

Performance Assessment Modernization at WIPP - 15077

Todd R. Zeitler*, Thomas B. Kirchner*, Glenn E. Hammond**, Heeho D. Park*

* Sandia National Laboratories, Carlsbad, NM 88220

** Sandia National Laboratories, Albuquerque, NM 87185

ABSTRACT

To aid in WIPP PA calculations, the US DOE has undertaken a broad modernization effort that includes: 1) updating the hardware and operating system used; 2) moving the current PA (performance assessment) codes to the new system; and 3) adding a database to store results essential for the analysis of assessments; 4) integrating run-control scripts with a web site to facilitate monitoring of performance assessment calculations; and 5) adding new capabilities via implementation of the PFLOTRAN code to replace multiple codes currently in use. The majority of codes were migrated from an Open VMS/Alpha processor system to a Unix/Intel processor system. Most codes were upgraded to use the dynamic allocation features of Fortran 95 and double precision variables. A few codes were migrated from Linux/Intel systems to the Unix/Intel systems and these code required nearly no changes to the source code. All of the migrated codes were requalified under Sandia's WIPP QA program. The implementation of some of the conceptual models using PFLOTRAN is being driven by the expectation that changes in the configuration of the repository will require a 3-dimensional representation of the mined areas. Symmetry in the current configuration of the repository made a 2-D representation adequate.

PFLOTRAN is an open-source, state-of-the-art massively parallel subsurface flow and reactive transport code. A version of this code is being developed specifically for WIPP PA in order to maintain the capabilities of the current flow and transport codes, as well as add the capabilities of a parallel code that will treat additional processes, including: 1) miscible multi-phase flow; and 2) heat transfer. PFLOTRAN will replace the two codes (i.e. BRAGFLO and NUTS) currently used for flow and transport calculations, simplifying the PA code structure and speeding up calculations by simultaneously solving for flow and transport with a single code. To ensure that PFLOTRAN reproduces the results of the current code suite, existing test cases have been used for verification purposes. Additional test cases have been developed to establish the new capabilities for WIPP PA. The result of the modernization effort will be a state-of-the-art subsurface flow and transport capability that will serve WIPP PA into the future.

INTRODUCTION

The WIPP, located in southeastern New Mexico, has been developed by the US DOE for the geologic (deep underground) disposal of TRU waste. Containment of TRU waste at the WIPP is regulated by the US EPA according to the regulations set forth in Title 40 of the Code of Federal Regulations (CFR), Part 191. The US DOE demonstrates compliance with the containment requirements according to the Certification Criteria in Title 40 CFR Part 194 by means of PA calculations. WIPP PA calculations estimate the probability and consequence of potential radionuclide releases from the repository to the accessible environment for a regulatory period of 10,000 years after facility closure.

To aid in WIPP PA calculations, the DOE has undertaken a broad modernization effort that includes updates to the hardware and software used to perform PA calculations used for compliance with EPA regulations [1]. The principal update to the hardware is a change from an Open VMS/Alpha processor platform to a modern UNIX (Solaris)/Intel processor platform (Figure 1). The Open VMS/Alpha system has been used for PA calculations since 1996 and has increasingly limited (and expensive) operating system and hardware support. In addition, the HP Alpha systems are no longer in production, thus making

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. This research is funded by WIPP programs administered by the Office of Environmental Management (EM) of the U.S Department of Energy.

expansion of the cluster nearly impossible. The UNIX/Intel system, on the other hand, has considerably lower costs for hardware and operating system support and is expandable with new hardware for increased future capabilities. The UNIX operating system Solaris is widely used in academia and industry, and thus training courses are readily available. The transition to a new platform necessitates the migration of PA computer codes and verification of their complete functionality on the new system.

