

GPU-based Approaches for Shape Diameter Function Computation and Its Applications Focused on Skeleton Extraction

Abstract

In this paper two approaches for the computation of the shape diameter function (SDF) on the GPU are outlined and compared. The SDF is a scalar function describing the local thickness of an object. It can be used for consistent mesh partitioning and skeletonization. In the first approach, we have reorganized the tracing of the rays to be well suited for the rasterization hardware. To the best of our knowledge, this is the first method to show how to compute the SDF using only the rasterization hardware and without the need of any acceleration data structures. The second approach uses parallel ray casting and an octree traversal using OpenCL. We demonstrate that the first method achieves similar results as the ray casting using OpenCL. In addition, it is faster for large meshes and it is simpler to implement. Furthermore, we extend the SDF computation by fast post-processing using texture-space diffusion. The fast SDF computation can be used in many applications such as the automatic skeleton extraction as we demonstrate in the article.

Keywords:

shape diameter function, OpenCL, depth peeling, skeleton extraction, skeleton texture mapping

1. Introduction

The shape diameter function (SDF) is a scalar function defined on the mesh surface which expresses a measure of the diameter of the object's volume in the neighborhood of each point on the surface. See Figure 1 for some examples. The original idea was proposed by Shapira et al. [1] and it is related to the concept of Medial Axis Transform (MAT) [2].




Figure 1: The shape diameter function (SDF) expresses a measure of the local thickness around each point on the surface. The SDF values are colored in a scale from low (red) to high (blue).

algorithms to accelerate the calculation of the SDF. In the first approach, instead of tracing multiple rays per vertex at once, we launch one ray for each vertex multiple times. In fact, our method is based on iteratively peeling two or three successive depth layers of the mesh from multiple views around the mesh. From the review of the current literature, to the best of our knowledge, we can state that our depth peeling approach is the first one that computes the SDF exploiting only the rasterization hardware and without the need of any acceleration data structures. In the second approach, we propose a parallel method for computing the SDF that is performed on the GPU using OpenCL, exploiting the independence of rays. Our approach starts by transferring the data containing indices of triangles and our acceleration structure to the GPU. We then calculate a diameter value for each ray independently. Finally, the computation of the final SDF value is performed for each vertex gathering its computed rays diameters.

Given the above tools for SDF computation, we illustrate a few practical applications of the fast SDF computation. The SDF values have to be computed repeatedly during an iterative process, thus we have to compute the SDF as fast as possible. Proposed applications are the computation of skeleton sheets for mesh parameterization and an improvement of mesh contraction weights to extract finer skeletons with small scale details. Additionally, we illustrate how the SDF can be used for the fast extraction of skeletons with linear topology.

2. Related Work

Shape diameter function The shape diameter function (SDF) was introduced by Shapira et al. [1]. Given a point on the mesh surface, the SDF is defined as the weighted average of

The SDF computation is based on creating a cone centered around the inward-normal direction of every point and sending several rays inside this cone to the other side of the mesh (see Figure 3), measuring the distance at the point of intersection. This process, despite being highly parallelizable, is computationally expensive. If processed on the CPU, the task becomes extremely inefficient. Therefore, this paper presents two

the lengths of rays traced from the given point inside the cone centered at its inward-normal. For each ray, the normal at the point of intersection is checked against the normal at the starting point. If the angle between the two normals is less than 90 degrees the ray is rejected. From a practical point of view, this test avoids taking into account rays pointing directly to the outside of the mesh. To make this method robust and also immune to non-watertight meshes, outliers removal is performed preserving only those rays whose lengths fall within one standard deviation from the median of all lengths. In [3], it is pointed out that the outliers removal technique proposed in [1] generates counter-intuitive results in some cases. For example, in the case of a mesh composed of two parallel (infinite) planes, the SDF value obtained by the Shapira et al. method is given by the average of one correct value that is the length of the ray cast along the normal and the lengths of all the other rays thrown inside the cone that systematically overestimates the correct diameter. The same problem occurs at the bifurcation of a Y-shape. In this case, the correct value would be the minimum of all rays lengths, thus increasing the opening of the cone has the only effect of adding noise to the diameter estimation when taking an average ray length. In [3], in order to resolve the dilemma between a small or large cone, a more conservative estimation of the SDF is introduced by using an adaptive cone size. In particular, the algorithm starts with a large cone aperture that gets decreased as long as the measure of the absolute growth of the minimum ray length remains within a certain threshold. For example, in the case of the infinite parallel planes this method converges to a small cone size giving a correct SDF value equivalent to the distance between the two planes. Another application of SDF was introduced in [4], where SDF is used for mesh segmentation based on strokes drawn by the user. Essentially, the main operation in the SDF calculation is given by tracing rays. In [1], such operation is performed on the CPU using an octree as an acceleration structure. As Table 1 shows, ray tracing on the CPU is costly and the need for an acceleration structure is mandatory. A fast approximation of the SDF is proposed in [5] where the SDF value is computed only for some Poisson-distributed points over the mesh surface and then the computed values are propagated across the whole mesh using the Poisson equation. The latter paper reports only the timing for a single low poly model of a horse where a performance in the order of seconds is achieved but at the cost of some approximation. Instead, our GPU method based on depth peeling computes the SDF values for each vertex without approximations and without the need of acceleration structures. Moreover, we achieve performance in the order of few seconds even for meshes with hundreds of thousands of vertices.

Depth peeling The depth peeling technique, introduced in [6], starts by rendering the scene getting the nearest fragment to the camera. In the subsequent passes, the previous depth buffer is bound to a texture in a fragment shader. The scene is rendered again and all fragments whose depth is less or equal to the corresponding depth from the previous pass are discarded. To avoid read-modify-write hazards, the technique ping-pongs the depth values between two different buffers. Depth peeling is a robust image-based solution to different kinds of problems,

specifically: order-independent transparency, layered depth image and ray tracing. The most severe drawback of this technique is that it requires the geometry to be drawn multiple times causing a strong limitation for GPU-bound applications. A more efficient implementation of the depth peeling algorithm has been proposed in [7].

