

# Advances in ArborX to support exascale applications

Andrey Prokopenko<sup>1</sup> , Daniel Arndt<sup>1</sup> ,  
Damien Lebrun-Grandié<sup>1</sup> , Bruno Turcksin<sup>1</sup>,  
Nicholas Frontiere<sup>2</sup> , J.D. Emberson<sup>2</sup>  and  
Michael Buehlmann<sup>2</sup> 

The International Journal of High  
Performance Computing Applications  
2025, Vol. 39(1) 167–176  
© The Author(s) 2024  
Article reuse guidelines:  
[sagepub.com/journals-permissions](https://sagepub.com/journals-permissions)  
DOI: 10.1177/10943420241298296  
[journals.sagepub.com/home/hpc](https://journals.sagepub.com/home/hpc)



## Abstract

ArborX is a performance portable geometric search library developed as part of the Exascale Computing Project (ECP). In this paper, we explore a collaboration between ArborX and a cosmological simulation code HACC. Large cosmological simulations on exascale platforms encounter a bottleneck due to the in-situ analysis requirements of halo finding, a problem of identifying dense clusters of dark matter (halos). This problem is solved by using a density-based DBSCAN clustering algorithm. With each MPI rank handling hundreds of millions of particles, it is imperative for the DBSCAN implementation to be efficient. In addition, the requirement to support exascale supercomputers from different vendors necessitates performance portability of the algorithm. We describe how this challenge problem guided ArborX development, and enhanced the performance and the scope of the library. We explore the improvements in the basic algorithms for the underlying search index to improve the performance, and describe several implementations of DBSCAN in ArborX. Further, we report the history of the changes in ArborX and their effect on the time to solve a representative benchmark problem, as well as demonstrate the real world impact on production end-to-end cosmology simulations.

## Keywords

Geometric search, clustering, kokkos, GPU, DBSCAN, cosmology

## 1. Introduction

ArborX is a performance portable geometric search library (Lebrun-Grandié et al., 2020). ArborX was developed as part of the Exascale Computing Project (ECP) (Messina, 2017), a multi-year US Department of Energy (DOE) program. The ECP aimed to provide an exascale computing ecosystem for DOE mission-critical applications from many different domains, including cosmology, combustion, material science, and additive manufacturing.

The complexity and the size of the DOE applications requires enormous computational resources. DOE funded supercomputers, such as Frontier (OLCF, 2022), Perlmutter (NERSC, 2022), and Aurora (ALCF, 2023), differ in their hardware architectures, providing GPU accelerators from distinct vendors (AMD, Nvidia, and Intel, respectively). Given that more than 95% of the performance of each system is coming from these accelerators, it is critical for a software library to be performance portable.

The primary goal of ArborX is to support DOE applications at scale. Initially developed within the Data-TransferKit library (Slattery et al., 2013), ArborX became a standalone library in early 2019 when it became clear that many other applications could benefit from its functionality. Since then, ArborX experienced a significant growth in the number of users and the diversity of use cases. ArborX now includes a wide variety of algorithms that depend on the spatial proximity in the data: range and nearest searches, clustering, ray tracing, and interpolation.

<sup>1</sup>Oak Ridge National Laboratory, Oak Ridge, TN, USA

<sup>2</sup>CPS Division, Argonne National Laboratory, Lemont, IL, USA

### Corresponding author:

Andrey Prokopenko, Oak Ridge National Laboratory, 1 Bethel Valley Rd,  
Oak Ridge, TN 37831, USA.

Email: [prokopenkoav@ornl.gov](mailto:prokopenkoav@ornl.gov)

In this paper, we explore a collaboration of ArborX with a selected exascale cosmology application that occurred over the 2020 – 2023 period.

The Hardware/Hybrid Accelerated Cosmology Code (HACC; Habib et al. (2016); Frontiere et al. (2023)) is an extreme-scale cosmological simulation framework. The HACC challenge problem includes simulations with over ten trillion particles to produce the most detailed synthetic sky maps ever made. In-situ analysis is crucial, as these simulations would generate more than 100PB of data. Consequently, the analysis suite has to scale, along with the code. ArborX helps to address computational bottlenecks in several analysis components. In this paper, we will focus on one such area: identification of clusters within the full particle set.

The paper is organized as follows. We present the importance and the impact of ArborX improvements on the performance of the analysis code in HACC in Section 2. We provide the background information on the algorithm and ArborX library in Section 3. We then describe performance and algorithm improvements made in ArborX to achieve these results in Section 4. Finally, we provide directions for further growth in Section 5.

## 2. ArborX impact on production simulations with HACC

Identification of halos (regions with a high density of dark matter particles) is one of the most important analysis steps in cosmology. A common technique for such identification is to use density-based clustering techniques FOF (Friends-of-Friends), which is a specific case of a general DBSCAN (Density-based Spatial Clustering of Applications with Noise) algorithm (Ester et al., 1996). DBSCAN relies on the quick identification of particle neighborhoods, grouping together points that are closely packed together while marking the points in the low-density regions as noise (outliers).