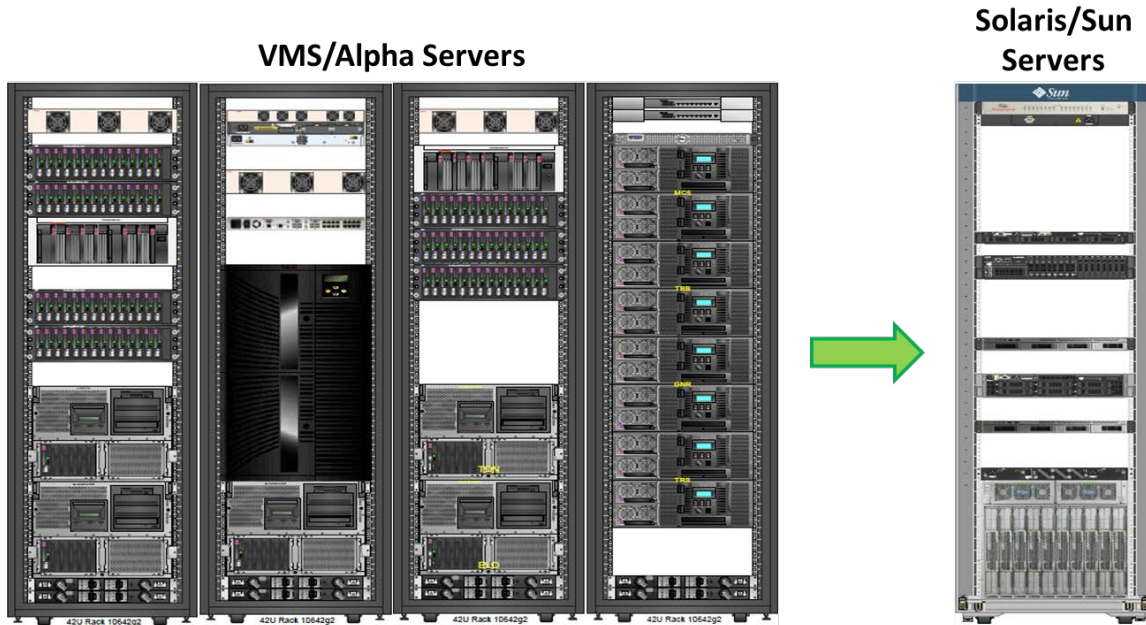


Figure 1: Hardware configuration for the VMS (left) and new Solaris (right) systems.

To date, PA calculations have relied on symmetry to justify the use of two-dimensional computational grids to describe the three-dimensional repository. Using these 2-D grids makes fluid and gas flow calculations highly efficient, which had been a necessity due to the relatively limited computational power of the PA computational cluster. The resulting system has an acceptable spatial resolution for compliance calculations. However, advances in computational power have sparked the creation of efficient 3-D subsurface flow codes that can take advantage of parallel computing platforms. PFLOTRAN is an open-source 3-D subsurface transport and flow code that will be used to replace the BRAGFLO (brine and gas flow) and NUTS (radionuclide transport) modules from the current PA calculation system [2]. The capabilities of PFLOTRAN will be expanded to include all BRAGFLO and NUTS capabilities used currently in PA calculations, as well as new capabilities such as miscible multi-phase flow and heat transfer. The use of the PFLOTRAN code represents a major change to PA software capabilities and is one that will allow simplification of the PA code structure (Figure 2), greater spatial resolution in the computational grids, as well as flexibility in possible repository reconfiguration.

