


Advanced Graphics



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Part II



Study Guide & Supplementary Material

Neil Dodgson, 2007

Computer Laboratory
15 J J Thomson Avenue
Cambridge CB3 0FD

E-mail: nad@cl.cam.ac.uk
<http://www.cl.cam.ac.uk/~nad/>

Contents

Advanced Graphics Lecture Notes

N. A. Dodgson (62 pages)

An introduction to the surface representations used in CAD

N. A. Dodgson (5 sheets)

Advanced Graphics 2006: subdivision curves & surfaces

N. A. Dodgson (10 sheets)

Marching cubes: a high resolution 3D surface construction algorithm

W. E. Lorensen & H. E. Cline (8 pages)

Mathematical elements for computer graphics (extract)

D. F. Rogers & J. A. Adams (17 sheets)

Chapter XI (“Radiosity”) from 3-D Computer Graphics

S. R. Buss (9 sheets)

Other useful material

The course web pages:

<http://www.cl.cam.ac.uk/Teaching/current/AdvGraph/>

Advanced Graphics Lecture Notes

Neil Dodgson*

University of Cambridge Computer Laboratory

Overview

This course provides students with a solid grounding in a variety of three-dimensional modelling mechanisms. It also provides an introduction to radiosity, animation, graphics cards, and current commercial uses of computer graphics.

Syllabus

- **Introduction.** Revision of the ray tracing and polygon scan conversion methods of making images from 3D models; the pros and cons of each approach. Current uses of computer graphics in animation, special effects, Computer-Aided Design and marketing. [0.75 lecture]
- **The polygon.** Drawing polygons. Graphics cards. Polygon mesh management: data structures. [0.75 lecture]
- **Ray tracing.** The primitive geometric shapes used in ray tracing: plane, polygon, sphere, cylinder, cone, box, disc, torus. Ray intersection calculations and normal calculations for these. Converting the primitives into polygons for use in polygon scan conversion. [1.5 lectures]
- **Splines for modelling arbitrary 3D geometry** (splines are the standard 3D modelling mechanism for Computer-Aided Design). Features required of surface models in a Computer-Aided Design package. Bézier curves and surfaces. B-splines, from uniform, non-rational B-splines through to non-uniform, rational B-splines (NURBS). [2.5 lectures]
- **Subdivision surfaces** (an alternative mechanism for representing arbitrary 3D geometry, now widely used in the animation industry). Introduction to subdivision. Pros and cons when compared to NURBS. [2 lectures]
- **Implicit surfaces and voxels.** 3D pixels and the marching cubes algorithm; medical applications of this. [1 lecture]

*Written 10/99, modifications made 09/00, 10/02, 09/04, 04/06, 03/07. ©1999, 2000, 2002, 2004, 2006, 2007 Neil A. Dodgson

- **Other ways to create complex geometry.** Generative models: extrusion, revolution, sweeping, generalised cylinders. Constructive solid geometry (CSG): set theory applied to solid objects; different implementations of this using ray tracing and polygons. [1 lecture]
- **Radiosity.** Accurate calculation of the diffuse inter-reflections in a scene. [1.5 lectures]
- **Computer animation.** A brief introduction to some techniques in animation. [1 lecture]

Objectives

On completing the course, students should be able to

- produce equations for each geometric primitive, derive a ray/primitive intersection algorithm for each, describe how each can be approximated by polygons
- define NURBS basis functions, understand the use of NURBS curves and surfaces in 2D and 3D modelling
- describe and explain how to use generative models, constructive solid geometry, implicit surfaces, voxel rendering and subdivision surfaces; describe how each representation can be converted to polygons
- be able to compare and contrast ray tracing with polygon scan conversion
- be able to explain the basic radiosity algorithm

Why *Advanced Graphics*? The title “Advanced Graphics” dates from the year in which the course was first proposed. At this time a 16 lecture course on various advanced topics in graphics was envisaged. The course is now 12 lectures long. Today, it is mainly concerned with 3D modelling techniques, so the course title is, perhaps, a little misleading. Computer Laboratory policy is, however, to minimize changes to course titles. 3D modelling is important because it underpins all of the practical uses of 3D computer graphics.

What’s examinable? Everything except where explicitly noted otherwise. This means that anything that is covered in the lectures is examinable, even if it is not in the notes, unless I say otherwise, and that anything that is in the notes is examinable, unless noted otherwise.

Lecture handouts and supervision material Some of the lecture course material is available on the web¹. This material is also printed out to provide these lecture notes. Other material is bound in with these notes (this material cannot be put on the web for copyright reasons). There are exercises scattered throughout the notes. These can usually be found at the end of sections. My thanks to Dr Jonathan Pfautz and Dr Andy Penrose for some of the exercises.

¹<http://www.cl.cam.ac.uk/Teaching/current/AdvGraph>

Book list and their abbreviations The following books are referred to in the course. Each is preceded by the abbreviation used in these notes to refer to that book.

- **FvDFH** Foley, J.D., van Dam, A., Feiner, S.K. & Hughes, J.F. (1990). *Computer Graphics: Principles and Practice*. Addison-Wesley (2nd ed.). The traditional “bible” of Computer Graphics. It tends to be terse but it has wide coverage of all of the basics.
- **F&vD** Foley J.D. & van Dam, A. (1984). *Fundamentals of Interactive Computer Graphics*. Addison-Wesley (1st ed.). The earlier version of **FvDFH**. It contains only about half of the material of the second edition, but is still comprehensive about the basics of computer graphics.
- **SSC** Slater, M., Steed, A. & Chrysanthou, Y. (2002). *Computer Graphics & Virtual Environments*. Addison Wesley. A more recent book which covers all the basics. Also has sections on Constructive Solid Geometry (Ch. 18), Quadrics (also Ch. 18), Radiosity (Ch. 15), and an introduction to Bézier and B-Spline curves and surfaces (Ch. 19).
- **Buss** Buss, S.R. (2003). *3-D Computer Graphics*. Cambridge University Press. Another recent book which has the best description of radiosity (Ch. XI) that I have ever read. It also contains chapters on Bézier curves (VII), B-Splines (VIII), ray tracing (IX and X) and animation (XII).
- **R&A** Rogers, D.F. & Adams, J.A. (1990). *Mathematical Elements for Computer Graphics*. McGraw-Hill (2nd ed.). A good coverage of the mathematics of the 2D and 3D representation of shape as it was understood in the year of publication. Explains Bézier, B-spline, and NURBS curves and surfaces in great detail. Also covers conics and quadrics.
- **Farin** Farin, G. (2002, 5th ed.; 1997, 4th ed.). *Curves and Surfaces for CAGD*. Morgan Kaufmann (5th ed.). Academic Press (4th ed.). A good alternative source for information on Bézier, B-Spline, NURBS, and conics. Regularly updated since its original publication in 1988.
- **W&W** Warren, J. & Weimer, H. (2002). *Subdivision Methods for Geometric Design*. Morgan Kaufmann. The only book on the market devoted entirely to subdivision methods.
- **GG I-V** *Graphics Gems I* (1990) to *Graphics Gems V* (1995). Academic Press. A collection of five books containing a wide variety of algorithms for use in computer graphics. A wide range of tips, tricks and techniques is included.

Other material in the handout In addition to the lecture notes that you are reading now, I have included copies of the following:

1. Lorensen and Cline’s 1987 SIGGRAPH paper “Marching cubes: a high resolution 3D surface construction algorithm”, *Proc SIGGRAPH ’87*, pages 163–169. This is relevant to the section on voxels and marching cubes (section 8.3.2).

