



A Complete Distance Field Representation

Jian Huang¹, Yan Li², Roger Crawfis¹, Shao-Chiung Lu³ and Shuh-Yuan Liou⁴

¹Computer and Information Science, The Ohio State University, Columbus, OH

²Electrical Engineering, The Ohio State University, Columbus, OH

³Visteon Inc., Dearborn, MI

⁴Ford Motor Company, Dearborn, MI

Abstract

Distance fields are an important volume representation. A high quality distance field facilitates accurate surface characterization and gradient estimation. However, due to Nyquist's Law, no existing volumetric methods based on the linear sampling theory can fully capture surface details, such as corners and edges, in 3D space. We propose a novel complete distance field representation (CDFR) that does not rely on Nyquist's sampling theory. To accomplish this, we construct a volume where each voxel has a complete description of all portions of surface that affect the local distance field. For any desired distance, we are able to extract a surface contour in true Euclidean distance, at any level of accuracy, from the same CDFR representation. Such point-based iso-distance contours have faithful per-point gradients and can be interactively visualized using splatting, providing per-point shaded image quality. We also demonstrate applying CDFR to a cutting edge design for manufacturing application involving high-complexity parts at un-precedented accuracy using only commonly available computational resources.

CR Categories: I.3.6 [Computer graphics]: Methodology and techniques - Graphics data structures; I.3.5 Computational Geometry and Object Modeling - Object modeling

Keywords: distance fields, volume modeling, polygonal surfaces, point-based models, graphics.

1. Introduction

A voxel-based volume, as a 3D raster, holds discrete sample points representing a certain multi-dimensional entity. In an alias-free volume discretization, only frequency components below half the Nyquist sampling rate would be stored. As a natural description of solid physical entities, volume representations have found applications in a variety of areas, including medicine, mechanical engineering, scientific computing and simulations. In order to utilize volume technologies, it has been common to convert surface models, such as a polygonal mesh exported by a CAD package, to a volume representation. In this process, first, one needs to voxelize the surface model into a hollow volume representing the surface shape [1][13][14]. Second, a distance transform is computed to construct a solid volume that encompasses a distance or thickness field recording distances to the surface. Euclidean distance has not been commonly used due to both efficiency concerns and the fact that accuracy is already compromised in the binary surface volume model. Instead, most applications use less accurate distance heuristics such as Manhattan or chessboard distance, or a Chamfer distance [3].

Voxelization techniques that convert surface shapes into binary volumes, with 1's representing occupancy and 0's repre-

senting empty space, have been developed in [1][13][14]. These methods are practical and commonly used nowadays. Surface shapes, however, are infinitely thin in space. To sample this thin shape, one needs infinitely high sampling rates. To avoid such a difficulty, Kaufman's algorithm [13][14] increases the thickness of a surface, and therefore, to some extent, band limits the frequency spectrum before the sampling, or scan-conversion, process. Huang et. al [11] discovered and proved the sufficient and necessary thickness of the surface shape that guarantees a correct discrete topology in the resulting volume representation. Unfortunately, all these methods are based on binary volume representations, which are highly susceptible to aliasing artifacts. To address this issue, Sramek and Kaufman initiated data representations in non-binary formats [21]. In their paper, they show one has to use higher order smoothing functions to pre-filter and band limit the spectrum of the volume. Later, an incremental voxelization method for non-binary volume is reported in [5].

Over the years, in addition to the search for optimal voxelization, the community has also been exploring other representations of surfaces. Distance fields are scalar fields, with each element in the 3D volume representing the minimal distance to a certain shape. It is common practice to use signed values to distinguish between interior and exterior of the shape. Compared to the surface shapes that correspond to impulses in 3D space, distance fields are much smoother. For shapes without sharp corners and edges, both the surface position and gradient can be reconstructed relatively accurately using a distance field [1][9]. When corners and sharp edges are introduced, high frequency components are also brought into the spectrum. To preserve such details, super-sampling with exceptionally high volume resolution, as well as low-pass filtering, is necessary to achieve an alias-free representation.

In this paper, we present a novel scheme for representing distance fields. We build our volumetric representation of distance fields based on a complete distance definition. Our approach is separate from the theory of linear sampling. We name our distance field representation a "complete distance field representation (CDFR)", because once the distance volume is constructed, we can extract any distance contour to any requirement of accuracy. As a comparison, conventional approaches based on a single valued distance field can only achieve higher accuracy by re-building the whole distance volume at an increased resolution. However, in most cases building a high resolution distance volume is non-trivial both in computational time and storage space. A recent work [15] stores the estimated local edge positions in x , y and z directions with each voxel and can extract triangle meshes from volume data at a much improved quality compared to conventional methods [2][10]. However, their distance field representation is not complete and the surface extraction is still based on estimation. Our method only relies on exact computations, and, different from the approach in [2][10][15], uses a point-based approach to represent the extracted contour surfaces.

Hierarchical data structures applied to represent distance fields efficiently have been reported in [7], where adaptively sampled distance fields (ADF) are introduced. ADFs help in reducing volume storage size when fewer details are locally present. The specific ADF implementation described in [7] relies on a single valued distance representation, therefore that implementation still depends on a band-limited spectrum that discards all details beyond the cut-off bandwidth supported by the leaf level in the tree structure. Our CDFR algorithm can be embedded into the ADF structure for an exact distance field representation that is efficient both in terms of storage and processing time.

Conventionally, distance fields are most often rendered with ray-casting approaches [1][5][7][9]. High interactive rates are often not attainable. CDFR can also be rendered with a slightly modified ray-casting scheme. However, for efficient rendering of distance contours, we also discuss a method to reconstruct a point-based contour of any distance value from CDFR. Such sparse point-based models can be efficiently rendered with splatting [12].