Here, we discuss the capability impact of incorporating DBSCAN algorithm implementation developed in ArborX within HACC by exploring a suite of full-scale end-to-end cosmology simulations. The developments in ArborX required to achieve this impact are described in Section 4.


The results were measured on the *Summit* supercomputer using 256 nodes, each equipped with 6 Nvidia V100 GPUs. The domain volume is  $(576\text{Mpc/h})^3$ . We present results from both a gravity-only simulation, featuring  $N = 2304^3$  particles, and a companion hydrodynamic simulation that evolved twice as many particles to account for both dark matter and gas particles individually. Both simulations trace the evolution from the early universe (over 13 billion years ago) to the present.

Figure 1 illustrates the complete measured computational time (excluding I/O) of the gravity-only simulation over 625 long-range force steps (red curve). As the simulation progresses, particles cluster, leading to an increase in solver time. ArborX is utilized in approximately 100 analysis steps during the simulation, with the most computationally demanding component being the FOF identification of dark matter halos within the full particle set.

In a series of downscaled test simulations at the same mass resolution, HACC has observed a significant speedup in performance, ranging from 10 to 12 times faster for FOF finding with ArborX when compared to a highly optimized CPU OpenMP threaded algorithm.

To visually illustrate the impact of such a performance boost in a production run, Figure 1 features a gray curve conservatively representing a tenfold increase in FOF execution time. This showcases, at a minimum, how slow the solver would have been without ArborX. The gray spikes corresponding to the 100 steps are now distinctly pronounced, clearly highlighting the computational cost of running analyses without GPU accelerated cluster finding. The speedup ratio in the bottom panel indicates that even with this conservative estimate, the full time-stepper shows an improvement of approximately a factor of 2.

The ArborX optimizations are so impactful that HACC can now perform analysis steps on every long-range force simulation timestep. This capability is particularly important for hydrodynamic simulations, where cluster finding needs to be executed at a much higher cadence. Furthermore, ArborX enables the capability of performing in-situ substructure finding within a simulation. Historically, this



**Figure 1.** Visualization of the performance impact of ArborX on analysis steps for a production gravity-only cosmology simulation.

task has necessarily been carried out in post-processing due to its computational expense.


In the case of hydrodynamic simulations, gas particles regularly evolve into stars, collectively forming galaxies. Each analysis step in HACC requires the identification of galaxies, achieved by running DBSCAN on all stellar particles using  $minPts = 10$ . Figure 2 visualizes star particles within one of the largest dark matter halos in the  $(576\text{Mpc}/h)^3$  simulation. The largest circle in the figure encloses the full halo, and each DBSCAN-identified galaxy is marked with a corresponding circle (using a radius equal to the farthest particle from the center of the galaxy). Full particle queries of this complexity would have been too expensive to perform during a simulation prior to the incorporation of ArborX into the analysis pipeline.

### 3. Background

#### 3.1. DBSCAN algorithm

Here, we give a short overview of the DBSCAN algorithm. For more details, we refer the readers to Ester et al. (1996).

Let  $X$  be a set of  $n$   $d$ -dimensional points to be clustered. Let  $\varepsilon$  and  $minPts$  be two given parameters. An  $\varepsilon$ -neighborhood of a point  $x$  is defined as  $N_\varepsilon(x) = \{y \in X \mid dist(x, y) \leq \varepsilon\}$ , where  $dist(\cdot, \cdot)$  is a distance metric (e.g., Euclidean). A point  $x$  with  $|N_\varepsilon(x)| \geq minPts$  is called a *core point*. A point  $y$  is *directly density-reachable* from a point  $x$  if  $x$  is a core point and  $y \in N_\varepsilon(x)$ . A point  $y$  is *density-reachable* from a point  $x$  if there is a chain of points  $x_1, \dots$ ,



**Figure 2.** Visualization of in-situ substructure finding of a large particle cluster from a hydrodynamic simulation using DBSCAN. Image credit: Azton Wells, Argonne National Laboratory.

$x_n, x_1 = x, x_n = y$ , such that  $x_{i+1}$  is directly density-reachable from  $x_i$ . Two points  $x$  and  $y$  are called *density-connected* if both are density-reachable from some point  $z \in X$ . Finally, any point that is not a core point but is density reachable from one is called *border point*. The remaining points, i.e., the points that are not core or border points, are called *noise*. Noise points are considered to be outliers not belonging to any cluster. Any cluster then consists of a combination of core points (at least one) and border points (possibly, none). Figure 3 provides an illustration for a set of points with  $minPts = 4$ .


The goal of a DBSCAN algorithm implementation is thus to find the individual clusters in a dataset efficiently.

#### 3.2. ArborX library

In this Section, we give a brief introduction to geometric search algorithms and the ArborX library, setting up a background for performance improvements in Section 4.

