



Minimal Representations of Polygons and Polyhedra

W. Randolph Franklin and Salles Viana Gomes de Magalhães

5.1 Introduction

This gem presents several different data structures to represent a geometric object. They facilitate efficient parallel computation of mass properties, especially when the object is the result of an intersection, union, or overlay. The object might be a polygon or polyhedron and may have multiple nested or disjoint components. It may be a partition of a polygon or polyhedron, such as a 2D and 3D triangulation. 2D triangulations are also known as planar graphs and Geographic Information Science (GIS) maps. These ideas have been validated with several implementations in Computer-Aided Design and GIS and tested on large datasets.

Each data structure is simply a set of objects of some fixed-size type. The set has no extra structure, that is, it is unordered. The element object type is a (ordered) tuple or a (unordered) set. This gem first reviews some classic data structures such as representing a polygon \mathcal{P} as a set of edges, where each edge is an ordered pair of its vertices. Then it presents several original data structures, created by the first author unless otherwise specified.

Each data structure is designed to optimize some operation. For example, with the above data structure, \mathcal{P} 's area is easily computable by summing the areas of the triangles subtended by each edge at the origin. Each operation is expressible in the *map-reduce* form and parallelizes quite well. We have implementations of a number of these in OpenMP on a multicore Intel Xeon.

Other operations include testing point inclusion in 2D and 3D, computing the volume of the union of many isothetic cubes, and computing the areas of all the nonempty intersections of the faces of two overlapping planar graphs. These

ideas are particularly suitable for computing mass properties of Boolean combinations because they require only incomplete information about the Boolean combination, and so are easier to implement and run faster. Nevertheless, the computation is exact, apart from the usual floating point roundoff (apart from our implementations that compute exactly with big rationals).

These data structures have no pointers or trees. Explicitly storing global information like loops of edges and shells of faces is not necessary (although that information could be deduced if another application needed it).

These techniques can often reduce the storage needed to represent geometric objects, compared to representations that store complete topologies with many pointers from each element to all its adjacent elements. Since a good laptop computer may have 32GB or more of main memory, and a good workstation 256GB, one might reasonably ask, what's the point? The answer is that the datasets are also getting larger. There is an enormous speed advantage to being able to store the dataset in real memory instead of in virtual memory or on another node of a cluster. On modern supercomputers, the cost of accessing data often dominates the cost of processing it. Reductions in storage lead directly to reductions in total wall-clock time. Efficient, compact data structures improve performance considerably.

Minimizing space usage is especially true when the target platform is a Graphics Processing Unit (GPU). GPUs are not just for gamers; as of June 2021, six of the ten fastest known supercomputers [Top500 2019] have NVIDIA GPUs. The fastest of these and the second fastest overall is the Summit, which has 27,648 NVIDIA Volta V100 GPUs [Oak Ridge National Lab 2019]. So, for maximum performance, we need to target GPUs. Although the Summit has over 10PB of memory, most of it is relatively slow to access. Each GPU has only up to 96KB of fast shared memory for each of its 14 streaming multiprocessors [NVIDIA 2017]. Data too large for the fast shared memory must spill over to the main memory on the GPU, which is 100 times slower. A similar situation obtains when the platform is a network of processing nodes programmed with the Message Passing Interface. Accessing data over the network is orders of magnitude slower than accessing local data.

These ideas were partly motivated by the philosophy that a designer knows he has achieved perfection not when there is nothing left to add but when there is nothing left to take away, and as a reaction to representations storing every possible topological relation. Some of these ideas could well date to the invention of analytic geometry, although we are unaware of such prior art.

5.2 Representing One Polygon

Here are several representations for one polygon \mathcal{P} . \mathcal{P} need not be simply connected, and nested holes and islands are allowed (see Figure 5.1(a)). Nonmanifold




Figure 5.1 General polygon figures. (a) Polygon with holes and island. (b) Jordan curve point inclusion testing on set of unoriented edges. (c) Areas subtended by edges.

vertices, which have more than two edges, are also allowed if they are considered to be several vertices whose coordinates coincide.

5.2.1 Unoriented Edges

Here, \mathcal{P} is represented as the set $\mathcal{U} = \{u_i\}$ where each $u_i = \{(x_{i0}, y_{i0}), (x_{i1}, y_{i1})\}$ is one unoriented edge of \mathcal{P} , represented as the set of its endpoints. For example, the triangle with vertices $(0, 0), (1, 0), (0, 2)$ would be represented as the set $\{\{(0, 0), (1, 0)\}, \{(0, 0), (0, 2)\}, \{(1, 0), (0, 2)\}\}$. Which side of u_i is on the inside of \mathcal{P} is not stored. Since \mathcal{U} is a set, its elements are not ordered.