The paper is organized in the following way. We give a brief introduction to single valued distance fields and the limitations of such representations in Section 2. In Section 3, we present a complete distance definition (CDD) and complete distance field representation (CDFR) based on CDD. We prove the correctness of our schemes in constructing CDFRs and extracting point-based contours from CDFRs in Section 4. Results on a variety of data sets are shown in Section 5. Finally, we conclude and discuss future work in Section 6.

2. Distance Fields

Traditionally, distance fields are defined as spatial fields of scalar distances to a surface geometry or shape. Each element in a distance field specifies its minimum distance to the shape. As long as the shape is represented by an oriented manifold, positive and negative distances can be used to distinguish outside and inside of the shape, for instance, using negative values on the outside and positive on the inside. Distance fields have a number of applications in constructive solid geometry [1][7], surface reconstruction and normal estimation [9] and morphing [1][3]. Distance fields are also applied to concurrent engineering [16] where simulations and analysis involving the interior of geometries, such as die-casting simulation or thickness analysis of parts [22], are routine.

For an alias-free sampling of a signal, Nyquist's Law dictates that the sampling rate must be at least two times the highest frequency component in the signal. In spatial domain, geometry is infinitesimally thin, and has an infinitely wide frequency spectrum. The sharp details on the surface, such as corners and edges, also reside on the high ends in the spectrum. Even with an overwhelmingly large volume resolution, one still needs extensive low-pass filtering to limit the bandwidth of the geometric shape. These low-pass filtering operations, with either simple box filters [11][13][14] or specifically designed higher order filters [21], inevitably cause a loss of the exact surface details. Converting the surface shape to a distance field, which is smoother, provides a way to exactly locate the surface [9] during reconstruction. But the underlying assumption of having a completely smooth surface that is free from sharp corners and edges is unrealistic for most scenarios.

Frisken et. al [7] developed a well analyzed framework for adaptively sampled distance fields (ADF), by which one can build hierarchies of distance fields at different levels of detail and be able to cross over different levels of detail as needed. They also vary sampling rates according to the amount of details that are

available locally. They used tri-linear interpolation to reconstruct distances, and were able to demonstrate a suite of applications with impressive rendering quality. However, ADF [7] does not fundamentally solve the problem of losing surface details in discrete representations. After the leaf level of ADF is constructed, the loss in surface details is final and irreversible. When the primary goal of an application shifts from visual quality to accuracy, ADFs with trilinear interpolation may not satisfy the accuracy needs with a guarantee, simply because the true distance fields are not linear where corners or edges are present. What the hierarchies provide is an ability to save computational and storage resources when less details are encountered. For models with fine details everywhere, the ADF eventually resorts to an extremely large voxelization. For the applications where accuracy is highly sought after, current ADFs based on single valued distances incur overwhelming costs, because most practical geometrical models are rich in details at a wide range of scale.

A high quality distance field should be accurate, efficiently stored and can be efficiently processed. There are two fundamental issues involved in building a high quality distance field. First, we need an accurate way to represent the distance from an arbitrary point in 3D space to an arbitrary shape. Second, how should we optimally organize the distance representations in space?

This first question is our focus in this paper. We intend to show CDFR as a fundamental fix that preserves all geometric details in the true distance field. Only exact computations are used. ADF, as a systematic framework, addresses the second fundamental issue, by providing an adaptive and smooth transition between resolution levels, depending on the amount of surface details available.

3. Complete Distance Definition (CDD) & Complete Distance Field Representation (CDFR)

In this paper, we propose a complete distance definition (CDD). Corresponding to different surface representations, such as parametric surfaces, implicit surfaces or subdivision/polygonal mesh surfaces, there could be different instantiations of CDDs. For this paper, we focus on a simple case where the surface is represented by polygonal meshes. When accuracy is paramount, Euclidean distances are preferred over other distance metrics, such as Chamfer distance or Manhattan distance. The CDD distances are true Euclidean distances. Before discussing CDD, we will discuss a few observations that motivated our research on CDFR.

3.1 Some Observations

Distance fields are very smooth in some simple scenarios. For instance, suppose in a 1-dimensional space, we have an impulse. It's frequency components extend to infinity. There is no way to use a finite sampling frequency to sample the impulse without aliasing. But on the other hand, as illustrated in Fig. 1, the signed distance field of the impulse is a linear function which extends from negative infinity to positive infinity. Sampling this linear function can be accurate with a relatively low sampling rate.

Unfortunately, this feature does not hold in higher dimensions where corners are present. As presented in [11], when extended into 2D or 3D, the discrepancies and discontinuity on corners makes the distance field non-smooth. For instance, in the triangle in Fig. 2, we have a rather faithful sampling in the light grey grids on the edges, because the geometry is locally linear*. But the sam-

*Locally linear: reconstruction is accurate with linear interpolations.

pling is not sufficient in the dark grey grids that have corners. The non-linear distance fields within the dark grey grids make it impossible to accurately recover the correct distance distribution from the grid samples.




Figure 1: Without low-pass filtering, it's impossible to sample the impulse (left), but we can sample its distance field (right).




Figure 2: Corners in the triangle cause complexity in the distance field, resulting in an aliased spectrum after sampled.

According to Nyquist's Law, to sample such complicated distance fields, one must low-pass filter the corners and smooth out the sharpness. Using ADF [7], for the grids on the corners, a higher resolution would be used, whereas in the light grey grids, a much lower resolution might suffice.

To capture the exact location of the impulse in Fig. 1, we do not have to use sampling. Alternatively, all one needs is to place an anchor point at some location, and record the signed distance from the anchor to the impulse. In this method, preserving the exact position of the impulse is made straight-forward. This observation motivated our research towards a new distance representation for distance fields.

