

From Merging Frameworks to Merging Stars: Experiences using HPX, Kokkos and SIMD Types

Gregor Daiss[†], Srinivas Yadav Singanaboina^{*}, Patrick Diehl^{†*}§, Hartmut Kaiser^{*}, Dirk Pflüger[†]

^{*}LSU Center for Computation & Technology, Louisiana State University, Baton Rouge, LA, 70803 U.S.A

[†]IPVS, University of Stuttgart, Stuttgart, 70174 Stuttgart, Germany

Email: Gregor.Daiss@ipvs.uni-stuttgart.de

[§] Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA, 70803 U.S.A.

Abstract—Octo-Tiger, a large-scale 3D AMR code for the merger of stars, uses a combination of HPX, Kokkos and explicit SIMD types, aiming to achieve performance-portability for a broad range of heterogeneous hardware. However, on A64FX CPUs, we encountered several missing pieces, hindering performance by causing problems with the SIMD vectorization. Therefore, we add `std::experimental::simd` as an option to use in Octo-Tiger’s Kokkos kernels alongside Kokkos SIMD, and further add a new SVE (Scalable Vector Extensions) SIMD backend. Additionally, we amend missing SIMD implementations in the Kokkos kernels within Octo-Tiger’s hydro solver. We test our changes by running Octo-Tiger on three different CPUs: An A64FX, an Intel Icelake and an AMD EPYC CPU, evaluating SIMD speedup and node-level performance. We get a good SIMD speedup on the A64FX CPU, as well as noticeable speedups on the other two CPU platforms. However, we also experience a scaling issue on the EPYC CPU.

Index Terms—HPX, Kokkos, SIMD, astrophysical simulation

I. INTRODUCTION

Astrophysical applications drive the need for high-performance computing. They require a lot of computational power for their simulation, as well as numerous developer hours to ascertain they run efficiently on the ever-changing set of current hardware platforms.

Octo-Tiger is an astrophysical simulation used to simulate binary star systems and their eventual outcomes [1]. Octo-Tiger is built on HPX [2] for task-based programming to scale to all cores within one compute node and, beyond that, to thousands of other compute nodes in distributed scenarios. Using this, Octo-Tiger achieved scalability on Cori [3], and more recently Piz Daint [4] and Summit [5].

Consequently, given these last two machines, the most recent work on Octo-Tiger was focused on porting the original CPU-only implementation of Octo-Tiger to GPUs. To have portable GPU kernels, much of this was done with Kokkos and an HPX-Kokkos integration to allow us to use Kokkos kernels as HPX tasks. To improve CPU performance of said Kokkos kernels, we used SIMD (Single Instruction Multiple Data) types provided by Kokkos [6]. While working well in Octo-Tiger, this was only used in the Kokkos kernels of one solver and only tested on older CPU platforms [7].

In turn, this work focuses not on running Octo-Tiger on GPUs, but instead on our development process modifying Octo-Tiger to run efficiently on Fujitsu A64FX™ CPUs. This is done to prepare Octo-Tiger for experimental runs on Fugaku.

Hence, we will use the current development snapshot of Octo-Tiger to investigate how well our utilized execution model (a mixture of HPX, Kokkos, and explicit vectorization) works on modern CPUs and share some additions we had to make when porting to A64FX. Although said mixture of frameworks already gained us a certain degree of portability, we had to supplement some missing ingredients in this work to make Octo-Tiger run more efficiently on A64FX CPUs and make use of their SVE SIMD instructions:

- 1) To have a wider range of SIMD backends accessible, we integrated `std::experimental::simd` within our Kokkos kernels (while maintaining compatibility with the Kokkos SIMD types).
- 2) We add `std::experimental::simd` compatible SVE types, allowing for explicit vectorization on A64FX CPUs.
- 3) Some of our Kokkos kernels simply did not yet use the explicit vectorization with SIMD types, as they have only been used on GPUs so far. Therefore, we supplement those kernels (two kernels in our new hydro solver) with these types and apply SIMD masking where required.

To test these new additions, we run node-level tests on an A64FX CPU, determining node-level scaling and SIMD speedup using various backends (both for the entire application and the most important CPU kernels). We investigate the speedup of both the new hydro SIMD implementation and the speedup of the surrounding hydro solver in a hydro-only compute scenario.

We further use this opportunity to run the same tests (using different SIMD types) on recent Intel and AMD CPUs, to evaluate the performance on current CPUs, as previous CPU results were gathered on now-outdated hardware with fewer cores [7]. Moreover, those previous results did not yet include the SIMD additions to the hydro solver from this work.

This paper is structured as follows: Section II briefly introduces the real-world astrophysics application: Octo-Tiger. Section III emphasizes the technology used in Octo-Tiger and how its components fit together to create a powerful yet portable execution model. Section IV describes the changes that we had to make for this to work for A64FX. Section V lists the performance measurements conducted with the real-world application. Section VI focuses on the related work. Finally, Section VII concludes the work and outlines the next steps.




Fig. 1. Flow on the surface of the two stars. In the center we see the aggregation belt with the mass transfer from the smaller star to the larger star.

II. SCIENTIFIC APPLICATION: OCTO-TIGER

From the astrophysics perspective, our interest are binary star systems and their eventual outcomes, especially white dwarf mergers and the contact binary v1309 Sco. For the latter, emission of the red light during the merger was observed. To gain more understanding in the process, self-consistent simulations with a very high resolution to resolve the star atmosphere to extract the simulated light curve are necessary. This would allow for a direct comparison with the observation data. Furthermore, with more understanding of the light curve in v1309 Sco, it will be possible to reliably simulate the light curve of other star mergers.

In the following, we show a brief overview of the solvers and data-structure Octo-Tiger uses to simulate these binary star systems.

a) Solvers: The star system is modeled as a self-gravitating, astrophysical fluid. To solve the system, Octo-Tiger uses a coupled hydro solver and gravity solver. The inviscid Navier-Stokes equations are solved using finite volumes in the hydro solver [8]. In turn, Newtonian gravity is solved using a modified Fast Multipole Method (FMM) in the gravity solver [9].

The most important Octo-Tiger compute kernels are part of those two solvers. In the hydro solver, there is the `reconstruct` kernel, using the piecewise-parabolic method computing the values of the variables that are being evolved at 26 quadrature points. Furthermore, there is the `flux` kernel, which uses those values to compute the final flux using Newton-Cotes quadrature. In turn, the gravity solver contains multiple compute kernels calculating the same-level interactions between close-by cells (the most compute-intensive FMM step): The `Monopole` kernel for non-refined cells and the `Multipole` kernel (and its `rho` variants for angular momentum correction and root variant for a specialization for the root sub-grid). For brevity, we refer to [1] for more details on the solvers. For more details on gravity solver kernels in particular, we refer to [10], [11]. Lastly, for a convergence study of a production run and more details about the hydro kernels, we refer to [12].

b) Data-Structure: Octo-Tiger uses adaptive mesh refinement (AMR) to focus on the area of interest, the atmosphere between the two stars with the aggregation belt, see Figure 1.

