

FFT Benchmark Performance Experiments on Systems Targeting Exascale

Alan Ayala
Stanimire Tomov
Piotr Luszczek
Sébastien Cayrols
Gerald Raghianti
Jack Dongarra

ICL Technical Report ICL-UT-22-02

Knoxville, March 1, 2022

Abstract

The Exascale Computing Project (ECP) of the United States, supports the design of software and hardware towards computation at exaflops rate. Many applications of this kind, rely on efficient and scalable Fast Fourier Transform (FFT) computation. In a previous work, we presented a interim benchmark report comparing almost a dozen state-of-the-art FFT libraries, analyzing their scalability and main features.

In this report, we present FFT Benchmark performance experiments of parallel multidimensional FFT libraries, with an emphasis on hardware of large-scale heterogeneous systems with hardware accelerators that are available in pre-exascale systems and the upcoming supercomputers within the scope of ECP. The results are obtained using an FFT Benchmark harness that was developed to easily benchmark and compare numerous FFT libraries on DOE exascale systems.

Contents

1	Introduction	7
2	Experimental Setup	9
2.1	Description of Hardware Resources	9
2.2	Description of Software Resources	10
2.3	Data Inputs and Outputs	10
3	FFT Performance on Summit and Pre-Frontier (Spock)	12
3.1	Bandwidth Analysis	12
3.2	Theoretical Peak Performance	13
3.2.1	FFT Performance Analysis for a Multi-Core CPU	13
3.2.2	FFT Performance Analysis for a single GPU	13
3.3	Comparison of Strong Scalability Results	15
3.3.1	CPU-based libraries	15
3.3.2	GPU-based libraries	16
4	Discussion	17
4.1	Impact of Vendors' GPUs and Software	17
4.2	Impact of the libraries setup time	18
4.3	The FFT communication bottleneck	19

5	Conclusions	20
6	Appendix: Tensor transposition cost for parallel FFTs	23
6.1	Multi-core CPU based systems	24
6.2	GPU based systems	26

List of Figures

2.1	In this report we evaluate the performance of parallel 3-D FFT libraries using the 1-D (slabs) decomposition or the 2-D (pencils) decomposition. In both cases we use the minimum number of reshapes and permutations. . . .	11
3.1	Evolution of bandwidth achieved on Summit and Spock of inter-node communication when increasing the communication volume and using one NIC. Summit has a theoretical peak of 50 GB/s (when using 2 NICs per node) while Spock's peak bandwidth is 25GB/s (using one NIC).	12
3.2	Percentage of theoretical peak achieved on a multi-core CPU of Summit and Spock.	13
3.3	Percentage of theoretical peak achieved on a single GPU of Summit and Spock.	14
3.4	Comparison of parallel FFT libraries on up to 4 Summit nodes, using 4 MPIs per node and 1 MPI per single core of IBM POWER 9.	15
3.5	Comparison of parallel FFT libraries on up to 4 Spock nodes, using 4 MPIs per node and 1 MPI per EPYC-7662 core.	15
3.6	Comparison of parallel FFT libraries with GPU support on up to 4 Summit nodes, using 4 MPIs per node and 1 GPU per MPI.	16
4.1	Performance comparison of a 256^3 FFT using heFFTe with vendor 1-D FFT libraries (NVIDIA and AMD) on Summit and Spock. The number of GPUs used per node is four.	17
4.2	Comparison of the average setup (FFT planning) time for all libraries from Figure 3.4 at 4 nodes.	18

4.3	Comparison of setup (planning) time for different libraries. Note that FFTADV MPI is only composed of two kernels, since the packing and unpacking is embedded into MPI_Alltoallw.	19
6.1	Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 168 IBM Power9 cores, 42 MPIs per node. We use heFFTe with FFTW backend and pipelined MPI_Isend and MPI_Irecv communication.	24
6.2	Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 168 IBM Power9 cores, 42 MPIs per node. We use heFFTe with FFTW backend and MPI_Alltoall communication.	25
6.3	Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 16 NVIDIA GPUs, 4 MPIs per node. We use heFFTe with CUFFT backend and pipelined MPI_Isend and MPI_Irecv communication.	26
6.4	Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 16 NVIDIA GPUs, 4 MPIs per node. We use heFFTe with CUFFT backend and MPI_Alltoall communication.	27

List of Tables

1.1	Multidimensional FFT Libraries in the FFT benchmark tests. All libraries were tested using their latest release version.	7
2.1	Software versions used on Summit.	10
2.2	Software versions used on Spock.	10
6.1	Available MPI routines in FFT libraries	23

Acknowledgment

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering and early tested platforms, in support of the nation's exascale computing imperative.

Introduction

In this report, we present results from our FFT Benchmark experiments in order to analyze the performance of 3-D FFT [1, 2, 3] on both the pre-exascale systems as well as on the hardware expected to arrive in time for upcoming exascale machines. To this end, we consider nine state-of-the-art software libraries listed in Table 1.1 that we had accessed to and consequently used throughout this report.

To the extend possible, the presented results were obtained from using a common harness that was developed especially to compare FFT libraries in a common environment. The harness has been developed as open source project [4] and includes both performance measurements and error checking.

Table 1.1: Multidimensional FFT Libraries in the FFT benchmark tests. All libraries were tested using their latest release version.

Library Name	Version	CPU Test	GPU Test
AccFFT	2.0	yes	yes
2Decomp&FFT	1.5.847	yes	no
FFTE	7.0	yes	yes
FFTW	3.3.8	yes	no
FFTMPI	1.0	yes	no
heFFTe	2.2	yes	yes
SWFFT	1.0	yes	no
P3DFFT	3	yes	no
FFTADVMPI	N/A	yes	no

Amongst the most recent libraries from Table 1.1, heFFTe [5] and FFTW [6] libraries are

used for error validation and their results served as the basis for comparisons with the other FFT libraries. This is, for a given input array $X \in \mathbb{R}^{m \times n \times q}$ and a parallel FFT library, we calculate the computation error as shown in Eqs. (1.1) and (1.2).

$$E_1 = \|X - \mathbf{IFFT}_{\text{heFFTe}}(\mathbf{FFT}_{\text{Library}}(X))\|_{\max}, \quad (1.1)$$

$$E_2 = \|X - \mathbf{IFFT}_{\text{FFTW}}(\mathbf{FFT}_{\text{Library}}(X))\|_{\max}, \quad (1.2)$$

This is, we compute an inverse transform (with heFFTe and FFTW) of the forward transform obtained by the given library. For the experimental part in Section 3, we verify that the errors are less than 10^{-14} , which is roughly in the order of $\mu\|X\|_2$, being the machine precision $\mu = O(10^{-16})$.