3.2 Complete Distance Definition (CDD)

CDD is a set of parameters describing both the distance from a 3D point to a surface geometry primitive and the geometry primitive itself. Specifically, when the shape is represented as a mesh of triangles, CDD reduces to a tuple that consists of a scalar canonical distance value, and a description of the triangle with a vertex list and an edge list:

$$\langle \text{distance}, \langle v_1, v_2, v_3 \rangle, \langle e_1, e_2, e_3 \rangle \rangle \quad (1)$$

The value *distance* is the true Euclidean distance from the voxel center to a finite triangle. This distance is defined in the pseudo code in Fig. 3.

While the return value is the CDD distance, the input parameters include a triangle, *tri*, and a 3D point, *pnt*. If *pnt* orthogonally projects into *tri* (case C_1), the return value is the orthogonal distance from *pnt* to the plane where *tri* lies. Otherwise, we check whether *pnt* orthogonally projects onto any of the three edges. If yes (case C_2), then the returned distance value is the shortest distance from *pnt* to an edge that *pnt* projects orthogonally onto. In case neither C_1 or C_2 applies (case C_3), the distance is the minimal distance from *pnt* to the three vertices. This definition of distance to a finite triangle is further illustrated in Fig. 4.

```
float CDD (triangle tri, vec3 pnt)
{
    float mindist = MAXIMUM;
    if (pnt projects orthogonally into tri's interior) // C1
        mindist = distance from pnt to tri's plane;
    else
    {
        for each edge of tri, e_i
            if (pnt projects orthogonally onto e_i) // C2
                mindist = min(mindist, distance from pnt to e_i);
        for each vertex of tri, v_i
            mindist = min(mindist, distance from pnt to v_i); // C3
    }
    return mindist;
}
```

Figure 3: Definition of distance from a point to a finite triangle




Figure 4: If *pnt* projects into *tri*, it's case C_1 . Otherwise, *pnt* is either C_2 or C_3 , depending on whether it's closer to an edge or a vertex. This diagram is drawn in 2D for ease of illustration.

We can still use positive and negative distance to distinguish inside and outside. We term the triangle that is the closest to *pnt* as the *base triangle* of *pnt*. If *pnt* is closest to a triangle and the distance is of case C_1 , then this triangle is *pnt*'s *base triangle*. If *pnt*'s distance is not case C_1 , rather, it's case C_2 or C_3 , looking for *pnt*'s *base triangle* is more complicated. For C_2 cases, let's label the projection point of *pnt* on that corresponding edge as, p' , and we record the vector pointing from p' to *pnt* as V . Between the two triangles sharing that edge, the triangle with a normal direction closer to V 's direction, i.e. larger absolute dot product value, $|V \bullet Normal|$, is *pnt*'s *base triangle*. Very similarly, in C_3 cases, among the triangles sharing that closest vertex, we can easily find out the *base triangle* of *pnt* by comparing dot product values. We are interested in finding out *pnt*'s *base triangle*, because by using the outward normal direction of the *base triangle* and the relative position of *pnt*, we can determine the sign of the distance at *pnt* without ambiguity.

To better illustrate the process in determining the distance sign, in Fig. 5a, we show several 3D examples, shown in 2D. p_1 through p_6 are 2D points. t_1 through t_6 are triangles that form the surface mesh. $p_2, 3, 4$ and 5 are all case C_1 . From the normal direction of t_2 , we can tell p_2 is outside, p_3 is inside. Similarly, using the normal of t_3 and t_5 , one can tell that p_4 is inside, and p_5 is outside, respectively. p_1 and p_6 are both C_2 . We show an enlarged view of these two cases in Fig. 5b. By comparing the dot products, we can tell p_1 's sign is determined by t_2 , and for p_6 , it is decided by t_1 .

Finally, to save space, we store the description of all triangles in a separate array and only keep a triangle index in a CDD tuple.




Figure 5: (a) 2D illustrations of the process to determine the sign of the distance of a point. The solid black arrows depict the outward normal direction of each triangle. Points p_2 through p_5 project into the triangles, i.e. case C_1 . The signs of the distances of p_2 through p_5 are determined by evaluating the normal direction of each point’s base triangle. p_1 and p_6 are examples of C_2 cases. (b) Enlarged view of p_1 and p_6 . Both p_1 and p_6 are outside.

3.3 A Complete Distance Field Representation (CDFR)

In this section, we show the process that uses CDD to build a complete distance field representation (CDFR), allowing exact capture of all geometric details, e.g. sharp corners and edges, to any level of accuracy.

Given a surface mesh, in the voxelization step, we store CDD tuples with each surface voxel, rather than single valued distances.




Figure 6: A 2D illustration of building a CDD tuple list for a surface voxel, Vox . There are 3 triangles intersecting Vox . The CDD tuple list is organized in ascending distance order, with the minimal distance of Vox being d_2 .

For each triangle touching a surface voxel, a CDD tuple is stored with that voxel. The end result of the voxelization step leaves all surface voxels with a list of CDD tuples, sorted in ascending order by distance values. Fig. 6 provides an example of voxelizing a single surface voxel, Vox . There are three triangles touching Vox . T_2 is case C_1 , with T_1 and T_3 being case C_2 or C_3 . The minimal distance of Vox , measured from the center of Vox is d_2 . As a result, Vox has a sorted list of 3 CDD tuples.

At the end of voxelization, we have a volume where each voxel which the surface intersects contains a list of polygons cutting through it.

