

# MINIMIZING THE IMPACT OF SOFTWARE EVOLUTION ON RADIOACTIVE WASTE MANAGEMENT

Glenn E. Hammond and Jennifer M. Frederick

*Sandia National Laboratories, PO Box 5800, MS 0747, Albuquerque, NM, 87185-747, gehammo@sandia.gov*

*This paper discusses software design choices made during the ongoing development of PFLOTRAN with the intent of improving the longevity and flexibility of the code while maintaining robustness. PFLOTRAN is a massively parallel reactive multiphase flow and transport code employed to simulate subsurface processes in support of radioactive waste management licensing and performance assessment studies. PFLOTRAN's open source licensing, software configuration management, object-oriented design through modern Fortran 2003/2008, approach to coupling scientific process models, and automated testing are discussed in detail.*

## I. INTRODUCTION

Software evolution poses a major challenge to simulation tools employed in support of radioactive waste management. With advances in understanding through fundamental scientific research comes the development of increasingly mechanistic process models and a growing demand to incorporate such process models within subsurface simulators. Modern numerical methods and programming paradigms may facilitate the implementation and solution of these new process models, but there is likely an arduous union to be made with some measure of legacy code. To make matters worse, the commodity computing hardware on which these simulators are executed has a relatively limited shelf life, while the life of a high-performance supercomputer is even shorter at three to five years. However, software evolution has a positive side. With robust software engineering and a plan for long-term maintenance, a simulator can evolve over time incorporating and leveraging advances in the computational and domain sciences. In this positive light, what practices in software engineering and code maintenance can be employed within radioactive waste management performance assessment (PA) modeling to maximize the positive aspects of software evolution while minimizing its negative side effects?

This work presents measures taken in the development of PFLOTRAN to better ensure sustainable software development within the context of software evolution and PA for radioactive waste management. PFLOTRAN<sup>1,2</sup> is an open source, massively-parallel code executed within

Sandia National Laboratories' Generic Disposal System Analysis (GDSA) framework<sup>3</sup> to simulate subsurface reactive multiphase flow and transport processes in the geologic subsurface in support of PA. By adhering to the code development practices outlined below, it is hopeful that PFLOTRAN will remain a robust and viable simulation tool for years into the future.

## II. SUSTAINABLE PFLOTRAN DEVELOPMENT

PFLOTRAN developers have attempted to keep the code flexible, extensible, and robust through deliberate choices in software design that should allow PFLOTRAN to evolve over time. These design choices include open source licensing, adherence to software configuration management, object oriented design, modular process model coupling, and comprehensive automated testing. The implementation of each of these design choices is discussed in the following sections.

### II.A. Open Source Development

Open source software development has gained favor over the past several decades. In short, open source software refers to code that is publicly available and can be modified and shared by anyone. Therefore, any developer may copy the code, make modifications and offer up those modifications either by merging the updates with the main source code repository or by distributing the code separately. Whether code modifications are merged with the original source code repository is up to the owners of that repository, but regardless, all code may be shared freely, and should be shared in the true spirit of open source licensing.

There are many benefits of open source development that go beyond the availability of source code and the ability to modify it. First, it encourages collaboration. In many cases, a group of independent software developers self-select and learn to work together to develop, debug and test software with little to no duplication of effort. Ideally these developers have diversified expertise and provide internal peer review. This self-selection process promotes increasingly compatible developer relationships and teaming environments. Open source provides transparency that exposes complete implementation detail critical to

scientific reproducibility, a transparency that is necessary for software quality assurance (QA) within radioactive waste management. This level of detail is most often excluded from publication in journals and formal reports. The developer can interrogate code at the finest or deepest level to verify actual implementation. Without the availability of source code, software becomes a black box and the user is dependent solely upon documentation which may be less thorough than desired. This transparency is important for resolving discrepancies between codes. It is the authors' experience that one of the largest impediments to inter-code comparisons are subtle details, such as undocumented averaging schemes employed to represent state variables (e.g. fluid density at cell interfaces) or customized conditional expressions designed to improve solver convergence. These details may be delineated in the documentation, but one must often interrogate source code to clarify the actual implementation.