The harness software that we prepared for testing the libraries from Table 1.1 uses hardware GPU acceleration whenever possible and allows mixed-interface to accommodate C, C++, and Fortran codes by wrapping and properly linking object files together with the required dependent runtime, system, and language. Some tests still had to be performed outside the software harness due to the cross-language and cross-platform portability issues.

In our previous benchmark report [7], we considered all libraries from Table 1.1, except FFTADVMPi, which is the name we give to an implementation that was recently integrated into our software harness based on the work from authors in [8].

Experimental Setup

2.1 Description of Hardware Resources

The performance experiments in this report are obtained using two supercomputers located at the Oak Ridge National Laboratory, namely Summit and Spock. The former is an IBM AC922 supercomputer featuring IBM POWER 9 CPUs and NVIDIA Volta V100 GPUs and is currently ranked number two on the TOP500 list of the largest supercomputing installations. The latter machine is Spock, which is a precursor of the upcoming Frontier supercomputer at ORNL. Frontier is expected to overcome the 1 Exa-flop/s barrier using double-precision data. Spock is composed of nodes with AMD CPUs and GPUs. Below we list their technical specifications relevant for this report:

- Summit has a total of 4,608 nodes. Each node consists of two sockets, each with a 22-core IBM POWER 9 CPU and 3 NVIDIA Volta V100 GPUs. The total of 6 GPU accelerators provide a theoretical double-precision peak performance of approximately 47 Tera-flop/s. Within the same socket, the computing units are connected with the NVIDIA NVLink node interconnect with a theoretical uni-directional bandwidth of 50 GB/s (or 100 GB/s bi-directional). Inter-node injection bandwidth is limited to two Mellanox Infiniband NICs totaling in uni-directional bandwidth of 25 GB/s (or 50 GB/s bi-directional).
- Spock has hardware and software in preparation for the upcoming much larger Frontier system. We use it as an early-access testbed available as part of the Exascale Computing Project (ECP) early access platforms. In total, it has a total of 36 compute nodes, each consisting of one 64-core AMD EPYC 7662 CPU equipped with to 256 GB of DDR4 memory and connected to 4 AMD MI100 GPUs. The CPU is connected to all GPUs via PCIe Gen4 with uni-directional bandwidth of 32 GB/s (or 64

GB/s of bi-directional bandwidth). The GPUs are connected in an all-to-all arrangement via the Infinity Fabric (xGMI) with uni-directional bandwidth of 46 GB/s (or 92 GB/s bi-directional bandwidth). The Spock nodes are connected with Slingshot-10 dragonfly network providing a node injection uni-directional bandwidth of 12.5 GB/s (25 GB/s bi-directional) from a single NIC card.

2.2 Description of Software Resources

For this report we use the libraries showed in Table 1.1; for their compilation and execution, we use the software listed below in Tables 2.1 and 2.2.

Table 2.1: Software versions used on Summit.

Software Module	Version Used in Tests
CUDA	11.0.3
FFTW	3.3.9
GNU compilers	9.3.0
Spectrum MPI	10.4.0.3-20210112
CMake	3.21.3
PGI	20.4
Scalasca	2.5
ScoreP	7.0
Vampir	9.0

Table 2.2: Software versions used on Spock.

Software Module	Version Used in Tests
ROCM	4.5.0
FFTW	3.3.9
GNU compilers	9.3.0
Cray-MPICH	8.1.12
CMake	3.21.3

2.3 Data Inputs and Outputs

We performed a variety of experiments to benchmark the nine state-of-the-art parallel FFT libraries mentioned earlier with only three of them having support for GPU accelerators to offload computation. More specifically, we used the following data sizes and distribution formats:

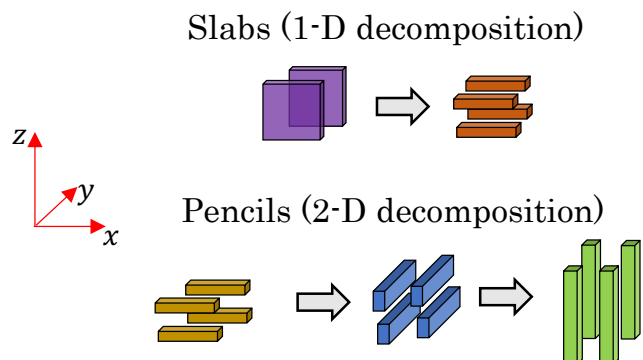


Figure 2.1: In this report we evaluate the performance of parallel 3-D FFT libraries using the 1-D (slabs) decomposition or the 2-D (pencils) decomposition. In both cases we use the minimum number of reshapes and permutations.

- A 3-D FFT transform of size 256^3 was used throughout for strong scaling purposes. The input was a set of random double-precision complex numbers. We chose this size because it is more often found in the applications we encountered, especially in comparison to larger sizes, such as 1024^3 in the previous report [7]. This of course presents a greater challenge for strong-scaling experiments as the data movement is likely to dominate at smaller node counts but is more representative of the FFT sizes in practical use.
- We used **pencil** shape as both the beginning and ending data distribution in order to reduce the number of reshape operations to a minimum. Also, this allows more software libraries to be included in this survey, as some of the tested libraries do not support any general **brick** decomposition. We also considered the case where the input **pencils** form a **slab** decomposition. The cases considered had one reshape operation for slabs decomposition, and two reshape operations for pencils decomposition. Figure 2.1 shows the distribution considered in our tests.
- In order to maintain statistical rigor, we report the average time of ten consecutive runs of single forward 3-D FFT transforms and use that value on the graphs directly or convert them into Gflop/s rate using the well-known FFT complexity: $5N \log_2(N)$, where N is the FFT size.
- For a fair comparison, we employed our benchmark harness [4], which ensures that all libraries used the same input and output data as well as MPI process grids or to perform the same type of communication exchange: either slabs or pencils. Also for fairness purposes, we disabled the tuning phases in case it was enabled by default in any given library: this was done to report out-of-the-box experimental results as some of the libraries do not include tuning as one of the phases included in their functional scope.

FFT Performance on Summit and Pre-Frontier (Spock)

3.1 Bandwidth Analysis

In Figure 3.1, we compare the practical inter-node peak bandwidth we obtained on either Summit or Spock. We used the same configuration of the Network Interface Cards (NIC) on both systems. Note that the inter-node bandwidth achieved on both systems gets close to the peak. Summit has two NICs per node and thus gets twice the overall bandwidth of Spock that was only equipped with a single network card.

