Multilevel Monte Carlo predictions of first passage times in threedimensional discrete fracture networks: a graphbased approach

Water Resources Research

TECHNICAL REPORTS: METHODS

10.1029/2019WR026493

Multilevel Monte Carlo Predictions of First Passage Times in Three-Dimensional Discrete Fracture Networks: A Graph-Based Approach

S. Berrone[1,2], **J. D. Hyman**[3], **and S. Pieraccini**[2,4]

[1]Dipartimento di Scienze Matematiche, Politecnico di Torino, Turin, Italy, [2]Member of the INdAM Research Group GNCS, INdAM, Rome, Italy, [3]Computational Earth Science (EES-16), Earth and Environmental Sciences Division, Los Alamos National Laboratory, Los Alamos, NM, USA, [4]Dipartimento di Ingegneria Meccanica e Aerospaziale, Politecnico di Torino, Turin, Italy

Key Points:

- Multilevel Monte Carlo is used to predict first passage times in flows through discrete fracture networks
- The method provides estimates of the mean and variance at several orders of magnitude lower computational cost than standard Monte Carlo
- We identify a hierarchy of subnetworks based on the shortest topological paths through the network using a graph-based method

Supporting Information:

- Supporting Information S1

Correspondence to:
J. D. Hyman,
jhyman@lanl.gov

Citation:
Berrone, S., Hyman, J. D., & Pieraccini, S. (2020). Multilevel Monte Carlo predictions of first passage times in three-dimensional discrete fracture networks: A graph-based approach. *Water Resources Research*, *56*, e2019WR026493. https://doi.org/10.1029/2019WR026493

Received 7 OCT 2019
Accepted 2 MAY 2020
Accepted article online 13 MAY 2020

Abstract We present a method combining multilevel Monte Carlo (MLMC) and a graph-based primary subnetwork identification algorithm to provide estimates of the mean and variance of the distribution of first passage times in fracture media at significantly lower computational cost than standard Monte Carlo (MC) methods. Simulations of solute transport are performed using a discrete fracture network (DFN), and instead of using various grid resolutions for levels in the MLMC, which is standard practice in MLMC, we identify a hierarchy of subnetworks in the DFN based on the shortest topological paths through the network using a graph-based method. While the mean of these ensembles is of critical importance, the variance is also essential in fractured media where uncertainty is an overarching theme, and understanding variability across an ensemble is a requirement for safety assessments. The method provides good estimates of the mean and variance at two orders of magnitude lower computational cost than MC.

1. Introduction

Estimating the first arrival time of a chemical species transported along with groundwater flow through fractured media at a control plane or well location is a common task in subsurface hydrology. However, the high degree of uncertainty associated with flow and transport through fractured media in the subsurface limits our ability to do so accurately and in turn impacts our predictive capabilities in a variety of civil and industrial endeavors including the long-term storage of spent nuclear fuel (Follin et al., 2014; Selroos et al., 2002), geo-sequestration of CO_2 (Jenkins et al., 2015), hydrocarbon extraction (Hyman et al., 2016), environmental restoration of contaminated fractured media (VanderKwaak & Sudicky, 1996), and aquifer storage and management (Kueper & McWhorter, 1991). These uncertainties are in part due to the cost of obtaining data to constrain our models as well as the scale/time dependence of observations (Carrera et al., 2005). The expansive range of relevant length scales in fractured media exacerbate these issues (Bonnet et al., 2001; Neuman, 2005).

Discrete fracture networks (DFN) (Cacas et al., 1990; Dershowitz & Einstein, 1988; Fidelibus et al., 2009) are a computational tool to model flow and transport through fractured media (cf. Berre et al. 2018, for a recent overview). The cost associated with DFN models is substantially higher than conventional continuum models, but the trade off is that DFN models can typically represent a wider range of transport phenomena with higher fidelity due to the inclusion of detailed features of the fracture network (Hadgu et al., 2017; Hyman et al., 2019a; Painter et al., 2002; Sweeney et al., 2019). The inclusion of these features introduces additional layers of uncertainty because more parameters have to be calibrated and/or sampled due to uncertainty in fracture location, size, orientation, and hydrological properties. In turn, these additional features both increase the cost of a DFN model run and the number of them required for reliable estimates of a quantity of interest (QoI). To reduce the computational cost associated with high-fidelity three-dimensional DFN simulations, modelers have developed methods to identify network backbones, which are subnetworks where the majority of flow and transport occurs (Abelin et al., 1991; Rasmuson & Neretnieks, 1986). Which fractures are in the backbone depends on the macro-scale structure of the network, for example, orientations and density (Hyman et al., 2019b); meso-scale hydrological attributes, for example, fracture permeability (De Dreuzy et al., 2002); and the flow direction relative to the background stress field (Kang et al., 2019). Once a backbone has been identified, the cost of performing a flow and transport simulation thereon is

reduced, and little loss of accuracy in the prediction of first arrival times is incurred (Viswanathan et al., 2018). These methods can be loosely broken up into two classes: The first of which is geometrical based (fractures smaller than a threshold size, Wellman et al., 2009, or dead-end fractures, Berkowitz & Scher, 1998, are removed), and the other is topologically based (sets of connected fractures are combined to identify potential backbones using a graph-based representation to characterize network connectivity Aldrich et al., 2017; Hyman et al., 2018; Osthus et al., 2020).

Literature lacks non-intrusive, reliable, and efficient methods capable to tackle uncertainty quantification analyses on such complex frameworks. Classical Monte Carlo (MC) is the baseline method to assess uncertainties in subsurface flow and transport scenarios and is the most common tool in use today (cf. Ballio & Guadagnini, 2004; Berrone et al., 2015; Gelhar, 1986; Guadagnini & Neuman, 1999a, 1999b; Tartakovsky, 2007; Zhang, 2001). In standard MC, an ensemble of possible scenarios/realizations is generated, and then estimates of a particular QoI, such as first arrival times, are determined. Standard MC can suffer from slow convergence if the system heterogeneity is quite large, as is the case within the context of fractured media, and a large ensemble may be required for MC integrals to converge to within an acceptable tolerance. Moreover, most applications of MC focus solely on the first moment of the distribution of a QoI, the mean, and higher-order moments, specifically the variance which is a critical value for uncertainty quantification in fractured media, are not estimated at all. There are currently no methods available to overcome the computational burdens associated with producing an ensemble of large-scale DFN model runs that are required for standard MC thus rendering its application to be relatively impracticable. However, there are a number of variants that outperform standard MC, namely, multifidelity MC (Canuto et al., 2019; Ng & Willcox, 2014; O'Malley et al., 2018; Peherstorfer et al., 2016) and multilevel MC (Berrone et al., 2018; Giles, 2015; Lu et al., 2016), that could be modified to overcome these issues. The general idea of these variants is to combine estimates with different levels of accuracy where the quality of the estimate depends on the computational cost; low cost simulations/models produce more data but with lower accuracy. So long as the estimations of the QoI at these various levels/fidelity are systematically related, one can combine them to provide better estimates than standard MC at an overall lower computational cost.

