Environmental Modelling & Software 55 (2014) 176–189

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

A high-performance workflow system for subsurface simulation

Vicky L. Freedman [a,*], Xingyuan Chen [a], Stefan Finsterle [b], Mark D. Freshley [a], Ian Gorton [d], Luke J. Gosink [a], Elizabeth H. Keating [c], Carina S. Lansing [a], William A.M. Moeglein [a], Christopher J. Murray [a], George S.H. Pau [b], Ellen Porter [a], Sumit Purohit [a], Mark Rockhold [a], Karen L. Schuchardt [a], Chandrika Sivaramakrishnan [a], Velimir V. Vessilinov [c], Scott R. Waichler [a]

[a] Pacific Northwest National Laboratory, 902 Battelle Blvd, Richland, WA 99352, USA
[b] Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA
[c] Los Alamos National Laboratory, 30 Bikini Atoll Rd, Los Alamos, NM 87545 94720, USA
[d] Carnegie Mellon Software Engineering Institute, 4500 Fifth Avenue, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Received 2 October 2013
Received in revised form
15 January 2014
Accepted 25 January 2014
Available online 14 February 2014

Keywords:
Workflows
ASCEM
Akuna
Amanzi
Model calibration
Uncertainty analysis
Contaminant transport
Vadose zone

ABSTRACT

The U.S. Department of Energy (DOE) recently invested in developing a numerical modeling toolset called ASCEM (Advanced Simulation Capability for Environmental Management) to support modeling analyses at legacy waste sites. This investment includes the development of an open-source user environment called Akuna that manages subsurface simulation workflows. Core toolsets accessible through the Akuna user interface include model setup, grid generation, sensitivity analysis, model calibration, and uncertainty quantification. Additional toolsets are used to manage simulation data and visualize results. This new workflow technology is demonstrated by streamlining model setup, calibration, and uncertainty analysis using high performance computation for the BC Cribs Site, a legacy waste area at the Hanford Site in Washington State. For technetium-99 transport, the uncertainty assessment for potential remedial actions (e.g., surface infiltration covers) demonstrates that using multiple realizations of the geologic conceptual model results in greater variation in concentration predictions than when a single model is used.

Software availability

Name of software: Akuna
Developers: Consortium of National Laboratories (PNNL, LBNL, LANL)
Contact: tim.scheibe@pnnl.gov, vicky.freedman@pnnl.gov
Hardware requirements: Desktop computer with at least 4 GB memory
Software requirements: 32-bit Windows XP/Windows 7, 64-bit Windows 7 and 64-bit Mac OS
Programming language: 64-bit Java 6
Availability: Web links for downloading executable and tutorial files at http://akuna.labworks.org/download.html
Cost: Free

1. Introduction

Significant complexity is involved in computational simulation, including preparing data for input, executing multiple simulations, visualizing results and tracking the data that evolve from multiple analyses. The overall process is not readily amenable to automation, since each step usually requires that the modeler examine the results before proceeding to the next step in the analysis. Moreover, the process of data preparation, execution, analysis and decision-making is often followed by even more data preparation, execution, analysis and decision-making as the investigation proceeds. This process can occur over long time periods, and can involve significant user interaction. For example, an environmental computational analysis can require a repetitive cycle of moving data to a supercomputer or workstation for analysis and simulation, launching the simulations, and managing the storage of the output results. To step through this workflow, modelers typically make extensive use of batch files, shell scripts and scripting-language

* Corresponding author.
E-mail address: vicky.freedman@pnnl.gov (V.L. Freedman).

1364-8152/$ – see front matter © 2014 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.envsoft.2014.01.030

programs to link the sequence of applications needed to complete the analysis. For large data sets, data reduction techniques and parallel visualization may be needed to analyze results generated from the simulations.

Numerical models are frequently used to assess future risks, support remediation and monitoring program decisions, and assist in design of specific remedial actions for complex systems. These decisions are often made with incomplete information, and the impacts of knowledge gaps need to be quantified. Subsurface science is not the only environmental discipline that faces the challenge of making management decisions in the presence of significant uncertainty. Modeling is used in policy and decision-making for other disciplines, such as climate change (e.g., Li et al., 2014; Stainforth et al., 2006), sustainable development (e.g., Mortberg et al., 2013; De Lara and Marinet, 2009) and future energy supplies (e.g., Arnette, 2013; Jebaraj and Iniyan, 2006). Given the importance of identifying uncertainty, several freely available software packages, such as PEST (Doherty, 2010a, 2010b) and UCODE (Poeter et al., 2005) have emerged for uncertainty quantification. Web-based distributed modeling architectures (Bastin et al., 2013) have also emerged to assist modelers in quantifying uncertainty. However, uncertainty quantification can be extremely computationally intensive, requiring many model runs for their implementation, thus making their deployment difficult without high performance computing (i.e., supercomputers).

The U.S. Department of Energy (DOE) has recognized the need for high performance computing and has recently made investments in developing computational tools that can be used to predict the long-term behavior of subsurface contaminant plumes. Remediation of legacy DOE wastes is one of the most complex and technically challenging cleanup efforts in the world, with costs over the next few decades projected to be $265-305 billion (USDOE, 2008). The Advanced Simulation Capability for Environmental Management (ASCEM) program currently underway uses state-of-the-art scientific tools for integrating data, scientific understanding and software. One of the key features of ASCEM is the user environment, Akuna, which is a customized interface for managing subsurface modeling workflows. Akuna provides users with a range of tools to manage environmental and simulator data sets, translate conceptual models to numerical models (including grid generation), execute simulations, and visualize results. Additional toolsets provide users with methods for sensitivity analysis, model calibration and uncertainty quantification.

Several different scientific workflow systems exist [e.g.,Triana (Churches et al., 2006), Pegasus (Deelman et al., 2005), Kepler (Ludascher et al., 2006) and Taverna (Oinn et al., 2004)]. Some of these systems target a particular scientific domain (e.g., Taverna) while others are more generic (e.g., Triana and Kepler). For example, the Kepler system (http://www.kepler-project.org), is used to create, coordinate and execute scientific workflows that can be customized to the user's needs. Typical domain scientists, however, do not have the programming expertise needed to customize Kepler to fit their workflows, and assistance from computer programmers is usually required.

In addition to standalone scientific workflow systems, the Linked Environments for Atmospheric Discovery (LEAD; Plale et al., 2006) project demonstrates how workflows can be used to solve problems specific to Earth system science by integrating different technologies such as web and grid services and workflow systems. Integration of data and model workflows is demonstrated in Turuncoglu et al. (2013), who discuss coupling an Earth System Modeling Framework (ESMF) with the Regional Ocean Modeling System (ROMS) and Weather Research and Forecasting Model (WRF). The focus of their work is on the development of portable and replicable simulation workflows to create self-describing models with common model component interfaces.

User interfaces are an important component of the workflow system. Commercial user interfaces (UIs) (e.g. GMS (2012), Visual MODFLOW (2012), and Groundwater Vistas (2012)) have been developed specifically for groundwater flow and transport using the MODFLOW (Harbaugh, 2005) family of codes, a U.S. Geological Survey simulator that is the de facto standard code for aquifer simulation. Akuna, however, is unique in four major aspects. The first is in its ability to facilitate both serial and high-performance computation (HPC) in a workflow environment already customized for subsurface modeling. Although it is specifically designed to work with the ASCEM simulator, Amanzi, it can be used with other simulators as long as it is set up to read and write that simulator's file formats. Second, unlike many of the UIs for MODFLOW, Akuna provides an interface for variably saturated and multiphase flow simulators, and is not restricted to groundwater only applications. A third distinguishing characteristic is that Akuna is an open-source, platform-independent UI that integrates with other open-source software (e.g., WorldWind (2012), VisIt (2012)) for providing the user with all of the tools needed to perform a complete modeling analysis from model setup, calibration and uncertainty quantification. Finally, Akuna provides a client-server architecture and collaborative user interface, enabling users to perform their modeling analysis cooperatively from disparate locations.

The primary objective of this paper is to demonstrate Akuna capabilities that have been developed to date. This is accomplished by using the BC Cribs Site as an example application for the workflow system. To this end, the large-scale disposal of liquid inorganic waste is simulated for this site, which is located at the Hanford Site in southeastern Washington State. These subsurface discharges were a byproduct of nuclear weapons production during the Cold War. The BC Cribs Site received nearly 140 Ci of technetium-99 (^{99}Tc) in approximately 39 million liters of water (Kincaid et al., 2006). To date, this contamination has migrated to approximately 70 m below ground surface (bgs) into a 107 m thick vadose zone. Remediation of the recalcitrant ^{99}Tc is receiving increased attention in recent years because of its long half-life (2.13×10^5 years), the difficulty posed by its location in the deep vadose zone, and its near-term threat to groundwater.

The Akuna software is used to demonstrate model setup, calibration and uncertainty analysis for the BC Cribs Site, and to develop a model that can be used for evaluating potential remediation alternatives. The impact of accounting for multiple geologic realizations in an analysis of future boundary conditions is evaluated, and could be potentially important for future remedial actions at the site. Throughout the example application, it is demonstrated that use of high-performance computing makes execution of multiple simulations feasible, and the Akuna toolset streamlines the process.

2. Akuna user environment

Akuna is an open-source, platform-independent user environment. It includes features for basic model setup, sensitivity analysis, parameter estimation, uncertainty quantification, launching and monitoring simulations, and visualization of both model setup and simulation results. Features of the model setup tool include visualizing wells and lithologic contacts, generating surfaces or loading surfaces produced by other geologic modeling software (e.g., Petrel (2012), EarthVision (2012)), and specifying material properties, initial and boundary conditions, and model output. Currently, the model setup tool is equipped with a rectilinear grid generator for generating structured grids (orthogonal elements with a uniform pattern). Currently, unstructured grids (orthogonal or non-orthogonal elements with either uniform or non-uniform pattern) can be incorporated via file read using an Exodus file format (Schoof and Yarberry, 1994). Partial integration with WorldWind (WW, 2012) has been completed thus far, and allows the user to import files that can be visualized in WorldWind.

After the model has been set up, Akuna facilitates launching either a single run, or multiple runs needed for sensitivity analyses, parameter estimation and uncertainty quantification. Automated job launching and monitoring capabilities allow a user to submit and monitor simulation runs on high-performance, parallel computers with batch queue systems. Visualization of large output files can be performed without moving the data back to local resources. These capabilities make high-performance computing easier for users who might not be familiar with batch queue systems and usage protocols on different supercomputers and clusters.

Akuna supports a common workflow needed for developing and applying a subsurface model (Fig. 1). Many elements of this workflow are repeatedly and iteratively performed as part of the modeling process. Typically, a conceptual understanding of the system to be analyzed is gained from site characterization efforts and monitoring data. This conceptual understanding is then translated into a mathematical model and implemented in a numerical model, which requires tools to describe the model domain with its salient hydrogeochemical features, associated material properties, initial and boundary conditions, forcing terms, as well as information on how space and time are discretized for numerical solution. These functions are supported by Akuna's Model Setup (MS) toolset.