Open source development promotes more optimal use of funding. Funds can be pooled across a diverse set of projects and budgets, and between funding agencies that have no ability or intention to collaborate. Funding that would have been spent on licensing fees can be redirected to development, and the funding entity may choose the exact capability to be implemented and the developer to complete the task, whereas with proprietary codes, the users may be relegated to the mercy and timeline of a few over-committed developers. For those who have limited funding (e.g. graduate students, startup businesses, third-world researchers), open source codes open doors to capability that would not be afforded otherwise.

The community that grows around an open source code can drive the code to evolve beyond the original intent of the code. This evolutionary process is evident in PFLOTRAN's history where process models outside subsurface flow and reactive transport (the original purpose for development) have been developed and coupled to PFLOTRAN (e.g. surface water flow, geomechanics, and geophysics). It should be noted that proprietary codes can also benefit from this same community feedback. However, a member of an open source community has the option to depart and retain the source code, using it as they so desire.

Finally, open source development promotes natural selection. The ability to self-select as a developer and user encourages adaptation to community needs or the size of the community declines over time. Ideally, the most fit codes survive while those that cannot remain viable, or refuse to adapt to change, disappear.

While funded by the US Department of Energy (DOE) SciDAC-II program, PFLOTRAN was licensed under the GNU Lesser General Public License (LGPL). This LGPL license allows for third parties to develop proprietary software around PFLOTRAN (e.g. linking to PFLOTRAN as a library or through a user interface), while any modifications to the original source code itself must be

documented and remain open source. The complete PFLOTRAN code base is freely accessible at <https://bitbucket.org/pfлотran/pfлотran-dev>.

## II.B. Software Configuration Management

Software configuration management (SCM) records all changes to software over its lifetime. The current PFLOTRAN code base is managed with Mercurial<sup>4</sup>, a distributed version control tool that stores code revisions in changesets. Distributed version control enables software developers to clone the original repository and work independently, merging changes with the base repository or other clones as needed. Without configuration management and version control, code modifications must occur sequentially (e.g. only one developer modifying a file at a time) or simultaneous edits must be merged by hand. Mercurial enables developers to modify the same file and merge these edits automatically. Should developers modify the same lines in a source file and attempt to merge the files, Mercurial marks the lines, including both versions from both changesets and forces the developer merging the edits to perform a manual merge, rectifying any discrepancies. Mercurial also allows developers to revert back to any changeset recorded in the past.

Anyone with access to the internet and Mercurial installed on their computer may add capability to PFLOTRAN by (1) forking the main repository on Bitbucket, (2) cloning the forked repository locally, (3) modifying the local source code, (4) ensuring that all tests pass, (5) committing the successful changes as a new changeset and pushing the changeset back to the forked repository, and (6) submitting a pull request to the administrators of the main repository.

In the context of radioactive waste management where quality assurance is paramount, software must often be QA'ed for licensing or performance assessment purposes. Software configuration management greatly facilitates this process by providing comprehensive control on source code recordkeeping. All code modifications are tracked. As additional capability is added to a code, the SCM documents the updates automatically. Specific versions of the code may be frozen. If using Mercurial, one may branch the code repository, splitting a frozen branch of the source code from the original default branch and preventing any new changes to the frozen branch. The frozen branch may exist forever within the repository, while the default branch continues to be developed. To a certain degree, bug fixes performed within one branch may be merged with the other, though this becomes more difficult over time as the two branches diverge. Furthermore, given identical code compilers, operating systems and computer hardware, one may always revert to the frozen branch, build an identical executable, and produce the same result.

## II.C. Modular Object Oriented Design

For the benefit of ensuring longevity of software developed for radioactive waste management, it is important that these tools be designed and coded in a manner that is modular and extensible. Modern programming paradigms help to facilitate modularity and extensibility and overall longevity of the software product. One common feature among most modern programming paradigms is object oriented design, where an object possesses the data and procedures needed to provide functionality and interfaces are set up for interaction between objects.

PFLOTTRAN is designed as a hierarchy of nested objects from the simulation object at the highest level to the low level, cell-centric auxiliary variable objects that locally hold all state variables for each grid cell. Consider the low level material object within PFLOTTRAN, which is composed of data describing material properties such as porosity, permeability, rock particle density, etc. at each grid cell. The material object contains procedures to initialize and update these properties over time and deliver the properties to other objects (e.g. process models) or operations (e.g. flux calculations). The material object also possesses pointers to sub-objects such as the characteristics curves object that encapsulates properties associated with variably saturated flow of a fluid within the porous media (e.g. residual saturations, air entry pressures). The material object itself is owned by a higher level object known as the realization object, which contains all the parameters (including all material objects) necessary to solve the governing mass and energy conservation equations associated with a subsurface problem.

