



RESEARCH ARTICLE

10.1002/2014WR016394

Key Points:

- We develop a new methodology to reconstruct breakthrough curves from tracked particles
- Error is evaluated as a function of the number of particles used
- The new method outperforms traditional reconstructions for risk metrics

Correspondence to:

E. R. Siirila-Woodburn,
erwoodburn@lbl.gov

Citation:

Siirila-Woodburn, E. R., D. Fernández-García, and X. Sanchez-Vila (2015), Improving the accuracy of risk prediction from particle-based breakthrough curves reconstructed with kernel density estimators, *Water Resour. Res.*, 51, 4574–4591, doi:10.1002/2014WR016394.

Received 11 SEP 2014

Accepted 13 MAY 2015

Accepted article online 20 JUN 2015

Published online 25 JUN 2015

Improving the accuracy of risk prediction from particle-based breakthrough curves reconstructed with kernel density estimators

Erica R. Siirila-Woodburn^{1,2}, Daniel Fernández-García¹, and Xavier Sanchez-Vila¹
¹Hydrogeology Group, Department of Geotechnical Engineering and Geosciences, Universitat Politècnica de Catalunya, Barcelona, Spain, ²Now at Earth Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California, USA

Abstract While particle tracking techniques are often used in risk frameworks, the number of particles needed to properly derive risk metrics such as average concentration for a given exposure duration is often unknown. If too few particles are used, error may propagate into the risk estimate. In this work, we provide a less error-prone methodology for the direct reconstruction of exposure duration averaged concentration versus time breakthrough curves from particle arrival times at a compliance surface. The approach is based on obtaining a suboptimal kernel density estimator that is applied to the sampled particle arrival times. The corresponding estimates of risk metrics obtained with this method largely outperform those by means of traditional methods (reconstruction of the breakthrough curve followed by the integration of concentration in time over the exposure duration). This is particularly true when the number of particles used in the numerical simulation is small ($<10^5$), and for small exposure times. Percent error in the peak of averaged breakthrough curves is approximately zero for all scenarios and all methods tested when the number of particles is $\geq 10^5$. Our results illustrate that obtaining a representative average exposure concentration is reliant on the information contained in each individual tracked particle, more so when the number of particles is small. They further illustrate the usefulness of defining problem-specific kernel density estimators to properly reconstruct the observables of interest in a particle tracking framework without relying on the use of an extremely large number of particles.

1. Introduction

The process of computing both environmental and human health risk begins with an examination of the environmental concentration statistics, all of which are typically derived from breakthrough curves (BTCs). Of specific importance is the temporally averaged concentration over the exposure duration, $ED [T]$, which is considered the true amount of exposure posing some adverse effect. The U.S. EPA Risk Assessment Guidance for Superfund [U.S. EPA, 1989, 2001] recommends that the exposure concentration, C , to be used in risk assessment calculations be defined as “the arithmetic average of the concentration that is contacted over the exposure period.”

In hydrogeology, work has been made to link risk with environmental processes through a number of physical parameters which would affect the concentration BTC signal in the modeling process. For example, Andričević and Cvetković [1996] showed that geologic heterogeneity and uncertainty in the sorption estimate are important factors in determining risk. Maxwell et al. [1998] showed that risk decreased as subsurface heterogeneity increased. Enzenhoefer et al. [2012] explained how the time between peak and bulk solute breakthrough, important in risk analysis, increase in the presence of fractured media, or as aquifer heterogeneity increases. Siirila and Maxwell [2012b] showed that the statistical anisotropy of the aquifer and the inclusion of small-scale processes such as kinetic sorption and local dispersion in contaminant transport modeling could potentially raise risk over the remediation action level. de Barros and Fiori [2014] found similar results regarding the importance of statistical anisotropy and transverse dispersion in determining the uncertainty of environmental concentration distributions. These works suggest that hydrogeologic parameters and processes are important in accurately assessing risk. But, despite this body of literature, it should be pointed out that the emphasis in many papers is in either the peak concentration or the shape of the tail in the BTC. Rather, the U.S. EPA and other agencies propose evaluating risk from some representative

concentration average over a predefined duration. This work focuses on this latter, less studied point. That is, what is the effect of BTC representation (and the associated uncertainty) on the calculation of the long-term average concentration, and how does this effect impact risk-driven decision analysis?

A proper analysis of time-averaged concentrations is heavily reliant on an accurate representation of the BTC. A key element to consider in BTC representation is the modeling method adopted to reconstruct the BTC from numerical simulations. Eulerian methods tend to produce numerical dispersion oversmoothing on both the solute plume and the subsequent BTC. Lagrangian methods by means of particle tracking (PT) are free of numerical dispersion but include other shortcomings, like the need to reconstruct concentrations from particle distributions in space or in time. In any case, PT methods are shown to be effective and computationally inexpensive compared to other methods [Salamon *et al.*, 2006; Boso *et al.*, 2013; Henri and Fernández-García, 2014]. PT methods have been successfully used for the last 60 years in physics [De Josselin de Jong, 1958; Scheidegger, 1954] and for 35 years in subsurface hydrology [Ahlstrom *et al.*, 1977; Prickett *et al.*, 1981]. Since the number of particles (np) used for tracking is a modeler's choice, two basic questions arise: (1) how can the different observables of interest be reconstructed from a finite number of particles, and (2) since in theory exact results can only be approached as np tends toward infinity, what is the minimum number of particles to meaningfully reconstruct such observables?

Kernel density estimators (KDEs) have been extensively used in a wide number of applications to provide smooth estimates of different variables of interest, including image reconstruction, robotics, astronomy, process optimization, and fluid mechanics, amongst others. Early examples in the groundwater literature are from *Tompson* [1993] and *Tompson et al.* [1996] who used a fixed kernel support proportional to the mesh size to estimate concentrations during operator-splitting reactive transport simulations. The smoothed-particle hydrodynamics (SPH) method [e.g., *Morris et al.*, 1997; *Liu et al.*, 2003; *Herrera et al.*, 2009] also introduced a kernel function to smooth the properties of particles within a radial distance, but in this case the set of particles carry all the concentration information. In any case, the support of the kernel is a free parameter that exhibits a strong influence on the resulting estimate of concentrations. Although the support of the kernel function can be fixed, these methods do not take advantage of the full power of the KDE, which can adapt the support to the number of particles and the intrinsic shape of concentrations. Alternative methods for reconstructing breakthrough curves from the arrival statistics of particles are also available in the literature, such as the inverse Gaussian method [e.g., *Hathhorn and Charbeneau*, 1994; *Enzenhofer et al.*, 2014]. Contrary to the KDE, this method is essentially parametric and therefore linked to a specific shape of the breakthrough curve.

Advancements in the use of kernels in reconstruction processes include obtaining optimal kernel support estimates that can lead to improved methods for reconstruction. Optimal KDE methods have been recently proposed in the PT literature to overcome oversmoothing and statistical artificial fluctuations in the BTC [Fernández-García and Sánchez-Vila, 2011; Pedretti and Fernández-García, 2013; Siirila-Woodburn *et al.*, 2015]. Optimal KDE methods allow reconstructing an unbiased estimator of the BTC, further providing a measure of the estimation error that adapts to the number of particles used in the simulation and the smoothness of the solution. Additionally, KDE methods show promise in diminishing tradeoffs traditionally thought to be necessary between simulation time and accuracy.

In this work, we focus on the benefits of using optimal KDEs to reconstruct the curves for both concentration and average concentration as a function of time. Such benefits are assessed by comparing the reconstructed curves with and without the use of an optimal KDE. Different hydraulic conductivity fields are simulated as a way of obtaining a range of typical BTC shapes of ranging complexity, and are then used as proxies of representative curves. Average concentration versus time curves for a given exposure duration (ED) are then derived. In both KDE and non-KDE cases, the curves are reconstructed as a function of the number of particles used. This comparison is used to assess how the number of particles can be reduced with the inclusion of optimal kernel estimators.

The paper is structured as follows. Section 2 introduces the main concepts used in the paper including curve reconstruction using kernels and the average daily dose. In section 3, we develop a direct methodology to compute the average concentration BTCs directly from the distribution of a finite number of particle arrival times, and as a function of exposure duration, without the need to reconstruct the full BTC. For this purpose, a (suboptimal) KDE being a function of exposure duration is derived directly from the optimal KDE

that would be obtained for BTC reconstruction. Section 4 outlines the methods and metrics used to assess the performance of the proposed approach. Finally, in section 5, we compare concentration and average concentration curve reconstructions utilizing traditional and new methods described here. The relative error of each method is quantified with the use of a high-resolution (in terms of number of particles) simulation. The implications of BTC characterization and reconstruction are discussed throughout the paper in the context of risk analysis and risk management.

2. Kernel Density Estimators and Application to Risk Assessment Metrics

The reconstruction of BTCs from a finite number of particles entails assigning to each individual particle a mass and a corresponding arrival time at a compliance surface. Since particle tracking methods involve the use of a finite number of particles, the compilation of arrival times can be seen as a subset obtained from random sampling of the actual distribution of arrival values. Thus, arrival times can be treated as a sample subset drawn from the full probability density function, allowing for the use of statistical reconstruction techniques. Once the arrival times are known, two main types of approaches for BTC reconstruction can be used, as described below.

2.1. Traditional Reconstruction of Concentration BTCs

The traditional reconstruction method consists of selecting a number of bins in time. Bins must be selected as a partition (i.e., not overlapping and without gaps) of the time domain. Particles are then univoquely classified into a bin. At a given time, the mass of solute reaching the compliance surface is then computed by adding the number of particles (actually adding their corresponding mass) associated with a given bin. Concentrations are derived from the knowledge of particle mass and flow rate. This method converges to the real solution for a bin size tending to zero, thus the need to use a very large number of particles, especially when the receding limb of the BTC is of interest. To avoid confusion, any reference to a “traditional” particle tracking method in this manuscript refers to this process. The acronym “PT” for particle tracking opposed to “KDE” for kernel density estimator will differentiate this method and the following one.

