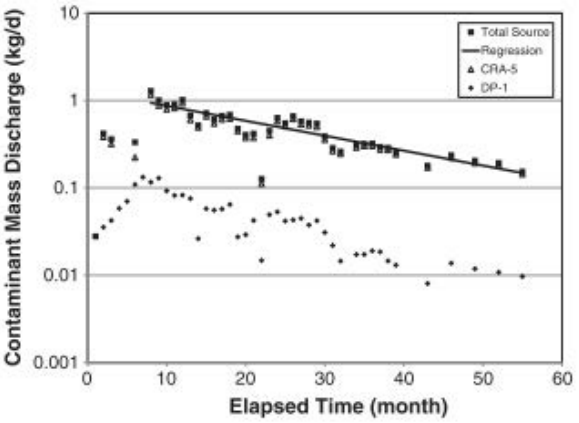

Figure 3.7: Observed contaminant mass discharge (or flux) during the 2007–2012 pump and treat operation at the TIAA site. Zero time corresponds to Autumn 2007. CRA-5 and DP-1 are the two abstraction wells used during the operation (figure 3.8). From Brusseau et al. (2013).

3.5.2 Representation of the DNAPL source zone processes for the TIAA site

The upper sand and gravel aquifer is confined, with the piezometric head 29 m below ground level, in the clay aquitard material. The sand and gravel is 3 m thick, under 34 m of mostly aquitard material. There is not, from the literature, an obvious vertical migration pathway by which chlorinated solvent DNAPL could have migrated to the sand aquifer through such a thickness of impermeable material. The presence of chlorinated solvent in the sand aquifer probably indicates either vertical migration pathways such as fractures or leaky borehole casings, disposal into wells or a laterally discontinuous geometry of sedimentary layers (Lawrence et al., 2006).

The aquitard clay layers are known to contain chlorinated solvent mass. This mass is probably as DNAPL in fractures, dissolved into pore water or intercalated into clay minerals (Matthieu et al., 2013). However, the aim of this study is to show that the

source term framework can produce realistic solvent fluxes to groundwater, so the model is simplified by not attempting to represent the clay portion of the source zone. The data for solvent flux is taken from the CRA-5 abstraction well (figure 3.7), screened in the sand and gravel aquifer. And groundwater flow in the aquitard material is probably negligible so that, while TCE concentration may be high, advected contaminant flux will be small. Therefore, the significance of the overlying aquitard for the aqueous flux in the sand and gravel aquifer is simply to reduce the amount of contaminant mass that reaches the aquifer. This may be considered by modifying the infiltration flux to the sand and gravel aquifer, where infiltration flux may be thought of as usage rate, J_{import}, times infiltration fraction, f_{i}, assuming that both f_{w} and f_{sl} are 1 (see table 3.3). As such, the source term model for this study uses just one layer, representing the sand and gravel aquifer with an infiltration flux giving the input during the period of solvent spillage, thought to be between the 1940s and 1970s (assumed to be from 1945 to 1975 in this study). (An alternative approach would be to represent the aquitard layer explicitly in one or more layers but, because in this case there is no monitoring in this geological layer, this would result in unnecessary complexity.)

The horizontal location of the source zone is not constrained. Even if the spill and disposal locations were known, there is no guarantee that the DNAPL source zone in the sand gravel aquifer would be directly below these, as DNAPL migration exploits unpredictable pathways and considerable lateral migration may occur before a favourable vertical migration pathway is found. Therefore three different source zone locations are tested (figure 3.8). The relevant difference between each location is the proximity of the two pumping wells, especially CRA-5, which influences the size of the impact of the wells on the Darcy velocity through the source zone (section 3.5.3). A source zone that is closer to CRA-5 will be flushed more efficiently.

This chapter argues that there is no consensus on the most appropriate upscaled empirical model to apply to transient DNAPL dissolution (section 3.2.2). Three stock models (constant Γ, converging Γ and dual-domain pool and ganglia) are provided within the `DNAPL` package, though custom models may be programmed (section 3.3). For this study, these three models are each used and an attempt to find a good model fit to data is attempted for each DNAPL distribution and dissolution model at each trialled source zone location. For each dissolution model, a set of parameters is used, which are similar between the three models (table 3.5). The parameter sets are designed to have an intuitive physical meaning, rather than to be minimal and, as such, there are direct trade-offs within some groups of parameters, which is discussed further in section 3.5.4.

The data used for calibration is the contaminant mass discharge, or flux, to well CRA-5 (figure 3.7), which is inferred as the product of concentration and abstraction rate. It is assumed that there is negligible travel time between the source zone and well. The flux to well DP-1 is much smaller and the well is both more distant and a weaker abstraction. It is likely that the flux to DP-1 is caused by abstraction of the pre-existing plume and does not reflect source term decay during the pump and treat operation. It is assumed that the source zone is in the capture zone of CRA-5 during the pump and

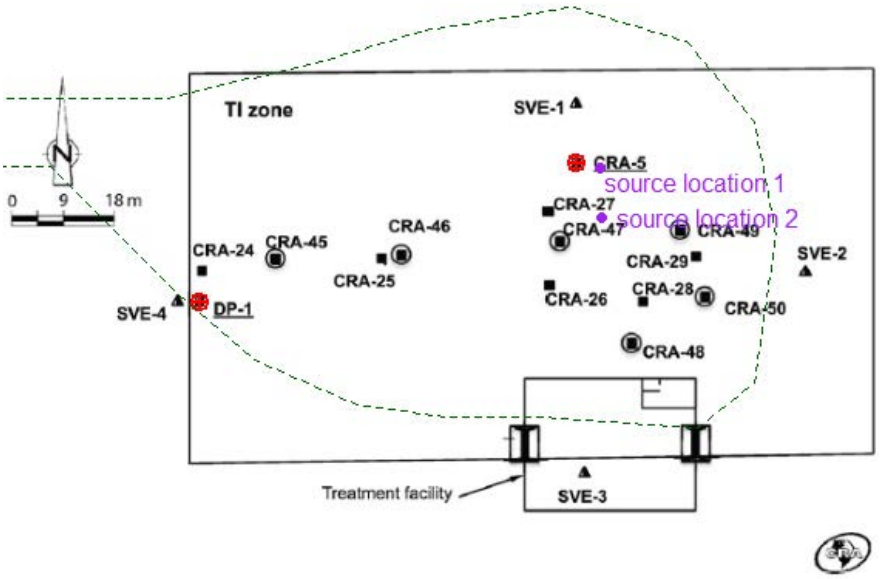

Figure 3.8: Map of the Central TIAA site, highlighting the two groundwater abstraction wells (red) and the three trialled source locations (purple). The observed 1000 µg l^{-1} TCE contour is also plotted. Adapted from Brusseau et al. (2013).

treat operation and therefore that the flux to this well shows the effects of transient source dissolution. Similarly, the earliest flux data for well CRA-5 is probably more related to pre-existing plume concentrations, so the calibrations place less importance on matching the low-concentration data from the first six months of the pump and treat.

After establishing which groups of parameters have very similar effect, some parameters are given set values. The remaining free parameters are varied for each of the nine source location-source dissolution model combinations until a satisfactory fit to data is achieved. No quantitative calibration goal was established *a priori* as it was unclear how good a calibration would be possible in each case. Instead, the best visual fit to data was sought in each individual case.

3.5.3 Transient flow rate through the source zone

There are two wells that are operated during the pumping test (figure 3.8). Aside from the groundwater flow induced by these wells, there is a background groundwater flow caused by the natural hydraulic gradient ∇h, for which a value of 5×10^{-3} is used with flow directed toward the west, inferring from the mapped plume (Brusseau et al., 2013). The magnitude of horizontal Darcy velocity through the source zone is determined by vector addition, with the components of Darcy velocity induced by a pumping well at position $\mathbf{r}_{\text{well}}$ given by:

Table 3.5: Parameters pertaining to the DNAPL distribution and dissolution models. All three models (constant Γ, converging Γ and dual-domain pool and ganglia) use similar parameter sets. Because some groups of parameters were found to have very similar effects on the source term result, some parameters are given set values in order to reduce the number of free parameters. Note that the dual-domain pool and ganglia model uses a constant Γ value for the dissolution of the pool and that the large maximum pool width of 100 m effectively imposes no limit on the amount of lateral migration of the pool.

parameter	set?	set value
infiltration flux, $J_{\text{import}} \times f_{\text{i}}$	no	
Γ or Γ_0	no	
porosity, ϕ	no	
ganglia width, w_{g}	yes	0.5 m
maximum pool width, $w_{\text{p,max}}$	yes	100 m
pool height h_{p}	yes	0.5 m
residual NAPL saturation, $S_{\text{r,n}}$	yes	0.1
residual water saturation, $S_{\text{r,w}}$	yes	0.1

$$\begin{aligned} \mathbf{q}_{\text{well}} &= \frac{Q_{\text{well}}}{2\pi \left|\mathbf{r}_{\text{SZ,well}}\right| b} \hat{\mathbf{r}}_{\text{SZ,well}} \\ \hat{\mathbf{r}}_{\text{SZ,well}} &= \frac{\mathbf{r}_{\text{SZ,well}}}{\left|\mathbf{r}_{\text{SZ,well}}\right|} \\ \mathbf{r}_{\text{SZ,well}} &= \mathbf{r}_{\text{well}} - \mathbf{r}_{\text{SZ}} \end{aligned} \tag{3.10}$$

where Q_{well} is the well's abstraction rate, $\mathbf{r}_{\text{SZ}}$ is the position of the source zone and b is the aquifer thickness. The vectors of the pumping well-induced Darcy velocities and the background Darcy velocity are summed to give the transient Darcy flow in the source zone (figure 3.9). A source zone closer to the pumping wells will experience a larger Darcy velocity during the pumping test. The background Darcy velocity is given by $K\nabla h$ (K is hydraulic conductivity, which is 9.46 m d^{-1} for the sand and gravel: Brusseau et al., 2013). Therefore, with the two wells CRA-5 and DP-1 operating, the Darcy velocity through the source zone is:

$$\mathbf{q} = K\nabla h + \mathbf{q}_{\text{CRA-5}} + \mathbf{q}_{\text{DP-1}} \tag{3.11}$$

3.5.4 Results and discussion

Although the list eight of parameters in table 3.5 seems to give many degrees of freedom with which to calibrate each DNAPL and source location test, in fact the parameters form groups that have almost identical impact on the result of the source term. For example, the ganglia width, pool height, porosity and residual saturations all control the flow-normal area that the source zone architecture presents to the groundwater

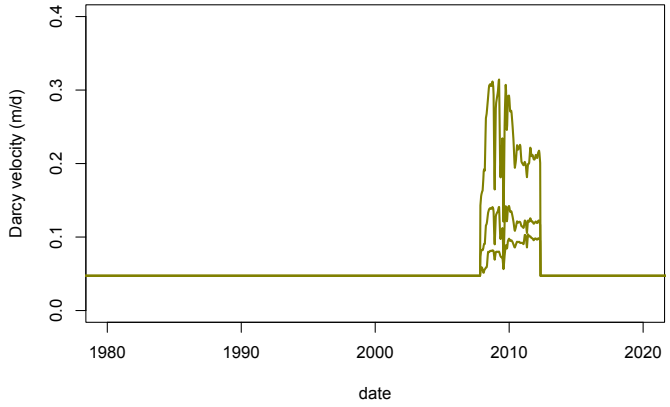

Figure 3.9: Magnitude of horizontal Darcy velocity for the source zones located as shown in figure 3.8 in the sand and gravel. Source zone location 1 experiences the strongest effect of abstraction and 3 the weakest.

flow. A higher porosity or lower residual water saturation allows DNAPL mass to be contained in a smaller aquifer volume, so that the flow-normal area and therefore effluent mass flux is smaller. Therefore, as recorded in table 3.5, the values of all but one of the parameters in this group (the porosity) are given set values which are realistic in the context but cannot be said to be constrained. The remaining parameter is then allowed to vary to achieve calibration. Similarly, the infiltration flux and the maximum pool width have similar effect on the source term result, by controlling the amount of solvent mass in the source zone by the end of the spillage period in 1975. These parameters have different effects at the early stage of the source zone, but there are no observations of this period. Therefore the maximum pool width is set at a value which effectively does not constrain the lateral migration of the pool because there is insufficient infiltrating solvent mass in all cases to exceed the maximum pool dimensions. The effect of varying Γ or Γ_0 is unique, with higher values giving rise to more effective reductions in flux from the pumping test. Parameters with non-unique influence on the source term cannot be said to be constrained by calibration. Both the pool height and the ganglia width, for example, have similar effects and multiple combinations of values for these parameters may give the same source term result.

Table 3.6 shows the parameter estimations arising from each source location-DNAPL dissolution model pair. A good calibration is achieved for each source location-DNAPL model pair using different parameter values. There is no available data pertaining to the early stages of the source term, so the estimated early source term flux cannot be said to be constrained. The results suggest that if the source zone was further from CRA-5, then the early stage source zone flux was probably higher. With regard to the parameter values in each calibration, the most obvious variation is in Γ or Γ_0, with more distant source locations preferring much higher values (1 is closest and 3 is furthest). Additionally, Γ_0 for the converging Γ model is generally higher because the

value of Γ in this model converges to 1 with progressive source zone depletion. The calibrated values of Γ, which are all above 1, indicate that the source term probably has favourable characteristics for a simple pump and treat operation, retaining high contaminant flux returns until not much source zone mass remains. The Γ and Γ_0 estimates need to be higher to achieve calibration for more distant source zones because the magnitude of influence on groundwater flow from the CRA-5 abstraction diminishes with distance. Lower Darcy velocity means that the effluent concentration needs to be higher in order to give the same effluent flux as the closer source zones and so Γ must be higher to sustain high effluent concentration several decades after the cessation of solvent spillage. The nine calibrated results are plotted in figure 3.10.

Table 3.6: Values of calibrated parameters for each source location with each DNAPL dissolution model.

DNAPL model and source location	infiltration flux (kg d^{-1})	Γ or Γ_0	porosity
constant Γ, location 1	1.29	2.3	0.065
converging Γ, location 1	1.48	8.0	0.058
dual-domain, location 1	1.30	2.0	0.050
constant Γ, location 2	2.17	7.5	0.045
converging Γ, location 2	1.53	80	0.090
dual-domain, location 2	1.30	7.0	0.080
constant Γ, location 3	2.15	11	0.050
converging Γ, location 3	1.40	150	0.100
dual-domain, location 3	1.90	11	0.052

With a fitted model of the TCE source term, it is possible to predict what would have happened without the pump and treat operation. This is done by re-running the source term model with the same parameters but without the variation in groundwater flow during the pump and treat operation. Figure 3.11 shows the predicted difference the pump and treat operation has had on the source term. Independent of the location of the source term, the source term model estimates that future effluent flux from the source term has been reduced by a little less than a factor of ten. Resulting abstracted concentrations will depend on abstraction rate. The model suggests that this has been an effective remediation operation because the source term decline is characterised by high Γ, giving efficient mass recovery from even an old source zone. However, the process of back-diffusion from TCE dissolved in clay pore water or intercalated in clay minerals has not been included in this source term model and may produce an additional persistent source that will be persistent later (Matthieu et al., 2013; Yang et al., 2015). This consideration highlights that different processes may be dominant at different stages of the source term and calibration of a source term to historical data is not a guarantee that all relevant processes for future water quality have been accounted for.

The TIAA site still has an extensive aqueous TCE plume that will mean that water quality improvements will not be immediate. However, DNAPL source zones are gen-

erally the most recalcitrant component of chlorinated solvent contamination and that this source of contamination has been significantly reduced is a beneficial result for the long-term water quality of the site and potential downstream receptors. Because of the non-uniqueness of the parameter fit, the experiment cannot be said to have fully characterised the source zone characteristics such as the dimensions of the DNAPL pool. However, the method was able to give a calibrated model of the transient evolution of dissolved TCE flux from the source zone which could be used to forecast future water quality risk.

This experiment shows that the DNAPL source term framework presented in this chapter, although empirical in nature, is able to capture realistic DNAPL source term characteristics, in particular dissolved solvent flux as it responds to transient groundwater flow and declines with time. The purpose of the framework is to generate realistic and consistent source term fluxes from multiple point sources at the regional scale, a context in which there will generally be less density of data available than there is for the TIAA site. However, the successful application of the framework to an intensely monitored contaminated land site gives confidence for its application to larger scale contexts, as in chapters 4 and 5. All three DNAPL distribution and dissolution models built into this chapter were able to give good calibrations with different calibration sets. This highlights that the most appropriate upscaled model to apply to DNAPL dissolution in aquifers is uncertain (section 3.2). A key component of the source term framework is flexibility with models as well as parameters, allowing this process uncertainty to be addressed.

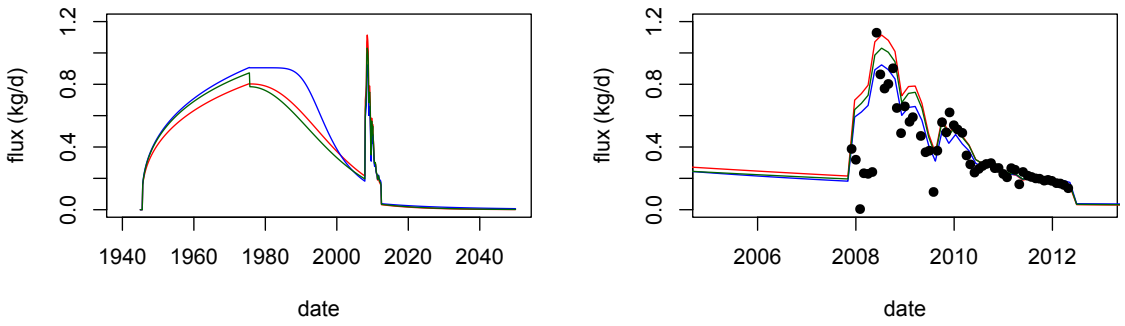

(a) source zone location 1

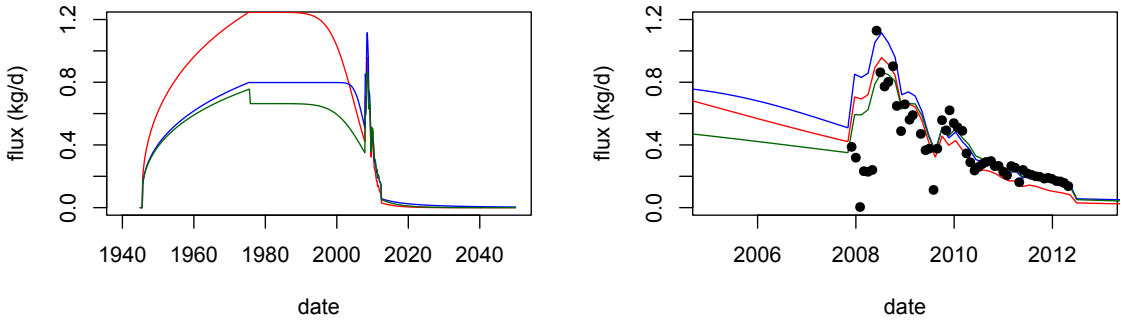

(b) source zone location 2

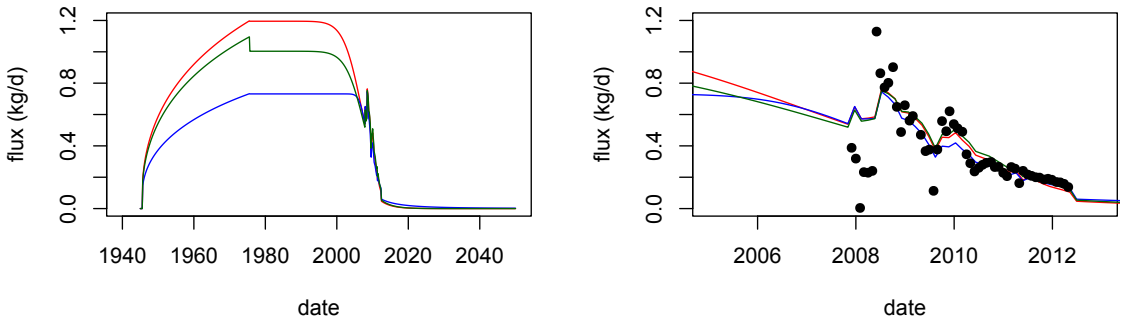

(c) source zone location 3

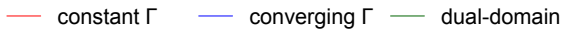

Figure 3.10: Source term flux results for each of the three trialled source zone locations, with each DNAPL dissolution model plotted. On the left the entire time period of the modelled source term is plotted; on the left the period of the pump and treat is plotted, with CRA-5 contaminant mass discharge data shown as solid black circles.

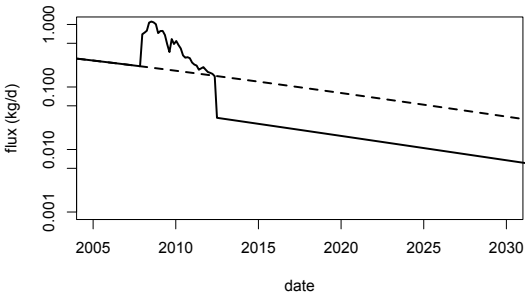

(a) source zone location 1

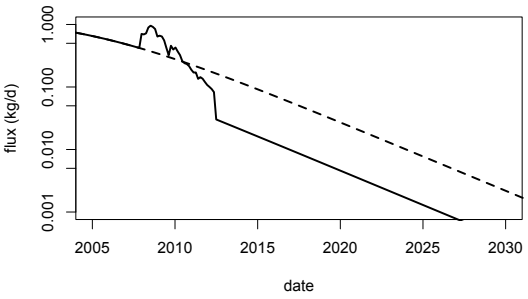

(b) source zone location 2

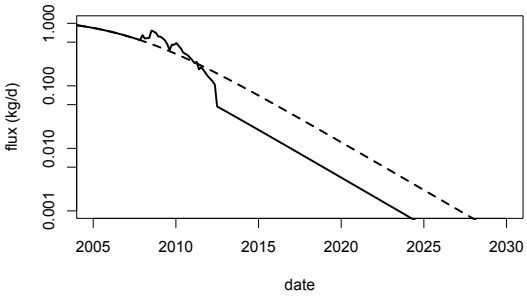

(c) source zone location 3

Figure 3.11: Impacts of the pump and treat on the future source term flux with the constant Γ DNAPL dissolution model (note log scale). Solid lines are the calibrated source terms with the influence of the pumping wells. The dashed line shows the equivalent models but with no pumping wells and so indicate the predicted source term flux if there had been no pump and treat operation.

Chapter 4

THE DNAPL SOURCE TERM FRAMEWORK AND DRW TRANSPORT SOLUTION APPLIED TO A REGIONAL AQUIFER CONTEXT: THE TAME VALLEY

The dense non-aqueous phase liquid (DNAPL) source term model framework (chapter 3; Barry, 2017b), allows water resource managers to generate regionally consistent source terms for the distribution of historic chlorinated solvent point sources. This DNAPL source term responds to transient variation in groundwater flow and is programmed to account for spatial variation in aquifer thickness. Source zones may dissolve at different rates in different locations, but they can be described with a consistent set of source term input parameters. This method has been demonstrated for a contaminated land site (section 3.5). However, this chapter demonstrates the usage of the DNAPL source term model framework in regional context. The Tame Valley region, Witton, Birmingham is chosen because this region has been surveyed on multiple occasions spanning about 20 years, giving a reasonable transient record, although the number of wells available for monitoring has declined with time (Rivett, 1989; Taylor, 1998; Botha, 2006; Murcott, 2008). Nonetheless, compared to the study in section 3.5, there is a lower density of data and greater uncertainty about spill locations.

The purpose of this investigation is primarily to show that the DNAPL source term model framework, in conjunction with the dynamic random walk (DRW) transport code (section 2.5; Barry, 2017c) is suitable for regional-scale water quality management and forecasting. Results from this experiment may also give transferable observations about the behaviour of chlorinated solvents in UK Permo-Triassic sandstone aquifers and give specific local insight into the present and future groundwater quality of the Tame Valley region.

The groundwater quality of the Tame Valley has been extensively studied, including in the aforementioned surveys, since the problem of widespread chlorinated solvent contamination was first being recognised. Findings about the contaminant hydrogeology of the region are overviewed in section 4.1, including a presentation of observed concentrations in figures 4.2, 4.3 and 4.4. Some studies have sought to increase understanding of the processes occurring in the area (and the Birmingham Aquifer more broadly), both natural and anthropogenic, through numerical modelling. These models of the Birmingham Aquifer are summarised in section 4.2. This section also details the extraction of a submodel using telescopic mesh refinement (TMR) for use in this study. In light of the body of previous research, the organic contamination of the Tame Valley sandstone aquifer presents a rare opportunity to compare model predictions with observations spanning a long period of time and covering a several different boreholes at the regional scale. Sections 4.3 and 4.4 describe the modelling experiments, involving a progressive stochastic calibration, conducted as part of this study to gain insight into likely source term and transport parameters and to test the use of the DNAPL source term model framework. Finally, section 4.5 presents results from the experiment, and section 4.6 states the arising conclusions and an evaluation of the methods used.

4.1 Hydrogeology and Organic Groundwater Quality in the Tame Valley

4.1.1 Geology and hydrogeology

The Tame Valley region lies on the Birmingham lower Permo-Triassic Sandstone aquifer, Kidderminster Formation. The Birmingham aquifer as a whole is composed principally of, in increasing order of age, the Bromsgrove Formation, Wildmoor Formation and Kidderminster Formation sandstones. The aquifer is divided into two halves by the Birmingham fault, east of which the aquifer is confined under the younger Mercia Mudstone Group, and west of which it is unconfined. The fault acts as a barrier to flow in some locations, but in other parts, notably close to the Tame Valley area, cross-fault flow is inferred from hydrochemical correlation (Jackson and Lloyd, 1983; Ford and Tellam, 1994; Daily and Buss, 2013). The Tame Valley region, Witton is west of the Birmingham fault. The River Tame is the most significant watercourse in the area in terms of control on the groundwater system and apart from the local influence of abstractions, groundwater flow is directed toward the River Tame.

The Kidderminster Formation consists of medium to coarse grained sandstone (British Geological Survey, 2016). In Birmingham, apart from a thin and discontinuous breccia layer (Hopwas Breccia: Daily and Buss, 2013), it forms the basal unit of the Permo-Triassic sandstone sequence, unconformably overlying the Carboniferous sequence, of which most units are aquitards and minor aquifers. The Permo-Triassic Sandstone of the West Midlands is not very cemented in general, enhancing both hydraulic conductivity and porosity (Allen et al., 1997; Daily and Buss, 2013). The Kidderminster Sandstone is, however, more heterogeneous than other sandstones in the area, with some much lower values of hydraulic conductivity and porosity (tables 5.1 and 5.2, Allen et al., 1997). Fractures are also likely to have a significant local effect on the hydraulic conductivity at local scales, although their influence at the scale of the whole study area is probably of limited extent (Allen et al., 1998). Due to the influence of fractures, the scale of measurement is a very important factor in the determination of hydraulic conductivity. Allen et al. (1997) record hydraulic conductivity values for the Kidderminster Formation significantly higher for pump test results, than for core samples (table 5.1). Large faults may have hydraulic influence on a greater scale of measurement than pump tests.

Variation in abstraction demand has had a large impact on the water table distribution through time in the Birmingham Aquifer. Major abstraction of groundwater for public water started around 1860 though all but ceased by 1905 because of the construction of the Elan Valley Reservoirs which have supplied the entire public water demand for the city until this day. However, abstraction demand did not have much of a hiatus because industrial use of groundwater increased and continued to do so until shortly after the second world war, whereupon abstraction declined and the water table rebounded (Greswell, 1992; Rivett et al., 2012). The water table rise has approached pre-abstraction level, causing concerns about flood damage to building

basements (Greswell, 1992; Lerner and Barrett, 1996), but on the other hand opening the opportunity of a plentiful water resource that could be used (Rivett et al., 2012).

Rainfall in the area is typically between 600 and 800 mm a^{-1}, with a fraction of that infiltrating to the aquifer as direct natural recharge, depending on land use and properties of superficial deposits. Through detailed spatial analysis and classification of these factors, Thomas and Tellam (2006) estimate an average recharge across the unconfined Birmingham Aquifer of about 100 mm d^{-1}, though this is very heterogeneously distributed in reality and excludes leakages from water mains and sewers, probably a significant input. A perhaps surprising suggestion from Thomas and Tellam (2006) is that paved areas do not seem to impede recharge more than vegetated areas, in fact often less so, perhaps attributable to their assumption that evapo-transpiration does not reduce recharge in paved areas. Considering, in addition, losses from the water distribution and sewer network, it is probably unfounded to assume that recharge in the Tame Valley is less than what it would be under natural conditions just because of it being urbanised.

4.1.2 Chlorinated solvents in groundwater

Birmingham has a long industrial history, growing rapidly with the arrival of the canal and railway networks in the early nineteenth century (Ford and Tellam, 1994). Early on, Birmingham's industries were predominantly metal working, tanning and gas works, but by the early twentieth century industrial activities had diversified to include car manufacture, jewellery and armaments, among others (Taylor, 1998; Rivett et al., 2005). Industrial growth was initially focussed in Smethwick and around the city centre. Later growth took factories north toward the Tame Valley and Aston and south to Selly Oak and Longbridge (Ford and Tellam, 1994; Rivett et al., 2005). Many of Birmingham's industry types, particularly those involved with metal working and engineering, became associated with the use of chlorinated solvents in the twentieth century (Department of the Environment, 1995). The Tame Valley industrial area has been dominated by large factories for armaments manufacture and to the west, a large dry cleaning site, an industrial activity that has been particularly associated with tetrachloroethene (PCE) (figure 4.1). There is significant localised PCE contamination close to the dry cleaning site (figure 4.2). Although remediation operations have been in place in the region, the contaminant sources have proved hard to remediate and reductions in aqueous concentration have been modest where seen at all (Taylor, 1998; Rivett et al., 2005).

Section 3.1.1 gives an overview of the organic groundwater quality surveys carried out in the Birmingham aquifer over the course of three decades, all of which have especially good coverage of the Tame Valley region. The collated results through time of these surveys are shown in figures 4.2, 4.3 and 4.4. Solvent spillages onto the ground were thought to have mostly stopped by the late 1990s due to environmental regulation and awareness (EPA, 1990; Taylor, 1998). However, it is clear that solvent contamination of groundwater in the region has persisted well beyond that period. Instead of a

decline that might be expected, 1,1,1–trichloroethane (1,1,1–TCA) and PCE concentrations have actually increased through successive surveys for the same boreholes and trichloroethene (TCE) has been roughly stable (Taylor, 1998; Botha, 2006; Murcott, 2008). However, the number of available boreholes for sampling in the Tame Valley region has been steadily declining, making cross-comparison difficult. Reasons for the continued increase or at least lack of decline in solvent concentrations are attributed to the persistent nature of DNAPL source zones, especially solvent pools, although in individual cases local factors such as abstraction changes may play a significant role. 1,1,1–TCA was generally adopted later by UK industry (figure 3.4), which may give expectation of a later observed concentration peak, as seems to be observed. PCE is observed in large concentration at just one well in the Tame Valley surveys, with a large increase between the 1988 and 2006 surveys (Botha, 2006). Comparisons of concentrations at close-by boreholes with different screened intervals suggest that the depth penetration of contamination is limited, at least in some cases (Botha, 2006). Section 3.1.2 discusses how the depth penetration of DNAPL may depend on the physical properties of the chemicals, geological heterogeneity as well as the quantity of DNAPL spilt. Ford and Tellam (1994) also conducted an inorganic water quality survey of the Birmingham aquifer. Although correlation between contamination and industrial land use was observed, inorganic water quality was generally better than organic and showed little correlation to the organic water quality, an observation they attribute to different chemical properties as the Permo-Triassic sandstone in Birmingham is much more sorptive of ionic substances than organics.

Degradation products from chlorinated solvents, cis–1,2–dichloroethene (cis–DCE) in particular, were also observed to be increasing (Murcott, 2008). This is perhaps surprising, because the biotic dechlorination of TCE to 1,1–dichloroethene (1,1–DCE) only occurs significantly in reducing conditions that do not characterise the Birmingham Permo-Triassic Sandstones (Rivett et al., 2007, 2012). Botha (2006) speculates that aquifer heterogeneity may play a role, with locally reducing zones causing this degradation pathway. Anthropogenic deposits and contaminants often create anaerobic conditions in shallow aquifer horizons (Rivett et al., 2005, 2012). Ellis and Rivett (2007) also found significant degradation product concentrations in groundwater discharging to the River Tame. However, lack of depth variation in concentrations observed from multi-level piezometer samples strongly suggested that degradation in riverbeds is not very significant and that these products would have been formed upstream in the aquifer. Conversely, however, organic carbon content of the riverbeds (where mean organic carbon fraction on the sediment, $f_{\mathrm{oc}} = 0.011$) is much greater than that in the aquifer (mean $f_{\mathrm{oc}} = 4.9 \times 10^{-4}$), resulting in a much greater retardation of dissolved solvents in the riverbed than in the sandstone aquifer. Retardation factor R_f is probably close to 1.4 for 1,1,1–TCA, 1.3 for TCE and 1.8 for PCE which is a more hydrophobic (Taylor, 1998; Ellis and Rivett, 2007). These values are estimated based on octanol-water partition coefficients (K_{ow}) from Piwoni and Banerjee (1989) and Taylor (1998), using the Piwoni and Banerjee (1989) empirical model relating K_{oc} to K_{ow} and the equation for retardation factor $R_f = 1 + K_{\mathrm{oc}} f_{\mathrm{oc}} \rho_{\mathrm{b}} / \phi$, with ρ_{b} and ϕ representing dry bulk density and porosity respectively. Results from these surveys appear to suggest that,

while significant degradation products have been found in groundwater samples, the degradation processes are probably confined to shallow aquifer layers where conditions are reducing due to anthropogenic influence with degradation in the wider aquifer minimal. 1,1,1–TCA may, however, abiotically degrade to 1,1–DCE in the wider aquifer. Abiotic hydrolysis and elimination of 1,1,1–TCA to acetate and 1,1–DCE respectively both occur in aerobic or anaerobic conditions. The half-life of 1,1,1–TCA from both these reactions together is 1–4 years at 20°C and around 10 years at 10°C (Scheutz et al., 2011). Slight sorptive retardation is likely, with $R_f < 2$, as the aquifer material has little organic material, although this property may well vary heterogeneously. Retardation factor and first order degradation half-life are tested stochastically in this study (section 4.3), with parameter ranges informed by these findings.

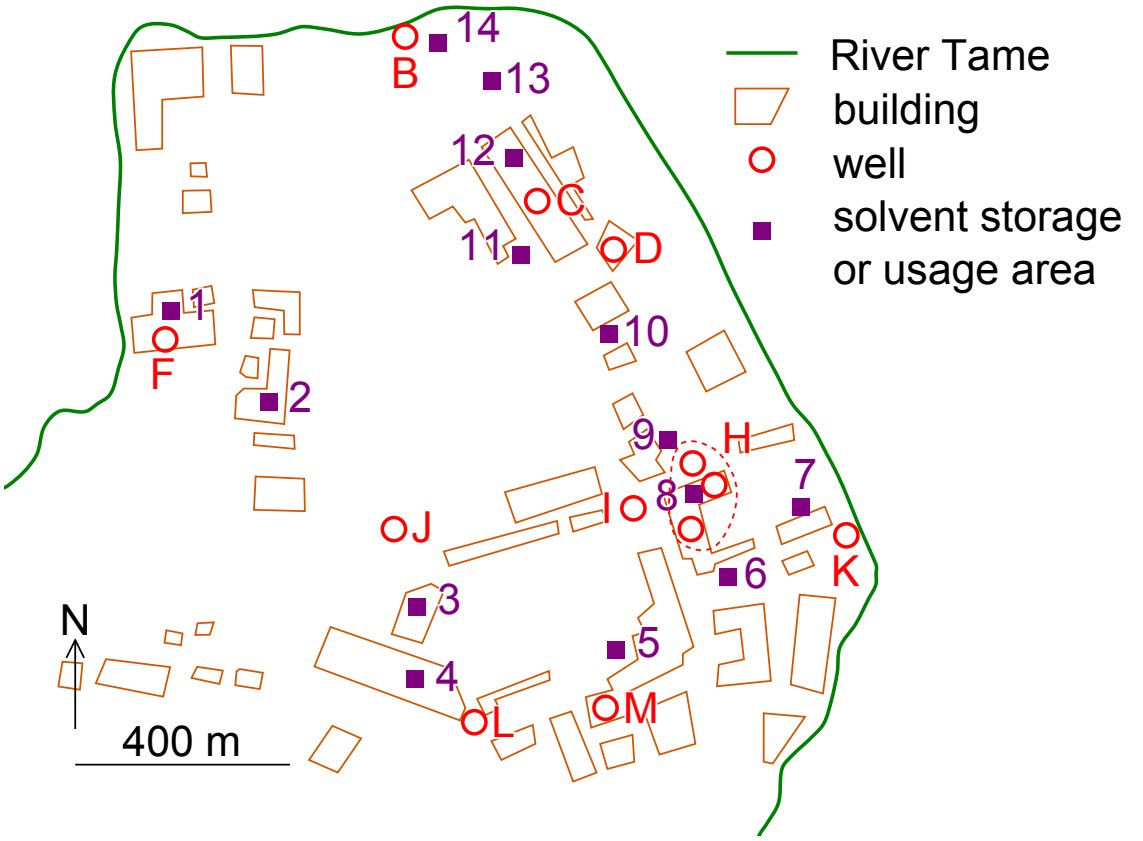

Figure 4.1: Map of industrial buildings, abstractions and likely solvent storage and disposal locations in the Tame Valley on the south side of the river, adapted from historical maps by Rivett (1989). Labels are attached to wells and potential source locations (solvent storage and usage locations) for reference in this study. The wells labelled H are aggregated because individual abstraction rates for each are not available and they are represented as one in Daly's (2005) model. The map is relevant for 1989 and more recent wells are not included in the figure.

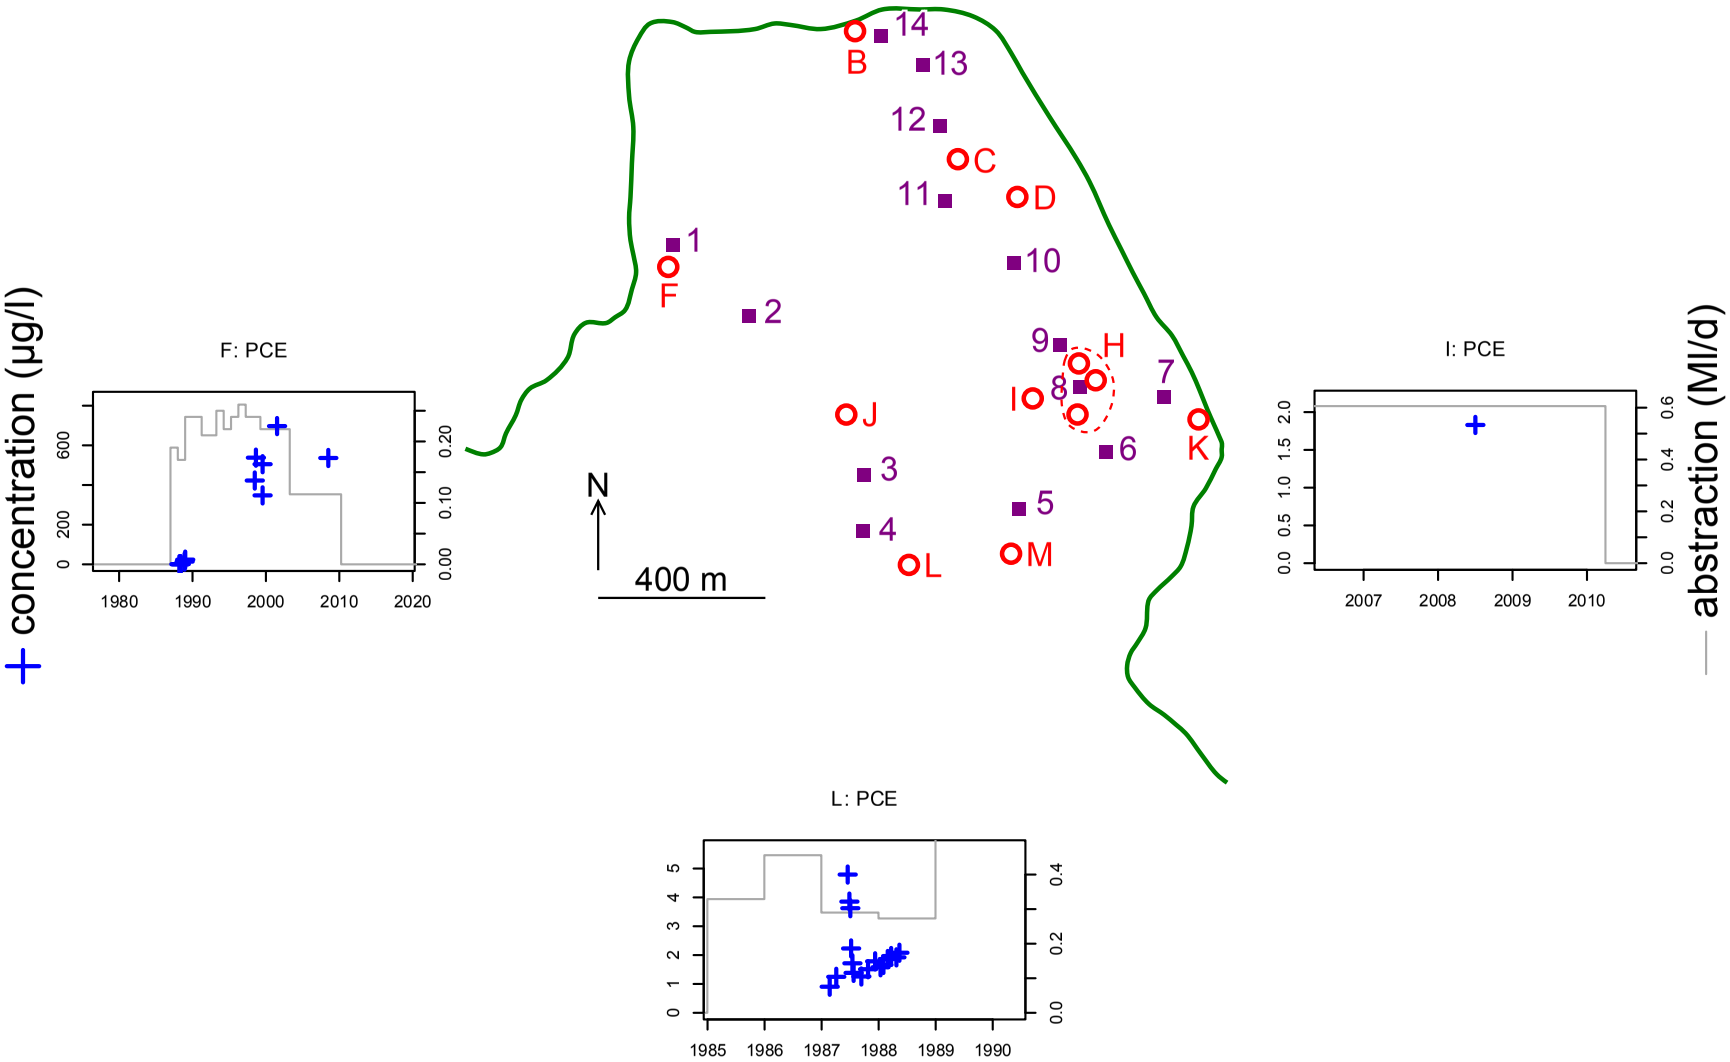

Figure 4.2: Recorded PCE concentrations in Tame Valley wells. Note different scales for each plot. A feature legend is in figure 4.1.

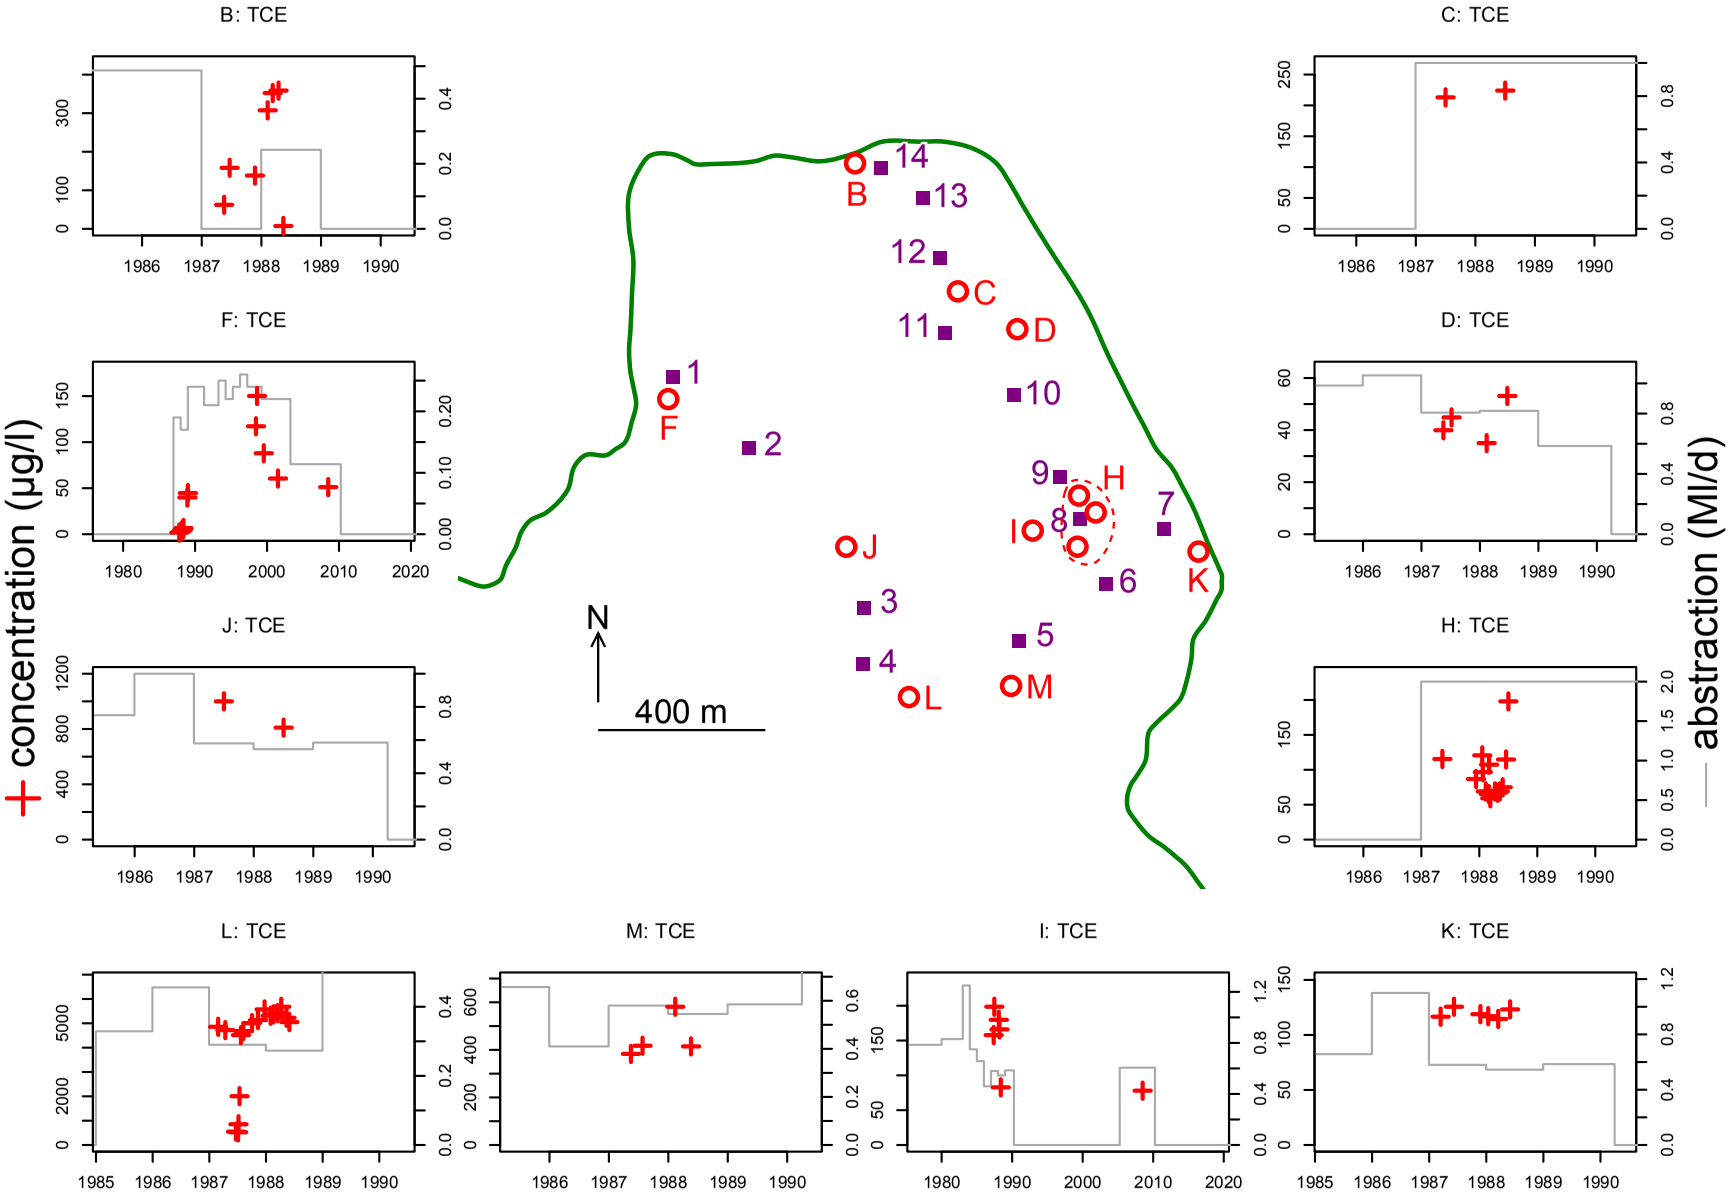

Figure 4.3: Recorded TCE concentrations in Tame Valley wells. Note different scales for each plot. A feature legend is in figure 4.1.