We propose a methodology that combines multilevel Monte Carlo (MLMC) with graph-based backbone identification procedures to address the current limitations of using DFN simulations for estimations first two moments of the distribution of first arrival times in fracture media. The method is distinguished from available methods in that a hierarchy of nested subnetworks, identified using a graph-based method, comprise the levels for use in MLMC rather than a hierarchy of model mesh resolutions. For an exposition of the method, we create an ensemble of semi-generic fracture networks and estimate the mean and variance of the distribution of first arrival times of a passive tracer transported through the networks. The method is significantly more computationally affordable than standard MC while providing estimates of similar accuracy. Thus, the method overcomes the aforementioned issues associated with using DFN simulations for uncertainty quantification of first passage times in fractured media and, in particular, addresses the issue of computational burden by combining graph-based techniques and MLMC.

2. Methodology

2.1. DFN Simulations

We construct an ensemble of semi-generic DFNs and simulate transport through the networks using the DFNWORKS computational suite (Hyman et al., 2015). The networks are semi-generic meaning that they are not meant to represent any particular field site, but the parameters used in their construction are loosely based on field observations. The networks are composed of circular fractures whose orientations are uniformly random and centers are uniformly distributed throughout the domain which is a cube with sides measuring 15 m. Fracture radii r [m] are sampled from a truncated power law distribution with exponent α and lower and upper cutoffs $(r_0; r_u)$,

$$r = r_0 \left[1 - u + u \left(\frac{r_0}{r_u} \right)^{\alpha} \right]^{-1/\alpha}, \tag{1}$$

where u is a random number sampled from the continuous uniform distribution on the closed interval [0,1]. We select a value of $\alpha = 2.6$, lower cut off $r_0 = 1$ m, and upper cut off $r_u = 5$ m. Isolated fractures and clusters that do not connect the inflow and outflow boundaries are removed from the domain because they

do not contribute to flow. The fracture networks contain around 500 fractures each. The connected network density (De Dreuzy et al., 2012) is about 10 times the critical percolation value (Bour & Davy, 1997). Thus, the networks are dense enough that there are multiple structural paths in the network between the inflow and outflow boundaries. The feature-rejection algorithm for meshing (FRAM) (Hyman et al., 2014) is used to generate a Delaunay triangulation of each DFN.

Variability in hydraulic properties is included into the network by correlating the hydraulic aperture (b) of a fracture to its radius using a positively correlated power law relationship $b = 5.0\times10^{-5}r^{1/2}$ (De Dreuzy et al., 2002; Hyman et al., 2016). Given the range of radii we consider, 1 m to 5 m, the apertures range from $5\cdot10^{-5}$ m to $1\cdot10^{-4}$ m. We do not include in-fracture aperture variability in these simulations, as was done in De Dreuzy et al. (2012), Frampton et al. (2019), and Makedonska et al. (2016). Fracture permeability is related to the aperture using the cubic law $k = b^2/12$ and is thus also positively correlated to the fracture radii. Given the limited range of fracture radii, less than one order of magnitude, the heterogeneity provided by these hydraulic properties is also small, and heterogeneity in flow field results primarily from the network structure.

Flow within the network is modeled using Reynolds equation (Zimmerman & Bodvarsson, 1996) which is numerically integrated using the parallelized subsurface flow and reactive transport code PFLOTRAN (Lichtner et al., 2015) to determine the pressure and volumetric flow rates throughout the entire network, which can be used to reconstruct the Eulerian velocity field $\mathbf{u}(\mathbf{x})$ within the DFN (Makedonska et al., 2015; Painter et al., 2012). A pressure difference of 1 MPa is prescribed along the x-axis in each of the networks. All other boundary faces are assigned a no-flow boundary condition.

We represent the spreading of a nonreactive conservative solute through each DFN by a cloud of passive tracer particles. We do not allow particles to interact with the matrix for simplicity (Hyman et al., 2019). Complete mixing is used to determine what direction particles exit out of fracture intersections (Berkowitz et al., 1994; Kang et al., 2015; Park et al., 2001, 2003; Swherman et al., 2018; Stockman et al., 1997). The imposed pressure gradient is aligned with the x-axis, and thus the primary direction of flow is in the x direction. Particle initial positions $\mathbf{a}$ are uniformly distributed along fracture intersections with the inlet plane $\mathbf{x}_0 = (0, y, z)$, that is, using resident injection (Hyman et al., 2015; Kreft & Zuber, 1978). The trajectory $\mathbf{x}(t;\mathbf{a})$ of a particle starting at $\mathbf{a}$ at time $t = 0$ is given by the advection equation

$$\frac{d\mathbf{x}(t;\mathbf{a})}{dt} = \mathbf{v}(t;\mathbf{a}), \qquad \mathbf{x}(0;\mathbf{a}) = \mathbf{a}, \tag{2}$$

where the Lagrangian velocity $\mathbf{v}(t;\mathbf{x})$ is given in terms of the Eulerian velocity $\mathbf{u}(\mathbf{x})$. The travel time $\tau(\mathbf{x}_1;\mathbf{a})$ of a particle to exit the domain via the outlet plane $\mathbf{x}_1 = (15, y, z)$ is

$$\tau(\mathbf{x}_1;\mathbf{a}) = \inf\{t | x(t;\mathbf{a}) > 15\} \quad . \tag{3}$$

The first passage time of all particles through a network realization ω is given by

$$Q(\omega) = \inf_{\{\mathbf{a}\}}\{\tau(\mathbf{x}_1;\mathbf{a})\} \ , \tag{4}$$

which is random variable that represents our *quantity of interest* (QoI).

2.2. MLMC

While referring the reader to the supporting information for a detailed description of MLMC, we present the fundamental elements below. From the ensemble of networks we can obtain a set of first passage times. Standard MC estimates the mean of this set, denoted $\mathbb{E}[Q]$, using the sample mean of N estimations

$$\bar{Q} = N^{-1}\sum_{i=1}^{N} Q(\omega_i) \quad . \tag{5}$$

When obtaining $Q(\omega_i)$ involves the numerical solution of a model problem, the accuracy of the solution must also be considered. For example, different mesh resolutions can result in different levels of accuracy in the numerical solution. Letting Q_h denote a numerical approximation of Q computed with a given accuracy, then the sample mean is given by

$$\bar{Q}_h = N^{-1}\sum_{i=1}^{N} Q_h(\omega_i) \quad . \tag{6}$$

The principal idea behind MLMC is to combine estimates of $\bar{Q}_h$ obtained with different levels of accuracy. The general stochastic behavior of the QoI is captured using a large number of samples obtained with a relatively inaccurate estimator that is computationally inexpensive. This estimates are refined using a few more accurate samples that are obtained with more costly estimators. In the framework of numerical solutions of partial differential equations, these levels may correspond to different meshes resolutions, referred to as *geometric* MLMC (see Berrone et al., 2018, for an example of use in the framework of fracture networks), or they may correspond to different models (see O'Malley et al., 2018, for an example of this methodology applied to fracture networks).