3.4 Distance Transform

For a distance transform, initially, we use an outside flooding algorithm to eliminate all outside voxels from our computation. For the remaining voxels, a contour by contour CDD propagation

is performed from the surface voxels to the interior. During this process, a voxel looks for CDD tuples that have been newly propagated to anyone of its 26-neighbors [11]. It inherits all new CDD tuples from its neighbors, and for each triangle, it computes the true Euclidean distance from its own position. An updated list of CDD tuples are then sorted into ascending order and the first CDD tuple in the list contains the current distance, $cur_distance$, of this voxel. All the CDD tuples that contain a distance value within the range:

$$[cur_distance, cur_distance + \sqrt{3} * voxel\ size]$$

are stored with that voxel. This is a sufficient range to guarantee correctness in the distance transform, as we will prove in Section 4. The CDD tuples out of this range are discarded. This process of distance transform iterates until no voxels find new CDD tuples from its 26-neighbors affecting its current CDD tuples list.

3.5 Reconstructing A Distance Contour

The most frequent way in which a distance field is used is by reconstructing or extracting an iso-distance contour. For instance, a user asks the following request, “show me the zero distance contour with an error tolerance of 0.5mm”. The conventional way of reconstructing sub-voxel distance is to trilinearly interpolate in-between voxels [7]. Often times this reconstruction step is embedded in ray-casting procedures at rendering time. While this works for some applications, there is no guarantee on the level of accuracy. From CDFR, we extract a distance contour with a fulfillment of an arbitrarily high accuracy requirement. We store the extracted distance contours as point-based models [6][20], so that we can render the contours at high interactive rates with splatting [4][12][18].

The extracting procedure works as following. Given a requested interior thickness, t ($t > 0$), we traverse these voxels with a distance value in the range:

$$\left(t - \frac{\sqrt{3}}{2} voxel\ size, t + \frac{\sqrt{3}}{2} voxel\ size \right) \quad (2)$$

The requested iso-contour will pass through the span of these voxels. Unlike marching cubes [10], We do not use conditions like minimal thickness $\leq t$ and maximal thickness $\geq t$, because the underlying assumption of having a linear function is not true in our case. There could be cases where the 8 corner voxels are just surrounding the maximal thickness point in the model, and none of the 8 voxels exactly captures that maximum.

After identifying the relevant voxels, we then subdivide the voxels into sub-voxels, or points [2]. We only extract the sub-voxels that are close to the desired surface into a point-based iso-distance contour represented the surface. In order to support the error tolerance, E , picked by users, the size of each sub-voxel must be:

$$\frac{\sqrt{3}}{2} sub\ voxel\ size < E \quad (3)$$

For each sub-voxel, or point, we compute the signed true distance for all the CDD tuples resident on each of the 8 cornering voxels. The points that have the minimal positive distance value within the range $[t - E/2, t + E/2]$ are extracted into the point-based iso-distance contour.

3.6 High Quality Gradients

Besides using the distance contour for analysis, visualizations of the distance contours are also highly desired in applications. For point-based models, having high quality normal information on each point is essential for high image quality.

Our CDFR offers an additional advantage in this perspective. When extracting the distance contour from the base triangle of

each sub-voxel, the normal of this point is computed. If this point is of case C_1 to its base triangle, then the normal of the base triangle is this point's true gradient. If the point is one of the cases C_2 or C_3 , the gradient is the vector V in Fig. 5. For instance, in a C_2 case, the 3D point, P , first gets projected onto the closest edge. The gradient is the vector connecting P and its projection. In C_3 cases, the gradient direction is obtained by connecting P and the closest vertex. Therefore, the normal vectors computed for the whole point-based model is continuous and accurate. High quality per point shading is thus supported.

4. Proof of Sufficiency

To prove the correctness of CDFR, we need a proof of sufficiency. That is, when we need to reconstruct the local distance field in the span of any voxel, all the surface primitives affecting this local area are present on that voxel.

A surface primitive, such as a triangle, affects a local field in 3D space by being the closest surface triangle to at least one position in this local area. Based upon this observation, we devise our proof of sufficiency with a proof by contradiction:

Suppose in the CDFR, R , there exists a local voxel, V , in whose span there exists at least one point, $P(x,y,z)$, whose base triangle, T , is not resident on the voxel, V .

Without loss of generality, we write the distance from P to T as D . All distance fields are continuous functions, although they may not have continuous derivatives. For a point, $P'(x+dx,y+dy,z+dz)$, that is closely neighboring P , the minimal distance from P' to T is bounded by:

$$[D - \sqrt{dx^2 + dy^2 + dz^2}, D + \sqrt{dx^2 + dy^2 + dz^2}] \quad (4)$$

Due to deduction, when P' incrementally moves from P towards V 's center point, it logically follows that the distance from V 's center point to T is bounded by:

$$[D - \int_p^v ds, D + \int_p^v ds] \quad (5)$$

Equation (5) can be rewritten as:

$$[D - \text{dist}(V, P), D + \text{dist}(V, P)] \quad (6)$$

However, the minimum distance to P , which is D , must also be smaller than $\text{min}D + \text{dist}(V, P)$, with $\text{min}D$ denoting the minimum distance of the surface to V . Therefore, the distance of T to V , must be within the following range:

$$[\text{min}D, \text{min}D + 2\text{dist}(V, P)] \quad (7)$$

Since P is in the span of V , the maximum possible distance is $(\sqrt{3}/2) \text{ voxel size}$, the range in (7) is actually a subset of:

$$[\text{min}D, \text{min}D + \sqrt{3}\text{voxelsize}] \quad (8)$$

Contradiction. Since during our distance propagation process, Equation (8) is exactly the range that we maintain on each voxel. Hence, triangle T must be resident on voxel V . The assumed case can not exist. Proof completed.