The benefit of oriented design is that when programmed correctly, modifications to an object's data and procedures will have little to no impact on other portions of the code. The following Fortran example illustrates this concept. Suppose there is procedure that calculates a constitutive relation or property as a function of permeability, porosity, and rock particle density, all of which are contain in a material object. Without a material object, these parameters must be passed into the procedure separately:

```
call ConstitutiveRelation(permeability, &
    porosity, soil_particle_density)
```

If additional parameters such as tortuosity and soil heat capacity are added to the constitutive relation calculation, the argument list must be expanded:

```
call ConstitutiveRelation(permeability, &
    porosity, soil_particle_density, &
    tortuosity, heat_capacity),
```

which can be laborious if argument lists are repeatedly changing for large numbers of procedures. When these properties are contained within a material object:

```
type :: material
    real*8 :: permeability(:)
    real*8 :: porosity
    real*8 :: soil_particle_density
    real*8 :: tortuosity
    real*8 :: heat_capacity
end type material,
```

no changes to the argument list are necessary:

```
call ConstitutiveRelation(material).
```

With the object oriented paradigm, one is free to alter the contents of the material object without impacting all the calls to procedures utilizing that material object.

Object oriented paradigms combined with modern programming languages can further increase the degree of encapsulation and extensibility. PFLOTTRAN employs Fortran 2003/2008 which supports classes and pointers to procedures. With pointers to procedures, a generic constitutive relation call can be redirected to any number of subroutines as long as they have identical argument lists. In other words, during setup, the constitutive relation procedure pointer is directed towards one of many available procedures:

```
if (one_way) then
    ConstitutiveRelation => &
        ConstitutiveRelationOneWay
else if (the_other) then
    ConstitutiveRelation => &
        ConstitutiveRelationTheOther
endif.
```

This approach provides flexibility in algorithmic design as no single, hardwired constitutive relation must be used. The user has the option to choose from any number of constitutive relations and the code developer can support all options without cluttering code with excessive conditions and/or select case statements strewn throughout the code. For example, the select case statement:

```
select case(i)
    case(one_way)
        call ConstitutiveRelationOneWay(material)
    case(the_other)
        call ConstitutiveRelationTheOther(material)
end select
```

that must be repeated every time the constitutive relation is called simplifies to

```
call ConstitutiveRelation(material).
```

For supporting large numbers of equations of state within PFLOTRAN that calculate fluid density as a function of pressure and temperature, pointers to procedures have greatly simplified the code.

Fortran classes provide additional flexibility. They are extensible, meaning one may extend a class by adding additional member variables to the class and altering class procedures to consider these new member variables. In Fortran, the constitutive relation example can be recast as a member procedure within a class. For instance, the material object can possess a procedure that evaluates the constitutive relation as a function of the material object's member variables.

```
type :: material
...
contains
  procedure :: ConstitutiveRelation
end type material.
```

One simply calls the member procedure and the compatible procedure pointer is chosen with the material object automatically passed in as the first argument:

```
call material%ConstitutiveRelation().
```

The constitutive relation procedure is set on the fly. The class essentially accomplishes the function of an object/pointer to procedure combination while providing greater flexibility, as the class object is more abstract, allowing for a much broader virtual argument list for the procedure.

Another example of where object oriented design has been valuable within PFLOTRAN is in the concept of a process sandbox. PFLOTRAN currently supports sandboxes for source/sinks and kinetically-formulated chemical reactions. Consider the chemical reaction sandbox. Many PFLOTRAN users have asked for a means of implementing custom kinetic rate expressions for chemistry. The reaction sandbox fulfills this purpose by isolating PFLOTRAN's chemistry and providing a simplified reaction framework within which the researcher may quickly implement a kinetic reaction without completely learning or understanding PFLOTRAN's reaction process model. The reaction sandbox also serves as a tool for testing kinetic reactions prior to acceptance and integration within the code.

To implement the reaction sandbox, a base reaction class is defined as essentially an empty shell with empty member procedures for reading, initializing, reacting, and destroying the object. A developer desiring to add a new kinetic reaction within PFLOTRAN can use the reaction sandbox to prototype a new reaction by extending the base reaction class, adding custom parameters within the extended class and chemistry algorithms within the extended procedures. For each call to the chemistry

process model at a grid cell, PFLOTRAN loops over a linked list of reaction sandboxes, evaluating the custom reactions. Since all code modifications are isolated to a single sandbox module, the developer may modify the respective sandbox reaction without understanding the full implementation of PFLOTRAN chemistry and without adversely affecting other portions of the chemistry code.