Letting $\ell = 0, 1, \ldots, L$ denote the level of accuracy, with the convention that $\ell = 0$ corresponds to the lowest accuracy, and $Q^\ell(\omega_i)$ denote a set of N samples of Q computed at level ℓ, the multilevel estimator for the mean value of Q is defined

$$m_1^{\mathrm{ML}} := \bar{Q}^0_{N_0} + \sum_{\ell=1}^{L} (\bar{Q}^\ell_{N_\ell} - \bar{Q}^{\ell-1}_{N_\ell}), \tag{7}$$

where

$$\bar{Q}^\ell_{N_\ell} = N_\ell^{-1} \sum_{i=1}^{N_\ell} Q^\ell(\omega_i), \tag{8}$$

for all $\ell = 0, \ldots, L$. While the first moment of this distribution is of critical importance, the second moment of the distribution, the variance $\sigma^2[Q]$, provides information about the variability across the ensemble and is a measure of the reliability of the estimates provided by the simulations. Bierig and Chernov (2015) extended the MLMC to compute $\sigma^2[Q]$ using

$$m_2^{\mathrm{ML}} := \sum_{\ell=0}^{L} (h_2(Q^\ell_{N_\ell}) - h_2(Q^{\ell-1}_{N_\ell})), \tag{9}$$

where

$$h_2(Q_N) = \frac{NS_2 - S_1^2}{N(N-1)}, \tag{10}$$

and $S_\alpha = \sum_{i=1}^{N} (Q(\omega_i))^\alpha$ and $h_2(Q^{-1}_{N_0}) = 0$. Note that $h_2(Q_N)$ is an unbiased estimator of the variance of Q. Recently, Pisaroni et al. (2017) generalized MLMC to compute the pth order central moments.

The estimators are applied with an optimal number of samples at each level (N_ℓ) chosen to reach a mean square error (MSE) smaller than a given accuracy ε^2. The effectiveness of MLMC is related to the rate of growth (with respect to ℓ) of the computational cost C_ℓ for the computation of one sample, relative to the rate of decrease of the numerical bias $|\mu_p(Q^\ell) - \mu_p(Q)|$ and of the variance of the correction added at each level where p denotes the moment of the distribution; $p = 1$ is the mean and $p = 2$ is the variance. In the case of the mean, for example, assuming that $|\mu_1(Q^\ell) - \mu_1(Q)| \leq c_\alpha 2^{-\alpha\ell}$, $\sigma^2[Q^\ell - Q^{\ell-1}] \leq c_\beta 2^{-\beta\ell}$, and $C_\ell \leq c_\gamma 2^{\gamma\ell}$ for suitable constants $\alpha, \beta, \gamma, c_\alpha, c_\beta, c_\gamma$, if one has $2\alpha \geq \min(\beta, \gamma)$, then the overall computational cost C needed for reaching the ε accuracy is bounded by

$$C \leq c \begin{cases} \varepsilon^{-2}(\log \varepsilon)^2 & \text{if} \quad \beta = \gamma \\ \varepsilon^{-\left(2 + \frac{\gamma - \beta}{\alpha}\right)} & \text{if} \quad \beta < \gamma \\ \varepsilon^{-2} & \text{if} \quad \beta > \gamma \end{cases}. \tag{11}$$

The target accuracy ε is closely related to the minimum value of L needed to achieve it.

2.3. Multilevel Graph Representations of a DFN

Field and laboratory experiments, as well as numerical simulations, of flow and transport through fractured media have revealed the existence of primary subnetworks where the majority of flow and transport occurs and the fastest transport passes through. In this section, we describe a methodology to identify a hierarchy of these primary subnetworks that will be used as levels in the MLMC estimator.

We use a graph-based approach to represent the structure of the networks. At the core of the DFN methodology is the conceptual model of a set of fractures, which are discrete entities, intersecting with one another

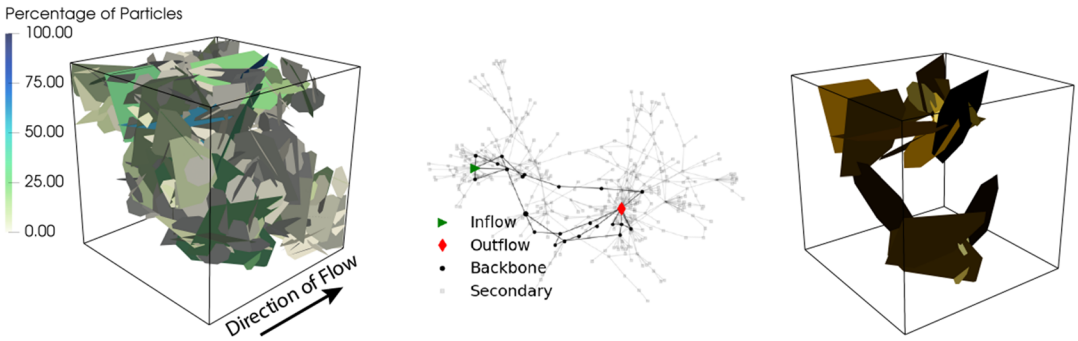

Figure 1. (Left) DFN composed of 380 fractures colored by the percentage of particles passing through each fracture. (Middle) Graph representation of the DFN shown to the left. The backbone corresponding to the 10 shortest paths is colored black. (Right) DFN backbone corresponding to the 10 shortest paths in the graph shown in the middle.

to form a network. The mathematical construct of a graph, a set of vertices connected by a set of edges, is a formal way to represent this system. We use a graph representation where vertices correspond to fractures in the DFN and edges indicate that two fractures intersect (Huseby et al., 1997; Hyman et al., 2018). This graph is augmented to include vertices corresponding to the inflow and outflow boundaries by adding source and target vertices. If a fracture intersects the inlet plane the corresponding vertex is connected to the source vertex, and if a fracture intersects the outlet plane the corresponding vertex is connected to the target vertex. The construction of this graph is purely topological; geometric and hydrological properties are not accounted for as vertex and edges weights.

We identify subgraphs corresponding to the union of the edges in n shortest paths from the source to target using the method presented by Hyman et al. (2017). The n shortest paths is a generalization of the shortest path to include n total paths (possibly overlapping) in order of nondecreasing length starting from the shortest path. We consider only simple paths (loopless) from the source to the target. The edge lengths are assigned unity, so the shortest paths are those with the fewest number of edges between the source and the target. The preimage of this subgraph, which is its equivalent fracture subnetwork, has the fewest number of intersections and thus connected fractures, spanning the inflow and outflow boundaries.