We do not claim our storage is minimal. We might have kept more CDD tuples on each voxel than necessary. However, enforcing that minimality would introduce more complexity. As long as we use a triangle index in CDD tuples instead of complete description of each triangle, the extra storage cost that we spend is low. We have traded for simplicity in implementation.

5. Results and Analysis

The resolution of CDFR volumes do not affect the accuracy of the distance field. Also, the CDFR construction step is indepen-

dent from the step that reconstructs iso-distance contours. Before we analyze the performance of our approach, we show images of distance contours on a few sample parts to demonstrate the accurate Euclidean distance fields obtained. All point-based models are rendered with splatting [12]. All results are collected on a SGI Octane with a 300MHz processor and 512MB memory. Table 1 provides a full description of the models used for test and analysis of our algorithm. We have also applied our algorithm to a very complicated part, Engine Cylinder Head, for heavy section detection. The details of the Engine Cylinder Head model is described in Section 5.5.

Model	No. Triangles	Bounding Box Size (x,y,z) (inch)	Maximal Thickness(inch)
Cube	12	(5, 5, 5)	2.5
Tetrahedron	4	(1, 1, 1)	0.2
1-Tooth	16	(1, 2, 2)	0.48
6-Star	48	(1, 3, 3.46)	0.49
Connector	242	(6.9, 2.0, 2.9)	0.50
Brevi	1812	(38.1, 34.9, 96.0)	13.00

TABLE 1. Physical Information of Test Models.

5.1 Experimental Data Sets

We first look at some simple cases. Cubes and tetrahedral cells are the simplest. They are convex and symmetric. With a true Euclidean distance field, the thickness contours of different values are in the exact shape of the outer surfaces, including the sharp edges and corners (Fig. 7).




Figure 7: A cube and tetrahedron, with the surface mesh shown in semi-transparency. The distance contours (shown in red, per-point shaded) are of thickness (a) 0.6 inch and (b) 0.1 inch.

Concavities cause additional complexity in distance fields. Two concave examples, a one-ended tooth and a six-pointed star are shown in Fig. 8. For the one-ended tooth, we choose a small thickness value to extract a contour close to the surface, while for the six-pointed star, a larger thickness is chosen. The evolving effects in Euclidean distance fields are interesting, with corners being smoothed out on the interior distance contours in both Fig. 8a and Fig. 8b. Small thickness contours closer to the surfaces retain more detail of the surface shape. The shape of the deeper contours manifest more global features of the shape (Fig. 8).

All four models have a CDFR at a low resolution of $32 \times 32 \times 32$. To reconstruct the thickness contours in Fig. 7 and 8, we set the accuracy to 1/500 of the longest dimension of the model. For all 4 models, there are less than 400K points in the point-based contour. We obtained 3 frames/second rendering rates of the accurate distance contours with per point shading.




Figure 8: Two concave examples, a 6-pointed star and a one-ended tooth. (a) For small thickness values (0.2 inch), the distance field retains most corners and edges, with little smoothing. (b) As the thickness increases (0.35 inch), the distance field evolves into the model, showing more smoothing.

We note that, for conventional volumes, interactive rendering is still an overwhelming challenge on most computing platforms for a volume of size 500^3 [17].

5.2 Real-World Models

Surface graphics based CAD/CAM systems often spend hours to perform jobs involving the interior of real world designs. Volume techniques have been considered. But due to the overwhelming costs incurred in high resolution volumes, the affordable accuracy is very limited with current computing systems when applied to design for manufacturing purposes [16]. The application of volume graphics to CAD is limited. We tested our approach on two industry production models, ‘connector’ and ‘brevi’, for which both accuracy and interactive frame rates are highly sought after.




Figure 9: Results of ‘connector’. The surface mesh is shown in semi-transparency, and the per-point shaded distance contours are at thicknesses (a) 0.2 inch and (b) 0.35 inch.

All contours in Fig. 9 and 10 are extracted to an accuracy of $1/1024$ of the length of each part. In Fig. 9, the thickness contours within the ‘connector’ part show crisp edges and corners, while at the same time, retaining topological features in the surface geometry at different levels of scale.

In Fig. 10, we present the thickness contours within the ‘brevi’ part at thicknesses of 10 and 4 inches. Fig. 10a demonstrates that even at the core of a part, sharp corners in the distance fields still exist at a variety of scales. The hole on the lower right corner of Fig. 10b is a discontinuous point in the distance field, and is one of the causes of failures during manufacturing. We are pleased that our algorithm captures such artifacts. (See color plate for greater details).




Figure 10: Sample images of ‘brevi’. With the per-point shaded contours at thicknesses (a) 10 inches and (b) 4 inches.

The contours in Fig. 9b and Fig. 10a are small, with 400K and 120K points, respectively. Around 2 to 5 frames/second have been recorded. However, Fig. 9a and Fig. 10b both have relatively large contours, with 1.5 to 2 million points, and only about 0.5 frames/second rates are obtained.

5.3 CDFR Size and Construction Time

The detailed storage structure on each voxel is shown in Fig. 11. Each voxel contains a 1-byte flag, $vCnt$. Empty voxels, i.e. outside voxels, have $vCnt$ set to zero. Surface voxels have a $vCnt$ value in the range between 1 and 127, denoting the number of triangles present on this voxel. Voxels entirely in the interior are distinguished by having a $vCnt$ larger than 127. The value ($vCnt - 127$) is the count of triangles present. For very complicated models at very low CDFR resolution, $vCnt$ may overflow. In that case, a larger number is needed for $vCnt$, or an adaptive subdivision scheme is required. By subdividing each voxel into an even number of sub-voxels, we also break up areas having a central curvature point, such as a sphere. In total, $(4 \cdot n + 5)$ bytes are needed per non-empty voxel, with n being the number of triangles present on that voxel.

```

CDD_voxel
{
  unsigned char vCnt; // in/suf/out, and counter of triangles
  float cur_distance; // current minimal distance on this voxel
  int triangles[triangle count]; // dynamic array of triangle indices
}

```

Figure 11: The storage on each voxel in a constructed CDFR.