The set of unoriented edges suffices to uniquely represent \mathcal{P} . From a point-set view, it determines the points contained in \mathcal{P} . If the more global topology of \mathcal{P} is required, such as the edge loops, then the two edges with a common vertex can be determined, so we can walk around each loop by traversing edge-vertex-edge-vertex.... The set of loops can be determined with a union-find algorithm or by repeating the previous traversal until every edge has been used. The hierarchical nesting of loops, such as shown in Figure 5.1(a), can be determined by inclusion testing of a sample vertex of each loop against each other loop. More efficient algorithms are available if speed is a problem. However, the point of this gem is that we usually don't need this information.

This representation supports testing whether a query point q is included inside \mathcal{P} . Using the classic Jordan curve theorem [Jordan 1887], first implemented in Shimrat [1962], we test whether q is below each u_i . It runs a semi-infinite ray up from q and computes what edges it intersects. q is inside \mathcal{P} iff the number of intersections is odd; see Figure 5.1(b), where the green point is inside but the red points outside.

5.2.2 Oriented Edges

Here, \mathcal{P} is represented as the set $\mathcal{E} = \{e_i\}$, where each $e_i = ((x_{i0}, y_{i0}), (x_{i1}, y_{i1}))$ is one oriented edge of \mathcal{P} , represented as the ordered pair of its two endpoints. For example, the triangle with vertices $(0, 0), (1, 0), (0, 2)$ would be represented as the set $\{((0, 0), (1, 0)), ((1, 0), (0, 2)), ((0, 2), (0, 0))\}$. The two vertices of each edge are ordered so that when traveling from the first vertex to the second the inside of the polygon is on the left. The edge $((0, 2), (0, 0))$ is different from the edge $((0, 0), (0, 2))$.

This representation supports computing mass properties of \mathcal{P} such as *area* by summing the signed areas subtended by each edge and the origin. See Figure 5.1(c).

With oriented edges, point inclusion becomes a little easier if the order of the edges intersecting the ray drawn up from the query point q toward $y = \infty$ is known, or alternatively, all the edges intersecting closer than a certain distance, or the closest intersecting edge, are known. This can be determined with, for example, a uniform grid, often in constant expected time. By inspecting which side of the closest inspecting edge that q is on, we immediately know whether q is inside \mathcal{P} .

This can also be used to locate a point in a planar graph, that is, to compute which face (aka polygon) of a planar graph contains q . That is a subroutine of OVERPROP [Franklin et al. 1994] and in 3D of PinMesh [Magalhães et al. 2016a]. The expected preprocessing time is linear in the number of vertices; the expected query time constant.

5.2.3 Augmented Axis-aligned Vertices

This is the special case of a polygon whose edges are all horizontal or vertical. It motivates the more general case to follow. The set of vertices is insufficient to uniquely define a polygon. However, augmenting each vertex with a little extra information suffices. For axis-aligned edges, as in Figure 5.2(a), one additional bit of information is sufficient to compute mass properties such as area.




Figure 5.2 Axis-aligned and diamond polygons. (a) Axis-aligned polygon. (b) The eight cases of s . (c) Diamond polygon.

Define an augmented vertex as $a = (x, y, s)$, where (x, y) is the vertex's coordinates and s is a property of that vertex's local geometry. There are eight cases; s is $+1$ for four and -1 for the other four, as shown in Figure 5.2(b). Then the area is $\sum x_i y_i s_i$, summed over the augmented vertices of the polygon. For example, the area of Figure 5.2(a) is

$$1 \cdot 1 \cdot 1 - 1 \cdot 4 \cdot 1 + 1 \cdot 4 \cdot 2 - 1 \cdot 2 \cdot 2 + 1 \cdot 2 \cdot 3 - 1 \cdot 1 \cdot 3 = 5$$

Testing whether point q is inside is possible by computing which quadrant of each vertex includes q , combining that with that vertex's s to form a characteristic function of \mathcal{P} , and summing. (The characteristic function $\chi_{\mathcal{P}}(q) = 1$ iff q is contained in \mathcal{P} and 0 otherwise.)