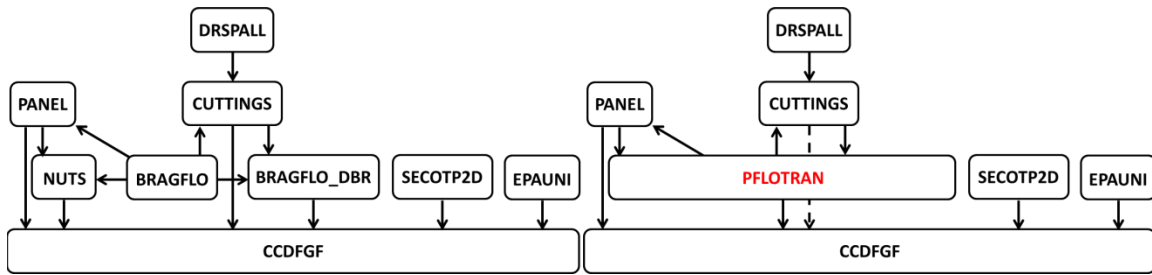


Figure 2: WIPP PA code structure: currently (left) and substituting PFLOTRAN for BRAGFLO and NUTS (right), (DRSPALL:spallings releases; CUTTINGS:cuttings and cavings releases; PANEL:radionuclide concentration calculation; NUTS:radionuclide transport in Salado; BRAGFLO:brine and gas flow in Salado; BRAGFLO_DBR:direct brine releases; SECOTP2D:radionuclide transport in Culebra; EPAUNI:radionuclide inventory; CCDFGF:summation of all release scenarios; PFLOTRAN:brine and gas flow and radionuclide transport in Salado).

DESCRIPTION/DISCUSSION

Software Migration

The migration of PA codes to the Solaris system was a relatively straightforward, although work-intensive, process. Most of the codes from the Open VMS/Alpha systems were upgraded to use the dynamic allocation features of Fortran 95 and double precision variables. Previously these codes were compiled using Fortran 77 conventions, with single precision variables and with dynamic allocation being implemented using a library of C functions. A few codes were also migrated from Linux/Intel processor systems and these codes required nearly no changes to the source code. The quality assurance (QA) process intrinsic to the PA work requires that each code be “requalified” following any change to the source code, or any significant change to the operating system or hardware system; in this migration process, all three changed, so requalification was necessary. Requalification requires recompiling each code and submitting each code to a series of test problems intended to confirm that the code implements requirements specified in its Requirements Document as quantified by criteria specified in its Verification and Validation Plan. In many cases a “regression test” can be performed, in which results from the current test are compared with the test results for the previous code version. If there are no significant differences, then the code passes the regression test.

In the current migration case, the floating point representation used in VMS on the Alpha processors was different than the IEEE standard employed on the Intel processors. In addition, some small differences were expected because single-precision variables in the VMS code versions have been changed to double-precision. Small differences due to the change in floating point representation can be magnified for codes in which there is a convergence on a solution. Differences are considered inconsequential if the relative percent differences between the VMS and Solaris outputs are less than 10^{-4} . If there are differences greater than this threshold, then the test results become part of a “validation test,” in which the results are compared to some results standard, previously determined based on results from another computational code or an analytical solution. If the new results are comparable to the expected values (usually within a few percent), then the code passes the validation test and is presumed to be working properly.

The code migration process is complete following documentation of all test case results, as well as updates to user’s manuals [3]. The process was lengthy (~12 months and involving many people), but now the codes are fully qualified to run on the Solaris system. As final tests, the complete calculations for the PABC-2009 and CRA-2014 calculations were run on Solaris and the results are nearly identical to

those from VMS (Figure 3). Additionally, the first new compliance calculation (for the impacts of an additional exhaust shaft) has been performed on the Solaris system [4].