The main index in ArborX is a bounding volume hierarchy (BVH). BVH is a tree structure created from a set of objects in a multi-dimension space. Each object is wrapped in a geometric format (*bounding volume*) to form the leaf nodes of the tree. Each node of a BVH is an aggregate of its children, with the node's bounding volume enclosing the bounding volumes of its children. The bounding volume around all objects, called *scene bounding volume*, is stored at the root of the hierarchy.

As was demonstrated in Karras (2012), binary BVH is a good choice for GPU-based searches, particularly for low-dimensional data typical in the scientific applications. Fast BVH construction algorithms use a space-filling curve (Z-curve) to improve the spatial locality of the user data, followed by a single bottom-up construction to produce a binary tree structure (hierarchy). While the resulting quality is somewhat worse than produced by the best available



**Figure 3.** Classification of points for DBSCAN with  $minPts = 4$ . Core points are shown in red, border in blue, and noise are in gray.

algorithms, the construction procedure is extremely fast and produces a tree of sufficient quality in most situations.

During the search (also called a traversal), each thread (a host or a GPU thread depending on the backend) is assigned a single query, *i.e.*, a spatial or  $k$ -nearest neighbor search problem. All the traversals are performed independently in parallel in a top-down manner. To reduce the data and thread divergence, the queries are pre-sorted with the goal to assign neighboring threads the queries that are geometrically close.

To reduce the software development cost, ArborX uses Kokkos (Trott et al., 2022) for on-node parallelism to allow running on a variety of commodity and HPC hardware, including Nvidia, AMD, and Intel GPUs. Kokkos abstracts common parallel execution patterns, such as parallel loops, reductions, and scans (prefix sums), from the underlying hardware architecture. In addition, Kokkos provides an abstraction for a multi-dimensional array data structure called `View`. It is a polymorphic structure, whose layout depends on the memory the data resides in (host or device).

Kokkos supports both CPUs and GPUs (Nvidia, AMD, Intel) through providing backends, *e.g.*, OpenMP, CUDA, SYCL, HIP. Using Kokkos allows running the same code on CPUs or GPUs by simply changing the backend through a template parameter, resulting in a higher developer productivity.

Now, we briefly describe the scope of different focus areas within ArborX.

**3.1.1. Core functionality.** ArborX supports two kinds of search types: range and nearest. The *range* search finds all objects that intersect with a query object. For example, finding all objects within a certain distance is a range search. The *nearest* search, on the other hand, looks for a certain number of the closest objects, regardless of their distance. Both searches support multi-dimensional data, with dimensions ranging 1–10.

These two search types require very different tree traversal algorithms. The range query has to explore all nodes in a tree that satisfy the given predicate. It can be implemented in a stackless manner (see Section 4.2.1). In contrast, the nearest search can terminate early while it has found the best possible candidates. Nearest search is more complicated to implement, and relies on a stack and a priority queue structures.

ArborX provides both on-node and distributed (through using MPI) implementations for both searches.

**3.2.2. Clustering functionality.** ArborX provides implementations for several clustering algorithms that depend on distance calculations: DBSCAN (the focus of this work) and Euclidean minimum spanning tree (Prokopenko et al., 2023b). There are ongoing efforts to provide an efficient implementation for the HDBSCAN\*<sup>1</sup> algorithm (Campello et al., 2015).

**3.2.3. Ray tracing functionality.** ArborX provides basic support for ray tracing.


**3.2.4. Interpolation functionality.** ArborX implements moving least squares interpolation algorithm (Quaranta et al., 2005). In this method, support and subsequently the interpolation operator are constructed through solving local least squares problems defined by compactly supported radial basis functions.

## 4. ArborX improvements

ArborX's primary goal is performance. For a user, this comes with an implicit agreement that ArborX will be diligent in implementing features in a way that does not slow down user applications. For this to happen, ArborX uses several benchmarks to test proposed new functionality.


One of these benchmarks is the DBSCAN algorithm for a cosmology problem. The challenge problem for the HACC project is a 12.2 trillion particle gravity-only simulation with a companion hydrodynamic simulation that models both gas and dark matter with  $2 \times 1.8$  trillion points. A downscaled version of the gas simulation was run using  $2 \times 1024^3$  particles with  $(256\text{Mpc/h})^3$  domain volume. For our dataset, we used a snapshot from the last step of the simulation performed with HACC, when the clusters are clearly formed. The data was taken from a single rank of the original 64 MPI rank job, consisting of  $\approx 37\text{M}$  dark matter particles (gas particles are excluded for cluster finding). Figure 4 shows a 3D visualization of the data sample. The value of  $\epsilon$  was set to 0.042<sup>2</sup>.