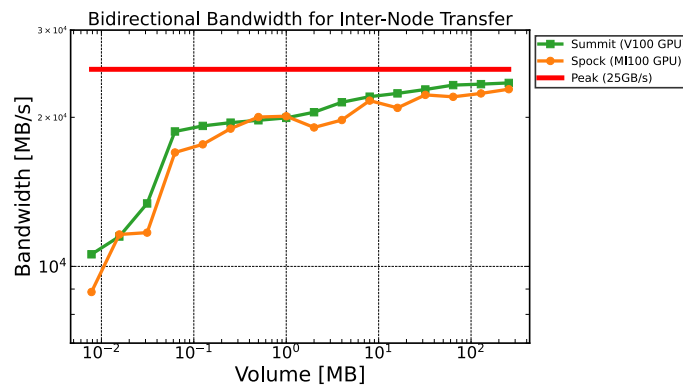


Figure 3.1: Evolution of bandwidth achieved on Summit and Spock of inter-node communication when increasing the communication volume and using one NIC. Summit has a theoretical peak of 50 GB/s (when using 2 NICs per node) while Spock's peak bandwidth is 25GB/s (using one NIC).

3.2 Theoretical Peak Performance

3.2.1 FFT Performance Analysis for a Multi-Core CPU

In Figure 3.2, we compare how close to the theoretical peak performance we can get on each CPU architecture. We tested the sizes that correspond to the ones that are applicable for GPUs in figures below even though we could have used much larger data sets due to the much larger size of the main memory accessible from the CPU.

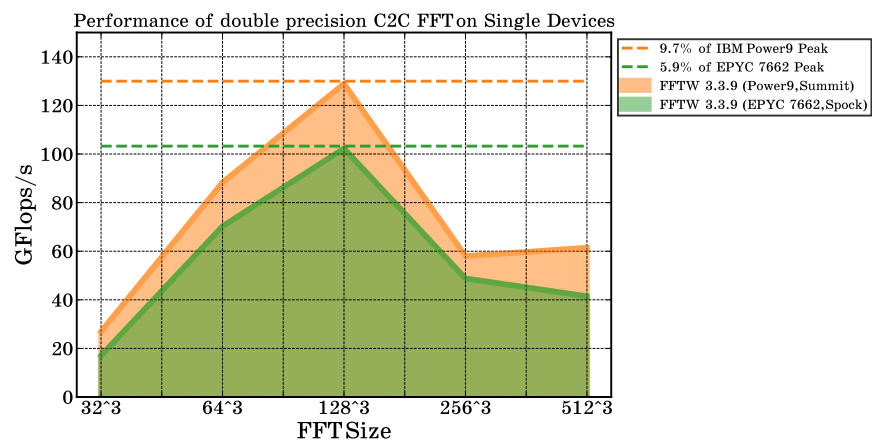


Figure 3.2: Percentage of theoretical peak achieved on a multi-core CPU of Summit and Spock.

3.2.2 FFT Performance Analysis for a single GPU

In Figure 3.3, we show performance for both GPU platforms. On a single GPU, we observed that the largest 3-D FFT we can compute using either cuFFT [9] or rocFFT [10] was 512³.

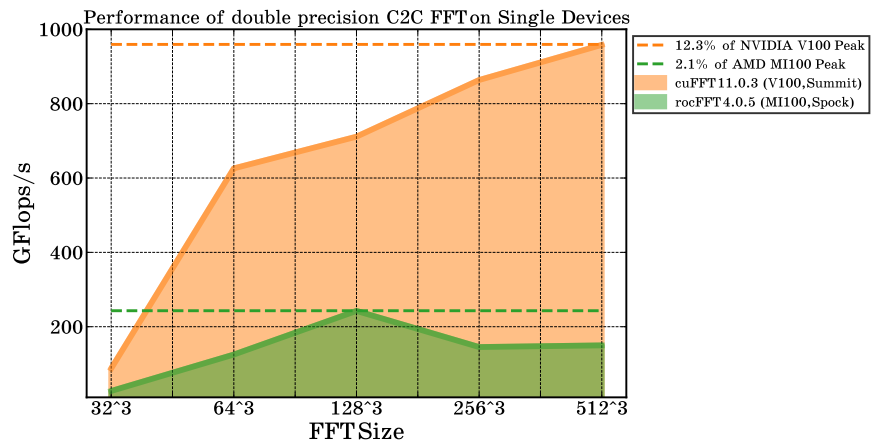


Figure 3.3: Percentage of theoretical peak achieved on a single GPU of Summit and Spock.

3.3 Comparison of Strong Scalability Results

3.3.1 CPU-based libraries

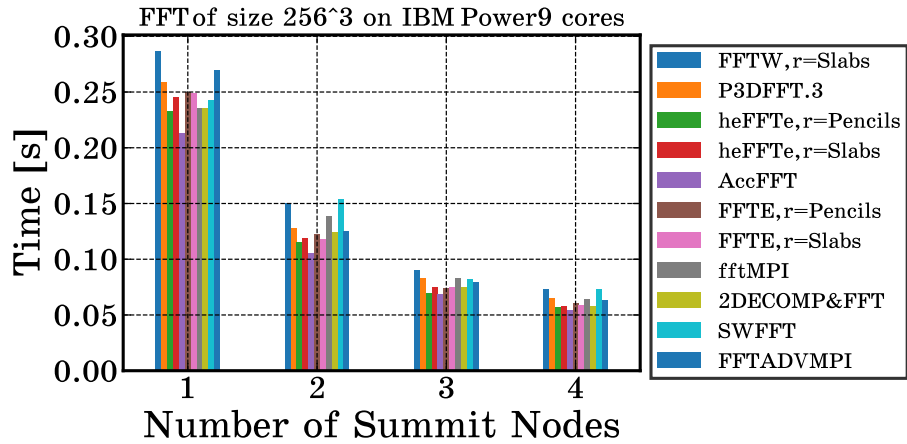


Figure 3.4: Comparison of parallel FFT libraries on up to 4 Summit nodes, using 4 MPIs per node and 1 MPI per single core of IBM POWER 9.

