

Lawrence Berkeley National Laboratory

LBL Publications

Title

AMReX and pyAMReX: Looking Beyond ECP

Permalink

<https://escholarship.org/uc/item/05h6j8zz>

Authors

Myers, Andrew

Zhang, Weiqun

Almgren, Ann

et al.

Publication Date

2024-03-18

AMReX and pyAMReX: Looking Beyond ECP

Journal Title
XX(X):1–12
©The Author(s) 2024
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Andrew Myers¹, Weiqun Zhang¹, Ann Almgren¹, Thierry Antoun^{1,2}, John Bell¹, Axel Huebl¹, Alexander Sinn³

Abstract

AMReX is a software framework for the development of block-structured mesh applications with adaptive mesh refinement (AMR). AMReX was initially developed and supported by the AMReX Co-Design Center as part of the U.S. DOE Exascale Computing Project, and is continuing to grow post-ECP. In addition to adding new functionality and performance improvements to the core AMReX framework, we have also developed a Python binding, pyAMReX, that provides a bridge between AMReX-based application codes and the data science ecosystem. pyAMReX provides zero-copy application GPU data access for AI/ML, in situ analysis and application coupling, and enables rapid, massively parallel prototyping. In this paper we review the overall functionality of AMReX and pyAMReX, focusing on new developments, new functionality, and optimizations of key operations. We also summarize capabilities of ECP projects that used AMReX and provide an overview of new, non-ECP applications.

Keywords

Adaptive mesh refinement; structured mesh/grids; particles; co-design; performance portability

Introduction

The AMReX software framework was initially developed as part of the U.S. DOE Exascale Computing Project (ECP) to support the development of block-structured adaptive mesh refinement (AMR) algorithms for multiphysics applications described by partial differential equations in simple or complex geometries (Zhang et al. 2021). Block-structured AMR uses a hierarchical representation of the solution at multiple levels of resolution where the solution on each level is defined on the union of data containers at that resolution. These data containers, which represent the solution over a logically rectangular subregion of the domain, can contain field data defined on a mesh, Lagrangian particles or combinations of both.

AMR reduces the computational cost and memory footprint compared to a uniform mesh while preserving accurate descriptions of different physical processes in complex multiphysics algorithms. AMReX builds on decades of development of the fundamental algorithmic underpinnings of adaptive mesh refinement, going back to work by Berger and Olinger (1984), Berger and Colella (1989) and Almgren et al. (1998). There are also a number of different open-source AMR frameworks; see the 2014 survey paper by Dubey et al. (2014) for a discussion of these frameworks and references to both the frameworks and application codes that use them.

The design of AMReX was informed by two major factors. First, AMReX should be able to support a wide range of multiphysics applications with different performance characteristics without imposing restrictions on how application developers construct their algorithms. Careful consideration was given to separating the design of the data structures and basic operations from the algorithms

that use those data structures. AMReX uses a layered design that provides a rich set of flexible tools that can be used by a wide range of applications. This layered design allows application developers to interact with the software at several different levels of abstraction. Applications can simply use the AMReX data containers and iterators and none of the higher-level functionality, both in stand-alone AMReX-based codes and in non-AMReX based codes that selectively use AMReX functionality. At an intermediate level of functionality, applications can use the data structures and iterators for single- and multi-level operations but retain complete control over the time evolution algorithm, i.e., the ordering of algorithmic components at each level and across levels. AMReX also provides developers the option to exploit additional functionality in AMReX that is designed specifically for AMR algorithms that subcycle in time. This layer provides stubs for the necessary operations such as advancing the solution on a level, correcting coarse grid fluxes with time- and space-averaged fine grid fluxes, averaging data from fine to coarse and interpolating in both space and time from coarse to fine.

The other major factor influencing the AMReX design was performance portability. Current HPC architectures typically include some type of GPU accelerator. Achieving high performance on these architectures requires efficient

¹Lawrence Berkeley National Laboratory, Berkeley, CA, USA

²Now with: CEA, Paris, France

³Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

Corresponding author:

Andrew Myers, MS 50A-3111, Lawrence Berkeley, National Laboratory, Berkeley, CA 94720.

Email: atmyers@lbl.gov

utilization of the accelerator. However, different accelerators have different capabilities and programming models. To isolate applications from a particular architecture and programming model without sacrificing performance, AMReX introduced a lightweight abstraction layer that effectively hides the details of the architecture from the application. This layer provides constructs that allow the user to specify what operations they want to perform on a block of data without specifying how those operations are carried out. AMReX then maps those operations onto the target hardware at compile time so that the hardware is utilized effectively. AMReX currently supports CUDA, HIP and SYCL for GPU acceleration and OpenMP for multi-core CPU architectures. AMReX-based applications have run successfully at scale on some of the largest supercomputers in the world, each of which utilizes quite different hardware: OLCF's AMD MI250X-based Frontier, NERSC's NVIDIA A100 machine Perlmutter, ALCF's Aurora, which uses Intel Xe GPUs, and Riken's Fugaku platform, which uses ARM A64FX CPUs, to name a few.

AMReX was successfully used by six of the ECP applications to meet their performance / capability goals. These applications represented a wide range of disciplines and computational paradigms. The particle-in-cell code WarpX (Vay et al. 2018; Fedeli et al. 2022) is used for simulations of particle accelerators, beams and laser-plasmas. The astrophysics code, Castro (Almgren et al. 2010; Almgren et al. 2020), uses high-fidelity explicit algorithms for compressible flow in combination with self-gravity and nuclear reaction networks. The cosmology code, Nyx (Almgren et al. 2013), models compressible flow in an expanding universe with self-gravity and a Lagrangian particle representation of dark matter. The MFIX-Exa code (Musser et al. 2022; Porcu et al. 2023) models multiphase particle-laden flows with reactions and heat transfer effects in complex geometries. The combustion codes, PeleC (de Frahan et al. 2023), based on a fully compressible formulation, and PeleLMex (Esclapez et al. 2023), based on a low Mach number formulation, model reacting flows with detailed kinetics and transport in complex geometries. The AMR-Wind code solves incompressible flow with additional physics to model atmospheric boundary layers, with coupling to an unstructured flow solver to model flow around turbine blades (Sharma et al. 2024).

In addition to the ECP applications, AMReX is also being used in applications ranging from beam dynamics (ImpactX: (Huebl et al. 2022a)) and plasma wakefield acceleration (HiPACE++: (Diederichs et al. 2022)) to geophysical fluid dynamics (ERF (Almgren et al. 2023) for atmospheres and REMORA for oceans) to epidemiology (ExaEpi), to name a few. These new applications, in combination with performance improvements and the addition of new features for existing applications, are driving the addition of new capabilities to AMReX.

