CLOVER

Scientific applications need to apply efficient and scalable implementations of numerical operations, such as matrix-vector products and Fourier transforms, in order to simulate their phenomena of interest. Software libraries are powerful means of sharing verified, optimized numerical algorithms and their implementations. The CLOVER project is delivering scalable, portable numerical algorithms to facilitate efficient simulations. To the extent possible, the team preserves the existing capabilities in mathematical libraries, while evolving the implementations to run effectively on the pre-exascale and exascale systems and adding new capabilities that may be needed by applications.

Mathematical libraries encapsulate the latest results from the mathematics and computer science communities, and many exascale applications rely on these numerical libraries to incorporate the most advanced technologies available in their simulations. Advances in mathematical libraries are necessary to enable computational science on exascale systems, as the exascale architectures introduce new complexities that algorithms and their implementations need to address in order to be scalable, efficient, and robust. The CLOVER project is ensuring the healthy functionality of the mathematical libraries on which these applications depend. The libraries supported by the CLOVER project, SLATE, heFFTe PEEKS, and Kokkos Kernels, span the range from lightweight collections of subroutines with simple application programming interfaces (APIs) to more "end-to-end" integrated environments and provide access to a wide range of algorithms for complex problems.

SLATE provides dense linear algebra operations for large-scale machines with multiple GPU accelerators per node. The team focuses on adding support to SLATE for the most critical workloads required by exascale applications: linear systems, least squares, matrix inverses, singular value problems, and eigenvalue problems.

heFFTe implements the fast Fourier transform used in many domain applications including molecular dynamics, spectrum estimation, fast convolution and correlation, signal modulation, and wireless multimedia applications. The team is designing and implementing a fast and robust 2D and 3D fast Fourier transform library that targets large-scale heterogeneous systems with multi-core processors and hardware accelerators.

PEEKS is delivering production-quality, next-generation latency-tolerant, and scalable preconditioned iterative solvers. The team is producing the design and infrastructure support required to effectively implement these solvers and delivering them using a standardized API.

Kokkos Kernels provides performance portable sparse and dense linear algebra and graph kernels on current and future heterogeneous architectures. The team is delivering architecture-aware, high-performance algorithms for performance-critical kernels to applications for use on pre-exascale and exascale architectures.

PI: Jack Dongarra, University of Tennessee–Knoxville

Collaborators: University of Tennessee–Knoxville, Sandia National Laboratories

Progress to date

- The CLOVER team produced a version of SLATE that supports BLAS 3, norms, linear solvers, mixed-precision linear solvers, and least-squares solvers and includes compatibility APIs for LAPACK and ScaLAPACK users.
- The team completed a design and implementation for heFFTe targeting distributed accelerated systems that includes various technologies for scheduling computation and communications, highly optimized GPU kernels, and CUDA-aware MPI routines.
- The team deployed and benchmarked the PEEKS implementation of the parallel generation of preconditioners based on incomplete factorization, developed a parallel threshold ILU factorization, and incorporated pipeline Krylov solvers.
- The team is porting Kokkos Kernels to run efficiently on the pre-exascale machines and is adding methods to address the needs identified by exascale applications.