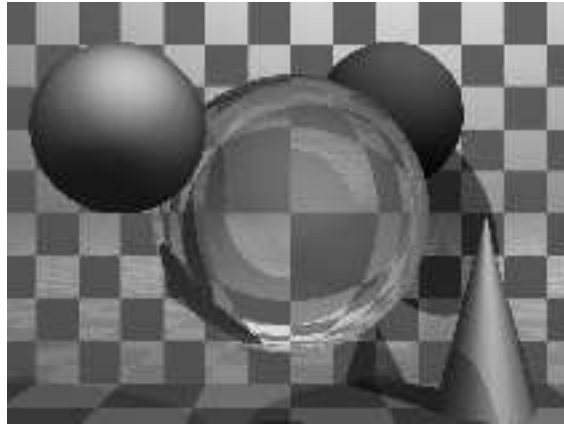


Figure 1: A basic ray traced model showing refraction and shadowing.

2. Extracts from Rogers and Adams (**R&A**): table 4-8 (conics), figure 6-18 (quadrics), and *parts* of sections 6-2 (surfaces of revolution), 6-3 (sweeps), 5-8 (Bézier curves), 5-9 (B-splines), and 5-13 (NURBS). These extracts are relevant to the corresponding sections of this course.
3. Chapter XI, “Radiosity” from **Buss**. This is used in the Radiosity section of this course.

Note on copyright material The items in the above list are copyrighted material provided under the University of Cambridge’s license from the Copyright Licensing Agency. This allows us to make one copy for each student and supervisor (“tutor”) on the course *within certain limits*. These are: no more than three works and no more than 5% or one whole article or chapter from each work. This material is provided solely for the student’s own study. Further copying of this handout is a breach of copyright.

Be warned: to fit inside these limits I have heavily edited the extracts from **R&A**. In particular, I have included none of the worked examples. To thoroughly understand the material I suggest that you read this extract and then borrow (or buy) a copy of **R&A** in order to go through the examples.

1 Basic 3D modelling

1.1 Ray tracing vs polygon scan conversion

These are the two standard methods of producing images of three-dimensional solid objects. They were covered in some detail in the Part IB course. If you want to revise them then check out **FvDFH** sections 14.4, 15.10 and 15.4 or **F&vD** sections 16.6 and 15.5. Line drawing is also used for representing three-dimensional objects in some applications. It is briefly covered later on.

1.1.1 Ray tracing

Ray tracing has the tremendous advantage that it can produce realistic looking images. The technique allows a wide variety of lighting effects to be implemented. It also permits a range of primitive shapes that is limited only by the ability of the programmer to write an algorithm to intersect a ray with the shape. It is considered by many to be the natural or obvious way to render 3D objects.

Ray tracing works by firing one or more rays from the eye point through each pixel. The colour assigned to a ray is the colour of the first object that it hits, determined by the object's surface properties at the ray-object intersection point, the illumination at that point, and contributions from any reflection or refraction that occurs at that point. The colour of a pixel is some average of the colours of all the rays fired through it. The power of ray tracing lies in the fact that secondary rays are fired from the ray-object intersection point to determine its exact illumination (and hence colour). This spawning of secondary rays allows reflection, refraction, and shadowing to be handled with ease. A simple raytraced image can be seen in Figure 1 and the basic algorithm in Figure 13.

Ray tracing's big disadvantage is that it is slow. It takes minutes, or hours, to render a reasonably detailed scene. Ray tracing was first implemented in hardware by a Cambridge company, Advanced Rendering Technologies², in the late 1990s. The quality of the images that they can produce is high compared with polygon scan conversion. This is their main selling point. However, ray tracing is so computationally intensive that it is not possible to produce images at the same speed as hardware assisted polygon scan conversion. Other researchers are trying to do this by using multiple processors (dozens to hundreds), but ray tracing will always be slower than polygon scan conversion.

Ray tracing therefore is only used where the visual effects cannot be obtained using polygon scan conversion. This means that it is, in practice, used by a minority of movie and television special effects companies, advertising companies, and enthusiastic amateurs.

1.1.2 Example

The kitchen in Figure 2 was rendered using the ray tracing program *Rayshade*³. *Rayshade* has not been updated for over a decade. An alternative ray tracer, which is kept up to date, is *POVray*⁴, with which you may like to experiment. It is worth visiting the *POVray* website to see the stunning imagery which has been produced using the ray tracer .

The close-ups of the kitchen scene in Figures 3 and 4 show some of the power of ray tracing. The kitchen sink reflects the wall tiles. The bench top in front of the kitchen sink has a specular highlight on its curved front edge. The washing machine door is a perfectly curved object (impossible to achieve with polygons). The inner curve is part of a cone, the outer curve is a cylinder. You can see the floor tiles reflected in the door. Both the washing machine door and the sink basin were made using CSG techniques (see section 8.2). The grill on the stove casts interesting shadows (there are two lights in the scene). This sort of thing is much easier to do with ray tracing than with polygon scan conversion.

²<http://www.artvps.com/>

³<http://graphics.stanford.edu/~cek/rayshade/rayshade.html>

⁴<http://www.povray.org/>



Figure 2: A ray traced model of a kitchen design.



Figure 3: A close up of the kitchen sink.




Figure 4: Close up views of the washing machine door and the grill on the stove.




Figure 5: A scan converted model of a city (courtesy of Jon Sewell).

1.1.3 Polygon scan conversion

This term encompasses a range of algorithms where polygons are rendered, normally one at a time, into a frame buffer. The term *scan* comes from the fact that an image on a CRT is made up of *scan lines*. Examples of polygon scan conversion algorithms are the painter's algorithm, the *z*-buffer, and the A-buffer (see your Part IB lecture notes, **FvDFH** chapter 15, or **F&vD** chapter 15). In this course we will generally assume that polygon scan conversion refers to the *z*-buffer algorithm or one of its derivatives, such as the A-buffer.

The advantage of polygon scan conversion is that it is fast. Polygon scan conversion algorithms are used in computer games, flight simulators, and other applications where interactivity is important. To give a human the illusion that they are interacting with a 3D model in real time, you need to present the human with animation running at 10 frames per second or faster for passive viewing on a monitor, TV, or movie screen. Research at the University of North Carolina⁵ has experimentally shown that for immersive virtual reality applications this is not high enough and at least 15 frames per second is a minimum. Polygon scan conversion is capable of providing this sort of speed. The NVIDIA⁶ GeForce8 graphics processing unit (GPU) architecture, for example, has up to 128 parallel stream processors running at 1.35GHz. It can render up to 36.8 billion textured pixels per second, and can render scenes containing several million triangles in real time. While we might hope that scientific or medical applications were considered important applications of computer graphics, it is the game industry that is driving the development of graphics card technology.

One problem with polygon scan conversion is that it can only support simplistic lighting models, so images do not necessarily look realistic. For example: transparency can be supported, but refraction requires the use of a texture-mapping technique called "refraction mapping"; reflections can be supported, at the expense of rendering a "reflection map" before rendering the scene; shadows can be produced using "shadow maps". All of these are more complicated methods than those used in ray tracing. Where ray tracing

⁵<http://www.cs.unc.edu>

⁶<http://www.nvidia.com/>




Figure 6: An SGI O2 computer and components drawn with and without texture maps.




Figure 7: Left: some floating objects in a simulated environment. Right: a close up of the red ball showing the reflection of the ball in the shiny floor.

is a clean and simple algorithm, polygon scan conversion uses a variety of tricks of the trade to get the desired results. The other limitation of polygon scan conversion is that it only has a single primitive: the polygon, which means that everything is made up of flat surfaces. This is especially unrealistic when modelling natural objects such as humans or animals, unless you use polygons that are no bigger than a pixel, which is indeed what happens these days. An image generated using a polygon scan conversion algorithm, even one which makes heavy use of texture mapping, will still tend to look computer generated.

1.1.4 Examples

Texture mapping is a simple way of making a polygon scan conversion (or a ray tracing) scene look better without introducing lots of polygons. The images in Figure 6 show a scene both with and without any texture maps. Obviously this scene was designed to be viewed with the texture maps turned on. This example shows that texture mapping can make simple geometry look interesting to a human observer.