Figure 5 shows the timeline of the improvements in ArborX's core search functionality together with the



**Figure 4.** Benchmark problem data sampled from a single rank. The clusters are clearly formed.





**Figure 5.** Timeline of the ArborX improvements for the DBSCAN benchmark problem ( $\epsilon = 0.042$ ,  $\text{minPts} = 2$ ) on an Nvidia A100 GPU.

algorithmic developments for this benchmark problem. The code was run on an Nvidia A100 GPU. Here, we highlight certain critical points which affected the performance:

- (1) Initial implementation (Section 4.3.1)
- (2) Initial introduction of FDBSCAN (Section 4.3.3) and use of the callbacks (Section 4.1.1)
- (3) Switching from using Karras construction algorithm to Apetrei's (Section 4.2.1)
- (4) Switching to using stackless traversal (Section 4.2.1)
- (5) Changing the callback in FDBSCAN to improve  $\text{minPts} > 2$  case
- (6, 9) Switching to using 64-bit indices (Section 4.2.2)
- (7) Using pair-traversal in FDBSCAN (Section 4.2.3)
- (8) Initial introduction of FDBSCAN-DENSEBOX (Section 4.3.4)

It is clear that the most significant improvements to the performance were the stackless traversal, the move from 32- to 64-bit Morton codes, and the pair traversal for FDBSCAN. Overall, the performance of the algorithm improved by factor  $\approx 9.2$ , with the latest version clustering the full  $\approx 37\text{M}$  benchmark problem in under 0.15s on an Nvidia A100. Performance improvements were mostly monotonic, with the exception for (5), which significantly improved the FDBSCAN callbacks algorithm for  $\text{minPts} > 2$  at the cost of a slowdown for  $\text{minPts} = 2$ .

Initially, FDBSCAN-DENSEBOX performed significantly better than FDBSCAN, however the latter became the faster one for this problem with the introduction of the pair traversal. There is ongoing research to integrate pair traversal into the FDBSCAN-DENSEBOX algorithm.

We will now describe each interface, performance and algorithmic improvement in more details.

## 4.1. Interface improvements

**4.1.1. Callbacks.** The original ArborX interface was designed to handle the tasks of interest to the Data-TransferKit library. The results were produced as a pair of Kokkos views (`offsets`, `values`), with the `values` view containing the values satisfying the predicates, and the `offsets` view containing the offsets into `values` associated with each query. However, it was observed that users are often interested in performing some operation on the results of each query and not the results themselves. For example, they may be only interested in the number of neighbors, average distance, or updating some quantity. In these cases, storing the results may be unnecessary, penalizing both performance through memory writes and memory usage. For some problems, storing the found objects results in running out of memory even for simple counting kernels.

To address this issue, we introduced the callback functionality in ArborX. It allows execution of a user-provided code on a positive match. Figure 6 demonstrates the interface. Here, `Predicate` represents a search query (e.g., an object to intersect with), and `Value` represents the data stored in the BVH. We support both *pure* callbacks that perform a user operation without storing any results, and callbacks that allow a user to modify the results before storing them.

**4.1.2. Early termination.** One of the steps in the DBSCAN algorithm is determination of the core points, i.e., the points that have a specified number of neighbors within the radius. To find them, we need to execute a traversal procedure.

It is, however, unnecessary to continue the procedure once the threshold has been achieved and it has been established that a point is a core point. Thus, we updated the interface to allow a callback to indicate (through a return

```

struct Callback {
    template<typename Predicate, typename Value>
    KOKKOS_FUNCTION
    RT operator() (Predicate const &predicate,
                  Value const &value) const;
};

struct CallbackWithOutput {
    template<typename Predicate, typename Value,
            typename OutputFunc>
    KOKKOS_FUNCTION
    void operator() (Predicate const &predicate,
                    Value const &value,
                    OutputFunc const &output) const;
};

```

**Figure 6.** Callbacks interface. The return type RT could either be void, or `enum CallbackTreeTraversalControl`. The latter affects the traversal, allowing early termination (see Section 4.1.2).

value) whether the traversal should continue. For low values of *minPts*, this saves a significant amount of time.


There are many other potential applications of this feature. In general, if it is known that the traversal will produce at most *m* answers, it can be stopped after achieving that value. For example, a search for a mesh cell containing a given particle in particle-in-cell simulations can be terminated after the first such cell has been identified.

## 4.2. Core algorithms improvements

In this section, we describe improvements to the interface and several optimization techniques to improve the performance of the implemented DBSCAN algorithms in Section 4.3.

**4.2.1. Stackless traversal.** Stackless traversal is a technique to avoid explicitly managing a stack of node pointers for each thread during the traversal. Usage of stacks is undesirable as it may lead to lower occupancy due to higher memory demands per thread. The approach in [Torres et al. \(2009\)](#) introduced a concept of *rope* (also called *escape index*), an index of a node where the traversal should proceed if the intersection test with the current node is not satisfied, or if the node is a leaf node. In [Figure 7](#), the references to the right child (denoted by dotted lines) of the internal nodes are removed, and the ropes (denoted curved arrows) are introduced for both internal and leaf nodes. For the nodes on the right-most path, the ropes point to the artificial terminal node called *sentinel*.

Originally, ArborX used Karras' approach ([Karras, 2012](#)) for the BVH construction. The way the Karras algorithm orders internal nodes makes it possible to set the ropes as part of the standard bottom-up hierarchy construction.