For efficient computations, an adaptive octree with a Cartesian sub-grid in each leaf is used. Each sub-grid has 512 cells within a $8 \times 8 \times 8$ cube, which is the default. Hence, each compute kernel invocation is operating on one sub-grid only, using HPX to launch many such kernels concurrently for the available sub-grids. The cube size of the sub-grids can be configured during compile time. The performance benefits of various sub-grid sizes were studied in [5].

III. HPX, KOKKOS AND SIMD TYPES

Octo-Tiger is built using a combination of frameworks that make up our execution model: We use HPX for task-based parallelization and distributed computing. We use Kokkos for writing compute kernels that are portable between various CPU and GPU platforms. Lastly, we use SIMD types (provided by Kokkos) for explicit vectorization on CPUs, whilst still supporting GPU execution (via instantiation of the SIMD template type with scalar types).

In this section, we briefly cover these utilized frameworks and how they can fit together to complement each other.

A. HPX

HPX is an Asynchronous Many-Task Runtime system that is implemented in C++ [2], [13]. The library implements all APIs related to concurrency and parallelism as mandated by recent ISO standards C++20 and C++23 in a conforming way. In this context, it implements all the (more than 100) parallel algorithms as described in the C++ standard. It has been described in detail in other publications, such as [14], [15].

In the context of this paper, HPX has been used as an underlying runtime platform for the Octo-Tiger astrophysics application used as the domain science driver (see [4] for more details), thus managing local and distributed parallelism while observing all necessary data dependencies. Data and task dependencies can be expressed with HPX futures, and chained together in an execution graph. This graph can be built asynchronously, with the HPX worker threads processing the tasks when their dependencies are fulfilled. This task-based programming model is particularly useful for parallel implementations of adaptive, tree-based codes like Octo-Tiger, as we can build the task graph quickly when traversing the tree to make concurrent work available to the system.

We also use the performance monitoring library APEX [16] that is well integrated with HPX. APEX can be applied to capture a combination of task-based events with hardware counter information for optimizing HPX on different hardware platforms. APEX can further measure the runtime of annotated HPX tasks, getting mean execution times for them, which can be used to determine speedups for specific parts of the code (such as compute kernels with and without SIMD).

B. HPX with Kokkos

For portability, we use Kokkos' [17] abstractions for memory and execution within Octo-Tiger. Kokkos allows us to easily write compute kernels that work on both CPU and

GPU, allowing us to run the same kernel implementation on NVIDIA[®] and AMD[®] GPUs, and even on the CPU if desired.

Usually, these kernels would be launched in a fork-join manner using Kokkos, for example using an OpenMP execution space on the CPU to achieve concurrency. However, there exist two HPX-Kokkos integrations, which allow us to avoid this and use Kokkos in a more task-based fashion.

The first one is the Kokkos HPX execution space. Kernels launched within this execution space are split into HPX tasks, which will be processed by the existing HPX worker threads (hence there is no need for multiple competing thread pools). However, this alone does not provide the functionality to launch Kokkos functions asynchronously and integrate them within HPX’s execution graph. For that, we use the second integration: HPX-Kokkos [7]. This integration (and its executors) allow us to obtain `hpx::futures` for asynchronous Kokkos kernel (and function) launches, facilitating their integration with HPX’s asynchronous execution graph.

Together, these integrations allow us to write portable Kokkos kernels that we can launch from arbitrary HPX tasks and treat as HPX tasks themselves. This way, we can express dependencies between kernels (and other tasks) using HPX futures, automatically triggering new tasks when an asynchronously launched kernel finishes. Furthermore, this way the Kokkos kernels make use of HPX resources for CPU execution (namely the existing worker threads).

C. Explicit Vectorization with the Kokkos SIMD types

To gain truly portable kernels that run well on both the CPU and GPU, we need to take SIMD vectorization into account. Modern CPUs offer a lot of the potential floating-point performance because of their ability to run instructions on multiple data items at once (SIMD). While the compiler can use this automatically in certain circumstances, it is often more reliable to use the appropriate instructions directly to ensure SIMD usage. Since different hardware platforms often use different instruction sets (AVX, AVX512, NEON and SVE to name some examples), abstractions using C++ types have been developed to increase portability and ease-of-use. Kokkos itself includes such SIMD types¹, notably capable of using vectorization on the CPU by compiling to the correct SIMD instructions, but also able to use scalar types in case the same compute kernel is compiled for GPU usage [6].

D. Utilizing the Frameworks within Octo-Tiger

Octo-Tiger was built to rely on HPX for parallelization and distributed computing, with each sub-grid in the octree representing an HPX component which can be placed on arbitrary compute nodes. However, its compute kernels changed significantly over the last two years. In the past, Octo-Tiger used separate compute kernels for CPU execution (using Vc for SIMD types) and for the GPU execution (using NVIDIA[®] CUDA[®]) [4], [11]. In an effort to unify these implementations into a single portable one for each kernel,

¹<https://github.com/kokkos/simd-math>




Fig. 2. The execution model used to launch compute kernels in Octo-Tiger: We launch Kokkos kernels concurrently from arbitrary HPX tasks (and thus threads), keeping track of the results via the HPX futures HPX-Kokkos returns. Kokkos in turn splits the kernel into HPX tasks (using the HPX execution space). We use SIMD types for explicit SIMD vectorization whilst keeping the kernel compatible with GPU execution.

we started to slowly switch to Kokkos when we developed and introduced the general idea and framework of the HPX-Kokkos integrations [7].

As mentioned in Section II, we have multiple kernels for each solver: The gravity solver contains `multipole` and `monopole` kernels (and various specializations that include angular momentum corrections (rho variant) and a specialization for the tree root subgrid to process remaining gravity interactions not handled by sub-grids down the tree. The hydro solver includes the `reconstruct` and the `flux` kernels.

Each kernel only operates on one sub-grid (and its ghost-layers) at a time. This means that during the solver execution, we traverse the tree, launching a multitude of different Kokkos compute kernels in the process. Using the HPX-Kokkos integration, these kernels can be launched (for different sub-grids) from multiple tasks simultaneously, with their execution status and results being integrated into the HPX execution graph using the futures returned by these asynchronous launches.

