

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

An extension to Voro++ for multithreaded computation of Voronoi cells

### Permalink

<https://escholarship.org/uc/item/05s6b85h>

### Authors

Lu, Jiayin

Lazar, Emanuel A

Rycroft, Chris H

### Publication Date

2023-10-01

### DOI

10.1016/j.cpc.2023.108832

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# An extension to VORO++ for multithreaded computation of Voronoi cells

Jiayin Lu<sup>a,b</sup>, Emanuel A. Lazar<sup>c</sup>, Chris H. Rycroft<sup>b,d,a</sup>

<sup>a</sup>*John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, United States*

<sup>b</sup>*Department of Mathematics, University of Wisconsin–Madison, Madison, WI 53711, United States*

<sup>c</sup>*Department of Mathematics, Bar-Ilan University, Ramat Gan 5290002, Israel*

<sup>d</sup>*Mathematics Group, Lawrence Berkeley Laboratory, Berkeley, CA 94720, United States*

---

## Abstract

VORO++ is a software library written in C++ for computing the Voronoi tessellation, a technique in computational geometry that is widely used for analyzing systems of particles. VORO++ was released in 2009 and is based on computing the Voronoi cell for each particle individually. Here, we take advantage of modern computer hardware, and extend the original serial version to allow for multithreaded computation of Voronoi cells via the OpenMP application programming interface. We test the performance of the code, and demonstrate that it can achieve parallel efficiencies greater than 95% in many cases. The multithreaded extension follows standard OpenMP programming paradigms, allowing it to be incorporated into other programs. We provide an example of this using the *VoroTop* software library, performing a multithreaded Voronoi cell topology analysis of up to 102.4 million particles.

*Keywords:* Voronoi tessellation, computational geometry, multi-threaded programming

---

## Program summary

*Program title:* VORO++

*Developer's repository link:* <https://github.com/chrlshr/voro>

*Licensing provisions:* BSD 3-clause (with LBNL modification)

*Programming language:* C++

*External routines/libraries:* OpenMP

*Nature of problem:* Multithreaded computation of the Voronoi tessellation in two and three dimensions

*Solution method:* The VORO++ library is built around several C++ classes that can be incorporated into other programs. The two largest components are the `container...` classes

that spatially sort input particles into a grid-based data structure, allowing for efficient searches of nearby particles, and the `voronoicell...` classes that represent a single Voronoi cell as an arbitrary convex polygon or polyhedron. The Voronoi cell for each particle is built by considering a sequence of plane cuts based on neighboring particles, after which many different statistics (*e.g.* volume, centroid, number of vertices) can be computed. Since each Voronoi cell is calculated individually, the Voronoi cells can be computed using multithreading via OpenMP.

## 1. Introduction

The Voronoi tessellation was originally introduced in 1907 [1] and is now a broadly used technique in computational geometry [2]. Consider a set of points in a domain. Each point has a corresponding Voronoi cell that is defined as the part of the domain that is closer to that point than to any other. In two dimensions (2D) with the Euclidean metric, the Voronoi cells are irregular polygons that perfectly partition the domain to create the Voronoi tessellation (Fig. 1(a)). Each edge in the tessellation is the perpendicular bisector between neighboring points. In three dimensions (3D) the Voronoi cells are irregular polyhedra (Fig. 1(b)). Voronoi cells can also be generalized to non-Euclidean geometries with different distance metrics [3, 4].

The Voronoi tessellation has been used in a remarkable number of different scientific fields. It has been extensively used to analyze systems of particles or atoms, where features of the Voronoi cells (*e.g.* volume, surface area, number of faces) provide insight in particle structure; examples include the analysis of granular materials [5, 6, 7], colloids [8], nanosphere systems [9], metallic glasses [10, 11], liquids [12], as well as active [13] and supercritical fluids [14, 15]. The Voronoi cells themselves, which form irregular polygons/polyhedra, have been used to model different physical phenomena, such as polycrystalline materials [16, 17, 18], solidification processes [19, 20], and biological cells [21]. The Voronoi tessellation has also been used to construct computational meshes on which to solve partial differential equations [22, 23], such as for climate modeling [24], groundwater flow [25, 26], and astrophysical flows [27]. Other applications include control of multi-robot systems [28], calculating snow aggregate scattering properties [29], and modeling animal territorial control [30]. There are many more examples than the ones given here, highlighting the ubiquity of this geometrical construction [31].

A variety of software packages are available for calculating the Voronoi tessellation. The Qhull library [32, 33] is widely used and incorporated into MATLAB (via the `voronoin` command) and Python (via the `scipy.spatial.Voronoi` command). The Computational Geometry Algorithms Library (CGAL) [34] provides a variety of functions for computing the Voronoi tessellation, and Triangle [35, 36] can compute the Voronoi tessellation in 2D. All of these libraries primarily focus on computing the Voronoi tessellation as an entire mesh, shown in blue in Fig. 1.

In 2009 Rycroft released VORO++ [37, 38], a software library written in C++ that takes the alternative approach of calculating the Voronoi cells individually, so that each cell, an example of which is shown in red in Fig. 1, is computed as a separate object. The library grew out of research on particulate granular flows, where the Voronoi cell volumes were useful for understanding particle packing structure [39, 40, 41]. The cell-based perspective has some




Figure 1: (a) An example two-dimensional Voronoi tessellation shown in blue generated by the black crosses. A single Voronoi cell in the tessellation is shown in red as an irregular polygon. (b) An example three-dimensional Voronoi tessellation shown in blue generated by the white spheres. A single Voronoi cell is shown in red as an irregular polyhedron.

advantages and drawbacks compared to the entire mesh approach (Subsec. 1.2). However, it has proven effective in a wide range in applications, particularly those involving rapid analysis of particle systems. VORO++ has a command-line utility that can perform a variety of different analyses, and it has a C++ application programming interface (API) that allows it to be called from user-written programs. It has been incorporated into other software such as LAMMPS [42, 43] and OVITO [44, 45].

### 1.1. Algorithms for computing the Voronoi tessellation

Since the 1970's, a wide variety of methods for computing the Voronoi tessellation have been proposed [2]. For computing the entire Voronoi mesh, some popular methods include the Fortune sweeping algorithm [46, 47] and the incremental approach whereby the mesh is continually updated as new particles are added [48, 49]. Another method is to use the lift-up mapping, projecting a point  $\mathbf{x} \in \mathbb{R}^n$  to a paraboloidal surface  $(\mathbf{x}, \|\mathbf{x}\|^2) \in \mathbb{R}^{n+1}$ . The hyperplanes tangential to the surface form facets that exactly match the Voronoi tessellation when projected back to  $\mathbb{R}^n$ . This can be efficiently computed in arbitrary dimensions using the quickhull algorithm [32], which forms the basis of Qhull [33]. Another approach involves introducing a computational grid, and sweeping out from each point with the fast marching method [50, 51] to construct Voronoi cells [52], which is computationally expensive but more flexible for calculations on spaces with non-Euclidean distance metrics.

### 1.2. The cell-based approach: advantages and drawbacks

The cell-based approach that VORO++ uses has also been explored in the literature [53, 54, 55, 56]. As discussed by Okabe *et al.* [2] it has a significant difficulty that is illustrated




Figure 2: (a) Typical case where three Voronoi cells (blue polygons) for three points (black crosses) meet at a vertex. The Voronoi cells would normally touch, but are spaced slightly apart from each other for illustrative purposes. (b) Special case where four Voronoi cells meet at a vertex that is equidistant from four points; floating-point errors could lead to additional small edges (red) for some cells. (c) Schematic representation adapted from Lazar *et al.* [57] where each dashed region represents a different Voronoi cell topology. Some common crystalline lattices (*e.g.* BCC) are located within a single region, whereas others (*e.g.* FCC, HCP) are located at junctions, and small perturbations from the ideal configuration sample different topologies.

in Fig. 2(a,b). In most cases, for randomly-distributed points, each vertex of the Voronoi tessellation will be common between three Voronoi cells as shown in Fig. 2(a). However, in certain situations a vertex may be equidistant from four particles as shown in Fig. 2(b). This could happen either because of a special arrangement of the particles (*e.g.* a crystalline formation), or for random arrangements when particles happen to be aligned within the limit of floating point truncation error. In the cell-based approach where the cells are computed independently, small floating point errors in one cell may lead to the creation of additional facets (shown by the red line in Fig. 2(b)) meaning that the topologies of the edges and faces of the Voronoi cells are not consistent.

Despite this difficulty, there are many situations where the cell-based approach is attractive. The entire Voronoi mesh is often not required, and the individual Voronoi cells can be analyzed independently. In VORO++ a typical workflow is to compute a Voronoi cell, calculate and store various statistics about the cell, delete the cell, and move onto the next point. Because only a single cell needs to be stored at any one time, this results in a large memory saving and an improvement in cache efficiency, creating an inherent performance boost over building the entire mesh. Furthermore, for many commonly used measurements, it is not necessary for the edge topology to agree perfectly. For example, Voronoi cell volumes and centroids (used in Lloyd’s algorithm [58, 59]) are not sensitive to small changes in edge topology.

For other measurements, such as the number of faces or edges, the precise topology of the Voronoi cell can have an appreciable effect. However, one can reasonably argue that measurements that rely too heavily on these small topological changes are problematic to begin with—it should not be the case that a diagnostic indicator of a physical characteristic be sensitive to truncation error (*i.e.* often around a factor of  $10^{-16}$  in double-precision arithmetic) particularly when experimental errors or simulation discretization errors are usually far larger. The recent work of Lazar, Srolovitz, and coworkers provides a useful theoretical framework in which to address this issue [57, 60, 61, 62, 63]. Voronoi cells can be pictured as residing in a phase space that is divided into discrete regions representing the complete face and edge topologies (Fig. 2(c)). Commonly studied crystalline lattices such as FCC and HCP lie at

intersections of these topologies, meaning that small perturbations (*e.g.* thermal vibrations) around the idealized lattice will push the Voronoi cells into a well-defined family of different topologies. Other lattices such as BCC may lie in the interior of a single region. Lazar released the *VoroTop* package [64, 65], which uses this framework to analyze ensembles of Voronoi cell topologies, classify different particle packings, and identify features such as grain boundaries. A key part of *VoroTop* is the computation of the Weinberg vector [66] for each Voronoi cell, which uniquely characterizes the cell’s vertex and edge topology.

### 1.3. Outline of this paper

Advances in supercomputing power have enabled simulations with very large numbers of particles [67, 68]. In addition, there is currently interest in developing data-driven approaches for screening large databases of materials and structures, where the Voronoi tessellation can be a useful analysis tool [69, 70, 71, 72]. Thus there is a need to compute the Voronoi tessellation at a large scale and in parallel. Currently, there are some parallel approaches in the literature for computing Voronoi cells in a distributed-memory model [73, 74]. Here, we consider parallelizing VORO++ using a shared-memory model with multithreading. Since modern consumer laptops and desktops contain CPUs with 4–8 cores and servers contain CPUs with upward of 16 cores, multithreading enables a large practical speedup without the additional complexity of using distributed-memory architectures.

The cell-based approach used by VORO++ is inherently amenable to parallelization, since each Voronoi cell can be computed independently. Here we develop a general multithreaded extension of VORO++ that provides good parallel performance across a range of different scenarios. An ideal basis for doing this is OpenMP, an API for shared-memory multiprocessing [75, 76]. The core component of OpenMP is a set of compiler directives beginning with `#pragma omp` that instruct the compiler to multithread certain lines and loops within a C++ code. A key feature of OpenMP is that if a program is compiled without OpenMP enabled, then the `#pragma omp` directives are ignored and the compiler will create a standard executable that runs in serial. For open-source scientific software, which is compiled and run on a wide range of different systems, this serial interoperability is a major advantage.

The extension to the VORO++ API is designed to make it as simple as possible for the user to incorporate multithreading into their programs. Multithreading a loop over all Voronoi cells requires adding a small number of `#pragma omp` directives that match typical OpenMP usage. This has required redesigning the mechanism for looping over Voronoi cells from previous versions of VORO++, but this is done so that most of the complexity is hidden from the user. Furthermore, we demonstrate that our extension is interoperable with standard OpenMP functionality for tuning and controlling the division of work between threads. With our extension, we show that we can achieve excellent parallel efficiency of above 95% across a range of cases.

*VoroTop* provides a particularly good example for our extension to VORO++. A typical *VoroTop* analysis requires computing all of the Voronoi cells, and then calculating the Weinberg vector for each one. The Weinberg vector is a relatively expensive calculation, requiring  $O(n^2)$  work for a Voronoi cell with  $n$  vertices. Thus, since the Voronoi cell and Weinberg vector computations can be processed independently and divided evenly among threads, this




Figure 3: Illustration of a 2D container covering the rectangular region  $[a_x, b_x] \times [a_y, b_y]$ . The container is further divided into a  $4 \times 4$  grid of blocks  $(i, j)$  for  $i, j \in \{0, 1, 2, 3\}$ . Each block  $(i, j)$  is indexed as  $k = i + 4j$ , so that  $k \in \{0, 1, \dots, 15\}$ . Points are spatially sorted into blocks. The yellow point is the second point in block 1.

represents an ideal scenario for multithreading. In Section 4 we demonstrate this on a 36 core server. We show that the time to process 102.4 million particles can be reduced from almost an hour to under two minutes, highlighting a dramatic practical performance benefit.

