

Geometric Data Structures for Computer Graphics *

Gabriel Zachmann & Elmar Langetepe

Informatik II/I

University of Bonn, Germany

email: {zach,langetep}@cs.uni-bonn.de

<http://web.cs.uni-bonn.de/~zach>

<http://web.cs.uni-bonn.de/I/staff/langetepe.html>

13th August 2003

Course Description

In recent years, methods from computational geometry have been widely adopted by the computer graphics community yielding elegant and efficient algorithms. This course aims at familiarizing practitioners in the computer graphics field with a wide range of data structures from computational geometry. It will enable attendees to recognize geometrical problems and select the most suitable data structure when developing computer graphics algorithms.

The course will focus on algorithms and data structures that have proven to be versatile, efficient, fundamental and easy to implement. Thus practitioners and researchers will benefit immediately from this course for their everyday work.

* You can find the slides and the multimedia material for this tutorial at <http://cg.cs.uni-bonn.de/course/>

Prerequisites

Participants of this course should be familiar with the basic principles of computer graphics and the type of problems in the area. Familiarity with computer programming is not required.

The intended audience are practitioners working in 3D computer graphics (VR, CAD/CAM, entertainment, animation, etc.) and students from both computer graphics and computational geometry, possibly working on a master or PhD thesis.

Syllabus

Our goal is to familiarize practitioners and researchers in computer graphics with some very versatile and ubiquitous geometric data structures, enable them to readily recognize geometric problems during their work, modify the algorithms to their needs, and hopefully make them curious about further powerful treasures to be discovered in the area of computational geometry.

In order to achieve these goals in an engaging yet sound manner, the general concept throughout the course is to present each geometric data structure in the following way: first, the data structure will be defined and described in detail; then, some of its fundamental properties will be highlighted; after that, one or more computational geometry algorithms based on the data structure will be presented; and finally, a number of recent, representative and practically relevant algorithms from computer graphics will be described in detail, showing the utilization of the data structure in a creative and enlightening way.

We have arranged the topics in roughly increasing degree of difficulty. The hierarchical data structures are ordered by increasing flexibility, while the non-hierarchical topics build on each other. Finally, the last topic presents a generic technique for dynamizing geometric data structures.

Detailed informations about the topics discussed and the corresponding presenters are given in the schedule below.

- Component 1. *Introduction*
Introduction – Zachmann/Langetepe.
- Component 2. *Quadtree/Octree*
Construction, complexity, balancing, navigation – Langetepe.
Terrain visualization, iso-surface generation – Zachmann.
- Component 3. *Bounding volume hierarchies*
Definition, construction, hierarchical collision detection – Zachmann.
- Component 4. *Voronoi diagrams / Delaunay triangulations*
Definition, construction, generalizations, applications – Langetepe.
NURBS tessellation, texture synthesis – Zachmann.
- Component 5. *Distance fields*
Definition, modelling, morphing – Zachmann.
- Component 6. *Dynamization of geometric data structures*
Definition, amortized Insert and Delete – Langetepe.

Contents

1	Introduction	5
2	Quadtrees and K - d -Trees	5
2.1	Quadtrees and Octrees	5
2.1.1	Definition	5
2.1.2	Complexity and Construction	6
2.1.3	Neighbor Finding	6
2.2	K - d -Trees	6
2.3	Height Field Visualization	7
2.4	Isosurface Generation	10
2.5	Ray Shooting	11
2.5.1	3D Octree	12
2.5.2	5D Octree	13
3	BSP Trees	15
3.1	Rendering Without a Z-Buffer	16
3.2	Representing Objects with BSPs	17
3.3	Boolean Operations	17
3.4	Construction Heuristics	19
3.4.1	Convex objects	20
3.4.2	Cost driven heuristic	20
3.4.3	Non-uniform queries	20
3.4.4	Deferred, self-organizing BSPs	21
4	Bounding Volume Hierarchies	22
4.1	Construction of BV Hierarchies	24
4.1.1	Bottom-up	24
4.1.2	Insertion	25
4.1.3	Top-down	25
4.1.4	Construction criteria	25
4.1.5	The criterion for collision detection	26
4.2	Collision Detection	28
5	Voronoi Diagrams	28
5.1	Definitions and Elementary Properties	29
5.1.1	Voronoi Diagram	29
5.1.2	Delaunay Triangulation	30
5.2	Computation	31
5.3	Generalization of the Voronoi Diagram	32
5.3.1	Voronoi Diagram and Delaunay Triangulation in 3D	32
5.3.2	Constrained Voronoi diagrams	34
5.3.3	Other Types of Generalizations	34
5.4	Applications of the Voronoi Diagram	34
5.4.1	Nearest Neighbor or Post Office Problem	34
5.4.2	Motion planning	35
5.4.3	Other Applications of the Voronoi Diagram in 2D	36
5.5	Texture Synthesis	36
5.6	Shape Matching	38

6	Distance Fields	39
6.1	Computation and representation of DFs	39
6.1.1	Propagation method	41
6.1.2	Projection of distance functions	42
6.2	Applications of DFs	43
6.2.1	Morphing	43
6.2.2	Modeling	44
7	Dynamization of Geometric Data Structures	44
7.1	Model of the Dynamization	45
7.2	Amortized Insert and Delete	46
7.2.1	Amortized Insert: Binary Structure	46
7.2.2	Amortized Delete: Occasional Reconstruction	47
7.2.3	Amortized Insert <i>and</i> Amortized Delete	48
7.3	Worst-Case sensitive Insert and Delete	48
7.4	A Simple Example	49

1 Introduction

In recent years, methods from computational geometry have been widely adopted by the computer graphics community. Many solutions draw their elegance and efficiency from the mutually enriching combination of such geometrical data structures with computer graphics algorithms.

This course imparts a working knowledge of a number of essential geometric data structures and explains their elegant use in several representative, diverse, and current areas of research in computer graphics such as terrain visualization, texture synthesis, modelling, and tessellation. Our selection of data structures and algorithms consists of well-known concepts, which are both powerful *and* easy to implement, ranging from quadtrees to BSP trees and bounding volume trees. We do not try to provide an exhaustive survey of the topics touched upon here — this would be far beyond the scope of this course. Neither do we try to present leading edge algorithms because we feel that for practical purposes a good trade-off between simplicity and efficiency is important.

Participants will learn to recognize geometric problems and develop a thorough understanding of suitable algorithms. As a prerequisite, they should already be familiar with the basic principles of computer graphics and the type of problems in the area.

The general concept throughout the course is to present each geometric data structure in the following way: first, the data structure will be defined and described in detail; then, some of its fundamental properties will be highlighted; after that, one or more computational geometry algorithms based on the data structure will be presented; and finally, a number of recent, representative and practically relevant algorithms from computer graphics will be described in detail, showing the utilization of the data structure in a creative and enlightening way.

2 Quadtrees and K - d -Trees

Within this section we will present some fundamental geometric data structures.

In Section 2.1, we introduce the quadtree structure. Its definition and complexity, the recursive construction scheme and a standard application are presented. Quadtrees and octrees have applications in mesh generation as shown in Section 2.3, 2.4, 2.5.

A natural generalization of the one-dimensional search tree to k dimensions is shown in Section 2.2. The k - d -tree is efficient for axis-parallel rectangular range queries.

The quadtree description was adapted from de Berg et al.²³ and the k - d -tree introduction was taken from Klein.⁵⁵

2.1 Quadtrees and Octrees

2.1.1 Definition

A *quadtree* is a rooted tree so that every internal node has four children. Every node in the tree corresponds to a square. If a node v has children, their corresponding squares are the four quadrants, as shown in Figure 1.




Figure 1: An example of a quadtree.

Quadtrees can store many kinds of data. We will describe the variant that stores a set of points and suggest a recursive definition. A simple recursive splitting of squares is continued until there is only one point in a square. Let P be a set of points.

The definition of a quadtree for a set of points in a square $Q = [x1_Q : x2_Q] \times [y1_Q : y2_Q]$ is as follows:

- If $|P| \leq 1$ then the quadtree is a single leaf where Q and P are stored.
- Otherwise let Q_{NE} , Q_{NW} , Q_{SW} and Q_{SE} denote the four quadrants. Let $x_{mid} := (x1_Q + x2_Q)/2$ and $y_{mid} := (y1_Q + y2_Q)/2$, and define

$$\begin{aligned} P_{NE} &:= \{p \in P : p_x > x_{mid} \wedge p_y > y_{mid}\}, \\ P_{NW} &:= \{p \in P : p_x \leq x_{mid} \wedge p_y > y_{mid}\}, \\ P_{SW} &:= \{p \in P : p_x \leq x_{mid} \wedge p_y \leq y_{mid}\}, \\ P_{SE} &:= \{p \in P : p_x > x_{mid} \wedge p_y \leq y_{mid}\}. \end{aligned}$$

The quadtree consists of a root node v , Q is stored at v . In the following, let $Q(v)$ denote the square stored at v . Furthermore v has four children: The X -child is the root of the quadtree of the set P_X , where X is an element of the set $\{NE, NW, SW, SE\}$.

2.1.2 Complexity and Construction

The recursive definition implies a recursive construction algorithm. Only the starting square has to be chosen adequately. If the split operation cannot be performed well, the quadtree is unbalanced. Despite this effect, the depth of the tree is related to the distance between the points.

Theorem 1 *The depth of a quadtree for a set P of points in the plane is at most $\log(s/c) + \frac{3}{2}$, where c is the smallest distance between any two points in P and s is the side length of the initial square.*

The cost of the recursive construction and the complexity of the quadtree depends on the depth of the tree.

Theorem 2 *A quadtree of depth d , which stores a set of n points, has $O((d+1)n)$ nodes and can be constructed in $O((d+1)n)$ time.*

Due to the degree 4 of internal nodes, the total number of leaves is one plus three times the number of internal nodes. Hence it suffices to bound the number of internal nodes.

Any internal node v has one or more points inside $Q(v)$. The squares of the node of a single depth cover the initial square. Therefore, at every depth we have at most n internal nodes that gives the node bound.

The most time-consuming task in one step of the recursive approach is the distribution of the points. The amount of time spent is only linear in the number of points and the $O((d+1)n)$ time bound holds.

The 3D equivalent of quadtrees are *octrees*. The *quadtree* construction can be easily extended to *octrees* in 3D. The internal nodes of *octrees* have eight sons and the sons correspond to boxes instead of squares.

2.1.3 Neighbor Finding

A simple application of the quadtree of a point set is *neighbor finding*, that is, given a node v and a direction, north, east, south or west, find a node v' so that $Q(v)$ is adjacent to $Q(v')$. Normally, v is a leaf and v' should be a leaf as well. The task is equivalent to finding an adjacent square of a given square in the quadtree subdivision.

Obviously, one square may have many such neighbors, as shown in Figure 2.




Figure 2: The square q has many west neighbors.

For convenience, we extend the neighbor search. The given node can also be internal, that is, v and v' should be adjacent corresponding to the given direction *and* should also have the same depth. If there is no such node, we want to find the deepest node whose square is adjacent.

The algorithm works as follows. Suppose we want to find the north neighbor of v . If v happens to be the *SE*- or *SW*-child of its parent, then its north neighbor is easy to find, it is the *NE*- or *NW*-child of its parent, respectively. If v itself is the *NE*- or *NW*-child of its parent, then we proceed as follows. Recursively find the north neighbor of μ of the parent of v . If μ is an internal node, then the north neighbor of v is a child of μ ; if μ is a leaf, the north neighbor we seek for is μ itself.

This simple procedure runs in time $O(d+1)$.

Theorem 3 *Let T be quadtree of depth d . The neighbor of a given node v in T a given direction, as defined above, can be found in $O(d+1)$ time.*

Furthermore, there is also a simple procedure that constructs a balanced quadtree out of a given quadtree T . This can be done in time $O(d+1)m$ and $O(m)$ space if T has m nodes. For details see Berg et al.²³

Similar results hold for *octrees* as well.

2.2 K - d -Trees

The k - d -tree is a natural generalization of the one-dimensional search tree.

Let D be a set of n points in \mathbb{R}^k . For convenience let $k = 2$ and let us assume that all X - and Y -coordinates are different. First, we search for a

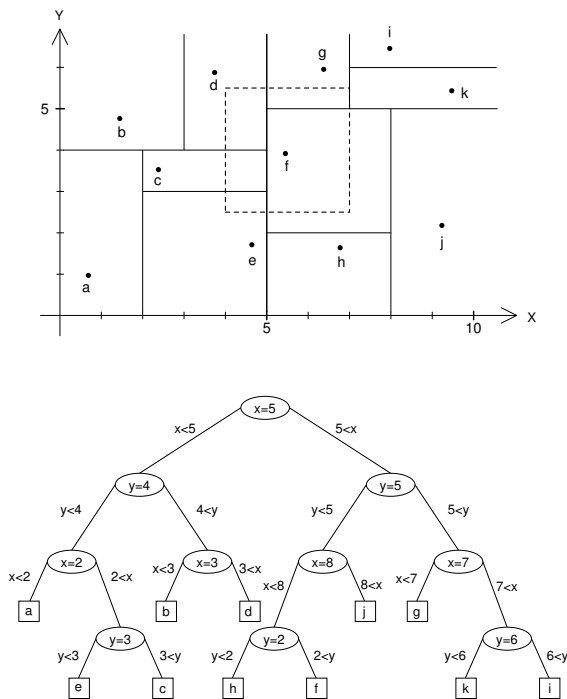


Figure 3: A k - d -tree for $k = 2$ and a rectangular range query. The nodes correspond to split lines.

split-value s of the X -coordinates. Then we split D by the *split-line* $X = s$ into subsets:

$$D_{<s} = \{(x, y) \in D; x < s\} = D \cap \{X < s\}$$

$$D_{>s} = \{(x, y) \in D; x > s\} = D \cap \{X > s\}.$$

For both sets we proceed with the Y -coordinate and *split-lines* $Y = t_1$ and $Y = t_2$. We repeat the process recursively with the constructed subsets. Thus, we obtain a binary tree, namely the 2 - d -tree of the point set D , as shown in Figure 3. Each internal node of the tree corresponds to a split-line. For every node v of the 2 - d -tree we define the rectangle $R(v)$, which is the intersection of halfplanes corresponding to the path from the root to v . For the root r , $R(r)$ is the plane itself; for the sons of r , say *left* and *right*, we produce to halfplanes $R(\textit{left})$ and $R(\textit{right})$ and so on. The set of rectangles $\{R(l) : l \text{ is a leaf}\}$ gives a partition of the plane into rectangles. Every $R(l)$ has exactly one point of D inside.

This structure supports range queries of axis-parallel rectangles. For example, if Q is an axis-parallel rectangle, the set of sites $v \in D$ with $v \in Q$

can be computed efficiently. We simply have to compute all nodes v with:

$$R(v) \cap Q \neq \emptyset.$$

Additionally we have to test whether the points inside the subtree of v are inside Q .

The efficiency of the k - d -tree with respect to range queries depends on the depth of the tree. A balanced k - d -tree can be easily constructed. We sort the X - and Y -coordinates. With this order we recursively split the set into subsets of equal size in time $O(\log n)$. The construction runs in time $O(n \log n)$. Altogether the following theorem holds:

Theorem 4 *A balanced k - d -tree for n points in the plane can be constructed in $O(n \log n)$ and needs $O(n)$ space. A range query with an axis-parallel rectangle can be answered in time $O(\sqrt{n} + a)$, where a denotes the size of the answer.*

2.3 Height Field Visualization

A special area in 3D visualization is the rendering of large terrains, or, more generally, of height fields. A height field is usually given as a uniformly-gridded square array $h : [0, N - 1]^2 \rightarrow \mathbb{R}$, $N \in \mathbb{I}$, of height values, where N is typically in the order of 16,384 or more (see Figure 4). In practice, such a raw height field is often stored in some image file format, such as GIF. A regular grid is, for instance, one of the standard forms in which the US Geological Survey publishes their data, known as the Digital Elevation Model (DEM).³²

Alternatively, height fields can be stored as *triangular irregular networks* (TINs) (see Figure 5). They can adapt much better to the detail and features (or lack thereof) in the height field, so they can approximate any surface at any desired level of accuracy with fewer polygons than any other representation.⁶⁴ However, due to their much more complex structure, TINs do not lend themselves as well as more regular representations to interactive visualization.

The problem in terrain visualization is a low viewpoint directed at the horizon: then, there are a few parts of the terrain that are very close, while the majority of the visible terrain is at a larger distance. Close parts of the terrain should be rendered with high detail, while distant parts should be rendered

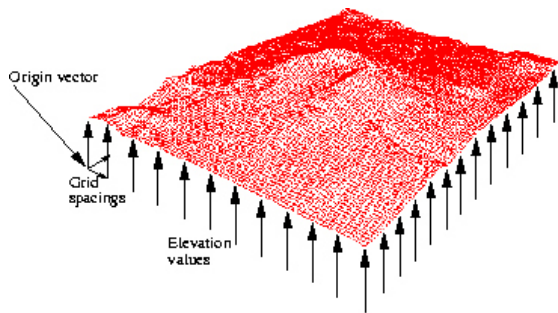


Figure 4: A height field approximated by a grid.¹⁵




Figure 5: The same height field approximated by a TIN.




Figure 6: In order to use quadtrees for defining a height field mesh, it should be balanced.




Figure 7: A quadtree defines a recursive subdivision scheme yielding a 4-8 mesh. The dots denote the newly added vertices. Some vertices have degree 4, some 8 (hence the name).

with very little detail in order to maintain a high frame rate.

In order to solve this problem, a data structure is needed that allows to quickly determine the desired level of detail in each part of the terrain. Quadtrees are such a data structure, in particular, since they seem to be a good compromise between the simplicity of non-hierarchical grids and the good adaptivity of TINs. The general idea is to construct a quadtree over the grid, and then traverse this quadtree top-down in order to render it. At each node, we decide whether the detail offered by rendering it is enough, or if we have to go down further.

One problem with quadtrees (and quadrangle-based data structures in general) is that nodes are not quite independent of each other. Assume we have constructed a quadtree over some terrain as

depicted in Figure 6. If we render that as-is, then there will be a gap (a.k.a. *crack*) between the top left square and the fine detail squares inside the top right square. The vertices causing this problem are called *T-vertices*. Triangulating them would help in theory, but in practice this leads to long and thin triangles which have problems on their own. The solution is, of course, to triangulate each node.

Thus, a quadtree offers a recursive subdivision scheme to define a triangulated regular grid (see Figure 7): start with a square subdivided into two right-angle triangles; with each recursion step, subdivide the longest side of all triangles (the hypotenuse) yielding two new right-angle triangles each (hence this scheme is sometimes referred to as "longest edge bisection").⁶⁵ This yields a mesh where all vertices have degree 4 or 8 (except the




Figure 8: The 4-8 subdivision can be generated by two interleaved quadtrees. The solid lines connect siblings that share a common father.

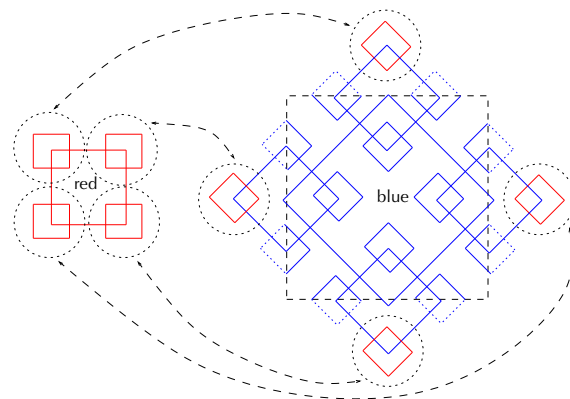


Figure 9: The red quadtree can be stored in the unused “ghost” nodes of the blue quadtree.

border vertices), which is why such a mesh is often called a 4-8 mesh.

This subdivision scheme induces a directed acyclic graph (DAG) on the set of vertices: vertex j is a child of i if it is created by a split of a right angle at vertex i . This will be denoted by an edge (i, j) . Note that almost all vertices are created twice (see Figure 7), so all nodes in the graph have 4 children and 2 parents (except the border vertices).

During rendering, we will choose cells of the subdivision at different levels. Let M^0 be the fully subdivided mesh (which corresponds to the original grid) and M be the current, incompletely subdivided mesh. M corresponds to a subset of the DAG of M^0 . The condition of being crack-free can be reformulated in terms of the DAGs associated with M^0 and M :

M is crack-free \Leftrightarrow

M does not have any T-vertices \Leftrightarrow

$$\forall (i, j) \in M : (i', j) \in M, \text{ where } \text{parent}(j) = \{i, i'\} \quad (1)$$

In other words: you cannot subdivide one triangle alone, you also have to subdivide the one on the other side. During rendering, this means that if you render a vertex, then you also have to render all its ancestors (remember: a vertex has 2 parents).

Rendering such a mesh generates (conceptually) a single, long list of vertices that are then fed into the graphics pipeline as a single triangle strip. The pseudo-code for the algorithm looks like this (simplified):

submesh(i,j)

```

if error(i) < τ then
    return
end if
if Bi outside viewing frustum then
    return
end if
submesh(j, cl)
V += pi
submesh(j, cr)
    
```

where $\text{error}(i)$ is some error measure for vertex i , and B_i is the sphere around vertex i that completely encloses all descendant triangles.