In this paper, we focus on capabilities that have been developed in the past four years. We first provide an brief overview of AMReX. We then discuss a number of enhancements to AMReX aimed at enabling new application codes and improving performance, particularly on GPU-accelerated architectures. Next, we discuss the python binding, pyAMReX, that provides access to AMReX data

structures and operations from Python. Finally, we present some conclusions and discuss future directions.

AMReX Overview

In this section, we provide an overview of the basic capabilities of AMReX; we refer the interested reader to Zhang et al. (2021) for details.

AMReX provides a set of data containers that can be used to store structured mesh and particle data on distributed memory systems. Although the focus is on supporting block-structured AMR codes, AMReX's flexible data structures and functions for operating on these data are also useful for building non-AMR applications.

The multi-dimensional array is one of the most fundamental data structures used by many scientific computing applications. AMReX is written in C++, which does not have native multi-dimensional array support. `amrex::FArrayBox` is an AMReX class for storing multi-dimensional mesh data, with a Fortran-like syntax for accessing the data. It holds the data for a logically rectangular spatial region with its integer index bounds specified by `amrex::Box`. The framework supports data that is cell-centered, face-centered, edge-centered or nodal.

In AMReX, the spatial index is global, thus the lower bounds of the array in `FArrayBox` do not need to be zero. In fact, they can be negative, which is useful for applications solving partial differential equations. For block-structured AMR, the computational domain consists of multiple logically rectangular subdomains that are distributed among MPI processes. `amrex::MultiFab` is a container for a collection of `FArrayBox` objects that are locally owned. The `amrex::MultiFab` also contains metadata for communication between MPI processes. The `amrex::MFIter` class iterates over a `amrex::MultiFab` object in parallel, with each MPI process iterating over the `FArrayBox` objects it owns. `amrex::MFIter` also supports a logical tiling option for cache-blocking and multi-threaded execution with OpenMP.

AMReX also provides a `amrex::ParIter` class for iterating over distributed particle data. Particles in AMReX are associated with a grid and a level. (We note that even in the absence of mesh data we do construct a mesh hierarchy to organize the particles.) In general, particle data can be stored in Array-of-Struct (AoS) style, Struct-of-Array (SoA) style, or in some combination of the two. The original implementation of particles in AMReX used AoS only. Later, support for optional particle components stored in SoA style was added, but the limitation that the required particle data (positions and id numbers) be stored in AoS-style remained.

In addition to the basic data types, AMReX supports a multilevel embedded boundary representation of complex geometry; linear solvers for cell-centered and nodal data; asynchronous I/O in a native format readable by ParaView, VisIt and yt; mesh pruning; interfaces to hypre and PETSc solvers; and level set functionality for particle-wall interactions. In addition, the capability to use the bintree (Dhruv 2024) algorithm for regridding has been added recently.

As noted above, AMReX uses MPI for communication between processes. Common communication patterns such

as ghost cell exchange, copying data from one distribution to another, and communication between AMR levels (including interpolation, restriction, etc.) are created automatically and cached for performance. AMReX also provide MPI functionality for common global reduction operations across MPI ranks. AMReX aggregates messages into communication buffers to reduce latency.

New features and optimizations

In this section, we discuss new capabilities added to AMReX to improve application performance and to support new application codes. We particularly focus on features added to support and improve the performance of post-ECP applications, including ImpactX, ERF, ExaEpi and REMORA, and on supporting the Python interfaces described below.

Performance Portability

Many modern supercomputers are hybrid computing systems built with both CPUs and GPUs, with the latter providing most of the computing power. There are several different GPU vendors and their GPUs have different architectures. Programming for these heterogeneous systems is challenging because there is no standard programming model for all of these different GPU architectures. To facilitate heterogeneous programming for AMReX applications and for our own development of AMReX, we have created a performance portability abstraction layer, which is built upon the native solutions provided by the vendors. That is, we use CUDA, HIP and SYCL for NVIDIA, AMD and Intel GPUs, respectively. AMReX’s portability layer is similar to those provided by Kokkos (Edwards et al. 2014) or RAJA (Beckingsale et al. 2019) but tailored to the needs of block-structured AMR applications.

AMReX provides a construct, `amrex::ParallelFor`, for launching a GPU kernel looping over a one- or multi-dimensional iteration space. Listing 1 shows an example of using `ParallelFor` on three `MultiFab` objects. The example code runs on CPUs and all major GPUs. On multicore CPU architectures, we decompose the loops over the individual `FArrayBox`’s into logical tiles to enable OpenMP parallelism and improve cache performance. The tile size is specified at runtime in the inputs file and can also be set to different values for different loops.

Here, the data type of `a`, `b` and `c` is `amrex::Array4`, which is a lightweight alias to `FArrayBox`. Unlike `FArrayBox`, this non-owning data structure can be used in GPU kernels. The separation of data ownership and access is a very important design pattern for GPU programming. For these types of relatively simple kernels, high efficiency can be achieved. For example, the triad kernel in this example is memory bandwidth limited and is able to reach more than 80% of the peak memory bandwidth on NVIDIA A100 GPUs, despite the cost associated with multi-dimensional arrays such as the conversion between 1D and 3D indices.

Kernel Fusing

AMR applications can have many relatively small patches when the regions identified as requiring higher resolution

```

1 // triad on MultiFabs: mfa = mfb + scalar*mfc
2 #pragma omp parallel if (Gpu::notInLaunchRegion())
3 for (MFIter mfi(mf, TilingIfNotGPU());
4      mfi.isValid(); ++mfi)
5 {
6     auto const& a = mfa.array(mfi);
7     auto const& b = mfb.const_array(mfi);
8     auto const& c = mfc.const_array(mfi);
9     ParallelFor(mfi.tilebox(),
10              [=] AMREX_GPU_DEVICE (int i, int j, int k) {
11                 a(i, j, k) = b(i, j, k) + scalar * c(i, j, k);
12             });
13 }

```

Listing 1: Example of `ParallelFor`. This code can be compiled to run on CPU with OpenMP or GPU with CUDA, HIP, or SYCL.

are small and cannot be merged to form larger patches. To launch a GPU kernel for each of the small patches would incur a significant cost due to GPU kernel launch latency. For example, a simple kernel like the triad in Listing 1 running on an AMD MI250X will only be able to achieve $\sim 10\%$ of its peak memory bandwidth on small boxes of 32^3 cells. To alleviate the latency issue of launching multiple small kernels, AMReX provides a kernel fusion capability. Listing 2 shows an example of rewriting the code in Listing 1 to use a single kernel launch for all the local data in the `MultiFab`.

```

1 // triad on MultiFabs: mfa = mfb + scalar*mfc
2 auto const& a = mfa.arrays;
3 auto const& b = mfb.const_arrays;
4 auto const& c = mfc.const_arrays;
5 ParallelFor(mfa,
6            [=] AMREX_GPU_DEVICE (int box,
7                                  int i, int j, int k) {
8                a[box](i, j, k) = b[box](i, j, k)
9                    + scalar * c[box](i, j, k);
10            });

```

Listing 2: Example of kernel fusion. Only one GPU kernel is launched even when there are multiple patches.