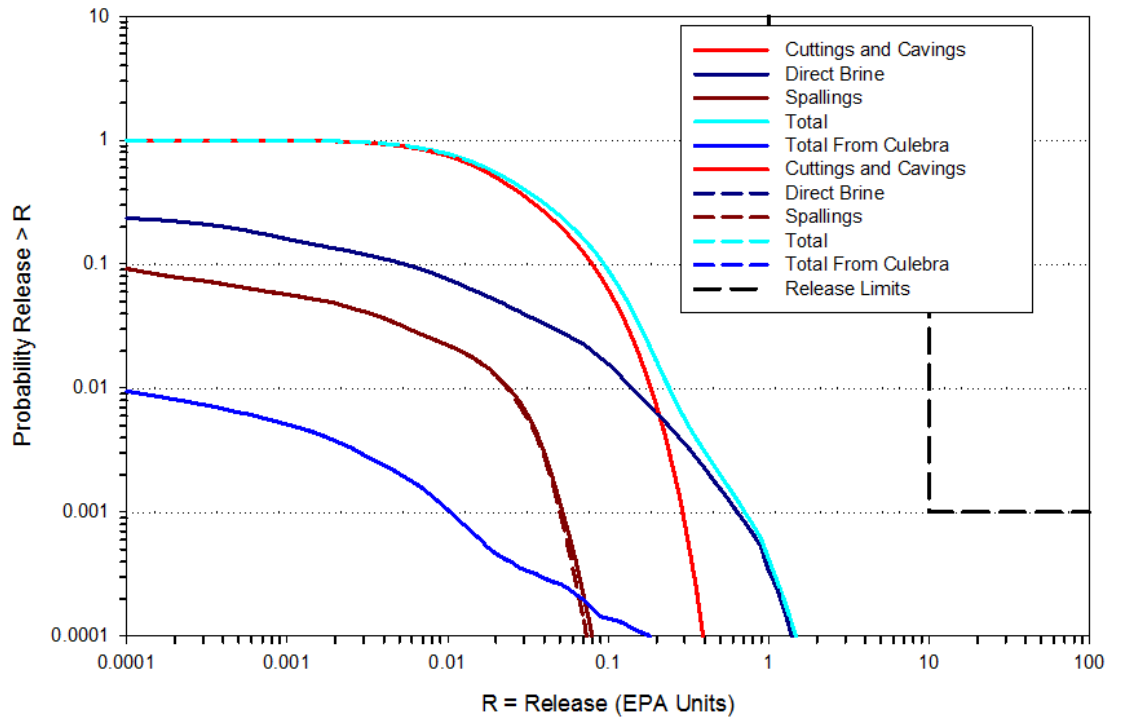


Figure 3: Comparison of releases for VMS codes (dashed lines) and Solaris codes (solid lines) for the PABC-2009 [3].

Improvements in Data Storage

A typical PA requires the generation of 300 CCDF curves for 37 outputs. The 300 CCDF curves express the uncertainty in important input parameters. The Latin Hypercube Sampling program is used to generate the 300 sets of parameters that give rise to the 300 CCDF curves. Each CCDF curve is derived from 10,000 simulations of possible futures where each future represents a unique combination of possible intrusion events, such as mining or drilling intrusions, coupled with other stochastic processes, such as the probability of encountering a pressurized brine pocket. The 10,000 values in a curve are reduced by binning the data to 161 values. Nevertheless, each analysis typically generates 1,787,100 records of output. These output records and the parameter values sampled by LHS were previously stored as text files but are now additionally inserted into database tables. Having these data in database tables greatly simplifies selecting and sorting the data for plotting and statistical analyses.

Improved Monitoring of Runs

The WIPP QA requirements specify that all PAs be run in a highly controlled, highly reproducible manner. All input files and codes must be stored in a secure repository and checked out for use in a given PA analysis and all important output files must also be stored in a secure repository. Concurrent Versioning System (CVS) repositories are used for this purpose. Scripts are used to control the execution of the codes in both VMS and Solaris. However, unlike the VMS scripts the scripts developed on Solaris implement a significantly greater degree of checking for potential errors prior to starting runs, such as

verifying that all specified input files exist, and also verifying that all expected output files and the directories in which they are to be placed are generated. Given that an analysis can generate more than 250,000 files, verifying the existence of input files and files generated by one code for use in another is important to ensuring that codes will not fail because of lack of inputs. The scripts also handle checking input files out of CVS repositories, output files into repositories, and ensuring that both input and output files are tagged as having been used in the analysis. The run control scripts also update a database as the analysis progresses. That database is polled by a web site every minute to display the progress of the various codes, identify codes that fail, and provide access to the log file generated by the scripts to help identify the cause of the failures. The web site can also be used to view information about all of the parameters, to view the QA documents for each code, and to view the source code.

New Software Development

Subsurface two-phase flow is currently calculated by BRAGFLO for PA calculations. The resulting flow fields are then passed to the radionuclide transport code NUTS. These codes are run using two 2-D spatial grids; one is a vertical north-south cross-section of the repository used for brine/gas flow and radionuclide transport and the other is a horizontal cross-section used for direct brine release scenarios, also calculated by BRAGFLO in a separate mode from the brine/gas flow calculations. Two-dimensional grids were implemented because of computational limitations at the time BRAGFLO was developed.