To analyze storage costs of a CDFR, we constructed CDFR volumes for each model at different resolutions. Cube, Tetrahedron, 1-Tooth (one-ended tooth) and 6-Star (six-pointed star) are simple models, for which we built CDFR volumes at 32 and 64 resolutions and show the results in Table 2 and Table 3, respectively. The ‘Resolution’ column shows the actual dimension of the CDFR volume. The construction time and final sizes of the CDFRs are shown under ‘Timing’ and ‘Size’. ‘In/Sur/Out’ indicates the distribution of interior, surface and outside voxels in the CDFR. Finally, the average number of triangles on surface and interior voxels are presented in the last two columns.

Model	Resolution	Timing (sec)	Size (KB)	In/Sur/Out (K voxels)	Avg No. Tri/Sur	Avg No. Tri/In
Cube	32x32x32	2.50	213	9.2/5.7/17.8	1.29	2.47
Tetra	32x32x32	0.59	64	1.0/2.3/29.5	1.14	2.19

TABLE 2. CDFR facts at 32 resolution.

Model	Resolution	Timing (sec)	Size (KB)	In/Sur/Out (K voxels)	Avg No. Tri /Sur	Avg No. Tri/In
1-Tooth	20x32x32	0.79	94	2.6/3.7/14.2	1.45	2.58
6-Star	15x29x32	0.99	52	0.5/2.1/11.3	2.17	4.30

TABLE 2. CDFR facts at 32 resolution.

We notice that, in Table 2, the average number of triangles on surface voxels is mostly less than or around 2. On the interior, the average number of triangles per voxel is about two times that. Solid and compact models, such as a cube, take more space because it has more non-empty voxels. Overall, considering the 32K voxels in the volume, CDFR is fairly efficient in storage for all the models.

Model	Resolution	Timing (sec)	Size (KB)	In/Sur/Out (K voxels)	Avg Tri/Sur	Avg Tri/In
Cube	64,64,64	34.24	2,438	148.9/29.4/83.8	1.13	2.23
Tetra	64,64,64	3.59	606	20.8/10.6/230.7	1.06	2.07
1-Tooth	36,64,64	11.36	966	50.0/18.5/80.0	1.20	2.27
6-Star	28,61,64	3.97	398	13.8/9.8/62.4	1.49	2.88

TABLE 3. CDFR facts at 64 resolution.

As we double the resolution (Table 3), the average number of triangles per voxel decreases and approaches the lower limit of the number of triangles affecting each voxel, which is 1 triangle/voxel. Doubling the resolution, the number of interior voxels increases faster than the number of surface voxels. Hence, both the size and construction time of CDFR increase by more than a factor of 8.

Model	Resolution	Timing (sec)	Size (KB)	In/Sur/Out (K voxels)	Avg Tri/Sur	Avg Tri/In
Connector	128,43,58	8.05	970	26.4/20.8/272.0	1.95	2.91
Connector	256,81,112	82.72	7,548	30.7/91.5/1,924.0	1.43	2.53
Brevi	56,52,128	51.53	3,459	100.4/50.5/221.7	3.01	4.66
Brevi	106,98,256	448.2	25,260	1,075.4/205.2/1,379	1.96	3.69

TABLE 4. CDFR facts for ‘connector’ and ‘brevi’.

For the two industry parts, ‘connector’ and ‘brevi’, we use higher resolutions, since there are more surface details. We chose 128 and 256 resolutions. A 128 resolution seems high enough to limit the average number of triangles on each voxel for ‘connector’. However, for ‘brevi’, a resolution of 256 is necessary to cut down the (number of triangles)/voxel. The increase in storage size and construction time by 8 to 10 times is still true, as the volume resolution is doubled.

5.4 Point-based Iso-Distance Contour Extraction Time

Using a higher resolution CDFR has no effect on the accuracy in the final representation. However, it causes an dramatic cubic increase in storage size and construction time. The main motivation in using higher CDFR resolutions is to have more efficient distance contour extraction, due to less CDD tuples in each voxel and more accurate localization of voxel spans that may contain the requested iso-contour. We tested all of the models on the time to extract iso-distance point-based contours from different CDFR resolutions. In Table 5, we show these timings, in seconds, to extract a contour from both 32-res and 64-res CDFRs of the four simpler models. The ‘Thickness’ column shows the iso-distance value chosen. For each CDFR resolution, we collect timings for 3 levels of accuracy and organize the results in regard to which conventional volume resolution the extracted contours would correspond to in accuracy. We list the three corresponding conventional volume resolutions, 128, 192, 384, under both ‘32 Res CDFR’ and ‘64 Res CDFR’.

Model	Thickness (inch)	Timing (32-Res CDFR) (sec)			Timing (64-Res CDFR) (sec)		
		128 res	192 res	384 res	128 res	192 res	384 res
Cube	0.6	1.13	3.96	29.66	0.84	1.95	14.48
Tetra	0.1	0.22	0.55	4.02	0.15	0.29	1.61
1-Tooth	0.2	0.73	2.15	14.79	0.43	1.61	7.75
6-Star	0.35	0.54	1.53	12.23	0.22	0.61	4.30

TABLE 5. Time to extract iso-distance contour from simple models. Two CDFR resolutions have been tested with 3 levels of accuracy.