Once an initial numerical model has been developed, Akuna's Simulation Run (SR) toolset can be used to launch and monitor a single simulation, the results of which can be analyzed and visualized. If this initial run is considered reasonable, a formal local or global sensitivity analysis can be performed using Akuna's Sensitivity Analysis (SA) toolset to identify the parameters that most strongly influence the system behavior, and to examine output variables that are sensitive to the parameters of interest. These parameters may include material properties, but also initial and boundary conditions, and any aspect of the conceptual model that can be suitably parameterized. If measurements of sufficient sensitivity and accuracy are available, the model can be automatically calibrated using Akuna's Parameter Estimation (PE) toolset. This step not only provides effective parameter values that can be considered consistent with the data collected at the site, but also provides estimates of uncertainty. Akuna's Uncertainty Quantification (UQ) toolset can then be used to evaluate the uncertainty of model predictions and provide the basis for a subsequent assessment of environmental and health risks.

In practical applications, the workflow is usually not as linear as described above. Hence, double arrows amongst all the steps in the workflow are shown in Fig. 1. The toolsets integrated into Akuna are transparent and can be flexibly invoked to accommodate any application's particular workflow.

2.1. Akuna architecture

Akuna's desktop UI provides a front end to the simulation workflow (Fig. 2). The cross-platform UI is written in Java and is built on the Velo knowledge management framework (Gorton et al., 2011), which provides a robust open-source content management system to manage workflow data and metadata. The Velo framework is a client-server architecture comprised of an extensible front-end user interface coupled with an extensible back-end content management system. All project data are stored on the server and protected with fine grained access controls. The user environment allows users to create groups and control permissions on their projects, enabling them to work in private workspaces or to enable collaborative modeling as appropriate. Velo's messaging system allows users to see changes made by others in real time.

The Velo user environment includes many reusable components such as a data browser that provides access to all the tools associated with the workflow. Shared as well as private workspaces are supported to enable collaborative modeling. Toolsets, such as the customized Akuna UI, run on top of the Velo framework. The Model Setup Toolset is executed within Akuna, and is an important interface used for setting up model input files, viewing the conceptual model, and staging multiple simulation runs, such as sensitivity analysis (SA) and uncertainty quantification (UQ) through the Toolset UIs. Other tools that can be used within Model Setup include LaGrit and Gridder (L/G, 2012) (for mesh generation) and WorldWind (2012) (for visualization within its geographical context).

Agni has a critical role in the Akuna architecture, as it accepts job launching requests from the Akuna client, executes them and reports information back to the UI. Agni also controls the local execution of the simulator for the four types of simulation tasks (SR, SA, PE and UQ). For example, Agni is responsible for sampling the parameter space and providing the parameter sets to the simulator for multiple simulation runs. Akuna is also responsible for the analysis toolsets for sensitivity analysis (SA), parameter estimation (PE) and uncertainty quantification (UQ). Once simulations are completed, a visualization toolset within Akuna can be used to plot

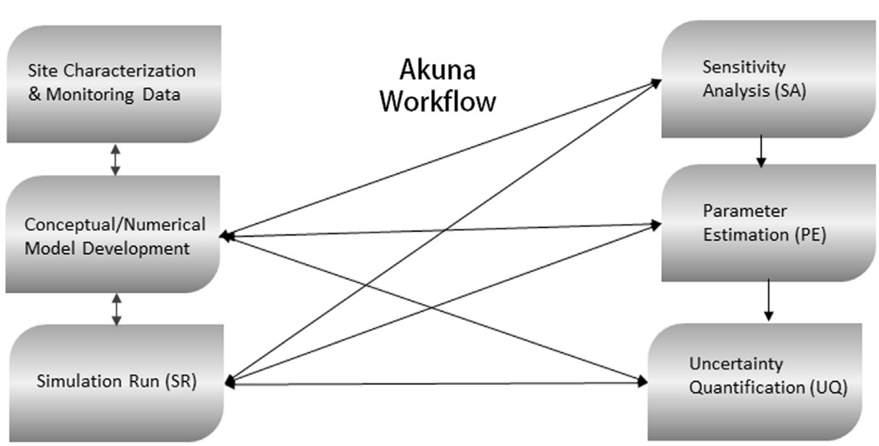

Fig. 1. Workflow using Akuna. Typical workflow starts with using site data to develop a conceptual and numerical model, followed by a simulation run. Considerable iteration is possible as the investigation proceeds, as shown by double arrows connecting steps within the workflow.

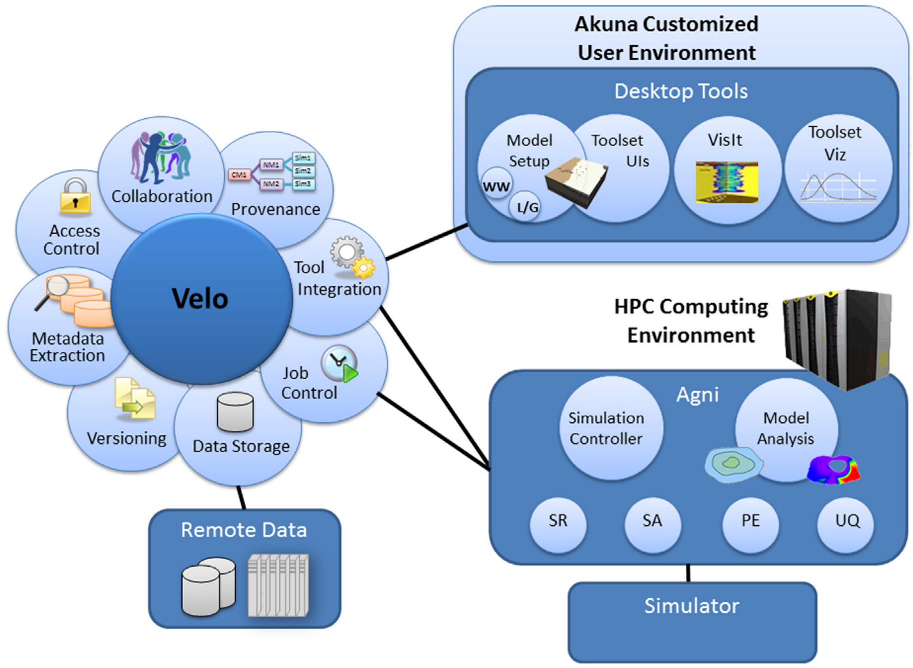

Fig. 2. Akuna architecture showing toolsets that interact with the Velo framework, and how the Akuna Desktop UI interfaces with other toolsets.

non-spatial output, such as the impact of different parameter values on model output (sensitivity analysis), or concentrations over time for a particular point in the simulation domain (uncertainty quantification). VisIt (2012) is an external plotting software package that can also be invoked to visualize spatial output, such as the concentration distribution throughout the entire simulation domain at a single (or multiple) points in time.

Amanzi is the main simulator supported by the Akuna platform. However, Akuna and Agni are designed to accommodate other simulators that can be plugged in using a set of defined interfaces. Currently, Akuna is also setup with STOMP/eSTOMP subsurface simulators (White and Oostrom, 2000, 2006), and will integrate other simulators in future releases.

ASCEM also has a remote data management system to import, organize, retrieve, and search across various types of observational datasets needed for environmental site characterization and numerical modeling. The framework provides capabilities to organize, interactively browse on maps, search by filters, select desired data, plot graphs, and save selected data for subsequent use in the modeling process. Further description of this capability is beyond the scope of this paper, but readers are encouraged to view the website at http://babe.lbl.gov/ascem/maps/SRDataBrowser.php.

3. Akuna example application

The example presented here demonstrates the model setup, calibration and uncertainty toolsets for the BC Cribs Site at Hanford. This use case was chosen because 1) it is an unsaturated flow problem that involves conservative (non-reactive) contaminant transport; 2) its sparse data set lent itself to a simplified use case during the toolset development; and 3) the recalcitrant technetium-99 (^{99}Tc) that resides in the deep vadose zone is one of the most challenging remediation problems in the DOE complex today. Because the contamination still largely resides in the vadose zone, its threat to groundwater has only recently been recognized as requiring remedial action. The simulation work presented here represents the first effort at simulating historical subsurface discharges at BC Cribs. This model will be used in the future to evaluate potential remedial actions at the site.

3.1. Site background

At BC Cribs (and other Hanford locations), large volumes of radiological wastes were released into the subsurface during the development and manufacture of nuclear weapons (see Fig. 3). BC Cribs received scavenged waste from the uranium and ferrocyanide recovery processes from 1956 to 1958 in six open 12.2 m square pits reinforced with wood framing at the bottom. The cribs received waste in large quantities (−42,000 L at a time) from a siphon tank that when full, automatically flushed its contents through a pipe to the crib (DOE/RL, 2008). This practice resulted in significant ^{99}Tc (and nitrate) contamination in the 107 m thick vadose zone. ^{99}Tc is a long lived radionuclide with a half-life of 2.13 × 10^5 years. Since the vadose zone at Hanford is oxidizing, the presumed technetium species is the pertechnetate anion, $TcO_4^{\ -}$, which exhibits high mobility under these conditions (Icenhower et al., 2008). To date, this contamination has migrated to approximately 70 m below ground surface (bgs), and has the potential to contaminate both groundwater and the nearby Columbia River. The remediation of ^{99}Tc at BC Cribs poses a unique challenge because conventional remediation technologies, such as pump and treat, are ineffective in the vadose zone, and the contamination is too deep for excavation.

The heterogeneous nature of the sediments in the vadose zone at BC Cribs also confounds the understanding of the distribution and extent of ^{99}Tc in the subsurface. Because the affected vadose zone is more than 100 m thick, thorough characterization using traditional field methods is prohibitive. The vadose zone sediments of the Hanford formation, a major stratigraphic unit at the site that spans nearly the entire vadose zone at BC Cribs, is known to contain relatively thin (0.5 m or less), fine-textured lenses that can extend laterally for tens of meters (Serne et al., 2009). These small-scale heterogeneities enhance lateral spreading of water and contaminants and reduce the vertical movement (Ward et al., 2009). However, the distribution of these fine- grained layers at BC Cribs is

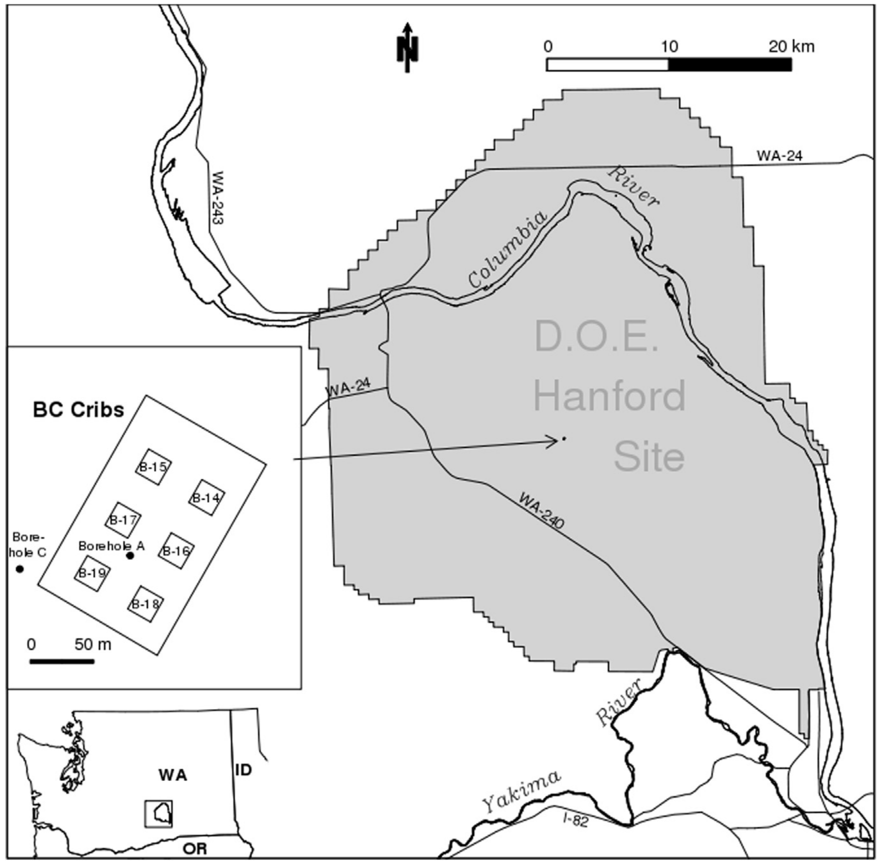

Fig. 3. Map of BC Cribs at the Hanford Site in Southeastern Washington State, showing crib and borehole locations.

largely unknown. Since flow and transport in porous media are determined by its structure and connectivity, this presents a large source of uncertainty in the geologic conceptual model at the BC Cribs site.