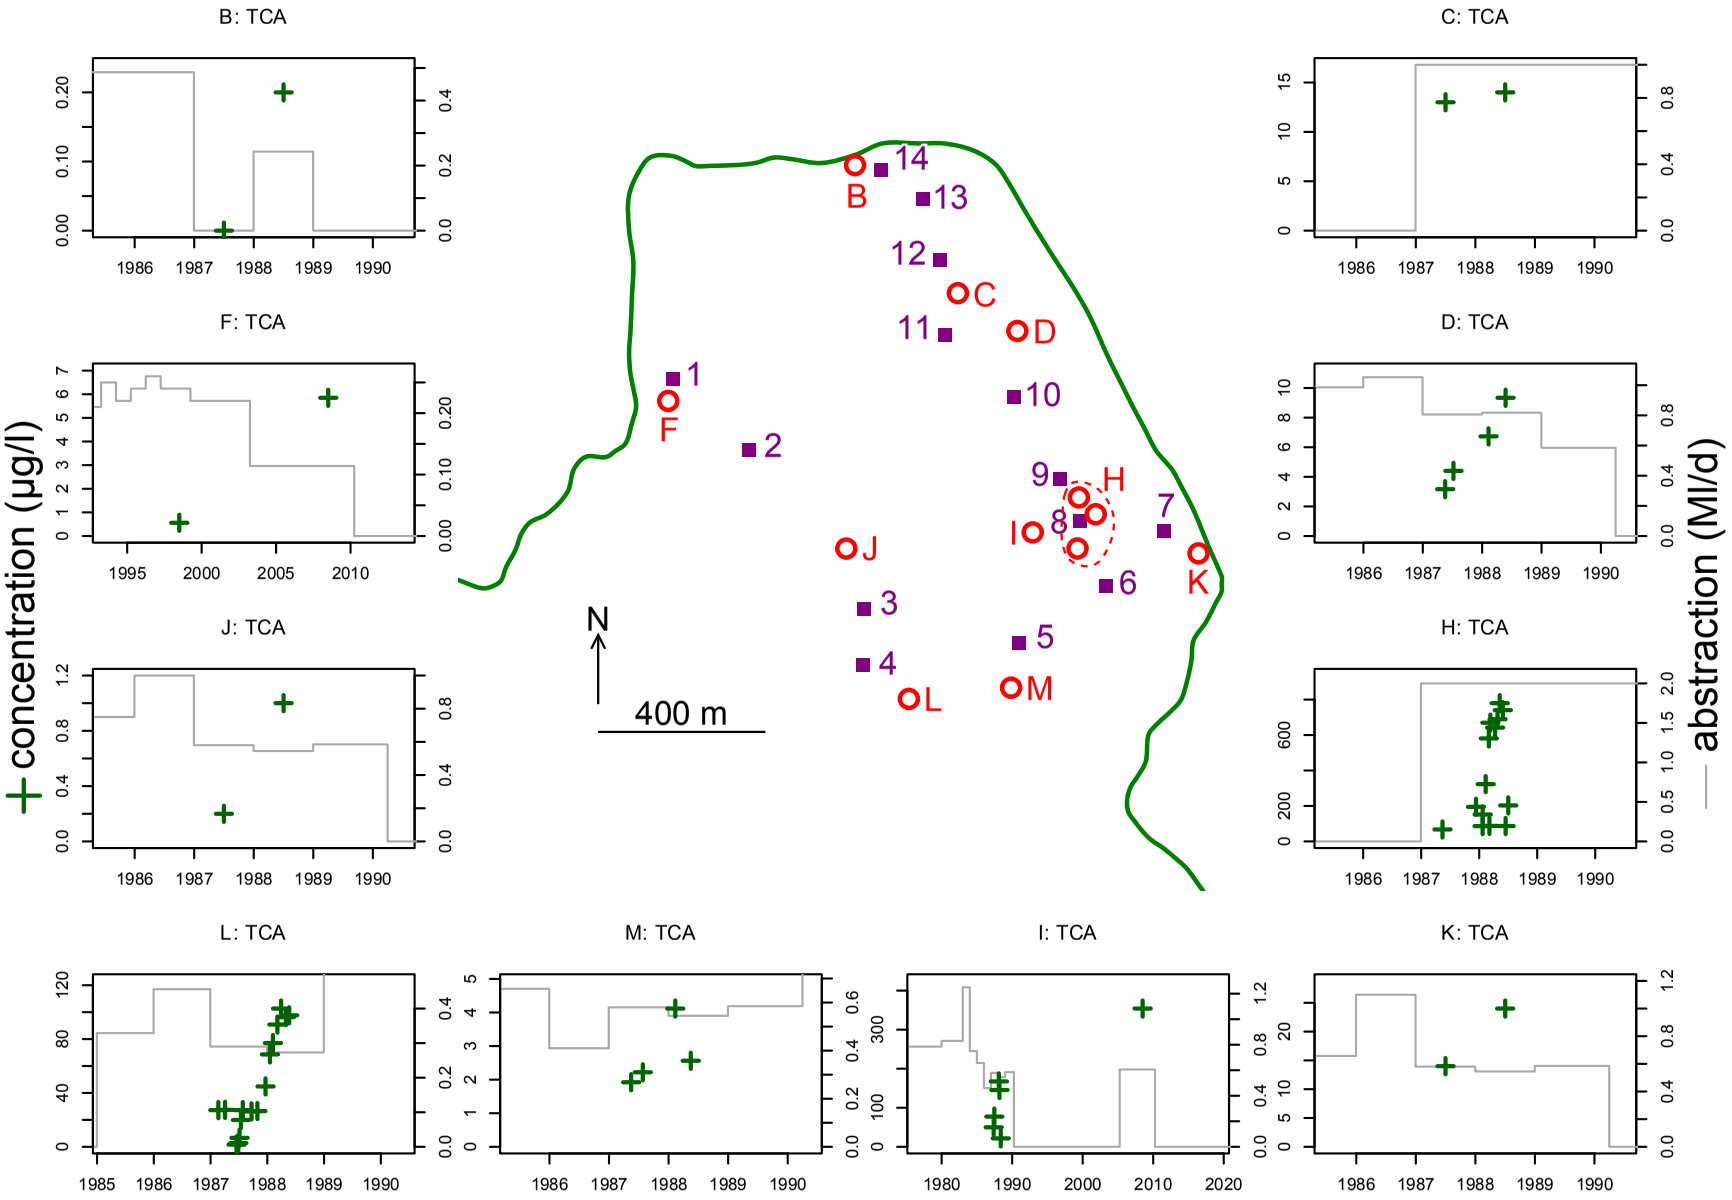

Figure 4.4: Recorded 1,1,1–TCA concentrations in Tame Valley wells. Note different scales for each plot. A feature legend is in figure 4.1.

4.2 Tame Valley Groundwater Model

4.2.1 Groundwater models of the Birmingham Aquifer

Three major modelling efforts for the Birmingham aquifer have been undertaken in the past, for different purposes. Greswell (1992) had the goal of forecasting further groundwater rise following the gradual decline in abstraction following 1945, which could have threatened local structures and potentially caused flooding. Daly (2005) focussed on the migration of contaminants in the changing flow regime of the Birmingham aquifer. The regional model developed for the Environment Agency (EA) focussed on the sustainable future use of the aquifer, aiming to enable informed resource licensing strategies while controlling groundwater head and protecting groundwater-dependent ecologies (Daily and Buss, 2013). Although later studies have referenced the Greswell (1992) model (Daily and Buss made no reference to Daly), each model development has been largely independent and this highlights that the purpose of a model is very important for the initial stages of model design and that repurposing of groundwater models is often more difficult than starting over.

Nonetheless, there are large consistent features of all three models, both in the inputs and outputs. The Birmingham Fault is a major divide between the confined and unconfined aquifers: Greswell used low transmissivity cells for this feature and Daily and Buss, similarly, used MODFLOW's horizontal flow barrier (HFB) module whereas Daly emplaced a head-dependent boundary along the fault and excludes the confined aquifer from the model region, which anyway has sparse data coverage. The River Tame was represented by a river boundary condition by Greswell, part of the stream boundary condition output from 4R by Daily and Buss and as a constant head boundary by Daly, who does not include aquifer north of the river in her model. Daly's assumption that there is no groundwater flow passing under the River Tame corroborates the modelled water table from Daily and Buss, where the River Tame is at a minimum of groundwater head in both sandstone layers of the model, and this is also the general picture given by hydrographs across the unconfined aquifer. However, for a more local-scale study of contaminant migration in the Tame Valley region, the implicit assumption that there will be no contaminant migration past the river must be acknowledged. Greswell adopted a distributed empirical approach to the recharge estimation, considering only direct recharge, by using attenuation factors for different classes of drift and urban cover, which were tested during calibration. Additionally, by analysing water distribution network data, he was able to estimate losses from water mains and sewers and add that recharge contribution to the model. Daly adopted this approach though slightly modified as she used a separate layer to represent groundwater flow in the drift deposits. Daily and Buss employed the 4R routed recharge and run-off numerical code which is designed as a pre-processor to MODFLOW for determining recharge and stream inputs (Heathcote et al., 2004). Abstractions were input into each of the models principally according to operators' abstraction returns held by licensing authorities and is generally considered to be a well constrained stress.

The discretisations of the models are very different however. Greswell (1992) uses just one layer for the sandstone, largely because at that stage there was insufficient testing data available to assign distinct properties to different sandstone horizons. Daly (2005) and Daily and Buss (2013) both divide the sandstone aquifer into two layers for the numerical model. However, Daly does so in order to define a depth-varying contaminant source, whereas Daily and Buss distinguish the hydraulic properties of the Bromsgrove Formation from the Wildmoor, Kidderminster and Hopwas Formations. Greswell and Daily and Buss, concerned with regional head distribution and flow balance, use very coarse grids (cell dimensions 250–1000 m and 200 m respectively). However, Daly uses variable refinement down to 4 m $\times$ 4 m grid cells, attempting to eliminate artificial oscillations from the transport solution (she was not completely successful in this and based her transport conclusions principally on particle-tracking results). Because of the grid refinement and the purpose of the model, Daly's model has the most precise locations of abstractions, of key importance for precise estimation of abstracted concentrations especially where contaminant sources are very close to wells (figure 4.1). Additionally, the coarse grid size used by the other models means that at the scale of the Tame Valley region, some nearby wells are aggregated into the same cell.

All the models are calibrated against groundwater level hydrographs and Daily and Buss attempt to match stream flow outputs from the stream package to stream gauge data. Figure 4.5 shows how the model results compare to three hydrographs, one of which is in the Tame Valley region. The IMI borehole, in the Tame Valley, seems to agree well with all three model results but the models disagree with each other about the head during the 1960s and 1970s, when the aquifer was intensely abstracted. The data from Constitution Hill (figure 4.5c), somewhat south of the Tame Valley, goes back further in time and corroborates best with Greswell's and Daly's results. The results of Daily and Buss do not reproduce the observed decrease in head during the 1960s and 1970s. Daily and Buss use a slightly greater value for specific yield in the Tame Valley region, but this is unlikely to account for the scale of the difference. All three models agree on a minimum head shortly before 1970, when abstraction was at its peak (Greswell, 1992), but disagree about the value of the minimum. However, all three models are roughly consistent with respect to unpumped head values in the Tame Valley region, to which heads are returning or perhaps have already done so. This value is around 92 metres above ordnance datum (mAOD), controlled strongly by the stage of the River Tame.

The study in this chapter uses a flow model based on that of Daly (2005), principally because the model was built for the most similar purpose (namely, contaminant transport) and the grid was refined such that locations, especially those of abstraction wells, are much more precise than for the other models. The calibration to historic groundwater level data in the Tame Valley region gives confidence that this is an appropriate model to use. The approximation of the River Tame as a fixed head is a significant assumption and not ideal for a study focussed on the Tame Valley area. One effect will be the misrepresentation of vertical flows in the vicinity of the river. Additionally, drawdown at well B (figure 4.1) may be underestimated as, in the model, it draws much of its water from the very nearby fixed head. However, given observed head dis-

tributions as discussed above it is probably true that groundwater flow passing under the river in either direction is less significant than groundwater discharge to the river.

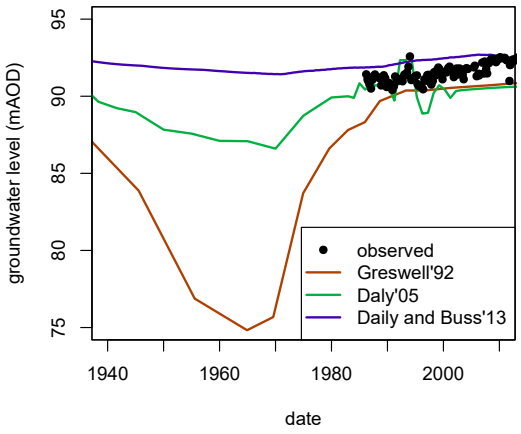

(a) IMI borehole, in the Tame Valley

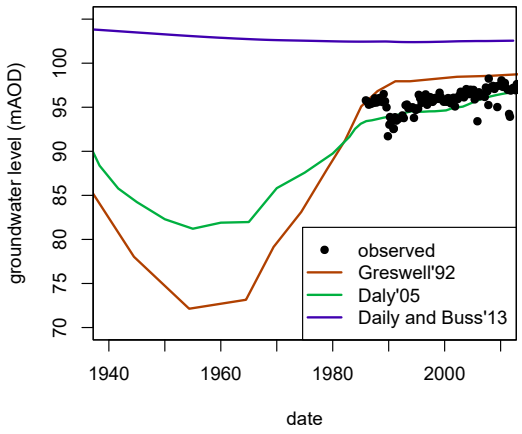

(b) HP Foods, or Tubes B borehole

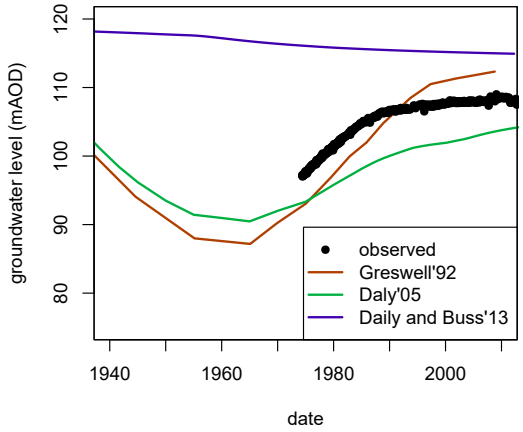

(c) Constitution Hill borehole

Figure 4.5: Comparison of modelled and observed hydrographs showing the calibration of the three Birmingham groundwater models discussed in this chapter. Note different head axes.

4.2.2 Sub-model extraction

The groundwater model is adapted from the MODFLOW (Harbaugh et al., 2000) model built for the Birmingham sandstone aquifer by Daly (2005). A nested model was extracted using the TMR tool available within Groundwater Vistas (Rumbaugh and Rumbaugh, 2011), over the relevant area, measuring 2688 m × 2016 m (figure 4.6). The resultant grid was in fact given a coarser grid than the original, with uniform 32 m × 32 m cells. The top layer was removed as this layer was is almost entirely dry in the Tame Valley region of the model and the TMR-determined fixed head boundaries used to specify the input from the larger model had head values below the

layer's bottom elevation. The constant fixed head boundary used to represent the River Tame required some manual correction as the automatic coarsening of the grid caused some gaps at the boundary. Comparing the water balances for the submodel and the equivalent region of Daly's model shows that the TMR extraction was successful (figure 4.7). The model time steps are the same as the original except with seven further year-long stress periods, up to 2011, and an extended final stress period, with the resulting model period from 1850 to 2050. Abstraction had mostly ceased in the Tame Valley region by 2004, but two abstractions remained or became active after 2004, whose abstraction rates are inferred from Murcott (2008) (abstraction returns were not available from the EA because the information is confidential).

Daly (2005) built and calibrated her model in two stages. She first built a steady-state model for the pre-abstraction regime, which she matched approximately to the previous simulation by Greswell (1992). Her transient model then included historic abstraction records (provided by the EA) and was calibrated against observed head data from eight monitoring wells. By modifying hydraulic conductivities, recharge attenuation through different surface geologies, storage properties and the conductance applied to the Birmingham fault, Daly (2005) was able to achieve a good match to data, with residual mean from simulation to observation within 5% of the total head range and residual standard deviation within 10% of total head range. Daly (2005) prioritised a good calibration in the Tame Valley region, where she also conducted contaminant transport studies. For the contaminant transport studies Daly (2005) used a constant concentration source zone, with MT3D (Zheng, 1990), but she had difficulty with artificial oscillations and numerical dispersion in the solution. The submodel extracted in this study is able to use a coarser grid because the DRW method, a fully Lagrangian method, is used to solve contaminant transport (chapter 2; Barry, 2017c).

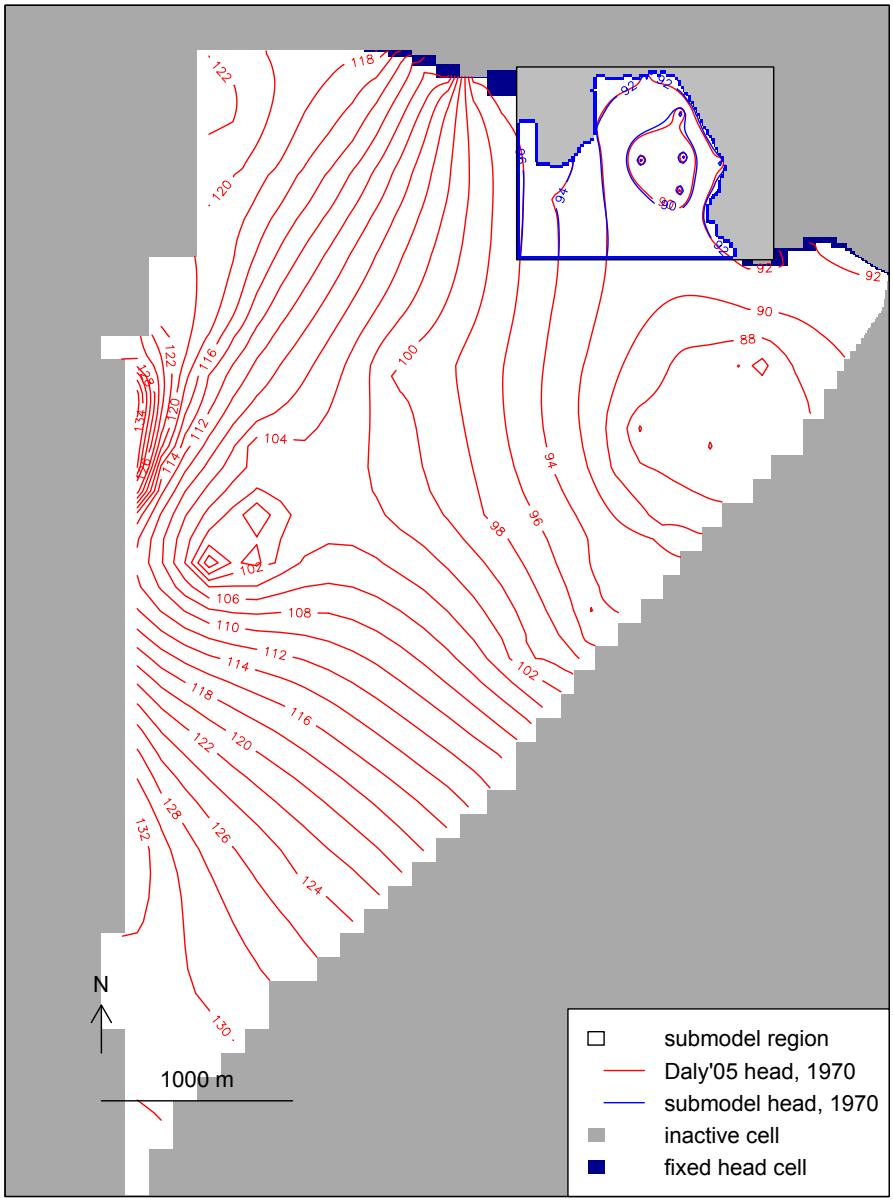

Figure 4.6: Submodel window extracted from larger-scale Birmingham model (Daly, 2005).

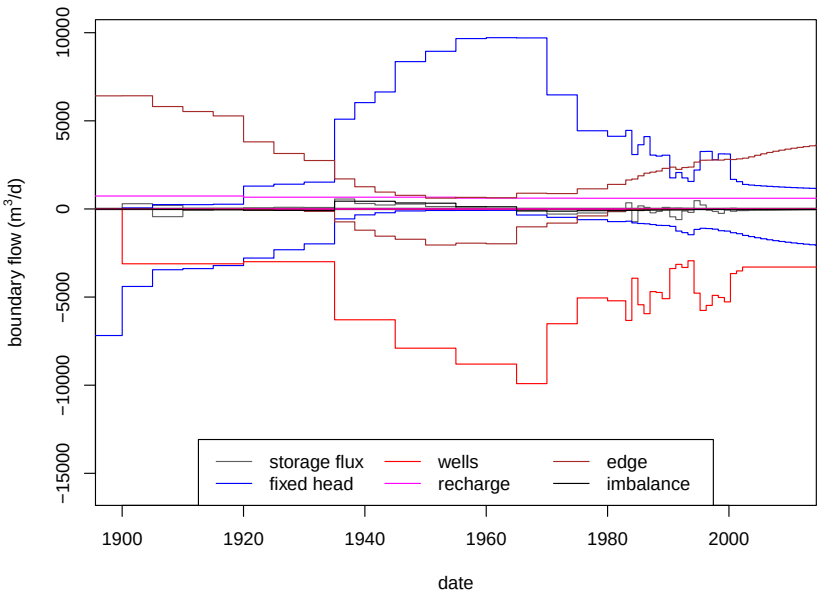

(a) region of submodel within larger model

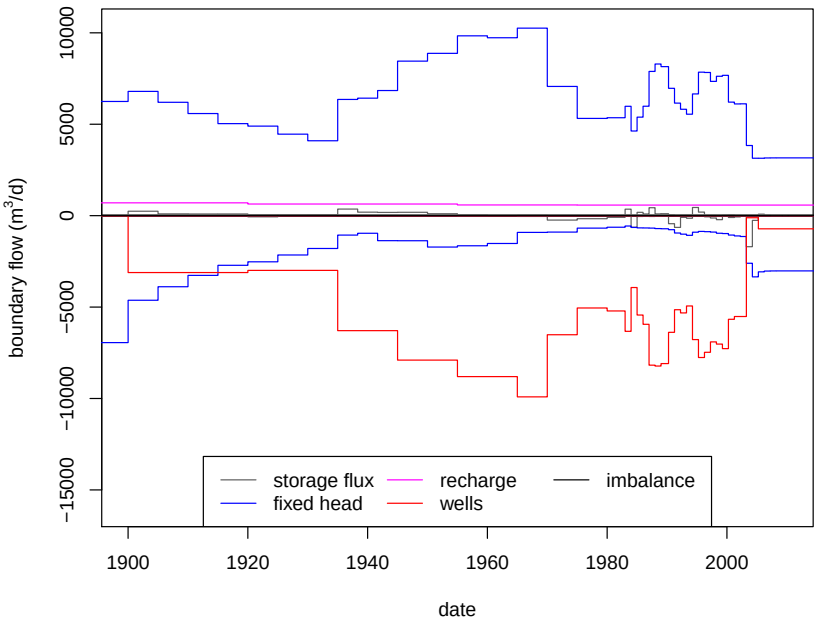

(b) submodel

Figure 4.7: Volumetric water flux balance, comparing (a) the submodel region out of the full Daly (2005) model and (b) the submodel itself. Inflows and outflows are separated, with negative values indicating outflows. The flux through the edge of the region in (a) is reproduced by transient fixed head boundaries along the edges of the submodel in the TMR method. There were some modifications made to abstraction wells after 2003 in the submodel.

4.3 A Tame Valley DNAPL Source Term and Aqueous Transport Model

The flux of dissolved mass from chlorinated solvent DNAPL is highly uncertain and has been the subject of many intensive studies at contaminated land sites. These studies are sometimes successful in characterising mass-flux relationships (for example: Brusseau et al., 2007; Rivett et al., 2014), but it has proven hard to infer source zone characteristics that are transferable from one source zone to another (DiFilippo and Brusseau, 2008). Clearly, in historically heavily industrialised centres an intensive source zone investigation at each former solvent-user site, which could take several years to produce meaningful results (for example: Brusseau et al., 2013), is not feasible. Additionally, many former solvent-users have since closed with the land now repurposed, precluding most site investigation, let alone remediation. Information on spills from these sites is inaccessible and any representation of their contribution to the regional contamination will need to be through a generalised source term model. Risk management tools that have attempted to quantify the persistent chlorinated solvent threat at the regional scale have tended to use a simple recharge-leached source term (Tait et al., 2004a; Troldborg et al., 2008). Although sorbed and aqueous phase solvent contamination may reside in the vadose zone after infiltration of the free phase, vaporisation of chlorinated solvents, which are usually volatile, is much more rapid than aqueous dissolution and as such the vadose zone contamination is short-lived compared to the phreatic zone contamination, which will be the more significant component of old sources and, importantly, not directly related to recharge.

Therefore, a recharge-leached source term representation cannot provide a satisfactory approximation for the solvent DNAPL source term. But neither does the body of research conducted into source zones at solvent-contaminated sites offer a generalisable DNAPL source representation that can be used at sites with minimal characterisation (table 3.2). This study aims to substantiate the use of the source term framework described in chapter 3. There is no direct observation of DNAPLs in the subsurface in the Tame Valley area, at least not published, so parameters defining the dimensions of DNAPL source zones are uncertain. All that can be discerned to date about the nature of chlorinated solvent source zones in the area is by indirect inference using abstracted aqueous concentrations. To attempt a deterministic source term model without uncertainty analysis would therefore be preposterous, and as such a progressively refined stochastic approach is adopted.

DNAPL source terms in this study are generated with the `DNAPL` package (Barry, 2017b), which is described and demonstrated in chapter 3. Each source term has five layers, with one in the upper MODFLOW model layer and four in the lower, evenly distributed through the water column. Recall that the groundwater model has two layers, the top layer of the parent groundwater model having been removed (section 4.2.2). Aquifer thickness and groundwater heads at each model location are read using the `RNetCDF` and `Rflow` packages (Michna and Woods, 2016; Barry, 2017e).

4.3.1 Spill locations

Rivett (1989) has identified solvent storage and disposal areas that may well have underlying DNAPL source zones (figure 4.1). However, different solvents will be associated with each location and some locations may not have leaked or been used much. Additionally, exact locations of DNAPL infiltration into the ground may differ slightly from the locations identified by Rivett (1989), especially where usage areas were identified within a large building. The initial stage of the source term refinement focuses on spill locations and which chlorinated solvents, if any, were spilt at each location. This is expanded upon in section 4.4.1.

4.3.2 Source term specifications

After a satisfactory estimate of the distribution of real source zones is achieved, a more detailed analysis of the parameters describing the source zone is conducted. This is done using Monte Carlo stochastic analysis, using the procedure described in section 4.4.2. The stochastic parameters describing each source zone realisation (table 4.2), as well as some parameters with fixed values, are input to the DNAPL source term framework (chapter 3).

Fixed input parameters to the DNAPL source term model mostly relate to chemical properties of the solvents. Other than those found in table 3.1, the residual saturations of water and non-aqueous phase liquid (NAPL) defined from maximum drainage followed by maximum rewetting (see figure 3.2) are set at 0.1. Porosity is also required by the DNAPL model. The total porosity is used for this purpose, a value of 0.24 as used by Daly (2005), approximately the median value determined for core samples of the Kidderminster Sandstone (table 5.2). Note that a different value, relating to the estimated fracture porosity, is used for effective porosity in the aqueous transport model (section 4.3.3).

Values for Γ describe different source depletion models, with higher values describing a source term that is stronger at early stages but depleted more quickly. Additionally, it is tested whether the value for Γ is allowed to converge to 1 with progressive source depletion, in line with the findings of Chen and Jawitz (2009), outlined in section 3.2. The equations for the source depletion models are, with a constant power law:

$$1 - \frac{C}{C_0} = \left(1 - \frac{m}{m_0}\right)^{\Gamma} \tag{4.1}$$

where Γ is either constant (Zhu and Sykes, 2004; Falta et al., 2005a) or converges to 1 as:

$$\Gamma = 1 + \frac{m}{m_0}\left(1 - \Gamma_0\right) \tag{4.2}$$

C and m are effluent concentration and source zone mass, in a given layer, respectively. m_0 is the historic maximum of mass in a contaminated layer and C_0 and Γ_0 represent concentration and source depletion parameter when mass was at this maximum. Flux is calculated from concentration as:

$$J = CA_{\mathrm{SZ}}q_{\mathrm{h}} \tag{4.3}$$

A_{SZ} is the source zone area perpendicular to flow, which depends on the dimensions of the pools and ganglia as well as the layer height, and q_{h} is the horizontal magnitude of the Darcy velocity, a transient value interpreted from the MODFLOW results data sets at the location of each source zone. Source zones nearer to an abstraction well will typically experience larger values of q_{h} and therefore produce a larger effluent flux J at early times and become depleted more quickly. The behaviour of DNAPL when it reaches the base of the aquifer is also tested, with the stochastic code choosing either to allow DNAPL to be lost from the bottom of the aquifer (invading fractures in the Carboniferous sequence) or not (spreading laterally on the contact with the Carboniferous sequence). In the latter case large persistent pools may form at the base with sufficient volume of spillage.

4.3.3 Aqueous transport model

Aqueous transport, which links the source term to the abstracted concentrations via advective-dispersive-reactive migration through the aquifer, is modelled using the Lagrangian DRW method (chapter 2; Barry, 2017c). This method is particularly well suited to this context as there are multiple contaminant point sources spread over a regional scale. The DRW does not exhibit numerical dispersion even on very coarse grids because mixing within model cells is not assumed at any stage of calculation. Another benefit is that contaminant release is not constrained to be located at the centre of grid cells, although with MODFLOW the point abstractions must be. The aqueous effluent flux output of the DNAPL source term model easily interfaces with the transport model, giving depth- and time-dependent input flux which is represented as particle releases in DRW.

For the Tame Valley study the following transport parameters were used initially for identifying the likely real source zones (section 4.4.1). Longitudinal and transverse dispersivity were 50 and 5 m respectively. A linear decay was specified with half-life 10 years and a retardation factor of 2 was specified assuming equilibrium sorption. Daly (2005) used a value of 0.24 for effective porosity. However, this value is more representative of matrix porosity, as it is based on core sample measurements. In Permo-Triassic sandstone fracture flow generally accounts for most of the groundwater flow, so the fracture porosity would be a more appropriate value. There are no direct measurements of the fracture porosity of the Birmingham sandstones, but values for specific yield, representing pore water that is easily drained, can give an indication of likely appropriate values. Specific yield measurements and calibrations in the West Midlands Permo-Triassic Sandstones generally give values in the range 0.1 to 0.16

(Allen et al., 1997; Cuthbert et al., 2001; Daly, 2005; Daily and Buss, 2013). 0.15, a value at the higher end of this range, is chosen for effective porosity, reflecting that flow is predominantly but not exclusively fracture-based. For the stochastic parameter estimation (section 4.4.2), longitudinal dispersivity, retardation factor and degradation were varied with values chosen from values given in table 4.2, with transverse dispersivity always a tenth of longitudinal dispersivity. A 200 day time step was used and for the DRW-specific solver parameters, 1 dispersion pair (N_{dp}) was used, horizontal coalescing radius r_{co} was set to $6\sqrt{\alpha_L}$ horizontally, where α_L is longitudinal dispersivity in metres because low-dispersivity models require denser particle swarms to describe the plume accurately, and vertical coalescing radius was set to 20 m.

4.4 Progressive Refinement of the DNAPL Source Term Model

4.4.1 Estimation of solvent spill locations

Section 4.3.1 indicates that Rivett (1989) gives an initial estimate of solvent spill locations in the Tame Valley region, but that it is likely that the real set of spill locations for each chlorinated solvent is a subset of these locations. In this initial stage of calibration, a more refined determination of the locations of spills for each solvent is sought. Free parameters (those which may be changed during the calibration process) for each solvent are whether each potential spill location has a spill of that solvent and the location of the spill, for which small reasoned adjustments are allowed (generally a few tens of metres). Set parameters for other source term and transport parameters are used to refine the real spill locations, because the existence and locations of spills will be the most sensitive influence on the modelled results. Parameters used at this stage of the calibration are given in table 4.1 and are generally in the middle of the range of stochastic variation used in the next stage of calibration (section 4.4.2; table 4.2).

Where modelled abstracted concentration results greatly exceed historical observed concentrations from the field it indicates that a simulated plume captured by the well is unlikely to be significant and therefore the originating source zone is probably not real. The calibration target was set that all modelled results should be within one order of magnitude of observed concentrations at the end of this stage of calibration. However, where observed concentrations are below 10 µg l^{-1}, which is low relative to most observed concentrations in the Tame Valley, modelled results of zero are accepted.

4.4.2 Calibration of source term and transport parameters

When the spill locations have been satisfactorily refined, a calibration of DNAPL source term and aqueous transport parameters is possible. Calibration is attempted for the parameters listed in table 4.2. Because of the large number of parameters and probable

Table 4.1: Initial guess DNAPL source term and aqueous transport parameters used in the first stage of calibration for refinement of the spill locations. Further discussion of these parameters is in table 4.2.

parameter	value
source depletion power Γ	1
peak usage rate related to a potential source	2 kg d^{-1}
total pool height per layer	0.5 m
maximum pool width	5 m
ganglia width	0.2 m
converging Γ?	no
impermeable aquifer base?	no
longitudinal and transverse dispersivity	50 and 5 m
retardation factor	2
degradation half-life	∞

similarities in the effects of certain groups of parameters, a stochastic approach was adopted. A stochastic approach to calibration enables sensitivity to many parameters to be tested simultaneously and is better able to identify where parameters are in trade-off, having very similar effects on the results. Each parameter is assigned, with some justification, a set of allowed values which are are sampled randomly for each model run. Table 4.2 gives the initial sets of values, with reasoning, that were assigned to each parameter.

The stochastic approach is executed in rounds, with each round consisting of 400 model runs. After each run, the results are analysed and visualised in order to identify the most sensitive parameters at each well and for each contaminant. The result aggregation and visualisation method used is demonstrated in figure 4.8. For parameters which are particularly sensitive, in general, optimum values are identified by comparing with observed concentrations. The optimum value for a parameter may well be a compromise between the matches to data at different wells. The tested parameters are not allowed to vary between source locations, apart from the peak usage (and hence spillage) rate that applies to each source location, or heterogeneously across the model. Some parameters, those which are probably significantly solvent-specific such as retardation factor, are allowed to vary between different solvents, but not others, such as Γ and dispersivity. These restrictions are used because, although parameters may actually vary heterogeneously, too many degrees of freedom would lead to a calibration which was largely curve-fitted and whose results would be less useful for application to other regions or more broadly in the Birmingham region. To substantiate the utility of the DNAPL source term framework, it must be demonstrated that it can produce reasonable results with regionally consistent parameters, because in most contexts for which it is intended there is insufficient data to support varying source term parameters such as the ganglia width.

The target for this calibration is to refine stochastic parameter ranges such that the interquartile range for modelled abstracted concentrations are spread over no more

than an order of magnitude and that observed concentrations are within the modelled interquartile range.

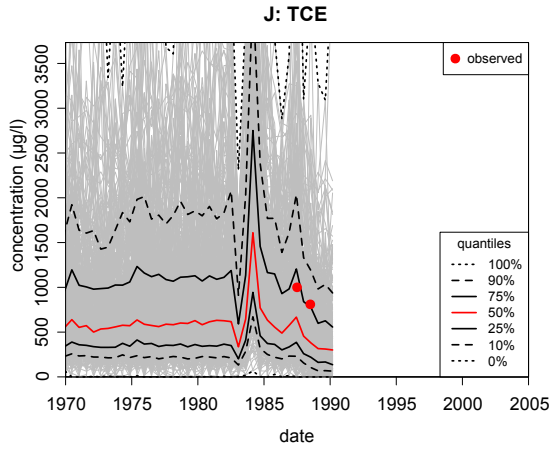

(a) raw results from 400 realisations, with quantiles

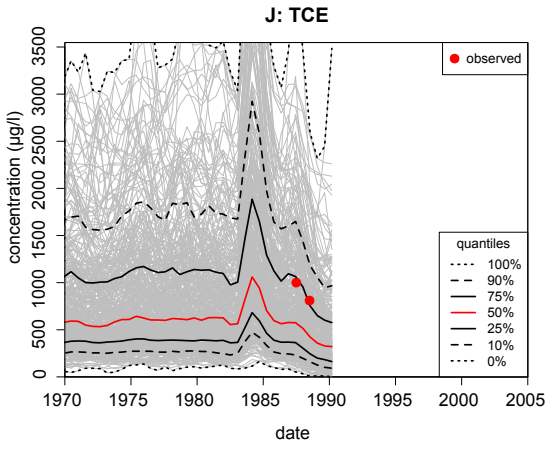

(b) results smoothed with Gaussian mask

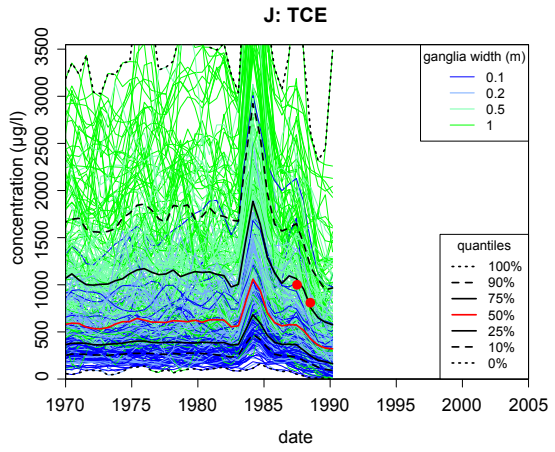

(c) results colour coded by a parameter value: ganglia width in this case

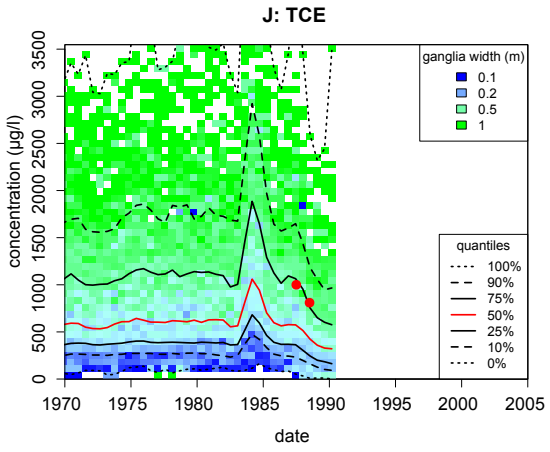

(d) colour flood of aggregated mean parameter value

Figure 4.8: Demonstration of analysis of stochastic results and visualisation of parameter sensitivity for one contaminant at one well (TCE at well J). The smoothed quantiles show the modelled range of likely concentrations and the colour flood allows the most sensitive parameters, as well as parameter values that best reproduce observed concentrations, to be identified. In this example, the observed concentrations plot on the 75^{th} percentile of the stochastic results (subfigure d) and suggest that a value for ganglia width between 0.5 and 1 m would best reproduce the data, at least for TCE at well J.

Table 4.2: Trialled parameter values used in stochastic model runs for TCE, PCE and 1,1,1–TCA source terms.

parameter	trialled values	justification
source depletion power Γ	0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20	range observed for contaminated land sites by DiFilippo and Brusseau (2008)
peak usage rate (kg d^{-1}) related to a potential source	0.5, 1, 2, 5, 10	For TCE for example, peak national usage is about 22,000 kg d^{-1}. Tait et al. (2004b) estimate there are around 70,000 potential chlorinated solvent user industries nationally. Divided equally this would give around 0.3 kg d^{-1} for each site, but the Tame Valley sites are large and not all sites would have used TCE.
total pool height per layer (m)	0.1, 0.2, 0.5, 1	Kueper et al. (2003) calculate pool heights for chlorinated solvents up to 1 m, depending on nature of capillary barrier, in sand or sandstone sequences.
maximum pool width (m)	1, 2, 5, 10	Field observations of lateral spreading by Kueper et al. (1993) and Brewster et al. (1995).
ganglia width (m)	0.1, 0.2, 0.5, 1	Hard to measure in field conditions. Model results of Parker and Park (2004) give a value of 0.3 m although this may be dependent on their model's cell size, which is similar.
converging Γ?	yes, no	Observed by Chen and Jawitz (2009) and Johnston et al. (2014), but this conclusion is not necessarily transferable.
impermeable base to aquifer?	yes, no	Parker et al. (2004) observe very little DNAPL penetration of a silt aquitard in Pleistocene glacial sediments. However, Carboniferous Coal Measures are older and likely more fractured.
longitudinal dispersivity α_L (m); transverse dispersivity always $\alpha_T = \alpha_L/10$	10, 20, 50, 100	Hydrodynamic dispersivity is often approximated as a tenth of characteristic transport length scale.
retardation factor R_f	1, 1.5, 2	section 4.1.2
degradation half-life (years)	10, 20, 200	section 4.1.2 (highest value is deliberately finite for easier analysis, but effectively represents negligible degradation)

4.5 Results and Discussion

4.5.1 Spill locations

The first stage of the source term investigation was to determine which spill locations are real and contain which solvents. When comparing the modelled abstracted concentrations to historic observed concentrations, there were many cases in which the modelled results obviously exceeded observations by orders of magnitude. Most obviously, PCE has hardly been observed in most of the wells in the Tame Valley apart from well F (figure 4.2). Therefore, when a PCE source term was included at all the other spill locations apart from spill 1, abstracted PCE was greatly overestimated by the modelling. TCE is the most extensive contaminant, with high observed concentrations at many of the wells (figure 4.3). 1,1,1–TCA is mostly present at large concentrations in the south-east of the region, in wells H, I and L (figure 4.4) and predicted spill locations that would contaminate these wells are chosen. The predicted spill locations for these three solvents are plotted in figure 4.9. These spill locations are used for results that follow.

Compared to figure 4.1, the modelled source zones have four spill sites moved slightly. Additionally, some of the identified potential spill locations are thought, on the basis of this stage of calibration, not have any significant spilt solvent. Source location 4 is moved slightly to the east, closer to well L, in order to achieve the magnitude of concentration observed at well L. The historic map indicates that solvent storage at location 4 would have been in a large building, so it feasible that the point of penetration of solvents into the subsurface may have been anywhere in the building. The proposed new location in figure 4.9 is still within that building. Source 7 is moved about 30 m to the north to give a less strong breakthrough to well K. And the modelled concentrations at wells C and D are very sensitive to the exact position of source 11, which sits close to the capture zone boundary between the two wells. It was found that moving the source west by 50 m to the corner of the nearby building produced the best results for C and D. Because concentrations at wells H and I were overestimated, source 9 was moved about 80 m to the north-west, to the neighbouring building (still on the same industrial site). These distances are small enough to be explained by surface drainage or subsurface lateral migration and do not necessarily imply that Rivett (1989) imprecisely identified potential spill locations.

With these revisions, it was possible to achieve a match to data within one order of magnitude for most observed concentrations above 10 µg l^{-1}. There were a few reasoned exceptions. Well F is thought to be anomalous with regard to TCE. Here, the TCE could be degradation product of PCE, as it follows a similar trend and the industrial activity at that location was dry cleaning, associated with PCE. To avoid attempting to match contamination which is in fact a degradation product, spill 1 is not included as a TCE source. Well L has overestimated 1,1,1–TCA concentration, but moving the 1,1,1–TCA source from source 4 to source 3 underestimates the concentration at well L and overestimates at well J. Because it would be tenuous to invoke

an additional source location between 3 and 4, the 1,1,1–TCA source at location 4 is retained. TCE observations at well I were also difficult to match, with modelled concentrations exceeding observed by a factor of about 15, and the reason for this is less clear.

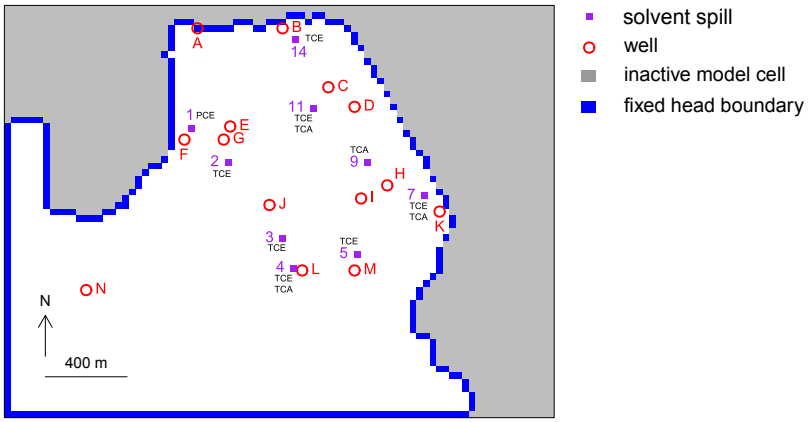

Figure 4.9: Suspected solvent spill locations (usage, storage or disposal locations, Rivett, 1989), labelled with solvents that are likely to be present as deduced from the first stage of calibration, identifying likely spill locations. Wells are labelled with capital letters and potential spills are labelled with numbers — these labels are used in the text to refer to specific features.

4.5.2 Contaminant plume migration

Figure 4.10 and the supplementary electronic material (appendix D.2) show the plume evolution for each solvent, using the source locations determined in section 4.5.1. Two of the time slices in figure 4.10 roughly coincide with organic groundwater quality surveys (Rivett, 1989; Murcott, 2008) and a predictive time slice at 2025 assuming that all abstraction ceases after 2010. The parameters used for these plumes are those used before the stochastic runs (see section 4.3.3), with the source depletion parameter Γ set at 1. Clearly, different plumes would arise from different parameters, especially transport parameters. However, general patterns in the predicted concentration distribution can be observed. The observed match between observations and the calculated plumes are encouraging. Recall that no attempt is made to model the TCE in well F (see figure 4.9 for labelled well locations) as it may be a degradation product from PCE. Wells H and I also show some mismatch for TCE, discussed further in section 4.5.3.