2.2. Kernel Density Estimator (KDE) Reconstruction of Concentration BTCs

In this approach, the key idea is that the arrival time of each particle is a realization from a statistical distribution, indicating that the particles used in the numerical simulation are just a subsample from a large population of all the possible particles that could have been selected a priori. The kernel function $K [T^{-1}]$ is a symmetric probability distribution function that accounts for the uncertainty associated with such subsampling. Denoting the arrival time of particle i as t_i , a Gaussian distribution centered at t_i is typically used for this purpose (a non-Gaussian kernel would be equally valid and are the subject of future analysis)

$$K(t-t_i; h) = \frac{1}{\sqrt{2\pi}h} \exp\left(-\frac{(t-t_i)^2}{2h^2}\right) \quad (1)$$

The parameter h is the standard deviation of the kernel function and controls the degree of uncertainty, and thus should be a value that relates to np . The integral of K over all time is by definition one. Given a random sample of travel times $\{t_1, \dots, t_{np}\}$ obtained from a numerical particle tracking simulation, the concentration BTC can be reconstructed from the superposition of the particle locations

$$C(t) \approx \frac{1}{Q} \sum_{i=1}^{np} m_i K(t-t_i; h) \quad (2)$$

where $C(t)$ is the flux-averaged concentration at time t , Q is the total flow rate at the control location, and m_i is the mass of the i th particle. For ease in notation, we refer to the duration $t-t_i$ as u below. The same formulation applies to resident concentrations by slightly changing the meaning of Q to incorporate porosity.

In (2), the parameter h of the KDE controls the degree of smoothing. Small h values produce noisy estimates and large h values produce biasedness and oversmoothing of the breakthrough curve [see, e.g., Pedretti and Fernández-García, 2013]. The optimal h balances these two effects by minimizing fluctuations and bias

in the estimates. An asymptotic optimal estimate of h , obtained in the limit when np tends to infinity while the product $np \cdot h$ tends to zero, can be determined by [Silverman, 1986; Hardle, 1990]

$$h_{opt} = \left(\frac{R(K_1)}{S(x) \sigma^4(K_1)} \right) np^{-1/5} \quad (3)$$

where $K_1(u) = K(u;1)$ is the unitary elemental kernel function (meaning that h is set to unity), $R(K)$ is the integral of K^2 over u , $\sigma(K_1)$ is the second moment of $K_1(u)$, and $S(x)$ is the kernel-based estimate of the second derivative of the normalized breakthrough curve (see, e.g., Engel et al. [1994] or Fernández-García and Sánchez-Vila [2011], for details on the calculation of $S(x)$ and the optimal h). Although this method offers a valid approach to estimate early-time behavior and peak of BTCs, it produces significant fluctuations at the tails, where fewer particles exist. In this context, Pedretti and Fernández-García [2013] proposed a locally adaptive method named UAB to improve tail estimation while maintaining the accuracy in both early-time and peak concentrations.

2.3. Risk Metric: Average Daily Dose (ADD)

In human health risk analysis under nonacute conditions, emphasis is placed on exposure to hazardous substances. Exposure is typically defined by means of a concept known as the average daily dose [$M_{\text{contaminant}}^{-1} M_{\text{human}}^{-1} T^{-1}$], defined as [U.S. EPA, 1989]:

$$ADD = \langle C \rangle \left[\frac{IN}{BW} \right] \frac{ED \times EF}{AT} \quad (4)$$

where IN/BW [$L^3 M_{\text{human}}^{-1} T^{-1}$] is the intake rate per unit body weight of a representative individual, AT is the averaging time or expected lifetime [T], EF is the standard exposure frequency [-], ED is the exposure duration and $\langle C \rangle$ is the average concentration an individual is exposed to during that particular ED value. Thus, the ED represents the true interaction time between an individual and the hazardous substance. A discussion on the effect of the averaging period of the concentration is discussed by the International Programme on Chemical Safety, IPCS [WHO, 2005]. For example, the U.S. EPA suggests a 30 year exposure duration for noncarcinogens and a 70 year exposure duration (the convention for a lifetime) for carcinogens [U.S. EPA, 1989].

The choice of the operational ED depends also on the given exposed population in the community. When children are known to be at special risk and considered as a different population cohort, it may be more appropriate to use an ED between 1 and 6 years [e.g., U.S. EPA, 2008]. While a smaller ED more accurately characterizes high-frequency concentration fluctuations [e.g., Maxwell et al., 2008; Maxwell and Kastenber, 1999; Siirila and Maxwell, 2012a], smaller ED values typically also result in larger $\langle C \rangle$ values. However, it is important to note that the mechanism for calculating exposure may balance the effect of the larger $\langle C \rangle$ with a smaller ED . As shown in equation (4), the ADD linearly decreases as the ED decreases, somehow balancing the effect of the larger associated $\langle C \rangle$ values, with the net effect of diminishing the overall sensitivity of the ED to the ADD . The sensitivity of the ED on the shape of the BTC for the different methods is also explored in this paper.

3. Evaluating $\langle C \rangle$ From a Lagrangian Prospective

In this section, we present a methodology for the evaluation of $\langle C \rangle$ over an ED and as a function of time, directly derived from the distribution of arrival times recorded from tracked particles at a compliance surface (in general, a plane perpendicular to the main flow direction or at a monitoring or pumping well). The methodology will be tested in further sections.

For the evolution of the concentration arriving at a compliance surface (expressed as $C(t)$, the definition of a BTC), the temporally averaged BTC for a fixed ED is given by

$$\langle C(t; ED) \rangle = \frac{1}{ED} \int_{t-ED/2}^{t+ED/2} C(\tau) d\tau \quad (5)$$

Once an actual $C(t)$ curve is available, $\langle C(t; ED) \rangle$ can be easily reconstructed from equation (5). Nevertheless, in a particle tracking framework, the average concentration versus time curve is not directly available, but

rather should be reconstructed from the values of arrival times associated with the ensemble of particles. A simple approach could be composed of two steps: first, reconstruct the BTC by means of any available methodology such as those discussed in section 2.1, and second, use the reconstructed curve to perform the integral in equation (5) for a number of discrete values of t .

The limitation of this method arises from the ability to properly reconstruct the initial BTC in the first step. As stated before, a direct reconstruction could imply the need for a large np value. Moreover, available KDE methods do not ensure a stable propagation of the residual statistical fluctuations associated with the calculation of concentrations in subsequent risk metric estimations. Here we propose an alternative method based on the KDE approach that overcomes this problem by providing a direct estimation of the averaged breakthrough curve (integral of concentrations over the ED) without the calculation of the initial BTC. Essentially, the proposed KDE approach expresses that in order to estimate the integral over a curve there is no need to know the entire curve in detail. A direct kernel density estimator of the arithmetic average of the concentration over the exposure duration can be written as

$$\langle C(t; ED) \rangle \approx \frac{1}{Q} \sum_{i=1}^{np} m_i \bar{K}(t - t_i; h, ED) \quad (6)$$

while in general the new KDE appearing in equation (6), $\bar{K}(t; h, ED)$, could be assumed, it can be related to that used for BTC reconstruction and provided in equation (2). Starting with equation (5) and using the function for C as given in equation (2) we obtain

$$\langle C(t; ED) \rangle \approx \frac{1}{Q ED} \int_{t-ED/2}^{t+ED/2} \sum_{i=1}^{np} m_i K(\tau - t_i; h) d\tau \quad (7)$$

By direct identification between equations (6) and (7)

$$\bar{K}(t - t_i; h, ED) = \frac{1}{ED} \int_{t-ED/2}^{t+ED/2} K(\tau - t_i; h) d\tau \quad (8)$$

If the local kernel is Gaussian, and given by equation (1), the final expression for the kernel is obtained as a combination of error functions

$$\bar{K}(u; h, ED) = \frac{1}{2 ED} \left[\operatorname{erf}\left(\frac{u+ED/2}{\sqrt{2} h}\right) - \operatorname{erf}\left(\frac{u-ED/2}{\sqrt{2} h}\right) \right] \quad (9)$$

This kernel is still a probability density function that integrates to one. Finally the expression for the average concentration at the compliance surface for a given exposure duration can be directly estimated from particle data with the following formula, where we have considered as a simplified case that all particles share the same mass, so that $m_i = M_T / np$

$$\langle C_{ANA}^{np}(t; ED) \rangle = \frac{M_T}{2Q np ED} \sum_{i=1}^{np} \left[\operatorname{erf}\left(\frac{t+(ED/2)-t_i}{\sqrt{2} h}\right) - \operatorname{erf}\left(\frac{t-(ED/2)-t_i}{\sqrt{2} h}\right) \right] \quad (10)$$

By including total mass instead of the individual mass of each particle in equation (10), we explicitly remark that the expression is np dependent. In a general case where the mass of particles is not constant, the mass of each particle would be included in the sum in equation (10), and np would not be explicitly present. The suffix ANA is kept for future comparisons and indicates that the kernel was analytically derived. Slight variations of equation (10) would be obtained if the limits of the integral in equation (5) were chosen so as the ED was not centered at a given time t_0 $\{t \in [t_0 - \frac{ED}{2}, t_0 + \frac{ED}{2}]\}$, but rather defined forward $\{t \in [t_0, t_0 + ED]\}$, or backward $\{t \in [t_0 - ED, t_0]\}$. While the actual results presented later would differ, the discussion and conclusions of the relevance of ED would hold.

The new kernel expression (equation (9)) is shown in Figure 1 as a function of normalized exposure duration ($F = ED/h$). For the purpose of comparison, the same figure shows the non-ED averaged kernel (equation (1), black line). For small F values, the average kernel from equation (9), $\bar{K}(u; h, ED)$, is very similar to Gaussian, and in the limit ($F \rightarrow 0$) it would converge to equation (1). On the other hand, a large F value




Figure 1. Shapes of the averaged kernel in equation (9) as a function of normalized exposure duration, $F = ED/h$ (see key).

would imply a flattening of the curve, thus extending the impact of the kernel to a longer duration in time, and reducing the value of the peak. When $F \rightarrow \infty$, the kernel function $\bar{K}(u; h, ED)$ approaches a box model with support equal to the ED . In such a case, the method is independent of h and equation (10) simply expresses a particle count per ED interval.

In general, when F is not a large number, the definition of a new estimator should involve the selection of a new optimal h value. In this case, the asymptotic solution is not known but we will demonstrate through numerical simulations that equation (3) still provides a good suboptimal solution for health risk analysis. Nevertheless some preliminary analysis on the value of h can be performed. We start by considering that the BTC displays a Gaussian-like shape with a standard deviation σ . In this case, equation (3) simplifies to

$$h_{opt} \cong 1.06 \sigma np^{-1/5} \quad (11a)$$

The time interval with nonnegligible values of concentrations (T) in the BTC can be estimated, for instance, as $T = 6.58\sigma$, which accounts for 99.9% of mass. Defining a constant dimensionless value $a = T/ED$ yields a standard deviation equivalent to $\sigma = a \times ED/6.58$, where equation (11a) becomes

$$h_{opt} \cong 0.16 a ED np^{-1/5} \quad (11b)$$

Then, the optimal ratio of ED to h_{opt} (equation (11b)) is defined by:

$$F_{opt} = \frac{ED}{h_{opt}} \cong 6.2 a^{-1} np^{1/5} \quad (12)$$

In human health risk analysis, the U.S. EPA [1989] recommends assuming that receptors are exposed to the long-term average water concentrations following a contamination period of at least 30 years. This typically leads to long BTCs with values of a ranging between 5 and 1500 (depending on the choice of ED) and to small values of the dimensionless ratio F ($F < 1$). From this, we conclude that, in practical situations, F is relatively small and therefore h_{opt} can be directly estimated through equation (3).