Note that this algorithm can produce the same vertex multiple times consecutively; this is easy to check, of course. In order to produce one strip, the algorithm has to copy older vertices to the current front of the list at places where it makes a “turn”; again, this is easy to detect, and the interested reader is referred to.⁶⁵

One can speed up the culling a bit by noticing that if B_i is completely inside the frustum, then we do not need to test the child vertices any more.

We still need to think about the way we store our terrain subdivision mesh. Eventually, we will want to store it as a single linear array for two reasons:

1. The tree is complete, so it really would not make sense to store it using pointers.
2. We want to map the file that holds the tree into memory as-is (for instance, with Unix’ `mmap` function), so pointers would not work at all.

We should keep in mind, however, that with current architectures, every memory access that can

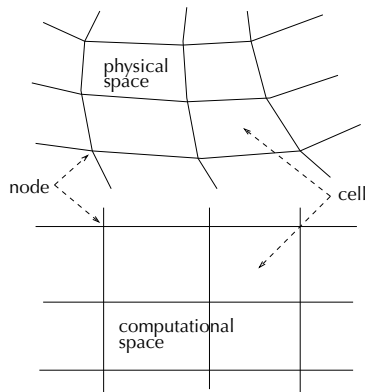


Figure 10: A scalar field is often given in the form of a curvilinear grid. By doing all calculations in computational space, we can usually save a lot of computational effort.

not be satisfied by the cache is extremely expensive (this is even more so with disk accesses, of course).

The simplest way to organize the terrain vertices is a matrix layout. The disadvantage is that there is no cache locality at all across the major index. In order to improve this, people often introduce some kind of blocking, where each block is stored in matrix and all blocks are arranged in matrix order, too. Unfortunately, Lindstrom and Pascucci⁶⁵ report that this is, at least for terrain visualization, worse than the simple matrix layout by a factor 10!

Enter quadtrees. They offer the advantage that vertices on the same level are stored fairly close in memory. The 4-8 subdivision scheme can be viewed as two quadtrees which are interleaved (see Figure 8): we start with the first level of the “red” quadtree that contains just the one vertex in the middle of the grid, which is the one that is generated by the 4-8 subdivision with the first step. Next comes the first level of the “blue” quadtree that contains 4 vertices, which are the vertices generated by the second step of the 4-8 subdivision scheme. This process repeats logically. Note that the blue quadtree is exactly like the red one, except it is rotated by 45°. When you overlay the red and the blue quadtree you get exactly the 4-8 mesh.

Notice that the blue quadtree contains nodes that are outside the terrain grid; we will call these nodes “ghost nodes”. The nice thing about them is that we can store the red quadtree in place of these ghost nodes (see Figure 9). This reduces the number of unused elements in the final linear array down to 33%.




Figure 11: Cells straddling the isosurface are triangulated according to a lookup table. In some cases, several triangulations are possible, which must be resolved by heuristics.

During rendering we need to calculate the indices of the child vertices, given the three vertices of a triangle. It turns out that by cleverly choosing the indices of the top-level vertices this can be done as efficiently as with a matrix layout.

The interested reader can find more about this topic in Lindstrom et al.,⁶⁴ Lindstrom and Pascucci,⁶⁵ Balmelli et al.,⁸ Balmelli et al.,⁷ and many others.

2.4 Isosurface Generation

One technique (among many others) of visualizing a 3-dimensional volume is to extract isosurfaces and render those as a regular polygonal surface. It can be used to extract the surfaces of bones or organs in medical scans, such as MRI or CT.

Assume for the moment that we are given a scalar field $f : \mathbb{R}^3 \rightarrow \mathbb{R}$. Then the task of finding an isosurface would “just” be to find all solutions (i.e., all roots) of the equation $f(\vec{x}) = t$.

Since we live in a discrete world (at least in computer graphics), the scalar field is usually given in the form of a *curvilinear grid*: the vertices of the *cells* are called *nodes*, and we have one scalar and a 3D point stored at each node (see Figure 10). Such a curvilinear grid is usually stored as a 3D array, which can be conceived as a regular 3D grid (here, the cells are often called *voxels*).

The task of finding an isosurface for a given value t in a curvilinear grid amounts to finding all cells of which at least one node (i.e., corner) has a value less than t and one node has a value greater than




Figure 12: Octrees offer a simple way to compute isosurfaces efficiently.

t . Such cells are then triangulated according to a lookup table (see Figure 11). So, a simple algorithm works as follows:⁶⁶ compute the sign for all nodes ($\oplus \triangleq > t$, $\ominus \triangleq < t$); then consider each cell in turn, use the eight signs as an index into the lookup table, and triangulate it (if at all).

Notice that in this algorithm we have only used the 3D array — we have not made use at all of the information exactly *where* in space the nodes are (except when actually producing the triangles). We have, in fact, made a transition from *computational space* (i.e., the curvilinear grid) to *computational space* (i.e., the 3D array). So in the following, we can, without loss of generality, restrict ourselves to consider only regular grids, that is, 3D arrays.

The question is, how can we improve the exhaustive algorithm. One problem is that we must not miss any little part of the isosurface. So we need a data structure that allows us to discard large parts of the volume where the isosurface is guaranteed to *not* be. This calls for octrees.

The idea is to construct a complete octree over the cells of the grid⁹⁶ (for the sake of simplicity, we will assume that the grid's size is a power of two). The leaves point to the lower left node of their associated cell (see Figure 12). Each leaf v stores the minimum v_{\min} and the maximum v_{\max} of the 8 nodes of the cell. Similarly, each inner node of the octree stores the min/max of its 8 children.

Observe that an isosurface intersects the volume associated with a node v (inner or leaf node) if and only if $v_{\min} \leq t \leq v_{\max}$. This already suggests how the algorithm works: start with the root and visit recursively all the children where the condition holds. At the leaves, construct the triangles as usual.

This can be accelerated further by noticing that if the isosurface crosses an edge of a cell, then that




Figure 13: Volume data layout should match the order of traversal of the octree.

edge will be visited exactly four times during the complete procedure. Therefore, when we visit an edge for the first time, we compute the vertex of the isosurface on that edge, and store the edge together with the vertex in a hash table. So whenever we need a vertex on an edge, we first try to look up that edge in the hash table. Our observation also allows us to keep the size of the hash table fairly low: when an edge has been visited for the fourth time, then we know that it cannot be visited any more; therefore, we remove it from the hash table.

2.5 Ray Shooting

Ray shooting is an elementary task that frequently arises in ray tracing, volume visualization, and in games for collision detection or terrain following. The task is, basically, to find the earliest hit of a given ray when following that ray through a scene composed of polygons or other objects.

A simple idea to avoid checking the ray against all objects is to partition the universe into a regular grid (see Figure 14). With each cell we store a list of objects that occupy that cell (at least partially). Then, we just walk along the ray from cell to cell, and check the ray against all those objects that are stored with that cell.

In this scheme (and others), we need a technique called *mailboxes* that prevents us from checking the ray twice against the same object:³⁸ every ray gets a unique ID (we just increment a global variable holding that ID whenever we start with a new ray); during traversal, we store the ray's ID with the object whenever we have performed an intersection test with it. But before doing an intersection test with an object, we look into its mailbox whether or not the current ray's ID is already there; if so, then




Figure 14: Ray shooting can be implemented efficiently with an octree.




Figure 15: The same scenario utilizing an octree.

we know that we have already performed the intersection test in an earlier cell.

In the following, we will present two methods that both utilize octrees to further reduce the number of objects considered.

2.5.1 3D Octree

A canonical way to improve any grid-based method is to construct an octree (see Figure 15). Here, the octree leaves store lists of objects (or, rather, pointers to objects). Since we are dealing now with polygons and other graphical objects, the leaf rule for the octree construction process must be changed slightly:

1. maximum depth reached; or,
2. only one polygon/object occupies the cell.

We can try to better approximate the geometry of the scene by changing the rule to stop only when there are no objects in the cell (or the maximum depth is reached).

How do we traverse an octree along a given ray? Like in the case of a grid, we have to make “horizontal” steps, which actually advance along the ray. With octrees, though, we also need to make “vertical” steps, which traverse the octree up or down.

All algorithms for ray shooting with octrees can be classified into two classes:

- Bottom-up: this method starts at that leaf in the octree that contains the origin of the ray; from there it tries to find that neighbor cell that is stabbed next by the ray, etc.
- Top-down: this method starts at the root of the octree, and tries to recurse down into exactly

those nodes and leaves that are stabbed by the ray.

Here, we will describe a top-down method.⁷⁹ The idea is to work only with the ray parameter in order to decide which children of a node must be visited.

Let the ray be given by

$$\vec{x} = \vec{p} + t\vec{d}$$

and a voxel v by

$$[x_l, x_h] \times [y_l, y_h] \times [z_l, z_h]$$

In the following, we will describe the algorithm assuming that all $d_i > 0$; later, we will show that the algorithm works also for all other cases.

First of all, observe that if we already have the line parameters of the intersection of the ray with the borders of a cell, then it is trivial to compute the line intervals half-way in between (see Figure 16):

$$t_\alpha^m = \frac{1}{2}(t_\alpha^l + t_\alpha^h), \quad \alpha \in \{x, y, z\} \quad (2)$$

So, for 8 children of a cell, we need to compute only three new line parameters. Clearly, the line intersects a cell if and only if $\max\{t_i^l\} < \min\{t_j^h\}$.

The algorithm can be outlined as follows:

```
|  |
| --- |
| traverse( $v, t^l, t^h$ ) |
| compute  $t^m$ |
| determine order in which sub-cells are hit by the ray |
| for all sub-cells  $v_i$  that are hit do |
| traverse( $v_i, t^l|t^m, t^m|t^h$ ) |
| end for |

```

where $t^l|t^m$ means that we construct the lower boundary for the respective cell by passing the appropriate components from t^l and t^m .




Figure 16: Line parameters are trivial to compute for children of a node.




Figure 17: The sub-cell that must be traversed first can be found by simple comparisons. Here, only the case $t_x^l > t_y^l$ is depicted.

In order to determine the order in which sub-cells should be traversed, we first need to determine which sub-cell is being hit first by the ray. In 2D, this is accomplished by two comparisons (see Figure 17). Then, the comparison of t_x^m with t_y^m tells us which cell is next.

In 3D, this takes a little bit more work, but is essentially the same. First, we determine on which side the ray has been entering the current cell by Table 1.

Next, we determine the first sub-cell to be visited by Table 2 (see Figure 18 for the numbering scheme). The first column is the entering side determined in the first step. The third column yields the index of the first sub-cell to be visited: start with an index of zero; if one or both of the conditions of the second column hold, then the corresponding bit in the index as indicated by the third column should be set.

Finally, we can traverse all sub-cells according to Table 3, where “ex.” means the exit side of the ray for the *current* sub-cell.

If the ray direction contains a negative component(s), then we just have to mirror all tables along the respective axis (axes) conceptually. This can be implemented efficiently by an XOR operation.

2.5.2 5D Octree

In the previous, simple algorithm, we still walk along a ray every time we shoot it into the scene. However, rays are essentially static objects, just like the geometry of the scene! This is the basic observation behind the following algorithm.^{3,6} Again, it




Figure 18: Sub-cells are numbered according to this scheme.

makes use of octrees to adaptively decompose the problem.

The underlying technique is a discretization of rays, which are 5-dimensional objects. Consider a cube enclosing the unit sphere of all directions. We can identify any ray’s direction with a point on that cube, hence it is called *direction cube* (see Figure 19). The nice thing about it is that we can now perform any hierarchical partitioning scheme that works in the plane, such as an octree: we just apply the scheme individually on each side.

Using the direction cube, we can establish a one-to-one mapping between direction vectors and points on all 6 sides of the cube, i.e.,

$$S^2 \leftrightarrow [-1, +1]^2 \times \{+x, -x, +y, -y, +z, -z\}$$

We will denote the coordinates on the cube’s side by u and v .

$\max\{t_i^l\}$	Side
t_x^l	YZ
t_y^l	XZ
t_z^l	XY

Table 1: Determines the entering side.

Side	condition	index bits
XY	$t_z^m < t_x^l$	0
	$t_y^m < t_x^l$	1
XZ	$t_x^m < t_y^l$	0
	$t_z^m < t_y^l$	2
YZ	$t_y^m < t_x^l$	1
	$t_z^m < t_x^l$	2

Table 2: Determines the first sub-cell.

current sub-cell	exit side		
	YZ	XZ	XY
0	4	2	1
1	5	3	ex
2	6	ex	3
3	7	ex	ex
4	ex	6	5
5	ex	7	ex
6	ex	ex	7
7	ex	ex	ex

Table 3: Determines the traversal order of the sub-cells.




Figure 19: With the direction cube, we can discretize directions, and organize them with any hierarchical partitioning scheme.




Figure 20: A uv interval on the direction cube plus a xyz interval in 3-space yield a beam.

Within a given universe $B = [0, 1]^3$ (we assume it is a box), we can represent all possibly occurring rays by points in

$$R = B \times [-1, +1]^2 \times \{+x, -x, +y, -y, +z, -z\} \quad (3)$$

which can be implemented conveniently by 6 copies of 5-dimensional boxes.

Returning to our goal, we now build six 5-dimensional octrees as follows. Associate (conceptually) all objects with the root. Partition a node in the octree, if

1. there are too many objects associated with it; *and*
2. the node's cell is too large.

If a node is partitioned, we must also partition its set of objects and assign each subset to one of the children.

Observe that each node in the 5D octree defines a beam in 3-space: the xyz -interval of the first three coordinates of the cell define a box in 3-space, and

the remaining two uv -intervals define a cone in 3-space. Together (more precisely, their Minkowski sum) they define a beam in 3-space that starts at the cell's box and extends in the general direction of the cone (see Figure 20).

Since we have now defined what a 5D cell of the octree represents, it is almost trivial to define how objects are assigned to sub-cells: we just compare the bounding volume of each object against the sub-cells 3D beam. Note that an object can be assigned to several sub-cells (just like in regular 3D octrees). The test whether or not an object intersects a beam could be simplified further by enclosing a beam with a cone, and then checking the objects bounding sphere against that cone. This just increases the number of false positives a little bit.

Having computed the six 5D octrees for a given scene, ray tracing through that octree is almost trivial: map the ray onto a 5D point via the direction cube; start with the root of that octree which is associated to the side of the direction cube onto which the ray was mapped; find the leaf in that oc-

tree that contains the 5D point (i.e., the ray); check the ray against all objects associated with that leaf.

By locating a leaf in one of the six 5D octrees, we have discarded all objects that do *not* lie in the general direction of the ray. But we can optimize the algorithm even further.

First of all, we sort all objects associated with a leaf along the dominant axis of the beam by their minimum (see Figure 21). If the minimum coordinate of an object along the dominant axis is greater than the current intersection point, then we can stop — all other possible intersection points are farther away.

Second, we can utilize ray coherence as follows. We maintain a cache for each level in the ray tree that stores the leaves of the 5D octrees that were visited last time. When following a new ray, we first look into the octree leaf in the cache whether it is contained therein, before we start searching for it from the root.

Another trick (that works with other ray acceleration schemes as well) is to exploit the fact that we do not need to know the *first* occluder between a point on a surface and a light source. Any occluder suffices to assert that the point is in shadow. So we also keep a cache with each light source which stores that object (or a small set) which has been an occluder last time.

Finally, we would like to mention a memory optimization technique for 5D octrees, because they can occupy a lot of memory. It is based on the observation that within a beam defined by a leaf of the octree the objects at the back (almost) never intersect with a ray emanating from that cell (see Figure 22). So we store objects with a cell only if they are within a certain distance. Should a ray not hit any object, then we start a new intersection query with another ray that has the same direction and a starting point just behind that maximum distance. Obviously, we have to make a trade-off between space and speed here, but when chosen properly, the cut-off distance should not reduce performance too much while still saving a significant amount of memory.

3 BSP Trees

BSP trees (short for *binary space partitioning trees*) can be viewed as a generalization of k -d trees. Like k -d trees, BSP trees are binary trees, but now the

orientation and position of a splitting plane can be chosen arbitrarily. To get a feeling for a BSP tree, Figure 23 shows an example for a set of objects.

The definition of a BSP (short for BSP tree) is fairly straight-forward. Here, we will present a recursive definition. Let h denote a plane in \mathbb{R}^d , h^+ and h^- denote the positive and negative half-space, resp.

Definition 1 (BSP tree)

Let S be a set of objects (points, polygons, groups of polygons, or other spatial objects), and let $S(v)$ denote the set of objects associated with a node v . Then the BSP $T(S)$ is defined by

1. If $|S| \leq 1$, then T is a leaf v which stores $S(v) := S$.
2. If $|S| > 1$, then the root of T is a node v ; v stores a plane h_v and a set $S(v) := \{x \in S \mid x \subseteq h_v\}$ (this is the set of objects that lie completely inside h_v ; in 3D, these can only be polygons, edges, or points). v also has two children T^- and T^+ ; T^- is the BSP for the set of objects $S^- := \{x \cap h_v^- \mid x \in S\}$, and T^+ is the BSP for the set of objects $S^+ := \{x \cap h_v^+ \mid x \in S\}$.

This can readily be turned into a general algorithm for constructing BSPs. Note that a splitting step (i.e., the construction of an inner node) requires us to split each object into two disjoint *fragments* if it straddles the splitting plane of that node. In some applications though (such as ray shooting), this is not really necessary; instead, we can just put those objects into both subsets.

Note that with each node of the BSP a convex cell is associated (which is possibly unbounded): the “cell” associated with the root is the whole space, which is convex; splitting a convex region into two parts yields two convex regions. In Figure 23, the convex region of one of the leaves has been highlighted as an example.

With BSPs, we have much more freedom to place the splitting planes than with k -d trees. However, this also makes that decision much harder (as almost always in life). If our input is a set of polygons, then a very common approach is to choose one of the polygons from the input set and use this as the splitting plane. This is called an *auto-partition* (see Figure 24).

While an auto-partition can have $\Omega(n^2)$ fragments, it is possible to show the following in 2D.^{23,75}




Figure 21: By sorting objects with in each 5D leaf, we can often stop checking ray intersection quite early.




Figure 22: By truncating the beam (or rather, the list of objects) we can save a lot of memory usage of a 5D octree, while reducing performance only insignificantly.




Figure 23: An example BSP tree for a set of objects.

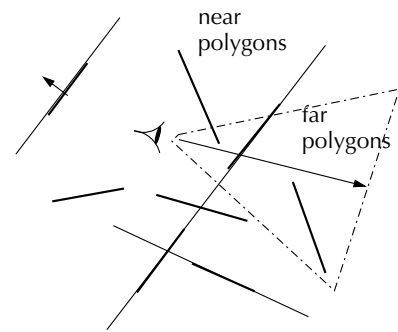


Figure 25: BSP trees are an efficient data structure encoding visibility order of a set of polygons.




Figure 24: Left: an auto-partition. Right: an example configuration of which any auto-partition must have quadratic size.

Lemma 1

Given a set S of n line segments in the plane, the expected number of fragments in an auto-partition $T(S)$ is in $O(n \log n)$; it can be constructed in time $O(n^2 \log n)$.

In higher dimensions, it is not possible to show a similar result. In fact, one can construct sets of polygons such that any BSP (not just auto-partitions) must have $\Omega(n^2)$ many fragments (see Figure 24 for a “bad” example for auto-partitions).

However, all of these examples producing quadratic BSPs violate the *principle of locality*: polygons are small compared to the extent of the whole set. In practice, no BSPs have been observed that exhibit the worst-case quadratic behavior.⁶⁹

3.1 Rendering Without a Z-Buffer

BSP trees were introduced to computer graphics by Fuchs et al.³⁶ At the time, hidden-surface removal was still a major obstacle towards interactive computer graphics, because a z-buffer was just too costly in terms of memory.

In this section, we will describe how to solve this problem, not so much because the application itself is relevant today, but because it nicely exhibits one of the fundamental “features” of BSP trees: they

enable efficient enumeration of all polygons in *visibility order* from any point in any direction.¹

A simple algorithm to render a set of polygons with correct hidden-surface removal, and without a z-buffer, is the *painter's algorithm*: render the scene from back to front as seen from the current viewpoint. Front polygons will just overwrite the contents of the frame buffer, thus effectively hiding the polygons in the back. There are polygon configurations where this kind of sorting is not always possible, but we will deal with that later.

How can we efficiently obtain such a visibility order of all polygons? Using BSP trees, this is almost trivial: starting from the root, first traverse the branch that does *not* contain the viewpoint, then render the polygon stored with the node, then traverse the other branch containing the viewpoint (see Figure 25).

For sake of completeness, we would like to mention a few strategies to optimize this algorithm. First of all, we should make use of the viewing direction by skipping BSP branches that lie completely behind the viewpoint.

Furthermore, we can perform back-face culling as usual (which does not cause any extra costs). We can also perform view-frustum culling by testing all vertices of the frustum against the plane of a BSP node.