3.2. Geologic conceptual model

The gravels, sand, and silt sediments in the vadose zone at BC Cribs were represented stochastically in the conceptual model. Using characterization data from five deep wells, a facies-based geologic conceptual model at BC Cribs was developed based on methods presented in Scheibe et al. (2006) and Murray (1994). Lithofacies used in mapping the BC Cribs area were identified based on analysis of spectral gamma ray well log data, primarily from the Th-232 and K-40 curves. Spherical variogram models (Goovaerts, 1997) were fit to all of the experimental variograms. A 10:1 horizontal to vertical anisotropy ratio was assumed so that the horizontal variogram models could be developed. Three lithofacies were identified by clustering of spectral gamma log data (Th-232 and K-40) from borehole wireline logging. Facies 1 was identified as dominantly sand, facies 2 as a sandy gravel, and facies 3 as a silty (muddy) sand.

Ten realizations of the geologic conceptual model were generated using sequential indicator simulation (Deutsch and Journel, 1998). Although additional simulations would provide a more complete analysis addressing conceptual model uncertainty, for the purpose of demonstrating Akuna workflow, ten realizations were sufficient. Cross-sections through three of the cribs, 216-B-19, 216-B-17 and 216-B-15, are shown in Fig. 4 for these realizations. Cribs are shown in yellow at the top of the domain. The cross-sections show commonality in the locations and thicknesses of the three different facies, but also demonstrate differences in small-scale heterogeneity among the different geostatistical realizations.

3.3. Model Setup Toolset

Akuna's Model Setup Toolset facilitates rapid creation of the simulation input file. Once a new model is started in Akuna, the user is led through steps to define the conceptual model and its associated mesh. Once the extent of the domain is defined, the mesh can be generated. For the BC Cribs use case, the domain was defined as 320 m in the x-direction, and 280 m in the y-direction (after rotation to place the grid in a Cartesian E–W and N–S oriented reference frame). The rectilinear grid was generated using Gridder, a structured mesh generation tool in the Model Setup Toolset, and was discretized at a 5 m resolution in both horizontal directions, and a 1.0 m resolution in the vertical. The domain thickness was set to 107 m, and the water table was set at the bottom. This discretization yielded a total of 383,488 nodes in the simulation domain.

The geologic conceptual model of the BC Cribs involved stochastic realizations of the lithofacies, as described in the previous section. Fig. 5 shows how the lithofacies, which were assigned on a cell-by-cell basis via a file read, can be viewed using slices within the Model Setup Viewer. Another option for defining the geologic model involves defining stratigraphic layers (surfaces), and the Model Setup Toolset fills in regions of the model between the surfaces. Although not applied to the BC Cribs problem, boundary-fitted meshes can also be generated that conform to these surfaces.

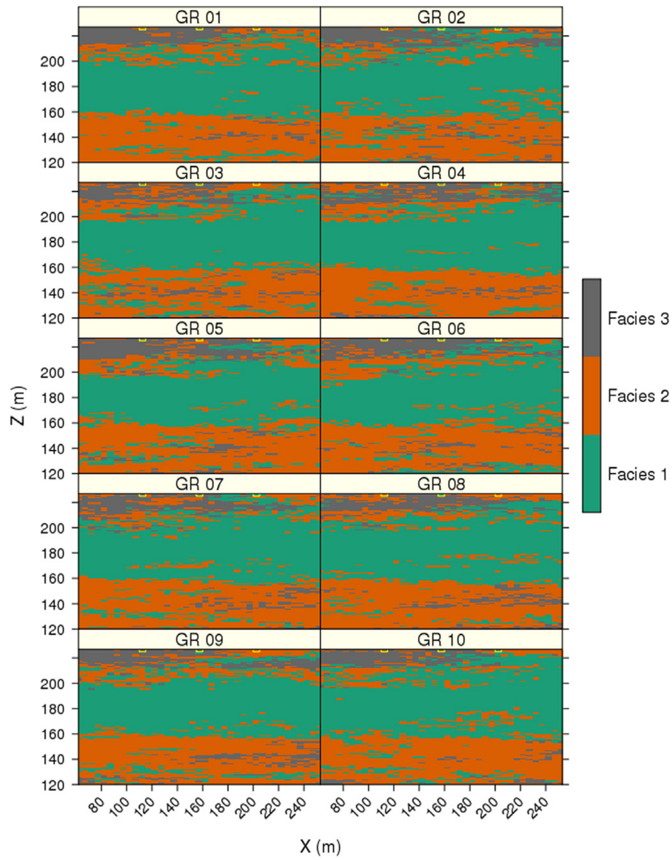

Fig. 4. Geologic cross-sections through Cribs 216-B-19, 216-B-17 and 216-B-15 for the 10 different geologic realizations (GR) of the conceptual model. Numbers at the top of each cross-section refer to the geologic realization number.

3.4. Model Parameter Estimation (calibration)

The simulation period for the calibration was from year 0–2008. The years 0–1956 were an initialization period to yield a steady-state flow field by 1956. Crib releases commenced in 1956, and flow and transport was simulated until 2008, the year in which borehole measurements of moisture content and concentration occurred at Boreholes A and C (Fig. 3). Estimates for vadose zone hydraulic parameters were determined using pedotransfer functions (Guber et al., 2006) developed from spectral gamma log, grain size, and hydraulic property data for Hanford sediments (Table 1). The calibration assumed constant properties within each facies.

Within the Parameter Estimation (PE) Toolset, parameters were selected for the model calibration. Porosity and permeability for each of the three facies were estimated, yielding a total of 6 parameters. Unsaturated hydraulic parameters were not used in the model calibration because initial simulations indicated that these parameters were relatively insensitive, which was likely due to the relatively dry state of the vadose zone when the moisture contents were measured in 2008, and only one measurement in time was available at each vertical borehole location. If transient measurements of moisture content and solute concentration had been available, simulation results would likely have been much more sensitive to unsaturated hydraulic parameters. An anisotropy ratio of 10:1 was assumed for the horizontal to vertical permeability based on convention, since direct measurements of anisotropy in permeability for Hanford vadose zone sediments were not available. Data measured at Boreholes A and C were extracted from the ASCEM data management system, and the PE toolset was used to load the measured data.

Model calibration was performed using the job launcher within the PE Toolset. A restart capability is available should the PE exceed allocated queue time. Each of the calibrations was executed on Hopper, a remote supercomputer at the National Energy Research Scientific Computing (NERSC) Center, using the job launching and monitoring capabilities in Akuna. PE execution control options were defined through the toolset, including the use of the Levenberg–Marquardt algorithm (Levenberg, 1944; Marquardt, 1963) to identify parameter sets that minimize the objective function.

Launching a PE simulation requires that the user define both the total number of processors required and the number of processors per task. For the BC Cribs PE, a total of 576 processors were used. Six simulations were executed simultaneously (in task-parallel computation) for evaluation of the sensitivity matrix, and each simulation run used 96 processors with an execution time of ~24 h.

Upon successful completion of a calibration, several options exist for examining the results. A tabular summary of parameter and error estimates is automatically generated in the PE Toolset, including the plot of the objective function value versus iteration number as shown in Fig. 6a. This shows that the least squares sum of differences between the simulated and measured moisture content and concentration at Boreholes A and C decreases significantly with the first few iterations, and then only modestly improves with successive iterations. In addition, the user can generate graphics that display the simulated and observed quantities (e.g., measured and observed concentrations as shown in Fig. 6a for Borehole A). VisIt software can also be directly launched from within Akuna to view spatial quantities, as shown by the concentration distribution of ^{99}Tc in 1960 in Fig. 7 (after subsurface discharges to the cribs ended in 1958) for geologic realization 01 (realization number referenced in Fig. 4). Horizontal cross-sections through the cribs are also shown, one through cribs 216-B-15, 216-B-17 and 216-B-19 and one through 216-B-14, 216-B-16 and 216-B-18. A cross-section through Borehole A, located between four of the cribs, is also shown in Fig. 7.

Model calibration was performed in the same manner for all 10 geologic realizations of the conceptual model, with each calibration on average performing 8 iterations. A summary of the parameter ranges estimated is presented in Table 2, and shows that significant changes in parameter estimates occurred from their initial estimates. The pedotransfer functions used to estimate initial parameter values were developed from regional data, not site-specific data. This factor, combined with the statistical nature of the pedotransfer functions, contributes to the resulting differences between initial estimates of porosity and permeability and optimized values.

Fig. 8 shows the match between measured and simulated data for all ten realizations of the conceptual models. Porosity and permeability typically are inversely correlated in fully water-saturated aquifer systems, but show less correlation for unsaturated systems. However, inclusion of unsaturated hydraulic parameters in the model calibration could impact the current estimates of porosity and permeability. The parameter estimation process was performed primarily for illustrative purposes, so estimation of additional parameters was not pursued. Further iteration on the number of lithofacies and their distribution, as well as inclusion of additional parameters would be needed to improve the calibration results.