The images in Figure 7 were generated using polygon scan conversion. Texture map-



Figure 8: An example of environment mapping: a silvered SGI O2 computer reflecting an environment map of the interior of a cafe.

ping has been used to make the back and side walls more interesting. All the objects are reflected in the floor. This reflection is achieved by duplicating all of the geometry, upside-down, under the floor, and making the floor partially transparent. The close-up shows the reflection of the red ball, along with a circular “shadow” of the ball. This shadow is, in fact, a polygonal approximation to a circle drawn on the floor polygon and bears no relationship to the lights whatsoever. You may need to look at the version of these images on the Lab’s website to see the images more clearly and in colour.

Environment mapping (Figure 8) is another clever idea which makes polygon scan conversion images look more realistic. In environment mapping we have a texture map of the environment which can be thought of as wrapping completely around the entire scene (you could think of it as six textures on the six inside faces of a big box). The environment map itself is not drawn, but if any polygon is reflective then the normal to the polygon is found at each pixel (this normal is needed for Gouraud shading anyway) and from this, and a vector pointing to the eye, the appropriate point (and therefore colour) on the environment map can be located. You may note that finding the correct point on the environment map is actually a simple (and easily optimised) piece of ray tracing.

1.1.5 Line drawing

An alternative to the above methods is to draw the 3D model as a wire frame outline. This is obviously unrealistic, but is useful in particular applications. The wire frame outline can be either see through or hidden lines can be removed (**FvDFH** section 15.3 or **F&vD** section 14.2.6). In general, the lines that are drawn will be the edges of the polygons which would be drawn by a polygon scan conversion algorithm.

Line drawing was historically faster than polygon scan conversion. However, modern graphics cards can handle both lines and polygons at about the same speed. Line drawing of 3D models is used in Computer Aided Design (CAD) and in 3D model design. The software which people use to design 3D models tends to use line drawing in its user



Figure 9: Screen shots from commercial flight simulators (circa 1995).



Figure 10: The real cockpit of a commercial flight simulator: an exact replica of the equivalent airplane's cockpit.

interface with polygon scan conversion providing preview images of the model. I find it interesting that, when **R&A** was first written in 1976, the authors had only line drawing algorithms with which to illustrate their 3D models. Only one figure in the entire book did not use exclusively line drawing: Fig. 6-52, which had screen shots of a prototype polygon scan conversion system. Technology has moved on enormously since then.

1.1.6 Applications of computer graphics

Visualisation generally does not require realistic looking images. In science we are usually visualising complex three dimensional structures, such as protein molecules, which have no “realistic” visual analogue. In medicine we generally prefer an image that helps in diagnosis over one which looks beautiful. Polygon scan conversion is therefore normally used in visualisation (although some data require voxel rendering – see section 8.3.2).

Simulation uses polygon scan conversion because it can generate images at interactive speeds. At the high (and very expensive) end a great deal of computer power is used. In 1990, the most expensive flight simulators (those with full hydraulic suspension and other fancy stuff) cost about £10M, of which £1M went on the graphics kit. Similar rendering power is available today on a graphics card which costs a hundred pounds and fits in a PC. Figure 9 shows screen shots from two commercial flight simulators in the mid-1990s; Figure 10 shows the simulator’s cockpit, which is an exact physical replica of the cockpit on a real aircraft. Although the cost of the graphics has dropped dramatically, the cost of the physical kit has not.

3D games (for example Quake⁷ and Unreal⁸) use polygon scan conversion because it gives interactive speeds. A lot of other “3D” games (for example SimCity⁹, Civilisation, Diablo¹⁰) use pre-drawn sprites (small images) which they simply copy to the appropriate position on the screen. This essentially reduces the problem to an image compositing operation, requiring much less processor time. The sprites can be hand drawn by an artist or created in a 3D modelling package and rendered to sprites in the company’s design office. Donkey Kong Country (mid-1990s), for example, was the first game to use sprites which were ray traced from 3D models.

You may have noticed that the previous sentence is the first mention of ray tracing in this section. It transpires that the principal uses of ray tracing, in the commercial world, are in producing a small quantity of super-realistic images for advertising and in producing a small proportion of the special effects for film and television. Despite what you may have expected, most special effects are done using sophisticated polygon scan conversion algorithms.

The first movie to use 3D computer graphics was *Star Wars*¹¹ [1977]. Graphics were not used for the space ships, animals or sets, however. You may recall that there were some line drawn computer graphics toward the end of the movie in the targeting interfere on the X-wing fighter. All of the spaceship shots, and all of the other fancy effects, were done using models, mattes (hand-painted backdrops), and hand-painting on the actual film. Computer graphics technology has progressed incredibly since then. The twenty-fifth anniversary re-release^{25th} of the Star Wars trilogy included a number of computer graphic enhancements, all of which were composited into the original movie.

Twenty years on we saw computer graphics effects of the kind found in movies such as the (rather bloodythirsty) *Starship Troopers*¹² [1997]. Most of the giant insects in

⁷<http://www.idsoftware.com/games/quake/quake3-gold/>

⁸<http://www.unreal.com>

⁹<http://www.simcity.com/>

¹⁰<http://www.blizzard.com/diablo2/>

¹¹<http://www.starwars.com/>

¹²<http://www.imdb.com/title/tt0120201/>

the movie are completely computer generated. The spaceships are a combination of computer graphic models and real models. The largest of these real models was 18' (6m) long: so it was obviously still worthwhile spending a lot of time and energy on the real thing.

Special effects are not necessarily computer generated. Compare *King Kong* [1933]¹³ with *King Kong* [2005]¹⁴. The plot has barely changed, but the special effects have improved enormously: changing from hand animation (and a man in a monkey suit) to swish computer generated imagery. Not every special effect you see in a modern movie is computer generated. In *Starship Troopers*, for example, the explosions are real. They were set off by a pyrotechnics expert against a dark background (probably the night sky), filmed, and later composited into the movie. In *Titanic*¹⁵ [1997] the scenes with actors in the water were shot in the warm Gulf of Mexico. In order that they look as if they were shot in the freezing North Atlantic, cold breaths had to be composited in later. These were filmed in a cold room over the course of one day by a special effects studio. Film makers obviously need to balance quality, ease of production, and cost. They will use whatever technology gives them the best trade off. This is increasingly computer graphics, but computer graphics is still not useful for everything by quite a long way.

In the three *Lord of the Rings* movies¹⁶, almost anything which could be shot in live action was shot this way. Computer graphics were used only where they were easier or cheaper or the only feasible way to do something. For example, in *Return of the King*, the lava was originally to be produced by computer graphics simulation. When the results were found to be not realistic enough, some of the shots were re-done using real gunk flowing down a real model of a mountainside. Helms Deep, in *The Two Towers*, consisted of some computer graphics, a small-scale model of the whole thing, a quarter-scale model of the wall and citadel and a full-scale model of parts of the citadel for real actors to perform on. Compositing all the component of any given shot is an interesting image processing task. In a typical movie, each frame (at 24 frames per second) will have anything from twenty to over a hundred separate elements which need to be composited to make the final image.

Completely computer-generated movies have been with us for over a decade. *Toy Story*¹⁷ [1995] was the world's first feature length computer generated movie. Two more were released in 1998 (*A Bug's Life*¹⁸ [1998] and *Antz*¹⁹ [1998]). These were followed by *Toy Story 2*²⁰ [1999], *Dinosaur* [2000], *Shrek*²¹ [2001], *Monsters Inc*²² [2001], *Ice Age*²³ [2002], *Finding Nemo* [2003], and *Shrek 2* [2004]. The genre is now well established and there are several recent examples, with more in the pipeline. Note the subject matter of these movies (toys, bugs, dinosaurs, monsters, sea life, fairytale characters). It is still very difficult to model humans realistically and much research is being undertaken

¹³<http://www.imdb.com/title/tt0024216/>

¹⁴<http://www.imdb.com/title/tt0360717/>

¹⁵<http://www.titanicmovie.com/>

¹⁶<http://www.lordoftherings.net/>

¹⁷<http://www.pixar.com/featurefilms/ts/index.html>

¹⁸<http://www.pixar.com/featurefilms/abl/index.html>

¹⁹<http://www.imdb.com/title/tt0120587/>

²⁰<http://www.pixar.com/featurefilms/ts2/index.html>

²¹<http://www.shrek.com/>

²²<http://www.pixar.com/featurefilms/inc/index.html>

²³<http://www.iceagemovie.com/>

in the field of realistic human modelling. *Final Fantasy*²⁴ [2001] was the first serious attempt to represent fully human characters in a fully computer-generated movie.

