

PETSc/TAO

Many application codes rely on high-performance mathematical libraries to solve the systems of equations that must be solved during their simulation. Because the solvers often dominate the computation time of such simulations, these libraries must be efficient and scalable on the upcoming complex exascale hardware architectures for the application codes to perform well. The PETSc/TAO project delivers efficient mathematical libraries to application developers for sparse linear and nonlinear systems of equations, time-stepping methods, and parallel discretization techniques and provides libEnsemble to manage the running of a large collection of related simulations needed for numerical optimization, sensitivity analysis, and uncertainty quantification (the so-called outer-loop).

Algebraic solvers, generally nonlinear solvers that use sparse linear solvers via Newton's method, and integrators form the core computation of many scientific simulations. The Portable Extensible Toolkit for Scientific Computations/Toolkit for Advanced Optimization (PETSc/TAO) is a scalable mathematical library that runs portably on everything from laptops to the existing pre-exascale machines. The PETSc/TAO project is extending and enhancing the library to ensure that it will be performant on exascale architectures, is delivering the libEnsemble tool to manage a collection of related simulation for outer-loop methods, and is working with exascale application developers to satisfy their solver needs.

There are no scalable "black box" sparse solvers or integrators that work for all applications, nor single implementations that work well for all scales of problem size. Hence, algebraic solver packages provide a wide variety of algorithms and implementations that can be customized for the application and range of problem sizes at hand. The PETSc/TAO team is currently focusing on enhancing the PETSc/TAO library

to include scalable solvers that efficiently utilize many-core and GPU-based systems. This work includes implementing reduced synchronization algorithms that scale to larger concurrency than solvers with synchronization points and performance and data structure optimizations for the basic data structures to better utilize many-core and GPU-based computing systems as well as provide scalability to the exascale.

The availability of systems with over 100 times the processing power of today's machines compels the use of these systems, not just for a single simulation but rather within a tight outer-loop of numerical optimization, sensitivity analysis, and uncertainty quantification. This requires the implementation of a scalable library for managing a dynamic hierarchical collection of running, possibly interacting, scalable simulations. The library libEnsemble directs such multiple concurrent simulations. In this area, our team is focused on the development of libEnsemble, the integration of libEnsemble with the PETSc/TAO library, and extension of the PETSc/TAO library to include new algorithms capable of using libEnsemble.

Progress to date

- The PETSc/TAO team delivered PETSc/TAO version 3.11, which includes pipeline Krylov implementations; improved GPU support for their algebraic multigrid solver; new data structures and a new communication module backend that uses the star forest paradigm to improve performance; portability support for algebraic multigrid on GPUs; and support for multithreading in third-party libraries that use OpenMP.
- The team also delivered libEnsemble version 0.4.1, which includes an option to run libEnsemble in central or distributed configurations; updated tests, examples, and documentation; and testing using the POUNDERS and APOSSM numerical optimization solvers.
- The team also completed preliminary testing and benchmarking to confirm that the GPU backends in PETSc/TAO are working correctly and delivering performance gains for single-node configuration cases on Summit and continued benchmarking libEnsemble on Summit.
- The PETSc/TAO team shares overlapping membership with the ECP Center for Efficient Exascale Discretizations (CEED) co-design center and is working closely with them on common issues including the use of high-order matrix-free discretizations (libCEED) and scalable mesh management techniques. In addition, the team is collaborating with the AMReX co-design center on using both the TAO optimization algorithms and the PETSc iterative solvers from AMReX applications. PETSc/TAO/libEnsemble is currently used by least nine software components being developed by the ECP application teams.

PI: Barry Smith, Argonne National Laboratory

Collaborators: Argonne National Laboratory