It is obvious, that finer accuracy results in longer extraction time. Using a higher resolution CDFR, this extraction time significantly drops, as shown by the 32-res versus 64-res CDFRs in Table 5. In Table 6, we show the speedups by using higher resolution CDFR’s. For any CDFR resolution, the number of points extracted for a certain distance contour stays roughly the same. Using a higher resolution CDFR, the average amount of time spent to extract a point on the contour is much less.

Model	Thickness (inch)	time/point (μs /point) 32-Res CDFR			time/point (μs /point) 64-Res CDFR		
		128 res	192 res	384 res	128 res	192 res	384 res
Cube	0.6	11.74	18.49	34.36	6.42	6.61	12.37
Tetra	0.1	11.98	12.91	19.63	4.72	4.99	6.92
1-Tooth	0.2	15.17	19.80	27.10	5.09	10.93	13.56
6-Star	0.35	34.03	43.83	78.36	10.37	12.80	20.41

TABLE 6. Comparing per-point extraction time (μs /voxel) at different CDFR resolutions.

On the two industrial parts, we chose a thickness of 0.3 inches for ‘connector’, and 4 inches for ‘brevi’. We used 128-res and 256 res CDFRs and the 3 levels of accuracy correspond to conventional volume resolutions of 512, 768 and 1024. The results in Table 7 confirm our findings from Table 5 and Table 6. Using higher resolution CDFR effectively cuts down iso-contour extraction time.

Model	128-Res CDFR			256-Res CDFR		
	512 res	768 res	1024 res	512 res	768 res	1024 res
Connector Timing (sec)	2.43	7.12	16.92	1.03	2.69	5.67
Connector: time/point (μs /point)	24.08	30.92	41.45	9.88	11.00	13.38
Brevi Timing (sec)	23.66	74.65	174.79	9.25	28.29	64.46
Brevi: time/point (μs /point)	32.51	45.27	59.72	12.23	16.63	21.36

TABLE 7. ‘Connector’ and ‘brevi’ iso-contour extraction timing (sec) and per-point extraction time (μs /point).

5.5 A Cutting Edge Application

We also applied our algorithm to a challenging design part, an engine cylinder head with 135,429 surface triangles. It is modified from a real design model for our research purposes. When built, it weighs 40Kg’s, and measures $266 \times 480 \times 157$ cm in size. Typically, the maximal thickness of engine cylinders is only about 9 to 10 mm. In engines blocks, heavy sections are an important source of physical failure. It is highly desired to be able to detect heavy sections at an accuracy higher than 0.15mm. Unfortunately, there has been no reliable and affordable way to perform such detection in the early design stages. Extensive resources have to be spent in the dreadfully long and cyclic design, prototyping and verification process. For conventional volume techniques to handle this task, one needs to build a volume having at least $1774 \times 3200 \times 1046$ voxels in floating-point numbers, amounting to 24GB. Even so, there is still no guarantee of an accurate Euclidean distance field, due to the binary surface volumes in voxelization, inaccurate distance map and linear interpolations used in reconstruction. Further-

more, constructing or rendering of data sets at such a size is overwhelmingly difficult.




Figure 12: The thickness contour of a cylinder model, at 8.5mm thickness and 0.137mm error tolerance.

With the CDFR representation, at an CDFR resolution of $142 \times 250 \times 87$, we are able to construct a CDFR of size 37MB in 30 minutes. From this CDFR, we extract the distance contour as a point-based representation with 0.137 mm accuracy, corresponding to a $1988 \times 3500 \times 1218$ conventional volume resolution. For a thickness contour of 8.5mm, the extraction stage takes about 632 seconds, and the resulting point-based model has 450K points and is rendered interactively at 2 frames/sec, with the surface triangles being sorted and rendered semi-transparently at the same time (see Fig. 12). This whole process is done on our 512MB memory, 300MHz processor SGI Octane.

This result is strong, in that it brings an unprecedented accuracy to commonly available computing platforms. It is a tool that provides a guarantee of correctness as well as an interactive capability for visual investigation of highly complex models. This specific model is courtesy of Ford Motor Company.

6. Discussions and Future Work

In this paper we proposed CDFR as an accurate description of distance fields resulting from a surface shape. CDFR volumes are not band-limited and exactly capture surface details in the 3D volume. Specifically, we demonstrate the advantages offered by CDFR on triangle surface meshes. High accuracy and visual quality is achievable with a point-based iso-distance contour extraction. In addition to providing a proof of correctness and visualization results on real data sets, performance and storage issues have been discussed. With CDFR, the initial resolution of the volume does not affect the accuracy in subsequent iso-surface extraction. The initial resolution, however, provides a trade off between storage and speed. CDFR can be rendered directly with ray-casting using an adapted method of ray-object intersection check. However, we propose in this paper to use point-based approaches to rendering distance contours from CDFR with genuine per-point shading.

One future extension is to use hierarchical data structures, such as ADF, to organize CDFR. During extraction, it is the complexity on each voxel that affects the overall performance. Using ADF to spatially manage CDFR would expedite the reconstruction process and reduce CDFR storage costs. Furthermore, hierarchically organized CDFR provides a general framework to compute the Euclidean distance from a 3D point to polygonal mesh. While the applications discussed in this paper strictly deal with the interior of a shape, this framework can also be applied to the exterior of complex models and scenes. All distance values are computed exactly at a low constant time. No complicated approximation schemes [8] based on progressive meshes etc. is necessary. Appli-

cations that may benefit from this framework include path planning with a guaranteed accuracy in highly complicated scenes, such as assembly of airplanes, large scale machinery. General graphics applications such as hypertexture renderings can also make use of this novel distance field representation.

7. Acknowledgment

This project has been funded by Ford Motor Company, Vis-
teon, Inc., and NSF CAREER Award #9876022. The connector, brevi and engine models are provided by Ford Motor. We would like to thank the reviewers of this paper for their insightful comments, too.