1.1.7 Polygon scan conversion or ray tracing for special effect?

While ray tracing gives a better range of lighting effects than polygon scan conversion, we usually get acceptable results with polygon scan conversion through the use of techniques such as environment mapping and the use of enormous numbers of tiny polygons. The special effects industry still dithers over whether to jump in and use ray tracing. Many special effects are done using polygon scan conversion, with maybe a bit of ray tracing for special things (giving a hybrid ray tracing/polygon scan conversion algorithm).

Toy Story [1995], for example, used Pixar's proprietary polygon scan conversion algorithm. It took between one and three hours to render each frame (these frames have a resolution of 1526×922 pixels) and over 800,000 CPU hours were absorbed in the making of the movie (roughly a CPU century). More expensive algorithms can be used in less time if you are rendering for television (I estimate that about one sixth of the pixels are needed compared to a movie) or if you are only rendering a small part of a big image for compositing into live action.

At the ACM SIGGRAPH²⁵ conference in 1998 I had the chance to hear about the software that some real special effects companies were using. Two of these companies used ray tracing and two were pretty happy using polygon scan conversion.

BlueSky—ViFX Ray traced everything using CGI-Studio.

Digital Domain²⁶ Used ray tracing provided by commercial software, except when the commercial software cannot do what they want. Used MentalRay²⁷ on *Fifth Element* [1997]; used Alias²⁸ models (NURBS) passed to Lightwave²⁹ (polygons) for one advertisement; used MentalRay³⁰ plus Renderman³¹ for another advertisement.

Rhythm + Hues³² Used a propriety renderer, which was about ten years old in 1998. It has been rewritten many times. They made only limited use of ray tracing.

Station X Used Lightwave³³ plus an internally developed renderer which is a hybrid between ray tracing and *z*-buffer.

At Eurographics 2002 and SIGGRAPH 2002, it was apparent that little had changed over the intervening four years: the computers had got faster and artists were produc-

²⁴<http://www.imdb.com/title/tt0173840/>

²⁵<http://www.siggraph.org/>

²⁶<http://www.d2.com/>

²⁷<http://www.mentalray.com/>

²⁸<http://www.aliaswavefront.com/>

²⁹<http://www.newtek.com/>

³⁰<http://www.mentalray.com/>

³¹<http://www.pixar.com/renderman/>

³²<http://www.rhythm.com/>

³³<http://www.newtek.com/>

ing more detailed work but polygon scan conversion is still the technology of choice for almost all commercial applications of computer graphics.

1.2 Exercises

1. Compare and contrast the capabilities and uses of ray tracing and polygon scan conversion.
2. In what circumstances is line drawing more useful than either ray tracing or polygon scan conversion.
3. (a) When is realism critical? (b) Give 5 examples of applications where different levels of visual realism are necessary and explain what sort of rendering is needed for each and why.
4. “The quality of the special effects cannot compensate for a bad script.” Discuss with reference to movies that you have seen.

2 The polyyon

2.1 Polygon mesh management

In order to do polygon scan conversion or line drawing we need to know how to handle polygon meshes.

2.1.1 Drawing polygons

In order to draw a polygon, you obviously need to know its vertices. To get the shading correct you also need to know its normal. The direction of the normal tells you which side is the front of the polygon and which is the back. Many systems assume one-sided polygons: the front side is shaded and the back side either is coloured matt grey or black or is not even considered. This is sensible if the polygon is part of a closed polyhedron. In many applications, all objects consist of closed polyhedra; but you cannot guarantee that this will always be the case, which means that you will get unexpected results if the back sides of polygons are actually visible on screen.

The normal vector does not need to be specified independently of the polygon’s vertices because it can be calculated from the vertices. As an example: assume a polygon has three vertices, A, B and C. The normal vector can be calculated as: $N = (C - B) \times (A - B)$.

Any three adjacent vertices in a polygon can be used to calculate the normal vector but the *order* in which the vertices are specified is important: it changes whether the vector points up or down relative to the polygon. In a right-handed co-ordinate system the three vertices must be specified anti-clockwise round the polygon as you look down the desired normal vector (i.e. as you look at the front side of the polygon). If there are more than three vertices in the polygon, they must all lie in the same plane, otherwise the shape will not be a polygon.

Thus, for drawing purposes, we need to know only the vertices and surface properties of the polygon. The vertices naturally give us both edge and orientation information.

The surface properties are such things as the specular and diffuse colours, and details of any texture mapping which may be applied to the polygon. These things are generally specified at the vertices (diffuse colour, specular colour, texture co-ordinates) for use in Gouraud or Phong shading.

2.1.2 Interaction with polygon mesh data

The above is fine for drawing but, if you wish to manipulate the polygon mesh (for example, in a 3D modelling package), then it is useful to know quite a lot more about the connectivity of the mesh. For example: if you want to move a vertex, which is shared by four polygons, you do not want to have to search through all the polygons in your data structure trying to find the ones which contain a vertex which matches your vertex data, you want some data structure which allows easy access to the relevant information.

The various versions of the *winged-edge data structure* are particularly useful for handling polygon mesh data. The version shown in Figure 11 contains explicit links for all of the relationships between vertices, edges and polygons, thus making it easy to find, for example, which polygons are attached to a given vertex, or which polygons are adjacent to a given polygon (by traversing the edge list for the given polygon, and finding which polygon lies on the other side of each edge).

The vertex object contains the vertex's co-ordinates, a pointer to a list of all edges of which this vertex is an end-point, and a pointer to a list of all polygons of which the vertex is a vertex. It also has a pointer to the vertex's surface properties (such as colour and texture coordinates).

The polygon object contains (a pointer to) the polygon's surface property information (such as its texture map), a pointer to a list of all edges which bound this polygon, and a pointer to an ordered list of all vertices of the polygon.

The edge object contains pointers to its start and end vertices, and pointers to the polygons which lie to the left and right of it.

Figure 11 shows just one possible implementation of a polygon mesh data structure. **FvDFH** section 12.5.2 describes another winged-edge data structure which contains slightly less information, and therefore requires more accesses than the one shown here to find certain pieces of information. The implementation that would be chosen depends on the needs of the particular application which is using the data structure. The trade-off is between ease of extracting information and ease of updating the data structure. **F&vD** section 13.2 and **SSC** pp. 170–172 also contain some information on polygon meshes.

In general, we will want a polygon mesh to form a manifold surface. This is where the surface is what a human would naturally think of as a surface, without any three-way joins or other peculiar features; a surface which you could flatten onto a plane given sufficiently many cuts and a bit of stretching here and there. Mathematically, a manifold surface is where the neighbourhood of every point is topologically equivalent to a disc (except at the edges of the manifold, where it is topologically equivalent to a half disc). The principal upshot of this is that each edge in the polygon mesh can be the edge of either one or two polygons, no more and no less.




Figure 11: One version of the winged edge data structure.