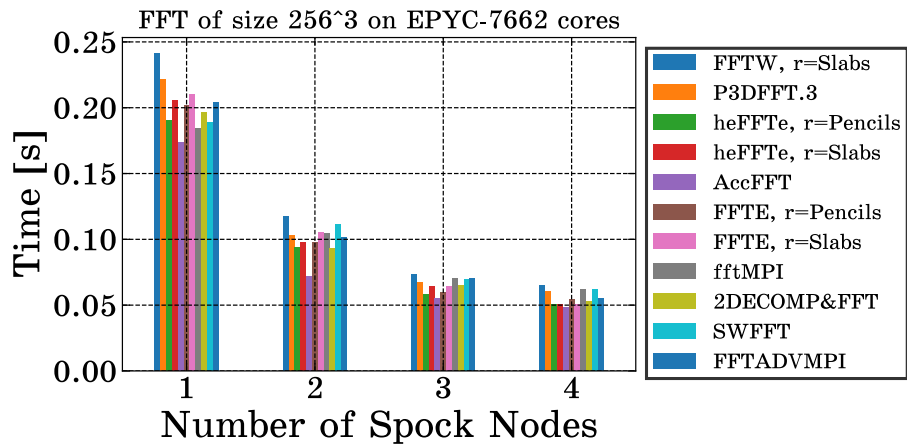


Figure 3.5: Comparison of parallel FFT libraries on up to 4 Spock nodes, using 4 MPIs per node and 1 MPI per EPYC-7662 core.

Figures 3.4 and 3.5 show the strong scalability timing of 256^3 FFTs on the multicore CPUs of Summit and Spock, respectively. These results feature all seven software libraries considered

for this report and hence these figures offer the most comprehensive comparison between them.

3.3.2 GPU-based libraries

Figure 3.6 shows the strong scaling timings of 256^3 FFTs on the GPUs of Summit spread among up to 4 nodes. Note that heFFTe, AccFFT, and FFTE are the only libraries compared in the figure as they are the only ones that provide support for distributed-memory systems with NVIDIA GPUs. Finally, heFFTe is the only library that provides support for distributed-memory systems with AMD GPUs. Therefore, we show in Section 4.1 performance results on Spock exclusively featuring the heFFTe library.

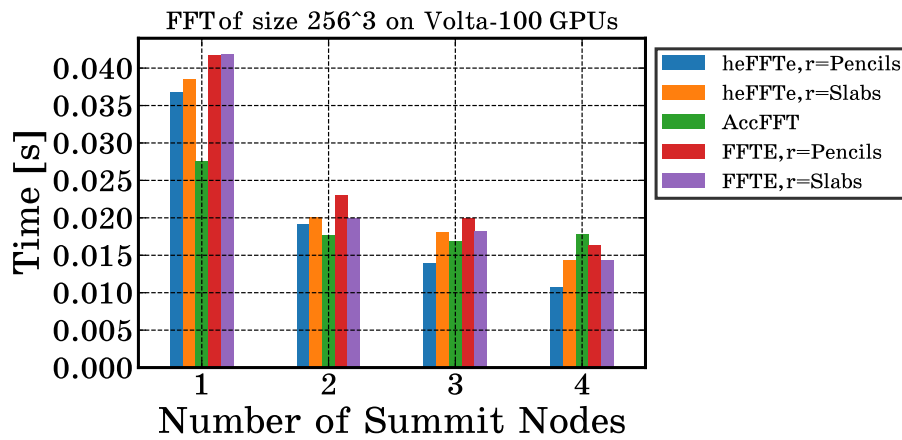


Figure 3.6: Comparison of parallel FFT libraries with GPU support on up to 4 Summit nodes, using 4 MPIs per node and 1 GPU per MPI.

Discussion

In this chapter, we analyze different scenarios that can play a fundamental role in tuning FFT software implementations to adapt them to run efficiently on exascale hardware components.

4.1 Impact of Vendors' GPUs and Software

At the time of this writing, heFFTe is the only library that provides multidimensional FFTs for distributed-memory systems with AMD and Intel GPUs; therefore, it is not currently possible to compare hybrid FFT performance on Spock-like systems.

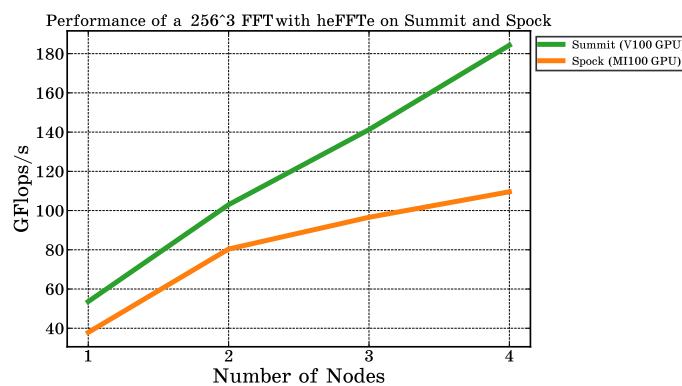


Figure 4.1: Performance comparison of a 256^3 FFT using heFFTe with vendor 1-D FFT libraries (NVIDIA and AMD) on Summit and Spock. The number of GPUs used per node is four.

Figure 4.1 shows the strong scalability timing of 256^3 FFTs of heFFTe on the GPUs of Summit in comparison with Spock. The performance shown is for heFFTe’s cuFFT backend on Summit and heFFTe’s rocFFT backend on Spock. This figure sheds light on the optimizations required for systems based on AMD GPUS with a single NIC configuration:

- The scaling deteriorates past 2 nodes indicating excessive increase of demand for communication bandwidth capacity as the all-to-all exchanges dominate quickly the completion time.
- The scaling dip may most likely be alleviated by additional network resources such as extra NIC cards, which are present on the Summit system for a higher injection bandwidth of the off-node traffic.
- Vendors such as AMD, Intel and NVIDIA need to further accelerate the inter-node GPU transfers. Efforts such as NCLL [11], would help with this task and must get a considerable attention towards exascale.

4.2 Impact of the libraries setup time

Most applications, such as LAMMPS [12], require several FFTs to compute calculations such as the energy of a system. When the same data structure is going to be used many times, the time for setting up the transform may not be relevant; however, when single or few transformations are needed then it may be relevant to analyze how much time a given library takes for planning the FFT. In Figure 4.2, we show how different setup timing varies amongst libraries, AccFFT being the slowest one almost $3\times$ worse than the fastest one (heFFTe).

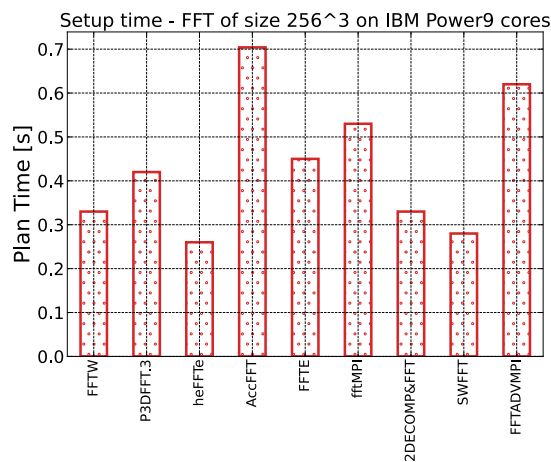


Figure 4.2: Comparison of the average setup (FFT planning) time for all libraries from Figure 3.4 at 4 nodes.

