On Implementing Straight Skeletons: Challenges and Experiences

Günther Eder 🆔
Universität Salzburg, FB Computerwissenschaften, Austria
geder@cs.sbg.ac.at

Martin Held 🆔
Universität Salzburg, FB Computerwissenschaften, Austria
held@cs.sbg.ac.at

Peter Palfrader 🆔
Universität Salzburg, FB Computerwissenschaften, Austria
palfrader@cs.sbg.ac.at

―――― **Abstract** ――――

We present CGAL implementations of two algorithms for computing straight skeletons in the plane, based on exact arithmetic. One code, named SURFER2, can handle multiplicatively weighted planar straight-line graphs (PSLGs) while our second code, MONOS, is specifically targeted at monotone polygons. Both codes are available on GitHub. We discuss algorithmic as well as implementational and engineering details of both codes. Furthermore, we present the results of an extensive performance evaluation in which we compared SURFER2 and MONOS to the straight-skeleton package included in CGAL. It is not surprising that our special-purpose code MONOS outperforms CGAL's straight-skeleton implementation. But our tests provide ample evidence that also SURFER2 can be expected to be faster and to consume significantly less memory than the CGAL code. And, of course, SURFER2 is more versatile because it can handle multiplicative weights and general PSLGs as input. Thus, SURFER2 currently is the fastest and most general straight-skeleton code available.

1 Introduction

Straight skeletons were introduced to computational geometry by Aichholzer et al. [2]. Suppose that the edges of a simple polygon P move inwards with unit speed in a self-parallel manner, thus generating mitered offsets inside of P. Then the *(unweighted) straight skeleton* of P is the geometric graph whose edges are given by the traces of the vertices of the shrinking mitered offset curves of P; see Figure 1, left. The process of simulating the shrinking offsets is called *wavefront propagation*. Straight skeletons are known to have applications in diverse fields, with the modeling of roof-like structures being one of the more prominent ones [19, 15, 16] We refer to Huber [17] for a detailed discussion of typical applications.

As a first generalization of straight skeletons, the *multiplicatively weighted straight skeleton* was introduced early on by Aichholzer and Aurenhammer [1] and then by Eppstein and Erickson [13]. In the presence of multiplicative weights, wavefront edges no longer move

at unit speed. Rather, they move at different speeds: Every edge of P is assigned its own constant speed; see Figure 1. Although weighted straight skeletons have been used in applications for years, their characteristics were studied only recently by Biedl et al. [5, 6].

■ **Figure 1** Left: The (unweighted) straight skeleton (in blue) plus a family of wavefronts (dashed) for the green polygon. Right: The weighted straight skeleton for the case that edges marked with ∗ have twice the weight and edges marked with ▽ have half the weight of the unmarked edges.

Held and Palfrader [15] define an *additively weighted straight skeleton* as the geometric graph whose edges are the traces of vertices of the wavefronts over the propagation period, with additive weights being assigned to the edges of P. The presence of additive weights means that the edges of P start to move at different points in time. In [15, 16] they study straight skeletons of planar straight-line graphs (PSLGs) with a combination of both additive and multiplicative weights. (A PSLG is an embedding of a planar graph such that all edges are straight-line segments that do not intersect pairwise except at common end-points.)

The straight-skeleton algorithms with the best worst-case bounds are due to Eppstein and Erickson [13] and Vigneron et al. [10, 23]. These algorithms seem difficult to implement. Indeed, progress on implementations has been rather limited so far. Felkel and Obdržálek [14] describe a simple straight-skeleton algorithm but it turned out to be flawed [24]. The first comprehensive code for computing straight skeletons was implemented by Cacciola [9] and is shipped with CGAL [22]. It is based on the algorithm by Felkel and Obdržálek [14], but was modified significantly to make it work correctly [8]. It handles polygons with holes as input.

The straight-skeleton code BONE by Huber and Held [18] handles PSLGs as input and runs in $O(n \log n)$ time and $O(n)$ space in practice. However, it is not capable of handling weighted skeletons. As a first step to extend BONE to weighted skeletons, Eder and Held [12] generalize motorcycle graphs to weighted motorcycle graphs. Still, while this is a fascinating research area of its own, this does not yet provide an avenue for an actual implementation.

2 Contribution

Palfrader et al. [21] describe a prototype implementation of a straight-skeleton code named SURFER. Since SURFER is the fastest implementation of unweighted straight skeletons known to us, it was the natural candidate for a re-implementation and extension to multiplicative weights. As the handling of degenerate cases had been fairly tricky for unweighted input, we decided to base the new implementation, SURFER2, on exact arithmetic. (However, SURFER2 has an option to run it on conventional floating-point arithmetic.) In order to be able to compare SURFER2 to a special-purpose algorithm for computing straight skeletons of monotone polygons, which is known to have an $\mathcal{O}(n \log n)$ worst-case complexity, we also implemented the algorithm by Biedl et al. [4]. This implementation was named MONOS.

We subjected our implementations to extensive practical tests. Summarizing, SURFER2 is slower than CGAL's straight-skeleton code for small data sets, of about the same speed for polygons that contain 1000 vertices, and 100 times faster for polygons with 10 000 vertices. The memory footprint of SURFER2 is linear while CGAL's code consumes a quadratic amount

of memory. For monotone polygons, our code MONOS achieves a speed-up of about 1.5 compared to SURFER2. Our tests showed clearly that engineering considerations do yield practical speed-ups and that there is a significant price to pay for the use of exact arithmetic.

3 Straight skeleton and wavefront propagation

While every input edge of a polygon emanates a wavefront edge only to one side, namely to the interior of the polygon, PSLG edges emanate wavefront edges on both sides. As in the case of polygons, all wavefront edges are joined-up along angular bisectors of the input edges. In order to handle input vertices of degree one, at each such vertex v one additional wavefront edge is created that moves away from v and its corresponding PSLG edge e in a direction orthogonal to e; see the left-most vertical segment in Figure 2a. This way, the wavefronts form closed polygons even in the presence of degree-one input vertices.

