ASCEM Phase I Demonstration

December 16, 2010

ASCEM-SITE-102010-1

United States Department of Energy

Environmental Management

safety ❖ performance ❖ cleanup ❖ closure

Susan Hubbard (LBNL)
Boris Faybishenko (LBNL)
Mark Freshley (PNNL)
Deb Agarwal (LBNL)
John Bell (LBNL)
Wes Bethel (LBNL)
Miles Denham (SRNL)
Greg Flach (SRNL)
Vicky Freedman (PNNL)
Glenn Hammond (PNNL)
David Higdon (LANL)
Jennifer Horsman (LBNL)
Elizabeth Keating (LANL)
Peter Lichtner (LANL)
Laura Monroe (LANL)
Phil Moore (SRNL)
David Moulton (LANL)
George Pau (LBNL)
Daniel Schep (SRNL)
Karen L Schuchardt (PNNL)
Roger Seitz (SRNL)
Arie Shoshani (LBNL)
Nic Spycher (LBNL)
Paul Weber (LANL)

CONTRIBUTORS:

Benjamin Andre (LBNL)
Erin Barker (PNNL)
Markus Berndt (LANL)
Gary Black (PNNL)
Mike Buksas (LANL)
Joann Campbell (LANL)
Neil Carlson (LANL)
Karl Castleton (PNNL)
Anthony Drummond (LBNL)
Rao Garimella (LANL)
Ian Gorton (PNNL)
Luke Gosink (PNNL)
Konstantin Lipnikov (LANL)
Richard Mills (ORNL)
William Perkins (PNNL)
Lori Pritchett-Sheats (LANL)
Mark Rockhold (PNNL)
Alex Romosan (LBNL
Douglas Sassen (LBNL)
Mudita Singhai (PNNL)
Chandrika Sivaramakrishnan (PNNL)
Kirsten Fagnan (LBNL)
Michael Lijewski (LBNL)
Mudita Singhai (PNNL)
Carl Steefel (LBNL)
Will Stringfellow (LBNL)
Glenn Taylor (SRNL)
Yun Wei (LLNL)
Mark Williams (PNNL)
Signe Wurstner (PNNL)
Steve Yabusaki (PNNL)

DISCLAIMER

Printed in the United States of America

Prepared for
U.S. Department of Energy

Authorization:

Mark Freshley, Pacific Northwest National Laboratory
ASCEM Site Applications Thrust Lead

December 16, 2010
Date

Dr. Juan Meza, Lawrence Berkeley National Laboratory
ASCEM Technical Integration Manager

December 16, 2010
Date

Dr. Paul Dixon, Los Alamos National Laboratory
Multi-Lab ASCEM Program Manager

December 16, 2010
Date

Concurrence:

Dr. Mark Williamson, ASCEM Program Manager
EM-32 Groundwater & Soil Remediation

01/11/11
Date

Kurt Gerdes, EM-32
Director for Groundwater & Soil Remediation

1/11/11
Date

Table of Contents

LIST OF ABBREVIATIONS

API	Application Program Interface
ARARs	Applicable or relevant and appropriate requirements
AMR	Adaptive Mesh Refinement
ASCEM	Advanced Simulation Capability for Environmental Management
ASC	Advanced Simulation & Computing of NNSA
ASCR	Advanced Scientific Computing Research
BER	Biological and Environmental Research
CGNS	CFD (Computational fluid dynamics) General Notation System
CSV format	Comma-separated-values format
DOE	Department of Energy
EM	Office of Environmental Management
GUI	Graphical User Interface
HPC	High-Performance Computing
I/O	Input / Output
LaGriT	Los Alamos Grid Toolkit
MDL	method detection limit
MFD	Mimetic Finite Difference
MNA	Monitored Natural Attenuation
MOAB	Mesh-Oriented database
MPC	Multi-Process Coordinator
NERSC	National Energy Research Scientific Computing Center
PA	Performance Assessment
PDF	Probability density function
PE	Parameter estimation
PETSc	Portable, Extensible Toolkit for Scientific Computing
PK	Process Kernel
PQL	practical quantification limit
SciDAC	Scientific Discovery through Advanced Computing
SRS	Savannah River Site
SFA	Scientific Focus Area of DOE-BER
STK	Sierra Toolkit
UI	User interface
UQ	Uncertainty quantification

OPEN SOURCE SOFTWARE LIBRARIES AND URLs:

VisIt	http://www.llnl.gov/visit
PEST	http://www.sspa.com/pest
PSUADE	https://computation.llnl.gov/casc/uncertainty_quantification
Google Maps	http://code.google.com/apis/maps/documentation/javascript/
FLOT	http://code.google.com/p/flot/
Mediawiki (Velo)	http://www.mediawiki.org
Semantic Mediawiki	http://semantic-mediawiki.org
Subversion	http://subversion.apache.org
Python	http://www.python.org/
Stk_Mesh	http://trilinos.sandia.gov/release_notes-10.2.html
PostegreSQL	http://www.postgresql.org/
Trilinos	http://trilinos.sandia.gov/
MOAB	http://trac.mcs.anl.gov/projects/ITAPS/wiki/MOAB
HDF5 format	http://www.hdfgroup.org/HDF5/doc/H5.intro.html

1. EXECUTIVE SUMMARY

The U.S. Department of Energy (DOE) Office of Environmental Management (DOE-EM), Technology Innovation and Development (EM-30), is supporting development of the Advanced Simulation Capability for Environmental Management Initiative (ASCEM). ASCEM is an emerging state-of-the-art scientific approach and software infrastructure for understanding and predicting contaminant fate and transport in natural and engineered systems. The modular and open-source high performance computing tool will facilitate integrated approaches that enable standardized assessments of performance and risk for EM cleanup and closure decisions. The initiative is organized into three Thrust Areas: 1) Multi-Process High Performance Computing (HPC), 2) Platform and Integrated Toolsets, and 3) Site Applications. As part of the development process, a series of demonstrations are planned to:

- Test several developing ASCEM components
- Engage end users in applications
- Illustrate ASCEM progress
- Provide feedback to developers.

The first (Phase I) demonstration, largely undertaken from September through December, 2010, focused on illustrating individual (stand-alone) ASCEM capabilities. Future demonstrations will focus on integrating ASCEM capabilities and using ASCEM to address DOE-EM remediation challenges.

The Phase I Demonstration was designed to provide an early snapshot of advances associated with four specific components of ASCEM: Data Management; Visualization; Uncertainty Quantification (UQ); and High Performance Computing (HPC). Leveraging and integration of existing open source software are central principles of ASCEM. This approach is expected to result in lower overall project costs, faster development times, and broad community acceptance. Leveraging occurs through incorporation of advances developed by the DOE Office of Science through their Biological and Environmental Research (BER) and Advanced Scientific Computing Research (ASCR) programs, as well as the DOE National Nuclear Security Agency (NNSA) Advanced Simulation and Computing (ASC) Program. Coordination of the Phase I effort was facilitated through a working group mechanism and the use of common datasets associated with a contaminated site, the SRS River Site (SRS) F-Area.

For the Phase I Demonstration, the Data Management component adapted and implemented open-source, web-based tools to allow users to easily import, browse, filter, graph, query, and output datasets common to environmental remediation investigations. Capabilities were developed to allow visualization of these and other datasets, including depositional information, hydrostratigraphic surfaces, and the evolution of contaminant plumes. ASCEM capabilities were also developed to allow a user to perform uncertainty quantification using a variety of different analysis approaches within a graphical user interface (GUI). Prototypes of selected toolsets within the ASCEM Multi-Process HPC simulator, called Amanzi, were developed and tested on laptops and desktops running both Linux and Mac OS X, and on several supercomputers including the Cray XT4 system at NERSC using 256 cores and the Hopper XE6 at NERSC using 2304 cores. Both unstructured and structured mesh approaches were used to simulate geochemical and hydrological processes using F-Area data and information.

Two supplementary efforts were also undertaken to advance new ASCEM capabilities and engage different end user communities. These advances include use of an adaptive mesh refinement approach to more efficiently and accurately simulate potential release from the degradation of closed tanks, and development of approaches to quickly visualize simulation and UQ output. An ASCEM model setup tool was developed in conjunction with data from the Hanford vadose zone to translate and visualize conceptual model information into numerical model input within the same computing environment.

This document illustrates the significant progress that has been realized in developing several critical ASCEM components during the first year of the project. An accompanying PowerPoint file provides additional figures and movies to support this Demonstration report. While many of the individual ASCEM components will lead to performance and flexibility that will exceed what is available today, the most significant contribution of ASCEM is expected to be its integrated framework and associated computationally efficient, open-source, modular, portable, and accessible software codes. The development of a process-based computational framework that can be easily and consistently used across the DOE EM complex is expected to improve cleanup efficacy and decrease overall costs associated with the DOE legacy waste stewardship obligation.

INTRODUCTION AND PROBLEM STATEMENT

The mission of DOE-EM is to complete the safe cleanup of the environmental legacy from the nation's five decades of nuclear weapons and nuclear energy research and production. Contamination has been introduced into complex subsurface environments by intentional disposal through injection wells, disposal facilities, and evaporation or seepage ponds; and by accidental spills and leaks from waste storage tanks, basins, and transfer lines. The subsurface environment is characterized by multiple hydrological, geochemical, and microbiological processes occurring at different scales, significant heterogeneity, and daunting measurement and observational constraints (DOE/SC-0123, 2010). This cleanup effort is one of the most complex and technically challenging in the world; it is projected to be ongoing for decades (DOE, 2000) and to cost between $265-305 billion to complete (DOE/CF-028, Volume 5).

Recent workshops and panels have concluded that gaps in the technical foundation supporting environmental decisions have led to ineffective remediation (Congress, 2006) and that the complexity and magnitude of the DOE environmental problem justifies a long-term investment in environmental remediation science and technology (DOE 2008a,b, DOE 2010, NRC 2009). Based on these and other workshop reports, the DOE EM-32 Office of Groundwater and Soil Remediation Program identified several key needs, including the development of numerical tools that can accurately predict the long-term behavior of subsurface contaminant plumes and degradation of engineered materials used for waste disposal. Currently, no single process-based computational framework is used across the DOE EM complex in a consistent manner, as is needed to enable standardized assessments of performance and risk associated with EM cleanup and closure activities.

1.1 ASCEM Goals and Organization

ASCEM is developing a state-of-the-art scientific tool and transformational approach for integrating data, software, and scientific understanding. The initiative is combining petaflop supercomputing capabilities, with new and open-source high performance computing modeling application platforms, data analysis and integration approaches, and evolving understanding of subsurface hydrological-biogeochemical processes in order to improve subsurface contaminant fate and transport simulations as needed to support risk-informed environmental remediation and waste management decisions. ASCEM will also provide other DOE programs and the scientific community with a powerful open-source code and approach that should be applicable to a variety of subsurface flow and transport problems. ASCEM is envisioned to be updated and augmented as new insights and approaches are developed through BER, ASCR, and other research programs.

ASCEM is organized into three Thrust Areas: 1) Multi-Process High Performance Computing (HPC), 2) Platform and Integrated Toolsets, and 3) Site Applications. The relationship between the Thrust Areas is illustrated in Figure 1. The HPC Thrust includes meshing approaches, new solvers for multi-physics coupled processes, advanced methods of discretization in time and space, capabilities to select and coordinate the use of problem-specific processes, and application-programming interfaces. The Platform Thrust includes an integrated software infrastructure to facilitate model setup and analysis, parameter estimation and uncertainty quantification, risk assessment and decision support, information and data management, and visualization in a consistent and flexible user interface and modeling workflow. The Site Application Thrust coordinates and implements demonstrations through "working groups" to provide data and feedback to developers and to ensure that the ASCEM software infrastructure is developed in a manner that will engage users and benefit DOE-EM's remediation obligations.

A series of demonstrations will be performed in a phased manner to correspond with development of the Platform and HPC components and with ASCEM releases. The ASCEM demonstrations are designed to advance, test, and illustrate ASCEM capabilities as well as engage end-users. This first demonstration is focused on illustrating individual (stand-alone) capabilities within the HPC and Platform Thrusts, whereas later demonstrations will focus on integrating ASCEM capabilities and illustrating how those capabilities can be used to address DOE-EM problems. While the most significant contribution of ASCEM is expected to be the integrated framework with its associated computationally-efficient, open-source, and modular characteristics, in many cases ASCEM will also advance individual components whose performance or flexibility will exceed what is available today. In this report, the Phase I ASCEM Demonstration is described.

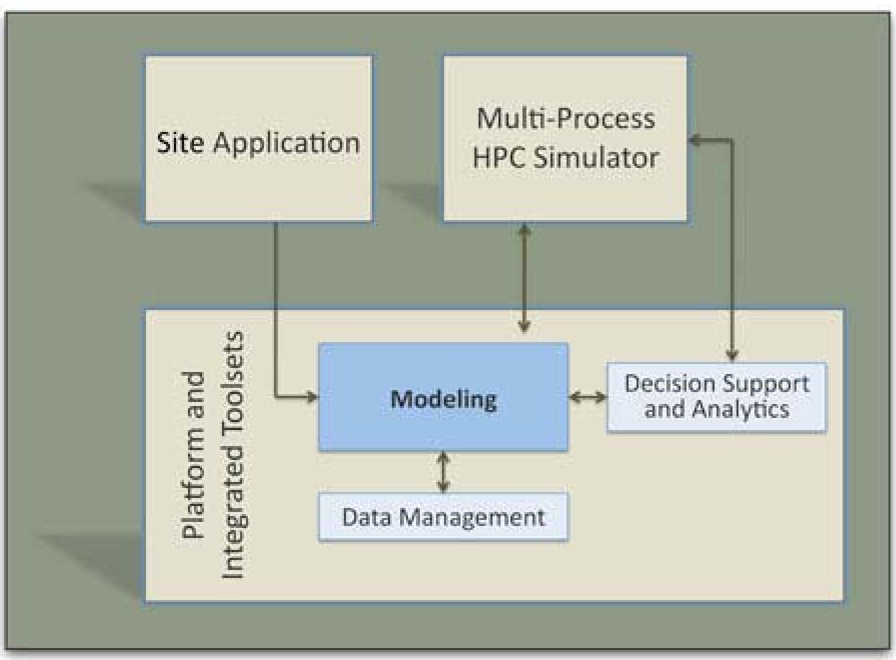

Figure 1. Illustration of the relationship between the three ASCEM Technical Thrust Areas: HPC Simulator, Platform and Toolsets, and Site Applications.

1.2 Phase I (2010) Demonstration

The Phase I Demonstration illustrates advances in individual capabilities within the HPC and Platform thrusts. Development of ASCEM capabilities was initiated in 2010 after the software requirements were defined based on DOE EM site needs; this demonstration focuses on only a subset of developed capabilities. Decisions about which early components of ASCEM to demonstrate were made by considering: (a) feasibility of code/tool advancement within a short time period (primarily October 1-December 10, 2010); (b) value as a vehicle for communicating ASCEM advances to a variety of stakeholders; (c) availability and characteristics of available datasets for testing and developing ASCEM components.

The salient features of candidate sites and the selection process for the Phase 1 Demonstration were summarized in the "*Site Selection*" Document (ASCEM-SITE-091310-01, 2010). The selection process led to a decision by the ASCEM team to focus the ASCEM Phase I demonstration on the contaminated SRS F-Area around the seepage basin and in particular on the following Platform and HPC components.

PLATFORM THRUST COMPONENTS:

- **Data Management**, including development of data import, organization, and query tools using a map-based interface, with an initial focus on contaminant concentration and hydrostratigraphic variables
- **Visualization** of the hydrostratigraphy, water table, topography, wells, and migration of contaminant plumes through aquifers over time
- **Uncertainty Quantification**, including evaluation of strategies and development of tools to evaluate the sensitivity of model output to parameter suites

HPC THRUST

- **HPC** capabilities, including the development of: a parallel, unstructured mesh capability; capabilities to perform three-dimensional parallel flow simulation; capabilities to simulate transport of a non-reactive contaminant**;** and development of a prototype of the Reaction Toolset that includes aqueous speciation, mineral precipitation and dissolution, and sorption.

Although the main contribution of ASCEM is envisioned to be the integrated nature of the framework, the Phase I demonstration focuses on advancing individual components using common datasets from a specific site. Future demonstrations will focus on illustrating integration of expanded capabilities into new versions of ASCEM. The Phase I

Demonstration was designed to be tractable and facilitate the eventual transfer of insights and methods being developed through other DOE EM and BER-supported projects at the SRS F Area into ASCEM.

The remainder of this report is organized by the demonstration components. Descriptions of the datasets used for the Phase I Demonstration, focused on the SRS F-Area, are given in Section 4. Descriptions of the specific goals, approach, and accomplishments for each of the four demonstration components (data management, visualization, uncertainty quantification, and high performance computing) that use F-Area data are provided in Sections 5–8. In addition to the F-Area ASCEM Phase I Demonstration, two supplementary Demonstration activities were initiated that were beyond the scope of the defined Phase I Demonstration deliverables: the Hanford Deep Vadose Zone and the Waste Tank Performance Assessment. These supplementary activities were viewed as opportunities to engage additional communities, to develop new ASCEM capabilities using datasets other than from the F-Area, and to strategically position ASCEM for the 2011 Phase II and future demonstrations. Progress made by these two additional working groups is described in Section 9, followed by a summary discussion provided in Section 10.

A PowerPoint file accompanies this document and provides additional figures and videos to illustrate ASCEM Phase I Demonstration advances.

1.3 SRS F-Area Contamination and Remediation

The SRS is located in south-central South Carolina, near Aiken, approximately 100 miles from the Atlantic Coast. It covers an area of approximately 800 square kilometers (300 square miles) and contains facilities constructed in the early 1950s to produce special radioactive isotopes (e.g., plutonium and tritium) for the U.S. nuclear weapons stockpile. SRS has ~172 × 10^6 m^3 of groundwater, soil, and debris contaminated with metals, radionuclides, and organics (NRC, 2000) as a result of on-site disposal practices.

The SRS F-Area Seepage Basins (located in the north-central portion of SRS) consist of three unlined, earthen surface impoundments that received ~7.1 billion liters (1.8 billion gallons) of acidic, low-level waste solutions. The acidic liquid waste (average influent pH of 2.9) originated from the processing of irradiated uranium in the F-Area Separations facility from 1950 through 1989. The plume currently extends from the basins to ~600 meters downgradient at a stream (Figure 2), and contains a large number of contaminants. Based on risk to potential receptors, the most hazardous contaminants are uranium isotopes, Sr-90, I-129, Tc-99, tritium, and nitrate. Groundwater is currently acidic, with pH values as low as 3.2 near the basins. As a result, the sediments that underlie the F-Area have been exposed to acidic solutions for many decades.

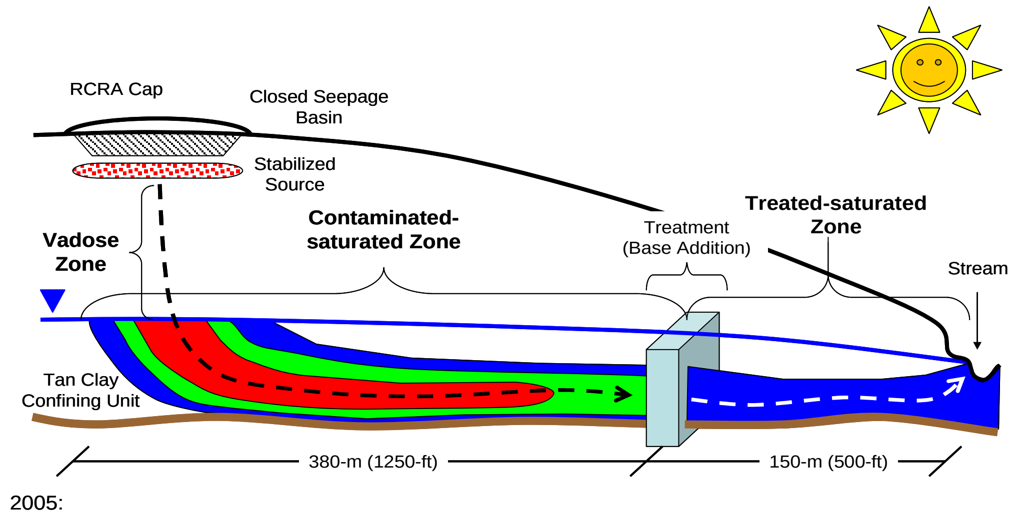

Figure 2. Schematic of the F-Area, showing the basins, the base addition treatment, the stream receptor, and the significant groundwater plume.

The basins were closed and capped in 1991. A pump-and-treat remediation system began operation in 1997, and it was replaced in 2004 by a hybrid funnel-and-gate system installed about 380 meters downgradient of the basin but upgradient from the receptor stream (see Figure 2). Alkaline solutions are now being injected into the gates in an attempt to neutralize the acidic groundwater downgradient of the seepage basins. Monitored Natural Attenuation (MNA) is a desired closure strategy for the site, based on the premise that rainwater will eventually neutralize the lingering mineral surface acidity, causing an increase in pH, and stimulating natural immobilization of U in the trailing end of the plume. If the natural pH neutralization upgradient from the treatment system is insufficient, additional enhanced neutralization will be required. Critical to assessing the *in situ* treatment requirements over a long time frame is the development of an understanding of the long-term H^+ and U sorption behavior at the site.

1.4 SRS F-Area Leveraging

Ongoing EM and BER-supported activities at SRS F-Area offer potential for significant leveraging to ASCEM. DOE EM-32 considers the F-Area Site to be a primary applied field research site for testing attenuation-based remedies for metals and radionuclides in groundwater. The primary goal is to develop tools, approaches to technical issues, and guidance that facilitate the use of attenuation-based remedies for metals and radionuclides in groundwater. The ongoing field research includes measurement of aquifer properties *in situ*, development of approaches to identifying reactive facies, and development and testing of attenuation-based remedies for I-129, Sr-90, and uranium. Research is conducted by members of the technical working group that oversees the EM-32 effort, as well as outside parties representing academia and industry.

The DOE Office of Science Subsurface Biogeochemistry Program of BER is supporting the basic research "Predicting Plume Mobility Challenge' of the LBNL Scientific Focus Area (SFA). This project, which is conducted jointly with SRNL at the F-Area, explores the impact of a pH gradient and the concept of reactive facies as an organizing principle to integrate laboratory and field information about transport-relevant properties and mechanisms for predictions of uranium and I-129 at the plume scale. The LBNL Challenge includes an extensive laboratory component focused on developing: (a) facies-specific surface complexation models (Dong et al., 2010); (b) approaches to identify and spatially distribute reactive facies using geophysical data (Sassen et al., 2010); (c) process model development, sensitivity analysis, and mechanistic reactive transport model development (Spycher et al, 2010); and (d) a formal evaluation of the benefit of increasing complexity on successful predictions of contaminant mobility over stewardship time frames.

Coordination between these ongoing activities with ASCEM offers the opportunity for significant and mutually beneficial leveraging. The ongoing EM-32 and BER-funded activities provide a conceptual understanding of relevant DOE problems and advanced insights and datasets that can be used by ASCEM Phase I as well as future demonstrations. In turn, ASCEM is developing HPC and Platform integrated infrastructure, which will provide capabilities critical for the success of the DOE-EM soil and groundwater remediation activities and that are beyond the reach of present-day technologies.

2. CONCEPTUAL HYDROGEOLOGICAL MODEL AND GEOCHEMICAL PROCESSES

Because of extensive site characterization and monitoring, the F-Area has many databases available that can provide data for developing and testing ASCEM capabilities. In particular, there is a wealth of hydrogeological and geochemical datasets available for use by the of which have been monitored since the 1980s (Savannah River Site, 2004a, 2004b). A Geographic Information System (GIS) database also exists for the site, and hydrofacies have already been identified for the F-Area (Jean et al., 2004). Examples of some of the available data at the F-Area are shown in Figure 3.

