


Accelerating X-Ray Tracing for Exascale Systems using Kokkos

Felix Wittwer 

*National Energy Research Scientific Computing Center,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
fwittwer@lbl.gov*

Derek Mendez

*Molecular Biophysics and Integrated Bioimaging Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
dermen@lbl.gov*

Aaron S. Brewster 

*Molecular Biophysics and Integrated Bioimaging Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
asbrewster@lbl.gov*

Michael E. Wall 


*Computer, Computational, and Statistical Sciences Division,
Los Alamos National Laboratory,
Mail Stop B256, Los Alamos, NM 87545, USA
mewall@lanl.gov*

Deborah J. Bard 


*National Energy Research Scientific Computing Center,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
djbard@lbl.gov*

Nicholas K. Sauter 

*Molecular Biophysics and Integrated Bioimaging Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
nksauter@lbl.gov*

Billy K. Poon 


*Molecular Biophysics and Integrated Bioimaging Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
bkpoon@lbl.gov*

James M. Holton 

*Molecular Biophysics and Integrated Bioimaging Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
jmholton@lbl.gov*

William E. Hart 

*Sandia National Laboratories,
Albuquerque, NM 87185, USA
wehart@sandia.gov*

Johannes P. Blaschke 

*National Energy Research Scientific Computing Center,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA
jpblaschke@lbl.gov*

Abstract—The upcoming exascale computing systems Frontier and Aurora will draw much of their computing power from GPU accelerators. The hardware for these systems will be provided by AMD and Intel, respectively, each supporting their own GPU programming model. The challenge for applications that harness one of these exascale systems will be to avoid lock-in and to preserve performance portability.

We report here on our results of using Kokkos to accelerate a real-world application on NERSC’s Perlmutter Phase 1 (using NVIDIA A100 accelerators) and the testbed system for OLCF’s Frontier (using AMD MI250X). By porting to Kokkos, we were able to successfully run the same X-ray tracing code on both systems and achieved speed-ups between 13 % and 66 % compared to the original CUDA code. These results are a highly encouraging demonstration of using Kokkos to accelerate production science code.

I. INTRODUCTION

The upcoming HPC systems Frontier and Aurora will be the first exascale computing systems. Both are capable of performing more than 10^{18} floating point operations per second. The majority of this computing power comes from the GPU accelerators of these systems. Due to their massive parallelism, GPUs are well suited for repetitive tasks.

Using GPUs requires vendor specific programming models, such as CUDA for NVIDIA or HIP for AMD. This makes portability between different systems challenging. As an alternative, programming models such as OpenMP offloading or Kokkos provide an abstraction layer between the source code and the GPU hardware [1]. With these abstraction layers, the same code can be compiled for different architectures, thus combining portability with high performance.

II. SCIENTIFIC BACKGROUND

X-ray crystallography is an indispensable tool to study the structure of molecules, with applications ranging from material science [2] to understanding the function of proteins [3]. In crystallography, small crystals of an unknown sample (e.g. a protein whose molecular structure is to be determined) are placed in an x-ray beam. By measuring how the x-rays are scattered, the position of each atom in the sample can be determined. A major application of x-ray crystallography is determining molecular structures of proteins. Knowing the protein structure is crucial to understand how a protein interacts. From this understanding, therapeutic drugs can be designed to, e.g., treat infections by blocking specific interactions.

Because the molecules inside the crystals are arranged in a long-range repeating order, the crystals scatter x-rays only into certain directions – ie. scattering only occurs at those angles, which allow for positive interference from the x-rays scattered by different atoms. X-rays scattered into these discrete directions form Bragg spots on the imaging detector (the bottom panels in Fig. 1 illustrate one possible Bragg spot pattern). Mathematically speaking, a molecular structure is represented using a list of real numbers called structure factors, where each Bragg spot is associated with one such structure factor. Reconstructing the protein structure is therefore a two step process: first, the structure factors are determined from the Bragg spots. Then, the structure factors are used to reconstruct the protein structure. If more structure factors are available for the reconstruction, the atomic positions can be determined more accurately. Therefore, increasing the signal of weak Bragg spots requires larger crystals or a more intense x-ray beam. However, many protein crystals are challenging to crystallize and grow only to microscopic size. Scientists therefore opt for increasing beam intensities. However, as proteins are sensitive to radiation damage, they are quickly destroyed by high intensity x-ray beams.