Discussion of programming paradigms and languages may seem beyond the scope of radioactive waste management. However, by considering such paradigms early in the code development process, the software tools that support radioactive waste management should be more flexible over a longer lifetime, reducing the amount of laborious refactoring of code required to keep pace with advances emanating from domain and computational science research.

## II.D. Process Model Coupling

A major challenge with code development is anticipating process models that will be needed in the future. Currently, most subsurface simulators developed in support of radioactive waste management consider process models for flow and energy transport in multiple fluid phases along with reactive transport. However, with little knowledge of new or more advanced process models that may emerge in the future, how does one consider these processes from an algorithmic standpoint? In other words, how does a code developer introduce a new process model, coupling it to existing process models without refactoring large portions of code in the integration process? In 2013, PFLOTRAN underwent a high-level overhaul where the single time stepping loop for flow and transport was removed in favor of separate time stepping loops for each process model. With separate, somewhat autonomous time stepping, process models were then coupled through a tree-structured linked list of process models.

Consider the process model coupler (PMC) class shown in Fig. 1.

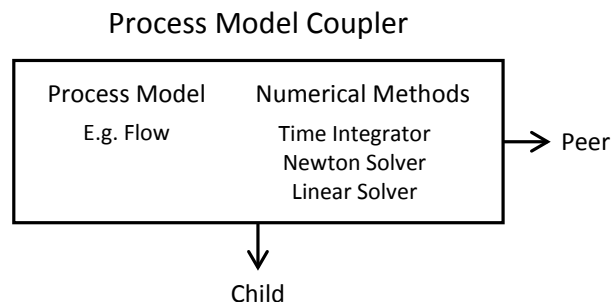


Fig. 1. PFLOTRAN process model coupler (PMC).

A PMC is composed of a process model (in this case flow), the data and numerical methods required to solve the process model, and pointers to two neighboring PMCs, a peer and a child. The PMC itself possesses all the information necessary to simulate a process, stepping forward in time until a synchronization point has been met.

The numerical methods employed within a PMC (i.e. time integrator for stepping, nonlinear equation solver, linear equation solver, etc.) could be considered identical to those employed within traditional subsurface simulators. A key difference is the linkage between the PMCs. Fig. 2 illustrates a somewhat complex linkage between six hypothetical process models (A, B, C, M, Y, Z). Here, PMC A is the master process model. As PMC A proceeds forward, taking a single time step, its child PMC B takes as many time steps as necessary to keep up with PMC A,




Fig. 2. Example linked list (coupling tree) of process model couplers.

whether that be lock step or sub-stepping (i.e. at multiple smaller time steps), but PMC B never passes PMC A in time. Similarly, for each PMC B time step, its child PMC C takes as many time steps as necessary to keep up with PMC B. A child PMC always plays catchup with the parent. Peer PMCs, on the other hand, time step to synchronization points. In this case, PMCs B, M, and Y are all peers; they are all children of PMC A. Their synchronization point is the PMC A time step, or the simulation time at the end of each PMC A time step. PMC B and its child PMC C step in time until they reach the synchronization point (PMC A's time), and then PMC M time steps as necessary to reach the synchronization point, and finally PMC Y (with child PMC Z following) time steps to the synchronization point, in that order. PMC A takes another step and the process repeats.

The benefit of using this tree of PMCs is that it provides great flexibility as long as process models can be loosely coupled (i.e. decoupled) between PMCs. If two process models are to be solved simultaneously in a fully-coupled manner (e.g. solute transport and chemical reaction), they must be solved within a single PMC. Also note that a PMC does not necessarily have to simulate a process at all. For instance, PFLOTRAN provides a PMC for salinity update that calculates the mass fraction of salt in the liquid phase when flow and transport process models

are employed to model nuclear waste disposal in saline aquifers (e.g. salt formations, brines in deep boreholes). The full tree of PMCs for this selection of process models is shown in Figure 3.




Fig 3. Coupling tree for multiphase flow, reactive transport, and salinity update process model couplers.