The open-source code PFLOTRAN has been used for subsurface flow and transport [2]. It is capable of using 3-D grids and parallel processing (Figure 4). It is currently being modified to support a WIPP-specific code that will preserve the current capabilities of BRAGFLO and NUTS while adding 3-D grids to expand the capabilities of PA calculations. PFLOTRAN also has two-phase miscible flow and heat transfer capabilities, neither of which exists in BRAGFLO. In order to add PFLOTRAN to the list of approved codes for WIPP PA calculations, it must go through the QA process of qualification.

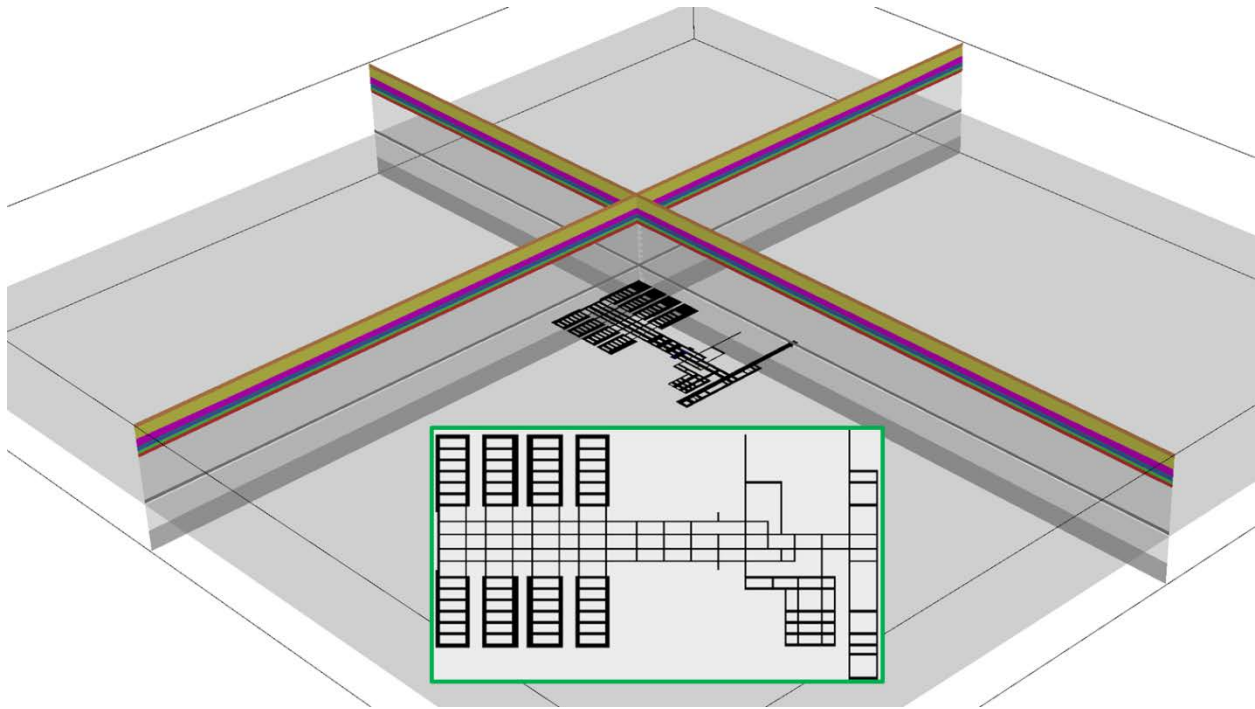


Figure 4: Representation of the 3-D grid describing the WIPP repository and surrounding area. Colored areas represent strata above the repository. The inset shows a 2-D representation of the waste panels and experimental and operations areas.

The qualification of PFLOTRAN necessitates it passing a series of test cases designed to test code requirements as laid out in a Requirements Document. At a minimum, the new tests must test all of the current capabilities of BRAGFLO and NUTS. The test cases already established for BRAGFLO and NUTS compare their results against outputs from other codes (e.g., TOUGH) and analytical solutions. Because PFLOTRAN can also be run in 1-D and 2-D modes, the BRAGFLO and NUTS test cases can be used directly to validate PFLOTRAN and a set of initial tests have been conducted early in the development phase to show that PFLOTRAN can mirror BRAGFLO (Figure 5) and NUTS (Figure 6) results.