Serial femtosecond crystallography (SFX) avoids the problem of radiation damage by using ultra-short x-ray pulses from x-ray free-electron lasers (XFELs), called “shots”. In essence, the pulses are so short that the Bragg spots can be measured before beam damage has had time to degrade the sample. Each x-ray pulse lasts only a few ten femtoseconds, fast enough to freeze all atomic motion and capture a snapshot of the crystal before any damage becomes visible. Still, this process will destroy the sample, so fresh crystals must be constantly fed in before each x-ray pulse, see Fig. 1. Measuring a complete dataset requires thousands of crystals, each producing only a single scattering image.

Most crystallography methods determine the structure factors by simply integrating the number of photons in each Bragg spot. Apart from the structure factors, the intensity of a Bragg spot is also influenced by the orientation of the crystal, how the crystal is composed from smaller crystals and how the energy spectrum of the photons fluctuates from shot to shot. By averaging multiple images, each under different conditions, these influences can be averaged out. For SFX, this requires




Fig. 1. *Top*: Experimental setup for SFX. The protein crystals are delivered to the x-ray beam with a water jet. For each x-ray pulse, the imaging detector records the scattered x-rays. *Bottom*: Two detector close-ups. Left: Small scattering angles close to the beam axis. Right: Wide angle scattering at the edge of the detector.

tens of thousands of scattering images.

Instead of simply integrating each Bragg spot, x-ray tracing aims to simulate the photon intensity of every detector pixel by creating a physical model of each crystal [4]. Modeling each crystal is an iterative process, in which the current parameter estimates are refined until the simulated scattering image matches the measured one. By recovering the unknown parameters from each crystal, x-ray tracing can accurately determine the structure factors from an order of magnitude fewer scattering images. As measurement time at XFELs is scarce, this allows scientists to study more samples or more experimental conditions during the same experiment.

However, while x-ray tracing may require fewer experimental resources, the computational effort to accurately model each pixel grows considerably in comparison to conventional methods. X-ray tracing therefore faces the challenge of performing a more complex data analysis in less time. At the same time, quick feedback on the quality of the collected data is critical for XFEL experiments [5]. Providing these experiments with quick-turnaround analysis of the terabytes of data requires high performance computing (HPC) at the Exascale.

III. GPU ACCELERATION WITH KOKKOS

To simulate scattering images from a physical model, we have developed the program *nanoBragg* [6], which is part of the CCTBX software suite [7], [8]. CCTBX consists of a high-level workflow written in Python and accelerated kernels written in C++ (interfacing with Python via *Boost.Python*). This approach has been proven to be highly successful in analyzing large data sets using HPC resources [5], [9]. By designing *nanoBragg* as a kernel compatible with the CCTBX workflow, we see it as the first building block of a larger pixel-level XFEL data analysis workflow.

In x-ray tracing, each detector pixel can be simulated nearly independently of all other pixels, making x-ray tracing ideally suited to exploit the massive parallelism of modern graphic processing units (GPUs). To this end, the original *nanoBragg* code was previously ported to NVIDIA CUDA, resulting in a $500\times$ speedup on an NVIDIA V100 over a multi-threaded CPU implementation on Intel Xeon Ivy Bridge. *nanoBragg* distributes tasks (batches of images to simulate) to multiple nodes and GPUs using MPI (using *mpi4py*). Each GPU then performs work independently. Solving for an unknown molecular structure therefore follows a fork-join parallelism, where each GPU simulates a batch of parameters independently, followed by a MPI reduction [4]. The *nanoBragg* benchmark discussed here forgoes the final reduction step, instead saving the simulated images to the file system.

