

ALEXA

Many scientific applications are written in Fortran and need to access scalable algorithms for efficiency, need to pass data between different grids with different parallel distributions, or need reduced representations of high-dimensional data, for example, to optimize storage. The Accelerated Libraries for Exascale (ALExa) project is providing technologies to address these needs for exascale applications.

Complex scientific applications may need to combine results from different computational grids to perform their required simulations, where each computational grid represents only part of the physics. Moreover, the simulations on each grid may be written in Fortran and require access to scalable solvers in C++. The ALExa project is developing three components to address these issues and enable applications to better use exascale systems: the Data Transfer Kit (DTK), Tasmanian, and ForTrilinos.

The DTK provides the ability to transfer computed solutions between grids with different layouts on parallel accelerated architectures, enabling simulations to seamlessly combine results from different computational grids to perform their required simulations. The team is focused on adding new features needed by applications and ensuring that the library is

performant on the pre-exascale and exascale architectures.

Tasmanian provides the ability to construct surrogate models with low memory footprint, low cost, and optimal computational throughput, enabling optimization and uncertainty quantification for large-scale engineering problems, as well efficient multi-physics simulations. The team is focused on reducing the GPU memory overhead and accelerating the simulation of the surrogate models produced.

ForTrilinos provides a software capability for easy automatic generation of Fortran interfaces to any C/C++ library, as well as a seamless pathway for large and complex Fortran-based codes to access the Trilinos library through automatically generated interface code.

Progress to date

- The ALExa team demonstrated the DTK performance portable search capability on multi-threaded platforms, which exhibited a 10–15× speedup over standard search libraries.
- The team enabled GPU-accelerated surrogate model simulations in TASMANIAN, developed new algorithms for asynchronous surrogate construction that exploit extreme concurrency, and demonstrated a 100× reduction of memory footprint in sparse representation of neutrino opacities for the ExaStar project.
- The team developed a SWIG/Fortran tool that automatically generates Fortran object-oriented interfaces and necessary wrapper code for any given C/C++ interface, demonstrated advanced inversion-of-control functionality that allows a C++ solver to invoke user-provided Fortran routines, and used this tool to provide Fortran access to a wide variety of linear and nonlinear solvers in the Trilinos library.

PI: John Turner, Oak Ridge National Laboratory

Collaborators: Oak Ridge National Laboratory