## 2. Methods

### 2.1. Overview of VORO++

We now provide an overview of design and the key methods of VORO++ [37]. The code is structured around several C++ classes for storing particles and computing Voronoi cells. The code can perform both 2D and 3D computations, and the classes responsible for these have “\_2d” and “\_3d” suffixes, respectively. The algorithmic principles are identical in both 2D and 3D. Therefore, for the rest of this section, we focus on the 2D implementation for simplicity.

#### 2.1.1. The *container\_2d* class

VORO++ has a variety of container types that represent rectangular domains holding all of the positions. The `container_2d` class holds particles in a 2D rectangle. The user can specify the coordinate ranges  $[a_x, b_x]$  and  $[a_y, b_y]$ , and indicate whether the container is periodic in each direction. The container is further divided into a rectangular grid of  $n_x \times n_y$  blocks of equal size, into which the particles are spatially sorted. In this way, every particle can be identified by its block index, and its point index inside the block. For example, in Fig. 3, the yellow point is the 2<sup>nd</sup> point in block 1.

The grid of blocks provides a large boost in performance, allowing the code to quickly locate neighboring particles during the Voronoi cell construction. See Sec. 2.1.4 for details on how the number of blocks is chosen. There is also a variant container class called `container_poly_2d` that stores polydisperse particle arrangements. Each particle has an associated radius, which can be used to compute the radical Voronoi tessellation [2].

### 2.1.2. Using a `voronoicell_2d` class to compute a Voronoi cell

The `voronoicell_2d` class represents a single Voronoi cell as a convex polygon, with a set of vertices connected by edges. The `voronoicell_2d` class contains routines for constructing the Voronoi cell, as well as routines for computing different statistics about it, such as its area or centroid.

The `voronoicell_2d` class uses a coordinate system where the origin is centered on the particle. Consider a specific particle  $P$  located at position  $\mathbf{p} = (p_x, p_y)$  within the container. To compute its Voronoi cell, the `voronoicell_2d` class is first initialized as a rectangle  $[c_x, d_x] \times [c_y, d_y]$  filling the entire container, without considering any other particles. Specifically, in the  $x$  direction,

$$[c_x, d_x] = \begin{cases} [a_x - p_x, b_x - p_x] & \text{if the } x \text{ direction is non-periodic,} \\ [-\frac{b_x - a_x}{2}, \frac{b_x - a_x}{2}] & \text{if the } x \text{ direction is periodic.} \end{cases} \quad (1)$$

For the periodic case, the maximum extent of the initial Voronoi cell is determined by the perpendicular bisectors of the periodic images of  $P$  that are displaced by  $\pm(b_x - a_x, 0)$ . Similarly, in the  $y$  direction,

$$[c_y, d_y] = \begin{cases} [a_y - p_y, b_y - p_y] & \text{if the } y \text{ direction is non-periodic,} \\ [-\frac{b_y - a_y}{2}, \frac{b_y - a_y}{2}] & \text{if the } y \text{ direction is periodic.} \end{cases} \quad (2)$$

To construct the Voronoi cell for  $P$  the code then considers the effect of neighboring particles. If a neighbor is located at  $\mathbf{q}$  relative to  $P$ , then that will remove the half-space

$$\mathbf{r} \cdot \mathbf{q} > \frac{\mathbf{q} \cdot \mathbf{q}}{2} \quad (3)$$

where  $\mathbf{r} = (x, y)$ . The boundary of the half-space, given by  $\mathbf{r} \cdot \mathbf{q} = \frac{1}{2}\mathbf{q} \cdot \mathbf{q}$ , is the perpendicular bisector between  $P$  and its neighbor. The `voronoicell_2d` class contains a routine called `plane()` that recomputes the vertices and edges of the Voronoi cell based on cutting by a plane. To compute the Voronoi cell for  $P$ , the code considers the neighboring particles and applies plane cuts based on removing the half spaces of the form given in Eq. (3), as illustrated in Fig. 4.

Hypothetically, if all plane cuts for all other particles are applied, then the `voronoicell_2d` class will precisely represent the Voronoi cell of  $P$ . In practice, it is only necessary to consider plane cuts from a small set of neighbors around  $P$ , as described in the following section.






Figure 4: Illustration of the action of the `plane()` routine. The Voronoi cell of a particle  $P$  located at  $\mathbf{p}$  is initialized as a large rectangular box filling the computational domain. The `plane()` routine repeatedly cuts down the rectangular box by planes that are the perpendicular bisectors between the particle and its neighbors, located at  $\mathbf{q}_1$  and  $\mathbf{q}_2$ .




Figure 5: Illustration of different computational costs for different Voronoi cells. (a) A partially computed Voronoi cell is shown in light purple, after considering half-space intersections from four neighboring particles. The maximum distance to a vertex is  $R$ , and thus particles that lie outside a circle of radius  $2R$  can be omitted from the computation; only a single additional particle needs to be tested. (b) If a particle does not have neighbors on all sides, its Voronoi cell may extend a long distance in one direction, so that  $R$  is much larger. Hence many more particles lie within the circle of radius  $2R$  and cannot be ruled out from the computation.

### 2.1.3. The `voro_compute` class

The `container_2d` class contains a member `vc` of class type `voro_compute_2d`, which holds the data structures for computing the Voronoi cells from the container's spatially sorted particles. For a given particle  $P$  the `voro_compute_2d` class computes the Voronoi cell following the procedure in the previous section, but using as few plane cuts as possible. As an example, consider the particle in the top left corner of Fig. 1(a): its Voronoi cell has two faces that adjoin cells for neighboring particles. Thus, if those two particles were considered first, then all remaining plane cuts would have no effect.

The `voro_compute_2d` class therefore computes Voronoi cells by considering plane cuts

from nearby points first, and then use bounds to terminate the computation as soon as possible. The class first considers particles in the same block as  $P$ , and then sweeps outwards to consider nearby blocks. When each block is considered, two bounds can be used to determine whether the Voronoi cell is complete or if more plane cuts are required:

- Radius bound – if  $R$  is the maximum distance of a Voronoi cell vertex to  $P$ , then no particles more than a distance  $2R$  away can possibly influence the cell. This bound is fast to compute, but it has no directional sensitivity: if a cell extends a long way in one direction then particles a long distance in other directions will still need to be tested.
- Block bound – a given block in the grid can be tested to see if any particle within it can possibly influence the Voronoi cell. This can be done by performing a sequence of half-space intersection tests based on the block’s corners. The code sweeps outward from  $P$ , testing blocks until it reaches those that cannot influence the cell. This computation is slower than the radius bound but it has directional sensitivity.

The `voro_compute_2d` class uses a combination of the two bounds. It begins by using the radius bound, which works effectively for particles in densely-packed regions with many close neighbors. This is illustrated in Fig. 5(a), where the `voroicell_2d` polygon is shown after considering four neighboring particles. At this point, the bounding circle of radius  $2R$  only contains a single additional particle. Thus, once this particle is considered, then the Voronoi cell will be complete and it will not be necessary to consider further particles. In contrast, Fig. 5(b) shows a case where the `voroicell_2d` polygon is extended in one direction because it is at the edge of a particle arrangement. We refer to such cases as extended Voronoi cells, where the maximum Voronoi vertex distance is large compared to a typical inter-particle separation length. In this case, the circle from the radius bound covers the whole domain, and no particles can be ruled out from the computation.

If the radius bound is not successful in rapidly terminating the computation, then the `voro_compute_2d` class switches over to the block bound. This can help cut down the number of particles to consider, but the search space that needs to be considered can still inherently be much larger. Hence extended Voronoi cells can take substantially longer to compute. It is important to consider this large difference in Voronoi cell computation time when designing the multithreaded extension.

#### 2.1.4. Choice of the block size

The size of the grid of blocks in the container affects the computation time. Let  $N$  be the total number of particles and define  $N_{\text{avg}} = N/(n_x n_y n_z)$  to be the average number of particles per block. If  $N_{\text{avg}}$  is too large, then each block contains many particles, and since particles are not spatially sorted within a block, the code must spend a long time looping through all of them. If  $N_{\text{avg}}$  is too small, then the code must search through many blocks to complete a Voronoi cell computation. The best performance is achieved by choosing  $N_{\text{avg}}$  as a balanced value between these two extremes.

In addition, the optimal performance is usually achieved when the blocks have roughly equal side lengths. Hence, for a given target average number of blocks,  $N_{\text{opt}}$ , the code chooses the number of blocks as follows:

- In 2D, set  $\lambda = \sqrt{N/(N_{\text{opt}}(b_x - a_x)(b_y - a_y))}$  and define  $n_x = \lceil \lambda(b_x - a_x) \rceil$  and  $n_y = \lceil \lambda(b_y - a_y) \rceil$ .
- In 3D, set  $\lambda = \sqrt[3]{N/(N_{\text{opt}}(b_x - a_x)(b_y - a_y)(b_z - a_z))}$  and define  $n_x = \lceil \lambda(b_x - a_x) \rceil$ ,  $n_y = \lceil \lambda(b_y - a_y) \rceil$ , and  $n_z = \lceil \lambda(b_z - a_z) \rceil$ .

The use of the ceiling operator  $\lceil \cdot \rceil$  ensures that the grid dimensions are always greater than zero. For homogeneous random particle arrangements, the best performance is achieved for  $N_{\text{opt}} = 3.4$  in 2D and  $N_{\text{opt}} = 4.6$  in 3D, although the code is not that sensitive to this choice and good performance is achieved across a wide range. These values are used by default in the code, although they can be overridden by the user. In some of the examples in this paper, which are chosen to highlight different scenarios, we determine different values of  $N_{\text{opt}}$  that improve performance.

### 2.1.5. Procedure for Voronoi cell computation

A typical usage of the library is as follows:

- I. Initialize the container. Insert particles into the container, spatially sorting them into the grid of blocks.
- II. Loop over the blocks in the container, and for each block:
  - i. Loop over the particles in the block, and for each particle:
    - a. Calculate the Voronoi cell of the particle.
    - b. Compute and store required statistics about the Voronoi cell.

Each Voronoi cell computation is independent of the others, making VORO++ highly suitable for parallel computation. A straightforward parallelization approach is to distribute particles to different threads, and compute their Voronoi cells simultaneously.

## 2.2. Multi-threaded extensions

### 2.2.1. Changes to code architecture

We now describe the changes required to make VORO++ multithreaded using OpenMP. The `vc` member within the container class, which is a `voro_compute_2d` class, is responsible for calculating the Voronoi cell of a particle. The `vc` member allocates workspace for searching through the blocks for neighboring particles. Thus it is not thread-safe, since if two threads used the same `vc` member, they would generate race conditions on the workspace.

Therefore we create copies of `voro_compute_2d` class object in the container, based on the number of threads being used. `vc` is no longer a single `voro_compute_2d` class, but becomes an array of them. The class constructor accepts an additional argument `num_t` that determines the number of `voro_compute_2d` classes to allocate. Thread `k` then uses `vc[k]` to compute its Voronoi cells. If needed, the function `change_number_thread(num_t)` can be used to reallocate the number of `voro_compute_2d` classes available.

Moreover, in a typical loop like the one in Sec. 2.1.5, the code creates a variable `c` that is a `voronocell_2d` class, for representing the Voronoi cell of the particle. Similar to `vc`,

different threads in computation cannot use the same `voronoicell_2d` object. This issue is solved by creating thread-private copies of `c` in the parallel computation.

### 2.2.2. Random access iterator and OpenMP parallelization

The OpenMP directive `#pragma omp parallel` creates a team of threads to execute a section of code. Within a parallel section, the directive `#pragma omp for` can be placed before a `for` loop, to distribute the iterations of the loop to the different threads. Listing 1 demonstrates how to use these two directives to parallelize the filling of an array `c` with square roots of the integers.

```
1 double c[256];
2 #pragma omp parallel
3 {
4 #pragma omp for
5     for(int i=0;i<256;i++) {
6         c[i]=sqrt(double(i));
7     }
8 }
```

Listing 1: Short example demonstrating basic OpenMP directives

Each thread will be assigned a subset of values of `i` to set in the array. Since each array entry can be set independently, this code can be multithreaded without resulting in a race condition. In basic usage like the example above, the `#pragma omp for` directive is placed before loops over integers. However, since version 3.0 of the OpenMP standard, it is possible to parallelize a `for` loop using any C++ random access iterator.

In C++, each class can have associated iterator classes that are designed to iterate over the elements of that class. There are several types of iterator, differentiated by how much functionality they offer. The simplest is the forward iterator, which supports basic operations for stepping forward sequentially. The iterator `a` represents an index of the associated class, and the forward iterator must support the operation `a++` to step forward to the next index. Thus, in the context of a `VORO++` container, an iterator would store the block index and point index of a particle, and the iterator would support the operation `a++` to step forward to the next particle, as illustrated in Fig. 6(a). Full requirements for the functionality of a forward iterator are available in the C++ documentation [77].

In the multithreaded extension of `VORO++`, we created random access iterators on all of the container classes to iterate over all of the particles. These iterators support all of the functionality of a forward iterator, but also contain additional functions for making arbitrary jumps in the particle indexing. A key ability required by the random access iterator is to evaluate `a+n` for an integer `n`, allowing the iterator to jump forward by `n` steps in the index. This additional functionality is required since the threads need to start at different points within a parallelized `for` loop (Fig. 6(b)). Figure 6(c) illustrates how this works for a `VORO++` container.