Another problem with the simple algorithm is that a pixel is potentially written to many times (this is exactly the *pixel complexity*), although only the last write “survives”. To remedy this, we must traverse the BSP from front to back. But in order to actually save work, we also need to maintain a 2D BSP for the screen that allows us to quickly discard those parts of a polygon that fall onto a screen area that is already occupied. In that 2D screen BSP, we mark all cells either “free” or “occupied”. Initially, it consists only of a “free” root node. When a new polygon is to be rendered, it is first run through the screen BSP, splitting it into smaller and smaller convex parts until it reaches the leaves. If a part reaches a leaf that is already occupied, nothing happens; if it reaches a free leaf, then it is inserted beneath that leaf, and this part is drawn on the screen.




Figure 26: Each leaf cell of BSP representation of an object is completely inside or completely outside.




Figure 27: Using BSPs, we can efficiently compute these boolean operations on solids.

3.2 Representing Objects with BSPs

BSPs offer a nice way to represent volumetric polygonal objects, which are objects consisting of polygons that are closed, i.e., they have an “inside” and an “outside”. Such a BSP representation of an object is just like an ordinary BSP for the set of polygons (we can, for instance, build an auto-partition), except that here we stop the construction process (see Definition 1) only when the set is empty. These leaves represent homogeneous convex cells of the space partition, i.e., they are completely “in” our “out”.

Figure 26 shows an example for such a BSP representation. In this section, we will follow the convention that normals point to the “outside”, and that the right child of a BSP node lies in the positive half-space and the left child in the negative half-space. So, in a real implementation that adheres to these conventions, we can still stop the construction when only one polygon is left, because we know that the left child of such a pseudo-leaf will be “in” and the right one will be “out”.

Given such a representation, it is very easy and efficient, for instance, to determine whether or not a given point is inside an object. In the next section, we will describe an algorithm for solving a slightly more difficult problem.

¹ Actually, the first version of Doom used exactly this algorithm to achieve its fantastic frame rate (at the time) on PCs even without any graphics accelerator.




Figure 28: The fundamental step of the construction is this simple operation, which merges a BSP and a plane.




Figure 29: The main building block of the algorithm consists of these four cases (plus analogous ones).




Figure 30: Computation of boolean operations is based on a general merge operation.

3.3 Boolean Operations

In solid modeling, a very frequent task is to compute the intersection or union of a pair of objects. More generally, given two objects A and B, we want to compute $C := A \text{ op } B$, where $\text{op} \in \{\cup, \cap, \setminus, \ominus\}$ (see Figure 27). This can be computed efficiently using the BSP representation of objects.^{68,69} Furthermore, the algorithm is almost the same for all of these operations: only the elementary step that processes two leaves of the BSPs is different.

We will present the algorithm for boolean operations bottom-up in three steps. The first step is a sub-procedure for computing the following simple operation: given a BSP T and a plane H , construct a new BSP \hat{T} whose root is H , such that $\hat{T}^- \triangleq T \cap H^-$, $\hat{T}^+ \triangleq T \cap H^+$ (see Figure 28). This basically splits a BSP tree by a plane and then puts that plane at the root of the two halves. Since we will not need the new tree \hat{T} explicitly, we will describe only the splitting procedure (which is the bulk of the work anyway).

First, we need to define some nomenclature:

T^-, T^+ = left and right child of T , resp.

$R(T)$ = region of the cell of node T (which is convex)

T^\oplus, T^\ominus = portion of T on the positive/negative side of H , resp.

Finally, we would like to define a node T by the tuple (H_T, p_T, T^-, T^+) , where H is the splitting plane, p is the polygon associated with T (with $p \subset H$).

The pseudo-code below is organized into 8 cases, of which the code details 4 (see Figure 29):

```

split-tree( $T, H, P$ )  $\rightarrow (T^\ominus, T^\oplus)$ 
{ $P = H \cap R(T)$ }
case  $T$  is a leaf :
    return  $(T^\ominus, T^\oplus) \leftarrow (T, T)$ 
case "anti-parallel" and "on" :
    return  $(T^\ominus, T^\oplus) \leftarrow (T^+, T^-)$ 
case "pos./pos." :
     $(T^{+\oplus}, T^{+\ominus}) \leftarrow \text{split-tree}(T^+, H)$ 
     $T^\ominus \leftarrow (H_T, p_T, T^-, T^{+\oplus})$ 
     $T^\oplus \leftarrow T^{+\ominus}$ 
case "mixed" :
     $(T^{+\oplus}, T^{+\ominus}) \leftarrow \text{split-tree}(T^+, H, P \cap R(T^+))$ 
     $(T^{-\oplus}, T^{-\ominus}) \leftarrow \text{split-tree}(T^-, H, P \cap R(T^-))$ 
     $T^\ominus \leftarrow (H_T, p_T \cap H^-, T^{-\oplus}, T^{+\oplus})$ 
     $T^\oplus \leftarrow (H_T, p_T \cap H^+, T^{-\oplus}, T^{+\oplus})$ 
    return  $(T^\ominus, T^\oplus)$ 
end case
    
```

This might look a little bit confusing at first sight, but it is really pretty simple. A few notes might be in order.

The polygon P is only needed in order to find the case applying at each recursion. Computing $P \cap R(T^+)$ might seem very expensive. However, it can be computed quite efficiently by computing $P \cap H_T^+$, which basically amounts to finding the two edges that intersect with H_T . Please see Chin¹⁸ for more details on how to detect the correct case.

It seems surprising at first sight that function `split-tree` does almost no work — it just traverses the BSP tree, classifies the case found at each recursion, and computes $p \cap H^+$ and $p \cap H^-$.

The previous algorithm is already the main building block of the overall boolean operation algorithm. The next step towards that end is an algorithm that performs a so-called *merge* operation on two BSP trees T_1 and T_2 . Let \mathcal{C}_i denote the set of elementary cells of a BSP, i.e., all regions $R(L_j)$




Figure 31: A graphical depiction of the merge step in the algorithm for boolean operations on objects represented by BSP trees.

of tree T_i where L_j are all the leaves. Then the merge of T_1, T_2 yields a new BSP tree T_3 such that $\mathcal{C}_3 = \{c_1 \cap c_2 | c_1 \in \mathcal{C}_1, c_2 \in \mathcal{C}_2, c_1 \cap c_2 \neq \emptyset\}$ (see Figure 30).

The merge operation consists of two cases. The first, almost trivial, case occurs when one of the two operands is a leaf: then at least one of the two regions is homogenous, i.e., completely inside or outside. In the other case, both trees are inhomogeneous over the same region of space: then, we just split one of the two with the splitting plane from the root of the other, and we obtain two pairs of BSPs, that are smaller, and still cover the same regions in space; those two pairs can be merged recursively (see Figure 31). The following pseudo-code describes this recursive procedure more formally:

```

merge( $T_1, T_2$ )  $\rightarrow T_3$ 
if  $T_1$  or  $T_2$  is a leaf then
    perform the cell-op as required by the boolean operation to be constructed (see below)
else
    ( $T_2^\ominus, T_2^\oplus$ )  $\leftarrow$  split-tree( $T_2, H_1, \dots$ )
     $T_3^- \leftarrow$  merge( $T_1^-, T_2^\ominus$ )
     $T_3^+ \leftarrow$  merge( $T_1^+, T_2^\oplus$ )
     $T_3 \leftarrow (H_1, T_3^-, T_3^+)$ 
end if
    
```

The function `cell-op` is the only place where the semantic of the general merge operation is specialized. When we have reached that point, then we know that one of the two cells is homogeneous, so we can just replace it by the other node's sub-tree suitably modified according to the boolean operation. The following table lists the details of this function (assuming that T_1 is the leaf):

Operation	T_1	Result
\cup	in	T_1
	out	T_2
\cap	in	T_2
	out	T_1
\setminus	in	T_2^c
	out	T_1
\ominus	in	T_2^c
	out	T_2

Furthermore, we would like to point out that the merge function is symmetric: it does not matter whether we partition T_2 with H_1 or, the other way round, T_1 with H_2 — the result will be the same.

3.4 Construction Heuristics

One can prove that in general, auto-partitions have the same complexity as other partitionings.^{23,75} In addition, it has been proven that “fat” objects (i.e., objects with bounded aspect ration) and “uncluttered” scenes admit a BSP with linear size.^{21,22}

However, for practical applications, the “hidden” constants do matter. So, the construction strategy should produce “good” BSPs. Depending on the application itself, however, the definition of exactly what a “good” BSP tree is can be quite different. Basically, there are two kinds of applications:

- **Classification:** These are applications where the BSP is used to determine the inside/outside status of a point, or whether or not a line intersects with an object.

In this case, we try to optimize the balancedness of the BSP.

- **Visibility:** These are applications where the BSP is used to sort the polygons in “visibility order” such as rendering without z-buffer.

Consequently, we try to minimize the number of splits, i.e., the size of the BSP.

3.4.1 Convex objects

As an example, consider an convex object. In that case, an auto-partition has size $O(n)$, takes time $O(n^2)$ to construct, and is a linear tree. This does not seem to be a very smart BSP (although it is perfectly suitable for visibility ordering).

If we allow arbitrary splitting planes, then we can balance the BSP much better. The construction time will be

$$T(n) = n + 2T\left(\frac{n}{2} + \alpha n\right) \in O(n^{1+\delta}), \quad 0 < \alpha < \frac{n}{2}$$

where α is the fraction of polygons that are split at each node. The following table shows the actual complexities for some values of α :

α	0.05	0.2	0.4
	$n^{1.15}$	n^2	n^7

As mentioned above, the question now is how to choose the splitting planes. We propose the following simple heuristic:² Compute a representative vertex for each polygon (barycenter, bbox center, etc.). Determine a plane, such that on both sides are about the same number of points, and such that all points are far away from the plane (obviously, this is an optimization).

3.4.2 Cost driven heuristic

In order to derive construction criteria, obviously, one needs to define the *quality* of a BSP. An abstract measure is the *cost* of a BSP T

$$C(T) = 1 + P(T^-)C(T^-) + P(T^+)C(T^+) \quad (4)$$

where $P(T^-)$ is the probability that the left subtree T^- will be visited under the condition that the tree

T has been visited (dito for $P(T^+)$). This probability depends, of course, on the application the BSP is going to be used for. For instance, in the case of inside/outside queries

$$P(T^-) = \frac{\text{Vol}(T^-)}{\text{Vol}(T)}$$

Obviously, trying to optimize Equation 4 globally is prohibitively expensive. Therefore, a local heuristic must be employed to guide the splitting process. A simple heuristic is the following:⁶⁹ estimate the cost of the subtrees by the number of polygons, and add a penalty for polygons that are split by the current node, so that

$$C(T) = 1 + |S^-|^\alpha + |S^+|^\alpha + \beta s \quad (5)$$

where S is the set of polygons associated with a node, s the set of polygons split by a node, and α, β are two parameters that can be used to make the BSP more balanced or smaller. (⁶⁹ reports that $\alpha = 0.8, \dots, 0.95$ and $\beta = \frac{1}{4}, \dots, \frac{3}{4}$ are usually good values to start from.) Again, this yields an optimization process, but now it only a local process.

If the BSP is to be an auto-partition, then a very quick approximation of the local optimum yields very good results: just sort the polygons according to their size and evaluate Equation 5 for the first k polygons. Then, choose the one that produces the least costly BSP (subtree), the rationale being that the probability that a polygon gets split is proportional to its size, so we try to get rid of them as early as possible.

An even simpler heuristic was proposed by.³⁶ Randomly choose k polygons from S and select the one that produces the least number of splits. They report that $k = 5$ yielded near-optimal BSPs (for visibility ordering).

3.4.3 Non-uniform queries

In the previous section, we assumed that all queries are uniformly distributed over a certain domain. This is a valid assumption if nothing is known about the distribution. On the other hand, if we know more about it, then we should make use of this and construct the BSP such that frequent queries can be answered most quickly.³

² For the case of convex objects, a heuristic has already been proposed by.⁸⁷ However, we believe that their heuristic has some flaws.

³ The same principle is underlying the Huffman encoding scheme.

Indeed, quite often we know more about the queries. For instance, in ray tracing, the starting points are usually not uniformly distributed in space; for instance, they usually do not emanate from the interior of objects. Also, the prominent polygons of an object are hit more frequently than those that are within cavities or completely inside.

According to,¹ we can estimate the costs of a query by

$$\begin{aligned} C(\text{query}) &= \# \text{ nodes visited} \\ &\leq \text{depth}(\text{BSP}) \cdot \# \text{ stabbed leaf cells} \end{aligned}$$

So, according to this, we should minimize the number of stabbed leaf cells before a polygon hit occurs. The factors influencing the probability that a ray hits a polygon are

1. If the angle between a ray and a polygon is large, then the probability is large.
2. If the polygon is large (relative to the total size of the object/universe), then the probability is large.
3. ...

Let $\omega(l)$ denote the density of all rays l over some domain D ; this could be measured, or it could be derived from the geometry. Let S be the set of polygons over which the BSP is to be built. Assign to each polygon p a score

$$\text{score}(p) = \int_D w(S, p, l) \omega(l) dl$$

where the weight w is defined as

$$w(S, p, l) = \sin^2(\mathbf{n}_p, \mathbf{r}_l) \frac{\text{Area}(p)}{\text{Area}(S)}$$

and \mathbf{n}_p is the normal of p and \mathbf{r}_l is the direction of l .

So the algorithm for constructing a BSP adapted to a given distribution is the following randomized greedy strategy: sort all polygons according to $\text{score}(p)$; choose randomly one out of the “top k ” and split the set S . Thus, polygons with a higher probability of getting hit by the ray tend to end up higher in the tree.

The BSP thus constructed has been named *customized* BSP by.¹ The authors report that it usually has about $2 \times$ as many polygons as its “oblivious” version, but the queries have to visit only $2 - 10 \times$ less polygons.

3.4.4 Deferred, self-organizing BSPs

Now, what should we do if we just don’t know the query distribution, or if it is too expensive to measure it by experiments? The answer is, of course, to defer the complete construction of the BSP, in other words, we only build as much of the BSP that is absolutely necessary. In addition, we keep a history (in some form) of all the queries we have seen so far, so whenever we have to build a new part of the BSP we base the construction on that history (as a best guess of the probability distribution of all queries to come).²

As an example, let us consider the problem of detecting intersections between a 3-dim. line segment and a set of polygons.⁴

Since a BSP is now no longer completely constructed before we use it, the nodes must store additional information:

1. the polygon P defining the splitting plane; or, if it is a preliminary leaf,
2. a list $\mathcal{L} \subseteq S$ of polygons associated with it, for which the subtree has not been built yet.

The algorithm for answering queries with a ray R now also triggers the construction of the BSP:

```

testray( $R, v$ )
if  $v$  is a leaf then
  for all  $P \in \mathcal{L}_v$  do
    if  $R$  intersects  $P$  then
      return hit
    end if
  end for
else
   $v_1 \leftarrow$  child of  $v$  that contains startpoint of  $R$ 
   $v_2 \leftarrow$  other child
  testray( $R, v_1$ )
  if no hit in  $v_1$  then
    testray( $R, v_2$ )
  end if
end if

```

Since the line segment is finite (in particular, if it is short), this algorithm can usually stop much earlier, but the details have been omitted here.

The initial BPS is just one node (the root) with $\mathcal{L} = S$, i.e., all polygons of the object or scene.

Now we need to fill in the open issues, namely:

⁴ Sometimes, this is also called “collision detection”, in particular in the game programming industry, because we are only interested in a yes/no answer, while in ray tracing we want to determine the *earliest* intersection.

1. when do we split a preliminary leaf? (and which one?)
2. how do we split?

For the “when” question: we keep an access counter per node, which gets incremented every time that node is traversed during a query. Whenever one of the counters is over a certain threshold, then we split that node (leaf). The threshold can be an absolute value or relative to the sum of all counters.

For the “how” question: we keep a counter per polygon $P \in \mathcal{L}_v$, which is incremented every time an intersection with P is found. We sort \mathcal{L}_v according to this counter; this can be done incrementally whenever one of the counters changes. If a split is to be performed for v , then we use the first polygon from \mathcal{L}_v as the splitting plane.

It turns out that many polygons are never hit by a line segment. Therefore, with this algorithm, a BSP subtree will never be “wasted” on these polygons, and they will be stored at the end of the lists at the leaves.

There are other ways to organize the polygon lists at the leaves: move the polygon currently being hit to the front of the list; or, swap it with the one in front of it. However, according to,² the strategy that sorts the list seemed to work best.

According to,² the performance gain in their experiments was a factor of about 2–20.

4 Bounding Volume Hierarchies

Like the previous hierarchical data structures, bounding volume hierarchies (BVHs) are mostly used to prevent performing an operation exhaustively on all objects. Like with previously discussed hierarchical data structures, one can improve a huge range of applications and queries using BVHs, such as ray shooting, point location queries, nearest-neighbor search, view frustum and occlusion culling, geographical data bases, and collision detection (the latter will be discussed in more detail below).

Often times, bounding volume (BV) hierarchies are described as the opposite of spatial partitioning schemes, such as quadtrees or BSP trees: instead of partitioning space, the idea is to partition the set of objects recursively until some leaf criterion is




Figure 32: Some of the most commonly used BVs, and some less often used ones.

met.⁵ Here, objects can be anything from points to complete graphical objects. With BV hierarchies, almost all queries, which can be implemented with space partitioning schemes, can also be answered, too. Example queries and operations are ray shooting, frustum culling, occlusion culling, point location, nearest neighbor, collision detection.

Definition 2 (BV hierarchy)

Let $O = \{o_1, \dots, o_n\}$ be a set of elementary objects. A bounding volume hierarchy for O , $BVH(O)$, is defined by

1. If $|O| = e$, then $BVH(O) :=$ a leaf node that stores O and a BV of O ;
2. If $|O| > e$, then $BVH(O) :=$ a node v with $n(v)$ children v_1, \dots, v_n , where each child v_i is a BV hierarchy $BVH(O_i)$ over a subset $O_i \subset O$, such that $\bigcup O_i = O$. In addition, v stores a BV of O .

The definition mentions two parameters. The threshold e is often set to 1, but depending on the application, the optimal e can be much larger. Just like sorting, when the set of objects is small, it is often cheaper to perform the operation on all of them, because recursive algorithms always incur some overhead.

Another parameter in the definition is the arity. Mostly, BV hierarchies are constructed as binary trees, but again, the optimum can be larger. And what is more, as the definition suggests, the out-degree of nodes in a BV hierarchy does not necessarily have to be constant, although this often simplifies implementations considerably.

Effectively, these two parameters, e and $n(v)$, control the balance between linear search/operation

⁵ However, we will argue at the end that BV hierarchies are just at the other end of a whole spectrum of hierarchical data structures.

(which is exhaustive), and a maximally recursive algorithm.

There are more design choices possible according to the definition. For inner nodes, it only requires that $\bigcup O_i = O$; this means, that the same object $o \in O$ could be associated with several children. Depending on the application, the type of BVs, and the construction process, this may not always be avoidable. But if possible, you should always split the set of objects into disjoint subsets.

Finally, there is, at least, one more design choice: the type of BV used at each node. Again, this does not necessarily mean that each node uses the same type of BV. Figure 32 shows a number of the most commonly used BVs. The difference between OBBs⁶ and AABBs is that OBBs can be oriented arbitrarily (hence “oriented bounding boxes”). DOPs^{53,57,99} are a generalization of AABBs: basically, they are the intersection of k slabs. Prisms and cylinders have been proposed by Barequet et al.⁹ and Weghorst et al.,⁹⁴ but they seem to be too expensive computationally. A spherical shell is the intersection of a shell and a cone (the cone’s apex coincides with the sphere’s center), and a shell is the space between two concentric spheres. Finally, one can always take the intersection of two or more different types of BVs.⁵²

There are three characteristic properties of BVs:

- tightness,
- memory usage,
- number of operations needed to test the query object against a BV.

Often, one has to make a trade-off between these properties: generally, the type of BV that offers better tightness also requires more operations per query and more memory.

Regarding the tightness, one can establish a theoretical advantage of OBBs. But first, we need to define tightness.⁴⁰

Definition 3 (Tightness by Hausdorff distance)

Let B be a BV, G some geometry bounded by B , i.e., $g \subset B$. Let

$$h(B, G) = \max_{b \in B} \min_{g \in G} d(b, g)$$

be the *directed Hausdorff distance*, i.e., the maximum distance of B to the nearest point in G . (Here,

d is any metric, very often just the Euclidean distance.) Let

$$\text{diam}(G) = \max_{g, f \in G} d(g, f)$$

be the *diameter* of G .

Then we can define *tightness*

$$\tau := \frac{h(B, G)}{\text{diam}(G)}.$$

See Figure 33 for an illustration.

Since the Hausdorff distance is very sensitive to outliers, one could also think of other definitions such as the following one:

Definition 4 (Tightness by volume)

Let $C(v)$ be the set of children of a node v of the BV hierarchy. Let $\text{Vol}(v)$ be the volume of the BV stored with v .

Then, we can define the tightness as

$$\tau := \frac{\text{Vol}(v)}{\sum_{v' \in C(v)} \text{Vol}(v')}.$$

Alternatively, we can define it as

$$\tau := \frac{\text{Vol}(v)}{\sum_{v' \in L(v)} \text{Vol}(v')}.$$

where $L(v)$ is the set of leaves beneath v .