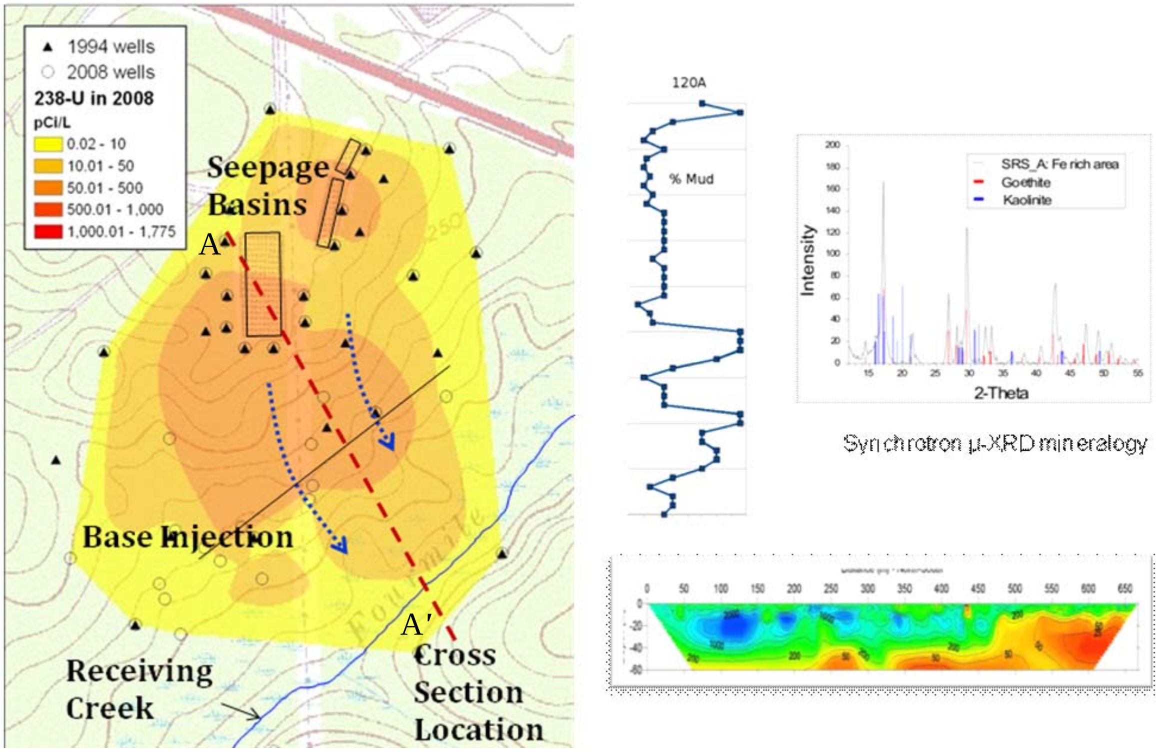

Figure 3. Left: Distribution of a U-238 plume in the groundwater near the F-Area Seepage Basins, where the distance between the basins and the receiving creek is ~0.5 km; Right: examples of available datasets beyond concentration data, including borehole lithology logs, geophysical data, and mineralogical characterization data. The dashed line on the map corresponds to the A-A' cross section shown in Figure 5.

2.1 Site Hydrogeology

The F-Area Seepage Basins site is located in the Tertiary Eocene sediments of the Atlantic Coastal Plain that underlie the F-Area. Figure 4 summarizes the main physical unit classifications at the F-Area, including:

- Lithostratigraphy, which is used to describe sediment units that are defined on the basis of distinctive and dominant sediment characteristics (such as sand, silt, or clay).
- Depositional Environments, which describe the environmental setting when the sediments were deposited (such as barrier beach or shoreface).
- Hydrostratigraphy, which define units that have common hydrological characteristics. The hydrostratigraphy at the site is conceptualized to be relatively simple, consisting of gently dipping Atlantic Coastal Plain unconsolidated and semi-consolidated units comprised of sands and clays. The shallowest hydrostratigraphic unit considered in this demonstration is the Upper Three Runs Aquifer, which consists of an upper aquifer zone (UAZ), a Tan Clay Confining Zone (TCCZ), and a

lower aquifer zone (LAZ). Beneath the LAZ is the Gordon confining unit (GCU) and beneath that, the Gordon aquifer.

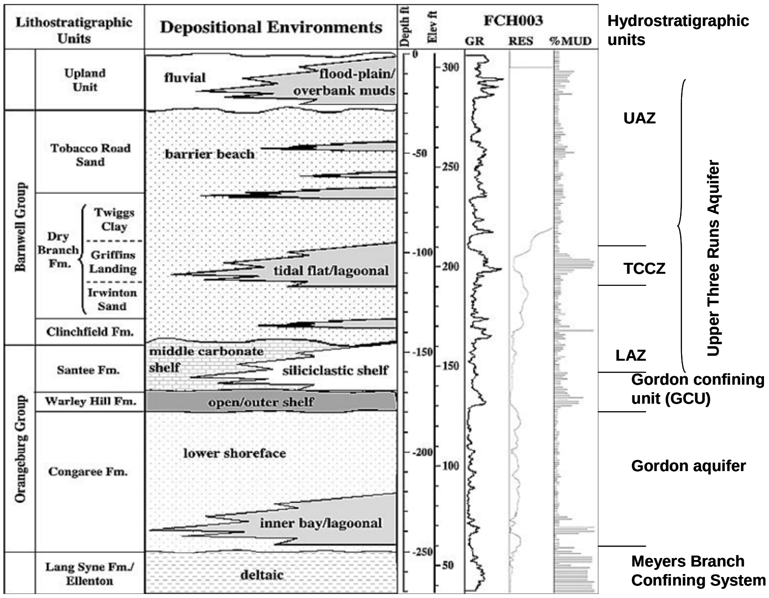

Figure 4. Lithostratigraphic and hydrostratigraphic units and associated depositional environments beneath the study area (Jean et al., 2004).

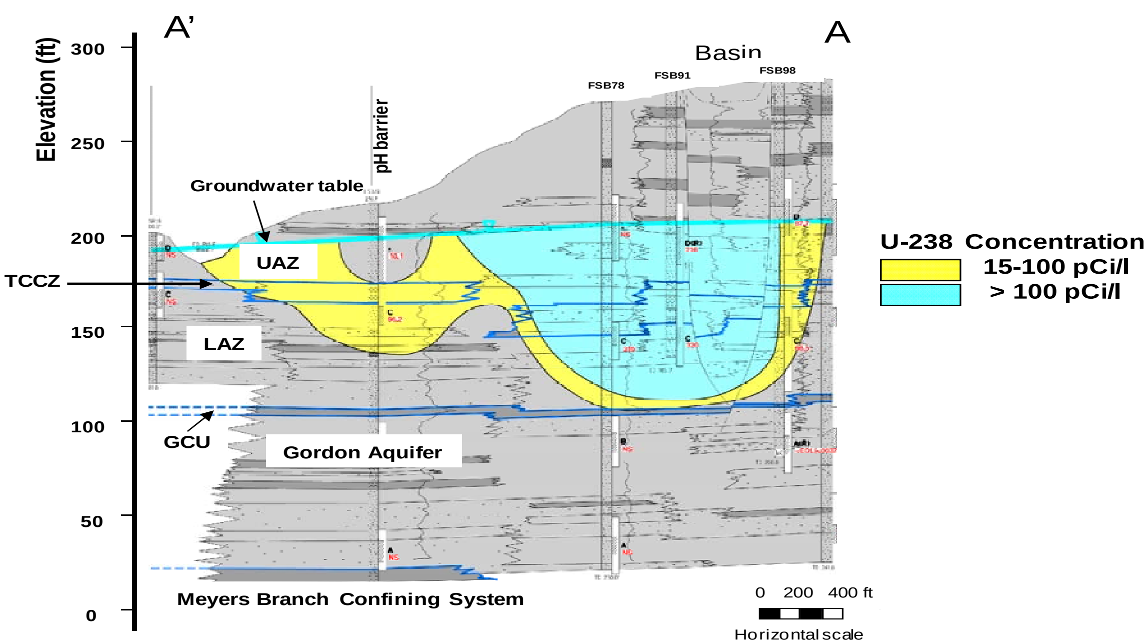

Figure 5. Schematic cross section along the axis of the plume shown in Figure 3, illustrating the topography, water table, hydrostratigraphic units, and U-238 distribution (modified from SRNS, 2010). The distance from A-A' is approximately 0.5 km.

Contamination in the F-Area primarily exists in the unconfined UAZ and in the LAZ, as is illustrated in Figure 5. The depth to the water table varies as a function of topography (Figure 5). Background waters (outside of plume region) have dilute Ca/Na bicarbonate (TDS < 200 mg/L) and naturally low pH (typically ~5.5 at shallow depths and increasing in deeper units).

2.2 Key Geochemical Processes in the F-Area

A conceptual model of the processes that exert the most control on contaminant migration within the F-Area has been summarized by Denham and Vangelas (2010) based on site characterization and monitoring data collected to date. This conceptual model includes four key zones (Figure 6). Zone A represents the interface between the vadose and saturated zones directly below the seepage basins. Zone B represents the entire vertical extent of the plume from the footprint of the source to the areas of base injection. Zone C represents the areas directly influenced by base injection. Finally, Zone D represents the area near the seepline and stream. Although some of the ASCEM components cover the entire region spanned by Zones A-D, many of the processes associated with the base injection and downgradient seepline/stream are not considered in this first demonstration.

The contaminant considered for this particular demonstration is uranium, occurring entirely as U(VI). In the Phase I Demonstration, the hydrological and geochemical processes that occur in Zones A and B were considered. Oxidizing conditions prevail in these acidic and low-organic content regions, and for this reason redox and microbial processes are not currently modeled in the HPC demonstration. The pH controls adsorption of uranium, dissolution/precipitation of the minerals of interest (kaolinite, goethite), and many aqueous reactions. Therefore, modeling the pH evolution was also part of the HPC demonstration.

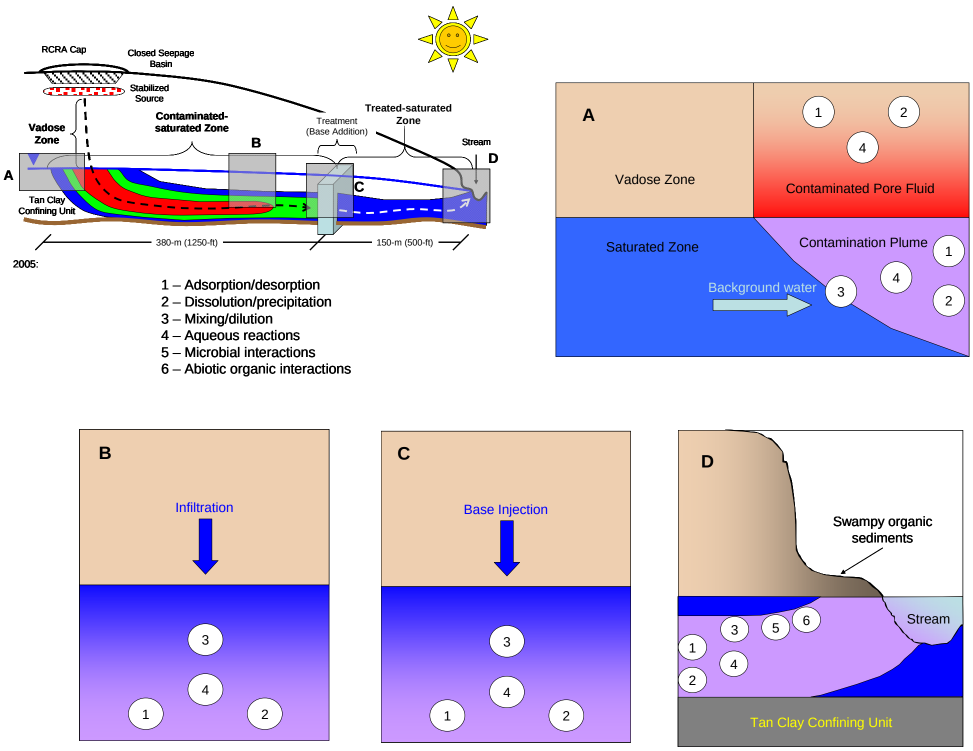

Figure 6. Conceptual model of attenuation processes in the entire F-Area Seepage Basins plume, including: (1) adsorption/desorption; (2) dissolution/precipitation; (3) mixing/dilution; (4) Aqueous reactions; (5) microbial interactions; and (6) abiotic organic interactions. The HPC simulations performed in this demonstration consider a subset of these processes.

The following processes affecting the attenuation of uranium at the F-Area were considered in the HPC demonstration:

- Adsorption/desorption: Adsorption and desorption are considered to be dominant natural attenuation mechanisms. As a leading pH gradient advances, adsorption of free protons onto mineral surfaces is an important process for neutralizing pH and slowing its advance. As a trailing gradient of higher pH values advances, desorption of free protons from mineral surfaces is important in slowing its advance.
- Dilution/mixing. Dilution is also considered to be an important attenuation mechanism. Dilution and mixing occur at the interfaces of the plume and uncontaminated water. Implicit in this model is that this process includes the neutralization of free protons by hydrolysis in the uncontaminated water.
- Mineral dissolution/precipitation: Mineral dissolution/precipitation processes occur throughout the plume. These are particularly important in slowing the advance of a leading pH gradient. The dissolution of kaolinite and goethite consumes free protons. In contrast, precipitation of iron and aluminum hydroxides releases free protons. This may be a factor in slowing the advance of a trailing pH gradient.

Likewise, it may occur where surface infiltration mixes with the plume and at the front of leading pH gradients.

- Aqueous reactions. Aqueous reactions that occur between dissolved species can both add and remove free protons from groundwater and occur throughout the plume. They tend to buffer changes in pH because the reactions themselves are driven by the concentration of free protons. The dominant inorganic species involved in these reactions are those of dissolved aluminum and ferric iron.

3. KEY DATASETS USED FOR THE ASCEM PHASE I DEMONSTRATION

Although the Phase I demonstration was performed to test four individual ASCEM components, common datasets were used to enhance illustration of the ASCEM capabilities and to poise the ASCEM components for integration and data sharing in Phase II and subsequent demonstrations. The key datasets used in the Phase I demonstration are briefly described below.

- Concentration database, which includes 44 measurements of ion concentration and other wellbore parameters collected during 1990–2009. These data were retrieved from 145 monitoring wells and from three different aquifers: the UTZ; the LTZ; and the Gordon aquifer. A list of the measured analytes and parameters is given in the Appendix.
- Depositional database (Smits et al., 1997), which includes information about lithology, well coordinates, well depths, well screened zones, depths of stratigraphic units, particle size distribution, and depositional environments.
- Hydrostratigraphic database, which includes the coordinates of the base of the units determined from wellbore data (Smits et al., 1997).
- GIS data, which includes digital elevation models (DEM) from digitized Barnwell 1:24,000 scale topographic maps (http://www.ncgc.nrcs.usda.gov/ http://ngmdb.usgs.gov/), satellite imagery (7.5 minute New Ellenton South West Digital Orthophoto Quad https://www.dnr.sc.gov/pls/gisdata) and additional GIS F-Area cultural feature data (SRNS-EM-2010-00055).
- Hydrologic data extracted from a large-scale flow model (Smits et al. 1997; Flach and Harris 1999) and subsequently ported to the PORFLOW code (Flach 2004) and extracted for use in the HPC component of this demonstration. Unsaturated soil properties (water retention and relative permeability curves) were taken from Phifer et al. (2006).

Although F-Area Site data (described above) and the geochemical processes (described in Section 3.2) are included in this demonstration, the results of this Phase I demonstration are intended to illustrate the advancement of ASCEM developed capabilities rather than to realistically represent system properties and behaviors. At this stage, no attempt was made to validate simulations or other ASCEM output based on F-Area datasets or process models. As such, the results shown in this document should not be used to guide remedial decisions, but instead to assess if the deliverables of the Phase I ASCEM demonstration, which focused on capability development, were met.

5. DATA MANAGEMENT (PLATFORM THRUST)

5.1 Objective and Specific Goals

The overarching objective for the ASCEM Database and Data Management component (hereafter called the ‘Data Management’ component) of the Phase I Demonstration is to develop the capabilities to import, organize, search, and manage various types of data commonly used for subsurface flow and transport investigations and numerical modeling. Although many data are available, the F-Area datasets are distributed among multiple spreadsheets, scientific reports and publications, and data files are stored in different formats and on different computers. The dispersed and heterogeneous nature of the datasets, which is common to many other contaminated DOE sites, hinders effective use of the data for advancing the EM clean-up effort. The ultimate goal of this ASCEM component is development of an integrated knowledge management environment that enables users to easily find, access, and add to the combined knowledge and data stored in the system.

Specific goals of the Data Management Phase I component are to illustrate the capabilities to manage two different types of data: (a) measured or simulated data, referred to as “transparent data,” and (b) documents, graphs, pictures, and similar data objects, referred to as “opaque data.” Transparent data can be searched and extracted while opaque data can only be accessed as a whole, although text documents can be indexed with keywords. In the Phase I Demonstration, the transparent data included concentration and depositional data while the opaque data included F-Area GIS data and various F-Area pdf reports.

5.2 Approach

A set of tools was developed/modified and implemented for the management of transparent and opaque data. A relational database called PostgreSQL was implemented in ASCEM to handle the management of transparent data: this is an open-source relational system that has flexible search capabilities. (Note that URLs for PostgreSQL and all other open-source software used in ASCEM are provided in the list given on Page 8). Web-based tools, such as Google maps and a JavaScript plotting package called FLOT were adapted for ASCEM to enable display of wells on maps, to query the data, and display results in terms of graphs and tables. A web-based knowledge management framework called Velo was customized for ASCEM opaque data management and is a domain independent framework developed from a number of open-source technologies, including Mediawiki, the Semantic Mediawiki extension, and the Subversion version control system (Gorton et al., 2010). Velo provides a rich file sharing, record management, and collaboration environment with the capability to incorporate new tools for viewing and processing scientific data. Velo was modified in

ASCEM to describe opaque data and to store and index metadata associated with wellbore data (such as descriptions of measurement variables or acquisition procedures).

Because inconsistencies in wellbore nomenclature, wellbore coordinate systems, and measurement units are common problems associated with DOE subsurface databases, an additional component of the Phase I Demonstration activities included the development of techniques to automatically render these attributes consistent (such as using a standard coordinate projection of UTM Zone 17, NAD83).

5.3. Accomplishments

5.3.1 Transparent data

The first step a user might take in developing a conceptual and numerical model of the subsurface would be to import and plot the location of wellbores and display associated measured parameters. Figure 7 illustrates the use of the ASCEM Mapping Tool (MAP) to display the F-Area well locations on a map, which could be either a satellite image or (as shown) a topographic map with roads. In Figure 7, a well or a cluster of wells can be selected for further investigation by clicking on a single well or by drawing a rectangle around a cluster of wells. Once selected, measurements from the concentration database (analytes and other attributes) are displayed.

Plotting the temporal variations in measurements collected from a single wellbore and querying the measurement database are also common procedures used with model development. Figure 8 shows how the ASCEM Plotting Tool (PLOT) can be used to illustrate the time series of concentration values for one or several of the analytes. This display also provides the MDL (method detection limit) and PQL (practical quantitation limit) for each analyte as is shown in Figure 8. An ASCEM-developed filter interface tool (FILTER) displays and queries the database measurements. Figure 9 shows a map resulting from using the interface to identify wells whose measurements fall within the minimum and maximum concentrations of a chosen analyte, based on a user-defined range.

The MAP tool described above can also be used to display data from the depositional database, and the PLOT and FILTER tools can also be used to assess characteristics associated with a single or a cluster of wells. Figure 10 illustrates use of the ASCEM MAP tool for displaying the percentage of mud, sand, and gravel as a function of depth. A user can also choose to only display wells and their attributes associated with a particular depositional environment. Clicking on one or more depositional environments icons will allow the user to display a map view of the wells that sample the selected depositional environments.

All information displayed from the concentration or depositional databases can be saved in a CSV (comma-separated-values)–formatted file, which includes a header (or a companion file) that records the parameters used to select the data in the file.

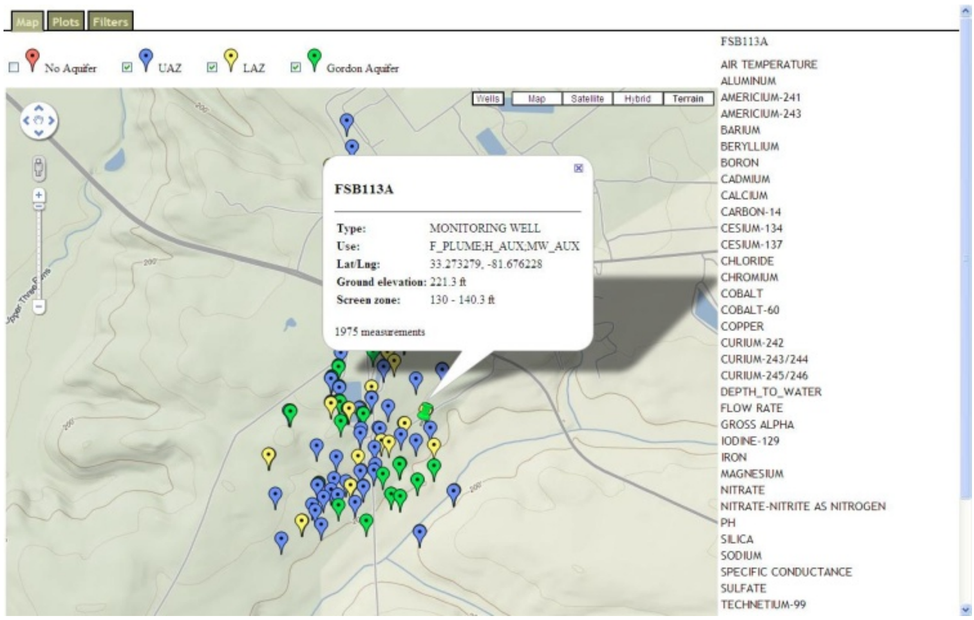

Figure 7. A map-based interface was developed for browsing concentration data using the developed ASCEM MAP tool. The analytes measured in the selected well are shown on the right. The wells are colored according to the aquifer where the wellbores are screened.

5.3.2 Opaque data

The Velo browsing interface was modified for ASCEM to permit description of fields and provenance associated with the depositional data. The user can search for terms in the available documents, and the terms are hyperlinked to the data to enable easy navigation between opaque and transparent data. Velo will be extended in future efforts to ingest and provide searchable metadata on any desired format of data, including Microsoft Word and Excel documents, formatted well logs, and seismic data.