5.2.3.1 Implementation Considerations

Although this data structure is provided for pedagogical purposes as a simple demonstration of the concepts, efficiently storing that extra bit per vertex may be of interest. First, we might use a separate bit vector for all the vertices' bits. This is easy and efficient with modern languages. Alternatively, if the vertex coordinates don't need the full precision of, say, a four-byte int, we might steal one bit from a coordinate field, reducing its precision to only 31 bits.

5.2.4 Augmented Diagonal Vertices

This is a similar special case of polygons whose edges all have slopes of ± 1 (see Figure 5.2(c)). As before, we supplement each vertex with one additional bit of information, used to compute mass properties such as area. Define an augmented vertex as $a = (x, y, s)$, where (x, y) is the vertex's coordinate and s is defined similarly to in Figure 5.2(b).

There is also a whole class of area functions:

$$\forall \alpha \forall \beta \forall \gamma \forall \delta \quad A = \sum_i (\alpha x_i^2 + (1 - \alpha) y_i^2 + \beta x_i + \gamma y_i + \delta) s_i.$$

Special cases include $A = \sum_i x_i^2 s_i$ and $A = \sum_i y_i^2 s_i$.

5.2.5 Augmented General Vertices, aka Vertex Neighborhoods

A similar algorithm exists for general polygons. Since the set of vertices is insufficient to uniquely represent a polygon, we will augment each vertex V with information on its local neighborhood by including a generalization of the above s . This will abstract the information about the directions that the edges leave the vertex and which sector is interior to the polygon.




Figure 5.3 Vertex neighborhood for general polygons. (a) χ function for a vertex with $\theta < \pi$. (b) χ function for a vertex with $\theta > \pi$. (c) The characteristic function of a triangle by summing χ for each vertex.

Specifically, for each vertex, v , define $\chi_v : E^2 \rightarrow \mathbb{R}$, that is, a scalar function on all points in the plane depending on how they relate to v 's neighborhood, as illustrated in Figure 5.3(a). For $\theta < \pi$:

$$\chi_v(p) = \begin{cases} \frac{1}{2} - \frac{\theta}{2\pi} & \text{if } p \text{ is in the angle between the two rays} \\ -\frac{\theta}{2\pi} & \text{if } p \text{ is in the other two wedges} \\ \frac{1}{4} - \frac{\theta}{2\pi} & \text{if } p \text{ is on either ray, but not on } v \\ 0 & \text{if } p = v \end{cases}.$$

For $\theta > \pi$:

$$\chi_v(p) = \begin{cases} \frac{\theta}{2\pi} - \frac{1}{2} & \text{if } p \text{ is in the angle between the two rays as shown in Figure 5.3(b)} \\ \frac{\theta}{2\pi} & \text{if } p \text{ is in the other two wedges} \\ \frac{\theta}{2\pi} - \frac{1}{4} & \text{if } p \text{ is on either ray, but not on } v \\ 0 & \text{if } p = v \end{cases}$$

$\theta = \pi$ does not occur since it would be a straight line with no vertex.

Adding the χ functions for all the vertices of a polygon gives a characteristic function for the polygon. Figure 5.3(c) shows an example of a triangle. This can be

extended to all polygons, including concave with multiple components and holes. Other formulae are presented in Franklin [1983, 1987] and Franklin et al. [1990]. With this characteristic function, we can do point inclusion tests.

In spirit, this slightly resembles Green's theorem. However, while Green's theorem computes a polygon's area by integrating along its boundary edges, this formula uses only the boundary vertices.

5.2.6 Vertex–Edge Direction Adjacencies, or Cusps

We can get even simpler than the above representation by splitting each vertex into two pieces called *cusps*, one for each adjacent edge. Note that we do not know the complete edge but only its angle and which side is inside. See Figure 5.4(a). Let v be the vertex's position. It is convenient to represent the extra info as two unit-vectors, t and n . t is a unit tangent vector along the edge, and n is a unit vector normal to t pointing to the inside side of the edge. n adds only one bit of new info. So, \mathcal{P} is now $\{(v, t, n)\}$. A k -gon will have $2k$ set elements.




Figure 5.4 Cusps and half-spaces. (a) 2D cusp. (b) One 3D cusp of a cube. (c) Polygon defined by half spaces of the edges.