(a) initial wavefront. **(b)** edge event. **(c)** split event.
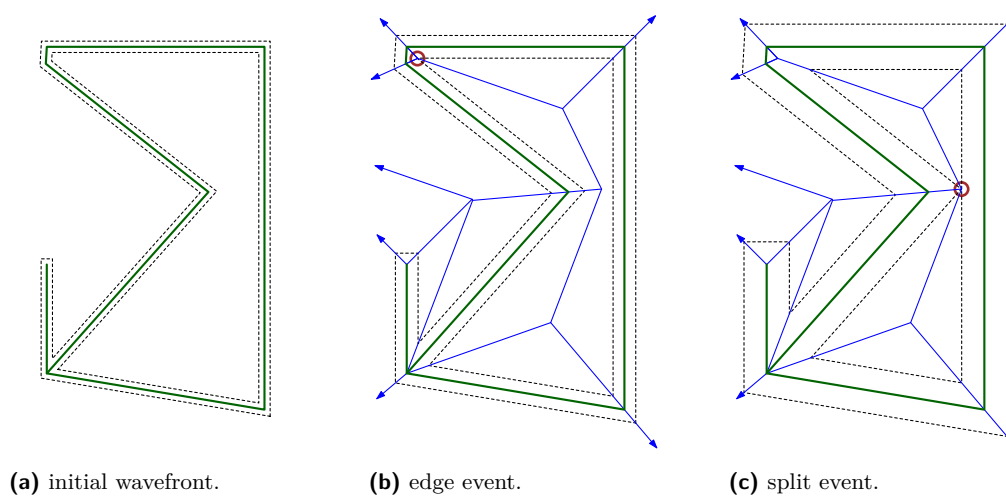
Figure 2 Wavefronts (dashed) at specific times for a PSLG (green), straight skeleton (blue).

Initially, the wavefront coincides with the input edges. If sufficiently little time has elapsed after the start of the wavefront propagation such that no event has occurred yet, then the wavefront forms a mitered offset of the PSLG. The polygons of this mitered offset partition the plane into two (possibly disconnected or multiply-connected) areas: one area that has been swept by the wavefront and another area that has not yet been reached. As time increases and the propagation continues, the area swept by the wavefront gets larger and larger, and wavefront edges will shrink to zero length, thus collapsing in *edge events*; see Figure 2b. Furthermore, vertices may move into the interior of non-incident wavefront edges in *split events*; see Figure 2c. Every split event splits one wavefront edge. Depending on the topology of the not-yet swept region affected, such an event will also split the region into two polygons or reduce the number of holes of the region.

This wavefront-propagation process continues until no more events happen. Then the straight skeleton is the planar graph whose edges are the traces of wavefront vertices during the propagation. Note that some wavefront vertices on the outer region will continue to escape to infinity; these vertices trace out rays that form unbounded arcs of the skeleton.

Simulating the wavefront propagation is a common building block of several straight-skeleton algorithms. These algorithms differ primarily in how they find the next event. All future events already known are usually maintained in a priority queue of event times. It is standard to use a min-heap to store the times of those future events.

Straight skeleton of a monotone polygon

4.1   Monotone algorithm

Biedl et al. [4] describe an $\mathcal{O}(n \log n)$ time algorithm to compute the straight skeleton of a simple n-vertex monotone polygon $\mathcal{P}$. Without loss of generality we assume $\mathcal{P}$ to be x-monotone. Their algorithm consists of two steps: (i) The polygon $\mathcal{P}$ is split into an upper and lower monotone chain, and the straight skeleton of each chain is computed individually. (ii) The final straight skeleton $\mathcal{S}(\mathcal{P})$ is obtained by merging these two straight skeletons.

We employ a classical wavefront propagation to obtain the straight skeleton of a monotone chain in $\mathcal{O}(n \log n)$ time [4]. The monotonicity of the chains guarantees that every change of the wavefront topology is witnessed by an edge event. That is, split events cannot occur.

In the second step, the skeletons of the upper and lower chains are merged to form $\mathcal{S}(\mathcal{P})$. This merge is based on a left-to-right traversal of the two chains and their respective straight-skeleton faces. Starting at the leftmost vertex of both chains, the angular bisector between the incident edges is constructed. It intersects arcs of both the top and bottom straight skeleton. We stop at the first intersection reached and modify the respective straight skeleton locally by creating a node. Then the next bisector is constructed between the two edges that induce the faces of the upper and lower skeleton we are currently in. This process is repeated until the rightmost vertex of $\mathcal{P}$ is reached, thus obtaining the final skeleton $\mathcal{S}(\mathcal{P})$.

4.2   Implementational details

Monos uses Cgal's `Exact_predicates_exact_constructions_kernel_with_sqrt` algebraic kernel, which is backed by Core's `Core::Expr` exact number type.

Intersection computations.   A major part of the merge step are intersection tests and intersection computations between bisectors and skeleton arcs: Careful engineering resulted in substantial performance improvements for these intersection tests. Initially, we applied Cgal's `do_intersect` and `intersection` rather naively to an arc segment (seen as a straight-line segment) and a bisector. However, explicitly deciding whether the end-points of an arc segment lie on different sides of the supporting line of a merge bisector is sufficient to decide whether an intersection occurs. Once we know that an intersection occurs then we apply Cgal's `intersection` routine to the supporting lines of the arc and the bisector. A test on more than one thousand input polygons with at least 10 000 vertices each shows average runtime savings of about 9 % when using the latter method.

Collinear edges.   Input that contains collinear edges needs special care. The crux is that no "direction" need be implied by a wavefront vertex that traces out a bisector. Assume that two collinear wavefront edges become adjacent during the computation of the straight skeleton of a monotone chain. At the time of the event we have a straight-skeleton node p at which the event occurs. Hence we define the bisector of the two collinear edges to be perpendicular to them and to start at p, thereby moving into the same direction as the wavefront edges.

During the merge step, it is less obvious how to proceed in such a case. In Figure 3, left, we see that the left-to-right merge has arrived at the point p. The next bisector, b, is defined by the input edges e_i and e_j. If the merge had progressed from right to left, then b would be horizontal and it would start at the point p'', in analogy to how this situation would be handled if it occurred during the construction of the straight skeleton of a monotone chain. See the horizontal red dotted segment, b, in Figure 3, left: This bisector need not intersect p

but can arrive anywhere along the boundary of the face of e_i or the face of e_j. Hence, we need to look ahead and find the point p' where it intersects one of the two boundaries. Note that this "look ahead" does not break the algorithm's complexity as we only have to walk along the boundaries of the two faces we stepped into. At some point, the boundaries intersect, which is the end-point of b. If b is not incident to p, then we need to add the start node p' of b as well; see Figure 3, right. Afterwards we can proceed with the standard merge.

Figure 3 Left: Constructing the merge chain $\mathcal{M}$ at a specific point p where collinearities have to be handled.   Right: Looking ahead lets us identify p' as start of the next horizontal bisector.