In this example, if there are 512 patches of 32^3 cells each, only one GPU kernel is launched to work on all 512 patches, which enables it to achieve similar performance as if operating on a single patch of 256^3 cells.

Compile-Time Kernel Specialization

The development of AMReX is often driven by the needs of AMReX-based applications. A number of these applications often use runtime parameters in GPU kernels such as shown in Listing 3, which shows an example with six possible options for execution with branches based on runtime conditions.

Because the runtime options are the same for all threads, thread divergence is not an issue. However, different branches may place drastically different demands on computing resources such as registers. The compiler could not know at compile time which branch will be taken. Even if the cheapest branch is taken at run time, the

```

1 int A_runtime_option = ...;
2 int B_runtime_option = ...;
3 enum A_options : int { A0, A1, A2, A3};
4 enum B_options : int { B0, B1 };
5 ParallelFor(N, [=] AMREX_GPU_DEVICE (int i)
6 {
7     // ...
8     if (A_runtime_option == A0) {
9         // ...
10    } else if (A_runtime_option == A1) {
11        // ...
12    } else if (A_runtime_option == A2) {
13        // ...
14    } else {
15        // ...
16    }
17    if (A_runtime_option != A3 &&
18        B_runtime_option == B1) {
19        // ...
20    }
21    // ...
22 });

```

Listing 3: Example of GPU kernel with runtime options.

performance could be severely affected by the assumption that the most expensive branch might be taken. To address this performance issue, AMReX provides a special version of the `ParallelFor` kernel-launching construct that is optimized for this type of case.

```

1 int A_runtime_option = ...;
2 int B_runtime_option = ...;
3 enum A_options : int { A0, A1, A2, A3};
4 enum B_options : int { B0, B1 };
5 ParallelFor(
6     TypeList<CompileTimeOptions<A0,A1,A2,A3>,
7         CompileTimeOptions<B0,B1>
8         >{},
9     {A_runtime_option, B_runtime_option},
10    N,
11    [=] AMREX_GPU_DEVICE (int i,
12        auto A_control,
13        auto B_control) {
14        // ...
15        if constexpr (A_control.value == A0) {
16            // ...
17        } else if constexpr (A_control.value == A1) {
18            // ...
19        } else if constexpr (A_control.value == A2) {
20            // ...
21        } else {
22            // ...
23        }
24        if constexpr (A_control.value != A3 &&
25            B_control.value == B1) {
26            // ...
27        }
28        // ...
29 });

```

Listing 4: Example of the optimized version of a GPU kernel with runtime options.

Using this form of `ParallelFor`, the example in Listing 3 can be rewritten as shown in Listing 4.

Here, the code is expanded into all combinations of the run time parameters at compile time using C++17 fold expressions. The resulting code does not incur an unnecessary performance penalty if the expensive branches are not taken.

Parallel Reductions

AMReX provides a flexible and performance portable approach for common node-level reductions used by applications. Listing 5 shows an example of computing the sum, minimum and maximum of a `MultiFab` on a parallel compute device (GPU or multi-core CPU). As mentioned earlier, global reduction primitives for multi-node reductions via MPI are also provided.

```

1 auto const& a = mf.const_arrays();
2 auto result = ParReduce(
3     TypeList<ReduceOpSum, ReduceOpMin, ReduceOpMax>{},
4     TypeList<Real, Real, Real>{},
5     mf,
6     [=] AMREX_GPU_DEVICE (int box,
7         int i, int j, int k)
8     {
9         auto v = a[box](i, j, k);
10        return {v, v, v};
11    }
12 );
13 // result is a tuple holding the sum, min and max.

```

Listing 5: Performing an on-device, parallel reduce on mixed reduction types.

These parallel reduction operations for mixed data types also work for particle data, which is discussed in more detail in the next section. This is useful, for example, in the AMReX-based epidemiological modeling code `ExaEpi`. Every day, the total number of infected, hospitalizations, people on ventilators, etc. need to be tracked for diagnostic purposes. Fusing these global reduction operations minimizes the number of GPU kernel launches needed.

MPI Communication of Mesh and Particle Data

We have added support to AMReX for GPU-aware MPI. If it is enabled, AMReX allocates the buffers from a dedicated memory arena to improve performance. Moving data between the buffers and data containers (e.g., `MultiFab`) could be costly if care is not taken to minimize the GPU kernel launch overhead. When packing/unpacking the buffer, we often need to copy data from/to hundreds of small regions, which could be as small as a single cell. We use kernel fusion to ameliorate this cost. For mesh data, we have also added support for function templates for user-defined index mapping to capture application-specific communication patterns, which can be useful for multi-block applications and non-Cartesian coordinates.

Pure Struct-of-Array Particles

Particles in AMReX are associated with a level and grid in the adaptive mesh refinement hierarchy based on their

position coordinates. For CPU-only applications, grids are typically further subdivided into tiles for OpenMP threading and efficient cache-blocking (Zhang et al. 2021).

As a result, particles in AMReX are required to have `AMReX_SPACEDIM` position components. In addition, a 64-bit unsigned integer id number that uniquely identifies each particle both on and across all MPI ranks is also required; currently, AMReX supports up to $\sim 5 \times 10^{11}$ unique particles on up to $\sim 16.7 \times 10^6$ MPI ranks. These required components can be extended by an arbitrary number of real- and integer-valued optional components that can be used to store application-specific data (for example, mass, momentum, moment of inertia, ionization level, etc.). These optional components can either be registered at compile time or added at run time.

As discussed earlier, AMReX initially required the particle coordinates and id number to be stored in an AoS-style. This layout is non-optimal in several ways. First, any time the particle positions were accessed, the id numbers for those particles would also be read into cache. If the id numbers (mostly used for postprocessing and to flag particles for later removal) were not needed, between 1/4 (double precision floats, 3 spatial dimensions) and 2/3 (single precision floats, 1 spatial dimension) of the available memory bandwidth was wasted reading in data that wouldn't be used. A similar point could be made for operations that only need to access and/or modify the id number and not the particle position coordinates. Second, the AoS layout is unable to take advantage of SIMD instructions (either manual or auto-vectorization), in which the exact same operations are applied to adjacent (and properly aligned) memory locations. Taking advantage of these instructions is crucial to achieving maximum performance on modern CPU architectures. An SoA data layout can also better take advantage of vectorized loads and stores, in which consecutive threads access consecutive memory locations, on GPU architectures as well.