4. Numerical Setup, Observables, and Performance Calculations

4.1. PT Simulations

To assess the performance of the proposed approach, the Lagrangian particle tracking model, SLIM-FAST [Maxwell, 2010], was used to simulate nonreactive solute transport for four different BTCs of varying solute behavior (see section 4.2). SLIM-FAST uses the traditional particle tracking (PT) method to reconstruct BTCs

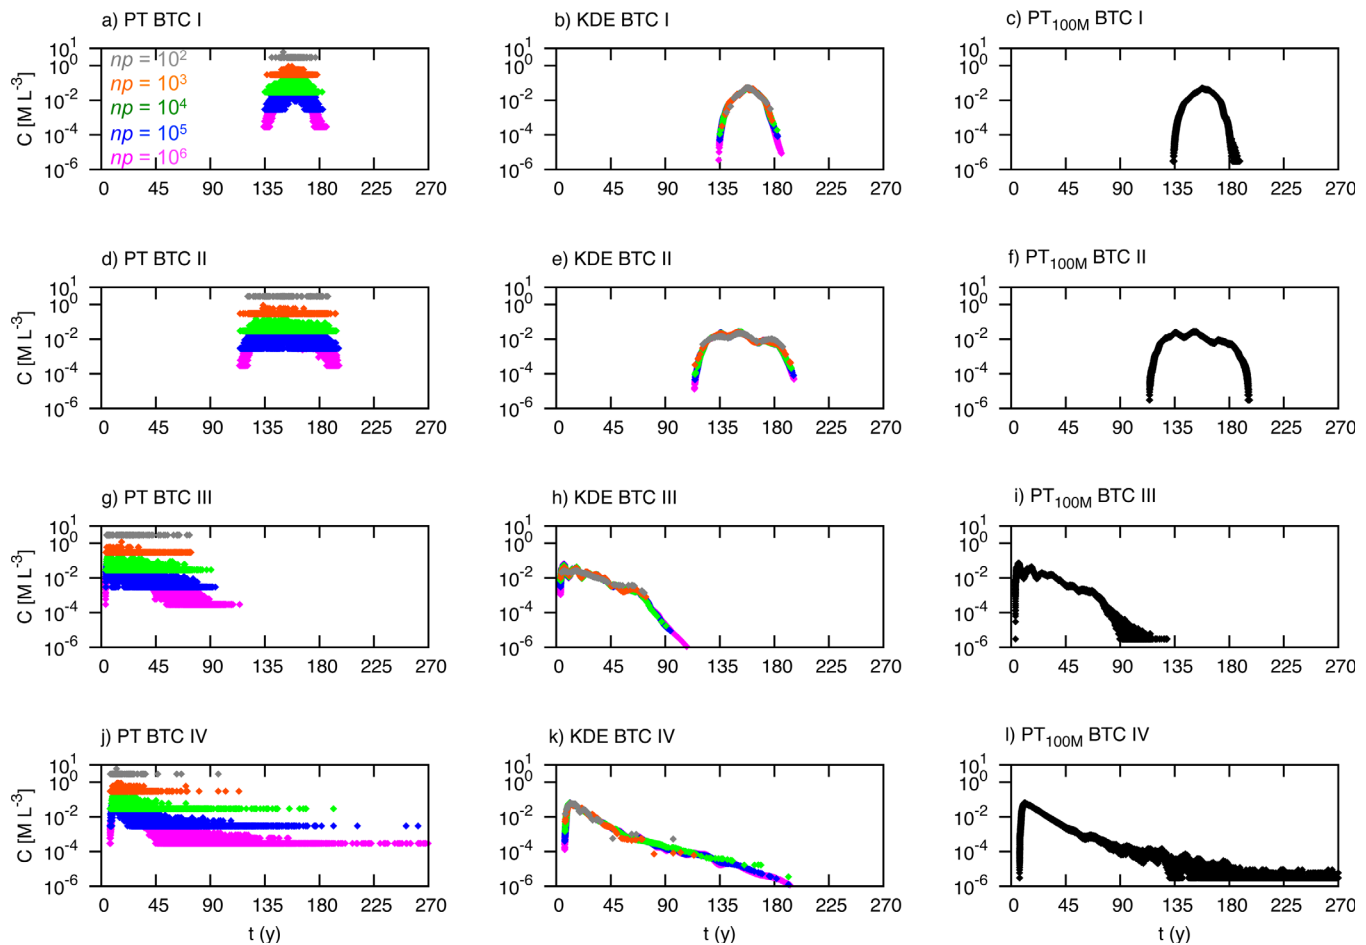


Figure 2. Representative BTCs I-IV (rows) simulated with variable np (colors, see a). For a given BTC, results are shown from left to right for PT simulations (a, d, g, j), KDE reconstructed curves (b, e, h, k), and for the corresponding 100M PT simulations (c, f, i, l).

with evenly distributed bins in time, as described in section 2.1. A full description of the flow and transport parameters can be found in Appendix A. For each of the four scenarios analyzed, solute transport is solved using the particle tracking model with varying np , spanning over 4 orders of magnitude ($np = 10^2 - 10^6$). Results are also compared to a $np = 10^8$ simulation of the same flow field to test for accuracy. An initial concentration of 1 $[ML^{-3}]$, equivalent to a total mass of 300 $[M]$, is used in each simulation regardless of np . Thus, greater histogram resolution of particle mass results for high np and low histogram resolution of particle mass results for low np . The particle tracking seed will pose more sensitivity on the results as np decrease due to the random nature of the method. Because our purpose is to illustrate the differences between PT and KDE methods for a single simulation, the seed is not altered in our simulations. Multiple simulations of differing seeds to quantify the uncertainty in the random walk of the particle tracking simulation may be an interesting topic of future studies.

For each value of np used ($10^2 - 10^6$), the raw PT BTC is compared with a reconstructed KDE BTC. Several metrics are used for this quantitative comparison as discussed in sections 4.3 and 4.4. The KDE approach used here incorporated a universal adaptive bandwidth (UAB) method as developed by Pedretti and Fernández-García [2013]. UAB is automatic, and locally adaptive, playing on the strengths of the established global and adaptive bandwidth methods especially to reconstruct heavy-tailed BTCs. Following the discussion in Pedretti and Fernández-García [2013], the singular dimensionless input parameter which controls sensitivity, $\alpha \in [0, 1]$, is fixed to 0.5. Further information on the role of α can be found in Pedretti and Fernández-García [2013].

4.2. Representative BTCs

The aforementioned four representative BTCs were chosen to span an array of realistic and typical BTCs that may result in natural settings so that they can be considered somewhat representative of BTC shapes. Here we focus on the effect of characterizing these curves, not on the processes that created them. As such,

a description of the parameterization of these curves is listed in Appendix A, but is not discussed in further detail. Figure 2 shows the four representative BTC cases (I–IV) as a function of np . Representative BTC I (Figure 2a) is nearly Gaussian, showing somewhat ideal Fickian conditions. Representative BTCs II–IV (Figures 2d, 2g, and 2j) show some degree of anomalous behavior, displaying a number of local peaks of different duration in time [Siirila-Woodburn *et al.*, 2015]. BTC IV (Figure 2j) resembles a power-law like curve. Although these curves were generated for subsurface solute transport, we should note that this comparison is applicable in many fields where particle tracking methods are used, such as biological processes, aerodynamics, and medical applications, to name a few.

4.3. Reconstruction of Unaveraged BTCs, C Versus t , and Associated Error

PT and reconstructed KDE BTCs simulated with $np = 10^2 - 10^6$ are compared for accuracy with reference to a high resolution PT simulation of the same flow field, simulated with $np = 10^8$ (referred to as PT_{100M} hereafter). The PT_{100M} BTC is considered the “true” transport solution, or as closely to the true solution which is numerically feasible at this time, although we acknowledge the true solution will only be achieved with an infinite number of particles. Following the approach described in Siirila-Woodburn *et al.* [2015], the estimated 95% confidence interval of the PT_{100M} BTC was found to be smaller than about 10^{-8} for all times. To quantify how much the PT and KDE BTCs differ from the high resolution result, an integrated root-mean-square error (IRMSE) metric is calculated given the solute concentration measures $\{C(t_1), \dots, C(t_n)\}$

$$IRMSE_i^{np} = \sqrt{\frac{\sum_{j=1}^n (C_i^{np}(t_j) - C_{PT}^{100M}(t_j))^2}{n}} \quad (13)$$

where i is the reconstruction scenario (either PT or KDE) and where n is the number of mass breakthrough times used in the reconstruction of the curve, and the superscript in C indicates the number of particles used in the reconstruction of the curve. Equation (13) is calculated for each of the three EDs and for $np = 10^2 - 10^6$. $IRMSE_i^{np}$ is a global metric of BTC error of the entire BTC without any temporal exposure duration averaging, illustrating the closeness of fit between the PT and KDE BTCs with the high resolution PT_{100M} BTC.

4.4. Calculation of Averaged BTCs, $\langle C \rangle$ Versus t

In addition to the estimation of average concentrations based on the proposed KDE method, a temporally centered, arithmetic moving average is also used to calculate the exposure concentration as a function of time:

$$\langle C_{PT}^{np}(t) \rangle = \frac{1}{ED} \sum_{j=1}^n C_{PT}^{np}(t_j) (t_{j+1} - t_j) \quad (14)$$

where the subscript PT denotes that a traditional BTC reconstruction method is employed to estimate concentrations and where the sum extends to all t values that are comprised between $[t_0 - \frac{ED}{2}, t_0 + \frac{ED}{2}]$. Equation (14) is calculated for $np = 10^2 - 10^6$ and for three different ED values, corresponding to 1, 10, and 70 (y) for each scenario. We note here that although risk is not explicitly calculated in this paper, the magnitude of $\langle C \rangle$ is. As shown in equation (4), the magnitude of $\langle C \rangle$ is linearly proportionate with the Average Daily Dose (ADD, see section 2.3), which in turn is linearly proportionate with noncarcinogenic risk, or exponentially proportionate with carcinogenic risk [U.S. EPA, 1989].