**Figure 7.** Hierarchy for stackless traversal. The references to right children (dotted arrows) of the internal nodes 0-6 (orange) are replaced with ropes (curved arrows). Additionally, each leaf node 7-14, (green) now has a rope. The ropes of the nodes on the right-most path point to the sentinel node S (yellow).

However, to improve the hierarchy construction time, we switched to using [Apetrei \(2014\)](#) as the latter is more efficient. The unfortunate side effect of the switch was that the ropes had to be set separately, as Apetrei's algorithm orders internal nodes differently. However, we were able to overcome this side-effect by finding a way to recover Karras' ordering from the Apetrei's one. For more details, see [Prokopenko and Lebrun-Grandié \(2024\)](#).

**4.2.2. 64-bit Morton codes.** BVHs are often used in computer graphics for ray tracing, with performance often compared in terms of rays per second casted for rendering a set of scenes presented as a set of surfaces (e.g., triangles). One of the challenges of applying these algorithms to scientific data is the difference in the range of scales. It is common for the ratio between the densest and the sparsest regions of scientific data to be several orders of magnitude. This requires certain adjustments of the algorithms.

Standard linear BVH (LBVH) ([Lauterbach et al., 2009](#)) algorithms rely on space-filling curves (typically, Z-order curve based on Morton codes) to organize the spatial locality of the data. The case of two objects happen to have the same Morton index resolved in an ad-hoc manner. However, the scientific data may degrade severely in this case. In our original implementation, we used 32-bit Morton codes, meaning that each dimension of the 3D problem was partitioned in 1024 bins. It turned out that due to the very high density of the particles in some regions, too many particles were assigned to the same bin, i.e., having the same Morton code. For example, for the benchmark problem at least 64% particles shared their Morton code with at least one other particle, with the maximum of 3,569 having the same index. This resulted in decreased performance of the algorithms due to the worse hierarchy quality.

We chose to address this problem by using the 64-bit Morton codes. As we can see in [Table 1](#), showing the

**Table 1.** An overview of the hierarchy construction differences between 32- and 64-bit Morton codes for the benchmark problem.

	32-bit	64-bit
#Duplicate codes (> 3 times)	1,311,912	0
#Points with duplicate code	23,539,027	528
max same code duplicates	3569	2
Number of hierarchy levels	43	49

statistics for the benchmark problem in, we eliminate almost all duplicate Morton codes for our problem when using 64-bit resolution. This becomes even more crucial for larger problems.


Using 64-bit codes has two minor drawbacks. It slightly increases the hierarchy construction cost due to the slower sorting of the 64-bit integers compared to the 32-bit ones. Additionally, increasing resolution in Morton codes will typically result in a deeper hierarchy (in this case, 43 to 49). However, we found the tradeoff worth it due to significant speedup in traversal algorithms.

**4.2.3. Pair traversal.** Many computational problems, including DBSCAN, can be seen as: given a radius  $\varepsilon$ , find all pairs  $(i, j)$  of points that are within distance  $\varepsilon$  of each other, and execute some operation  $F(i, j)$  on the pair. In DBSCAN, that operation is Union. In other applications, the operation could be to compute a force (e.g., in molecular dynamics) or to increase the value of the count (e.g., computing 2-point correlations).

For this class of problems, each pair needs to be processed only once. Regular traversal would process each pair twice, once for each thread, resulting in unnecessary work. Thus, a new approach is desired.

We implemented a new hierarchy traversal algorithm. For a given point  $i$ , instead of starting the traversal from the root node 0, we start it with the leaf node corresponding to this point. Combined with the ropes structure described earlier, it guarantees that that only pairs  $(i, j)$  with  $i < j$  will be found, thus processing each pair exactly once. Figure 8 demonstrates the traversal for a thread corresponding to index 4. It is clear that a thread would need to examine fewer nodes compared to the regular traversal. This results in fewer memory accesses used during the traversal, reduced number of distance computations, and reduced number of applying the operation  $F$ .

We explained the algorithm above for a tree with ropes. It is possible to achieve a similar effect for a regular tree with left and right children by masking subtrees. This would impose certain requirement on the ordering of the internal nodes, and an initial top-down traversal.



**Figure 8.** An example of the tree traversal mask for a thread corresponding to a point with index 4. The thread skips all internal and leaf nodes in gray, and starts with the corresponding leaf node 14 (pink), which is the 4-th leaf node. The traversal then follows the standard procedure.

### 4.3. DBSCAN algorithm improvements

**Algorithm 1** Disjoint-set DBSCAN algorithm

```

1: procedure DSDBSAN( $X, minPts, \varepsilon$ )
2:   for each point  $x \in X$  do
3:      $N \leftarrow \text{GetNeighbors}(x, \varepsilon)$ 
4:     if  $|N| \geq minPts$  then
5:       mark  $x$  as core point
6:       for each  $y \in N$  do
7:         if  $y$  is marked as a core point then
8:           Union( $x, y$ )
9:         else if  $y$  is not a member of any cluster then
10:          mark  $y$  as a member of a cluster
11:          Union( $x, y$ )

```

The original DBSCAN algorithm in Ester et al. (1996) was hard to parallelize due to its breadth-first manner of encountering new points. Improvements to the algorithm in Patwary et al. (2012) broke with its breadth-first nature. The authors used the UNION-FIND (Tarjan, 1979) approach to maintain a disjoint-set data structure. The approach relies on two main operations: UNION and FIND. Find( $x$ ) determines the representative of a set that a point  $x$  belongs to, while Union( $x, y$ ) combines the sets that  $x$  and  $y$  belong to.