Getting back to the tightness definition based on the Hausdorff distance, we observe a fundamental difference between AABBs and OBBs:⁴⁰

- The tightness of AABBs depends on the orientation of the enclosed geometry. What is worse is that the tightness of the children of an AABB enclosing a surface of small curvature is almost the same as that of the father.

The worst case is depicted in Figure 34. The tightness of the father is $\tau = h/d$, while the tightness of a child is $\tau' = \frac{h'}{d/2} = \frac{h/2}{d/2} = \tau$.

- The tightness of OBBs does not depend on the orientation of the enclosed geometry. Instead, it depends on its curvature, and it decreases approximately linearly with the depth in the hierarchy.

Figure 35 depicts the situation for a sphere. The Hausdorff distance from an OBB to an enclosed




Figure 33: One way to define tightness is via the directed Hausdorff distance.




Figure 34: The tightness of an AABB remains more or less constant throughout the levels of a AABB hierarchy for surfaces of small curvature.




Figure 35: The tightness of an OBB decreases for deeper levels in a OBB hierarchy for small curvature surfaces.

spherical arc is $h = r(1 - \cos \phi)$, while the diameter of the arc is $d = 2r \sin \phi$. Thus, the tightness for an OBB bounding a spherical arc of degree ϕ is $\tau = \frac{1 - \cos \phi}{2 \sin \phi}$, which approaches 0 linearly as $\phi \rightarrow 0$.

This makes OBBs seem much more attractive than AABBs. The price of the much improved tightness is, of course, the higher computational effort needed for most queries per node when traversing an OBB tree with a query.

4.1 Construction of BV Hierarchies

Essentially, there are 3 strategies to build BV trees:

- bottom-up,
- top-down,
- insertion

From a theoretical point of view, one could pursue a simple top-down strategy, which just splits the set of objects into two equally sized parts, where the objects are assigned randomly to either subset. Asymptotically, this usually yields the same query time as any other strategy. However, in practice, the query times offered by such a BV hierarchy are by a large factor worse.

During construction of a BV hierarchy, it is convenient to forget about the graphical objects or primitives, and instead deal with their BVs and consider those as the atoms. Sometimes, another

simplification is to just approximate each object by its center (barycenter or bounding box center), and then deal only with sets of points during the construction. Of course, when the BVs are finally computed for the nodes, then the true extents of the objects must be considered.

In the following we will describe algorithms for each construction strategy.

4.1.1 Bottom-up

In this class, we will actually describe two algorithms.

Let B be the set of BVs on the top-most level of the BV hierarchy that has been constructed so far.⁸⁰ For each $b_i \in B$ find the nearest neighbor $b'_i \in B$; let d_i be the distance between b_i and b'_i . Sort B with respect to d_i . Then, combine the first k nodes in B under a common father; do the same with the next k elements from B , etc. This yields a new set B' , and the process is repeated.

Note that this strategy does not necessarily produce BVs with a small “dead space”: in Figure 36, the strategy would choose to combine the left pair (distance = 0), while choosing the right pair would result in much less dead space.

The second strategy is less greedy in that it computes a tiling for each level. We will describe it first in 2D.⁶² Again, let B be the set of BVs on the top-most level so far constructed, with $|B| = n$.




Figure 36: A simple greedy strategy can produce much “dead space”.

The algorithm first computes the center c^i for each $b^i \in B$. Then, it sorts B along the x-axis with respect to c_x^i . Now, the set B is split into $\sqrt{n/k}$ vertical “slices” (again with respect to c_x^i). Now, each slice is sorted according to c_y^i and subsequently split into $\sqrt{n/k}$ “tiles”, so that we end up with k tiles (see Figure 37). Finally, all nodes in a tile are combined under one common father; its BV is combined, and the process repeats with a new set B' .

In \mathbb{R}^d it works quite similarly: we just split each slice repeatedly by $\sqrt[d]{n/k}$ along all coordinate axes.

4.1.2 Insertion

This construction scheme starts with an empty tree. Let B be the set of elementary BVs. The following pseudo-code describes the general procedure:

```

1: while  $|B| > 0$  do
2:   choose next  $b \in B$ 
3:    $v := \text{root}$ 
4:   while  $v \neq \text{leaf}$  do
5:     choose child  $v'$ ,
       so that insertion of  $b$  into  $v'$  causes minimal
       increase in the costs of the total tree
6:      $v := v'$ 
7:   end while
8: end while

```

All insertion algorithms only vary step 2 and/or 5. Step 2 is important because a “bad” choice in the beginning can probably never be made right afterwards. Step 5 depends on the type of query that is




Figure 37: A less greedy strategy combines BVs by computing a “tiling”.

to be performed on the BV tree. See below for a few criteria.

Usually, algorithms in this class have complexity $O(n \log n)$.

4.1.3 Top-down

This scheme is the most popular one. It seems to produce very good hierarchies while still being very efficient, and usually it can be implemented easily.

The general idea is to start with the complete set of elementary BVs, split that into k parts, and create a BV tree for each part recursively. The splitting is guided by some heuristic or criterion that (hopefully) produces good hierarchies.

4.1.4 Construction criteria

In the literature, there is a vast number of criteria for guiding the splitting, insertion, or merging, during BV tree construction. (Often, the authors endow the thus constructed BV hierarchy with a new name, even though the BVs utilized are well known.) Obviously, the criterion depends on the application for which the BV tree is to be used. In the following, we will present a few of these criteria.

For ray tracing, if we can estimate the probability that a ray will hit a child box when it has hit the father box, then we know how likely it is, that we need to visit the child node when we have visited the father node. Let us assume that all rays emanate from the same origin (see Figure 38). Then,




Figure 38: The probability of a ray hitting a child box can be estimated by the surface area.

we can observe that the probability that a ray s hits a child box v' under the condition that it has hit the father box v is

$$P(s \text{ hits } v' | s \text{ hits } v) = \frac{\theta_{v'}}{\theta_v} \approx \frac{\text{Area}(v')}{\text{Area}(v)} \quad (6)$$

where Area denotes the surface area of the BV, and θ denotes the solid angle subtended by the BV. This is because for a convex object, the solid angle subtended by it, when seen from large distances, is approximately proportional to its surface area.³⁹ So, a simple strategy is to just minimize the surface area of the BVs of the children that are produced by a split.⁶

A more elaborate criterion tries to establish a *cost function* for a split and minimize that. For ray tracing, this cost function can be approximated by

$$C(v_1, v_2) = \frac{\text{Area}(v_1)}{\text{Area}(v)} C(v_1) + \frac{\text{Area}(v_2)}{\text{Area}(v)} C(v_2) \quad (7)$$

where v_1, v_2 are the children of v . The optimal split $B = B_1 \cup B_2$ minimizes this cost function:

$$C(B_1, B_2) = \min_{B' \in \mathcal{P}(B)} C(B', B \setminus B')$$

where B_1, B_2 are the subsets of elementary BVs (or objects) assigned to the children. Here, we have assumed a binary tree, but this can be extended to other arities analogously.

Of course, such a minimization is too expensive in practice, in particular, because of the recursive definition of the cost function. So, Fussell and Subramanian,³⁷ Müller et al.,⁶⁷ and Beckmann et al.¹⁰ have proposed the following approximation algorithm:

⁶ For the insertion scheme, the strategy is to choose that child node whose area is increased least.³⁹

```
for  $\alpha = x, y, z$  do
  sort  $B$  along axis  $\alpha$  with respect to the BV centers
  find
```

$$k^\alpha = \min_{j=0 \dots n} \left\{ \frac{\text{Area}(b_1, \dots, b_j)}{\text{Area}(B)} j + \frac{\text{Area}(b_{j+1}, \dots, b_n)}{\text{Area}(B)} (n - j) \right\}$$

```
end for
choose the best  $k^\alpha$ 
```

where $\text{Area}(b_1, \dots, b_j)$ denotes the surface area of the BV enclosing b_1, \dots, b_j .

If the query is a point location query (e.g., is a given point inside or outside the object), then the volume instead of the surface area should be used. This is because the probability that a point is contained in a child BV, under the condition that it is contained in the father BV, is proportional to the ratio of the two volumes.

For range queries, and for collision detection, the volume seems to be a good probability estimation, too.

A quite different splitting algorithm does not (explicitly) try to estimate any probabilities. It just approximates each elementary BV/object by its center point. It then proceeds as follows. For a given set B of such points, compute its principal components (the Eigenvectors of the covariance matrix); choose the largest of them (i.e., the one exhibiting the largest variance); place a plane orthogonal to that principal axis and through the barycenter of all points in B ; finally, split B into two subsets according to the side on which the point lies. (This description is a slightly modified version of Gottschalk et al.⁴⁰) Alternatively, one can place the splitting plane through the median of all points, instead of the barycenter. This would lead to balanced trees, but not necessarily better ones.

4.1.5 The criterion for collision detection

In the following, we will describe a general criterion that can guide the splitting process of top-down BV hierarchy construction algorithms, such that the hierarchy produced is good in the sense of fast collision detection¹⁰⁰ (see Section 4.2).

Let $C(A, B)$ be the expected costs of a node pair (A, B) under the condition that we have already determined during collision detection that we need to traverse the hierarchies further down. Assuming




Figure 39: By estimating the volume of the Minkowski sum of two BVs, we can derive an estimate for the cost of the split of a set of polygons associated with a node.

binary trees and unit costs for an overlap test, this can be expressed by

$$C(A, B) = 4 + \sum_{i,j=1,2} P(A_i, B_j) \cdot C(A_i, B_j) \quad (8)$$

where A_i, B_j are the children of A and B , resp., and $P(A_i, B_j)$ is the probability that this pair must be visited (under the condition that the pair (A, B) has been visited).

An optimal construction algorithm would need to expand (8) down to the leaves:

$$\begin{aligned} C(A, B) = & P(A_1, B_1) + P(A_1, B_1)P(A_{11}, B_{11}) \\ & + P(A_1, B_1)P(A_{12}, B_{11}) + \dots + \\ & P(A_1, B_2) + P(A_1, B_2)P(A_{11}, B_{21}) \\ & + \dots \end{aligned} \quad (9)$$

and then find the minimum. Since we are interested in finding a local criterion, we approximate the cost function by discarding the terms corresponding to lower levels in the hierarchy, which gives

$$C(A, B) \approx 4(1 + P(A_1, B_1) + \dots + P(A_2, B_2)) \quad (10)$$

Now we will derive an estimate of the probability $P(A_1, B_1)$. For sake of simplicity, we will assume in

the following that AABBs are used as BVs. However, similar arguments should hold for all other kinds of convex BVs.

The event of box A intersecting box B is equivalent to the condition that B 's "anchor point" is contained in the Minkowski sum $A \oplus B$. This situation is depicted in Figure 39.⁷ Because B_1 is a child of B , we know that the anchor point of B_1 must lie somewhere in the Minkowski sum $A \oplus B \oplus \mathbf{d}$, where $\mathbf{d} = \text{anchor}(B_1) - \text{anchor}(B)$. Since A_1 is inside A and B_1 inside B , we know that $A_1 \oplus B_1 \subset A \oplus B \oplus \mathbf{d}$. So, for arbitrary convex BVs the probability of overlap is

$$\begin{aligned} P(A_1, B_1) &= \frac{\text{Vol}(A_1 \oplus B_1)}{\text{Vol}(A \oplus B \oplus \mathbf{d})} \\ &= \frac{\text{Vol}(A_1 \oplus B_1)}{\text{Vol}(A \oplus B)} \end{aligned} \quad (11)$$

In the case of AABBs, it is safe to assume that the aspect ratio of all BVs is bounded by α . Consequently, we can bound the volume of the Minkowski sum by

$$\begin{aligned} \text{Vol}(A) + \text{Vol}(B) + \frac{2}{\alpha} \sqrt{\text{Vol}(A) \text{Vol}(B)} &\leq \\ \text{Vol}(A \oplus B) &\leq \\ \text{Vol}(A) + \text{Vol}(B) + 2\alpha \sqrt{\text{Vol}(A) \text{Vol}(B)} \end{aligned} \quad (12)$$

So we can estimate the volume of the Minkowski sum of two boxes by

$$\text{Vol}(A \oplus B) \approx 2(\text{Vol}(A) + \text{Vol}(B))$$

yielding

$$P(A_1, B_1) \approx \frac{\text{Vol}(A_1) + \text{Vol}(B_1)}{\text{Vol}(A) + \text{Vol}(B)} \quad (13)$$

Since $\text{Vol}(A) + \text{Vol}(B)$ has already been committed by an earlier step in the recursive construction, Equation 10 can be minimized only by minimizing $\text{Vol}(A_1) + \text{Vol}(B_1)$. This is our criterion for constructing restricted boxtrees.

Construction algorithm. According to the criterion derived above, each recursion step will try to split the set of polygons so that the cost function (10) is minimized. This is done by trying to find a good splitting for each of the three coordinate

⁷ In the figure, we have chosen the lower left corner of B as its anchor point, but this is arbitrary, of course, because the Minkowski sum is invariant under translation.

axes, and then selecting the best one. Along each axis, we consider three cases: both subsets form lower boxes with respect to its parent, both are upper boxes, or one upper and one lower box.

In each case, we first try to find a good “seed” polygon for each of the two subsets, which is as close as possible to the outer border that is perpendicular to the splitting axis. Then, in a second pass, we consider each polygon in turn, and assign it to that subset whose volume is increased least. In order to prevent “creeping greediness”,⁸ we run alternately through the array of polygons. After good splitting candidates have been obtained for all three axes, we just pick the one with least total volume of the subsets.

The algorithm and criterion we propose here could also be applied to construct hierarchies utilizing other kinds of BVs, such as OBBs, DOPs, and even convex hulls. We suspect that the volume of AABBs would work fairly well as an estimate of the volume of the respective BVs.

This algorithm has proven to be geometrically robust, since there is no error propagation. Therefore, a simple epsilon guard for all comparisons suffices.

It can be shown that this algorithm is, under certain assumptions,⁹ in $O(n)$, where n is the number of polygons.¹⁰⁰

4.2 Collision Detection

Fast and exact collision detection of polygonal objects undergoing rigid motions is at the core of many simulation algorithms in computer graphics. In particular, all kinds of highly interactive applications such as virtual prototyping need exact collision detection at interactive speed for very complex, arbitrary “polygon soups”. It is a fundamental problem of dynamic simulation of rigid bodies, simulation of natural interaction with objects, and haptic rendering.

Bounding volume trees seem to be a very efficient data structure to tackle the problem of collision detection for rigid bodies. All kinds of different types of BVs have been explored in the past:

sphere,^{46,74} OBBs,⁴⁰ DOPs,^{57,99} AABBs,^{58,88,98} and convex hulls,³¹ to name but a few.

Given two hierarchical BV volume data structures for two objects A and B , almost all hierarchical collision detection algorithms implement the following general algorithm scheme:

```

traverse(A,B)
if A and B do not overlap then
    return
end if
if A and B are leaves then
    return intersection of primitives
        enclosed by A and B
else
    for all children A[i] and B[j] do
        traverse(A[i],B[j])
    end for
end if

```

This algorithm quickly “zooms in” on pairs of close polygons. The characteristics of different hierarchical collision detection algorithms lie in the type of BV used, the overlap test for a pair of nodes, and the algorithm for construction of the BV trees.

The algorithm outlined above is essentially a simultaneous traversal of two hierarchies, which induces a so-called *recursion tree* (see Figure 40). Each node in this tree denotes a BV overlap test. Leaves in the recursion tree denote an intersection test of the enclosed primitives (polygons); whether or not a BV test is done at the leaves depends on how expensive it is, compared to the intersection test of primitives.

During collision detection, the simultaneous traversal will stop at some nodes in the recursion tree. Let us call the set of nodes, of which some children are not visited (because their BVs do not overlap), the “bottom slice” through the recursion tree (see the dashed lines in Figure 40).

One idea is to save this set for a given pair of objects.⁶³ When this pair is to be checked next time, we can start from this set, going either up or down. Hopefully, if the objects have moved only a little relative to each other, the number of nodes that need to be added or removed from the bottom slice is small. This scheme is called *incremental hierarchical collision detection*.

5 Voronoi Diagrams

For a given set of sites inside an area the Voronoi diagram is a partition of the area into regions of

⁸ This happens, for instance, when the sequence of polygons happens to be ordered such that each polygon increases one subset’s volume just a little bit, so that the other subset never gets a polygon assigned to.

⁹ These assumption have been valid in all practical cases we have encountered so far.

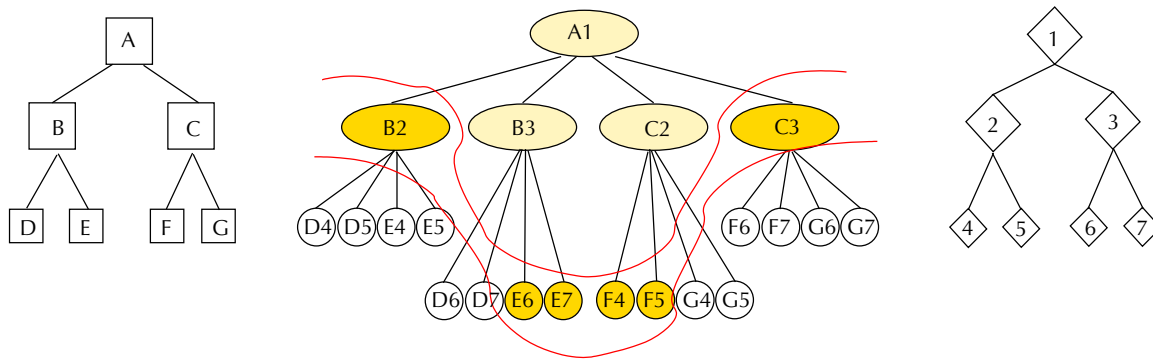


Figure 40: The recursion tree is induced by the simultaneous traversal of two BV trees.




Figure 41: Hierarchical collision detection can discard many pairs of polygons with one BV check. Here, all pairs of polygons in A_1 and B_2 can be discarded.

the same neighborhood. The Voronoi diagram and its dual have been used for solving numerous problems in many fields of science.

We will concentrate on its application to geometric problems in 2D and 3D. For an overview of the Voronoi diagram and its dual in computational geometry one may consult the surveys by Aurenhammer,⁴ Bernal,¹¹ Fortune³⁴ and Aurenhammer and Klein.⁵ Additionally, chapters 5 and 6 of Preparata and Shamos⁷⁷ and chapter 13 of Edelsbrunner²⁹ could be consulted.

We start in Section 5.1 with the simple case of the Voronoi diagram and the Delaunay triangulation of n points in the plane, under the Euclidean distance. Additionally we mention some of the elementary structural properties that follow from the definitions.

In Section 5.2 different algorithmic schemes for computing the structures are mentioned. We present

a simple *incremental construction* approach which can be easily generalized to 3D (see Section 5.3.1).

We present generalizations of the Voronoi diagram and the Delaunay Triangulation in Section 5.3. In Section 5.3.1 transformations to three dimensions are given and in Section 5.3.2 the concept of *constrained Voronoi diagrams* is introduced. A collection of other interesting generalizations is presented in Section 5.3.3.

In Section 5.4 applications of the Voronoi diagram and the Delaunay triangulation in 2D and 3D are shown. First, in Section 5.4.1 we discuss the famous post office problem. Then a simple application of the Voronoi diagram of line segments for motion planning is presented in Section 5.4.2. Finally, in Section 5.4.3 a collection of 2D-applications is shown.

Note, that we can only sketch many of the subjects here. For further details and further literature see one of the surveys mentioned above. The figures are taken from Aurenhammer and Klein.⁵

5.1 Definitions and Elementary Properties

5.1.1 Voronoi Diagram

Let S a set of $n \geq 3$ point sites p, q, r, \dots in the plane. In the following we assume that the points are in *general position*, i.e., no four of them lie on the same circle and no three of them on the same line.

For points $p = (p_1, p_2)$ and $x = (x_1, x_2)$ let $d(p, x)$ denote their Euclidean distance. By \overline{pq} we denote the line segment from p to q . The closure of a set A will be denoted by \overline{A} .

Definition 5 For $p, q \in S$ let

$$B(p, q) = \{x \mid d(p, x) = d(q, x)\}$$

be the bisector of p and q . $B(p, q)$ is the perpendicular line through the center of the line segment \overline{pq} . It separates the halfplane

$$D(p, q) = \{x \mid d(p, x) < d(q, x)\}$$

containing p from the halfplane $D(q, p)$ containing q . We call

$$VR(p, S) = \bigcap_{q \in S, q \neq p} D(p, q)$$

the Voronoi region of p with respect to S . Finally, the Voronoi diagram of S is defined by

$$V(S) = \bigcup_{p, q \in S, p \neq q} \overline{VR(p, S)} \cap \overline{VR(q, S)}.$$

An illustration is given in Figure 42. It shows how the plane is decomposed by $V(S)$ into Voronoi regions. Note that it is convenient to imagine a simple closed curve Γ around the “interesting” part of the Voronoi diagram.