Depending on which device the Kokkos kernels use for execution, it will internally use either scalar types (on GPU) or types compiling down to the appropriate SIMD instructions (for example AVX512 for Intel) – provided the kernel has been implemented with the SIMD types. Of course, we can control which SIMD types are being used at compile time, meaning we can easily use scalar types on the CPU as well if we want to gauge the speedup we gain by using SIMD.

This model of an HPX application launching and synchronizing compute kernels using HPX / HPX-Kokkos, executing them via Kokkos on the correct device in the correct block configuration and, finally, having them adapted to the target device with SIMD types is exemplified in Figure 2.

IV. SOFTWARE CHANGES AND ADDITIONS FOR THIS WORK

Preparing Octo-Tiger to run on Fugaku, we contributed multiple changes recently to improve performance on CPU platforms generally, but on A64FX CPUs particularly.

A. Add SIMD types to Kokkos hydro kernels

At the start of this work, the hydro compute kernels (`reconstruct` and `flux`) within Octo-Tiger were using Kokkos, but did not yet contain SIMD types, as the amount of branching in these kernels proved to be problematic for the SIMD implementation in the past. We added a first implementation with the SIMD types to Octo-Tiger², reducing branching where possible and using SIMD masks where not possible. This implementation was later refined³ and is currently still in the process of being further optimized.

To improve the execution time of small compute kernels (very much like the two hydro kernels) launched with the Kokkos HPX execution space, we added an optimization that if a kernel is small enough to be launched as one task, it will be executed immediately on the same core (as to not waste the hot cache). We hope to upstream this change to Kokkos soon.

B. Integrate `std::experimental::simd` SIMD types

`std::experimental::simd` is an explicit vectorization API which was proposed for target Parallelism TS 2 for C++20 standard. `std::experimental::simd` is inspired by the `Vc` library (now deprecated) and provides support for explicit vectorization using data-parallel types. These are high-level C++ abstract types which are inlined to vector intrinsics of the underlying CPU architecture. This helps in writing portable high-level vectorized code using data-parallel types (and associated helper functions), without having to deal with intrinsics and low-level code, whilst still obtaining the same performance benefits from the vectorization.

Octo-Tiger supports vectorization using Kokkos SIMD and thus adding a `std::experimental::simd` vectorization backend was straightforward as it was just a change of the namespace from Kokkos SIMD to `std::experimental::simd`. Additionally, we only needed one wrapper for the `choose` function due to a difference in the method name in between the two frameworks. With this small integration⁴, we can essentially still use exactly the same Kokkos compute kernels, but instantiate them with the SIMD types from `std::experimental::simd` instead.

However, we have not completely switched to `std::experimental::simd` because the Kokkos SIMD types also work for GPU execution (being instanced as scalar double types). This does not seem to be supported yet in `std::experimental::simd`, as GPU device-side compilation with their scalar types fails. As we continue to use both frameworks, we are limited to using only the functionality of `std::experimental::simd` that has a counterpart in Kokkos SIMD (or has a fallback implementation we provide).

Integrating `std::experimental::simd` still opens our Kokkos kernels up to a more diverse set of possible SIMD backends (as we can now use all SIMD types within `std::experimental::simd` inside the Kokkos kernels).

²Part of <https://github.com/STELLAR-GROUP/octotiger/pull/426>

³<https://github.com/STELLAR-GROUP/octotiger/pull/428>

⁴<https://github.com/STELLAR-GROUP/octotiger/pull/403>

C. Add `std::experimental::simd-compatible` SVE types

Whilst the Kokkos SIMD contains SVE types itself already, these seem to rely on autovectorization (unlike different SIMD instructions sets in the library that are implemented with explicit vectorization). Furthermore, this particular SVE backend did not compile for Octo-Tiger initially, and only yielded scalar performance after having been modified to compile.

To still benefit from the performance boost SVE offers, we implemented a `std::experimental::simd-compatible` SVE backend (not relying on vectorization by the compiler) offering SVE SIMD types to our Octo-Tiger Kokkos kernels and to other applications.

Our new backend offers a portable high-level API which defines most of the functionalities of `std::experimental::simd` with the help of the ARM C Language Extensions (ACLE) for SVE defined in `arm_sve.h`. This is a standard header provided by ARM[®] that contains all the types and functions for SVE vectorization currently supported by GCC and Clang. For benefiting from SVE vectorization, we used explicit vectorization with the help of fixed length vector registers defined by the compiler. The vector length is set during compile time using the flag `-msve-vector-bits=N`.

Integrating this SVE backend within Octo-Tiger is straightforward given our `std::experimental::simd` integration: One only needs to include the header and change the SIMD namespace accordingly. Then the Kokkos kernels can use our new SVE types. Using our experimental implementation of the SVE types enables Octo-Tiger to use SVE in all Kokkos compute kernels.

We made our new SVE backend available on GitHub⁵. Notably, it can be used independent of Octo-Tiger in other applications. These applications simply need to use its SIMD SVE types to achieve explicit vectorization on A64FX CPUs.

V. RESULTS

We use three different CPU platforms for our tests:

- A two-socket AMD[®] EPYC[™] 7H12 CPU @ 2.60GHz with 128 cores.
- A two-socket Intel[®] Xeon[®] Platinum 8358 CPU @ 2.60GHz with 64 cores (Intel[®] Icelake node).
- A Fujitsu A64FX[™] CPU @ 1.80 GHz on Stony Brook University's Ookami cluster with 48 cores (ARM[®] node).

It is worth mentioning that the AMD and Intel nodes can and will run at higher frequencies as their boost is enabled. We will run three different tests on each CPU. For all tests, we use double precision.

For brevity, we use `SES` to stand for `std::experimental::simd` in the following. For example, with `SES AVX512` the Kokkos kernels use the AVX512 types within `std::experimental::simd`, with `KOKKOS NEON` they use the NEON types within the Kokkos SIMD library instead. `SES SVE` makes them use our new SVE types.