Min-heap.   Initially, MONOS employed C++'s `std::set` as a priority queue due to the requirement to update elements. This structure tends to be implemented as a red-black tree and provides all operations required at appropriate asymptotic costs. However, profiling MONOS showed that a significant amount of time was spent in queuing operations. Therefore, we switched to a self-developed binary min-heap already in use in SURFER2. Performance tests showed an average performance gain of 5%.

5 Weighted straight skeleton of a PSLG

5.1 Triangulation-based algorithm

Aichholzer and Aurenhammer [1] describe an algorithm for computing the straight skeleton of general PSLGs in the plane. It carries over to multiplicatively weighted input in a natural way, provided that all weights are positive. (Biedl et al. [5] show that most of the theory falls apart if negative weights are allowed.) Their approach constructs the straight skeleton by simulating the wavefront propagation. As the wavefront sweeps the plane, they maintain a kinetic triangulation of that part of the plane which has not yet been swept. This triangulation is obtained by triangulating the area inside the convex hull of all wavefronts. Furthermore, all edges of this convex hull are linked with a dummy vertex at infinity.

The area of each triangle of this kinetic triangulation changes over time as its vertices, which are vertices of the wavefront, move along angular (straight-line) bisectors of the input edges. Since each vertex moves at constant velocity, the area of each triangle is a quadratic function in time. As long as no triangle collapses to zero area these triangles serve as certificates for the validity of the kinetic triangulation: The wavefront does not change combinatorially while no triangle collapses. Every change in the topology of the wavefront, i.e., every edge event and every split event, is witnessed by a triangle collapse. Of course, the roots of the quadratic functions give the collapse times of the triangles.

Note that not all triangle collapses witness a corresponding event of the wavefront. Some triangle collapses correspond to internal events only, where the triangulation changes as a wavefront vertex moves over a triangulation diagonal. Aichholzer and Aurenhammer [1] call this type of event a *flip event* because it merely requires the flip of a triangulation diagonal.

The number of edge and split events is linear because the straight skeleton has a linear combinatorial complexity. However, no better bound than $O(n^3)$ is known for the maximum number of flip events for an n-vertex PSLG, with a theoretical worst-case lower bound of $\Omega(n^2)$ [17]. Consequently, the algorithm's worst-case runtime is bounded by $\Omega(n^2 \log n)$ and by $\mathcal{O}(n^3 \log n)$. Our previous work [21] showed that this algorithm runs in $\mathcal{O}(n \log n)$ time in practice when using IEEE 754 arithmetic. In the sequel we present our implementation, SURFER2, of this algorithm based on exact arithmetic, and with support for weighted input.

5.2   Implementational details

SURFER2 uses some of CGAL's geometric primitives, such as Points and Segments. It defaults to using CGAL's `Exact_predicates_exact_constructions_kernel_with_sqrt` algebraic kernel, which is backed by CORE's `Core::Expr` exact number type. SURFER2 uses none of CGAL's advanced data structures, with the exception of CGAL's `DCEL` arrangement [20], which is used for the purpose of presenting the straight skeleton to the user. In addition to a library API, SURFER2, like MONOS, provides both a command line interface and a GUI. GraphML [7] is our input and output format, as it allows weighted input.

Data structures.   To represent the wavefront and the area not yet swept by the wavefront, our implementation explicitly has objects of kinetic triangles, wavefront edges, and wavefront vertices. For each kinetic triangle, we store pointers to the three incident wavefront vertices, and on each side we either have a pointer to the neighboring triangle or a pointer to the incident wavefront edge. Every wavefront edge has references to its two incident wavefront vertices as well as a pointer to the incident kinetic triangle. Wavefront vertices trace out straight skeleton arcs during the propagation period, moving along the bisector of input edges. Every vertex stores pointers to its incident wavefront edges, i.e., the edges that emanate from the input edges defining the bisector along which the vertex moves. Furthermore, it stores the time when it came into existence, and, if it has already been stopped by an event, the time when it was stopped. Additionally, each vertex caches its velocity (as a function of its incident edges), its start location, and also its stop location (if applicable).

We start out by constructing a constrained Delaunay triangulation of the input graph using CGAL's 2D-Triangulation package [25]. From this initially static triangulation we construct the kinetic triangulation just outlined: Neighborhood relations between triangles from the initial triangulation are copied to the kinetic triangulation for triangles that are not split by a constraint. At each constraint, we initialize two wavefront edges, one moving in each direction, that are incident to one kinetic triangle each. Lastly, we initialize the wavefront vertices and compute their velocities based on their incident wavefront edges.

Event types and event queue.   The event queue is a binary min-heap whose elements are kinetic triangles, with an order imposed by the next event each triangle witnesses. An unbounded triangle, which is defined by the vertex at infinity and one edge of the convex hull, serves not only as a witness for its edge collapsing but also for one of its vertices leaving the convex hull. For each bounded triangle, we store the time of its next collapse.

Finding the collapse time of a triangle is straightforward in theory. (Recall that it is a root of a second-degree polynomial given by the signed determinant of its moving vertices.) Checking when a vertex leaves the convex hull can be achieved by considering triangles of

three subsequent convex-hull vertices. In practice we would like to know not only the time but also the type of event. This helps to ensure we make progress with the wavefront propagation. Progress is guaranteed if we process an event that is a split event or an edge event: There is only a linear number of such "real" events because each event effectively causes a collapsed triangle to be removed from the triangulation. The crux comes to light when investigating flip events. A flip event happens when a wavefront vertex becomes incident on a triangulation diagonal. We consider the quadrilateral formed by the two triangles along this diagonal, and flip the diagonal within this quadrilateral. If the other triangle also collapsed at precisely the same time, then after this flip we still have two triangles that collapse right now, and we might do another flip that re-establishes the previous configuration; cf. Figure 4.


Figure 4 Potential flip-event loop when, at time t, the vertices v_1 through v_4 all become collinear.

To handle multiple events that occur at the same time we apply a twofold strategy: First, we always prioritize "real" events over flip events. (This can be done with no additional cost by using a secondary key for the priority queue.) Second, we need to impose an order on flip events. Consider a maximal set $\mathcal{T}$ of neighboring triangles that all collapse at the same time to line segments on the same supporting line, without any triangulation diagonal shrinking to zero length. Note that a flip in a triangle will always cause its longest edge to be flipped. Let T be the triangle that has the longest edge e among the triangles of $\mathcal{T}$. If the other triangle, T', that is incident at e does not belong to $\mathcal{T}$, then it has positive area, and performing the flip will reduce the number of elements in $\mathcal{T}$ and we have made progress. If, however, $T' \in \mathcal{T}$ then performing the flip will cause the longest edge within $\mathcal{T}$ to become strictly smaller. Hence, a monotonicity argument implies that we make progress, too.