To address these limitations, we have added to AMReX the ability to store particle data, including the position and id components, in pure SoA style. Figure 1 illustrates the differences in data layout for the original AMReX implementation where position and id numbers were stored in AoS form and the new pure SoA-style. Our implementation of this feature makes heavy use of template metaprogramming with C++17 features such as `constexpr if` to maintain compatibility with the legacy AoS-style particle layout and to ensure that migrating codes from the old to the new layout is minimally invasive. Listing 6 demonstrates how to use the new capabilities. The code at the top is AoS-only, while the code on the bottom works for both AoS-style and SoA-style particle data.

Since these changes, several AMReX-based codes have made the transition to pure SoA particle data. In particular, we studied the effect of this change on the performance of ImpactX (Huebl et al. 2022a), a beam dynamics code including space charge effects that uses AMReX for its particle and mesh data structures as well as GPU offloading. ImpactX is the successor of the IMPACT-Z code for beam dynamics modeling in particle accelerators (Qiang et al. 2000).




Figure 1. Top: Old particle data layout with particle positions x , y , and z and the unique id number stored AoS-style, along with three additional components, a , b , and i stored SoA-style, where a and b are real-valued and i is integer-valued. This example shows the case of 3 spatial dimensions and double precision floating point data. **Bottom:** New particle data layout where all quantities are stored SoA-style. Note that, for GPU execution, consecutive threads access adjacent memory locations when reading in a component (say, the id numbers) for the new layout but not the old. This enables vectorized loads and stores and does not waste memory bandwidth on data that will not be used. The new layout also enables better cache reuse and vectorization on CPU architectures as well.

```

1 // old style based on AoS data layout
2 auto& tile = pti.GetParticleTile();
3 const auto np = tile.numParticles();
4 auto pdata = tile.getParticleTileData();
5
6 amrex::ParallelFor(np,
7 [=] AMREX_GPU_DEVICE (int i)
8 {
9     ParticleType& p = pdata.m_aos[i];
10    if (p.id().is_valid()) {
11        some_function(pdata, i);
12    }
13 });

```

```

1 // new style using pure SoA data layout
2 auto& tile = pti.GetParticleTile();
3 const auto np = tile.numParticles();
4 auto pdata = tile.getParticleTileData();
5
6 amrex::ParallelFor(np,
7 [=] AMREX_GPU_DEVICE (int i)
8 {
9     auto p = pdata[i];
10    if (p.id().is_valid()) {
11        some_function(pdata, i);
12    }
13 });

```

Listing 6: Example of migrating a ParIter loop from the old layout to the new one. In both cases, we loop over all our particles and apply an operation to only the valid ones (those with ids greater than 0). The first code snippet illustrates the old style in which the particle id was always stored as AoS. The second snippet will work for both pure SoA and legacy-style data layouts.

Table 1 summarizes the improvement benefits seen in ImpactX from transitioning its particles to the new, pure SoA format. As a benchmark, we use a standard test problem that propagates a particle beam through a single FODO cell, i.e. through one cell of lattice constructed of alternating focusing and defocusing quadrupolar lenses. The performance improvements are shown for a few select routines and also for the overall runtime of the test. For the GPU test, we used a single NVIDIA A100 SMX with 80 GB of HBM. The CPU tests were run on a single core of a 12th Gen Intel(R) Core(TM) i9-12900H. Both single precision (SP) and double precision (DP) results are shown. Overall, most routines see at least a small improvement from the change in data structure and bandwidth savings, with improvements of 2-3x not being uncommon due to synergies with improved GPU occupancy, compiler auto-vectorization for simple functions, and overhauled particle ID functions. Finally, as an added benefit, the new layout also enables the

	GPU		CPU	
	SP	DP	SP	DP
Push::Drift	2.32	1.80	2.17	1.76
Push::Quad	2.32	1.32	1.05	1.00
collect_lost	2.78	3.65	2.48	3.24
add_particles	1.05	1.06	1.02	1.02
overall	2.34	1.65	1.28	1.21

Table 1. ImpactX speedup from transition to pure SoA data layout and optimized particle id handling. DP: double precision, SP: single precision.

zero-copy GPU Python interface described in a later section for the particle position and id numbers.

Memory Management

On heterogeneous systems, there are a variety of memory types: device memory, managed memory, pinned host memory, and pageable host memory. AMReX provides a number of memory allocators so that the user can choose the type of memory for allocation. To minimize the cost of data movement, AMReX by default allocates data in device memory, and a runtime parameter is provided for switching to managed memory. Allocating and deallocating GPU memory can be very costly. To minimize memory allocation, AMReX pre-allocates a large chunk of device memory during initialization and stores it in a memory arena. This significantly improves the performance for operations where temporary memory space is needed for example, in our implementation of computing the sum of the data in a MultiFab. It takes 0.0001 seconds on NVIDIA A100 for a MultiFab with 256^3 cells using the memory arena; Without the memory arena, the cost is 0.0006 seconds.

A common mistake in GPU programming allows the memory used in a GPU kernel to be freed before the kernel finishes execution, due to the asynchronous nature of GPU kernel launches. One could add a synchronization call after the GPU kernel launch, but this could add unnecessary synchronization to the code. In AMReX, we have implemented an asynchronous-safe memory arena that does not require synchronization. An example is shown in Listing 7, in which the code is asynchronous-safe because we use `amrex::The_Async_Arena`.

```
1 {
2   FArrayBox tmp(..., The_Async_Arena());
3   ParallelFor(...); // Async kernel using tmp
4 } // Scope ends: the destructor of tmp is called.
```

Listing 7: Example of using the asynchronous-safe memory arena.

The implementation for CUDA uses a stream ordered memory allocator that guarantees the free happens after the kernel finishes. For HIP and SYCL, we use a host callback function to implement the memory safety guarantee.

These memory arenas also work with particle data. In fact, with the inherently dynamic data structures necessitated by particle methods, the use of the arenas is particularly desirable. The memory arenas also allow flexibility in placing particle data in different types of memory for GPU execution. By calling the appropriate arena, users can place data in device memory, pinned host memory, managed memory, etc, based on how they plan to use this data in their application workflow. For example, for off-line or for in-situ analysis, it often makes sense to perform the analysis on the host side, to avoid over-subscribing device memory. Both WarpX and the AMReX-based biological cell modeling code BoltzmannMFX (Palmer et al. 2023) frequently copy data into temporary host-side particle containers that use pinned memory to speed up the device-to-host data transfers for analysis.

pyAMReX

AMReX has seen wide adoption as a software library used in pre-compiled applications. With the popularity of interactive computing in recent years, e.g., via Project Jupyter (Kluyver et al. 2016), and the rise of scripting languages as productive runtime “glue code” to connect and steer software components, using AMReX from a scripting layer for tasks like integration, data processing and/or prototyping is desirable.