Figure 42: A Voronoi diagram of points in the Euclidean plane.

The common boundary of two Voronoi regions belongs to $V(S)$ and is called a *Voronoi edge*, if it contains more than one point. If the Voronoi edge e borders the regions of p and q then $e \subset B(p, q)$ holds. Endpoints of Voronoi edges are called *Voronoi*

vertices; they belong to the common boundary of three or more Voronoi regions.

There is an intuitive way of looking at the Voronoi diagram. For any point x in the plane we can expand the circle $C(r)$ with center x and radius r by increasing r continuously. We detect three cases depending on which event occurs first:

- If $C(r)$ hits one of the n sites, say p , then $x \in VR(p, S)$.
- If $C(r)$ hits two sites p and q simultaneously x belongs to the Voronoi edge of p and q .
- If $C(r)$ hits three sites p, q and r simultaneously x is the Voronoi vertex of p, q and r .

We will enumerate some of the significant properties of Voronoi diagrams.

1. Each Voronoi region $VR(p, S)$ is the intersection of at most $n - 1$ open halfplanes containing the site p . Every $VR(p, S)$ is open and convex. Different Voronoi regions are disjoint.
2. A point p of S lies on the convex hull of S iff its Voronoi region $VR(p, S)$ is unbounded.
3. The Voronoi diagram $V(S)$ has $O(n)$ many edges and vertices. The average number of edges in the boundary of a Voronoi region is less than 6.

The Voronoi diagram is a simple linear structure and provides for a partition of the plane into cells of the same neighborhood. We omit the proofs and refer to the surveys mentioned in the beginning.

Note, that the Voronoi edges and vertices build a graph. Therefore the diagram normally is represented by a graph of linear size. For example the diagram can be represented by a *doubly connected edge list DCEL*, see de Berg et al.,²³ or with the help of an *adjacency matrix*.

5.1.2 Delaunay Triangulation

We consider the dual graph of the Voronoi diagram, the so called *Delaunay triangulation*. In general, a *triangulation* of S is a planar graph with vertex set S and straight line edges, which is maximal in the sense that no further straight line edge can be added without crossing other edges. The triangulation of a point set S has not more than $O(|S|)$ triangles.

Definition 6 The Delaunay triangulation $DT(S)$ is the dual Graph of the Voronoi diagram. The edges of $DT(S)$ are called Delaunay edges.

Obviously, the Delaunay triangulation $DT(S)$ is a triangulation of S , an example is shown in Figure 43.

We present two equivalent definitions of the Delaunay triangulation. They are applied for the computation of the diagram and give also rise to generalization, for example if the dual of a Voronoi diagram is no longer well-defined.

1. Two points p, q of S give rise to a Delaunay edge iff a circle C exists that passes through p and q and does not contain any other site of S in its interior or boundary.
2. Three points of S give rise to a Delaunay triangle iff their circumcircle does not contain a point of S in its interior.




Figure 43: Voronoi diagram and Delaunay triangulation.

5.2 Computation

The construction of the Voronoi diagram has time complexity $\Theta(n \log n)$. The lower bound $\Omega(n \log n)$ can be achieved by the following reductions.

- A reduction to the convex hull problem is given by Shamos.⁸⁵
- A reduction to the ϵ -closeness problem is given by Djidjev and Lingas²⁶ and by Zhu and Mirzaiian.¹⁰¹

The well-known computation paradigms

- *Incremental construction*,
- *Divide-and-Conquer* and

- *Sweep*

are convenient for the construction of the Voronoi diagram or the Delaunay triangulation, respectively. They can also be generalized to other metrics and sites other than points, for example line segments or polygonal chains. The result of the algorithms is stored in a graph of linear size, see above.

All these approaches run in deterministic time $O(n \log n)$. We explain a simple *Incremental construction* technique which runs in $O(n \log n)$ expected time and computes the Delaunay triangulation. The presentation is adapted from Klein and Aurenhammer.⁵ The technique can easily be generalized to the three dimensional case as we will see in Section 5.3.1.

Simple incremental construction: The insertion process is described as follows: We construct $DT_i = DT(\{p_1, \dots, p_{i-1}, p_i\})$ by inserting the site p_i into DT_{i-1} . We follow Guibas and Stolfi⁴¹ and construct DT_i by exchanging edges, using Lawson's⁶⁰ original edge flipping procedure, until all edges invalidated by p_i have been removed.

It is helpful to extend the notion of triangle to the unbounded face of the Delaunay triangulation. If \overline{pq} is an edge of the convex hull of S we call the supporting outer halfplane H not containing S an *infinite triangle* with edge \overline{pq} . Its circumcircle is H itself, the limit of all circles through p and q whose center tend to infinity within H . As a consequence, each edge of a Delaunay triangulation is now adjacent to two triangles.

Those triangles of DT_{i-1} whose circumcircles contain the new site, p_i , are said to be *in conflict* with p_i . According to the (equivalent) definition of the DT_i , they will no longer be Delaunay triangles.

Let \overline{qr} be an edge of DT_{i-1} , and let $T(q, r, t)$ be the triangle adjacent to \overline{qr} that lies on the other side of \overline{qr} than p_i ; see Figure 44. If its circumcircle $C(q, r, t)$ contains p_i then each circle through q, r contains at least one of p_i, t . Consequently, \overline{qr} cannot belong to DT_i , due to the (equivalent) definition. Instead, $\overline{p_i t}$ will be a new Delaunay edge, because there exists a circle contained in $C(q, r, t)$ that contains only p_i and t in its interior or boundary. This process of replacing edge \overline{qr} by $\overline{p_i t}$ is called an *edge flip*.

The necessary edge flips can be carried out efficiently if we know the triangle $T(q, s, r)$ of DT_{i-1} that contains p_i , see fig. Figure 45. The line segments connecting p_i to q, r , and s will be new Delaunay edges, by the same argument from above.




Figure 44: If triangle $T(q, r, t)$ is in conflict with p_i then former Delaunay edge \overline{qr} must be replaced by $\overline{p_i t}$.

Next, we check if e. g. edge \overline{qr} must be flipped. If so, the edges \overline{qt} and \overline{tr} are tested, and so on. We continue until no further edge currently forming a triangle with, but not containing p_i , needs to be flipped, and obtain DT_i .

Two tasks have to be considered:

1. Find the triangle of DT_{i-1} that is in conflict with p_i .
2. Perform all flips starting from this triangle.

It can be shown that the second task is bounded by the degree of p_i in the new triangulation. If the triangle of DT_{i-1} containing p_i is known, the structural work needed for computing DT_i from DT_{i-1} is proportional to the degree d of p_i in DT_i .

So we yield an obvious $O(n^2)$ time algorithm for constructing the Delaunay triangulation of n points: we can determine the triangle of DT_{i-1} containing p_i within linear time, by inspecting all candidates. Moreover, the degree of p_i is trivially bounded by n .

The last argument is too crude. There can be single vertices in DT_i that do have a high degree, but their *average* degree is bounded by 6.

With a special implementation using a directed acyclic graph (DAG), also called *Delaunay tree* due to Boissonnat and Teillaud,¹² we can detect the triangles of DT_{i-1} which are in conflict with p_i in $O(\log i)$ expected time.

Altogether we get the following result:

Theorem 7 *The Delaunay triangulation of a set of n points in the plane can be easily incrementally constructed incrementally in expected time $O(n \log n)$, using expected linear space. The average is taken over the different orders of inserting the n sites.*

5.3 Generalization of the Voronoi Diagram

5.3.1 Voronoi Diagram and Delaunay Triangulation in 3D

We will see that incremental construction is also appropriate for the 3D case. The following description was adapted from Aurenhammer and Klein.⁵

Let S be a set of n point sites in 3D. The *bisector* of two sites $p, q \in S$ is the perpendicular plane through the midpoint of the line segment \overline{pq} . The region $VR(p, S)$ of a site $p \in S$ is the intersection of halfspaces bounded by bisectors, and thus is a 3-dimensional convex polyhedron. The boundary of $VR(p, S)$ consists of *facets* (maximal subsets within the same bisector), of *edges* (maximal line segments in the boundary of facets), and of *vertices* (endpoints of edges). The regions, facets, edges, and vertices of $V(S)$ define a *cell complex* in 3D.

This cell complex is face-to-face: if two regions have a non-empty intersection f , then f is a face (facet, edge, or vertex) of both regions. As an appropriate data structure for storing a 3-dimensional cell complex we mention the facet-edge structure in Dobkin and Laszlo.²⁷

Complexity: The number of facets of $VR(p, S)$ is at most $n - 1$, at most one for each site $q \in S \setminus \{p\}$. Hence, by the Eulerian polyhedron formula, the number of edges and vertices of $VR(p, S)$ is $O(n)$, too. This shows that the total number of components of the diagram $V(S)$ in 3D is $O(n^2)$. In fact, there are configurations S that force each pair of regions of $V(S)$ to share a facet, thus achieving their maximum possible number of $\binom{n}{2}$; see, e.g., Dewdney and Vranich.²⁵ This fact sometimes makes Voronoi diagrams in 3D less useful compared to 2-space. On the other hand, Dwyer²⁸ showed that the *expected* size of $V(S)$ in d -space is only $O(n)$, provided S is drawn uniformly at random in the unit ball. This result indicates that high-dimensional Voronoi diagrams will be small in many practical situations.

In analogy to the 2-dimensional case, the Delaunay triangulation $DT(S)$ in 3D is defined as the geometric dual of $V(S)$. It contains a tetrahedron for each vertex, a triangle for each edge, and an edge for each facet, of $V(S)$. Equivalently, $DT(S)$ may be defined using the empty sphere property, by including a tetrahedron spanned by S as Delaunay iff its circumsphere is empty of sites in S . The circumcenters of these empty spheres are just the vertices




Figure 45: Updating DT_{i-1} after inserting the new site p_i . In (ii) the new Delaunay edges connecting p_i to q, r, s have been added, and edge \overline{qr} has already been flipped. Two more flips are necessary before the final state shown in (iii) is reached.




Figure 46: The Voronoi diagram of a point set in L_1 and L_2 . Note, that there are structural differences.




Figure 47: An Euclidean Voronoi diagram of line segments.

of $V(S)$. $DT(S)$ is a partition of the convex hull of S into tetrahedra, provided S is in general position. Note that the edges of $DT(S)$ may form the complete graph on S .

Simple incremental construction: Among the various proposed methods for constructing $V(S)$ in 3D, *incremental insertion* of sites (compare Section 5.2) is most intuitive and easy to implement. Basically, two different techniques for integrating a new site p into $V(S)$ have been applied. The more obvious method first determines all facets of the region of p in the new diagram, $V(S \cup \{p\})$, and then deletes the parts of $V(S)$ interior to this region; see e.g. Watson,⁹³ Field,³³ and Tanemura et al.⁸⁶ Inagaki et al.⁴⁷ describe a robust implementation of this method.

In the dual environment, this amounts to detecting and removing all tetrahedra of $DT(S)$ whose circumspheres contain p , and then filling the 'hole' with empty-sphere tetrahedra with p as apex, to obtain $DT(S \cup \{p\})$. An example of an edge flip in 3D is shown in Figure 48. Joe,⁴⁹ Rajan,⁷⁸ and Edelsbrunner and Shah³⁰ follow a different and numerically more stable approach. Like in the planar case, after having added a site to the current Delaunay triangulation, certain flips changing the local tetrahedral structure are performed in order to achieve local "Delaunayhood". The existence of such a sequence of flips is less trivial, however. Joe⁴⁸ demonstrated that no flipping sequence might exist that turns an arbitrary tetrahedral triangulation for S into $DT(S)$.

A complete algorithm with run time $O(n^2)$ can be found in Shah.³⁰




Figure 48: Two-into-three tetrahedra flip for five sites.

5.3.2 Constrained Voronoi diagrams

For many applications it is not useful to consider neighborhood relations without any constraints. For example, a small distance between two cities may become unimportant if there is a (natural or artificial) border line between them and the border line itself should not serve as a new site. To overcome such problems the concept of constrained Voronoi diagrams was introduced.⁶¹

Let S be a set of n points in the plane, and let L be a set of non-crossing line segments spanned by S , thus we have $|L| \leq 3n - 6$. The segments in L may be considered as obstacles. The *bounded distance* between two points x and y in the plane is defined as follows:

$$b(x, y) = \begin{cases} d(x, y) & \text{if } \overline{xy} \cap L = \emptyset \\ \infty & \text{otherwise} \end{cases}$$

where d stands for the Euclidean distance. In the resulting *constrained Voronoi diagram* $VD(S, L)$, regions of sites that are close but not visible from each other are separated by segments in L , an example is shown in Figure 49.

The exact dual of $VD(S, L)$ may be no longer a full triangulation of S (even if S is included) but we can modify $VD(S, L)$ in order to dualize it into a "near to Delaunay" triangulation $DT(S, L)$ that also includes L .

For every line segment l we proceed as follows: All sites of the clipped regions to the left of a segment l build a special Voronoi diagram to the right and vice versa. The neighborhood within these special diagrams lead to additional edges inserted into the dual of $VD(S, L)$. Note that the endpoints of the line segments are always neighbors in the new diagrams and therefore the line segments itself are inserted (see Figure 49).

Algorithms for computing the constrained Voronoi diagram $VD(S, L)$ and the constrained Delau-

nay triangulation $DT(S, L)$ have been proposed in Lee and Lin,⁶¹ Chew,¹⁶ Wang and Schubert⁹² and Wang.⁹¹ In Seidel⁸¹ and Kao and Mount⁵¹ one will find good implementation schemes. An application of $DT(S, L)$ to quality mesh generation can be found in Chew.¹⁷

5.3.3 Other Types of Generalizations


We simply list some of the generalization schemes and show examples of some intuitive ones.

- Different metrics
 - L_1 , a comparison of L_1 and L_2 is shown in Figure 46
 - L_∞
 - Convex distance functions
- Different space
 - On trees and graphs
 - Higher dimensions
- Weights
- More general sites
 - Line segments (see Figure 47).
 - Polygonal chains
- Farthest point Voronoi diagram
- K-th order Voronoi diagram
- Colored objects

5.4 Applications of the Voronoi Diagram

5.4.1 Nearest Neighbor or Post Office Problem


We consider the well-known post office problem. For a set S of sites in the plane and an arbitrary query point x we want to compute the point of S closest to x efficiently.

Figure 49: The constrained Voronoi diagram $VD(S, L)$ and its dual.

In the field of computational geometry there is a general technique for solving such query problems. One tries to decompose the query set into classes so that every class has the same answer. Now for a single answer we only have to determine its class. This technique is called *locus approach*.

The Voronoi diagram represents the locus approach for the post office problem. The classes correspond to the regions of the sites. For a query point x we want to determine its class/region and return its owner.

To solve this task a simple technique can be applied. We draw a horizontal line through every vertex of the diagram and sort the lines in $O(n \log n)$ time, see Figure 50. The lines decompose the diagram into *slabs*. For every *slab* we sort the set of crossing edges of the Voronoi diagram in linear time. Altogether we need $O(n^2)$ time for the simple construction.

Figure 50: After constructing the slabs, a query point x can be located quickly.

For a query point x we locate its *slab* in $O(\log n)$ time and afterwards its region in $O(\log n)$ time by binary search.

Theorem 8 Given a set S of n point sites in the plane, one can, within $O(n^2)$ time and storage, construct a data structure that supports nearest neighbor queries: for an arbitrary query point x , its nearest neighbor in S can be found in time $O(\log n)$.

The simple technique can be easily extended to 3D. There are also more efficient approaches, i.e., Edelsbrunner²⁹ constructs a $O(\log n)$ search structure for the Voronoi diagram in linear time and with linear space.

5.4.2 Motion planning

We present a simple application of Voronoi diagrams for computing collision-free paths in polygonal environments. Let us assume that a two dimensional environment with polygonal obstacles is given. For simplicity we assume that a circle-shaped robot is given and a collision-free path from s to t should be computed (see Figure 51).

We can solve this problem by computing the Voronoi diagram $VD(S)$ of the corresponding line segments. This gives rise to a reasonable solution. Let us assume that the robot moves between to obstacle line segments l_1 and l_2 , for safety it should maximize its distance to both segments. That is, for every position x we want to maximize

$$|xl_i| = \min_{y \in l_i} |xy|, \quad i = 1, 2.$$

This goal is reached if the robot moves along the bisector of l_1 and l_2 . Altogether the robot should try to move along the Voronoi diagram of the line segments of the polygonal scene. For every segment of $VD(S)$ we can easily compute whether the robot fits between the corresponding obstacles.




Figure 51: Computing a collision-free path from s to t for a circle-shaped robot in the presence of polygonal obstacles.

Now we can solve our problem as follows:

- Compute the Voronoi diagram $VD(S)$ of the set of line segments S .
 - Consider the Voronoi regions of s and t in $VD(S)$.
 - Consider the closest point $s' \in VD(S)$ to s , respectively $t' \in VD(S)$ to t .
 - Delete the segments of $VD(S)$ where the robot does not fit between the obstacles.
 - Consider the graph G of the remaining diagram together with s' and t' .
 - Compute a path from s' to t' in G by standard graph algorithms.
- Nearest Neighbor Search
 - $O(n)$ for all nearest neighbors of the sites
 - $O(k \log^2 n)$ expected time for k -th nearest neighbors of query point x
 - Minimum Spanning Tree and TSP-Heuristic, $O(n \log n)$
 - Largest empty circle, $O(n)$
 - Smallest enclosing circle (square with fixed orientation), $O(n)$
 - Smallest color spanning circle (square with fixed orientation), $O(nk)$, where k is the number of colors
 - Localization problems, see Hamacher⁴²
 - Clustering of objects, see Dehne and Noltmeier²⁴

Altogether we have the following result:

Theorem 9 *There is a collision-free path of a circle-shaped robot with radius r in a polygonal environment of a set line segments S if and only if the radius r is smaller than $|ss'|$ and $|tt'|$ and a collision-free movement from s' to t' along the segments of $VD(S)$ exists.*

5.4.3 Other Applications of the Voronoi Diagram in 2D

There are many different geometrical applications of the Voronoi diagram and its dual. Here we simply list some of them, together with some performance results, provided that the diagram is given:

- Closest Pair of sites, $O(n)$

All these results stem more or less from the linear complexity of the diagram. As we have already mentioned the complexity of the diagrams in three dimension is also linear in many practical situations. Thus many of the presented problems can be solved in three dimensions with almost the same time bound. We will present some special applications for 3D.

5.5 Texture Synthesis

Textures are visual detail of the rendered geometry, which have become very important in the past few years, because the cost of rendering with texture is the same as the cost without texture. Virtually all




Figure 52: The texture synthesis algorithm proceeds in scan line order through the texture and considers only the neighborhood around the current pixel as shown.




Figure 53: Using an image pyramid, the texture synthesis process becomes fairly robust against different scales of detail in the sample images.



Figure 54: Some results of the texture synthesis algorithm.⁹⁵ In each pair, the image on the left is the original one, the one on the right is the (partly) synthesized one.

real-world objects have texture, so it is extremely important to render them in synthetic worlds, too.

Texture synthesis generally tries to synthesize new textures, either from given images, from a mathematical description, or from a physical model. Mathematical descriptions can be as simple as a number of sine waves to generate water ripples, while physical models try to describe the physical or biological effects and phenomena that lead to some texture (such as patina or fur). In all of these “model-based” methods, the knowledge about the texture is in the model and the algorithm. The other class of methods starts with one or more images; then they try to find some statistical or stochastic description (explicitly or implicitly) of these, and finally it generates a new texture from the statistic.

Basically, textures are images with the following properties:

1. Stationary: if a window with the proper size is moved about the image, the portion inside the window always appears the same.
2. Local: each pixel’s color in the image depends only on a relatively small neighborhood.

Of course, images not satisfying these criteria can be used as textures as well (such as façades), but if you want to synthesize such images, then a statistical or stochastic approach is probably not feasible.

In the following, we will describe a stochastic algorithm that is very simple, very efficient, and works remarkably well.⁹⁵ Given a sample image, it does not, like most other methods, try to compute explicitly the stochastic model. Instead, it uses the sample image itself, which implicitly contains that model already.

We will use the following terminology:

- I = Original (sample) image
- T = New texture image
- p_i = Pixel from I
- p = Pixel from T to be generated next
- $N(p)$ = Neighborhood of p (see Figure 52)

Initially, T is cleared to black. The algorithm starts by adding a suitably sized border at the left and the top, filled with random pixels (this will be thrown away again at the end). Then, it performs the following simple loop in scan line order (see Figure 52):

- 1: for all $p \in T$ do
- 2: find the $p_i \in I$ that minimizes $|N(p) - N(p_i)|^2$
- 3: $p := p_i$
- 4: end for

Well, the search in line 2 is exactly a nearest-neighbor search! This can be performed efficiently with the algorithm presented in Section 5.4.1: if $N(p)$ contains k pixels, then the points are just $3k$ -dimensional vectors of RGB values, and the distance is just the Euclidean distance.

Obviously, all pixels of the new texture are deterministically defined, once the random border has been filled. The shape of the neighborhood $N(p)$ can be chosen arbitrarily, it must just be chosen such that all but the current pixel are already computed. Likewise, other “scans” of the texture are possible and sensible (for instance a spiral scan order), they must just match the shape of $N(p)$.