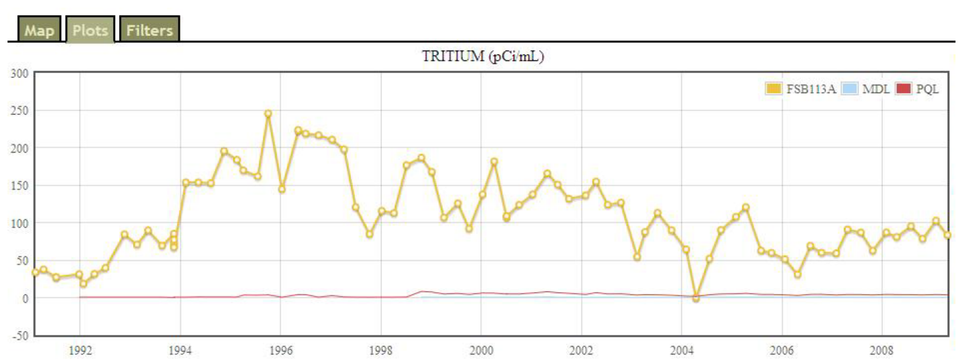

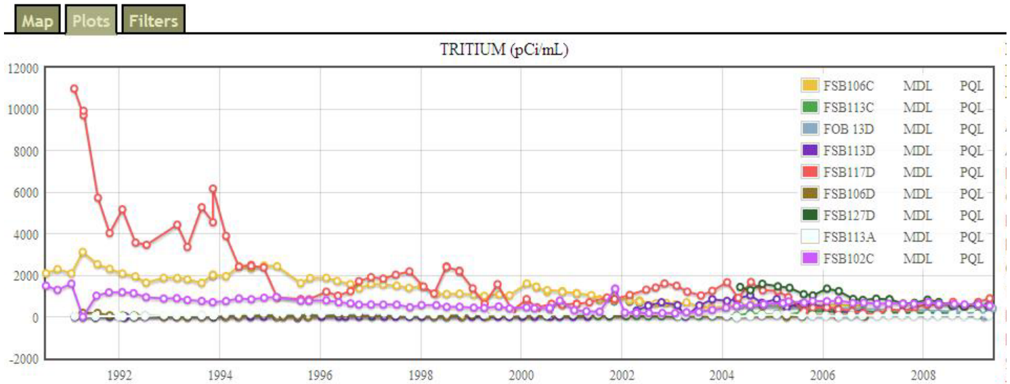

Figure 8. The ASCEM PLOT tool allows visualization of time series measurements collected from a single or cluster of wells. Upper figure: tritium concentrations over time from one well compared with the MDL and PQL concentrations. Lower figure: tritium concentrations over time in eight wells.

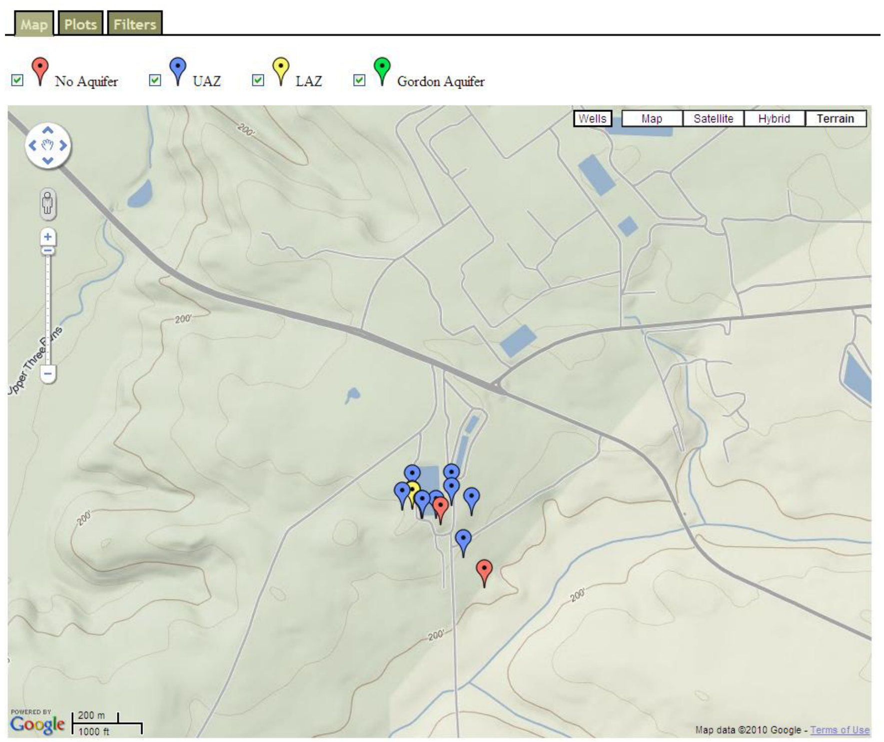

Figure 9. Illustration of the capabilities of the ASCEM FILTER tool, which allows a user to select for display a specific analyte and range. This example shows how the tool was used to display only wells having measured tritium concentrations above 20,000 pCi/mL.

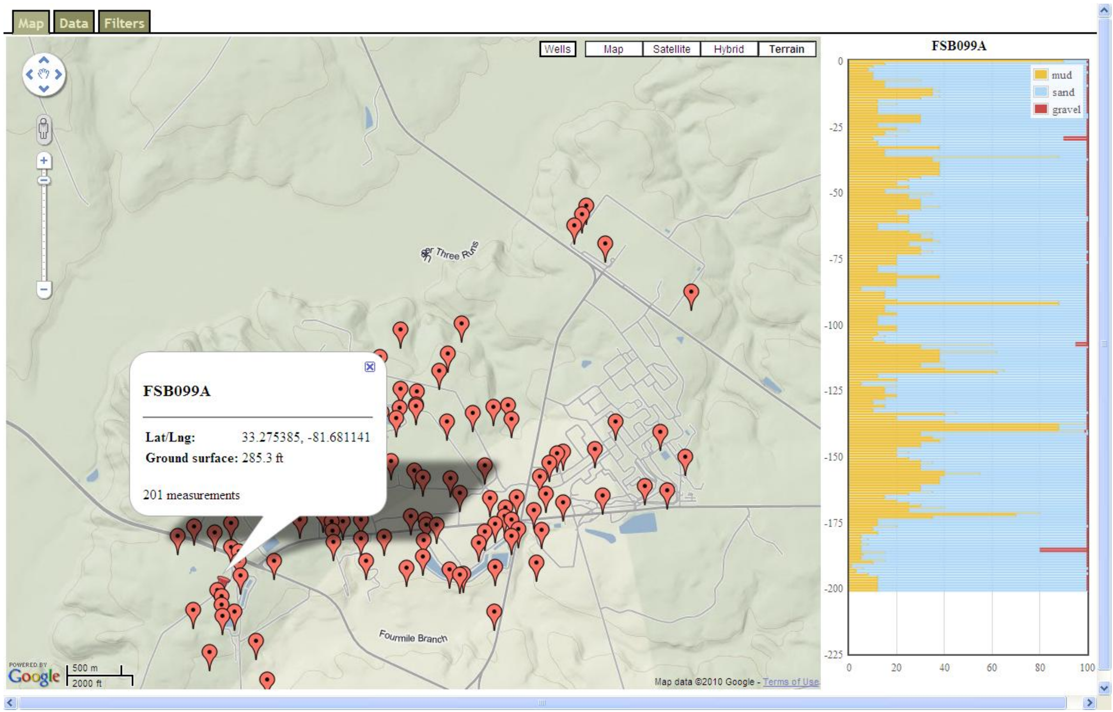

Figure 10. Illustration of the MAP tool, which can be used for browsing lithological data and displaying specific associated parameters, such as the percent of mud, sand, and gravel as a function of depth (shown on right).

5.4. Discussion

The Phase I deliverable of the Data Management component included development of data import, organization, and query tools, with an initial focus on contaminant concentration and hydrostratigraphic variables. The Phase I ASCEM Data Management efforts have met these objectives through development of a coordinated framework that includes: a) a relational database PostgreSQL, with flexible search capabilities to manage transparent data; b) Web-based tools, including Google maps and a JavaScript plotting package FLOT, to enable mapping of well layout, to query the data across multiple databases of various formats, and display results as graphs and tables; and c) a web-based management framework Velo, which has been customized to search across opaque data, to store and index metadata associated with site characterization and monitoring data, and to provide hyperlinks to connect directly the opaque data, interactive maps of well layout, and the transparent data. A particularly useful characteristic of the tools is that they can be used iteratively to display and query different subsets or characteristics of the databases. The tools and retrieved data are expected to be used for a variety of purposes, including: probing suites of wells that have particular characteristics; displaying various attributes associated with a single well; comparing measured and simulated data; or identifying outliers and erroneous data—all of which are common procedures of site characterization, modeling, and risk assessment activities.

During the Phase II Demonstration, the Data Management ASCEM component will concentrate on further development of an integrated approach for managing both transparent and opaque site characterization and modeling data, including the development of advanced scientific database and data management software. The Phase II demonstration will also include implementation of the data management software needed to support the ASCEM visualization, HPC, and UQ components. Data management capabilities will also be extended to develop a uniform schema description and User Interface (UI) properties of the dataset to be ingested, generate UIs automatically from such descriptions, and develop tools to integrate data from multiple overlapping datasets. In addition, the utility of data mining and analysis, such as the use of data cubes, will be explored to facilitate the following: the site characterization data analysis; development of a conceptual model; preparation of input data sets for HPC and UQ modeling; and analysis of modeling outputs.

The overall approach of using a coordinated framework for managing transparent and opaque site characterization and modeling data is essential for the development of an integrated knowledge management environment. Such a system is currently unavailable for easy use with subsurface data, and development within ASCEM will allow end-users and scientists to easily access, manage, and query heterogeneous datasets.

6. VISUALIZATION (PLATFORM THRUST)

6.1 Objective and Specific Goals

The overall objective of the Visualization component of the ASCEM project is to develop and demonstrate the ability to perform visual exploration and analysis of a diverse range of conceptual and numerical model data common to environmental management problems. The goals of the Phase I Demonstration are to show progress toward that larger objective by focusing on exploration of the SRS F-Area site characterization data, including the layout of observation wells, surface topography, depositional layers, water table, main hydrostratigraphic units, and spatial and temporal variations of radionuclide concentrations in groundwater.

The Visualization component of the Phase I Demonstration focused on a combination of technology development, extension, and application to meet the following specific objectives:

- Three-dimensional navigation through hydrostratigraphic units
- Temporal navigation of contaminant plume migration within the physical framework
- Visual exploration of depositional and other subsurface data
- Visual display and exploration of how variations of model input parameters influence simulation output.

In this section, the first three objectives are addressed using data from the F-Area, including the concentration, depositional, and GIS databases. The fourth objective is described in association with the waste tank supplementary problem in Section 9.

6.2 Approach

The ASCEM Visualization approach leverages the capabilities of the VisIt visual data analysis and exploration application. This approach allows ASCEM to take advantage of an open-source, production-quality, petascale-capable visual-data-analysis application that supports a diverse set of visualization and data analysis operations, and that can run effectively on diverse platforms ranging from laptops to the world's largest supercomputers. VisIt was initially developed as part of the NNSA ASC program, now benefits from technology contributions from visualization researchers and developers around the world. VisIt has proven effective at displaying a diverse set of data, including scalar and vector fields defined on two- and three-dimensional structured and unstructured meshes, over 300 different operators for manipulating and analyzing data, and robust infrastructure to support execution on large problems on petascale-class platforms. VisIt is the primary deployment vehicle for research and development at the DOE Visualization and Analytics Center for Enabling Technology (VACET), which is part of the Scientific Discovery through Advanced Computing (SciDAC) program within ASCR.

6.3 Accomplishments

The primary focus of Phase I consisted of: a) implementing extensions to VisIt in support ASCEM requirements; and b) testing the modified visualization capabilities using data from the SRS F-Area. The Phase 1 Demonstration required conversion of environmental datasets from F-Area to a form that could be readily loaded into VisIt. The data preparation was performed manually for this demonstration, but will eventually be formalized as part of the ASCEM data management process. The manual data preparation processes involved averaging (to overcome challenges resulting from uneven temporal sampling of wellbore concentration data) and manual manipulation of GIS data (including roads, buildings, surface topography, and aerial imagery). A custom data loader was constructed to import continuous well log information such as depositional environment data to VisIt.

Figure 11 illustrates use of the modified VisIt tool to image the physical framework of the SRS F-Area, Figure 12 shows a different view of the site including the locations of monitoring wells, GIS data (roads and buildings), surface topography, depositional environment, uranium-238 (U-238) concentrations in the (LAZ) depicted on a white-to-red scale (low to high concentration) at a single point in time, and position of the GCU as a green-colored surface. The figure shows two different perspectives of the plume distribution in 1994; on screen, the viewpoint can be changed by clicking and dragging the mouse.

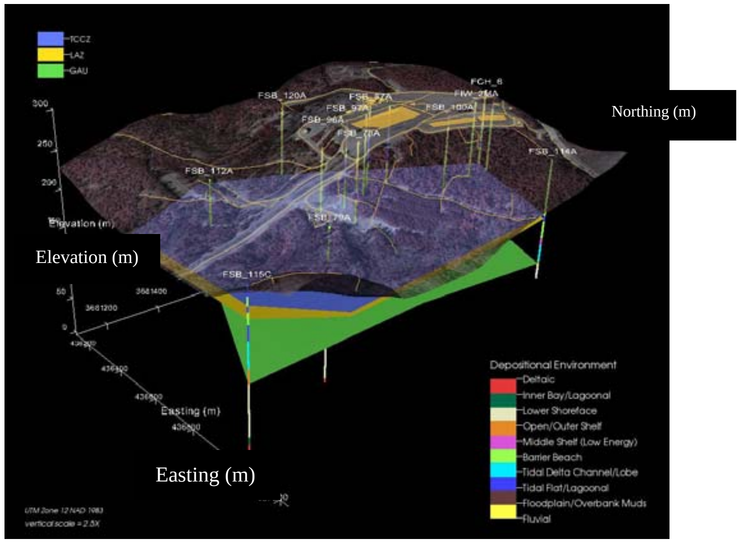

Figure 11. The use of ASCEM capabilities to visualize the physical framework of the F-Area, including the top surfaces of the key stratigraphic layers and depositional environments. The illustration also shows the layout of the monitoring wells, a 3-D image of the surface topography along with roads and building footprints and the distribution of the depositional environments along the well depths.

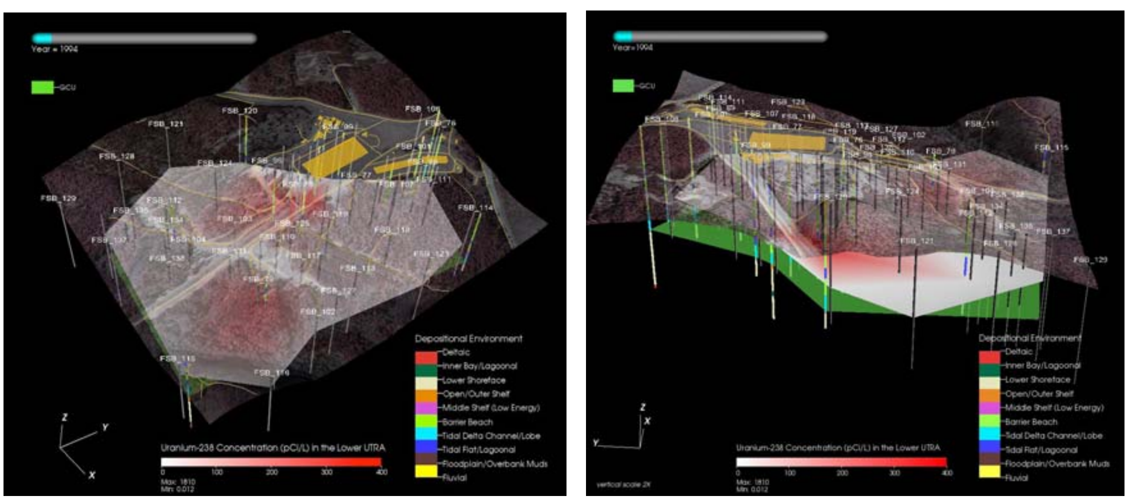

Figure 12. Examples of ASCEM Visualization output showing the U-238 concentration in the LAZ for the year 1994 from different perspectives (the *y*-coordinate is oriented to the North, and the axes scales are the same as in Figure 11). The illustration also shows the layout of the monitoring wells, a 3-D image of the surface topography along with roads and building footprints, the distribution of the depositional environments along the well depths, and the depth of the GCU.

As part of the development process for The Phase I demonstration, the visualization team experimented with several different algorithms and techniques aimed at producing the best possible results and to streamline future work. For example, visualizations of contaminant data were generated using two different algorithms. The first was to interpolate the discretely sampled concentration data onto a structured mesh using VisIt's built-in resample operator, which is based upon an inverse-distance method. Results of this approach are shown in Figure 13. Because this approach does not take into account that contaminants may not move across aquifers, an alternative approach was pursued. The second approach involved the construction of a DeLaunay triangulation using VisIt's built-in operator to produce two- and three-manifold surfaces/volumes using the screening zone locations within observation wells. An example of this approach is shown in Figure 14. In the longer term, new capabilities will be added to VisIt that support more robust, geostatistically meaningful interpolation.

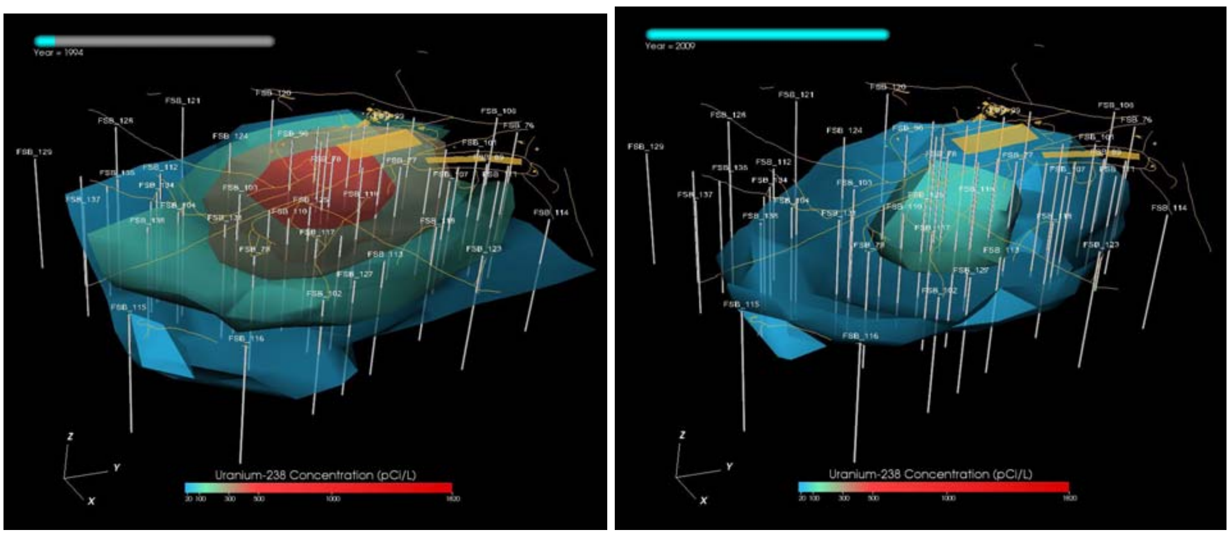

Figure 13. The use of an inverse-distance approach to visualize the evolution of four isosurfaces of U-238 concentration over time steps (left-1994, and right-2009). Also displayed are monitoring wells, buildings, and roads. (The *y*-coordinate is oriented to the North).

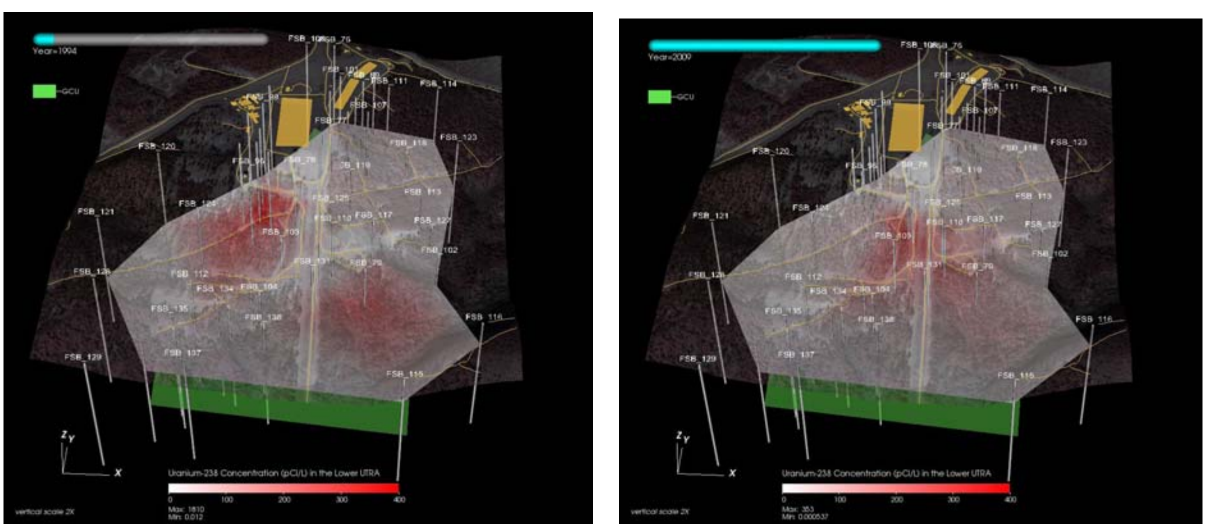

Figure 14. Use of a DeLaunay triangulation approach to visualize the evolution of U-238 concentration over time (left-1994, and right-2009). Also displayed are the surface topography,

buildings/roads, depositional environment, monitoring wells, and the depth of GCU unit. (The *y*-coordinate is oriented to the North).

VisIt's internal architecture allows users to quickly assemble visualizations that use common environmental datasets and different mapping operators, including DeLaunay triangulated surfaces/volumes and interpolated structured meshes. Additionally, the ability to visualize observational wells, surface topography, GIS data, etc., concurrent with both methods for displaying contaminant data was demonstrated. The VisIt interface supports three dimensional navigation through the data, as well as temporal navigation of underlying, time-varying data. A gnomon (a 3D-axis icon) in the lower-left corner of the figures indicates the position/orientation of the scene, and the current time value appears as a slider bar in the upper left corner of the images. In the case of continuous data (e.g., concentration measurements), the levels appear on a color legend shown at the bottom of each image. In the case of discrete data (e.g., the hydrostratigraphic data) a color legend with meaningful labels appears in the lower left.

6.4 Discussion

The Phase I Visualization demonstration included development of approaches to visualize the hydrostratigraphy, water table, topography, wells, and migration of contaminant plumes through aquifers over time. These objectives were met by extending the DOE-developed VisIt software package and testing the modifications using data from the F-Area. For the Phase I Demonstration, visualization was conducted using a desktop machine; future efforts will illustrate the use of VisIt capabilities on platforms ranging from the desktop to petascale-class machines.

ASCEM tools were developed that allow users to visualize the following: key surface features of a contaminated site; wellbores and associated depositional environments; stratigraphy and topography; and evolution of contaminant plumes (depicted as a surface associated with an individual layer) or as a volume. Particularly useful tools or characteristics of the approach include the ability to visualize many different types of data (point, surfaces, volumes) in an uncluttered fashion, the joint visualization of physical features and contaminant concentrations, and the use of slider bars to navigate through temporal datasets, such as concentration values. Additional ASCEM visualization advances, which illustrate developed capabilities beyond the proposed Phase I Demonstration deliverables, are described in Section 9.2

To streamline future efforts, several VisIt Python scripts were developed to implement new visual exploration capabilities for ASCEM. Future development will focus on two broad types of functional capabilities with respect to these scripts. First, a set of pre-defined processing sequences that produce specific types of visual output have been defined. An example of this would be observation wells for displaying contaminant concentrations. Second, different types of processing sequences will be needed, for different site data (e.g.

at F Area) and other sites. These will be based on the concept of codifying processing sequences and associated parameter values.

Future advances for the Visualization component of the Platform Toolsets will continue to leverage VisIt as a general purpose visual data analysis and exploration infrastructure that can be tailored to specific problems and application domains. It can be run on platforms that range from laptops to petascale-class machines. No existing visual data analysis capabilities have this flexibility.

For the next phase of development, the Visualization component will focus on closer integration with Data Management. Specifically, a data loader will be designed and implemented that can communicate directly with other ASCEM data management infrastructure. The visualization capabilities will be closely linked with data exploration infrastructure to enable analysis. The visualization component will also be more closely integrated with the Multi-Processor HPC simulation capability to facilitate analysis of model output.