2.2 Hardware polygon scan conversion quirks

A piece of polygon scan conversion hardware, such as the Silicon Graphics³⁴ Reality Engine or the NVIDIA³⁵ GeForce family of graphics cards, has generally consisted of a *geometry* engine and a *rendering* engine. The geometry engine will handle the transformations of all vertices and normals, and some of the shading calculations. The rendering engine will implement the polygon scan conversion algorithm on the transformed data. Modern graphics cards allow for user programming in both the geometry and rendering engine. Machine instructions are provided for the usual operations (addition, multiplication), and also for such necessary things as taking the dot product of two vectors. The geometry and rendering engines both have multiple copies of the same hardware to allow for multiple vertices and polygons to be processed in parallel. These are generally built with a SIMD (single instruction, multiple data) parallel processor architecture. The architecture is optimised for processing graphics, so the user is somewhat limited in what he or she can program. However, the most recent graphics cards allow for a good deal of flexibility. Early generations of cards allowed a limited number of instructions. For example, the NVIDIA GeForce 3 card (2001) had a maximum of 256 instructions in the whole program, no more than twelve working registers, no jumps or loops, no access to general memory. The latest NVIDIA GeForce 8 cards (2007) have thousands of registers and allow up thousands of instructions, with jumps and loops. The introduction of jumps and loops causes interesting issues with the SIMD architecture, requiring

³⁴<http://www.sgi.com>

³⁵<http://www.nvidia.com/>




Figure 12: Left: a triangle strip set. Right: a triangle fan set.

different pipes to be able to choose whether or not to execute any given instruction.

To give you an idea of the complexity which is possible, on the GeForce4 generation of NVIDIA cards (2002), the information that is passed to the geometry engine, for a single vertex, is position, weight, normal, primary and secondary colour, fog coordinate, and eight texture coordinates; all sixteen of these are floating-point four-component vectors. The output from the geometry engine is homogeneous clip space position, primary and secondary colours for front and back faces of the polygon, fog coordinate, point size, and texture coordinate set; where they are all again floating-point four-component vectors except for the output fog coordinate and the point size³⁶. Both geometry and rendering engines have read access to the texture buffers. Graphics cards are now so powerful that they are being used as general purpose co-processors for a variety of mathematics-intensive computation tasks, using texture buffers for storing intermediate results.

The latest (2007) generations of NVIDIA GeForce graphics cards (the GeForce 8 family) and ATI chips (the Xenos chip used in the Xbox 360) have progressed (or reverted?) to a *unified shader* model of processing, where any processor can handle either the geometry processing or the pixel processing. This allows more efficient distribution of the processing load as appropriate to the objects being rendered on the screen.

2.2.1 Triangles only

When making a piece of hardware to render a polygon, it is much easier to make the hardware handle a fixed number of vertices per polygon, than a variable number. Most hardware implementations thus implement only triangle drawing. This is not a serious drawback. Polygons with more vertices are simply split into triangles.

2.2.2 The triangle strip set and triangle fan set

In addition to simple triangle drawing, rendering hardware may also implement either or both of the triangle strip set and triangle fan set to speed up processing through the geometry engine (see Figure 12). Each triangle in the set has two vertices in common with the previous triangle. Each vertex is transformed only once by the geometry engine, giving a factor of three speed up in geometry processing.

³⁶You are not expected to remember all of these input and output registers, but they give you an idea of the complexity of the processing which can go on inside a graphics card.

For example, assume we have triangles **ABC**, **BCD**, **CDE** and **DEF**. In naïve triangle rendering, the vertices would be sent to the geometry engine in the order **ABC BCD CDE DEF**; each triangle's vertices being sent separately. With a triangle strip set the vertices are sent as **ABCDEF**; the adjacent triangles' vertices overlapping.

A triangle fan set is similar. In the four triangle case we would have triangles **ABC**, **ACD**, **ADE** and **AEF**. The vertices would again be sent just as **ABCDEF**. It is obviously important that the rastering engine be told whether it is drawing standard triangles or a triangle strip set or a triangle fan set.

2.2.3 The vertex cache

The triangle strip and fan sets work because there is a vertex cache which can hold all the relevant data about two vertices. Around 2000, a vertex cache was introduced to graphics cards. On the NVIDIA family of cards, the initial version held the twenty most recently used vertices, hence obviating the need to be explicitly specify fan sets and strip sets, although you still need to send the triangles to the card in some reasonably coherent order and you do need to let the graphics card know that the triangles form a set with the same surface properties. You also need to index the vertices so that you refer to each by its index rather than by sending the (x, y, z) coordinates again.

2.3 Exercises

1. Calculate both surface normal vectors (left-handed and right-handed) for a triangle with points $(1, 1, 0)$, $(2, 0, 1)$, $(-1, -2, -1)$.
2. Confirm that the following statements provide a definition of a polygon mesh which represents a manifold surface:
 - (a) A vertex belongs to at least two edges.
 - (b) A vertex is a vertex of at least one polygon.
 - (c) An edge has exactly two end points.
 - (d) An edge is an edge of either one or two polygons.
 - (e) A polygon has at least three vertices.
 - (f) A polygon has at least three edges.
3. Work out the algorithm that is required to modify a winged-edge data structure when an edge is split. You may ignore surface property information for the polygons and you may assume that the edge that is split is split exactly in half. The algorithm could be called by the function call:

```
split_edge( vertex_list v, edge_list e,
            polygon_list p, edge edge_to_split )
```

where the winged-edge data structure is made up of the three linked lists of objects (vertices, edges, and polygons).

4. [2002/7/9] Describe the situations in which it is sensible to use a winged-edged data structure to represent a polygon mesh and, conversely, the situations in which a winged-edged data structure is not a sensible option for representing a polygon mesh. What is the minimum information which is required to successfully draw a polygon mesh using Gouraud shading? [4 marks]

3 Ray tracing primitives

A *primitive* is a shape for which a ray-shape intersection routine has been written. More complex objects can be built out of the primitives. Most ray tracers will have a variety of primitives. They are limited only by the ability of the programmer to write a function to analytically intersect a ray with the surface of the shape. The basic raytracing algorithm is given in Figure 13. Figure 15 shows some of the common primitives.

For practical experience, download and play with either *POVray*³⁷ or *Rayshade*³⁸.

3.1 Mathematical preliminaries

3.1.1 Vector arithmetic

It is helpful to remember your vector arithmetic. A 3D vector is represented thus:

$$\mathbf{V} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad (1)$$

For ease of writing such definitions in text we may say $\mathbf{V} = (x, y, z)$, where we understand that this ordered triple is equivalent to the vector.

The magnitude of vector \mathbf{V} is:

$$|\mathbf{V}| = \sqrt{x^2 + y^2 + z^2} \quad (2)$$

The dot product of two vectors, $\mathbf{A} = (x_A, y_A, z_A)$ and $\mathbf{B} = (x_B, y_B, z_B)$, is:

$$\mathbf{A} \cdot \mathbf{B} = x_A x_B + y_A y_B + z_A z_B \quad (3)$$

which could also be written as a matrix multiplication:

$$\mathbf{A} \cdot \mathbf{B} = \mathbf{A}^T \mathbf{B} = [x_A y_A z_A] \begin{bmatrix} x_B \\ y_B \\ z_B \end{bmatrix} \quad (4)$$

The dot product is a scalar value equal to:

$$\mathbf{A} \cdot \mathbf{B} = |\mathbf{A}| |\mathbf{B}| \cos \theta \quad (5)$$

where θ is the angle between the two vectors. Note that this means that:

$$\mathbf{V} \cdot \mathbf{V} = |\mathbf{V}|^2 \quad (6)$$

³⁷<http://www.povray.org/>

³⁸<http://graphics.stanford.edu/~cek/rayshade/>

```

Select eye point ( $E$ ), look point ( $L$ ), and up vector
Set screen plane to be centred on  $L$ , perpendicular to vector  $EL$ 
Set screen size and number of pixel
For each pixel {
    Define  $D$  to be the vector from  $E$  to the centre of the pixel
    Raytrace( $E, D$ )
}

function Raytrace( $E, D$ ) returns Colour {
    nearest_ $t$  =  $\infty$ 
    nearest_object = NULL
    for each object {
        find  $t$ , the smallest, non-negative real solution of
        the ray/object intersection equation
        if  $t$  exists {
            if  $t < \text{nearest\_}t$  { nearest_ $t$  =  $t$  ; nearest_object = current object }
        }
    }
    colour = black
    if nearest_object exists {
        find normal vector,  $N$ , at intersection point
        if object is reflective {
            reflected_colour = Raytrace(intersection point, reflection vector)
            colour += reflection_coeff * reflected_colour
        }
        if object is refractive {
            refracted_colour = Raytrace(intersection point, refracted vector)
            colour += refraction_coeff * refracted_colour
        }
        for each light {
            if shadow_ray(intersection point, light position) returns No_Shadow {
                calculate light's colour contribution by doing the illumination calculations
                using  $D, N$ , the current light, and the object properties
                colour += lights colour contribution
            }
        }
    }
    return colour
}