The upcoming exascale systems Frontier (OLCF) and Aurora (ALCF) are essential to obtain rapid feedback during XFEL experiments. However, the GPUs for these systems will be provided by AMD, and Intel, respectively. As both vendors supply their own alternative to CUDA, *nanoBragg* must be adapted to each system. To avoid code duplication and site-specific optimizations, we decided to use Kokkos [1].

A. Kokkos

Kokkos is a C++ programming model that allows a single code base to target different HPC platforms [1]. To achieve this, Kokkos implements abstract memory and execution spaces. Only during the compilation are the abstract spaces and commands converted to CUDA, HIP, OpenMP, etc. This way, the same Kokkos code can be used on systems with different GPUs or even systems with no GPU.

Each execution space is associated with a particular memory space. While the execution space defines where the computation is done, the computation itself is defined by *execution patterns*, which are the Kokkos equivalent to kernels in CUDA. Kokkos implements three patterns:

- `parallel_for`: Corresponds to a for-loop where each iteration is executed independently, for example element-wise addition of two vectors.
- `parallel_reduce`: A reduction that collects and combines the result of all iterations, for example finding the longest word in a list.
- `parallel_scan`: A combination of the former two that uses multiple reductions, for example calculating a histogram of an image.

The typical use case for execution patterns is to work on large data arrays. Kokkos provides its own data structure called *View*. A *View* is a multi-dimensional array which can be transferred between different memory spaces and layouts. The last one is important, e.g. in matrix multiplications, the best performance on a multi-core CPU might be achieved with a row-major layout and on a GPU with a column-major layout. Kokkos automatically manages the conversion between layouts when a *View* is transferred between different memory spaces. This makes it easy to first initialize a *View* on the CPU and then transfer it to GPU memory.

B. Porting nanoBragg to Kokkos

The task of porting *nanoBragg* can be split into two parts: The first part is converting the CUDA kernels to Kokkos' execution patterns, the second part is replacing all CUDA arrays with Kokkos' *Views*. As there is no interaction between the individual detector pixels, all kernels can be replaced with the `parallel_for` pattern. Most of the computational work in *nanoBragg* is done in three Kernels:

- *nanoBraggSpots*: The most complex kernel, containing about 350 lines of code. This kernel simulates the Bragg spots. It takes into account crystal orientation, mosaicity and photon energy.
- *addBackground*: The second most complex kernel with about 120 lines of code. This kernel simulates the background scattering from the air and water that surround the crystal.
- *addArray*: A simple kernel that adds the results from the other two kernels.

We illustrate the porting process by using a simplified version of the *addArray* kernel:

```
1 __global__
2 void addArray(double* lhs,
3               float* rhs,
4               int total_pixels) {
5     int j = blockDim.x * blockIdx.x + threadIdx.x;
6     if (j < total_pixels) {
7         lhs[j] = lhs[j] + (double) rhs[j];
8     }
9 }
```

Listing 1. The *addArray* kernel in CUDA

For the Kokkos version, we replaced the *lhs* and *rhs* arrays with *Views* and the body of the kernel with a `parallel_for` pattern. As the memory management of *Views* is done by Kokkos, we could remove the custom-written CUDA memory management. The `parallel_for` execution pattern takes three arguments: an individual name for debugging and profiling purposes, an execution policy and a functor that holds the computational work. In simple cases like this, the execution policy is simply to iterate over all *total_pixels*. The functor can be integrated into the pattern by using lambda expressions, this way the structure mirrors how kernels are defined:

```

1 void addArray(Kokkos::View<double*> lhs,
2              Kokkos::View<float*> rhs,
3              int total_pixels) {
4   Kokkos::parallel_for(
5     "addArray", total_pixels,
6     KOKKOS_LAMBDA (const int& j) {
7       lhs(j) = lhs(j) + (double) rhs(j);
8     }
9   );
10 }

```

Listing 2. The *addArray* kernel in Kokkos

The other kernels were converted much in the same way, with a few challenges. One of these is an edge case when using lambda expressions for device code in member functions of a class. Suppose *init()* is a member function that initializes the view data with a certain value:

```

1 void Container::init() {
2   Kokkos::parallel_for("init", size,
3   KOKKOS_LAMBDA (const int& j) {
4     data(j) = m_value;
5   }
6   );
7 }

```

Listing 3. Lambda capture

Since they are class members, the compiler replaces *data* and *m_value* during the compilation with *this->data* and *this->m_value*. The problem arises from *this* being a pointer in CPU memory. If CUDA is used, this creates an illegal memory access when the GPU tries to access the pointer.

