

# ECP SOLLVE Validation and Verification OpenMP Offloading Testsuite

Thomas Huber<sup>1</sup>, Swaroop Pophale<sup>3</sup>, Nolan Baker<sup>1</sup>, Nikhil Rao<sup>1</sup>, Michael Carr<sup>1</sup>, Jaydon Reap<sup>1</sup>, Kristina Holsapple<sup>1</sup>, Joshua Hoke Davis<sup>4</sup>, Tobias Burnus<sup>5</sup>, Seyong Lee<sup>3</sup>, David E. Bernholdt<sup>3</sup>, Sunita Chandrasekaran<sup>1,2</sup>  
<sup>1</sup>University of Delaware, <sup>2</sup> Brookhaven National Laboratory, <sup>3</sup>Oak Ridge National Laboratory, <sup>4</sup>University of Maryland, <sup>5</sup>Siemens Electronic Design Automation

## ABSTRACT

Scaling OpenMP with LLVM for Exascale performance and portability (SOLLVE) aims to scale OpenMP by leveraging LLVM for exascale performance and portability of applications. A Verification & Validation (V&V) testsuite tests various OpenMP directives to evaluate system & vendor compliance. The V&V suite is ran on various systems, including NERSC's Perlmutter system & Oak Ridge National Laboratory's Summit & Crusher systems.

## BACKGROUND INFO

OpenMP, a parallel-based programming model, allows for performance optimization in C, C++ & Fortran, with its features called "directives" listed in the OpenMP specification (spec). The purposes of the SOLLVE Verification and Validation (V&V) testsuite are to:

- Evaluate compiler's compliance with the specification
- Identify ambiguities in the specification
- Illustrate a system's ability to run OpenMP directives & utilize offloading parallel directives on GPUs
- Demonstrate the use & purpose of new OpenMP directives to application developers

OpenMP is useful for many application developers working on HPC systems to ensure their code is running at maximum efficiency. Our testsuite ensures that the OpenMP specification, compiler vendors & system operators are implementing OpenMP effectively.

## APPROACH

- Tests are usually written to ensure the test would fail if the directive were to not work properly.
- If the test **fails**, analysis is done to determine the issue:
  - The test was written improperly
  - The directive has not yet been implemented
  - The specification is too ambiguous
- If the test **passes**, it is merged into the repository.
- Results for each system & compiler are then uploaded to the website.

## TEST EXAMPLE

```
#pragma omp parallel num_threads(threads)
while(1){
    int tot;
    #pragma omp atomic read
    tot = total;
    if (tot <= 0)
        break;
    #pragma omp masked
    {
        OMPVV_TEST_AND_SET_VERBOSE(errors, omp_get_thread_num() !=
0); // primary thread
        ct++;
        #pragma omp atomic
        total = total-1;
    }
    OMPVV_TEST_AND_SET_VERBOSE(errors, ct != 10);
}
```

Figure 1: test\_masked.c code segment

This test, which focuses on the **omp masked** directive which only runs on the master thread, demonstrates:

- OMPVV\_TEST\_AND\_SET\_VERBOSE function, which is part of the V&V suite and is a primary tool in reporting errors for test results
- Errors reported if the running thread is not the master (first) thread
- Errors reported if masked code segment does not run 10 times