Significant variability occurred for the permeability estimates for both Facies 1 and 2. For Facies 1, a two order-of-magnitude difference exists, whereas for Facies 2, the estimates vary by nearly four orders of magnitude. The large variability in permeability estimates for Facies 2 is likely due to its insensitivity to the existing data. This facies is primarily located at the bottom of the

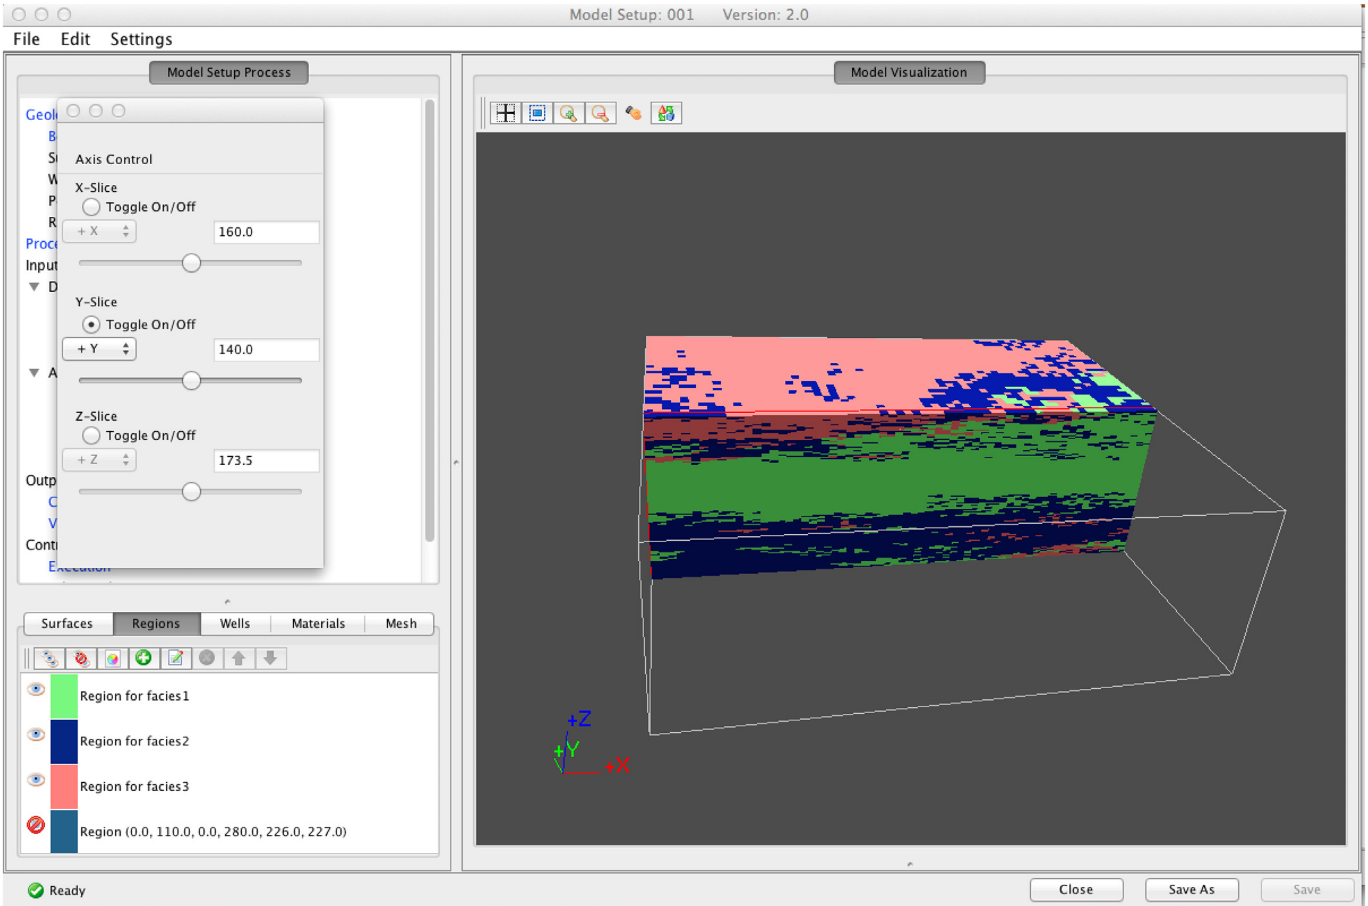

Fig. 5. Model setup visualization tool with the viewer window displaying a cutaway of the lithofacies distribution. Facies 1 (dominantly sand) is shown in green, Facies 2 in blue (sandy gravel), and Facies 3 in pink (silty sand).

domain, but the bulk of the cribs' releases have not yet reached this depth by the year 2008 when the measurements occurred. This is an example of the measured data being too sparse to uniquely determine hydraulic parameters, and underscores the importance of considering uncertainty in predictions of mass transport at BC Cribs.

3.5. Uncertainty quantification

The objective of the Uncertainty Quantification (UQ) was to evaluate the impact of a range of future net infiltration (recharge) conditions at BC Cribs. One hundred different recharge rates were applied in model runs as a constant boundary condition for the years 2012–3000. The 100 values of recharge rate were randomly sampled from a uniform distribution of 0.1–75 mm/yr, which provided an adequate sampling of the parameter space. The range in water recharge rates represented potential impacts from site operation and management actions that influence the net infiltration rate, such as the emplacement of an infiltration barrier (lower recharge rates) or a no-action alternative consisting of monitored natural attenuation (higher recharge rates).

The impact of uncertainty in the future recharge rate was represented by metrics related to ^{99}Tc concentration in the capillary fringe. The simulated "observation" points were located directly beneath Boreholes A and C and represent vadose zone concentrations. Consequently, concentrations were much higher than they would be if they were diluted by groundwater and sampled over the screened interval of a well. For the purposes of demonstration, these concentrations were analyzed within the context of a threshold concentration, a metric analogous to a maximum concentration level (MCL). A value of 100,000 pCi/L was arbitrarily selected as the threshold concentration in this analysis. The primary metrics used in the UQ were peak concentration, the amount of time from present to the first exceedance of the threshold concentration, and the duration of time that the threshold concentration was exceeded.

The transition from a successful calibration to an uncertainty analysis is accomplished within Akuna by selecting the UQ Toolset, and specifying that the parameter estimates from the PE should be used in the simulations for the uncertainty analysis. After identifying the recharge parameter, number of simulations (100), and the range of values (0.1–75 mm/yr), a histogram of samples is generated (Fig. 9), and input files are generated. The same recharge distribution was used for each of the 10 geologic realizations. The ten different sets of hydraulic parameters estimated from the calibration were used in the Monte Carlo simulation, for a total of 1000 simulations.

Table 1
Parameter estimates for each facies using pedotransfer functions.

	Horizontal permeability (m^2)	Porosity	Brooks & Corey Entry Head (m)	Brooks & Corey λ
Facies 1	1.99×10^{-13}	0.408	0.413	0.283
Facies 2	6.93×10^{-12}	0.220	0.039	0.261
Facies 3	2.07×10^{-10}	0.240	0.037	0.387