The reconstructed curve based on KDE, $\langle C_{ANA}^{np}(t) \rangle$, is also reported for the different scenarios, exposure duration, and number of particles. The curves are obtained from equation (10) using the optimal h value estimated during the reconstruction of BTCs, equation (3). Because $\langle C_{ANA}^{np}(t) \rangle$ results in a continuous solution, some points in the $\langle C_{ANA}^{np}(t) \rangle$ BTC are affected by more original PT data than others. Particularly, when the magnitude of $\langle C \rangle$ is low, points affected by less than 10 particles can be removed from the averaged-BTC (i.e., when $\langle C \rangle \leq 10/(np \times ED)$) to improve simulation results.

4.5. $\langle C \rangle$ Comparison Metrics

Similar to the $IRMSE_i^{np}$ metric for unaveraged BTCs (equation (13)), the global difference between $\langle C \rangle$ -averaged BTCs and the $\langle C \rangle$ -averaged PT_{100M} curve can be calculated as

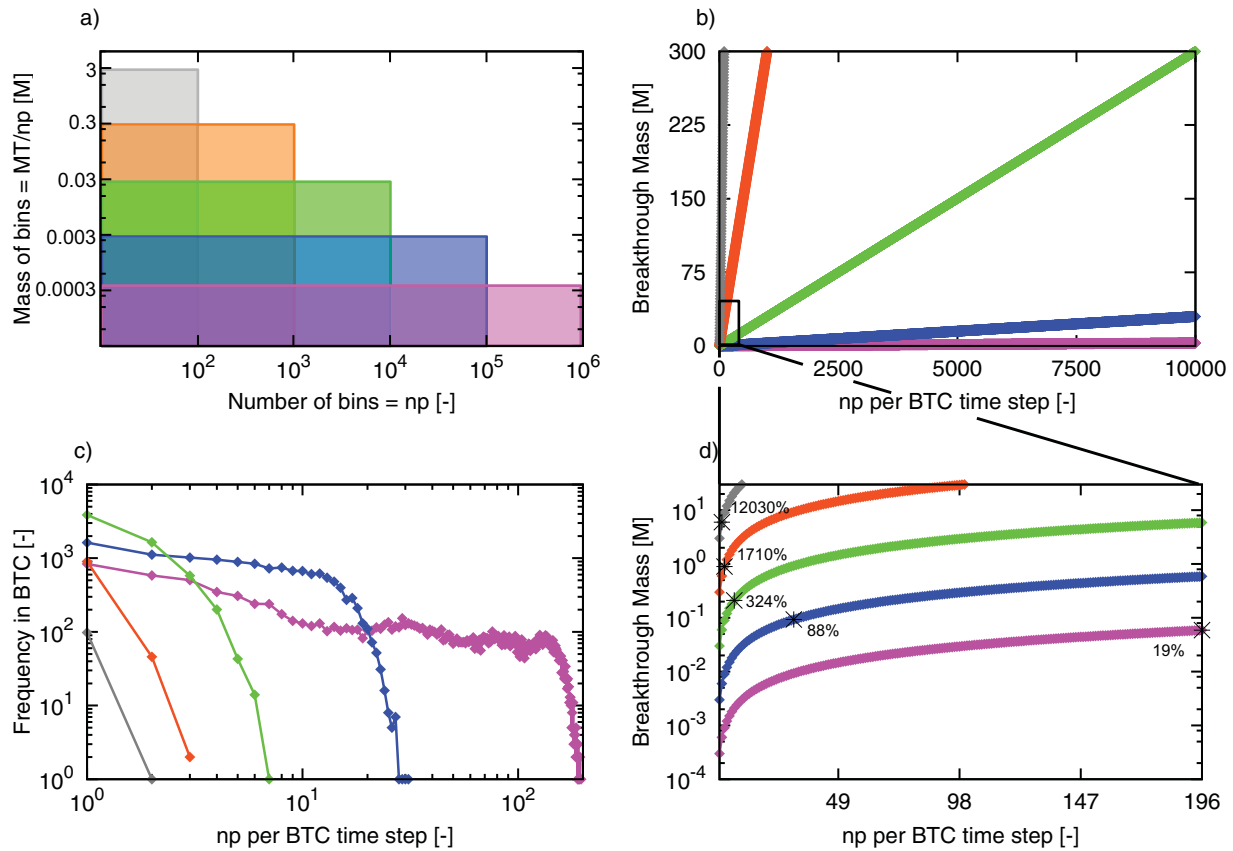


Figure 3. a) Demonstration of mass conservation in particle tracking simulations as a function of the number of particles used (see color key in Figure 2a), where total mass is equal to 300 [M] in these simulations. b) Mass breakthrough corresponding to the number of particles crossing a plane (i.e. the resolution of the histogram bin) for the different np scenarios. c) An example of histogram or bin resolution frequency of the histogram resolution for representative BTC. d) An inset of the histogram range from 0-196 particles per BTC step and the corresponding peak of each curve and resultant percent error.

$$IRMSE_{i < C}^{np} = \sqrt{\frac{\sum_{j=1}^n (\langle C_i^{np}(t_j) \rangle - \langle C_{PT}^{100M}(t_j) \rangle)^2}{n}} \quad (15)$$

$IRMSE_{i < C}^{np}$ is thus a global metric of the error of the entire BTC with exposure duration averaging, illustrating the closeness of fit between the PT and KDE BTCs with the high resolution PT_{100M} BTC.

There is also interest in assessing the impact of np and the KDE on the distribution of $\langle C \rangle$ for a given curve, as it directly influences the ADD. For a given percentile of $\langle C \rangle$, sensitivity in BTC reconstruction with variable np and with the use of the optimal KDE is assessed by calculating a percent error, PE (%):

$$PE_{i < C}^{np}(p) = \frac{\langle C_i^{np}(p) \rangle - \langle C_{PT}^{100M}(p) \rangle}{\langle C_{PT}^{100M}(p) \rangle} * 100 \quad (16)$$

where p (-) is the percentile of the $\langle C \rangle$ cumulative distribution function. Results are shown for the 100th percentile (i.e., maximum) of $\langle C \rangle$ as this value is often used in the calculation of risk as a conservative estimate [U.S. EPA, 1989, 2001].

5. Results and Discussion

5.1. Analysis of Unaveraged BTCs, C Versus t

As shown in Figure 2, PT BTCs (first column) show a large bias over several orders of magnitude when np is low (note the log y axis in Figure 2). This bias is due to the discontinuous nature of the histogram reconstruction in particle tracking methods, which yield inaccurate results when the resolution of mass is too

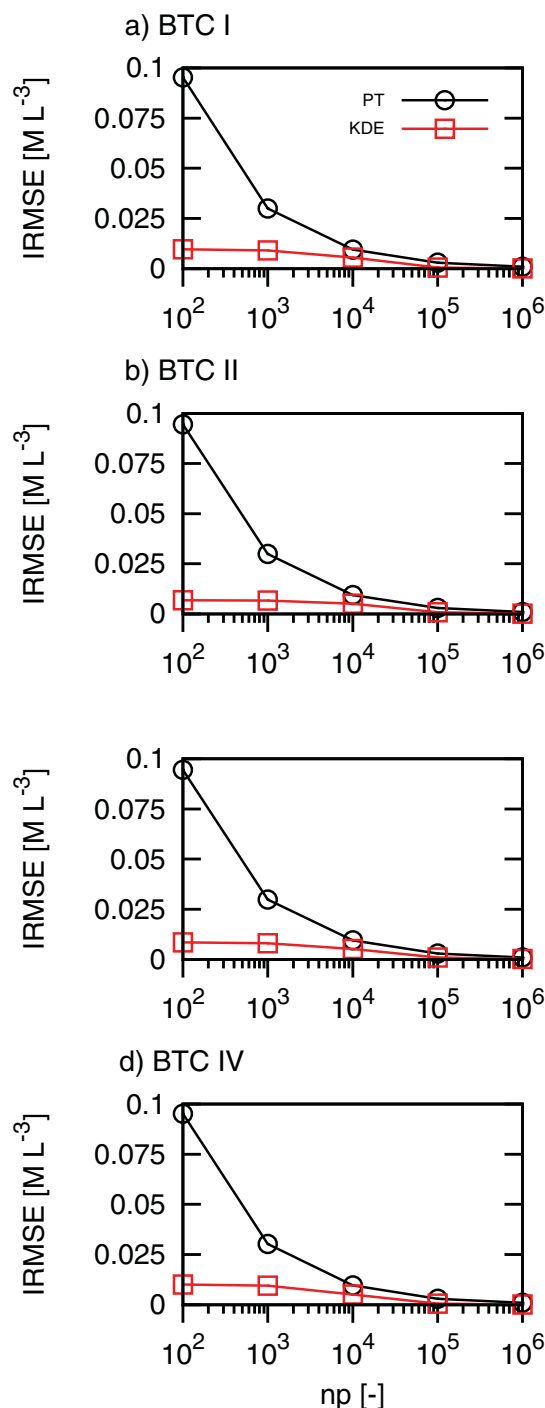


Figure 4. Integrated root mean square error for PT and KDE curves (see key in a) for representative BTCs I-IV (a-d, respectively), and as function of np .

different histogram bins or particle resolutions in the curve. An inset of this histogram range from 0 to 196 particles in the BTC per time step is shown in Figure 3d to exemplify the differences in mass breakthrough for the different np scenarios. The peak mass in the PT BTC I, $PT_{pk}^{np} [M]$, is equivalent to the maximum number of crossed particles at any given time step, and is marked with a black star on each curve. The corresponding percent error of the peak with the high-resolution curve is equivalent to $(PT_{pk}^{np} - PT_{pk}^{np=100M}) / PT_{pk}^{np=100M} * 100$, and is also noted for each curve using BTC I as an example. The low histogram resolution of

coarse. Recall that total mass in the system is constant regardless of np in these simulations; thus when np is varied over 4 orders of magnitude, so is the resolution of the particle tracking breakthrough curve. To explain further, one can relate the resolution of a particle to a bin size of a histogram. Figure 3a shows that for a total mass, MT , of 300 $[M]$, the mass of each individual particle or “bin” of the histogram is equivalent to MT/np . As the number of bins or particles increases, the resolution of BTC also increases. The first important point here is that regardless of np in these numerical experiments, MT is constant and conserved (as denoted by a constant area under each curve in Figure 3a). Second, as np increases, the accuracy of transport behavior is improved because parcels of mass are effectively split to resolve more locations and ultimately breakthrough times.