⁵<https://github.com/srinivasyadav18/sve>

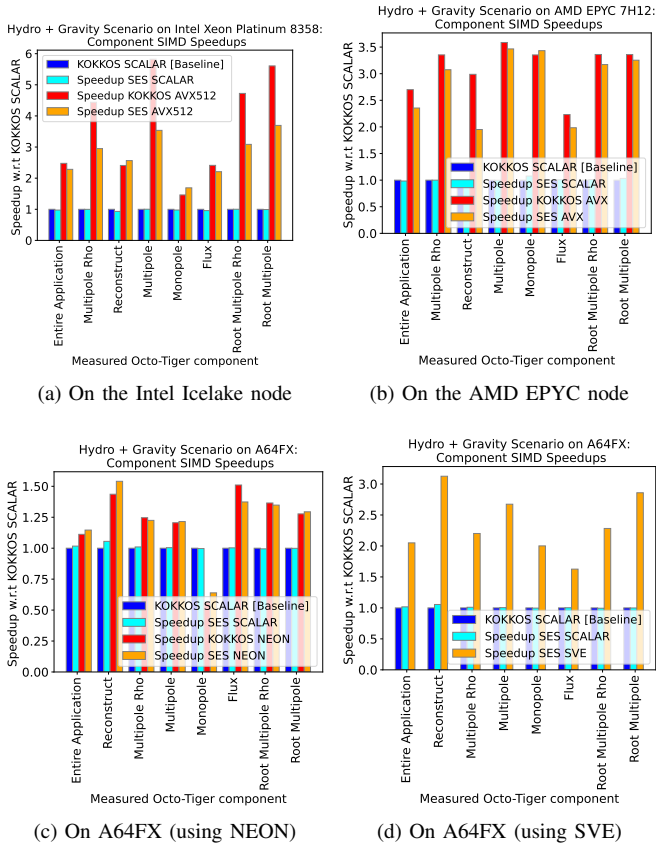


Fig. 3. Single Core SIMD Speedups with respect to the runtime with scalar types. The speedups are shown for both the entire application and for the relevant compute kernels.

A. Test 1: Single-Core SIMD Speedups

First, we take a look at the single-core speedups using SIMD. For this, we use the “rotating star” scenario [1], simulating a rotating star in equilibrium. This is one of Octo-Tiger’s basic tests, as it uses both the gravity and the hydro solver, helping to uncover even small errors if the simulation runs long enough. In the context of this work, it is a useful benchmark, as it includes both solvers, and the additional work we would encounter in a production scenario (tree-traversals, exchange of ghost cells, determination of time-step size, and so on). We disable the adaptive refinement here, to make the runtimes more comparable to a later run of a hydro-only scenario: Without AMR, both runs build a full tree for a given tree depth, resulting in the same amount of overall cells. We choose a low tree level to only have 512 leaf sub-grids, as this is realistic for large distributed runs, where we also only have a limited number of sub-grids available per compute node.

Given this scenario, we would like to know how both the total application speedup, as well the speedups for all major compute kernels when using SIMD. By using APEX [16], we get the mean runtime for each compute kernel at the end of the run. Therefore, we simply need to compile Octo-Tiger with the SIMD (or scalar) types we want to use and run the simulation to obtain the runtime data. Over successive

runs with different types, we can easily get the speedups for both the entire application and the compute kernels themselves this way. The baseline for our speedups will be the runtime using the Kokkos scalar types (KOKKOS SCALAR). Note that we do not disable the compiler autovectorization, as we want a fair estimate of the benefit of our explicit vectorization using SIMD types compared to what the compiler is able to autovectorize.

Besides these baseline runtimes, we run with the `std :: experimental :: simd` scalar types (SES SCALAR). They should give no speedup, but can be useful to uncover unexpected bottlenecks. Finally, we run with the appropriate SIMD types on each CPU, using both `std :: experimental :: simd` and the Kokkos SIMD types, calculating the speedup using the previous runtimes of the Kokkos SCALAR runs.

Figure 3 shows the speedups for each compute kernel and the entire application runtime when switching to the SIMD types. Each set of bars in a subfigure represents on Octo-Tiger component, with the bars themselves representing the speedup using one of the possible types. There are four subfigures, one for the Intel CPU (Fig. 3a), one for the AMD CPU (Fig. 3b), and two for the A64FX CPU. Here, on A64FX, we test the SIMD implementation both with NEON (Fig. 3c) and with our new experimental SVE types (Fig. 3d). While there are also SVE types within the Kokkos SIMD library, we were unable to make them work beyond scalar performance. Hence, these Kokkos SVE types are missing in Figure 3d, and we only use the experimental SES SVE types introduced in this work.

Overall, we get a noticeable SIMD speedup for our real-world application. While not yet as good as the more compute-intensive gravity compute kernels (i.e. multipole and others), the two new hydro SIMD kernels `flux` and `reconstruct` still provide noticeable speedups, especially on the A64FX machine.

On one core of the Intel Icelake, the Kokkos SIMD types (KOKKOS AVX512) perform consistently better than `std :: experimental :: simd` (SES AVX512), which exemplifies that differences in the backend implementation can lead to significantly different performance outcomes. Using the Kokkos SIMD types for AVX and AVX512 is generally preferable, at least for our use-case. However, the `std :: experimental :: simd` interface and the SVE types we have added (SES SVE) are essential on the A64FX machine, giving us a good overall application speedup as we can now use explicit SVE vectorization.

Unfortunately, the speedups we obtained fell short of perfect SIMD speedups. However, given that our compute kernels are many hundred lines of code each, and this is a work in progress, the results are still promising (and useful for our production-runs as they already reduce the amount of compute time). Furthermore, the flexibility of being able to simply switch between SIMD backends using types is already extremely useful here. Not only to be able to run on all these different platforms without changing any code to begin with, but also to switch to `std :: experimental :: simd` where required (i.e. on the A64FX node).




Fig. 4. Parallel efficiency and node-level scaling, going from one to all cores of the respective three platform. Each graph includes one run with the SIMD types for the current platform (two graphs for A64FX to show both NEON and SVE), as well the same run using the scalar types for comparison.

B. Test 2: Node-Level Scaling and Parallel Efficiency

Moving beyond one core, we now take a look at the node-level scaling. For this, we are once again using the same Octo-Tiger scenario as for the previous test, looking at the overall computation time for 10 timesteps (each timestep consists of 6 gravity and 3 hydro solver iterations). We are using the SIMD types with the best results from the previous test on each machine: Kokkos AVX512 on Intel, Kokkos AVX on AMD, and SES SVE on the A64FX ARM node. We further include a run using Kokkos NEON on this ARM node for comparison. For all runs, we also compare with scalar runs, to see if (and how) the SIMD speedup is influenced when scaling to the complete node: The SIMD speedup can change for a multitude of reasons (for example: frequency changes, work starvation, memory bottlenecks).