Another type of event that shows up in practice if the input need not be in general position is given by non-incident vertices that coincide at some point in time. Topologically, this is equivalent to a split event as previously non-incident wavefront portions become incident. From the viewpoint of a kinetic triangulation this event looks more like an edge event, with the exception that it is not a wavefront edge that collapsed but a triangulation diagonal. Handling this event is thus similar to handling two edge events, with one event per triangle incident at the collapsed diagonal. Figure 5, left, shows such an event.

When parallel wavefront elements move into each other, events happen as triangles collapse. The wavefront edges that comprise these parallel wavefront elements are incident to different kinetic triangles. As we process these collapsed triangles, one after the other, we want to leave the kinetic triangulation in a consistent state between events. Generally, this means creating new kinetic vertices to replace old ones as edges collapse and splits happen, which involves computing the velocity of the new kinetic vertex. The velocity is such that the wavefront vertex remains on the intersection of the wavefront edges as they propagate, i.e., it points along the angular bisector of the incident edges and has the appropriate speed. However, what is the velocity of a wavefront vertex between opposing wavefront edges that
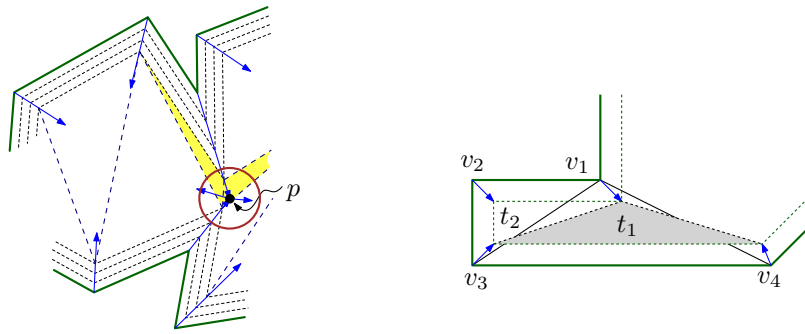
■ **Figure 5** Left: The two shaded triangles are witnesses for an event at p, where two reflex vertices move into each other.   Right: Infinitely fast vertices get created when the triangles collapse. If t_1 is processed first, an infinitely fast vertex is created at the locus of v_1 moving to the left, and one at the locus of v_2 and v_3 moving to the right, both incident to t_2. If t_2 is processed first, only one infinitely fast vertex is created at the locus of v_2 and v_3 moving to the right.

overlap and span an angle of zero? The limit, as the angle approaches zero, is infinite speed with a direction along the wavefront edges. Therefore, in our implementation we have the concept of infinitely fast vertices. These get created when opposing wavefront edges become incident and share a vertex; cf. Figure 5, right. Whenever we create an infinitely fast vertex, the triangle that is incident to this vertex is processed next as we consider it to have an event immediately. If more than one triangle is incident to an infinitely fast wavefront vertex v, then we pick the one that has the shorter constraint incident to v.

During the wavefront propagation, instances may occur where three or more wavefront vertices become collinear and remain so in the course of the propagation. If these vertices are adjacent on the convex hull, then the determinant check for whether a vertex leaves the convex hull would show a zero at all times. When encountering an always-zero determinant, depending on the specific configuration, scheduling an event for "right now" may resolve the issue, e.g., by flipping a triangulation diagonal such that the three vertices involved no longer are incident to the same kinetic triangles. If three vertices on the convex hull are collinear, then such a flip would not be advisable because removing the middle vertex from the convex hull (and keeping it on the books as a non-convex-hull vertex) would cause a bounded triangle with an always-zero determinant.

Collapse time and order of events.   While the actual collapse time is one of the roots of a quadratic polynomial, solving quadratics is not always necessary. Avoiding root finding for quadratic polynomials will increase accuracy when working with limited-precision data types, and it will result in less complex expression trees when working with exact numbers as provided by CORE's `CORE::Expr`. In particular, it will avoid the computation of another square root. Hence, we try to employ geometric knowledge derived from local combinatoric properties as much as this is possible. For instance, a triangle with two incident wavefront edges will never see a flip event: If it sees an event at all, it will be the collapse of either one or both of its wavefront edges. (In the latter case the entire triangle collapses.)

Triangles with exactly one incident wavefront edge can encounter all three types of events – edge, split, and flip events – but we can determine each such event without computing the roots of a determinant: The times of split and flip events can be found by considering the distance between the vertex opposite the wavefront edge to the supporting line of the

wavefront edge. This distance is linear in time and when it passes through zero, we either have a flip or split event as the vertex comes to lie on the supporting line of the wavefront edge. Likewise, an unbounded triangle witnessing a wavefront edge collapse needs to consider only the linear function that models the wavefront edge's length. Unbounded triangles will not need to witness collapses of a triangulation diagonal because this event will be witnessed by the neighboring (bounded) triangle. Finding the time when a vertex leaves the convex hull may require determining the roots of a quadratic polynomial if no incident wavefront edge is on the convex hull. Otherwise, we employ the distance-to-supporting-line method again. The only bounded triangles for which we need to consider quadratic polynomials are triangles not incident to a wavefront edge: These triangles can cause flip events.

We process the events in order of their event times, breaking ties based on the event type: Events with infinitely fast vertices come first, then the other real events, and finally we process flip events and convex-hull-update events. Ties among flip events are broken by handling the event with the longest edge first. Since our implementation uses exact arithmetic provided by CORE, all these comparisons can be done exactly. Event handling generally is straightforward and consists of updating the kinetic triangulation to remain valid after the event, for real events dropping collapsed triangles, and stopping kinetic vertices and creating new ones. (Recall that we know the type of event.)

Once no event is left in the queue with a finite event time, the list of kinetic vertices with their start and stop times and positions yield the arcs of the final straight skeleton. As post processing, we construct a CGAL DCEL (doubly-connected edge list) structure from the kinetic vertices. Each kinetic vertex and each input segment corresponds to two DCEL half-edges, and each DCEL face is incident to one input segment or terminal input vertex.

Pre-emptive flips.   Flip events are internal events of the kinetic triangulation, and as such they are not directly needed for the wavefront propagation process. They are always the result of a reflex vertex moving over the opposite diagonal in a triangle. As already outlined, flip events can lead to loops and thus require special handling. We tested whether we can reduce the number of flip events by restructuring triangles as soon as they are created, either initially or as the result of another event, rather than wait for a triangle to collapse before we do a flip. The very simple approach that we chose was to immediately flip a triangulation diagonal that is incident to two convex vertices if the flip would make it incident to a reflex vertex, in the assumption that this reflex vertex is then less likely to cause a flip event later.