Based on this tree, the salinity update is a peer to the reactive transport and updates the salt mass fraction after the reactive transport PMC has caught up with the flow step and immediately prior to the next flow step. There is no solution of a nonlinear system of equations or time stepping involved. It is solely an update to the state variable storing salt mass fraction.

For radioactive waste management, the PMC tree has facilitated the incorporation of multiple waste form and waste package process model classes and the isotope solubility, partitioning, decay and ingrowth (SPDI) capability in PFLOTRAN. Both of these PMCs can be a child of PFLOTRAN's reactive transport process model. When both are employed, the waste form/package PMCs are the child of the reactive transport process model while the isotope SPDI PMC is a peer of the waste form/package PMC as shown in Figure 4. The use of process model couplers within PFLOTRAN has significantly eased the level of effort required to implement new scientific processes that support radioactive waste management within the code.




Figure 4. Coupling for multiphase flow, solute transport, waste form/package and isotope solubility, partitioning, decay and ingrowth (SPDI) process model couplers.

## II.E. Automated Testing

Testing plays an important role in establishing the credibility and reliability of a code, especially in the closely regulated licensing and performance assessment

environment common to radioactive waste management. As open source licensing, software configuration management, and a flexible, extensible object oriented coding paradigm can foster an efficient, distributed code development environment, programmers from around the world may rapidly add, modify and merge new capability to a code base. This potential for rapid progress engenders the need for better testing, especially for software that supports radioactive waste management where QA is of significant importance.

PFLOTRAN employs automated unit and regression testing to ensure that changes merged into the base PFLOTRAN repository do not alter known solutions of existing process models. From a QA standpoint, this testing capability is vital. However, it should be emphasized that these test suites are only as good as their code coverage. In other words, if no tests exist for an encoded process model, it is impossible to determine if an error has been introduced into the process model by recent changes merged by configuration control.

### *II.E.1. Unit Testing*

Unit tests are designed to test a small subset of the code, usually contained within a single subroutine. PFLOTRAN has unit tests primary for constitutive relations such as fluid equations of state and characteristic curves (e.g. saturation and relative permeability functions). These tests are implemented through the pFUnit<sup>5</sup> testing framework developed by NASA. Within this framework, the developer adds code to call subroutines with a specific set of arguments (e.g. parameters to be fed into the constitutive relation). The result returned by the subroutine is then compared to a known solution and an error is reported if the difference between the calculated and known solution is outside an allowed tolerance. These tests are conducted in isolation. In other words, a PFLOTRAN executable is not required to complete a unit test; just the subroutine being executed as it would be called from a library.

### *II.E.2. Regression Testing*

Regression tests execute an entire simulation, comparing sampled simulation results against a gold standard. PFLOTRAN's regression testing suite is written in Python by PFLOTRAN developers. The ability to sample and export simulation results at specific cells and/or at every Nth cell was developed in a regression module written in Fortran. The code writes a regression (.regression) file for every simulation executed in the suite. The regression testing script recursively traverses the suite's directory structure reading in custom test configurations and tolerances from configuration files specific to each directory. The configuration files and the gold files (.regression.gold), against which the .regression

results are compared, are stored in the code repository. There are currently more than 200 regression tests in the PFLOTRAN repository that complete in approximately three minutes.

Whenever the PFLOTRAN code base is modified, it is expected that the code developer will execute all unit and regression tests in the repository. These can be executed from the command line by typing 'make test'. If any tests fail, the developer is responsible for pinpointing the cause for the change in solution and determining whether the new solution is more accurate or an error. Variability in solution between operating systems and compilers can be problematic when comparing to a gold standard. For this reason, tolerances are employed to allow for slight differences in solution. However, the default tolerance is very tight at  $1.e-12$  (absolute difference).

### *II.E.3. Buildbot*

The PFLOTRAN developers employ Buildbot<sup>6</sup> for continuous integration (automated building and testing). Any modifications to the main Bitbucket code repository cause the Buildbot master to spawn jobs on workers that download, configure and build all the source code and libraries needed to compile a PFLOTRAN executable. Once the executable is successfully built, the workers run the unit and regression test suites, reporting back to the master the results of the build and test. These results are updated on a public website and the developers are notified of the results by email. The benefit of utilizing automated building and testing through Buildbot is that code covered by unit and regression tests cannot be adversely changed without the developer intentionally updating the gold files. The developers can carefully track the state of PFLOTRAN.