The `VORO++` iterators contain all the functions listed in the random access iterator standard [78]. This includes the dereference operators `*a` and `a[n]`, which resolve to the member of the class the iterator points to. However, in the current context, dereferencing




Figure 6: (a) Illustration of using a forward iterator to represent particles in the `container_2d` class of VORO++. The forward iterator only allows the particle index to be stepped forward one at a time. If the iterator `a` is pointing at the yellow particle, then `a++` will step forward to the red particle. (b) Requirement of iterator `a` to represent particles in the container of VORO++. The iterator needs to be able to access particles at any arbitrary offset position relative to the current particle that the iterator is pointing at. For ten particles in the container and two parallel threads, thread 1 starts at `a` and thread 2 starts at `a+5`. (c) Illustration of using a random access iterator to represent particles in the container of VORO++. The random access iterator can access particles of any arbitrary offset positions relative to the current particle that it is pointing at. Here, `a+8` steps forward eight particles to the red particle.

is conceptually unclear, since the Voronoi cell associated with the iterator does not exist in memory and must be subsequently computed. Because of this, the dereference functions within the C++ iterators simply emit errors if they are called. For VORO++, the iterators are used to index into the container's particles and loop through them, and thus the dereference operators are not required in normal usage.

In addition to the standard iterator that loops through the particles, two variations are provided: `iterator_subset` that can loop over a subset of particles, and `iterator_order` that can loop over an ordered list of particles.

### 2.2.3. Example implementation

An example of the multi-threaded version of VORO++ is provided in Listing 2. The example demonstrates how to compute the Voronoi cells of a random 2D particle arrangement, and then calculate their average perimeter:

```

1 #include <stdio>
2 #include <stdlib>
3
4 #include "voro++.hh"
5 using namespace voro;
6
7 // Returns floating point number uniformly distributed over [0,1)

```

```

8 inline double rnd() {return (1./RAND_MAX)*static_cast<double>(rand());}
9
10 int main() {
11
12     // Number of parallel threads
13     int num_t=4;
14
15     // Number of particles to use
16     int N=100000;
17
18     // Construct a 2D container as a periodic unit square divided into a
19     // 160x160 grid of blocks. Each block initially holds up to 8 particles. The
20     // final argument sets the number of voro_compute objects for use by the
21     // threads.
22     container_2d con(0.0,1.0,0.0,1.0,160,160,true,true,8,num_t);
23
24     // Add particles to the container
25     for(int i=0;i<N;i++) con.put(i,rnd(),rnd());
26
27     // Declare iterator
28     container_2d::iterator cli;
29
30     // Parallel Voronoi computation to compute the average Voronoi cell
31     // perimeter
32     double tperim=0.;
33 #pragma omp parallel num_threads(num_t)
34     {
35         // Thread-private Voronoi cell object and perimeter counter
36         voronoicell_2d c(con);
37         double perim=0.;
38
39         // Iterate through the particles
40 #pragma omp for
41         for(cli=con.begin();cli<con.end();cli++)
42             if(con.compute_cell(c,cli))
43                 perim+=c.perimeter();
44
45         // Add local perimeter counter to global perimeter counter using atomic
46         // operation to prevent race condition
47 #pragma omp atomic
48         tperim+=perim;
49     }
50
51     // Print average Voronoi cell perimeter
52     printf("Average Voronoi cell perimeter is %.12g\n",tperim/N);
53 }

```

Listing 2: Example code of multi-threaded VORO++

After creating the container on line 22, we then add  $N = 10^5$  particles at random on line 25. We then use OpenMP on the random access iterator, and distribute particles to different

	Assign workloads to threads evenly	Assign workloads to threads unevenly
(a) Even workloads		
(b) Uneven workloads		

Table 1: Schematic illustration of load balancing among threads. (a) If iterations have identical workloads, they should be assigned to threads evenly for higher efficiency. (b) If iterations have uneven workloads, they should be assigned to threads unevenly, following some strategies, so that the workloads among threads are balanced.

threads for parallel computation. Lines 27–49 demonstrate using the iterator for parallel computation. On line 28 we declare an iterator of the container. We then use the `for` loop on the iterator to loop through all the particles in the container, as shown in lines 40 & 41. For each particle, we compute its Voronoi cell (line 42) and then add its perimeter to a local counter, `perim`. Note that in certain situations (*e.g.* when occluded by a wall) a Voronoi cell may not exist. In this case the `compute_cell` function returns false, and subsequent computation should be skipped.

For parallelization, we follow the standard OpenMP syntax. We mark the start of the parallel block with `#pragma omp parallel`, and add the `#pragma omp for` directive before the `for` loop. This simple implementation following standard OpenMP and C++ practices is beneficial to the user. At the end of the parallel block, on line 48, each thread adds its local perimeter counter to a global counter, `tperim`. Since the threads are all writing to the global counter, this operation is performed atomically to avoid a race condition. Once the parallel block has terminated, the average perimeter is reported as  $\approx 0.01265$ . If the program is run with different particle numbers  $N$ , one finds that the perimeter asymptotically behaves like  $4/\sqrt{N}$ . Surprisingly, even though this is for a random particle arrangement, this asymptotic value matches the perimeter of a  $\sqrt{N} \times \sqrt{N}$  grid of squares covering the domain.

#### 2.2.4. Load balancing

Load balancing is important for achieving good parallel performance. If the iterations have even workloads, as shown in Table 1(a), they should be assigned to threads evenly. In this way, all threads are busy working during the full duration of the computation, and all threads finish at approximately same time. In comparison, if the iterations are assigned unevenly to threads, some threads will finish early and remain idle waiting for others to finish, resulting in inefficient performance. On the contrary, if iterations have uneven workloads, as shown in Table 1(b), then assigning an equal number of iterations to each thread may result in substantial idle time. Better performance can be achieved by distributing the iterations so that the total workloads are balanced.

Parallel strategy	Traits	Schematic illustration
<code>schedule(static)</code>	<ol style="list-style-type: none"> <li>1. Iterations are pre-assigned evenly to threads.</li> <li>2. Good for identical workloads.</li> <li>3. Low overhead costs.</li> </ol>	
<code>schedule(dynamic)</code>	<ol style="list-style-type: none"> <li>1. Thread takes work as available.</li> <li>2. Good for imbalanced workloads.</li> <li>3. Higher overhead costs.</li> </ol>	
<code>schedule(guided)</code>	<ol style="list-style-type: none"> <li>1. “Mixed” strategy: Iterations are divided into chunks with decreasing chunk sizes, for threads to grab.</li> <li>2. Some imbalanced workloads.</li> <li>3. Higher overhead costs.</li> </ol>	

Table 2: Three basic built-in OpenMP parallel strategies, their traits and schematic illustrations of the work assignment process.

OpenMP provides a number of strategies for load balancing. Here we explain in detail the three most basic strategies, `schedule(static)`, `schedule(dynamic)`, `schedule(guided)`, with Table 2 showing an overview of their properties. For the `schedule(static)`, iterations (which for VORO++ correspond to particles) are pre-assigned to the threads by dividing the entire set of work units into equal sections. This strategy is good for identical workloads, and it has low overhead costs, since once the work is assigned, each thread can operate independently. For the `schedule(dynamic)`, each thread takes a unit of work as available and is never idle. If a thread finishes its work, it immediately takes the next iteration available. This strategy comes with higher overhead costs since the threads must coordinate with each other when taking more work. `schedule(guided)` is a mixed strategy. Threads take work as available and are never idle, but rather than taking one unit of work per time, each thread takes a chunk. Moreover, the chunk sizes decrease as the assignment progresses. This is good for when there is some imbalance in workloads, but does not offer as much flexibility as `schedule(dynamic)`. It also has high overhead costs.

Any of these strategies can be used with the multi-threaded VORO++ extension, by adding the corresponding keyword after the `#pragma omp for` directive. For example, in Listing 2, changing line 40 to `#pragma omp for schedule(dynamic)` will enable the dynamic assignment. In addition, there are two modified strategies that accept an additional integer `chunk` option. In `schedule(static,chunk)` the threads each take `chunk` work units at once to process, rather than dividing up the total work into even sections. This still does not require any thread coordination, but affects the work distribution.<sup>1</sup> In `schedule(dynamic,chunk)`

<sup>1</sup>For example, consider two threads operating on the integers  $(0, 1, \dots, 15)$ . With `schedule(static)` the work is given to the two threads as  $(0, 1, \dots, 7)$  and  $(8, 9, \dots, 15)$ . With `schedule(static,3)` the first thread receives the chunks  $(0, 1, 2)$ ,  $(6, 7, 8)$ ,  $(12, 13, 14)$  and the second thread receives the chunks  $(3, 4, 5)$ ,  $(9, 10, 11)$ ,  $(15)$  with the final chunk being truncated.



the threads take `chunk` units of available work at once. This cuts down on the overhead costs, since there are fewer points where the threads must coordinate. Depending on the arrangement of particles, one strategy may be preferred over another. In particular, Sec. 2.1.3 showed that for inhomogeneous particle distributions, the Voronoi cell computation can vary by a large factor. In this case, we expect that the dynamic or guided scheduling strategy will be advantageous.

### 2.2.5. Parallel insertion of particles

In Listing 2 the time spent on computing the Voronoi cells is much larger than the time spent on inserting the particles into the container. Nevertheless, once the particle system grows larger, the time for particle insertion becomes substantial. Because of this, we also developed an approach for inserting particles into the container in parallel.

As shown on line 25 of Listing 2, the `put(...)` function is used to insert a particle into the container in serial. Internally, this function first computes the block index `k` that the particle lies within. The container class has two array entries associated with the block: `co[k]` is the total number of particles currently in the block, and `mem[k]` is the total number of memory slots allocated for the block, typically initialized to 8 at the start of the computation. The `put(...)` function therefore does the following:

- I. If `co[k]==mem[k]` then all memory slots are full. In this case, dynamically reallocate the memory for this block, and double the available slots.
- II. Add the particle information in the `co[k]` slot, and then increment `co[k]`.

This routine is not thread-safe for two reasons: (a) if two threads try to reallocate the memory simultaneously it will be corrupted, and (b) if two threads try and increment `co[k]` simultaneously this will create a race condition.

In the multithreaded extension, we created a function `put_parallel(...)` that can insert particles into the container in a thread-safe manner. This function also requires an overflow buffer for storing some particles. The `put_parallel(...)` does the following:

- I. Use the `#pragma omp atomic capture` directive on the operation `l=co[k]++`. This function will atomically increment `co[k]` while storing its previous value in the variable `l`.
- II. If `l>=mem[k]` there is no available memory slot. Store the particle in the overflow buffer.
- III. If `l<mem[k]`, store the particle into slot `l`.

In step I., the atomic capture directive ensures that each thread will obtain a unique slot number to write into. For example, if `co[k] = 2` and two threads insert a particle into this block, then one is guaranteed to obtain `l = 2` and the other is guaranteed to obtain `l = 3`. Afterward `co[k] = 4`.

Within `put_parallel(...)`, it is not possible to dynamically reallocate the memory, since this could interfere with other threads. Hence, if a block has no available slots, it must

be stored into the overflow buffer. Assuming that the block memory was roughly allocated correctly beforehand, the number of particles going to the overflow buffer should be a small fraction of the total. The overflow buffer operations are done within an OpenMP `critical` section of code, so that only one thread operates on the buffer at one time. If the overflow buffer runs out of space, this thread can dynamically extend the buffer without causing a race condition. Once all of the particles have been added with `put_parallel(...)`, it is necessary to call a function `put_reconcile_overflow()`. This function operates serially on the overflow buffer, and dynamically extends the block memory as needed to ensure there are enough available slots to insert any remaining particles.

Using parallel insertion can cause the code to produce slightly different results when it is repeated. Depending on the exact timing of the threads, which differs from run to run, the particles may be inserted into the blocks in a different order. This reordering can have a minor effect on the floating point errors that are incurred during the Voronoi cell construction. The code also contains a function `add_parallel(pt_list,num,nt_)` that takes in a pointer `pt_list` to array of `num` particle positions stored as  $(x, y, z)$  triplets, with an additional radius argument for the polydisperse containers. It performs a multi-threaded insertion of the particles into the container using `nt_` threads.

### 3. Parallel Performance

#### 3.1. Voronoi cell computation

We first examine the performance of the multithreaded code for three representative non-periodic particle arrangements, using both a 2D domain  $[0, 1]^2$  and 3D domain  $[0, 1]^3$ . They are:

- I. *Homogeneous* –  $10^8$  particles randomly and homogeneously distributed throughout the domain.
- II. *Localized* – We define a square of area 0.1 in 2D, or a cube of volume 0.1 in 3D, positioned in the center of the domain. We take the  $10^8$  particles from the previous case, and keep only those within this region, resulting in approximately  $10^7$  particles in total.
- III. *Extreme clustering* – We use an uneven distribution clustered around the regions  $x = \frac{1}{2}$ ,  $y = \frac{1}{2}$ , and  $z = \frac{1}{2}$  (for 3D only). This is achieved by generating each particle coordinate as a non-uniformly distributed random variable. In the  $x$  direction, the number  $X$  is chosen to be uniformly distributed over  $[0, 1]$ , and then the particle coordinate is set to  $x = \frac{1}{2} + 4(X - \frac{1}{2})^3$ , resulting in clustering near  $x = \frac{1}{2}$ . The same procedure is used in the other directions.