The concentration distribution is heterogeneous, meaning that, while general organic water quality in this area is clearly bad, not all samples, especially passive samples, would return a high concentration. Sampling of unpumped wells can produce useful data, notwithstanding consideration of the origin of sampled groundwater in passive techniques (for example: McMillan et al., 2014). However, it must be recognised that the concentration from unpumped boreholes may be very different to nearby locations.

Although monitoring boreholes are often considered superior for mapping aquifer concentrations (Graber et al., 2008), abstracted concentrations may in fact be more useful unless a very dense monitoring network can be established (unlikely beyond the boundaries of a contaminated site), because abstracted concentrations represent something akin to the average concentration of a large volume of surrounding groundwater.

Rivett et al. (2012) suggested a decline in abstraction from the Birmingham aquifer over the last half century may have released contaminant plumes that were previously largely captured by these abstractions to the wider aquifer, potentially causing a delayed contaminant problem that will hinder resource development. This phenomenon is predicted by the calculated plumes in this study, but the plumes migrate to the nearby River Tame after being released. The River Tame water quality could be derogated by this, but Ellis and Rivett (2007) found that in fact the concern is likely to be only local, at points of the river where plumes emerge, and that downstream concentrations will be very dilute. Therefore, the contamination of groundwater in the Tame Valley region, although high in concentration, is not likely to give concern to other areas. However, it is possible that some contaminant migration passing under the River Tame to areas to the north may occur contrary to the groundwater model assumption that there is no flow across the River Tame (section 4.2.1).

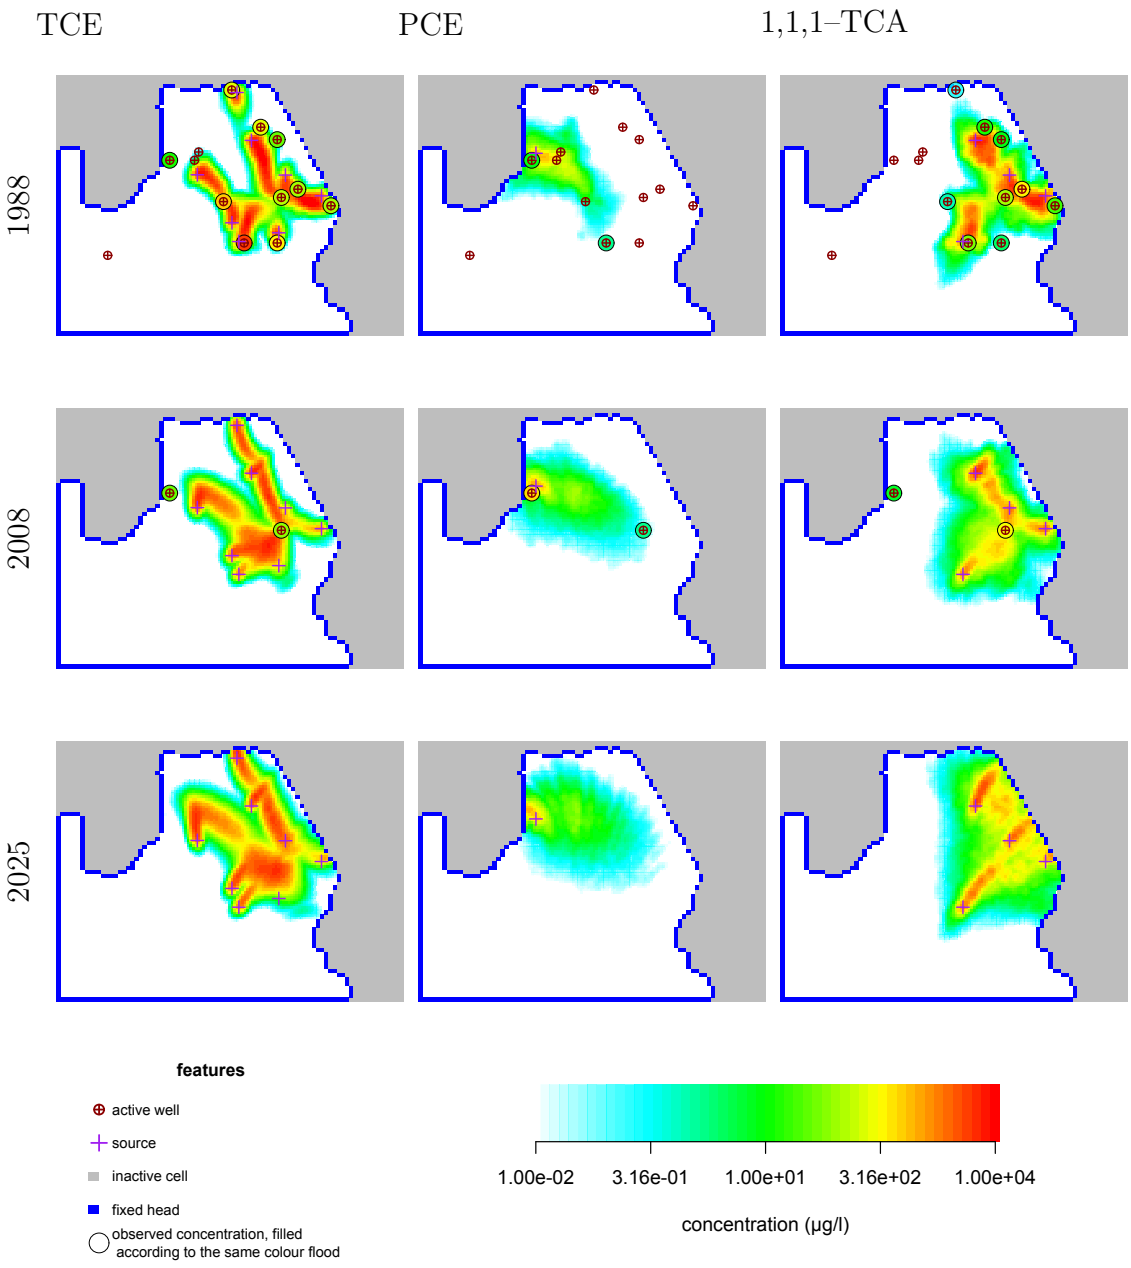

Figure 4.10: Time slices of modelled contaminant plumes arising from sources of different chlorinated solvents. Observed concentrations are plotted as coloured circles for the earlier times; the values plotted are the means of data collected from a well within a year of the middle of the displayed year. Note that the colour flood is spread logarithmically, which makes the plume appear more spread than a linear colour flood. Animated plots showing the complete transient evolution of plumes are included in the supplementary electronic material (appendix D.2). See table 3.1 for chlorinated solvent solubility limits, for comparison.

4.5.3 Source term and transport parameters

A series of stochastic model runs were used to try to find the optimum source term parameters for the region. For each contaminant, the set of 400 model runs with stochastically selected parameters were analysed as a collection and the most sensitive parameters sought. The modelled results were smoothed and aggregated to show transient quantiles and sensitivity to each tested parameter (section 4.4.2; figure 4.8).

The first round of stochastic runs clearly showed that the peak usage rate applied to the source term model was the most sensitive parameter (figure 4.11). This is unsurprising, as this parameter governs the amount of contaminant mass entering the subsurface. Values for peak usage were revised for each active source location in light of the results from round 1 of calibration (table 4.3). Peak usage is allowed to vary between source locations because it is unreasonable to assume that all spills were of equal magnitude. Note that, because the conservative assumption is taken that no spilt solvent volatilises (section 3.3.1), the actual usage rate of solvent on site which contributed to each spill site may in fact be greater than the values concluded here.

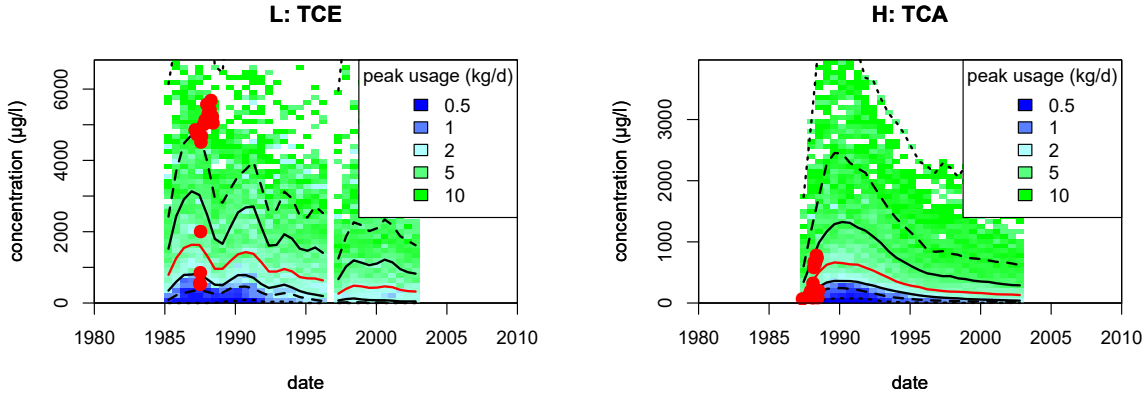

(a) TCE at well L

(b) 1,1,1–TCA at well H

Figure 4.11: Example stochastic results from round 1 of stochastic analysis, showing sensitivity to peak usage applied to sources. Red circles are observed concentrations.

Round 2 of the stochastic calibration was run after the peak usage was modified (table 4.3). Analysis of the second round showed that the parameters which were most sensitive varied between wells and contaminants. The ganglia width was generally the most sensitive parameter for PCE and TCE (for example: figure 4.12a), but not for 1,1,1–TCA. The results at wells which were further from source locations tended to be more sensitive to transport parameters (figure 4.12d) whereas source term parameters were, broadly speaking, equally important for all wells. Although ganglia width was generally sensitive, it was not possible to find a value that would produce better matches to data for all wells and contaminants. Unlike peak usage rate, this parameter is likely to be much more uniform between source zones and as such the protocol for

Table 4.3: Modified peak usage rates (in kg d^{-1}) applied to each source location after analysing the results from the first stochastic round. Source location numbers are as in figure 4.9. The initial range of trialled values at each active source location is in table 4.2.

source location	PCE	TCE	1,1,1–TCA
1	2	–	–
2	–	4	–
3	–	4	–
4	–	10	0.7
5	–	3	–
6	–	–	–
7	–	3	2
8	–	–	–
9	–	1	2
10	–	–	–
11	–	2	2
12	–	–	–
13	–	–	–
14	–	5	–

parameter refinement precludes source-specific values for source zone dimension parameters (section 4.4.2). Therefore the range of values trialled for ganglia width was not modified after round 2. The pool width and height parameters were generally less sensitive than ganglia width (for example: figure 4.12b), but analysis of results showed that most observed concentrations favoured higher values for these parameters, and as such the ranges for these parameters were slightly refined (table 4.4). For the transport parameters, most observed concentrations favoured dispersivity values in the middle of the the initial range (figure 4.12c, for example) and it was found that predicted retardation factors from section 4.1.2 were consistent with observed concentrations. The dispersivity and retardation parameter ranges were modified accordingly (table 4.4).

Table 4.4: Refined parameter ranges after analysing round 2 of stochastic calibration. Initial parameter ranges are in table 4.2 and the peak usage rate at each source location was modified according to table 4.3 after round 1.

parameter	refined range
pool height (m)	0.2, 0.5, 1 (0.1 removed)
max. pool width (m)	2, 5, 10 (1 removed)
longitudinal dispersivity (m)	20, 30, 40, 50
retardation factor ()	PCE: 1.8; TCE: 1.3; 1,1,1–TCA: 1.4

Analysis of round 3 of the stochastic calibration indicated that most datasets favoured a higher value of ganglia width. Some revisions were also found to be necessary for source locations (source 14) and peak usage rates (TCE at source 7). These are summarised in table 4.5.

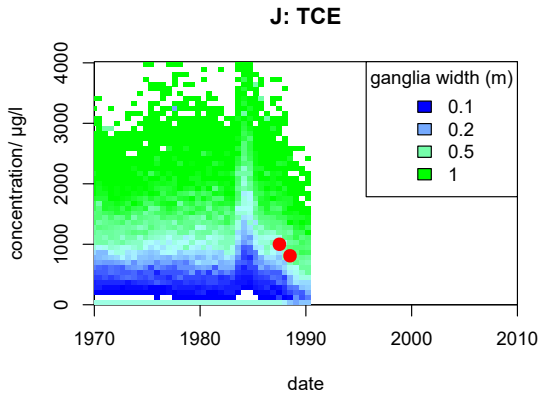

(a) TCE at well J; sensitivity to ganglia width

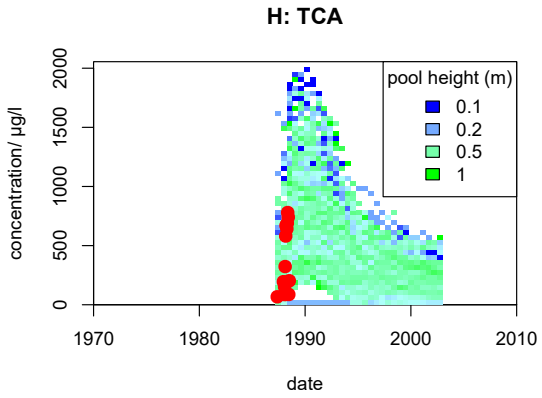

(b) 1,1,1–TCA at well H; sensitivity to pool height

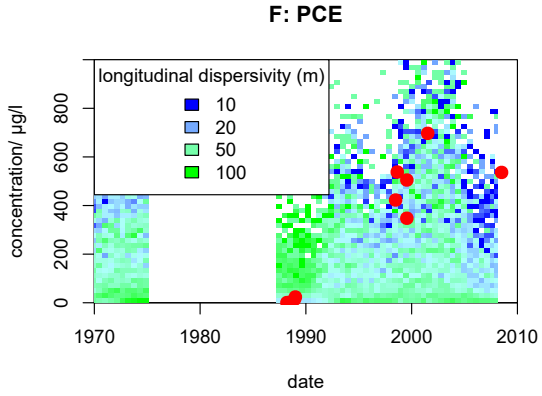

(c) PCE at well F; sensitivity to dispersivity

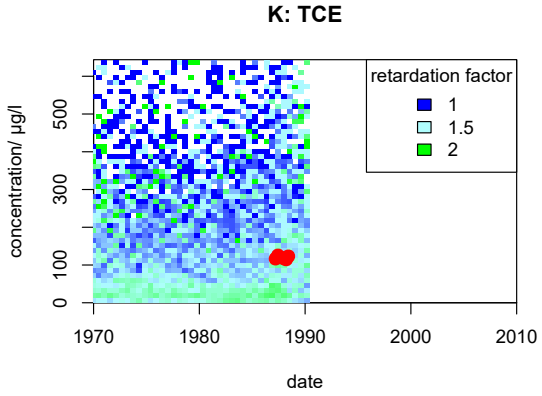

(d) TCE at well K; sensitivity to retardation factor

Figure 4.12: Example parameter sensitivities from round 2 of the stochastic calibration.

After the fourth round of stochastic calibration, it was judged that further improvements with the prescribed restrictions on degrees of freedom (section 4.4.2) were unlikely. Of the concentration time series for which matches were attempted (excluding TCE at well F which is likely to be from degradation of PCE), five plotted within the interquartile range (examples: figure 4.13) and a further seven were considered close (examples: figure 4.14). Four attempted matches continued to show problems (figure 4.15). All interquartile ranges spanned less than an order of magnitude, or the 75^{th} percentile was less than 50 µg l^{-1} if the 25^{th} percentile was less than 5 µg l^{-1}, so satisfactory parameter refinement was achieved. For TCE and 1,1,1–TCA at wells L and H, the data sets show a rapid breakthrough responding to the start of pumping (figures 4.3 and 4.4). The transient pumping variation in each case is on time-scales shorter than the stress periods in the groundwater model, so in these cases calibration was only attempted for the higher data at the end of the transient responses (figures 4.13 b and d).

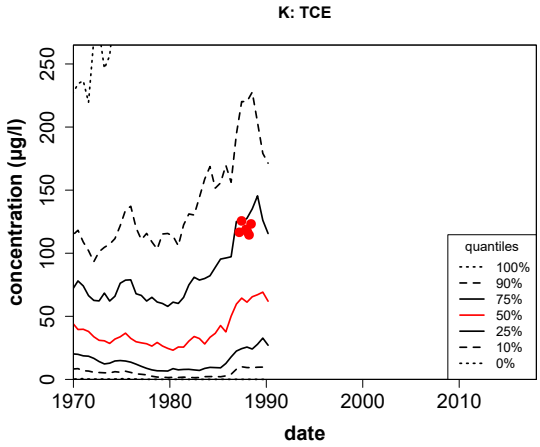

(a) TCE at well K

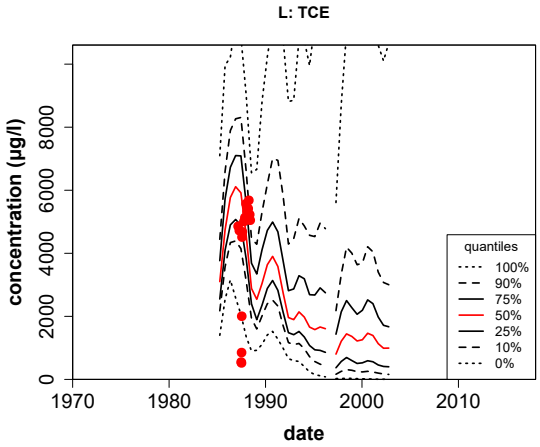

(b) TCE at well L

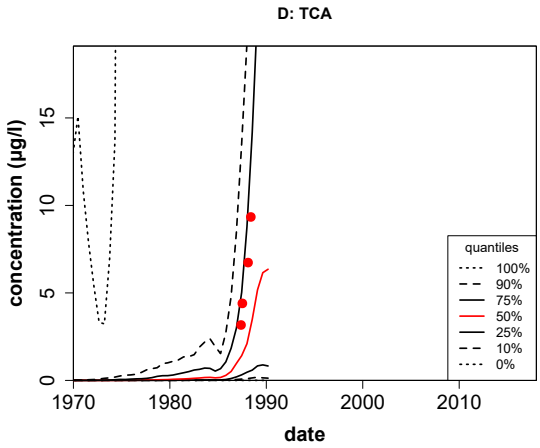

(c) 1,1,1–TCA at well D

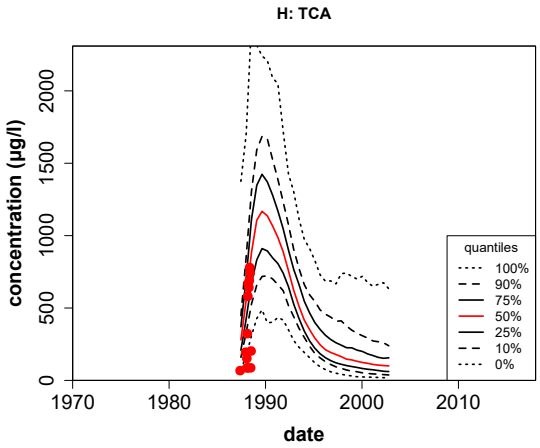

(d) 1,1,1–TCA at well H

Figure 4.13: Transient quantiles of model results which achieved a good match to data after the fourth round of stochastic calibration.

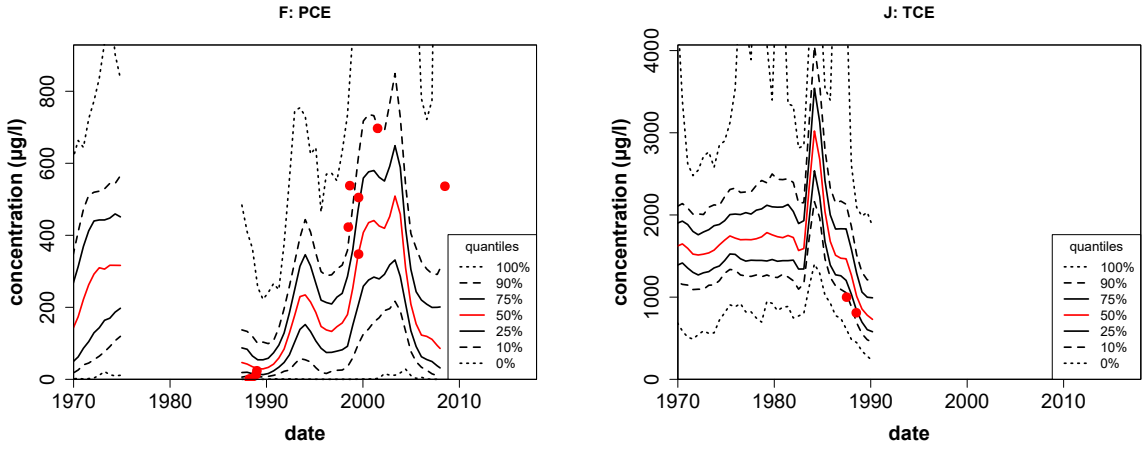

(a) PCE at well F

(b) TCE at well J

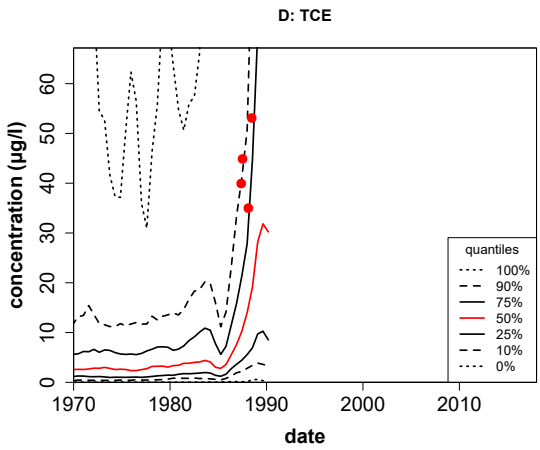

(c) TCE at well D

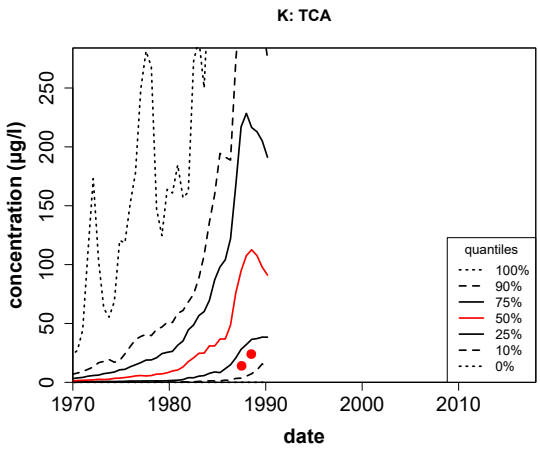

(d) 1,1,1–TCA at well K

Figure 4.14: Examples of transient quantiles of model results which did not satisfy the calibration criteria of having data within the interquartile range, but nonetheless did not deviate too greatly from observations.

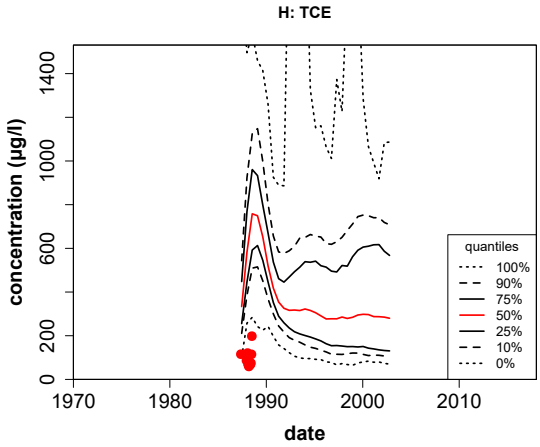

(a) TCE at well H

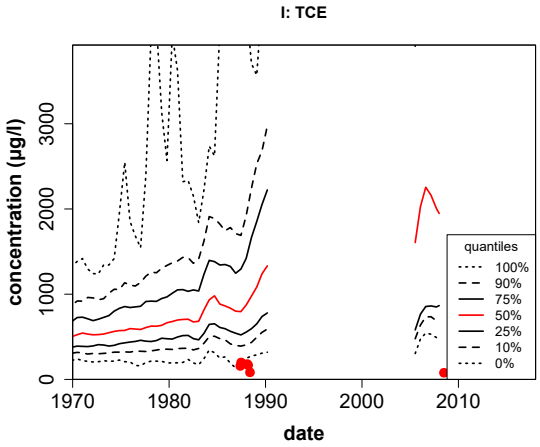

(b) TCE at well I

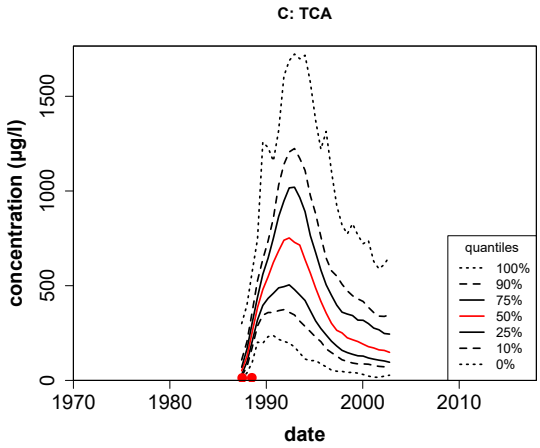

(c) 1,1,1–TCA at well C

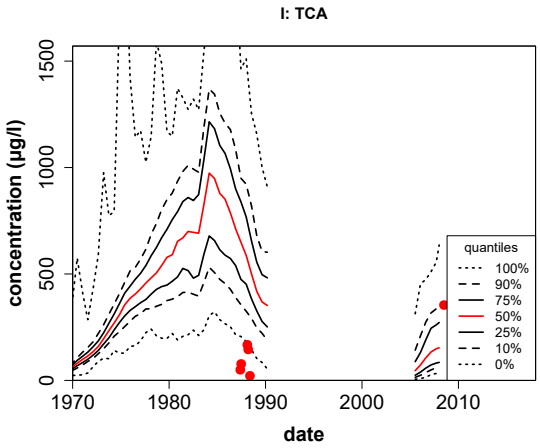

(d) 1,1,1–TCA at well I

Figure 4.15: Transient quantiles of model results for which there remained problems with the calibration.

Table 4.5: Modifications and parameter refinements after analysing round 3 of the stochastic calibration.

parameter	modification
ganglia width	refined to a value of 0.7 m
source locations	source 14 moved 40 m to the south-east
peak usage	changed the peak usage rate for source 7 to 1 kg d^{-1}

It is encouraging the coupled source term and dissolved phase transport models can be calibrated to reproduce most of the observed historical data within an interquartile range. However, the cases for which the matches to data are less good merit discussion. Wells F and I show some of the few data collected after 2000. The late stage data for PCE at well F (figure 4.14a) and for 1,1,1–TCA at well I (figure 4.15d) are unexpectedly high, plotting close to the 90^{th} percentile even though modelled results either agree well with or overestimate earlier data (the same is not true for TCE at well I: figure 4.15b). Although the model results for these cases do show a late-stage increase, the magnitude is less than that observed. It is possible that water table rise (figure 4.5) results in remobilisation of solvent contamination in the vadose zone, a process that is not accounted for in the source term model. Some of the other mismatches may be attributed to mistimed breakthroughs. For example, modelled results of 1,1,1–TCA at well C (figure 4.15c) correctly show that the initial concentrations are relatively low, less than 50 µg l^{-1}, but then the modelled rise in concentration is too quick to closely match the observed concentrations.

The restrictions on degrees of freedom stipulated that most parameters should not vary heterogeneously or between source zones (section 4.4.2). This gives confidence that the achieved calibration is not simply curve-fitted but indicates representative source term descriptions for generic DNAPL spills within this region. The applicability of these calibrated parameter ranges to regions outside the Tame Valley will depend on how different the aquifer properties are.

The parameter range was not refined to single values for each parameter. Further refinement of parameters would be unlikely to result in a net improvement in calibration. It is perhaps predictable that it is not possible to find a single parameter set that works for all source zones and it is certainly true that in all regional settings like this, or larger scale, the DNAPL source term must be trialled with a range of different parameter values and dissolution models, not just to find the ranges of parameter values which best reproduce the data, but also to test the range of possible results for predictive simulations which are of not directly constrained by any observations. For the abstracted concentrations, up to the latest data from 2008, little sensitivity was found to the source depletion parameter Γ, nor the assumption of a perfectly impermeable aquifer base. This does not, however, mean that these inputs are unimportant, as they may become sensitive parameters in predictive scenarios beyond the time frame of data. This issue is discussed further in section 4.5.4.

4.5.4 The chlorinated solvent source term

Figure 4.16 shows example results for a TCE, PCE and 1,1,1–TCA source term. The results used are from the fourth stochastic round (section 4.5.3) and show the complete ensemble of 400 model realisations in each case. The source terms were generated up to the year 2050, with the modelled flow rate beyond 2010 assuming a cessation of abstraction in the region and a single stress period was used for the predictive period of the model beyond this time. The results show that the model is able to capture the persistence of source terms well beyond the period of solvent spillage, especially for TCE and PCE. PCE forms the most persistent source term because of its low solubility (table 3.1). TCE is thought to have been used in greater quantities (section 4.5.3) and over a longer period (figure 3.4) than 1,1,1–TCA, so that there is simply a much larger mass of TCE released to the subsurface, which accounts for the greater persistence of modelled TCE source terms than 1,1,1–TCA.

The plots in figure 4.16 are colour coded by Γ and the impermeable aquifer base assumption, two source term parameters that are particularly sensitive in the later stages of the source term. Pool height and width, which are not plotted, also show a correlation with source term flux, though more so in the earlier stages. Whether or not the value of Γ converged to 1 (section 4.3.2) was not found to be very important. The sensitivity of the former two parameters is particularly interesting because the values of these parameters were fairly inconsequential when attempting to match observed concentration data, the latest of which is from 2008 and most of which is from the late 1980s. This implies that the behaviour of source terms in the early stages, during and shortly following the period of usage and spillage, is very different from the behaviour of source terms in the late stages, decades after the cessation of solvent spillage. This presents the problem that the important parameters which concern the future risk to supply wells and other receptors cannot be easily assessed using historic data. This emphasises, therefore, that it is very important to test a wide range of parameter values and perhaps especially for the source depletion parameter Γ and the impermeable base assumption.

The relationship of the transient source term flux to Γ in the lowest layers, especially for TCE and PCE, seems surprising on first consideration (figures 4.16a and 4.16c). Higher values of Γ describe source terms that are stronger initially, but less persistent. This effect can be seen, albeit subtly, in the lower layers in the 1,1,1–TCA source term (figure 4.16e). The explanation is that for TCE and PCE source terms, a large amount of mass accumulates in the pools that is very persistent. At some time beyond the range of the modelled source terms, the high Γ source zones with impermeable aquifer base will be depleted before those generated with low Γ.

Whether or not the underlying aquifer is an effective barrier to vertical DNAPL migration is a significant parameter for the future source term strength and persistence, though of less importance than Γ. If it is, then extensive DNAPL pools may exist on top of the boundary which could take more than a century to fully deplete. If however, DNAPL has naturally sequestered itself in the underlying Carboniferous Coal Measures, then the chlorinated solvent source to the sandstone is predicted be less strong

and deplete quicker. However, it should be noted that the effect of back-diffusion, which can produce a persistent source from contaminated aquitard materials (Yang et al., 2016), has not been modelled in this study.

Another factor that affects the long-term source term strength is the presence of nearby abstractions. Figure 4.17 compares two TCE source terms, of which one was close to a strong abstraction at well M. This source term (spill 5) shows a stronger flux in the early stages, but is all but depleted by the year 2000. This is remarkable considering the persistence of most other TCE source terms and indicates how effective a simple pump and treat-based remedial solution can be. Well M was, however, operating while the source zone was forming and it would not necessarily be so effective to position a strong abstraction close to an old source zone now. This consideration is represented empirically in this method by the Γ parameter which, especially but not exclusively for low values of Γ, would yield a lower effluent concentration from old source zones — with an equivalent groundwater flow rate, the effluent flux is stronger from an early-stage source zone than a late-stage source zone. Therefore, a pump an treat operation near a late-stage source zone would not lead to such an effective reduction in source mass.

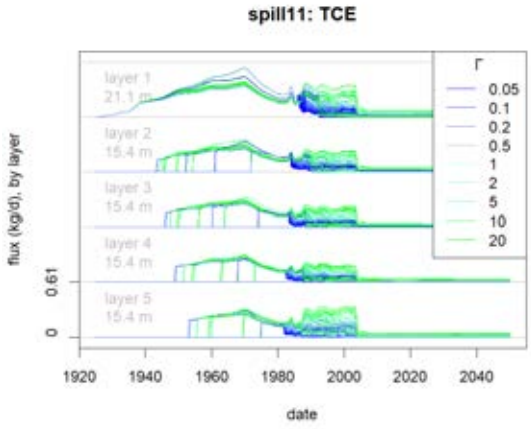

(a) TCE, spill 11

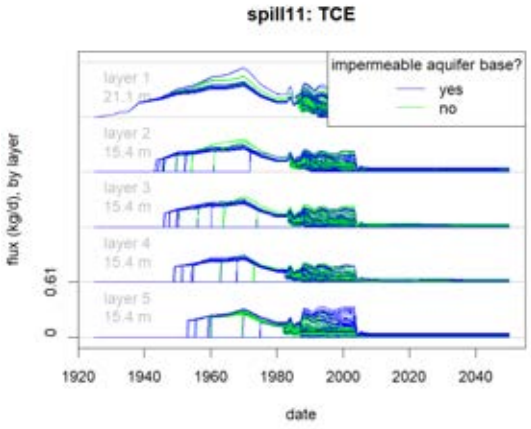

(b) TCE, spill 11

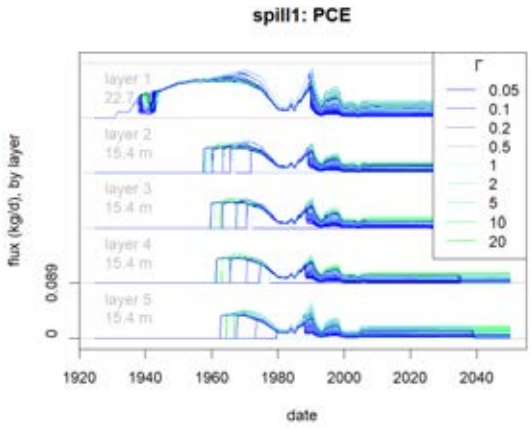

(c) PCE, spill 1

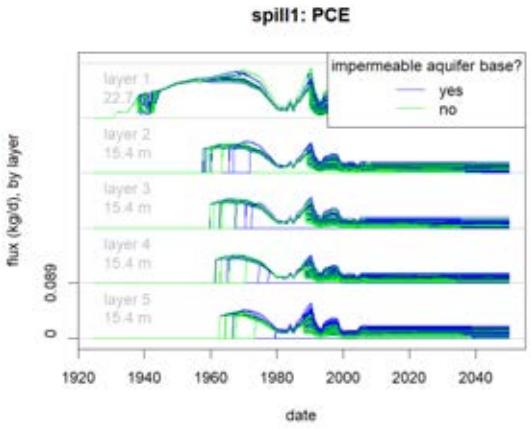

(d) PCE, spill 1

(continued on next page)

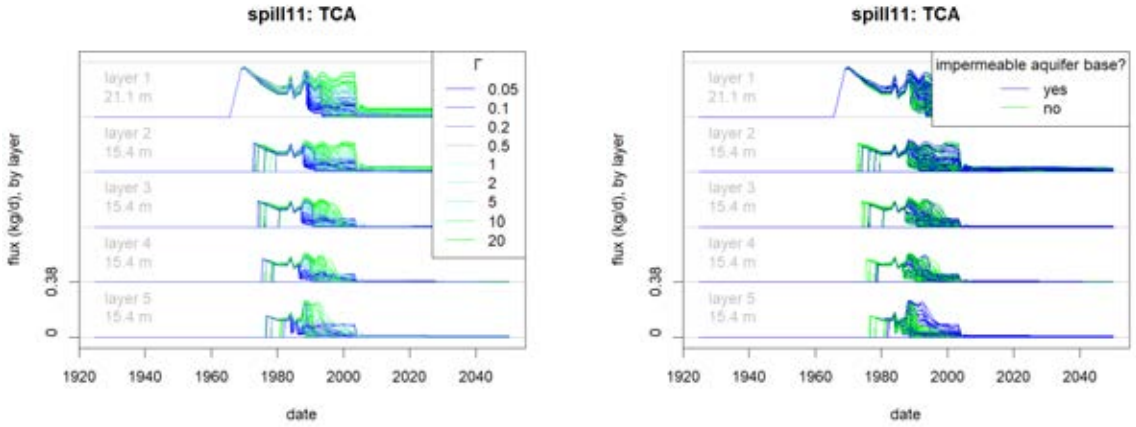

(e) 1,1,1–TCA, spill 11

(f) 1,1,1–TCA, spill 11

Figure 4.16: Example source terms for TCE, PCE and 1,1,1–TCA, showing the complete set of stochastic results from the fourth stochastic round, colour coded by the source depletion parameter Γ (right) and whether the assumption of a perfectly impermeable aquifer base was in place (left). The results are plotted by layer of the DNAPL source term model, of which layer 1 is in layer 1 of the groundwater model (upper sandstone), and layers 2 to 5 are evenly distributed in layer 2 of the groundwater model (lower sandstone). Layer thicknesses are shown in grey. The same flux axis applies to all model layers.

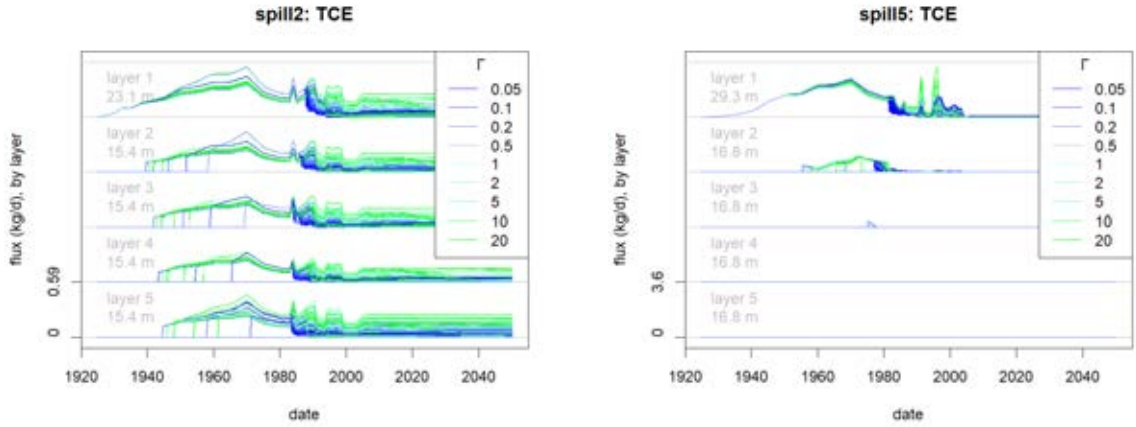

(a) TCE, spill 2

(b) TCE, spill 5

Figure 4.17: Comparison of two TCE DNAPL source terms, from spill locations 2 and 5 (see figure 4.9). Spill 5 is very close to well M whereas spill 2 is not close to any major abstraction (note that abstraction at well G was considerably weaker and shorter in duration than at M). Note different flux axis scales. Further explanations of these plots is in the caption to figure 4.16.

4.6 Conclusions

4.6.1 Chlorinated solvents in the Tame Valley region and Permo-Triassic Sandstones

This computational study of the organic contamination of the Tame Valley corroborates the very high but heterogeneous dissolved solvent concentrations that have been observed in the region during various surveys, for TCE, PCE and 1,1,1–TCA. It is likely that this contamination will persist in the groundwater for several decades to come. Key factors in determining this are recalcitrant DNAPL source zones that form a persistent source to groundwater (although less so for 1,1,1–TCA), high aquifer porosity and low degradation rates, although the effects of sorption are minor. The results have enabled better estimation of the locations of signficant DNAPL spills through model iteration and improvement. PCE is mostly found in the west of the region, probably originating from a single spill location and TCE and 1,1,1–TCA sources are distributed around the east of the region. 1,1,1–TCA was probably used in lower quantities and for a shorter period than TCE (figure 3.4; table 4.3) and will therefore be a less persistent contaminant. PCE, on account of its lower solubility (table 3.1) and slightly greater sorption potential, will be the most persistent although spatially more limited (figure 4.10).

However, it is not expected that groundwater contamination in the Tame Valley will cause problems to other water bodies. The future migration of the plumes will take them to the nearby River Tame and the magnitude of flow in the River Tame combined with likely volatilisation in turbulent surface waters will serve to greatly reduce concentrations downstream (Ellis and Rivett, 2007).

The DNAPL source terms in Permo-Triassic sandstone aquifers more generally could be very persistent, perhaps for more than a century depending of quantities of solvent used and spilt, and distributed deep in the aquifer. The most appropriate value for the Γ source depletion parameter is important for estimating the long-term persistence of sources, but this factor is difficult to estimate from inspecting historical data. The operation of historical abstractions in the region is also an important feature to constrain. Not only do abstractions have a strong influence on the transient groundwater flow field and therefore contaminant migration, they can also play a key role in the flushing of DNAPL source zones (figure 4.17).

Perhaps the most important message for using DNAPL source term method is that a range of parameters must be tested. It is almost certainly infeasible to obtain a deterministic set of parameters that can be applied confidently to just one region, much less that are transferable to other contexts. Therefore a range of parameters to test must be established, guided by field and laboratory studies, local conditions and human judgment, which pertain to the source term and aqueous transport.

4.6.2 The DNAPL source term framework and DRW solution coupled as a risk methodology

The `DNAPL` and `DRW` packages, facilitated by the `Rflow` package in `R` (R Core Team, 2014; Barry, 2017b,c,e) provide a holistic regional chlorinated solvent risk methodology, tracking mass from initial use through DNAPL source terms and dissolved phase transport to abstraction by receptors. The packages are comprehensively documented, robustly tested and seamlessly linked, with functions for output analysis provided. Open-source geographic information system tools available within `R` additionally enable the linking of land-use models for the generation of source terms, which is demonstrated in chapter 5. Therefore, the methodology is simple to implement for regions where there exists a transient groundwater flow model. The methods, in particular the DRW code, are specifically designed for the regional context and to be used in conjunction with a regional-scale contaminant transport model. The `DNAPL` package is designed so that additional processes and models may be programmed as specific contexts or advancing research into DNAPL behaviour demands, providing it with flexibility to be applied broadly and adaptability such that it will not become obsolete.

This study demonstrates the power of these combined tools to model the historic and future solvent contamination of groundwater across a large region consistently and comprehensively. Users will need to be thoroughly aware of the assumptions and limitations of the various components of model, many of which are merely default options that can be changed if necessary. In particular, the empirical and uncertain nature of the source term must be recognised and a large number of parameter combinations should be trialled, ideally stochastically. Nonetheless, the methodology is able to produce realistic results that generally corroborate well with observations and, where they do not, indicate areas for which the understanding needs further refinement, for example with abstraction histories and source locations. With a good understanding of the model's components and processes, the DNAPL source term framework and DRW transport code can provide a platform with which conceptual models can be consistently numerically represented and which can identify key areas of future concern and uncertainty.

Chapter 5

FORECASTING THE TRANSIENT ABSTRACTION OF CHLORINATED SOLVENTS IN A MULTI-LAYERED REGIONAL AQUIFER: STOURBRIDGE

The town of Stourbridge has a long industrial history. Under Stourbridge is a multi-layered Permo-Triassic Sandstone aquifer system that is an important water resource exploited by Severn Trent Water (STW) and South Staffordshire Water (SSW). The presence of dissolved chlorinated solvent contamination in abstracted groundwater has become a problem for public water supply (PWS) wells in the area. Because of the poor constraint on the location and transient evolution of chlorinated solvent spills, the future groundwater quality from abstraction wells is uncertain, which is a challenge for asset management planning. Section 5.1 reviews the hydrogeology, industrial history and water quality of the Stourbridge Permo-Triassic Sandstone. Treatment works such as air stripping towers have a threshold raw concentration above which the concentration reduction will not be sufficient to achieve safe drinking water, so models of future abstracted water quality are vital for determining whether these measures will be effective in the long term.

The application of the dynamic random walk (DRW) method (chapter 2) and the dense non-aqueous phase liquid (DNAPL) source term framework (chapter 3) allow a quantitative and internally consistent analysis of chlorinated solvents in the region, tracking solvent mass from industrial import and usage to abstraction in the dissolved phase by receptors. The study combines consideration of industrial history, analysis of the local geological layering and lithology, representation of DNAPL migration and dissolution and numerical modelling of advective-dispersive aqueous transport. Inevitably, there is uncertainty in trying to represent some features quantitatively. However, by framing the whole problem in a 3-dimensional and transient framework, the methods give the means for visual and quantitative analysis for asset planning.

The method requires a regional groundwater flow model and Stourbridge is covered by the Environment Agency (EA)'s West Midlands Worfe regional model. This is reviewed in section 5.2. The calibration to observed groundwater hydrographs in the Stourbridge area is displayed and boundary condition modifications are tested to assess the sensitivity of the groundwater flow field to local boundary condition assumptions. Using the historic industrial site locations identified in section 5.1.2 as source locations, the transient transport of potential contaminant plumes through the aquifer is assessed in section 5.3. Additionally, seven scenarios for the future operation of a key STW supply well, Winnie Road[1], are trialled, to predict the variation in concentration with abstraction rate. Section 5.4 develops consistent DNAPL source term models which reflect the characteristics of the geological sequence. The sensitivity of the source terms to solvent usage rates, source zone dimensions and source zone depletion model (using the empirical parameter Γ: section 3.3) are tested. The results of coupling the dissolved transport model to the DNAPL source term models and their implications for the future management of Winnie Road are then discussed in section 5.5. Finally, conclusions about the future groundwater quality of the Stourbridge Permo-Triassic Sandstone and about the utility of the DRW and DNAPL source term methods for regional groundwater quality forecasting are presented in section 5.6.

[1] Names of current and historic abstraction wells in this chapter have been changed to protect their security. A random error is also applied to their positions in map figures.

5.1 Hydrogeology and Organic Groundwater Quality in the Stourbridge area

5.1.1 Geology and hydrogeology

Stourbridge straddles the faulted boundary between Carboniferous Coal Measures (to the east) and Permo-Triassic Sandstone (to the west), between the Stafford and Worcester Basins (Allen et al., 1997). In the Stourbridge area, the sandstone sequence extends in age from the Bridgnorth Formation (lower Permian) to the Wildmoor Formation (Anisian, lower Triassic) and the younger Bromsgrove Formation outcrops close by to the south.

The Permo-Triassic sandstone in the West Midlands does not feature any major aquitard units which isolate the formations, but apart from the Bridgnorth Formation which is massive and relatively isotropic, each formation exhibits stratigraphical layering on varying scales which diminishes the vertical hydraulic conductivity (Allen et al., 1997; Walker and Jameson, 2010). There is a general fining-upward trend throughout the sequences, with layered variations superimposed on this trend (Tyler-Whittle et al., 2002). The Bridgnorth Formation is the coarsest, formed of sandstones, breccias and conglomerates (Allen et al., 1997; Tait et al., 2004b) and as such is the most permeable (Cuthbert et al., 2001). The Bridgnorth Formation is also poorly cemented and friable and the Winnie Road abstraction, screened in this formation, has suffered from sediment ingress and even the formation of voids at higher pumping rates (Allen et al., 1997; Anderson and Gamble, 2009; Robeznieks, 2009). Deep borehole logs show some of the layer characteristics (figure 5.1), particularly the relative heterogeneity of the Kidderminster Formation, which probably causes strong hydraulic anisotropy in that layer.

To (2016) performed permeability measurements on outcrops of Kidderminster and Bridgnorth Formations. Occasional very low-permeability layers were found in the Kidderminster Formation, which was formed in a fluvial depositional environment; these low-permeability layers are likely to have been formed under standing water, in flood plains or lagoons. By contrast, the permeability range of the Bridgnorth Formation, which was formed in a high-energy aeolian environment, is much narrower and no low-permeability layers were found.