### *II.E.4. Verification and Validation Testing for QA*

Verification and validation (V&V) are important for establishing correctness of simulation results. Verification usually involves a comparison with an established simulated or analytical solution, while validation checks the correctness of a simulated solution against experimental results or data. V&V testing for PFLOTRAN is currently under development as part of PFLOTRAN's QA test suite. Currently, individual process models for multiphase flow such as liquid and gas phase fluid flow, and thermal conduction are covered by more than 50 tests within the suite. The QA test suite is not launched through PFLOTRAN's makefile, as is the case for unit and regression tests. Hence, it is currently not automated through Buildbot. However, anyone can run the QA test suite while building PFLOTRAN documentation. There are currently plans to enlarge the coverage of the QA test suite to include verification of individual process models against analytical solutions for solute transport and

biogeochemical reaction. In addition, PFLOTRAN simulations composed of multiple nonlinear process models, for which analytical solutions do not exist, will be verified against other simulators.

### III. CONCLUSION

Although software tends to have a limited lifespan, careful design and planning in the development of a code can significantly lengthen the duration of an application's viable existence. A subsurface simulator's ability to evolve with new paradigms in hardware and software engineering or to accommodate new scientific processes relies heavily upon choices made early in its development. Within the context of software that supports radioactive waste management (licensing, performance assessment, etc.), this ability for codes to adapt and evolve is important from both the fiscal and scientific investment viewpoint. Funding is limited, and researchers who tend to cling to legacy codes also desire the incorporation of new capability (e.g. process models and constitutive relations). Though no software product is perfect, this work summarizes design choices and methodologies employed in the development of the subsurface reactive multiphase flow and transport code PFLOTRAN in an attempt to facilitate software evolution.

PFLOTRAN is licensed under an open source (GNU LGPL) license which encourages community involvement ranging from requests for capability to development and debugging. Software configuration control for PFLOTRAN maintained with Mercurial, a distributed source control management tool, and the entire code base is freely available through a Bitbucket repository with support provided through the website [www.pflotran.org](http://www.pflotran.org).

The developers of PFLOTRAN have attempted to leveraged object oriented design in an effort to keep the code flexible, extensible, and maintainable through the use of Fortran classes and procedure pointers, a capability enabled through the adoption of Fortran 2003/2008. Fortran classes have enabled the incorporation of a cascading or nested, tree-style linked list paradigm for coupling PFLOTRAN's process models, and this approach greatly facilitates the addition of new process models and alternate implementations of existing process models (i.e. the developer can swap process models on the fly).

Finally, PFLOTRAN's robustness has been strengthened over the past several years through automated unit and regression testing. All modifications checked into the main PFLOTRAN Bitbucket repository are tested through the Buildbot continuous integration framework with errors immediately reported to the developers. These design choices were made with the intent of lengthening PFLOTRAN's life, making it a more viable tool for simulating subsurface processes well into the future.

### ACKNOWLEDGMENTS

Research presented in this manuscript was funded by the U.S. Department of Energy Office of Environmental Management (Waste Isolation Pilot Plant), Nuclear Energy (Spent Fuel and Waste Disposition Program) and Science (Biological and Environmental Research – Subsurface Biogeochemical Research) Programs.

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 paper is Sandia publication: SAND2016-12458 C.

### REFERENCES

1. G.E. HAMMOND, P.C. LICHTNER and R.T. MILLS "Evaluating the Performance of Parallel Subsurface Simulators: An Illustrative Example with PFLOTRAN", *Water Resources Research*, **50**, doi:10.1002/2012WR013483 (2014).
2. P.C. LICHTNER, G.E. HAMMOND, C. LU, S. KARRA, G. BISHT, B. ANDRE, R.T. MILLS, J. KUMAR, and J.M. FREDERICK, "PFLOTRAN Webpage and User Manual", <http://www.pflotran.org> (2016).
3. P.E. MARINER, E.R. STEIN, J.M. FREDERICK, S.D. SEVOUGIAN, G.E. HAMMOND, and D.G. FASCITELLI, "Advances in Geologic Disposal System Modeling and Application to Crystalline Rock", FCRD-UFD-2016-000440, SAND2016-9610R, Sandia National Laboratories, Albuquerque, NM (2016).
4. <https://www.mercurial-scm.org>.
5. <http://pfunit.sourceforge.net>.
6. <http://buildbot.net>.