The results are mixed and can be seen in Figure 4. The overall best runtime is achieved on the Intel Icelake machine (Fig. 4a), reaching 3.65s. Runtimes on the A64FX machine improved drastically with our new SVE backend (Fig. 4d), both over using the scalar types and over using the Kokkos NEON types (Fig. 4c). However, the performance still falls short of the Icelake CPU. While this is to be expected to a certain degree given the lower processor frequency and fewer cores, the difference is noticeable. Looking at the APEX runtime data and comparing them to runs on different machines, it seems that every component in Octo-Tiger runs a bit slower, with no obvious new hotspot on the A64FX CPU. The slowdown is most severe in the gravity solver, though, probably due to the better SIMD speedups we get using AVX512. Furthermore, the frequency penalty on the Intel CPU when using AVX512 is lower than we expected; in fact, the CPU seems to run faster than the base frequency

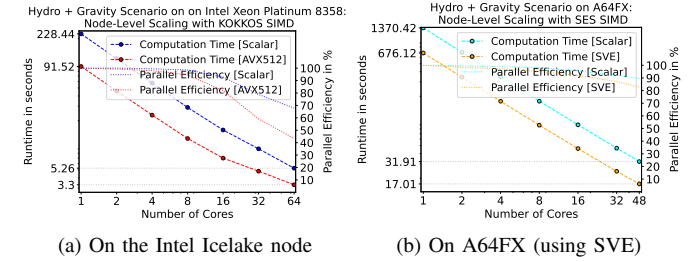


Fig. 5. Improving the overall A64FX and Intel runtimes shown in Figure 4 by using more tasks per multipole kernel launch.

(staying in a range of 2.6-3.2 Ghz during test runs with our scenario).

Given that our compute kernels are working on one sub-grid each, we usually just launch one HPX task per Kokkos kernel, as this has been shown to work best with Octo-Tiger in previous work [7]. However, these past tests were done on CPUs with lower core counts. On the CPUs used in this work with larger core counts, we can improve the performance by using more tasks, thus having more, albeit smaller work packages available during the tree traversal, which helps to keep all cores busy. As we use the HPX execution space within Kokkos, this is a simple change in the kernel launch where we have to specify a different tiling configuration. Hence, we can easily split the multipole kernel (the most compute-intensive compute kernel per sub-grid) into more tasks and see how the performance is impacted by this. By splitting this multipole kernel into more than one task per sub-grid using the Kokkos HPX execution space (into 16 tasks to be precise), we can achieve a better runtime, as can be seen in Figure 5. Here, we can reduce the overall computation runtime of the scenario from 20.67s (see Fig. 4d) to 17.01s (see Fig. 5b) on the A64FX ARM node just with this multipole change. This also helped on the Intel CPU, reducing the runtime to 3.3s (Fig. 5a). However, this yielded no noticeable benefit on the AMD EYPC machine, as the bottleneck here appears to be different.

Scalability on the AMD EPYC machine (Fig. 4b) proved to be troublesome: Beyond 64 cores, the runtime deteriorates. At first, we suspected that given our small scenario, the CPU is simply starved, but this is not the case; we see a similar parallel efficiency when running a larger scenario. Looking again at the mean kernel runtimes themselves instead, we noticed that the mean runtime of some kernels stay the same when moving to 128 cores (as it should be, since each kernel is executed by one core by default). However, the mean runtime for other kernels (notably the `flux` kernel) increases. This might be caused by inter-socket memory accesses hitting a bandwidth limit; however, at the time of submission, the profiling effort here is still ongoing and subject to future work. For now, the EPYC shows Octo-Tiger’s current node-level scaling limit.

To summarize, using our SIMD changes, the node-level scaling looks acceptable on the Intel and good on the ARM node. Changing the launch configuration of the Multipole Kokkos kernel to use more tasks yielded additional benefits.

C. Test 3: Hydro-only Speedups

As a last test, we briefly investigate the performance with a hydro-only scenario, as the most recent change to the compute kernels was the introduction of the SIMD types to the `flux` and `reconstruct` kernels. This ensures that the only SIMD speedup we will see in this test will be due to our new implementation of those kernels, and not due to some existing SIMD implementation in the gravity solver.

We use the Sedov Blast Wave scenario, another one of Octo-Tiger’s standard tests, as this scenario also has an analytical solution. As we turned off the AMR, this scenario again contains 512 leaf sub-grids, leading to the same overall grid as we had for the previous ”rotating star” tests. This ensures that the computation times shown here should be relatable to the time the hydro solver part required in the previous two tests. The only parts with explicit SIMD vectorization in this hydro test are the `reconstruct` and `flux` kernels, the rest of the solver and surrounding application do not benefit from switching to the SIMD types. Consequently, the computation time speedups themselves are lower than the speedups of the individual kernels as seen previously in Figure 3. We further include the old non-Kokkos hydro implementation here, to make sure that there are no performance regressions. This previous hydro implementation was only relying on HPX for multicore usage, but was incompatible with GPUs, hence the conversion to Kokkos in the first place.

The results of this test can be found in Figure 6. As expected, the best runs were those with SIMD types. There is no significant penalty when switching from the legacy kernels to Kokkos, even when just using the scalar types. Using SIMD the best overall speedup over the old implementation is reached on the A64FX machine. On the Intel Icelake CPU, we took a look at the APEX data to investigate the speedup. Here, we spend a lot of time in the handling of the ghost cells and comparatively lesser percentage of the runtime in the compute kernels, resulting in the lower overall SIMD speedup shown here.




Fig. 6. Running the Hydro-Only Sedov Blast Wave scenario using legacy hydro implementation, and the Kokkos hydro implementation using either scalar types or SIMD types (AVX512, AVX and SVE respectively). Left is the computation runtime on only one core, on the right is the same scenario when scaled to 64 cores (48 on A64FX).

VI. RELATED WORK

A. Asynchronous many-task systems (AMT)

HPX uses light-weight user threads to expose concurrency on top of system threads. Other notable solutions with this approach are: Chapel [18], Charm++ [19], PaRSEC [20], Uintah [21] and Legion [22]. For a more detailed comparison of AMT’s, like Cilk Plus, OpenMP, Intel TBB, Qthreads, StarPU, GASPI, Chapel, Charm++ and HPX, we refer to [23].

B. C++ Performance Portability

There are several options for performance portability using C++. First, there is Raja [24] providing CUDA, HIP, OpenMP and TBB backends. Second, there is SYCL providing hipSYCL, neoSYCL, triSYCL and TBB. Finally, there is Thrust providing CUDA, TBB and OpenMP. For a comparison between Kokkos and SYCL, we refer to [25]. However, none other than Kokkos provides an HPX backend.

C. SIMD

To achieve vectorization, one can rely on the auto-vectorization provided by the compiler [26], [27]. However, this is more like a black box, and the programmer does not have full control. More user control is provided by explicit SIMD, e.g. UME::SIMD [28], Vc [29] or `std::simd`. Here, the programmer uses types, like `std::experimental::simd`, which represent vector registers. On the other hand, there is implicit SIMD, where the programmer adds directives (such as pragmas) to indicate if code should get vectorized by the compiler [30]. A well known example of this is are the SIMD pragmas in OpenMP. We use explicit SIMD to have the most control.