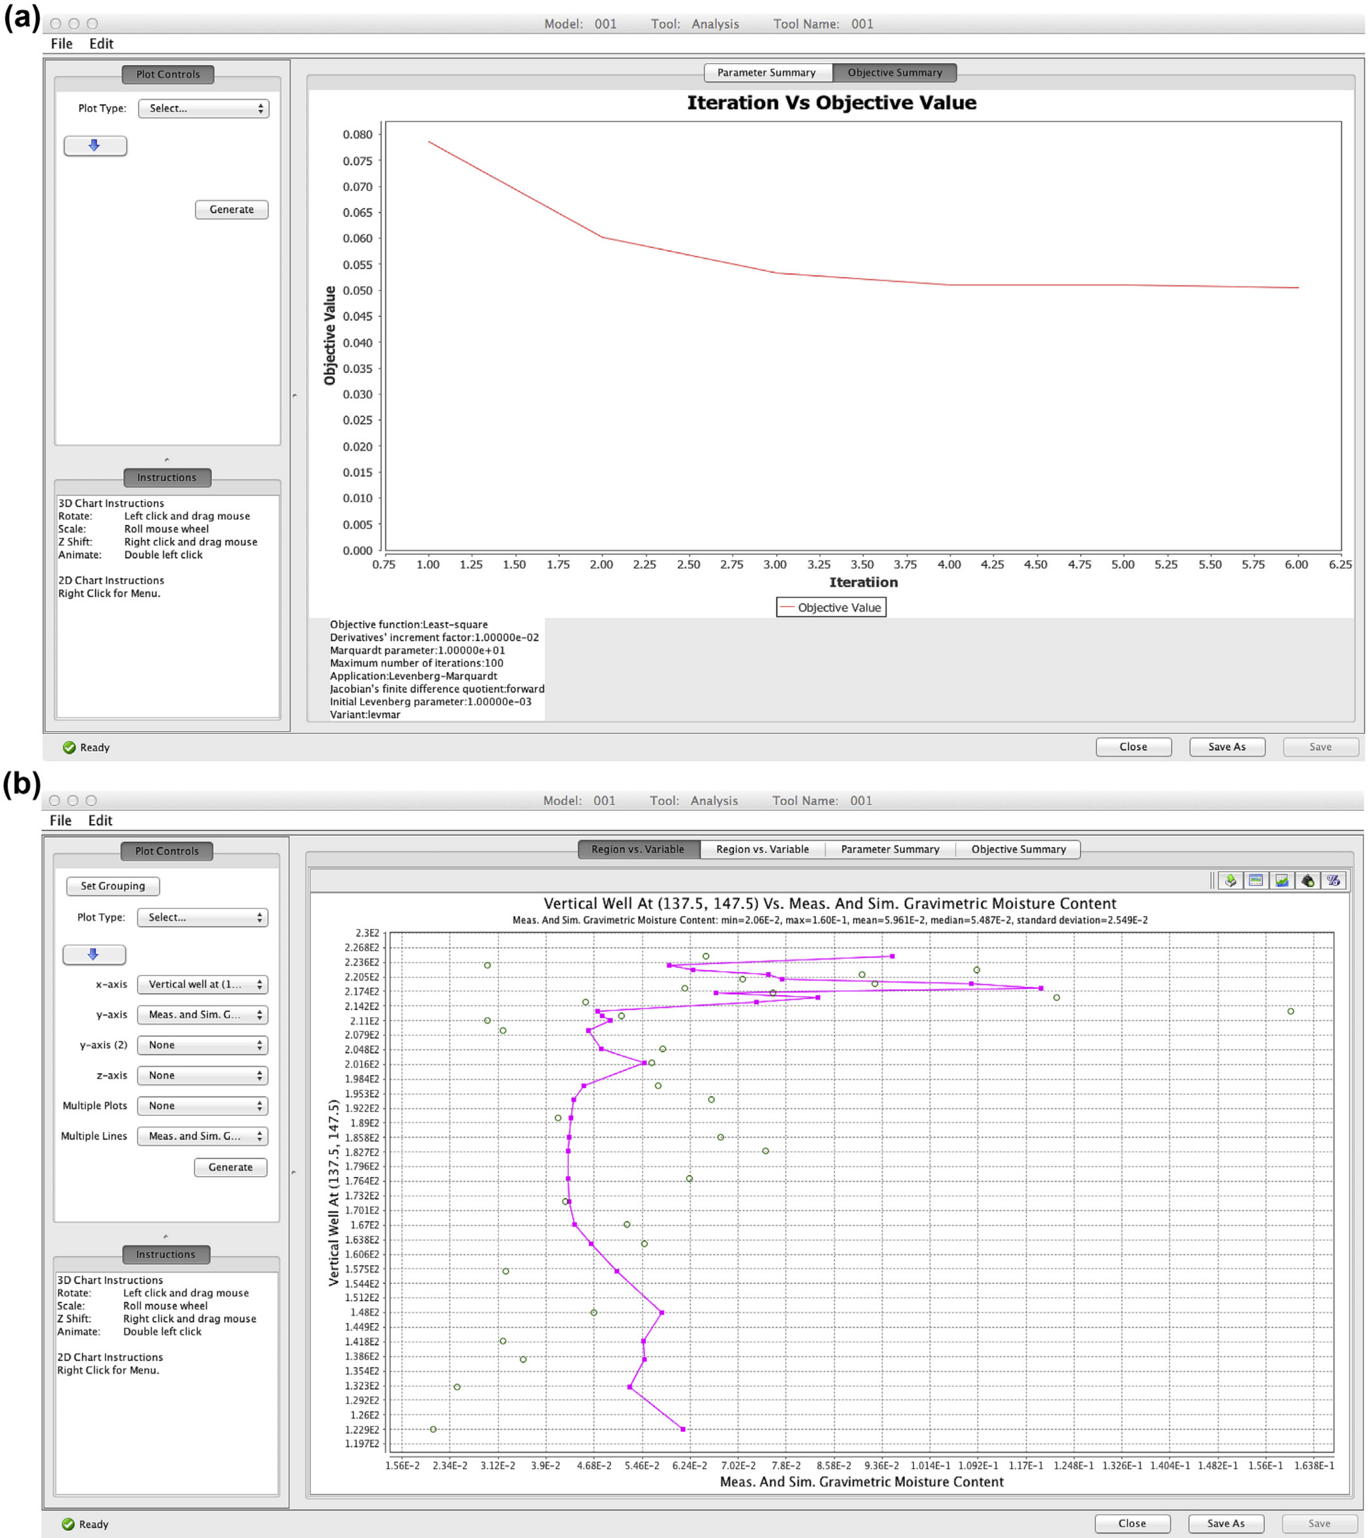

Fig. 6. Screenshots from Akuna showing a) the value of the objective function decreasing with the number of iterations; and b) elevation vs. measured (open green circles) and simulated gravimetric moisture content (lines and points in pink) at Borehole A.

Simulations in the UQ Toolset are launched in a similar manner to the PE Toolset. The UQ was launched using a total of 9600 cores, with 96 cores per model run. All 100 simulations for one geologic realization were completed within 6 h, yielding a total execution time of 60 h for the 1000 simulations. Once the simulations were complete, breakthrough curves, histograms and scatter plots were generated to interpret results of the analysis using the Visualization Toolset in Akuna. Fig. 10a shows a screenshot from Akuna that plots

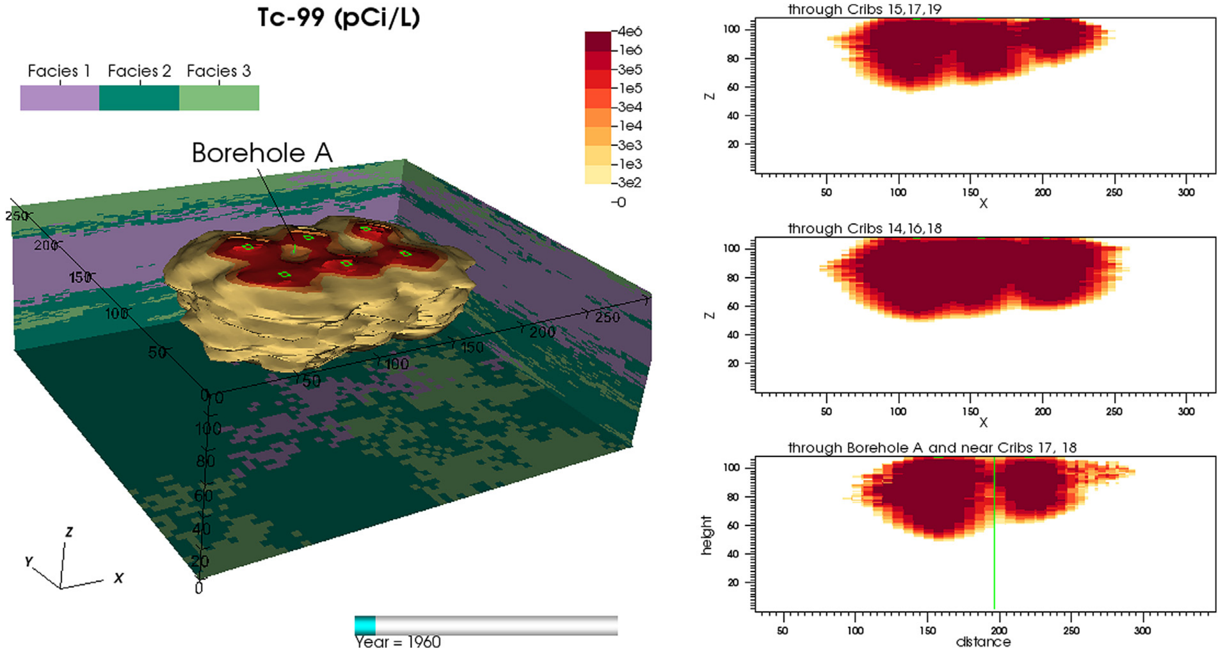

Fig. 7. Spatial distribution of ^{99}Tc in the year 1960 using VisIt software (subsurface discharges to the cribs ended in 1958.) for geologic realization number 01(see Fig. 4). Two horizontal cross-sections are shown through the cribs, and one cross section through Boreholes A is also shown.

the mean and 95% confidence intervals for ^{99}Tc over time at borehole locations A and C. A histogram showing the time to reach the peak concentration at these same locations is shown in Fig. 10b.

To analyze the uncertainty results across the 10 geologic realizations (1000 simulation runs), the breakthrough curves (BTCs) for all runs are compared to the BTCs shown for a single realization (01). The 95% confidence intervals are wider when all 10 realizations are considered (Fig. 11). For example, the upper bound on the confidence interval is approximately 85% higher at Borehole A for all 10 models than just for 01. A similar increase in the 95% confidence interval is shown for Borehole C.

The variability across all runs is also noted in the scatter plot depicting the number of years that ^{99}Tc is above the arbitrary threshold concentration of 100,000 pCi/L (Fig. 12a). At Borehole A, the trend demonstrates that lower recharge rates increase the amount of time the concentrations are above the threshold concentration, whereas higher recharge rates generally translate into shorter periods of time that exceed the threshold concentration. In some cases, like GR01 at Borehole C, the post-2012 recharge rate has no impact on the number of years to exceedance. Even with recharge rates close to zero (e.g., <10 mm/yr), the plume is close enough to the water table in the year 2012 that the threshold concentration is exceeded within 50 years. With other conceptual model realizations, a lower recharge rate increases the number of years required to exceed the threshold concentration. Fig. 12b is a histogram that compares the number of years of exceedance for the single and multiple realizations. Greater variability occurs at both locations for all realizations of the conceptual model, although the variability is more significant for shorter periods of exceedance at Borehole C and more significant for longer periods of exceedance at Borehole A. These results have important implications for remediation technologies that reduce the recharge rate, such as soil desiccation and placement of infiltration barriers (Wellman et al., 2011; Truex et al., 2013), technologies currently being considered at the BC Cribs Site. A reduction in the recharge rate may delay the arrival of peak concentrations to the water table, but it may also prolong the duration at which the concentrations are above the threshold concentration.

One of the primary advantages of high-performance computing is the reduction in computational time, which means that multiple models can be analyzed. On average, only 30 h were required to complete the model calibration (24 h) and execute 100 simulations in an uncertainty analysis (6 h) for a single geologic realization. Equivalent simulation without parallel processing would have taken 90 days to complete a single model calibration, assuming sufficient memory was available. The uncertainty analysis would have required 22 days, assuming that 100 computers were available to simultaneously execute the 100 simulations for just one of the geologic realizations in the uncertainty analysis.

Table 2
Minimum and maximum parameter estimates among the different realizations of the conceptual model.

Parameter	Min(m^2)	Max(m^2)
Horizontal Permeability – Facies 1	1.68×10^{-12}	1.38×10^{-10}
Horizontal Permeability – Facies 2	1.17×10^{-14}	1.89×10^{-10}
Horizontal Permeability – Facies 3	1.00×10^{-14}	1.45×10^{-13}
Porosity – Facies 1	0.132	0.266
Porosity – Facies 2	0.165	0.283
Porosity – Facies 3	0.243	0.342

4. Summary and conclusions

4.1. BC Cribs multiple conceptual model results

In addition to demonstrating Akuna Toolsets, this paper provides insight on the relative roles of recharge rates and lithofacies distributions on predictions of ^{99}Tc transport at the Hanford BC Cribs site. This analysis represents the first field-scale modeling effort at the BC Cribs site that establishes baseline conditions for a "no-action" alternative, as well as a preliminary assessment of uncertainties associated with potential remedial actions, such as the placement of surface infiltration covers. Although the parameter estimation results provided a far from perfect match between

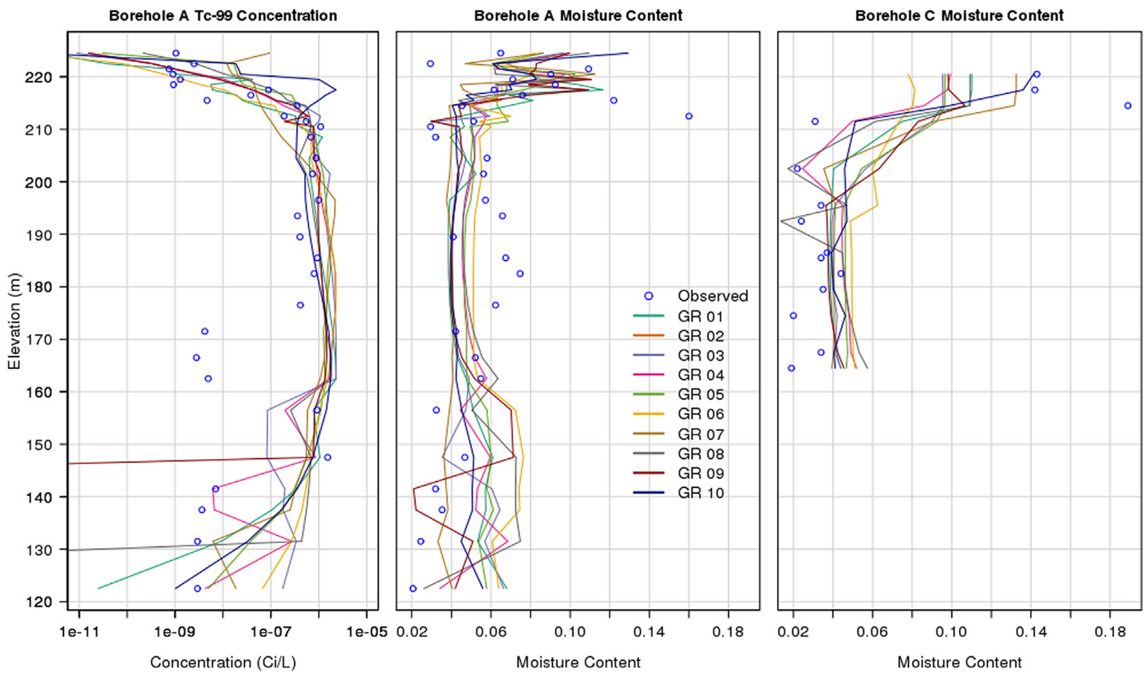

Fig. 8. Simulated and measured moisture contents and concentrations at Boreholes A and C for all 10 geologic conceptual model realizations.