Then the polygon's area is

$$\frac{1}{2} \sum_i (v_i \cdot t_i)(v_i \cdot n_i),$$

where \cdot is the scalar vector product here. The perimeter (boundary) length is

$$- \sum_i (v_i \cdot t_i).$$

For example, consider a rectangle with vertices $(0, 0), (2, 0), (2, 3), (0, 3)$. $(0, 0)$ has two adjacent edges. The edge from $v = (0, 0)$ to $(2, 0)$ has tangent vector $t = (1, 0)$ and normal vector $n = (0, 1)$. The edge from $v = (0, 0)$ to $(0, 3)$ has tangent vector $t = (0, 1)$ and normal vector $n = (1, 0)$. So the two cusps created from vertex $(0, 0)$ are $((0, 0), (1, 0), (0, 1))$ and $((0, 0), (0, 1), (1, 0))$. Those two cusps contribute $0 \cdot 0 + 0 \cdot 0 = 0$ to the area and $0 + 0 = 0$ to the length.

5.2.7 Halfplanes

If we extend each edge of \mathcal{P} to an infinite line and then to a halfplane of the inside side, then we may represent \mathcal{P} 's interior, in Constructive Solid Geometry (CSG) style, as a conjunctive normal form Boolean expression of them, with each line used only once [Dobkin et al. 1988]. See Figure 5.4(c), for example, where

$$\mathcal{P} = abc(d + e + f)(g + h + i) + jkl$$

may be used to test point inclusion. In more detail, a , for example, represents a half plane defined by that edge extended to infinity, with the half plane interior on the inside side of the edge. In the above formula, $+$ means union and an omitted multiplication operator, intersection.

This concept does not extend to E^3 ; the formula may have to use some lines more than once. Also, computing volume is an open question.

5.3 Representing One Polyhedron

This section presents several data structures to represent a 3D polyhedron \mathcal{P} . A complete description would include the loops of edges and shells of faces, with many explicit adjacency relations such as pointers from each face to its adjacent vertices and edges. However, that is unnecessary for the computation of many properties, and not storing it saves a lot of space. To simplify the following descriptions, assume that the faces of \mathcal{P} are triangles.

5.3.1 Unoriented Faces

To test whether \mathcal{P} contains a test point q , it suffices to run a semi-infinite ray up from q and count how many faces of \mathcal{P} it crosses (the Jordan curve algorithm). When testing q against a face of \mathcal{P} , we do not need to know which side of the face is on the interior of \mathcal{P} . That is, the set of *unoriented faces* suffices.

5.3.2 Oriented Faces

The volume of \mathcal{P} can be computed by decomposing \mathcal{P} into tetrahedra and summing their signed volumes. Each tetrahedron is defined by one face of \mathcal{P} and the coordinate origin (or any other fixed point). The volume of a tetrahedron is a determinant whose entries are simple expressions of the tetrahedron's vertices. The volume may be negative. This algorithm requires knowing the inside side of each face, that is, the faces are *oriented*.

5.3.3 Axis-aligned Faces

Perhaps the faces of \mathcal{P} are all parallel with one of the xy , xz , or yz planes. This may not be that common, but when it does happen the volume formula is particularly simple. The vertices of such a polyhedron may be either convex or concave and may open out in any of eight directions, making 16 possible types of vertices. Each type of vertex is assigned a type bit, s , in a 3D version of Figure 5.2(b). Then the volume of \mathcal{P} is $\sum_i x_i y_i z_i s_i$, summed over all the vertices.

5.3.4 Vertex Neighborhoods

For v , a trihedral vertex in E^3 , $\chi_v(p)$ is defined using the vertex angles of the adjacent faces. Let the three faces be f_1, f_2 , and f_3 , and their angles at v be α, β , and γ , respectively. Then

$$\chi_v(p) = \begin{cases} \frac{2\pi - \alpha - \beta - \gamma}{4\pi} & \text{if } p \text{ is inside all three face planes.} \\ \frac{\alpha - \beta - \gamma}{4\pi} & \text{if } p \text{ is outside } f_1 \text{ but inside } f_2 \text{ and } f_3. \\ - - - & \text{etc. for two similar cases.} \\ \frac{\alpha + \beta - \gamma}{4\pi} & \text{if } p \text{ is outside } f_1 \text{ and } f_2 \text{ but inside } f_3. \\ - - - & \text{etc. for two similar cases.} \\ \frac{\alpha + \beta + \gamma - 2\pi}{4\pi} & \text{if } p \text{ is outside all three face planes.} \end{cases}$$