There are multiple ways around this issue: One way is to avoid lambda expressions in this situation and instead explicitly define the functors. Another option is to copy the specific member variables into local variables before the execution pattern is called. And finally, in C++17 lambda expressions have been extended so they can explicitly capture the object, not the pointer. Unfortunately, the CCTBX code is currently not compatible with C++17. Therefore, we choose to use local variables in these situations and otherwise minimize the use of kernels in member functions.

IV. PERFORMANCE AND PORTABILITY

The goal of porting the *nanoBragg* code to Kokkos is to achieve portability between the upcoming Frontier and Aurora systems as well as Perlmutter. We test the portability on Perlmutter Phase 1 and Crusher, the latest Frontier testbed system. These systems are not production resources, the performance numbers given here were not always reproducible and will be different from the final systems.

Increasing the portability of a code can potentially reduce the performance. We therefore also assess on Perlmutter the performance of the Kokkos port in comparison to the original CUDA version. On Perlmutter, we use the CUDA backend of Kokkos and on Crusher accordingly the HIP backend.

Each node of Perlmutter Phase 1 is equipped with an AMD EPYC 7763 CPU and four NVIDIA A100 GPUs. On Crusher each node is equipped with an AMD EPYC 7A53 CPU and




Fig. 2. Strong scaling of *nanoBragg* for the simulation of 100 000 scattering images. The dashed line shows ideal scaling where doubling the number of nodes halves the simulation time. On Perlmutter Phase 1, Kokkos is used with the CUDA backend. On Crusher, Kokkos is used with the HIP backend.

four AMD MI250X GPUs. Each MI250X GPU is equipped with two graphic compute dies (GCDs) totalling eight GCDs per node. These eight GCDs function effectively like eight separate GPUs per node.

As a benchmark, we used a test scenario that simulates 100 000 scattering patterns. Due to performance fluctuations of the file system (a consequence of Perlmutter and Crusher not being production systems), we did not include the saving of the simulated images into the benchmark. The tests were performed using 32, 64, and 128 nodes, all results are plotted in Fig. 2. Every test case shows almost ideal scaling behavior, highlighting how x-ray tracing benefits from parallel computing. On Perlmutter, the Kokkos version is about 13% faster than the original CUDA version, indicating no loss in performance due to the port to Kokkos. Concerning portability, Kokkos allows *nanoBragg* to exploit the hardware of Crusher for a more than 60% faster performance per node, compared to Perlmutter Phase 1.

Next we re-enable file I/O. This introduces workflow latency and dependence on a shared resource (the file system). We observe that the strong scaling in shown in Fig. 2 is preserved – albeit shifted up due to the increased time spent on I/O. CCTBX seeks to hide workflow latency by assigning more than one MPI rank to each GPU. We observe that this is a largely successful strategy (as total runtime increases with increasing number of MPI ranks per GPU). Fig. 3 shows the total runtime to trace 100 thousand images using 64 Perlmutter nodes. We see that increasing the number of ranks per GPU increases the total throughput. Interestingly, the CUDA implementation achieves the highest throughput when many (approx. 7 or greater) ranks are sharing the same GPU, whereas the Kokkos implementation achieves the maximum throughput when two MPI ranks share the same GPU. Furthermore, comparing Fig. 2 and Fig. 3, we see that the increased latency due to I/O can be hidden by sharing GPUs across ranks.