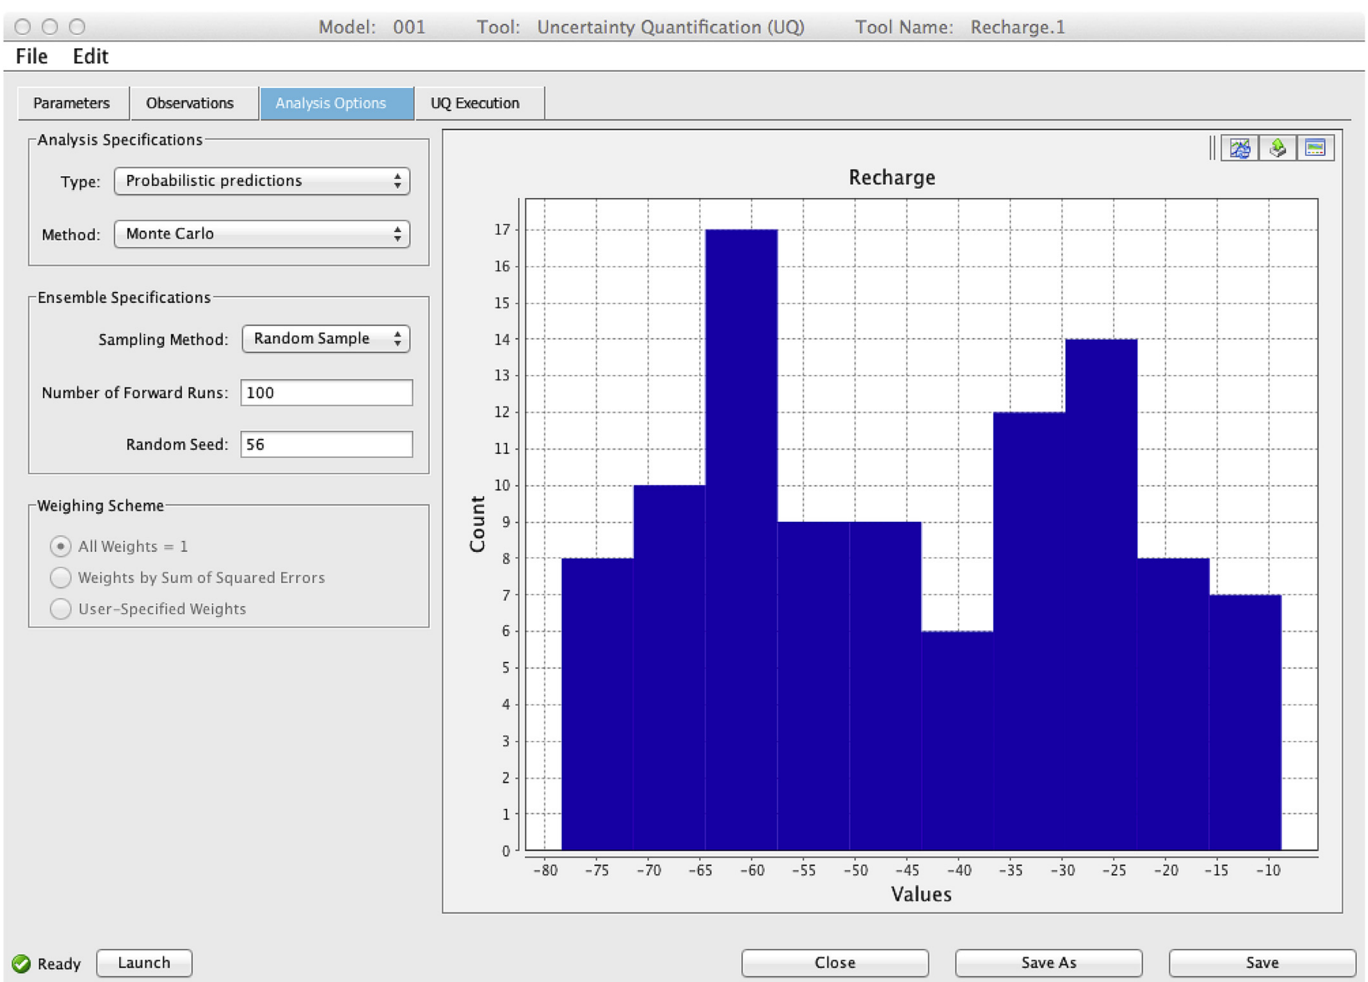

Fig. 9. Screenshot from the UQ Toolset showing a histogram of the recharge rates sampled from a uniform distribution. The horizontal axis displays the recharge rates (negative numbers to represent downward direction), whereas the vertical axis displays the number of realizations in that interval.

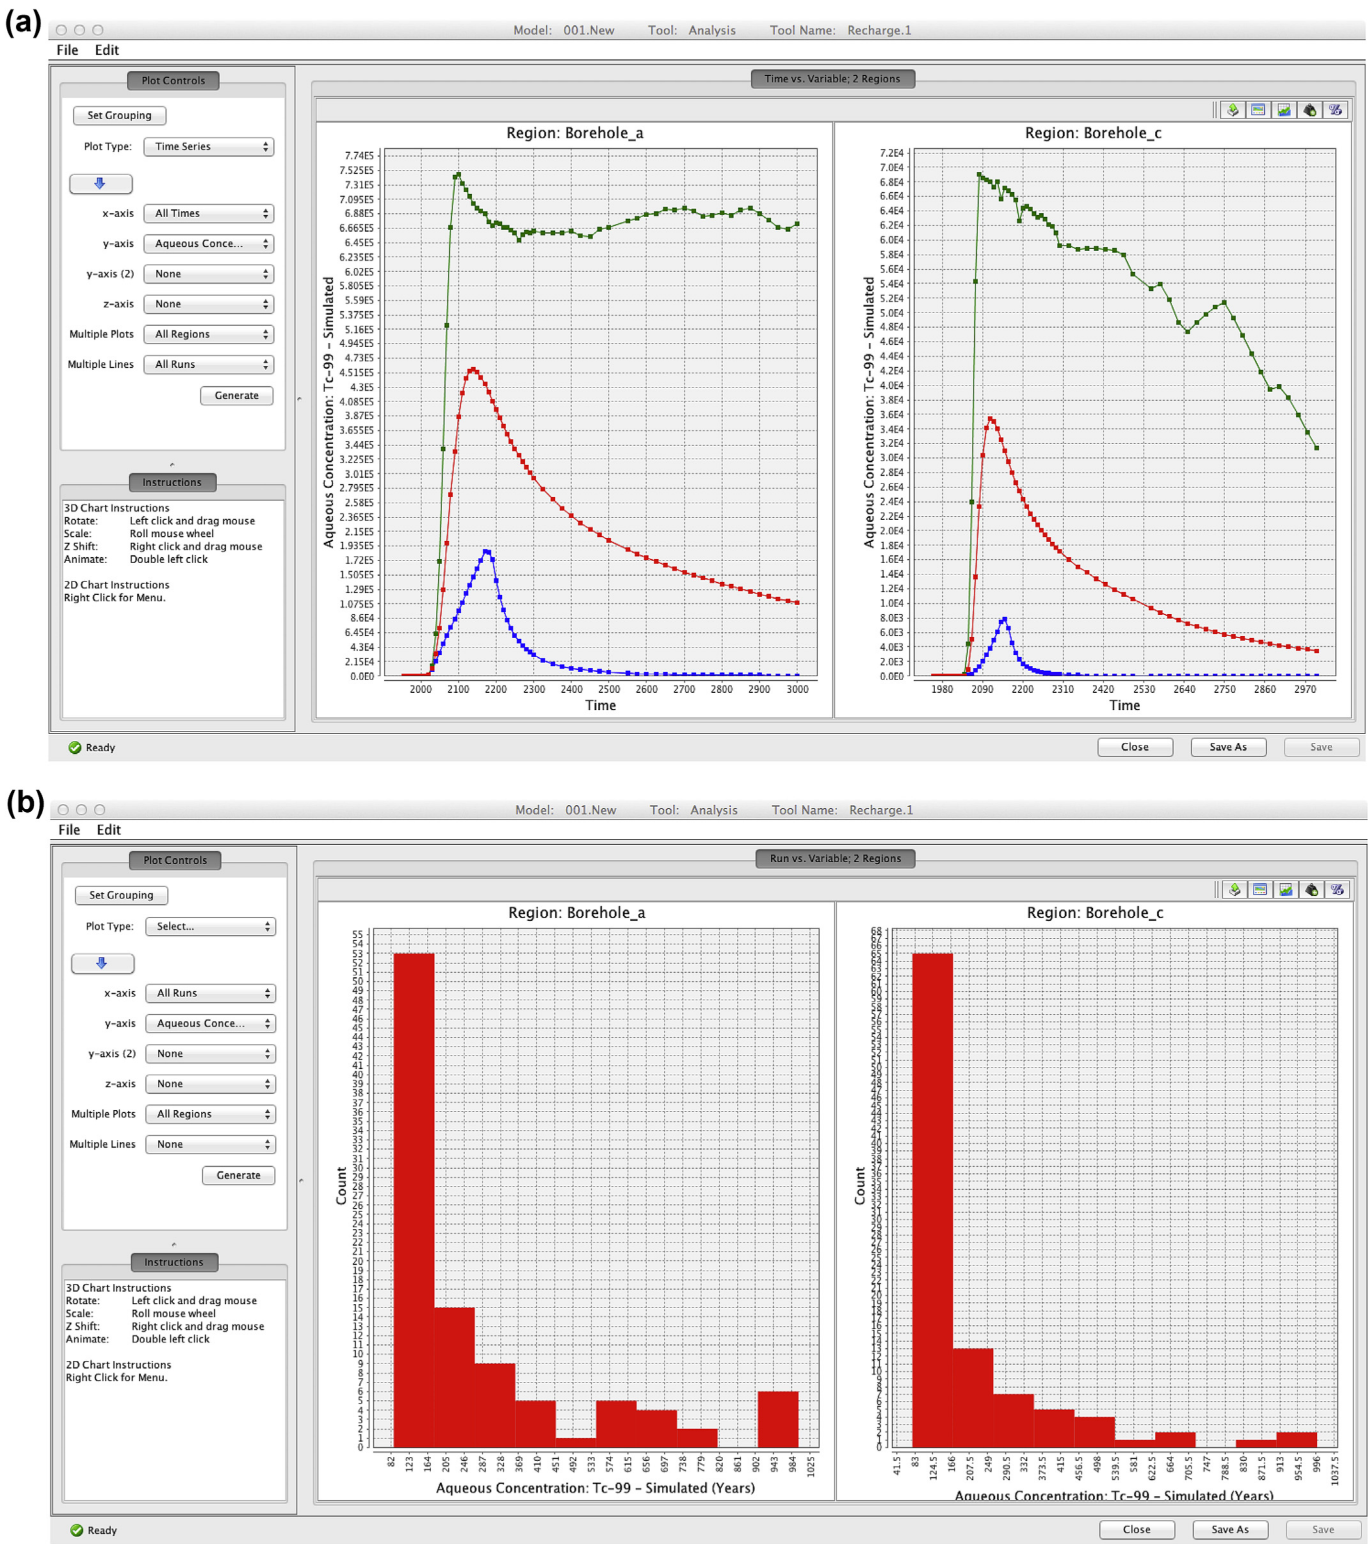

Fig. 10. Screenshot from UQ Toolset showing a) mean and 95% confidence intervals for the ^{99}Tc breakthrough curve at monitoring locations beneath Boreholes A and C; and b) histogram for time to reach the peak concentration at monitoring locations beneath Boreholes A and C.

measured and simulated values of moisture content and concentration, the new parameter estimates showed a significant improvement in matching historical data over initial parameter estimates. Modeling is an iterative process. Improvements in historical data matching are expected as the conceptual model is refined (e.g., boundary conditions, lithofacies distributions), unsaturated hydraulic properties are included in the calibration, and as new capabilities are incorporated into the Akuna toolsets.