The quality of the texture depends on the size of the neighborhood $N(p)$. However, the optimal size itself depends on the “granularity” in the sample image. In order to make the algorithm independent, we can synthesize an image pyramid (see Figure 53). First, we generate a pyramid I^0, I^1, \dots, I^d for the sample image I^0 . Then, we synthesize the texture pyramid T^0, T^1, \dots, T^d level by level with the above algorithm, starting at the coarsest level. The only difference is that we extend the neighborhood $N(p)$ of a pixel p over k levels as depicted by Figure 53. Consequently, we have to build a nearest-neighbor search structure for each level, because as we proceed downwards in the texture pyramid, the size of the neighborhood grows.

Of course, now we have replaced the parameter of the best size of the neighborhood by the parameter of the best size per level and the best number of levels to consider for the neighborhood. However, as Wei and Levoy⁹⁵ report, a neighborhood of 9×9 (at the finest level) across 2 levels seems to be sufficient in almost all cases.

Figure 54 shows two examples of the results that can be achieved with this method.

5.6 Shape Matching

As the availability of 3D models on the net and in databases increases, searching for such models becomes an interesting problem. Such a functionality is needed, for instance, in medical image databases, or CAD databases. One question is how to specify a query. Usually, most researchers pursue the “query

by content” approach, where a query is specified by providing a (possibly crude) shape, for which the database is to return best matches.¹⁰ The fundamental step here is the matching of shapes, i.e., the calculation of a *similarity* measure.

Almost all approaches perform the following steps:

1. Define a transformation function that takes a shape and computes a so-called *feature vector* in some high dimensional space, which (hopefully) captures the shape in its essence. Naturally, those transformation functions are preferred that are invariant under rotation and/or translation and tessellation.
2. Define a *similarity measure* d on the feature vectors, such that if $d(f_1, f_2)$ is large, then the associated shapes s_1, s_2 do not *look* similar. Obviously, this is (partly) a human factors issue. In almost all algorithms, d is just the Euclidean distance.
3. Compute a feature vector for each shape in the database and store them in a data structure that allows for fast nearest-neighbor search.
4. Given a query, i.e., a shape, compute its feature vector, and retrieve the nearest neighbor from the database. Usually, the system also retrieves all k nearest neighbors. Often times, you are not interested in the exact k nearest neighbors but only in *approximate* nearest neighbors (because the feature vector is an approximation of the shape anyway).

The main difference among most shape matching algorithms is, therefore, the transformation from shape to feature vector.

So, fast shape retrieval essentially requires a fast (approximate) nearest neighbor search. We could stop our discussion of shape matching here, but for sake of completeness, we will describe a very simple algorithm (from the plethora of others) to compute a feature vector.⁷¹

The general idea is to define some *shape function* $f(P_1, \dots, P_n) \rightarrow \mathbb{R}$, which computes some geometrical property of a number of points, and then evaluate this function for a large number of random points that lie on the surface of the shape. The resulting distribution of f is called a *shape distribution*.

¹⁰This idea seems to originate from image database retrieval, where it was called QBIC = “query by image content”.




Figure 55: The shape distribution of a number of different simple objects.

For the shape function, there are a lot of possibilities (your imagination is the limit). Examples are:

- $f(P_1, P_2) = |P_1 - P_2|$;
- $f(P_1) = |P_1 - P_0|$, where P_0 is a fixed point, such as the bounding box center;
- $f(P_1, P_2, P_3) = \angle(\overline{P_1P_2}, \overline{P_1P_3})$;
- $f(P_1, P_2, P_3, P_4) = \text{volume of the tetrahedron between the four points.}$

Figure 55 shows the shape distributions of a few simple objects with the distance between two points as shape function.

6 Distance Fields

Distance fields can be viewed as a way to represent surfaces (and solids). They are a very powerful data structure, because they contain a huge amount of information (compared to just the surface itself). This is why they usually take fairly long to compute. Consequently, this computation can usually be performed only as a preprocessing step. Thus, they are difficult to adapt for deformable geometry.

Definition 5 (Distance field)

Let S be a surface in \mathbb{R}^3 . Then, the distance field of a surface S is a scalar function $D_S : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that for all $p \in \mathbb{R}^3$,

$$D_S(p) = \text{sgn}(p) \cdot \min\{d(p, q) | q \in S\}, \quad (14)$$

where

$$\text{sgn}(p) = \begin{cases} -1, & \text{if } p \text{ inside} \\ +1, & \text{if } p \text{ outside} \end{cases}$$

In other words, D_S tells for any point p the distance to the closest point on the surface S .

We can augment the DF further to obtain a *vector distance field* V_S by storing a vector to the closest point $p \in S$ for each point x .⁵⁰ So, $D_S(x) = |V_S(x)|$. This is a vector field. Another vector field that is often used in the context of DFs is the *gradient of the distance field*, or just *gradient field*.

Figure 56 shows a distance field for a simple polygon in the plane. The distance of a point from the surface is color-coded (the distance for points inside the polygon is not shown). Figure 57 shows a vector distance field (for the same surface).

Apparently, distance fields (DFs) have been “invented” in many different areas, such as computational physics,^{82,84} robotics,⁵⁴ GIS (see Figures 58 and 59 for an example), and image processing.⁷³ Not surprisingly, DFs come under many other names, such as *distance maps* and *potential fields*. The process of, or the algorithm for computing a DF is sometimes called *distance transform*.¹¹

Distance fields have close relationships to isosurfaces and implicit functions. When regarded as an implicit function, the isosurface for the iso-value 0 of a DF is exactly the original surface (see Figure 56). However, the converse is not true in general, i.e., the DF of a surface defined by an implicit function is not necessarily identical to the original implicit function.

There is another data structure related to distance fields, namely Voronoi diagrams (see Section 5). Given a vector distance field for a set of points, edges, and polygons (not necessarily connected), then all points in space, whose vectors point to the same feature (point, edge, or polygon), are in the same Voronoi cell (see Figures 56 and 57). (We could also regard the vector of the DF as a kind of ID of the respective Voronoi cell.)

6.1 Computation and representation of DFs

For special surfaces S , we may be able to compute D_S analytically. In general, however, we have to discretize D_S spatially, i.e., store them in a 3D voxel grid, octree, or other space partitioning data structure. Voxel grids and octrees are the most commonly used data structures for storing DFs, and

¹¹Sometimes, this term bears the connotation of producing *inaccurate* distance fields.

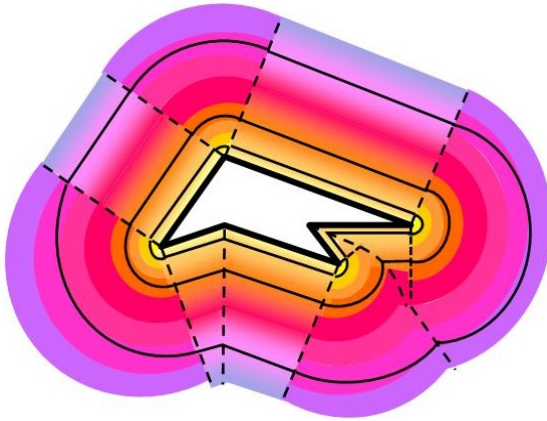


Figure 56: Example of a distance field in the plane for a simple polygon (thick black lines). The distance field inside the polygon is not shown. The dashed lines show the Voronoi diagram for the same polygon. The thin lines show isosurfaces for various isovalues.

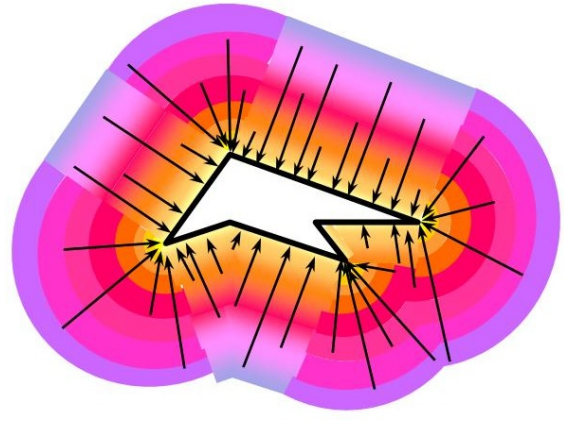


Figure 57: The vector distance field for the same polygon as shown on the left. Only a few vectors are shown, although (in theory) every point of the field has a vector.




Figure 58: A network of roads described in the plane by a set of edges.




Figure 59: The distance map of these roads. The distance of each point in the plane from a road is color coded. It could be used, for example, to determine the areas where new houses can be built.

we will describe algorithms for them in the following. More sophisticated representations try to store more information at each voxel in order to be able to extract distances quickly from the field.⁴⁵

Since each cell in a discrete distance field stores one signed distance to the surface for only one “representative”, the distance of other points must be interpolated from these values. One simple method is to store exact distances for the nodes (i.e., the corners of the voxels), and generate all other distances in the interior of voxels by trilinear in-

terpolation. Other interpolation methods could, of course, be used as well.

The simplest discrete representation of a DF is the 3D voxel grid. However, for most applications this is too costly not only in memory utilization but also in computational efforts. So, usually DFs are represented using octrees (see Figure 60), because they are simple to construct and offer fairly good adaptivity and allow for easy algorithms.³⁵ This representation is called an *adaptively sampled distance field* (ADF). Actually, the kind of hierarchical




Figure 60: Representing distance fields by octrees has both memory and computational advantages.

space partitioning is a design parameter of the algorithm; for instance, BSPs (see Section 3) probably offer much more efficient storage at the price of more complex interpolation.

Constructing ADFs works much like constructing regular octrees for surfaces, except that here, we subdivide a cell if the distance field is not well approximated by the interpolation function defined so far by the cell corners.

The following two algorithms for computing DFs produce only flat representations (i.e., grids), but they can be turned into an ADF by coalescing cells in a bottom-up manner if the resulting cell still describes the DF well enough.¹²

6.1.1 Propagation method

This method bears some resemblance to region growing and flood filling algorithms and is also called *chamfer method*. It can produce only approximations of the exact DF.

The idea is to start with a binary DF, where all voxels intersecting the surface are assigned a distance of 0, and all others are assigned a distance of ∞ . Then, we somehow “propagate” these known values to voxels in the neighborhood, taking advantage of the fact that we already know the distances between neighboring voxels.¹³ Note that these distances only depend on the local neighborhood constellation, no matter where the “propagation front” currently is.

¹²Usually, it is faster to compute an ADF top-down, though.

¹³This idea has been proposed as early as 1984,¹³ and maybe even earlier.

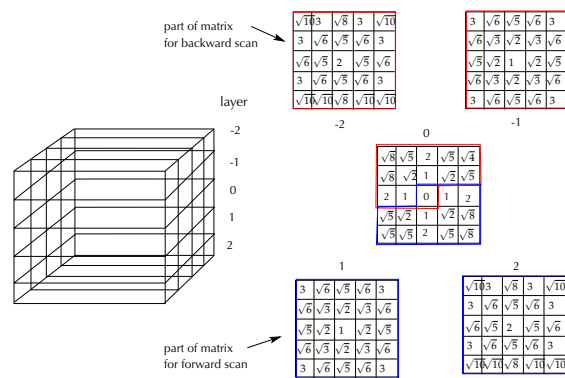


Figure 61: By convoluting a “distance matrix” with an initially binary distance field (0 or ∞), one gets an approximate distance field. Here a $5 \times 5 \times 5$ example of such a matrix is shown.

More formally speaking, since we want to compute the DF only for a discrete set of points in space, we can reformulate Equation 14 as follows

$$\tilde{D}_S(x, y, z) = \min_{i, j, k \in \mathbb{Z}^3} \{ D(x + i, y + j, z + k) + d_M(i, j, k) \} \quad (15)$$

where d_M is the distance of a node (i, j, k) from the center node $(0, 0, 0)$. Actually, this is already a slight approximation of the exact DF.

We can now further approximate this by not considering the infinite “neighborhood” $(i, j, k) \in \mathbb{Z}^3$, but instead only a local one $\mathbb{I} \subset \mathbb{Z}$. Since a practical computation of \tilde{D}_S would compute each node in some kind of scan order, we choose \mathbb{I} such that it contains only nodes that have already been computed by the respective scan. Thus, d_M can be pre-computed and stored conveniently in a 3-dimensional matrix (see Figure 61).

This process looks very similar to conventional convolution, except that here we perform a minimum operation over a number of sums, while the conventional convolution performs a sum over a number of products.

Obviously, a single scan will not assign meaningful distance values to all nodes, no matter what order the scan performs. Therefore, we need at least two scans: for instance, first, a 3D scanline order from top to bottom (“forward”), and second, backwards from bottom to top. In each pass, only one half of the 3D matrix d_M needs to be considered (see Figure 61). It could still happen, of course, that there are still nodes with distance ∞ in the respective part of the chamfer matrix.

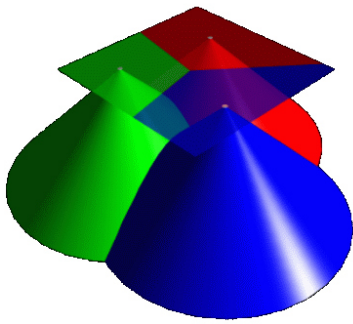


Figure 62: The distance function of a point site in the plane is a cone.⁴⁴




Figure 63: More complex sites have a bit more complex distance functions.




Figure 64: The distance function of sites in 3D is different for the different slices of the volume.

The time complexity of this method is in $O(m)$, m = number of voxels. However, the result is only an approximation of the distance field; in fact, the accuracy of the distance values can be pretty low.¹⁴

In order to further improve accuracy, one can make a few more passes, not only from top to bottom and back, but also from left to right, etc. Another option is to increase the size of d_M .

Analogously, an approximate vector distance field can be computed.⁵⁰ In that case, the matrix d_M contains vectors instead of scalars. In order to improve the accuracy, we can proceed in two steps: first, for each voxel in the neighborhood of the surface (usually just the 3^3 -neighborhood) we compute vectors to the *exact* closest point on the surface; then, this *shell* is propagated by several passes with the vector-valued matrix d_M . The time for computing the exact distance shell can be optimized by utilizing an octree or BV hierarchy for closest point computation, and initializing each search with the distance (and vector) to the closest point of a neighbor voxel.

6.1.2 Projection of distance functions

The previous method can be applied to all kinds of surfaces S . However, it can produce only very gross approximations. In the following, we describe a method that can produce much better approximations. However, it can be applied only to polygonal surfaces S .

The key technique applied in this approach is to embed the problem into more dimensions than are inherent in the input/output itself. This is a very general technique that often helps to get a better and simpler view on a complex problem.

The method described in the following was presented by Hoff et al.,⁴⁴ and, from a different point of view, by Sethian.⁸³

Let us consider the distance field of a single point in 2D. If we consider this a surface in 3D, then we get a cone with the apex in that point. In other words, the distance field of a point in a plane is just the orthogonal projection of suitable cone onto the plane $z = 0$.

Now, if there are n point sites in the plane, we get n cones, so each point in the plane will be “hit” by n distance values from the projection (see Figure 62). Obviously, only the smallest one wins and gets “stored” with that point — and this is exactly what the z -buffer of graphics hardware was designed for.

These cones are called the *distance function* for a point. The distance function for other features (line segments, polygons, curves) is a little bit more complicated (see Figure 63).

So, computing a discrete DF for a set of sites in the plane can be done by rendering all the associated distance functions (represented as polygonal meshes) and reading out the z -buffer (and possibly the frame buffer if we also want the site IDs, i.e., discretized Voronoi regions). If we want to compute a three-dimensional DF, then we proceed slice-by-slice. One noteworthy catch, though, is that the distance functions of the sites *changes* from slice to slice. For instance, the distance function for a point

¹⁴Furthermore, the distance values are not even an upper bound on the real distance.

that lies *not* in the current slice is a hyperboloid, with the point site being coincident with the vertex of the hyperboloid (see Figure 64).

6.2 Applications of DFs

Due to their high information density, distance fields have a huge number of applications. Among them are motion planning in robotics,^{54,59} collision detection, shape matching,⁷⁰ morphing,¹⁹ volumetric modeling,^{14,35,97} navigation in virtual environments,⁹⁰ reconstruction,⁵⁶ offset surface construction,⁷⁶ and dynamic LODs, to name but a few. In the following, we will highlight two easy ones of them.

6.2.1 Morphing

One interesting application of DFs is morphing, i.e., the problem of finding a “smooth” transition from a shape $S \subset \mathbb{R}^d$ to a shape $T \subset \mathbb{R}^d$, i.e., we are looking for a shape transformation $M(t, \mathbb{R}^d)$, $t \in [0, 1]$, such that $M(0, S) = S$, $M(1, S) = T$.

Sometimes, there is a bit of confusion about terms: sometimes, morphing refers to any smooth transition, sometimes it refers only to those transitions that are bijective, continuous, and have a continuous inverse. In the latter case, the term morphing is equivalent to the notion of *homeomorphism* in the area of topology, and the term *metamorphosis* is used to refer to the broader definition. The difference between a homeomorphism and a metamorphosis is that the former does not allow to cut or glue the shape during the transition (this would change the topology).

In the following, we will describe a simple method for metamorphosis of two shapes given as volumetric models; it was presented by Cohen-Or et al.²⁰ The nice thing about a volumetric representation is that it can naturally handle genus changes (i.e., the number of “holes” in S and T is different).

In the simplest form, we can just interpolate linearly between the two DFs of S and T yielding a new DF

$$M(t, S) = J(t, S, T) = tD_S + (1 - t)D_T. \quad (16)$$

In order to obtain a polygonal representation, we can compute the isosurface for the value 0.

This works quite well, if S and T are “aligned” with each other (in an intuitive sense). However, when morphing two sticks that are perpendicular




Figure 65: By specifying correspondences, the user can identify feature points that should transition into each other during the morph.²⁰

to each other, there will be in-between models that almost vanish.

Therefore, in addition to the DF interpolation, we usually want to split the morphing into two steps (for each t): first, S is warped by some function $W(t, x)$ into S' , then $D_{S'}$ and D_T are interpolated, i.e., $M(t, S) = tD_{W(t,S)} + (1 - t)D_T$.

The idea of the warp is that it encodes the expertise of the user who usually wants a number of prominent feature points from S to transition onto other feature points on T . For instance, she wants the corresponding tips of nose, fingers, and feet of two different characters transition into each other.

Therefore, the user has to specify a small number of *correspondences* $(p_{0,i}, p_{1,i})$, where $p_{0,i} \in S$, and $p_{1,i} \in T$ (see Figure 65). Then, we determine a warp function $W(t, x)$ such that $W(1, p_{0,i}) = p_{1,i}$.

Since the warp function should distort S as little as possible (just “align” them), we can assemble it from a rotation, then a translation, and finally an *elastic* part that is not a linear transformation:

$$W(t, \mathbf{x}) = ((1 - t)I + tE)(R(\mathbf{a}, t\theta)\mathbf{x} + t\mathbf{c}) \quad (17)$$

where E is the elastic transformation, $R(\mathbf{a}, \theta)$ is a rotation about axis \mathbf{a} and angle θ , and \mathbf{c} is the translation.

The rotation and translation of the warp function can be determined by least-squares fitting, while the elastic transformation can be formulated as a




Figure 66: Example of difference of two DFs. On the left, the two DFs for tool and object are shown superimposed; on the right, the result is shown.




Figure 67: Example of a 3D difference operation.¹⁴

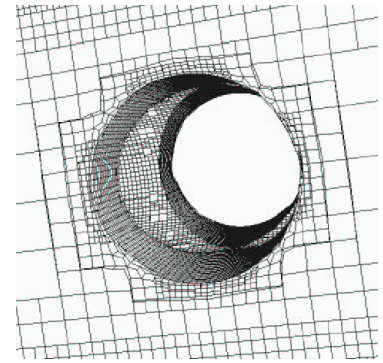


Figure 68: Close-up view showing the adaptivity.

scattered data interpolation problem. One method that has proven quite robust is the approach using radial basis functions. The interested reader will find the details in Cohen-Or et al.²⁰

6.2.2 Modeling

A frequently used modeling paradigm in CAD is *volumetric modeling*, which means that objects are, conceptually, defined by specifying for each point whether or not it is a member of the object. the advantage is that it is fairly simple to define Boolean operations on such a representation, such as union, subtraction, and intersection. Examples of volumetric modeling is *constructive solid geometry* (CSG) and *voxel-based modeling*. The advantage of using distance fields for volumetric modeling is that they are much easier to implement than CSG, but provide much better quality than simple voxel-based approaches.

Computing the Boolean operation on two DFs basically amounts to evaluating the appropriate operator on the two fields. The following tables gives these operators for a field D_O (for the object) and a field D_T (for the tool):

Boolean op.	operator on DF
Union	$D_{O \cup T} = \min(D_O, D_T)$
Intersection	$D_{O \cap T} = \max(D_O, D_T)$
Difference	$D_{O - T} = \max(D_O, -D_T)$

Figure 66 shows an example of the difference operation.