Figure 1 presents a visual explanation of our workflow for identifying network backbones. The DFN shown on the left is composed of 380 fractures, which are colored by the percentage of particles passing through each fracture. The middle sub-figure shows the graph representation of the DFN. The black nodes correspond to fractures in the backbone, here identified using the 10 shortest paths, and the gray nodes are the secondary structure, fractures not in the 10 shortest paths. The inflow node is a green triangle and the outflow node a red diamond. The sub-figure on the right is the corresponding backbone of the 10 shortest paths in the graph shown in the middle. Note that fractures in the backbone closely match those with larger percentages of particles in the left sub-figure even though no particle attributes are used to identify the backbone. The subnetwork contains 23 of the original 380 fractures, and while it makes up less than 20% of the original network by surface area, over 45% of the total transport occurs on this subnetwork. While the number of particles passing through a fracture is not directly related to the travel time of particles, it does indicate that the flow is highly channelized and that there are primary subnetworks. In particular, particles traverse along larger fractures, which have lower resistance to flow due to the positive correlation between fracture size and aperture and thus where the fastest transport occurs due to both low structural and hydraulic resistance.

Table 1 reports the mean number of fractures in the subnetworks along with the number of degrees of freedom at each level. We consider four values of $n = 1, 5, 10, 20$. We also consider the 2-core of the graph, which is an upper bound on the union of loopless paths from source to target. The n-core of a graph is the maximal subgraph that contains vertices of degree n or greater (Seidman, 1983). In terms of MLMC, we thus have five levels: $\ell = 0$: shortest path ($n = 1$), $\ell = 1 \rightarrow 3$: $n = 5$, 10, and 20 and $\ell = 4$: 2-core. We generate 10,000 samples at $\ell = 0$; 4,000 at $\ell = 1$; 1,000 at $\ell = 2$; 1,000 at $\ell = 3$; and 100 at $\ell = 4$. The lowest level $\ell = 0$ (the shortest paths) is only about 1% of the network by number of fractures (N_f) and degrees of freedom (dofs). In all other levels, except $\ell = 4$ (the 2-core of the network), the number of dofs is less than 10% of

Table 1
Computational Cost at MLMC Levels—Mean Values Over All Realizations

Level	N_f	N_f %	dofs	dofs %	N_ℓ - $\mathbb{E}$	N_ℓ - σ^2
$\ell = 0$	4.96	1.12	5,029.47	1.26	2,961	8,622
$\ell = 1$	12.27	2.78	14,810.21	3.72	1,255	885
$\ell = 2$	17.72	4.01	21,982.93	5.52	326	688
$\ell = 3$	25.77	5.83	32,265.56	8.11	35	134
$\ell = 4$	249.06	56.33	270,874.91	68.05	1	2
Full network	442.15	100.00	398024.75	100.00	—	—

Note. N_f, number of fractures; N_f%, percentage of entire network; dofs, number of degrees of freedom in linear system for pressure; dofs %, percentage of entire network; N_ℓ, number of samples required at each level for the mean $\mathbb{E}$ and variance σ^2.

the entire network. The 2-core, in contrast, is over 50% of the network. Thus, simulations preformed on the subnetworks are significantly cheaper computationally than those performed using the entire network.

These reductions in computational cost are balanced by the accuracy in predicting the first arrival time; as n increases, so does the accuracy for a number of reasons. Foremost, the fastest particle might not transverse the shortest topological path, and thus the smaller values of n may provide a poor estimate of first arrival time. As the size of the subnetwork grows, the likelihood of the fastest path being included in the subnetwork increases and the predictions improve. The fractures that make up these subnetworks tend to be large, which leads to higher aperture in the adopted ensemble setup and have orientations aligned with the primary direction of flow on which the fastest transport tends to occur. A detailed discussion of the relationship between the size of the subnetwork and predictions of first arrival times is provided in Hyman et al. (2017).

The proposed methodology is broken up into three steps: (1) For a given DFN, we identify a hierarchy of primary subnetworks using the graph-based methods outlined above. (2) Each subnetwork is meshed, steady-state flow determined, and the first arrival time of particles through each subnetwork is obtained. (3) We combine the values obtained across the hierarchy of subnetworks in the MLMC framework to obtain a multilevel estimation of the first moments of the QoI. Note that the proposed method does not depend on the particular graph-based method to identify the subnetworks, only that the accuracy in predicting the first arrival times increases with the MLMC level, which here corresponds to different subnetworks identified using the n shortest paths and is a requirement verified in the following section. Given the low degree of hydraulic heterogeneity in the networks, this purely topological characterization is adequate for identifying where the fastest transport occurs in this ensemble. However, in networks with larger hydraulic heterogeneity, different graph representations that account for geometric and hydraulic attributes could provide better estimations.

3. MLMC Results

In this section we report the results of using MLMC to estimate the mean and variance of first arrival times across an ensemble of fracture networks. Figure 2 shows the numerical bias $|\Delta_\ell h_p|$, and standard error $V_{\ell,p}$ for the estimations of the mean (left) and variance (right) at each level, the definition of these measurements are in the supporting information. As the level increases, these values decreases, thereby confirming that the selected hierarchy of subnetworks is appropriate for MLMC estimations. Next, we need to determine the number of samples required at each level. To do so we first estimate the parameters α, β, γ that are mentioned in section 2.2 by fitting the data plotted in Figure 2 with regression values for α and β

$$|\Delta_\ell h_p| \simeq c_\alpha 2^{-\alpha\ell} \qquad V_{\ell,p} \simeq c_\beta 2^{-\beta\ell} \quad . \tag{12}$$