The new pyAMReX project is part of the AMReX software ecosystem and builds directly on the AMReX C++ library. The Python language bindings of pyAMReX bridge the compute role of AMReX in block-structured simulation codes with the data science software ecosystem. With pyAMReX providing zero-copy GPU data access for workflows in AI/ML and in situ analysis, applications can be coupled and extended from Python, and rapid, massively parallel prototyping is enabled.

pyAMReX is implemented as a C++ library via pybind11 (Jakob et al. 2017), which targets the Python C API to expose standardized C entry points of Python modules. Due to the evolution of the C++ language from C++11 to C++17, type introspection and meta-programming inside pybind11 make adding new bindings for AMReX functions, classes, etc. simple and maintainable.

```

1 #include "pyAMReX.H"
2 #include <AMReX_Box.H>
3
4 // ...
5 py::class_< Box >(m, "Box")
6     // constructor(s)
7     .def(py::init<
8         IntVect const &,
9         IntVect const &
10        >(),
11         py::arg("small"),
12         py::arg("big")
13     )
14     // ...
15
16     .def("intersects",
17         &Box::intersects)
18     .def_property_readonly(
19         "ok", &Box::ok)
20 ;

```

Listing 8: pyAMReX uses plain C++ code to register Python types. This snippet defines the AMReX class `Box` in Python with `pybind11`.

Listing 8 shows how to expose the C++ type `amrex::Box` to Python; once it is registered other classes and methods can use this type in their Python interfaces. Following the same logic, applications that are built on AMReX just need to add minimal Python bindings themselves, because they can rely on all the Python-registered AMReX types from pyAMReX.

Python modules written with `pybind11` are pre-compiled to shared libraries, which automatically expose only a set of standardized, C API symbols for `import` into the Python interpreter, that ultimately call AMReX C++ routines. Additional type extensions that are more effectively written in pure Python are added at `import` time. In the future, pyAMReX might transition to `nanobind` (Jakob 2022), which is a recent reimplementation of `pybind11` APIs with improved compile-, link- and runtime of the binding code.

Zero-Copy APIs

Exchanging large, heap/GPU-allocated data across API interfaces in Python, without the need to create copies, benefits from the permissive memory access rights in operating systems, since the Python interpreter and all its dynamically loaded (imported) libraries reside in the same process. Thus, any process-allocated memory can be accessed across Python modules and loaded C++ libraries by exchanging a pointer and the respective meta information that describe data placement and layout.

Luckily, in the past years the exchange of multi-dimensional (ND) array data in Python has been standardized in CPU code via the array protocol (Consortium for Python Data API Standards 2021) driven by NumPy (Harris et al. 2020), based on the earlier buffer protocol implementations in Python. Adding support for sharing a Python type’s data with read/write access requires implementation of a `__array_interface__` dictionary property for AMReX Python types, which describes primarily shape, stride, data types, description

and data pointer, among other details. The standardized `__cuda_array_interface__` uses nearly the same metadata, with an optional CUDA stream added. Since unified memory pointers are available in CUDA, a pointer to GPU data can be exchanged between AMReX and CuPy (Okuta et al. 2017), Numba (Lam et al. 2015), or PyTorch (Paszke et al. 2019) to create views into their respective host-controlled GPU data container that works the same way as it works with NumPy CPU array data.

A new iteration on the CPU/GPU array interfaces was pursued with the deep learning pack standard (DLPack). DLPack uniformly wraps CPU or GPU data and brings improvements in supported accelerated programming models, more supported packages in the data science ecosystem, and in lifetime (data ownership) handling. A DLPack implementation in pyAMReX is ongoing* and will enable ROCm and SYCL GPU support.

Listing 9 shows how the array interface is used in practice. Line 7 creates a CuPy ND array that is a readable and writable data view into an existing AMReX `Array4`. CuPy reads the `Array4.__cuda_array_interface__` dictionary and extracts data pointer and layout information to create a non-owning `cupy.ndarray`. Equivalently, NumPy ND arrays and PyTorch tensors are created.

```

1 import amrex.space3d as amr
2 import cupy as cp
3
4 mfab = amr.MultiFab(<...>)
5 for mfi in mfab:
6     array4 = mfab.array(mfi)
7     cp_arr = cp.array(array4, copy=False)

```

Listing 9: Standardized interfaces on Python types can be used to create views from one datatype into another. Here, a CuPy ND array is used to read-write access the field data in an AMReX `Array4`.

Convenient methods are added to types like `Array4`, e.g., `.to_numpy()` and `.to_cupy()`, which spare the user from the extra line and also keep the ND index order in AMReX-typical Fortran order (indexed `x, y, z, component` as `F_CONTIGUOUS`, although `C_CONTIGUOUS` is the default for Python). Notably, in benchmarks with CuPy, Fortran indexing caused performance drawbacks. Thus, experienced users have the choice to set a `.to_cupy(order="C")` argument to index as C-ordered `component, z, y, x` and perform optimally with CuPy.†

Additional arguments such as `.to_numpy(copy=True)` are supported and perform the expected device-to-host copy if called on a device memory container, and a device-to-device copy if performed as `.to_cupy(copy=True)`. For users that use the managed memory arena in AMReX, it is worth noting that `.to_numpy(copy=False)` works for read and write on

*<https://github.com/AMReX-Codes/pyamrex/issues/9>

†<https://github.com/cupy/cupy/issues/7783>

GPU memory as expected, but includes the performance penalty of implicit data transfers.

Computing in Python

Currently, `pyAMReX` does not expose `amrex::ParallelFor/Reduce/...` computing primitives, but can call pre-compiled functions from Python that make use of those. In order to be able to generate new compute kernels at runtime from the Python scripting language, `pyAMReX` makes use of the aforementioned zero-copy APIs to pass data for compute into a third-party Python CPU/GPU framework of choice, e.g., NumPy, CuPy, Numba, PyTorch, which then can support just-in-time (JIT) GPU kernel compilation.

Listing 10 shows how to generate accelerated array computations on an ND field (`AMReX MultiFab`). As is typical in `AMReX`, the outermost loops iterate over mesh-refinement levels and data blocks on the current device. Lines 22ff. create a CuPy or NumPy array view, followed by a device-side assignment of the value `42.0` to all index points using the `()` ellipsis, independent of array rank.

```

1 # finest active MR level, get from a
2 # simulation's AmrMesh object, e.g.:
3 # finest_level = sim.finest_level
4
5 # iterate over mesh-refinement levels
6 for lev in range(finest_level + 1):
7     # get an existing MultiFab, e.g.,
8     # from a simulation:
9     # mfab = sim.get_field(lev=lev)
10    # Config = sim.extension.Config
11
12    # grow (aka guard/ghost/halo) regions
13    ngv = mfab.n_grow_vect
14
15    # get every local block of the field
16    for mfi in mfab:
17        # global index box w/ guards
18        bx = mfi.tilebox().grow(ngv)
19
20        # numpy representation: non-
21        # copying view, w/ guard/ghost
22        field = mfab.array(mfi).to_cupy() \
23                if Config.have_gpu else \
24                mfab.array(mfi).to_numpy()
25
26        field[()] = 42.0

```

Listing 10: An example Python script showing accelerated array operations on a field with `pyAMReX v24.03`.