Parallel ray casting There are several ways to perform ray casting. Carr et al. [8] proposed to generate rays and perform traversal of acceleration structures on the CPU, then store the results and perform the ray-triangle intersections on the GPU. Purcell et al. [9] proposed to store the scene geometry and the acceleration structure on the GPU and to use the same device to perform both traversal and intersection. Recent efforts to optimize the algorithms for GPUs have demonstrated to obtain better results on traversal of acceleration structures than a single-core execution on CPU. Therefore, the second option is more suitable for the SDF algorithm.

3. GPU-based Computation of the Shape Diameter Function

We propose two different methods for computing the SDF using the GPU. The first approach exploits the OpenGL programmable pipeline via shaders. The second approach relies on OpenCL to parallelize ray casting for each vertex. We provide a comparison of two methods in terms of their computed values and performances.

3.1. Parallelized Ray Casting using Depth Peeling

The SDF calculation using depth peeling exploits the inherently parallel task of tracing independent rays in a different way than standard ray tracing. In particular, we start by generating a set of cameras by uniformly sampling a sphere around the 3D model and placing an orthogonal camera at each sampled position. The camera view direction is set to look down along the sphere normal at each sampled point. We call the set of cameras \mathbf{C} . As a preprocessing step, we also pack all vertices in a single texture that we call \mathbf{T} . Now, for each camera, we alternate drawing front or back faces performing a depth peeling until no more fragments are written or after a maximum amount of iterations is reached. We keep track of the result of the last 3 render operations using a circular array of 3 framebuffers: \mathbf{F}_0 , \mathbf{F}_1 and \mathbf{F}_2 . At every odd iteration, we bind the last three generated depth textures (two at the first pass) and we perform the computation of one diameter value along the current view direction. We accumulate the diameters computed from all directions in a texture array. From this array, we calculate the final SDF value by first removing outliers and then by performing a weighted sum of the inliers. As in the original algorithm, the weight of each ray is given by the inverse of the angle between the ray and the center of the cone around the inward-normal of each vertex.

To simplify the comprehension of the algorithm, we refer to Figure 2 during our explanation. An outline of this method is reported in Algorithm 1 and Algorithm 2. Given a camera $c \in \mathbf{C}$, we bind \mathbf{F}_0 and we render the front-most part of the model as depicted in Figure 2(a). We write to both the color and the

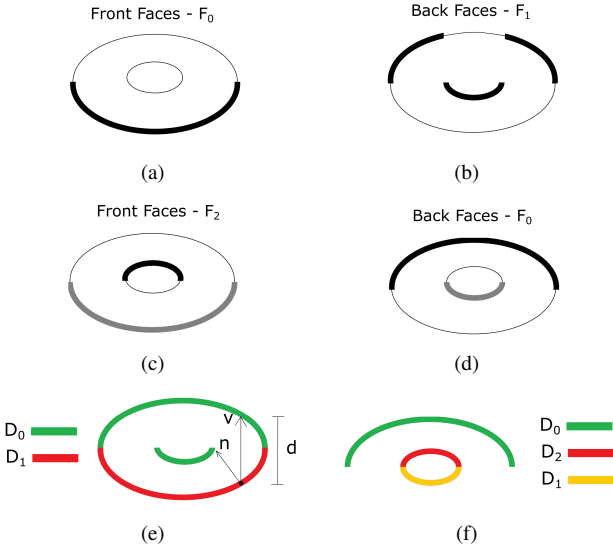


Figure 2: Figures (a), (b), (c) and (d) show the first four iterations of our depth peeling ray tracing applied to a torus. The images show the front-most surfaces as bold black lines, hidden surfaces as thin black lines, and peeled away surfaces as light grey lines. Figures (e) and (f) illustrate the diameter calculation launched at each odd iteration.

155 depth buffer, using a 32 bit floating point texture for the RGB
 156 color channel and a 32 bit depth buffer. In the color channel we
 157 store the normals of the vertices. At the next iteration, we draw
 158 the back faces of the model to F_1 as depicted in Figure 2(b).
 159 Given F_0 and F_1 , we can now calculate the diameter value for
 160 the vertices that belong to the first depth layer D_0 , which is the
 161 one stored in F_0 . The depth layer D_0 is highlighted in red in
 162 Figure 2(e). We bind a framebuffer of the same dimension of
 163 the texture T and we draw a full-screen quad to launch a frag-
 164 ment program for each vertex. The fragment program fetches
 165 the position for each vertex from T . The vertex position is trans-
 166 formed by the current camera matrix $c \in C$ and its projection in
 167 the image space v_p is found along with its distance v_z with re-
 168 spect to the camera. At this point, we obtain the depth values
 169 z_0 and z_1 by sampling respectively the textures D_0 and D_1 (red
 170 and green lines in Figure 2(e), respectively) at the position v_p .
 171 If $v_z < z_1$, we conclude that the vertex must belong to the first
 172 depth layer D_0 and thus its diameter value along this view di-
 173 rection is given by: $d = \text{abs}(v_z - z_1)$. Otherwise, if $v_z > z_1$ we
 174 discard the fragment.

175 In the third iteration, we render to F_2 the front faces of the
 176 3D model but discarding all fragments whose depth is less than
 177 the one stored in D_1 . See Figure 2(c) for the resulting depth
 178 buffer D_2 . At the fourth iteration, we render to F_0 (overwriting
 179 all data previously stored) the back faces of the 3D model dis-
 180 carding all fragments whose depth is less than the one stored in
 181 D_2 . See Figure 2(d) for the resulting depth buffer D_0 . At this
 182 point, we can calculate the diameter value for the depth layer
 183 D_2 depicted in red in Figure 2(f). The computation of the di-
 184 ameter value proceeds exactly as in the previous pass except for
 185 the depth comparison performed between successive layers. As
 186 a matter of fact, at this pass we obtain the depth values z_0, z_2

187 and z_1 by sampling respectively the textures D_0, D_2 and D_1 at
 188 the projected vertex position v_p , where the order of the values
 189 is consistent with the use of our circular array.