7. UNCERTAINTY QUANTIFICATION (UQ; PLATFORM THRUST)

7.1. Objective and Specific Goals

One of the ASCEM goals is to develop tools that can be used to quantify the uncertainty in the models and simulation outputs resulting from EM performance and risk assessments. The objective of the Uncertainty Quantification (UQ) component is to develop capabilities within ASCEM for understanding and assessing key uncertainties in analyses. The range of capabilities being considered includes standard tools such as parametric analysis and forward propagation of uncertainties as well as advanced methods such as Markov chain Monte Carlo for highly parameterized models accounting for structural uncertainty. While the Phase I Demonstration included scoping work on advanced methods, most of the effort was spent on developing a software framework for UQ that can incorporate a wide variety of tools, both existing and future developments.

Specific goals of the UQ demonstration included the development of tools, interfaces, and approaches to: a) select parameters to consider for the UQ analysis; b) select model outputs for the UQ analysis; and c) perform UQ analyses.

7.2 Approach

The development supporting the Phase I Demonstration for UQ included the ability to easily access three different methodologies and several analysis types within the ASCEM UQ GUI, including: a) local sensitivity analysis; b) global sensitivity analysis (including Sobol variance decomposition and Morris one-at-a-time (Morris, 1991); and c) probabilistic predictions (including Monte Carlo and Null-space Monte Carlo (Tonkin et al [2007]). To support the Phase I demonstration, several open-source UQ software packages, PEST

(Doherty, 2010) and PSUADE (Tong, 2007), were linked with the GUI. For all approaches, the UQ analysis process is based on the assumption that a physically-realistic conceptual model exists with specified parameter inputs and model outputs.

In addition, the UQ development team worked closely with the platform core, model setup, and parameter estimation (PE) tasks to develop a prototype user interface (UI) for UQ. The interface is being developed in Java and will be deployed as a desktop application.

7.3 Accomplishments

A snapshot of the developed ASCEM UQ GUI is shown in Figure 15. This GUI has three main worksheets, which are associated with Parameter Selection, Model Output, and Analysis Options.

Parameter selection allows the user to choose from input parameters specified in the conceptual model setup. The GUI includes the option to specify upper and lower limits for each parameter, and can also specify an appropriate transformation of the input parameters. The user must also select the model outputs that will be considered for the UQ analyses. The Model Output page of the GUI includes the option of selecting a subset of the available model outputs for study. The Analysis Option worksheet (shown in Figure 15) illustrates how the user can choose an analysis type (local, global, probabilistic) as well as an analysis method after having selected input parameters and model outputs for this analysis.

In the Phase I Demonstration, the ASCEM UQ capabilities were illustrated using a Morris one-at-a-time method, which produces a measure of sensitivity for each input parameter. The example problem is described in Figure 16. It is important to realize that the geochemical reactions used for the Phase I Demonstration (and for HPC simulations) are preliminary, and the model results and associated uncertainty analysis should also be considered as such. To carry out the analysis, the GUI was designed to access UQ routines and subsequently produce a prescribed ensemble of forward model runs that the HPC core will be able to carry out. The user can inspect this ensemble design and make plots to ensure the ensemble is covering the input parameter space as desired before the design is transferred to the HPC core. In the Phase I Demonstration, TOUGHREACT (Xu et al, 2004) was used to test the ASCEM UQ capabilities because the HPC simulator, Amanzi, was still under development.

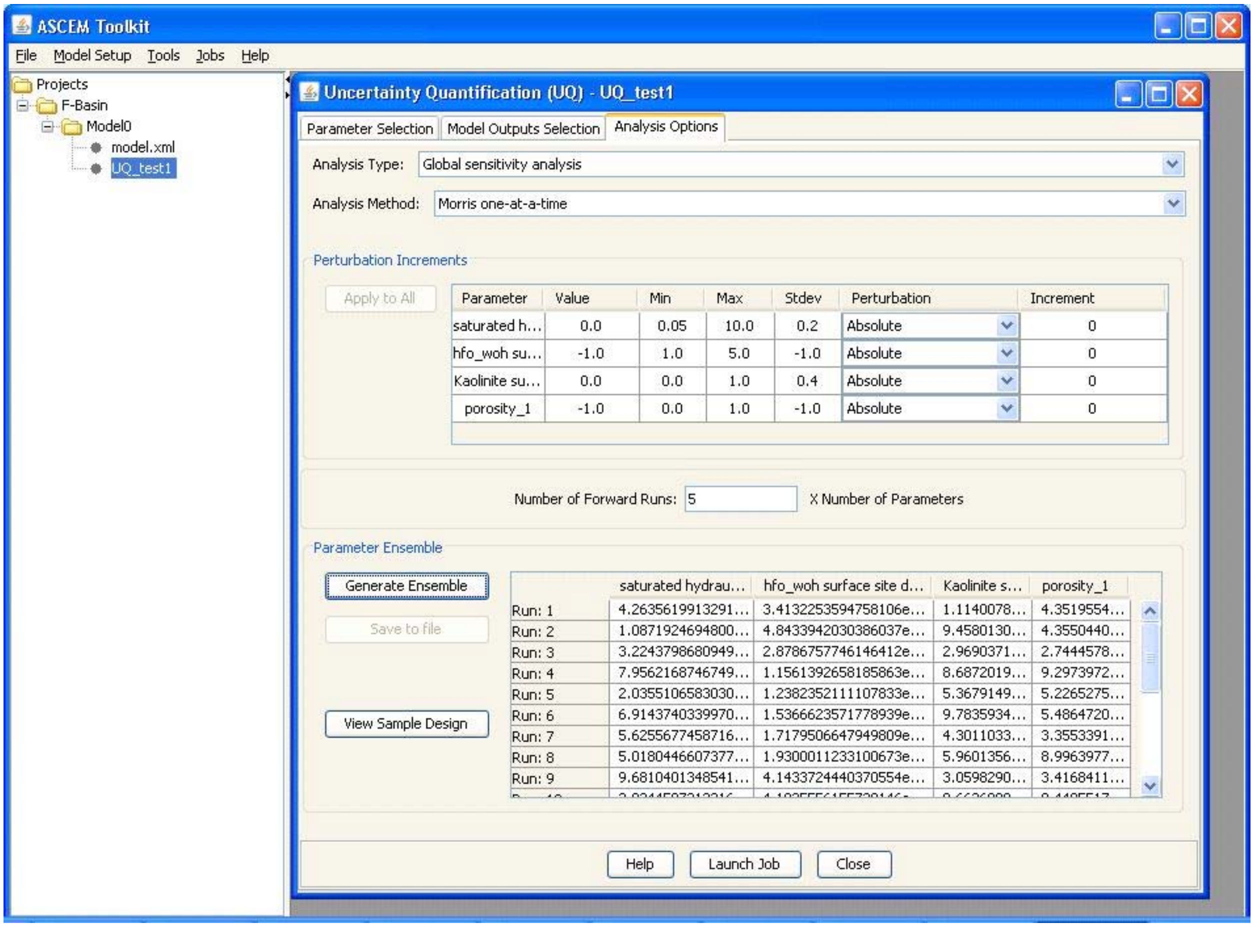

Figure 15. Screenshot of GUI for carrying out a global sensitivity analysis using the one-at-a-time method of Morris (1991). Here, the user has already carried out parameter selection and model output selection. After choosing "Global sensitivity analysis" and then "Morris one-at-a-time," the user is prompted to select the size of forward run ensemble to carry out for the analysis. The user chooses an ensemble size of 5 times the number of input parameters. The UQ tool then generates the ensemble of forward model runs, each of which is determined by its input parameter settings. The user can view the ensemble before sending this ensemble of forward model runs out to be computed (eventually, via the HPC core).

Box UQ1: F-Site Forward Model

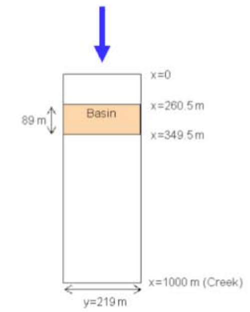

Stylized 1-D reactive transport model of the Savannah River F-Site.

The 1-D TOUGHREACT model (Xu et al, 2004) for the Savannah River test site is our test case. This model simulates the migration of basin infiltration through 1000m of aquifer, where it participates in a relatively large number of chemical reactions. While this model is relatively simple hydrologically, it is relatively complex geochemically. We decided to vary the following parameters: basin infiltration rate (constant in time) from 1955 – 1989; basin infiltration water chemistry, sorption site surface area, and mineral precipitation/dissolution surface area. The resulting 21 individual parameters are shown in Table 1 below. For the purpose of this demo, the outputs considered are shown in Table 2. For some of the outputs, physical observations are available – for others, such as well concentrations 100 years from now, have no observations available.

Table UQ1: Input parameters for forward model

	Parameter type	Parameter	Initial guess	Lower bound	Upper bound
1	Chemistry of basin infiltration water	H+	5.000E-03	5.000E-07	5.000E-02
2		Al+++	1.000E-08	1.000E-09	1.000E-07
3		Ca++	0.586E-04	0.586E-05	0.586E-03
4		Cl-	0.282E-03	0.282E-04	0.282E-03
5		Fe+++	1.75E-06	1.75E-07	1.75E-05
6		HCO3-	0.590E-04	0.590E-05	0.590E-03
7		K+	0.309E-04	0.309E-05	0.309E-03
8		Mg++	0.382E-04	0.382E-05	0.382E-03
9		Na+	0.248E-03	0.248E-04	0.248E-02
10		SiO2(aq)	1.14E-04	1.14E-05	1.14E-03
11		SO4--	0.384E-03	0.384E-04	0.384E-02
12		Sr++	0.149E-06	0.149E-07	0.149E-05
13		NO3-	0.01000	0.001000	0.08000
14		UO2++	0.31E-04	0.80E-05	0.80E-04
15	Basin infiltration rate	Basin Infil (kg/s)	.3816E-1	.3816E-3	.3816
16	Surface sorption site surface area	hfo_woh_sa	1.150E+06	1.150E+05	1.150E+07
17		k1_AlOH_sa	23.8E+04	23.8E+03	23.8E+05
18		k1_X-_sa	4.76E+04	4.76E+03	4.76E+05
19	Mineral surface area	SiOH_sa	0.14E+04	0.14E+03	0.14E+05
20		kaol_sa	26.E+04	26.E+03	26.E+05
21		gibbsite_sa	120.E+04	120.E+03	120.E+05

Table UQ2: Model outputs – concentrations of different solutes at different wells and times

Well	Year	Solute	Target	Reference
95DR	2000	pH	3.5	(1)
126	2000	pH	4.0	(1)
95DR	2008	pH	3.5	(1)
126	2008	pH	4.0	(1)
95DR	2000	SO_4	20	(1)
126	2000	SO_4	20	(1)
95DR	2008	SO_4	/	N/A
126	2008	SO_4	/	N/A
95DR	2000	Al	18	(1)
126	2000	Al	20	(1)
95DR	2008	Al	10	(1)
126	2008	Al	1.	(1)
95DR	2000	U	1.52	(2)
126	2000	U	0.8	(2)
95DR	2008	U	/	N/A
126	2008	U	/	N/A

Note: (1) Spycher, N., S. Hubbard and B. Faybishenko, ASCEM Phase I Demo Site Selection F-Area, Savannah River, ASCEM Program Meeting June 30, 2010.
(2) Metal and Rads Project Database, Savannah River Laboratory, Aiken, South Carolina.

Figure 16. Example problem to test the developed ASCEM UQ capabilities, showing the forward model set up and the associated model inputs and outputs. In the figure, model inputs in Table UQ1 are the labels shown on the x-axis of the plots in Figure 17; the model outputs in Table UQ2 are used for the sensitivity analyses summarized in Figure 17.

The outputs from an ensemble were extracted from each of the forward model output files and used to compute sensitivities. In the example shown in Figure 16, the output is taken to be the sum of concentrations recorded at the two wells, 95DR and 126, and the analysis required an ensemble of 220 forward model runs. For each output, sensitivities are estimated—one for each model input parameter. Shown in Figure 17 is the sensitivity of each model input to the concentration of Al, pH, UO_2 and SO_4. The vertical axes are the mean sensitivity coefficients of the computed concentrations related to the 21 sampled parameters. These coefficients give a measure of the change in concentration as the input parameter is varied across its allowed range. Currently, the graphics are created using an interface to a plotting language used by PSUADE. Eventually, the graphics will be produced by interfacing with the visualization utilities in ASCEM (Section 6).

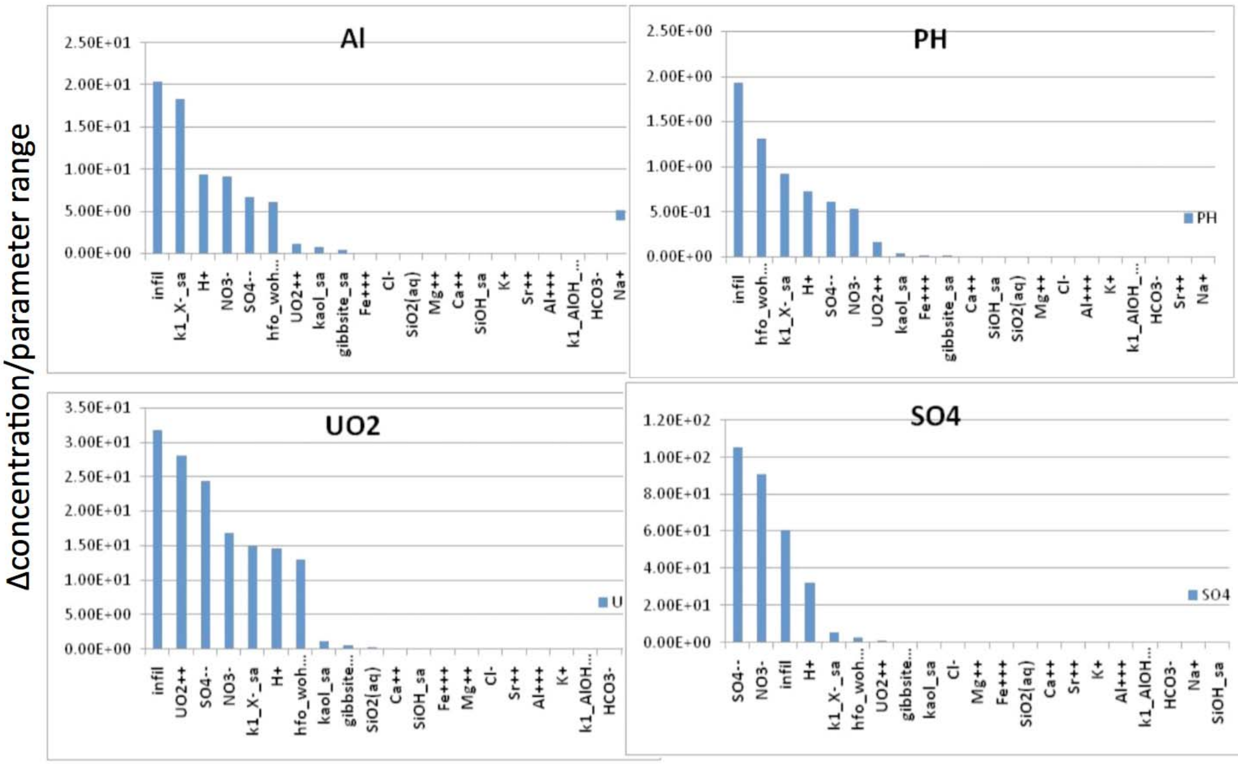

Figure 17. Example of ASCEM UQ output showing Morris one-at-a-time sensitivities for four different outputs. The computed sensitivities give the change in output over the range for each parameter.

Although 21 input parameters were considered in this study, each of the outputs is controlled by a much smaller subset of input parameters. This suggests that one can manage the different outputs by focusing on the most sensitive parameters. This analysis also indicates that many of the input parameters have no impact at all on the concentrations of the solutes. Hence, the uncertainty regarding these parameters is not constrained by measurements of solute concentration at these two wells.

While only results from the Morris one-at-a-time analysis are shown, the UQ Toolset developed for the Phase I Demonstration also allows sensitivity analyses using local derivatives and a Sobol decomposition. The results of these analyses led to a very similar conclusion, that the outputs are controlled by the same input parameters. The GUI allows the results of these different analyses to be collected and viewed as shown in Figure 18.

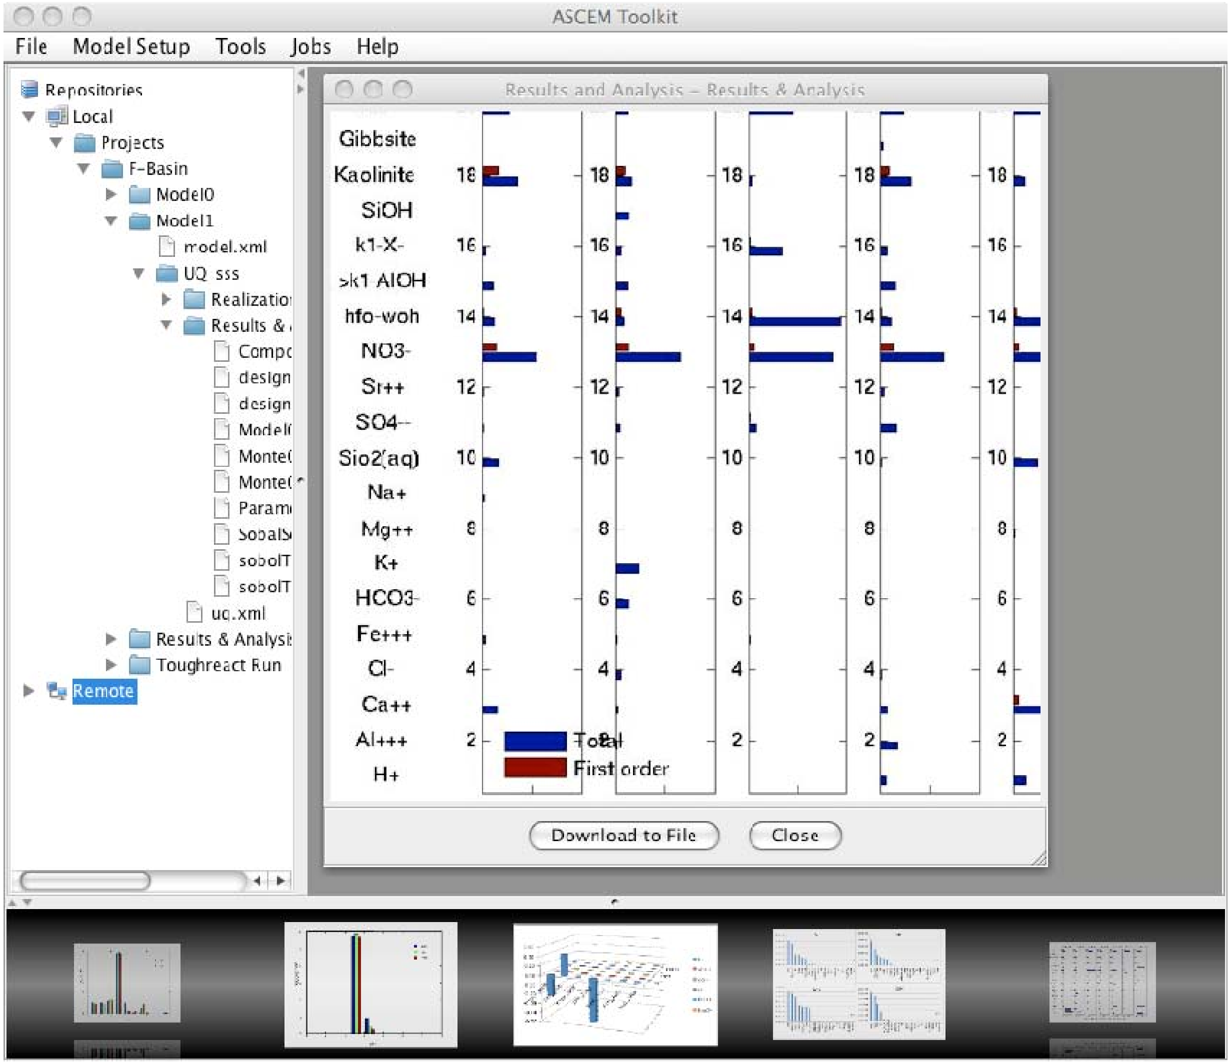

Figure 18. The ASCEM UI persistently tracks the inputs and outputs of the UQ process and provides tools not only for setting up and running the UQ analysis, but also for managing and viewing the data. The figure shows the "imageflow" tool providing a summary view of the generated images.

7.4. Discussion

The ASCEM Phase I Demonstration deliverables associated with this component included evaluation of UQ strategies and development of tools to evaluate sensitivity of model output to parameter suites. These objectives were met through the development of an ASCEM UQ GUI, which allows the user to select model parameter and outputs and to perform UQ analysis using a different analysis approaches. While forward model runs were implemented on a parallel computer, each of the UQ analyses could have been performed on a laptop. This GUI was successfully implemented and tested using data from the SRS F-Area.

Several advances to the ASCEM UQ tool are planned for the Phase II demonstration. One major activity for the Phase II demonstration will be to further define and develop the software interfaces and standard data formats that will effectively link the UQ algorithms to the rest of the ASCEM Platform and HPC. Other open source algorithms (e.g., DAKOTA, StatLib) as well as ASCEM-developed codes will be incorporated in future versions, and integration among UQ, HPC, and other ASCEM components will be performed. Over the longer term of the ASCEM project, the UQ and Parameter Estimation tasks will incorporate a number of capabilities, such as:

- Response surface methods for approximating a forward model response
- Searching the input parameter space for input settings that may lead to extreme or adverse outcomes
- Analyses and utilities that can handle highly parameterized models
- Assessment/evaluation of simplified, or reduced, forward models
- Conceptual model uncertainty.

Although many open-source UQ analysis approaches already exist, the various methods are not available in a single software analysis package. Bringing them together under the ASCEM UQ tool represents a major advance that is expected to improve EM performance and risk-based decisions.

8. HIGH PERFORMANCE COMPUTING (HPC THRUST)

8. 1 Objective and Specific Goals

The overarching objective of the ASCEM Multi-Process HPC Simulator, Amanzi (which means “water” in Zulu), is to provide a flexible and extensible computational engine to simulate the various model scenarios created through Platform toolsets. This component of the ASCEM Demonstration focuses on highlighting progress on early prototypes of selected Toolsets within the HPC Simulator, using a conceptual model of the SRS River F-Area site.

The specific goals of Phase I for HPC included: development of a parallel, unstructured mesh capability to illustrate the flexible and effective treatment of complex geometries; capabilities to perform three-dimensional parallel flow simulation; capabilities to simulate transport of a nonreactive contaminant; and a prototype of the Reaction Toolset that includes aqueous speciation, mineral precipitation and dissolution, and sorption (formulated as multi-component ion exchange and/or surface complexation). Using an operator-split approach, these reactions will be coupled to transport and used to simulate the concentration of uranium, and other important chemical species in the groundwater at the SRS F-Area.

In addition to the initial requirements for the Phase 1 Demonstration described above, the ASCEM team undertook the extra challenge of prototyping key elements of a structured mesh capability for the HPC Simulator on a simplified model of the SRS River F-Area. The specific goals were to model leakage from a seepage basin and track the migration of the contaminant through the vadose zone and the subsequent evolution of the plume. The simulations are based on a multiphase flow formulation to represent both the flow of water and of gas within the vadose zone. Two cases were considered: 1) simulation of a passive tracer (Na+); and 2) a more complex case using a geochemical model with 17 primary species. The simulations were performed using an adaptive mesh refinement capability (AMR) to track the location of the plume.

The approach of the Phase I Demonstration for the HPC component was intended to highlight the capabilities to handle realistic geometries and site characterization data. As with the UQ demonstration, the conceptual model for flow and geochemical reactions used in the HPC demonstration is preliminary, and thus simulation results should not be interpreted as being representative for this site. For ease of discussion, the approaches and accomplishments associated with the structured and unstructured mesh components are described separately below.