Reference

- [1] D. Breen, S. Mauch and R. Whitaker, "3D scan conversion of CSG models into distance volumes", Proc. 1998 IEEE Symposium on Volume Visualization, pp. 7-14, 1998.
- [2] H. Cline, W. Lorensen, S. Ludke, C. Crawford, B. Teeter, "Two algorithms for the three-dimensional reconstruction of tomograms", Medical Physics, 15(3), May/June, 1988, pp. 320-327.
- [3] D. Cohen-Or, D. Levin, and A. Solomovici, "Three-dimensional distance field metamorphosis", ACM Transactions on Graphics, Vol. 17, No. 2, pp. 116-141, April 1998.
- [4] Crawfis, R., Max, N., Texture Splats for 3D Scalar and Vector Field Visualization, IEEE Visualization'93 Proceedings, October, 1993, 261-267.
- [5] F. Dachille, A. Kaufman, "Incremental triangle voxelization", Proc. of Graphics Interface 2000, pp. 205-212, May 2000.
- [6] H. Pfister, J. Barr, M. Zwicker, M. Gross, "Surfel: surface elements as rendering primitives", Proc. of Siggraph 2000, New Orleans, 2000.
- [7] S. Frisken, R. Perry, A. Rockwood, T. Jones, Adaptively Sampled Distance Fields: A General Representation of Shape for Computer Graphics, Proc. of SIGGRAPH'2000, New Orleans, LA, July, 2000.
- [8] A. Gueziec, "Meshsweeper: dynamic point-to-polygonal mesh distance and applications", IEEE Transactions on visualization and computer graphics, vol. 7, no. 1, pp. 47 - 61, Jan - Mar, 2001.
- [9] S. Gibson, Using Distance Maps for smooth surface representation in sampled volumes, Proc. 1998 IEEE Volume Visualization Symposium, pp. 23-30, 1998.
- [10] W. Lorensen, H. Cline, Marching Cubes: a high resolution 3D surface construction algorithm, Computer Graphics (SIGGRAPH 87 Proceedings), 1987, 163-169.
- [11] J. Huang, R. Yagel, V. Fillipov, Y. Kurzion, An Accurate Method to Voxelize Polygonal Meshes, Proc. of IEEE/ACM Symposium on Volume Visualization, October, 1998, Chapel Hill, NC.
- [12] J. Huang, K. Mueller, N. Shareef, R. Crawfis, "FastSplats: optimized splatting on rectilinear grids", Proc. of IEEE Conference on Visualization, October, 2000, Salt Lake City, Utah
- [13] A. Kaufman, "An algorithm for 3D scan-conversion of polygons", Proc. of Eurographics '87, pp. 197-208, North Holland, August, 1987.
- [14] A. Kaufman, "An algorithm for 3D scan-conversion of parametric curves, surfaces, and volumes", Proc. of SIGGRAPH'87, pp. 171-179, July, 1987.
- [15] L. Kobbelt, M. Botsch, U. Schwanecke, H. Seidel, "Feature sensitive surface extraction from volume data", Proc. of SIGGRAPH'2001, July, 2001.
- [16] S. Lu, A. Rebello, R. Miller, G. Kinzel, R. Yagel, "A simple visualization tool to support concurrent engineering design", Journal of Computer-Aided Design, Vol. 29, No. 10, pp. 727-735.
- [17] Meissner, M., Huang, J., Bartz, D., Mueller, K., Crawfis, R., A Practical Evaluation of Popular Volume Rendering Algorithms, Proc. of Symposium of Volume Graphics 2000, Salt Lake City, Utah.
- [18] Muller, K., Shareef, N., Huang, J., Crawfis, R., High-Quality Splatting on Rectilinear Grids with Efficient Culling of Occluded Voxels, IEEE Transaction on Visualization and Computer Graphics, Vol. 5, No. 2, pp 116-135, 1999.
- [19] S. Parker, M. Parker, Y. Livnat, P. Sloan, C. Hansen, and P. Shirley, "Interactive ray tracing for volume visualization" IEEE Transactions On Visualization and Computer Graphics, Vol. 5 (3), pp. 238-250, 1999.
- [20] S. Rusinkiewicz, M., Levoy, "QSplat: a multi-resolution point rendering system for large meshes", Proc. of Siggraph 2000, New Orleans, 2000.
- [21] M. Sramek, A. Kaufman, "Alias-free voxelization of geometric objects", IEEE Transactions On Visualization and Computer Graphics, Vol. 5, (3), pp. 250-267, 1999.
- [22] R. Yagel, S. Lu, A. Rubello, R. Miller, "Volume-based reasoning and visualization of dicastability" In Proc. IEEE Visualization '95, pp. 359-362, 1995.




Figure 7. A cube and tetrahedron, with the surface mesh shown in semi-transparency. The distance contours (shown in red, per-point shaded) are extracted with an error tolerance of 1/500 of the longest dimension of the models. The thickness in (a) is 0.6 inch and 0.1 inch in (b).




Figure 8. Two concave examples, a 6-pointed star and a one-end tooth. (a) At small thickness (0.2 inch), the distance field retains most corners and edges, with little smoothing. (b) As thickness increases (0.35 inch), the distance field evolves into the core of the model, showing more smoothing and losing more surface details but still preserving the overall shape.




Figure 9. Results of 'connector'. The surface mesh is shown in semi-transparency, and the per-point shaded distance contours are at thicknesses (a) 0.2 inch and (b) 0.35 inch.




Figure 10. The sample images of 'brevi'. With the per-point shaded contours at thickness (a) 10 inches and (b) 4 inches.




Figure 12. The thickness contour of a modified engine cylinder head model, at 8.5mm thickness and 0.137mm error tolerance.