D. SVE

SVE is the new vectorization standard [31] introduced in ARM v8 in the Fugaku Fujitsu Supercomputer. This follows the VLA (vector length agnostic) programming model, where the vector length is set dynamically from 128 bits to 2048 bits in increments of 128 bits. SVE vectorization can be utilized in multiple ways as well, by using compiler auto-vectorization or by using explicit vectorization using fixed sized SIMD types. For more details, we refer to [32].

VII. CONCLUSION

The presented execution model, a combination of HPX, Kokkos and explicit vectorization, proves to be versatile in its use in Octo-Tiger. We can now easily switch both the SIMD library (Kokkos SIMD or `std::experimental::simd`) and the used SIMD extensions. Moreover, adjusting the number of HPX tasks that a Kokkos kernels gets split into further helps in coping with larger CPU core counts during the FMM tree traversals.

We tested this execution model (and our changes) on three CPU architectures using Octo-Tiger as a real-world benchmark. Overall, Octo-Tiger benefits from the changes presented here: Introducing the new SIMD implementation for the two hydro kernels, we experienced a speedup on all platforms. On

A64FX, our target platform for the changes presented here, we measured the largest improvements. Here, switching the SIMD types to our new SVE types and increasing the number of HPX tasks clearly improved the performance. Both of these options can be simply changed during compile time. This allows us to run more efficiently on A64FX using SIMD instead of the scalar types. However, taking a closer look at the node-level scaling, we unfortunately uncovered a pre-existing bottleneck on the AMD EPYC platform when going from 64 to 128 cores.

It is noteworthy that the presented execution approach is application-independent. The added SVE SIMD types in particular can be useful for other applications and will benefit the community beyond Octo-Tiger.

Taking a look at Octo-Tiger's current development snapshot in this work, for our next steps, we plan to examine the node-level scalability limit observed on the AMD EPYC node, and to improve the multitude of compute kernels in Octo-Tiger with SVE to achieve better SIMD speedups for each of them. Afterwards, we plan to run Octo-Tiger on NERSC's Permuter, Riken's Fugaku, and CSC's LUMI to gather more performance measurements, both node-level and distributed, on a diverse set of machines.

VIII. ACKNOWLEDGMENT

The authors would like to thank Stony Brook Research Computing and Cyberinfrastructure, and the Institute for Advanced Computational Science at Stony Brook University for access to the innovative high-performance Ookami computing system, which was made possible by a \$5M National Science Foundation grant (#1927880).

COPYRIGHT NOTICE

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

REFERENCES

- [1] D. C. Marcello *et al.*, "Octo-Tiger: a new, 3D hydrodynamic code for stellar mergers that uses HPX parallelization," *Monthly Notices of the Royal Astronomical Society*, vol. 504, no. 4, pp. 5345–5382, 2021.
- [2] H. Kaiser *et al.*, "HPX-the C++ standard library for parallelism and concurrency," *Journal of Open Source Software*, vol. 5, no. 53, p. 2352, 2020.
- [3] T. Heller *et al.*, "Harnessing billions of tasks for a scalable portable hydrodynamic simulation of the merger of two stars," *The International Journal of High Performance Computing Applications*, vol. 33, no. 4, pp. 699–715, 2019.
- [4] G. Daiß *et al.*, "From Piz Daint to the stars: Simulation of stellar mergers using high-level abstractions," in *Proceedings of the international conference for high performance computing, networking, storage and analysis*, 2019, pp. 1–37.
- [5] P. Diehl *et al.*, "Octo-Tiger's New Hydro Module and Performance Using HPX+ CUDA on ORNL's Summit," in *2021 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2021, pp. 204–214.
- [6] D. Sahasrabudhe *et al.*, "A portable SIMD primitive using Kokkos for heterogeneous architectures," in *International Workshop on Accelerator Programming Using Directives*. Springer, 2019, pp. 140–163.

- [7] G. Daiß *et al.*, "Beyond fork-join: Integration of performance portable Kokkos kernels with HPX," in *2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2021, pp. 377–386.
- [8] A. Kurganov *et al.*, "New High-Resolution Central Schemes for Non-linear Conservation Laws and Convection-Diffusion Equations," *Journal of Computational Physics*, vol. 160, no. 1, pp. 241–282, 2000.
- [9] D. C. Marcello, "A Very Fast and Angular Momentum Conserving Tree Code," *Astronomical Journal*, vol. 154, p. 92, 2017.
- [10] D. Pfander *et al.*, "Accelerating Octo-Tiger: Stellar mergers on Intel Knights Landing with HPX," in *Proceedings of the International Workshop on OpenCL*, ser. IWOCCL '18. New York, NY, USA: ACM, 2018, pp. 19:1–19:8.
- [11] G. Daiß, "Octo-Tiger: Binary star systems with HPX on Nvidia P100," Master's thesis, 2018.
- [12] P. Diehl *et al.*, "Performance Measurements Within Asynchronous Task-Based Runtime Systems: A Double White Dwarf Merger as an Application," *Computing in Science & Engineering*, vol. 23, no. 3, pp. 73–81, 2021.
- [13] H. Kaiser *et al.*, "STELLAR-GROUP/hpx: HPX V1.8.1: The C++ Standards Library for Parallelism and Concurrency," Jul. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.598202>
- [14] —, "Higher-level Parallelization for Local and Distributed Asynchronous Task-based Programming," in *Proceedings of the First International Workshop on Extreme Scale Programming Models and Middleware*, ser. ESPM '15. New York, NY, USA: ACM, 2015, pp. 29–37.
- [15] T. Heller *et al.*, "Closing the Performance Gap with Modern C++," in *High Performance Computing*, ser. Lecture Notes in Computer Science, M. Tauber *et al.*, Eds., vol. 9945. Springer International Publishing, 2016, pp. 18–31.
- [16] K. Huck *et al.*, "An Early Prototype of an Autonomic Performance Environment for Exascale," in *Proceedings of the 3rd International Workshop on Runtime and Operating Systems for Supercomputers*, ser. ROSS '13. New York, NY, USA: Association for Computing Machinery, 2013. [Online]. Available: <https://doi.org/10.1145/2491661.2481434>
- [17] C. R. Trott *et al.*, "Kokkos 3: Programming Model Extensions for the Exascale Era," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 4, pp. 805–817, 2022.
- [18] B. L. Chamberlain *et al.*, "Parallel programmability and the Chapel language," *The International Journal of High Performance Computing Applications*, vol. 21, no. 3, pp. 291–312, 2007.
- [19] L. V. Kale *et al.*, "Charm++ a portable concurrent object oriented system based on C++," in *Proceedings of the eighth annual conference on Object-oriented programming systems, languages, and applications*, 1993, pp. 91–108.
- [20] G. Bosilca *et al.*, "Parsec: Exploiting heterogeneity to enhance scalability," *Computing in Science & Engineering*, vol. 15, no. 6, pp. 36–45, 2013.
- [21] J. D. d. S. Germain *et al.*, "Uintah: A massively parallel problem solving environment," in *Proceedings the Ninth International Symposium on High-Performance Distributed Computing*. IEEE, 2000, pp. 33–41.
- [22] M. Bauer *et al.*, "Legion: Expressing locality and independence with logical regions," in *SC'12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2012, pp. 1–11.
- [23] P. Thoman *et al.*, "A taxonomy of task-based parallel programming technologies for high-performance computing," *The Journal of Supercomputing*, vol. 74, no. 4, pp. 1422–1434, 2018.
- [24] D. A. Beckingsale *et al.*, "RAJA: Portable performance for large-scale scientific applications," in *2019 IEEE/ACM international workshop on performance, portability and productivity in hpc (P3HPC)*. IEEE, 2019, pp. 71–81.
- [25] J. R. Hammond *et al.*, "A comparative analysis of Kokkos and SYCL as heterogeneous, parallel programming models for C++ applications," in *Proceedings of the International Workshop on OpenCL*, 2019, pp. 1–2.
- [26] D. Naishlos, "Autovectorization in GCC," in *Proceedings of the 2004 GCC Developers Summit*. Citeseer, 2004, pp. 105–118.
- [27] H. Finkel, "Autovectorization with LLVM," in *The LLVM Compiler Infrastructure 2012 European Conference*, 2012.
- [28] P. Karpinski *et al.*, "A high-performance portable abstract interface for explicit SIMD vectorization," in *Proceedings of the 8th International*