8.2 Approach

The Multi-Process HPC simulator takes as input a conceptual model, which describes a set of coupled processes such as flow and reactive transport. A conceptual model is expressed mathematically by a system of differential equations that represent the relevant conservation laws, constitutive laws, equations of state, and reactions. Various parameters required for the model are specified, along with initial and boundary conditions. To represent this system of equations on a computer, a mesh (grid) is provided with the model. A mesh may be thought of as a collection of discrete cells or grid blocks that fill the domain of interest. For a given mesh, a relationship between variables (e.g., pressure), parameters (e.g., permeability), and mesh geometry is developed. This process is referred to as discretization, and gives rise to a system of equations that represent the model. This discrete system is often nonlinear and must be solved to determine the quantities of interest, such as the concentration of particular contaminants.

The hierarchical and modular design of the Multi-Process HPC Simulator reflects the steps in translating a conceptual model to a numerical model producing output for analysis. At the highest level, the Multi-Process Coordinator (MPC) and the Process Kernels (PKs) represent the conceptual model. The PKs are high-level objects that represent tangible processes such as flow and transport. Mathematically a PK represents a specific set of differential equations. The Multi-Process Coordinator (MPC) manages the coupling of all the PKs that comprise the conceptual model, as well as the data associated with the conceptual model.

At the next level of design the HPC Toolsets includes Mesh Infrastructure, Discretization, Reactions, and Solvers. The Mesh Infrastructure Toolset provides interfaces and supporting routines to leverage existing mesh representation libraries. The Discretization Toolset provides the procedures that generate the discrete system of equations from a given continuum model on a given mesh. The Reaction Toolset implements geochemical reactions such as aqueous speciation and sorption. At the lowest level, the HPC Core Infrastructure provides low-level services such as data structures to operate on parallel computers, interfaces to other frameworks, input/output, and error handling.

The first stage included developing prototypes for the components outlined in the design. To achieve the flexibility and extensibility that users ultimately need, modularity was maintained as a central theme throughout the development effort. Modularity was realized by developing an Application Programming Interface (API) that defines how a component has access to data and services in the code. APIs or interfaces in Amanzi facilitate leveraging existing HPC frameworks, libraries, and tools. The main steps in development of Amanzi associated with the Phase I Demonstration were as follows:

- Developed support for parallel unstructured hexahedral meshes that leverages capabilities in the Mesh-Oriented datABase (MOAB) library. This consisted of an existing mesh library that is part of the Interoperable Technologies for Advanced Petascale Simulations (ITAPS) Center with the SciDAC program of ASCR. The mesh was partitioned by Zoltan.
- Collaborated with the mesh generation team under Platform to generate meshes with LaGriT (Los Alamos Grid Toolkit) in Exodus II format. This consisted of an existing unstructured grid format developed at Sandia National Laboratories that allows important features of the problem to be identified on the mesh, such as particular boundaries and the basin locations.
- Developed the Multiprocess Coordinator (MPC) to manage the system and control the coupling of processes as well as its evolution in time.
- Developed Process Kernels (PKs) for flow, transport, and chemical reactions.
- Developed the Discretization Toolset to support the flow and transport Process Kernels. The discretization of flow is particularly challenging on unstructured grids with tensor permeability. Advances in the Mimetic Finite Difference (MFD) discretization methods (Brezzi et al., 2007), which have been developed under the Applied Mathematics program of ASCR by members of the ASCEM team, were leveraged for this development.
- Developed a prototype of the Geochemistry Toolset to support selected processes relevant to the SRS F-Area. Development of the Geochemistry Toolset leveraged the experience and expertise of ASCEM Computational Geoscientists, who not only pioneered the geochemical algorithms, but have implemented them in codes such as CrunchFlow (Steefel, 2009; Steefel et al. 2003), and PFLOTRAN (Hammond and Lichtner, 2010).
- Solved the discrete system of equations using the Trilinos nonlinear and linear solvers.
- Developed the ability to drive the simulation with an input file based on the Extensible Markup Language (XML). For Amanzi, the capabilities in the Teuchos package of Trilinos were leveraged.

- Developed an output API and interface to the CFD General Notation System (CGNS) library because this format can be read by the VisIt application (see discussion of VisIt in Section 6).
- Developed a prototype of the error handling service.

A prototype of Amanzi comprised of these components was used to simulate flow and reactive transport using a conceptual model of the SRS F-Area.

To deliver a flexible and efficient capability that supports a graded and iterative approach to conceptual model development, it is imperative that Amanzi support both structured and unstructured meshes. To take advantage of advanced features and emerging architectures within each meshing paradigm, toolsets (such as discretizations and solvers) need to have implementations that are targeted to each particular meshing strategy. Thus, a block-structured adaptive mesh methodology was selected as the primary basis for development. In this approach, a coarse logically-rectangular mesh is prescribed to cover the computational domain. Refinement criteria are then used to identify parts of the domain where additional resolution is required. A collection of boxes is defined that cover the points identified for refinement. Each of these boxes is then used to define a grid patch at a finer resolution to represent the solution in that region. This procedure is applied at increasing finer levels of resolution until the desired resolution is obtained. In the block-structured refinement approach considered for Phase I, the finer resolution data are organized in large aggregate grids containing a number of grid points. Thus, the irregular work associated with adaptive refinement is focused on the relationship of larger grid patches that tile the region of interest, in contrast to the need to identify the relationship between individual cells in a cell-by-cell refinement approach.

The structured adaptive mesh methodology was implemented using a BoxLib software framework. BoxLib contains a collection of C++ data abstractions designed to support the implementation and parallelization of this type of Adaptive Mesh Refinement (AMR) algorithm. The discretization algorithms and approach to adaptive refinement used for Phase I are discussed by Pau et al. (2009).

8.3 Accomplishments

To guide development of the conceptual and numerical models (including the computational domain, boundary conditions, and model parameters), two- and three-dimensional scoping studies were conducted using an existing simulator PFLOTRAN (Hammond and Lichtner, 2010). Based on these scoping studies, the computational domain was selected to cover the region shown in Figure 18, and unstructured hexahedral meshes were generated that captured the topography and hyrdostratigraphy. Specifically, meshes with different resolutions were generated with LaGriT (the Los Alamos Grid Toolkit) and written in the Exodus II format. One of these meshes with horizontal resolution of approximately 16 m, is shown in Figure 19.

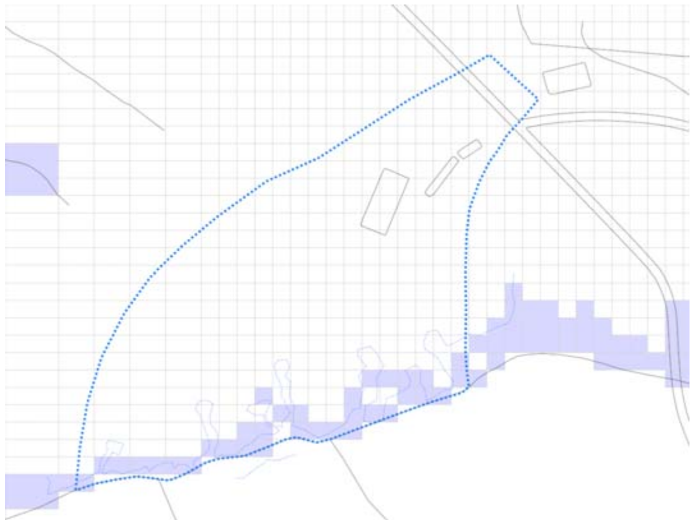

Figure 18. The sub-region of the F-Area encompassed within the HPC simulation is shown. The selection of the side boundaries is based on estimates of the no-flow lines obtained from a large-scale simulation of the General Separations Area (Flach, 2004), which includes the F-Area seepage basins, and the scoping studies with PFLOTRAN.

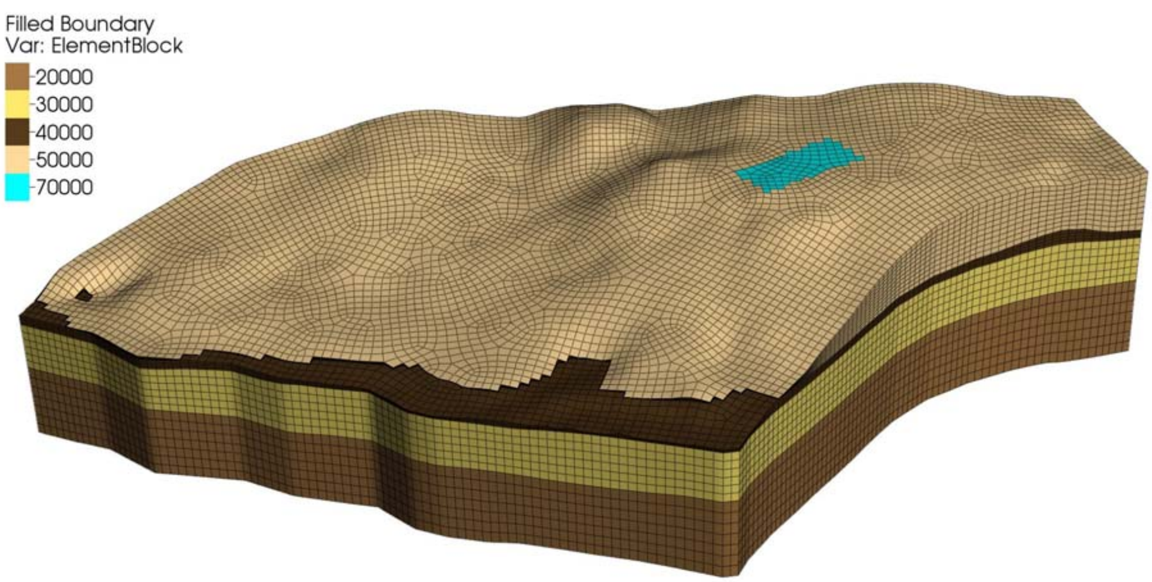

Figure 19. An unstructured hexahedral mesh of the sub-region described in Figure 18 is shown along with the four major hydro-stratigraphic units that are included in this model. The horizontal resolution of this mesh is approximately 16m. The top layer (ID: 5000) is the Upper Aquifer Zone (UAZ), followed by the Tan-Clay Confining Zone (ID:4000), the Lower Aquifer Zone (ID:3000), and finally the Gordon Confining Unit (ID:2000). The largest F-basin is shown in cyan (ID:7000) and is where the contaminant source was positioned in this model.

A Mesh Class and an API were developed to support the use of different mesh database libraries within Amanzi. The Mesh API provides important services for the MPC, process kernels, and discretizations. An important feature of the mesh libraries is their explicit representation and efficient access to the faces of cells. These entities play an important role in discretizations that achieve local mass conservation. A simple structured mesh database, dubbed "simple_mesh," was developed to enable unit tests of basic Toolset capabilities. For unstructured hexahedral meshes, the interface service routines to the MOAB API had to be developed. The MOAB API was used to read the mesh and to partition it across multiple processors, along with the auxiliary data (i.e., side sets). In a partitioned mesh, each processor has mesh entities that it owns, such as cells, as well as copies (ghosts) of mesh entities that neighboring processors own. An example of an F-Basin mesh with 64 partitions is shown in Figure 20. Each color represents local data that a processor owns. MOAB allows the user to specify the type and extent of ghosting; entire ghost elements with faces are included in Amanzi. In addition, the maps used by the Trilinos framework to provide the communication of mesh and field data between processors were developed.

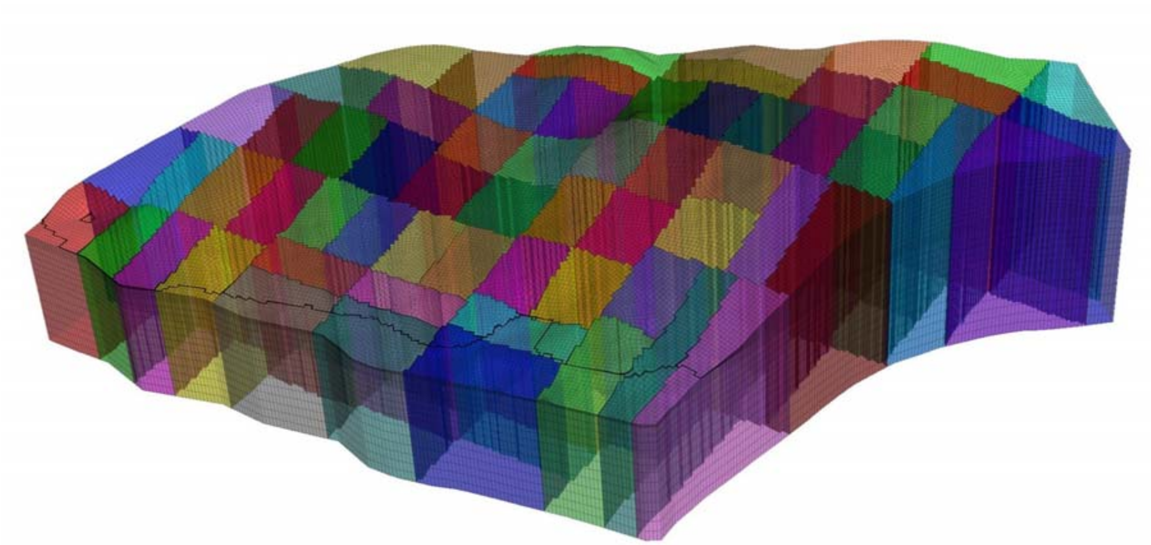

Figure 20. Meshes are read by MOAB and partitioned by Zoltan. In this image, each color represents a partition or subdomain of the mesh, and there are 64 partitions. The surface is translucent to show that horizontal partitioning dominates with a very shallow domain. Typically in a parallel run, each partition is assigned to a processor core.

The interface for the Discretization Toolset was developed, as well as the Mimetic Finite Difference discretization of the differential equations that model flow on unstructured hexahedral meshes (Morel, 2001). This prototype of the Discretization Toolset was used to develop a process kernel for single-phase Darcy flow on hexahedral unstructured meshes. Similarly, a standard Finite Volume discretization of the differential equations that model transport on unstructured meshes was developed. Process Kernels for both flow and transport were developed that used these discretizations. In both cases unit and verification tests were developed and run.

To simulate chemical reactions, a Reaction Toolset was developed that provides classes to depict chemical species and reactions. In addition, classes were developed that support process models (e.g., activity coefficient calculations) and a simple nonlinear solver. Included in this set of classes is the Beaker object within which a chemical reaction is solved, much like a batch reaction in a laboratory beaker. To support the modular design of Amanzi, a simple API for the Reaction Toolset was developed. Essentially, the necessary chemical constraints and parameters (e.g., time step size, tolerances, etc.) are all that is required for the Beaker object to solve a geochemistry step. To facilitate the use of the Reaction Toolset, I/O routines have been developed for reading and writing geochemical data, debugging, and verification/validation exercises.

To perform flow and reactive-transport simulations for the Phase I Demonstration, a MPC was implemented. Driven by an XML input file, this prototype MPC can selectively run individual processes or execute the steps to couple processes sequentially. This sequential coupling is often referred to as operator splitting, and it is commonly used for reactive-transport. In this simplified model a steady-state flow field is assumed, hence a complete time-step is composed of a transport step, which advects the total component concentrations, followed by a nonlinear solve for geochemical reactions locally on each mesh cell.

To test the MPC in this setting and facilitate comparisons of the new Reaction Toolset with the geochemistry modules in an existing one-dimensional simulation, a one-dimensional model problem with prescribed flow and component concentrations at the left boundary, and zero initial concentration across the domain was specified. The length of the system was 100 m, and the flow velocity of 19 m/y, equal to the discharge rate from the F-Area seepage basin, was used. A Geochemistry PK for a five-component system was developed consisting of the primary species UO_2^{+2}, Al^{+3}, $SiO_{2(aq)}$, HPO_4^{-2}, and H^+, with a number of secondary complexes. Kaolinite and quartz were used for the primary mineral assemblage, while the phosphate-bearing mineral $(UO_2)_3(PO_4)_{2.4H2O}$ formed as a minor secondary phase. The MPC managed the Process Kernels for transport and reaction to simulate the propagation of the reacting front across the domain. The computed concentrations of pH and total UO_2^{+2} are shown in Figure 21 at a sequence of times. Uranium was slightly retarded through several surface complexation reactions, although the extent of retardation was reduced by the uranium-bearing complexes $UO_2H_2PO_4^+$ and $UO_2HPO_{4(aq)}$.

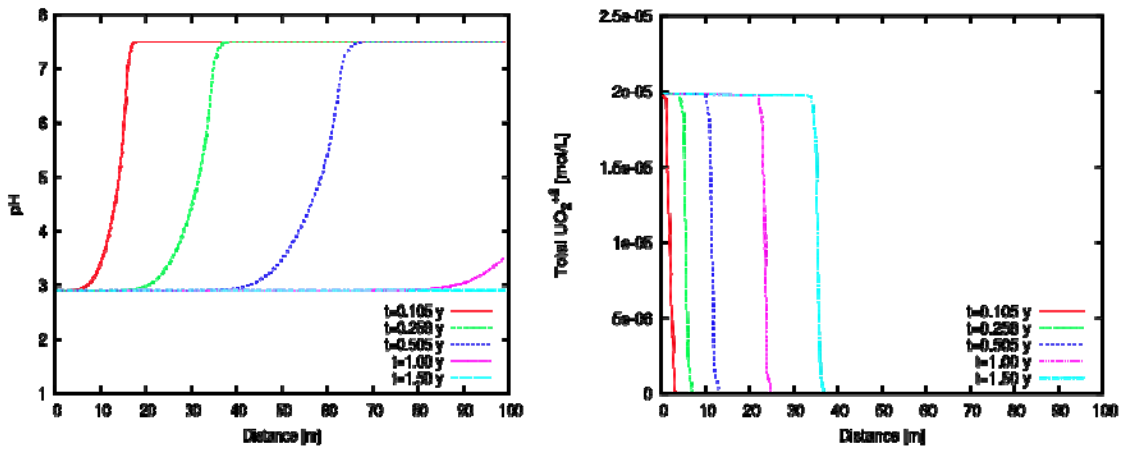

Figure 21. Illustration of Amanzi Reaction Toolset output, showing one-dimensional reactive-transport with a prescribed flow velocity of 19m/y and fixed concentration at the left boundary is shown for the five-component geochemistry system. These results are in excellent agreement with the existing simulators PFLOTRAN and Crunchflow.

To demonstrate the full capabilities of the Amanzi prototype, a numerical model of the F-Area seepage basins was constructed. The domain identified in Figure 18 was used. For the simulation, a refinement of the hexahedral mesh shown in Figure 19 with a horizontal resolution of approximately 8m was used (the coarser mesh is shown to facilitate visualization). The vertical resolution of the mesh was approximately 2.7m, and it contained 562,887 hexahedral cells.

The values of model parameters, initial and boundary conditions were based on the data provided by the F-Area Working Group (Flach, 2010). In the flow model, the top surface used a boundary condition with prescribed infiltration rates. Specifically, the infiltration rate outside seepage pond was 38cm/y, and inside the seepage pond was 19m/y. No flow boundary conditions were used on the sides and bottom of the domain. Scoping studies with an existing simulator, PFLOTRAN, which used Richards' equation to model flow in the vadose zone, indicated that the water table had a head difference of approximately 12m from the side above the seepage pond to the Fourmile Branch. This difference was prescribed in the model to ensure that the background flow present in the saturated zone would be approximated, as implementation of Richards' equation in Amanzi was not fully realized. Using this configuration, a single-phase steady-state Darcy flow problem was solved using 256 cores of the Cray XT4 system at NERSC. It is important to note that since the entire domain is saturated in this simplified model, there is significant flow in the Upper Aquifer Zone (UAZ) toward Fourmile Branch.

For the model of the F-Area seepage basin, a Geochemistry PK with 17 primary aqueous (or basis) species that were considered: Na^{+}, Ca^{+2}, Fe^{+2}, K^{+}, Al^{+3}, H^{+}, $N_{2(aq)}$, NO_3^{-}, HCO_3^{-}, Cl^{-}, SO_4^{-2}, HPO_4^{-2}, F^{-}, $SiO_{2(aq)}$, UO_2^{+2}, $O_{2(aq)}$, tracers along with 12 secondary aqueous complexes, 11 kinetically reacting minerals, and 8 equilibrium surface complexes on 3 surface sites (>SiOH, >FeOH, >AlOH). All reactions but mineral precipitation-dissolution are equilibrium based. For the low pH (2.9) basin discharge fluid, phosphate complexes $UO_2H_2PO_4^+$ and $UO_2HPO_{4(aq)}$, and the fluoride complex (UO_2F^+) were found to be important in reducing uranium retardation, although the geochemistry in the model is considered preliminary.

The results of the complete flow and transport simulation using the unstructured mesh are shown in Figure 22. The simulation of this numerical model was run on 256 cores of the Cray XT4 system at NERSC for a time of 10 years. In this simulation the flow PK ran first, and computed the steady-state saturated flow field with a parallel linear solve across the 256 cores. This flow field was used in the reactive-transport by the Transport PK to advect the total component concentrations of the 17 primary species and one non-reactive tracer in each time step. After this advective step, the Geochemistry PK performed a nonlinear solve to invoke the reactions and update the concentrations. All of the required reaction processes were executed in this simulation. Figure 22 shows that the tracer has advanced from the seepage basin to the Fourmile Branch, and is beginning to spread laterally toward the far boundary. The concentrations of the non-reactive tracer and UO_2 are the same at the seepage basin; the isosurfaces at 1.5E-6 clearly show retardation of the uranium plume as a result of geochemical processes.

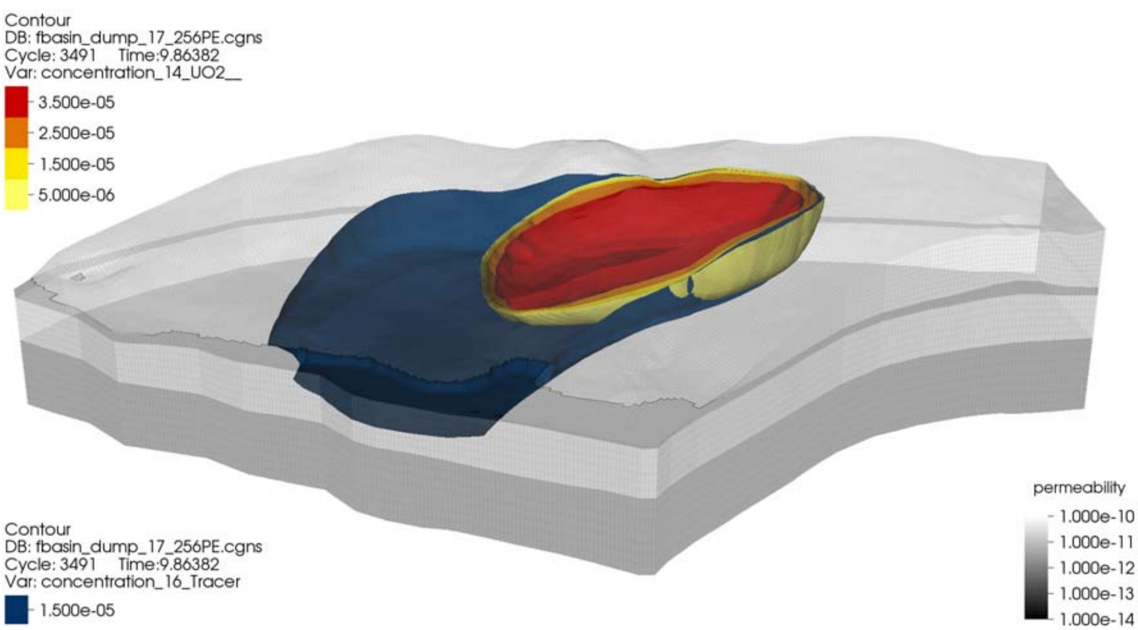

Figure 22. Iso-surfaces of the Uranium plume (yellow and red) and a non-reactive tracer (blue) are shown at 9.86 years for the unstructured mesh F-Area seepage basin model described above. The simulation was run on 256 cores of the Cray XT4 system at NERSC. The 17-component chemistry model was used and the retardation of the Uranium plume relative to the non-reactive tracer is evident, as the tracer has already reached the Fourmile Branch.