5.3.5 Cusps

We can also represent polyhedron \mathcal{P} as a set of 3D *cusps*, where each cusp is an adjacency relation between a vertex, an edge, and a face. However, neither the complete edge nor the face is stored—only the definitions of the infinite lines or planes that include them plus a face’s inside side. Specifically, v is a vertex’s coordinates, t is a unit tangent vector from v along an edge, and n is a unit vector normal to t in the plane of a face of \mathcal{P} that points into the interior of that face. (Unlike in 2D, n adds more than one bit of information.) Define also one more unit vector, b , perpendicular to both n and t , pointing into the interior of \mathcal{P} . b adds only one bit of information. \mathcal{P} is represented as the set of its cusps. For example, a cube has 48 cusps since each of eight vertices is incident on three edges, each of which is incident on two faces (see Figure 5.4(b)). In general, a vertex incident on k faces will be part of $2k$ cusps.

Some mass properties may be computed as follows.

- The total edge length, $L = -\frac{1}{2} \sum_i v_i \cdot t_i$. The $\frac{1}{2}$ is needed because each edge is counted twice.
- The surface area, $A = \frac{1}{2} \sum_i (v_i \cdot t_i) (v_i \cdot n_i)$.
- The volume $V = -\frac{1}{6} \sum_i (v_i \cdot t_i) (v_i \cdot n_i) (v_i \cdot b_i)$.

The volume, area, and length formulae may be derived by dropping a perpendicular from \mathcal{O} onto each face and onto each edge and joining \mathcal{O} to each vertex. All these edges, plus the faces joining them, partition \mathcal{P} into one simplex for each cusp. The volume of each simplex is one term of the volume formula for \mathcal{P} .

5.4 Don’t We Always Know the Edges?

One might argue that, in practice, we always know the edges, and so, vertex-based representations are pointless.

However, perhaps the polygon is the output of a function, such as the Boolean intersection of two other polygons. One application is computing whether two polygons interfere or collide with each other, which happens iff their intersection area is nonzero. In this case, computing only the output vertices and their neighborhoods may be much easier than computing the output edges. Indeed, there are two types of output vertices here: (a) some input vertices and (b) some intersections of input edges. Since each output vertex must be inside an input polygon, we filter the candidate output vertices with a point inclusion function. Note that this application requires the input polygons’ edges; its output is a different representation from its input.

Additionally, computing the output edges would require partitioning the input edges at their intersections with each other, requiring sorting the intersections along each edge and then testing whether to include each resulting segment. That is much more work and also harder to express in a map-reduce form.

5.5 Applications

The above representations permit certain problems to be solved much more efficiently, and in parallel, such as computing mass properties of geometric intersections and unions. Here are two example problems, for which we have designed, implemented, and tested algorithms on large datasets. The current implementations are on multicore Intel Xeons, but the techniques would be even more useful on manycore NVIDIA GPUs. Smaller test implementations are now being developed on that platform.

5.5.1 Volume of Union of Cubes

The problem is to compute the volume of the union of tens of millions of identical isothetic cubes; the idea is shown with 30 cubes in Figure 5.5(a). The algorithm optimizes the composition of the union and the volume operations so that it is not necessary to compute the explicit output union polyhedron. It is sufficient to compute only the set of output vertices together with their neighborhoods. A vertex's neighborhood consists of the directions of any adjacent edges and faces, and for a face, which side is interior. The neighborhood does not include the other vertices on those adjacent edges and faces. Not needing to compute that is a considerable simplification. Components of the algorithm include finding all the face-face intersections and edge-face intersections among the input cubes. If the input is independently and identically distributed uniform random, then there is an equation for the expected output volume, and it agrees with what is computed. Theoretical analysis and implementation and test results are in Franklin [2004, 2005].

5.5.2 Overlay

From Franklin and de Magalhães [2019],

Suppose that a polygon \mathcal{P} has been meshed, or partitioned, into smaller faces, in two different ways, M_0 and M_1 for two different applications. Each mesh is optimal for some application and would be suboptimal for the other application. Assume that the faces of M_0 have some property, such as mass, that would be useful for the faces of M_1 . (If the density varies, then the mass is not simply the area.) One quick approximation for the mass of f' ,