For the data considered here, values of $\alpha = 1.86$ and $\beta = 1.52$ for the mean and $\alpha = 1.51$ and $\beta = 2.25$ for the variance are obtained. Next, we estimate the minimum value of L needed for achieving a fixed accuracy ε (see supporting information for the precise relationship between L and ε). We consider a relative accuracy ε_r such that

$$\varepsilon = \varepsilon_r \bar{Q}, \tag{13}$$

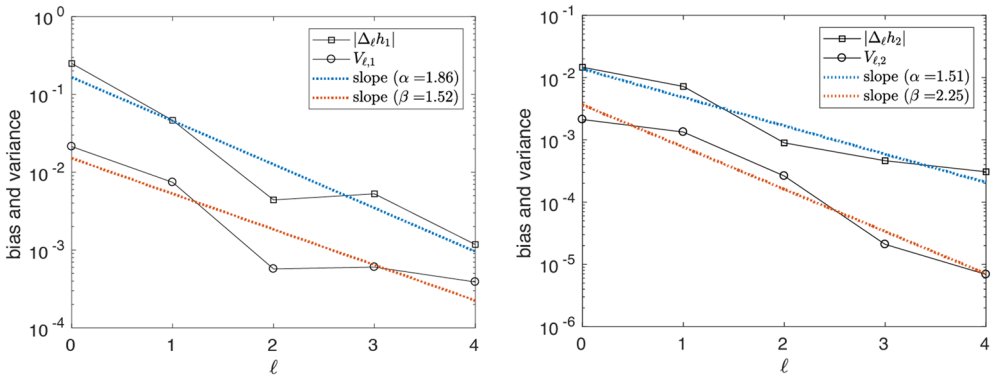

Figure 2. Numerical bias and standard error, along with the regression values of α and β. Left: mean value; right: variance. As the level increases, these values decreases, thereby confirming that the selected hierarchy of subnetworks is appropriate for MLMC estimations.

where $\bar{Q}$ is a rough estimate of $\mu_p(Q)$, which can be obtained using a few samples at level L and standard MC. For the results presented here, we consider a value of $\varepsilon_r = 0.05$ for the mean and $\varepsilon_r = 0.25$ for the variance. At these values, the minimum L value needed for reaching the corresponding accuracy is 3 for $p = 1$ and 4 for $p = 2$. The number of samples needed at each level can be computed starting from an estimate of $V_{\ell,p}$ and the computational cost at each level C_ℓ. The largest part of C_ℓ is solving the linear system for pressure, which we approximate as the number of dofs$^{\log_2(7)}$. Using these values, we determine that grows exponentially with a rate parameter of $\gamma = 2.99$. The values obtained for N_ℓ are reported in Table 1, and the formula used is detailed in the supporting information. Notice that there is an inverse relationship between the number of dofs relative to the full network and the number of samples required at each level which allows the MLMC method to provide more efficient estimates of the QoI than using a single level alone. For example, while the 2-core would provide very good estimations of the QoI at a reduced computational cost relative to the full networks it would still require a large number of samples to obtain the desired accuracy if used independently of the other levels. In contrast, only one or two samples are needed in the MLMC framework thereby drastically reducing the computational cost. Also, the cost of computing the variance is much larger than the mean, and we are limited by the number of samples generated at the lowest level, even though we considered a larger value of ε_r.

Figure 3 provides a comparison of the estimations obtained using MLMC (blue) and standard MC applied to the full network (orange), which is reported here for reference. The ordinate is the value of the first arrival time, normalized by the median value of the entire particle plumes computed using the full networks. The abscissa is the relative computational cost, which is normalized with respect to the cost for a single simulation at $\ell = 0$. The left sub-figure reports estimates of the mean, and the right reports estimates of the variance. The partial sums in (9) that are obtained at levels $\ell = 0, \ldots, 4$ are shown in blue. The reported values are the mean of 10 MLMC runs (blue circle), each of which is a random sampling of N_ℓ at each level. Note that it is possible that a single MLMC run can produce an estimate for the mean or variance that is rather large. Thus, in practice, one constructs a small set of MLMC runs and takes the average value of that ensemble in a manner similar to standard MC but computed over the MLMC runs. By definition the average of different MLMC estimations will naturally converge to the value being estimated because the MLMC estimators are unbiased. However, this convergence is obtained at much lower cost than standard MC, which is one strength of the MLMC method. In this case we found that 10 MLMC runs to be sufficient, but under other scenarios, the number of members required in the ensemble may vary. The value obtained with standard MC, the first arrival times obtained in the entire fracture networks, is plotted as an orange square, and a dashed orange line is included for comparison with those estimates obtained using MLMC. We use 180 samples of the full network for the standard MC estimate of the mean, which corresponds to the expected number of samples needed to achieve the considered accuracy ε; for the variance we use 400 samples. The insets show the mean square error defined in the supporting information, plotted as a function of the number of levels used and confirms that the observed behavior is in accordance with that predicted by the theory. In the case of the estimation of the mean value, as the number of levels increases, the mean square error decreases and the MLMC estimate converges to that obtained using standard MC rapidly. Note

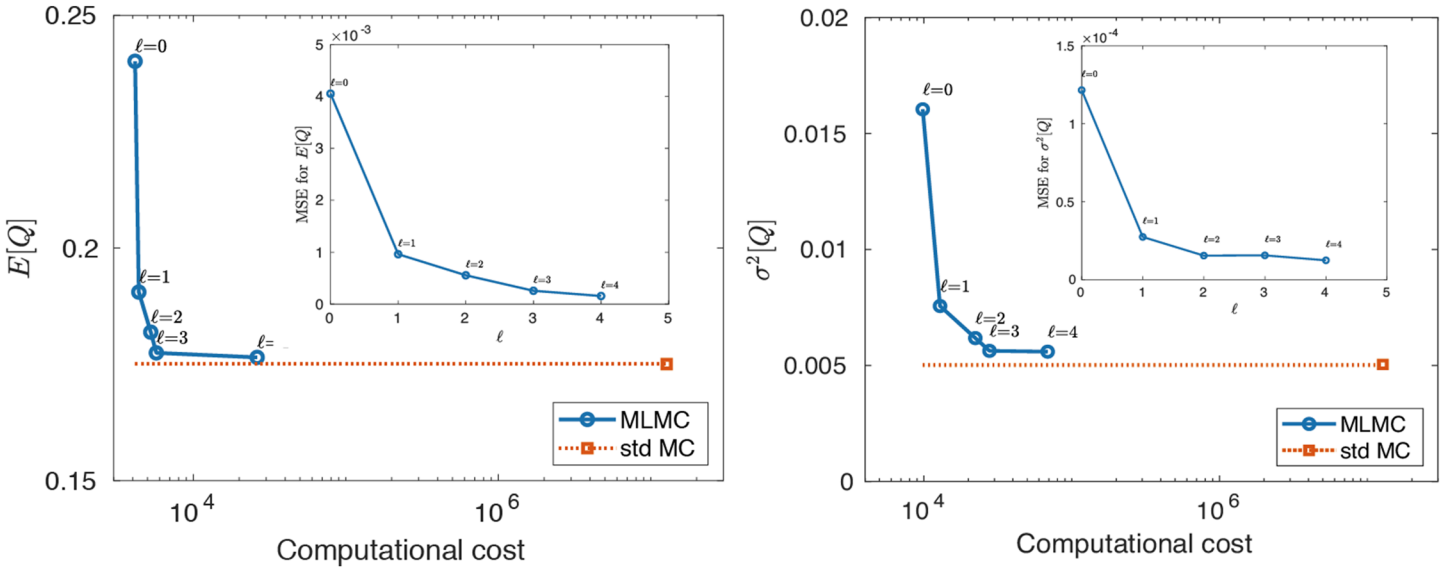

Figure 3. Comparison of the estimations obtained using MLMC (blue) and standard Monte Carlo (orange) plotted as a function of computational cost (left: mean; right: variance). The insets show the decrease of the MSE plotted as a function of the number of levels used. As the number of levels increases, the mean square error decreases, and the mean values converges to that obtained using standard Monte Carlo rapidly.

that the value reached by the MSE at $\ell = 4$ is approximately $1.7 \cdot 10^{-4}$, which is close to the target accuracy ε^2 of $7.6 \cdot 10^{-5}$.