Fig. 3. Time to simulate 100 thousand diffraction patterns using 64 Perlmutter nodes. In contrast to Fig. 2, this benchmark includes the time to save each image to the file system. For comparison, the previous times without file saving are indicated with dashed lines. As the filesystem performance on Perlmutter can vary (due to it not being a production system) at the time of writing, this graph shows the best performance over several repeated runs.

TABLE I
KERNEL RUN-TIMES.

	nanoBraggSpots	addBackground	addArray
CUDA	8.28 ms	1.87 ms	0.13 ms
Kokkos	6.98 ms	1.76 ms	0.12 ms
Speed-up	+15.7%	+5.9%	+7.7%

V. KERNEL PROFILING

To analyze how the Kokkos version performs better on Perlmutter than the original CUDA version, we used Nsight Systems and Nsight Compute to profile both versions. For the three main kernels, the execution times on a NVIDIA A100 GPU are given in Table I. For all three kernels, Kokkos achieves a speed-up of more than five percent. As the *nanoBraggSpots* kernel accounts for most of the GPU computation time, the speed-up of more than 15% for this kernel has a significant impact on the overall run time. On the other end of the complexity spectrum is the *addArray* kernel. While it is one of the smallest kernels and any improvements will only have a small influence in the grand scheme, achieving a speed-up of nearly eight percent even for short kernels is remarkable.

Using Nsight Compute we studied the critical *nanoBragg-*

TABLE II
NANOBRAAGGSPOTS DETAILS.

	CUDA	Kokkos
Run-time	8.28 ms	6.98 ms
Compute Throughput	65.05 %	77.42 %
Memory Throughput	21.05 %	21.35 %
Registers	130	116
Theoretical Occupancy	18.75 %	25 %
Achieved Occupancy	16.8 %	24.74 %

Spots kernel in detail. The results are summarized in Table II. In the table, the compute and memory throughput give an overview how much of the compute and memory resources of the GPU are utilized by the *nanoBraggSpots* kernel. The performance difference between CUDA and Kokkos is mostly a result of the higher compute throughput of the Kokkos kernel, the memory throughput is nearly identical. The increased compute throughput in Kokkos is made possible by using fewer registers. Each multiprocessor on an A100 GPU has a total of 65 536 32 bit-registers for a maximum of 2048 simultaneous threads (64 warps per multiprocessor times 32 threads per warp). However, the GPU can only run at full occupancy if each thread uses less than $65536/2048 = 32$ registers. With 130 registers, the GPU can run at most 12 warps out of the 64 available. The Kokkos kernel uses 116 registers, which is below the critical threshold of 128 registers, and can therefore run a maximum of 16 warps, 33% more.

The situation for the other kernels is similar. The *addBackground* kernel goes from 96 to 80 registers and the *addArray* kernel goes from 25 to 16 registers. Kokkos uses consistently fewer registers than the original CUDA implementation, allowing a higher occupancy and throughput.

VI. CONCLUSION

In this paper, we have reported on our work to port *nanoBragg* to Kokkos. Starting from a CUDA code base, most of the port was straightforward. Arising difficulties from the port could be solved. We have demonstrated that the performance of the code did not suffer from this port and has even increased by 13%. We have also successfully demonstrated performance portability between Perlmutter Phase 1 (NERSC) and Crusher (OLCF). We are currently in the process of testing the Aurora test system (ALCF). Until now, we have not used advanced Kokkos features such as nested parallelism, which should further increase the performance.

As *nanoBragg* is representative of many scientific codes, our reported findings indicate that Kokkos is a powerful tool to not only achieve performance portability in real-world applications, but also to accelerate existing codes. Moreover, we are able to integrate Kokkos into existing scientific workflows without needing to port an entire code base. This allows scientists to make progress, even if staff time is limited.