One difference when operating on NumPy/CuPy views compared to `AMReX Array4` is the index offset of each local data block. The former use zero-based indexing on local blocks while `AMReX` uses global indexing.

In Listing 10, only the select-all array operation `[()] = 42.0` represents the “hot loop” of operations (or on GPU, the kernel function). Other Python array ellipsis selectors such as operating on an index range with an offset (e.g., `[1:, 1:, 1:]`) or a slice (e.g., `[0, ...]`), etc. can be used as usual to form more complex expressions such as stencils. In frameworks like CuPy and Numba, multiple

subsequent array operations can also be fused into one kernel, using their respective framework JIT syntax.

Particle operations follow the same logic, as shown in Listing 11. Both the legacy AoS + SoA and the new pure SoA particle layout are supported in `pyAMReX`. The pure SoA particle layout provides better performance as well as write access for GPU particle positions and indices in CuPy.[‡]

```

1 # code-specific getter function, e.g.:
2 # pc = sim.get_particles()
3 # Config = sim.extension.Config
4
5 # iterate over mesh-refinement levels
6 for lvl in range(pc.finest_level + 1):
7     # loop local tiles of particles
8     for pti in pc.iterator(pc, level=lvl):
9         # compile-time and
10        # runtime attributes
11        soa = pti.soa().to_cupy() \
12                if Config.have_gpu else \
13                pti.soa().to_numpy()
14
15        x = soa.real["x"]
16        y = soa.real["y"]
17
18        # write to all particles in tile
19        x[:] = 0.30
20        y[:] = 0.35
21        soa.real["z"][:] = 0.40
22
23        soa.real["a"][:] = x[:]**2
24        soa.real["b"][:] = x[:] + y[:]
25        # ...
26
27        # all int attributes
28        for soa_int in soa.int.values():
29            soa_int[:] = 12

```

Listing 11: An example Python script showing accelerated particle operations with `pyAMReX v24.03`.

As for the ND field arrays, the “hot loop” of computation is best placed in array operations such as `[:] = value`. Selective indexing and masking with temporary arrays are the usual patterns in high-performance Python, and various Python frameworks support JIT compilation of explicit indices when needed.

Project Integration

`pyAMReX` was initially designed with workflows in mind: (i) to enhance an existing `AMReX` application with Python code, data-science and/or AI/ML capabilities, (ii) to write a standalone application or test of `AMReX`, rapidly prototyped in Python.

Enhancing, steering and modifying HPC applications with routines written in Python is a desirable feature in exascale-ready codes such as `WarpX`, which already benefited from a runtime modularity in the (CPU-only) predecessor code `WARP` (Grote et al. 2005). Spearheading the adoption of `pyAMReX` for exascale modeling codes, the `AMReX`-based `WarpX` and `ImpactX` simulation codes (Huebl et al. 2022b)

[‡]<https://github.com/cupy/cupy/issues/2031>

use user-defined C++ to Python callbacks and manipulate AMReX-managed GPU data via pyAMReX: users adjust the treatment of selected particle species, prototype new physical accelerator elements or replace numerical solvers, load external data, and steer simulation behavior in a way that would be too custom or complex to define in a key-value based inputs file syntax.

An integration with a machine learning (ML) framework has recently been reported in Sandberg et al. (2024). In this work, the ImpactX code (Huebl et al. 2022a) is augmented using pyAMReX to couple analytical particle beam updates and data-driven particle updates using a pre-trained neural network. The ML training data was created with a high-fidelity WarpX simulation and used to train a neural network offline. The resulting model is then integrated with the CUDA bindings of pyAMReX to update the GPU-accelerated ImpactX particle simulation state with GPU-accelerated neural network inference using PyTorch. For particle accelerator physicists, replacing traditionally costly sections of a beamline simulation with such specialized surrogate models enables the study of whole-device models in start-to-end simulations.

Finally, pyAMReX also supports the design of standalone benchmarks, prototype examples, new AMReX-powered command line tools and even applications, purely from Python. AMReX tutorials have been updated to include examples of the Python usage of the `MultiFab` class and to solve a multi-dimensional heat equation, providing easy accessibility of AMReX to new community members and for use in education. Additionally, being able to integrate AMReX functionality in scripts will enable users and developers to build standalone data processing tools, e.g., for post-processing, and will enable rapid prototyping and unit testing of AMReX functionality and numerical solvers based on AMReX.

Conclusions and Future Work

In this paper, we have reviewed the basic features of AMReX and discussed a number of new features that have been added to the framework. One of these new features supports the aggregation of a number of small computational kernels into a single kernel launch. A related feature incorporates an efficient kernel launching mechanism that supports selection of different kernel options at runtime. We have also developed a flexible reduction mechanism that performs a number of different reductions simultaneously.

We have made several enhancements to the particle functionality in AMReX. We have introduced a more flexible approach for representing particle data that can combine both SoA and AoS representations. In particular, AMReX now supports a pure SoA representation that has led to significant performance improvements for a number of applications. We have also developed a Python binding of AMReX, pyAMReX, that provides a bridge between AMReX simulation methodology and data science software.

There are a number of additional capabilities that we are currently developing within AMReX to support the requirements of new applications. One area of active development is to extend interoperability using a multiblock paradigm to couple different applications efficiently. In a

multiblock approach, each application has its own index space and the framework provides the mapping between index spaces and between (possibly) different solution representations. The applications can then be run in an MPMD mode using separate MPI communicators with AMReX handling the communication between the different codes needed for the overall simulation. This facilitates coupling different applications without significant recoding. Use cases here range from solving a system of PDEs on a mapped multiblock domain that cannot be represented with a global index space to coupling different PDE solvers in different domains to coupling a stochastic algorithm such as kinetic Monte Carlo to a fluctuating hydrodynamics solver to model a catalytic surface interacting with a fluid. A related area of development is to support algorithm refinement approaches in which the representation of the underlying physics changes as a function of resolution. A hybrid algorithm that dynamically switches from a coarse-grained continuum model to a fine-grained particle model depending on flow conditions would be a typical use case for this capability. We are also extending the python interface to AMReX to support integration of ML concepts into AMReX-based applications. Use cases here would include developing ML-based surrogates for outer loop operations and developing ML methodologies for coupling between different scales.