Algorithm 1 Depth Peeling Ray Tracing

Require: set of cameras C
Require: circular array of 3 framebuffers F
Require: diameters texture array A
Require: maximum iterations N
Ensure: SDF value at each vertex

```

i ← 0
j ← 0
k ← 0
for each  $c$  in  $C$  do
  while  $i < N$  do
    setCamera( $c$ )
    if  $i == 0$  then
      bindShader(nopeeling)
    else if  $j > 0$  then
      bindShader(peeling)
      bindDepthTexture( $F[j - 1]$ )
    else
      bindShader(peeling)
      bindDepthTexture( $F[2]$ )
    end if
    if  $i \% 2 == 0$  then
       $F[j] \leftarrow \text{renderModel}(\text{GL.FRONT})$ 
    else
       $F[j] \leftarrow \text{renderModel}(\text{GL.BACK})$ 
    end if
    if  $i \% 2 \neq 0$  then
      if  $i > 0$  then
         $pB \leftarrow (j + 1) \% 3$ 
         $pF \leftarrow (j == 0) ? 2 : (j - 1)$ 
         $A[k] \leftarrow \text{calcDiameter}(F[pF], F[j], F[pB])$ 
      else
         $A[k] \leftarrow \text{calcDiameter}(F[j - 1], F[j], \text{NULL})$ 
      end if
       $k \leftarrow (k + 1)$ 
    end if
     $j \leftarrow (j + 1) \% 3$ 
     $i \leftarrow (i + 1)$ 
  end while
end for
 $SDF \leftarrow \text{weightedSum}(A)$ 

```

Algorithm 2 Calculate Diameter

Require: front faces buffer FRONT
Require: back faces buffer BACK
Require: back faces from previous step PREVBK
Require: texture pack with vertices normals T_n
Require: texture pack with vertices positions T_v
Require: current view direction D
Ensure: diameter value at current vertex

```

 $N \leftarrow \text{texelFetch}(T_n, \text{fragCoord.xy})$ 
 $\text{angle} \leftarrow \text{acos}(N \cdot D)$ 
if  $\text{angle} > 60$  then
  discard fragment and return
end if
 $V \leftarrow \text{texelFetch}(T_v, \text{fragCoord.xy})$ 
 $P \leftarrow \text{project}(V)$ 
if isFalseIntersection() then
  discard fragment and return
end if
 $z_{\text{Front}} \leftarrow \text{texture2D}(\text{FRONT}, P.xy)$ 
 $z_{\text{Back}} \leftarrow \text{texture2D}(\text{BACK}, P.xy)$ 
 $z_{\text{PrevBack}} \leftarrow \text{texture2D}(\text{PREVBK}, P.xy)$ 
if  $z_{\text{PrevBack}} < P.z$  and  $P.z < z_{\text{Back}}$  then
  return  $z_{\text{Back}} - z_{\text{Front}}$ 
else
  discard fragment and return
end if

```

190 In the fragment shader for SDF computation, we thus accept
 191 only vertices whose depth v_z satisfies the following inequality:
 192 $z_1 < v_z < z_0$. The comparison using the three layers is necessary

193 because vertices whose depth satisfies the condition $\mathbf{v}_z < \mathbf{z}_0$ but
194 not the condition $\mathbf{z}_1 < \mathbf{v}_z$ already got their values computed from
195 the previous passes. The algorithm described proceeds until no
196 more fragments get rasterized or after a fixed number of iterations.
197 To test if any fragment is drawn to the framebuffer, we
198 simply start an occlusion query before each depth peeling iteration.
199 We now point out some implementation details neglected
200 until now for the sake of clarity. When evaluating the SDF function,
201 the corresponding fragment program is early discarded if
202 the current view direction does not form an angle of less than
203 60 degrees with the inward-normal of the vertex under consid-
204 eration. Moreover, to perform the false intersection test from
205 [1], we also pack all the vertex normals in a texture and we test
206 the vertex normal against the normal at the point of intersec-
207 tion stored in the color texture drawn during the depth peeling
208 pass. Finally, once a single diameter value is computed for a
209 vertex, it is stored using the *imageStore* command in a texture
210 array bound to the SDF fragment program as an *image2DArray*.
211 Similarly, a counter contained in an unsigned integer texture is
212 increased using the *imageAtomicAdd* function. At the end of
213 the whole process, we have an unsigned integer texture con-
214 taining the number of rays traced per vertex and a texture array
215 containing all the ray lengths for each vertex. To gather the fi-
216 nal SDF value, we bind a framebuffer of the same dimension of
217 the texture \mathbf{T} and we draw a fullscreen quad. For each vertex,
218 we now loop over the diameter values contained in the texture
219 array and we sort them using a naive bubble sort. Then, we
220 calculate mean and variance of the values and we calculate the
221 weighted sum of all rays lengths that are less than one standard
222 deviation away from the median (which is taken as the middle
223 element of the sorted values). We can perform SDF calculation
224 both for vertices and faces. The aforementioned explanation as-
225 sumes that values are calculated per vertex. For the face case,
226 we trace rays through the barycenter of the triangle.

227 3.2. Parallelized Ray Casting using OpenCL

228 Our second approach for the computation of the SDF ex-
229 tends the original implementation of the algorithm by porting
230 it to the GPU using OpenCL. The computation of the SDF is
231 a three-step process that consists of ray generation, ray casting
232 and traversal of the chosen acceleration structure and finally the
233 computation of the distance at the point of intersection.




Figure 3: Rays are cast in cones from each vertex. The cast rays are displayed as red lines.

234 In the first step, we have to generate rays in a cone centered
235 around the inward-normal of a given point and their associated

236 weights (Figure 3). As in the original algorithm, the weight of
237 each ray is given by the inverse of the angle between the ray
238 and the center of the cone. The generation of rays on the CPU
239 and their transfer to the GPU create unnecessary overhead be-
240 cause we must store every ray and its associated weight. There-
241 fore, the rays should be generated on the GPU. However, ran-
242 dom generation of rays would imply having a pseudo-random
243 generator in the OpenCL. Rolland [3] has tackled this problem
244 by defining a cone sampling strategy consisting of random rays
245 that are uniformly generated inside the cone. Our algorithm fur-
246 ther expands the former method by evenly distributing the rays
247 in a cone using the spherical Fibonacci point set [10]. Further-
248 more, a uniform distribution of rays helps prevent unnecessary
249 bias in the values on the mesh. A comparison of random ray
250 generation with respect to our approach can be seen in Figure
251 4. The algorithm generates rays restricted to a given sphere cap.