Figure 7 shows examples of the three particle arrangements in both 2D and 3D, but with a reduced number of particles for ease of visualization. The grid of blocks is set according to the method described in Sec. 2.1.4, but with specific choices of  $N_{\text{opt}}$  are shown in Table 3. For the homogeneous case we use the default choices of  $N_{\text{opt}}$ . For the other two cases, we scan

Container initialization	2D		3D	
	$N_{\text{opt}}$	$n_x, n_y$	$N_{\text{opt}}$	$n_x, n_y, n_z$
(a) Homogeneous	3.4	5,423	4.6	279
(b) Localized	30	1,825	49	126
(c) Extreme clustering	0.5	14,142	0.5	584

Table 3: Choices of  $N_{\text{opt}}$  and the size of the grid of blocks for the three different example particle configurations, in 2D and 3D.

over a range of values of  $N_{\text{opt}}$  and choose the one that results in the fastest execution. Since the structure of these two cases differs considerably from the homogeneous arrangement, we find that the best values of  $N_{\text{opt}}$  are quite different. For example, for the extreme clustering case we use the much smaller value of  $N_{\text{opt}} = 0.5$ , since this provides better spatial sorting of the clustered particles.

Here, and throughout Sec. 3 all tests were performed on a Ubuntu Linux computer with dual Intel Xeon E5-2650L v4 processors with 14 low-power cores and a 1.7 GHz base clock speed. We measure the time to compute all of the Voronoi cells, without performing any analyses on them. For each case, we compare the five OpenMP load balancing strategies introduced in Sec. 2.2.4. For the two chunk-based strategies, we first scan over a range of chunk sizes from 1 to 100 with the 28-thread parallel code, and choose the one with the lowest computation time. We report this chunk size in our results.

The relevant measures of parallel performance are:

- I. The wall-clock time  $t_p$  required to compute all Voronoi cells using  $p$  threads.
- II. The parallel efficiency calculated as

$$T_e(p) = \frac{t_1}{p \cdot t_p}, \quad (4)$$

which measures the effective slowdown from the hypothetically perfect parallel scaling.

The computation time and parallel efficiency are affected by Intel’s Turbo Boost technology, which boosts the clock speed based on the number of active cores. For the E5-2650L v4 chip, the clock speed is boosted by 0.8 GHz when a single core is in use, but this is reduced to 0.3 GHz when six or more cores are used. This creates a clear signature in the parallel efficiency data, since the wall-clock times for many threads are increased because the CPU slows down. We explore this issue in detail in Appendix A. Since we aim to examine the performance of the code independent of hardware intricacies, we work with the *adjusted parallel efficiency*, which factors out the Turbo Boost effect. This is computed as

$$T_e^{\text{adj}}(p) = T_e(p)A(p), \quad (5)$$

where

$$A(p) = \frac{\text{average clock speed for 1 thread}}{\text{average clock speed for } p \text{ threads}}. \quad (6)$$

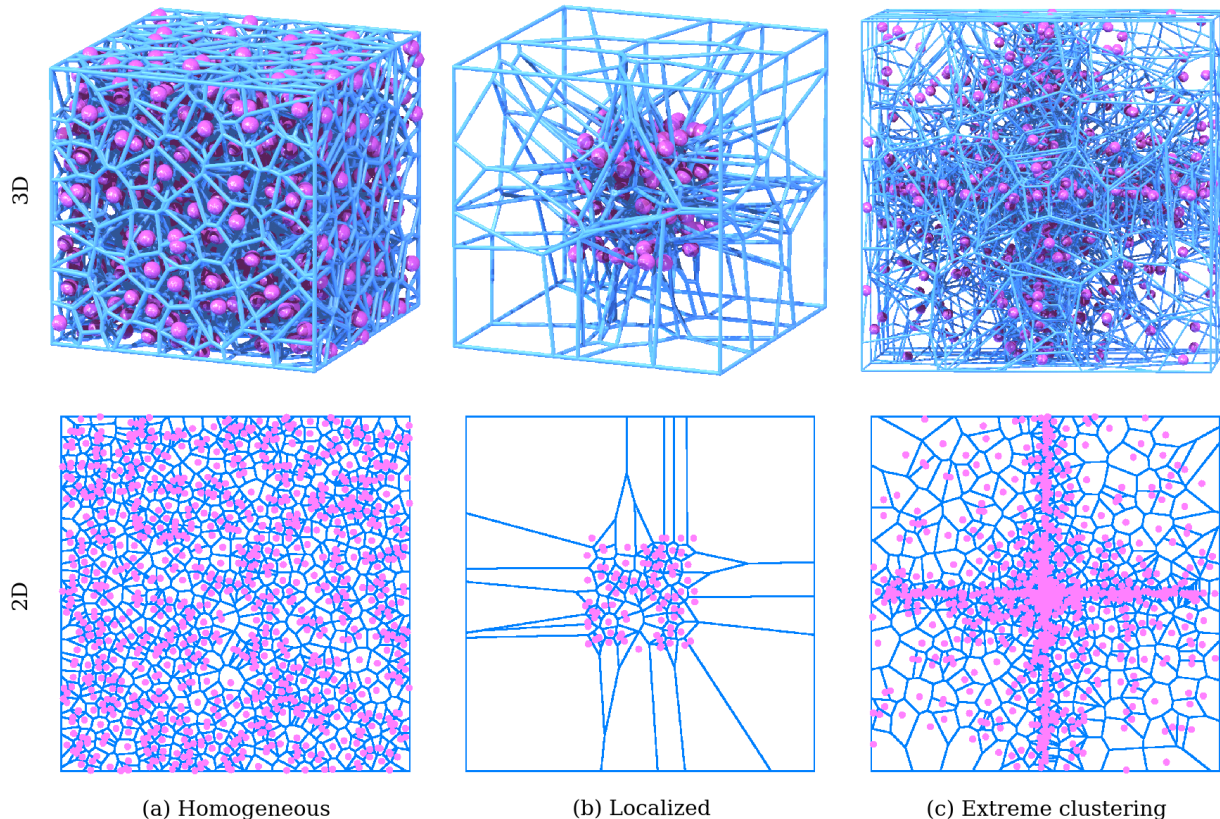


Figure 7: Illustration of three particle distribution cases in both 2D and 3D, but scaling down the particle number by  $1/10^5$  for ease of visualization. (a) Homogeneous and random particle distribution with 1000 particles. (b) Localized and random particle distribution with 95 particles. (c) Extreme clustering particle distribution with 1000 particles.

The details of how the average clock speeds are computed are provided in Appendix A. We account for Turbo Boost in the parallel efficiency, but we report the original wall-clock computation times, since they give a reasonable estimate for the Voronoi cell computation in a real world setting where Turbo Boost is enabled by default.

### 3.1.1. Comparison of the three example cases

Figure 8 shows the adjusted parallel efficiency of the five strategies as a function of the number of threads, for the three different particle arrangements in both 2D and 3D. Since the computer has 28 physical cores, we find that the parallel efficiency drops substantially beyond 28 threads in all cases, and thus we limit our discussion to the results for the first 28 threads. Table 4 shows the computation time using the serial code, and the optimal 28-thread parallel code for the three cases in both 2D and 3D. For each 28-thread time, we report the parallel strategy that achieves this optimal result.

*3.1.1.1. Comparison of different particle arrangements.* Figure 8 shows that for the homogeneous particle distribution the parallel efficiencies are similar among the strategies, and




Figure 8: Adjusted parallel efficiency against number of threads of the five load balancing strategies, for the three example cases in both 2D and 3D. (a) Homogeneous and random particle distribution. (b) Localized and random particle distribution. (c) Extreme clustering particle distribution.

Wall-clock time (s)	2D			3D		
	1 thread	28 threads	Opt. strategy	1 thread	28 threads	Opt. strategy
(a) Homogeneous	150.53	7.03	(guided)	2099.72	93.34	(guided)
(b) Localized	39.88	2.03	(dynamic,38)	589.62	26.85	(dynamic)
(c) Extreme clustering	1030.12	45.45	(dynamic,60)	10368.70	464.96	(dynamic,18)

Table 4: Voronoi computation time with the serial code and the 28-thread parallel code for the three example cases in both 2D and 3D, using each case’s optimal strategy.

all are in the high range above 80% in 2D and 95% in 3D. In comparison, for the localized particle distribution, the parallel efficiencies are very different, with `schedule(guided)` and `schedule(static)` performing significantly worse than the others. This large difference is also true for the extreme clustering case. For a container with random and homogeneous particle arrangement, we expect that the computation time for each Voronoi cell to be similar, as the Voronoi cells are of similar size. Therefore, particles have similar workloads. No matter what strategies we use to distribute particles to threads, the threads have similar workloads as well, thus resulting in the similar performance among the five strategies.

However, for the random and localized distribution of particles, where the sizes of Voronoi cells are only restricted by the container, many particles in the outer layer have extended Voronoi cells that take much longer to compute as described in Sec. 2.1.3. The imbalance in the workload causes `schedule(static)` and `schedule(guided)` to perform substantially worse, with parallel efficiencies dropping below 50%. `schedule(dynamic)` performs much better, achieving efficiencies above 90% in 2D and 95% in 3D. In addition, both `schedule(dynamic,chunk)` and `schedule(static,chunk)` can have similar parallel performance to `schedule(dynamic)`, if optimal chunk sizes are chosen. Table 4 shows that the total computation times for the localized case are less than the homogeneous case, although the localized case only has a tenth of the particles of the homogeneous case. The localized case takes considerably longer per particle.

For the extreme clustering case, there are many more particles concentrating in the grid blocks in the domain center. The Voronoi computation for these particles is substantially more expensive than the Voronoi computation for particles in less dense regions. When `VORO++` computes Voronoi cell of a particle, it will loop through all neighboring particles in the nearby grid blocks (Sec. 2.1.3), and thus more particles will be considered. In addition, the uneven distribution creates imbalanced workload since the particles in the center will take longer to compute than those in the periphery. Therefore, similar to the localized case, we see that `schedule(dynamic)` has huge benefits and performs much better than `schedule(static)` and `schedule(guided)`. Also, `schedule(static,chunk)` and `schedule(dynamic,chunk)` are able to perform as well as `schedule(dynamic)`, if optimal chunk sizes are chosen. Table 4 shows that the total time for the extreme clustering case is considerably longer than for the homogeneous case.

*3.1.1.2. Comparison of 2D and 3D Voronoi computation.* 3D Voronoi cells are much more complicated than 2D Voronoi cells. A 2D Voronoi cell is a simple polygon, where every vertex is connected to two others. By contrast, representing a 3D Voronoi cell requires an edge table with varying connectivity between the vertices. This difference is clear in the timing data for the homogeneous case in Table 4: the serial code computes 664000 cells per second in 2D, but only 47600 cells per second in 3D.

Since each 3D Voronoi cell computation involves a larger amount of work, the code can achieve higher threading efficiency in 3D, because comparatively less time is lost to overhead in parallelizing the loop. In Fig. 8 for the homogeneous case with `schedule(dynamic)`, the 2D computations only achieve efficiencies of  $\approx 75\%$  for large thread counts, whereas the 3D computations achieve efficiencies over 95%. The data in Fig. 8 shows that the




Figure 9: Adjusted parallel efficiency as function of the number of particles. The data obtained is from the homogeneous particle distribution case in both 2D and 3D, both using `schedule(guided)`. While parallel efficiency must theoretically be less than or equal to one, small random fluctuations in timing cause several data points in this graph to exceed this value.

`schedule(dynamic, chunk)` strategy substantially improves the efficiency of the 2D computations, since proportion of time spent on thread assignment will be reduced. With this strategy, the 2D computations can achieve thread efficiencies approaching 95%, as in the 3D case.

### 3.1.2. Code performance for different system sizes

We now examine the performance of the code for a random homogeneous particle system when the total number of particles  $N$  is varied. For each case, the grid of blocks is again set according to the method described in Sec. 2.1.4, and the values of  $N_{\text{opt}}$  remain 3.4 for 2D and 4.6 for 3D. We use the `schedule(guided)` strategy, which was previously shown to achieve very good parallel efficiency when  $N = 10^8$ , in both 2D and 3D (Fig. 8). We compute parallel efficiencies by comparing the performance of a single thread to the performance of 28 threads.

As shown in Fig. 9, the parallel efficiencies remain roughly constant over the range from  $N = 10^5$  to  $N = 10^8$ . The Voronoi cell computations are well suited to parallel computation, and even for  $N = 10^5$  there is sufficient work for the particles to be parallelized effectively. As expected, the parallel efficiencies seen in 2D are lower than those in 3D. Some small fluctuations in parallel efficiency are observed in Fig. 9, particularly for lower particle counts in 2D. Because the wall clock times for these computations are shorter, they are more susceptible to small variations in computation time.

### 3.1.3. Effects of clipping extended Voronoi cells

To illustrate further the effect of extended Voronoi cells on the computation cost, we examine another case where the Voronoi cells are clipped by a bounding volume. This is demonstrated in Fig. 10 where the Voronoi cells for the 3D localized arrangement of Fig. 7




Figure 10: Illustration of the 3D localized particle distribution test from Fig. 7, but where the Voronoi cells are each clipped by a bounding cube of side length  $h$  as described in the text. Here, the total number of particles has been scaled down by a factor of  $10^5$  from the computational test for ease of visualization. Correspondingly, the clipping length  $h$  has been scaled up by a factor of  $10^{5/3}$ .

are clipped by a bounding cube centered on each particle. Clipped Voronoi cells are useful in some practical situations [21, 29], such as when particles have a finite interaction range, and extended Voronoi cells provide no practical information.