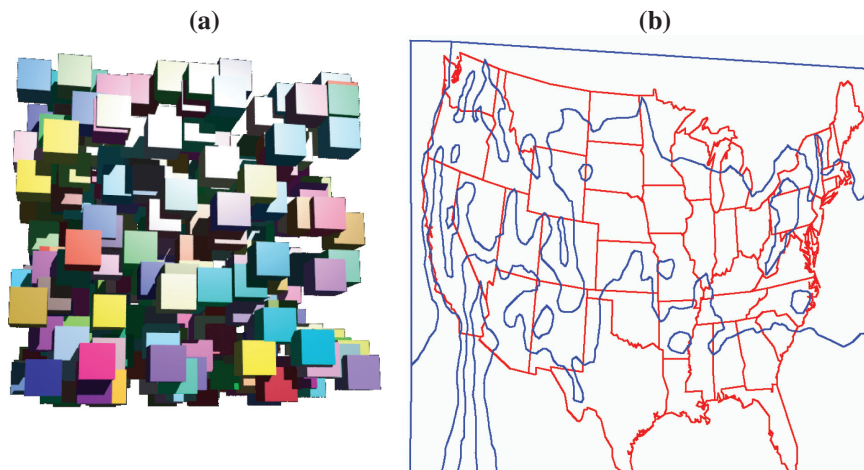


Figure 5.5 Applications. (a) Union of many cubes. (b) Overlaying state and isotherm maps.

a face of M_1 is a weighted sum of the masses of the overlapping faces of M_0 , with the weights being the areas of the intersections of f' with the overlapping faces of M_0 . The compute-bound component is to identify all the nonempty intersections of any face of M_0 with any face of M_1 , and compute their areas. PAROVER2 is an algorithm and preliminary parallel implementation for that, whose execution time is linear in the number of intersections, which is generally linear in the number of faces. PAROVER2 uses local topological formulae to compute the areas of the output faces (aka outfaces) from their vertices and half-edges.

This builds on ideas and implementations first described in Franklin et al. [1994] and Franklin and Wu [1987]. Figure 5.5(b) illustrates overlaying maps (planar graphs) for US coterminous states and July isotherms. Magalhães et al. [2015] reports on using EPUG-Overlay to intersect the surface drainage system of the United States (21,652,410 vertices, 21,060,354 edges, and 219,831 faces) with the 2010 United States Census block groups (32,762,740 vertices, 32,103,306 edges, and 518,837 faces). Computing the overlaid planar graph, excluding I/O, took 322 elapsed seconds on the dual 8 core Xeon. The parallel speedup was a factor of 11. This also used multiprecision rational numbers to avoid roundoff errors.

Franklin et al. [2018] reports on intersecting two 3D meshes using these algorithms, in addition to big rationals and other techniques. Extensions that add one or more of 3D, rational numbers to avoid roundoffs, and Simulation of Simplicity to handle geometric degeneracies include Franklin and Wu [1987], Franklin [1990, 1992], Franklin and Sivaswami [1990], Wu and Franklin [1990], Franklin

and Kankanhalli [1993], Audet et al. [2013], Magalhães and Franklin [2014], Magalhães et al. [2015, 2016b], Franklin et al. [2017a, 2017b, 2018], and Franklin and de Magalhães [2019].

5.6 Checking Polytope Validity

All these formulae require that \mathcal{P} be a legal polytope; otherwise, the volume is not even defined. Indeed, in that case, the computed volume will vary with respect to many rigid transformations. Therefore, we may test \mathcal{P} 's legality with a Monte Carlo algorithm as follows. Perform several random transformations, such as translations, scalings, or rotations, whose effect on \mathcal{P} 's volume is known. If \mathcal{P} 's volume does not always change correctly, then \mathcal{P} is illegal. Otherwise, \mathcal{P} probably does not have any errors such as missing facets. However, this does not detect errors such as \mathcal{P} 's boundary, including a double loop.

5.7 Summary

These equations are reductions; they sum a function applied to each element of a set. Therefore, they are easily parallelizable. One concern is that, especially if \mathcal{O} is well outside \mathcal{P} , many significant digits may be lost from summing offsetting positive and negative terms. Therefore, double precision, at least, is recommended.

Determining the explicit edges and faces from any of these representations is similar to determining those properties of a polytope represented as a CSG tree.

Extensions to higher E^d and also extensions to curved facets appear feasible. For example, the key to the cusp representation is that either vertex and tangent of an edge uniquely determine the infinite line, including the edge. There is a unique point on the line, the foot of a perpendicular dropped from \mathcal{O} , which is used to partition the edge into two pieces, and hence to partition \mathcal{P} into simplices. Any point that is a property only of the infinite line would serve as well as the foot of the perpendicular. For curved lines, it suffices to define a mapping from each line to a unique point on the line, and the rest of the theory follows. These minimal representations have many advantages.