VII. ACKNOWLEDGMENTS

N.K.S, J.P.B., F.W. and D.B. acknowledge support from the Exascale Computing Project (grant 17-SC20-SC), a collaborative effort of the Department of Energy (DOE) Office of Science and the National Nuclear Security Administration. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of

REFERENCES

- [1] C. R. Trott, D. Lebrun-Grandie, D. Arndt, J. Ciesko, V. Dang, N. Ellingwood, R. Gayatri, E. Harvey, D. S. Hollman, D. Ibanez, N. Liber, J. Madsen, J. Miles, D. Poliakoff, A. Powell, S. Rajamanickam, M. Simberg, D. Sunderland, B. Turcksin, and J. Wilke, "Kokkos 3: Programming model extensions for the exascale era," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, pp. 805–817, Apr. 2022.
- [2] E. A. Schriber, D. W. Paley, R. Bolotovskiy, D. J. Rosenberg, R. G. Sierra, A. Aquila, D. Mendez, F. Poitevin, J. P. Blaschke, A. Bhowmick, R. P. Kelly, M. Hunter, B. Hayes, D. C. Popple, M. Yeung, C. Pareja-Rivera, S. Lisova, K. Tono, M. Sugahara, S. Owada, T. Kuykendall, K. Yao, P. J. Schuck, D. Solis-Ibarra, N. K. Sauter, A. S. Brewster, and J. N. Hohman, "Chemical crystallography by serial femtosecond X-ray diffraction," *Nature*, vol. 601, no. 7893, pp. 360–365, 2022.
- [3] S. M. Keable, A. Kölsch, P. S. Simon, M. Dasgupta, R. Chatterjee, S. K. Subramanian, R. Hussein, M. Ibrahim, I.-S. Kim, I. Bogacz, H. Makita, C. C. Pham, F. D. Fuller, S. Gul, D. Paley, L. Lassalle, K. D. Sutherland, A. Bhowmick, N. W. Moriarty, I. D. Young, J. P. Blaschke, C. d. Lichtenberg, P. Chernev, M. H. Cheah, S. Park, G. Park, J. Kim, S. J. Lee, J. Park, K. Tono, S. Owada, M. S. Hunter, A. Batyuk, R. Oggenfuss, M. Sander, S. Zerdane, D. Ozerov, K. Nass, H. Lemke, R. Mankowsky, A. S. Brewster, J. Messinger, N. K. Sauter, V. K. Yachandra, J. Yano, A. Zouni, and J. Kern, "Room temperature XFEL crystallography reveals asymmetry in the vicinity of the two phyloquinones in photosystem I," *Scientific Reports*, vol. 11, no. 1, p. 21787, 2021.
- [4] D. Mendez, R. Bolotovskiy, A. Bhowmick, A. S. Brewster, J. Kern, J. Yano, J. M. Holton, and N. K. Sauter, "Beyond integration: Modeling every pixel to obtain better structure factors from stills," *IUCrJ*, vol. 7, pp. 1151–1167, Nov. 2020.
- [5] J. P. Blaschke, A. S. Brewster, D. W. Paley, D. Mendez, A. Bhowmick, N. K. Sauter, W. Kröger, M. Shankar, B. Enders, and D. Bard, "Real-time XFEL data analysis at SLAC and NERSC: A trial run of nascent exascale experimental data analysis," in *CUG2021 Proceedings*, p. 10, May 2021.
- [6] N. K. Sauter, J. Kern, J. Yano, and J. M. Holton, "Towards the spatial resolution of metalloprotein charge states by detailed modeling of XFEL crystallographic diffraction," *Acta Crystallogr. D*, vol. 76, pp. 176–192, Feb. 2020.
- [7] R. W. Grosse-Kunstleve, N. K. Sauter, N. W. Moriarty, and P. D. Adams, "The Computational Crystallography Toolbox: Crystallographic algorithms in a reusable software framework," *Journal of Applied Crystallography*, vol. 35, no. 1, pp. 126–136, 2002.
- [8] N. K. Sauter, J. Hattne, R. W. Grosse-Kunstleve, and N. Echols, "New Python-based methods for data processing," *Acta Crystallographica Section D*, vol. 69, pp. 1274–1282, Jul 2013.
- [9] A. Giannakou, J. P. Blaschke, D. Bard, and L. Ramakrishnan, "Experiences with cross-facility real-time light source data analysis workflows," in *2021 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC)*, pp. 45–53, 2021.