Figure 4: (Left) random ray generation, (right) uniform ray generation using spherical Fibonacci point set.

251 Each ray is then transformed into the world coordinates by mul-
252 tiplying it by the tangent space matrix specified by the point's
253 inward-normal and by the two orthogonal unit vectors spanning
254 the tangent plane of the point. It is important to note that only
255 valid points are used to generate the rays, where a point is valid
256 if it has a non-zero normal, tangent and binormal vectors. In the
257 case of non-manifold models, the use of triangle barycenters in-
258 stead of vertices is more reliable as the normal of some points
259 cannot be properly determined. As a side effect, this also helps
260 reduce unnecessary data transfer thanks to the fact that normal,
261 tangent and binormal can be easily calculated.

262 In the second step, for each ray, we traverse our acceleration
263 structure. The acceleration structure is one of the most impor-
264 tant parts of ray tracing. For the purpose of comparison with the
265 original article, we built an octree structure, even though a BVH
266 or kdTree could be traversed faster. Our octree implementation
267 is based on Laine's stack-based approach [11].

268 In the third step, we compute ray-triangle intersections be-
269 tween the cast ray and the triangles contained in the found oc-
270 tree node. For this purpose, we use the Moller and Trumbore
271 algorithm [12] because it mainly uses the dot and cross prod-
272 ucts which can be evaluated very efficiently on current graph-
273 ics hardware. Once the values per ray are calculated, they are
274 transferred back to the CPU, where the final outliers removal
275 takes place. The rays with lengths that do not fall within one
276 standard deviation from the median of all lengths are removed.
277 The inliers are then averaged using the weights calculated in
278 the first step. To overcome the errors in the measure caused by
279 pose changes, a smoothing operation is necessary. Therefore,
280 we perform a k -ring neighborhood Gaussian smoothing, where
281

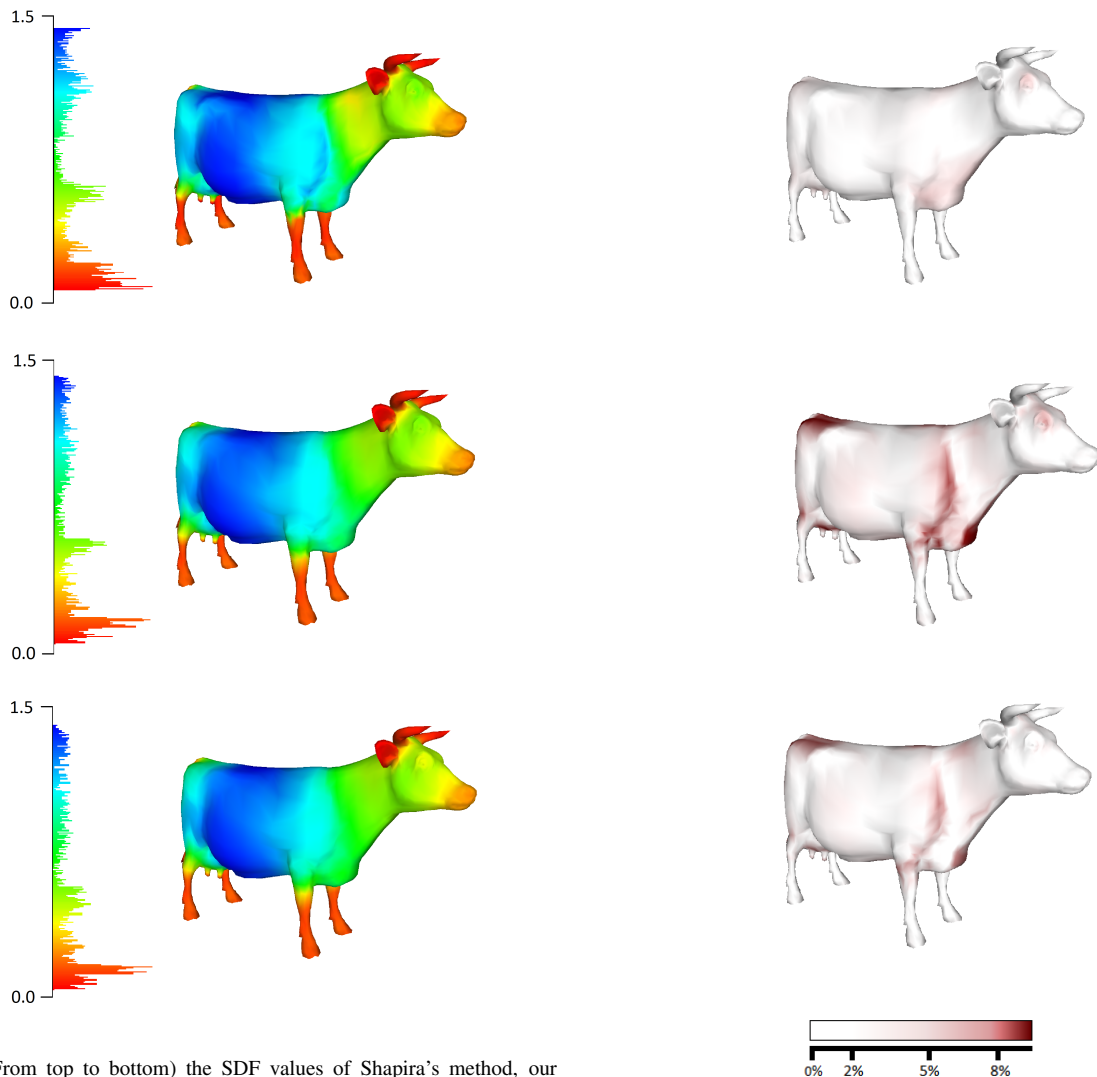


Figure 5: (From top to bottom) the SDF values of Shapira’s method, our OpenGL depth peeling and our OpenCL raycasting. The SDF values are colored in a scale from low (red) to high (blue) and the histogram of computed SDF values is on the left side of the image.

Figure 6: (From top to bottom) the difference of SDF values of OpenGL depth peeling vs. OpenCL raycasting, Shapira vs. OpenCL raycasting and Shapira vs. OpenGL depth peeling. (Bottom) the color scale of differences from white to red.