The uncertainty analysis presented here was designed for a future condition that would not impact parameter estimates

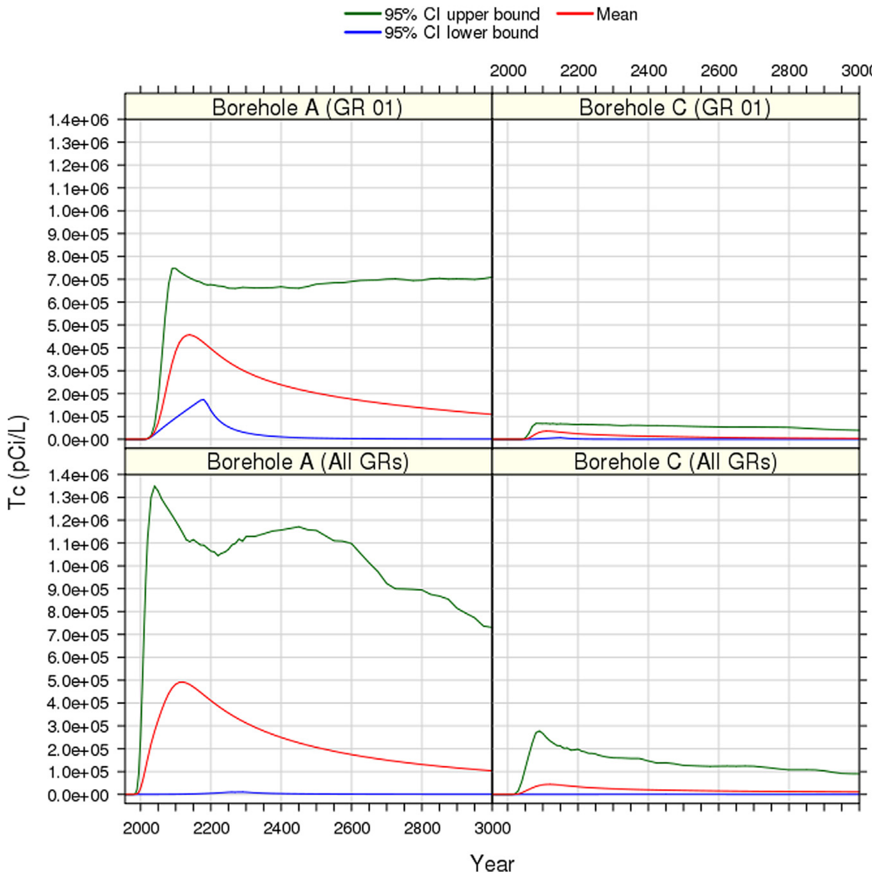

Fig. 11. Breakthrough curves showing mean and 95% confidence intervals at boreholes A and C. Two figures at bottom plot concentration vs. time for all ten geologic realizations (GR), whereas the two figures at top plot the same quantities for a single conceptual model.

obtained from the model calibration. Source terms at BC Cribs also represent a large source of uncertainty, but were not examined in this analysis, since changes in the liquid discharges represented another condition that would likely generate different hydraulic parameter estimates. Identifying uncertainty in the source terms will also be important to identifying potential impacts at the site. Future analyses will examine the individual contributions of source term, geologic conceptual model and parameter uncertainty.

The results of this analysis provide insight on risk associated with different remedial designs impacting the net infiltration rate at the site. The greater range in response for all metrics examined in this analysis emphasizes the importance of examining uncertainty with respect to subsurface heterogeneities, as well as any other sources of uncertainty that may impact mass transport to the water table. Using Akuna to generate breakthrough curves, histograms and scatter plots for UQ once simulations were completed facilitated a rapid analysis and identification of trends.

4.2. Akuna Toolsets

The Akuna modeling framework demonstrated in this paper provides new capabilities – initially targeted at remediation of legacy DOE waste sites, but applicable to many other areas where subsurface flow and transport modeling is needed – for model setup, execution, and analysis, from model calibration through uncertainty analysis. The use of high performance computing, and the accessibility to it that is facilitated by Akuna, allow a user to rapidly develop conceptual and numerical models of a site, and to perform numerical simulations and analyses. A primary focus of this paper was on illustrating Akuna's integrated set of tools that support the full workflow that is needed for subsurface flow and transport modeling, which includes a tightly coupled set of analysis and job launching and monitoring tools that can be used in both serial and parallel computing environments.

Akuna is open-source, cross-platform, and designed to support multiple simulators. It supports seamless exploitation of supercomputing resources and yet can run on a user's desktop. Akuna provides complete tracking of the workflow and can also support collaborative modeling.

The first user release of the ASCEM software will occur in 2013. The ASCEM software will be updated annually, with capabilities that are largely dictated by simulation requirements within the DOE complex. Currently, plans for the 2014 release include a user-friendly UI for reactive geochemistry, unstructured grid generation, and expansion of the capabilities for both the PE and UQ Toolsets. Further integration with WorldWind (2012) is also planned. When integration is fully completed, users will be able to develop a model based on the initial visualization of their site in its actual geographic context, with displays of surface topography and geomorphic features. Interactive placement of a bounding polygon on the map via WorldWind will allow for rapid delineation of the lateral extent of the model domain. By 2015, unstructured grid generation within Akuna will be possible using LaGriT (2012), and Akuna will also provide toolsets to aide in regulatory decisions at waste sites. This will include a Risk Assessment (RA) toolset to assess environmental and health risks, and a Decision Support (DS) toolset, to evaluate and optimize performance measures.

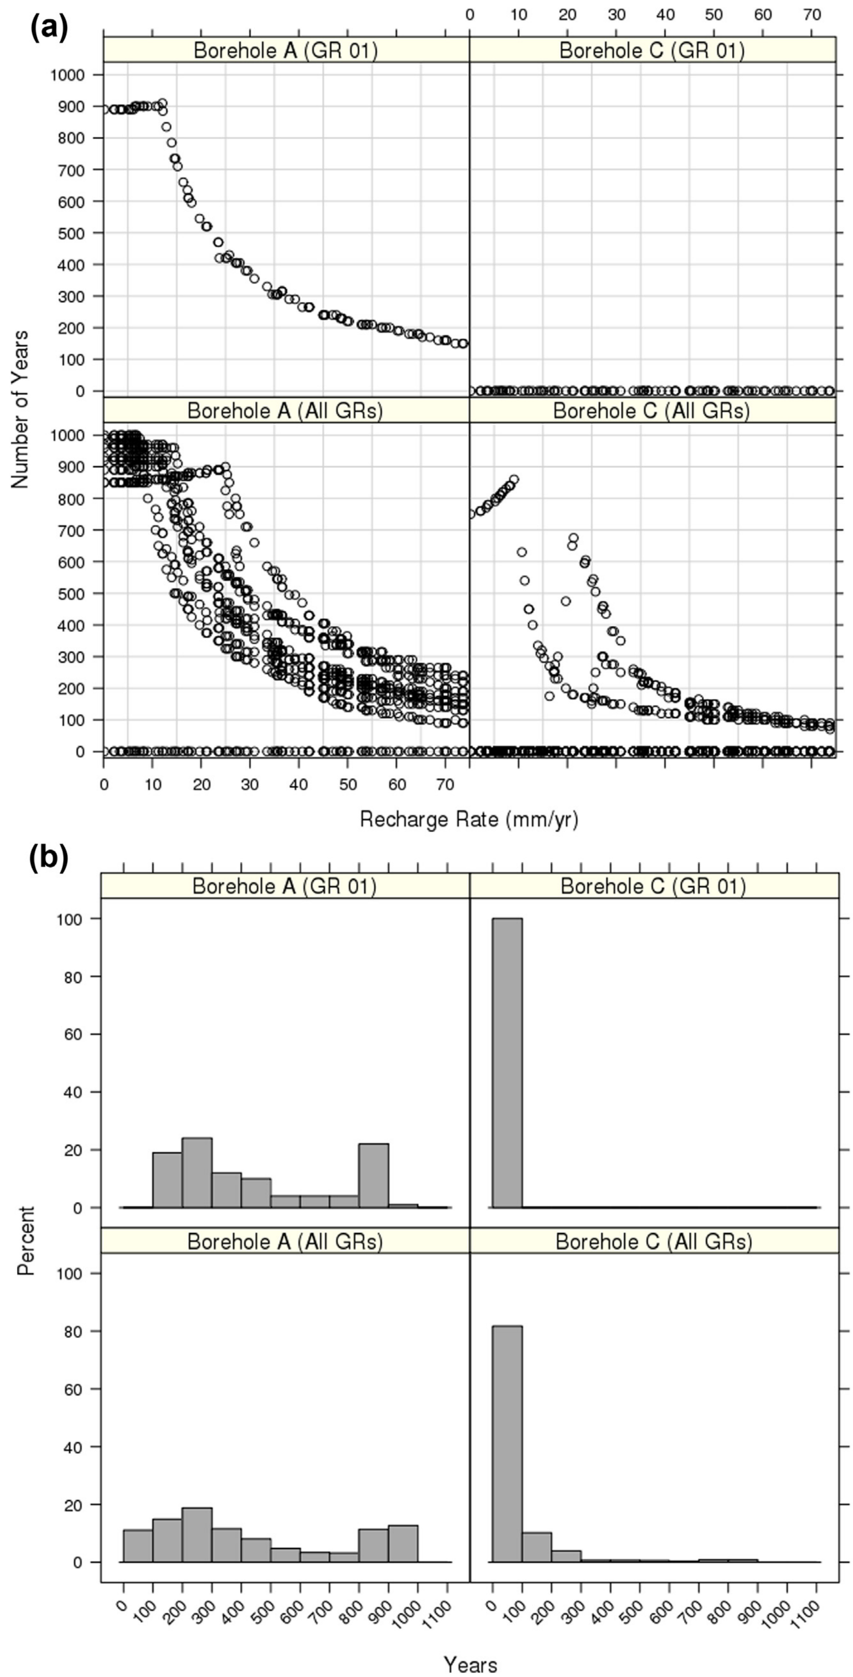

Fig. 12. Number of years above the threshold concentration at Boreholes A and C for the single and all geologic realizations (GR) of the conceptual model for a) recharge rate vs. years; and b) histogram of years.