If ADFs are used for representation of the DFs, then a recursive procedure has to be applied to the object's ADF in order to compute the resulting ADF. It is similar to the generation of ADFs (see above), except that here the subdivision is governed by the compliance of new object ADF with the tool ADF: a cell is further subdivided if it contains cells from the tool ADF that are smaller, and if the trilinear interpolation inside the cell does not approximate well enough the ADF of the tool.

7 Dynamization of Geometric Data Structures

We present a generic approach for the dynamization of an *arbitrary* static geometric data structure. Often a simple *static* data structure is sufficient if the set of represented geometric objects will have few changes over time. Once created, the static structure mostly has to cope with data queries due to its geometric intention. If the set of objects varies very much over time, there is need for more complex dynamic structures which allow efficient *insertion and deletion* of objects.

For example a one dimensional sorted array of a fixed set M is sufficient for *x is Element of M* queries. But if the set M has many changes over time, a dynamic AVL-tree would be more likely. The AVL-tree implementation is more complex

since rotations of the tree has to be considered for insertion and deletion of objects. Additionally, the AVL-tree dynamization was *invented* for the special case of the one-dimensional search. We want to show that it is possible to dynamize a simple static data structure indirectly but also efficiently in a general setting. Once this generic approach is implemented it can be used for many static data structures.

The generic approaches presented here are not optimal against a dynamization adapted directly to a single data structure, but they are easy to implement and efficient for many applications.

In Section 7.1 we formalize the given problem and define some requirements. In Section 7.2 we present methods allowing insertion and deletion in amortized efficient time. For many applications this is already efficient enough. Within this section the dynamization technique is explained in detail and the amortized cost of the new operations are shown. Similar ideas for the worst-case sensitive approach are sketched in Section 7.3. The *effort* of the dynamization itself is amortized over time. For details see Klein⁵⁵ or the work of Overmars⁷² and van Kreveld.⁸⁹ We present a simple example in Section 7.4.

7.1 Model of the Dynamization

Let us assume that *TStat* is a static abstract (geometric) data type. Since we have a geometric data structure, we assume that the essential motivation of *TStat* is a query operation on the set of stored objects *D*. i.e., for a query object *q* the answer is always a subset of *D* which might be empty.

We want to define the generic dynamization by a module that imports the following operations from *TStat*:

- build(V, D)*: Build the structure *V* of type *TStat* with all data objects in the set *D*.
- query(V, q)*: Gives the answer (objects of *D*) to a query to *V* with query object *q*.
- extract(V, D)*: Collects all data objects *D* of *V* in a single set and returns a pointer to this set.
- erase(V)*: Delete the complete data structure *V* from the storage.

The dynamization module should export a dynamic abstract (geometric) data type *TDyn* with the following operations:

- Build(W, D)*: Build the structure *W* of type *TDyn* with data objects in the set *D*.
- Query(W, q)*: Gives the answer (objects of *D*) to a query to *W* with query object *q*.
- Extract(W, D)*: Collects all data objects *D* of *W* in a single set and returns a pointer to this set.
- Erase(W)*: Delete the complete data structure *W* from the storage.
- Insert(W, d)*: Insert object *d* into *W*.
- Delete(W, d)*: Delete *d* out of *W*.

Note, that the *new* operations *Delete* and *Insert* are necessary since we have a dynamic data type now.

Additionally, we introduce some cost functions for the operations of the abstract dynamic and the abstract static data type. For example let $B_V(n)$ denotes the time function for the operation *Build(V, D)* of *TStat*. The notations are fully presented in Figure 69. The cost functions depend on the implementation of *TStat*. Note, that the cost function of *TDyn* will depend on the cost functions of *TStat* together with the efficiency of the our general dynamization.

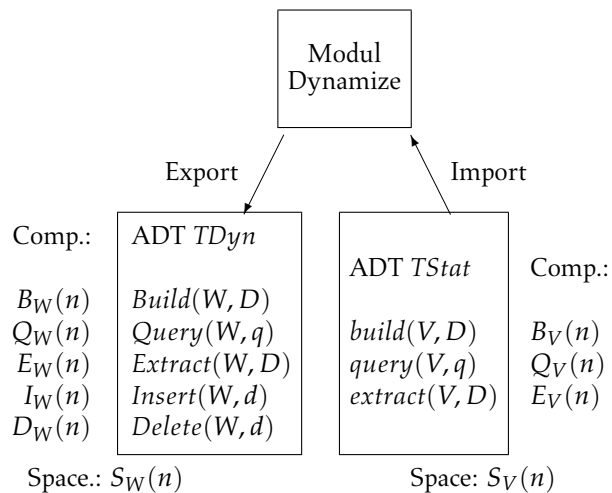


Figure 69: Dynamization in the generic sense.

In order to guarantee some bounds for the corresponding cost functions of *TDyn* the cost functions

of $TStat$ must not increase arbitrarily. On the other hand for the proof of some time bounds we need some kind of monotonic behavior in the functions, they should not oscillate. Altogether we define the following requirements which are fulfilled in many cases:

1. $Q_V(n)$ and $E_V(n)$ increase monotonically in n ; examples: $1, \log n, \sqrt{n}, n, n \log n, n^2, 2^n$.
2. $\frac{B_V(n)}{n}$ and $\frac{S_V(n)}{n}$ increase monotonically in n ; examples $n, n \log n, n^2, 2^n$.
3. For all $f \in \{Q_V, B_V, E_V, S_V\}$ there is a constant $C \geq 1$, so that $f(2n) \leq C f(n)$; examples: $1, \sqrt{n}, n, n^2$, and also $\log n$ with $n > 1$, as well as the products of this functions, but not 2^n .
4. $E_V(n) \leq 2 \cdot B_V(n)$.

Moreover, we assume that the that the *query* operation can be decomposed, i.e., for a decomposition $V = V_1 \cup V_2 \cup \dots \cup V_j$ of the data set V the results of the single operations $query(V_i, d)$ lead to the solution of $query(V, d)$. This is true for many kinds of range queries.

7.2 Amortized Insert and Delete

7.2.1 Amortized Insert: Binary Structure

The very first idea is to implement $Insert(W, d)$ and $Delete(W, d)$ directly by

$Insert(W, d) : Extract(W, D); Build(W, D \cup \{d\})$

$Delete(W, d) : Extract(W, D); Build(W, D \setminus \{d\})$.

This *throw-away* implementation is not very efficient. Therefore we distribute the n data objects of the static structure V among several structures V_i . If a new element has to be inserted we hope that only a single structure V_i may be concerned. Let

$$n = a_l 2^l + a_{l-1} 2^{l-1} + \dots + a_1 2 + a_0 \text{ mit } a_i \in \{0, 1\}.$$

Then $a_l a_{l-1} \dots a_1 a_0$ is the *binary representation* of n . For every $a_i = 1$ we build a structure V_i which has 2^i elements. The collection of these structures is a representation of W_n which is called *binary structure* (see Figure 70). To build up the binary structure W_n we proceed as follows:

$Build(W, D)$:

Compute *binary representation* of $n = |D|$.

Decompose D into sets D_i with

$$|D_i| = 2^i \text{ w.r.t. the representation of } n.$$

Compute $build(V_i, D_i)$ for every D_i .

In principle, the binary structure W_n can be constructed as quick as the corresponding structure V .

Lemma 10

$$B_W(n) \in O(B_V(n)).$$

Computing the binary representation of n and the decomposition into D_i can be done in linear time $O(n)$.

The operation $build(V_i, D_i)$ needs $B_V(2^i)$ time. We have $i \leq l = \lfloor \log n \rfloor$ and therefore we conclude:

$$\begin{aligned} \sum_{i=0}^{\lfloor \log n \rfloor} B_V(2^i) &= \sum_{i=0}^{\lfloor \log n \rfloor} 2^i \frac{B_V(2^i)}{2^i} \\ &\leq \sum_{i=0}^{\lfloor \log n \rfloor} 2^i \frac{B_V(n)}{n} \\ &\leq 2^{\log n} \frac{B_V(n)}{n} \\ &\in O(B_V(n)). \end{aligned}$$

We used the fact that $\frac{B_V(n)}{n}$ increases monotonically.

Altogether we have

$$B_W(n) \in O(n + B_V(n)) = O(B_V(n))$$

since $B_V(n)$ is at least linear.

Similar results hold for some other operations. We can prove

$$E_W(n) \leq \log n E_V(n)$$

the results of $extract(V_i)$ for at most $\log n$ structures V_i .

Additionally,

$$Q_W(n) \leq \log n Q_V(n)$$

holds if we assume that the query can be decompose as well, see the requirements.

It is also easy to see that

$$S_W(n) \leq \sum_{i=0}^{\lfloor \log n \rfloor} S_V(2^i) \in O(S_V(n)).$$

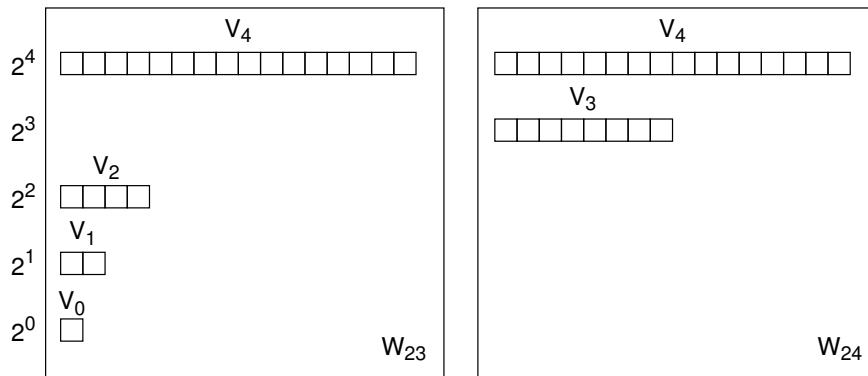


Figure 70: The binary structure W_n contains the structure V_i if $a_i = 1$ holds for the binary representation of n . For examples see $n = 23$ (left) and $n = 24$ (right).

Therefore it remains to analyse $I_W(n)$.

As we have seen from Figure 70 sometimes the whole structure of W_n is destroyed when W_{n+1} was build up. In this example we had to perform the following tasks:

$extract(V_0, D_0); extract(V_1, D_1); extract(V_2, D_2);$
 $D := D_0 \cup D_1 \cup D_2 \cup \{d\};$
 $build(V_3, D);$

In general we have to build up V_j and extract and erase $V_{j-1}, V_{j-2}, \dots, V_0$ only if $a_i = 1$ holds for $i = 0, 1, \dots, j-1$ and $a_j = 0$ holds (in the binary representation of the current n).

In this special case we have

$$\begin{aligned} I_W(n) &\leq \left(\sum_{i=0}^{j-1} E_V(2^i) \right) + Cj + B_V(2^j) \\ &\leq E_V(2^j) + Cj + B_V(2^j) \\ &\in O\left(B_V(2^j)\right). \end{aligned}$$

For a long sequence of insertions many of them are performed without extreme reconstructions. Thus the effort for all $Insert(W, d)$ is amortized over time.

Generally, let $Work$ be an arbitrary operation with cost function W . For a sequence of s different operations let $Work$ be applied k times. If

$$\frac{\text{total cost of } k \text{ } Work \text{ operationen}}{k} \leq \bar{W}(s)$$

holds for a monotonically increasing cost function \bar{W} , we say that the operation $Work$ is performed in *amortized time* $\bar{W}(s)$.

Note, that this is not an expected value and that \bar{W} is a function of s , i.e., the length of the operation sequence. The current data set may have a number of elements $n \leq s$.

For the $Insert(W, d)$ operation one can prove

$$\bar{I}_W(s) \in O\left(\frac{\log s}{s} B_V(s)\right).$$

Note, that except insertions there are only queries in s . Queries do not change the size of the data set and so we can also replace s by the number of insertions here.

Altogether, the results are presented in the following theorem.

Theorem 11 *A static abstract data type as presented in Figure 69 can be dynamized by means of the binary structure in a dynamic abstract data type $TDyn$ so that the operation $Insert(W, d)$ is performed in amortized time*

$$\bar{I}_W(s) \in O\left(\frac{\log s}{s} B_V(s)\right).$$

Let n be the size of the current data set. We have

$$\begin{aligned} S_W(n) &\in \log n S_v(n) \\ B_W(n) &\in B_V(n) \\ Q_W(n) &\in \log n Q_v(n). \end{aligned}$$

7.2.2 Amortized Delete: Occasional Reconstruction

Assume that we did not have implemented the $Insert$ operation, yet.

If we have to delete an object we can not choose its location beforehand. Therefore the deletion of an objects is much more difficult than the insertion. Deletion may cause fundamental reconstruction.

For many data structures it is easier to simply mark an object as deleted. Physically the object remains in the structure but does no longer belong to the data set D . These objects have bad influence on the running time of all operations although they are no longer necessary. Therefore from time to time we have to reconstruct the data structure for the actual data set.

First of all, for $TStat$ we introduce an additional operation $weak.delete(V, d)$ with cost function $WD_V(n)$. We simply want to construct a *strong* delete function with an acceptable amortized time bound for $TDyn$.

Therefore we use $weak.delete(V, d)$ until D has only the half size of V . Then we erase V and build a new structure V out of D . The cost of the *occasional reconstruction* is amortized over the preceding *delete*-operations. This gives the following result.

Theorem 12 *A static abstract data type as presented in Figure 69 with an additional operation $weak.delete(V, d)$ and with additional cost function $WD_V(n)$ can be dynamized by means of occasional reconstruction in a dynamic abstract data type $TDyn$ so that*

$$\begin{aligned} B_W(r) &= B_V(r) \\ E_W(r) &\in O(E_V(r)) \\ Q_W(r) &\in O(Q_V(r)) \\ S_W(r) &\in O(S_V(r)) \\ \bar{D}_W(s) &\in O\left(WD_V(s) + \frac{B_V(s)}{s}\right), \end{aligned}$$

holds. The size of the current actual data set is denoted with r and s denotes the length of the operation sequence.

We omit the proof here.

7.2.3 Amortized Insert and Amortized Delete

In the preceding sections we have discussed Insert and Delete separately. Now we want to show how to combine the two approaches.

A static abstract data type with a weak delete implementation is given. As in Section 7.2.1 we use

the binary structure for the insertion. The operation *weak.delete* is only available for the structures V_i and we have to extend it to W in order to apply the result of Section 7.2.2. If *Weak.Delete*(W, d) is applied, d should be marked as deleted in W . But we do not know in which of the structures V_i the element d lies. Therefore in addition to the binary structure we construct a balanced *searchtree* T that stores this information. For every $d \in W$ there is a pointer to the structure V_i with $d \in V_i$, an example is shown in Figure 71.

The additional cost of the search tree T is covered as follows. *Query* operations are not involved. For *Weak.Delete*(W, d) there is an additional $O(\log n)$ for searching the corresponding V_i and for marking d as deleted in V_i .

If an object d has to be inserted we have to update T . The object d gets an entry for its structure V_j in T , this is done in time $O(\log n)$ and it will not affect the time bound for the insertion. But furthermore if V_0, \dots, V_{j-1} has to be erased the corresponding objects should point to V_j afterwards. This can be efficiently realized by collecting the pointers of T to V_i in a list for every V_i . We collect the pointers and change them to " V_j ". This operation is already covered by time $O(B_V(2^j))$ for constructing V_j .

Altogether we conclude:

Theorem 13 *A static abstract data type as presented in Figure 69 with an additional operation $weak.delete(V, d)$ and with additional cost function $WD_V(n)$ can be dynamized by means of binary structure, searchtree T and occasional reconstruction in a dynamic abstract data type $TDyn$ so that the amortized time for insertion reads*

$$\bar{I}_W(s) \in O\left(\log s \frac{B_V(s)}{s}\right),$$

and the amortized time for deletion reads

$$\bar{D}_W(s) \in O\left(\log s + WD_V(s) + \frac{B_V(s)}{s}\right).$$

For the rest of the operations we have

$$\begin{aligned} B_W(r) &= B_V(r) \\ E_W(r) &\in O(\log r E_V(r)) \\ Q_W(r) &\in O(\log r Q_V(r)) \\ S_W(r) &\in O(S_V(r)). \end{aligned}$$




Figure 71: A structure W_{15} with a searchtree T storing pointers to V_i for every $d \in V_i$.

The size of the current actual data set is denoted with r and s denotes the length of the operation sequence.

7.3 Worst-Case sensitive Insert and Delete

In the last section we have seen that it is easy to amortize the *cost* of Insert and Delete analytically over time. The main idea for the construction of the dynamic data structure was given by the *binary structure* of W_n which has fundamental changes from time to time but the corresponding costs were amortized. Now we are looking for the worst-case cost of Insert and Delete. The idea is to distribute the construction of V_j *itself* over time, i.e., the structure V_j should be finished if $V_{j-1}, V_{j-2}, \dots, V_0$ has to be erased.

We only refer to the result of this approach. The ideas are very similar to the ideas of the preceding sections. Technically, a *modified binary representation* is used in order to distribute the effort of the reconstruction over time. For the interested reader we refer to Klein⁵⁵ or van Kreveld⁸⁹ and Overmars.⁷²

Theorem 14 *A static abstract data type as presented in Figure 69 with an additional operation $weak.delete(V, d)$ and with additional cost function $WD_V(n)$ can be dynamized in a dynamic abstract data type $TDyn$ so that*

$$\begin{aligned} \text{Build}(W, D) &\in O(B_V(n)) \\ \text{Query}(W, q) &\in O(\log n \cdot Q_V(n)) \\ \text{Insert}(W, d) &\in O\left(\frac{\log n}{n} B_V(n)\right) \\ \text{Delete}(W, d) &\in O\left(\log n + WD_V(n) + \frac{B_V(n)}{n}\right) \\ \text{Space} &O(S_V(n)). \end{aligned}$$

Here n denotes the number of relevant, stored data objects.

7.4 A Simple Example

For convenience, we take a simple example from Section 2 and apply Theorem 14, thus implementing worst-case sensitive insertion and deletion.

In Section 2.2 an easy implementation of the *static $k-d$ -tree* was presented with $S_{k-d}(n) = O(n)$ and query time $Q_{k-d}(n) = O(\sqrt{n} + a)$, where a represents the size of the answer (see Theorem 4). Obviously, *weak.delete*($k-d, x$) can be implemented in $O(\log n)$ time, thus we obtain $WD_{k-d}(n) = O(\log n)$. Additionally we have $B_{k-d}(n) = O(n \log n)$.

Let *Dyn*($k-d$) denote the dynamic variant based upon the *statically* implemented $k-d$ -tree.

Application of Theorem 14 results in:

$$\begin{aligned} \text{Build}(\text{Dyn}(k-d), D) &\in O(n \log n) \\ \text{Query}(\text{Dyn}(k-d), q) &\in O(\sqrt{n} \log n + a) \\ \text{Insert}(\text{Dyn}(k-d), d) &\in O\left(\log^2 n\right) \\ \text{Delete}(\text{Dyn}(k-d), d) &\in O(\log n) \\ \text{Space} &O(n). \end{aligned}$$

References

- [1] S. AR, B. CHAZELLE, AND A. TAL, *Self-Customized BSP Trees for Collision Detection*, Computational Geometry: Theory and Applications, 15 (2000), pp. 91–102.
- [2] S. AR, G. MONTAG, AND A. TAL, *Deferred, Self-Organizing BSP Trees*, in Eurographics, Sept. 2002, pp. 269–278. <http://www.ee.technion.ac.il/~ayellet/papers.html>.
- [3] J. ARVO AND D. B. KIRK, *Fast Ray Tracing by Ray Classification*, in Computer Graphics (SIGGRAPH '87 Proceedings), M. C. Stone, ed., vol. 21, July 1987, pp. 55–64.
- [4] F. AURENHAMMER, *Voronoi diagrams: A survey of a fundamental geometric data structure*, ACM Comput. Surv., 23 (1991), pp. 345–405.