We benchmarked this approach and found that it results in fewer flip events for about 98 % of our test runs. In the worst example we got 4 % more flip events, and in the best case flip events were reduced by slightly more than 10 %, with a median change of 3 % fewer flips. The benefit is less clear-cut when looking at runtime improvements, though. We noticed an improvement of about 1 % in the median and a surprising 50 % in the best case, but also a 30 % slowdown in the worst case. Still, pre-emptive flips are the default set-up for SURFER2.

Component-based priority queues.   Closed input loops partition the plane into connected components. (And in practice many PSLGs contain such loops.) The basic algorithm deals with events, witnessed by triangle collapses, without any consideration of the component the triangle is in. But the straight skeleton within one component is entirely independent of the straight skeleton within any other component: The different wavefronts never interact during their respective propagation processes. Hence, it would suffice to ensure that events are ordered correctly within each component but we could handle events from different components independently. Of course, in the perfect world of the Real RAM model it does not matter whether k items are to be stored in one or a few comparison-based priority queues, since handling k items takes $\Theta(k \log k)$ time one way or the other.

Once we talk about an implementation the Real RAM model is no longer applicable, though. Our implementation uses Core's `CORE::Expr` number type as shipped with Cgal. This gives us the luxury of actually having exact arithmetic, thus making it easy to know which events happen simultaneously. The significant price to be paid is that comparisons are no longer unit-cost: Each Core expression variable carries with it its expression tree as the entire history of how it was derived, as well as a numerical approximation with error bounds. Whenever we need to compare two numbers, Core can calculate more digits should this be needed to ascertain the exact relationship. This takes extra time and memory. In general, such comparisons are near-instantaneous when the values are (sufficiently) different. However, they can take a very long time if the numbers are actually equal. (Profiling showed that one such comparison may even take many seconds.)

To avoid comparing event times from different components we assign an integer identification number to each component and use this number as a secondary key for the event ordering in the priority queue. We added instrumentation to our implementation to count how often we actually compare equal event times during the wavefront propagation process. For random input in general position, the likelihood of events in different parts of the plane to happen at the same time is zero. Our tests confirmed that for random input there is nothing to gain by per-component event ordering. However, avoiding comparing event times from different components may yield substantial savings for some input classes: A reduction of the number of simultaneous events by a quarter may result in a reduction of the runtime by a factor of five in extreme cases! Savings can be expected to occur for inputs with vertices on integer grids, or, e.g., if two components both contain tight polygonal approximations of circular arcs with the same radius.

6 Results

Setup. All runtime tests were carried out on a 2015 Intel Core i7-6700 CPU. For most of our tests, memory consumption was limited to 10 GiB. The codes were compiled with `clang++`, version 7.0.1, against Cgal 5.0 except where stated otherwise.

The default setting for Cgal's straight-skeleton package [9] is to use the exact predicates but inexact constructions kernel that ships with Cgal. The use of this kernel ensures that the straight skeleton computed is combinatorially correct, even if the locations of the nodes need not be correct. Cgal's straight skeleton package can also be run with the exact predicates and exact constructions kernel. However this causes the runtimes to increase by a factor of roughly 100. Our codes, Monos and Surfer2, use the Cgal exact predicates and exact constructions with square-root kernel by default, as we construct wavefront vertices with velocities, and these computations involve square roots.

We tested both Cgal and Surfer2 on many different classes of polygons. Our test data consists of real-world (multiply-connected) polygons as well as of synthetic data generated by Rpg [3] and similar tools provided by the Salzburg Database [11]. For a polygon that has holes we used Cgal's straight-skeleton code that supports holes. Otherwise we used the implementation which only supports simple polygons.

Additionally, we tested both of our codes, Monos and Surfer2, with large monotone polygons, for up to 10^6 vertices. (We did not run Cgal on these inputs due to memory constraints.) Given the lack of a sufficiently large number of monotone real-world inputs, we used Rpg [3] to automatically generate thousands of monotone input polygons. We also ran Surfer2 on hundreds of real-world PSLGs. Most of our data came from GIS sources and

represents road and river networks, contour lines and the like. The runtimes on those inputs are quite comparable to the runtimes for polygons. That is, the test results presented here are also representative for SURFER2's performance on real-world data.
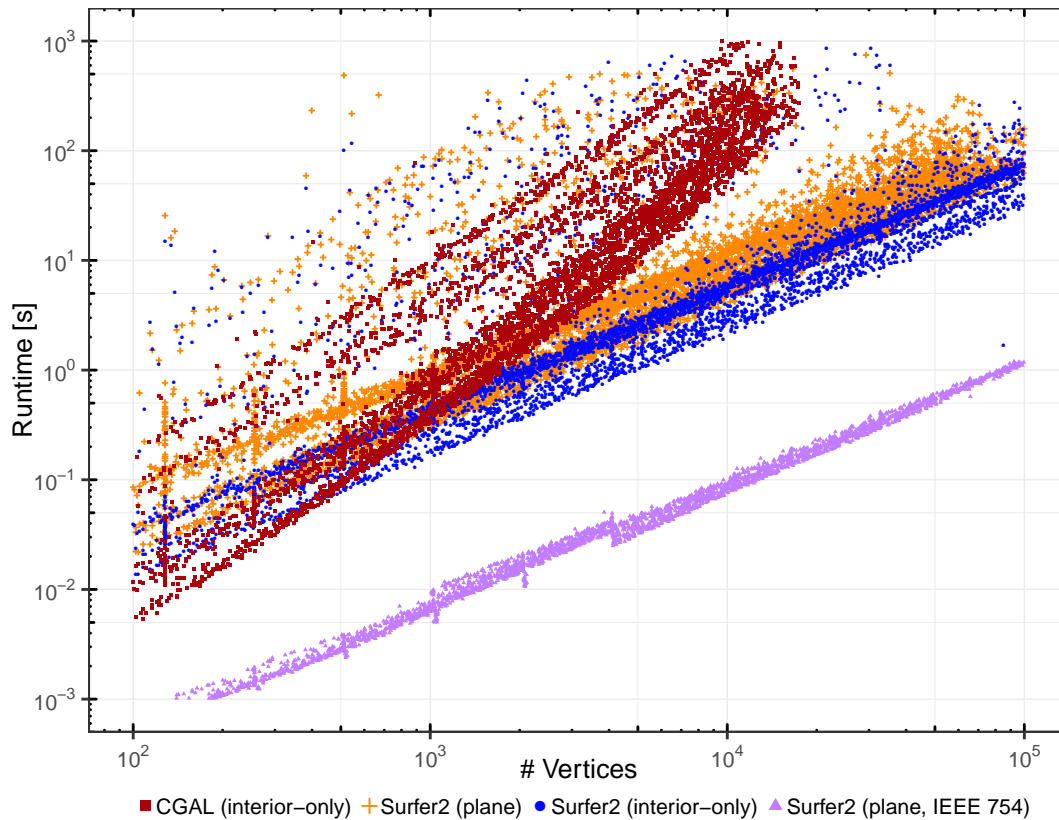
Figure 6 Runtimes of CGAL's code and different variants of SURFER2.