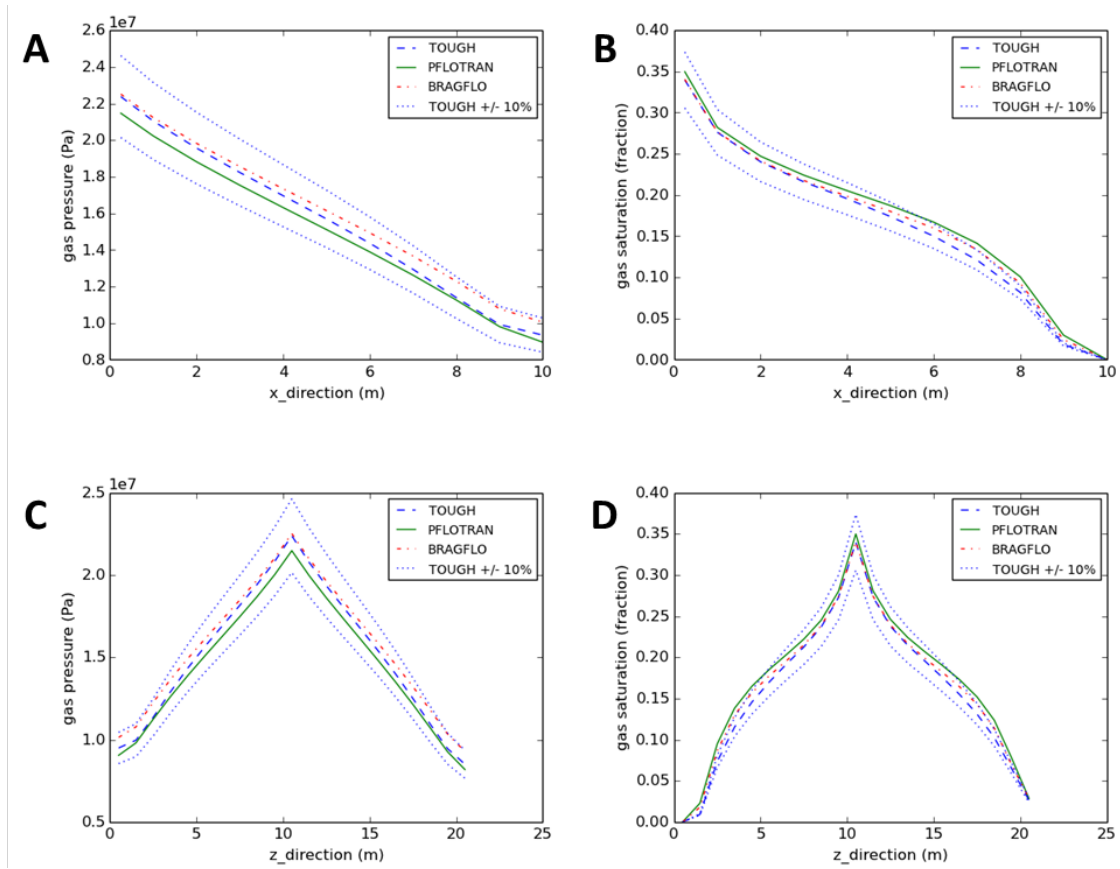


Figure 5: Preliminary comparison of PFLOTTRAN with BRAGFLO and TOUGH solutions for a test case that tests the ability to model gas injection and the subsequent migration of gas in a two-phase system.