VORO++ can clip Voronoi cells by any convex polyhedron centered on each particle. This is achieved by starting each Voronoi cell computation using a fixed shape, instead of making it fill the computational domain as described in Sec. 2.1.2. Here, we consider the 3D test using the localized particle arrangement shown in Fig. 7, with the domain  $[0, 1]^3$ . This configuration was made by generating  $N' = 10^8$  candidate particle positions at random, and then retaining the  $N \approx 10^7$  in a central cube covering 10% of the domain volume. We clip each Voronoi cell using a cube with side length  $h = 1.5\sqrt{3}d_0$ , where  $d_0 = (N')^{-1/3}$  represents a typical inter-particle separation length in the container. Figure 10 shows an illustration of the clipped Voronoi cells.

As shown in Fig. 11, for the same localized particles in the container, using `schedule(dynamic)`, the computation time is significantly lower for clipped Voronoi cells. For the original localized case of extended Voronoi cells, the computation time is 589.62s for one thread and 26.85s when using 28 threads. In comparison, when the Voronoi cells are clipped, the computation time decreases to 218.63s with one thread and 10.01s with 28 threads. Therefore, compared to the extended Voronoi cells, the computational costs for clipped Voronoi cells are significantly lower, and in both serial and parallel cases, there is a three-fold speedup in computation time. These timing results will depend on the size of the clipping region, with larger  $h$  requiring more time, and smaller  $h$  requiring less time. With the clipped cells, the size of the search space in the Voronoi cell construction (Sec. 2.1.3) is substantially reduced, improving the performance. Furthermore, using clipped cells removes the large disparities in computation time between different Voronoi cells, so the `schedule(static)` parallelization strategy will give good performance.



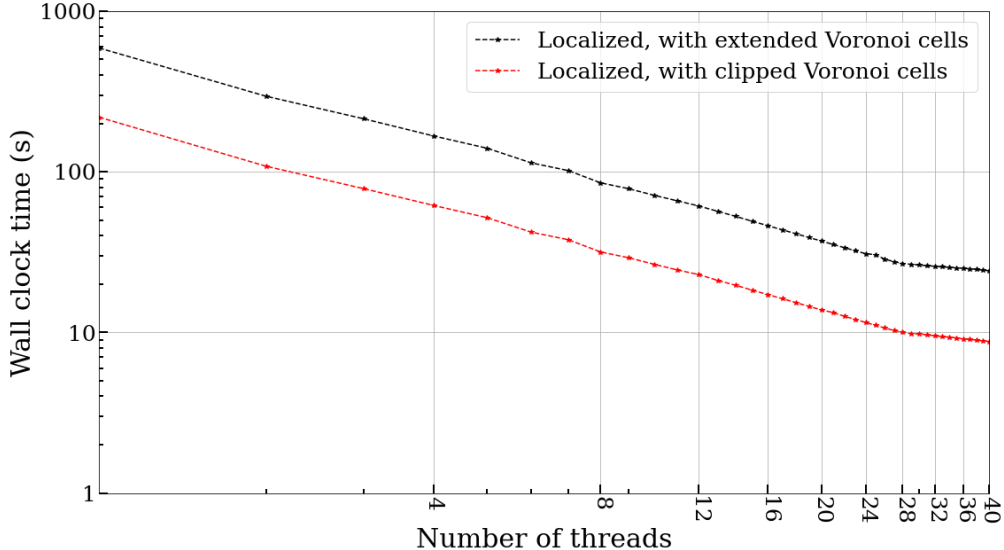


Figure 11: Computation time using different numbers of threads for the 3D localized particle distribution test using (a) the original Voronoi tessellation with many extended Voronoi cells, and (b) clipping each Voronoi cell by a bounding cube of side length  $h$  as described in the text. In both cases, the optimal strategy of `schedule(dynamic)` was used.

### 3.2. Performance of the parallel insertion routine

We now test the performance of the parallel insertion routine that was described in Sec. 2.2.5. We test in 2D and 3D using the domains  $[0, 1]^2$  and  $[0, 1]^3$ , respectively, with  $10^8$  homogeneous random particle positions. We use the default values of  $N_{\text{opt}} = 3.4$  and  $N_{\text{opt}} = 4.6$  for the 2D and 3D calculations, respectively. In each case we use an initial assignment of `mem[k]` equal to 8. The total number of particles in each block is well-approximated by a Poisson distribution,  $\text{Pois}(N_{\text{opt}})$ . This results in 0.83% and 4.51% of blocks exceeding the initial memory allocation in 2D and 3D, respectively. We expect 0.36% and 1.66% of the total particles will be placed into the overflow buffer in 2D and 3D, respectively.

We note that the performance of parallel insertion will depend substantially on many factors. Due to the small difference in  $N_{\text{opt}}$  between 2D and 3D, this results in a 4.6-fold difference in the number of particles going into the overflow buffer. Since the overflow buffer is processed serially at the end of the insertion, this will affect parallel efficiency. This will also depend upon the particle configuration; for example, molecular dynamics simulations often employ a repulsive potential at a short, well-defined length scale (*e.g.* the Lennard-Jones potential [79]). This causes particles to be spatially distributed more evenly than uniform random samples, so that the overflow buffer may not be needed. In contrast, particle configurations like the extreme clustering case of Fig. 7 may have a large fraction of particles going to the overflow buffer.

In many practical uses of VORO++, the Voronoi tessellation may be computed at multiple time points in a simulation, where the particle positions only vary slightly between successive snapshots. To handle this scenario, VORO++ has the ability to re-use the container data

Wall-clock time (s), <code>add_parallel(...)</code>	2D		3D	
	1 thread	28 threads	1 thread	28 threads
With overflow buffer	72.76	3.43	72.65	5.34
No overflow buffer	72.37	3.03	73.65	3.06

Table 5: Average computation time for the parallel particle insertion test. For each case, the insertion is performed once, when the overflow buffer will be used. The container is cleared and the insertion is repeated. Since block memory was extended during the first test, the overflow buffer will not be used on the second insertion. Each test was performed ten times, and the average wall-clock times are shown for the serial code and 28-thread parallel code in both 2D and 3D.

structure, by clearing the particles for inserting a new configuration. This is achieved by zeroing out the block-based particle counters `co[k]`, but leaving any extended memory allocation `mem[k]` in place. Thus, assuming only small movements in particles, the previously extended block memory should be well-matched to the new particle configuration and only a small number of particles will need to be placed in the overflow buffer.

To illustrate this and isolate the effect of the overflow buffer, we perform a sequence of two parallel insertion tests, in both 2D and 3D. We generate a list of  $N$  uniformly distributed particle positions, and call the `add_parallel(...)` function to add these particles into the container. On the first time, some particles will be placed into the overflow buffer, and will then be reconciled and store serially via the `put_reconcile_overflow()` function. This will result in `mem[k]` being adjusted to accommodate the additional particles. We then clear the container and make another `add_parallel(...)` call to reinsert the particles. Since the block memory was already extended to accommodate this particle configuration, zero particles will be placed into the overflow buffer on this second test. Each test was performed ten times with between 1 and 28 threads, and the average wall-clock time was recorded to calculate the average parallel efficiency. Since the parallel insertion routine involves both parallel and serial components, the changing clock frequency complicates the performance analysis. We therefore turn off Turbo Boost, so that all cores run at the same clock frequency regardless of the number of active threads. The average wall-clock time for the serial code and the 28-thread parallel code in both 2D and 3D are reported in Table. 5.

### 3.2.1. Efficiency against number of threads

Figure 12 show the results of parallel insertion in 2D and 3D using  $N = 10^8$  particles with a varying number of threads. On the first time, when the overflow buffer is used, the parallel efficiency in 2D is approximately 76% using 28 threads. In 3D, parallel efficiency is lower, and gradually decreases to 29% when 28 threads are used.

On the second time, without using the overflow buffer, the 28-thread parallel efficiency is approximately 85% in both 2D and 3D. These results indicate how the overflow buffer substantially affects the efficiency. The worse parallel performance in 3D is due to more particles being placed into the overflow buffer. The number of particles being placed into the overflow buffer averaged over ten tests are 356,044 and 1,675,330 for 2D and 3D, respectively, which are consistent with the aforementioned calculations using the Poisson distribution.




Figure 12: Average parallel efficiency for the parallel particle insertion test of  $10^8$  particles using a varying number of threads. For each case, the insertion is performed once, when the overflow buffer will be used. The container is cleared and the insertion is repeated. Since block memory was extended during the first test, the overflow buffer will not be used on the second insertion. Each test was performed ten times, and the average wall-clock times were recorded to calculate the average parallel efficiency. While parallel efficiency must theoretically be less than or equal to one, small random fluctuations in timing cause several data points in the 2D graph (left) to exceed this value.




Figure 13: Average parallel efficiency for the parallel particle insertion test of a varying number of particles, using 28 threads. For each case, the insertion is performed once, when the overflow buffer will be used. The container is cleared and the insertion is repeated. Since block memory was extended during the first test, the overflow buffer will not be used on the second insertion. Each test was performed ten times, and the average wall-clock times were recorded to calculate the average parallel efficiency.

### 3.2.2. Efficiency against number of particles

Figure 13 shows the parallel efficiency using 28 threads as number of particles  $N$  is varied. For both 2D and 3D, we observe a clear trend of increasing parallel efficiency for larger  $N$ , demonstrating that parallel insertion becomes more effective for larger particle configurations. We also see that with overflow buffer, the 3D case has worse parallel efficiency. Without overflow buffer, the parallel efficiencies are comparable in 2D and 3D, consistent with the results of Sec. 3.2.1.




Figure 14: A central red atom and its Voronoi cell, surrounded by gold-colored neighboring atoms.

#### 4. Application: Topological analysis using *VoroTop*

Voronoi tessellations are commonly used, among other applications, in the classification of local, structural features of atomistic data sets [57, 80, 81, 82]. Computational material scientists and physicists, for example, are often interested in identifying localized defects in materials [57], or else in characterizing the structure of nominally disordered systems such as liquids [83, 84], granular packings [39, 85], and soft glassy materials [86, 87]. Voronoi cells provide a proxy for describing local arrangements through consideration of their geometric and combinatorial properties [31, 88, 89].

Voronoi topology provides a robust method of classifying local structure [57, 60]. In particular, the types of faces of a Voronoi cell and the manner in which they are arranged describe the manner in which neighbors are arranged relative to a central particle and to one another. Figure 14 illustrates a central red atom and its Voronoi cell, surrounded by gold-colored neighboring atoms. Each face of the Voronoi cell corresponds to a unique neighboring particle, and so the total number of faces can be considered as a count of neighbors. Furthermore, the number of edges of a face counts the number of common neighbors between the central particle and the associated neighbor. A five-sided face, for example, indicates that the central particle shares five neighbors in common with the associated particle. In finite-temperature systems, particles are located in “general position” and small perturbations of the coordinates do not change features such as numbers of faces and edges, making them particularly useful for analyzing noisy data. Voronoi topology is thus particularly well-suited for studying high-temperature crystalline systems, where conventional methods such as centrosymmetry and common-neighbor analysis are typically ineffective [90].

Fully characterizing the topological structure of a Voronoi cell can be done efficiently using an algorithm of Weinberg [66] designed to determine whether two planar graphs are isomorphic. Although the general graph-isomorphism problem is not known to be solvable in polynomial time, the edge graphs associated with Voronoi cells, as convex polyhedra, are always 3-connected and planar [91]. This enables the use of Weinberg’s graph-tracing algorithm to provide a systematic description of Voronoi cell topologies. In particular, Weinberg’s algorithm traces the edge network of a graph to produce a unique “code”; two Voronoi cells are topologically equivalent if and only if their codes are identical. This algorithm takes as input a representation of a planar graph and outputs a code whose length is twice

the number of edges of the graph. Constructing each code is done in  $O(n^2)$  time, where  $n$  is the number of faces of the Voronoi cell. Voronoi cells in uniformly distributed particle systems typically have a narrow range of faces (between 10 and 20), but this number can grow for non-uniformly distributed ones.

Calculating the Voronoi cell topology for large particle systems is a good test case for a multi-threaded version of VORO++ because the Weinberg code for each cell is computed individually after the Voronoi cell itself has been constructed. A program called *VoroTop*, developed by Lazar, uses VORO++ to calculate the Voronoi cells and then implements Weinberg’s algorithm to produce a code for each particle, enabling further structural analysis [64]. When limited to running with a single-thread, the ability to analyze large systems is limited. While smaller systems can be analyzed separately in parallel using batch runs, this approach is ill-suited for large systems where memory constraints may limit the number of systems that can be computed simultaneously. We have thus incorporated the multithreaded version of VORO++ into an updated version of *VoroTop* to take advantage of its parallel architecture.