Runtimes and memory consumption.   While CGAL's code computes the straight skeleton either in the interior or the exterior of a polygon, SURFER2 can do both in one run because it treats a polygon as a PSLG. But SURFER2 can also be restricted to just the interior or the exterior of a polygon. Hence, for the plot in Figure 6 we ran SURFER2 twice, once applied to the entire plane and once for just the interiors of the polygons that were also handled by CGAL. Our tests make it evident that SURFER2 is significantly faster than CGAL for a large fraction of the inputs. In particular, its runtime seems to exhibit an $n \log n$ growth, compared to the clearly quadratic increase of the runtime of CGAL's code. Figure 7 shows that SURFER2's memory consumption grows (mostly) linearly while CGAL's code requires a clearly quadratic amount of memory.

The results for CGAL's straight skeleton package were to be expected because (at least back in 2010) it computed potential split events for each pair of reflex vertex and wavefront edge [17, Section 2.5.4]. Indeed, tests carried out in 2010 indicated that it requires $\mathcal{O}(n^2 \log n)$ time and $\Theta(n^2)$ space for n-vertex polygons, as discussed in [18]. Our test results suggest that the same algorithm and implementation are applied in the current CGAL 5.0, which is the version used in our tests. The theoretical upper bound on the runtime complexity of SURFER2 is given by $O(n^3 \log n)$. Thus, its very decent practical performance is noteworthy.

Figure 7 Memory use of CGAL, SURFER2 variants, and MONOS.

We also ran SURFER2 with IEEE 754 double as a number type instead of CORE's `CORE::Expr`. We admit, though, that SURFER2 is not (yet) fully prepared to deal with finite-precision arithmetic as it assumes that the order of the events is given reliably. With sufficient engineering effort, however, such requirements can be relaxed and the code could be made to run reliably even with finite-precision arithmetic. Hence, the current version of SURFER2 with the IEEE 754 backend works well on input in general position, but fails for some inputs when several events happen at exactly the same time and position. Nevertheless, we wanted to see the runtime and memory characteristics in order to get a better impression of the practical costs of using `CORE::Expr`: Therefore the plots in Figures 6 and 7 also show results for SURFER2 with IEEE 754 arithmetic.

The $\mathcal{O}(n \log n)$ bound for the complexity of MONOS is apparent in Figure 8, left. Given that MONOS is a special-purpose code designed specifically for handling monotone polygons, it had to be expected that it outperforms SURFER2 consistently for such input.

Dependence on input characteristics. There are outliers both in the runtime as well as in the memory consumption of SURFER2 which are clearly visible in Figures 6 and 7. (To a lesser extent this noise is visible for CGAL, too.) Given the fact that such outliers do not show up for the IEEE 754-based version of SURFER2, there is reason to assume that this behavior is not intrinsic to SURFER2's algorithm but that it has its roots in the use of exact arithmetic. To probe this issue further we investigated which input classes trigger these outliers. Figure 9 shows the runtimes of both codes for three different inputs classes.
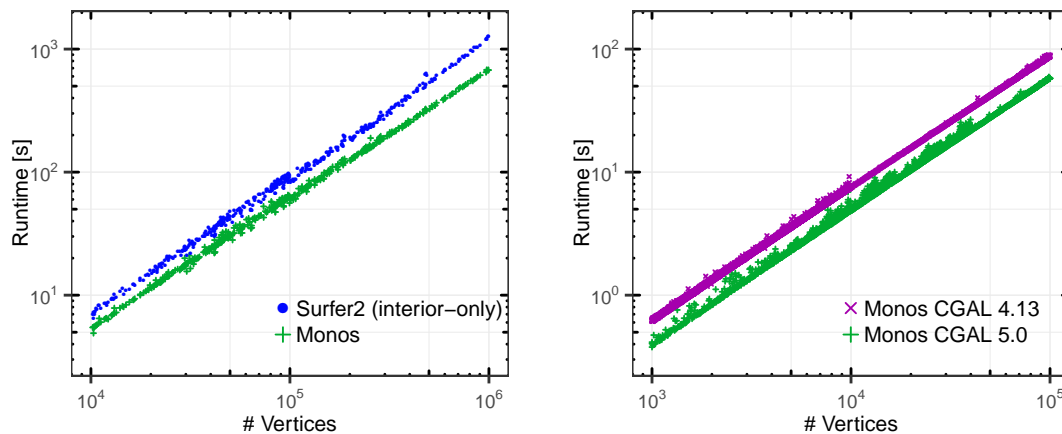
Figure 8 Monos vs. Surfer2 on monotone inputs, and Monos with Cgal 4 vs. Cgal 5.

The class of input that was most time-consuming to handle for all implementations were our random octagonal polygons. All polygons out of this group have interior angles that are multiples of 45°. We were surprised to see that axis-parallel input is not troublesome per se. The key difference between both groups of polygons is given by the fact that all octagonal polygons have their vertices on an integer grid while the vertices of the orthogonal polygons are (random) real numbers. Hence, for the octagonal polygons many events tend to happen at exactly the same time when opposite edges become incident in different parts of the polygon. It is apparent that the proper time-wise ordering of these events incurs a significant cost if `CORE::Expr` is used. The presence of parallel edges does not automatically increase the runtime of Surfer2, though. The likely explanation is that many events cause infinitely-fast vertices, which are then, in turn, easy to sort as they are handled immediately. For comparison purposes we ran the code on random simple polygons as a third class of input, with random vertex coordinates. Having any kind of co-temporal event is highly unlikely for those inputs, as are infinitely-fast vertices.

The excessive amounts of time consumed by ordering simultaneous events became even more apparent when we studied the impact of multiplicative weights on Surfer2's runtime. As expected, a difference in the timings for weighted and unweighted random polygons was hardly noticeable. For our randomly weighted octagonal polygons we expected reduced run-times and fewer outliers even when using exact arithmetic, due to very few truly simultaneous events. And, indeed, this was confirmed impressively by our tests; see Figure 9.