282 k specifies the blur radius in terms of edges in neighbourhood
283 radius.

284 4. Results and Comparison

285 We now compare the results from both our methods for par-
286 allel computation of the SDF. Finally, at the end of this section,
287 a texture-space diffusion method is proposed in order to smooth
288 the values across the surface. A performance comparison of the
289 two approaches is presented in Table 1. Furthermore, we have
290 compared the values of OpenGL depth peeling, OpenCL ray-
291 casting and the original Shapira’s implementation. The results
292 of all three sets of SDF values with their related histograms
293 can be seen in Figure 5. The differences among the methods
294 mapped into a red scale are visualized in Figure 6. To ensure
295 the correctness of our implementation, several basic tests have
296 been performed on different shapes with known diameters. Fig-
297 ure 7 shows all three methods applied on a circle swept along
298 a bezier curve with a diameter linearly interpolated between 2

299 and 20 units. All three methods give almost the same results
300 and the highest deviation (up to 10%) from the real diameter has
301 been measured in the thickest parts of the model using the orig-
302 inal Shapira’s method. Tests on all 3D models were performed
303 using 30 rays for each point and with a cone aperture of 120.
304 For the depth peeling method, we used 256 camera positions
305 around each model. The differences between the SDF values of
306 our two methods can be seen in Figure 9. All tests have been run
307 on a desktop PC with the following configuration: Intel Core i5,
308 2,67GHz with 4GB RAM and AMD Radeon R9 290. As you
309 can see in Table 1, the OpenCL raycasting is more efficient for
310 3D models with a modest number of vertices. For 3D models
311 with a high vertex count (above 150k), the depth peeling ap-
312 proach is instead definitely faster. The processing times for the
313 CPU approach using Shapira’s method [1] and the processing
314 times for the smoothing phase are also presented in Table 1.
315 To perform the smoothing, each vertex and his neighbors are




Figure 7: The SDF computed on a model with a known diameter: a circle swept along a bezier curve, where the diameter is linearly interpolated between 2 and 20 units. Results using (a) Shapira’s method, (b) OpenGL depth peeling and (c) OpenCL raycasting. The values are visualized from low (red) to high (blue). The histogram on the left of each method shows the distribution of the computed diameter values.

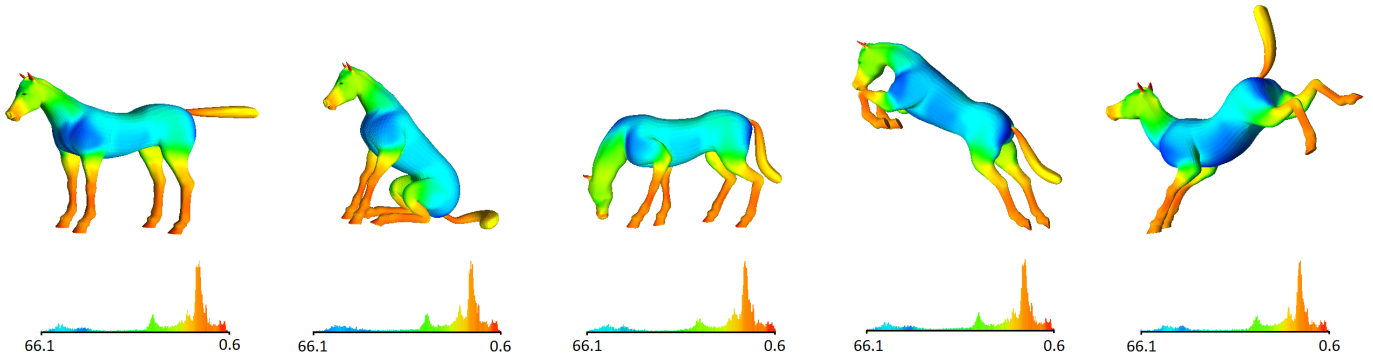


Figure 8: The SDF computed on a model in different poses using OpenGL depth peeling. The histogram below the images show that the SDF values computed using our technique are stable under pose changes.

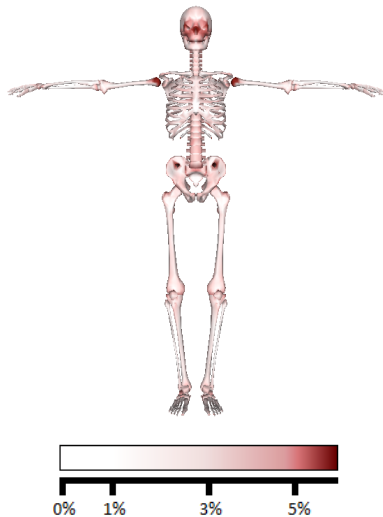


Figure 9: Comparison of computed values using our two GPU approaches. (Top) absolute difference of values between approaches. White vales are below 1%, a high difference is around 5%. (Bottom) the color scale of differences from white to red.

Mesh	Vertices	Faces	OGL	OCL	CPU	Smooth
Cow	2903	5804	392	31	878	43
C-shape	3554	7104	297	32	1017	46
S-shape	10034	20064	485	141	1984	140
Skeleton	8168	16040	687	78	735	141
Bunny	35947	69451	661	624	10376	530
Budha	144647	293232	1796	2262	39082	2808
Armadil.	172974	345944	2515	2901	89004	3557
Candela.	249976	499942	2418	4027	130324	4665
Gargoyle	500002	1000000	3957	8206	262695	7301

Table 1: Processing time comparisons of our two approaches and original Shapira’s approach. The “OGL” are times for the depth peeling technique, “OCL” are times for the parallel ray casting, “CPU” are times for the Shapira’s implementation and “Smooth” are times for the postprocessing step that makes Gaussian diffusion of SDF values on the mesh surface. For this smoothing step, we used neighborhood radius parameter $k = 2$. Times are in milliseconds [ms].

316 taken into account and the average of their values is calculated.
 317 Smoothing is computed on the CPU, because it cannot be im-
 318 plemented easily on the GPU as we need the k-ring connectivity
 319 information of the neighborhood of each vertex. As you can see
 320 in Table 1, this is the slowest step in our implementation.