An alternative thought experiment is to consider the number of particles that could cross the plane at a given time. This number ranges from zero particles (i.e., zero mass) to np particles (i.e., all particles cross at the same time). Figure 3b shows the number of particles crossing a plane and the corresponding mass breakthrough as a function of np . Another way to think about the concept of the x axis in this graph, (np per BTC time step) is as the resolution of the breakthrough curve; if this number is very low, the histogram of the particle tracking breakthrough curve is coarse, and transport may not be accurately resolved. Figure 3c shows the frequency of the histogram resolution for representative BTC I as an illustrative example. Here it is evident that if np is too low, the histogram resolution of the particle tracking mass is very coarse, resulting in a large bias in the BTC results (Figure 2, first column). In this example, $np = 10^2$ results in only two particle resolutions; similarly $np = 10^3$ results in only three particle resolutions. A true histogram shape begins to evolve as np increases past 10^4 , with the $np = 10^6$ curve displaying a variance of 196

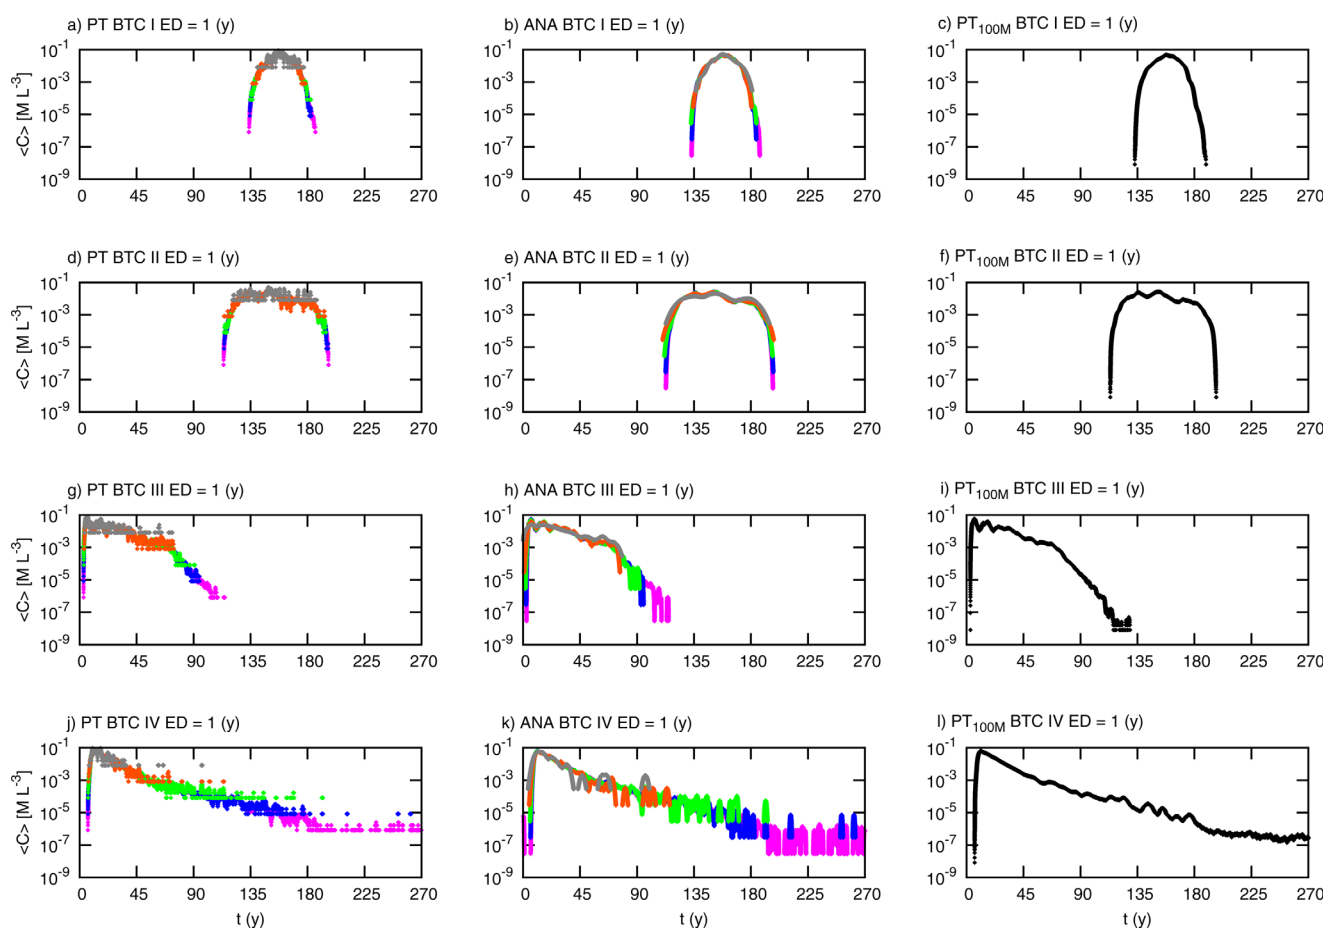


Figure 5. Example $\langle C \rangle$ for $ED = 1$ (y) for representative BTCs I–IV (rows) and for different i (columns). Variable np is shown as colors (see key in Figure 2a).

the $np = 10^2$ simulations translates to up to 1200% error whereas the $np = 10^6$ simulations translate to only 19% error. This exemplifies how underestimation of the particle tracking histogram resolution, while still conserving total mass may grossly overestimate a transport solution. It also shows how a sufficient number of particles may accurately capture the overall transport behavior and how increasing np results in a nonlinear error reduction.

KDE reconstructed BTCs (equation (2)) are shown in Figure 2 (middle column). For even low np , KDE BTCs collapse into the same C versus t space when compared to the PT counterpart (Figure 2, left column), which vary over a wide C parameter space. Furthermore, the KDE BTCs show excellent visual agreement with the PT_{100M} BTC (Figure 2, right column), even for more complex shaped BTCs such as representative BTCs III and IV. These results are consistent with past findings showing the strength of optimal KDEs for BTC reconstruction [Pedretti and Fernández-García, 2013]. Remarkably, for low np values, the use of the optimal KDE reduces the large bias in C shown in the PT approach, sometimes by as much as 2 orders of magnitude in C . That is, the smaller the np value, the better the improvement obtained using KDE. Note here that to aide in the best visual comparison of the two methods, points in the KDE breakthrough curve are plotted *only* at times where PT breakthrough data were obtained. As mentioned earlier, technically the KDE yields a continuous BTC, so that mass may be distributed in breakthrough times that do not correspond to the times the raw PT breakthrough data passed the plane. Specifically, KDE mass may be distributed to times neighboring the raw PT breakthrough data following the Gaussian kernel centered at each PT data point within a span of h . Thus, only a subset of the mass distribution is visually apparent in KDE BTCs (Figure 2, middle column). Integration of the continuous KDE BTCs is nonetheless equivalent to the integration of the other PT graphs in Figure 2.




Figure 6. Example $\langle C \rangle$ for $ED = 70$ (y) for representative BTCs I–IV (rows) and for different i (columns). Variable np is shown as colors (see key in Figure 2a).

Figure 4 shows $IRMSE_{KDE}^{np}$ (red lines) and $IRMSE_{PT}^{np}$ (black lines) for representative BTCs I–IV (Figures 4a–4d, respectively). For even very low np , the global error between the KDE and PT_{100M} BTCs is very low. In contrast, the PT global error is much higher, especially for simulations with $np < 10^4$. $IRMSE_{PT}^{np}$ does not approach $IRMSE_{KDE}^{np}$ until approximately $np = 10^4 - 10^5$; in other words, the global error in an optimal KDE BTC with only $np = 10^2$ has the same global error as a PT BTC with $np = 10^4 - 10^5$. This suggests that if unaveraged BTCs were to be used in any type of analysis, a substantial reduction in C versus t bias would be achieved in using the KDE. Alternatively, several orders of magnitude less np are required to obtain the same error in global BTC shape. This is true for all representative BTCs analyzed here.

5.2. Analysis of Averaged BTCs, $\langle C \rangle$ Versus t

Following the methodology in sections 2 and 3, averaging of representative BTCs (Figure 2) is conducted. Two methods are used, direct reconstruction from equation (14) and reconstruction using the analytical kernel, equation (10), with optimal support, equation (3). Figures 5 and 6 show $\langle C \rangle$ -averaged BTCs for the end-member ED s, $ED = 1$ (y) and $ED = 70$ (y), respectively. Averaged BTCs for $ED = 10$ (y) are not shown here for brevity. Representative BTCs I–IV are shown as rows, different i are shown in the first two columns (PT, ANA), and the last column shows the PT_{100M} averaged BTC. A few general trends are evident. First, while both i are visually similar to the shape of the PT_{100M} BTC (and this similarity increases as np increases), each poses a different characteristic shape after the averaging process. The PT averaged-BTCs are most discretized and noisy, especially in the tail of the BTC (e.g., Figure 5j).

Visually, a much better fit of the PT_{100M} averaged-BTCs is obtained with the ANA method (Figures 5 and 6, second column). This is especially true for smaller ED s, such as $ED = 1$ (y), Figure 5. With ANA averaged-