Simplicity: A set of tuples is a much leaner representation than representing a hierarchy of nested face shells and edge loops.

Ease of computing the data: If \mathcal{P} is being computed as the output of a Boolean operation, determining these local data structures is easier than determining \mathcal{P} 's global information. We save perhaps half the lines of code, and time, and reduce the number of special cases.

Fewer special cases: Multiple disconnected or nested components are automatically handled without ever recognizing their existence. Therefore, there is also a smaller opportunity for erroneously handling the special cases, increasing reliability.

Faster to code: This means that fewer lines of code are necessary, reducing coding time and allowing more time for debugging and testing. Also, the data structures are fixed size and don't use pointers. This avoids slow memory management operations, which also fragment memory.

Faster to execute, especially on parallel machines: The simplicity and regular data structures allow faster execution on both serial and parallel machines. Virtual memory is not thrashed. Current data structures are too complex for easy parallelization, which prefers simple regular data structures, such as structures of arrays of plain old datatypes. If the platform is an NVIDIA GPU, then warps of 32 threads are required to execute the same instruction (or be idle). Ideally, the data used by adjacent threads is adjacent in memory. That disparages pointers, linked lists, and trees.

Efficient operations when intersecting planar graphs: Instead of processing each face separately, the fact that two adjacent faces share an edge is utilized.

References

- S. Audet, C. Albertsson, M. Murase, and A. Asahara. 2013. Robust and efficient polygon overlay on parallel stream processors. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL'13*. ACM, New York, NY, 304–313. DOI: <https://doi.org/10.1145/2525314.2525352>.
- S. V. G. de Magalhães and W. R. Franklin. 2014. Exactly computing map overlays using rational numbers. In *Autocarto 2014*, Pittsburgh, PA. Cartography and Geographic Information Society (abstract only).
- S. V. G. de Magalhães, M. V. A. Andrade, W. R. Franklin, and W. Li. 2015. Fast exact parallel map overlay using a two-level uniform grid. In *4th ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data (BigSpatial)*, Bellevue, WA. DOI: <https://doi.org/10.1145/2835185.2835188>.
- S. V. G. de Magalhães, M. V. A. Andrade, W. Randolph Franklin, and W. Li. 2016a. PINMESH—Fast and exact 3D point location queries using a uniform grid. *Comput. Graph.* 58, 1–11. Special Issue on Shape Modeling International 2016. Awarded a reproducibility stamp. DOI: <https://doi.org/10.1016/j.cag.2016.05.017>.
- S. V. G. de Magalhães, M. V. A. Andrade, W. R. Franklin, W. Li, and M.-C. G. Gruppi. 2016b. Exact intersection of 3D geometric models. In *Geoinfo 2016, XVII Brazilian Symposium on GeoInformatics*. Instituto Nacional de Pesquisas Espaciais (Brasil), Campos do Jordão, SP, Brazil.