4.3 The FFT communication bottleneck

In the Fig. 4.3, we show a breakdown of the kernel components within a 3-D FFT running on multi-core CPUs for five of the libraries shown in Fig. 3.4. We can clearly observe a similar behavior for all the implementations, and note that optimizations of MPI_Alltoallw could potentially yield to getting much faster FFTs via the approach of the FFT library with advance MPI routines [8], which moves the packing and unpacking kernels into the network. This, however, is at an early stage, since MPI_Alltoallw is far from being optimized in comparison to MPI_Alltoall or MPI_Alltoallv, and in distributions such as SpectrumMPI, it is not even GPU-aware [13].

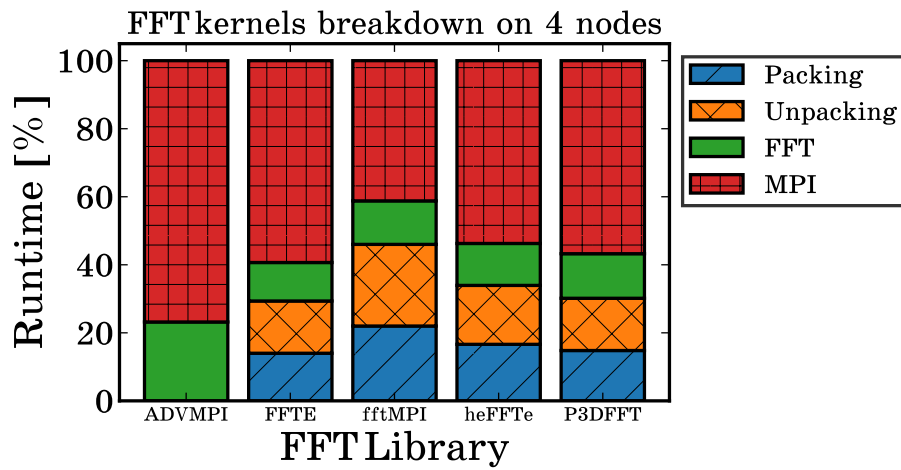


Figure 4.3: Comparison of setup (planning) time for different libraries. Note that FFTADVMPI is only composed of two kernels, since the packing and unpacking is embedded into MPI_Alltoallw.

Conclusions

This report introduced an FFT Benchmark harness that was developed to easily benchmark and compare numerous FFT libraries on DOE exascale systems. The harness is open source [4] and facilitates the addition of other FFT libraries and FFT benchmarking problems from FFT as well as application developers and users.

The FFT benchmark harness was used to benchmark and compare the FFT libraries from [7], by adding the Spock ECP system at ORNL that features AMD MI100 GPUs. Spock is an early-access system and precursor to Frontier, so only 4 nodes were available for now to run.

Results show that all libraries can be easily ported and run on the multicore CPU parts of these two systems, while support for GPUs is provided only in the heFFTe, AccFFT, and FFTE libraries for Nvidia GPUs, and only heFFTe supports AMD GPUs.

Single GPU 3-D FFTs on AMD GPUs need further optimizations. Currently, rocFFT on MI100 is up to $6\times$ slower than cuFFT on V100 GPUs. For 512^3 FFTs cuFFT reaches 12.3% of the V100 GPU compute peak vs. rocFFT reaches 2.1% of the MI100 GPU compute peak. Regardless of that, the 3-D FFTs are communication-bound and therefore the local (single GPU) performance is less important in the distributed computing setting. Indeed, Summit has about $2\times$ faster inter-node communication bandwidth, and our benchmark results quantify that this translates to about $2\times$ faster 3-D FFTs. Frontier is expected to have the same inter-node communication bandwidth as Summit.

Bibliography

- [1] H. Jagode, J. Hein, and A. S. Trew, “Task placement of parallel multi-dimensional FFTs on a mesh communication network,” Computer Science Dept., University of Tennessee Knoxville, Tech. Rep. UT-CS-08-613, 2008, <http://www.cs.utk.edu/library/2008.html>.
- [2] H. Jagode and J. Hein, “Custom assignment of MPI ranks for parallel multi-dimensional FFTs: evaluation of BG/P versus BG/L,” in *IEEE International Symposium on Parallel and Distributed Processing with Applications, ISPA 2008*, Sydney, NSW, Australia, December 10-12 2008, doi:10.1109/ISPA.2008.136.
- [3] H. Jagode, “Fourier transforms for the BlueGene/L communication network,” Master’s thesis, The University of Edinburgh, 2006, <http://www.epcc.ed.ac.uk/msc/dissertations/2005-2006/>.
- [4] “HPFFT: A Parallel FFT Benchmark Harness.” 2022. [Online]. Available: <https://github.com/icl-utk-edu/fiber>
- [5] A. Ayala, S. Tomov, A. Haidar, and J. Dongarra, “heFFTe: Highly Efficient FFT for Exascale,” in *ICCS 2020. Lecture Notes in Computer Science*, 2020.
- [6] M. Frigo and S. G. Johnson, “The design and implementation of FFTW3,” *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216–231, 2005, special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [7] A. Ayala, S. Tomov, P. Luszczek, S. Cayrols, G. Raghianti, and J. Dongarra, “Interim report on benchmarking FFT libraries on high performance systems,” Innovative Computing Laboratory, University of Tennessee, Tech. Rep. ICL-UT-21-03, Jun. 2021. [Online]. Available: <https://www.icl.utk.edu/publications/interim-report-benchmarking-fft-libraries-high-performance-systems>
- [8] L. Dalcin, M. Mortensen, and D. E. Keyes, “Fast parallel multidimensional FFT using advanced MPI,” *Journal of Parallel and Distributed Computing*, vol. 128, pp. 137–150, 2019.
- [9] “cuFFT library,” 2018. [Online]. Available: <http://docs.nvidia.com/cuda/cufft>