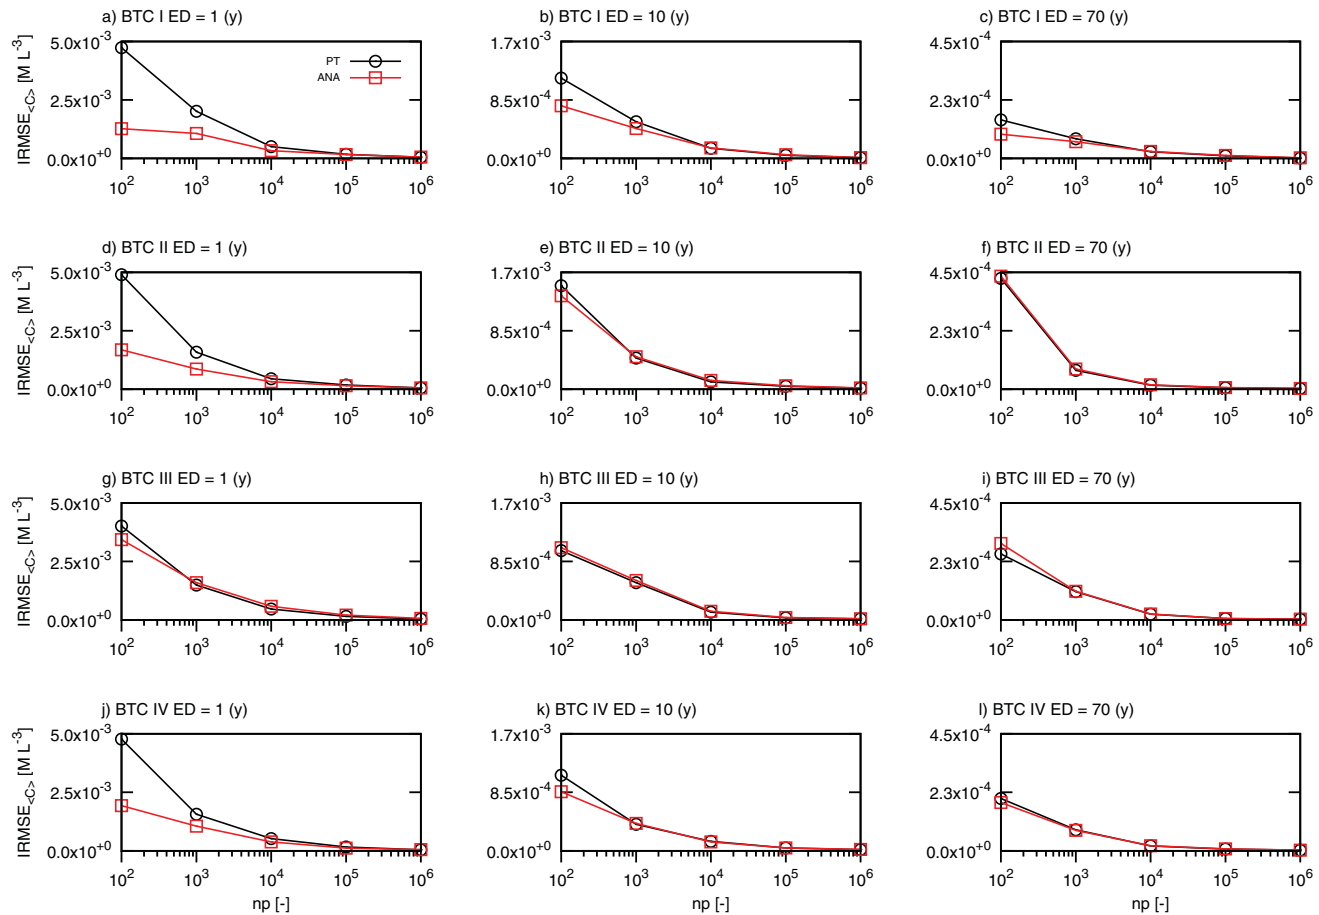


Figure 7. IRMSE of averaged, $\langle C \rangle$ vs. t BTCs as a function of np (see key in a). Representative BTCs are shown in rows, ED scenarios are shown in columns. Note that the y axis decreases as ED increases.

BTCs, the peak of the curve is visually reproduced much better than the PT averaged-BTC as the curve does not spuriously oscillate as much. Furthermore, the ANA solution also visually reproduces local peaks in the BTC at the correct magnitude and breakthrough times. For example, an obvious triple local peak exists in the PT_{100M} averaged representative BTC II; even with only $np = 10^2$ the ANA averaging solution visually reproduces these peaks (Figure 5e). As shown in Figure 6, the importance of these local peaks diminishes as the ED increases. The trends visually evident from examination of the averaged-BTCs are quantified below.

Figure 7 shows $IRMSE_{ANA\langle C \rangle}^{np}$ (red lines) and $IRMSE_{PT\langle C \rangle}^{np}$ (black lines) for representative BTCs I–IV (rows) and for $ED = 1, 10$, and 70 (y) (columns). Note that the y axis scale decreases as ED increases. For small ED (left column) and for simple BTC shapes (rows 1 and 2), $IRMSE_{ANA\langle C \rangle}^{np}$ is lower than $IRMSE_{PT\langle C \rangle}^{np}$, meaning that there is a benefit in the use of the new optimal kernel formulation, equation (10), especially for low np . The point varies at which np is sufficiently large so that the reconstruction method does not impact results. For most representative BTCs, and for most ED s, a benefit exists in using the ANA solution when $np \leq 10^4$. Only in a few cases are different np threshold values found; for example, in BTC IV, $ED = 10$ (y) (Figure 7k) the benefit of using the ANA approach is only clear for $np \leq 10^2$. An exception exist for BTC III at large ED s (Figure 7j), where the ANA method underperforms the PT one at $np = 10^2$. In any case, we point out that at these large ED values the $IRMSE_i^{np}$ is very small in comparison to that of smaller ED s, thus making these differences relatively insignificant.

Figure 8 shows the percent error in the peak, or $p = 100$ th percentile, of averaged-BTCs, $PE_i^{np\langle C \rangle}(100)$, as calculated by equation (16). $PE_{ANA\langle C \rangle}^{np}(100)$ are shown in red lines and $PE_{PT\langle C \rangle}^{np}(100)$ are shown in black lines for representative BTCs I–IV (rows) and for $ED = 1, 10$, and 70 (y) (columns). A negative percent error indicates an underestimation of the peak C , whereas a positive percent error indicates an overestimation.

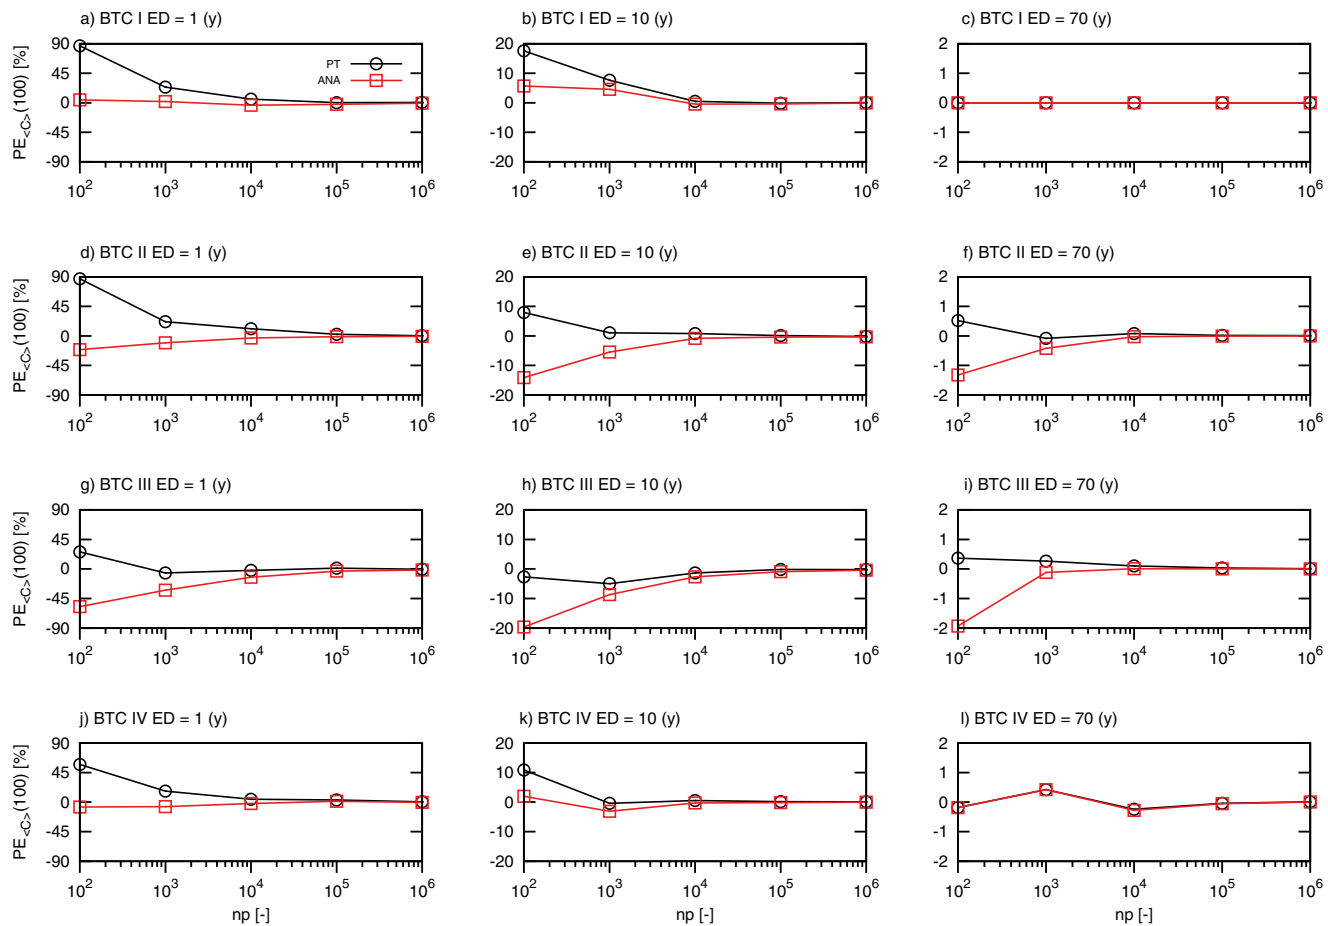


Figure 8. PE of averaged, $\langle C \rangle$ vs. t BTCs as a function of np (see key in a) for the 100th percentile ($p = 100$). Representative BTCs are shown in rows, ED scenarios are shown in columns. Note that the y axis decreases as ED increases.

For all $np \geq 10^5$ scenarios, $PE_{ANA \langle C \rangle}^{np}(100)$ and $PE_{PT \langle C \rangle}^{np}(100)$ are approximately equivalent with zero error. For scenarios with $np < 10^5$, the percent errors vary between the two scenarios, and in general vary more drastically when the ED is small (e.g., Figure 8, left column; note that the y axis decreases as the ED increases). For these scenarios ($np < 10^5$ and $ED < 10$ y), the ANA solution provides a much better fit as compared to the PT_{100M} peak C , considered as the true solution. The ANA solutions of representative BTCs I and IV yield an almost identical match with the PT_{100M} peak C , even for low np , where remarkably $PE_{ANA \langle C \rangle}^{np}(100)$ is nearly zero (Figures 8a and 8j). The ANA solution to BTC II also provides a relatively good fit with a low PE even at low np where again the ANA error is much smaller than the PT solution, which in this case is substantially large, at approximately 90% (Figure 8d).

Interestingly, the PT solution always overpredicts the peak C , whereas the ANA solution always underpredicts the peak C , albeit with a lower magnitude of error than the PT solution. The reason is that the use of kernels greatly reduces any bias in individual breakthrough points. Thus, while the PT solution tends to overpredict PT_{100M} peak C due to the noisy (discretized) nature of the original PT BTC, the ANA solution tends to slightly underpredict due to the smoothness imposed by the method.