```

Figure 13: A simplistic pseudocode version of the basic ray tracing algorithm. The function *shadow_ray* can be found in Figure 14.

```

function shadow_ray(point1, point2) returns Shadow or No_Shadow {
  ray defined with  $E=point1$ ,  $D=point2-point1$ 
  nearest_t =  $\infty$ 
  nearest_object = NULL
  for each object {
    find  $t$ , the smallest, non-negative real solution of
    the ray/object intersection equation
    if  $t$  exists {
      if  $t < nearest\_t$  {  $nearest\_t = t$  }
    }
  }
  if  $nearest\_t < 1$  return Shadow
  else return No_Shadow
}

```

Figure 14: A simplistic pseudocode version of the function *shadow_ray* used by the ray tracing algorithm in Figure 13.

and hence:

$$|\mathbf{V}| = \sqrt{\mathbf{V} \cdot \mathbf{V}} \quad (7)$$

Also note that the dot product between two perpendicular vectors is always zero. The cross product (vector product) of these two vectors is:

$$\mathbf{A} \times \mathbf{B} = \begin{bmatrix} y_A z_B - z_A y_B \\ z_A x_B - x_A z_B \\ x_A y_B - y_A x_B \end{bmatrix} \quad (8)$$

The cross product is a vector that is perpendicular to both \mathbf{A} and \mathbf{B} . This is a very handy way to make a new vector which is guaranteed perpendicular to a given vector. It also means that:

$$\mathbf{A} \cdot (\mathbf{A} \times \mathbf{B}) = 0 \quad (9)$$

$$\mathbf{B} \cdot (\mathbf{A} \times \mathbf{B}) = 0 \quad (10)$$

though this may be getting a bit esoteric and so let's move on.

3.1.2 Points and displacements

Both points and displacements are three-tuples of real numbers. However, they are different beasts and must be treated differently. A point is an absolute position in space relative to some fixed origin, $\mathbf{O} = (0, 0, 0)$. A displacement is an offset, specifying a distance and direction but *not* an absolute position.

Only certain arithmetic operations are permitted on points and displacements. Let $\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3, \dots$ represent points and $\mathbf{D}_1, \mathbf{D}_2, \mathbf{D}_3, \dots$ represent displacements. Displacements can be added,

$$\mathbf{D}_3 = \mathbf{D}_1 + \mathbf{D}_2, \quad (11)$$

but points cannot be added. The difference of two displacements is a displacement,

$$\mathbf{D}_3 = \mathbf{D}_1 - \mathbf{D}_2, \quad (12)$$

The difference of two points is also a displacement:

$$\mathbf{D}_1 = \mathbf{P}_1 - \mathbf{P}_2. \quad (13)$$

You can add a displacement to a displacement,

$$\mathbf{D}_3 = \mathbf{D}_1 + \mathbf{D}_2, \quad (14)$$

or a displacement to a point,

$$\mathbf{P}_2 = \mathbf{P}_1 + \mathbf{D}_1, \quad (15)$$

but *not* a point to a point. It is possible to take a weighted average of points,

$$\mathbf{P}_{\text{new}} = \alpha_1 \mathbf{P}_1 + \alpha_2 \mathbf{P}_2 + \cdots + \alpha_k \mathbf{P}_k, \quad \sum_{i=1}^k \alpha_i = 1 \quad (16)$$

provided that the weights sum to one, otherwise the operation makes no sense. To understand why this is so, add an offset, \mathbf{D}_o , to each point, $\mathbf{P}_1 \dots \mathbf{P}_k$. Only if the α_i sum to one does \mathbf{P}_{new} also get offset by the same amount, as it should be.

If you transform your coordinate system, then all points scale, rotate, and translate as you would expect. Displacements also scale and rotate but they do not translate. To see why this is the case, consider $\mathbf{D}_1 = \mathbf{P}_1 - \mathbf{P}_2$. Let the two points be translated: $\mathbf{P}'_1 = \mathbf{P}_1 + \mathbf{D}_o$ and $\mathbf{P}'_2 = \mathbf{P}_2 + \mathbf{D}_o$. Then:

$$\mathbf{D}'_1 = \mathbf{P}'_1 - \mathbf{P}'_2 \quad (17)$$

$$= (\mathbf{P}_1 + \mathbf{D}_o) - (\mathbf{P}_2 + \mathbf{D}_o) \quad (18)$$

$$= (\mathbf{P}_1 - \mathbf{P}_2) + (\mathbf{D}_o - \mathbf{D}_o) \quad (19)$$

$$= (\mathbf{P}_1 - \mathbf{P}_2) \quad (20)$$

$$= \mathbf{D}_1 \quad (21)$$

3.2 Equation of a ray

A ray is defined by an origin or eye point, $\mathbf{E} = (x_E, y_E, z_E)$, and an offset displacement, $\mathbf{D} = (x_D, y_D, z_D)$. The equation for the ray is:

$$\mathbf{P}(t) = \mathbf{E} + t\mathbf{D}, t \geq 0 \quad (22)$$

This is equivalent to the three equations:

$$\left. \begin{aligned} x(t) &= x_E + tx_D \\ y(t) &= y_E + ty_D \\ z(t) &= z_E + tz_D \end{aligned} \right\} t \geq 0 \quad (23)$$

When finding a ray-object intersection point, we are looking for the intersection point with the *lowest non-negative* value of t .




Figure 15: Some ray traced primitives: sphere, cylinder, cone, disc, torus, polygon

3.3 Equations for the primitives

3.3.1 Sphere

The sphere is the simplest finite object with which to intersect a ray. Practically any ray tracing program will include the sphere as a primitive. Scaling a sphere by different amounts along the different axes will produce an ellipsoid: a squashed or stretched sphere. There is thus no need to include the ellipsoid as a primitive provided that your ray tracer contains the usual complement of transformations (it would be a poor ray tracer if it did not). There is, however, some subtlety in how normals are transformed in such an anisotropically scaled object, which we discuss below in section 3.4.2.

The unit sphere, centred at the origin, has the implicit equation:

$$x^2 + y^2 + z^2 = 1 \quad (24)$$

In vector arithmetic, this becomes:

$$\mathbf{P} \cdot \mathbf{P} = 1 \quad (25)$$

To find the intersection between this sphere and an arbitrary ray, substitute the ray equation (Equation 23) in the sphere equation (Equation 24):

$$(x_E + tx_D)^2 + (y_E + ty_D)^2 + (z_E + tz_D)^2 = 1 \quad (26)$$

$$\Rightarrow t^2(x_D^2 + y_D^2 + z_D^2) + t(2x_Ex_D + 2y_Ey_D + 2z_Ez_D) + (x_E^2 + y_E^2 + z_E^2 - 1) = 0 \quad (27)$$

$$\Rightarrow at^2 + bt + c = 0 \quad (28)$$

$$\Rightarrow t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \quad (29)$$

where $a = x_D^2 + y_D^2 + z_D^2$, $b = 2x_E x_D + 2y_E y_D + 2z_E z_D$, and $c = x_E^2 + y_E^2 + z_E^2 - 1$. This gives zero, one, or two real values for t . If there are zero real values then there is no intersection between the ray and the sphere. If there are either one or two real values then chose the *smallest, non-negative* value of t . This gives the intersection point (at $P(t)$). If there are only negative values of t , then the line (of which the ray is a part) *does* intersect the sphere, but the intersection point is not on the part of the line which constitutes the ray. In this case there is again no intersection point between the ray and the sphere.

