heFFTe

The fast Fourier transform (FFT) is used in many domain applications—including molecular dynamics, spectrum estimation, fast convolution and correlation, signal modulation, and wireless multimedia applications. For example, distributed 3-D FFT is one of the most important kernels used in molecular dynamics computations, and its performance can affect an application's scalability on larger machines. Similarly, the performance of the first principle calculations depends strongly on the performance of the FFT solver. Specifically, for the US Department of Energy (DOE), we found that more than a dozen Exascale Computing Project (ECP) applications use FFT in their codes. To address these needs, ICL released the Highly Efficient FFTs for Exascale (heFFTe) library. The heFFTe v2.4 library release features very good weak and strong scalability, and performance that is close to 90% of the roofline peak.

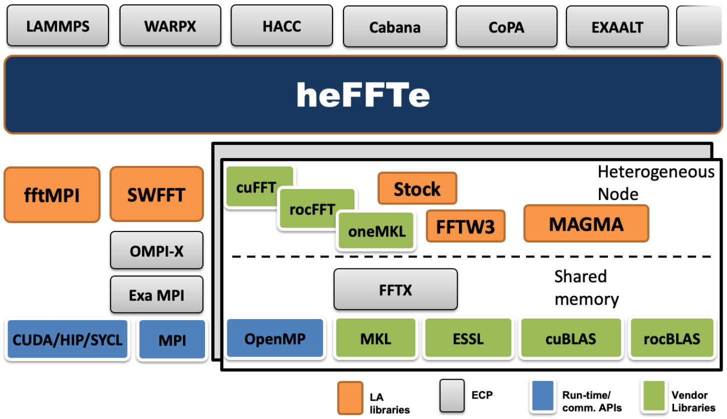

LAMMPS WARPX HACC Cabana CoPA EXAALT

heFFTe

fftMPI SWFFT OMPI-X Exa MPI CUDA/HIP/SYCL MPI

Heterogeneous Node

cuFFT rocFFT oneMKL Stock FFTW3 MAGMA

Shared memory

FFTX

OpenMP MKL ESSL cuBLAS rocBLAS

🟧 LA libraries ⬜ ECP 🟦 Run-time/comm. APIs 🟩 Vendor Libraries

The current state-of-the-art FFT libraries are not scalable on large heterogeneous machines with many nodes or even on one node with multiple high-performance GPUs (e.g., several NVIDIA V100 GPUs). Furthermore, these libraries require large FFTs in order to deliver acceptable performance on one GPU. Efforts to simply enhance classical and existing FFT packages with optimization tools and techniques—like autotuning and code generation—have so far not been able to provide the efficient, high-performance FFT library capable of harnessing the power of supercomputers with heterogeneous GPU-accelerated nodes. In particular, ECP applications that require FFT-based solvers might suffer from the lack of fast and scalable 3-D FFT routines for distributed heterogeneous parallel systems, which is the very type of system that will be used in upcoming exascale machines.

PUBLICATIONS

https://icl.utk.edu/fft/#papers

S. Cayrols, J. Li, G. Bosilca, S. Tomov, A. Ayala, J. Dongarra
Lossy all-to-all exchange for accelerating parallel 3-D FFTs on hybrid architectures with GPUs
2022 IEEE International Conference on Cluster Computing (CLUSTER), 2022.

A. Ayala, S. Tomov, M. Stoyanov, and J. Dongarra
Scalability Issues in FFT Computation
Parallel Computing Technologies (PaCT 2021), Lecture Notes in Computer Science, vol 12942.

A. Ayala, S. Tomov, P. Luszczek, S. Cayrols, G. Ragghianti, and J. Dongarra
Interim Report on Benchmarking FFT Libraries on High Performance Systems
Innovative Computing Laboratory Technical Report, no. ICL-UT-21-03: University of Tennessee, July 2021.

D. Sharp, M. Stoyanov, S. Tomov, and J. Dongarra
A More Portable HeFFTe: Implementing a Fallback Algorithm for Scalable Fourier Transforms
2021 IEEE High Performance Extreme Computing Conference (HPEC'21).

SPONSORED BY

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP

NNSA National Nuclear Security Administration

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.