Two time-dependent, multiphase flow simulations of the F-Area seepage basin were also performed using the structured mesh approach based on BoxLib. A region of 1.6 km by 1.6 km to a depth of 100m was used in the model. Capillary pressure was used to define the saturation in the vadose zone. Hydrostatic equilibrium with a prescribed water table location was specified up gradient of the seepage basin with a lower water table downstream representing Four Mile Creek. An infiltration rate of 6 x 10^{-7} m/s was specified over the surface of the domain. The seepage basin, which provides the source of the contaminant, was modeled with a higher infiltration rate. For this example, a geostatistical realization of the subsurface was used with a mean permeability of 833 mD. For the simulation, a base mesh of 512x512x64 was used with adaptive refinement used to increase the resolution around the contaminant plume resulting in an effective resolution of 1.5 meter around the plume. The simulations of nonreacting and reacting tracers were performed on 264 and 2304 processors on the Hopper XE6 at NERSC.

In the first structured mesh simulation, the contaminant was represented as passive tracer, Na+. Figure 23 shows the plume after 39, 255 and 441 days. The first image shows the density of Na+ at early time as the tracer has reached migrated through the vadose zone and is beginning to propagate into the saturated zone. At the later times, the Na+ is seen being transported through the saturated zone. The high flow rate out of the seepage basin causes groundwater flow in all directions away from the trench. The second simulation includes the detailed geochemistry described above. Images of the contaminant as it migrates through the domain are shown in Figure 24 relative to the tracer. Again, the simulation shows the impact of geochemical processes on the uranium plume.

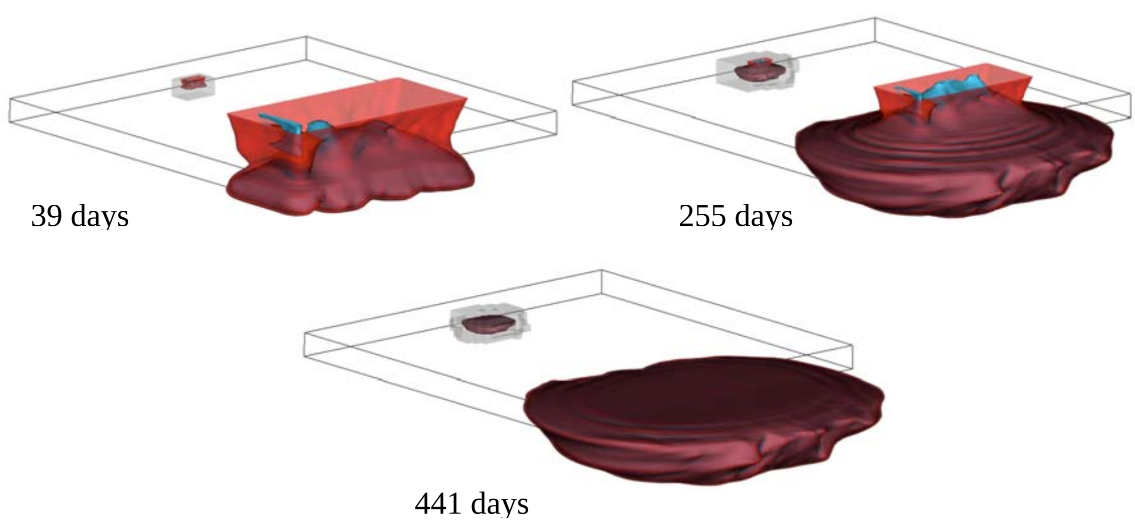

Figure 23. Examples of the simulation using the structured grid approach, showing the density of Na+ after 39, 255 and 441 days. The grey highlighted domain indicates where the grid was refined; the density of Na+ near the basin is shown as well as an enlargement of the Na+ plume. The contours are shown for values of 0.0275 and 0.0325 mol/L.

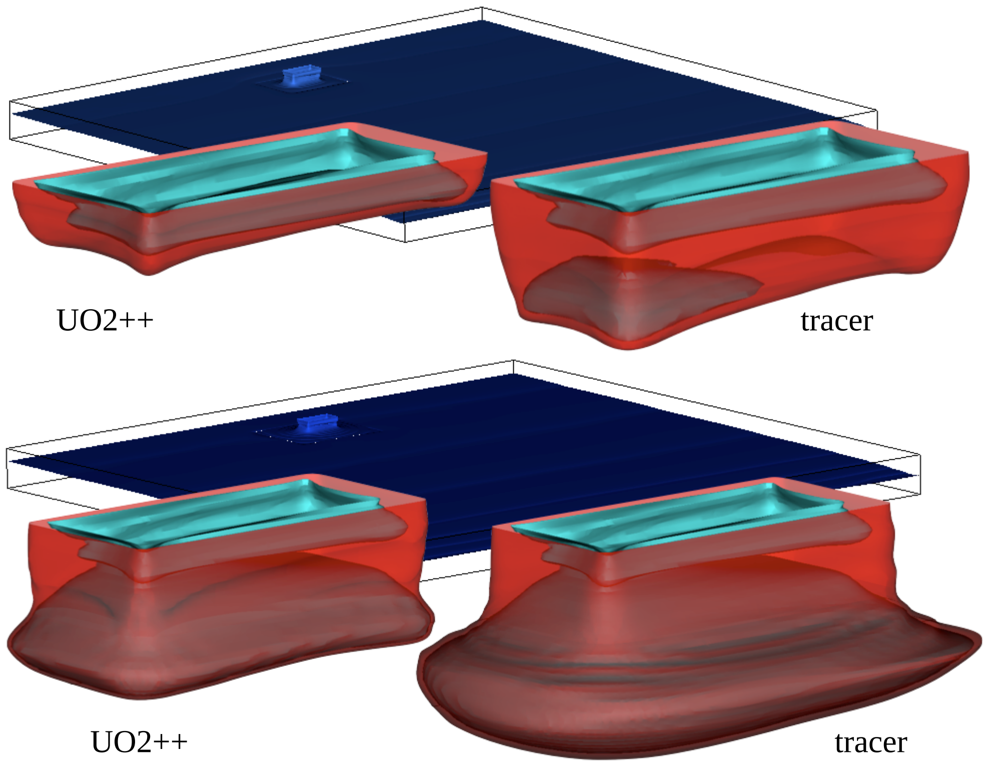

Figure 24. Examples of simulation output using the structured grid approach, showing the density of UO2++(aq) and a passive tracer(aq) after 24 days (top) and 105 days (bottom). The density near the basin, as well as enlarged images of the basin region area shown. The isosurfaces are 3.2e-5 mol/L (blue) and 2e-5 mol/L (red). The chemical reaction effectively retards the propagation of UO2++ into the domain. The background images show the saturation of water at the given times.

8.4 Discussion

In the Phase I Demonstration, the Multi-Process HPC team made significant progress towards its goals. Components for prototype toolsets of Amanzi were developed quickly and implemented. Parallel, unstructured hexahedral mesh and parallel single-phase flow and reactive transport capabilities were developed. All of the targeted geochemistry was implemented and run for the demonstration. The Trilinos framework was used to reduce the code development time. In the future, development of techniques to seamlessly use capabilities from multiple frameworks will be explored.

The Multi-Process HPC development team overcame a number of challenges during this first phase of development. For example, the Stk_Mesh database (from Sandia National Laboratories) was targeted as the mesh library because it is part of the Trilinos framework. However, it turned out that the current Stk_Mesh design targeted a very low-level infrastructure and failed to provide the higher-level functionality that was expected from a mesh database library. The lack of functionality caused a delay in the Mesh Infrastructure that ultimately led to a compromise in the flow process kernel for the demonstration. As a result, single-phase flow was implemented, instead of the full Richards' model for flow in the vadose zone.

The simulations performed using structured AMR methodology represent accomplishments beyond the defined scope of the Phase I Demonstration. These simulations demonstrated the ability of this methodology to model time-dependent multiphase flows with representative geochemistry on high-performance parallel architectures.

Developing an open-source community code with a complete open-source tool chain has many advantages, including reduced costs because the tools are freely available to the team and potential collaborators. In addition, the team was able to leverage expertise on algorithms and architectures, such as the Trilinos framework. Good examples of joint code development with ongoing projects include: refactoring the Exodus II writer in LaGrit to properly handle auxiliary data (i.e., side sets), and contributing bug fixes to the MOAB project. In turn, the lead developer of MOAB (on very short notice) added an important ghosting feature that was needed for Amanzi. This is an excellent example of how community codes can work, reducing both costs and development time for everyone in the long run.

A number of improvements and advances are under consideration for future development of Amanzi. For the next development phase, the major areas of emphasis will include integration with the Platform, enhancing flexibility and robustness of Amanzi, increased efforts on verification and validation, and developing advanced algorithms. Integration with the Platform Thrust will consist of being able to execute simulations created by a user through the Platform Toolsets and returning output for visualization and analysis. For the Phase 1 demonstration, this linkage was not yet operational.

For the short time allocated for the Phase I development effort, only essential features of Amanzi were included and non-essential features were scheduled for future development. For example, to simplify the current implementation, degenerate hexahedral cells (e.g., pinched-out cells) were not supported. This capability will be added over the next year using recent advances in finite difference methods. Similarly, while the Phase I development focused on an unstructured mesh capability, structured meshes proved to be a valuable addition. A more advanced structured mesh capability is targeted for implementation in Amanzi over the next year. In addition to the improvements listed above, Richards' model for unsaturated flow will be fully implemented. This development will provide a unique capability for Amanzi, namely a parallel unstructured grid capability for modeling flow in the vadose zone. In future development phases, prototyping of advanced algorithms that are anticipated to have significant long-term transformational benefits will continue. For example, adaptive mesh refinement (AMR) is anticipated to play a significant role in Amanzi as it matures. Similarly, capabilities to automatically generate the derivatives needed for parametric sensitivity analysis and optimization will also be explored.

Significantly more verification testing and additional benchmarking of Amanzi with existing simulators will be performed to ensure that the code is functioning properly. In addition, automated tests will be established to ensure robust testing for developers.

9. SUPPLEMENTAL PHASE I DEMONSTRATION ACTIVITIES

Two other working groups initiated efforts during Phase I to demonstrate additional ASCEM capabilities. These efforts were defined as opportunities to be tackled only if feasible, and if such activities would not detrimentally impact the development of capabilities required to meet the defined ASCEM Phase I goals associated with the F-Area Demonstration (Sections 5-8). These supplemental development and demonstration activities were performed to engage other communities and to advance ASCEM capabilities using datasets beyond that available for the SRS F-Area. This work helped to initiate the integration of ASCEM capabilities and lay the groundwork for Phase II demonstrations.

The two supplemental efforts focused on Waste Tank Performance Assessment (PA) and Deep Vadose Zone problems. The Waste Tank PA demonstrations addressed two identified needs: 1) effective interrogation and visualization of Monte Carlo uncertainty quantification results from transient 3D HPC simulations; and 2) explicit resolution of fine-scale features in man-made structures such as fractures, other discrete flow paths, and thin liners. The Deep Vadose Zone demonstration initiated development of model setup capabilities needed to translate conceptual models into computational grids.

9.1 Waste Tank Demonstration

Waste tank closures and disposal of residual waste in engineered containment systems are common components of remediation and operations across the DOE EM complex. Examples include tank closures, salt waste disposal, solid waste disposal vaults, grout encapsulated components, and *in situ* stabilization of structures under decontamination and decommissioning. Engineered barriers present unique process and simulation requirements in the form of geometries, materials and associated properties, and physical and chemical processes in comparison to purely geologic systems. To be responsive to these needs, ASCEM established a working group to pursue demonstrations that specifically address challenges related to PA modeling for systems involving engineered features.

ASCEM actively engaged the PA user community to seek input regarding areas of greatest need (see, for example, the ASCEM User Suggestions Document, October 2010 [http://ascemdoe.org/About/about-docs.html] and the Performance Assessment Community of Practice Technical Exchange, held April 13–14, 2010, in Richland, Washington [http://srnl.doe.gov/copexchange/agenda.htm]). Two key areas of need were selected to be the focus of Phase 1 demonstrations:

- Need for improved capabilities to interrogate, interpret, and explain results and uncertainties for performance assessment simulations
- Need for improved simulation capability to efficiently represent systems with highly contrasting material properties often involving thin features.

The initial demonstrations responding to the PA needs involve three components of ASCEM development: visualization; uncertainty quantification; and HPC. Visualization and UQ are described in Section 9.1.1, where new ASCEM visualization tools were developed and demonstrated using 3-D Monte Carlo simulation results obtained from a representative model. In Section 9.1.2, an adaptive algorithm is used to refine the computational mesh for the HPC simulator in the vicinity of a discrete flow path defined by a large contrast in permeability, resulting in great computational savings and high accuracy.

9.1.1 Visualization Component: Interrogation of UQ Results from 3-D Simulations

The PA process used to provide technical underpinning for DOE-EM remediation and disposal actions places high value on capabilities for analysis and synthesis of model outputs. Interrogation and interpretation of the vast amounts of simulation output associated with Monte Carlo uncertainty quantification for complex systems (3-D, engineered features, reactive transport) can be a time-consuming exercise and a source of potential errors. Efficient means for interrogating and interpreting model results are needed, because reviewers, regulators, and stakeholders expect the user to effectively identify the key parameters and processes and explain how they influence system behavior and the decision being made.

9.1.1.1 Objective and Specific Goals

The objective of this demonstration activity was to use visual data exploration and analysis techniques to interrogate and interpret transient 3-D simulation results with uncertainty quantification for an illustrative problem of interest to EM: a waste tank performance assessment. The specific goals pursued in this demonstration were to:

- Provide a high level visualization of the contaminant concentration for individual and the median of all Monte Carlo realizations
- Visualize the velocity field for individual realizations and the median of all the realizations
- Precisely visualize median values of any variable to perform selections on the data for further analysis
- Provide further tools for analysis of an output and its relation to the input parameters, based upon the selections performed in the previous objective.

9.1.1.2 Approach

A simple three-dimensional engineered barrier system, motivated by performance assessment of tank closures at SRS and Hanford, was defined (see vertical slice illustrated in Figure 25). In order to proceed with visualization development in parallel with HPC simulator development, PFLOTRAN was used to simulate flow and transport for the demonstration scenario. For this hypothetical model, 200 simulations were conducted with varying permeability, porosity, and flux parameters. The results of these runs consist of 200 input files containing contaminant concentration, velocity, and liquid saturation data for each simulation.

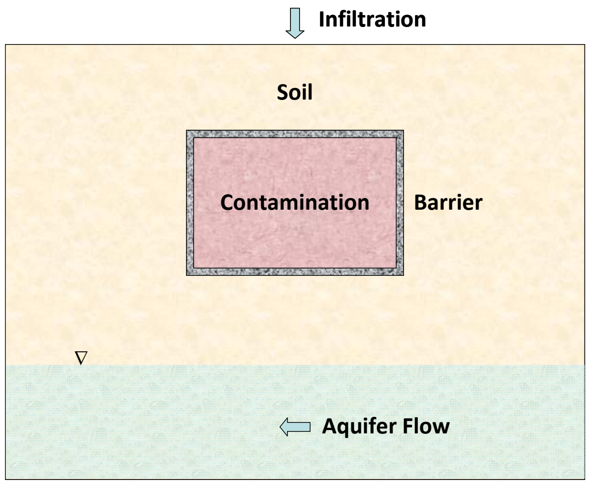

Figure 25. Vertical cross section of hypothetical moderate conductivity (*K*) waste form embedded in a low-hydraulic-conductivity engineered barrier.

Using the approach described in Section 6.2 of this report, the ASCEM team leveraged the VisIt visual data exploration and analysis applications and scripts created to facilitate that process. To apply the tool to the waste tank problem, modifications had to be made to these scripts. These changes were accepted by the VisIt development team and incorporated into VisIt Version 2.1. A post-processing script computed the median and other statistical quantities with the same file data structure as PFLOTRAN output files. For streamline and isosurface visualizations, PFLOTRAN data was loaded directly into VisIt. For interactive analysis of the data Python, Numerical Python (NumPy), and VisIt's Python scripts were used to automate reading 200 realizations and create additional plots such as histograms and scatter plots. These prototype visualization tools are designed to be integrated with ASCEM HPC simulators and uncertainty quantification tools in the next phase of development.

9.1.1.3 Accomplishments

For the Phase I waste tank demonstration, the team:

- Implemented enhancements to the VisIt PFLOTRAN file loader, and submitted these code changes to the VisIt development team. These changes were released to the broader scientific community as part of VisIt 2.1.

- Developed software for computing the median, mean, and other statistical data sets from the multiple realizations. In addition, they used multiple visualization techniques to produce images from individual or median of all realizations, such as visual display of isosurface of contaminant concentration (Figure 26a) and flux steamlines (Figure 26b).
- Designed and developed an ASCEM Graphical User Interface (GUI) leveraging VisIt's Python APIs to automate much of the visualization to perform further analysis of input parameters and output data (Figure 27).

The user interface window in Figure 27 illustrates how a user can select the realization or statistical output to display (e.g., median), choose the output variable (e.g., concentration), and select the time and location for diagnostic plots (e.g., histograms, crossplots). For both input and output variable histograms in Figure 27, the brown histogram indicates the full range of values across all realizations. For the output, this is the contaminant concentration at a selected (*x*, *y*, *z*, *t*) point; for the input, this is the initial infiltration flux. The overlain histogram (in green) represents the selected realizations. In this case, they are the realizations for the contaminant concentration that is above the 95% quantile at the selected spatial point. The scatter plot shows the input flux of each realization plotted against its contaminant concentration. These plots clearly show that high infiltration flux is strongly associated with high concentration for this system. On the lower right is a plot of the contaminant concentration at a (*x*,*y*,*z*) point over time. The four different lines represent the 5%, 50%, 95% quantiles and the mean.

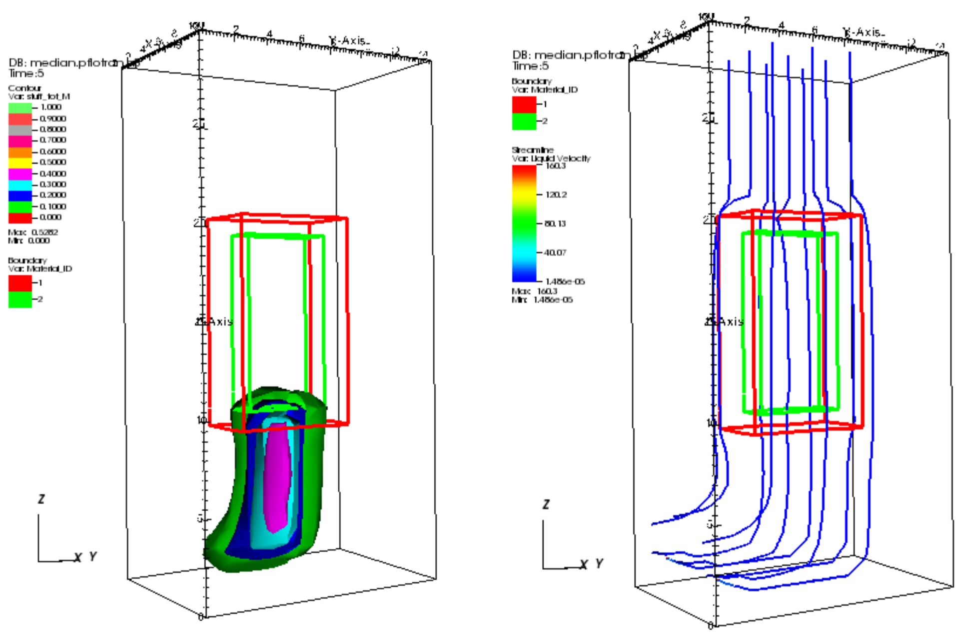

a) Concentration Isosurfaces **b) Streamlines**

Figure 26. Visualizations of (a) contaminant concentration isosurfaces and (b) flux streamlines.

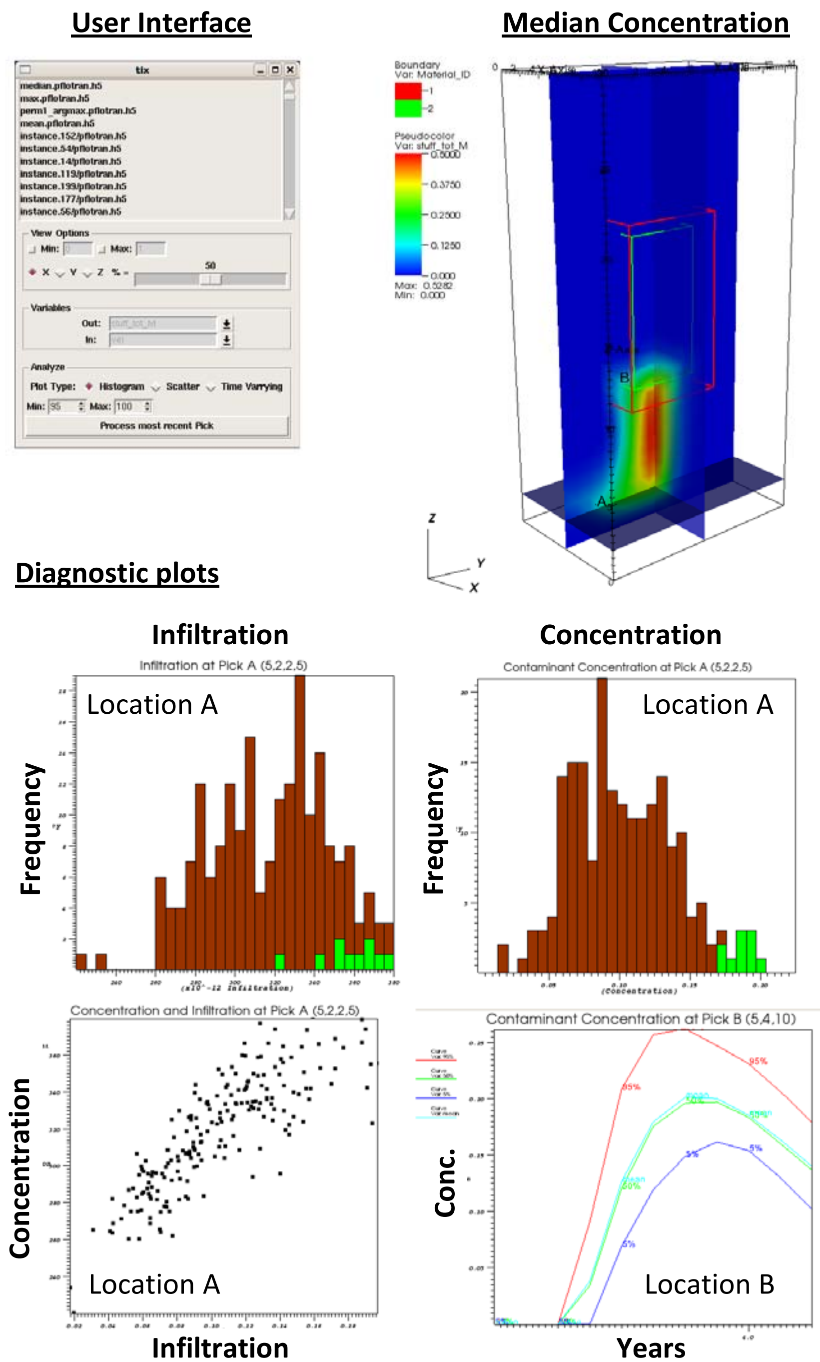

Figure 27. Plots generated from prototype UQ visualization GUI, including concentration slice, input and output histograms, scatter plot, and central tendencies and selected quantiles of time varying concentration at a selected location.