Using information up to $\ell = 3$ appears sufficient to obtain a good estimation of the mean, and the inclusion of the highest level only slightly improves the estimate; the relative error from the MC estimate is less than 1% when all levels are used. These results indicate the lower levels adequately capture stochastic variation in the QoI, and the few samples required at higher levels merely refine the estimation. The MLMC estimate of the mean is obtained at over two orders of magnitude lower computational expense than standard MC. In the case of the variance, the rate of convergence is slower, but the approximation of the variance is correctly caught nonetheless. When passing from $\ell = 1$ to $\ell = 2$ to $\ell = 3$ there is a significant improvement in the approximation with only moderate additional computational cost. Again, slight improvement is obtained by including $\ell = 4$; the relative distance is ≈10%. Recall that we considered a larger value of ϵ_r, and thus the relative error is larger. These observations underscore the previous conclusion that the lower levels adequately capture stochastic variation in the QoI and can provide estimations of not just the mean but the variance as well at reduced computational expense.

4. Discussion

We have presented a method that combines MLMC and a graph-based primary subnetwork identification algorithm to provide estimates of the mean and variance of the distribution of first passage times in fracture media at significantly lower computational cost than standard MC, which is typically computational prohibitive for use in large-scale uncertainty quantification. Simulations of solute transport are performed using a DFN, and instead of using various grid resolutions, which is common practice in MLMC, we identify a hierarchy of subnetworks for levels in the MLMC. The method provides a good estimation of the mean and variance at two orders of magnitude lower computational cost than MC. While the mean of these ensembles is of critical importance, the variance is also essential in fractured media where uncertainty is an overarching theme and understanding variability across an ensemble is a requirement for safety assessments. This study reports the first use of MLMC to approximate the variance of an ensemble in fractured media in a multi-model setting.

The adopted graph representation only accounts for the topology of the network, which is a limitation of the proposed version of the method. However, this is not a fundamental limitation as other graph representations could also be used to construct the hierarchy of nested subnetworks. Additional geometric and hydraulic features could also be accounted for using node and edge weights, as has been done in other

graph-based methods. While the chosen method works well for the ensemble of networks considered here, there could be other ensembles where more advanced graph representations would be warranted and could provide better estimations for the first arrival times. Additionally, if a different QoI was considered, for example, peak arrival or tailing behavior of a solute plume, then a different graph representation could be more appropriate. Regardless of these further considerations, the proposed framework to combine graph-based methods with MLMC provided a methodology to address a wide range of scenarios.

Most MLMC methods define levels using different refinement levels of the computational mesh. We presented an alternative methodology where the levels correspond to sub-domains where the relevant physics occurs. Such a methodology is not limited to DFN simulations and, in principal, is applicable to disciplines outside of subsurface hydrology. So long as sub-domains can be ranked in terms of their influence on a quantity of interest, the modification of MLMC presented here can be tailored to use those sub-domains as levels in a MLMC work flow and provide approximations of the quantity of interest at lower computational cost than MC.

Acknowledgments
J. D. H. is thankful for support from the U.S. Department of Energy through the Los Alamos National Laboratory Laboratory-Directed Research and Development Program Grants 20180621ECR and 20170103DR. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). S. B. acknowledges the support of the Italian MIUR Award "Dipartimento di Eccellenza 2018–2022" (E11G18000350001) to the Department of Mathematical Sciences, Politecnico di Torino. S. P. acknowledges the support of the Italian MIUR PRIN Project 201752HKH8_003. The software used for the simulations, DFNWORKS, is open source and can be obtained online (https://github.com/lanl/dfnWorks). This is a theoretical work for which no data need be made available LA-UR-19-29755.

References

Abelin, H., Birgersson, L., Moreno, L., Widén, H., Ågren, T., & Neretnieks, I. (1991). A large-scale flow and tracer experiment in granite: 2. Results and interpretation. *Water Resources Research*, *27*(12), 3119–3135.

Aldrich, G., Hyman, J. D., Karra, S., Gable, C. W., Makedonska, N., Viswanathan, H., et al. (2017). Analysis and visualization of discrete fracture networks using a flow topology graph. *IEEE Transactions on Visualization and Computer Graphics*, *23*(8), 1896–1909. https://doi.org/10.1109/tvcg.2016.2582174

Ballio, F., & Guadagnini, A. (2004). Convergence assessment of numerical Monte Carlo simulations in groundwater hydrology. *Water Resources Research*, *40*, W04603. https://doi.org/10.1029/2003WR002876

Berkowitz, B., Naumann, C., & Smith, L. (1994). Mass transfer at fracture intersections: An evaluation of mixing models. *Water Resources Research*, *30*(6), 1765–1773. https://doi.org/10.1029/94WR00432

Berkowitz, B., & Scher, H. (1998). Theory of anomalous chemical transport in random fracture networks. *Physical Review E*, *57*, 5858–5869. https://doi.org/10.1103/PhysRevE.57.5858

Berre, I., Doster, F., & Keilegavlen, E. (2018). Flow in fractured porous media: A review of conceptual models and discretization approaches. *Transport Porous Media*, *130*, 215–236.

Berrone, S., Canuto, C., Pieraccini, S., & Scialò, S. (2015). Uncertainty quantification in discrete fracture network models: Stochastic fracture transmissivity. *Computers and Mathematics with Applications*, *70*(4), 603–623.

Berrone, S., Canuto, C., Pieraccini, S., & Scialò, S. (2018). Uncertainty quantification in discrete fracture network models: Stochastic geometry. *Water Resources Research*, *54*, 1338–1352. https://doi.org/10.1002/2017WR021163

Bierig, C., & Chernov, A. (2015). Convergence analysis of multilevel Monte Carlo variance estimators and application for random obstacle problems. *Numerische Mathematik*, *130*, 579–613.

Bonnet, E., Bour, O., Odling, N. E., Davy, P., Main, I., Cowie, P., & Berkowitz, B. (2001). Scaling of fracture systems in geological media. *Reviews of Geophysics*, *39*(3), 347–383.

Bour, O., & Davy, P. (1997). Connectivity of random fault networks following a power law fault length distribution. *Water Resources Research*, *33*(7), 1567–1583.