An alternative formulation is to use the vector versions of the equations (Equations 22 and 25):

$$(\mathbf{E} + t\mathbf{D}) \cdot (\mathbf{E} + t\mathbf{D}) = 1 \quad (30)$$

$$\Rightarrow t^2(\mathbf{D} \cdot \mathbf{D}) + t(2\mathbf{E} \cdot \mathbf{D}) + (\mathbf{E} \cdot \mathbf{E} - 1) = 0 \quad (31)$$

$$\Rightarrow at^2 + bt + c = 0 \quad (32)$$

$$\Rightarrow t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \quad (33)$$

Where $a = \mathbf{D} \cdot \mathbf{D}$, $b = 2\mathbf{E} \cdot \mathbf{D}$, and $c = \mathbf{E} \cdot \mathbf{E} - 1$. In other words, exactly the same result, expressed in a more compact way. *Graphics Gems I* (p. 388) describes yet another way of arriving at the same result.

You will remember, from Part IB, that we need to know the normal vector (which is, of course, a displacement) at the intersection point in order to calculate the illumination and/or find the reflection ray. For a sphere centred at the origin, the normal is a vector passing through the origin and the point of intersection.

$$\mathbf{N} = (x, y, z) - (0, 0, 0) \quad (34)$$

$$= (x, y, z) \quad (35)$$

Remember that points and displacements are different beasts. Equation 34 is the subtraction of one point from another while Equation 35 is a displacement. So (x, y, z) in Equation 34 is a point while (x, y, z) in Equation 35 is a displacement.

To simplify the ray tracing algorithm, we may choose to assume that the normal vector points out of the surface on the same side of the surface as the eye point. In this case, we need to be careful about whether the eye point lies inside or outside of the sphere. If the origin of the ray lies inside the sphere then the normal at the intersection point is the negative of that in Equation 35, that is, for the case of \mathbf{E} inside the sphere,

$$\mathbf{N} = (-x, -y, -z). \quad (36)$$

To check whether \mathbf{E} is inside the unit sphere we simply need to check whether $|\mathbf{E}| < 1$.

3.3.2 Cylinder

Intersecting a ray with an infinitely long cylinder is practically as easy as intersecting one with a sphere. The tricky bit, if it can be called that, is to intersect a ray with

a more useful finite length cylinder. This is achieved by intersecting the ray with the appropriate infinitely long cylinder and then ascertaining where along the cylinder the intersection lies. If it lies in the finite length in which you are interested then keep the intersection. If it does not then ignore the intersection. Note that the ray tracer used to render the cylinder in Figure 15 has cylinders without end caps. This is the correct result if you follow the procedure outlined above. Adding end caps to your cylinders requires extra calculations. We consider both cases in more detail below.

The *infinite* unit cylinder aligned along the z -axis is defined as:

$$x^2 + y^2 = 1 \quad (37)$$

To intersect a ray with this, substitute Equation 23 in Equation 37.

$$(x_E + tx_D)^2 + (y_E + ty_D)^2 = 1 \quad (38)$$

$$\Rightarrow t^2(x_D^2 + y_D^2) + t(2x_Ex_D + 2y_Ey_D) + (x_E^2 + y_E^2 - 1) = 0 \quad (39)$$

$$\Rightarrow at^2 + bt + c = 0 \quad (40)$$

$$\Rightarrow t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \quad (41)$$

where $a = x_D^2 + y_D^2$, $b = 2x_Ex_D + 2y_Ey_D$, and $c = x_E^2 + y_E^2 - 1$.

The *finite* open-ended unit cylinder aligned along the z -axis is defined as:

$$x^2 + y^2 = 1, z_{\min} < z < z_{\max} \quad (42)$$

The only difference between this and Equation 37 being the restriction on z . To handle this finite length cylinder, solve Equation 41 above. This gives, at most, two values of t . Call these t_1 and t_2 . Calculate z_1 and z_2 using Equation 23 ($z_1 = z_E + t_1z_D$ and $z_2 = z_E + t_2z_D$) and then check $z_{\min} < z_1 < z_{\max}$ and $z_{\min} < z_2 < z_{\max}$. Whichever intersection point passes this test and, if both pass the test, has the smallest non-negative value of t , is the closest intersection point of the ray with the open-ended finite cylinder.

If we wish the finite length cylinder to be closed we must formulate an intersection calculation between the ray and the cylinder's end caps. The end caps have the formulae:

$$z = z_{\min}, \quad x^2 + y^2 \leq 1 \quad (43)$$

$$z = z_{\max}, \quad x^2 + y^2 \leq 1 \quad (44)$$

and we could find explicit intersections between the ray and these two discs. However, for the cylinder, there is a more efficient trick: once you have calculated the solutions to Equation 41 you will either know that there are no intersections with the infinite cylinder or you will know that there are one or two real intersection points represented by t_1 and t_2 , which may be negative at this point in the algorithm. The previous paragraph explained how to ascertain whether these correspond to points on the finite length open-ended cylinder. Now, if z_1 and z_2 lie either side of z_{\min} we know that the ray intersects the z_{\min} end cap, and can calculate the intersection point as:

$$t_3 = \frac{z_{\min} - z_E}{z_D} \quad (45)$$

A similar equation holds for the z_{\max} end cap. Note that the ray may intersect both end caps, for example when $z_1 < z_{\min}$ and $z_2 > z_{\max}$. Also, note that a ray parallel to the z -axis is a special case and needs to be handled separately. It is left as an exercise how to check for and how to handle this case (see exercises below).

The normal vector of this cylinder, at intersection point (x, y, z) on the surface of the infinite cylinder, will be $(x, y, 0)$ or $(-x, -y, 0)$ depending on whether the ray hits the inside or outside surface first. For a finite cylinder, if the intersection point is on an end cap, then the normal vector will be $(0, 0, -1)$ or $(0, 0, 1)$ depending on which end cap is hit and whether the origin of the ray is inside or outside the finite cylinder. The details of how to ascertain whether the ray's origin, E , is inside or outside of the cylinder in both the finite and infinite cases are left as an exercise.

3.3.3 Cone

Cones are very like cylinders, mathematically. Like the infinite cylinder, there is a simple mathematical definition of an infinite cone which makes it easy to write a ray-cone intersection algorithm. Note that a cone does not need to have a point – it can be truncated short of its ‘top’, as illustrated in Figure 15. The particular ray tracer used does not add end caps to cones.

The *infinite* double cone³⁹ aligned along the z -axis and having unit slope is defined as:

$$x^2 + y^2 = z^2 \quad (46)$$

To intersect a ray with this, substitute Equation 23 in Equation 46.

$$(x_E + tx_D)^2 + (y_E + ty_D)^2 = (z_E + tz_D)^2 \quad (47)$$

$$\Rightarrow t^2(x_D^2 + y_D^2 - z_D^2) + t(2x_Ex_D + 2y_Ey_D - 2z_Ez_D) + (x_E^2 + y_E^2 - z_E^2) = 0 \quad (48)$$

$$\Rightarrow at^2 + bt + c = 0 \quad (49)$$

$$\Rightarrow t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \quad (50)$$

where $a = x_D^2 + y_D^2 - z_D^2$, $b = 2x_Ex_D + 2y_Ey_D - 2z_Ez_D$, and $c = x_E^2 + y_E^2 - z_E^2$.

The *finite* open-ended cone aligned along the z -axis is defined as:

$$x^2 + y^2 = z^2, z_{\min} < z < z_{\max} \quad (51)$$

The only difference between this and Equation 46 being the restriction on z . Note that if z_{\min} and z_{\max} are both positive or both negative then you get a single cone with its top truncated. If either z_{\min} or z_{\max} is zero you get a single cone with its apex at the origin.

To handle this finite length cone you proceed as for the finite length cylinder, with several subtle modifications. In particular, it transpires that it is simpler to calculate the intersection of the ray with the two end-cap discs rather than use the trick we used for cylinders. This is because of the increased number of special cases which would need to be considered for cones. It is left as an exercise to work out how to handle this. It would be instructive to try to do this both ways: first by using explicit ray-disc intersection calculation and second by modifying the trick we used for cylinders, taking all special cases into account.