For completeness, Algorithm 1 shows the disjoint-set DBSCAN (DSDBSAN) algorithm as proposed in Patwary et al. (2012) (Algorithm 2). Each point only computes its own neighborhood (Line 3). If it is a core point, its neighbors are assigned to the same cluster (Lines 8 and 11).

We will now describe our initial implementation, and then briefly describe two newly developed algorithms (for more information, see Prokopenko et al. (2023a)).

**4.3.1. Initial DBSCAN implementation.** Our original approach was to solve the special case of DBSCAN of  $minPts = 2$ . In cosmology literature, it is usually called Friends-of-Friends (FOF). This case is simpler, as each point either belongs to a cluster as a core point, or is noise. It is equivalent to finding connected components in the undirected adjacency graph, with each pair of vertices within  $\varepsilon$  of each other have a corresponding edge in the graph.

Prior to introduction of the callbacks (see Section 4.1.1), ArborX only produced an explicit adjacency graph. Afterwards, we used the ECL-CC (Jaiganesh and Burtscher, 2018) algorithm to compute the connected components.

The major drawback of this approach was storing the full adjacency graph in memory. It imposed a severe restriction on the size of the problems one could run. The memory usage depended not only on the size of the problem  $n$ , but also on the parameter  $\varepsilon$ .

**4.3.2. Reformulated DBSCAN algorithm.** Both the original Ester et al. (1996) and Patwary et al. (2012) DBSCAN formulations do not expose enough parallelism for GPU implementations with thousands of threads. To address that, we reformulated the algorithm to consist of two phases. In the first phase (*preprocessing*), the algorithm determines the core points. In the second phase (*main*), it merges the pairs of close neighbors as they are being discovered.

---

**Algorithm 2** Parallel disjoint-set DBSCAN algorithm

---

```

1: procedure PDSDBSCAN( $X, minPts, \varepsilon$ )
2:   if  $minPts > 2$  then
3:     for each point  $x \in X$  in parallel do
4:       determine whether  $x$  is a core point
5:   for each pair of points  $x, y$  such that  $dist(x, y) \leq \varepsilon$  in parallel do
6:     if  $x$  is a core point then
7:       if  $y$  is a core point then
8:         Union( $x, y$ )
9:       else if  $y$  is not yet a member of any cluster then
10:        critical section:
11:          mark  $y$  as a member of a cluster
12:          Union( $x, y$ )


```

---

The pseudocode for the reformulated DBSCAN (PDSDBSCAN) algorithm is shown in Algorithm 2. The preprocessing phase is executed on Lines 3-4. The UNION-FIND algorithm is performed on Lines 8 and 11.

The main advantage of the reformulated algorithm is that it executes the neighbor searches completely in parallel for all data points, allowing it to take advantage of the underlying parallel index. In addition, it allows processing each found neighbor as soon as it is found and immediately discarding afterwards.

**4.3.3. FDBSCAN.** FDBSCAN (“fused” DBSCAN) constructs a search index (BVH) over all the points of the datasets. Several ArborX improvements described in Section 4.2 play a crucial role in achieving a good performance. We use callbacks to fuse tree traversal with the counting and UNION-FIND kernels, avoiding neighbor storage. This makes the algorithm use  $O(n)$  memory. We use early termination in the counting kernel, stopping a thread once  $minPts$  threshold has been achieved. In the main kernel constructing the clusters, we then use the pair traversal technique.



**Figure 9.** Regular 2D grid with grid size  $\varepsilon/\sqrt{2}$  superimposed over the dataset. The dense cells for  $minPts = 5$  are shown in red.

**4.3.4. FDBSCAN-DenseBox.** FDBSCAN-DENSEBOX is a modification to FDBSCAN accommodating regions with high density (with respect to  $\varepsilon$ ). These regions are characterized by the number of neighbors within  $\varepsilon$ -neighborhood far exceeding the  $minPts$  value. In this scenario, many of the distance computations may be avoided.

To achieve this goal, we superimpose a regular grid with the grid cell length of  $\varepsilon/\sqrt{d}$  (with  $d$  being the data dimension) on top of the data. The choice of the length parameter guarantees that each cell’s diameter does not exceed  $\varepsilon$ . Thus, any grid cell with at least  $minPts$  points in it will only contain core points, and the distance calculations among them can be eliminated. We call these cells *dense*. Figure 9 demonstrates an example of such grid over a set of points, with dense cells for  $minPts = 5$  marked in red.