Lastly, for the scenarios where percent errors between the PT and ANA scenarios differ ($np < 10^5$ and ED 1 and 10 y), the ANA solution results in less error than the PT solution. An exception exists for representative BTC III (all ED s) and very slightly for BTC II (only $ED = 10$ y), as discussed below. The exception to this trend is examined in Figure 9 for representative BTC III. Shown here is a close inspection of the averaged BTC III during the peak, and for only the first 40 years. PT, ANA, and PT_{100M} scenarios are compared (see key in c) for $np = 10^2 - 10^5$ (rows) and for $ED = 1, 10$, and 70 (y) (columns). From this close inspection, the

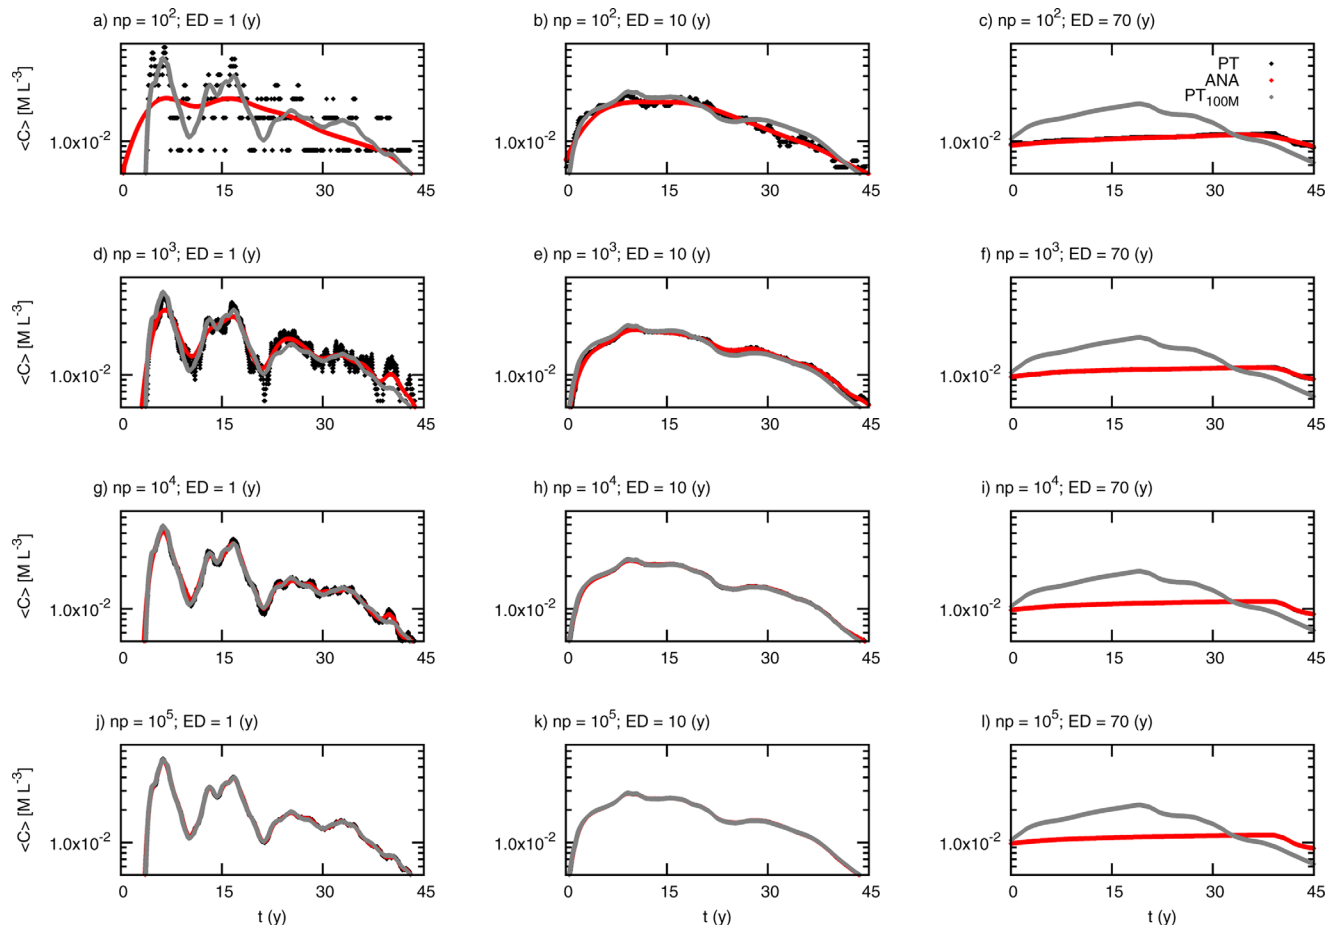


Figure 9. Averaged $\langle C \rangle$ vs. t for representative BTC III, $np = 102 - 105$ (rows) and for $ED = 1-70$ (y) (columns). PT, ANA, and PT100M scenarios are compared for close inspection of the early, peak breakthrough times before 45 years (see key in c).

underestimation of the PT_{100M} peak (gray line) due to the ANA solution and overestimation of the PT solution are obvious.

Furthermore, it is visually apparent why the ANA solution is underperforming in terms of matching the peak concentration with this type of curve reconstruction. As shown more evidently when np is low and for small ED (Figure 9a), when multiple oscillations in the concentration are present during a short time frame (note that concentration is in log-space in Figure 9), the ANA solution, which recall is based on the arrival times from the PT solution, reproduces a curve somewhere in between the two extremes of the highly discretized and noisy BTC from the PT solution. Once more particles are simulated, this effect is lessened, where the ANA solution reproduces a curve much more similar to that of the PT_{100M} curve. Nevertheless, there is still an underprediction of the peak due to the effect of adjacent peaks in the BTC. Thus, while the ANA solution provides an excellent fit to the PT_{100M} curve for most scenarios, caution must be taken in this approach when large oscillations in the BTC are temporally adjacent.

5.3. Analysis of Optimal h Values Relative to Total BTC and Exposure Durations

Figure 10 shows the solution to the dimensionless ratio $F = ED/h_{opt}$ (-) as solved by expression (12). Results are shown as a function of np for scenarios where the dimensionless ratio of $a = T/ED$ is equivalent to 1, 10, 100, and 1000 (-), gray lines. Recall that values of a typically range between 5 and 500, depending on the choice of ED , where the larger the a value, the larger the simulation period considered for a given ED . Also recall that as shown in section 5.2, typically risk analyses do not require more than $np = 10^4$ (-) to achieve accurate results. Figure 10 shows that for $a > 10$ (-) and $np < 10^4$ (-) the ratio F is smaller than 1, a typical limit for F in practical situations since $F > 1$ may result in a severe underestimation of the optimal h associated with the average concentration. The results of F from the numerical simulations presented here, in




Figure 10. The dimensionless ratio, $F = ED/h_{opt}$ derived from equation (12) as a function of np and for different a values (gray lines), where the larger the a value, the larger the simulation period considered for a given ED . Results from numerical simulations in this study are shown as points (symbols differentiate representative BTC I-IV and colors differentiate the simulated ED).

addition to simulations with an $ED = 30$ y, are shown as points in Figure 10, where colors denote the ED and symbols denote representative BTCs I-IV. Within the parameter space $a > 10$ (-) and $np < 10^4$ (-), F values from our numerical simulations are very small, well below $F = 1$ (-). In fact, all numerical results for F from our simulations are below $F = 1$ (-). From Figure 1, this implies that $\bar{K}(t; h, ED)$ is similar (but not equal) to $K(t; h)$. A result that explains why, in most situations, the optimal solution of h obtained during the reconstruction of BTCs based on the KDE method, equation (3), is still approximately valid for risk analysis.

6. Conclusions

In this study, how a breakthrough curve is represented and the effects on risk-driven analysis are analyzed. Particle tracking (PT) techniques, often considered favorable over Eulerian techniques due to their absence of numerical dispersion, are evaluated in terms of optimal breakthrough curve reconstruction as a function of the number of particles (np) used. While in reality the true solution can only be achieved as np approaches infinity, recent work has shown that given a relatively few number of particles, PT methods can yield reliably well-constructed breakthrough curves with the use of optimal kernel density estimators (KDEs), showing improvement in not only simulation time but also accuracy. This advancement shows promise in many fields such as human health risk analysis where breakthrough curves are often used. However, quantifying exposure relies on not simply breakthrough curves but those averaged over a given exposure duration (or the actual time an individual is in contact with the hazardous substance). We derive a new averaging method incorporating the optimal h from the KDE method (equation (10), referred to the ANA method) and compare the results with traditional PT breakthrough curves as a function of different exposure durations.

This study utilizes representative breakthrough curves of typical behavior in groundwater transport, simulated as a function of np ($10^2 - 10^6$) and averaged as a function of the exposure duration (1, 10, and 70 y). A number of metrics (both on averaged and unaveraged breakthrough curves) are used to quantify the error between different methods, including a comparison for accuracy with reference to a high resolution PT simulation of the same flow field, simulated with $np = 10^8$ (i.e., considered the “true” solution).

For unaveraged breakthrough curves, traditional PT results are compared to a KDE method. Unaveraged PT breakthrough curves show a large bias in breakthrough concentration (as high as 2 orders of magnitude).

Table A1. Flow and Transport Parameters

Parameter	Value	Units
Domain Size (x, y, z)	$1002 \times 501 \times 99$	(m)
Cell discretization ($\Delta x, \Delta y, \Delta z$)	$3.0 \times 3.0 \times 0.3$	(m)
Number of cells (n_x, n_y, n_z)	$334 \times 167 \times 330$	
Centered location of source (x, y, z)	250.0, 250.0, 50.0	(m)
Porosity	$\theta = 0.3$	
Change in head and gradient	$\Delta h = 2.5$ $J = 0.0025$	(m)
Time step	$dt = 1.0$	(d)

This bias appears especially large when compared to the optimal KDE results, which yield very similar breakthrough as the high resolution simulation, remarkably even when np is as low as 10^2 . The same holds for all representative breakthrough curves of varying shapes and complexities. For all representative cases, the global error of unaveraged KDE break-

through curves with only $np = 10^2$ has the same global error as PT breakthrough curves with $np = 10^4 - 10^5$. This suggests that (1) for a fixed np a substantial reduction in bias would be achieved by using the optimal KDE for unaveraged breakthrough curve analysis and (2) several orders of magnitude less np are required to obtain the same error in global breakthrough curve shape.

For exposure duration averaged breakthrough curves necessary in risk analysis, traditional PT results are compared with the newly derived ANA method that incorporates a KDE (equation (10)). Averaged breakthrough curves utilizing the ANA method best fit the high-resolution simulation, especially for smaller exposure durations. Given a large np , the choice in breakthrough curve reconstruction becomes less sensitive; the point at which np is sufficiently large varies depending mainly on the duration of exposure, but typically holds for simulations with $np \leq 10^4 - 10^5$. Percent error in the peak of averaged breakthrough curves is especially of interest in a risk assessment framework. For all $np \geq 10^5$ scenarios, ANA and PT percent error are approximately equivalent with zero error. For scenarios with $np < 10^5$, the percent errors vary between the two scenarios, and in general vary more drastically when the exposure duration is small. For these scenarios, the ANA solution yields an almost identical match with the high-resolution simulations, even for low np , with remarkably nearly zero percent error. In contrast, the PT error can be substantial, as high as 90%. An exception exists for certain breakthrough curve shapes when multiple oscillations in the concentration are present during a short time frame.