- [10] “rocFFT library,” 2021. [Online]. Available: <https://github.com/ROCmSoftwarePlatform/rocFFT>
- [11] NVIDIA, “NCCL library,” 2019. [Online]. Available: <https://github.com/NVIDIA/nvcl>
- [12] “Large-scale atomic/molecular massively parallel simulator,” 2018, available at <https://lammps.sandia.gov/>.
- [13] “Release notes on IBM SpectrumMPI 10.4,” 2021, Available at <https://www.ibm.com/docs/en/smpi/10.4?topic=release-notes>.
- [14] K. Czechowski, C. McClanahan, C. Battaglini, K. Iyer, P.-K. Yeung, and R. Vuduc, “On the communication complexity of 3D FFTs and its implications for exascale,” 06 2012.
- [15] A. Gholami, J. Hill, D. Malhotra, and G. Biros, “Accfft: A library for distributed-memory FFT on CPU and GPU architectures,” *CoRR*, vol. abs/1506.07933, 2015.
- [16] D. Takahashi, “FFTE 7.0: A fast Fourier transform package,” <http://www.ffte.jp/>, 2021.
- [17] “parallel 2d and 3d complex ffts,” 2018, available at <http://www.cs.sandia.gov/~sjplimp/download.html>.
- [18] A. Ayala, S. Tomov, M. Stoyanov, A. Haidar, and J. Dongarra, “Accelerating Multi-Process Communication for Parallel 3-D FFT,” in *2021 Workshop on Exascale MPI (ExaMPI)*, 2021, pp. 46–53.
- [19] A. Ayala, S. Tomov, M. Stoyanov, and J. Dongarra, “Scalability Issues in FFT Computation,” in *Parallel Computing Technologies*. Springer International Publishing, 2021.
- [20] P. Balaji, D. Buntinas, D. Goodell, W. Gropp, S. Kumar, E. Lusk, R. Thakur, and J. L. Träff, “MPI on a Million Processors,” in *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, M. Ropo, J. Westerholm, and J. Dongarra, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 20–30.

Appendix: Tensor transposition cost for parallel FFTs

In this appendix section, we further analyze the communication schemes available in most state-of-the-art FFT libraries. The effect that the parallel tensor transposition bottleneck has on scalability at large-scale has been extensively studied, see e.g., [7, 14].

Given the upcoming exascale ecosystem of software and hardware, we consider relevant to analyze the different MPI routines in which state-of-the-art libraries with GPU support, c.f., Table 6.1, rely to perform tensor transpositions. For the experiments below, we use heFFTe library (due to their wide range of communication options) and Vampir for visualization, c.f., Table 2.1.

Table 6.1: Available MPI routines in FFT libraries

Library	Communication Type ¹	
	<i>AlltoAll</i>	<i>Point-to-Point</i>
AccFFT [15]	MPI_Alltoall	MPI_Isend / MPI_Irecv MPI_Sendrecv
FFTE [16]	MPI_Alltoall MPI_Alltoallv	-
fftMPI [17]	MPI_Alltoallv	MPI_Send / MPI_Irecv
heFFTe [5]	MPI_Alltoall MPI_Alltoallv	MPI_Send / MPI_Isend MPI_Irecv

¹Refer to [18] for a survey of different MPI routines used in modern FFT libraries

6.1 Multi-core CPU based systems

Systems like Fugaku, an exascale supercomputer at the Riken Center for Computational Science in Kobe, Japan, rely on multi-core chips with high speed interconnections. In [19], authors studied FFT scalability issues in such systems. In Figure 6.1, we observe that for a small count of cores (168) the communication cost (tensor transpose) is $\approx 85\%$ even when overlapping the packing and unpacking with the MPI exchange.

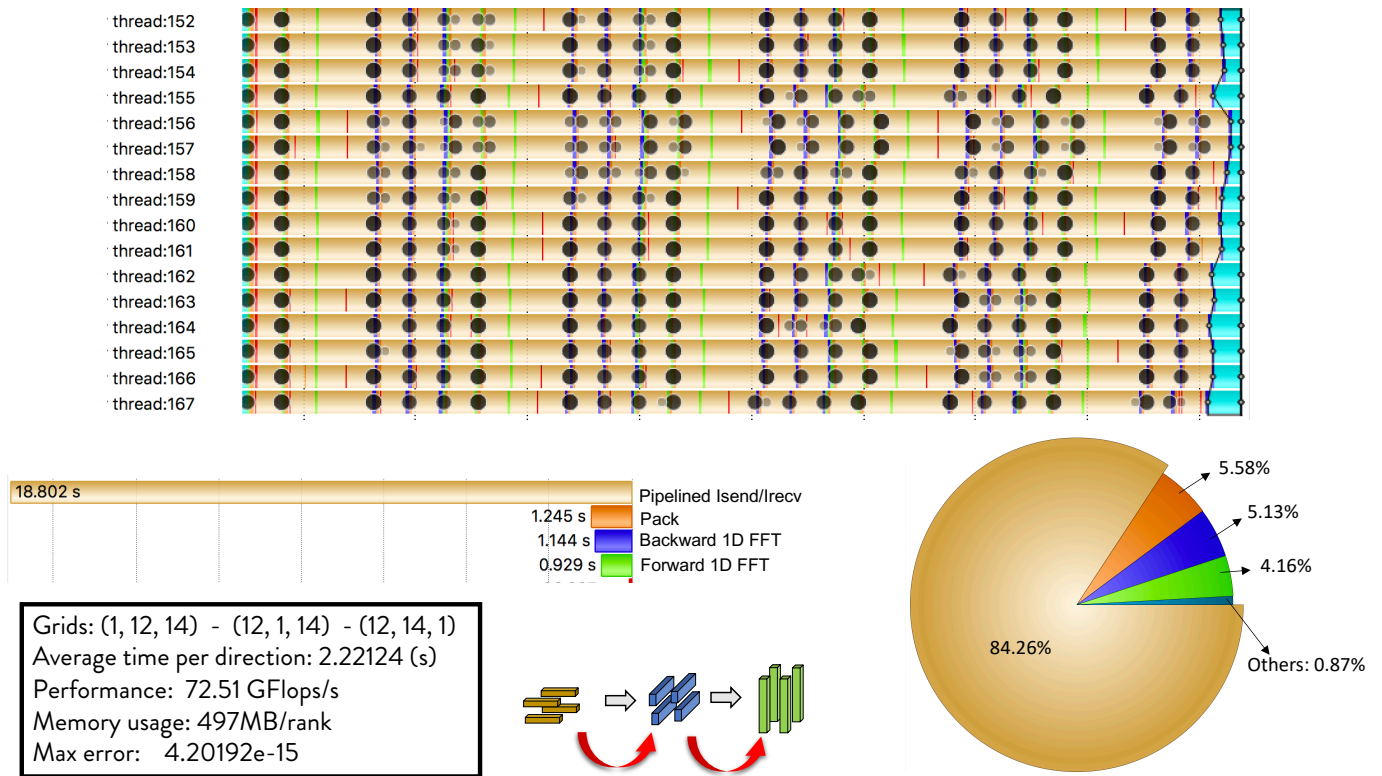


Figure 6.1: Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 168 IBM Power9 cores, 42 MPIs per node. We use heFFTe with FFTW backend and pipelined MPI_Isend and MPI_Irecv communication.