The formational hydraulic conductivities have a smaller logarithmic range than the core samples, which is explained by the averaging of properties at the larger scales measured by pump tests (table 5.1). This feature is seen most strongly in the Kidderminster Formation. This corroborates with the general observation that the Kidderminster Formation is lithologically the most heterogeneous formation. Similarly, the porosity measurements (table 5.2), which have a similar mean for all formations, show a much more heterogeneous range for the Kidderminster Formation than the Bridgnorth Formation. Allen et al. (1997) comment that in fact most of the porosity measurements for the Bridgnorth Formation are close to 30%, consistent with the observation

Table 5.1: Hydraulic conductivity measurements (m d^{-1}) at the formation scale (pump test) and core sample scale in the West Midlands Permo-Triassic Sandstone, from Allen et al. (1997). Means are geometric.

Formation	Formation range	Formation mean	Formation count	core range	core mean
Bromsgrove	0.014 – 486	1.58	19	2.6×10^{-4} – 16.4	0.50
Wildmoor	0.77 – 62.6	8.06	8	3.1×10^{-4} – 12.0	0.37
Kidderminster	0.14 – 69.2	4.93	13	4.6×10^{-6} – 17.8	0.49
Bridgnorth	NA	NA	0	2.5×10^{-4} – 9.4	0.95

that there is very little cementation in this formation, in contrast to the Kidderminster Formation.

The Permo-Triassic sandstones in the West Midlands typically occur in fault-bounded basins initiated at the onset of the Variscan Orogeny at the end of the Carboniferous period. Later phases of rifting have also caused the formation of other large faults crossing the sandstone layers (Allen et al., 1997). The large faults may either enhance or reduce transmissivity, and there are examples of both, but the general effect is a reduction of transmissivity across the fault (Allen et al., 1997, 1998). In the Stourbridge area the Western Boundary Fault separates the sandstone from the Carboniferous Coal Measures and the Stapenhill Fault offsets sandstone formations from each other (Robeznieks, 2009). It is not clear whether the Stapenhill fault acts as a hydraulic barrier and it is not represented in the regional model (Cuthbert et al., 2001). This question is addressed further in sections 5.2 and 5.3.4. Smaller scale fractures tend to enhance transmissivity. The generally higher values for formational than for core hydraulic conductivity (table 5.1), especially the extreme maximum for the Bromsgrove Formation, indicate that the fracture flow is significant in much of the West Midlands Permo-Triassic sandstone (core samples do not include fractures). The significance of fractures for hydraulic properties and contaminant migration may be reduced to an extent by sediment infill, however (Wealthall et al., 2001). Allen et al. (1998) suggest that the effects of fracture flow are most pronounced at the pumping test scale, but that at the regional scale matrix flow is the most important mechanism.

Table 5.2: Core porosity measurements (%) in the West Midlands Permo-Triassic Sandstone, from (Allen et al., 1997). Means are arithmetic.

formation	range	mean
Bromsgrove	8.0 - 36.2	26.8
Wildmoor	17.6 - 35.3	26.4
Kidderminster	3.6 - 33.8	23.8
Bridgnorth	16.6 - 33.6	28.4

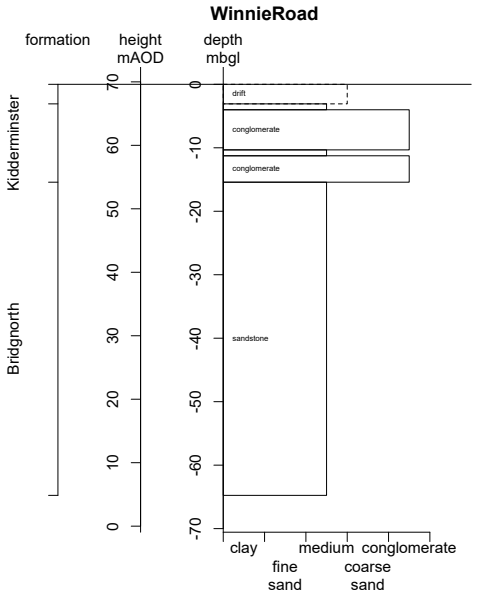

(a) Winnie Road

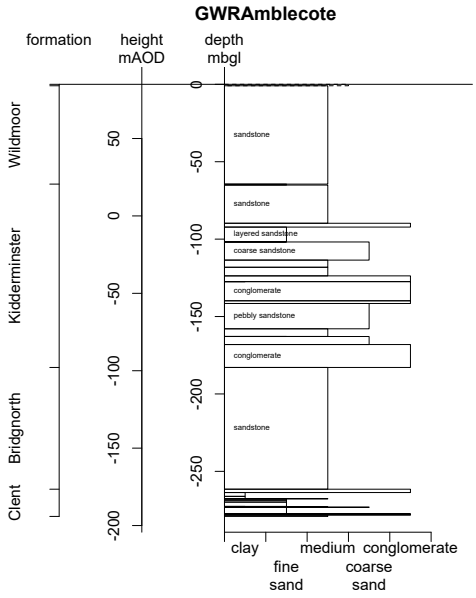

(b) GWR Amblecote

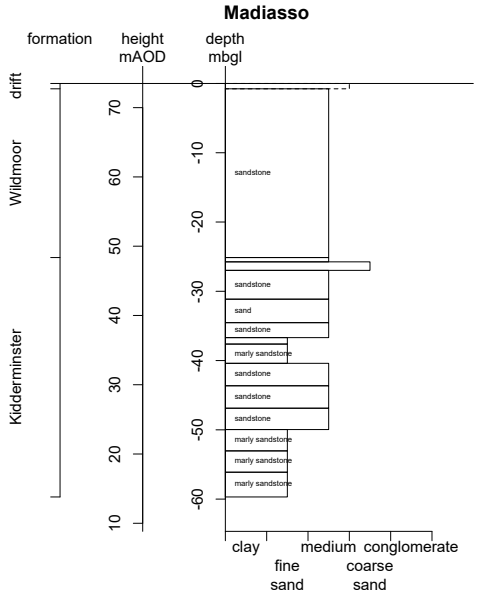

(c) Madiasso

Figure 5.1: Selected synthesised geological logs from the BGS database, demonstrating characteristic lithologies and layering of the sandstone formations and drift in the Stourbridge area. Logs do not tend to give quantitative textural analysis, so texture is represented qualitatively from lithological description. Dashed lines indicate that the lithological description was too vague to infer texture. Based on BGS (2017), with the permission of British Geological Survey.

5.1.2 Organic groundwater quality

Stourbridge sits at the western edge of the West Midlands conurbation and the transition from urban to rural is fairly sharp. The Winnie Road PWS is near this transition and as such its general capture zone surface projection straddles both rural and urban land use.

Much of the southern West Midlands Worfe catchment is affected by chlorinated solvents, largely from historical industrial activities in Kidderminster and Stourbridge (Tyler-Whittle et al., 2002; Tait et al., 2004b; Davies and Gordon, 2013). Stourbridge is an historic industrial centre especially well known for its glass manufacturing industry, and 14 former and current glass manufacturing sites have been identified in the area from historical maps (figure 5.2; table 5.3; appendix C). There are also a number of metals and engineering works which are more strongly associated with chlorinated solvent usage (Rivett et al., 2012).

Abstracted concentrations of tetrachloromethane (TeCM) and, to an extent, trichloroethene (TCE) have been rising since the 1990s, when it was first observed (figure 5.3). The specific source of the contamination has not been identified and anyway it is generally recognised that a remediation of the source zone would still leave a large amount of aqueous contaminant mass in the aquifer (Davies and Gordon, 2013). Two observations about the abstracted concentration and its response to the abstraction rate give some clues about the source location, however. There is a negative correlation between abstraction and concentration, with a decrease in abstraction rate between 2006 and 2007, just before the cessation of pumping, producing a sharp increase in abstracted TeCM concentration (figure 5.3a). And the transient response of concentration during the periods of testing after 2007 shows high initial concentrations followed by a quick reduction (figure 5.3b). Both these observations point to a dissolved-phase contaminant plume that passes directly through the well even in unpumped conditions, such that stronger or more prolonged abstraction draws in cleaner water from further away. Therefore, it is likely that there is a significant contaminant source impacting Winnie Road that is directly upstream of Winnie Road.

TeCM is not a prevalent solvent in groundwater compared to TCE, tetrachloroethene (PCE) and 1,1,1–trichloroethane (1,1,1–TCA), which could have been taken to suggest that it is associated with the glass industry, a particular feature of Stourbridge. However, there is no direct evidence from the literature that would link TeCM to glass manufacture. The historic use of chlorinated solvents in the glass industry is not well constrained, but Arnold et al. (1996) state that they may have been used for degreasing glass and the tools which were used and Cook (1990) indicate that the presence of hydrocarbons (grease) severely inhibits glass polishing, which may have given cause for the use of degreasing solvents. Records from the Catalyst Science Discovery Centre state some of the common uses of TeCM in the 1930s, but these don't include glass (I.C.I., 1931; Solvay, 1938). In the absence of any direct evidence to connect TeCM specifically with the glass industry, then, it is more likely that TeCM happened to be chosen for use by one or a small number of industries in Stourbridge than that TeCM is connected to the glass industry.

Asides from likely source locations and DNAPL behaviour of TeCM, naturally occurring chemical reactions, degradation and sorption, in the aquifer are a key factor controlling its distribution in groundwater. TeCM degradation may occur by three pathways: abiotic hydrolysis, reductive dehalogenation or reductive hydrolysis. Of these, only abiotic hydrolysis occurs in aerobic conditions at all, producing dichlorocarbon sulphide, carbon disulphide and carbon dioxide. This pathway degrades TeCM with a reported half life of around 40 years (Davis et al., 2003) rendering first-order degradation of TeCM in unconfined Permo-Triassic aquifers, which are thought to be generally aerobic (Rivett et al., 2007), very minor. Degradation of TeCM can be much faster in reducing conditions (Davis et al., 2003; Santharam et al., 2014) and produces n-chloromethane (n in 0–3), methanoic acid or carbon monoxide. Sorption may be justifiably represented with a linear isotherm at low aqueous concentration (Schwarzenbach and Westall, 1981; Roberts et al., 1986), meaning that, at equilibrium, the sorbed mass fraction is proportional to the aqueous concentration by a constant K_d. K_d is expressed as $f_\mathrm{oc}K_\mathrm{oc}$ where f_oc is the organic carbon fraction of the sediment and K_oc expresses the sorption capacity of the solute, related to the octanol-water partitioning coefficient (K_ow) with an empirical model. The Piwoni and Banerjee (1989) model used by Taylor (1998) is:

$$\log_{10} K_\mathrm{oc} = 0.691 \log_{10} K_\mathrm{ow} + 0.22 \tag{5.1}$$

With $\log_{10} K_\mathrm{ow}$ for TeCM at 2.7 (Roberts et al., 1986), using $f_\mathrm{oc} = 5 \times 10^{-4}$ and $\rho_\mathrm{b} = 2$ g cm^{-3} (dry bulk density) from Rivett et al. (2005) and porosity $\phi = 0.264$ (section 5.1.1), a retardation factor $R_f = 1 + \rho_\mathrm{b} K_\mathrm{d}/\phi$ is estimated at about 1.5. Roberts et al. (1986) show that when f_oc is very small, sorption onto inorganic solid material can make a significant contribution and inflate R_f to a greater value than one estimated with a literal f_oc. Therefore, $R_f = 2$ is used in this study, similar to the value estimated for PCE (which has similar K_ow: Roberts et al., 1986) in the Tame Valley (section 4.5). Degradation half-life is set at 40 years, representing abiotic hydrolysis, on account of probable aerobic conditions. However, it is acknowledged that some reductive dechlorination is feasible in reducing zones of the aquifer; trichloromethane (TrCM) is observed up to 4 µg l^{-1} at Winnie Road, which may either be from a TrCM or from dechlorination of TeCM.

Table 5.3: Historic industrial sites in the Stourbridge area. Note that opening and closing years may be somewhat approximate, often somewhere between the dates of historical maps where an industrial site appeared on one but not the other. The years 1900 and 2000 were often used if it was reasonably clear that a site opened before or closed after the period of solvent usage. Appropriate unique generic names are given where the site name was not evident from historical maps.

site name	opening year	closing year	site area (m^2)	industry
Albert Glass	1900	1945	3600	glass
Amblecote Works 1	1940	2000	20630	engineering
Amblecote Works 2	1930	2000	10530	galvanising
Amblecote Works 3	1940	2016	2390	metal
Amblecote Works 4	1900	1960	2880	glass
Amblecote Works 5	1940	1970	1940	metal
Amblecote Works 6	1930	2000	5230	iron
Amblecote Foundry	1940	1965	930	iron
Audnam Foundry	1900	1985	9460	engineering
Audnam Glass	1935	2000	5810	glass
Birmingham St. Foundry	1900	1985	4850	metal, engineering
Chain Works	1910	1960	7360	metal, tools
Coalbournhill	1900	2000	3000	glass
College Glass	1900	2000	3280	glass
Dennis Glass	1900	1990	6870	glass
Dial Glass	1886	2016	5930	glass
Engineering Works	1940	1985	2140	engineering
Footbridge Works	1940	1985	5680	metal, engineering
Glass Works 1	1910	1970	1280	glass
Glass Works 2	1910	1930	2480	glass
Glass Works 3	1900	2000	1890	glass
Leather Works	1900	1975	7430	leather
Mill St. Foundry	1900	2000	1760	metal
Old Wharf Saw Mills	1910	2000	6980	metal, tools
Platts Glass	1910	1970	2720	glass
Platts Road	1900	2000	960	engineering, iron
Red House	1788	1936	2880	glass
Stewkins Dial	1940	1990	4180	iron
Stourbridge Iron	1810	2000	16610	iron
Stourbridge Rolling Mills	1908	2000	8680	metal
Titan Works	1910	2000	10010	engineering
Victoria Works	1910	1960	1770	metal, tools
Wollaston Mills	1900	1960	11190	metal, tools
Wordsley	1900	1980	1030	glass
Wordsley Flint	1900	1965	2870	glass
total	1788	2016	180550	

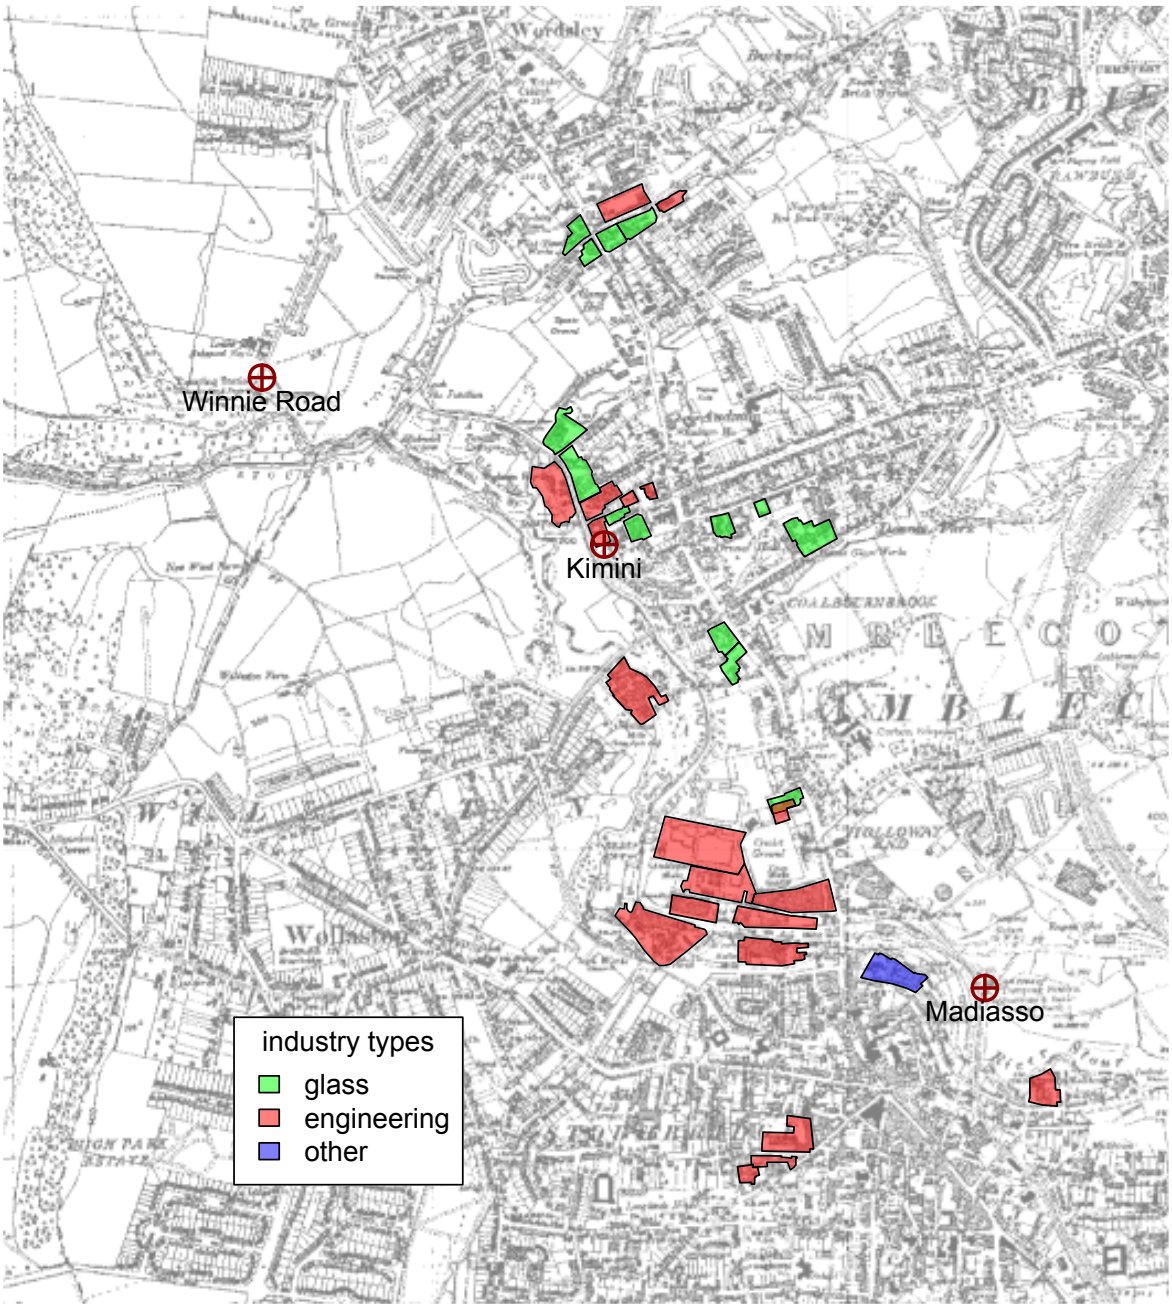

Figure 5.2: Site boundaries of historic and present-day industrial sites and PWS abstractions in the Stourbridge area, superimposed on a map of the area from 1938 (appendix C).

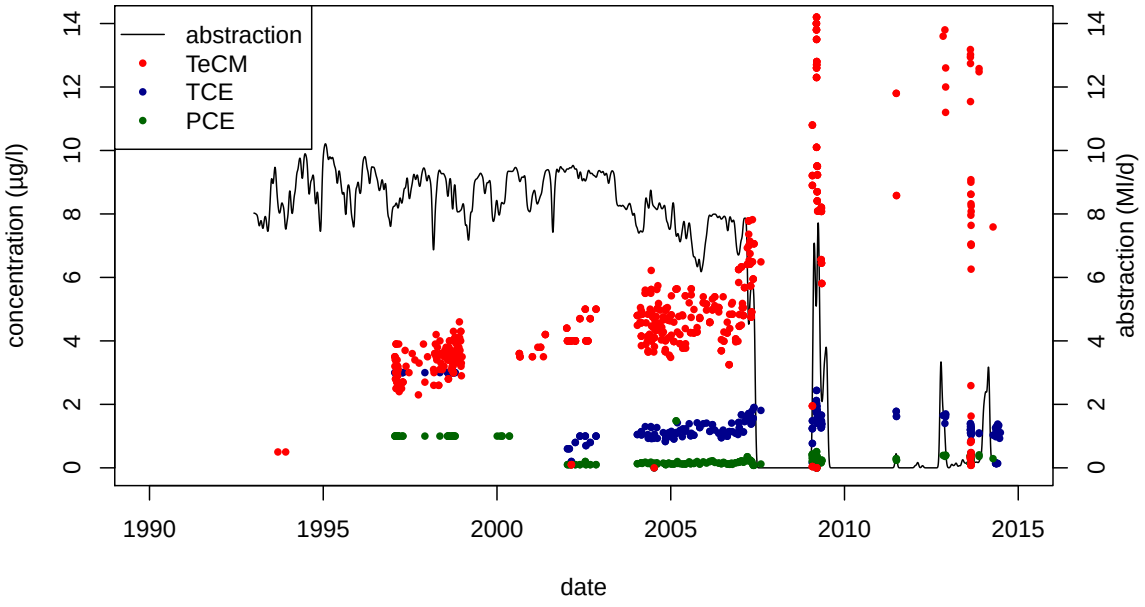

(a)

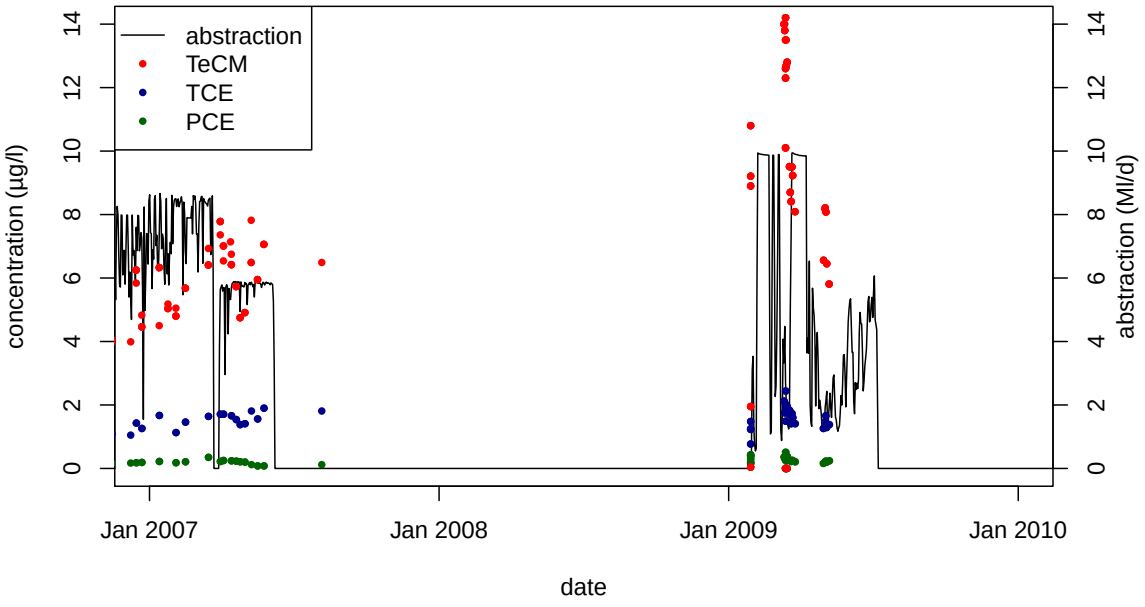

(b)

Figure 5.3: Observed chlorinated solvent concentrations at Winnie Road, compared to abstraction rate. In (a), the abstraction rate is smoothed with a Gaussian mask of standard deviation 10 days. There is a discrepancy in the timings of the abstraction and concentration data during the early 2009 testing, with the concentration data falling in mid-2008. It is assumed, from Davies and Gordon (2013), that the abstraction data is correctly timed (and this data collection is more automated). In these plots, the timing of this period of concentration data is moved, but it cannot be guaranteed that their timing is exactly correct.

5.2 The West Midlands Worfe Regional Model and Possible Adaptations

The Stourbridge aquifer is covered by the West Midlands Worfe regional groundwater model. This groundwater model was developed for the purpose of understanding and managing the water resource systems of the catchments of the Rivers Worfe and Stour, with a particular focus on management strategies during low flow periods (Cuthbert et al., 2001). A summary plot of the model and its boundaries is shown in figure 5.4.

Therefore, as with all the EA's regional groundwater models, the focus of the model is on water resources and distribution of flow. To use this model for the simulation of contaminant transport is a significant repurposing. The regional model covers a much greater area than the Stourbridge area, with the implications that calibration effort is spread broadly and not focussed on Stourbridge and that the grid cells are very large (500×500 m horizontally). The DRW method (chapter 2) is designed to work well with the restrictions encountered for regional-scale transport modelling, but even so there will be limitations to the precision of results in areas of variable flow direction and because wells in MODFLOW must be located at the centres of grid cells. Local grid refinement strategies are also available. Telescopic mesh refinement (TMR) (Rumbaugh and Rumbaugh, 2011) allows a finer-gridded submodel to be extracted from a regional model with edge boundary conditions fixed by the regional model results and MODFLOW-USG (Panday et al., 2013) allows local refinement to be made to a model without altering the grid in other areas. However, a key purpose of this project is to show that the DNAPL source term framework and the DRW method can be used meaningfully with unaltered regional resource models for water quality management, so these refinements are not performed in this study.

Focusing on Stourbridge, it must then be addressed what are the local data constraints that have been used in the model. Figure 5.5 shows a map of the observation boreholes and flow gauges in the vicinity of Stourbridge, some of which have been used to constrain the regional model. Hydrographs from these piezometers are plotted in figure 5.6. There are a good number of observation boreholes within 5 km, although only three screened within the Bridgnorth Formation which Winnie Road abstracts from (Davies and Gordon, 2013). Within 2 km there are three observation points, of which two are boreholes (only recently installed) and one is a flow gauge.

The hydrographs very close to Winnie Road (Kinver Street Shallow and Deep) start recently, in 2010, after the cessation of Winnie Road's abstraction (apart from tests) and roughly coincident with the end of the "recent actual" numerical model (figure 5.6). The hydrographs start in agreement with the numerical simulations and the simulated deep response to the pumping tests at Winnie Road (actual abstraction records for Winnie Road up to 2014 have been used to update the well package for all predictive simulations). The simulated shallow response is stronger than the data implies, probably because hydraulic conductivity is represented as isotropic in the regional model. The pumping tests conducted in early 2014 are captured in the deep piezometer, but not the shallow, indicating some hydraulic separation between the Bridgnorth and

Kidderminster Sandstones. The other piezometers screened in the Bridgnorth Formation, Kinver Old Waterworks and Kings Lodge, corroborate the simulation moderately but not excellently and both records show a long-term decrease in head which is not captured by the model. Of the other hydrographs, Check Hill and Prestwood show good matches to data, although the water table is below the bottom of the Prestwood borehole from about 1990 (personal communication with Matthew Weston, EA). Iverley House Farm and Heath Barn[2] show moderate fits, but the simulation results do not capture the observed water table decline since the 1980s. Heads at Norton 1 and Swinford are overestimated and at Roundhills are underestimated, which is curious because these three piezometers are in similar locations. Hydraulic flow barriers (Harbaugh et al., 2000) are used to represent some faults in the West Midlands Worfe model, but not in the Stourbridge area. However, the Stapenhill Fault runs between the overestimated locations and the underestimated locations (figure 5.5; Tyler-Whittle et al., 2002), which may explain the different hydraulic observations between the two areas. Horizontal hydraulic separation seems to be more important than vertical separation in this case because Swinford and Roundhills are both screened in the Wildmoor Formation and there is about 30 m difference between their observed heads.

Of the flow gauges near Stourbridge, the stream flow results from MODFLOW match well at Prestwood, close to Winnie Road, and moderately at Swindon, further away to the north (Streetly and Rahman, 2011). Fit to data for the nearby Stourton flow gauging station is not presented in the model report.

General groundwater flow in the Stourbridge area is towards the south-west, as estimated by the regional groundwater model in the Stourbridge region. Cuthbert et al. (2001) indicate that, from 1970 to 2001, the average recharge (37 Ml d^{-1}) is exceeded by the average total groundwater abstraction (44 Ml d^{-1}), with the difference being supplied principally by net flow from other regions (3 Ml d^{-1}) and flow from the head-dependent boundary that is used to represent connection to the coal mine works (3 Ml d^{-1}). In this period, groundwater was not depleted according to the model, but the observed hydrographs show a general decline in head that was not reproduced by the model (figure 5.6).

Although, in general, satisfactory calibration was achieved for the West Midlands Worfe regional model, it is not clear from the report why the Stapenhill Fault was not represented as horizontal flow barrier and how the conductance of the head-dependent boundary along the edge of the Coal Measures was decided. Therefore, two plausible adaptations of the West Midlands Worfe model are tested to investigate the effect of certain boundary condition representations on the model calibration and, later, on the plume transport. The adaptations are described in sections 5.2.1 and 5.2.2. The effects of these adaptations on simulated hydrographs around Stourbridge are plotted in figure 5.6. It was found that the adaptations did not result in an improvement on the calibrations to hydrographs. The effects of each adaptation were respectively to

[2] Heath Barn is located in the inactive area of the MODFLOW model, so is compared with results from the adjacent cell, one column to the east. As such its results may be sensitive to the model boundary and less well related to the observations.

change the groundwater flow direction and to reduce the rate of groundwater flow. The implications of this for plume transport to Winnie Road are discussed in section 5.3.4.

5.2.1 Adaptation 1

The Stapenhill Fault, striking approximately north-south and located west of Winnie Road and east of the Prestwood hydrograph (figure 5.5), is represented in this adaptation consistently with other faults in the north of the model, as a horizontal flow barrier with conductance per unit length 2.1×10^{-5} m d^{-1}, using the HFB package in MODFLOW (Harbaugh et al., 2000). The boundary was built and added to the existing horizontal flow barriers for the model in Groundwater Vistas (Rumbaugh and Rumbaugh, 2011) and starting heads for the recent and predictive models were updated using `Rflow` (Barry, 2017e).

Because modelled groundwater flow in the Stourbridge area is generally towards the south-west without this adaptation, the inclusion of the Stapenhill Fault as a flow barrier causes higher hydraulic head on the east side of the fault and lower on the west. This can be seen in the modelled hydrograph results, of which Swinford Common, Norton Borehole 1 and the Kinver Street head results (east of the fault) are increased by the adaptation while the others are all reduced (figure 5.6). This effect gives a slightly closer match to observations at Swinford Common and Roundhills-Roadside, which are separated by the Stapenhill Fault, but otherwise there is not significant improvement.

5.2.2 Adaptation 2

The head-dependent boundary condition along much of the edge of the sandstone aquifer is used to represent indirect recharge through old coal mine works. Across the model, mass balance contribution from the head-dependent boundaries which represent old mine works varies between 25 and 28 Ml d^{-1}. Compared to the wells, abstracting 180 to 280 Ml d^{-1} total, and recharge with average contribution about 480 Ml d^{-1}, this is small, but the local effect on the flow in the Stourbridge area of the model is significant, with the hydraulic gradient descending almost perpendicularly from the Coal Measures. It is therefore worth assessing the dependence of calibration and transport results on this poorly constrained boundary which has a strong influence on the area in question. The second trialled adaptation therefore has the conductance of the head-dependent boundary throughout the model, except for the extreme northern and southern boundaries where the head-dependent boundaries serve a different purpose (figure 5.4), reduced by a factor of ten everywhere (from 3.5 m^2 d^{-1} in the Stourbridge area down to 0.35 m^2 d^{-1}), in order to decrease the importance of these boundaries. This adaptation decreases the transient total flux from mine-related head-dependent boundaries to between 3 and 3.5 Ml d^{-1}, a bit more than a tenth of the original range. The boundary condition modification and starting heads updates were done with `Rflow` (Barry, 2017e).

By increasing the hydraulic resistance to this recharging boundary, the piezometric head across the model was reduced by this adaptation (figure 5.6). This generally resulted in a worse match between simulated and observed heads. Only slightly better matches were achieved where the heads are overestimated (Heath Barn and Roundhills-Roadside). Despite the the decrease in recharging water, the observed general groundwater decline was not reproduced. Therefore it cannot be said that this adaptation better reproduces observations.

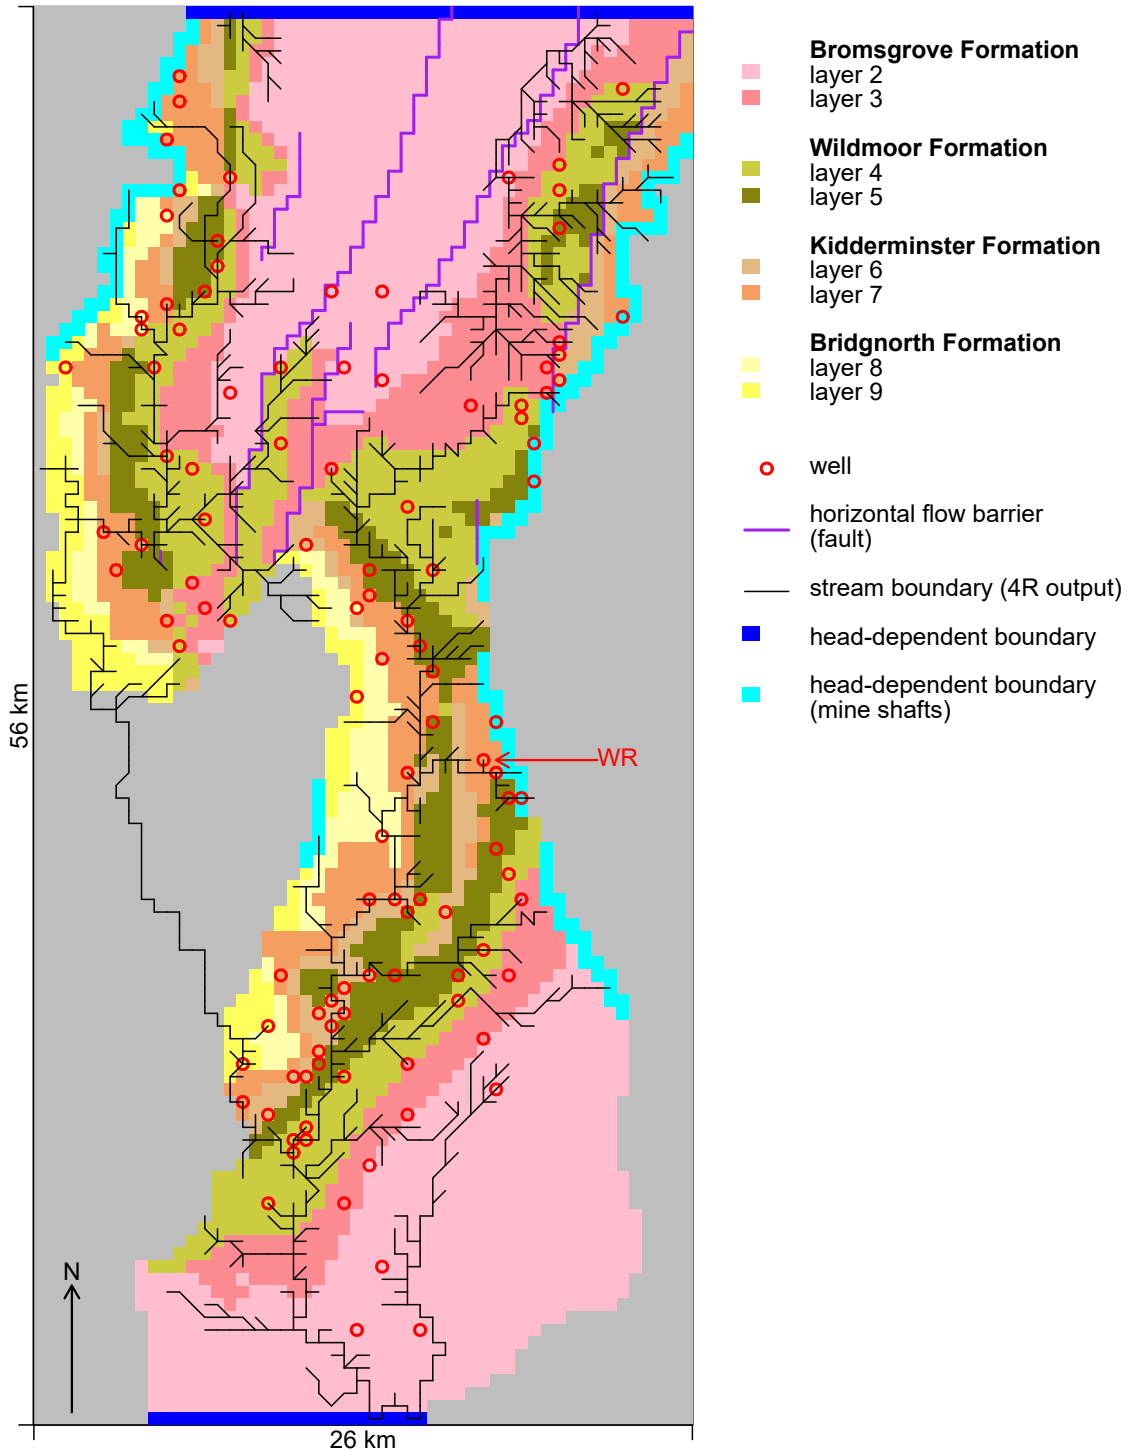

Figure 5.4: Summary of the layers (over-plotted) and boundaries (excluding distributed recharge) in the West Midlands Worfe regional model. Winnie Road PWS well is indicated (WR).

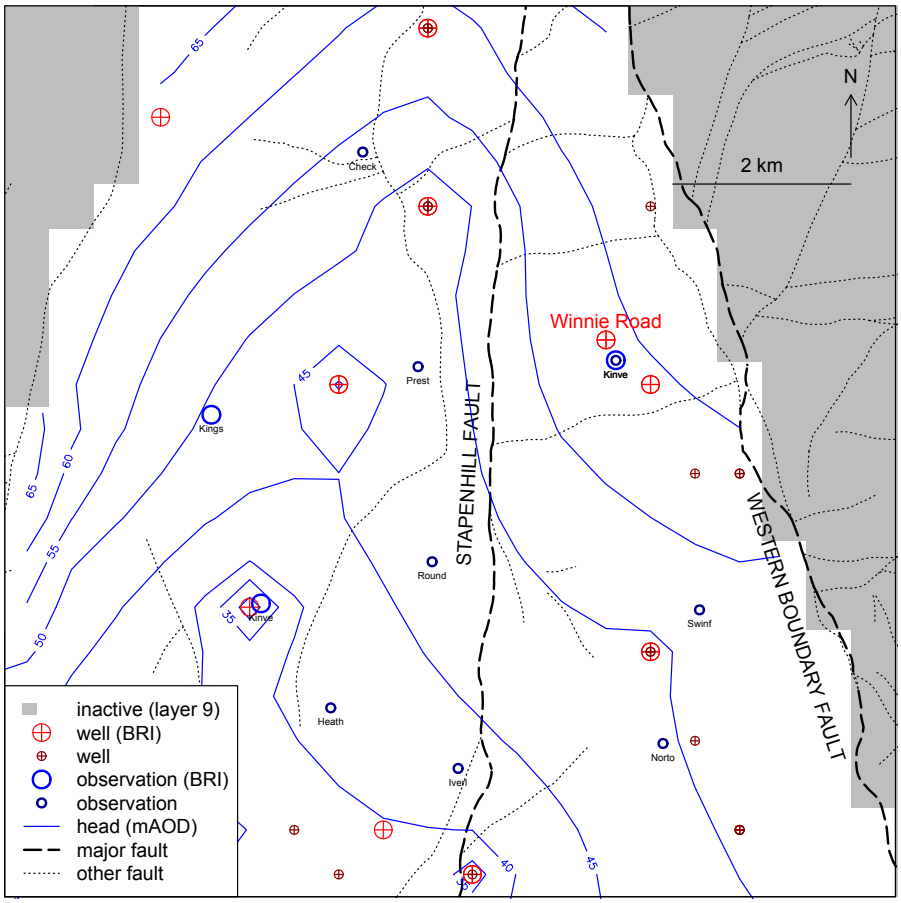

Figure 5.5: Distribution of abstraction wells (red) and observation boreholes (blue, with abbreviated names) as they are located in the MODFLOW models. There are more observations than are used in the regional model (Cuthbert et al., 2001; Streetly and Rahman, 2011). Wells and piezometers screened in the Bridgnorth Sandstone (BRI) are highlighted, because that is the formation that the Winnie Road PWS abstracts from. The two piezometers near Winnie Road are Kinver Street Shallow and Deep, which are recent hydrographs. The other "Kinve" label refers to Kinver Old Waterworks. Head contours (simulated) are for the end of 2010 (that is, the transition from the recent to the predictive simulation). Faults are from (British Geological Survey, 2013) but are not all represented in the model.

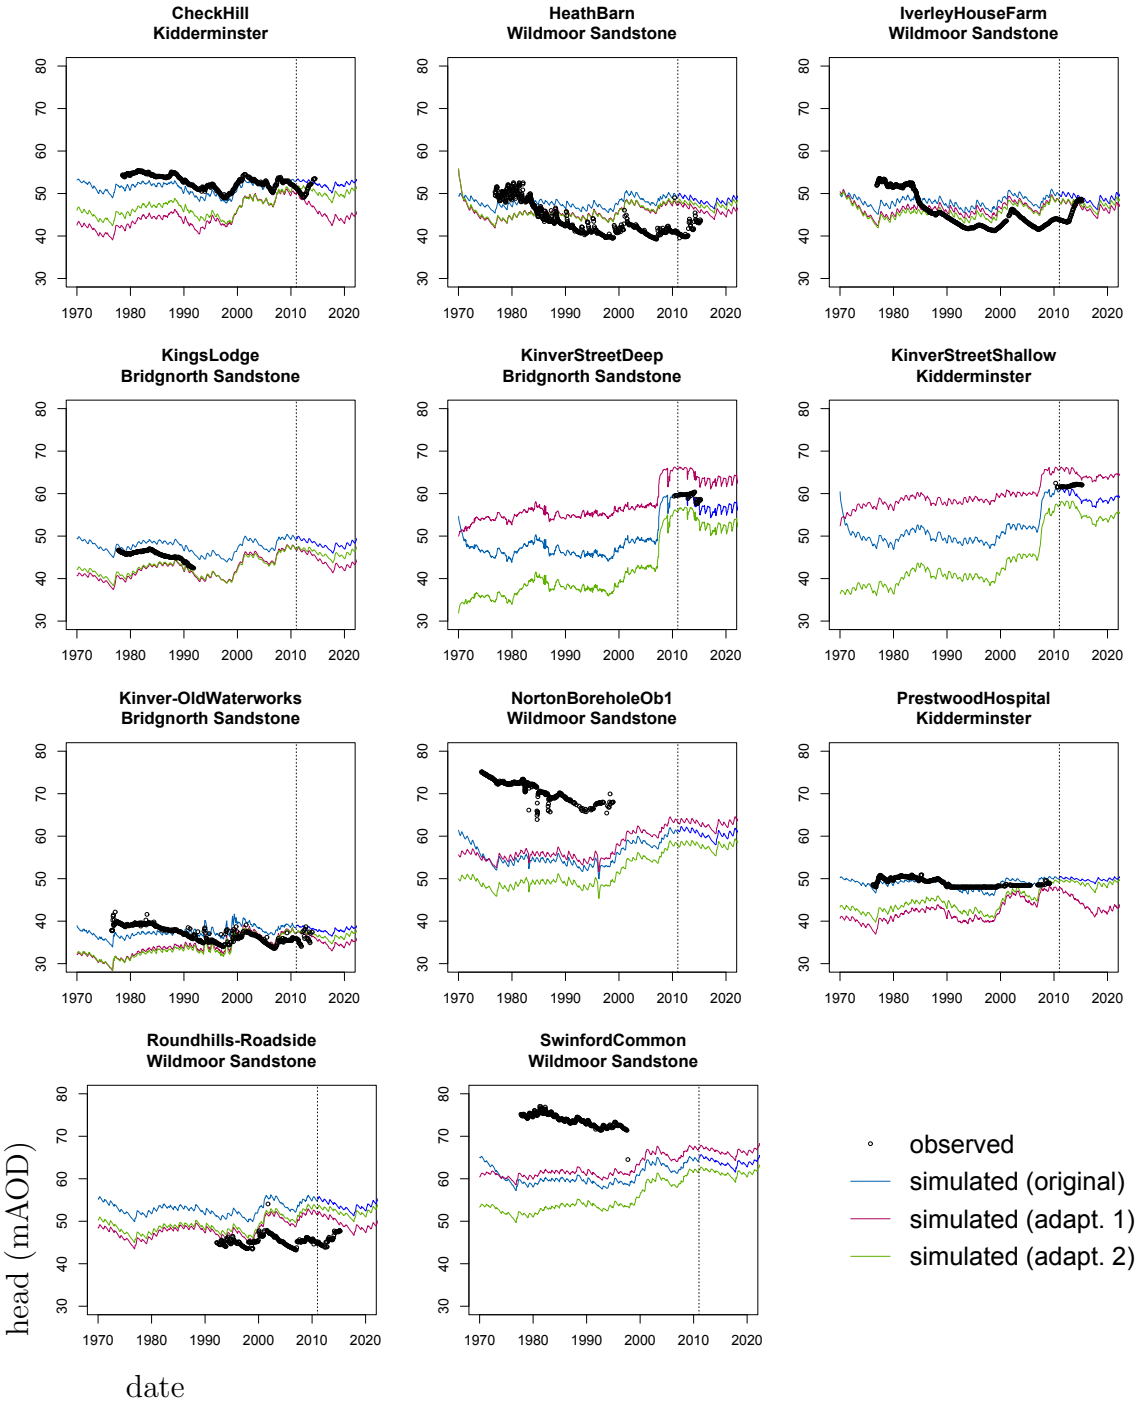

Figure 5.6: Hydrographs of piezometric head in the Stourbridge area (see figure 5.5). The formation in which each piezometer is screened has been indicated. Lines are simulated results for the original regional model and the two boundary condition adaptations and points are observed values. Dashed lines show the time dividing the recent actual models from the predictive models. The abstraction at Winnie Road up to 2014 has been updated to actual values in all versions of the predictive scenario.

5.3 Migration of Potential Contaminant Plumes

5.3.1 Abstraction scenarios

The regional model predictive scenario simulates groundwater flow from January 2011 to December 2052. Well abstractions are cyclical with time and based on recent abstraction patterns (to 2010). In this study, the well package for the predictive model was updated according to actual abstractions at Winnie Road up to March 2014, as far as data was held. During this period, Winnie Road is almost entirely unpumped. Beyond this time period, seven future pumping scenarios were set up, one with Winnie Road pumping cyclically as in the EA's baseline predictive simulation (average abstraction about 3 Ml d^{-1}), and six with Winnie Road given constant abstraction rates between 0 and 10 Ml d^{-1} (figure 5.7)[3]. Note that Winnie Road is represented as a well in both layers 8 and 9 of the regional model (the upper and lower Bridgnorth Formation), and the abstraction is split equally between the layers. These scenarios will test the relative severity of solvent contamination at Winnie Road with different abstraction rates, and the likely transient response as the well is started. Lower abstraction rates are less likely to result in the capture of a given plume, but conversely will result in less dilution of any contaminant that is captured.