321 4.1. Postprocessing using Texture-space Diffusion

322 We can use the extracted skeleton and a skeleton texture
 323 mapping (STM) [13] to blur the SDF values on the surface of
 324 the object. The idea is to map the SDF values into texture space,
 325 make the Gaussian diffusion there and map the smoothed values
 back to the model surface, as shown in Figure 10. Texture-




Figure 10: Computed SDF values (left) are stored in the texture and blurred (right) using texture-space diffusion with STM parameterization. Here, a kernel with radius 10px was used.

326 space diffusion using STM guarantees that local neighborhood
 327 of each point on the model surface is homotopic to local neigh-
 328 borhood in the texture space. Comparison of smoothing on the
 329 mesh using various k -ring areas and smoothing in the texture
 330 using Gaussian filter with various radius settings can be seen in
 331 Table 2.
 332

Radius	k-ring	CPU Smoothing	STM Smoothing
64	6	14,882	0,140
32	5	10,764	0,093
16	4	7,566	0,057
8	3	5,024	0,051
4	2	3,120	0,048
2	1	1,762	0,041

Table 2: Smoothing in a 2048 x 2048 texture using various radius settings. Times shown are in seconds [s].

333 5. Further Applications and Extensions

334 In this section, we present some applications of computed
 335 SDF values focusing on skeleton extraction. These applications
 336 need to compute SDF values at each step during an iterative
 337 process. Therefore, these applications of SDF are a motivation
 338 for fast parallel SDF computation on the GPU. In the skele-
 339 ton texture mapping approach presented in [13], we need to
 340 compute an internal skeleton from which every mesh vertex is
 341 visible. There are automatic techniques for curve-skeleton ex-
 342 traction [14, 15, 16], but extracted curve-skeletons using these
 343 techniques do not satisfy the above mentioned condition. Be-
 344 cause of this, we have extended the Laplacian-based smoothing
 345 method presented in [14] with an SDF term that moves skeleton
 346 vertices near the medial axis transform (MAT). In addition, the
 347 SDF values can be used to modify parameters as weights and
 348 grouping distances during Laplacian-based skeleton extraction
 349 to obtain better results with skeleton extraction. For further de-
 350 tails, we refer the reader to the original Laplacian-based mesh
 351 contraction [14]. Finally, we show an application of the SDF
 352 for the fast linear skeleton extraction.

353 5.1. Proposed SDF Modifications for Laplacian-based Skele- 354 ton Extraction

355 In order to parameterize a mesh using a skeleton texture
 356 map (STM) [13], we have to guarantee that the skeleton is reli-
 357 able. Reliability refers to the property of the curve-skeleton that
 358 every boundary point (i.e., the point on object’s surface) is vi-
 359 sible from at least one skeleton location [17]. Some basic ideas
 360 behind the modifications of skeleton extraction algorithm are
 361 shown here. However, the in-depth discussion of this topic is
 362 out of the scope of this paper. For a more detailed description,
 363 we refer the reader to [18].

364 5.1.1. Skeleton Sheets using SDF

In most cases, a mesh can be parameterized by a curve-
 skeleton, but not in general. As a matter of fact, in some cases,
 the skeleton must contain sheets, similarly to the medial axis
 transform (MAT). A skeleton consisting of a mixture of curves
 and surface sheets was defined in [19] and called a “meso-
 skeleton”. In [19], authors used Voronoi poles to guide the sur-
 face toward the medial axis. Here we show that this is the place
 where the SDF can be efficiently used. Shifted vertices into
 the central sheet using SDF are less noisy than Voronoi poles
 (or MAT) and SDF can be computed faster using our GPU ap-
 proaches. Furthermore, our method creates skeleton sheets only
 in the regions where the curve-skeleton reliability is not satis-
 fied. Meso-skeletons created using Voronoi poles [19] present
 sheet surfaces also in the regions which could be parameterized
 using just the curve-skeletons. This is clearly not efficient for
 STM because skeleton sheet takes much more space in the final
 texture. Therefore, we want to have the skeleton sheets only in
 the regions where they are inevitable. We extended the mesh




Figure 11: Examples of skeleton sheets of C-shape and S-shape models. Classical Laplacian-based skeleton extraction fails here because of non-convex cross-sections along the skeleton.

contraction method for skeleton extraction [14] with an SDF centering term as follows:

$$\begin{bmatrix} \mathbf{W}_L \mathbf{L} \\ \mathbf{W}_H \\ \mathbf{W}_C \end{bmatrix} \mathbf{V}' = \begin{bmatrix} \mathbf{0} \\ \mathbf{W}_H \mathbf{V} \\ \mathbf{W}_C \mathbf{C} \end{bmatrix}, \quad (1)$$

365 where matrix \mathbf{C} is composed of vertices shifted by the SDF
 366 values in the opposite of normal direction $\mathbf{c}_i = \mathbf{v}_i - (SDF_i/2) \cdot$
 367 \mathbf{n}_i . \mathbf{v}_i , \mathbf{n}_i and SDF_i are respectively the position, normal and
 368 SDF value of the vertex i . For a more detailed description of

369 the original contraction weights (\mathbf{W}_L , \mathbf{W}_H) and the introduced
 370 weights \mathbf{W}_C we refer reader to [14] and to [18], respectively.

371 During the iterative contraction process, we tend to collapse
 372 the vertices into the medial axis composed of vertices shifted by
 373 half of the SDF values in the opposite of the normal direction.
 374 Using this extended version, we are able to create the curve-
 375 skeleton in the parts of the model where this is reliable and
 376 skeleton sheets where the curve-skeleton would not satisfy the
 377 reliability condition (see Figure 11).

378 5.2. Modification of Iterative Parameters

379 At each iteration of the iterative contraction process, the
 380 SDF values are computed and used to drive the modification
 381 of the contraction weights and grouping distance parameters.
 382 This results in skeletons which reflect better small-scale details
 383 of the 3D model (see Figure 14).