- [5] F. AURENHAMMER AND R. KLEIN, *Voronoi Diagrams*, in Handbook of Computational Geometry, J.-R. Sack and J. Urrutia, eds., Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000, pp. 201–290. <http://wwwpi6.fernuni-hagen.de/Publikationen/tr198.pdf>.
- [6] J. AVRO AND D. KIRK, *A survey of ray tracing acceleration techniques*, in An Introduction to Ray Tracing, A. Glassner, ed., Academic Press, San Diego, CA, 1989, pp. 201–262. ISBN 0-12-286160-4.
- [7] L. BALMELLI, J. KOVACEVIC, AND M. VETTERLI, *Quadtrees for Embedded Surface Visualization: Constraints and Efficient Data Structures*, in Proc. of IEEE International Conference on Image Processing (ICIP), vol. 2, Oct. 1999, pp. 487–491.
- [8] L. BALMELLI, T. LIEBLING, AND M. VETTERLI, *Computational Analysis of 4-8 Meshes with Application to Surface Simplification using global Error*, in Proc. of the 13th Canadian Conference on Computational Geometry (CCCG), Aug. 2001.
- [9] G. BAREQUET, B. CHAZELLE, L. J. GUIBAS, J. S. B. MITCHELL, AND A. TAL, *BOXTREE: A Hierarchical Representation for Surfaces in 3D*, Computer Graphics Forum, 15 (1996), pp. C387–C396, C484. ISSN 0167-7055.
- [10] N. BECKMANN, H.-P. KRIEGEL, R. SCHNEIDER, AND B. SEEGER, *The R*-tree: An efficient and robust access method for points and rectangles*, in Proc. ACM SIGMOD Conf. on Management of Data, 1990, pp. 322–331.
- [11] J. BERNAL, *Bibliographic notes on Voronoi diagrams*, tech. rep., National Institute of Standards and Technology, Gaithersburg, MD 20899, 1992.
- [12] J.-D. BOISSONNAT AND M. TEILLAUD, *On the randomized construction of the Delaunay tree*, Theoret. Comput. Sci., 112 (1993), pp. 339–354. http://www.inria.fr/cgi-bin/wais_ra_sophia?question=1140.
- [13] G. BORGEFORS, *Distance transformations in arbitrary dimensions*, in Computer Vision, Graphics, Image Processing, vol. 27, 1984, pp. 321–345.
- [14] P.-T. BREMER, S. D. PORUMBESCU, F. KUESTER, B. HAMANN, K. I. JOY, AND K.-L. MA, *Virtual Clay Modeling using Adaptive Distance Fields*, in Proceedings of the 2002 International Conference on Imaging Science, Systems, and Technology (CISST 2002), H. R. A. et al., ed., vol. 1, Athens, Georgia, 2002.
- [15] P. J. C. BROWN, *Selective Mesh Refinement for Rendering*, PhD dissertation, Emmanuel College, University of Cambridge, Feb. 1998. <http://www.cl.cam.ac.uk/Research/Rainbow/publications/pjcb/thesis/>.
- [16] L. P. CHEW, *Constrained Delaunay triangulations*, Algorithmica, 4 (1989), pp. 97–108.
- [17] ———, *Guaranteed-quality mesh generation for curved surfaces*, in Proc. 9th Annu. ACM Sympos. Comput. Geom., 1993, pp. 274–280.
- [18] N. CHIN, *Partitioning a 3D Convex Polygon with an Arbitrary Plane*, in Graphics Gems III, D. Kirk, ed., Academic Press, 1992, chapter V.2, pp. 219–222.
- [19] D. COHEN-OR, A. SOLOMOVICI, AND D. LEVIN, *Three-Dimensional Distance Field Metamorphosis*, ACM Transactions on Graphics, 17 (1998), pp. 116–141. <http://visinfo.zib.de/EVlib/Show?EVL-1998-152>.
- [20] ———, *Three-Dimensional Distance Field Metamorphosis*, ACM Transactions on Graphics, 17 (1998), pp. 116–141. ISSN 0730-0301.
- [21] M. DE BERG, *Linear Size Binary Space Partitions for Fat Objects*, in Proc. 3rd Annu. European Sympos. Algorithms, vol. 979 of Lecture Notes Comput. Sci., 1995, pp. 252–263.
- [22] ———, *Linear size binary space partitions for uncluttered scenes*, Algorithmica, 28 (2000), pp. 353–366.
- [23] M. DE BERG, M. VAN KREVELD, M. OVERMARS, AND O. SCHWARZKOPF, *Computational Geometry: Algorithms and Applications*, Springer-Verlag, Berlin, Germany, 2nd ed., 2000.
- [24] F. DEHNE AND H. NOLTEMEIER, *A computational geometry approach to clustering problems*, in Proc. 1st Annu. ACM Sympos. Comput. Geom., 1985, pp. 245–250.
- [25] A. K. DEWDNEY AND J. K. VRANCH, *A convex partition of R^3 with applications to Crum's problem and Knuth's post-office problem*, Utilitas Math., 12 (1977), pp. 193–199.
- [26] H. DJIDJEV AND A. LINGAS, *On computing the Voronoi diagram for restricted planar figures*, in Proc. 2nd Workshop Algorithms Data Struct., vol. 519 of Lecture Notes Comput. Sci., 1991, pp. 54–64.
- [27] D. P. DOBKIN AND M. J. LASZLO, *Primitives for the manipulation of three-dimensional subdivisions*, Algorithmica, 4 (1989), pp. 3–32.
- [28] R. A. DWYER, *Higher-dimensional Voronoi diagrams in linear expected time*, Discrete Comput. Geom., 6 (1991), pp. 343–367.

- [29] H. EDELSBRUNNER, *Algorithms in Combinatorial Geometry*, vol. 10 of EATCS Monographs on Theoretical Computer Science, Springer-Verlag, Heidelberg, West Germany, 1987.
- [30] H. EDELSBRUNNER AND N. R. SHAH, *Incremental topological flipping works for regular triangulations*, *Algorithmica*, 15 (1996), pp. 223–241.
- [31] S. A. EHMANN AND M. C. LIN, *Accurate and Fast Proximity Queries Between Polyhedra Using Convex Surface Decomposition*, in *Computer Graphics Forum*, vol. 20, 2001, pp. 500–510. ISSN 1067-7055.
- [32] A. A. ELASSAL AND V. M. CARUSO, *USGS digital cartographic data standards - Digital Elevation Models*, Technical Report Geological Survey Circular 895-B, US Geological Survey, 1984.
- [33] D. A. FIELD, *Implementing Watson's algorithm in three dimensions*, in *Proc. 2nd Annu. ACM Sympos. Comput. Geom.*, 1986, pp. 246–259.
- [34] S. FORTUNE, *Voronoi diagrams and Delaunay triangulations*, in *Computing in Euclidean Geometry*, D.-Z. Du and F. K. Hwang, eds., vol. 1 of *Lecture Notes Series on Computing*, World Scientific, Singapore, 1st ed., 1992, pp. 193–233.
- [35] S. F. FRISKEN, R. N. PERRY, A. P. ROCKWOOD, AND T. R. JONES, *Adaptively Sampled Distance Fields: A General Representation of Shape for Computer Graphics*, in *Siggraph 2000, Computer Graphics Proceedings*, K. Akeley, ed., Annual Conference Series, 2000, pp. 249–254. <http://visinfo.zib.de/EVlib/Show?EVL-2000-61>.
- [36] H. FUCHS, Z. M. KEDEM, AND B. F. NAYLOR, *On Visible Surface Generation by a Priori Tree Structures*, in *Computer Graphics (SIGGRAPH '80 Proceedings)*, vol. 14, July 1980, pp. 124–133.
- [37] D. FUSSELL AND K. R. SUBRAMANIAN, *Fast Ray Tracing Using K-D Trees*, Technical Report TR-88-07, U. of Texas, Austin, Dept. Of Computer Science, Mar. 1988.
- [38] A. S. GLASSNER, ed., *An Introduction to Ray Tracing*, Academic Press, 1989.
- [39] J. GOLDSMITH AND J. SALMON, *Automatic Creation of Object Hierarchies for Ray Tracing*, *IEEE Computer Graphics and Applications*, 7 (1987), pp. 14–20.
- [40] S. GOTTSCHALK, M. LIN, AND D. MANOCHA, *OBB-Tree: A Hierarchical Structure for Rapid Interference Detection*, in *SIGGRAPH 96 Conference Proceedings*, H. Rushmeier, ed., ACM SIGGRAPH, Aug. 1996, pp. 171–180. held in New Orleans, Louisiana, 04-09 August 1996.
- [41] L. J. GUIBAS AND J. STOLFI, *Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams*, *ACM Trans. Graph.*, 4 (1985), pp. 74–123.
- [42] H. W. HAMACHER, *Mathematische Lösungsverfahren für planare Standortprobleme*, Verlag Vieweg, Wiesbaden, 1995.
- [43] I. HOFF, KENNETH E., T. CULVER, J. KEYSER, M. LIN, AND D. MANOCHA, *Fast Computation of Generalized Voronoi Diagrams Using Graphics Hardware*, pp. 375–376.
- [44] K. E. HOFF III, T. CULVER, J. KEYSER, M. LIN, AND D. MANOCHA, *Fast Computation of Generalized Voronoi Diagrams Using Graphics Hardware*, LA, California, Aug. 8–13 1999, pp. 277–286. <http://www.cs.unc.edu/~geom/voronoi/>.
- [45] J. HUANG, Y. LI, R. CRAWFIS, S. C. LU, AND S. Y. LIOU, *A complete distance field representation*, in *Proceedings of the conference on Visualization 2001*, 2001, pp. 247–254. ISBN 0-7803-7200-X.
- [46] P. M. HUBBARD, *Real-Time Collision Detection and Time-Critical Computing*, in *SIVE 95, The First Workshop on Simulation and Interaction in Virtual Environments*, no. 1, Iowa City, Iowa, July 1995, University of Iowa, pp. 92–96.
- [47] H. INAGAKI, K. SUGIHARA, AND N. SUGIE, *Numerically robust incremental algorithm for constructing three-dimensional Voronoi diagrams*, in *Proc. 4th Canad. Conf. Comput. Geom.*, 1992, pp. 334–339.
- [48] B. JOE, *3-Dimensional Triangulations from Local Transformations*, *SIAM J. Sci. Statist. Comput.*, 10 (1989), pp. 718–741.
- [49] ———, *Construction of Three-Dimensional Delaunay Triangulations Using Local Transformations*, *Comput. Aided Geom. Design*, 8 (1991), pp. 123–142.
- [50] M. W. JONES AND R. A. SATHERLEY, *Shape representation using space filled sub-voxel distance fields*, pp. 316–325.
- [51] T. C. KAO AND D. M. MOUNT, *Incremental construction and dynamic maintenance of constrained Delaunay triangulations*, in *Proc. 4th Canad. Conf. Comput. Geom.*, 1992, pp. 170–175.

- [52] N. KATAYAMA AND S. SATOH, *The SR-tree: An Index Structure for High-Dimensional Nearest Neighbor Queries*, in Proc. ACM SIGMOD Conf. on Management of Data, 1997, pp. 369–380.
- [53] T. L. KAY AND J. T. KAJIYA, *Ray Tracing Complex Scenes*, in Computer Graphics (SIGGRAPH '86 Proceedings), D. C. Evans and R. J. Athay, eds., vol. 20, Aug. 1986, pp. 269–278.
- [54] R. KIMMEL, N. KIRYATI, AND A. M. BRUCKSTEIN, *Multi-Valued Distance Maps for Motion Planning on Surfaces with Moving Obstacles*, IEEE Transactions on Robotics and Automation, 14 (1998), pp. 427–436.
- [55] R. KLEIN, *Algorithmische Geometrie*, Addison-Wesley, Bonn, 1997. <http://www.oldenbourg.de/cgi-bin/rotitel?T=24382>.
- [56] R. KLEIN, A. SCHILLING, AND W. STRAER, *Reconstruction and simplification of surfaces from contours*, in Pacific Graphics, 1999, pp. 198–207. <http://cg.cs.uni-bonn.de/publications/publication.asp?id=31>.
- [57] J. T. KLOSOWSKI, M. HELD, J. S. B. MITCHELL, H. SOWRIZAL, AND K. ZIKAN, *Efficient Collision Detection Using Bounding Volume Hierarchies of k -DOPs*, IEEE Transactions on Visualization and Computer Graphics, 4 (1998), pp. 21–36.
- [58] T. LARSSON AND T. AKENINE-MLLER, *Collision Detection for Continuously Deforming Bodies*, in Eurographics, 2001, pp. 325–333. short presentation.
- [59] J.-C. LATOMBE, *Robot Motion Planning*, Kluwer Academic Publishers, Boston, 1991.
- [60] C. L. LAWSON, *Software for C^1 surface interpolation*, in Math. Software III, J. R. Rice, ed., Academic Press, New York, NY, 1977, pp. 161–194.
- [61] D. T. LEE AND A. K. LIN, *Generalized Delaunay Triangulation for Planar Graphs*, Discrete Comput. Geom., 1 (1986), pp. 201–217.
- [62] S. LEUTENEGGER, J. EDGINGTON, AND M. LOPEZ, *STR: A Simple and Efficient Algorithm for R-Tree Packing*, in Proceedings of the 13th International Conference on Data Engineering (ICDE'97), Washington - Brussels - Tokyo, Apr. 1997, pp. 497–507. ISBN 0-8186-7807-0.
- [63] T.-Y. LI AND J.-S. CHEN, *Incremental 3D Collision Detection with Hierarchical Data Structures*, in Proc. VRST '98, Taipei, Taiwan, Nov. 1998, ACM, pp. 139–144.
- [64] P. LINDSTROM, D. KOLLER, W. RIBARSKY, L. F. HUGHES, N. FAUST, AND G. TURNER, *Real-Time, Continuous Level of Detail Rendering of Height Fields*, in SIGGRAPH 96 Conference Proceedings, H. Rushmeier, ed., ACM SIGGRAPH, Aug. 1996, pp. 109–118. held in New Orleans, Louisiana, 04-09 August 1996.
- [65] P. LINDSTROM AND V. PASCUCCI, *Visualization of Large Terrains Made Easy*, in Proc. IEEE Visualization, San Diego, 2001.
- [66] W. E. LORENSEN AND H. E. CLINE, *Marching Cubes: A High Resolution 3D Surface Construction Algorithm*, in Computer Graphics (SIGGRAPH '87 Proceedings), M. C. Stone, ed., vol. 21, July 1987, pp. 163–169.
- [67] G. MLLER, S. SCHFER, AND W. D. FELLNER, *Automatic Creation of Object Hierarchies for Radiosity Clustering*, Computer Graphics Forum, 19 (2000). ISSN 0167-7055.
- [68] B. NAYLOR, J. AMANATIDES, AND W. THIBAUT, *Merging BSP Trees Yields Polyhedral Set Operations*, in Computer Graphics (SIGGRAPH '90 Proceedings), E. Baskett, ed., vol. 24, Aug. 1990, pp. 115–124.
- [69] B. F. NAYLOR, *A Tutorial on Binary Space Partitioning Trees*, ACM SIGGRAPH '96 Course Notes 29, (1996).
- [70] M. NOVOTNI AND R. KLEIN, *A Geometric Approach to 3D Object Comparison*, in International Conference on Shape Modeling and Applications, May 2001, pp. 167–175.
- [71] R. OSADA, T. FUNKHOUSER, B. CHAZELLE, AND D. DOBKIN, *Matching 3D models with shape distributions*, in Proceedings of the International Conference on Shape Modeling and Applications (SMI-01), B. Werner, ed., Los Alamitos, CA, May 7–11 2001, pp. 154–166.
- [72] M. H. OVERMARS, *The Design of Dynamic Data Structures*, vol. 156 of Lecture Notes Comput. Sci., Springer-Verlag, Heidelberg, West Germany, 1983.
- [73] D. W. PAGLIERONI, *Distance Transforms: Properties and Machine Vision Applications*, CVGIP: Graphical Models and Image Processing, 54 (1992), pp. 56–74.
- [74] I. J. PALMER AND R. L. GRIMSDALE, *Collision Detection for Animation using Sphere-Trees*, Computer Graphics Forum, 14 (1995), pp. 105–116. ISSN 0167-7055.
- [75] M. S. PATERSON AND F. F. YAO, *Efficient binary space partitions for hidden-surface removal and solid modeling*, Discrete Comput. Geom., 5 (1990), pp. 485–503.

- [76] B. A. PAYNE AND A. W. TOGA, *Distance field manipulation of surface models*, IEEE Computer Graphics and Applications, 12 (1992), pp. 65–71.
- [77] F. P. PREPARATA AND M. I. SHAMOS, *Computational Geometry: An Introduction*, Springer-Verlag, New York, NY, 1985.
- [78] V. T. RAJAN, *Optimality of the Delaunay triangulation in R^d* , in Proc. 7th Annu. ACM Sympos. Comput. Geom., 1991, pp. 357–363.
- [79] J. REVELLES, C. URENA, AND M. LASTRA, *An Efficient Parametric Algorithm for Octree Traversal*, in WSCG 2000 Conference Proceedings, University of West Bohemia, Plzen, Czech Republic, 2000. <http://visinfo.zib.de/EVlib/Show?EVL-2000-307>.
- [80] N. ROUSSOPOULOS AND D. LEIFKER, *Direct spatial search on pictorial databases using packed R-trees*, in Proceedings of ACM-SIGMOD 1985 International Conference on Management of Data, May 28–31, 1985, LaMansion Hotel, Austin, Texas, S. Navathe, ed., New York, NY 10036, USA, 1985, pp. 17–31. ISBN 0-89791-160-1. <http://www.acm.org/pubs/articles/proceedings/mod/318898/p17-roussopoulos/p17-roussopoulos.pdf>; <http://www.acm.org/pubs/citations/proceedings/mod/318898/p17-roussopoulos/>.
- [81] R. SEIDEL, *Constrained Delaunay triangulations and Voronoi diagrams with obstacles*, Technical Report 260, IIG-TU Graz, Austria, 1988.
- [82] J. A. SETHIAN, *An Analysis of Flame Propagation*, PhD thesis, Dept. of Mathematics, University of California, Berkeley, 1982. <http://visinfo.zib.de/EVlib/Show?EVL-1982-2>.
- [83] ———, *Level Set Methods*, Cambridge University Press, 1996. <http://visinfo.zib.de/EVlib/Show?EVL-1996-149>.
- [84] ———, *Level set methods and fast marching methods*, Cambridge University Press, 1999.
- [85] M. I. SHAMOS, *Computational Geometry*, ph.D. thesis, Dept. Comput. Sci., Yale Univ., New Haven, CT, 1978.
- [86] M. TANEMURA, T. OGAWA, AND W. OGITA, *A New Algorithm for Three-Dimensional Voronoi Tessellation*, J. Comput. Phys., 51 (1983), pp. 191–207.
- [87] E. TORRES, *Optimization of the Binary Space Partition Algorithm (BSP) for the Visualization of Dynamic Scenes*, in Eurographics '90, C. E. Vandoni and D. A. Duce, eds., Sept. 1990, pp. 507–518.
- [88] G. VAN DEN BERGEN, *Efficient Collision Detection of Complex Deformable Models using AABB Trees*, Journal of Graphics Tools, 2 (1997), pp. 1–14.
- [89] M. J. VAN KREVELD, *New Results on Data Structures in Computational Geometry*, ph.D. dissertation, Dept. Comput. Sci., Utrecht Univ., Utrecht, Netherlands, 1992.
- [90] M. WAN, F. DACHILLE, AND A. KAUFMAN, *Distance-field based skeletons for virtual navigation*, in Proceedings of the conference on Visualization 2001, 2001, pp. 239–246. ISBN 0-7803-7200-X.
- [91] C. A. WANG, *Efficiently updating the constrained Delaunay triangulations*, BIT, 33 (1993), pp. 238–252.
- [92] C. A. WANG AND L. SCHUBERT, *An optimal algorithm for constructing the Delaunay triangulation of a set of line segments*, in Proc. 3rd Annu. ACM Sympos. Comput. Geom., 1987, pp. 223–232.
- [93] D. F. WATSON, *Computing the n-Dimensional Delaunay Tessellation with Applications to Voronoi Polytopes*, Comput. J., 24 (1981), pp. 167–172.
- [94] H. WEGHORST, G. HOOPER, AND D. P. GREENBERG, *Improved Computational Methods for Ray Tracing*, ACM Transactions on Graphics, 3 (1984), pp. 52–69.
- [95] L.-Y. WEI AND M. LEVOY, *Fast Texture Synthesis Using Tree-Structured Vector Quantization*, in Siggraph 2000, Computer Graphics Proceedings, K. Akeley, ed., Annual Conference Series, 2000, pp. 479–488. <http://visinfo.zib.de/EVlib/Show?EVL-2000-87>.
- [96] J. WILHELMS AND A. V. GELDER, *Octrees for Faster Isosurface Generation Extended Abstract*, in Computer Graphics (San Diego Workshop on Volume Visualization), vol. 24, Nov. 1990, pp. 57–62.
- [97] C. YOUNGBLUT, R. E. JOHNSON, S. H. NASH, R. A. WIENCLAW, AND C. A. WILL, *Different Applications of Two-Dimensional Potential Fields for Volume Modeling*, Tech. Rep. UCAM-CL-TR-541, University of Cambridge, Computer Laboratory, 15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom, Aug. 2002. <http://www.cl.cam.ac.uk/users/nad/pubs/>.
- [98] G. ZACHMANN, *Real-time and Exact Collision Detection for Interactive Virtual Prototyping*, in Proc. of the 1997 ASME Design Engineering Technical Conferences, Sacramento, California, Sept. 1997. Paper no. CIE-4306.

- [99] ———, *Rapid Collision Detection by Dynamically Aligned DOP-Trees*, in Proc. of IEEE Virtual Reality Annual International Symposium; VRAIS '98, Atlanta, Georgia, Mar. 1998, pp. 90–97.
- [100] ———, *Minimal Hierarchical Collision Detection*, in Proc. ACM Symposium on Virtual Reality Software and Technology (VRST), Hong Kong, China, Nov.11–13 2002, pp. 121–128. <http://www.gabrielzachmann.org/>.
- [101] B. ZHU AND A. MIRZAIAN, *Sorting does not always help in computational geometry*, in Proc. 3rd Canad. Conf. Comput. Geom., Aug. 1991, pp. 239–242.