Next, we present an experiment using MPI_Alltoall for the tensor transposition on 168 CPU cores. Libraries that support this type of communication require to have a padding step before calling the MPI routine. In Fig. 6.4, we can observe that the pack and unpack kernels take a considerable amount of time (around 40%). In Fig. 6.4 we will see that when moving those kernels into the GPUs could reduce their cost to under 7%. In Section 4.3 we also commented on how moving these kernels into the network could help accelerate the FFT computation.

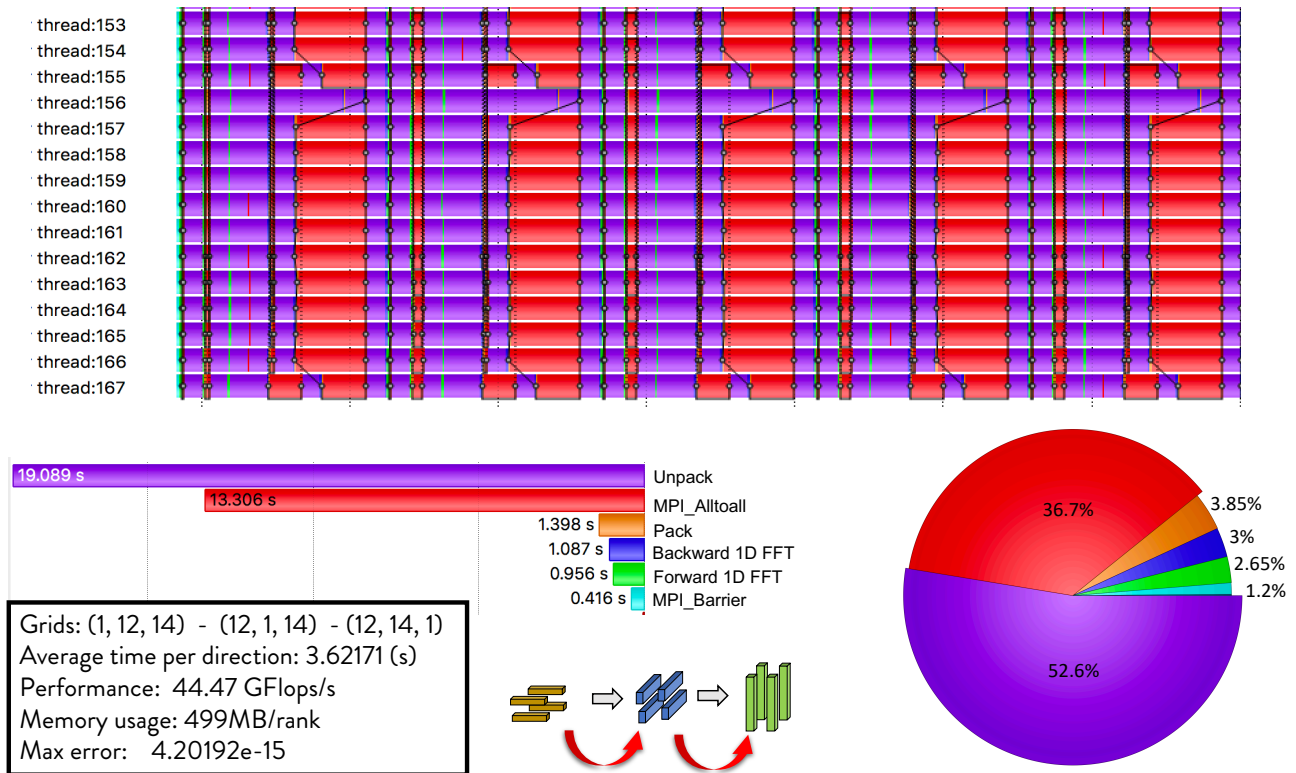


Figure 6.2: Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 168 IBM Power9 cores, 42 MPIs per node. We use heFFTe with FFTW backend and MPI_Alltoall communication.

6.2 GPU based systems

Systems like Summit, and the upcoming Frontier, rely on hybrid CPU-GPU systems. The accelerators have proven to considerable speedup the computation of local kernels, such as the batched 1-D FFTs, packing and unpacking, c.f., [5, 15, 16]. In Figures 6.3 and 6.4 we observe that the non-blocking point-to-point exchange yields to faster computation than the all-to-all approach; however, in both cases MPI covers around 90% of the runtime, and this was achieved only with 16 GPUs, which is just 0.07% of the total number of GPUs available in Summit.