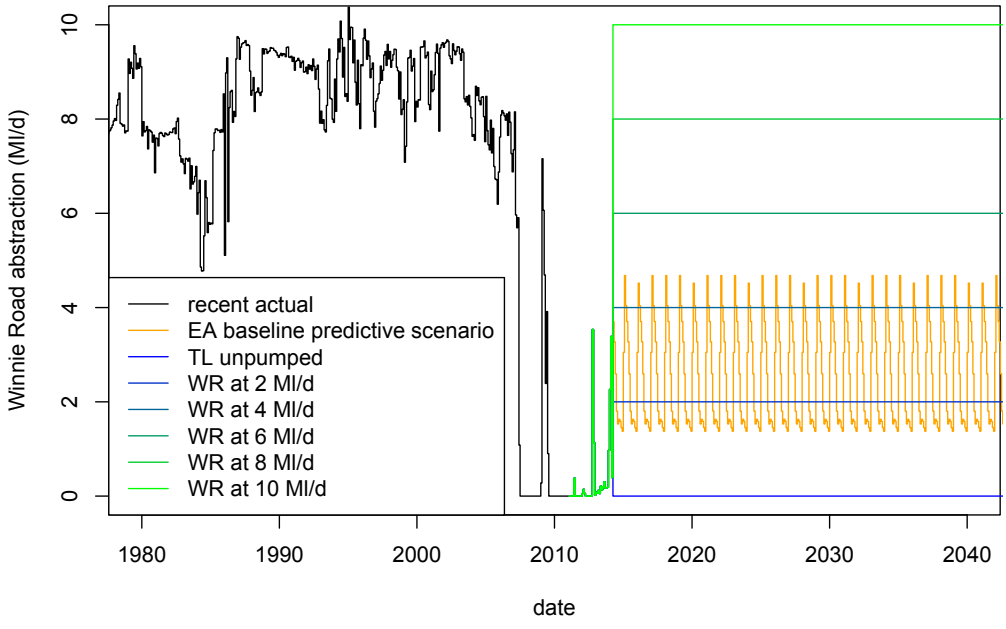

Figure 5.7: Modelled abstraction at Winnie Road in the recent actual model and the predictive scenarios. In all the predictive scenarios, abstraction up to 2014 has been updated with actual abstraction data.

[3] The predicted abstractions used across the EA model baseline predictive scenario is based on monthly average abstraction returns from the period 2006 to 2010 (Streetly and Rahman, 2011).

5.3.2 Conceptual sources

The contaminant transport model is initially run without a representation of the DNAPL source term using a constant, conceptual source term. This experiment is designed to assess the how the advective plume transport is influenced by the different groundwater model adaptations (section 5.2) and the seven Winnie Road future abstraction scenarios (section 5.3.1). The results from these sources do not represent real concentrations and are discussed in terms of conceptual flux and conceptual concentration in units of d^{-1} and m^{-3} respectively (the units do not have a mass component as there is no specific substance being discussed at this stage). The experiment in section 5.5 links aqueous contaminant transport with the DNAPL source term model results to estimate the distribution and abstraction of real concentration (in kg m^{-3}).

The sources for this experiment are based on the locations of historic industries (section 5.1.2). The constant conceptual source terms are constant releases from directly below the centroid of the sites (Bivand and Rundel, 2015) in each active layer which contains phreatic groundwater, with z-offset 0.5 (that is, vertically in the saturated midpoint of the layer). The conceptual source fluxes are all 1 d^{-1} and start for each site when the site started operating (table 5.3). The models are run with sources under each historic industrial site and also under sites belonging to the northern cluster of industrial sites (highlighted in figure 5.8) individually in order to isolate the characteristics of this potential dissolved phase plume. The conceptual concentration plumes will show the general movement of contaminants through the aquifer in the vicinity of Winnie Road and reveal how certain flow regimes and pumping patterns may protect or threaten the water quality at Winnie Road.

Site opening (and closing) are determined through examination of historical maps (appendix C) and searching online records if available. This was often quite approximate and usually an error of 10 years should be assumed. However, there is no systematic bias from this error (estimated dates may equally be too early or too late) and this study considers the evolution of contamination over many decades, so this uncertainty is not severe.

5.3.3 Transport parameters

The aqueous transport component of the model introduces more uncertain parameters. Longitudinal and transverse dispersivity were set at 200 and 20 m respectively, as in the Birmingham regional experiment (section 2.6). As discussed in section 5.1.2, retardation factor was set to 2 and decay constant to $\log 2/40$ a $= 4.74 \times 10^{-5}$ d^{-1}. Effective porosity for aqueous transport is set at 0.15, less than the total porosity and comparable the specific yield of the aquifer to reflect that a large proportion of the flow is transmitted through the fracture network (see section 4.3.3 for a more detailed discussion). The DRW method is used to solve aqueous transport (chapter 2; Barry, 2017c). DRW is a pure Lagrangian method that specialises in point source transport problems on coarse flow grids. The aqueous transport is solved with a constant time

step of 500 days, a coalescing search radius of 60 m horizontally and 40 m vertically and with the dispersion step producing 2 pairs of dispersed particles for each particle (see section 2.5 for explanation of these solver parameters).

5.3.4 Results of the abstraction scenarios

When sources under all the industrial sites are included, the distribution of dissolved phase plumes may be seen across the area (figure 5.8). It is seen that the first adaptation, adding the Stapenhill Fault as a horizontal flow barrier, changes the direction of flow from south-west to south-by-south-west. The second adaptation, reducing the influence of the head-dependent boundary, slows groundwater flow without noticeably changing its direction (this explains the higher but less dispersed normalised concentrations seen in figure 5.8, because the sources are specified as constant mass loadings). However, the abstracted concentration response to different abstraction scenarios is broadly similar when all potential point sources are included. With all sources included, larger abstractions result in a larger abstracted concentration. This is due to the large plume originating from the industrial area south-east of Winnie Road, which is increasingly drawn in by larger abstraction rates (figure 5.8).

Both the two boundary condition adaptations of the regional flow model profoundly change the predicted concentration response to abstraction at Winnie Road when just the plume from the northernmost industrial cluster is modelled. Whereas with the original flow model, the resulting normalised concentration is lower for higher future abstraction rates, with the first adaptation the correlation is reversed (figures 5.9 a and b). This is because this adaptation, by changing the general flow direction in the Stourbridge area from south-west to south-by-south-west, causes the most northerly contaminant plume to miss Winnie Road unless a high abstraction rate is used, which draws the plume in. By contrast, with the original model and the second adaptation, the northern contaminant plume passes through the well even if it is unpumped (figure 5.8). In the former "drawn-in" case, the impact of the plume is seen only with larger abstraction rates which capture the plume. In the latter "direct hit" case, the well abstracts the high-concentration plume centre but higher abstraction rates serve to dilute the concentration by drawing in cleaner water from further away. With the second adaptation, there appears to be no correlation between abstraction and concentration until beyond 2030, whereupon a negative correlation emerges as for the original model. The reasons for this behaviour are less clear. It is likely that with the second adaptation in the period up to 2030, Winnie Road draws in the longitudinally dispersed plume front, with the core of the plume migrating more slowly than in with the original flow model, resulting in a correlation somewhere between the "direct hit" and "drawn-in" cases. As the plume progresses, the situation settles to a negative correlation indicating a "direct hit" situation. The differences in plume migration, in terms of direction and speed, between the original model and the adaptations are visible (although subtle) in figure 5.8.

The observed abstracted concentration at Winnie Road, as discussed in section 5.1.2,

shows a negative correlation between abstraction and concentration. Therefore, as the model which best reproduces this and which has the best calibration to hydrographs (figure 5.6), the original regional model is probably the most appropriate representation of the Stourbridge groundwater flow field. This version of the model is used for the remainder of this chapter. Inspection of the modelled plume distributions (figure 5.8; appendix D.2) also reveals that not many of the potential spill locations are likely to impact Winnie Road. The most northerly potential plume is the most likely to be impacting the well and a plume originating from the large industrial region east of Winnie Road (figure 5.2) may be drawn in at larger abstraction rates, if it exists. A future abstraction rate of 10 Ml d^{-1} or lower will not draw in any of the other potential plumes.

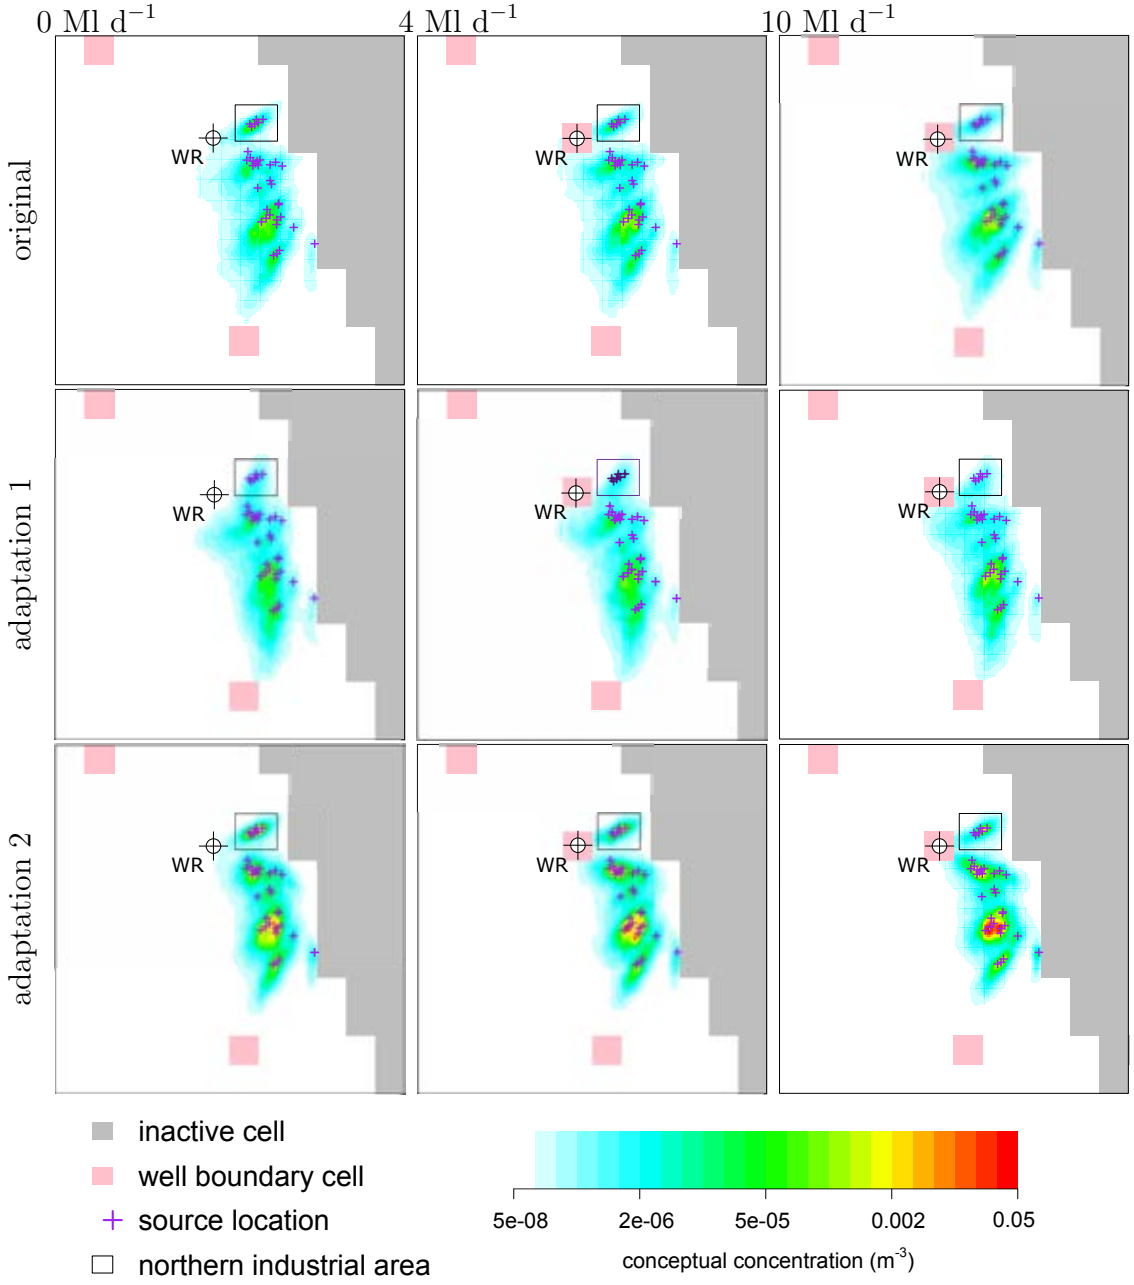

Figure 5.8: Contaminant plume time slices for 2026 in the Bridgnorth Formation originating from all potential source locations. The results use the original West Midlands Worfe groundwater model and the two trialled adaptations (section 5.2) and each is shown with three different future abstraction scenarios for Winnie Road. Note in particular the directions of contaminant plumes in the Winnie Road unpumped future scenario and how this is changed by the model adaptations. Note that the colour flood is spread logarithmically, which makes the plume appear more spread than a linear colour flood. Animated plots of migrating contaminant plumes are included in the supplementary electronic material (appendix D).

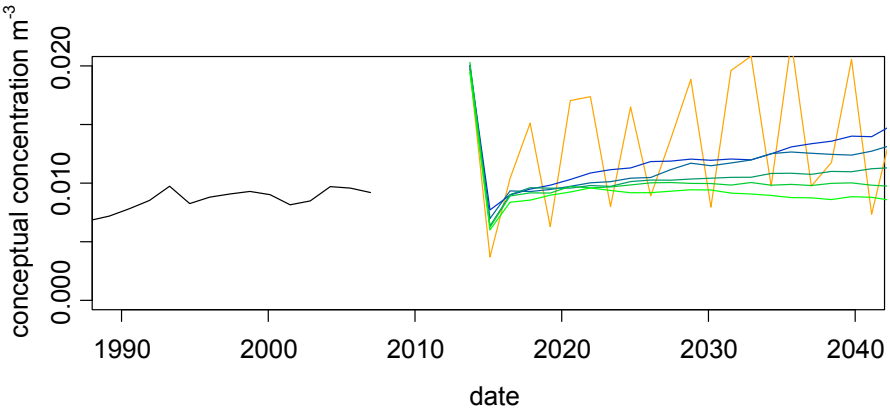

(a) original model

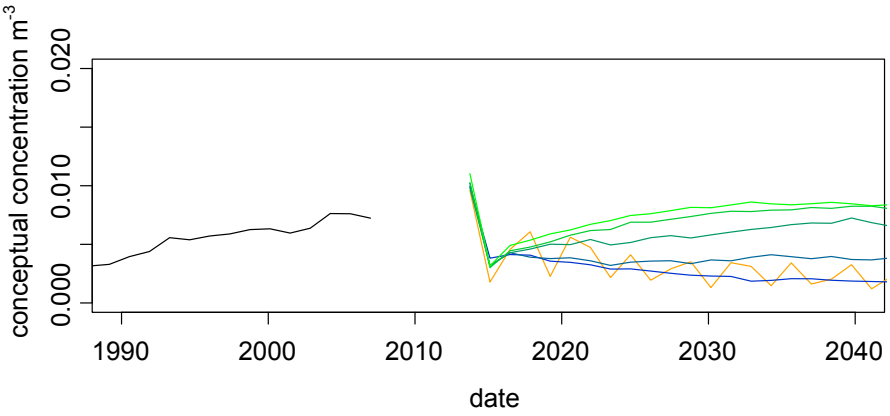

(b) adaptation 1

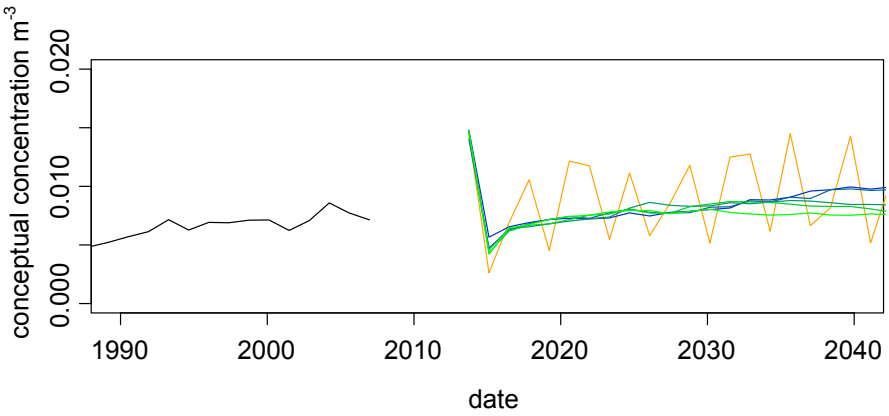

(c) adaptation 2

Figure 5.9: Conceptual concentration results at Winnie Road for each abstraction scenario (see figure 5.7 for legend) with the unchanged West Midlands Worfe regional groundwater model and the two adaptations. Just sources in the northernmost industrial area, highlighted in figure 5.8 are used for these results.

5.4 The Regional DNAPL Source Term

A source term for the Stourbridge region is calculated based on the DNAPL source term framework method (chapter 3; Barry, 2017b). The source term framework tracks chlorinated solvent mass from import to a factory site and routes it through spillage, infiltration, subsurface migration and finally dissolution into flowing groundwater. The method uses a conceptual model that is appropriate to the dissolution of DNAPLs in groundwater. It also allows consistent input parameters to be applied to multiple source zones across a region in order that consistent transient and depth-dependent effluent fluxes may be calculated that respond to local groundwater flows and geological layer thicknesses which are read from a groundwater flow model.

5.4.1 Spill locations

Spill locations are set to be at the centroids of the polygons which define the former industrial site areas (figure 5.2). In reality, solvents may have been disposed of anywhere on site, or even somewhere off the site boundaries, but in a regional-scale study, it is sensible to take the centroids as being the release locations, on average, especially when the distance from industrial sites to receptor is large compared to the sizes of the industrial sites.

The simplifying assumption that there is one spill location for each individual industrial site is used. Because the rate of solvent import (and hence spillage) is modelled to be related to site area, this assumption does not affect the total rate of chlorinated solvent spillage into the aquifer. However, one large spill will tend to migrate deeper into the aquifer than several smaller spills. Therefore, if multiple nearby industries had shared one disposal pit, then it is more likely that deeper aquifer units would be contaminated. Conversely, if one site used multiple storage and disposal areas (as was likely the case in the Tame Valley area: chapter 4), chlorinated solvent DNAPL migration would be more held up in shallower aquifer units, as hydraulic head in the non-aqueous phase liquid (NAPL) phase would be less able to build up in one location in order to overcome entry pressure requirements posed by geological layers acting as capillary barriers.

5.4.2 Layer divisions

The layers in the DNAPL source term model used for each spill location are based on the geological sequence under each location, as rendered in the layer elevations in the regional flow model (Cuthbert et al., 2001). One source term model layer is used for each MODFLOW layer that is not dry, with an additional 1 m thick layer to represent a basal pool that may accumulate at the base of the Bridgnorth Formation, the lowermost aquifer unit (excluding the Clent Formation, which in this area is thin and very heterogeneous and is not represented in the regional model: figure 5.1b). The number of layers used for each site's source term model varies according to which layer

contains the water table. Pool heights in this study represent the total height of pools in a layer, and are defined as a fraction of layer height (section 5.4.3), so the number of layer subdivisions does not affect the total amount of mass that may be held in the pool domain. The phreatic thickness and hence the uppermost active layer at a location was based on the average water table elevation. An example of the layer divisions for a source term is shown in table 5.4.

Table 5.4: DNAPL source term layer thicknesses under the Amblecote Works 1 site as an example, for which the water table is in the Wildmoor Formation.

description	thickness (m)	MODFLOW layer and z-offset of corresponding point release
Wildmoor Formation (WRS), up to water table	57.5	5, 0.5
Kidderminster Formation (KDM)	30.1	6, 0.5
Kidderminster Formation (KDM)	31.0	7, 0.5
Bridgnorth Formation (BRI)	73.0	8, 0.5
Bridgnorth Formation (BRI)	72.0	9, 0.507
basal pool at base of Bridgnorth Formation	1.0	9, 0.007

5.4.3 Source term inputs

TeCM source terms were generated from each spill site, using consistent source term parameters for one complete simulation of the source terms in the area.

The input spillage flux of chlorinated solvents were applied using the generic usage history discussed in chapter 3 (figure 3.4), with environmental legislation reducing the proportion of imported solvent that was spilt unsafely between 1974 and 1990. For each site, this transient spillage rate was clipped to the operational period of the factory, where relevant, and scaled to some peak usage per unit area of the site. The peak usage per unit area prescribes the import rate of chlorinated solvent to an industrial site of a given area at the peak of national chlorinated solvent consumption, in the 1950s and 1960s for TeCM. There is almost no quantitative constraint on this parameter and it requires human intuition to consider how much solvent a factory is likely to have needed during a working day. Multiple values spanning orders of magnitude are tested

in this study, which would correspond to a total peak usage rate for the region between 34 and 3400 tonnes per year if all sites were indeed using TeCM, which is unlikely. The upper limit of this range is greater than the total predicted UK consumption of TeCM in the 1930s (1200 tonnes per year: figure 3.4; I.C.I., 1931). The consideration of solvent usage rate as an input to the transient source term is a novel feature for chlorinated solvent regional risk models, where a constant recharge-leached source is a more common paradigm (for example: Tait et al., 2004a).

The dimensions of DNAPL pools will have a strong influence on the DNAPL capacity of each layer and therefore the penetration of DNAPL into lower layers. Additionally, a layer with greater mass of DNAPL will give a more persistent flux, simply because of greater source mass to be depleted. More heterogeneous rock formations would be expected to hold more DNAPL pools as the more pronounced layering will result in a greater hydraulic barrier to the downward migration of DNAPL (Kueper et al., 2003; Lawrence et al., 2006). Therefore the pool height is varied to represent varying layer heterogeneity. In this case, pool height should be understood to represent the total height of pools in a layer, rather than of a single pool at the bottom of a layer as figure 3.3 might imply. Pool height, as a proportion of layer thickness is specified for the Wildmoor, Kidderminster and Bridgnorth Formations with a ratio of 2 : 5 : 1 respectively; a pool height multiplier is used to scale this ratio to values for the proportion of layer thicknesses (table 5.5). The basal layer has a pool height of 1 m. A greater pool height may reflect that the layer has pools of greater height due to greater entry pressure contrasts between source term model layers, or a greater number of pools due to more capillary barriers within the source term model layer, or both. The effects of the maximum pool width — the width that pools reach before the layer is considered to be filled, reflective of the lateral continuity of layers — are somewhat similar to the effects of pool height, so this parameter is given a constant value of 20 m in all layers apart from the basal pool layer where the pool is allowed to spread without limit if the impermeable aquifer base assumption is taken, or 20 m otherwise (see table 5.5). Ganglia width is set at 0.5 m as for the Tame Valley study (chapter 4) and similar to the model results from Parker and Park (2004).

The empirical constant power law relationship between mass reduction and flux reduction is used to represent the NAPL dissolution (section 3.2; Zhu and Sykes, 2004; Falta et al., 2005b; Jawitz et al., 2005). The empirical parameter Γ is uncertain and varied to test possible values, with lower values representing more persistent source terms with lower flux at early times (DiFilippo and Brusseau, 2008).

Table 5.5 shows the key uncertain parameters that are thought to have a strong influence on the source term results. Source terms for all uncertain parameter combinations were simulated (150 in total) for TeCM from 1925 to 2050. Trialling all uncertain parameter combinations, it is necessary to have a small number of possible values for each parameter. However, the advantage is that it is much easier to isolate the effect of varying each parameter. For example, the trialled values of peak usage per unit area (table 5.5) are intended to approximate a log-uniform probability distribution function and it is deemed sufficient to allow only five values within that distribution so that all combinations can feasibly be trialled.

The water solubility and density of the chlorinated solvents from table 3.1 are used. The total sandstone porosity was set at 0.264, similar to the Tame Valley study (chapter 4). Allen et al. (1997) report mean values for porosity of core samples similar to this from all West Midlands sandstones (table 5.2). The residual saturations of NAPL and water are both set at 0.1. The source term evolution is solved with a time step of 100 days from 1925 to 2050, using the predictive flow field from the regional model predictive scenario based on recent abstractions, except that abstraction at the Winnie Road PWS was updated to actual abstractions up to March 2014, as in section 5.3.1.

Table 5.5: Trialled parameters relating to the DNAPL source term model.

parameter	trialled values				qualitative effect
peak usage per unit area (kg d^{-1} m^{-2})	5.0×10^{-4} 1.6×10^{-3} 5.0×10^{-3} 1.6×10^{-2} 5.0×10^{-2}				direct effect on spill and infiltration rate, which strongly influence depth penetration and persistence of source zone
Γ	0.2, 0.5, 1, 2, 5				smaller values of Γ result in higher effluent flux at early times and a longer tailing period
pool height as proportion of layer thicknesses		WRS	KDM	BRI	greater pool height results in source zones that penetrate less deep
	(i)	0.04	0.1	0.02	
	(ii)	0.1	0.25	0.05	
	(iii)	0.2	0.5	0.1	
impermeable aquifer base?	no, yes				whether DNAPL mass may be lost through the base of the aquifer (if not, there is no limit to the mass in the basal pool)

5.4.4 Source term results

The results mostly show the expected variations with the trialled inputs (table 5.5). Only source terms generated with the greater usage (and hence spillage) rates penetrate significantly to the deepest layers (figure 5.10). The Kidderminster Formation, which is more heterogeneous (section 5.1.1) and assigned greater total pool heights, acts as a substantial barrier to the downward migration of DNAPL. As such, the Bridgnorth Formation only becomes directly contaminated by chlorinated solvent DNAPL under large sites or if a large peak usage is specified. Pool height also has a large influence on depth penetration, with a layer able to contain more pooled DNAPL mass tending to retain rather than leak DNAPL (figure 5.12). The effects of varying Γ, by contrast,

have no influence on the depth penetration of DNAPL. Rather, the importance of Γ is seen in the late-stage source term, after the spillage of chlorinated solvent has ceased (figure 5.11).

Source term strength and persistence are, logically, in conflict, as a source term with greater dissolution flux will be depleted more quickly. This trade-off is controlled most directly by Γ, with lower values resulting in lower peak flux at early times, persisting longer than for higher Γ values. Figure 5.11 most clearly distinguishes the effect of different Γ values, although generally this was not found to be a very sensitive parameter. Almost all calculated source zones are persistent, enduring a long time beyond the cessation of TeCM spillage. The higher Γ source terms appear to sustain a greater effluent flux after the cessation of spillage, but those with lower Γ will persist longer into the future beyond the date ranges of the modelled source terms because of less efficient mass removal at earlier dates.

The Kidderminster Formation is allowed to contain more NAPL mass relative to its thickness than the Bridgnorth Formation (apart from the basal pool) to reflect its greater heterogeneity, and this behaviour is controlled using the pool height parameter in the DNAPL distribution and dissolution model. The Bridgnorth Formation under the College Glass Works site, for example, is effectively protected from DNAPL migration by the overlying Kidderminster Formation (figure 5.12). It is only directly contaminated by DNAPL when lower pool heights values are input. However, dissolved phase contamination can still reach the Bridgnorth Formation even when DNAPL does not penetrate by flow with a downward component, for example induced by abstraction from the Bridgnorth Formation.

If the impermeable aquifer base assumption is used, the basal pool layer (lowest) is specified differently, allowing a limitless pool to form (section 5.4.3). The impermeable base assumption greatly increases the predicted effluent flux from the basal pool in some cases. However, this is only seen where there is sufficient spilled mass to create a source zone penetrating to the base of the Bridgnorth Formation, at large sites or with large usage rates per unit area (figure 5.13). In these cases, with the impermeable base assumption, the effluent flux from the basal pool is significant but not dominant. Without the impermeable base assumption, the source contribution from the basal pool is very small. For moderate and small spill quantities, there is insufficient mass to form a source zone at the base of the aquifer and in these cases the impermeable base assumption is obviously inconsequential.

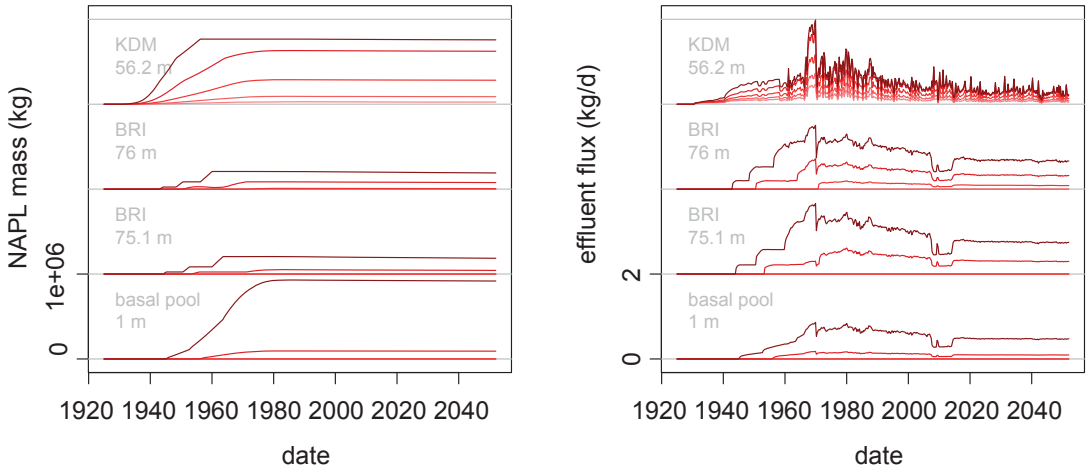

Figure 5.10: Source term mass and flux by layer for the Footbridge Works site, a medium-sized industrial site relative to others in the area (table 5.3), aggregated by peak usage. This figure shows how depth penetration of DNAPL mass and strength of effluent flux is influenced by the usage rate parameter. In this study, usage rate is given indirectly using a peak usage per unit site area parameter and using the estimated national consumption trend to give transient variation (figure 3.4).

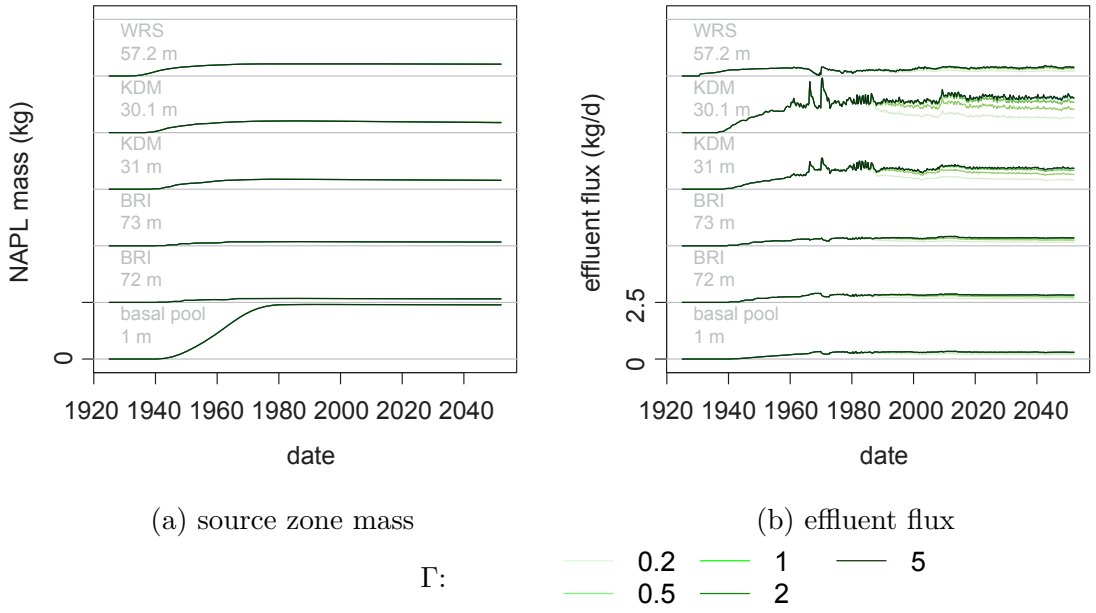

(a) source zone mass

(b) effluent flux

Figure 5.11: Source term mass and flux by layer for the Amblecote 1 site, a large industrial site relative to others in the area (table 5.3), aggregated by Γ. This figure shows how decline in effluent flux is influenced by Γ. Lower Γ gives a faster reduction in flux but in the very long term will result in a more persistent low-flux source term. Of all the four parameters tested, however, Γ is the least sensitive.

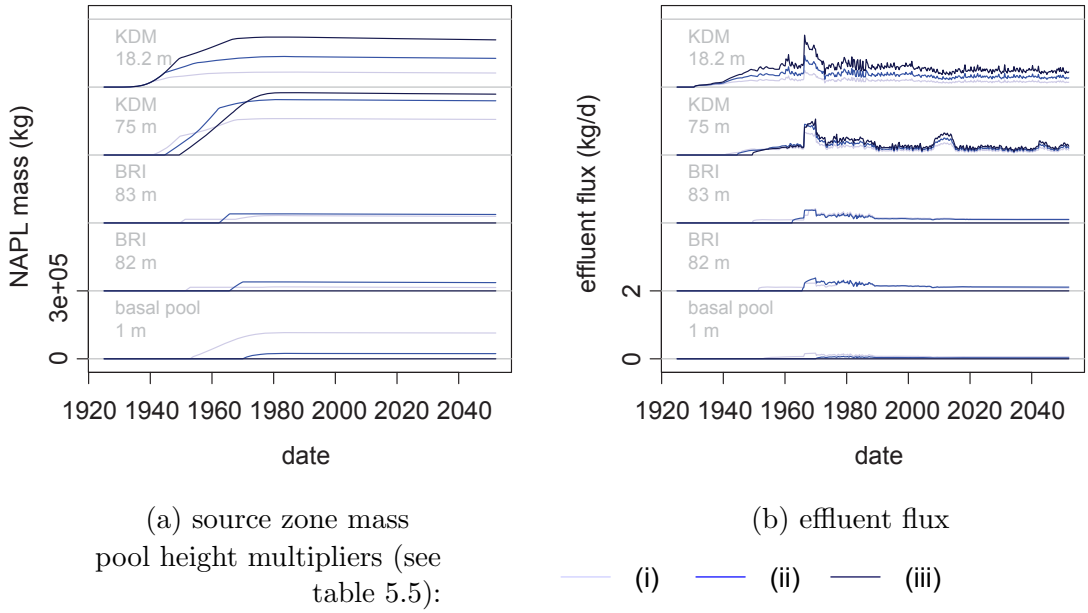

(a) source zone mass pool height multipliers (see table 5.5):

(b) effluent flux

Figure 5.12: Source term mass and flux by layer for the College Glass site, a medium-sized industrial site relative to others in the area (table 5.3) aggregated by pool height multiplier (see table 5.5). This figure shows how depth penetration of DNAPL mass and therefore flux is influenced by the pool height parameter. Note also that the Kidderminster Formation (KDM) layers retain considerably more DNAPL mass than the Bridgnorth Formation (BRI) layers. In all source term realisations, the pool height in the Kidderminster Formation relative to saturated thickness was greater than for the Bridgnorth Formation to reflect its greater heterogeneity.

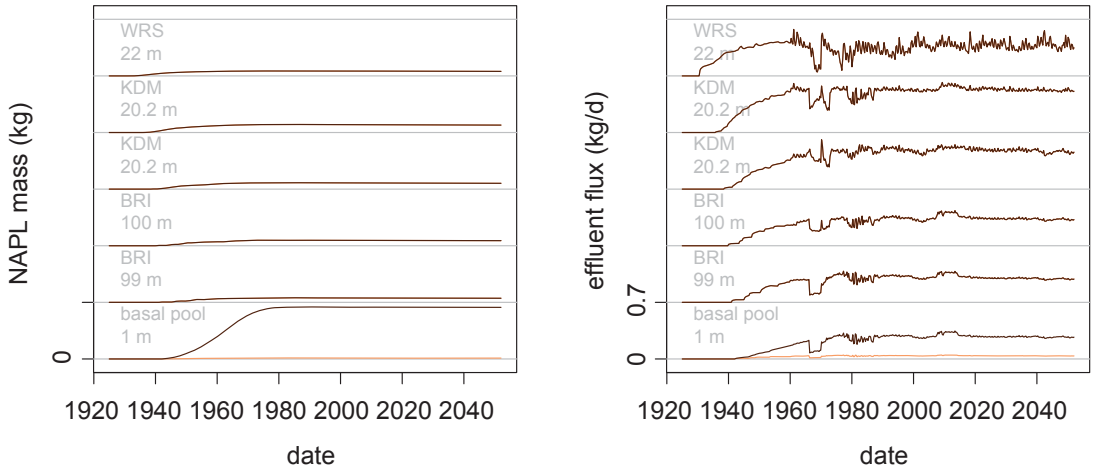

(a) source zone mass

(b) effluent flux

Figure 5.13: Source term mass and flux by layer for the Wollaston Mills site, a fairly large industrial site relative to others in the area (table 5.3), aggregated by whether an impermeable aquifer base is assumed. This figure shows how the impermeable aquifer base assumption affects the accumulation of mass in and hence effluent flux from the basal pool. This assumption is less significant for lower-mass source zones because DNAPL mass does not penetrate to the base in these cases anyway.

5.5 The Coupled DNAPL Source Term and Plume Migration Model of Chlorinated Solvent Concentration

In order to model abstracted concentrations of TeCM, the DNAPL source term model (section 5.4) was coupled to the DRW transport model (section 5.3). Only source terms for which the impermeable aquifer base assumption was in place were used for this experiment in order to save computational time and because not much mass was predicted to be lost from the base (section 5.4.4). The source terms were coupled to flow scenarios with Winnie Road pumping at 2, 6 and 10 Ml d^{-1} (section 5.3.1) and source terms needed to be recalculated for each flow scenario because the dissolution flux from each source term is directly related to the Darcy flux at the source zone location (section 3.3).

The DNAPL source terms account for the depth-penetration of the contaminant sources and the transient decline of the source terms following the period of solvent spillage. Because Winnie Road is screened in the Bridgnorth Formation, DNAPL source terms that do not penetrate through the Kidderminster Formation are likely to cause lower concentrations in abstracted water at Winnie Road. And consideration of the transient decline of the DNAPL source will enable forecasting of future water quality risk to the well. Aqueous transport parameters are the same as in section 5.3.3.

5.5.1 Results

The concentration results are large (figures 5.14), up to two orders of magnitude greater than observed concentrations at Winnie Road (figure 5.3) with larger usage rates. Concentrations are always well below TeCM solubility though. The most likely reason for the high modelled concentrations is that only a small number of the modelled source zones are actually real. TeCM is an uncommon solvent and so there is unlikely to be many real source terms for this contaminant. The results favour a smaller peak usage per unit area, with the smallest of the trialled values for this parameter resulting in concentration about a factor of ten too large. Faster degradation of dissolved TeCM in anaerobic pockets of aquifer may also play a role (section 5.1.2). It is probable that there are many fewer actual TeCM sources than included in the initial conceptual model.

The results, when colour-coded by parameter (figure 5.15), reveal that for the abstracted concentration of TeCM, the most sensitive parameter by far is the usage rate. Although individual source terms show sensitivity to other parameters (figures 5.11, 5.12 and 5.13), it is apparent that for the collective regional source term, the spillage rate is by far the most important with the source term conceptual model used in this study. It is perhaps surprising especially that the pool height should not be more significant, because it governs the amount of NAPL mass help up in higher geological layers, especially the Kidderminster Formation in this case (figure 5.12). At the early stages

of the breakthrough, between 1950 and 1970, there is some distinction between source terms with different pool heights (figure 5.15c). Greater total pool height means that, initially, less mass penetrates to the Bridgnorth Formation, so the early-stage source term is less strong. Later though (and fairly soon if the spill rate is large), greater overall pool heights means that more DNAPL mass may be held in the Bridgnorth Formation, so these source terms become stronger.

The scenario in which Winnie Road pumps at 10 Ml d^{-1} shows rising flux and concentration shortly after abstraction is resumed (figure 5.14c). This is not observed in the other pumping scenarios which were tested. It is not obvious what this rise might be attributed to, as all source terms are declining at this stage with no fresh solvent spillage beyond 1995. Examination of the plume migration (appendix D.2) suggests that this transient feature of the concentration forecast is due to the late arrival of a large plume. This plume originates from the cluster of industrial sites south-east of Winnie Road. According to the simulations, this plume would have been diverted from Winnie Road by abstraction at Kimini[4], south-east of Winnie Road, between 1970 and 1990. The model may not precisely represent the impact of the historic Kimini abstraction on Winnie Road because of the large grid cell size, meaning that the wells are not precisely located in space, which could be addressed with further study using grid refinement. However, it is likely that Kimini would have captured some of the solvent plumes in this area, inadvertently protecting Winnie Road for a while. The full impact of any plumes from this industry cluster, if real, may be seen some time in the future.

The variation of forecast concentration in the predictive abstraction scenarios is less clear than in the experiment without the DNAPL source term model (section 5.3.4). As Winnie Road comes back into operation in 2014, for a short period the 10 Ml d^{-1} abstraction rate achieves the lowest concentration (figure 5.14). However, beyond 2025 the correlation is reversed with the arrival of the later plume, and the 10 Ml d^{-1} abstraction yields the highest TeCM concentration by 2030. This difference is because the later plume originates from larger industrial sites (modelled as source zones with greater mass) than the earlier plume and therefore dominates the abstracted concentration response. Additionally, the earlier plume originated from industrial sites that closed, and therefore stopped spilling solvent, in general earlier than the industrial cluster producing the later plume, meaning that there is much less source mass for the earlier plume. The later plume, from larger sites, is drawn in from the periphery of Winnie Road's capture zone, so larger abstraction rates will result in a greater concentration, in contrast to the earlier plume, which travels directly towards Winnie Road (a "direct hit" situation: section 5.3.4). Therefore, whichever plume is dominant will dictate whether the abstracted concentration will correlate positively or negatively to abstraction rate. From these considerations and given the observed correlation between TeCM concentration and abstraction (section 5.1.2), it is likely that there is a contaminant source in the most northerly industry cluster. The existence of a DNAPL source zone under the former industrial area to the south-east of Winnie Road is purely speculative.

[4]see footnote 1

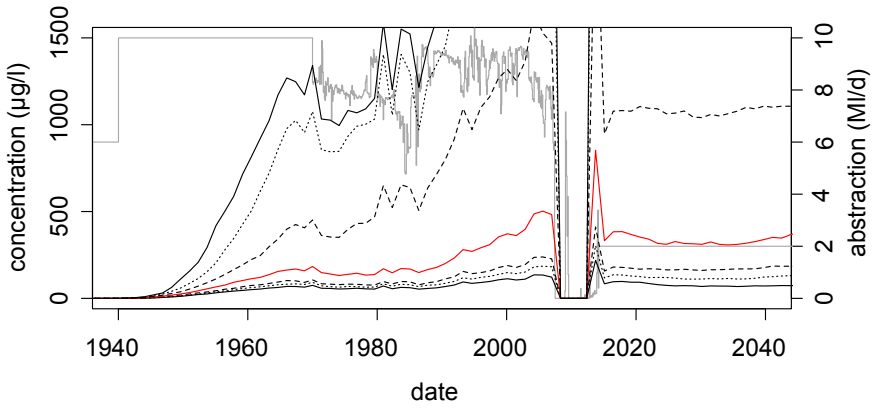

(a) Winnie Road pumped at 2 Ml d^{-1} after 2014

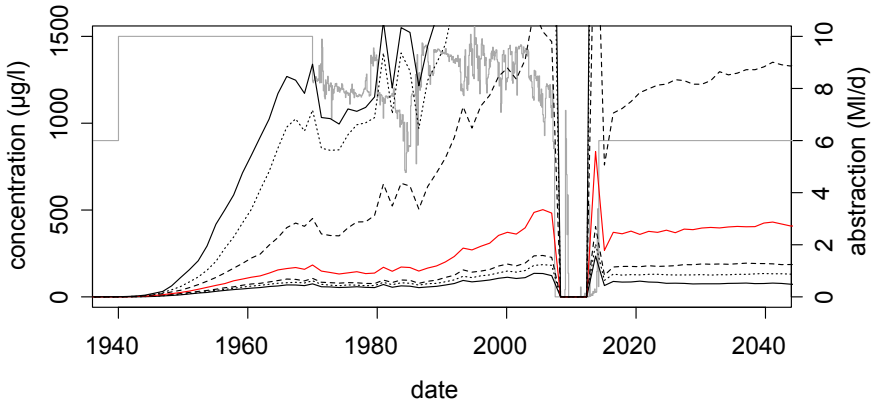

(b) Winnie Road pumped at 6 Ml d^{-1} after 2014

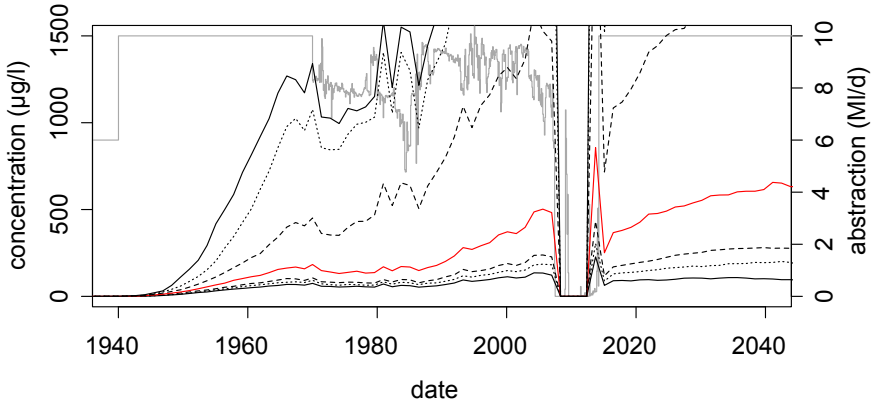

(c) Winnie Road pumped at 10 Ml d^{-1} after 2014

Figure 5.14: Quantile plots of the 75 modelled results for three Winnie Road predictive abstraction scenarios. The 0%, 10%, 25%, 50% (red), 75%, 90% and 100% quantiles are shown. The grey line shows the modelled abstraction rate.

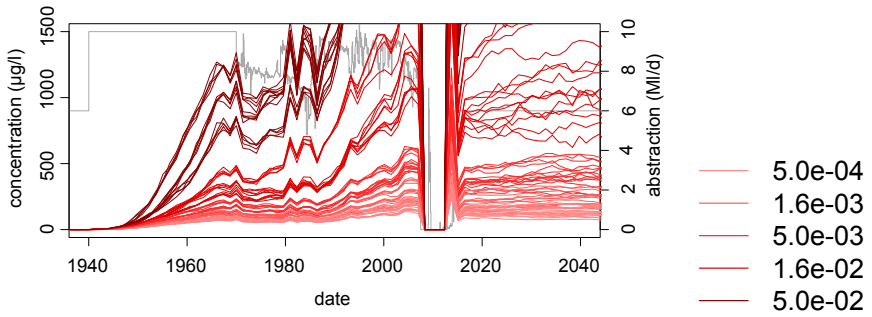

(a) peak usage per unit area (kg d^{-1} m^{-2})

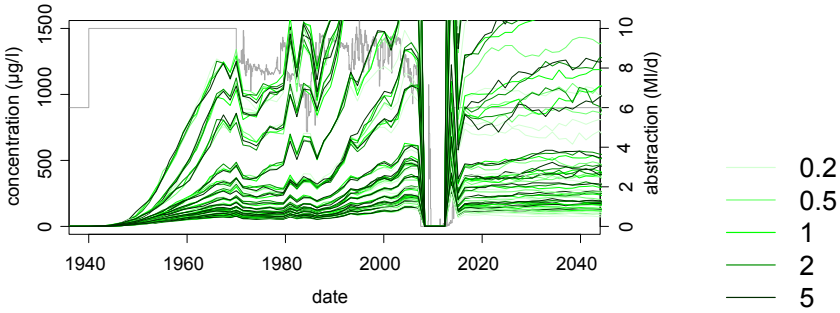

(b) Γ

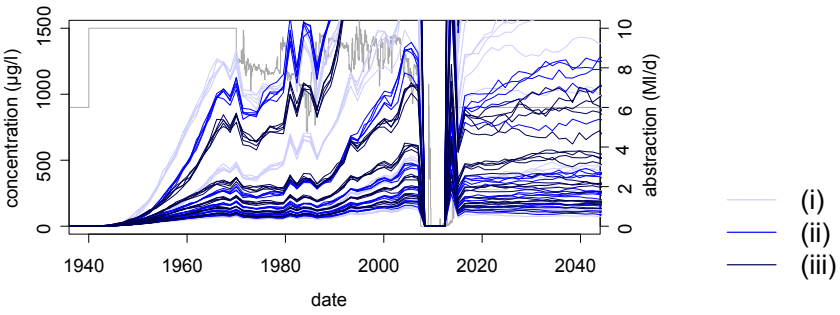

(c) pool height multiplier

Figure 5.15: Abstracted concentration results at Winnie Road for the ensemble of 75 realisations. Each plot shows the same results colour-coded by the value of a different input parameter, shown in the legends, in order to show the sensitivity to different parameters. The coloured lines show concentration and the dark grey line in each plot gives the modelled abstraction rate; the results with the 6 Ml d^{-1} Winnie Road abstraction predictive scenario are shown. Refer to table 5.5 for further explanation of the parameters.