Another driver for AMReX development is extending the performance portability constructs to work optimally on new architectures or new features in the software stack on existing architectures. One particularly interesting direction in this direction is the use of AMReX as the basis for the co-design of custom architectures. The developments of chiplet technology have the potential to support the design of performant, energy efficient specialized architectures for select applications.

Acknowledgements

We acknowledge the AMReX open source community for their invaluable contributions. In addition, we acknowledge and thank Shreyas Ananthan, Ryan Sandberg and all other [pyAMReX contributors](#). 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; the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Exascale Computing Project under contract DE-AC02-05CH11231; the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) via the Scientific Discovery through Advanced Computing (SciDAC) program FASTMath Institute; and the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under and resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

- Almgren A, Lattanzi A, Haque R, Jha P, Kosovic B, Mirocha J, Perry B, Quon E, Sanders M, Wiersema D, Willcox D, Yuan X and Zhang W (2023) Erf: Energy research and forecasting. *Journal of Open Source Software* 8(87): 5202. DOI:10.21105/joss.05202. URL <https://doi.org/10.21105/joss.05202>.
- Almgren A, Sazo MB, Bell J, Harpole A, Katz M, Sexton J, Willcox D, Zhang W and Zingale M (2020) CASTRO: A Massively Parallel Compressible Astrophysics Simulation Code. *Journal of Open Source Software* 5(54): 2513. DOI: 10.21105/joss.02513.
- Almgren AS, Beckner VE, Bell JB, Day MS, Howell LH, Joggerst CC, Lijewski MJ, Nonaka A, Singer M and Zingale M (2010) CASTRO: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-gravity. *Astrophysical Journal* 715(2): 1221–1238. DOI:10.1088/0004-637X/715/2/1221.
- Almgren AS, Bell JB, Colella P, Howell LH and Welcome M (1998) A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *Journal of Computational Physics* 142: 1–46. DOI:10.1006/jcph.1998.5890.
- Almgren AS, Bell JB, Lijewski MJ, Lukić Z and Andel EV (2013) Nyx: A massively parallel AMR code for computational cosmology. *The Astrophysical Journal* 765(1): 39. DOI: 10.1088/0004-637x/765/1/39.
- Beckingsale DA, Burmark J, Hornung R, Jones H, Killian W, Kunen AJ, Pearce O, Robinson P, Ryuji BS and Scogland TR (2019) RAJA: Portable Performance for Large-Scale Scientific Applications. In: *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*. pp. 71–81. DOI:10.1109/P3HPC49587.2019.00012.
- Berger MJ and Colella P (1989) Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics* 82: 64–84. DOI:10.1016/0021-9991(89)90035-1.
- Berger MJ and Oliger J (1984) Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics* 53: 484–512. DOI:10.1016/0021-9991(84)90073-1.
- Consortium for Python Data API Standards (2021) Python array API standard. URL <https://data-apis.org/array-api>.
- de Frahan MTH, Rood JS, Day MS, Sitaraman H, Yellapantula S, Perry BA, Grout RW, Almgren A, Zhang W, Bell JB and Chen JH (2023) Pelec: An adaptive mesh refinement solver for compressible reacting flows. *The International Journal of High Performance Computing Applications* 37(2): 115–131. DOI: 10.1177/10943420221121151.
- Dhruv A (2024) AMReX-Bittree-Performance. URL <https://github.com/Lab-Notebooks/AMReX-Bittree-Performance>.
- Diederichs S, Benedetti C, Huebl A, Lehe R, Myers A, Sinn A, Vay JL, Zhang W and Thévenet M (2022) Hipace++: A portable, 3d quasi-static particle-in-cell code. *Computer Physics Communications* 278: 108421. DOI:10.1016/j.cpc.2022.108421.
- Dubey A, Almgren A, Bell J, Berzins M, Brandt S, Bryan G, Colella P, Graves D, Lijewski M, Löffler F, O’Shea B, Schnetter E, Van Straalen B and Weide K (2014) A survey of high level frameworks in block-structured adaptive mesh refinement packages. *Journal of Parallel and Distributed Computing* 74(12): 3217 – 3227. DOI:10.1016/j.jpdc.2014.07.001.
- Edwards H, Trott CR and Sunderland D (2014) Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing* 74(12): 3202 – 3216. DOI:10.1016/j.jpdc.2014.07.003. Domain-Specific Languages and High-Level Frameworks for High-Performance Computing.
- Esclapez L, Day M, Bell J, Felden A, Gilet C, Grout R, de Frahan MH, Motheau E, Nonaka A, Owen L, Perry B, Rood J, Wimer N and Zhang W (2023) Pelelmex: an amr low mach number reactive flow simulation code without level sub-cycling. *Journal of Open Source Software* 8(90). DOI: 10.21105/joss.05450.
- ExaEpi (2023) *ExaEpi documentation*. URL <https://exaeppi.readthedocs.io>.
- Fedeli L, Huebl A, Boillod-Cerneux F, Clark T, Gott K, Hillairet C, Jaure S, Leblanc A, Lehe R, Myers A, Piechurski C, Sato M, Zaïm N, Zhang W, Vay JL and Vincenti H (2022) Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers. In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC ’22*. IEEE Press. ISBN 9784665454445. DOI:10.1109/SC41404.2022.00008.
- Grote DP, Friedman A, Vay J and Haber I (2005) The WARP Code: Modeling High Intensity Ion Beams. *AIP Conference Proceedings* 749(1): 55–58. DOI:10.1063/1.1893366.
- Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C and Oliphant TE (2020) Array programming with NumPy. *Nature* 585(7825): 357–362. DOI:10.1038/s41586-020-2649-2.
- Huebl A, Lehe R, Mitchell CE, Qiang J, Ryne RD, Sandberg RT and Vay JL (2022a) Next Generation Computational Tools for the Modeling and Design of Particle Accelerators at Exascale. In: *Proc. 5th Int. Particle Accel. Conf. (NAPAC’22)*, number 5 in International Particle Accelerator Conference. Albuquerque, NM, USA: JACoW Publishing, Geneva, Switzerland. ISBN 978-3-95450-232-5, pp. 302–306. DOI:10.18429/JACoW-NAPAC2022-TUYE2.
- Huebl A, Lehe R, Zoni E, Shapoval O, Sandberg RT, Garten M, Formenti A, Jambunathan R, Kumar P, Gott K, Myers A, Zhang W, Almgren A, Chad E M, Qiang J, Sinn A, Diederichs S, Thevenet M, Grote DP, Fedeli L, Clark T, Zaim N, Vincenti H and Vay JL (2022b) From Compact Plasma Particle Sources to Advanced Accelerators with Modeling at Exascale. In: *Proceedings of the 20th Advanced Accelerator Concepts Workshop (AAC’22)*, Advanced Accelerator Concepts Workshop. Hauppauge, NY, USA: arXiv. URL <https://arxiv.org/abs/2303.12873>. In print.
- Jakob W (2022) nanobind: tiny and efficient C++/Python bindings. <https://github.com/wjakob/nanobind>.
- Jakob W, Rhinelander J and Moldovan D (2017) pybind11 – Seamless operability between C++11 and Python. <https://github.com/pybind/pybind11>.