Workshop on Programming Models and Applications for Multicores and Manycores, 2017, pp. 21–28.

- [29] M. Kretz *et al.*, “Vc: A C++ library for explicit vectorization,” *Software: Practice and Experience*, vol. 42, no. 11, pp. 1409–1430, 2012.
- [30] A. Pohl *et al.*, “An evaluation of current SIMD programming models for C++,” in *Proceedings of the 3rd Workshop on Programming Models for SIMD/Vector Processing*, 2016, pp. 1–8.
- [31] arm, “Introduction to SVE (Version 1),” <https://developer.arm.com/documentation/102476/latest/>, Release information, 2022.
- [32] N. Stephens *et al.*, “The ARM Scalable Vector Extension,” *IEEE Micro*, vol. 37, no. 2, pp. 26–39, 2017.

APPENDIX A

ARTIFACT DESCRIPTION APPENDIX: FROM MERGING FRAMEWORKS TO MERGING STARS: EXPERIENCES USING HPX, KOKKOS AND SIMD TYPES

A. Abstract

This description contains the instructions needed to run the experiments described in the ESPM2 2022 Paper “From Merging Frameworks to Merging Stars: Experiences using HPX, Kokkos and SIMD Types”. Here, we briefly describe the utilized real-world application (Octo-Tiger), its required dependencies and the respective versions that we used to obtain the results in the paper. We further describe the utilized hardware and Octo-Tiger scenarios. Additionally, we provide details regarding how to build Octo-Tiger (and all required dependencies), where to find the required experiment scripts and how to run them to recreate our results (and plots).

B. Description

1) Check-list (artifact meta information):

- **Algorithm:** Octo-Tiger primarily uses the **Fast-Multipole Method for the gravity solver and finite volumes for the hydro solver**
- **Program:** Octo-Tiger
- **Compilation:** g++11.2 (buildscripts available)
- **Hardware:** Intel Xeon CPU, AMD EYPC CPU, A64FX CPU
- **Output:** Python script to recreate the plots is available
- **Experiment workflow:** Iterate over cores and SIMD extensions (Bash scripts to re-run these experiments are available)
- **Experiment customization:** Modified Kokkos and Octo-Tiger (commits and patches are available)
- **Publicly available?:** Yes (GitHub)

2) *How software can be obtained (if available):* Octo-Tiger can be obtained on GitHub:

<https://github.com/STELLAR-GROUP/octotiger>

We used two Octo-Tiger commits:

- [2461163996c576a5d43b4cd1a7f8c295d446f925](https://github.com/STELLAR-GROUP/octotiger/commit/2461163996c576a5d43b4cd1a7f8c295d446f925) (for using more tasks in the multipole kernels to generate Figure 5)
- [9437e172264a4c7b3f24ec4ccd4c380feba26755](https://github.com/STELLAR-GROUP/octotiger/commit/9437e172264a4c7b3f24ec4ccd4c380feba26755) (for all other experiments)

3) *Hardware dependencies:* We run our experiments (tests) on three different machines (nodes):

- 1) One Fujitsu A64FX™ CPU @ 1.80 GHz on Ookami (48 cores).
- 2) One university server (at LSU) containing a two-socket Intel® Xeon® Platinum 8358 CPU @ 2.60GHz (64 cores)
- 3) One university server (at LSU) containing a two-socket AMD® EPYC™ 7H12 CPU @ 2.60GHz (128 cores)

Memory requirements for the scenarios we use are minimal and should not exceed 4 GB for each scenario itself.

4) *Software dependencies:* Octo-Tiger has multiple, mandatory software dependencies:

- HPX: 1.8.0
- Kokkos (on the Intel and AMD nodes):
[2640cf70de338618a7e4fe10590b06bc1c239f4c](https://github.com/STELLAR-GROUP/hpx-kokkos)
- Kokkos (on Ookami):
[596bb0b1b](https://github.com/STELLAR-GROUP/hpx-kokkos)
- HPX-Kokkos:
[20a44967c742f5a7670b4dff9658d9973bf849f2](https://github.com/STELLAR-GROUP/hpx-kokkos)
URL: <https://github.com/STELLAR-GROUP/hpx-kokkos>
- CPPuddle (on the Intel and AMD nodes):
[6127562897dc2940869d744d490d0eb7b6fa37bc](https://github.com/STELLAR-GROUP/hpx-kokkos)
- CPPuddle (on Ookami):
[f1eed375685981b59723e5592961c5c774789a20](https://github.com/STELLAR-GROUP/hpx-kokkos)
URL: <https://github.com/STELLAR-GROUP/hpx-kokkos>
- Vc: 1.4.1
- Jemalloc: 5.2.1
- Silo: 4.10.2
- HDF5: 1.8.12
- Boost: 1.75 (on the Intel and AMD nodes)
- Boost: 1.78 (on Ookami)
- SVE types library:
[e3b2fb8d7bfda5d6eda90efee29fafdcbe895a25](https://github.com/srinivasyadav18/sve)
URL: <https://github.com/srinivasyadav18/sve>

All dependencies are available online! We have added links for the lesser known ones though, to ease reproducibility. We use slightly different versions (as outlined in the list above) on Ookami to get the toolchain compiled on this different system. For ease of use, there are buildscripts available at this URL: <https://github.com/STELLAR-GROUP/OctoTigerBuildChain>

We use two additional patchfiles, one for Kokkos and one for HPX-Kokkos. Kokkos has been customized based on the commits specified here, to allow the single-task optimization mentioned in Section IV-A). This has not yet been upstreamed. However, the buildscripts contain the patchfile for this (kokkos-single-task.patch) and apply the patch automatically during building. Accordingly, HPX-Kokkos requires a patch adapting to this which is also included in the buildscripts (sync.patch) and applied automatically for building. These changes will be upstreamed eventually, but for the given commits above these two patches are required.