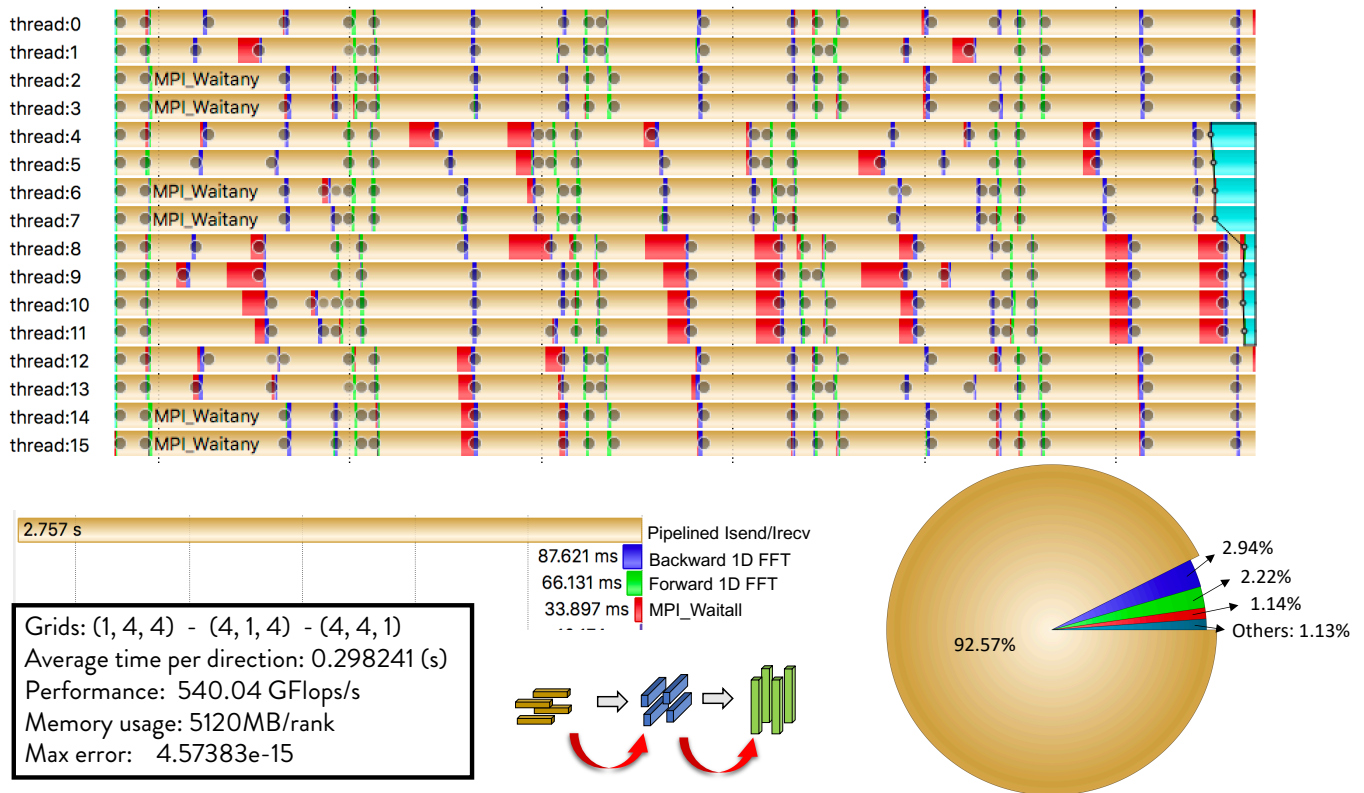


Figure 6.3: Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 16 NVIDIA GPUs, 4 MPIs per node. We use heFFTe with CUFFT backend and pipelined MPI_Isend and MPI_Irecv communication.

Figure 6.4 also shows how the sub-communicators were created for the all-to-all transfers (see the black polygons). All libraries from Table 1.1, set the MPI_groups and communicators during the FFT planning, trying to ensure good load-balancing. In [7], it was shown how tuning grids can ensure linear scaling to over forty thousand processes, and how a bad choice of grids could lead to scalability failure. And even though tuning FFT parameters helps to achieve better performance on several thousand of processes; our experiments also suggest that at very large processes count (e.g., millions), it is the associated latency of MPI_Alltoall(v) what will produce scaling failures, refer to [20] for an external analysis. Therefore, efforts to further optimize all-to-all MPI routines, mainly for small data-volume exchanges, are critical and will have a great impact on the performance of FFT libraries at the exascale level.

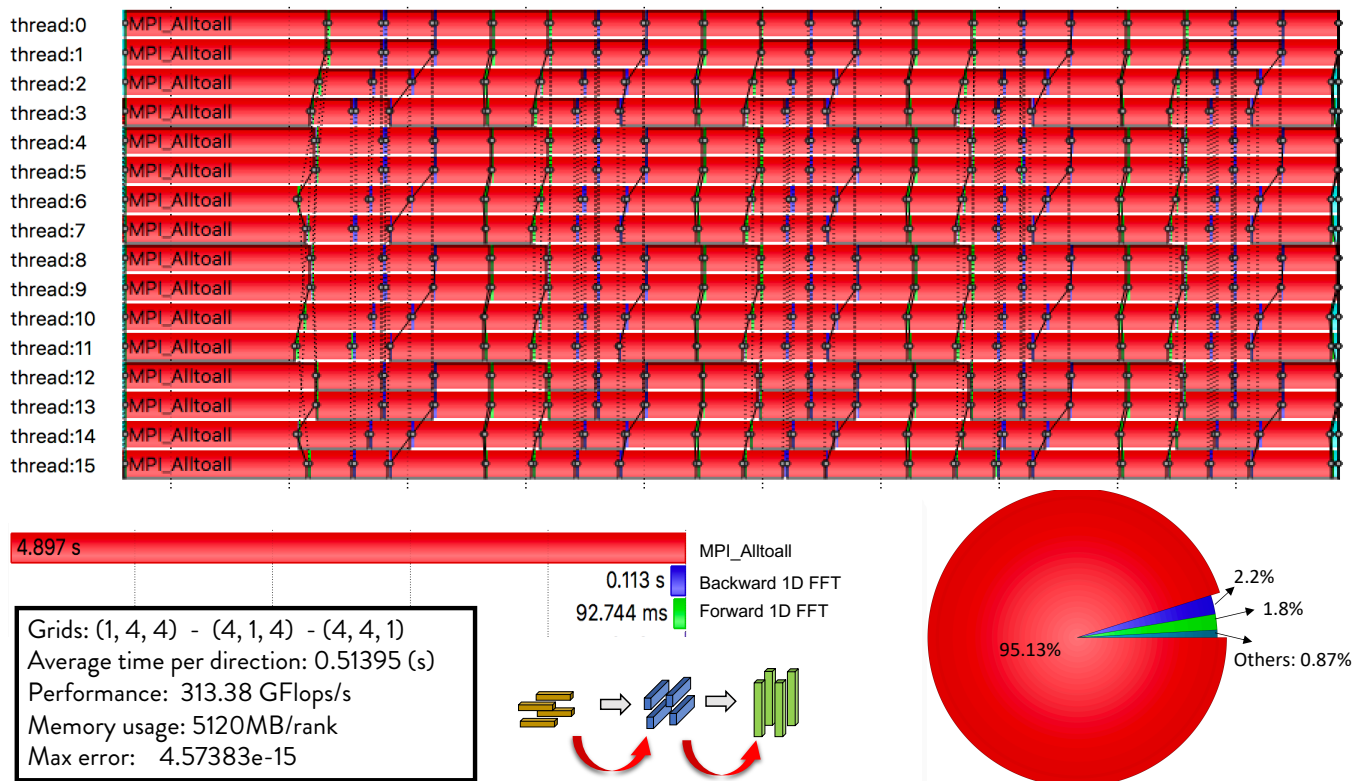


Figure 6.4: Vampir trace of back-to-back 3-D FFTs of size 1024^3 (5 forward + 5 backward), using 4 Summit nodes with 16 NVIDIA GPUs, 4 MPIs per node. We use heFFTe with CUFFT backend and MPI.Alltoall communication.