- Kluyver T, Ragan-Kelley B, Pérez F, Granger B, Bussonnier M, Frederic J, Kelley K, Hamrick J, Grout J, Corlay S, Ivanov P, Avila D, Abdalla S, Willing C and development team J (2016) Jupyter notebooks - a publishing format for reproducible computational workflows. In: Loizides F and Schmidt B (eds.) *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Netherlands: IOS Press, pp. 87–90.
- Lam SK, Pitrou A and Seibert S (2015) Numba: a llvm-based python jit compiler. In: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450340052. DOI:10.1145/2833157.2833162.
- Musser J, Almgren AS, Fullmer WD, Antepara O, Bell JB, Blaschke J, ott KG, Myers A, Porcu R, Rangarajan D, Rosso M, Zhang W and Syamlal M (2022) Mfix-exa: A path toward exascale cfd-dem simulations. *The International Journal of High Performance Computing Applications* 36(1): 40–58. DOI:10.1177/10943420211009293.
- Okuta R, Unno Y, Nishino D, Hido S and Loomis C (2017) CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations. In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*.
- Palmer BJ, Almgren AS, Johnson CGM, Myers AT and Cannon WR (2023) BMX: Biological modelling and interface exchange. *Scientific Reports* 13(1): 12235. DOI:10.1038/s41598-023-39150-1.
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J and Chintala S (2019) PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: Wallach H, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox E and Garnett R (eds.) *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Porcu R, Musser J, Almgren AS, Bell JB, Fullmer WD and Rangarajan D (2023) MFIX-Exa: CFD-DEM simulations of thermodynamics and chemical reactions in multiphase flows. *Chemical Engineering Science* 273: 118614. DOI:10.1016/j.ces.2023.118614.
- Qiang J, Ryne RD, Habib S and Decyk V (2000) An object-oriented parallel particle-in-cell code for beam dynamics simulation in linear accelerators. *Journal of Computational Physics* 163(2): 434–451. DOI:10.1006/jcph.2000.6570.
- REMORA (2023) *REMORA documentation*. URL <https://remora.readthedocs.io>.
- Sandberg RT, Lehe R, Mitchell CE, Garten M, Myers A, Qiang J, Jean-Luc V and Huebl A (2024) Synthesizing particle-in-cell simulations through learning and gpu computing for hybrid particle accelerator beamlines. In: *Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '24*. New York, NY, USA: Association for Computing Machinery. Under review.
- Sharma A, Brazell MJ, Vijayakumar G, Ananthan S, Cheung L, deVelder N, Henry de Frahan MT, Matula N, Mullaney P, Rood J, Sakievich P, Almgren A, Crozier PS and Sprague M (2024) ExaWind: Open-source CFD for hybrid-RANS/LES geometry-resolved wind turbine simulations in atmospheric flows. *Wind Energy* DOI:10.1002/we.2886.
- Vay J, Almgren A, Bell J, Ge L, Grote D, Hogan M, Kononenko O, Lehe R, Myers A, Ng C, Park J, Ryne R, Shapoval O, Thévenet M and Zhang W (2018) Warp-X: A new exascale computing platform for beam-plasma simulations. *Nuclear Instruments and Methods in Physics Research Section A Accelerators Spectrometers Detectors and Associated Equipment* DOI: 10.1016/j.nima.2018.01.035.
- Zhang W, Myers A, Gott K, Almgren A and Bell J (2021) AMReX: Block-structured adaptive mesh refinement for multiphysics applications. *The International Journal of High Performance Computing Applications* 35(6): 508–526. DOI:10.1177/10943420211022811.

Author biography

Andrew Myers is a Computer Systems Engineer in the Center for Computational Sciences and Engineering at Lawrence Berkeley National Laboratory. His work focuses on scalable particle methods for emerging architectures in the context of adaptive mesh refinement. He is an active developer of the AMReX framework and of the electromagnetic Particle-in-Cell code WarpX.

Weiqun Zhang is a Computer Systems Engineer at Lawrence Berkeley National Laboratory. He is interested in high-performance computing, computational physics, and programming in general. Currently, he works on the AMReX software framework and WarpX, an advanced electromagnetic Particle-in-Cell code.

Ann Almgren is a Senior Scientist at Lawrence Berkeley National Laboratory, and the Department Head of the Applied Mathematics Department in LBL's Applied Mathematics and Computational Research Division. Her primary research interest is in computational algorithms for solving PDE's in a variety of application areas. Her current projects include the development and implementation of new multiphysics algorithms in high-resolution adaptive mesh codes that are designed for the latest hybrid architectures. She is a Fellow of the Society of Industrial and Applied Mathematics and was the Deputy Director of the ECP AMReX Co-Design Center.

Thierry Antoun pursued his engineering studies at ENSTA Paris, followed by an MSc in Applied Mathematics, specializing in Modeling and Simulation at the Institut Polytechnique de Paris. He was a summer intern at Lawrence Berkeley National Laboratory, focusing on optimizing AMReX particle data structures and the PIC codes WarpX & ImpactX. Since the end of 2023, he is a member of the CExA team at the French Alternative Energies and Atomic Energy Commission (CEA), working on the Kokkos library to facilitate the transition of codes to GPUs.

John Bell is a Senior Scientist at Lawrence Berkeley National Laboratory. His research focuses on the development and analysis of numerical methods for partial differential equations arising in science and engineering. He is a Fellow of the Society of Industrial and Applied Mathematics and a member of the National Academy of Sciences, and was the Director of the ECP AMReX Co-Design Center.

Axel Huebl is a Research Scientist and computational physicist at Lawrence Berkeley National Laboratory. He works at the interfaces

of HPC, laser-plasma physics, and advanced particle accelerator research. He drives open standards in his community, is the software architect of the Beam, Plasma & Accelerator Simulation Toolkit (BLAST) and is a lead-developer of Exascale simulation software such as WarpX (2022 Gordon Bell Prize Winner) and ImpactX. He is an active developer of AMReX and leads pyAMReX.

Alexander Sinn is an undergraduate student at Deutsches Elektronen-Synchrotron (DESY). His research interest are in high-performance computing and modeling of particle accelerators. He is one of the maintainers of the BLAST simulation software HiPACE++ for wakefield accelerator modeling.