Experiences with CGAL. While running our tests, we also compared Cgal versions 4.13 and 5.0, as the latter was released only recently. We witnessed an improvement in the performance of our codes for the newer version; see Figure 8, right.

We also spent a considerable amount of time on debugging Surfer2, just to find out in the end that some of our inputs trigger a bug in the `CORE::Expr` number type. This bug occurs for both Cgal 4.13 and Cgal 5.0. In specific reproducible cases the floating-point filter fails. As a result predicates are evaluated incorrectly, such as Cgal/Core claiming that $0.49770 < 0.01047$. We have reported this bug to the Cgal project (Bug#4296).
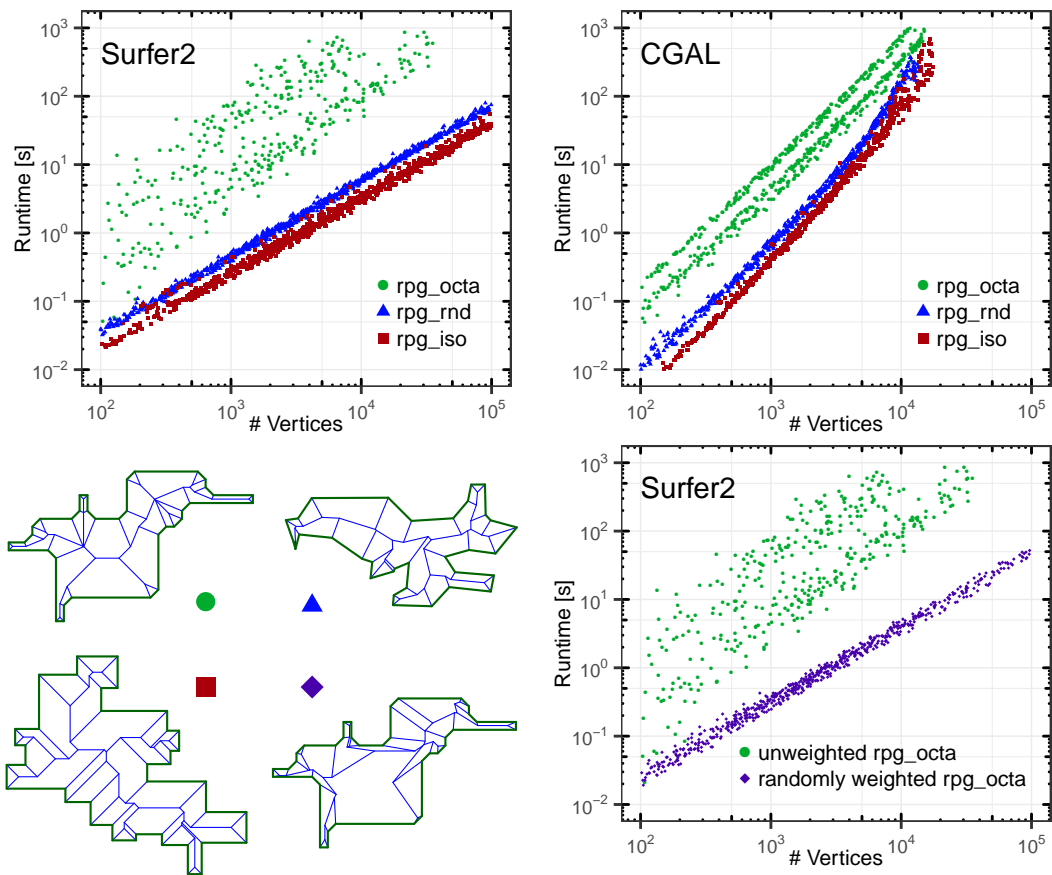
Figure 9 Top: The effect of different input classes on the runtime of Surfer2 and on Cgal. Bottom-left: Samples for different input classes (left to right, top to bottom): Octagonal input (on integer grid), random polygon, orthogonal polygon not on integer grid, and weighted octagonal input. Bottom-right: Surfer2 runtimes for unweighted and randomly weighted octagonal input.

7   Discussion and conclusion

Our implementations Monos and Surfer2 show that an $\mathcal{O}(n \log n)$ runtime can be achieved for the computation of straight skeletons of n-vertex monotone polygons and PSLGs, even if the edges of the PSLG are weighted by positive multiplicative weights. Engineering considerations do not improve the asymptotic behavior but help to improve the practical efficiency. Hence, our implementational efforts can be regarded as successful.

But our tests also make it evident that the use of the `CORE::Expr` number type forces one to abandon the concept of unit-cost comparisons: Multiple events that occur simultaneously have a significant impact on the practical runtime of Surfer2 if the `CORE::Expr` number type is used, while no impact is visible for the standard IEEE 754 arithmetic. It is obvious that it would pay off to detect groups of simultaneous events in a way that does not involve comparing the event times, as it is done by our current implementation. As discussed, a per-component ordering of the events in different priority queues reduces the burden that simultaneous events put on the efficiency. In our current version, we apply per-component computations only for different components induced by the input PSLG. If we could cheaply detect the splitting of a wavefront into two components and afterwards use new component

numbers in future updates of the priority queue, then we might be able to exploit this idea even further. How to carry out a dynamic maintenance of the component information with little computational overhead is on our agenda for future R&D work.

A second avenue for practical improvements is given by a fine tuning of the priority queue used for storing future events. It does not come as a surprise that not all real events computed and registered in the priority queue do indeed correspond to topological changes of the actual wavefront at the times when they are predicted to occur. As a consequence, we spend time on ordering future events even if only a fraction of them will actually happen. Hence, our tests suggest that it might be better to use some form of lazy priority queue. (Standard lazy binomial heaps or Fibonacci heaps still would result in $\mathcal{O}(\log n)$ many comparisons for a delete-min operation.) After all, we do not really require a proper min-heap that obeys the heap property on every layer. Rather, all we need to know is the next event.

An alternative to constructing collapse times and then comparing them is developing predicates for ordering triangle collapses. However, this might be more involved than it appears at first glance: A triangle consists of kinetic vertices that move along the (weighted) bisectors of pairs of input edges, which, in general, uniquely define the position of their wavefront vertex as a function of time. However, in the case of parallel, collinear wavefront edges stemming from one or more parallel, collinear input edges, their bisector is not a uniquely defined one-dimensional line. Rather, the exact trajectory of the kinetic vertex depends on the wavefront propagation that has happened so far; see Figure 10. Nevertheless, having and using exact predicates and resorting to explicit exact computations and comparisons of the collapse times only in degenerate cases might speed up computing the straight skeleton.