384 5.2.1. Updating of Grouping Distance using SDF

385 Here we show how SDF values can be used to create a skele-
 386 ton from a not fully contracted mesh. During the simplification
 387 of the contracted mesh, we add another condition upon which a
 388 half-edge collapse is applied. This new condition is satisfied if
 389 an Euclidean distance between two nodes is less than a distance
 390 threshold dt_i . The distance threshold for each node is computed
 391 as: $dt_i = ggt \cdot diag + (SDF_i/2) \cdot gmt$, where ggt is a global group-
 392 ing tolerance term, SDF_i is an SDF value for vertex i , gmt is
 393 global multiplication term and $diag$ is a diagonal of the model's
 394 bounding box. We obtained the best results with values set to
 395 $ggt = 0.05$ and $gmt = 1.25$. These parameters enable all the
 396 uncontracted parts of the mesh to collapse while the contracted
 397 parts of the mesh remain unchanged (see Figure 12).




Figure 12: Grouping distance using SDF. (a) Contraction of the mesh when the volume is too big for skeleton extraction, (b) using lower global grouping distance, (c) using higher global grouping distance, (d) using local grouping distance based on SDF.

398 5.2.2. Updating of Contraction Weights

399 During the iterative contraction process, we change the hold-
 400 ing weight \mathbf{W}_{H_i} after each iteration based on the SDF values.
 401 We scale \mathbf{W}_{H_i} according to the normalized SDF value $SDFnorm_i$.
 402 First, we update \mathbf{W}_{H_i} according to the surface change of one

403 ring areas as $\mathbf{W}_{H_i}^{t+1} = \mathbf{W}_{H_i}^0 \sqrt{A_i^0/A_i^t}$, where A_i^0 and A_i^t are the
 404 original and current areas of adjacent faces for vertex i , respec-
 405 tively. Then we normalize the SDF values to the interval $(0, 1)$
 406 as $SDFnorm_i = (SDF_i - minSDF)/(maxSDF - minSDF)$,
 407 where $minSDF$ and $maxSDF$ are the minimal and the maxi-
 408 mal SDF values. Afterwards, we change \mathbf{W}_{H_i} according to the
 409 volume change using SDF as $\mathbf{W}_{H_i} = \mathbf{W}_{H_i}/(SDFnorm_i * (1.0 -$
 410 $bias) + bias)$, where $bias$ is an offset value because we want to
 411 avoid division by zero. We set $bias$ as 0.01 in our tests, which
 412 produces high \mathbf{W}_{H_i} values for $SDFnorm_i = 0$, while avoiding a
 413 division by zero. Using our iterative modification of contraction
 414 weights (Figure 13), we were able to preserve small-scale details
 415 also in cases where the thickness of the 3D model changes
 416 rapidly (see Figure 14). More details concerning this applica-
 417 tion can be found in [18].




Figure 13: Contraction weights using SDF. The mesh is colored according to the SDF values. Red regions with higher SDF values are more contracted and blue regions with lower SDF value are less contracted.




Figure 14: Difference of small-scale details as hands, ears and tail without SDF (left) and with the SDF term (right).

418 5.3. Fast Linear Skeleton Extraction

419 The SDF values have been used for the fast linear skeleton
 420 extraction from particle simulation. We used our fast parallel
 421 GPU implementation for linear skeleton extraction from parti-
 422 cle simulation of the worm body [20]. The extracted skeleton

was used to compress the simulation and to visualize it using a skinning algorithm.

In order to compute the SDF of an input point cloud, a surface must be reconstructed. The reconstructed surface must be closed, but it must not be 2D-manifold. Therefore, we construct a local Delaunay triangulations for the neighborhood of each vertex in its tangent plane. In the next step, all edges in the local triangulations are copied into the global triangulation. This results in a graph that is a non-manifold orientable closed global triangulation.

Having unoriented point clouds, a normal can be estimated using PCA. However, the orientation of the normal remains unknown posing a serious problem for a correct SDF calculation. Therefore, we orient the normals in two steps. First, we choose a vertex whose normal orientation can be easily decided. For this purpose, we choose the outermost vertex in one direction, for example the topmost vertex. For this vertex, we are able to choose the normal orientation, because we know that its y coordinate has to be positive, facing out of the mesh bounding box. We then propagate the orientation over the mesh surface, by comparing normal orientations between neighboring vertices. At the end of this process we obtain an oriented point cloud.

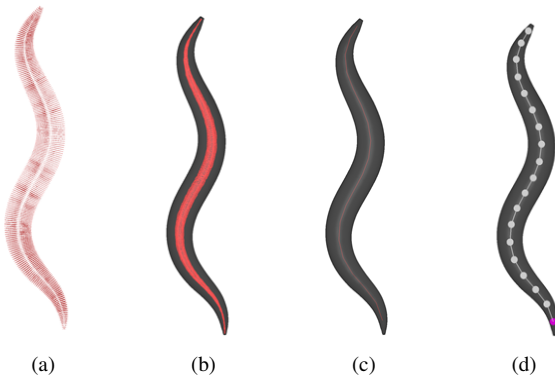


Figure 15: The algorithm for skeleton extraction from the worm cuticle based on SDF. (a) Mesh of cuticle, (b) SDF half-vectors, (c) Shifted mesh vertices, (d) poly line obtained by agglomerative clustering and (e) extracted skeleton structure.

The cuticle of the worm has a simple topology without branching and a fast skeleton extraction algorithm is needed. The algorithm first computes the SDF values and then constructs the final skeleton using agglomerative clustering with the distance threshold t (see Figure 15). From clustered polyline a more subdivided skeleton can be easily constructed. The method is faster than all known robust methods and works well for the worm topology.

In order to match skeletons from all the timesteps and calculate the rotations between different timesteps, we need to ensure that each skeleton has the same number of segments. Thus, we find a clustering threshold t using binary search between the minimal and maximal SDF values computed on the surface. The computed threshold t is then used for agglomerative clustering during the skeleton construction.

6. Discussion

The way in which the parallelized OpenCL raycasting proceeds is similar to Shapira’s approach [1]. The OpenGL depth peeling method proposes another point of view to approach the SDF calculation. Both proposed approaches achieved similar results in average, even though the depth peeling method is a bit faster in comparison to the parallelized version of ray casting technique, especially when the input model has more than 150k vertices. This is due to the fact that the whole SDF processing is calculated in a buffer space. Therefore, after the rendering phase, the calculation is independent of the mesh complexity. Both GPU methods are faster than former method as can be seen in Table 1. Unfortunately, we were not able to compare times with the Rolland’s method [3], because computational times were not mentioned in their paper. The comparison of our two implementations against the Shapira’s version shows that the results obtained are very similar. For this test, we set up our implementations to follow the details outlined in the original paper. The differences in the calculated values are very likely due to implementation details which were not mentioned in the original paper. Finally, as demonstrated in Figure 8, our depth peeling method is robust to pose changes with essentially no differences to the original Shapira’s method with given a sufficient sampling of the view sphere around the model.