9.1.1.4 Discussion

The UQ demonstration associated with a waste tank performance assessment was pursued as an optional opportunity to develop additional tools that enable the user to efficiently visualize and interrogate UQ results. This demonstration includes the ability to study statistical trends and key characteristics of ensemble output generated by Monte Carlo simulation runs. Using PFLOTRAN for a hypothetical problem of a waste tank, initial concepts for ASCEM visualization were used to perform visual data exploration and analysis of Monte Carlo simulation runs. Beyond the capability to visualize individual instances afforded by VisIt, a user can apply these ASCEM enhancements to quickly determine the input parameters that have the most impact on concentration.

Creating an interactive application for efficiently visualizing UQ output provides flexibility in applying different techniques to interrogate extensive amounts of output. This leads to a better understanding of the factors that may affect performance of a disposal system. Moreover, the visualization capabilities complement the sensitivity analysis and uncertainty quantification tools described in Section 7. Future efforts will include refinement and addition of functionality to these tools and integration of visualization concepts with ASCEM tools for uncertainty quantification.

9.1.2 HPC Waste Tank Component: Adaptive Meshing to Resolve Fine Scale Features

In addition to modeling large-scale hydrological flows and the resultant transport of contaminants, the HPC component of ASCEM also needs to be able to efficiently simulate man-made near-field features such as waste tanks and similar engineered barriers. These structures often contain small-scale features such as steel, high-density polyethylene, or geosynthetic liners, and discrete “fast flow” paths formed by conduits, separation gaps, or cracks in concrete or grout. Efficient and accurate simulations of these systems require a high mesh resolution in the vicinity of small-scale components, and lower resolution elsewhere to minimize the total number of computational cells. Standard meshing approaches are inefficient or impractical for these applications.

9.1.2.1. Objective and Specific Goals

The purpose of this task was to demonstrate that the HPC simulation software, Amanzi, being developed under ASCEM can be used to effectively treat resolution of fine-scale features. The specific case considered is a simplified model for a degraded waste tank, depicted in Figure 28. In this example, high permeability regions indicated in the figure are used to represent a postulated fast flow conduit through the tank. The specific goal of the simulation was to demonstrate the capability to accurately resolve the infiltration of water into the tank. To accomplish this task, the capability to resolve the finest feature of the model domain was needed, namely, a 1 cm tank liner.

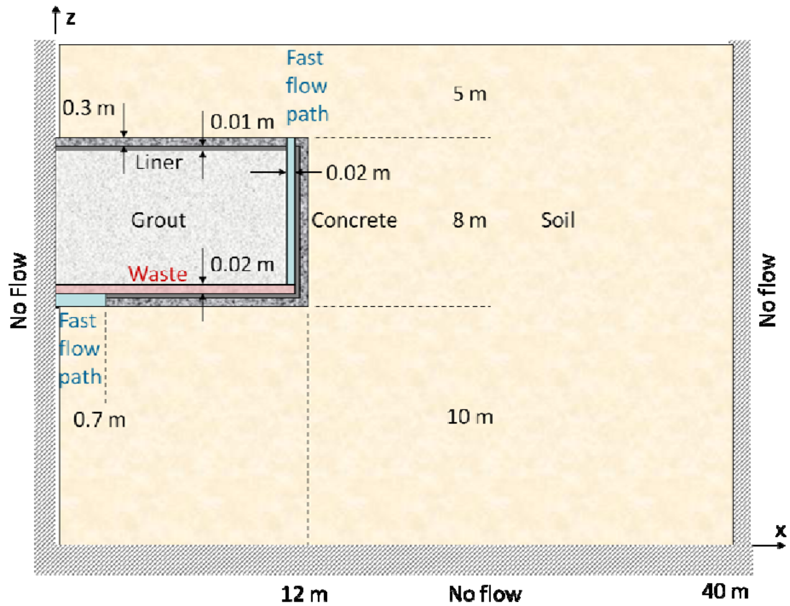

Figure 28. Illustration of geometry for tank problem. The permeabilities of soil (clay), grout, concrete, liner and fast flow path are modeled by 1 mD, 10 mD, 0.1 mD, 0.05 mD and 2000 mD, where mD stands for milliDarcy.

9.1.2.2. Approach

Accurate modeling of leakage into the tank requires that certain features be resolved on the scale of 1 cm in a domain on the scale of tens of meters. There are several potential strategies for how to treat this problem. The approach used for the Phase I Demonstration is based on a block-structured adaptive mesh refinement (AMR) methodology, BoxLib (https://ccse.lbl.gov/Software) that leverages work funded under the DOE/ASCR program. The prototype HPC simulator developed as part of this demonstration is envisioned to form the basis for the structured mesh solvers under Amanzi. Using this methodology, a coarse uniform mesh was initially defined to cover the computational domain. Refinement criteria were then used to identify parts of the domain where additional resolution was required. A collection of boxes was defined to include the points identified for refinement. Each of these boxes was used to define a grid patch at a finer resolution to represent the solution in that region. This procedure was applied recursively at increasing finer levels of resolution until the desired resolution was obtained. In the block-structured refinement approach considered here, the finer resolution data were organized in large aggregate grids containing a number of grid points. Thus, the irregular work associated with adaptive refinement was focused on the relationship of larger grid patches that tile the region of interest, in contrast of the need to identify the relationship between individual cells in a cell-by-cell refinement approach. The overall approach to adaptive refinement described here is discussed in detail by Pau et al. (2009).

9.1.2.3. Accomplishments

For the Phase I demonstration, a time-dependent, multiphase flow simulation of the infiltration of water into the waste tank was performed. The simulation was intended as a proof-of-concept rather than a realistic analysis of a waste tank. For that reason, synthetic properties were used to represent the different materials and a number of simplifying assumptions were used. In particular, the same relative permeabilities were assumed to represent all of the materials and variations of capillary pressure were ignored. Another simplifying assumption included use of a simple passive tracer to represent the waste at the bottom of a tank. The waste tank geometry was also represented as a simple Cartesian geometry.

The overall simulation was performed in a domain that was 20 m x 20 m. This region was initially represented with a coarse mesh of 256 x 256. Adaptive mesh refinement was used to refine around the different geometric features resulting in a mesh spacing of 0.5 cm at the finest resolution. For the simulation, conditions for the waste tank and surrounding soil were initialized at residual water saturation. This represents a situation in which the waste tank had previously drained, leaving only residual water with a thin layer of contaminant at the bottom of the tank. A high water load was imposed at the surface of the domain that induced water infiltration through the soil to the waste tank. Figure 29 shows images of water densities after 13 and 25 days, including the large-scale features of the simulation. The top image shows results at an early time. Early in time, the simulation predicts that water accumulates on the top of the tank where the infiltration encounters a low permeability region. In this image, water began to leak down the high permeability channel into the waste tank. Because of the large permeability contrast, the water flowing into the tank has a fairly low saturation. The bottom image shows the solution later in time. In this image, the water has reached the bottom of the tank and has begun to pool in its lower corner. The image also shows the accumulated water beginning to flow off the top of the tank.

The range of scales illustrated in this example makes it difficult to see the finer-scale feature of the flow. In Figure 30, an enlargement of the region in the neighborhood of the top right corner of the tank is shown. In this image, the flow of water into the high permeability channel and the spillage of water across the top of the tank are readily apparent. In Figure 31, an enlarged region near the lower-right-hand corner of the tank is shown. This image shows the water that has flowed in along the high permeability path beginning to accumulate at the bottom of the tank after 25 days. In the left image of Figure 31, the concentration of tracer after 25 days at the bottom of the high permeability zone is shown. It illustrates how the pooling water is beginning to mobilize the contaminant remaining in the tank.

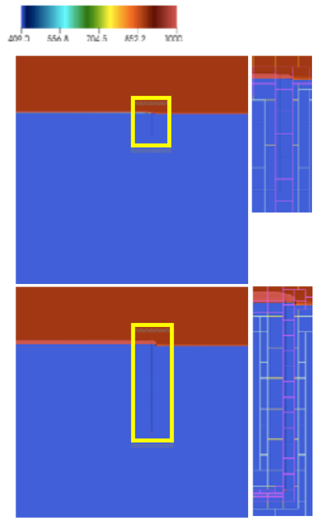

Figure 29. Results of simulation using adaptive mesh refinement showing the component densities of water after 13 days (top) and 25 days (bottom). Boxes in the top image show the finest and second finest levels of grids. Regions highlighted by the rectangular boxes are enlarged in images on the right, showing the multilevel grid structure

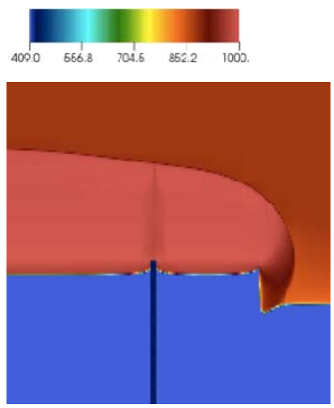

Figure 30. An enlargement of the right corner of the tank, showing the ability of the adaptive mesh approach to capture the details in the component density of water after 25 days.

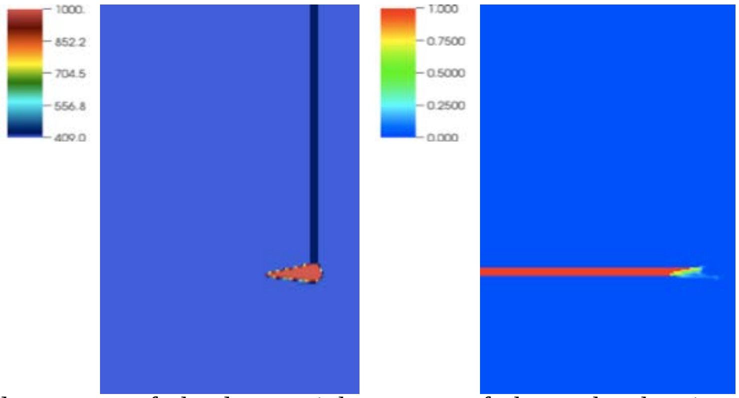

Figure 31. An enlargement of the lower right corner of the tank, showing the ability of the developed adaptive approach to capture the details of the component density of water (left) and concentration of tracer (right) after 25 days.

9.1.2.3. Discussion

The Phase I Demonstration of a waste tank performance assessment illustrated capabilities targeted for future integration with the Multi-Process HPC simulator, Amanzi. The problem was challenging in a number of respects. First, the simulation required resolution of features on the scale of one centimeter in a simulation that was performed on a 20 x 20 m domain. This required an adaptive mesh refinement approach to accurately represent the engineered system. A uniform mesh with equivalent resolution would have required 16 million cells compared to only 1 million cells using adaptive mesh refinement. The adaptive mesh refinement approach avoided numerical convergence issues common with this type of problem. The other challenging feature of this problem was the sharp contrast in permeability. In the simulations, there was a contrast in permeability of 40,000 between the high permeability channel and the liner, which normally causes numerical convergence problems. The linear solver for the pressure equation in the multiphase flow formulation was able to handle this without any significant convergence issues. In the next phase of ASCEM development, the capability to treat these types of engineered systems in three dimensions and with more realistic descriptions of the physical processes will be targeted.

9.2. Model Setup Tool

A prototype Model Setup tool was developed as part of another Phase I Demonstration opportunity. This supplemental application was initiated to align with efforts at the BC Cribs and Trenches at the Hanford Site in the arid southeastern portion of Washington State (Figure 32). The BC Cribs and Trenches is a component of the Deep Vadose Zone Applied Field Research Center, a collaborative effort between DOE EM-32 and DOE Richland Operations. The deep vadose zone is a focus of remediation efforts at the Hanford Site.

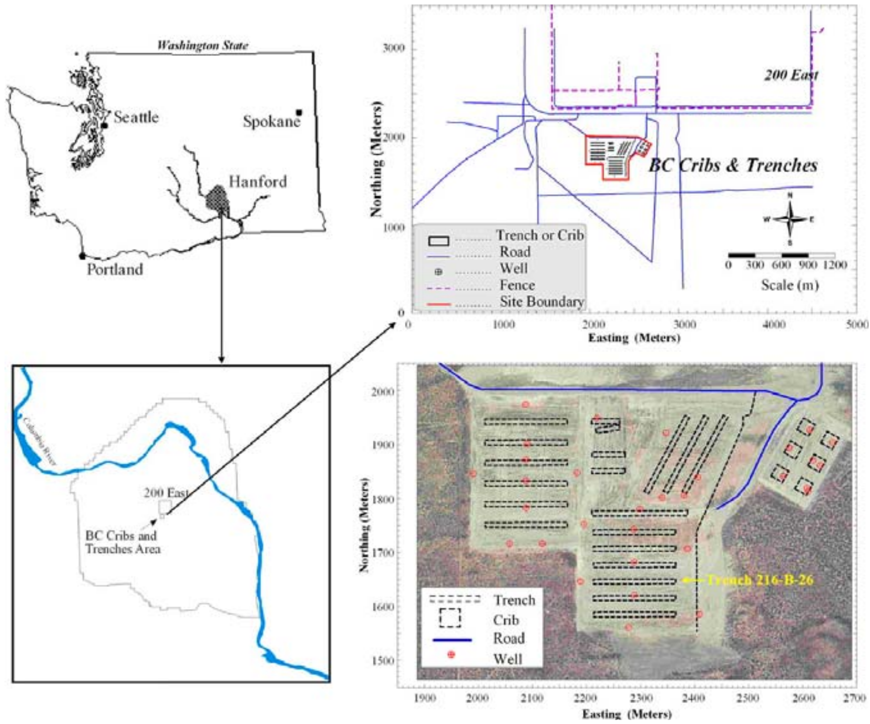

Figure 32. Location of the BC Cribs at the Hanford Site is Southeastern Washington State

BC Cribs and Trenches consists of 26 waste sites, including 20 unlined disposal trenches and 6 concrete disposal cribs utilized for the disposal of ~115,000 m^3 (30 million gal) of liquid radioactive mixed waste scavenged from the processing of uranium. The sodium nitrate rich mixed liquid waste (radiological and hazardous waste) was discharged to the trenches and cribs between 1956 and 1958, based on a concept specific retention, which assumes that the that the vadose zone would hold the liquid waste in place and prevent it from migrating to the water table.

However, wellbore data has confirmed that both vertical and horizontal migration of contaminants has occurred in the 100 m thick vadose zone beneath the BC Cribs Site. Core samples taken from the center of the 216-B-26 Trench (Figure 32) have confirmed the presence of chlorides, nitrates, sulfates, and other salts down to 44 m below ground surface (Serne and Mann, 2004). Because the contaminant distribution is associated with the distribution of fine-grained sediments, lateral migration (rather than specific retention) is now considered to be the primary process governing the limited vertical distribution of contaminants in the vadose zone. Critical for simulating lateral migration at BC Cribs is the correct representation of the BC Cribs subsurface stratigraphy.

9.2.1. Objective and Specific Goals

The overall objective of the Model Setup Tool is to develop functionality for efficiently translating between measured data, conceptual models, and computational grids. The process of developing a computational grid is a prerequisite for running a numerical flow and transport simulation; it is usually a labor-intensive process for model developers, one that requires visual inspection and comparison of the generated grid with data, as well as use of scripts and third-party software and visualization packages. Key to meeting this objective is the development of ASCEM capabilities that translate and visualize the conceptual model to the input required for the numerical model (grid specification, sediment layering, material property assignment) in the same environment as the numerical model, without the need to iteratively export data to external software packages.

9.2.2. Accomplishments

Efforts in Phase I focused on development of Model Setup tools for visualizing properties or features that are common components of conceptual and numerical model development, such as the stratigraphic layers. This was accomplished through the development of a Grid Viewer tool and integration of the viewer with the ASCEM UQ GUI described in Section 7.2. Three suites of functions were developed in the Grid Viewer to facilitate visualization of computational grid. The “Grid Control” functions provide options such as toggling (i.e., displaying) on and off axes, bounding boxes and dual views; the “Slice Control” functions permit visualization of the interior of the volume; and the “Cell-Render Control” functions control the visualization of the individual stratigraphic layers or material properties, and allows cutaway views using slider planes to visualize the interior domain. To visualize the stratigraphy, cell-based rendering is used, which assigns colors based on integer values representing the material types without interpolation. This method offers a distinct advantage over other software packages that assume that the integer assignment of layers is a continuous function. To facilitate seamless iteration between measured data, conceptual model, and the numerical model grid, the Model Setup Tool will eventually be able to read a variety of different formats files that define the grid and associated properties, including: Amanzi grid definition files, stratigraphic information (from borehole picks or ASCII gridded files from Petrel or Earthvision), and measured material properties (such as unsaturated hydraulic properties, or moisture content).

The different Grid Viewer panels permit inspection of the exterior and interior of the domain. Figure 33, developed using the Slice Control panel, shows the stratigraphy of the BC Cribs domain (note that the grid is not shown in this figure). This figure illustrates a three-dimensional view of the BC Cribs Site key sedimentary layers. Rapid 360^{o} rotational views can be obtained by clicking and dragging the mouse, rendering the exterior view useful for inspecting material assignments at the domain boundaries. Figure 34 shows how the interior of the domain can be visualized using the Slice Control functions. Using the slider bars, a user can toggle through successive planes in the x, y, or z direction. The screen-shot shown in Figure 35 shows three different planes selected with different transparency and color schemes (i.e., transfer functions); a user can choose the degree of transparency using the slider bar on the Slice Control Panel, or change the color scheme by accessing the drop-down box containing the ~40 different transfer functions (Figure 36).

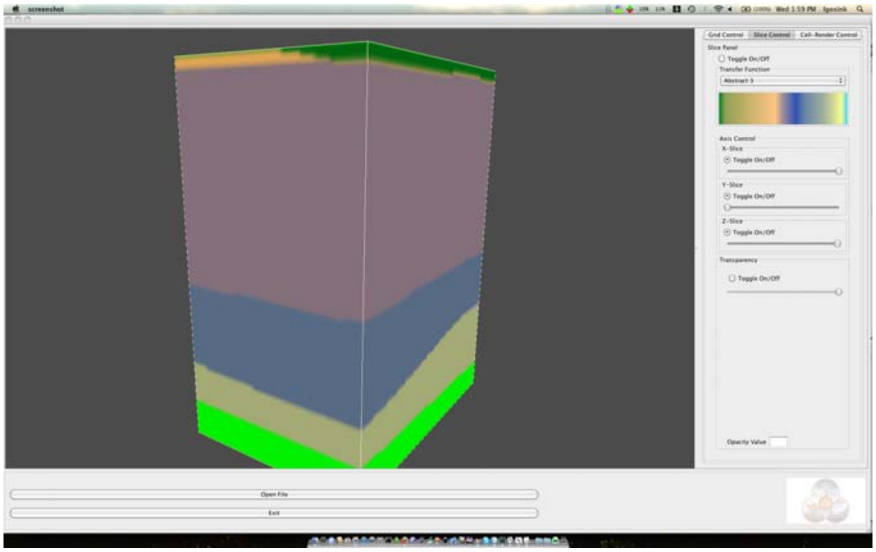

Figure 33. Three-dimensional view of the BC Cribs stratigraphy developed using the Model Setup Tool Grid Viewer Functions.

Examples of the Cell-Render Control functionality are shown in Figures 36a-c. A screen shot of the BC Cribs stratigraphy mapped to a 100 × 75 × 103 grid, where discretization is an increment of 1 m in each coordinate direction, is shown in Figure 36a. The cell based rendering approach clearly shows each of the layers mapped to the grid that the user has selected for viewing. At the top of the domain, however, the distinction between layers is obscured by the selected coloring scheme (i.e., transfer function). In Figure 36b, the screen shot demonstrates the capability for assigning a specific color to a material type or sedimentary layer. In this example, red is selected to represent the backfill layer at the top of the BC Cribs Site. The resulting view is shown in Figure 36c. Figure 36d demonstrates the ability to create cutaway views using the slider bars in the Cell-Render Control panel that allow for inspection of the material assignment to the grid. An example of the dual view functionality that is accessed through the Grid Control panel is shown in Figure 37. In this example, views were initially selected to examine the right and left domain boundaries, but were then rotated to inspect the top and bottom boundary views.

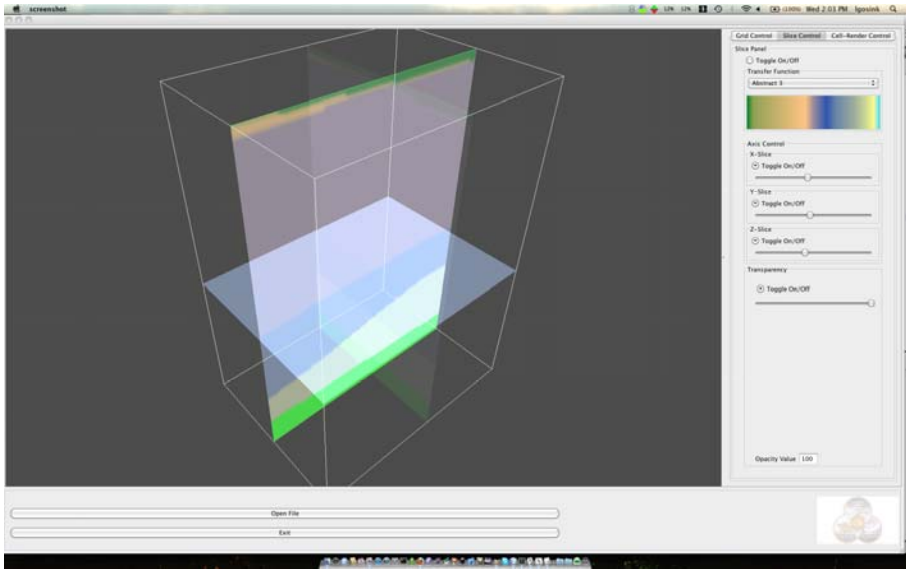

Figure 34. Cross-sectional views of the BC Cribs stratigraphy developed using the Slice Control Functions of the Model Setup tool, where transparency has been chosen to visualize the interior of the domain.