To test the efficiency of the multithreaded VORO++ in *VoroTop*, we considered systems with between 100,000 particles and 102.4 million particles. These systems were analyzed on a single computer with two Intel Xeon Gold 6240 CPUs running at 2.60 GHz. Each of the two processors has 18 cores, allowing us to run up to 36 threads while limiting each core to a single thread. Running *VoroTop* and VORO++ with more threads than physical cores can still reduce total runtime, even while decreasing overall parallel efficiency. Running *VoroTop* and VORO++ with more than 72 threads, however, led to decreased total performance, as the overhead cost associated with running more than two threads per core outweighed gains from multithreading. To better understand the impact of the parallel version of VORO++ on this topological analysis, no further computations, such as structure classification, were performed.

As in Sec. 3.2, we turned off Turbo Boost for this test. This reduces total performance but enables a simpler comparison of efficiency of the parallel versions of VORO++ and *VoroTop*. Furthermore, we only considered the `schedule(dynamic)` strategy. We constructed systems with each given number of particles, distributed randomly and uniformly in the unit cube with periodic boundary conditions. Each system was analyzed ten times with between 1 and 50 threads, and the average wall-clock runtime was recorded for computation of the Voronoi cells and the computation of their Weinberg codes.

Figure 15 shows the average wall-clock time for each system as a function of number of threads. It can immediately be seen that the total average wall-clock time decreases with additional cores, as expected. Figure 16 shows the unadjusted parallel efficiency  $T_e$  computed via Eq. (4) for different systems sizes and numbers of threads. Parallel efficiency is over 95% when the number of threads is at most equal to the number of physical cores, in this case 36. Figure 17 illustrates the parallel efficiency as a function of system size, for different numbers of threads. This efficiency is close to 95% in systems with at least half a million particles. Voronoi topology analysis through *VoroTop* thus provides an example in which implementation of a multithreaded version of VORO++ facilitates a significant speedup in practical applications.




Figure 15: Computation time against number of threads for the different systems. The two guidelines are graphs of functions proportional to  $1/\text{threads}$ .




Figure 16: Parallel efficiency as a function of threads for systems with different numbers or particles. For systems with at least half a million particles, *VoroTop* runs at over 95% parallel efficiency when the number of threads is at most the number of physical cores.




Figure 17: Parallel efficiency as a function of number of particles for different numbers of threads. For systems with at least half a million particles, *VoroTop* runs at over 95% parallel efficiency when the number of threads is at most the number of physical cores.

## 5. Conclusion

In this paper we have introduced a multithreaded extension to the VORO++ library using the OpenMP standard. Since VORO++ constructs each Voronoi cell individually, it is ideally suited for multithreaded computations, and we demonstrate parallel efficiencies above 95% across a range of different examples in 2D and 3D. Our extension to VORO++ is designed to follow standard OpenMP programming methods, making it straightforward for users to incorporate into their own programs. In Sec. 4, we presented an example of this using the *VoroTop* software package, where both Voronoi cell calculations and topology analysis were performed within an OpenMP loop, again resulting in efficiencies above 95%.

There are a number of extensions to consider. First, as described in Sec. 2.1.1, VORO++ is based upon dividing the computational domain into a regular grid of blocks. This simple data structure is efficient to compute on. It is also well-suited to roughly even particle distributions, since each block is responsible for a similar number of particles. One of the primary usages of VORO++ has been in analysis of molecular dynamics (MD) or discrete-element method (DEM) simulation, where particles often have a fixed interaction length, and therefore tend to have roughly even distributions, leading to good performance. However, there are other scenarios, such as computing the Voronoi tessellation for gravitationally interacting particles, where particles may be heavily clustered in small parts of the domain. In this case, the particles are unevenly distributed, resulting in worse performance. This was demonstrated in the extreme clustering example, which ran substantially slower than the homogeneous example despite having the same number of particles (Table 4). One possibility to improve

this would be to sort the particles into a structure like a  $k$ -d tree [92] or quad/octree [93], which would allow clustered regions of particles to be adaptively refined. This would entail a more complicated search through the blocks when computing a Voronoi cell, since the blocks would no longer be of uniform size. Applying a restriction in the tree construction (*e.g.* using a graded grid [94, 95]) may help simplify the implementation and reduce the number of cases that need to be considered. Since the tree data structure would even out the particles among the blocks, it would likely improve the parallel efficiency of the OpenMP implementation.

Another possible extension is to explore distributed memory calculations, which would allow the library to handle even larger particle systems, such as those generated by supercomputing facilities. Furthermore, many codes that are designed for supercomputers already use distributed parallel programming, such as via the Message Passing Interface (MPI). Thus if VORO++ could work in a distributed parallel setting, it could directly interface with the parallel simulation, rather than having to save the data and reload it for analysis. The issue of saving and reloading data is an important performance bottleneck in distributed parallel computing, and has previously been addressed in related work in the context of GPU parallelization [82].

Computing Voronoi cells in a distributed parallel environment is considerably more challenging than the multithreaded case and requires different algorithms. It is necessary to communicate to neighboring processors to obtain some of their particle positions. For dense particle arrangements, it should usually be sufficient for each processor to obtain particles in a region of fixed width beyond its domain, and then use those to apply plane cuts during the Voronoi cell construction. This is the approach taken by the `voronoi` command in LAMMPS [43]. However, in some situations, such as when a Voronoi cell extends out by a long way in one direction (Fig. 5(b)), the fixed-width region will be insufficient to completely determine the Voronoi cells. Because of this the parallel implementation would also have to allow for the possibility of non-local communication, where a processor could potentially talk to any other to obtain its particle information. These are interesting parallel computation challenges that can be explored in future work.

## Acknowledgements

This research was supported by a grant from the United States–Israel Binational Science Foundation (BSF), Jerusalem, Israel through grant number 2018/170. C. H. Rycroft was partially supported by the Applied Mathematics Program of the U.S. DOE Office of Advanced Scientific Computing Research under contract number DE-AC02-05CH11231. Additional support from the Data Science Institute in Bar-Ilan University is also gratefully acknowledged.

## Appendix A Turbo Boost and parallel efficiency

Modern Intel CPUs make use of the Turbo Boost technology, which boosts the CPU clock speed from its base value, depending on many factors (*e.g.* CPU temperature and available power). Most significantly, Turbo Boost depends on the number of cores in use. Our test system has dual Intel Xeon E5-2650L v4 processors with a base clock speed of 1.7 GHz.



Cores in use	Normal frequency each core (GHz)	Turbo Boost frequency each core (GHz)
1-2	1.7	2.5
3		2.3
4		2.2
5		2.1
6-14		2.0

Table 6: Normal clock frequency and Turbo Boost frequency for Intel Xeon E5-2650L: v4 microprocessor used in the performance tests.

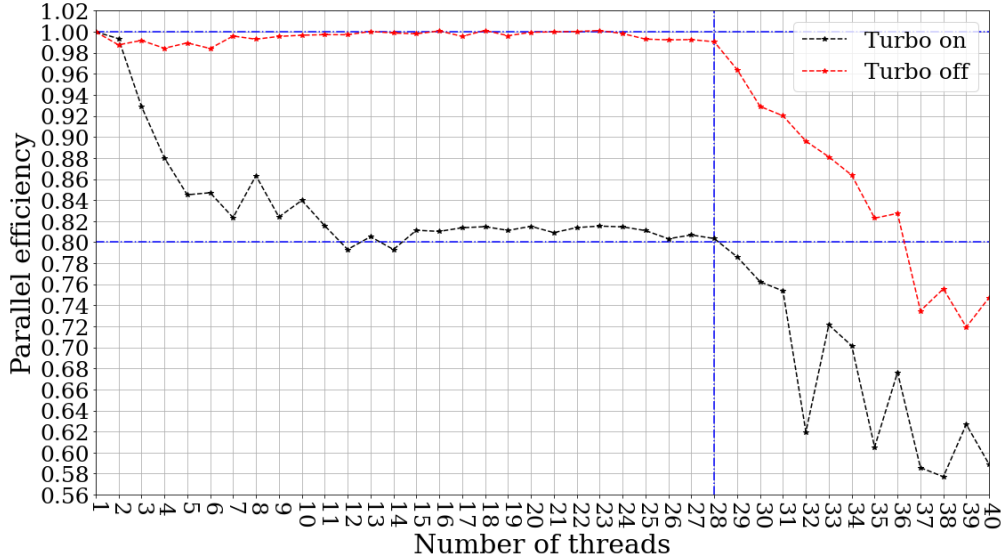


Figure A.1: Comparison of parallel efficiencies against number of threads, with Turbo Boost on, and with Turbo Boost off. The data obtained is from the homogeneous particle distribution case in 3D, using its optimal strategy `schedule(guided)`. The data with Turbo Boost off is obtained from four runs of the same test. We take the minimum computation time of the four runs against number of threads, and calculate the parallel efficiency using the minimum time. This avoids small timing fluctuations and gives a more representative picture in parallel performance.

Table 6 shows the boosted clock speeds in terms of the number of cores in use, ranging from 2.5 GHz for a single core to 2.0 GHz for six or more cores.

Turbo Boost has a noticeable effect on the measurement of parallel efficiency. Figure A.1 shows the parallel efficiency for the random homogeneous test case with  $10^8$  particles with the `schedule(guided)` strategy, evaluated using Eq. (4). Once the number of threads reaches 12, the parallel efficiency plateaus at  $\approx 80\%$ . This is precisely in agreement with the ratio of clock speeds from 2.5 GHz to 2.0 GHz. This suggests that the reduction in efficiency is almost entirely explained by Turbo Boost, rather than effects of the VORO++ parallelization. To verify this, we switched off Turbo Boost so that the clock speed remains at 1.7 GHz regardless of the number of cores in use. Figure A.1 shows the parallel efficiency for this case. Even though the simulation is slower overall, the parallel efficiencies remain close to 100%,

confirming that Turbo Boost is responsible for the loss in parallel efficiency. Across many of the timing results, Turbo Boost has a clear signature in the unadjusted parallel efficiency, with plateaus at 80% visible (Fig. A.2).

For the timing tests, we aim to use a measure of parallel efficiency that is not closely tied to the computer hardware. One approach is to switch off Turbo Boost completely, but this results in an overall loss of performance, and is not reflective of typical computer usage. We therefore decided to adjust the efficiency calculation to factor out the Turbo Boost as given in Eq. (5), scaling the usual measure by a factor of  $A(p)$  for  $p$  threads. Our first approach to construct the adjustment was by using the frequencies of Table 6, so that

$$A(p) = \frac{\text{boost frequency for a single thread}}{\text{boost frequency for } p \text{ threads}}. \quad (7)$$

However, this approach is not suitable for the less efficient strategies. For example, the `schedule(guided)` strategy of the 2D localized case has low efficiencies. During this computation, some threads become idle while waiting for other threads to finish their computation, and therefore the Turbo Boost frequency changes depend on the number of active cores. Since the adjustment in Eq. (7) assumes that all threads remain active, it can overestimate the parallel efficiency.

We therefore used a more general and accurate way to calculate the adjusted parallel efficiency by replacing the boost frequencies in Eq. (7) with average clock speeds of the computation. Average clock speed represents the true computational speed, and it is computed by dividing the total unhalted core cycles by the total task-clock time. Unhalted core cycles is a hardware performance counter, and can be measured for each thread using `perf`, a Linux performance analyzing tool [96]. The number of unhalted core cycles for a thread captures the CPU cycles run on the thread when the thread is active, which is an estimate of the work that the thread contributes to the parallel computation of the program. Specifically, we use `perf_event_open` with the `PERF_COUNT_HW_CPU_CYCLES` configuration option for measurement [97]. Task-clock time is a software performance counter, and it can be measured for each thread using the `PERF_COUNT_SW_TASK_CLOCK` configuration option with `perf_event_open` [97]. The task-clock time of a thread reports a clock count specific to the task that is running on that thread.

We then measure unhalted core cycles  $C_{\text{serial}}$  and the task-clock time  $T_{\text{serial}}^{\text{task}}$  that the computation uses for the serial code. For the parallel code, we measure the unhalted cycles  $C_i$  and the task-clock time  $T_i^{\text{task}}$  for each thread  $i = 1, \dots, p$ . Then we compute

$$\text{average clock speed for 1 thread} = \frac{C_{\text{serial}}}{T_{\text{serial}}^{\text{task}}}, \quad (8)$$

$$\text{average clock speed for } p \text{ threads} = \frac{\sum_{i=1}^p C_i}{\sum_{i=1}^p T_i^{\text{task}}}, \quad (9)$$

which are used in Eq. (6) to compute the adjustment factor.