Cacas, M. C., Ledoux, E., De Marsily, G., Barbreau, A., Calmels, P., Gaillard, B., & Margritta, R. (1990). Modeling fracture flow with a stochastic discrete fracture network: Calibration and validation: 2. The transport model. *Water Resources Research*, *26*(3), 491–500.

Canuto, C., Pieraccini, S., & Xiu, D. (2019). Uncertainty quantification of discontinuous outputs via a non-intrusive bifidelity strategy. *Journal of Computational Physics*, *398*, 108885. https://doi.org/10.1016/j.jcp.2019.108885

Carrera, J., Alcolea, A., Medina, A., Hidalgo, J., & Slooten, L. J. (2005). Inverse problem in hydrogeology. *Hydrogeology Journal*, *13*(1), 206–222.

De Dreuzy, J.-R., Davy, P., & Bour, O. (2002). Hydraulic properties of two-dimensional random fracture networks following power law distributions of length and aperture. *Water Resources Research*, *38*(12), 2065–2078.

De Dreuzy, J.-R., Méheust, Y., & Pichot, G. (2012). Influence of fracture scale heterogeneity on the flow properties of three-dimensional discrete fracture networks (dfn). *Journal of Geophysical Research*, *117*, B11207. https://doi.org/10.1029/2012JB009461

Dershowitz, W. S., & Einstein, H. H. (1988). Characterizing rock joint geometry with joint system models. *Rock Mechanics and Rock Engineering*, *1*, 21–51.

Fidelibus, C., Cammarata, G., & Cravero, M. (2009). Hydraulic characterization of fractured rocks. In M. Abbie, J. S. bedford (eds.), *Rock Mechanics: New research.* New York: Nova Science Publishers Inc.

Follin, S., Hartley, L., Rhén, I., Jackson, P., Joyce, S., Roberts, D., & Swift, B. (2014). A methodology to constrain the parameters of a hydrogeological discrete fracture network model for sparsely fractured crystalline rock, exemplified by data from the proposed high-level nuclear waste repository site at Forsmark, Sweden. *Hydrogeology Journal*, *22*(2), 313–331.

Frampton, A., Hyman, J. D., & Zou, L. (2019). Advective transport in discrete fracture networks with connected and disconnected textures representing internal aperture variability. *Water Resources Research*, *55*(7), 5487–5501. https://doi.org/10.1029/2018WR024322

Gelhar, L. W. (1986). Stochastic subsurface hydrology from theory to applications. *Water Resources Research*, *22*(9S), 135S–145S.

Giles, M. B. (2015). Multilevel Monte Carlo methods. *Acta Numerica*, *24*, 259–328. https://doi.org/10.1017/S096249291500001X

Guadagnini, A., & Neuman, S. P. (1999a). Nonlocal and localized analyses of conditional mean steady state flow in bounded, randomly nonuniform domains: 1. Theory and computational approach. *Water Resources Research*, *35*(10), 2999–3018.

Guadagnini, A., & Neuman, S. P. (1999b). Nonlocal and localized analyses of conditional mean steady state flow in bounded, randomly nonuniform domains: 2. Computational examples. *Water Resources Research*, *35*(10), 3019–3039.

Hadgu, T., Karra, S., Kalinina, E., Makedonska, N., Hyman, J. D., Klise, K., et al. (2017). A comparative study of discrete fracture network and equivalent continuum models for simulating flow and transport in the far field of a hypothetical nuclear waste repository in crystalline host rock. *Journal of Hydrology*, *553*, 59–70.

Huseby, O., Thovert, J., & Adler, P. (1997). Geometry and topology of fracture systems. *Journal of Physics A: Mathematical and Theoretical*, *30*(5), 1415.

Hyman, J. D., Aldrich, G., Viswanathan, H., Makedonska, N., & Karra, S. (2016). Fracture size and transmissivity correlations: Implications for transport simulations in sparse three-dimensional discrete fracture networks following a truncated power law distribution of fracture size. *Water Resources Research*, *52*, 6472–6489.

Hyman, J. D., Dentz, M., Hagberg, A., & Kang, P. (2019a). Emergence of stable laws for first passage times in three-dimensional random fracture networks. *Physical Review Letters*, *123*(24), 248,501.

Hyman, J. D., Dentz, M., Hagberg, A., & Kang, P. (2019b). Linking structural and transport properties in three-dimensional fracture networks. *Journal of Geophysical Research: Solid Earth*, *124*, 1185–1204. https://doi.org/10.1029/2018JB016553

Hyman, J. D., Gable, C. W., Painter, S. L., & Makedonska, N. (2014). Conforming Delaunay triangulation of stochastically generated three dimensional discrete fracture networks: A feature rejection algorithm for meshing strategy. *SIAM Journal of Scientific Computing*, *36*(4), A1871–A1894.

Hyman, J. D., Hagberg, A., Osthus, D., Srinivasan, S., Viswanathan, H., & Srinivasan, G. (2018). Identifying backbones in three-dimensional discrete fracture networks: A bipartite graph-based approach. SIAM Multiscale Modeling and Simulation.

Hyman, J. D., Hagberg, A., Srinivasan, G., Mohd-Yusof, J., & Viswanathan, H. (2017). Predictions of first passage times in sparse discrete fracture networks using graph-based reductions. *Physical Review E*, *96*(1), 013304. https://doi.org/10.1103/PhysRevE.96.013304

Hyman, J. D., Jiménez-Martínez, J., Viswanathan, H., Carey, J., Porter, M., Rougier, E., et al. (2016). Understanding hydraulic fracturing: A multi-scale problem. *Philosophical Transactions of the Royal Society A*, *374*(2078), 20150426.

Hyman, J. D., Karra, S., Makedonska, N., Gable, C. W., Painter, S. L., & Viswanathan, H. S. (2015). dfnWorks: A discrete fracture network framework for modeling subsurface flow and transport. *Computers & Geosciences*, *84*, 10–19.

Hyman, J. D., Painter, S. L., Viswanathan, H., Makedonska, N., & Karra, S. (2015). Influence of injection mode on transport properties in kilometer-scale three-dimensional discrete fracture networks. *Water Resources Research*, *51*, 7289–7308.

Hyman, J. D., Rajaram, H., Srinivasan, S., Makedonska, N., Karra, S., Viswanathan, H., & Srinivasan, G. (2019). Matrix diffusion in fractured media: New insights into power-law scaling of breakthrough curves. *Geophysical Research Letters*, *46*, 13,785–13,795. https://doi.org/10.1029/2019GL085454

Jenkins, C., Chadwick, A., & Hovorka, S. D. (2015). The state of the art in monitoring and verification—Ten years on. *International Journal of Greenhouse Gas Control*, *40*, 312–349.