5.5.2 The likely DNAPL source term impacting Winnie Road

In light of the model results from all 35 potential spill locations, an estimate of the real TeCM source term can be made. Because of the observed negative correlation between abstraction and concentration and the observed high initial concentrations when pumping is restarted (section 5.1.2), the dissolved-phase plume is mostly likely to originate from the most northerly cluster of industrial sites, which is directly upstream of Winnie Road. And because the modelled concentrations are much higher than observed, the peak solvent usage rate per unit of site area is likely to be at the lower end of the trialled values and there is perhaps just one real spill location. For this final model, a single source within the northern cluster is simulated with a peak usage rate of 0.4 kg d^{-1}. It is stressed that there is no direct evidence for a solvent spill at this exact location, but model results and observed concentrations seem to support the presence of a TeCM DNAPL source somewhere in the vicinity of this source location. Pool heights, as fractions of layer heights, of 0.05 and 0.01 are used for the Kidderminster and Bridgnorth Formations respectively (the water table is below the Wildmoor Formation at this location); these values maintain the (2 :)5 : 1 ratio (section 5.4.3) and were arrived at after some trial and calibration. The 6 Ml d^{-1} abstraction scenario is used for the future abstraction of Winnie Road (figure 5.7). A time step of 500 days is used generally and of 30 days in the 2007–2015 period in order to capture the short time-scale transient variations during the pump tests, though this is limited by the stress period lengths (1 month) in the MODFLOW model.

The results from this source term are plotted in figure 5.16. The DNAPL in the source term model does not actually penetrate into the Bridgnorth formation. The Bridgnorth Formation is contaminated by downward flow from the Kidderminster Formation, likely induced by the operation of Winnie Road. The evolution of source term mass (figure 5.16a) highlights that this source zone will deplete very slowly, likely maintaining a similar effluent flux beyond this century. The modelled abstraction shows, in general, a slow breakthrough which agrees well with the observed concentrations before 2007 (figure 5.16e). The scale of concentration response in the period 2007–2015 is well matched (figure 5.16f), though with monthly stress periods the shape of the modelled response to pump tests, which are up to two weeks in duration, is poorly resolved. The projected future concentration, using the 6 Ml d^{-1} abstraction scenario, shows a long-term peak concentration of about 10 µg l^{-1} in 2030, followed by a slow decrease.

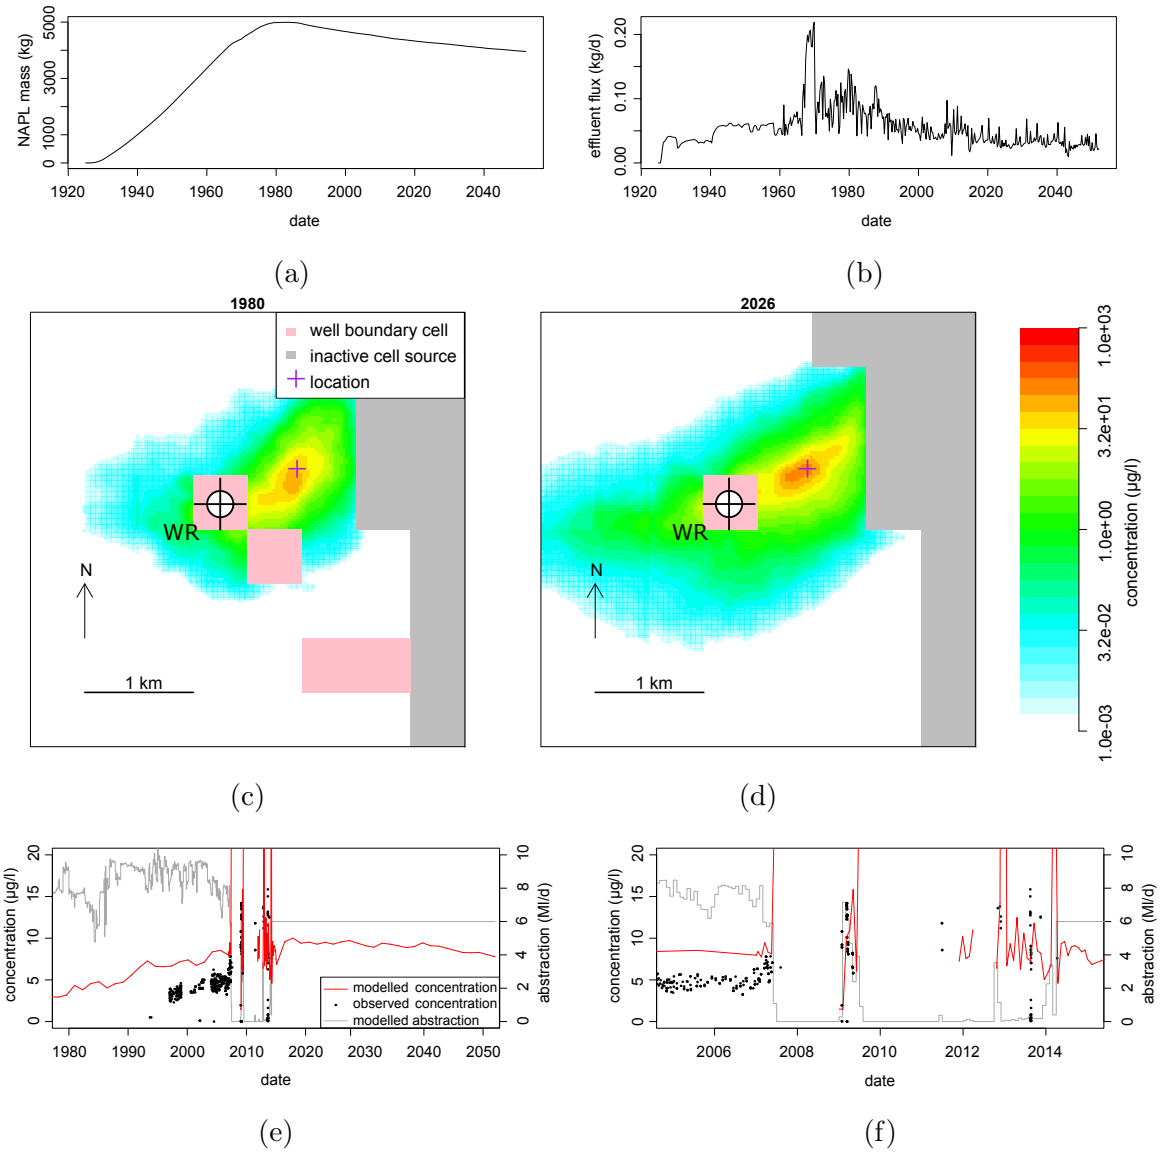

Figure 5.16: Likely source term and contaminant plume impacting Winnie Road. DNAPL source term mass and flux are shown in (a) and (b) respectively. Dissolved phase plume concentrations in the Bridgnorth Formation are colour-flooded in (d) and (e) (note log colour scale). The source term does not penetrate below the Kidderminster Formation, but dissolved contaminant is drawn into the Bridgnorth Formation by flow with a vertical component. Modelled and observed abstracted concentrations at Winnie Road are shown in (e) and (f).

5.6 Conclusions

This study has demonstrated the combined use of the DNAPL source term model framework (chapter 3) and the DRW aqueous transport solution (chapter 2) to produce an internally consistent transient regional risk assessment for a sandstone aquifer under a historic industrial area. This combined methodology uniquely couples a quantitative DNAPL-appropriate source term representation with a transport model for a collection of generic, poorly constrained potential solvent spill sites. Many inputs to the methodology require quantitative treatment of processes that are only poorly constrained, especially relating to historical chemical usage and spillage and to upscaled DNAPL behaviour. But by presenting a consistent framework for modelling and coupling these processes, the testing of model assumptions and parameter choices becomes much more streamlined and transparent. The coupled framework gives a much more efficient tool for trialling ranges of parameter values covering the likely range of values, as this work demonstrates.

For the future operation of Winnie Road, the transport modelling suggests that the most likely plume to be impacting the well has its source under the historic industrial cluster located about a kilometre to the north-east (the highlighted region in figure 5.8). This plume is predicted to travel directly towards Winnie Road even when Winnie Road is unpumped, a hypothesis corroborated by the observed transient concentration response (figure 5.3). Another plume from the historic industrial cluster to the east of Winnie Road could potentially impact Winnie Road. Such a plume would be drawn into the well with stronger abstraction rates. However, given current transient concentration behaviour and the match to data achieved with a single source to the north-east, it is likely that a single source to the north-east is impacting Winnie Road, that with a moderate 6 Ml d^{-1} abstraction rate concentration will peak around 10 µg l^{-1} (figure 5.16) and that slightly cleaner water can be achieved with higher abstraction rates (figure 5.9a). Although the possibility of a larger plume from the east or south-east should be considered, there is no direct evidence for this and so treatment works that can treat high volumes of water with up to 15 µg l^{-1} will likely be adequate for the long-term future[5]. In the past, Winnie Road was protected by the operation of Kimini. However, this well has been irrevocably decommissioned (Davies and Gordon, 2013) so any proposed scheme involving protection of Winnie Road with a scavenger well would require a new well. Most likely, this would be best placed to the north-east of Winnie Road. If all historic industrial sites are assumed to have solvent source zones, then future operation of Winnie Road will require moderate to low abstraction rates to avoid increasing abstracted concentrations (figure 5.14).

The results from the DNAPL source term model and the source term model coupled to the aqueous transport model demonstrate that, while DNAPL mass is likely to be significantly attenuated by the strongly layered Kidderminster Formation, somewhat

[5] Water from Winnie Road is blended in a district supply reservoir, meaning that in fact the greatest concern may be with the contaminant flux, rather than the concentration, depending on what proportion of total flow to the reservoir is supplied by Winnie Road. Water from Winnie Road PWS is also treated to remove chlorinated hydrocarbons before it enters the blending reservoir.

protecting the Bridgnorth Formation, large spills may still penetrate to the base of the aquifer. However, with downward hydraulic gradients induced by pumping from the Bridgnorth Formation, the Bridgnorth Formation may incur dissolved solvent contamination even without full penetration of DNAPL. The overestimated concentration suggests that there is probably just a few real sources, or even just one source, impacting Winnie Road. The combined source term and transport methodology gives a strong indication of the industrial cluster from which a plume is generated, but cannot be said to pinpoint a DNAPL source zone. An attempt to reproduce the observed abstracted concentrations at Winnie Road (section 5.5.2) indicates that a single source term gives results much more closely matching data (figure 5.16). With due consideration to uncertainty, the combined DNAPL source term and aqueous transport methods give a powerful tool for visualising and forecasting the future water quality at a PWS well.

The DNAPL source term to the Permo-Triassic Sandstone is likely to persist multiple decades into the future. Therefore, attempts to restore the groundwater quality of the area must address the issue of chlorinated solvent DNAPLs. However, because of the uncertain location of solvent spills, the likely depth of DNAPL in the aquifer and the fact that much of the historic industrial land in the area has changed ownership and been redeveloped, source zone remediation would be an expensive task with no guarantee of success. Therefore it is argued that future use of groundwater for PWS in the Stourbridge area would be better advised to use groundwater risk forecasting tools appropriate for the regional scale, such as the source term framework and transport solution developed for this project. This coupled source term and transport methodology enables asset managers to predict future abstracted concentrations and to experiment with different pumping scenarios. This will greatly help with assessing the future viability of assets, possible operational solutions and likely future treatment requirements.

Chapter 6

CONCLUSION

6.1 The Future of the Chlorinated Solvent Risk to Groundwater Resources in Permo-Triassic Sandstone

The dense non-aqueous phase liquid (DNAPL) source term results for the case studies indicates that dissolved effluent flux into groundwater from old solvent spills is likely to be significant for many decades into the future. Attempts at plume remediation would give only temporary benefit because persistent DNAPL sources would regenerate them. However, there is variation in the predicted persistence of dissolved solvents in groundwater. Of all the chlorinated solvents, tetrachloroethene (PCE) will probably be the last to disappear from groundwater on account of its much lower solubility (table 3.1). Its slightly higher sorption potential (section 4.1.2) will also cause greater persistence of dissolved phase contamination. 1,1,1–trichloroethane (1,1,1–TCA), on the other hand, is likely to disappear the soonest. Its peak concentration arrived later than for trichloroethene (TCE) (figures 4.4 and 4.3), but on account of likely lower volumes of the solvent used and that at a time of improved solvent disposal (figure 3.4; section 3.3), observed concentrations are predicted to decline more quickly (section 4.5.4). TCE was the solvent used in the greatest volume over the twentieth century and as a result is the most prevalent contaminant in groundwater; this is likely to remain the case in the Tame Valley in the coming decades (figure 4.10). In Stourbridge, the tetrachloromethane (TeCM) source will probably persist at least to the end of the century. Although it is likely that peak concentration at the Winnie Road public water supply (PWS) will be reached soon, future reduction in abstracted concentration may be very slow (figure 5.16).

Variation in the persistence of chlorinated solvent contamination also varies with other local effects. Source zones that were near a historic abstraction while they were forming may be much less of a problem today. DNAPL mass in early-stage source zones is generally more accessible to flow and if there is fast groundwater flow at the time of DNAPL migration then the contaminant mass may be more efficiently flushed. This phenomenon can be observed in the Tame Valley by comparing source terms of the same solvent in different locations (figure 4.17). Across Birmingham and likely other former industrial centres in the UK it was common for factories to operate their own abstraction boreholes, resulting in what has been described as inadvertent pump and treat operations (Daly, 2005; Rivett et al., 2012). In fact it is likely that these inadvertent pump and treat operations, functioning contemporaneously with solvent spillage and DNAPL source zone formation, were considerably more efficient than more recent pump and treat remediations which have had to remove a greater proportion of inaccessible solvent mass. The solvent legacy in UK aquifers could be worse than it actually is were it not for this accidental effect.

Persistence of solvent sources from today will depend strongly on how efficiently solvent mass is dissolved by passing groundwater. Mechanistically, and qualitatively, the amount of DNAPL mass in laterally extensive stagnant pools, in zones of lower hydraulic conductivity and the amount of dissolved phase mass sequestered into aquitard

layers which will diffuse out over a long period will govern this. The DNAPL source term framework recommends that these features are considered empirically with the empirical parameter Γ (section 3.3; equations 3.5 and 3.6). Lower values of Γ describe source terms with the most protracted tailing of effluent flux. In both the Tame Valley and Stourbridge regions, Γ has been difficult to estimate from currently existing data, because it is too soon after the cessation of solvent spillage. Source term predictions in the Tame Valley (figure 4.16) and Stourbridge (figure 5.11) indicate that the distinction between source terms characterised by different Γ values will emerge in the future. Therefore, Γ will prove significant for future management of groundwater quality. Future studies would benefit greatly from more recent data as DNAPL source zones transition from early stage (during spillage and shortly after) to late stage (the tailing period) characteristics.

A key focus of this project was attempting to produce transferable methods and results. Due to the empirical nature of the equations and conceptual models used in the DNAPL source term framework, transferability is a challenge. Transferability of methods is discussed in section 6.2 and this paragraph considers whether results from the Tame Valley and Stourbridge case studies could be defensibly applied to other situations. Similar case studies, especially in terms of scale and aquifer properties are more likely to be able to use the results of these studies as working hypotheses. For example, calibrated parameter values or ranges may be reasonably used as estimates in other Permo-Triassic Sandstone case studies, although parameter sensitivity testing is still recommended, but should be taken with more caution for different aquifers. Results which were seen in both the Tame Valley and Stourbridge studies are more likely to apply generally. These include the predicted persistence of chlorinated solvent sources for many decades (apart from 1,1,1–TCA), the emerging importance of the source depletion parameter Γ, the high sensitivity to peak usage rates at sites and the propensity for DNAPL spills to migrate to depth in even strongly layered aquifers.

6.2 Evaluation of the DRW and the DNAPL Source Term Framework Methods

The project aimed to develop new methods for the modelling of DNAPL source terms and aqueous contaminant transport appropriate for the regional scale. The first attempt to develop the contaminant transport model, the pathline flux transfer (PFT) method, made the assumption that transverse dispersion could reasonably be neglected in most cases (section 2.3). However, this approximation proved to have too great an impact on results and the solution was susceptible to pathline misdirection errors where there was complex interaction between capture zones (sections 2.4 and 2.6). The subsequent model attempt, the dynamic random walk (DRW) Lagrangian numerical transport method, is able to model both longitudinal and transverse dispersion both temporally and spatially as well as first-order reactive processes. By decoupling calculations from the flow model grid, it avoids the numerical dispersion and instability

incurred by Eulerian or part-Eulerian methods (those within MT3D) on coarse grids. Benchmarking DRW against the analytical Domenico (1987) solution and MT3D's Method of Characteristics (MOC) solution (Zheng, 1990) shows that it produces accurate results, correctly responding to variation in porosity, dispersivity and reaction parameters (section 2.5.3). Applied to regional case study, DRW is able to reproduce advective-dispersive transport from multiple sources to multiple receptors, comparing favourably with PFT, whose neglect of transverse dispersion results in at-risk receptors not being noticed and MT3D methods, whose numerical dispersion results in no distinction between receptors and whose results show artificial oscillations. The DRW method has greatly facilitated the two subsequent regional-scale case studies (chapters 4 and 5) by allowing coarse-gridded groundwater flow models to be repurposed for contaminant transport simulation.

Testing of the DNAPL source term framework was more complicated than for the DRW transport model. Whereas DRW is a numerical solution to an equation, the source term framework is, effectively, the solution to a conceptual model for which the equations (of which there may be multiple) are input by the user as one of the key uncertain components. The DNAPL source term framework was therefore tested by showing that it could be calibrated to observed data. Using effluent flux data from the Tucson International Airport Area (TIAA) site, as published by Brusseau et al. (2013), it was possible to calibrate a source term model based on the framework that responded correctly to transient groundwater flow and produced appropriate source depletion characteristics. The subsequent case studies in this project (chapters 4 and 5), also served to substantiate the DNAPL source term framework, demonstrating its ability to simulate regionally consistent source terms that corroborate observations.

It was important that the methods developed were simple to use to enable them to be transparent and accessible to decision makers who are not necessarily modelling experts. Both DRW and the DNAPL source term framework are programmed in the `R` functional language (R Core Team, 2014), which, together with all the external packages used, is open source. DRW also uses MODPATH 5 (Pollock, 1994), also available with a free license; the MODPATH 5 executable is incorporated into the DRW package. The flow field used as the basis for DRW is a MODFLOW groundwater flow model (Harbaugh et al., 2000), which needs to be converted to NetCDF using the `Rflow` package (Rew et al., 2016; Barry, 2017e). Unfortunately the DRW code therefore requires groundwater flow data in three formats: the original MODFLOW binary data files, NetCDF and the composite budget file required for MODPATH 5. DRW could be streamlined by programming it to use MODPATH 6 or later, which does not require a composite budget file (Pollock, 2012), or better, by writing `R` code that can replicate MODPATH functionality using the NetCDF results. Some knowledge of `R` is required to use the models, but they are programmed so that the user interface is as basic as possible. After installing and loading the packages (appendix D.1), DRW and the DNAPL source term framework can each be run with a single function call and the results from the DNAPL source term framework can be passed directly to the DRW code as the source term without any processing by the user. However, for users who do have more confidence with `R`, it is possible to use the methods in combination

with a large selection of other `R` packages to link with other processes and models. For example, the Stourbridge case study makes use of geographic information system (GIS) packages to link a GIS data set of industrial land use with the source term framework (section 5.1.2; Bivand and Lewin-Koh, 2015; Bivand and Rundel, 2015). Additionally it is relatively simple to use iterative loops in `R`'s base functionality to program stochastic model runs. The DRW package also includes a set of tailored post-processing functions for easy result analysis. The `DRW` and `DNAPL` packages have clear documentation and instructions for use accessible with `R`'s `help` function.

The model development aimed to produce models that did not require large numbers of conceptually obscure parameters. For the DRW method, the only parameters that are exclusively solver-related are the coalescing search radius, the number of dispersion pairs per time step, the maximum number of particles and the time stepping (section 2.5.1). For the DNAPL source term framework, the only solution parameter is the time stepping. The DNAPL framework does have more real-world parameters, relating both to physical behaviour and to historical practices. The required parameters vary depending on what source zone distribution and dissolution model is used, but with the supplied models the only parameter without a clear physical meaning is the empirical source depletion parameter Γ. However it is true that some other parameters, such as those pertaining to source zone dimensions, are almost impossible to measure at the regional scale and ought to be subjected to parameter sensitivity analysis.

A limitation for the DRW program is model run time, which is generally slower than the equivalent problem using MT3D. Model run times for the benchmark tests in section 2.5.3, for example, ranged from less than 1 minute to just over 1 hour, strongly dependent on the value for coalescing search radius, which is the most direct control on particle count. On the other hand, DRW model result files are generally smaller than MT3D result files because there is no need to save information for every MODFLOW grid cell. The DNAPL source term framework is much faster and generally does not exceed 2 seconds for a single source location.

Method transferability was a key goal of the project. Some general statements about present and future groundwater contamination by solvents can be made from the case studies, but this is limited by the empirical nature of particularly the source term model (section 6.1). However, the methods have not been developed specifically for the case studies and there is nothing to prevent their application to other aquifers experiencing chlorinated solvent contamination. It was recognised early that there is no universally applicable upscaled model for the DNAPL source term (section 3.2), so the method was programmed to allow different equations to be used. During this project the source term framework has been successfully applied to a sand and gravel aquifer (section 3.5) and to Permo-Triassic Sandstone aquifers. Different aquifers may require different conceptual models — for example a chalk case study may require a case study appropriate to DNAPL in fractures (Kueper et al., 2003) — and the DNAPL source term framework is able to accommodate this.

6.3 Review of Project Results in Light of the Aims

The aims of this project were stated in full in section 1.4.2 and this section considers whether these aims were achieved. Models for aqueous contaminant transport and the DNAPL source term were developed and their application, both independently and linked, was demonstrated on regional Permo-Triassic case studies.

6.3.1 Aim 1: developing a method for modelling regional-scale contaminant transport

The DRW method for advective, dispersive and reactive aqueous contaminant transport developed for this project has been successfully tested against analytical and numerical (MT3D) solutions. It is able to accurately reproduce all these processes (section 2.5.3). When applied to the regional scale using coarse-gridded flow models, it compares favourably with other methods. Without neglecting any component of physical dispersion, it does not incur numerical dispersion, artificial oscillations or instability (section 2.6). It is specifically designed for transient and multi-layered source terms such as dissolving DNAPL source zones.

The DRW method, however, is unable to calculate reactive processes apart from first-order decay or first-order equilibrium sorption and is not able to simulate multi-species transport. Non-equilibrium sorption and multi-species transport could be developed for later versions, but the DRW method, which tracks mass rather than concentration, likely precludes non-linear reaction processes. Additionally, it is questionable whether the hydrodynamic model of dispersion used in DRW is the most appropriate (section 2.1.4). And although the DRW method avoids many of the problems associated with transport modelling with coarse-gridded flow models, the imprecise locations of receptors, especially wells, and the poorly resolved flow field will always compromise accuracy unless grid refinement is conducted.

Therefore, the first project aim has been achieved with an effective method specially designed for the context of point sources at the regional scale. Further development of transport process solutions would benefit some contexts and grid refinement of flow models may still be necessary depending on the level of precision sought from results.

6.3.2 Aim 2: developing a quantitative representation of the chlorinated solvent DNAPL source term

The DNAPL source term framework developed for this project includes full accounting of contaminant mass from the point at which it was brought onto site for use to its dissolution into flowing groundwater. Estimated historic usage and safe disposal rates for the UK have been incorporated into the framework in a way which is simple to use. An extensive literature survey was carried out to determine the most appropriate

representation of the transient DNAPL dissolution source term in sandstone aquifers (sections 3.1 and 3.2). It was determined that upscaled models of DNAPL dissolution are very context-specific and that the DNAPL source term framework should allow flexibility in which upscaled model was used in order allow this level of conceptual uncertainty to be tested. The framework includes conceptual models of DNAPL distribution and dissolution that are based on literature-derived conceptual models with ganglia and pool domains, for example, but different conceptual models can be programmed. Additionally, a successful calibration of the source term framework method was demonstrated using data from a published pump and treat remediation operation (section 3.5). The framework interfaces with the DRW code as its source term very simply, with no processing required by the user.

One aim for the DNAPL method was that it should not require many parameters that are difficult to estimate or measure. There are, however, some parameters, particularly relating to the source zone dimensions (ganglia width for example) and the empirical source depletion parameter Γ which are difficult to estimate and impossible to measure without intense investigation. This project has attempted to present guidance from literature about sensible ranges for these parameters, but case studies will probably always require a stochastic parameter sensitivity analysis using observed concentrations for comparison to find reasonable values for these parameters, as conducted in chapters 4 and 5. Because of the empirical nature of the upscaled models proposed from the literature, transfer of conclusions and methods must be done with caution. Most intensive studies on DNAPL source zones in published literature have come from the USA. In UK contexts, differences in geology and industrial history could mean literature results from the USA are less applicable.

A quantitative framework for the calculation of regionally consistent chlorinated solvent DNAPL source zones, accounting for historic solvent usage and handling patterns, has been developed successfully. There are many remaining uncertainties about the generic distribution and dissolution behaviour of DNAPLs in the subsurface and so the framework has been developed with conceptual model flexibility. The framework is based on an extensive literature survey, but transferability of many results and empirical models is questionable. There are some required parameters which are not easily determined and thorough calibration will be necessary before the method can be responsibly used for predictions and in some cases parameter sensitivities, especially those pertaining to source depletion, will require future data to constrain (section 6.1).

6.3.3 Aim 3: substantiating the aqueous contaminant transport and DNAPL source term methods

The Tame Valley and Stourbridge case studies were used to substantiate the source term and transport methods. In both cases, for which there were regional groundwater models already developed, the methods could be applied. The models were able to link to the groundwater models and to each other with a simple interface. Model results were able to give a better indication about probable source locations in both

studies. Stochastic parameter sensitivity analysis was performed to refine the regional source term representation and to constrain likely future concentrations in the aquifer and abstraction wells. The results were able to reproduce observed concentrations in most cases, although the stress period length in both groundwater models was too long to precisely capture short time-scale transient abstraction responses. As well as parameter testing, conceptual models were tested. Whether or not the aquifer base should be assumed impermeable and whether the value for Γ should converge to 1 was tested in the DNAPL source term models, demonstrating the flexibility of this method. Adaptations to the flow model boundary conditions and the source locations were also tested. Results from the case studies were used to make predictions and recommendations about future management, summarised in sections 4.6, 5.6 and 6.1.

The calibration in both case studies was limited by the amount of data available. Although the density of abstracted concentration data in the Tame Valley was relatively good, both temporally and spatially, a diminishing number of active wells over time led to a lack of data constraint in more recent times. A long period of abstracted concentration data from Winnie Road available for Stourbridge, but concentration data at other locations was not available. Multi-species capability would have been useful for the Tame Valley case study so that the cis–1,2–dichloroethene (cis–DCE) concentration data could have been used for additional constraint and because TCE in one location was thought to be a degradation product of PCE. Additionally, transfer of conclusions from the case studies to other contexts must be done with caution, as discussed in section 6.1.

Overall, the case studies demonstrate the utility of the combined DNAPL source term framework and DRW methods, with a regional model of groundwater flow, to study the transient chlorinated solvent risk over a large area. The results provided further constraint on source locations, allowed predictions of likely future evolution of concentrations and were used to make recommendations for future operation and development of groundwater assets.

6.4 Application of Project Results

The intended application of the methods developed for this project is for forecasting the transient evolution of regional-scale risk to groundwater resources from chlorinated solvent spills. This application is demonstrated in chapters 4 and 5. The methods and results will be useful for water resource managers and environmental regulators for producing informed asset management plans or prioritising groundwater bodies that are of particular concern. However, it is worth considering whether the concepts and methods developed could have other applications.

The DRW method, based on the concept of dynamic particle allocation by continuous splitting, displacing and coalescing is designed for regional scale contaminant transport modelling from point sources. It was designed for use with a DNAPL source term but the code is not specific to such a source description. The method could be applied

to other point source problems such as light non-aqueous phase liquid (LNAPL) fuel spills or landfill leaks. Additionally, modification of the DRW code to allow it to track back in time could allow it to model dispersed capture zones. This paradigm is demonstrated by Frind et al. (2006) and Black and Foley (2013) using an Eulerian method and modifying DRW to include this capability would allow this approach to be used at the regional scale without incurring numerical artifacts commonly encountered with coarse grids. Beyond groundwater contaminant transport, the DRW concept may also be applicable to the migration of other entities in different media.

MassTrack, a key component of the DRW method (sections 2.2 and 2.5) removes the problem of how to deal with weak sinks for particle tracking. Outside of DRW, MassTrack could be used to enhance particle tracking experiments, for example by allowing the estimation of contaminant flux to streams that are represented as boundary conditions in the groundwater flow model.

The DNAPL source term framework is designed specifically for DNAPLs. It is a transferable method, not assuming a particular conceptual model, so that it can be applied to different contexts, because different geological settings require very different DNAPL conceptual models (Kueper et al., 2003). A similar paradigm is in place with MODFLOW, whose modular structure allows modellers to develop new modules to describe different types of boundary condition if none of the provided modules meet the needs of their problem. In a similar way incorporation of different types of DNAPL architectures and upscaled models is possible for a user with a good knowledge of `R` without needing to rewrite the `DNAPL` package code.

6.5 Recommendations for Further Research

The project has largely fulfilled its aims, with some limitations (section 6.3). Some of these limitations show obvious areas where future research would be beneficial, but the project has also opened new questions which would be both interesting and useful to address.

The DRW method has proved robust when benchmarked against other solutions, but the model is not able to simulate all processes. Both the physical and empirical basis for the popular hydrodynamic model of dispersion used in the DRW code have been questioned (section 2.1.4). Future developments should consider incorporating a dual-domain mass-transfer model of dispersion (Bianchi and Zheng, 2016). Sorption is included only as a linear isotherm equilibrium sorption (section 2.5.1), but first-order rate-limited sorption could also be calculated without extensive modifications to the code. Finally, multi-species reactive transport would also be compatible with the DRW algorithm, provided reactions are first-order. This would be particularly useful as it would enable concentrations of degradation products, for which there is often data, to be modelled, providing further parameter constraint.

The run time of the DRW code is quite slow and the number of stochastic runs that could be executed in the Stourbridge case study was limited by this (section 5.5). This

problem was somewhat alleviated later in the project by the use of parallel programming (R Core Team, 2014), but performance enhancements could also help project flow and allow a greater number of stochastic runs to be used. Profiling reveals that model run time is dominated roughly equally by the advection (principally MODPATH 5) and the dispersion step, with about 35% of the execution time for each. The MassTrack and coalescing modules are less expensive, each requiring about 12%. Speculatively, a more closely integrated interface with MODPATH could achieve some speed up and translating the DRW into a compiled programming language such as `Fortran` may achieve very significant run time reductions.

The DNAPL source term framework method has been designed to be flexible, not assuming a particular conceptual model for DNAPL distribution and dissolution. Although some conceptual models are supplied for easy use, users have the option to develop new conceptual models according to clearly documented instructions. The application of the source term framework to other case studies and particularly those in different geological settings would give the opportunity to enhance the framework with more context-specific conceptual models, resulting in a truly transferable chlorinated solvent DNAPL risk tool.

The methods developed for this project were designed for use with coarse-gridded regional groundwater flow models. Although this aim has largely been successful (section 6.3), there are still limitations to studies using coarse grids. Refinement of a detailed groundwater flow model can be a difficult task to do correctly without inadvertently disturbing boundary conditions. The emergence of more advanced modelling tools such as telescopic mesh refinement (TMR) (Rumbaugh and Rumbaugh, 2011) and MODFLOW-USG (Panday et al., 2013) offers more effective, targeted and less disruptive grid refinement options. Further development of the work of this project could usefully seek to build an automated grid refinement tool making use of one of these options that was specifically designed for use with regional flow models. The greatest difficulty in this would be correctly refining the shapes and locations of boundary conditions, especially well locations, which certainly should not be assumed to be at the centres of the original MODFLOW cells.

Both case studies for this project have highlighted the importance of continuing to collect concentration data. The late-stage depletion characteristics of chlorinated solvent source zones (parameterised as Γ) are distinct from the early-stage characteristics (during and shortly after the period of solvent spillage). Therefore it is difficult to forecast source dissolution rates using current datasets. Continued data collection and parameter sensitivity analysis will allow progressively refined predictions of future chlorinated solvent risk to receptors.

6.6 Summary

Chlorinated hydrocarbon solvents, released into the environment in industrial regions during much of the twentieth century, have formed a persistent source of contamination

with a particular impact on groundwater resources in the UK among other countries. Water supply assets exploiting groundwater from the Permo-Triassic Sandstones have come under pressure from this threat to water quality. The issue cannot be entirely addressed at contaminated land sites and regional-scale management tools are required to predict future risk and plan effective treatment operations and abstraction schemes. Such tools developed to date have partially succeeded in fulfilling this requirement but were not able to incorporate transient groundwater flows and did not adequately represent the chlorinated solvent DNAPL source term (chapter 1).

This project was initiated to develop regional-scale risk forecasting methods, building on previous work and addressing these shortcomings. The DRW method for advective, dispersive and reactive aqueous contaminant transport is fully transient and has been robustly tested. By using a Lagrangian solution algorithm, artifacts often associated with numerical solutions on coarse grids are avoided (chapter 2). A review of published research into DNAPL behaviour in aquifers has informed a selection of upscaled transient source term models which have been incorporated into a source term framework. The source term framework also uses estimates of historic solvent usage rates and handling practices. The results from the framework can be calibrated to observed transient effluent flux (chapter 3).

Regional case studies set in Permo-Triassic Sandstone aquifers were selected based on asset management requirements, data availability and existing regional groundwater flow models (chapters 4 and 5). The Tame Valley region in Birmingham is a historically heavily industrialised area for which there have been a number of organic groundwater quality surveys. Data from these surveys presented an unusually comprehensive spatial and temporal picture of the changing groundwater quality in the area. This gave a good opportunity to test the capability and accuracy of the DNAPL source term framework, linked to the DRW solution, in calculating the transient chlorinated solvent risk to an area using regionally consistent calibrated parameters. The study was able to find a reasonable calibration for most cases, whilst maintaining parameter consistency between different spill sites. Results indicated that poor groundwater quality in the region is likely to persist for some decades, preventing potable use, but that it is unlikely to cause significant contamination of other water bodies. The Stourbridge study demonstrates how the methods can be applied for the water quality forecasting and management of a PWS asset. By trialling different groundwater flow regimes and future abstraction scenarios and analysing observed concentration response to abstraction, source and plume characteristics could be deduced. From this, it was possible to estimate what would be the best option for continued use of the asset, in this case to abstract strongly and plan treatment works able to treat water with up to 15 µg l^{-1}.

The DNAPL source term framework and DRW aqueous contaminant transport methods, individually and combined, are transferable methods which, with awareness of limitations and parameter sensitivity analysis, allow users to gain a quantitative understanding of the regional chlorinated solvent risk as it varies with depth and time. Chlorinated solvent sources are predicted to persist well into the current century and

other water resource pressures are likely to increase. Application of regional risk management tools such as these have an important part to play in securing future water supply security.

REFERENCES

Akladiss, N., Asreen, R., Cohen, A., Hadley, P. W., MacDonald, A., Naugle, A., Waldron, J., Scheer, D., Smith, M. B., Syverson, L., Black, I., Nuttall, H. E., Sale, T. C., Pennell, K. D., Fiedler, L., Holmes, F., Lebrón, C., Osgerby, I., Ruiz, N., Stroo, H., 2011. Technical / Regulatory Guidance Integrated DNAPL Site Strategy. Tech. Rep. November, Interstate Technology & Regulatroy Council.

Alexandra, R., Gerhard, J. I., Kueper, B. H., 2012. Hydraulic displacement of dense nonaqueous phase liquids for source zone stabilization. Ground water 50 (5), 765–74.
URL http://www.ncbi.nlm.nih.gov/pubmed/22276594

Allen, D. J., Bloomfield, J. P., Gibbs, B. R., Wagstaff, S. J., Barker, J. A., Robinson, N., 1998. Fracturing and the hydrogeology of the Permo-Triassic sandstones in England and Wales. Tech. rep., British Geological Survey; Environment Agency.

Allen, D. J., Brewerton, L. J., Coleby, L. M., Gibbs, B. R., Lewis, M. A., MacDonald, A. M., Wagstaff, S. J., Williams, A. T., 1997. The physical properties of major aquifers in England and Wales. Tech. rep., Environment Agency; British Geological Survey.

Aller, L., Bennett, T., Lehr, J. H., Petty, R. J., Hackett, G., 1987. DRASTIC: a standardized system for evaluating groundwater pollution potential using hydrgeologic settings. Tech. rep., Robert S. Kerr Environmental Research Laboratory, Ada, Oklahoma.

Anderson, M., Gamble, D., 2009. Severn Trent Water AMP5 Water Quality Improvement Scheme - Winnie Road Critical Review October 2009 Document Control Sheet. Tech. Rep. October, MWH.

Anderson, M. R., Johnson, R. L., Pankow, J. F., mar 1992. Dissolution of Dense Chlorinated Solvents into Ground Water: 1. Dissolution from a Well-Defined Residual Source. Ground water 30 (2), 250—-256.
URL http://doi.wiley.com/10.1111/j.1745-6584.1992.tb01797.x

Arnold, J., Coggins, J., Palin, M., 1996. Profile of miscellaneous industries. Tech. rep., Department of the Environment; Arup Environmental; Photo Marketing Association International.

Bakker, M., Strack, O. D. L., may 1996. Capture Zone Delineation in Two-Dimensional Groundwater Flow Models. Water Resources Research 32 (5), 1309–1315.
URL http://doi.wiley.com/10.1029/96WR00198

Barry, C. J., 2017a. coalesce: Coalesce a field of weighted points. R Package.

Barry, C. J., 2017b. DNAPL: DNAPL Source Term Framework. R Package.

Barry, C. J., 2017c. DRW: Dynamic Random Walk. R Package.

Barry, C. J., 2017d. MassTrack: Track Particle Mass along Calculated Pathlines in Groundwater. R Package.

Barry, C. J., 2017e. Rflow: Tools for analysing and writing input and output files for USGS MODFLOW. R Package.
URL https://www.github.com/CJBarry/Rflow

Barry, D. A., Prommer, H., Miller, C. T., Engesgaard, P., Brun, A., Zheng, C., 2002. Modelling the fate of oxidisable organic contaminants in groundwater. Advances in Water Resources 25 (8-12), 945–983.

Basu, N. B., Fure, A. D., Jawitz, J. W., jul 2008. Predicting dense nonaqueous phase liquid dissolution using a simplified source depletion model parameterized with partitioning tracers. Water Resources Research 44 (7), W07414.
URL http://doi.wiley.com/10.1029/2007WR006008

Bear, J. J., 1979. Hydraulics of Groundwater. McGraw-Hill, New York, NY.

BGS, 2017. Borehole Scans.
URL http://bgs.ac.uk/data/boreholescans/home.html

Bianchi, M., Zheng, C., 2016. A lithofacies approach for modeling non-Fickian solute transport in a heterogeneous alluvial aquifer. Water Resources Research 52 (1), 552–565.

Bianchi, M., Zheng, C., Tick, G. R., Gorelick, S. M., 2011. Investigation of small-scale preferential flow with a forced-gradient tracer test. Ground Water 49 (4), 503–514.

Bivand, R., Lewin-Koh, N., 2015. maptools: Tools for Reading and Handling Spatial Objects. R package.
URL https://cran.r-project.org/package=maptools

Bivand, R., Rundel, C., 2015. rgeos: Interface to Geometry Engine - Open Source (GEOS). R Package.
URL https://cran.r-project.org/package=rgeos

Black, A., Foley, C., 2013. FlowSource: A program to efficiently delineate volumetric capture areas, pathways and source areas in groundwater models. In: Translating Science into Practice.

Blandford, T. N., Huyakorn, P. S., 1991. A Modular Semi-Analytical Model for the Delineation of Wellhead Protection Areas. Tech. Rep. 68, HydroGeoLogic, Herndon, VA.
URL https://www.epa.gov/water-research/

Blum, P., Annable, M. D., nov 2008. Partial source zone removal. Journal of Contaminant Hydrology 102 (1-2), 1–2.
URL http://www.ncbi.nlm.nih.gov/pubmed/19000587

Botha, S., 2006. An Assessment of the Magnitude and Temporal Variability of Volatile Organic Compounds (VOCs) in the Birmingham Aquifer. Msc, University of Birmingham.

Bredehoeft, J. D., Pinder, G. F., 1973. Mass transport in flowing groundwater. Water Resources Research 9 (1), 194–210.

Brewster, M. L., Annan, A. P., Greenhouse, J. P., Kueper, B. H., Olhoeft, G. R., Redman, J. D., Sander, K. A., 1995. Observed Migration of a Controlled DNAPL Release by Geophysical Methods. Ground water 33 (6), 977—-987.

British Geological Survey, oct 2013. 1:50 000 Geology [SHAPE geospatial data], Scale 1:50000, Tiles: ew139,ew140,ew153,ew154,ew167,ew168,ew182,ew183.
URL http://digimap.edina.ac.uk

British Geological Survey, 2016. The BGS Lexicon of Named Rock Units: Kidderminster Formation.

Brodie, A., 2012. Groundwater Model of the Leek Permo-Triassic Sandstone Aquifer Scoping Study. Tech. Rep. December, Environment Agency.

Broholm, K., Feenstra, S., Cherry, J. A., 1999. Solvent Release into a Sandy Aquifer. 1. Overview of Source Distribution and Dissolution Behavior. Environmental Science & Technology 33 (5), 681—-690.
URL http://pubs.acs.org/doi/10.1021/es980097d

Brooks, R. H., Corey, A. T., 1964. Hydraulic Properties of Porous Media. Hydrological Paper 3. Tech. rep., Civil Engineering Department, Colorado State University, Fort Collins, Colorado.

Brusseau, M. L., Guo, Z., may 2014. Assessing Contaminant-Removal Conditions and Plume Persistence through Analysis of Data from Long-term Pump-and-Treat Operations. Journal of Contaminant Hydrology 164, 16–24.
URL http://linkinghub.elsevier.com/retrieve/pii/S0169772214000643

Brusseau, M. L., Matthieu, D. E., Carroll, K. C., Mainhagu, J., Morrison, C., McMillan, A., Russo, A., Plaschke, M., jun 2013. Characterizing long-term contaminant mass discharge and the relationship between reductions in discharge and reductions in mass for DNAPL source areas. Journal of Contaminant Hydrology 149, 1–12.
URL http://www.ncbi.nlm.nih.gov/pubmed/23528743

Brusseau, M. L., Nelson, N. T., Zhang, Z., Blue, J. E., Rohrer, J., Allen, T., feb 2007. Source-zone characterization of a chlorinated-solvent contaminated Superfund site in Tucson, AZ. Journal of Contaminant Hydrology 90 (1-2), 21–40.
URL http://www.ncbi.nlm.nih.gov/pubmed/17049404

Burston, M. W., Nazari, M. M., Bishop, P. K., Lerner, D. N., aug 1993. Pollution of groundwater in the Coventry region (UK) by chlorinated hydrocarbon solvents. Journal of Hydrology 149 (1-4), 137–161.
URL http://linkinghub.elsevier.com/retrieve/pii/002216949390104H

Carey, M. A., Marsland, P. A., Smith, J. W. N., 2006. Remedial Target Methodology; Hydrogeological Risk Assessment for Land Contamination.
URL http://scholar.google.com/scholar?hl=en{&}btnG=Search{&}q=intitle:Remedial+Targets+Methodology+Hydrogeological+Risk+Assessment+for+Land+Contamination{#}0

Chen, X., Jawitz, J. W., dec 2009. Convergence of DNAPL source strength functions with site age. Environmental Science & Technology 43 (24), 9374–9.
URL http://www.ncbi.nlm.nih.gov/pubmed/19928794

Chisala, B. N., Tait, N. G., Lerner, D. N., apr 2007. Evaluating the risks of methyl tertiary butyl ether (MTBE) pollution of urban groundwater. Journal of Contaminant Hydrology 91 (1-2), 128–45.
URL http://www.ncbi.nlm.nih.gov/pubmed/17141916

Christ, J. A., Lemke, L. D., Abriola, L. M., jan 2005. Comparison of two-dimensional and three-dimensional simulations of dense nonaqueous phase liquids (DNAPLs): Migration and entrapment in a nonuniform permeability field. Water Resources Research 41 (1).
URL http://doi.wiley.com/10.1029/2004WR003239

Christ, J. A., Pennell, K. D., Abriola, L. M., jan 2012. Quantification of experimental subsurface fluid saturations from high-resolution source zone images. Water Resources Research 48 (1), W01517.
URL http://www.agu.org/pubs/crossref/2012/2011WR010400.shtml

Christ, J. A., Ramsburg, C. A., Pennell, K. D., Abriola, L. M., nov 2006. Estimating mass discharge from dense nonaqueous phase liquid source zones using upscaled mass transfer coefficients: An evaluation using multiphase numerical simulations. Water Resources Research 42 (11), W11420.
URL http://doi.wiley.com/10.1029/2006WR004886

Christ, J. A., Ramsburg, C. A., Pennell, K. D., Abriola, L. M., may 2010. Predicting DNAPL mass discharge from pool-dominated source zones. Journal of Contaminant Hydrology 114 (1-4), 18–34.
URL http://www.ncbi.nlm.nih.gov/pubmed/20227132

Chrysikopoulos, C. V., Lee, K. Y., Harmon, T. C., 2000. Dissolution of a well-defined trichloroethylene pool in saturated porous media: Experimental design and aquifer characterization. Water Resources Research 36, 1687—-1696.

Cook, L. M., apr 1990. Chemical processes in glass polishing. Journal of Non-Crystalline Solids 120 (1-3), 152–171.
URL http://www.sciencedirect.com/science/article/pii/0022309390902006

COPA, 1974. Control of Pollution Act.

Cuthbert, M. O., Faherty, J., Soley, R. W. N., Sutton, S., 2001. West Midlands Permo-Triassic Sandstones Water Resources Study Final. Tech. Rep. February, Entec UK.

Daily, P. J. J., Buss, S. R., 2013. Birmingham Permo-Triassic Sandstone Aquifer Investigation. Tech. rep., ESI.

Daly, G. E., 2005. Birmingham Abstraction Decline: Effects on Contaminant Distribution. Msc, University of Birmingham.

Davies, L., Gordon, R., 2013. AMP5 Catchment Management NEP Programme: Winnie Road Solvents. Tech. rep., ESI.

Davis, A., Fennemore, G. G., Peck, C., Walker, C. R., McIlwraith, J., Thomas, S., 2003. Degradation of carbon tetrachloride in a reducing groundwater environment: Implications for natural attenuation. Applied Geochemistry 18 (4), 503–525.

Dela Barre, B. K., Harmon, T. C., Chrysikopoulos, C. V., 2002. Measuring and modeling the dissolution of nonideally shaped dense nonaqueous phase liquid pools in saturated porous media. Water Resources Research 38 (8), 1133.

Delshad, M., Pope, G. A., Sepehrnoori, K., aug 1996. A compositional simulator for modeling surfactant enhanced aquifer remediation, 1. Formulation. Journal of Contaminant Hydrology 23 (4), 303–327.
URL http://www.sciencedirect.com/science/article/pii/0169772295001069

Department of the Environment, 1995. Engineering works, vehicle manufacturing works Industry Profile. Tech. rep., Department of the Environment.