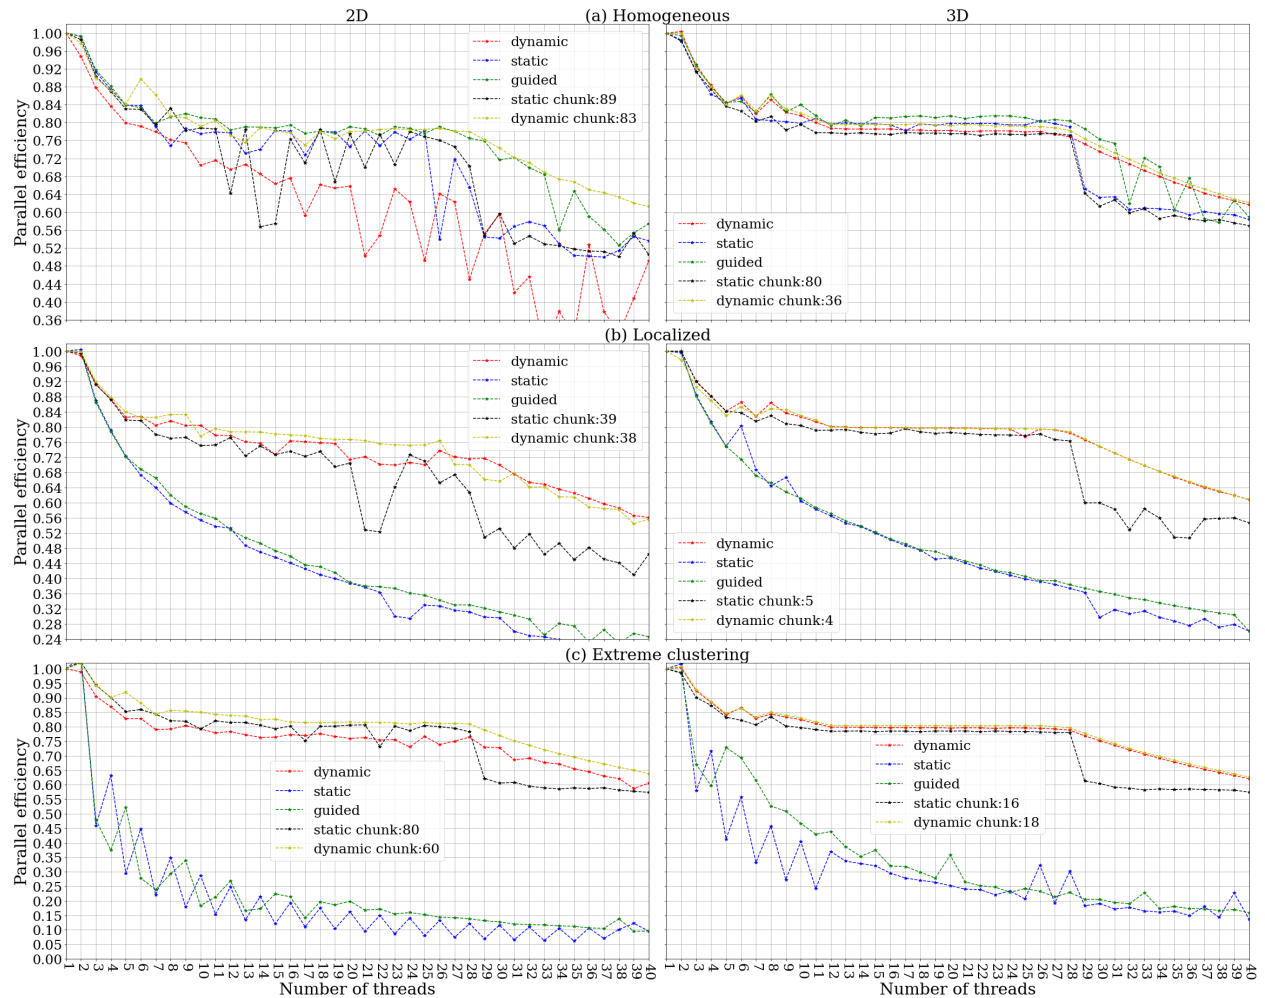


Figure A.2: Unadjusted version of the parallel efficiency shown in Fig. 8, for the three particle arrangement cases in both 2D and 3D. (a) Homogeneous and random particle distribution. (b) Localized and random particle distribution. (c) Extreme clustering particle distribution.

## References

- [1] G. Voronoi, Nouvelles applications des paramètres continus à la théorie des formes quadratiques, *Journal für die Reine und Angewandte Mathematik* 133 (1907) 97–178.
- [2] A. Okabe, B. Boots, K. Sugihara, S. N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, 2009.
- [3] L. P. Chew, R. L. S. Dyrsdale, Voronoi diagrams based on convex distance functions, in: *Proceedings of the First Annual Symposium on Computational Geometry, SCG '85*, Association for Computing Machinery, New York, NY, USA, 1985, pp. 235–244. [doi:10.1145/323233.323264](https://doi.org/10.1145/323233.323264).
- [4] R. Klein, D. Wood, Voronoi diagrams based on general metrics in the plane, in: R. Cori,

- M. Wirsing (Eds.), STACS 88, Springer Berlin Heidelberg, Berlin, Heidelberg, 1988, pp. 281–291. doi:10.1007/BFb0035852.
- [5] J. G. Puckett, F. Lechenault, K. E. Daniels, Local origins of volume fraction fluctuations in dense granular materials, *Phys. Rev. E* 83 (2011) 041301. doi:10.1103/PhysRevE.83.041301.
- [6] N. Guo, J. Zhao, Local fluctuations and spatial correlations in granular flows under constant-volume quasistatic shear, *Phys. Rev. E* 89 (2014) 042208. doi:10.1103/PhysRevE.89.042208.
- [7] C. H. Rycroft, A. Dehbi, T. Lind, S. Güntay, Granular flow in pebble-bed nuclear reactors: Scaling, dust generation, and stress, *Nuclear Engineering and Design* 265 (2012) 69–84. doi:10.1016/j.nucengdes.2013.07.010.
- [8] B. Rajaram, A. Mohraz, Steady shear microstructure in dilute colloid-polymer mixtures, *Soft Matter* 8 (2012) 7699–7707. doi:10.1039/C2SM25936B.
- [9] C. L. Phillips, C. R. Iacovella, S. C. Glotzer, Stability of the double gyroid phase to nanoparticle polydispersity in polymer-tethered nanosphere systems, *Soft Matter* 6 (2010) 1693–1703. doi:10.1039/B911140A.
- [10] R. Kramb, L. Ward, K. Jensen, R. Vaia, D. Miracle, Structural property comparison of Ca–Mg–Zn glasses to a colloidal proxy system, *Acta Materialia* 61 (18) (2013) 6911–6917. doi:10.1016/j.actamat.2013.08.003.
- [11] X. Z. Gao, M. H. Müser, L. T. Kong, J. F. Li, Atomic structure and energetics of amorphous–crystalline CuZr interfaces: a molecular dynamics study, *Modelling and Simulation in Materials Science and Engineering* 22 (6) (2014) 065007. doi:10.1088/0965-0393/22/6/065007.
- [12] C. Ruscher, J. Baschnagel, J. Farago, The Voronoi liquid, *Europhysics Letters* 112 (6) (2015) 66003. doi:10.1209/0295-5075/112/66003.
- [13] A. Wysocki, R. G. Winkler, G. Gompper, Cooperative motion of active Brownian spheres in three-dimensional dense suspensions, *Europhysics Letters* 105 (4) (2014) 48004. doi:10.1209/0295-5075/105/48004.
- [14] T. J. Yoon, M. Y. Ha, E. A. Lazar, W. B. Lee, Y.-W. Lee, Topological characterization of rigid–nonrigid transition across the Frenkel line, *The Journal of Physical Chemistry Letters* 9 (22) (2018) 6524–6528. doi:10.1021/acs.jpcllett.8b02715.
- [15] T. J. Yoon, M. Y. Ha, W. B. Lee, Y.-W. Lee, E. A. Lazar, Topological generalization of the rigid–nonrigid transition in soft-sphere and hard-sphere fluids, *Physical Review E* 99 (5) (2019) 052603. doi:10.1103/PhysRevE.99.052603.

- [16] I. Benedetti, M. H. Aliabadi, A three-dimensional grain boundary formulation for microstructural modeling of polycrystalline materials, *Computational Materials Science* 67 (2013) 249–260. doi:10.1016/j.commatsci.2012.08.006.
- [17] J. Orend, F. Hagemann, F. Klose, B. Maas, H. Palkowski, A new unified approach for modeling recrystallization during hot rolling of steel, *Materials Science and Engineering: A* 647 (2015) 191–200. doi:10.1016/j.msea.2015.08.085.
- [18] V. Gulizzi, C. Rycroft, I. Benedetti, Modelling intergranular and transgranular microcracking in polycrystalline materials, *Computer Methods in Applied Mechanics and Engineering* 329 (Supplement C) (2018) 168–194. doi:10.1016/j.cma.2017.10.005.
- [19] P. Budkewitsch, P.-Y. Robin, Modelling the evolution of columnar joints, *Journal of Volcanology and Geothermal Research* 59 (3) (1994) 219–239. doi:10.1016/0377-0273(94)90092-2.
- [20] Y. Feng, M. Založnik, B. Thomas, A. Phillion, A 3D discrete-element model for simulating liquid feeding during dendritic solidification of steel, *IOP Conference Series: Materials Science and Engineering* 529 (2019) 012031. doi:10.1088/1757-899x/529/1/012031.
- [21] W. Murphy, C. Carroll, M. Keidar, Simulation of the effect of plasma species on tumor growth and apoptosis, *Journal of Physics D: Applied Physics* 47 (47) (2014) 472001. doi:10.1088/0022-3727/47/47/472001.
- [22] A. Guittet, M. Lepilliez, S. Tanguy, F. Gibou, Solving elliptic problems with discontinuities on irregular domains – the Voronoi interface method, *Journal of Computational Physics* 298 (2015) 747–765. doi:10.1016/j.jcp.2015.06.026.
- [23] R. Loubère, P.-H. Maire, M. Shashkov, J. Breil, S. Galera, ReALE: A reconnection-based arbitrary-Lagrangian–Eulerian method, *Journal of Computational Physics* 229 (12) (2010) 4724–4761. doi:10.1016/j.jcp.2010.03.011.
- [24] T. Ringler, L. Ju, M. Gunzburger, A multiresolution method for climate system modeling: application of spherical centroidal Voronoi tessellations, *Ocean Dynamics* 58 (5–6) (2008) 475–498. doi:10.1007/s10236-008-0157-2.
- [25] C. M. Freeman, K. L. Boyle, M. Reagan, J. Johnson, C. Rycroft, G. J. Moridis, MeshVoro: a three-dimensional Voronoi mesh building tool for the TOUGH family of codes, *Computers & Geosciences* 70 (0) (2014) 26–34. doi:10.1016/j.cageo.2014.05.002.
- [26] L. Blanco-Martín, R. Wolters, J. Rutqvist, K.-H. Lux, J. T. Birkholzer, Thermal–hydraulic–mechanical modeling of a large-scale heater test to investigate rock salt and crushed salt behavior under repository conditions for heat-generating nuclear waste, *Computers and Geotechnics* 77 (2016) 120–133. doi:10.1016/j.compgeo.2016.04.008.
- [27] P. Camps, M. Baes, W. Saftly, Using 3D Voronoi grids in radiative transfer simulations, *Astronomy & Astrophysics* 560 (2013) A35. doi:10.1051/0004-6361/201322281.

- [28] E. Teruel, R. Aragues, G. López-Nicolás, A practical method to cover evenly a dynamic region with a swarm, *IEEE Robotics and Automation Letters* 6 (2) (2021) 1359–1366. [doi:10.1109/LRA.2021.3057568](https://doi.org/10.1109/LRA.2021.3057568).
- [29] R. Honeyager, G. Liu, H. Nowell, Voronoi diagram-based spheroid model for microwave scattering of complex snow aggregates, *Journal of Quantitative Spectroscopy and Radiative Transfer* 170 (2016) 28–44. [doi:10.1016/j.jqsrt.2015.10.025](https://doi.org/10.1016/j.jqsrt.2015.10.025).
- [30] R. Votel, D. Barton, T. Gotou, T. Hatanaka, M. Fujita, J. Moehlis, Equilibrium configurations for a territorial model, *SIAM Journal on Applied Dynamical Systems* 8 (3) (2009) 1234–1260. [doi:10.1137/070710123](https://doi.org/10.1137/070710123).
- [31] E. A. Lazar, J. Lu, C. H. Rycroft, Voronoi cell analysis: The shapes of particle systems, *American Journal of Physics* 90 (6) (2022) 469–480. [doi:10.1119/5.0087591](https://doi.org/10.1119/5.0087591).
- [32] C. B. Barber, D. P. Dobkin, H. Huhdanpaa, The quickhull algorithm for convex hulls, *ACM Trans. Math. Softw.* 22 (4) (1996) 469–483. [doi:10.1145/235815.235821](https://doi.org/10.1145/235815.235821).
- [33] <http://www.qhull.org/>.
- [34] <https://www.cgal.org>.
- [35] J. R. Shewchuk, Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator, *Applied Computational Geometry: Towards Geometric Engineering*, First ACM Workshop on Applied Computational Geometry, *Lecture Notes in Computer Science* 1148 (1996) 203–222. [doi:10.1007/BFb0014497](https://doi.org/10.1007/BFb0014497).
- [36] <https://www.cs.cmu.edu/~quake/triangle.html>.
- [37] C. H. Rycroft, Voro++: A three-dimensional Voronoi cell library in C++, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 19 (2009) 041111. [doi:10.1063/1.3215722](https://doi.org/10.1063/1.3215722).
- [38] <http://math.lbl.gov/voro++/>.
- [39] C. H. Rycroft, G. S. Grest, J. W. Landry, M. Z. Bazant, Analysis of granular flow in a pebble-bed nuclear reactor, *Phys. Rev. E* 74 (2006) 021306. [doi:10.1103/PhysRevE.74.021306](https://doi.org/10.1103/PhysRevE.74.021306).
- [40] C. H. Rycroft, Multiscale modeling in granular flow, Ph.D. thesis, Massachusetts Institute of Technology (2007).
- [41] C. H. Rycroft, Y. L. Wong, M. Z. Bazant, Fast spot-based multiscale simulations of granular drainage, *Powder Technology* 200 (1-2) (2010) 1–11. [doi:10.1016/j.powtec.2010.01.009](https://doi.org/10.1016/j.powtec.2010.01.009).
- [42] S. Plimpton, Fast parallel algorithms for short-range molecular dynamics, *Journal of Computational Physics* 117 (1) (1995) 1–19. [doi:10.1006/jcph.1995.1039](https://doi.org/10.1006/jcph.1995.1039).