5) *Datasets:* No additional datasets need to be downloaded! The inputfile required for the “rotating star” scenario can be created using the binary `gen_rotating_star_init` which will be build automatically when building Octo-Tiger (to be found inside the tools subdirectory within the Octo-Tiger build). Further configuration files (blast.ini and rotating_star.ini) are part of the Octo-Tiger source repository (in the test_problems subdirectory) and supply the exact configuration for each scenario.

Using these configuration files, the rotating star scenario can be started with the following runtime parameters (adapt the paths to the ini and binary files if required):

```
$ ./build/octotiger/build/octotiger \  
--config_file=rotating_star.ini \  
--unigrid=1 --max_level=3 --stop_step=10 \  
--correct_am_hydro=0 --theta=0.34 --disable_output=1
```

Respectively, for the Sedov Blast Wave scenario:

```
$ ./build/octotiger/build/octotiger \  
--config_file=blast.ini \  
--unigrid=1 --max_level=3 --stop_step=10 \  
--correct_am_hydro=0 --disable_output=1
```

For each scenario, we additionally configure whether the KOKKOS or LEGACY kernels are used with the following parameters:

```
--hydro_host_kernel_type=KOKKOS  
--multipole_host_kernel_type=KOKKOS  
--monopole_host_kernel_type=KOKKOS
```

Lastly, the number of utilized worker threads (and thus CPU cores) can be steered with the parameter `--hpx:threads`.

C. Installation

There are buildscripts available to ease building Octo-Tiger with the software dependencies given earlier:

<https://github.com/STELLAR-GROUP/OctoTigerBuildChain>

These scripts will automatically download and build the required dependencies, as well as Octo-Tiger itself.

1) Configuring on the Intel Node:

- To install on the LSU Intel machine, use the buildscripts commit: `ca909534c3c93f7fddb77f48d15c42dba558f0e0` (or the branch `espm2_icelake_build`)
- Make sure the modules `gcc/11.2.1` and `hwloc 2.4.1` are loaded!

2) Configuring on the AMD Node:

- To install on LSU AMD machine, use the buildscripts commit: `a53f1413105dd543560e30c77c710d9f32513fa0` (or the branch `espm_epyc`)
- Make sure the modules `gcc/11.2.1` and `hwloc 2.4.1` are loaded!

3) Configuring on the Oookami A64FX Node:

- To install on Oookami, use the buildscripts commit: `6c9f6361c0942cb5ceb2dde986c42614cc1102c0` (or the branch `espm2_oookami`)
- Make sure the modules `gcc/11.2.0` `cmake/3.22.1` and `hwloc/2.4.1` are loaded!

4) Building (on all machines):

- First, load the modules and configure the buildscripts as outlined in the last three subsections (depending on your machine).
- Afterwards, run the following command within the root directory of the buildscripts:

```
$ ./build-all.sh Release with-CC \  
without-cuda without-mpi without-papi \  
with-apex with-kokkos with-simd \  
with-hpx-backend-multipole \  
with-hpx-backend-monopole \  
with-hpx-cuda-polling without-otf2 \  
boost hdf5 silo jemalloc vc hpx \  
kokkos cppuddle octotiger
```

5) Changing the SIMD types:

- To switch SIMD types and libraries manually, the following Octo-Tiger CMAKE flags can be used (modify them in `build-octotiger.sh` in the buildscripts):

- `OCTOTIGER_KOKKOS_SIMD_LIBRARY`
- `OCTOTIGER_KOKKOS_SIMD_EXTENSION`

- Note: The experiment bash scripts will do this automatically, so there should be no need to do this manually!

D. Experiment workflow

The bash scripts to perform the experiments on the Intel and AMD node can be found on GitHub: <https://github.com/G-071/octotiger-performance-tests/tree/master/rostam/node-level-scaling-tests>

The experiment `test_cpu_performance.sh` gathers all the runtime data when being launched (after building Octo-Tiger) within the root directory of the buildscripts. This experiment iterates over all core and SIMD configurations, generating a (csv) log file containing the runtime data for the application and the compute kernels. There are multiple scenario files available within the GitHub repository. These scenario files mostly contain the core configuration and specify the SIMD types to be used on each machine. To make the experiments work, one needs to pass one of these scenario files (the rotating star one or the blast wave one for our experiments) to the bash script as the first parameter. Afterwards, run the second experiment script `test_hydro_speedup.sh` (also within the buildscripts root directory) to obtain the hydro-only runtime data required for the last figure in the paper.

Due to a change in the output format (requiring adaptations to the bash scripts), the respective scripts for Oookami can also be found on GitHub but in a different directory: <https://github.com/G-071/octotiger-performance-tests/tree/master/oookami/node-level-scaling-tests>

However, the process of running the tests is the same as for the Intel and AMD nodes otherwise.

E. Evaluation and expected result

To evaluate the data and plot the results, there is a python script (`plot_node_level_scaling.py`) available at: <https://github.com/G-071/octotiger-performance-tests/tree/master/rostam/node-level-scaling-tests/plot>

This python script can be used to generate the graphs in the paper, using the runtime data obtained in the last subsection! The script requires `pandas`, `numpy` and `matplotlib` to be installed. To run it, one needs to pass the file containing the csv results from a machine via the parameter `--filename` (pass the output file generated by `test_cpu_performance.sh`). Additionally, pass the utilized SIMD extension via `--simd_key` (for example `AVX512`). This should plot both the node-level scaling graphs (Figure 4) and the SIMD speedup graphs (Figure 3) and immediately store them as PDFs in the same directory. To plot the last figure (Figure 6), the script further expects the hydro-only results (obtained by running `test_hydro_speedup.sh`) to be present in

the following three files: `icelake_legacy_test.data`,
`epyc_legacy_test.data`, `arm_legacy_test.data`
The expected results are plots similar to what is shown in the
paper!

F. Experiment customization

Kokkos gets modified with a patch as outlined earlier. To
generate Fig. 5, a different Octo-Tiger commit (as mentioned)
has to be used before running the experiment scripts!