- D. P. Dobkin, L. J. Guibas, J. Hershberger, and J. Snoeyink. 1988. An efficient algorithm for finding the CSG representation of a simple polygon. *Comput. Graph. Proc. SIGGRAPH'88*, 22, 4, 31–40. DOI: <https://doi.org/10.1145/378456.378472>.
- W. R. Franklin. 1983. Rays—New representation for polygons and polyhedra. *Comput. Graph. Image Process.*, 22, 327–338. DOI: [https://doi.org/10.1016/0734-189X\(83\)90080-4](https://doi.org/10.1016/0734-189X(83)90080-4).
- W. R. Franklin. 1987. Polygon properties calculated from the vertex neighborhoods. In *Proceedings of 3rd Annual ACM Symposium on Computational Geometry*. ACM, 110–118. DOI: <https://doi.org/10.1145/41958.41969>.
- W. R. Franklin. 1990. Calculating map overlay polygon' areas without explicitly calculating the polygons—Implementation. In *4th International Symposium on Spatial Data Handling*, Zürich. International Geographical Union, 151–160.
- W. R. Franklin. 1992. Map overlay area animation and parallel simulation. In D. H. Douglas (Ed.), *Proceedings, SORSA'92 Symposium and Workshop*. Spatially Oriented Referencing Systems Association, 200–203.
- W. R. Franklin. 2004. Analysis of mass properties of the union of millions of polyhedra. In M. L. Lucian and M. Neamtu (Eds.), *Geometric Modeling and Computing: Seattle 2003*. Nashboro Press, Brentwood, TN, 189–202.
- W. R. Franklin. 2005. Mass properties of the union of millions of identical cubes. In R. Janardan, D. Dutta and M. Smid (Eds.), *Geometric and Algorithmic Aspects of Computer Aided Design and Manufacturing*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 67. American Mathematical Society, 329–345.
- W. R. Franklin and S. V. G. de Magalhães. 2019. Computing intersection areas of overlaid 2D meshes. In *IGS2019 International Geometry Summit Posters' Proceedings*. Solid Modeling Association, Vancouver, Canada.
- W. R. Franklin and M. S. Kankanhalli. 1993. Volumes from overlaying 3-D triangulations in parallel. In D. Abel and B. C. Ooi (Eds.), *Advances in Spatial Databases: Third Intl. Symp., SSD'93*, Vol. 692: Lecture Notes in Computer Science. Springer, Berlin, 477–489. DOI: https://doi.org/10.1007/3-540-56869-7_27.
- W. R. Franklin and V. Sivaswami. 1990. OVERPROP—Calculating areas of map overlay polygons without calculating the overlay. In *Second National Conference on Geographic Information Systems*. Canadian Institute of Surveying and Mapping, Ottawa, 1646–1654.
- W. R. Franklin and P. Y. F. Wu. 1987. A polygon overlay system in prolog. In *Autocarto 8: Proceedings of the Eighth International Symposium on Computer-Assisted Cartography*, Baltimore, Maryland. American Society for Photogrammetry and Remote Sensing, 97–106.
- W. R. Franklin, N. Chandrasekhar, M. Kankanhalli, V. Akman, and P. Y. F. Wu. 1990. Efficient geometric operations for CAD. In M. J. Wozny, J. U. Turner, and K. Preiss (Eds.), *Geometric Modeling for Product Engineering*, North-Holland, 485–498. Selected and Expanded Papers from the IFIP WG 5.2/NSF Working Conference on Geometric Modeling, Rensselaerville, NY, September 18–22, 1988.
- W. R. Franklin, V. Sivaswami, D. Sun, M. Kankanhalli, and C. Narayanaswami. 1994. Calculating the area of overlaid polygons without constructing the overlay. *Cartogr. Geogr. Inf. Syst.* 21, 2, 81–89. DOI: <https://doi.org/10.1559/152304094782610260>.

- W. R. Franklin, S. V. G. de Magalhães, and M. V. A. Andrade. 2017a. 3D-EPUG-Overlay: Intersecting very large 3D triangulations in parallel. In *2017 SIAM Conference on Industrial and Applied Geometry*, Pittsburgh, PA (talk). SIAM.
- W. R. Franklin, S. V. G. de Magalhães, and M. V. A. Andrade. 2017b. An exact and efficient 3D mesh intersection algorithm using only orientation predicates. In *S3PM-2017: International Convention on Shape, Solid, Structure, & Physical Modeling, Shape Modeling International (SMI-2017) Symposium*. Solid Modeling Association, Berkeley, CA, 19–23 (poster).
- W. R. Franklin, S. V. G. de Magalhães, and M. V. A. Andrade. 2018. Exact fast parallel intersection of large 3-D triangular meshes. In *27th International Meshing Roundtable*, Albuquerque, NM. Springer, 365–383. DOI: https://doi.org/10.1007/978-3-030-13992-6_20.
- M. C. Jordan. 1887. *Cours d'analyse de l'École Polytechnic, Tome Troisième, Calcul Intégral, Équations différentielles*. Gauthier-Villars, Paris. Retrieved November 1, 2019, from <https://www.maths.ed.ac.uk/~v1ranick/jordan/jordan.pdf>.
- NVIDIA. 2017. NVIDIA Tesla V100 GPU Architecture, the World's Most Advanced Data Center GPU, wp-08608-001v1.1. Retrieved November 1, 2019, from <https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf>.
- Oak Ridge National Lab. 2019. *Summit—Oak Ridge Leadership Computing Facility*. Retrieved November 1, 2019, from <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>.
- M. Shimrat. 1962. Algorithm 112: Position of point relative to polygon. *Commun. ACM* 5, 8, 434. DOI: <https://doi.org/10.1145/368637.368653>.
- Top500. 2019. Top 500 Supercomputer Sites. <http://www.top500.org>.
- P. Y. F. Wu and W. R. Franklin. 1990. A logic programming approach to cartographic map overlay. *Comput. Intell.* 6, 2, 61–70. DOI: <https://doi.org/10.1111/j.1467-8640.1990.tb00290.x>.