- [43] <https://www.lammps.org>.
- [44] A. Stukowski, Visualization and analysis of atomistic simulation data with OVITO – the Open Visualization Tool, *Modelling Simul. Mater. Sci. Eng.* (18) (2010) 015012. doi:10.1088/0965-0393/18/1/015012.
- [45] <https://www.ovito.org/>.
- [46] S. Fortune, A sweepline algorithm for Voronoi diagrams, in: *Proceedings of the Second Annual Symposium on Computational Geometry, SCG '86*, ACM, New York, NY, USA, 1986, pp. 313–322. doi:10.1145/10515.10549.
- [47] S. Fortune, A sweepline algorithm for Voronoi diagrams, *Algorithmica* 2 (1-4) (1987) 153–174. doi:10.1007/BF01840357.
- [48] P. J. Green, R. Sibson, Computing Dirichlet tessellations in the plane, *The Computer Journal* 21 (2) (1978) 168–173. doi:10.1093/comjnl/21.2.168.
- [49] D. Lee, B. Schachter, Two algorithms for constructing a Delaunay triangulation, *International Journal of Computer & Information Sciences* 9 (3) (1980) 219–242. doi:10.1007/BF00977785.
- [50] R. Kimmel, J. A. Sethian, Computing geodesic paths on manifolds, *Proceedings of the National Academy of Sciences* 95 (15) (1998) 8431–8435. doi:10.1073/pnas.95.15.8431.
- [51] R. Kimmel, J. A. Sethian, Optimal algorithm for shape from shading and path planning, *Journal of Mathematical Imaging and Vision* 14 (3) (2001) 237–244. doi:10.1023/A:1011234012449.
- [52] R. Kimmel, J. A. Sethian, Fast Voronoi diagrams and offsets on triangulated surfaces, in: P. Laurent, P. Sablonniere, L. Schumaker (Eds.), *International Conference on Curves and Surfaces*, Vanderbilt Press, Saint-Malo, France, 2001, pp. 193–202.
- [53] D. Rhynsburger, Analytic delineation of Thiessen polygons, *Geographical Analysis* 5 (2) (1973) 133–144. doi:10.1111/j.1538-4632.1973.tb01003.x.
- [54] J. L. Bentley, B. W. Weide, A. C. Yao, Optimal expected-time algorithms for closest point problems, *ACM Trans. Math. Softw.* 6 (4) (1980) 563–580. doi:10.1145/355921.355927.
- [55] B. Boots, D. Murdoch, The spatial arrangement of random Voronoi polygons, *Computers & Geosciences* 9 (3) (1983) 351–365. doi:10.1016/0098-3004(83)90006-7.
- [56] M. P. Quine, D. F. Watson, Radial generation of  $n$ -dimensional Poisson processes, *Journal of Applied Probability* 21 (3) (1984) 548–557. doi:10.2307/3213616.

- [57] E. A. Lazar, J. Han, D. J. Srolovitz, Topological framework for local structure analysis in condensed matter, *Proceedings of the National Academy of Sciences* 112 (43) (2015) E5769–E5776. doi:[10.1073/pnas.1505788112](https://doi.org/10.1073/pnas.1505788112).
- [58] S. Lloyd, Least squares quantization in pcm, *Information Theory, IEEE Transactions on* 28 (2) (1982) 129–137. doi:[10.1109/TIT.1982.1056489](https://doi.org/10.1109/TIT.1982.1056489).
- [59] Q. Du, V. Faber, M. Gunzburger, Centroidal Voronoi tessellations: Applications and algorithms, *SIAM Review* 41 (4) (1999) 637–676. doi:[10.1137/S0036144599352836](https://doi.org/10.1137/S0036144599352836).
- [60] E. A. Lazar, J. K. Mason, R. D. MacPherson, D. J. Srolovitz, Complete topology of cells, grains, and bubbles in three-dimensional microstructures, *Phys. Rev. Lett.* 109 (2012) 095505. doi:[10.1103/PhysRevLett.109.095505](https://doi.org/10.1103/PhysRevLett.109.095505).
- [61] E. A. Lazar, J. K. Mason, R. D. MacPherson, D. J. Srolovitz, Statistical topology of three-dimensional Poisson-Voronoi cells and cell boundary networks, *Phys. Rev. E* 88 (2013) 063309. doi:[10.1103/PhysRevE.88.063309](https://doi.org/10.1103/PhysRevE.88.063309).
- [62] H. Leipold, E. A. Lazar, K. A. Brakke, D. J. Srolovitz, Statistical topology of perturbed two-dimensional lattices, *Journal of Statistical Mechanics: Theory and Experiment* 2016 (4) (2016) 043103. doi:[10.1088/1742-5468/2016/04/043103](https://doi.org/10.1088/1742-5468/2016/04/043103).
- [63] E. A. Lazar, A. Shoan, Voronoi chains, blocks, and clusters in perturbed square lattices, *Journal of Statistical Mechanics: Theory and Experiment* 2020 (10) (2020) 103204. doi:[10.1088/1742-5468/abb6e3](https://doi.org/10.1088/1742-5468/abb6e3).
- [64] E. A. Lazar, VoroTop: Voronoi cell topology visualization and analysis toolkit, *Modelling and Simulation in Materials Science and Engineering* 26 (1) (2017) 015011. doi:[10.1088/1361-651x/aa9a01](https://doi.org/10.1088/1361-651x/aa9a01).
- [65] <https://www.vorotop.org>.
- [66] L. Weinberg, On the maximum order of the automorphism group of a planar triply connected graph, *SIAM Journal on Applied Mathematics* 14 (4) (1966) 729–738. doi:[10.1137/0114062](https://doi.org/10.1137/0114062).
- [67] J. C. Phillips, D. J. Hardy, J. D. C. Maia, J. E. Stone, J. V. Ribeiro, R. C. Bernardi, R. Buch, G. Fiorin, J. Hénin, W. Jiang, R. McGreevy, M. C. R. Melo, B. K. Radak, R. D. Skeel, A. Singharoy, Y. Wang, B. Roux, A. Aksimentiev, Z. Luthey-Schulten, L. V. Kalé, K. Schulten, C. Chipot, E. Tajkhorshid, Scalable molecular dynamics on CPU and GPU architectures with NAMD, *The Journal of Chemical Physics* 153 (4) (2020) 044130. doi:[10.1063/5.0014475](https://doi.org/10.1063/5.0014475).
- [68] N. Kondratyuk, V. Nikolskiy, D. Pavlov, V. Stegailov, Gpu-accelerated molecular dynamics: State-of-art software performance and porting from Nvidia CUDA to AMD HIP, *The International Journal of High Performance Computing Applications* (April 2021). doi:[10.1177/10943420211008288](https://doi.org/10.1177/10943420211008288).



- [69] T. F. Willems, C. H. Rycroft, M. Kazi, J. C. Meza, M. Haranczyk, Algorithms and tools for high-throughput geometry-based analysis of crystalline porous materials, *Microporous and Mesoporous Materials* 149 (1) (2012) 134–141. doi:10.1016/j.micromeso.2011.08.020.
- [70] M. Pinheiro, R. L. Martin, C. H. Rycroft, A. Jones, E. Iglesia, M. Haranczyk, Characterization and comparison of pore landscapes in crystalline porous materials, *Journal of Molecular Graphics and Modelling* 44 (2013) 208–219. doi:10.1016/j.jmngm.2013.05.007.
- [71] M. Pinheiro, R. L. Martin, C. H. Rycroft, M. Haranczyk, High accuracy geometric analysis of crystalline porous materials, *CrystEngComm* 15 (2013) 7531–7538. doi:10.1039/C3CE41057A.
- [72] R. Jalem, M. Nakayama, Y. Noda, T. Le, I. Takeuchi, Y. Tateyama, H. Yamazaki, A general representation scheme for crystalline solids based on Voronoi-tessellation real feature values and atomic property data, *Science and Technology of Advanced Materials* 19 (1) (2018) 231–242, pMID: 29707064. doi:10.1080/14686996.2018.1439253.
- [73] D. Starinshak, J. Owen, J. Johnson, A new parallel algorithm for constructing Voronoi tessellations from distributed input data, *Computer Physics Communications* 185 (12) (2014) 3204–3214. doi:10.1016/j.cpc.2014.08.020.
- [74] R. González, PARAVT: Parallel Voronoi tessellation code, *Astronomy and Computing* 17 (2016) 80–85. doi:10.1016/j.ascom.2016.06.003.
- [75] L. Dagum, R. Menon, OpenMP: an industry standard API for shared-memory programming, *IEEE Computational Science and Engineering* 5 (1) (1998) 46–55. doi:10.1109/99.660313.
- [76] <https://www.openmp.org>.
- [77] <http://www.cplusplus.com/reference/iterator/RandomAccessIterator>.
- [78] <https://cplusplus.com/reference/iterator/ForwardIterator/>.
- [79] J. E. Lennard-Jones, Cohesion, *Proceedings of the Physical Society* 43 (5) (1931) 461–482. doi:10.1088/0959-5309/43/5/301.
- [80] W. Brostow, M. Chybicki, R. Laskowski, J. Rybicki, Voronoi polyhedra and Delaunay simplexes in the structural analysis of molecular-dynamics-simulated materials, *Physical Review B* 57 (21) (1998) 13448.
- [81] A. Poupon, Voronoi and Voronoi-related tessellations in studies of protein structure and interaction, *Current Opinion in Structural Biology* 14 (2) (2004) 233–241.
- [82] M. Bernaschi, M. Lulli, M. Sbragaglia, GPU based detection of topological changes in Voronoi diagrams, *Computer Physics Communications* 213 (2016) 19–28. doi:10.1016/j.cpc.2016.11.005.

- [83] J.-P. Shih, S.-Y. Sheu, C.-Y. Mou, A Voronoi polyhedra analysis of structures of liquid water, *The Journal of Chemical Physics* 100 (3) (1994) 2202–2212.
- [84] F. W. Starr, S. Sastry, J. F. Douglas, S. C. Glotzer, What do we learn from the local geometry of glass-forming liquids?, *Physical Review Letters* 89 (12) (2002) 125501.
- [85] M. Shahinpoor, Statistical mechanical considerations on the random packing of granular materials, *Powder Technology* 25 (2) (1980) 163–176.
- [86] L. Derzsi, D. Filippi, G. Mistura, M. Pierno, M. Lulli, M. Sbragaglia, M. Bernaschi, P. Garstecki, Fluidization and wall slip of soft glassy materials by controlled surface roughness, *Physical Review E* 95 (5) (2017) 052602.
- [87] M. Lulli, R. Benzi, M. Sbragaglia, Metastability at the yield-stress transition in soft glasses, *Physical Review X* 8 (2) (2018) 021031.
- [88] M. A. Klatt, S. Torquato, Characterization of maximally random jammed sphere packings: Voronoi correlation functions, *Physical Review E* 90 (5) (2014) 052120.
- [89] S. Torquato, F. H. Stillinger, Jammed hard-particle packings: From Kepler to Bernal and beyond, *Reviews of Modern Physics* 82 (3) (2010) 2633.
- [90] A. Stukowski, Structure identification methods for atomistic simulations of crystalline materials, *Modelling and Simulation in Materials Science and Engineering* 20 (4) (2012) 045021. doi:10.1088/0965-0393/20/4/045021.
- [91] E. Steinitz, Polyeder und raumeinteilungen, *Encyk. der Math. Wiss.* 12 (1922) 38–43.
- [92] J. L. Bentley, Multidimensional binary search trees used for associative searching, *Commun. ACM* 18 (9) (1975) 509—517. doi:10.1145/361002.361007.
- [93] R. A. Finkel, J. L. Bentley, Quad trees a data structure for retrieval on composite keys, *Acta Informatica* 4 (1) (1974) 1–9. doi:10.1007/BF00288933.
- [94] H. Chen, C. Min, F. Gibou, A supra-convergent finite difference scheme for the poisson and heat equations on irregular domains and non-graded adaptive Cartesian grids, *Journal of Scientific Computing* 31 (1) (2007) 19–60. doi:10.1007/s10915-006-9122-8.
- [95] M. Theillard, C. Rycroft, F. Gibou, A multigrid method on non-graded adaptive octree and quadtree Cartesian grids, *Journal of Scientific Computing* (1) (2012) 1–15. doi:10.1007/s10915-012-9619-2.
- [96] Perf wiki, perf: Linux profiling with performance counters, [https://perf.wiki.kernel.org/index.php/Main\\_Page](https://perf.wiki.kernel.org/index.php/Main_Page) (2015).
- [97] The Linux man-pages project, perf\_event\_open(2) — Linux manual page, [https://man7.org/linux/man-pages/man2/perf\\_event\\_open.2.html](https://man7.org/linux/man-pages/man2/perf_event_open.2.html) (2021).