DiFilippo, E. L., Brusseau, M. L., may 2008. Relationship between Mass Flux Reduction and Source-Zone Mass Removal: Analysis of Field Data. Journal of Contaminant Hydrology 98 (1-2), 22–35.
URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3029099{&}tool=pmcentrez{&}rendertype=abstract

DiFilippo, E. L., Brusseau, M. L., apr 2011. Assessment of a simple function to evaluate the relationship between mass flux reduction and mass removal for organic-liquid contaminated source zones. Journal of Contaminant Hydrology 123 (3-4), 104–13.
URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3061350{&}tool=pmcentrez{&}rendertype=abstract

Domenico, P., 1987. An analytical model for multidimensional transport of a decaying contaminant species. Journal of Hydrology 91 (1), 49–58.
URL http://www.sciencedirect.com/science?{_}ob=ArticleListURL{&}{_}method=list{&}{_}ArticleListID=-1181942238{&}{_}sort=r{&}{_}st=13{&}view=c{&}md5=2e6b5eb27d94b969e8ef7ec54f39c5be{&}searchtype=a

Dowle, M., Srinivasan, A., Short, T., Lianoglou, S., Saporta, R., Antonyan, E., 2015. R Package 'data.table'.
URL https://cran.r-project.org/web/packages/data.table/data.table.pdf

Duong, T., 2016. ks: Kernel Smoothing. R Package.
URL http://cran.r-project.org/package=ks

Ellis, P. A., Rivett, M. O., apr 2007. Assessing the impact of VOC-contaminated groundwater on surface water at the city scale. Journal of Contaminant Hydrology 91 (1-2), 107–27.
URL http://www.ncbi.nlm.nih.gov/pubmed/17182150

Environment Agency, 2003. ConSim release 2. Contamination impacts on groundwater: simulation by Monte Carlo method. R&D Publication 132. Tech. rep., Golder Associates (UK) Ltd, Nottingham.

Environment Agency, 2015. Groundwater Source Protection Zones.
URL http://apps.environment-agency.gov.uk/wiyby/37833.aspx

EPA, 1990. Environmental Protection Act.
URL http://www.e-laws.gov.on.ca/html/regs/english/elaws{_}regs{_}950214{_}e.htm{#}BK50

Falta, R. W., Basu, N. B., Rao, P. S. C., 2005a. Assessing impacts of partial mass depletion in DNAPL source zones: II. Coupling source strength functions to plume evolution. Journal of Contaminant Hydrology 79, 45–66.

Falta, R. W., Rao, P. S. C., Basu, N. B., aug 2005b. Assessing the impacts of partial mass depletion in DNAPL source zones I. Analytical modeling of source strength functions and plume response. Journal of Contaminant Hydrology 78 (4), 259–80.
URL http://www.ncbi.nlm.nih.gov/pubmed/16019108

Fernàndez-Garcia, D., Sanchez-Vila, X., mar 2011. Optimal reconstruction of concentrations, gradients and reaction rates from particle distributions. Journal of contaminant hydrology 120-121, 99–114.
URL http://www.sciencedirect.com/science/article/pii/S0169772210000483

Fetter, C. W., 1999. Contaminant Hydrogeology. Prentice Hall.
URL https://books.google.co.uk/books?id=5B1SAAAAMAAJ

Fielding, M., Gibson, T., James, H., dec 1981. Levels of trichloroethylene, tetrachloroethylene and pdichlorobenzene in ground waters. Environmental Technology Letters 2 (12), 545–550.
URL http://www.tandfonline.com/doi/abs/10.1080/09593338009384086{#}.WG-S5UpAsuw.mendeley

Folkard, G. K., 1986. The significance, occurrence and removal of volatile chlorinated-hydrocarbon solvents in groundwaters. Water Pollution Control 85 (1), 63—-70.

Ford, M., Tellam, J. H., 1994. Source, type and extent of inorganic contamination within the Birmingham urban aquifer system, UK. Journal of Hydrology 156, 101–135.

Franklin, D., Hawke, N., Lowe, M., 1995. Pollution in the U.K. Sweet & Maxwell, London.

Frind, E. O., Molson, J. W., Rudolph, D. L., 2006. Well vulnerability: a quantitative approach for source water protection. Ground Water 44 (5), 732–42.
URL http://www.ncbi.nlm.nih.gov/pubmed/16961495

Frind, E. O., Molson, J. W., Schirmcr, M., Guiguer, N., 1999. Dissolution and mass transfer of multiple organics under field conditions: The Borden emplaced source. Water Resources Research 35 (3), 683–694.

Gelhar, L. W., Collins, M. A., 1971. General Analysis of Longitudinal Dispersion in Nonuniform Flow. Water Resources Research 7 (6), 1511–1521.

Gerhard, J. I., Kueper, B. H., aug 2003a. Capillary pressure characteristics necessary for simulating DNAPL infiltration, redistribution, and immobilization in saturated porous media. Water Resources Research 39 (8), n/a–n/a.
URL http://doi.wiley.com/10.1029/2002WR001270

Gerhard, J. I., Kueper, B. H., 2003b. Relative permeability characteristics necessary for simulating DNAPL infiltration, redistribution, and immobilization in saturated porous media. Water Resources Research 39 (8), 8–1—-8–16.
URL http://dx.doi.org/10.1029/2002WR001490

Gerhard, J. I., Pang, T., Kueper, B. H., Hecox, G. R., mar 2007. Time scales of DNAPL migration in sandy aquifers examined via numerical simulation. Ground water 45 (2), 147–57.
URL http://doi.wiley.com/10.1111/j.1745-6584.1998.tb01094.xhttp://www.ncbi.nlm.nih.gov/pubmed/17335479

Gibbs, L. M., Levine, M., 1981. Love Canal: my story. State University of New York Press.

Graber, E. R., Laor, Y., Ronen, D., dec 2008. Aquifer contamination by chlorinated-VOCs: the case of an urban metropolis megasite overlying the Coastal Plain aquifer

in Israel. Hydrogeology Journal 16 (8), 1615–1623.
URL http://link.springer.com/10.1007/s10040-008-0366-2

Greswell, R. B., 1992. The Modelling of Groundwater Rise in the Birmingham Area. Msc, University of Birmingham.

Grischek, T., Nestler, W., Piechniczek, D., Fischer, T., jan 1996. Urban Groundwater In Dresden, Germany. Hydrogeology Journal 4 (1), 48–63.
URL http://link.springer.com/10.1007/s100400050088

Guarnaccia, J., Pinder, G. F., Fishman, M., 1997. NAPL Simulator. Tech. rep., USEPA, Ada, Oklahoma.

Hadley, P. W., Newell, C. J., 2014. The New Potential for Understanding Groundwater Contaminant Transport. Groundwater 52 (2), 174–186.
URL http://onlinelibrary.wiley.com/doi/10.1111/gwat.12135/abstract

Harbaugh, A. W., Banta, E. R., Hill, M. C., McDonald, M. G., 2000. MODFLOW-2000, THE U. S. GEOLOGICAL SURVEY MODULAR GROUND-WATER MODEL USER GUIDE TO MODULARIZATION CONCEPTS AND THE GROUND-WATER FLOW PROCESS. Tech. rep., U.S. Geological Survey.
URL http://www.gama-geo.hu/kb/download/ofr00-92.pdf

Harr, J., 1995. A Civil Action. Century Books, London.

Heathcote, J. A., Lewis, R. T., Soley, R. W. N., may 2004. Rainfall routing to runoff and recharge for regional groundwater resource models. Quarterly Journal of Engineering Geology and Hydrogeology 37 (2), 113–130.
URL http://qjegh.lyellcollection.org/cgi/doi/10.1144/1470-9236/03-029

Held, R. J., Illangasekare, T. H., may 1995. Fingering of Dense Nonaqueous Phase Liquids in Porous Media: 1. Experimental Investigation. Water Resources Research 31 (5), 1213–1222.
URL http://doi.wiley.com/10.1029/95WR00428

Hunt, J. R., Sitar, N., Udell, K. S., 1988. Nonaqueous phase liquid transport and cleanup. 1. Analysis and mechanisms. Water Resources Research 24, 1247—-1258.

I.C.I., 1931. Carbon Tetrachloride Summary. Tech. rep.

Imhoff, P. T., Jaffé, P. R., Pinder, G. F., 1994. An experimental study of complete dissolution of a nonaqueous phase liquid in saturated porous media. Water Resources Research 30 (2), 307.

Jackson, D., Lloyd, J. W., 1983. Groundwater chemistry of the Birmingham Triassic Sandstone aquifer and its relation to structure. Quarterly Journal of Engineering Geology1 16 (2), 135—-142.
URL http://qjegh.geoscienceworld.org/content/16/2/135

Jamin, P., Dollé, F., Chisala, B. N., Orban, P., Popescu, I.-C., Hérivaux, C., Dassargues, A., Brouyère, S., jan 2012a. A regional flux-based risk assessment approach for multiple contaminated sites on groundwater bodies. Journal of contaminant hydrology 127 (1-4), 65–75.
URL http://www.ncbi.nlm.nih.gov/pubmed/22036196

Jamin, P., Dujardin, J., Crèvecoeur, S., Hérivaux, C., Dejonghe, W., Dollé, F., Orban, P., Colombano, S., Diels, L., Thomé, J.-P., Batelaan, O., Canter, F., Dassargues, A., Brouyère, S., 2012b. Flux-based Risk Assessment of the impact of Contaminants on Water resources and Ecosystems: FRAC-WECO. Tech. rep., Science for Sustainable Development.
URL http://www.belspo.be/belspo/ssd/science/Reports/FRACWECO{_}FinRep{_}AD.pdf

Jawitz, J. W., Fure, A. D., Demmy, G. G., Berglund, S., Rao, P. S. C., oct 2005. Groundwater contaminant flux reduction resulting from nonaqueous phase liquid mass reduction. Water Resources Research 41 (10), n/a–n/a.
URL http://doi.wiley.com/10.1029/2004WR003825

Jiang, J., Wu, J., 2013. Continuous time random walk in homogeneous porous media. Journal of Contaminant Hydrology 155, 82—-86.
URL http://www.sciencedirect.com/science/article/pii/S0169772213001241

Johnson, R. L., Pankow, J. F., 1992. Diussolution of Dense Chlorinated Solvents into Groundwater. 2. Functions for Pools of Solvent. Environmental Science & Technology 26 (5), 896–901.
URL http://pubs.acs.org/doi/abs/10.1021/es00029a004?journalCode=esthag{&}quickLinkVolume=26{&}quickLinkPage=896{&}selectedTab=citation{&}volume=26

Johnston, C. D., Davis, G. B., Bastow, T. P., Woodbury, R. J., Rao, P. S. C., Annable, M. D., Rhodes, S., jun 2014. Mass Discharge Assessment at a Brominated DNAPL Site: Effects of known DNAPL source mass removal. Journal of Contaminant Hydrology.
URL http://linkinghub.elsevier.com/retrieve/pii/S0169772214000850

Konikow, L. F., 2011. The Secret to Successful Solute-Transport Modeling. Ground Water 49 (2), 144—-159.

Kueper, B. H., Frind, E. O., 1991a. Two-Phase Flow in Heterogeneous Porous Media 1. Model Development. Water Resources Research 27 (6), 1049–1057.

Kueper, B. H., Frind, E. O., 1991b. Two-phase flow in heterogeneous porous media: 2. Model application. Water Resources Research 27 (6), 1059–1070.
URL http://www.hydrosys.uni-stuttgart.de/institut/hydrosys/publikationen/paper/2011/DA{_}KrausDavid.pdf{%}5Cnhttp://onlinelibrary.wiley.com/doi/10.1029/91WR00267/full

Kueper, B. H., Redman, J. D., Starr, R. C., Reitsma, S., Mah, M., 1993. A Field Experiment to Study the Behavior of Tetrachloroethylene Below the Water Table: Spatial Distribution of Residual and Pooled DNAPL. Ground water 31 (5), 756—-766.

Kueper, B. H., Wealthall, G. P., Smith, J. W. N., Leharne, S. A., Lerner, D. N., 2003. An illustrated handbook of DNAPL transport and fate in the subsurface, 1st Edition. Environment Agency.

LaBolle, E. M., Fogg, G. E., Tompson, A. F. B., 1996. Random-walk simulation of transport in heterogeneous porous media: Local mass-conservation problem and implementation methods. Water Resources Research 32 (3), 583–593.

LaBolle, E. M., Quastel, J., Fogg, G. E., Gravner, J., mar 2000. Diffusion processes in composite porous media and their numerical integration by random walks: Generalized stochastic differential equations with discontinuous coefficients. Water Resources Research 36 (3), 651–662.
URL http://doi.wiley.com/10.1029/1999WR900224

Lawrence, A. R., Stuart, M., Cheney, C. S., Jones, N., Moss, R. H., aug 2006. Investigating the scale of structural controls on chlorinated hydrocarbon distributions in the fractured-porous unsaturated zone of a sandstone aquifer in the UK. Hydrogeology Journal 14 (8), 1470–1482.
URL http://link.springer.com/10.1007/s10040-006-0068-6

Lerner, D. N., Barrett, M. H., 1996. Urban Groundwater Issues in the United Kingdom. Hydrogeology Journal 4 (1), 80–89.

Lerner, D. N., Harris, B., dec 2009. The relationship between land use and groundwater resources and quality. Land Use Policy 26, S265–S273.
URL http://linkinghub.elsevier.com/retrieve/pii/S0264837709001306

Liu, H., Zhang, Y., Valocchi, A. J., 2015. Lattice Boltzmann simulation of immiscible fluid displacement in porous media: Homogeneous versus heterogeneous pore network. Physics of Fluids 27 (5).

Lyne, F. A., McLachlan, T., 1949. Contamination of water by trichloroethylene. The Analyst 74 (882), 513.

Matthieu, D. E., Brusseau, M. L., Johnson, G. R., Artiola, J. L., Bowden, M. L., Curry, J. E., 2013. Intercalation of trichloroethene by sediment-associated clay minerals. Chemosphere 90 (2), 459–463.
URL http://www.sciencedirect.com/science/article/pii/S0045653512010041

McMillan, L. A., Rivett, M. O., Tellam, J. H., Dumble, P., Sharp, H., may 2014. Influence of vertical flows in wells on groundwater sampling. Journal of Contaminant Hydrology.
URL http://linkinghub.elsevier.com/retrieve/pii/S0169772214000655

Mercer, J. W., Cohen, R. M., 1990. A Review of Immiscible Fluids in the Subsurface: Properties, Models, Characterization and Remediation. Journal of Contaminant Hydrology 6, 107–163.

Mersmann, O., 2015. microbenchmark: Accurate Timing Functions. R Package.
URL https://cran.r-project.org/package=microbenchmark

Michna, P., Woods, M., 2016. RNetCDF: Interface to NetCDF Datasets. R Package.
URL https://cran.r-project.org/package=RNetCDF

Murcott, P., 2008. Birmingham abstraction decline: effects of groundwater levels and VOC contaminant extraction. Msc, University of Birmingham.

National Center for Biotechnology Information, 2017a. PubChem Compound Database; CID=31373.
URL https://pubchem.ncbi.nlm.nih.gov/compound/31373

National Center for Biotechnology Information, 2017b. PubChem Compound Database; CID=5943.
URL https://pubchem.ncbi.nlm.nih.gov/compound/5943

National Center for Biotechnology Information, 2017c. PubChem Compound Database; CID=6278.
URL https://pubchem.ncbi.nlm.nih.gov/compound/6278

National Center for Biotechnology Information, 2017d. PubChem Compound Database; CID=6575.

Neupauer, R. M., Wilson, J. L., 2004. Numerical Implementation of a Backward Probabilistic Model of Ground Water Contamination. Ground Water 42 (2), 175–189.
URL http://scholar.google.com/scholar?hl=en{&}btnG=Search{&}q=intitle:Numerical+Implementation+of+a+backward+probabilistic+model+of+ground+water+contamination{#}0

Neupauer, R. M., Wilson, J. L., 2005. Backward probability model using multiple observations of contamination to identify groundwater contamination sources at the Massachusetts Military Reservation. Water Resources Research 41 (2), 1–14.

Oostrom, M., Dane, J. H., Wietsma, T. W., 2006. A Review of Multidimensional, Multifluid Intermediate-Scale Experiments: Nonaqueous Phase Liquid Dissolution and Enhanced Remediation. Vadose Zone Journal 5 (2), 570.
URL https://www.soils.org/publications/vzj/abstracts/5/2/570

Panday, S., Langevin, C. D., Niswonger, R. G., Ibaraki, M., Hughes, J. D., 2013. MODFLOW USG Version 1: An Unstructured Grid Version of MODFLOW for Simulating Groundwater Flow and Tightly Coupled Processes Using a Control Volume Finite-Difference Formulation. Tech. rep., U.S. Geological Survey, Reston, Virginia.
URL https://pubs.usgs.gov/tm/06/a45/

Pankow, J. F., Cherry, J. A., 1996. Dense Chlorinated Solvents and other DNAPLs in Groundwater, 1st Edition. Waterloo Press.

Parker, B. L., Cherry, J. A., Chapman, S. W., 2004. Field study of TCE diffusion profiles below DNAPL to assess aquitard integrity. Journal of Contaminant Hydrology 74 (1), 197–230.
URL http://www.sciencedirect.com/science/article/pii/S0169772204000415

Parker, B. L., Cherry, J. A., Chapman, S. W., Guilbeault, M. A., may 2003. Review and Analysis of Chlorinated Solvent Dense Nonaqueous Phase Liquid Distributions in Five Sandy Aquifers. Vadose Zone Journal 2 (2), 116–137.
URL http://vzj.geoscienceworld.org/content/2/2/116.full

Parker, J. C., Park, E., may 2004. Modeling field-scale dense nonaqueous phase liquid dissolution kinetics in heterogeneous aquifers. Water Resources Research 40 (5), W05109.
URL http://doi.wiley.com/10.1029/2003WR002807

Pearce, A. E., Voudrias, E. A., Whelan, M. P., 1994. Dissolution of TCE and TCA pools in saturated subsurface systems. Journal of Environmental Engineering 120, 1191—-1206.

Piwoni, M. D., Banerjee, P., apr 1989. Sorption of volatile organic solvents from aqueous solution onto subsurface solids. Journal of Contaminant Hydrology 4 (2), 163–179.
URL http://linkinghub.elsevier.com/retrieve/pii/0169772289900193

Pollock, D. W., 1994. User's Guide for MODPATH / MODPATH-PLOT, Version 3 : A particle tracking post-processing package for MODFLOW, the U . S . Geological Survey finite-difference ground-water flow model. Tech. rep., U. S. Geological Survey, Reston, Virginia.

Pollock, D. W., 2012. User Guide for MODPATH Version 6A Particle-Tracking Model for MODFLOW. In: Section A, Groundwater Book 6, Modeling Techniques. U.S. Geological Survey, p. 58 p.

Power, C., Gerhard, J. I., Karaoulis, M., Tsourlos, P., Giannopoulos, A., jul 2014. Evaluating four-dimensional time-lapse electrical resistivity tomography for monitoring DNAPL source zone remediation. Journal of Contaminant Hydrology 162-163, 27–46.
URL http://www.ncbi.nlm.nih.gov/pubmed/24854903

R Core Team, 2014. R: A Language and Environment for Statistical Computing.
URL http://www.r-project.org/

Rahbaralam, M., Fernàndez-Garcia, D., Sanchez-Vila, X., 2015. Do we really need a large number of particles to simulate bimolecular reactive transport with random walk methods? A kernel density estimation approach. Journal of Computational Physics 303.

Rew, R., Davis, G., Emmerson, S., Davies, H., Hartnett, E., Heimbigner, D., Fisher, W., 2016. NetCDF documentation.
URL http://www.unidata.ucar.edu/software/netcdf/docs/index.html

Rivett, M. O., 1989. The Organic Contamination of the Birmingham Aquifer. Phd, University of Birmingham.

Rivett, M. O., Chapman, S. W., Allen-King, R. M., Feenstra, S., Cherry, J. A., nov 2006a. Pump-and-treat remediation of chlorinated solvent contamination at a controlled field-experiment site. Environmental Science & Technology 40 (21), 6770–81.
URL http://www.ncbi.nlm.nih.gov/pubmed/17144309

Rivett, M. O., Dearden, R. A., Wealthall, G. P., 2014. Architecture, persistence and dissolution of a 20 to 45 year old trichloroethene DNAPL source zone. Journal of contaminant hydrology 170, 95–115.
URL http://www.sciencedirect.com/science/article/pii/S0169772214001594

Rivett, M. O., Feenstra, S., jan 2005. Dissolution of an emplaced source of DNAPL in a natural aquifer setting. Environmental Science & Technology 39 (2), 447–55.
URL http://www.ncbi.nlm.nih.gov/pubmed/15707043

Rivett, M. O., Feenstra, S., Clark, L., 2006b. Influence of the First Publication on Groundwater Contamination by Trichloroethene. Environmental Forensics 7 (4), 313–323.
URL http://dx.doi.org/10.1080/15275920600996180

Rivett, M. O., Lerner, D. N., Lloyd, J. W., Clark, L., feb 1990. Organic contamination of the Birmingham aquifer, U.K. Journal of Hydrology 113 (1-4), 307–323.
URL http://linkinghub.elsevier.com/retrieve/pii/002216949090181V

Rivett, M. O., Shepherd, K. A., Keeys, L. L., Brennan, A. E., 2005. Chlorinated solvents in the Birmingham aquifer, UK: 1986–2001. Quarterly Journal of Engineering Geology and Hydrogeology 38, 337–350.

Rivett, M. O., Smith, J. W. N., Buss, S. R., Morgan, P., 2007. Nitrate occurrence and attenuation in the major aquifers of England and Wales. Quarterly Journal of Engineering Geology and Hydrogeology 40 (4), 335—-352.
URL http://qjegh.geoscienceworld.org/content/40/4/335

Rivett, M. O., Turner, R. J., Glibbery Née Murcott, P., Cuthbert, M. O., oct 2012. The legacy of chlorinated solvents in the Birmingham aquifer, UK: observations spanning three decades and the challenge of future urban groundwater development. Journal of contaminant hydrology 140-141, 107–23.
URL http://www.ncbi.nlm.nih.gov/pubmed/23022878

Roberts, P. V., Goltz, M. N., Mackay, D. M., 1986. A Natural Gradient Experiment on Solute Transport in a Sand Aquifer 3. Sorption of Organic Solutes and its Influence

on Mobility. Water Resources Research 22 (13), 2059–2067.
URL http://onlinelibrary.wiley.com/doi/10.1029/WR022i013p02059/full

Robeznieks, S., 2009. Winnie Road - Water Quality Testing. Tech. Rep. July, Grontmij.

Rowe, B. L., Toccalino, P. L., Moran, M. J., Zogorski, J. S., Price, C. V., nov 2007. Occurrence and potential human-health relevance of volatile organic compounds in drinking water from domestic wells in the United States. Environmental health perspectives 115 (11), 1539–46.
URL http://www.ncbi.nlm.nih.gov/pubmed/18007981

Rumbaugh, J. O., Rumbaugh, D. B., 2011. Groundwater Vistas.

Salamon, P., Fernandez-Garcia, D., Gomez-Hernnndez, J. J., Fernàndez-Garcia, D., Gómez-Hernández, J. J., 2006. A review and numerical assessment of the random walk particle tracking method. Journal of Contaminant Hydrology 87 (3-4), 277–305.
URL http://www.sciencedirect.com/science/article/pii/S0169772206000957

Sale, T. C., McWhorter, D. B., 2001. Steady state mass transfer from single-component dense nonaqueous phase liquids in uniform flow fields. Water Resources Research 37 (2), 393–404.

Santharam, S., Davis, L., Erickson, L., jul 2014. Biodegradation of carbon tetrachloride in simulated groundwater flow channels. Environmental Progress & Sustainable Energy 33 (2), 444–453.
URL http://doi.wiley.com/10.1002/ep.11808

Scheutz, C., Durant, N. D., Hansen, M. H., Bjerg, P. L., 2011. Natural and enhanced anaerobic degradation of 1,1,1-trichloroethane and its degradation products in the subsurface A critical review. Water Research 45 (9), 2701–2723.
URL http://www.sciencedirect.com/science/article/pii/S0043135411000947

Schwarzenbach, R. P., Westall, J., 1981. Transport of Nonpolar Organic Compounds from Surface Water to Groundwater. Environmental Science & Technology 15 (8), 1360–1367.
URL http://scholar.google.com/scholar?hl=en{&}btnG=Search{&}q=intitle:Transport+of+Nonpolar+Organic+Compounds+from+Surface+Water+to+Groundwater.{#}9

Schwille, F., 1988. Dense Chlorinated Solvents in Porous and Fractured Media: Model Experiments, english la Edition. Lewis Publishers, Michigan.

Solvay, 1938. Tetrachloromethane: General techniques and commercial uses [French]. Tech. rep., Solvay.

Stewart, M., North, L., oct 2006. A borehole geophysical method for detection and quantification of dense, non-aqueous phase liquids (DNAPL) in saturated soils. Journal of Applied Geophysics 60 (2), 87–99.
URL http://www.sciencedirect.com/science/article/pii/S0926985106000048

Streetly, H. R., Rahman, M. A., 2011. West Midlands Worfe Groundwater Model 2011 update and recalibration of the Cosford and Worfield area. Tech. rep., ESI.

Suchomel, E. J., Pennell, K. D., oct 2006. Reductions in contaminant mass discharge following partial mass removal from DNAPL source zones. Environmental Science & Technology 40 (19), 6110–6.
URL http://www.ncbi.nlm.nih.gov/pubmed/17051808

Sweijen, T., Hartog, N., Marsman, A., Keijzer, T. J. S., jun 2014. The transport behaviour of elemental mercury DNAPL in saturated porous media: Analysis of field observations and two-phase flow modelling. Journal of contaminant hydrology 161, 24–34.
URL http://www.ncbi.nlm.nih.gov/pubmed/24748026

Tait, N. G., Davison, R. M., Leharne, S. A., Lerner, D. N., may 2008. Borehole Optimisation System (BOS) — A case study assessing options for abstraction of urban groundwater in Nottingham, UK. Environmental Modelling & Software 23 (5), 611–621.
URL http://linkinghub.elsevier.com/retrieve/pii/S1364815207001739

Tait, N. G., Davison, R. M., Whittaker, J. J., Leharne, S. A., Lerner, D. N., may 2004a. Borehole Optimisation System (BOS)A GIS based risk analysis tool for optimising the use of urban groundwater. Environmental Modelling & Software 19 (5), 1111–1124.
URL http://linkinghub.elsevier.com/retrieve/pii/S1364815203002676http://linkinghub.elsevier.com/retrieve/pii/S1364815207001739

Tait, N. G., Lerner, D. N., Smith, J. W. N., Leharne, S. A., feb 2004b. Prioritisation of abstraction boreholes at risk from chlorinated solvent contamination on the UK Permo-Triassic Sandstone aquifer using a GIS. The Science of the total environment 319 (1-3), 77–98.
URL http://www.ncbi.nlm.nih.gov/pubmed/14967503

Taylor, L., 1998. The Distribution of Organic Contaminants in the Birmingham Aquifer: a comparison of two surveys a decade apart. Msc, University of Birmingham.

Tellam, J. H., Barker, R. D., jan 2006. Towards prediction of saturated-zone pollutant movement in groundwaters in fractured permeable-matrix aquifers: the case of the UK Permo-Triassic sandstones. Geological Society, London, Special Publications 263 (1), 1–48.

URL http://sp.lyellcollection.org/cgi/doi/10.1144/GSL.SP.2006.263.01.01

Thomas, A., Tellam, J. H., may 2006. Modelling of recharge and pollutant fluxes to urban groundwaters. The Science of the total environment 360 (1-3), 158–79.
URL http://www.ncbi.nlm.nih.gov/pubmed/16325236

To, B. X., 2016. Correlation between English Permo-Triassic Sandstone Lithofacies and Permeability and its importance for Groundwater. Phd, University of Birmingham.

Troldborg, M., Lemming, G., Binning, P. J., Tuxen, N., Bjerg, P. L., oct 2008. Risk assessment and prioritisation of contaminated sites on the catchment scale. Journal of contaminant hydrology 101 (1-4), 14–28.
URL http://www.ncbi.nlm.nih.gov/pubmed/18768238

Tyler-Whittle, R., Brown, S., Shand, P., 2002. Baseline Report Series: 3. The Permo-Triassic Sandstones of South Staffordshire and North Worcestershire. Tech. rep., British Geological Survey; Environment Agency.

UK Groundwater Forum, 2011. UK Groundwater Forum.
URL http://www.groundwateruk.org

Walker, C., Jameson, C., 2010. Winnie Road PWS - Catchment Investigation. Tech. Rep. 0, Grontmij.

Wang, F., Annable, M. D., Jawitz, J. W., sep 2013. Field-scale prediction of enhanced DNAPL dissolution based on partitioning tracers. Journal of contaminant hydrology 152, 147–58.
URL http://www.ncbi.nlm.nih.gov/pubmed/23911784

Wealthall, G. P., Steele, A., Bloomfield, J. P., Moss, R. H., Lerner, D. N., jul 2001. Sediment filled fractures in the Permo-Triassic sandstones of the Cheshire basin: observations and implications for pollutant transport. Journal of contaminant hydrology 50 (1-2), 41–51.
URL http://www.ncbi.nlm.nih.gov/pubmed/11475160

Westrick, J. J., Mello, J. W., Thomas, R. F., 1984. The Groundwater Supply Survey. Journal of the American Water Works Association 76 (5), 52–59.
URL http://www.awwa.org/publications/journal-awwa/abstract/articleid/11091.aspx

Whelan, M. P., Voudrias, E. A., Pearce, A., 1994. DNAPL pool dissolution in saturated porous media: Procedure development and preliminary results. Journal of Contaminant Hydrology 15, 224—-237.

Yang, M., Annable, M. D., Jawitz, J. W., 2015. Back Diffusion from Thin Low Permeability Zones. Environmental Science & Technology 49 (September), 415–422.

Yang, M., Annable, M. D., Jawitz, J. W., 2016. Solute source depletion control of forward and back diffusion through low-permeability zones. Journal of Contaminant Hydrology 193 (September), 54–62.

Zhang, C., Yoon, H., Werth, C. J., Valocchi, A. J., Basu, N. B., Jawitz, J. W., nov 2008. Evaluation of simplified mass transfer models to simulate the impacts of source zone architecture on nonaqueous phase liquid dissolution in heterogeneous porous media. Journal of Contaminant Hydrology 102 (1-2), 49–60.
URL http://www.ncbi.nlm.nih.gov/pubmed/18579257

Zheng, C., 1990. MT3D. Tech. rep., S.S. Papadopulos & Associates.

Zheng, C., 1993. Extension of the Method of Characteristics for Simulation of Solute Transport in Three Dimensions. Ground water 31 (3), 456–465.

Zheng, C., Gorelick, S. M., 2003. Analysis of solute transport in flow fields influenced by preferential flowpaths at the decimeter scale. Ground Water 41 (2), 142—-155.

Zheng, C., Wang, P. P., 1999. A Modular Three - Dimensional Multispecies Transport Model. Tech. rep., The University of Alabama, Tuscaloosa.
URL http://www.geo.tu-freiberg.de/hydro/vorl{_}portal/gw-modellierung/MT3DMSRefManual.pdf

Zhu, J., Sun, D., 2016. Significance of groundwater flux on contaminant concentration and mass discharge in the nonaqueous phase liquid (NAPL) contaminated zone. Journal of Contaminant Hydrology 192, 158–164.
URL http://dx.doi.org/10.1016/j.jconhyd.2016.08.002

Zhu, J., Sykes, J. F., 2004. Simple screening models of NAPL dissolution in the subsurface. Journal of Contaminant Hydrology 72 (1), 245–258.
URL http://www.sciencedirect.com/science/article/pii/S0169772203002389

Appendix A

A COALESCING ALGORITHM FOR DYNAMIC ORGANISATION OF PARTICLES

The dispersion solution used by the dynamic random walk (DRW) method (section 2.5.2) involves splitting particles into multiple new particles. Left unchecked, this would lead to an exponential increase in particle count. Clearly such a scheme is unmanageable without another algorithm being applied which manages the number of particles. It is also clear that within a plume many of the particles generated would have very similar locations to other particles. Therefore it is reasonable to group, or coalesce, the coincident particles. In this way a roughly uniform particle density may be used to represent the entire plume in an efficient manner.

Given a random set of points, the coalescing algorithm will find, for each point, any other points that are within a given radius, which is the coalescing search radius, a user-defined parameter. However, any points that are already grouped are not then rechecked, reducing the number of calculations that must be performed and avoiding long chains of grouped points. The grouping is done in three dimensions and a different search radius may be specified vertically, defining a search ellipsoid. When a group is coalesced, the resulting master particle is assigned the total mass from all the grouped particles and positioned in the weighted average position. In this way both mass and centre of mass are conserved globally and locally on scales greater than the coalescing radius.

For a large number of particles (n), this can require vast numbers of distance calculations to be executed (up to the square of the number of particles). However, most of these calculations will be comparing particles that are clearly very distant from each other when in fact each particle will only likely have a few close neighbours. It is difficult to tell a computer what is obvious to the human eye, but significant time-saving may be had by splitting the particle field into subregions. This sub-algorithm sorts the particles by x and then splits them into four (or more generically $\sqrt{s}$ where s is the number of subregions used) roughly equally populated groups. Each x grouping is

sorted further split into four ($\sqrt{s}$) groups by y value. Thus 16 (or s) subregions are defined with determined index ranges, such that for each particle, the distance from only n/s particles must be evaluated. So the algorithm's growth rate goes from $O(n^2)$ to $O(n^2/s)$. And when the number of subregions s is set to vary with n such that there should be at least, but as close as possible to 256 particles in each subregion (see paragraph below for discussion of this value), then s tends to $n/256$, so the algorithm growth rate tends to $O(n)$, although actual run time variation will depend on the specific code's indexing operation and how efficiently it can perform calculations on arrays of numbers. Some potential coalescences may be missed where they cross subregion boundaries, but this is not a problem for the algorithm's application in DRW. Figure A.1 shows the huge run time reductions that can be achieved by using subregions.

A minimum particle count of 256 per subregion is chosen as a general optimum for large particle counts, according to performance testing (figure A.1). This value is a compromise between minimising the number of groups of calculations (subregions) and minimising the number of calculations within each group (distance calculations between particles in a subregion). `R` benefits from sophisticated vectorisation, meaning that the growth rate of simple operations such as addition is much better than $O(n)$ when performed on vectors or arrays of numbers. Therefore, although a very small value for particles per subregion would minimise the total number of calculations to be performed, it also increases the number groups of calculations such that each group of calculations is less efficient. Additionally, very small numbers of particles per subregion make the coalescing function less effective as many potential groupings across subregion boundaries are missed. Therefore 256 is the minimum number of particles per subregion used in the current version of DRW (Barry, 2017c), as it is close to the fastest for large particle counts but will not inhibit the coalescing behaviour excessively.

Finally, a minimum particle mass may optionally be provided to save run time by avoiding tracking particles of negligible mass. In this case, particles below the minimum mass are coalesced with their nearest neighbours (even if the neighbours are also below the minimum mass) iteratively until all particles have more mass than the minimum mass. This option can reduce run time by further restricting the number of active particles, but can have the undesirable effect of suppressing the fringes of the plume, thus under-estimating dispersion.

The effects of the coalescence step are demonstrated in figure A.2, using a randomly distributed particle field. As the coalescing radius is increased, then more particles are absorbed into each group resulting in a lower final particle count. On the other hand, the precision of the concentration distribution is reduced. Note that a few of the final particles are within the coalescing radius of other particles (for example around $(80, 60)$ in figure A.2c), indicating that potential groupings have been missed. This is because the algorithm labels grouped particles as already assigned to a group and therefore not available for other groups. This feature both saves time (by reducing the number of necessary distance determinations) and avoids grouping long chains of particles. These groupings would be made if the coalescing algorithm was allowed a second pass, but would not generally be worth the extra computational effort.

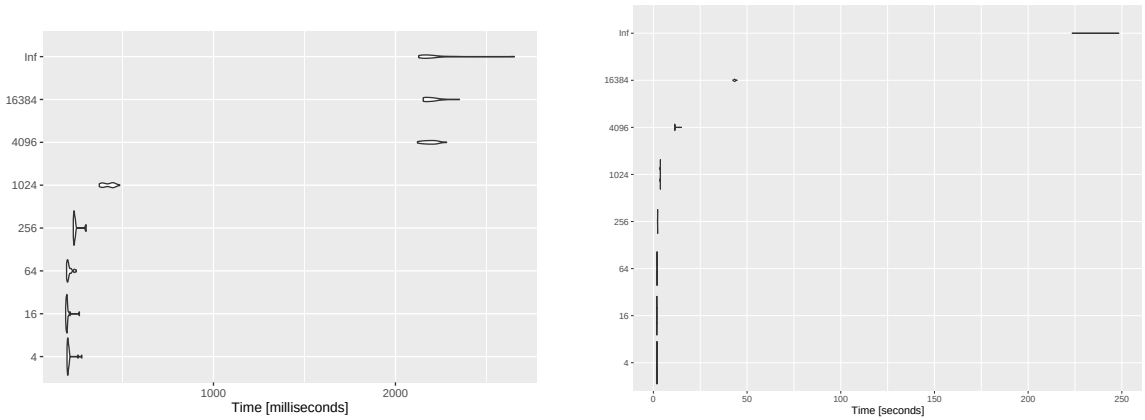

(a) $10,000$ particles (b) $100,000$ particles

Figure A.1: Performance testing of the coalescing program with different minimum numbers of particles per subregion. "Inf" implies subregions are not used and the subregions feature is automatically switched off if the number of particles per subregion would be greater than a quarter of the total starting number of particles. Testing was conducted using the `microbenchmark` package (Mersmann, 2015).

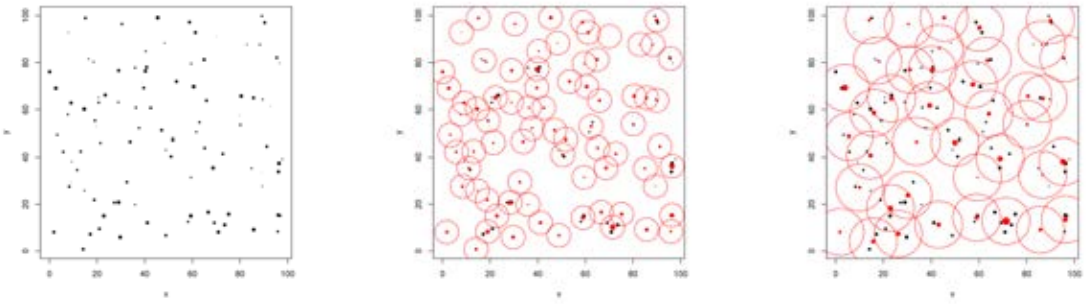

(a) random particle field, 100 particles (b) coalescing radius = 5, final particle count = 75 (c) coalescing radius = 10, final particle count = 40

Figure A.2: Demonstration of the effect of coalescing on a random field of 100 particles in two dimensions, for which mass is represented by the particle size. (a) shows the original particle field and (b) and (c) show the final particle field in red plotted over the original in black after coalescing with search radii 5 and 10 respectively. (b) and (c) also show the search radii as circles around each final particle.

Appendix B

SUMMARY INFORMATION AND PLOTS FOR THE TESTING OF THE DYNAMIC RANDOM WALK SCHEME

This appendix shows summary plots for all the benchmark simulations, for which parameters are given in table B.1. For each figure, the subfigures represent the following:

(a) a time slice (time = 3000) of the plume solutions from DRW, Method of Characteristics (MOC) and the analytical Domenico (1987) equation (2.25); note that the contour values are logarithmically spaced

(b) a transverse section of the interpreted concentration from the three solutions at $x = 350$ and time = 3000

(c) the solute flux with time to the downstream receptor (on the right, figure 2.19) for the three solutions

(d) the particle count with time for the DRW solution

The legend in figure B.1 shows which colour is used for each solution. Additionally, the blue lines at the left and right of the plume plots represent the fixed head boundaries. For subfigures (a) and (b) of each plot, the concentration fields for the DRW solution have been estimated from a weighted kernel density estimate of the particle swarms, using `R` package `ks` (R Core Team, 2014; Duong, 2016). This procedure convolves the point masses with bivariate normal distributions in the x and y dimensions, with standard deviations for the normal distributions of 5 in each direction, to produce a smoothed concentration field. The breakthrough flux (c) requires no post-processing as MassTrack saves mass removed at sinks by cell reference (Barry, 2017d).

Table B.1: Parameter combinations used in each of the benchmark simulations. Grid descriptions refer back to table 2.2.

simulation	longitudinal dispersivity	first-order degradation constant	retardation factor	effective porosity	time step	coalescing search radius	number of dispersion pairs	grid
1	1	0	1	0.10	50	1	4	default
2	3	0	1	0.10	50	2	4	default
3	10	0	1	0.10	50	5	4	default
4	30	0	1	0.10	50	5	4	default
5	10	0.000346	1	0.10	50	5	4	default
6	10	0.001386	1	0.10	50	5	4	default
7	10	0	2	0.10	50	5	4	default
8	10	0	5	0.10	50	5	4	default
9	10	0	1	0.03	50	5	4	default
10	10	0	1	0.30	50	1	4	default
11	10	0	1	0.10	10	5	4	default
12	10	0	1	0.10	20	5	4	default
13	10	0	1	0.10	100	5	4	default
14	10	0	1	0.10	50	2	4	default
15	10	0	1	0.10	50	10	4	default
16	10	0	1	0.10	50	50	4	default
17	10	0	1	0.10	50	5	1	default
18	10	0	1	0.10	50	5	2	default
19	10	0	1	0.10	50	5	10	default
20	10	0	1	0.10	50	5	4	coarse
21	10	0.000346	2	0.10	50	5	4	default
22	10	0.001386	5	0.10	50	5	4	default
23	10	0	1	0.10	50	2	10	default
24	10	0	1	0.10	50	10	1	default
25	10	0	1	0.10	50	5	4	very coarse

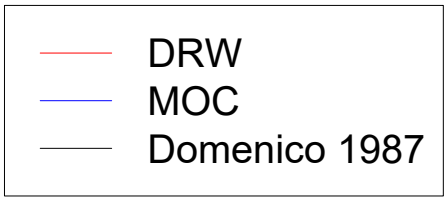

Figure B.1: Legend for plots in figures B.2–B.26, (a–c), showing the colour coding used for each transport solution

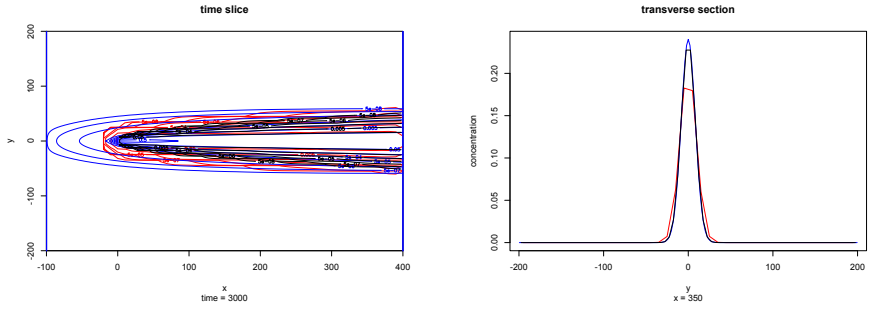

(a) (b)

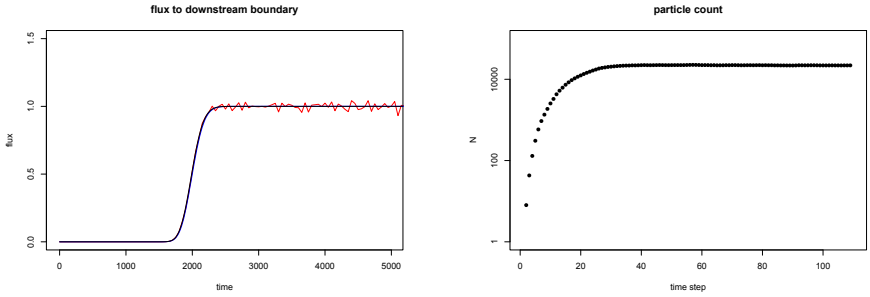

(c) (d)

Figure B.2: benchmark simulation 1

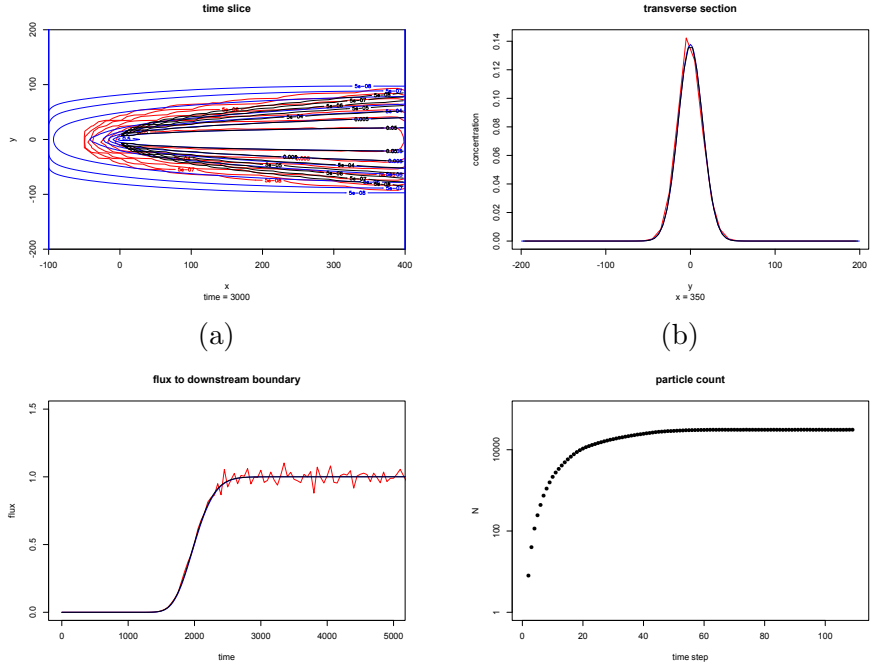

(a) (b)

(c) (d)

Figure B.3: benchmark simulation 2

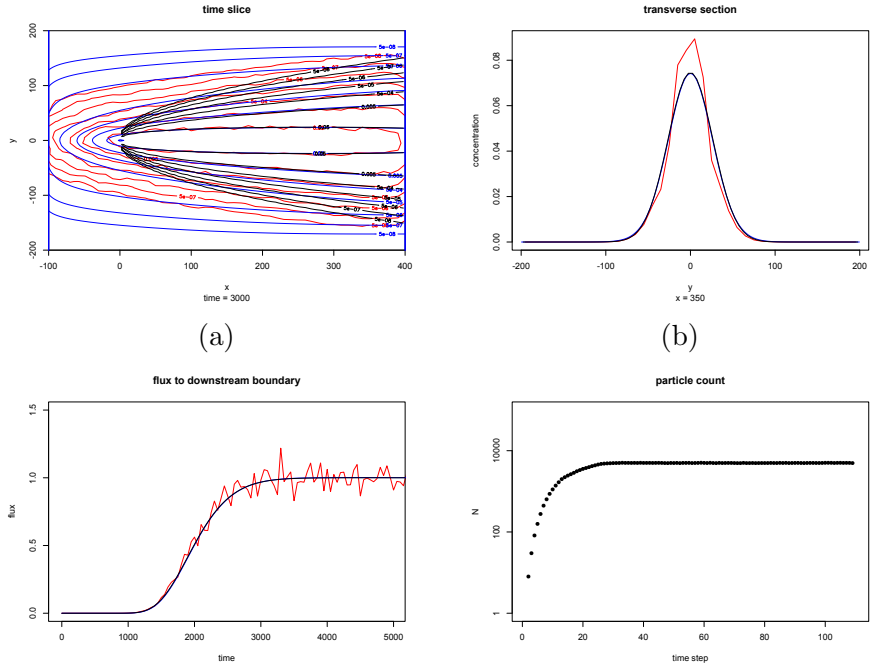

(a) (b)

(c) (d)