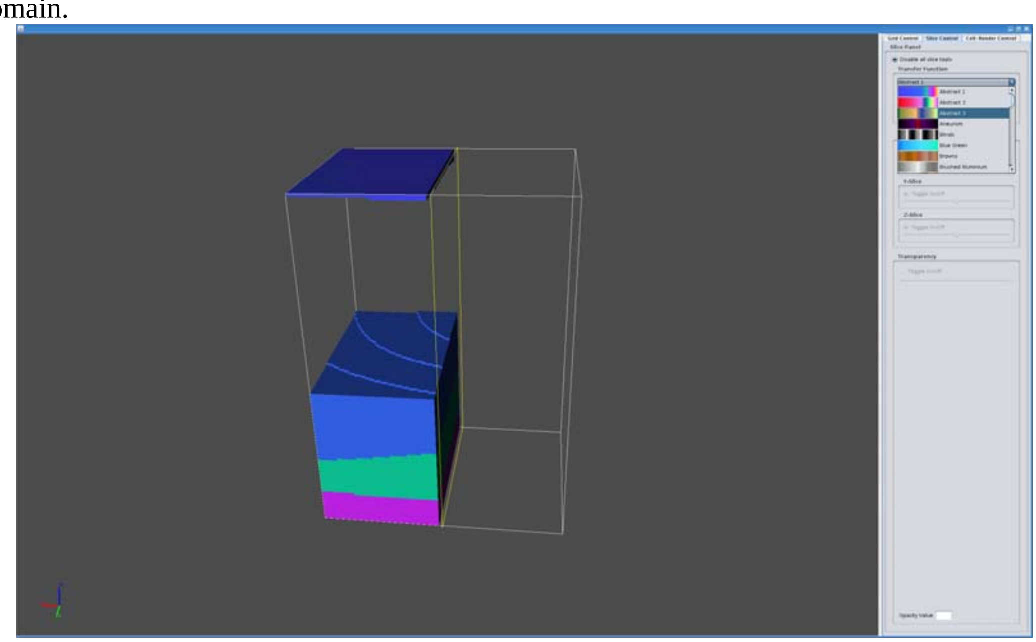

Figure 35. Cross-sectional view of BC Cribs stratigraphy showing the drop-down box in the Slice Control Panel that is used to access the different color schemes (i.e., transfer function).

a)

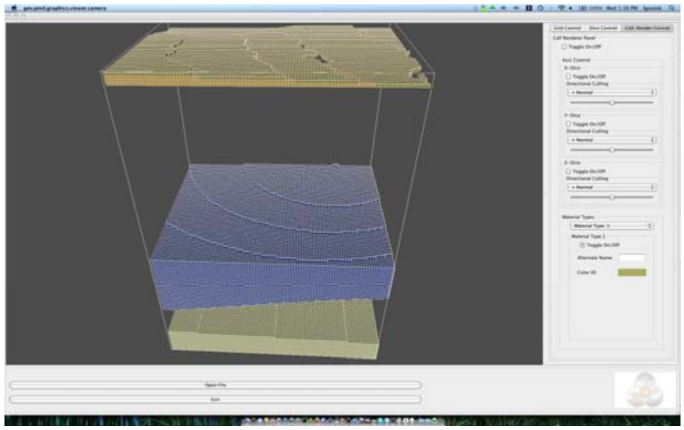

b)

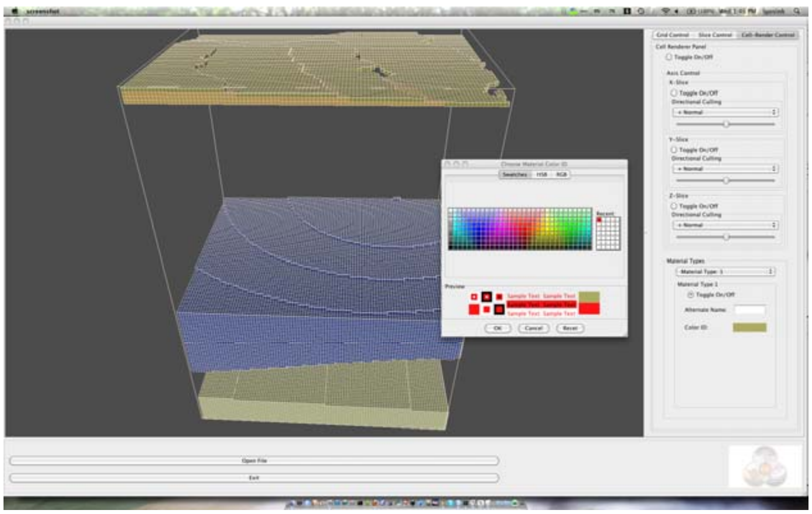

c)

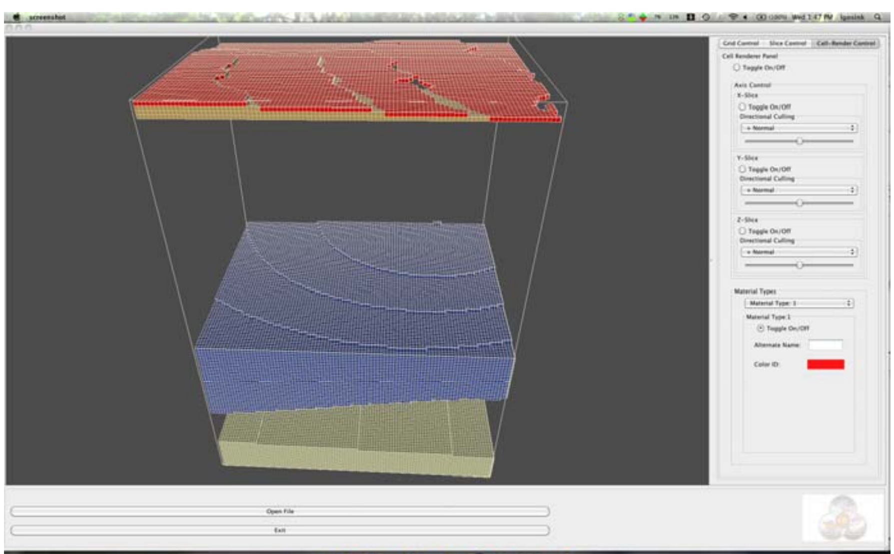

d)

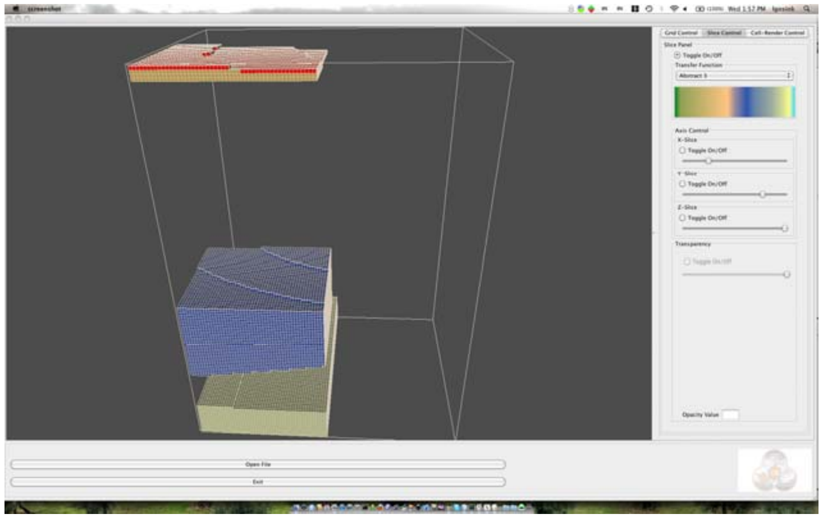

Figure 36. BC Cribs site geology showing: a) stratigraphy mapped to grid using existing color schemes; b) user selection of red for the type of a backfill material at the top of the domain to enhance visualization; c) the resulting color assignment; and d) a cutaway view of the stratigraphy mapped to the grid.

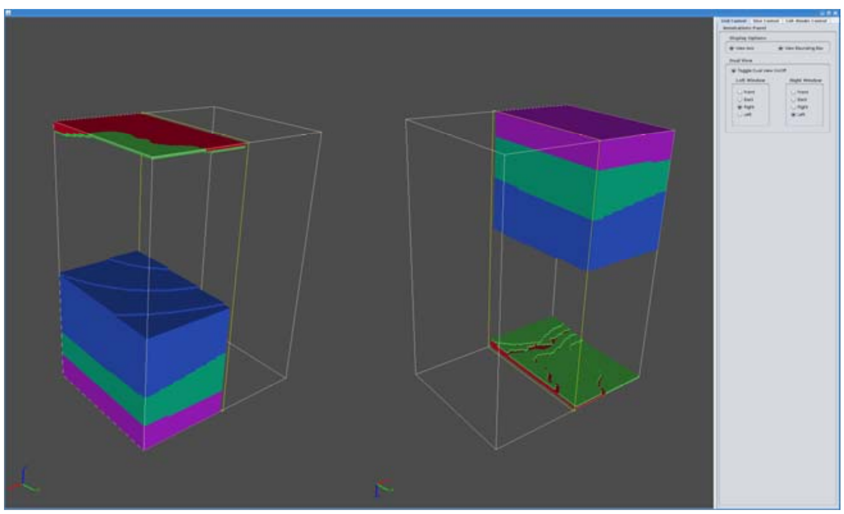

Figure 37. BC Cribs stratigraphy showing dual view functionality. Initially, views were selected to examine the right and left domain boundaries, but are now shown as rotated to inspect the top and bottom boundary views.

9.2.3. Discussion

For the Phase I Demonstration, a Model Setup Tool was developed and implemented. This tool is linked to the ASCEM UQ GUI and can also be used to visualize and link the subsurface layers and their associated properties to the computational grid. Several functions were developed to enhance the inspection of the interior and exterior of the domain. The principal advantage of this ASCEM tool is that the multistep process used for translating the conceptual model to the input required for the numerical model can be done in the same environment, without the need to export data to external software packages. The modular tool can be accessed through a web-based application and is written in Java so that it can be run on any operating system.

The initial Model Setup Tool advances will be enhanced in subsequent demonstrations. In particular, the tool will provide the functionality for reading conceptual model and measured data files and writing input files for the ASCEM simulator. Worksheets will be developed that will store conceptual model inputs (both qualitative and quantitative site data) that define the conceptual model. This will include a variety of different file formats for measured data, (Amanzi) grids, and stratigraphic definitions. Improvements in user options for the Grid Viewer will include viewing multiple planes in a single coordinate direction, and the ability to use isosurfaces to visualize material properties. New capability development will include the ability to compare different conceptual models and grid discretizations using a Model Setup Visualization Tool.

10. SUMMARY

The Phase I Demonstration, primarily performed from September 1-December 10, 2010, focused on advancing four of the several ASCEM components that are currently under development: Data Management; Visualization; Uncertainty Quantification; and HPC. Specific Phase I deliverables were developed for each of the four Demonstration components, with a plan to test ASCEM capabilities using common datasets derived from the contaminated SRS River F-Area Seepage Basins site. Use of common datasets facilitated coordination of Demonstration activities and is also expected to poise ASCEM for data sharing and integration across components, which will be the focus of the Phase II Demonstration. Leveraging SciDAC-, BER-, ASCR- and Advanced Simulation and Computing- (ASC-) developed advances and integration of existing open-source software are central tenets of ASCEM that are expected to lead to lower overall project costs and higher community acceptance.

Two supplemental Phase I Demonstration problems focused on Waste Tank Performance Assessment (PA) and Deep Vadose Zone problems. These efforts were identified to do the following: engage a broader set of working groups and end users than represented by the SRS F-Area demonstration; initiate studies focused on linking one or more ASCEM capabilities; lay the groundwork for ASCEM Phase II demonstration; and to consider simulation problems distinct from subsurface flow and transport, but critical to the intended ASCEM capabilities. The supplementary demonstration activity illustrates advances beyond those obligations identified in the Phase I Demonstration deliverables.

The Phase I Demonstration activities were coordinated by working groups housed under the Site Application Thrust. The working group members provided earth science expertise to the computationally-oriented ASCEM developers. Working group members helped to define the Phase I Demonstration goals, assemble the necessary input (conceptual models, data, process models, and other expert input), provide feedback during the Phase I development phase, and coordinate development of the individual components in the Phase I Demonstration.

Significant progress in advancing all four of the defined ASCEM capabilities was realized during the Phase I Demonstration. The Data Management component adapted and implemented a relational database as well as other open-source, web-based tools to allow users to easily ingest, browse, filter, graph, query, and output various types of data common to subsurface investigations. Tools were developed to handle both transparent data (such as wellbore concentration and lithology data) and opaque data (such as historical documents). A useful characteristic of the Data Management tools is that they can be used iteratively to display and query different subsets of the database based on subregion, characteristic, or parameter range specifications. The Visualization component of the Phase I Demonstration modified and extended open source VisIt software to facilitate visualization of data or features common to environmental remediation efforts, such as wellbore geometry; depositional information; hydrostratigraphic surfaces and topography; and the evolution of contaminant plumes. An advantage of these tools is the ability to visualize many different types of data (point, surfaces, volumes) in an uncluttered fashion, the joint visualization of physical features and contaminant concentrations, and the use of slider bars to navigate through temporal datasets (for example, to view the evolution of a subsurface contaminant plume within a physical framework).

For the Uncertainty Quantification component, ASCEM capabilities were developed to allow a user to choose model parameters and model outputs for a study, and to perform UQ analysis using a variety of different analysis approaches and types. An ASCEM GUI was developed as a framework for the UQ capabilities, which takes advantage of many open source UQ analysis approaches. A novelty of ASCEM UQ is its integration of the various methods, described in this report, which are available nowhere else in a single software analysis package. Bringing them together as a single ASCEM UQ tool thus represents a major advance.

Substantial progress was made on development of early prototypes of selected toolsets within the ASCEM Multi-Process HPC Simulator, now called Amanzi. A parallel unstructured hexahedral mesh capability, which can capture complex topography and hydrostratigraphy, was developed (meeting a Phase I deliverable). Building on the unstructured mesh capability, parallel single-phase flow and reactive transport capabilities were developed. Parallel simulations were run on 256 processors for a period of 10 years, exceeding the Phase I Demonstration goal of a two-year simulation run on 100 processors. All of the targeted geochemical processes were implemented in Amanzi's Reaction Toolset (aqueous speciation, mineral precipitation and dissolution, and sorption). A one-dimensional reactive-transport simulation with a five-component model of the F-Area geochemistry was performed and found to be in agreement with two existing codes. A more complex 17-component geochemistry model was run in the full F-Area reactive-transport simulation. Advances were realized with a structured mesh approach as well, which went beyond the Phase I requirements. A structured AMR capability, which used Amanzi's Reaction Toolset, was explored, exceeding goals established for the Phase I Demonstration. Parallel simulations of the F-Area seepage basins run on 2304 processors demonstrated the potential of this approach to model time-dependent multiphase flows with enhanced fidelity near engineered systems. For both the unstructured and structured simulations, results were visualized with the VisIt tool, which was also used in the Visualization demonstration.

ASCEM advances were also realized as part of Waste Tank and Deep Vadose Zone supplemental demonstration problems. The Tank Waste project demonstrated the use of new ASCEM tools to efficiently visualize and interrogate uncertainty results associated with 3-D Monte Carlo simulations of potential contamination due to the degradation of closed waste tanks, and an adaptive algorithm was developed using the unstructured approach described above to refine the computational mesh in the vicinity of a discrete tank flow path, which was controlled by a large contrast in permeability. Conceptual models associated with the Hanford Deep Vadose Zone formed the basis for development of an ASCEM Model Setup Tool, which was linked to the ASCEM UQ GUI. The tool can be used to visualize and link the subsurface layers and their associated properties to the computational grid. The principal advantage of this ASCEM tool is that the multistep process used for translating the conceptual model to the input required for the numerical model can be done in the same environment, without the need to export data to external software packages.

Collaboration hurdles were (and continue to be) surmounted as the large, multidisciplinary ASCEM team develops optimal mechanisms and shared vocabulary needed to coordinate and communicate efforts, respectively. The working group mechanism was found to be useful for aiding communication about the demonstration and for ensuring that the development of ASCEM capabilities are relevant to DOE-EM’s remediation effort. During subsequent demonstrations, the working groups will also be valuable for engaging site personnel in the ASCEM effort and for beta testing developing ASCEM capabilities. During Phase I, various collaborative tools were used to assist ASCEM development teams. Because subsequent demonstrations will focus on data sharing and component integration, it will be critical to develop a common framework for collaboration across the ASCEM team.

The Phase I Demonstration was designed to provide an early snapshot of specific ASCEM capabilities through describing advances associated with the defined Phase I F-Area effort, as well as two additional supplemental problems. It is the first of a series of ASCEM demonstrations that will be performed in a phased manner to correspond with development of the Platform and HPC Thrust components and with ASCEM releases. During FY11, the Site Applications Thrust will develop a long-range plan that will outline future demonstrations and consider associated demonstrations, tutorials, and documentation. Plans for specific component developments intended for FY11 were described throughout Sections 5–9. However, an integrated Phase II Demonstration plan will not be initiated until early 2011, allowing time to incorporate feedback from DOE and from the ASCEM team on the Phase I effort and experience, respectively.

While the most significant contribution of ASCEM is expected to be its integrated framework and associated computationally-efficient, open-source, modular, portable, and accessible characteristics, many of the individual ASCEM components already demonstrate performance or flexibility that exceeds what is available today. The development of an open source process-based computational framework that can be easily and consistently used across the DOE EM complex is expected to improve cleanup efficacy and decrease overall costs associated with the DOE legacy waste stewardship obligation.

10. REFERENCES

ASCEM-SITE-091310-01 (2010), ASCEM Site Applications Thrust Site Selection Task "Select Phase I Demonstration" Milestone, October 1, 2010.

Congress (2006), Energy and Water Development Appropriations Act. Conference Report on H.R. 2419.

Denham, M. and K. Vangelas (2010) Information Document for EM-32 Applied Science Initiative "Attenuation-Based Remedies for Metal and Radionuclide Contaminated Groundwater"—February 2010, SRNL-TR-2010-00047.

DOE (2000), *Status Report on Paths to Closure*; DOE/EM-0526; U.S. DOE Office of Environmental Management, Washington, DC.

DOE (2008a), Department of Energy, FY 2009 Congressional Budget Request, Environmental Management, Defense Nuclear Waste Disposal Nuclear Waste Disposal, DOE/CF-028, Volume 5.

DOE (2008b), Department of Energy (DOE), Engineering and Technology Roadmap: Reducing the Uncertainty in the EM Program (2008).

DOE/SC-1023 (2010), Complex Systems Science for Subsurface Fate and Transport. Report from the August 2009 Workshop, http://www.sc.doe.gov/ober/subsurfacecomplexity_03-05-10.pdf

DOE-ASCEM (2010), U.S. Department of Energy Advanced Simulation Capability for Environmental Management (ASCEM) Baseline Management Risk and Change Control Directive.

Doherty, J. (2010), PEST: Model-Independent Parameter Estimation. Australia: Watermark Numerical Computing. http:// www.pesthomepage.org

Dong, W., J. Wan, M. Denham, J. Seaman, S. Rakshit, T.K. Tokunaga, N. Spycher, and S.S. Hubbard (2010), B51C-0375: Geochemical Characteristics of the Contaminant Waste Plume in the F-Area of the Savannah River Site: From Kilometer to Micrometer Scales. Fall AGU, B51C-0375, San Francisco.

Flach, G.P. and M.K. Harris (1999), Integrated Hydrogeological Model of the General Separations Area (U); Volume 2: Groundwater Flow Model *(U)*. WSRC-TR-96-0399, Rev. 1.

Flach, G.P. (2004), Groundwater Flow Model of the General Separations Area Using Porflow (U). WSRC-TR-2004-00106.

Gorton I., G.D. Black, K.L. Schuchardt, C. Sivaramakrishnan, S.K. Wurstner, and P.S.Y. Hui (2010), GS3: A Knowledge Management Architecture for Collaborative Geologic Sequestration Modeling. 2010. In: *Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS 2010)*. IEEE Computer Society Press, New York, NY. doi:10.1109/ HICSS.2010.217

Hammond, G.E. and Lichtner, P.C. (2010) Field-Scale Model for the Natural Attenuation of Uranium at the Hanford 300 Area using High Performance Computing, W. Resour. Res. 46, W09527, doi:10.1029/2009WR008819, 1--31.

Jean, G.A., J.M. Yarus, G.P. Flach, M.R. Millings, M.K. Harris, R.L. Chambers, and F. H. Syms (2004), Three dimensional geologic model of southeastern Tertiary coastal-plain sediments, Savannah River Site, South Carolina: An applied geostatistical approach for environmental applications, Environmental Geosciences, 11 (4), 205–220.

Morel, J., M. Hall and M. Shashkov. (2001), A Local Support-Operators Diffusion Discretization Scheme for Hexahedral Meshes, J. Comp. Phys., 170, pp 338-372.

Morris, M.D. (1991), Factorial Sampling Plans for Preliminary Computational Experiments. Technometrics, 33, 161–174.

NRC (2000), U.S. Department of Energy's Environmental Management Science Program: Research Needs in Subsurface Science. National Academy of Sciences: Washington, DC.

NRC (2009), National Research Council (NRC), Advice on the Department of Energy's Cleanup Technology Roadmap, Gaps and Bridges. The National Academies Press.

Pau, G.S.H., A.S. Almgren, J.B. Bell, and M.J. Lijewski (2009), A parallel second-order adaptive mesh algorithm for incompressible flow in porous media. Phil. Trans. R. Soc., A 367, 4633–4654.

Phifer, M.A., M.R. Millings and G.P. Flach (2006), Hydraulic Property Data Package for the E-Area and Z-Area Soils, Cementitious Materials, and Waste Zones. WSRC-STI-2006-00198, 344 p.

Sassen, D., S.S. Hubbard, N. Spycher, J. Wan, and M. Denham (2010), Utilizing geophysics to identify reactive facies and to spatially distribute reactive transport parameters. Fall AGU, H53K-08, invited.

Serne, R.J. and F.M. Mann (2004), Preliminary data from 216-B-26 borehole in BC cribs area. RPP-20303, Rev. 0. CH2M HILL Group, Richland, WA.

Smits, A.D., M.K. Harris, K.L. Hawkins and G.P. Flach (1997), Integrated Hydrogeological Model of the General Separations Area (U); Volume 1: Hydrogeologic Framework (U). WSRC-TR-96-0399, Rev. 0.

Spycher N**.**, S. Mukhopadhyay, D. Sassen, J. Wan, J. Seaman and M. Denham (2010), On Modeling pH Buffers and Rebound at an Acid-Contaminated Site. Soil Science Society of America Annual Meeting, Long Beach, CA, Oct 31–Nov 3.

SRNS (2010), Annual Corrective Action Report for the F-Area Hazardous Waste Management Facility, the H-Area Hazardous Waste Management Facility, and the Mixed Waste Management Facility (U). SRNS-RP-2010-00172, Savannah River Nuclear Solutions, Aiken, SC 29808.

Steefel, C.I. (2009), CrunchFlow, Software for Modeling Multicomponent Reactive Flow and Transport, User's Manual, 2009. http://www.csteefel.com/CrunchPublic/CrunchFlowIntroduction.html. Steefel, C.I., Carroll, S., Zhao, P., and Roberts, S. (2003), Cesium migration in Hanford sediment: A multi-site cation exchange model based on laboratory transport experiments. J. of Contaminant Hydrology, 67, 219-246.

Tong, C. (2007), The PSUADE Software Package Version 1.0, LLNL code release UCRLCODE-235523.

Tonkin, M., J. Doherty, and C. Moore (2007), Efficient nonlinear predictive error variance for highly parameterized models. Water Resour.Res., 43, W07429, doi:10.1029/2006WR005348.

Xu, T. E. Sonnethal, N. Spycher, and K. Pruess. (2004), TOUGHREACT User's Guide: A Simulation Program for Nonisothermal Multiphase Reactive Geochemical Transport in Variably Saturated Geologic Media. LBNL-55460.

APPENDIX.

Analytes and parameters included in the concentration database used for the Data Management Activities, which were collected at the F-Area between January 1, 1990 and September 15, 2009.

AIR TEMPERATURE	FLOW RATE
ALUMINUM	GROSS ALPHA
AMERICIUM-241	IODINE-129
AMERICIUM-243	IRON
BARIUM	MAGNESIUM
BERYLLIUM	NITRATE
BORON	NITRATE-NITRITE AS NITROGEN
CADMIUM	NITRITES
CALCIUM	pH
CARBON-14	SILICA
CESIUM-134	SODIUM
CESIUM-137	SPECIFIC CONDUCTANCE
CHLORIDE	SULFATE
CHLOROPRENE	TECHNETIUM-99
CHROMIUM	TOTAL ALKALINITY (AS CACO3)
COBALT	TOTAL ORGANIC CARBON
COBALT-60	TRITIUM
COPPER	TURBIDITY
CURIUM-242	URANIUM-233/234
CURIUM-243/244	URANIUM-235
CURIUM-245/246	URANIUM-238
DEPTH_TO_WATER	WATER TEMPERATURE

Subset of the above table showing parameters used in the Visualization Task.

DEPTH TO WATER	TRITIUM
NITRATE-NITRITE AS NITROGEN	URANIUM-233/234
pH	URANIUM-235
	URANIUM-238