We have presented and validated the use of different optimal KDEs in breakthrough curve reconstruction when curves are averaged and unaveraged as a function of both the exposure duration and np . In addition to the presentation of a new, low-error prone breakthrough curve averaging solution applicable in a risk assessment framework, we illustrate that obtaining a representative average exposure concentration is heavily reliant on an accurate representation of the breakthrough curve, and should be considered in environmental analysis, especially when data is scarce. Although these results are presented in a groundwater example, these results are applicable to signal processing of any media.

Appendix A

To obtain the representative BTCs, saturated aquifer flow is simulated using the parallel, three-dimensional groundwater model ParFlow [Ashby and Falgout, 1996; Jones and Woodward, 2001; Kollet and Maxwell, 2006]. Two constant head boundaries at $x = 0.0$ (m) and $x = 1002$ (m) drive a groundwater gradient, while all other boundaries are no-flow. The 3-D model is discretized into $334 \times 167 \times 330$ (1.84×10^7) cells. A 10.0 (m) cubic source of np particles is instantaneously released, advected without the inclusion of local dispersion or diffusion, then tracked as a function of time at a down gradient plane perpendicular to the mean flow direction (at $x = 750.0$ m). All the values used in the simulations for the different parameters are displayed in Table A1.

Table A2. Representative BTC Heterogeneity Parameters

BTC	K_G (m d ⁻¹)	σ^2 (-)	λ_{3D} λ_{2D} λ_{1D} (m)
Curve I	1.0	0.25	3.0, 3.0, 0.3
Curve II	1.0	0.25	15.0, 15.0, 1.5
Curve III	1.0	4.0	120.0, 120.0, 1.5
Curve IV	5.0	9.0	15.0, 15.0, 1.5

Heterogeneity of $\ln(K)$ is simulated using the turning bands algorithm [Tompson et al., 1989] which enforces a semivariogram function through rotation of 250 one-dimensional lines through space, and where each value in the random field is a weighted average of values contained within each band. An exponential model is used to define spatial correlation of

$\ln(K)$. Parameters for representative BTCs a-d are listed in Table A2, where K_G (m d^{-1}) is the geometric mean of K , σ^2 (–) is the variance $\ln(K)$, and where λ_x , λ_y , and λ_z (m) are the directional correlation lengths of K . Representative BTCs I–III follow a Gaussian distribution whereas BTC IV follows a truncated Gaussian distribution with lower and upper cutoff values of 0.01 (m d^{-1}) and 35.0 (m d^{-1}).

Acknowledgments

The authors acknowledge the financial support provided by the Spanish Ministry of Science and Innovation, projects SCARCE Consolider-Ingenio 2010 (reference CSD2009-00065) and FEAR (CGL2012-38120), by the EU (project MARSOL, FP7-ENV-2013, grant 619120), and by the ICREA Academia Program. We would also like to thank Wolfgang Nowak and two anonymous reviewers for their constructive comments on this paper. Data can be obtained by contacting the authors.

References

- Ahlstrom, S. W., H. P. Foote, R. C. Arnett, C. R. Cole, and R. J. Serne (1977), Multi-component mass transport model: Theory and numerical implementation (discrete parcel random walk version), *Rep. BNWL-2127*, Battelle Pac. Northwest Lab., Richland, Wash.
- Andrićević, R., and V. Cvetković (1996), Evaluation of risk from contaminants migrating by groundwater, *Water Resour. Res.*, 32(3), 611–621.
- Ashby, S. F., and R. D. Falgout (1996), A parallel multigrid preconditioned conjugate gradient algorithm for groundwater flow simulations, *Nucl. Sci. Eng.*, 124(1), 145–159.
- Boso, F., A. Bellin, and M. Dumbser (2013), Numerical simulations of solute transport in highly heterogeneous formations: A comparison of alternative numerical schemes, *Adv. Water Resour.*, 52, 178–189.
- de Barros, F. P. J., and A. Fiori (2014), First-order based cumulative distribution function for solute concentration in heterogeneous aquifers: Theoretical analysis and implications for human health risk assessment, *Water Resour. Res.*, 50, 4018–4037, doi:10.1002/2013WR015024.
- De Josselin de Jong, G. (1958), Longitudinal and transverse diffusion in granular deposits, *Trans. AGU*, 39(1), 67–74, doi:10.1029/TR039i001p00067.
- Engel, J., E. Herrmann, and T. Gasser (1994), An iterative bandwidth selector for kernel estimation of densities and their derivatives, *J. Non-parametr. Statist.*, 4, 21–34.
- Enzenhoefer, R., W. Nowak, and R. Helmig (2012), Probabilistic exposure risk assessment with advective-dispersive well vulnerability criteria, *Adv. Water Resour.*, 36, 121–132.
- Enzenhoefer, R., T. Bunk, and W. Nowak (2014), Nine steps to risk-informed wellhead protection and management: A case study, *Ground Water*, 52(S1), 161–174, doi:10.1111/gwat.12161.
- Fernández-García, D., and X. Sánchez-Vila (2011), Optimal reconstruction of concentrations, gradients and reaction rates from particle distributions, *J. Contam. Hydrol.*, 120–121, 99–114.
- Hathhorn, E., and R. Charbeneau (1994), Stochastic Fluid Travel Times in Heterogeneous Porous Media, *J. Hydraul. Eng.*, 120(2), 134–146.
- Härdle, W. (1990), Smoothing techniques with implementation in S, Springer-Verlag, N. Y.
- Henri, C. V., and D. Fernández-García (2014), Toward efficiency in heterogeneous multispecies reactive transport modeling: A particle-tracking solution for first-order network reactions, *Water Resour. Res.*, 50, 7206–7230, doi:10.1002/2013WR014956.
- Herrera, P., M. Massabó, and R. Beckie (2009), A meshless method to simulate solute transport in heterogeneous porous media, *Adv. Water Resour.*, 32, 413–429.
- Jones, J. E., and C. S. Woodward (2001), Newton-Krylov-multigrid solvers for large-scale, highly heterogeneous, variably saturated flow problems, *Adv. Water Resour.*, 24(7), 763–774.
- Kollet, S. J., and R. M. Maxwell (2006), Integrated surface-groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model, *Adv. Water Resour.*, 29(7), 945–958.
- Liu, M. B., G. R. Liu, and K. Y. Lam (2003), Constructing smoothing functions in smoothed particle hydrodynamics with applications, *J. Comput. Appl. Math.*, 155(2), 263–284, doi:10.1016/S0377-0427(02)00869-5.
- Maxwell, R., S. Carle, and A. Tompson (2008), Contamination, risk, and heterogeneity: On the effectiveness of aquifer remediation, *Environ. Geol.*, 54(8), 1771–1786.
- Maxwell, R. M. (2010), *SLIM-FAST: A User's Manual*, vol. 4, Int. Ground Water Model. Cent., Golden, Colo.
- Maxwell, R. M., and W. E. Kastenberg (1999), Stochastic environmental risk analysis: An integrated methodology for predicting cancer risk from contaminated groundwater, *Stochastic Environ. Res. Risk Assess.*, 13(1–2), 27–47.
- Maxwell, R. M., S. D. Pelmulder, A. F. B. Tompson, and W. E. Kastenberg (1998), On the development of a new methodology for groundwater-driven health risk assessment, *Water Resour. Res.*, 34(4), 833–847.
- Morris, J. P., P. J. Fox, and Y. Zhu (1997), Modeling low Reynolds number incompressible flows using SPH, *J. Comput. Phys.*, 36(1), 214–226.
- Pedretti, D., and D. Fernández-García (2013), An automatic locally-adaptive method to estimate heavily-tailed breakthrough curves from particle distributions, *Adv. Water Resour.*, 59, 52–65.
- Prickett, T. A., T. G. Naymik, and C. G. Longquist (1981), *A Random Walk Solute Transport Model for Selected Groundwater Quality Evaluations*, p. 103, Ill. State Water Surv., Champaign, Ill.
- Salamon, P., D. Fernández-García, and J. J. Gómez-Hernández (2006), A review and numerical assessment of the random walk particle tracking method, *J. Contam. Hydrol.*, 87, 277–305.
- Scheidegger, A. E. (1954), Statistical hydrodynamics in porous media, *J. Geophys. Res.*, 66(10), 3273–3278.
- Siirila, E. R., and R. M. Maxwell (2012a), A new perspective on human health risk assessment: Development of a time dependent methodology and the effect of varying exposure durations, *Sci. Total Environ.*, 431, 221–232.
- Siirila, E. R., and R. M. Maxwell (2012b), Evaluating effective reaction rates of kinetically driven solutes in large-scale, statistically anisotropic media: Human health risk implications, *Water Resour. Res.*, 48, W04527, doi:10.1029/2011WR011516.
- Siirila-Woodburn, E. R., X. Sánchez-Vila, and D. Fernández-García (2015), On the formation of multiple local peaks in breakthrough curves, *Water Resour. Res.*, 51, 2128–2152, doi:10.1002/2014WR015840.
- Silverman, B. W. (1986), Density estimation for statistics and data analysis, Chapman & Hall, CRC.
- Tompson, A. F. B. (1993), Numerical simulation of chemical migration in physically and chemically heterogeneous porous media, *Water Resour. Res.*, 29(11), 3709–3726.
- Tompson, A. F. B., R. Ababou, and L. W. Gelhar (1989), Implementation of the three-dimensional turning bands random field generator, *Water Resour. Res.*, 25(10), 2227–2243.
- Tompson, A. F. B., A. L. Schafer, and R. W. Smith (1996), Impacts of physical and chemical heterogeneity on cocontaminant transport in a sandy porous medium, *Water Resour. Res.*, 32(4), 801–818.
- U.S. EPA (1989), *Risk Assessment Guidance for Superfund, vol. I, Human Health Evaluation Manual (Part A)*, Washington, D. C.
- U.S. EPA (2001), *Risk Assessment Guidance for Superfund, vol. III, Process for Conducting Probabilistic Risk Assessment (Part A)*, Washington, D. C.
- U.S. EPA (2008), *Child-Specific Exposure Factors Handbook*, Rep. EPA/600/R-06/096F, Washington, D. C.
- WHO (2005), *Principles of Characterizing and Applying Human Exposure Models, Harmonization Proj. Doc. 3*, Geneva.