Figure B.4: benchmark simulation 3

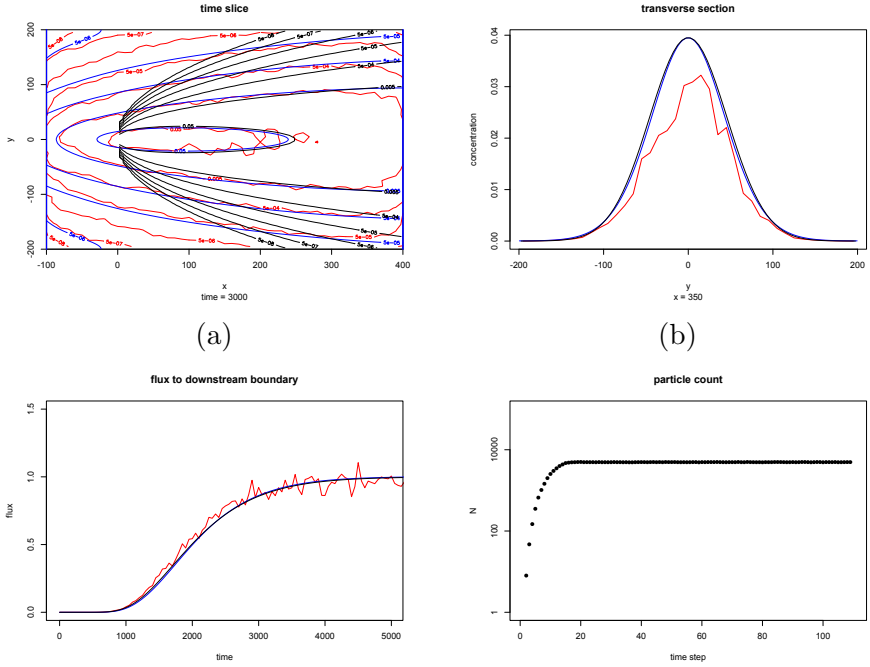

(a) (b)

(c) (d)

Figure B.5: benchmark simulation 4

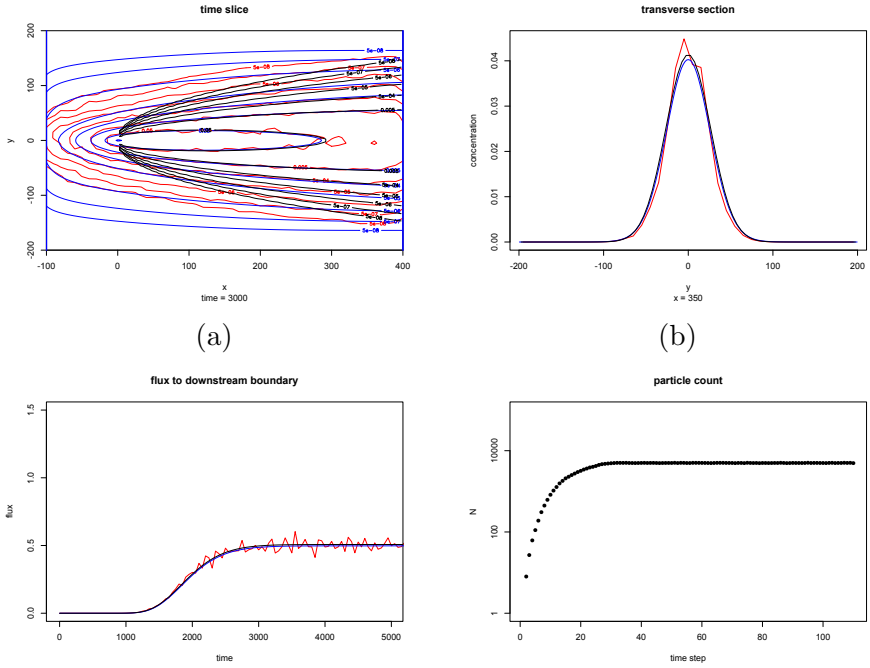

(a) (b)

(c) (d)

Figure B.6: benchmark simulation 5

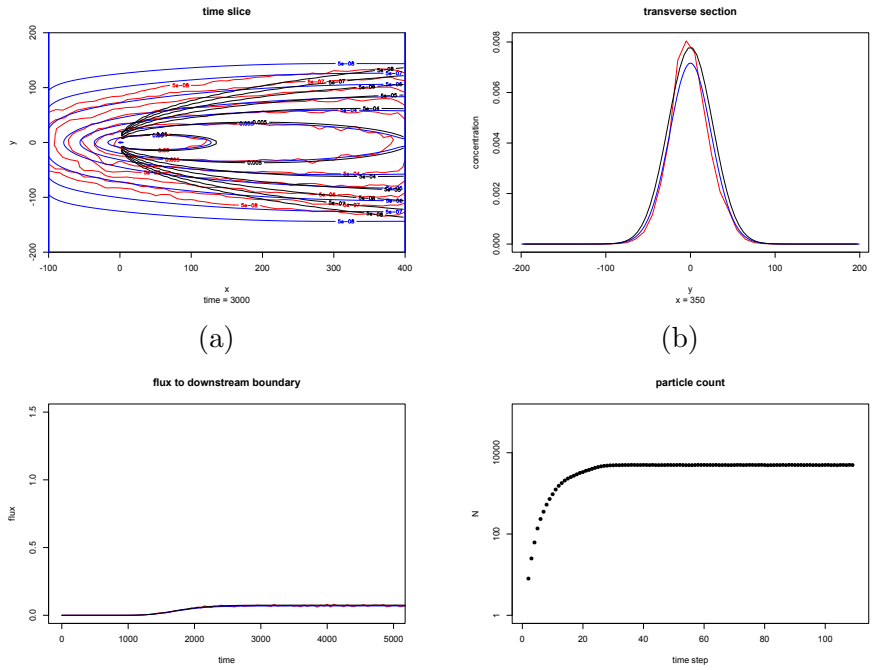

(a) (b)

(c) (d)

Figure B.7: benchmark simulation 6

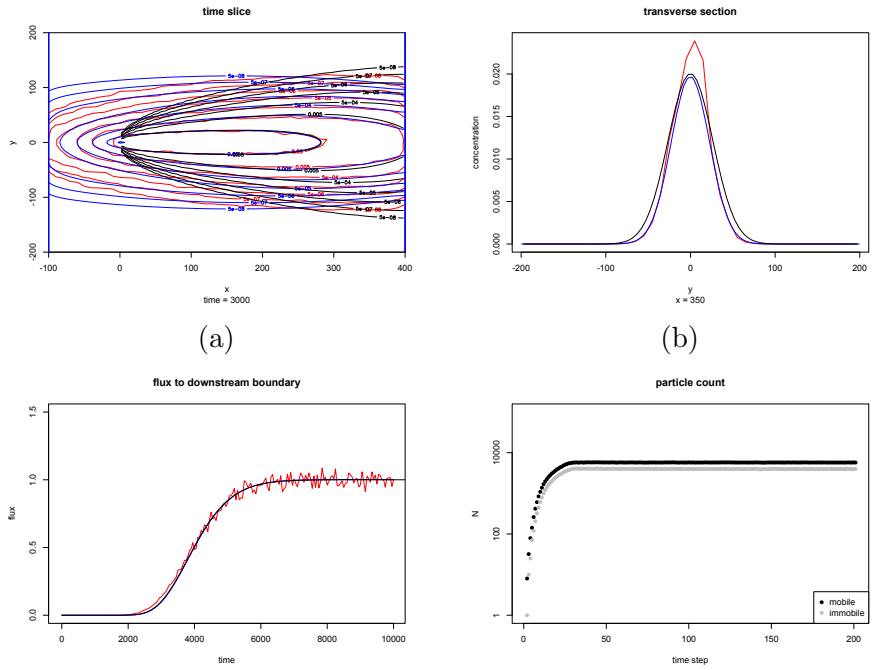

(a) (b)

(c) (d)

Figure B.8: benchmark simulation 7

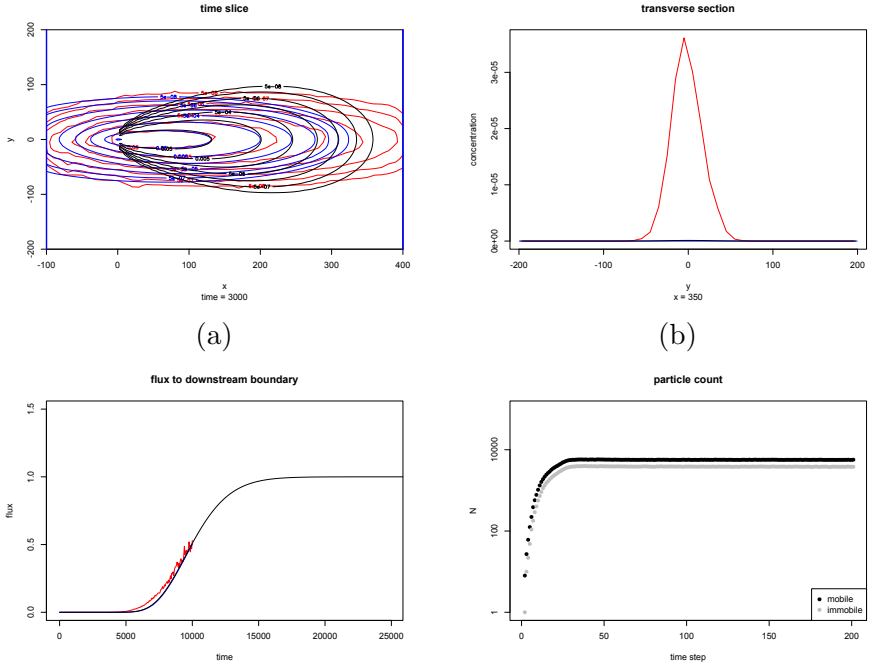

(a) (b)

(c) (d)

Figure B.9: benchmark simulation 8

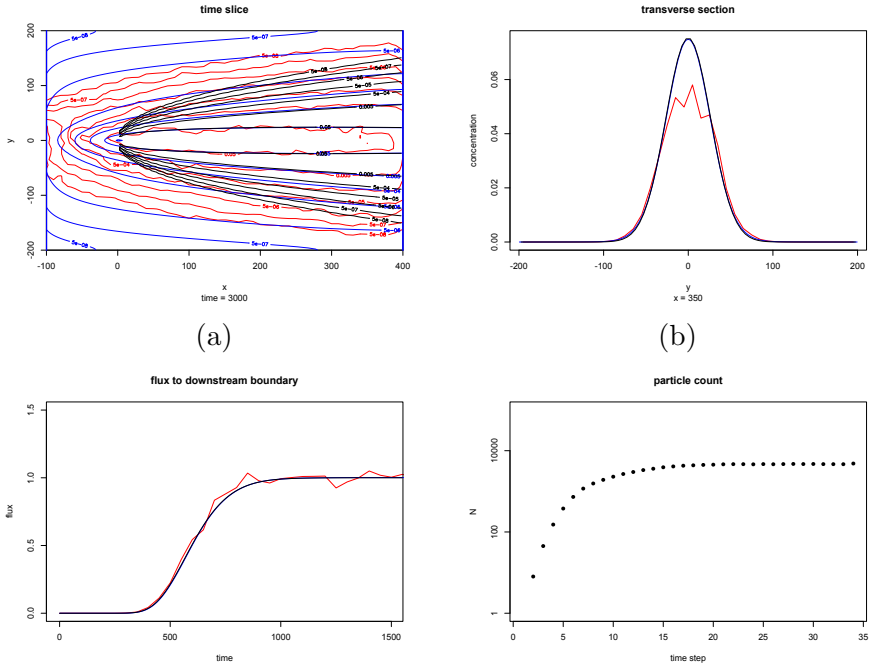

(a) (b)

(c) (d)

Figure B.10: benchmark simulation 9

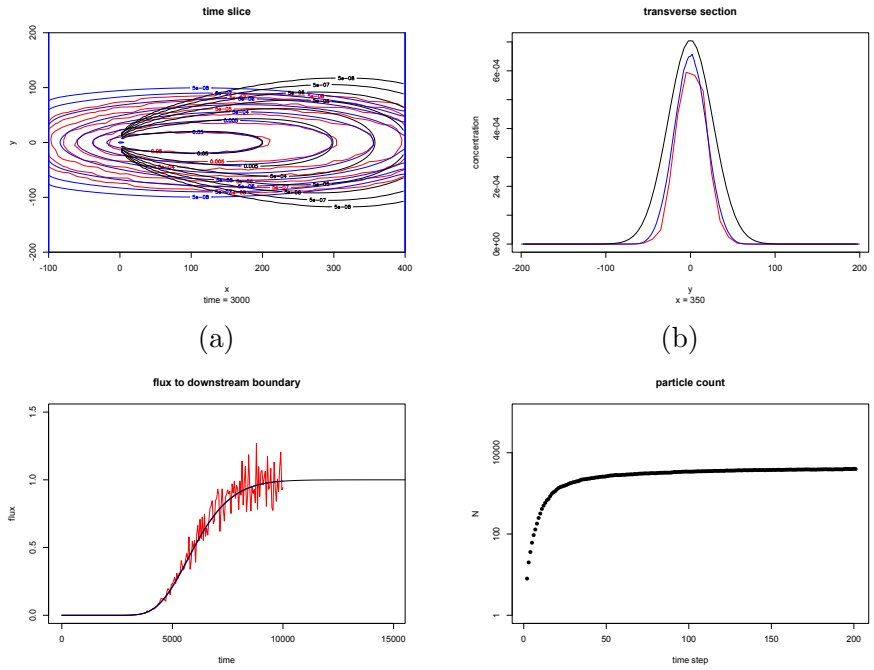

(a) (b)

(c) (d)

Figure B.11: benchmark simulation 10

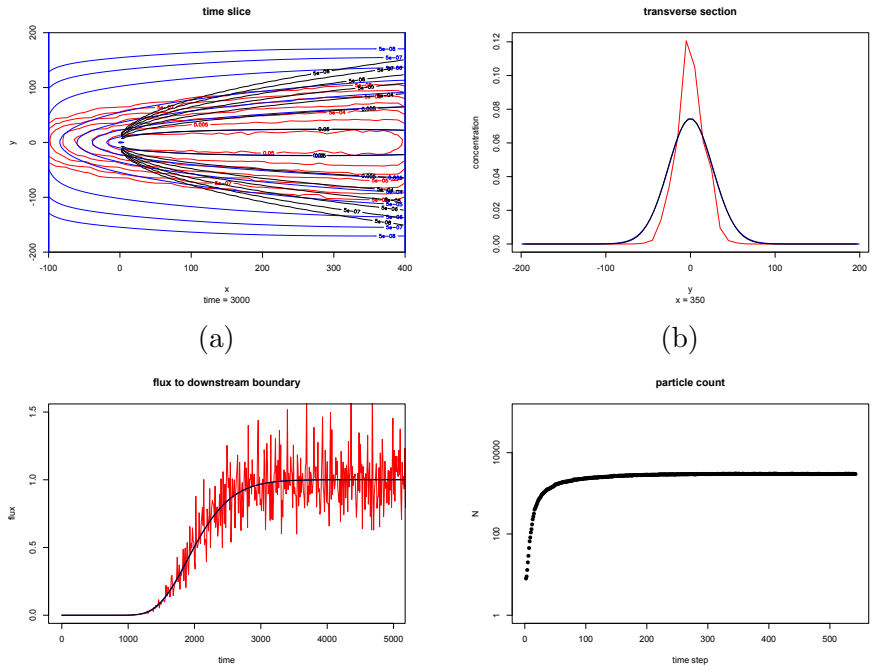

(a) (b)

(c) (d)

Figure B.12: benchmark simulation 11

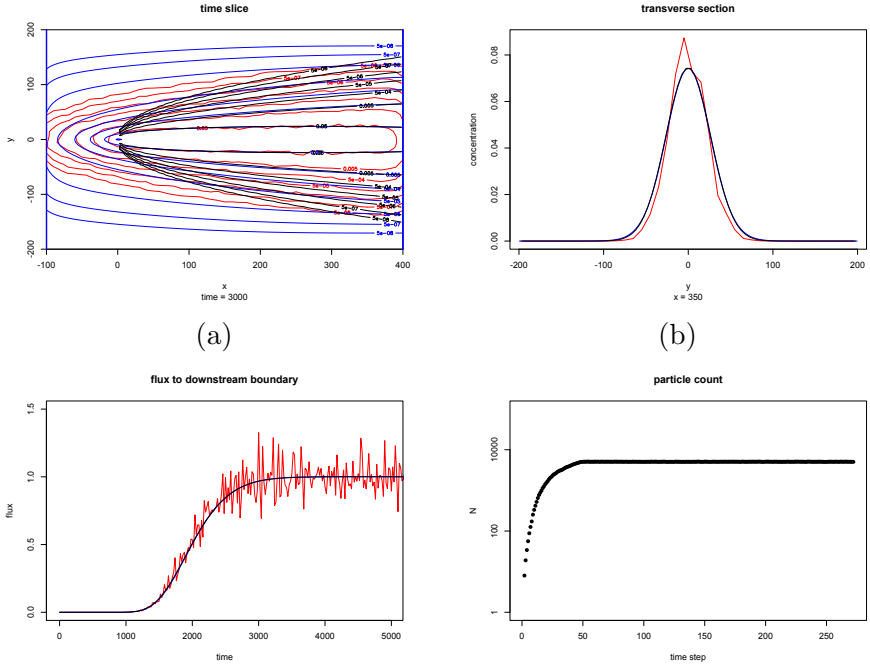

(a) (b)

(c) (d)

Figure B.13: benchmark simulation 12

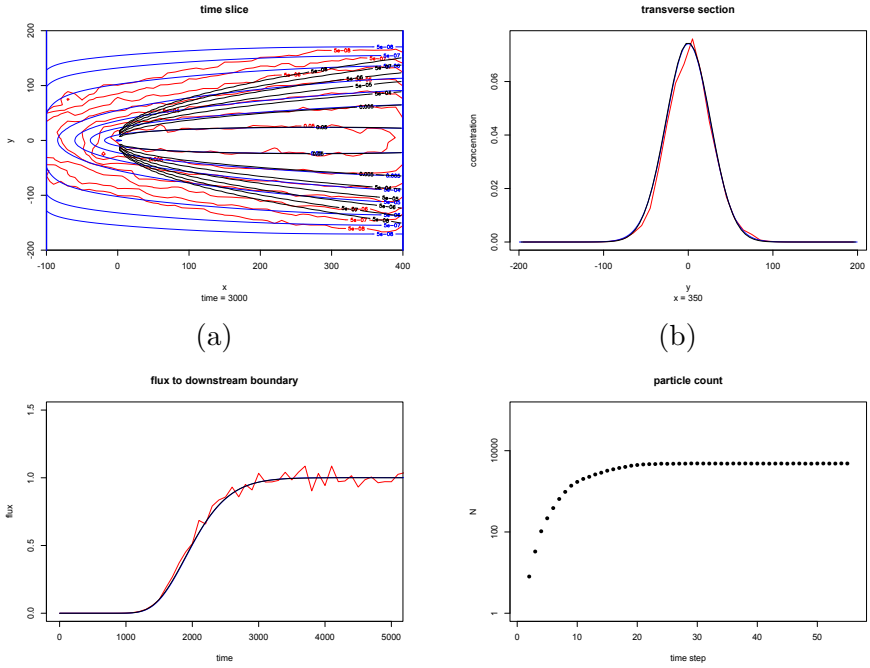

(a) (b)

(c) (d)

Figure B.14: benchmark simulation 13

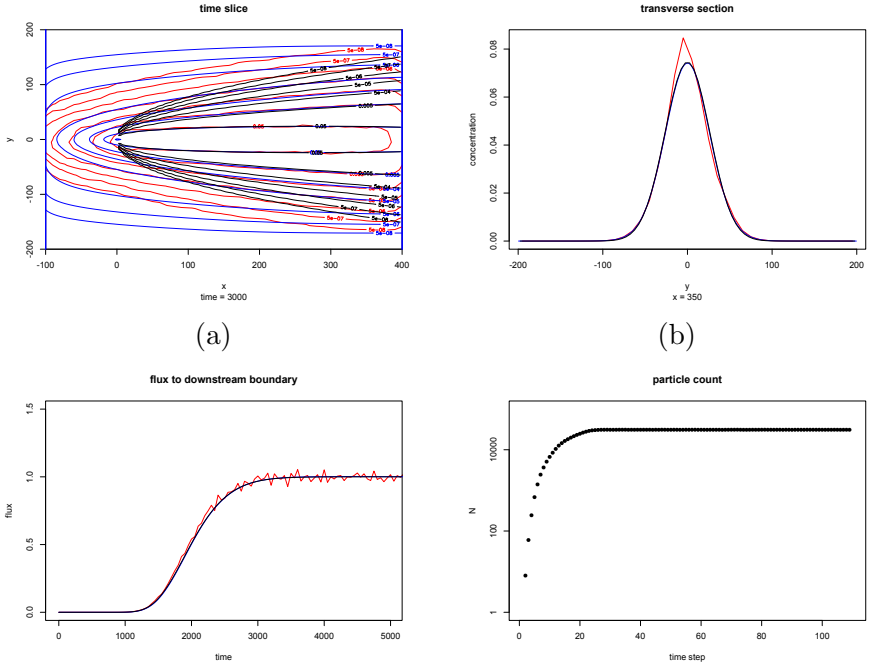

(a) (b)

(c) (d)

Figure B.15: benchmark simulation 14

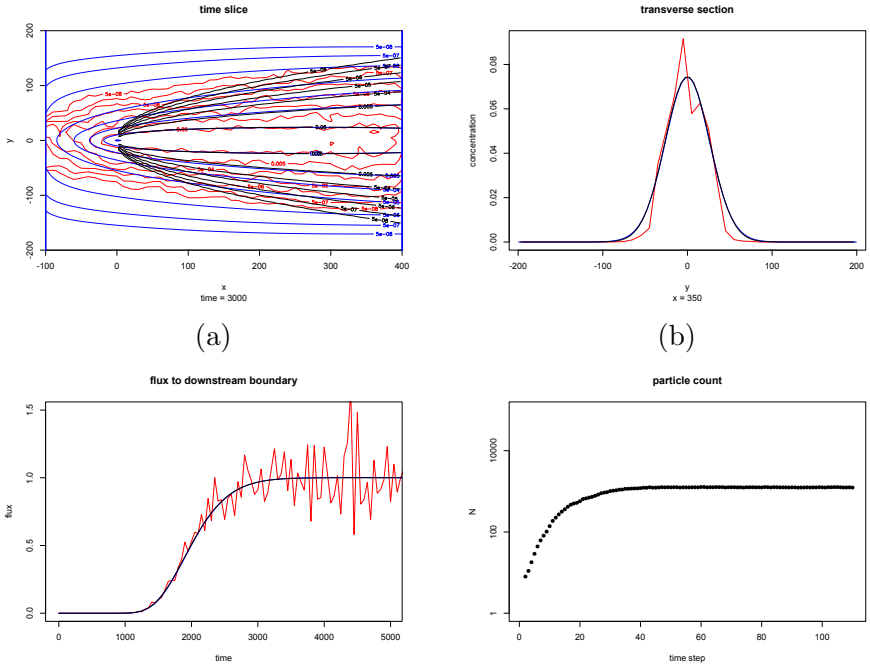

(a) (b)

(c) (d)

Figure B.16: benchmark simulation 15

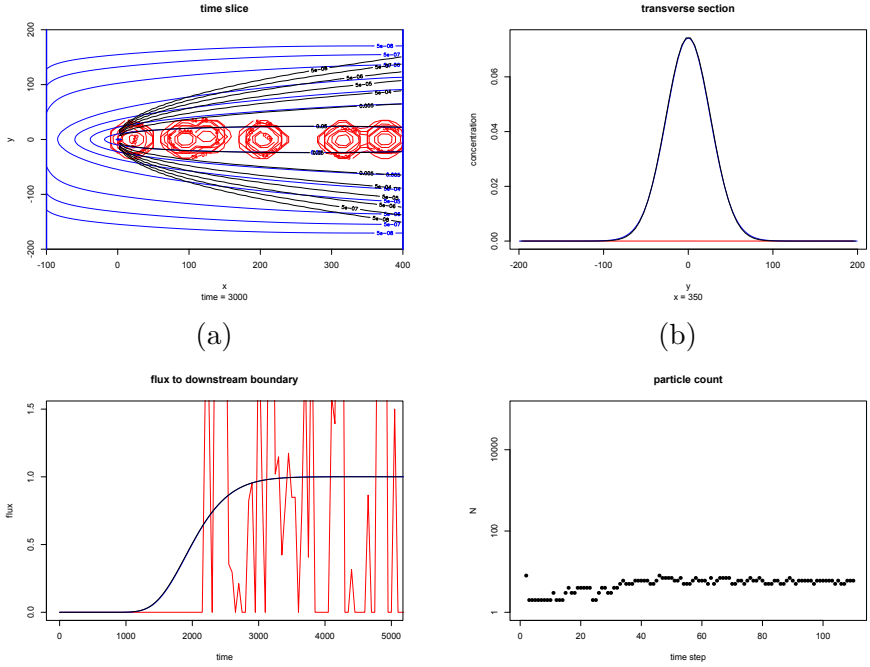

(a) (b)

(c) (d)

Figure B.17: benchmark simulation 16

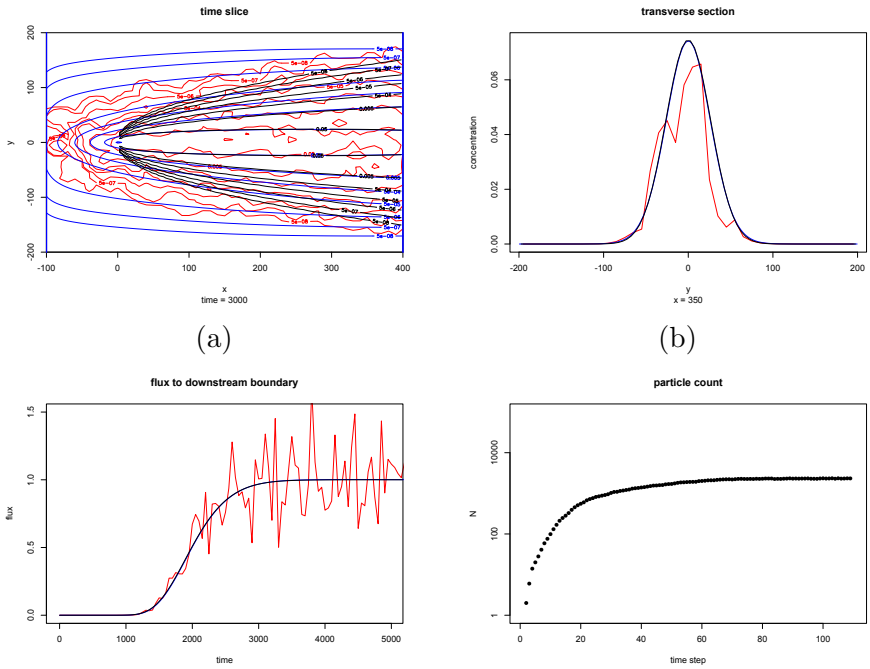

(a) (b)

(c) (d)

Figure B.18: benchmark simulation 17

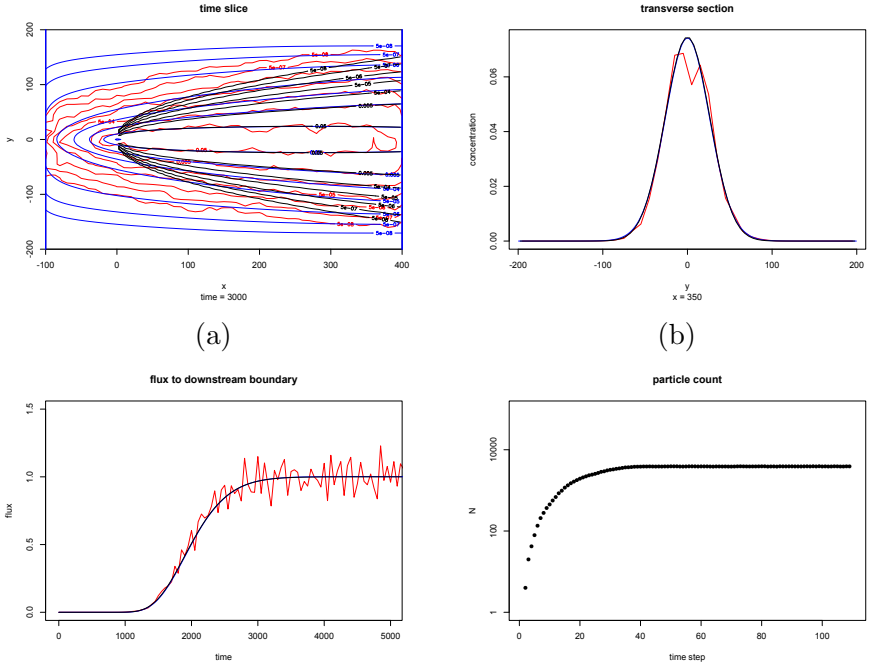

(a) (b)

(c) (d)

Figure B.19: benchmark simulation 18

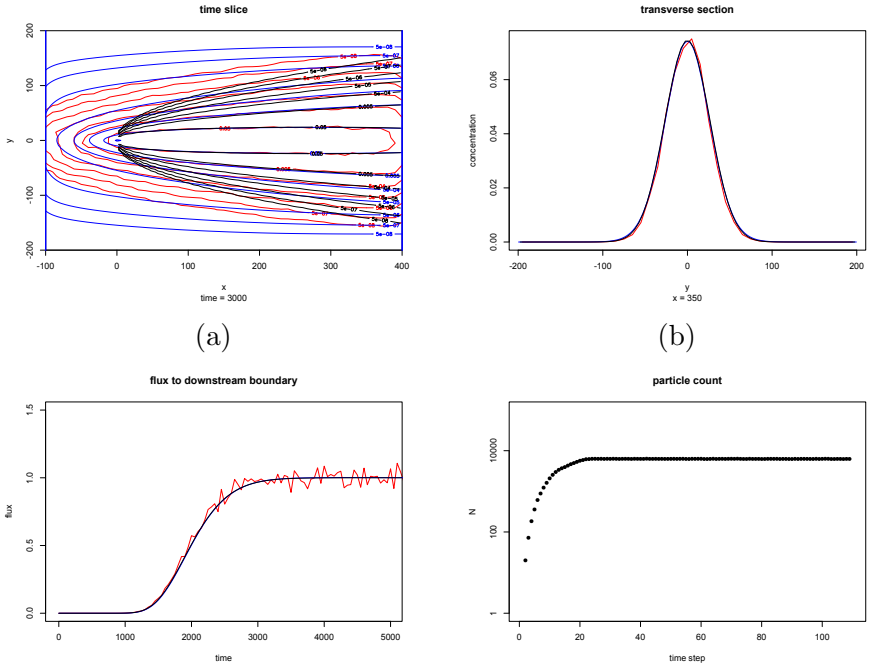

(a) (b)

(c) (d)

Figure B.20: benchmark simulation 19

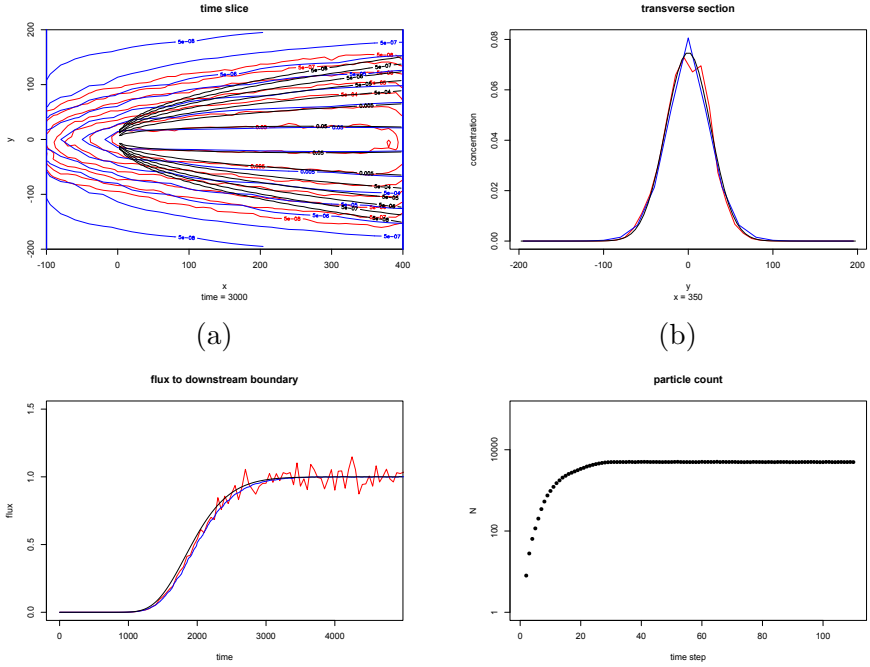

(a) (b)

(c) (d)

Figure B.21: benchmark simulation 20

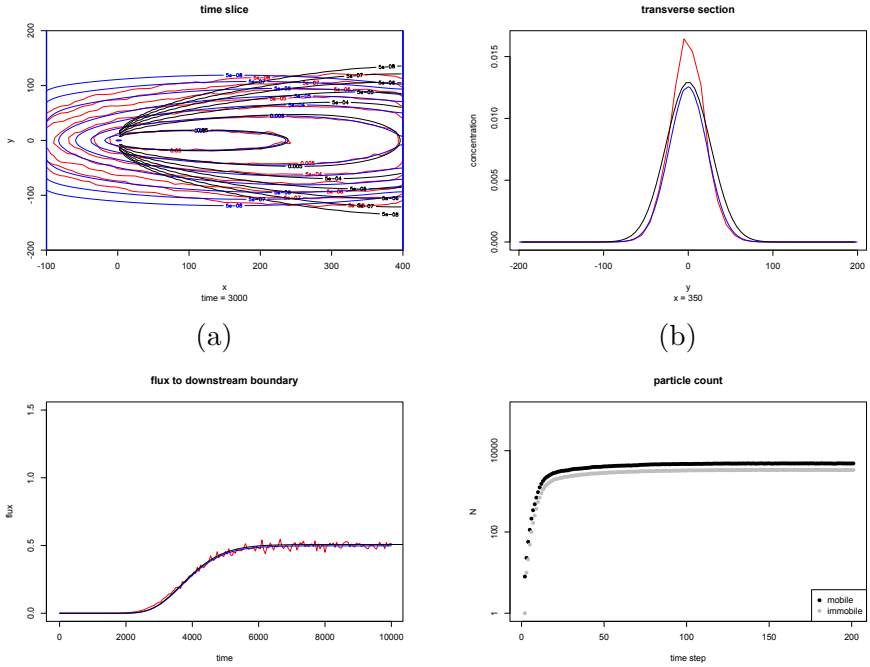

(a) (b)

(c) (d)

Figure B.22: benchmark simulation 21

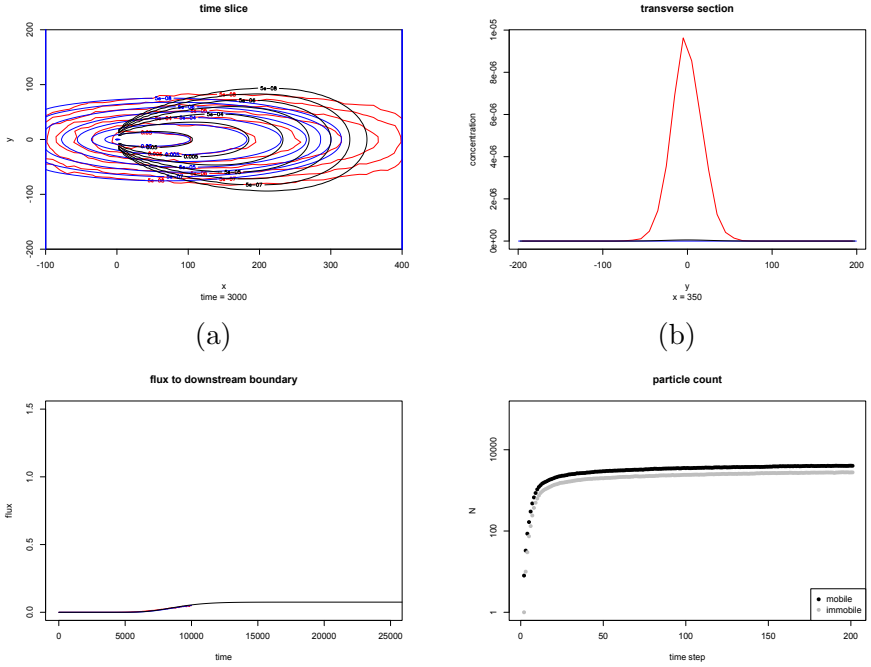

(a) (b)

(c) (d)

Figure B.23: benchmark simulation 22

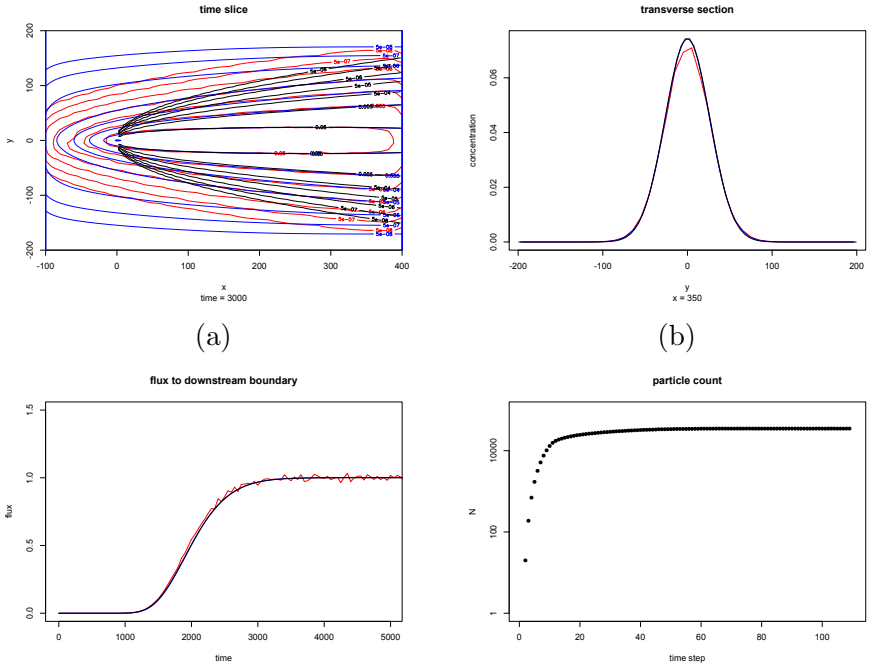

(a) (b)

(c) (d)

Figure B.24: benchmark simulation 23

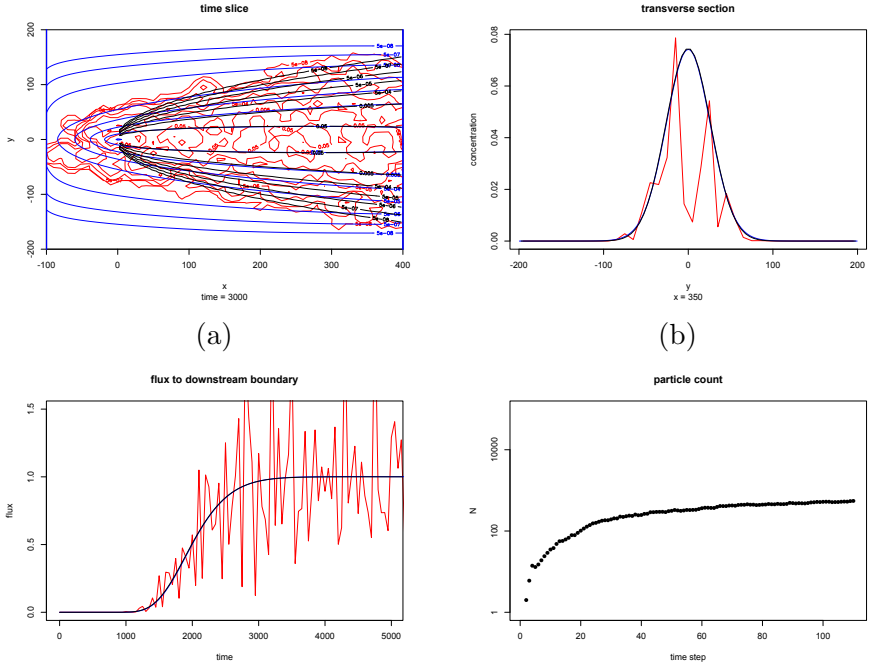

(a) (b)

(c) (d)

Figure B.25: benchmark simulation 24

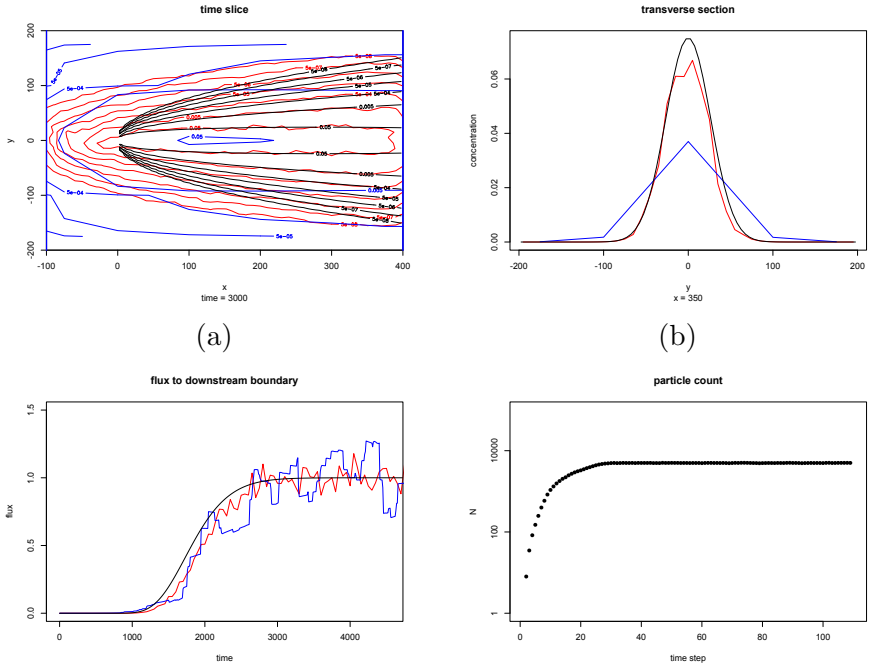

(a) (b)

(c) (d)

Figure B.26: benchmark simulation 25

Appendix C

HISTORICAL MAP REFERENCES

The historic industrial land use models for Birmingham and Stourbridge were constructed using the EDINA Historic Digimap Service. The service includes a collection of historical maps through the twentieth century and it was possible to identify historic industrial sites, their site boundaries and the rough date of opening and closing in each case. The type of industrial activity at each site was normally evident from labels on the map and sometimes was found through internet searches. The service may be found at `http://digimap.edina.ac.uk`. Original source data published 1906–1996 at various scales. The complete list of datasets, for which further details can be found at `http://digimap.edina.ac.uk/webhelp/historic/historicdigimaphelp.htm#about_historic_maps/map_summary.htm` is shown in table C.1.

Table C.1: Landmark map data sets used in this study.

Landmark epoch	scales	publishing date range
Epoch 3	1:2500	1906–1939
	1:10,560	1903–1949
Epoch 4	1:2500	1924–1949
	1:10,560	1922–1969
Epoch A5	1:1250	1943–1993
	1:2500	1943–1995
Epoch B6	1:1250	1944–1993
	1:2500	1949–1992
Epoch C7	1:1250	1946–1993
	1:2500	1954–1993
Epoch D8	1:1250	1951–1992
	1:2500	1960–1992
Epoch E9	1:1250	1953–1992
	1:2500	1967–1992
Epoch i5	1:10,560	1948–1977
Epoch i6	1:10,560	1949–1981
Epoch i7	1:10,560	1959–1982
Epoch i8	1:10,560	1967–1976
Epoch m5	1:10,000	1969–1996
Epoch m6	1:10,000	1973–1996
Epoch m5	1:10,560	1958–1996
	1:10,000	

Appendix D

SUPPLEMENTARY ELECTRONIC MATERIAL

D.1 Software packages

`R` packages developed for this project are included as zipped folders in the `SEM/packages/` subdirectory. To install them, `R`, which has a free license, must first be installed from `https://cran.r-project.org/`. RStudio is also highly recommended as an interactive environment for `R`, but is not essential (`https://www.rstudio.com/`). The included packages may then be installed by copying the zipped folders into a personal directory, for example `C:/<user>/Documents/`, and running the following commands in an `R` session:

```
# external packages from CRAN repository - requires internet connection
# - dependencies will be automatically installed
# -- fast data table sorting and grouping, the workhorse of DRW
install.packages("data.table")
install.packages("RNetCDF") # interface to NetCDF data sets
# -- smoothed concentration field post-processing for DRW models
install.packages("ks")
install.packages("pracma") # for the error function and its inverse
install.packages("abind") # tools for multi-dimensional arrays
install.packages("sp")
install.packages("stringr")
install.packages("rlist")
install.packages("plyr")

# included packages with this thesis
install.packages("~/Rflow.zip", repos = NULL)
install.packages("~/td.zip", repos = NULL)
install.packages("~/DNAPL.zip", repos = NULL)
install.packages("~/MassTrack.zip", repos = NULL)
install.packages("~/coalesce.zip", repos = NULL)
install.packages("~/DRW.zip", repos = NULL)
```

The order is important because some packages depend on others. "~" refers to the user's home directory, which is usually `C:/<user>/Documents/`. Packages then need to be loaded to be used, using the `library` function (`library(Rflow)`, for example). Table D.1 contains brief package descriptions. Further documentation is contained within the packages. Use the `help` function to access in-package documentation (`help(package = "Rflow")` or `help("GW.nc")`, for example).

D.2 Plume animations

Animated graphics of transient plume evolution are included in GIF format. For the Tame Valley case study in the `SEM/TameValley/` subdirectory, animated plumes for tetrachloroethene (PCE), trichloroethene (TCE) and 1,1,1–trichloroethane (1,1,1–TCA) are included, as indicated by the file names. A summary figure is in figure 4.10. The colour flood concentration legends are:

PCE:

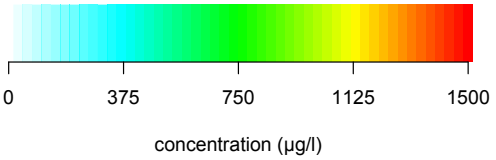

TCE and 1,1,1–TCA:

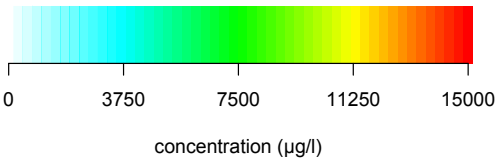

For the Stourbridge case study in the `SEM/Stourbridge/` subdirectory, animated plumes of normalised concentration (see section 5.3.2) are included for the original West Midlands Worfe regional model and the two tested adaptations (see section 5.2) and for the seven future Winnie Road abstraction scenarios (see section 5.3.1), 21 in total, with file names clearly labelled. The colour flood normalised concentration legend is:

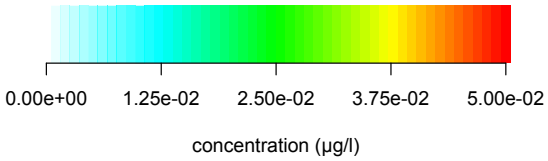

A legend of plotted features applicable to all animated figures is:

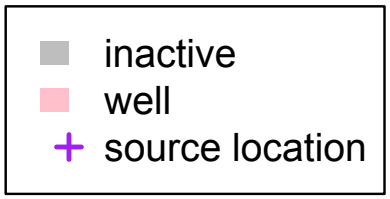

Table D.1: Description of packages

`Rflow`	Tools for analysing and writing input and output files for USGS MODFLOW. This package was used for example in adapting the West Midlands Worfe regional model boundary conditions (section 5.2). Also required by the `DNAPL` and `DRW` packages.
`td`	Tool for consistent referencing of dates to day numbers, useful for models spanning multi-year time frames. 30/12/1899 is set at 0, and each day counts as one unit. It helps to ensure that linking models are consistently time-referenced. It is consistent with Microsoft Excel's day numbering scheme except that it correctly does not include 29/02/1900.
`DNAPL`	The DNAPL source term framework. See chapter 3 especially section 3.4.
`MassTrack`	Track Particle Mass along MODPATH 5-Calculated Pathlines in Groundwater. See section 2.2. A key component of `DRW` but also useful as a standalone package.
`coalesce`	Coalesce a field of weighted points. Highly optimised. See appendix A. A key component of `DRW`.
`DRW`	The Dynamic Random Walk model. See section 2.5.