We modify the BVH construction algorithm of FDBSCAN to accommodate dense boxes. Instead of constructing the hierarchy only on the data points, we construct it out of a mix of dense cells and points outside of the dense cells. This poses no challenge to the BVH construction, as it only requires bounding volumes for a set of objects. During the traversal, if a found object is a dense cell, we perform a search over all of its contained points.

## 5. Conclusions and future work

We presented the progress made in ArborX in the past few years to support the HACC cosmology application. We identified an analysis problem of high importance to the overall cosmology simulation, and developed and implemented multiple algorithms to improve the performance of the dbscan algorithm. We presented a timeline of the core and algorithmic changes in ArborX using a benchmark problem. We have also shown that the ArborX speedup, which is of an order of magnitude when compared to a highly optimized CPU OpenMP algorithm in HACC,



resulted in a factor of 2 improvement in the full time-stepper performance.

Our ongoing research has several thrusts. We are developing an efficient implementation of an advanced density-based algorithm HDBSCAN<sup>1</sup> (Campello et al., 2015), which is an improvement over the DBSCAN algorithm. We are working on improving the SYCL based implementation to target the upcoming Aurora supercomputer. In addition, there are ongoing efforts to include the auto-tunable interface to select best occupancy parameters during the simulation runtime.

## Acknowledgement

We would further like to acknowledge the work of the ExaSky team and the development and testing efforts therein.

## Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

## Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: 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. This research used 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. Argonne National Laboratory's work was supported under the U.S. Department of Energy contract DE-AC02-06CH11357. Additionally, this study utilized resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

## ORCID iDs

Andrey Prokopenko  <https://orcid.org/0000-0003-3616-5504>  
 Daniel Arndt  <https://orcid.org/0000-0001-8773-4901>  
 Damien Lebrun-Grandié  <https://orcid.org/0000-0003-1952-7219>  
 Nicholas Frontiere  <https://orcid.org/0009-0005-8598-4292>  
 J.D. Emberson  <https://orcid.org/0000-0003-1406-0744>  
 Michael Buehlmann  <https://orcid.org/0000-0002-8469-4534>

## Notes

1. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, world-wide license to publish or reproduce the

published form of this manuscript, or allow others to do so, for United States Government purposes.

2. Following the formula  $\varepsilon = b(V/n)^{1/3}$ , where with  $b = 0.168$  is the linking length,  $V$  is the simulation volume ( $256^3$ ), and  $n$  being the number particles ( $1024^3$ ).

## References

- ALCF (2023) Aurora. <https://www.alcf.anl.gov/aurora/>.
- Apetrei C (2014) Fast and simple agglomerative LBVH construction. In: Borgo R and Tang W (eds) *Computer Graphics and Visual Computing (CGVC)*. The Eurographics Association. DOI: [10.2312/cgvc.20141206](https://doi.org/10.2312/cgvc.20141206).
- Campello RJGB, Moulavi D, Zimek A, et al. (2015) Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Transactions on Knowledge Discovery from Data* 10(1): 5:1–5:51. DOI: [10.1145/2733381](https://doi.org/10.1145/2733381).
- Ester M, Kriegl HP, Sander J, et al. (1996) A density-based algorithm for discovering clusters in large spatial databases with noise *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD '96*. AAAI Press, 226–231.
- Frontiere N, Emberson J, Buehlmann M, et al. (2023) Simulating hydrodynamics in cosmology with CRK-HACC. *The Astrophysical Journal Supplement Series* 264(2): 34.
- Habib S, Pope A, Finkel H, et al. (2016) HACC: simulating sky surveys on state-of-the-art supercomputing architectures. *New Astronomy* 42: 49–65.
- Jaiganesh J and Burtscher M (2018) A high-performance connected components implementation for GPUs *Proceedings of the 27th International Symposium on High-Performance Parallel and Distributed Computing*. New York, NY, USA: ACM, 92–104. HPDC '18. DOI: [10.1145/3208040.3208041](https://doi.org/10.1145/3208040.3208041).
- Karras T (2012) Maximizing parallelism in the construction of BVHs, octrees, and K-d trees *Proceedings of the Fourth ACM SIGGRAPH / Eurographics Conference on High-Performance Graphics, EGGH-HPG'12*. Goslar Germany, Germany: Eurographics Association, 33–37. DOI: [10.2312/EGGH/HPG12/033-037](https://doi.org/10.2312/EGGH/HPG12/033-037).
- Lauterbach C, Garland M, Sengupta S, et al. (2009) Fast BVH construction on GPUs. *Computer Graphics Forum* 28(2): 375–384.
- Lebrun-Grandié D, Prokopenko A, Turcksin B, et al. (2020) ArborX: a performance portable geometric search library. *ACM Transactions on Mathematical Software* 47(1): 1–2. DOI: [10.1145/3412558](https://doi.org/10.1145/3412558).
- Messina P (2017) The exascale computing project. *Computing in Science & Engineering* 19(3): 63–67.
- NERSC (2022) Perlmutter. <https://www.nersc.gov/systems/perlmutter/>.
- OLCF (2022) Frontier. <https://www.nersc.gov/systems/perlmutter/>.
- Patwary MMA, Palsetia D, Agrawal A, et al. (2012) A new scalable parallel DBSCAN algorithm using the disjoint-set data structure. In: *SC '12: Proceedings of the International*

*Conference on High Performance Computing, Networking, Storage and Analysis*, 1–11. DOI: [10.1109/SC.2012.9](https://doi.org/10.1109/SC.2012.9).