During the last three years, significant investment has been made in development of Akuna. It is now operational, and the personal-computer software can be downloaded using the information provided in the Software Availability Section of this paper. Current documentation has been gathered as a series of tutorials, which accompanies the software download.

Acknowledgments

This document was prepared by the Advanced Simulation Capability for Environmental Management (ASCEM) Project. Funding for this work was provided by the U.S. Department of Energy Office of Environmental Management to Pacific Northwest National Laboratory, operated by Battelle Memorial Institute for the Department of Energy (DOE) under Contract DE-AC05-76RL01830. The authors are also grateful to three anonymous reviewers, whose comments significantly enhanced the readability of this paper.

References

Arnette, A., 2013. Integrating rooftop solar into a multi-source energy planning optimization model. Appl. Energy 111, 456–467.
Bastin, L., Cornford, D., Jones, R., Heuvelink, G.B.M., Pebesma, E., Stasch, C., Nativi, S., Mazzetti, P., Williams, M., 2013. Managing uncertainty in integrated environmental modelling: the UncertWeb framework. Environ. Model. Softw. 39, 116–134.
Churches, D., Gombas, G., Harrison, A., Maassen, J., Robinson, C., Shields, M., Taylor, I., Wang, I., 2006. Programming scientific and distributed workflow with Triana services. Comput. Pract. Exper. 18 (10), 1021–1037.
De Lara, M., Marinet, V., 2009. Multi-criteria dynamic decision under uncertainty: a stochastic viability analysis and an application to sustainable fishery management. Math. Biosci. 217, 118–124.
Deelman, E., Singh, G., Su, M., Blythe, J., Gil, Y., Kesselman, C., Mehta, G., Vahi, K., Berriman, G.B., Good, J., Laity, A., Jacob, J.C., Katz, D.S., 2005. Pegasus: a framework for mapping complex scientific work- flows onto distributed systems. Sci. Program. 13 (3), 219–237.
Deutsch, C.V., Journel, A.G., 1998. GSLIB: Geostatistical Software Library and User's Guide, second ed. Oxford University Press, New York.
U.S. Department of Energy (DOE), 2008. DOE-EM Engineering and Technology Roadmap: Reducing Technical Uncertainty and Risk in the EM Program. U.S. Department of Energy, Office of Environmental Management, Washington, D.C.
Doherty, J., 2010a. PEST, Model-independent Parameter Estimation—User Manual. with slight additions, fifth ed. Australia, Watermark Numerical Computing, Brisbane.
Doherty, J., 2010b. Addendum to the PEST Manual. Australia, Watermark Numerical Computing, Brisbane.
(accessed 30.11.12.) EarthVision, 2012. http://www.dgi.com/earthvision/evmain.html.
Goovaerts, P., 1997. Geostatistics for Natural Resources Evaluation. Oxford University Press, New York.
Gorton, I., Sivaramakrishnan, C., Black, G., White, S., Purohit, S., Madison, M., Schuchardt, K., 2011. Velo: riding the knowledge management wave for simulation and modeling. In: Proceedings of the 4th International Workshop on Software Engineering for Computational Science and Engineering. Association for Computing Machinery, New York, pp. 32–40.
(accessed 30.11.12.) Groundwater Modeling System (GMS), 2012. http://chl.erdc.usace.army.mil/gms.
Guber, A.K., Pachepsky, Y.A., Van Genuchten, M.Th., Rawls, W.J., Simunek, J., Jacques, D., Nicholson, T.J., Cady, R.E., 2006. Field-scale water flow simulations using ensembles of pedotransfer functions for soil water retention. Vadose Zone J. 5, 234–247. http://dx.doi.org/10.2136/vzj2005.0111.
Harbaugh, A.W., 2005. MODFLOW-2005, the U.S. Geological survey modular ground-water model – the ground-water flow process. U.S. Geol. Surv. Tech. Methods 6-A16, 6.
Icenhower, J.P., Qafoku, N., Martin, W.J., Zachara, J.M., 2008. The Geochemistry of Technetium: a Summary of the Behavior of an Artificial Element in the Natural Environment. Report PNNL-18139. Pacific Northwest National Laboratory, Richland, Washington.
Jebaraj, S., Iniyan, S., 2006. A review of energy models. Renew. Sustain. Energy Rev. 10 (4), 281–311.
Kincaid, C.T., Eslinger, P.W., Aaberg, R.L., Miley, T.B., Nelson, I.C., Strenge, D.L., Evans Jr., J.C., 2006. Inventory Data Package for Hanford Assessments. PNNL-15829. Pacific Northwest National Laboratory, Richland, WA.
(accessed 30.11.12.) LaGriT, 2012. https://lagrit.lanl.gov/.
Levenberg, K., 1944. A method for the solution of certain non-linear problems in least squares. Q. Appl. Math. 2, 164–168.
Li, Y.P., Huang, G.H., Li, M.W., 2014. An integrated optimization modeling approach for planning emission trading and clean-energy development under uncertainty. Renew. Energy 61, 31–46.
Ludascher, B., Altintas, I., Berkley, C., Higgins, D., Jaeger, E., Jones, M., Leen, E.A., Tao, J., Zhao, Y., 2006. Scientific workflow management and the Kepler system. Concurr. Comput. Pract. Exper. 18 (10), 1039–1065.
Marquardt, D.W., 1963. An algorithm for least-squares estimation of non-linear parameters. SIAM J. Appl. Math. 11 (2), 431–441.
Mortberg, U., Haas, J., Zetterberg, A., Franklin, J.P., Jonsson, D., Deal, B., 2013. Urban ecosystems and sustainable urban development-analysing and assessing interacting systems in the Stockholm region. Urban Ecosyst. 16 (4), 763–782.
Murray, C.J., 1994. Identification and 3-D modeling of petrophysical rock types. In: Yarus, J.M., Chambers, R.L. (Eds.), Stochastic Modeling and Geostatistics. Am. Assoc. Pet. Geol., 323–337. Tulsa, Okla., AAPG Computer Applications in Geology, No. 3.
Oinn, T., Addis, M., Ferris, J., Marvin, D., Senger, M., Greenwood, M., Carver, T., Glover, K., Pocock, M.R., Wipat, A., Li, P., 2004. Taverna: a tool for the composition and enactment of bioinformatics workflows. Bioinformatics 20 (17), 3045–3054.
Petrel, 2012. Petrel E&P Software Platform (accessed 30.11.12.). http://www.slb.com/services/software/geo/petrel.aspx.
Plale, B., Gannon, D., Brotzge, J., Droegemeier, K., Kurose, J., McLaughlin, D., Wilhelmson, R., Graves, S., Ramamurthy, M., Clark, R.D., Yalda, S., Reed, D.A., Joseph, E., Chandrasekar, V., 2006. CASA and LEAD: adaptive cyber infrastructure for real-time multiscale weather forecasting. Computer 39 (11), 56–64.
Poeter, E.P., Hill, M.C., Banta, E.R., Mehl, S., Christensen, S., 2005. UCODE_2005 and six other computer codes for Universal sensitivity analysis, calibration, and uncertainty evaluation. U.S. Geol. Surv. Tech. Methods 6-A11, 283p.
Scheibe, T.D., Fang, Y., Murray, C.J., Roden, E.E., Chen, J., Chien, Y.-J., Brooks, S.C., Hubbard, S.S., 2006. Transport and biogeochemical reaction of metals in a physically and chemically heterogeneous aquifer. Geosphere 2 (4), 220–235.
Schoof, L.A., Yarberry, V.R., 1994. EXODUS II: a Finite Element Data Model. Sandia Report No. SAND92-2137. Sandia National Laboratories, Albuquerque, New Mexico.
Serne, R.J., Ward, A.L., Um, W., Bjornstad, B.N., Rucker, D.F., Lanigan, D.C., Benecke, M.W., 2009. Electrical Resistivity Correlation to Vadose Zone Sediment and Pore-water Composition for the BC Cribs and Trenches Area. Report PNNL-17821. Pacific Northwest National Laboratory, Richland, Washington.
Stainforth, D.A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D.J., Kettleborough, J.A., Knight, S., Martin, A., Murphy, J.M., Piani, C., Sexton, D., Smith, L.A., Spicer, R.A., Thorpe, A.J., Allen, M.R., 2006. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. Nature 433 (7024), 403–406.
Truex, M.J., Johnson, T.C., Strickland, C.E., Peterson, J.E., Hubbard, S.S., 2013. Monitoring vadose zone desiccation with geophysical methods. Vadose Zone J.. http://dx.doi.org/10.2136/vzj2012.0147.
Turuncoglu, U.U., Dalfes, N., Murphy, S., DeLuca, S., 2013. Toward self-describing and workflow integrated Earth system models: a coupled atmosphere-ocean modeling system application. Environ. Model. Softw. 39, 247–262.
U.S. Department of Energy (DOE/RL), 2008. Excavation-based Treatability Test Plan for the BC Cribs and Trenches Area Waste Sites. U.S. Department of Energy, Richland Operations, Richland, Washington. DOE/RL-2007-15, Rev. 0.
VisIt, 2012 (accessed 30.11.12.). https://wci.llnl.gov/codes/visit/.
Vistas, Groundwater, 2012 (accessed 30.11.12.). http://www.groundwater-vistas.com/index.html.
Visual MODFLOW, 2012 (accessed 30.11.12.). http://www.swstechnology.com/groundwater-modeling-software/visual-modflow-flex.
Ward, A.L., Serne, R.J., Benecke, M.W., 2009. Development of a conceptual model for vadose zone transport of Tc-99 at Hanford's BC Cribs and the screening of remedial alternatives. In: Waste Management 2009: Waste Management for the Nuclear Renaissance, March 1–5, Tucson, Arizona.
Wellman, D.M., Triplett, M.B., Freshley, M.D., Truex, M.J., Gephart, R.E., Johnson, T.C., Chronister, G., Gerdes, K.D., Chamberlain, S., Marble, J., Ramirez, R., Charboneau, B.L., Morse, J.G., Eberline, S., 2011. Deep vadose zone applied field Research Center: transformational technology development for environmental remediation. In: Waste Management Conference, February 27 – March 3, 2011, Phoenix, AZ.
White, M.D., Oostrom, M., 2000. STOMP – Subsurface Transport Over Multiple Phases: Theory Guide. Report PNNL-12030. Pacific Northwest National Laboratory, Richland, Washington.
White, M.D., Oostrom, M., 2006. STOMP – Subsurface Transport Over Multiple Phases, Version 4: User's Guide. Report PNNL-15782. Pacific Northwest National Laboratory, Richland, Washington.
WorldWind, 2012 (accessed 30.11.12.). http://worldwind.arc.nasa.gov/java/.