Figure 10 The (supporting line of) straight skeleton arc a (dashed, purple) depends not only on e_1, e_2 but also on the weights (e.g., ●), positions, and orientations of other, non-incident input edges.

8 Source code

Our source codes are provided on GitHub and can be used freely under the GPL(v3) license: `https://github.com/cgalab/monos` and `https://github.com/cgalab/surfer2`.

References

1. Oswin Aichholzer and Franz Aurenhammer. Straight Skeletons for General Polygonal Figures in the Plane. In *Voronoi's Impact on Modern Sciences II*, volume 21, pages 7–21. Institute of Mathematics of the National Academy of Sciences of Ukraine, 1998. `doi:10.1007/3-540-61332-3_144`.
2. Oswin Aichholzer, Franz Aurenhammer, David Alberts, and Bernd Gärtner. A Novel Type of Skeleton for Polygons. *Journal of Universal Computer Science*, 1(12):752–761, 1995. `doi:10.1007/978-3-642-80350-5_65`.

3   Thomas Auer and Martin Held. Heuristics for the Generation of Random Polygons. In *Proceedings of the 8th Canadian Conference on Computational Geometry (CCCG)*, pages 38–44, 1996.

4   Therese Biedl, Martin Held, Stefan Huber, Dominik Kaaser, and Peter Palfrader. A Simple Algorithm for Computing Positively Weighted Straight Skeletons of Monotone Polygons. *Information Processing Letters*, 115(2):243–247, 2015. `doi:10.1016/j.ipl.2014.09.021`.

5   Therese Biedl, Martin Held, Stefan Huber, Dominik Kaaser, and Peter Palfrader. Weighted Straight Skeletons in the Plane. *Computational Geometry: Theory and Applications*, 48(2):120–133, 2015. `doi:10.1016/j.comgeo.2014.08.006`.

6   Therese Biedl, Stefan Huber, and Peter Palfrader. Planar Matchings for Weighted Straight Skeletons. In *Proceedings of the 25th International Symposium on Algorithms and Computation (ISAAC)*, pages 117–127, 2014. `doi:10.1007/978-3-319-13075-0_10`.

7   Ulrik Brandes, Markus Eiglsperger, Ivan Herman, Michael Himsolt, and M. Scott Marshall. GraphML Progress Report Structural Layer Proposal. In *Proceedings of the 9th International Symposium on Graph Drawing*, pages 501–512, 2001. URL: `http://graphml.graphdrawing.org/`.

8   Fernando Cacciola. Private email correspondence, 2010.

9   Fernando Cacciola. 2D Straight Skeleton and Polygon Offsetting. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0 edition, 2019. URL: `https://doc.cgal.org/5.0/Manual/packages.html#PkgStraightSkeleton2`.

10   Siu-Wing Cheng, Liam Mencel, and Antoine Vigneron. A Faster Algorithm for Computing Straight Skeletons. *ACM Transactions on Algorithms*, 12(3):44:1–44:21, 2016. `doi:10.1145/2898961`.

11   Günther Eder, Martin Held, SteinÞór Jasonarson, Philipp Mayer, and Peter Palfrader. On Generating Polygons: Introducing the Salzburg Database. In *Proceedings of the 36th European Workshop on Computational Geometry*, pages 75:1–7, 2020.

12   Günther Eder and Martin Held. Computing Positively Weighted Straight Skeletons of Simple Polygons based on Bisector Arrangement. *Information Processing Letters*, 132:28–32, 2018. `doi:10.1016/j.ipl.2017.12.001`.

13   David Eppstein and Jeff Erickson. Raising Roofs, Crashing Cycles, and Playing Pool: Applications of a Data Structure for Finding Pairwise Interactions. *Discrete & Computational Geometry*, 22(4):569–592, 1999. `doi:10.1145/276884.276891`.

14   Petr Felkel and Š. Obdržálek. Straight Skeleton Implementation. In *Proceedings of the 14th Spring Conference on Computer Graphics (SCCG)*, pages 210–218, 1998.

15   Martin Held and Peter Palfrader. Straight Skeletons with Additive and Multiplicative Weights and Their Application to the Algorithmic Generation of Roofs and Terrains. *Computer-Aided Design*, 92(1):33–41, 2017. `doi:10.1016/j.cad.2017.07.003`.

16   Martin Held and Peter Palfrader. Skeletal Structures for Modeling Generalized Chamfers and Fillets in the Presence of Complex Miters. *Computer-Aided Design and Applications*, 16(4):620–627, 2019. `doi:10.14733/cadaps.2019.620-627`.

17   Stefan Huber. *Computing Straight Skeletons and Motorcycle Graphs: Theory and Practice*. Shaker Verlag, 2012. ISBN 978-3-8440-0938-5.

18   Stefan Huber and Martin Held. Theoretical and Practical Results on Straight Skeletons of Planar Straight-line Graphs. In *SoCG '11: Proceedings of the twenty-seventh Annual Symposium on Computational Geometry*, 2011. `doi:10.1145/1998196.1998223`.

19   Tom Kelly and Peter Wonka. Interactive Architectural Modeling with Procedural Extrusions. *ACM Transactions on Graphics*, 30(2):14:1–14:15, April 2011. `doi:10.1145/1944846.1944854`.

20   Lutz Kettner. Halfedge Data Structures. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0 edition, 2019. URL: `https://doc.cgal.org/5.0/Manual/packages.html#PkgHalfedgeDS`.

21 Peter Palfrader, Martin Held, and Stefan Huber. On Computing Straight Skeletons by Means of Kinetic Triangulations. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, pages 766–777, 2012. `doi:10.1007/978-3-642-33090-2_66`.

22 The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0 edition, 2019. URL: `https://doc.cgal.org/5.0/Manual/packages.html`.

23 Antoine Vigneron and Lie Yan. A Faster Algorithm for Computing Motorcycle Graphs. *Discrete & Computational Geometry*, 52(3):492–514, 2014. `doi:10.1007/00454-014-9625-2`.

24 Evgeny Yakersberg. *Morphing Between Geometric Shapes Using Straight-Skeleton-Based Interpolation*. MSc thesis, CS Dept., Technion, Haifa, Israel, 2004.

25 Mariette Yvinec. 2D Triangulation. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0 edition, 2019. URL: `https://doc.cgal.org/5.0/Manual/packages.html#PkgTriangulation2`.