## RESULTS ON PERLMUTTER, SUMMIT & CRUSHER

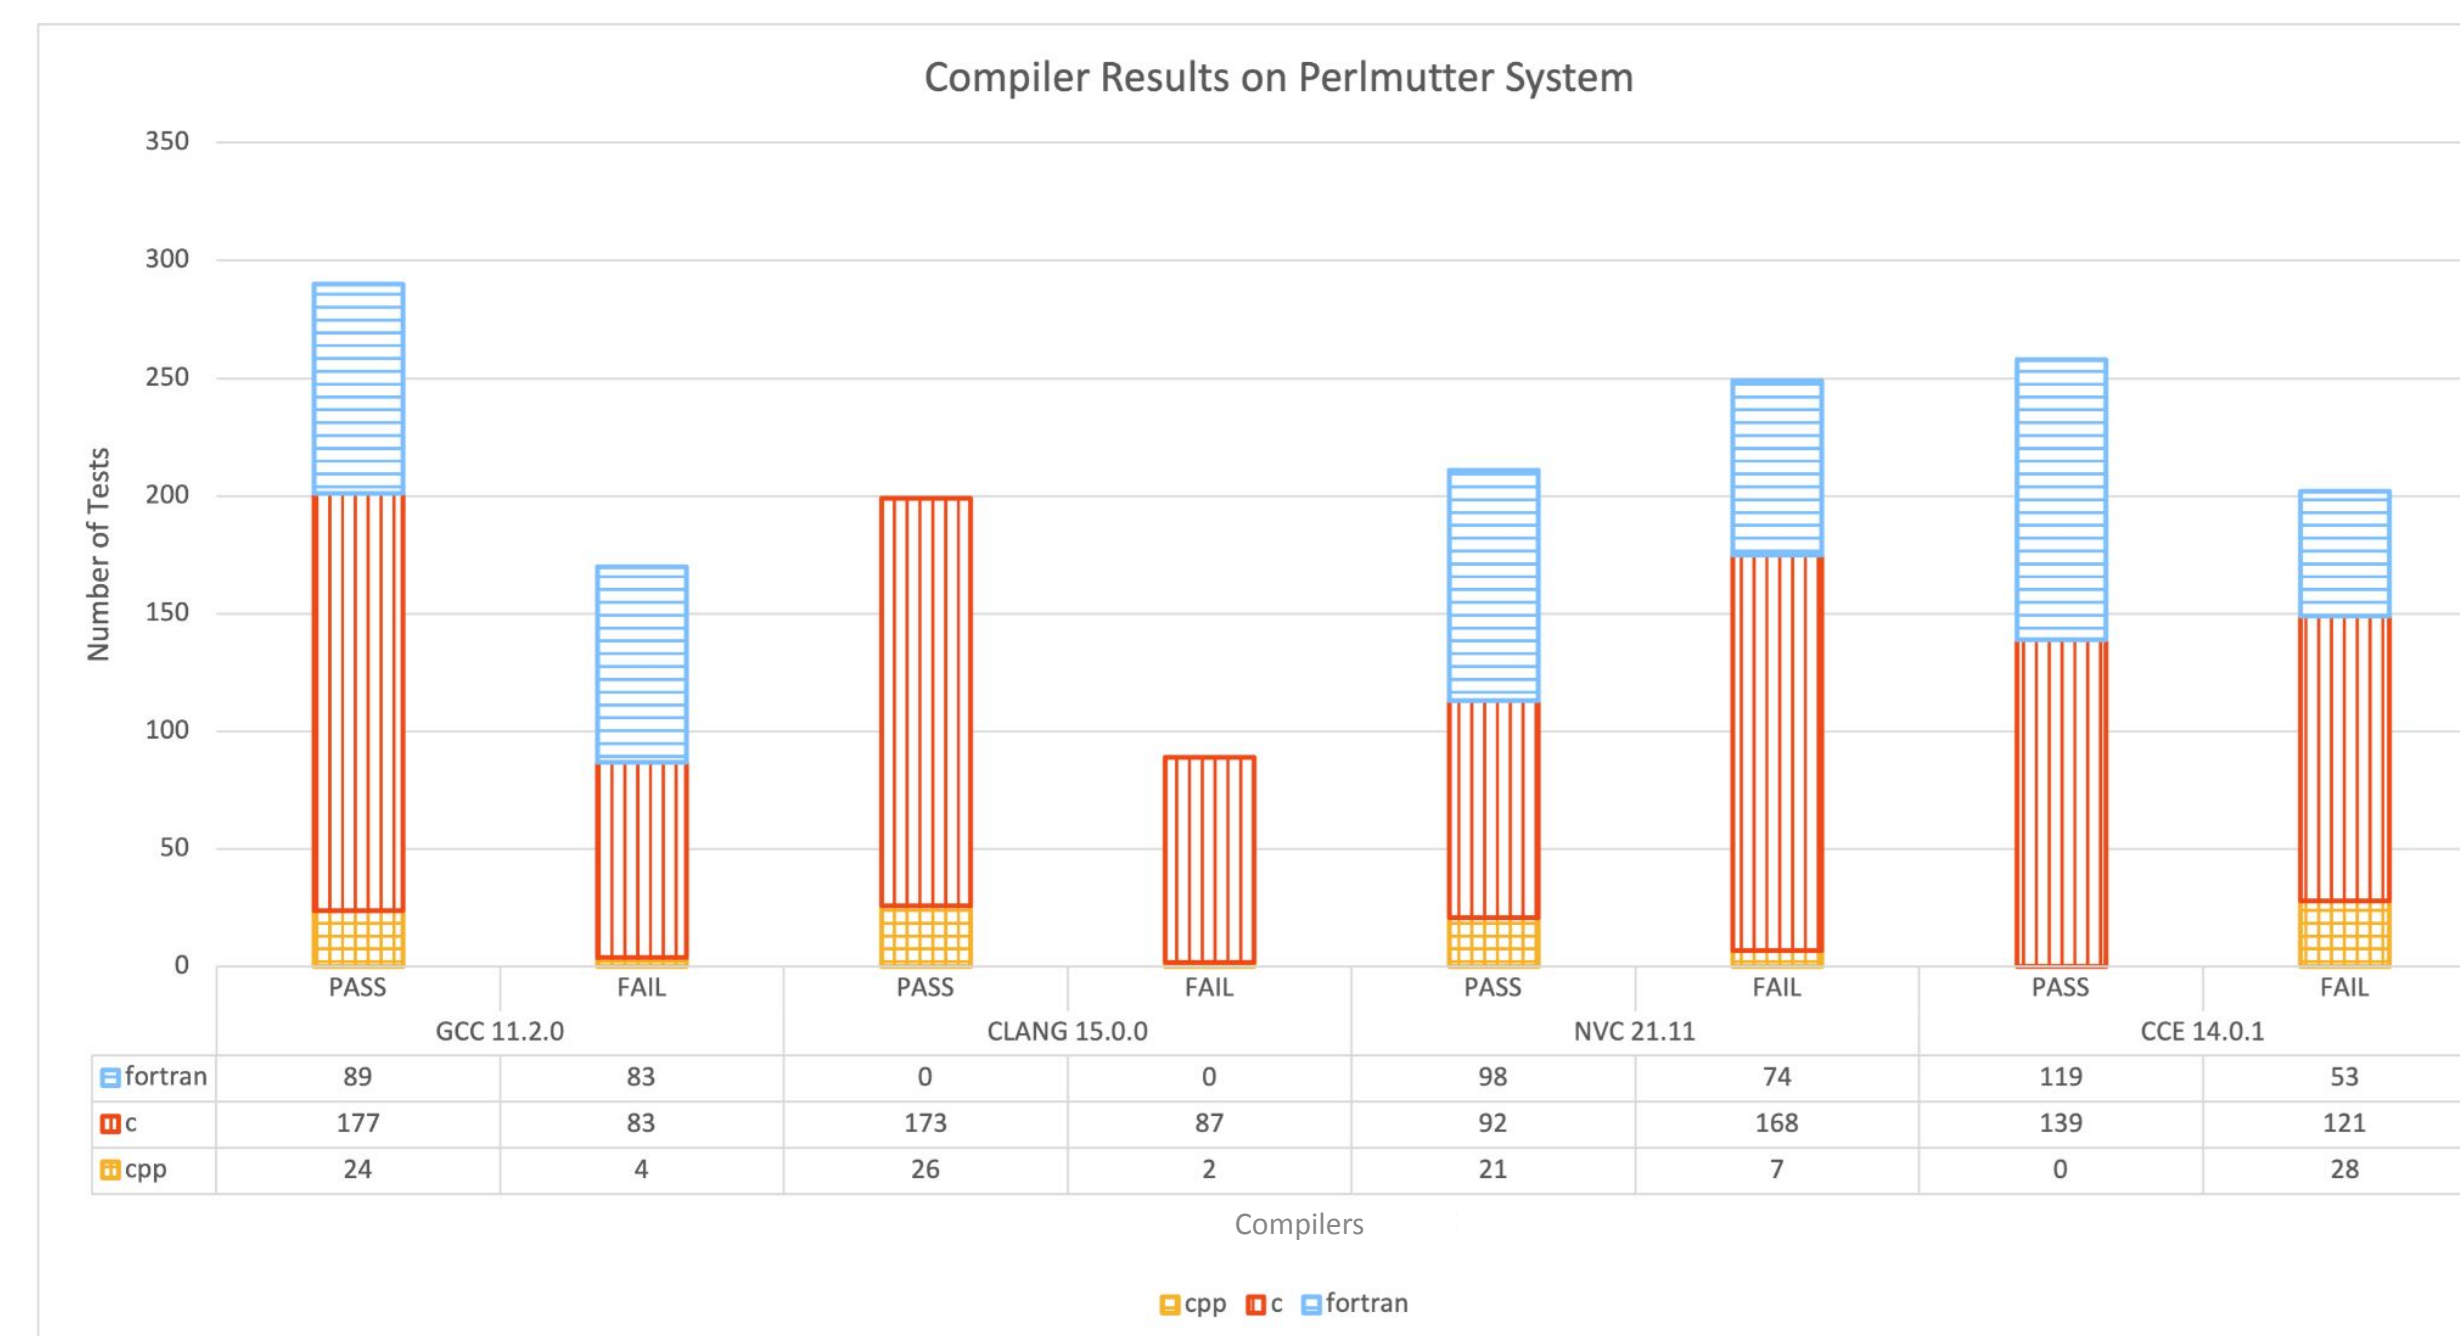


Figure 2: Results for GCC, Clang, NVC & Cray on NERSC's Perlmutter system.

Figure 2 shows GCC & Clang perform best on Perlmutter, while NVC performs the worst, with more failing tests than passing. This is interesting, as the NVHPC compiler performs much better on other systems, such as Summit (Figure 3). It is important to note that LLVM's Clang does not include a fortran compiler.




Figure 3: Results for various Gnu's GCC, LLVM's Clang & Nvidia's NVHPC compiler versions on the ORNL Summit system.

Figure 3 shows GCC performing the best on the Summit system. Compared to Perlmutter, GCC has over 100 more passing tests on the system, despite it using a slightly downgraded version of the compiler. In contrast to perlmutter, LLVM's clang compiler performs the worst on Summit.




Figure 4: Results for ROCm compiler & Cray CCE compiler on ORNL's Crusher system.

Figure 4 shows results on Crusher, which is a pre-exascale system, performs quite well with both CCE & ROCm compiler. Compared to Perlmutter, the best-performing run on Crusher has around 50 more passing tests. It is interesting to note that previous versions of CCE, such as version 13.0.0, cannot run OpenMP code as it requires dependencies from both ROCm 5 & ROCm 4.

## CONCLUSIONS

- In total, we have so far written 289 C, 26 C++ and, 172 Fortran tests
- The V&V suite targets up to OpenMP 5.2 specification
- The suite runs on pre-exascale systems (spock, crusher), Perlmutter and Argonne JLSE test beds