Kang, P. K., Dentz, M., Le Borgne, T., & Juanes, R. (2015). Anomalous transport on regular fracture networks: Impact of conductivity heterogeneity and mixing at fracture intersections. *Physical Review E*, *92*(2), 022148.

Kang, P. K., Lei, Q., Dentz, M., & Juanes, R. (2019). Stress-induced anomalous transport in natural fracture networks. *Water Resources Research*, *55*(5), 4163–4185. https://doi.org/10.1029/2019WR024944

Kreft, A., & Zuber, A. (1978). On the physical meaning of the dispersion equation and its solutions for different initial and boundary conditions. *Chemical Engineering Science*, *33*(11), 1471–1480.

Kueper, B. H., & McWhorter, D. B. (1991). The behavior of dense, nonaqueous phase liquids in fractured clay and rock. *Groundwater*, *29*(5), 716–728.

Lichtner, P. C., Hammond, G. E., Lu, C., Karra, S., Bisht, G., Andre, B., et al. (2015). PFLOTRAN user manual: A massively parallel reactive flow and transport model for describing surface and subsurface processes (*Tech. Rep.*): (Report No.: LA-UR-15-20403) Los Alamos National Laboratory.

Lu, D., Zhang, G., Webster, C., & Barbier, C. (2016). An improved multilevel Monte Carlo method for estimating probability distribution functions in stochastic oil reservoir simulations. *Water Resources Research*, *52*, 9642–9660. https://doi.org/10.1002/2016WR019475

Makedonska, N., Hyman, J. D. D., Karra, S., Painter, S. L., Gable, C. W. W., & Viswanathan, H. S. (2016). Evaluating the effect of internal aperture variability on transport in kilometer scale discrete fracture networks. *Advance Water Resources*, *94*, 486–497.

Makedonska, N., Painter, S. L., Bui, Q. M., Gable, C. W., & Karra, S. (2015). Particle tracking approach for transport in three-dimensional discrete fracture networks. *Computational Geology*, *19*, 1123–1137.

Neuman, S. P. (2005). Trends, prospects and challenges in quantifying flow and transport through fractured rocks. *Hydrogeology Journal*, *13*(1), 124–147.

Ng, L. W., & Willcox, K. E. (2014). Multifidelity approaches for optimization under uncertainty. *International Journal for numerical methods in Engineering*, *100*(10), 746–772.

O'Malley, D., Karra, S., Hyman, J., Viswanathan, H. S., & Srinivasan, G. (2018). Efficient Monte Carlo with graph-based subsurface flow and transport models. *Water Resources Research*, *54*, 3758–3766. https://doi.org/10.1029/2017WR022073

Osthus, D., Hyman, J. D., Karra, S., Panda, N., & Srinivasan, G. (2020). A probabilistic clustering approach for identifying primary subnetworks of discrete fracture networks with quantified uncertainty. *SIAM/ASA Journal on Uncertainty Quantification*, *8*(2), 573–600.

Painter, S., Cvetkovic, V., & Selroos, J.-O. (2002). Power-law velocity distributions in fracture networks: Numerical evidence and implications for tracer transport. *Geophysical Research Letters*, *29*(14), 1676. https://doi.org/10.1029/2002gl014960.

Painter, S. L., Gable, C. W., & Kelkar, S. (2012). Pathline tracing on fully unstructured control-volume grids. *Computational Geology*, *16*(4), 1125–1134.

Park, Y. J., Lee, K. K., & Berkowitz, B. (2001). Effects of junction transfer characteristics on transport in fracture networks. *Water Resources Research*, *37*(4), 909–923.

Park, Y., Lee, K., Kosakowski, G., & Berkowitz, B. (2003). Transport behavior in three-dimensional fracture intersections. *Water Resources Research*, *39*(8), 1215. https://doi.org/10.1029/2002WR001801

Peherstorfer, B., Willcox, K., & Gunzburger, M. (2016). Optimal model management for multifidelity Monte Carlo estimation. *SIAM Journal on Scientific Computing*, *38*(5), A3163–A3194.

Pisaroni, M., Krumscheid, S., & Nobile, F. (2017). Quantifying uncertain system outputs via the multilevel Monte Carlo method - Part I: Central moment estimation (*Tech. Rep.*): MATHICSE.

Rasmuson, A., & Neretnieks, I. (1986). Radionuclide transport in fast channels in crystalline rock. *Water Resources Research*, *22*(8), 1247–1256.

Seidman, S. B. (1983). Network structure and minimum degree. *Social Networks*, *5*(3), 269–287.

Selroos, J.-O., Walker, D. D., Ström, A., Gylling, B., & Follin, S. (2002). Comparison of alternative modelling approaches for groundwater flow in fractured rock. *Journal of Hydrology*, *257*(1-4), 174–188.

Sherman, T., Hyman, J. D., Bolster, D., Makedonska, N., & Srinivasan, G. (2018). Characterizing the impact of particle behavior at fracture intersections in three-dimensional discrete fracture networks. *Physical Review E*, *99*, 013110.

Stockman, H. W., Li, C., & Wilson, J. L. (1997). A lattice-gas and lattice Boltzmann study of mixing at continuous fracture junctions: Importance of boundary conditions. *Geophysical Research Letters*, *24*(12), 1515–1518.

Sweeney, M. R., Gable, C. W., Karra, S., Stauffer, P. H., Pawar, R. J., & Hyman, J. D. (2019). Upscaled discrete fracture matrix model (udfm): An octree-refined continuum representation of fractured porous media. *Computers & Geosciences*, *24*, 293–310.

Tartakovsky, D. M. (2007). Probabilistic risk analysis in subsurface hydrology. *Geophysical Research Letters*, *34*, L05404. https://doi.org/10.1029/2007GL029245

VanderKwaak, J., & Sudicky, E. (1996). Dissolution of non-aqueous-phase liquids and aqueous-phase contaminant transport in discretely-fractured porous media. *Journal of Contaminant Hydrology*, *23*(1-2), 45–68.

Viswanathan, H. S., Hyman, J., Karra, S., O'Malley, D., Srinivasan, S., Hagberg, A., & Srinivasan, G. (2018). Advancing graph-based algorithms for predicting flow and transport in fractured rock. *Water Resources Research*, *54*, 6085–6099. https://doi.org/10.1029/2017WR022368

Wellman, T. P., Shapiro, A. M., & Hill, M. C. (2009). Effects of simplifying fracture network representation on inert chemical migration in fracture-controlled aquifers. *Water Resources Research*, *45*, 6085–6099. https://doi.org/10.1029/2017WR022368

Zhang, D. (2001). *Stochastic methods for flow in porous media: Coping with uncertainties.* Amsterdam: Elsevier.

Zimmerman, R. W., & Bodvarsson, G. S. (1996). Hydraulic conductivity of rock fractures. *Transport Porous Media*, *23*(1), 1–30.