- Prokopenko A and Lebrun-Grandié D (2024) *Revising Apetrei's bounding volume hierarchy construction algorithm to allow stackless traversal*. Oak Ridge, TN (United States): Oak Ridge National Laboratory (ORNL). Technical Report ORNL/TM-2024/3259. DOI: [10.2172/2301619](https://doi.org/10.2172/2301619).
- Prokopenko A, Lebrun-Grandié D and Arndt D (2023a) Fast tree-based algorithms for DBSCAN for low-dimensional data on GPUs. In: *Proceedings of the 52nd International Conference on Parallel Processing, ICPP '23*. New York, NY, USA: Association for Computing Machinery, 503–512. DOI: [10.1145/3605573.3605594](https://doi.org/10.1145/3605573.3605594).
- Prokopenko A, Sao P and Lebrun-Grandié D (2023b) A single-tree algorithm to compute the Euclidean minimum spanning tree on GPUs *Proceedings of the 51st International Conference on Parallel Processing, ICPP '22*. New York, NY, USA: Association for Computing Machinery, 1–10. DOI: [10.1145/3545008.3546185](https://doi.org/10.1145/3545008.3546185).
- Quaranta G, Masarati P, Mantegazza P, et al. (2005) A conservative mesh-free approach for fluid-structure interface problems. In: *International Conference for Coupled Problems in Science and Engineering, Greece*.
- Slattery S, Wilson P and Pawlowski R (2013) The data transfer kit: a geometric rendezvous-based tool for multiphysics data transfer. In: *International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering (M&C 2013)*, 5–9.
- Tarjan RE (1979) A class of algorithms which require nonlinear time to maintain disjoint sets. *Journal of Computer and System Sciences* 18(2): 110–127. DOI: [10.1016/0022-0000\(79\)90042-4](https://doi.org/10.1016/0022-0000(79)90042-4).
- Torres R, Martín PJ and Gavilanes A (2009) Ray casting using a roped BVH with CUDA. In: *Proceedings of the 25th Spring Conference on Computer Graphics, SCCG '09*. New York, NY, USA: Association for Computing Machinery, 95–102. DOI: [10.1145/1980462.1980483](https://doi.org/10.1145/1980462.1980483).
- Trott CR, Lebrun-Grandié D, Arndt D, et al. (2022) Kokkos 3: programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems* 33(4): 805–817. DOI: [10.1109/TPDS.2021.3097283](https://doi.org/10.1109/TPDS.2021.3097283).

## Author biographies

*Andrey Prokopenko* is a computational scientist at Oak Ridge National Laboratory. His research interests include multigrid algorithms, geometric search, and extreme scale computing.

*Daniel Arndt* received his PhD degree in mathematics from University of Göttingen. He is currently a computational scientist with Oak Ridge National Laboratory. He worked on multiple ECP projects, including Kokkos with a primary focus on the SYCL backend, and is one of the main authors of the deal.II finite element library.

*Damien Lebrun-Grandié* is a Senior Computational Scientist with Oak Ridge National Laboratory. His research focus on the development of algorithms and enabling technologies for the solution of large-scale complex engineering and scientific problems. He is a high-performance computing expert, co-leader of the Kokkos performance portability project, and member of the ISO C++ Standards Committee.

*Bruno Turcksin* is a Computational Scientist at Oak Ridge National Laboratory. His work spans multiple disciplines including nuclear engineering, finite element, high performance computing, and additive manufacturing. He is one of the main developers of deal.II (finite element library), Kokkos (programming model library), and adamantine (additive manufacturing simulator).

*Nicholas Frontiere* is a computational scientist at Argonne National Laboratory, specializing in numerical cosmology and computational fluid dynamics. His research focuses on simulating the large-scale structure formation of the universe and developing high-performance computing applications for modern exascale machines.

*J. D. Emberson* is a computational cosmologist in the Computational Science Division at Argonne National Laboratory. He has a broad interest in the development and application of numerical techniques for large-scale cosmological structure formation simulations. This includes both N-body as well as hydrodynamic methods in order to gain insight into various astrophysical phenomena including dark energy, neutrinos, and primordial non-Gaussianity. Emberson is one of the main developers of HACC with an emphasis on integrating robust hydrodynamical modeling for use on next-generation exascale platforms.

*Michael Buehlmann* is a computational scientist at Argonne National Laboratory, specializing in numerical cosmology. His work involves the development and execution of large-scale cosmological simulations, with a focus on subgrid modeling and post-processing pipelines for exascale systems.