IN COLLABORATION WITH

OAK RIDGE National Laboratory

heFFTe

A 3-D FFT Execution Trace with Main FFT Components

(80 MPI processes on Intel Xeon E5-2650 v3 cluster, 1K × 1K × 1K grid)

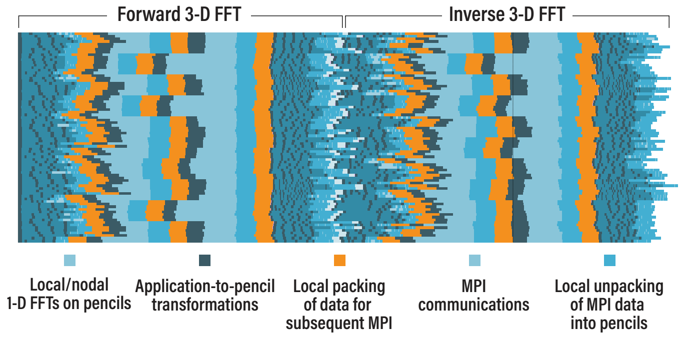

heFFTe Framework Design with Flexible API

(Need flexible FFT API for application-specific input and output)

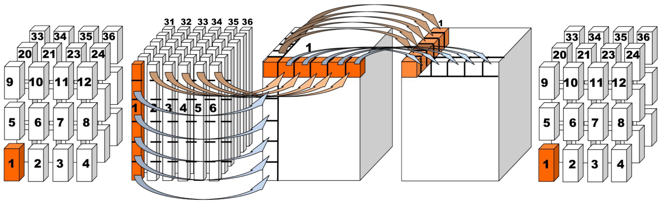

- Main objective is the design and implementation of a sustainable FFT library for exascale platforms;
- FFT-ECP's goal is to help key ECP applications and provide an efficient and flexible FFT API to take application-specific input and outputs;
- Approach is to: (1) collect and leverage existing FFT capabilities to build a sustainable FFT library (instead of creating a new and independent software stack); (2) optimize data movements and overlap computations with communications; (3) autotuning.

heFFTe Acceleration on GPUs

Local CPU kernels presented on Section 2 are typical on state-of-the-art parallel FFT libraries, heFFTe provide new GPU kernels for these tasks achieving over 40x speedup.

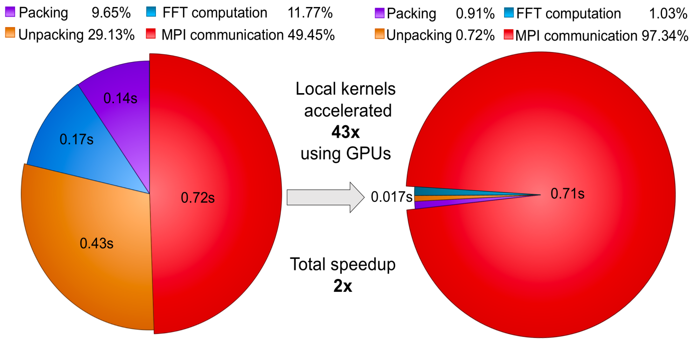

Local kernels accelerated **43x** using GPUs

Total speedup **2x**

Scalability

In the following figures, we showcase the scalability of heFFTe on both the Frontier and Summit supercomputers. Note that the observed speedups surpass the theoretical expectation of 4x, reaching 5.7x for strong scaling and 5.3x for weak scaling.

Strong scaling of heFFTe on 1024^3 FFT problem

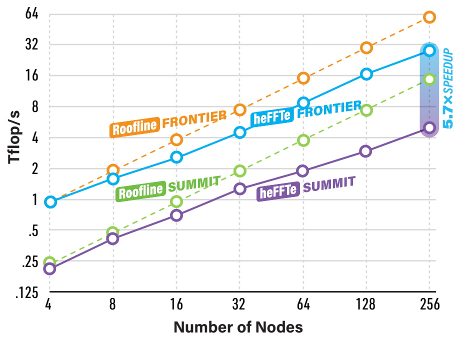

Weak scaling of heFFTe

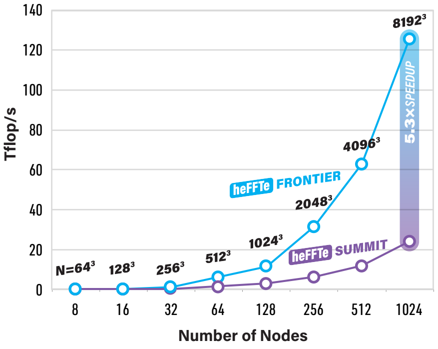

SPONSORED BY

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.

IN COLLABORATION WITH