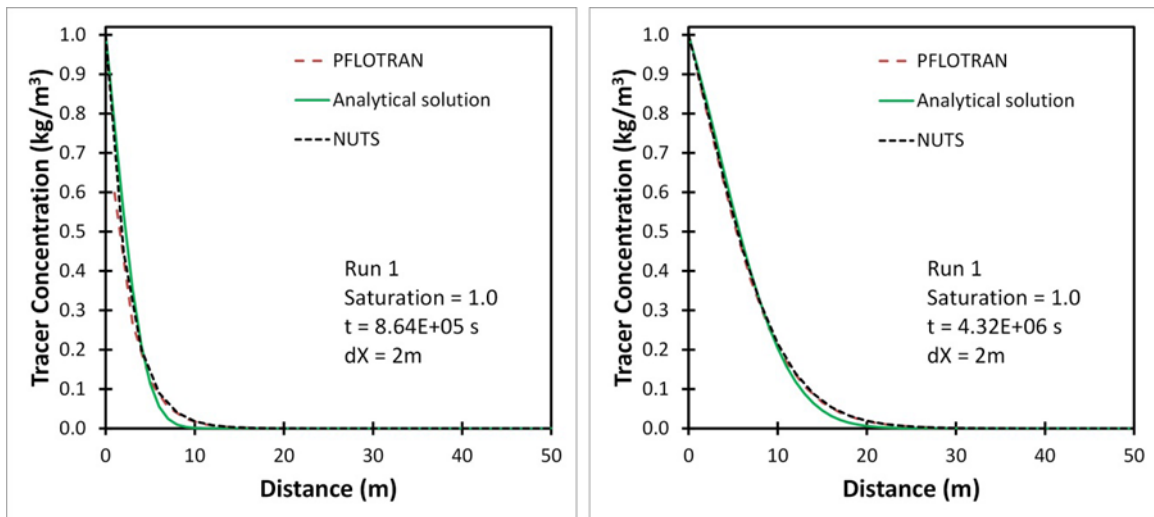


Figure 6: Preliminary comparison of PFLOTTRAN with NUTS and analytical solutions for a test case that verifies decay and dispersion in a constant velocity flow field and with constant porosity and saturation.

Any new capabilities that will be used in PA calculations will require new test cases. These new capabilities will obviously include using the 3-D mode of PFLOTTRAN, but will also include capabilities

that expand the potential for using the code to answer new questions. For example, BRAGFLO does not consider temperature effects, but PFLOTRAN has implemented heat transfer. The new code could thus be used to model the impacts of heat-generating waste if evaluating a future repository. New test cases will be developed to test these capabilities, either through comparison with other codes or analytical solutions where available. Although heat transfer is not an issue in the WIPP repository, an additional test will nonetheless be constructed to model heat transfer for a simple case, thus demonstrating that the heat transfer capability is properly implemented. A final critical step to the qualification of PFLOTRAN will be submission of the code to a peer review process, in which the conceptual models implemented in PFLOTRAN will be reviewed by a panel of experts.

CONCLUSIONS

The hardware and software platforms used for WIPP PA calculations are in the process of modernization. The new computational cluster runs on a Solaris operating system on an Intel processor platform capable of parallel processing. The migration of PA software codes from VMS to Solaris has been completed, required validation test cases successfully run, a complete PA run showing few or no differences in the projected release CCDF curves between VMS and Solaris, and code-specific QA documentation updated. Data are now stored directly into database tables to facilitate analyses and a new run control system helps to ensure that all expected input and output files exist and are properly stored in secure repositories. Finally, the implementation of PFLOTRAN as a new PA code to replace BRAGFLO and NUTS for flow and transport calculations is currently underway. When finished, this implementation will allow for 3-D grid functionality, as well as other added capabilities.

REFERENCES

1. Kirchner, T. AP-162, Revision 0: Analysis Plan for Migration of the Performance Assessment Codes to the Sun Solaris Blade Server Running with Intel Processors. Sandia National Laboratories, Carlsbad, NM. ERMS #557765 (2012).
2. Hammond, G. E.; P. C. Lichtner, and R. T. Mills. Evaluating the Performance of Parallel Subsurface Simulators: An Illustrative Example with PFLOTRAN. *Water Resources Research*, 50 (1), 208-228 (2014).
3. Kirchner, T., A. Gilkey, and J. Long. Summary Report on the Migration of the WIPP PA Codes from VMS to Solaris, AP-162, Revision 1. Sandia National Laboratories, Carlsbad, NM. ERMS #561757 (2014).
4. Camphouse, C. Analysis Plan for the Impact Assessment of an Additional WIPP Shaft. Sandia National Laboratories, Carlsbad, NM. ERMS #562714 (2014).

ACKNOWLEDGEMENTS

The authors wish to thank Sungtae Kim for providing Figure 5.