7. Conclusions and Future Work

We have implemented two different approaches for the computation of the SDF and compared their results and processing times. Both approaches are implemented on the GPU. The first one is calculated using shaders and depth peeling while the second one is based on parallelized ray casting using OpenCL. As mentioned in the discussion, we found out that it is more efficient to use OpenCL raycasting for relatively small meshes and OpenGL depth peeling for relatively large ones.

Furthermore, we presented several applications of the SDF. We have extended Laplacian smoothing-based skeleton extraction by an SDF term to obtain skeletons with finer details. We also extended curve skeletons to skeleton sheets which are contained in models with non-convex cross-sections. Finally, we used our customized SDF linear skeleton extraction for the fast compression of particle simulation based on skinning.

In both of our GPU implementations for the SDF computation, we maintain the original outliers removal approach proposed in [1], leaving the one proposed by [3] for future work. As a matter of fact, the adaptive cone size proposed in [3] helps with the estimation of the SDF only in some special cases, but it is computationally more expensive than original method as the ray casting must be iterated several times to find the correct cone size.

References

- [1] L. Shapira, A. Shamir, D. Cohen-Or, Consistent mesh partitioning and skeletonisation using the shape diameter function, *Vis. Comput.* 24 (2008) 249–259.

- 512 [2] H. I. Choi, S. W. Choi, H. P. Moon, Moon: Mathematical theory of medial
513 axis transform, *Pacific J. Math.*
- 514 [3] X. Rolland-Nevière, G. Doërr, P. Alliez, Robust diameter-based thickness
515 estimation of 3d objects, *Graphical Models* 75 (6) (2013) 279–296.
- 516 [4] L. Fan, L. Lic, K. Liu, Paint mesh cutting, *Computer Graphics Forum*
517 30 (2) (2011) 603–612.
- 518 [5] M. Kovacic, F. Guggeri, S. Marras, R. Scateni, Fast approximation of the
519 shape diameter function, in: *GraVisMa 2010 Proceedings*, 2010, p. 1724.
- 520 [6] C. Everitt, Interactive order-independent transparency, White paper,
521 nVIDIA 2 (6) (2001) 7.
- 522 [7] L. Bavoil, K. Myers, Order independent transparency with dual depth
523 peeling, Tech. rep., NVIDIA Developer SDK 10 (Feb. 2008).
- 524 [8] N. A. Carr, J. D. Hall, J. C. Hart, The ray engine, in: *Proceedings of*
525 *the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hard-*
526 *ware, HWWS '02*, Eurographics Association, Aire-la-Ville, Switzerland,
527 Switzerland, 2002, pp. 37–46.
- 528 [9] T. J. Purcell, I. Buck, W. R. Mark, P. Hanrahan, Ray tracing on pro-
529 grammable graphics hardware, *ACM Trans. Graph.* 21 (3) (2002) 703–
530 712.
- 531 [10] R. Marques, C. Bouville, M. Ribardire, L. P. Santos, K. Bouatouch,
532 Spherical fibonacci point sets for illumination integrals, *Computer Graph-*
533 *ics Forum* 32 (8) (2013) 134–143.
- 534 [11] S. Laine, T. Karras, Efficient sparse voxel octrees, in: *Proceedings of*
535 *the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and*
536 *Games, I3D '10*, ACM, New York, NY, USA, 2010, pp. 55–63.
- 537 [12] T. Möller, B. Trumbore, Fast, minimum storage ray-triangle intersection,
538 *J. Graph. Tools* 2 (1) (1997) 21–28.
- 539 [13] M. Madaras, R. Đurikovič, Skeleton texture mapping, in: *Proceedings*
540 *of the 28th Spring Conference on Computer Graphics, SCCG '12*, ACM,
541 New York, NY, USA, 2013, pp. 121–127.
- 542 [14] O. K.-C. Au, C.-L. Tai, H.-K. Chu, D. Cohen-Or, T.-Y. Lee, Skeleton
543 extraction by mesh contraction, *ACM Trans. Graph.* 27 (3) (2008) 44:1–
544 44:10.
- 545 [15] G. Aujay, F. Hétry, F. Lazarus, C. Depraz, Harmonic skeleton for re-
546 alistic character animation, in: *Proceedings of the 2007 ACM SIG-*
547 *GRAPH/Eurographics Symposium on Computer Animation, SCA '07*,
548 Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2007,
549 pp. 151–160.
- 550 [16] T. K. Dey, J. Sun, Defining and computing curve-skeletons with medial
551 geodesic function, in: *Proceedings of the 4th EG symposium on Geom.*
552 *processing*, 2006, pp. 143–152.
- 553 [17] N. D. Cornea, D. Silver, P. Min, Curve-skeleton properties, applica-
554 tions, and algorithms, *IEEE Transactions on Visualization and Computer*
555 *Graphics* 13 (3) (2007) 530–548.
- 556 [18] R. Đurikovič, M. Madaras, *Mathematical Progress in Expressive Im-*
557 *age Synthesis II: Extended and Selected Results from the Symposium*
558 *MEIS2014*, Springer Japan, Tokyo, 2015, Ch. Controllable Skeleton-
559 *Sheets Representation Via Shape Diameter Function*, pp. 79–90.
- 560 [19] A. Tagliasacchi, I. Alhashim, M. Olson, H. Zhang, Mean curvature skele-
561 tons, *Comp. Graph. Forum* 31 (5) (2012) 1735–1744.
- 562 [20] M. Piovračić, M. Madaras, R. Đurikovič, Physically inspired stretching
563 for skinning animation of non-rigid bodies, in: *Proceedings of the 31st*
564 *Spring Conference on Computer Graphics, SCCG '15*, ACM, New York,
565 NY, USA, 2015, pp. 47–53.