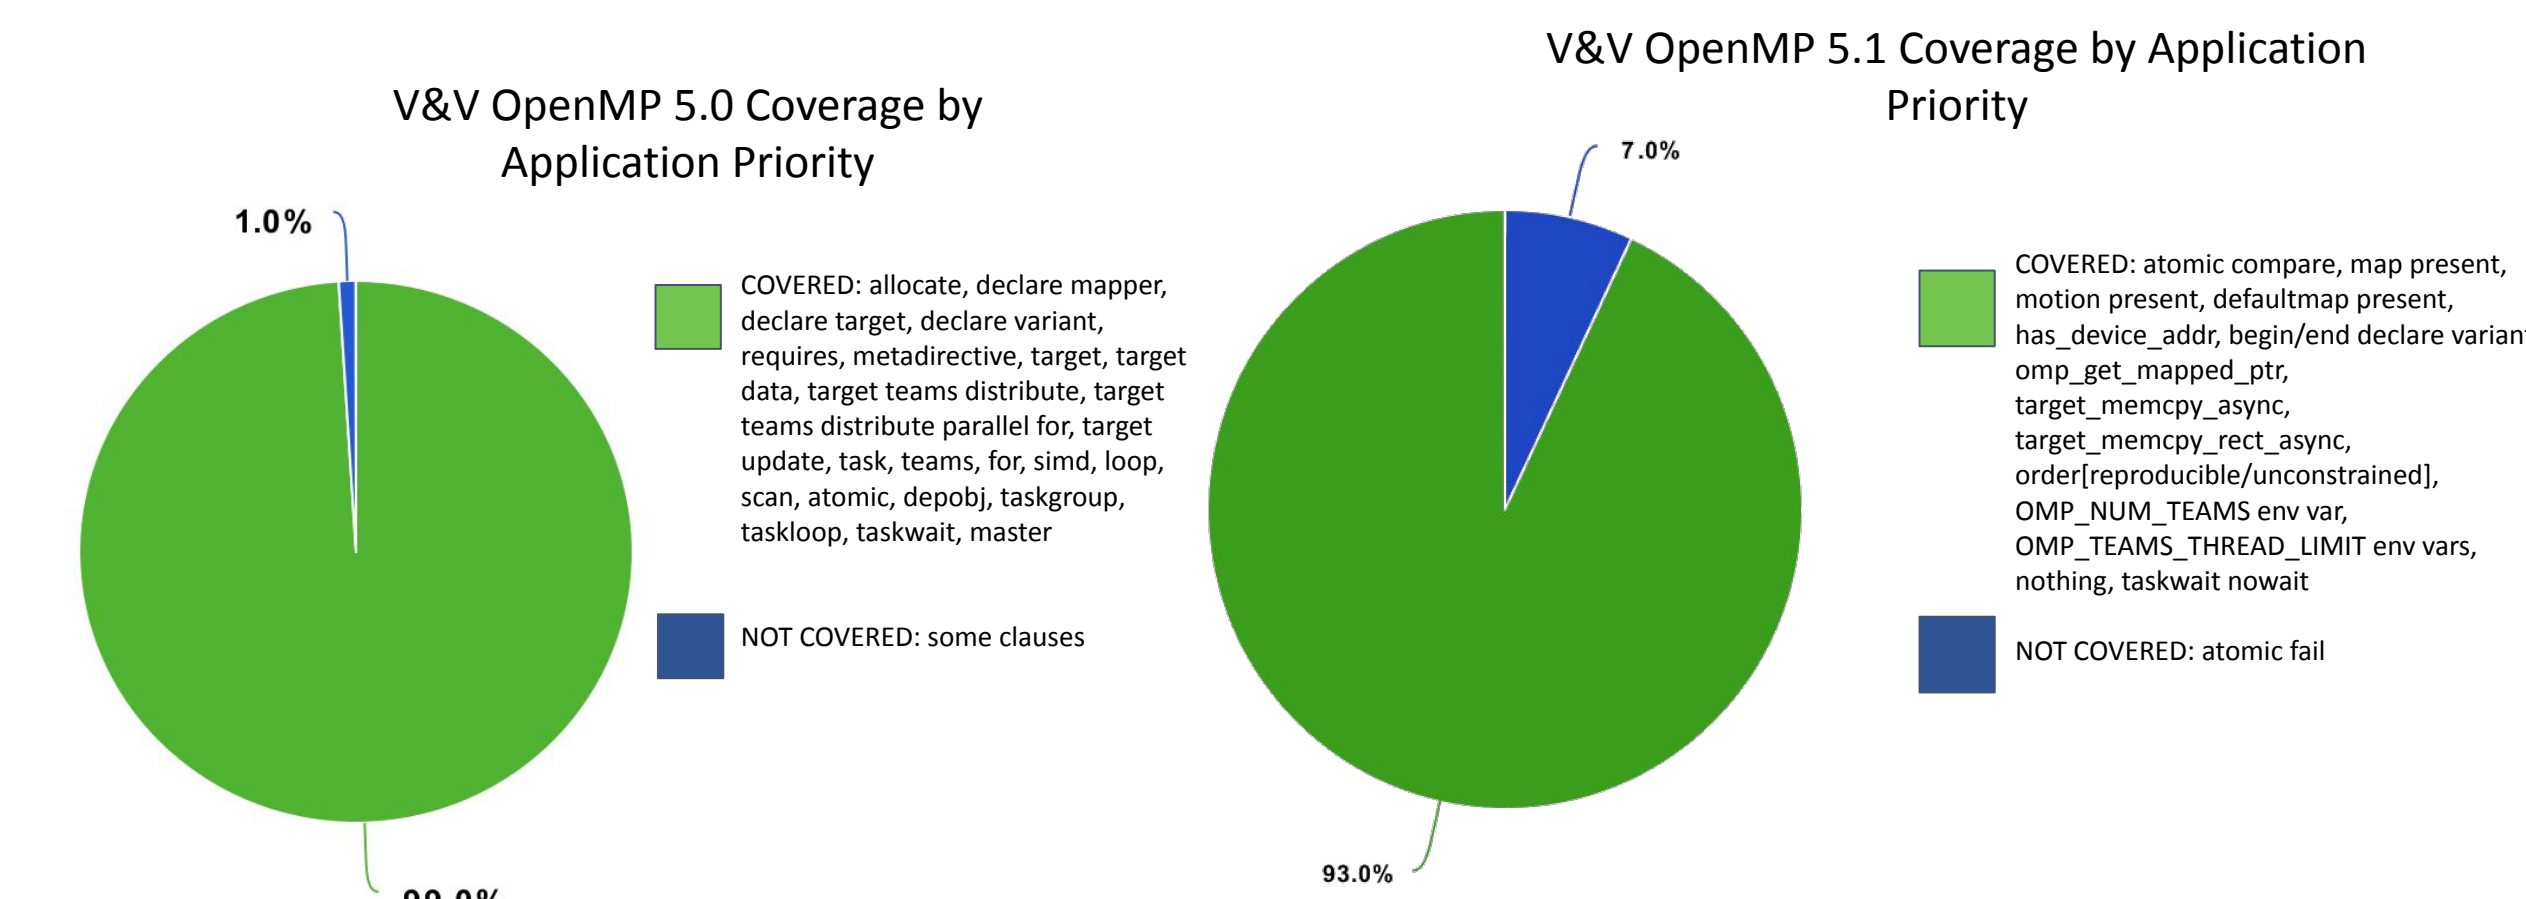


Figure 5: Pie charts showing percentage of testsuite's implementation of directives marked as "Medium" priority or higher by Application Developers at Oak Ridge National Laboratory.

GCC 12.2.1	NVHPC 22.5	LLVM 16.0.0
Atomic Compare	Pass	Fail
Atomic Compare (device)	Pass	Fail
Begin/End Declare Variant	Fail	Pass
get_mapped_ptr	Pass	Pass
Loop Order Reproducible	Pass	Fail
Loop Order Unconstrained	Pass	Fail
Metadirective (Nothing)	Fail	Pass
omp_num_teams env variable	Pass	Pass
omp_teams_thread_limit env variable	Pass	Fail
Parallel For Order Reproducible	Pass	Fail
Parallel For Order Unconstrained	Pass	Fail
Target Defaultmap(present)	Fail	Fail
Defaultmap(present) w/ scalar	Fail	Fail
Target has_device_addr	Pass	Fail
Target memcopy_async_depobj	Pass	Fail
Target memcopy_async_no_obj	Pass	Fail
Target memcopy_rect_async_depobj	Pass	Fail
Target memcopy_rect_async_no_obj	Pass	Fail
Targetupdate_to_present	Fail	Fail
Task await	Pass	Fail
Taskloop Simd Order Reproducible (device)	Pass	Fail
Taskloop Simd Order Unconstrained (device)	Pass	Fail

Figure 6: Table displaying "medium" or higher priority tests implemented by GCC 12.1.1, NVHPC 22.5 & LLVM 16.

Figure 5 shows that our suite has implemented the majority of 5.0 & 5.1 features of high priority. Figure 6 illustrates that GCC has by far implemented more key features in OMP 5.1 compared to NVHPC & LLVM.

## FUTURE WORK

- Working on new tests for the latest OpenMP specification (5.2)
- Expand our Fortran tests so that more C/C++ tests have Fortran counterparts
- Running/Testing the suite on more machines to further test the specification

## CALLS TO ACTION

GitHub V&V repository ([https://github.com/SOLLVE/sollve\\_vv](https://github.com/SOLLVE/sollve_vv)) is open for download, logging issues, suggesting tests, etc.

Scan QR code to see results!



## ACKNOWLEDGEMENTS

Thanks to OLCF: This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Thanks to ECP: This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

Thanks to NERSC: This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231