³⁹*double cone* means that it is two “standard” cones joined at their apices.

3.3.4 Torus

Toroids are reasonably rare in real life (doughnuts and tyre inner tubes notwithstanding). They are somehow alluring to the kinds of people who implement ray tracers and, having a reasonably straightforward mathematical definition, are reasonably simple to implement. They thus appear as primitives in many ray tracers. They become more useful when combined with Constructive Solid Geometry (see section 8.2).

A torus is defined by two parameters: the radius of the torus (that is the radius of the torus's defining circle, measured from the origin) and the radius of the tube (the perpendicular distance from the defining circle to the surface of the torus). These are R and r respectively. Normally $R > r$.

An implicit definition of the torus is:

$$\left(\sqrt{x^2 + y^2} - R\right)^2 + z^2 = r^2 \quad (52)$$

The torus can also be defined parametrically in terms of two angles, θ and ϕ , where θ can be thought of as the angle around the defining circle and ϕ the angle around the inside of the tube:

$$x = (R + r \cos \phi) \cos \theta \quad (53)$$

$$y = (R + r \cos \phi) \sin \theta \quad (54)$$

$$z = r \sin \phi \quad (55)$$

Substituting these three equations into Equation 52 will show that they are correct and is a useful exercise in algebraic manipulation.

To find the intersection points of a ray with a torus you need to substitute Equation 23 in Equation 52. Equation 56 is that substitution with the $(\sqrt{x^2 + y^2} - R)^2$ term expanded, the resulting square root term placed on one side of the equals sign, and all other terms placed on the other side:

$$\begin{aligned} & 2R\sqrt{x_E^2 + 2tx_Ex_D + t^2x_D^2 + y_E^2 + 2ty_Ey_D + t^2y_D^2} \\ = & R^2 + x_E^2 + 2tx_Ex_D + t^2x_D^2 + y_E^2 + 2ty_Ey_D + t^2y_D^2 + z_E^2 + 2tz_Ez_D + t^2z_D^2 - r^2 \end{aligned} \quad (56)$$

If we now square both sides we will get a quartic equation in t . This can be solved using a standard quartic root finder to find the four roots of the equation⁴⁰ (there are up to four intersection points between a torus and an arbitrary ray). A quartic root finder is described in *Graphics Gems V* (p. 3).

Finding the normal to this torus requires finding the intersection point on the surface, $(x, y, z) = ((R + r \cos \phi) \cos \theta, (R + r \cos \phi) \sin \theta, r \sin \phi)$, and the nearest point on the ring through the middle of the torus, $(R \cos \theta, R \sin \theta, 0)$. Subtracting one from the other gives the normal as $\mathbf{N} = (r \cos \phi \cos \theta, r \cos \phi \sin \theta, r \sin \phi)$.

⁴⁰The equation itself is rather fierce and you would not be expected to do the full expansion in an exam. For those who are interested, it looks like this:
 $t^4(x_D^4 + y_D^4 + z_D^4 + 2x_D^2y_D^2 + 2x_D^2z_D^2 + 2y_D^2z_D^2) + t^3(4x_D^3x_E + 4y_D^3y_E + 4z_D^3z_E + 4x_D^2y_Dy_E + 4x_D^2z_Dz_E + 4x_Dx_Ey_D^2 + 4y_D^2z_Dz_E + 4x_Dx_Ez_D^2 + 4y_Dy_Ez_D^2) + t^2(-2R^2x_D^2 - 2R^2y_D^2 + 2R^2z_D^2 - 2r^2x_D^2 - 2r^2y_D^2 - 2r^2z_D^2 + 6x_D^2x_E^2 + 2x_E^2y_D^2 + 8x_Dx_Ey_Dy_E + 2x_D^2y_E^2 + 6y_E^2y_D^2 + 2x_E^2z_D^2 + 2y_E^2z_D^2 + 8x_Dx_Ez_Dz_E + 8y_Dy_Ez_Dz_E + 2x_D^2z_E^2 + 2y_D^2z_E^2 + 6z_D^2z_E^2) + t(-4R^2x_Dx_E - 4R^2y_Dy_E + 4R^2z_Dz_E - 4r^2x_Dx_E - 4r^2y_Dy_E - 4r^2z_Dz_E + 4x_Dx_E^3 + 4x_E^2y_Dy_E + 4x_Dx_Ey_E^2 + 4y_Dy_E^3 + 4x_E^2z_Dz_E + 4y_E^2z_Dz_E + 4x_Dx_Ez_E^2 + 4y_Dy_Ez_E^2 + 4z_Dz_E^3) + (R^4 - 2R^2x_E^2 - 2R^2y_E^2 + 2R^2z_E^2 + r^4 - 2r^2R^2 - 2r^2x_E^2 - 2r^2y_E^2 - 2r^2z_E^2 + x_E^4 + y_E^4 + z_E^4 + 2x_E^2y_E^2 + 2x_E^2z_E^2 + 2y_E^2z_E^2) = 0$

3.3.5 Plane

The infinite plane is a simple object with which to intersect a ray. On its own it can represent boundary objects such as the ground or the sky or perhaps an infinite wall. Intersection with the infinite plane is a useful building block in a ray tracing system as it forms the initial step in polygon and disc calculations. It is also useful in Computational Solid Geometry (section 8.2).

A plane can be defined by a normal vector, \mathbf{N} and a point on the plane, \mathbf{Q} . A point, \mathbf{P} , is on the plane if:

$$\mathbf{N} \cdot (\mathbf{P} - \mathbf{Q}) = 0 \quad (57)$$

To find the ray/plane intersection substitute Equation 22 in Equation 57:

$$\mathbf{N} \cdot (\mathbf{E} + t\mathbf{D} - \mathbf{Q}) = 0 \quad (58)$$

$$\Rightarrow t = \frac{\mathbf{N} \cdot (\mathbf{Q} - \mathbf{E})}{\mathbf{N} \cdot \mathbf{D}} \quad (59)$$

If $\mathbf{N} \cdot \mathbf{D} = 0$ then the ray is parallel to the plane, and there is no intersection point. If $t < 0$ then the plane is behind the eye point and there is no intersection. If $t \geq 0$ then the intersection point is $\mathbf{E} + t\mathbf{D}$. The normal vector is obviously either \mathbf{N} or $-\mathbf{N}$ depending on whether \mathbf{E} is on the side of the plane pointed at by \mathbf{N} or not.

3.3.6 Polygon

Having the polygon as a ray tracing primitive allows a ray tracer to render anything that a polygon scan conversion algorithm could. To find the intersection of a ray with a polygon, first find the intersection of the ray with the infinite plane in which the polygon lies. Then ascertain whether the intersection lies inside or outside the polygon: this is a reasonably straightforward two dimensional graphics operation.

A polygon can be defined by an ordered set of vertices: $(\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3, \dots)$. To find the intersection point between a polygon and a ray, we first find the intersection point between the polygon's plane and the ray, and then ascertain whether this intersection point lies inside the polygon or not.

The normal of the polygon's plane can be found by the simple cross product: $\mathbf{N} = (\mathbf{P}_3 - \mathbf{P}_2) \times (\mathbf{P}_1 - \mathbf{P}_2)$. A point on the plane's surface is even easier to find: $\mathbf{Q} = \mathbf{P}_1$. The intersection calculation then proceeds as above, for the polygon.

If an intersection point between the ray and the plane is found then we can check whether or not the point lies inside the polygon in the following manner. First we project the intersection point and the polygon to two dimensions by simply throwing away one coordinate. The most obvious thing to do is throw away the coordinate in which the polygon has the smallest extent. We then test to see if the intersection point lies inside the two dimensional polygon using the odd/even test.

The odd/even test checks to see whether or not an arbitrary point lies inside an arbitrary polygon in two dimensions. This is done by drawing an arbitrary ray from the point to infinity. If the ray crosses an even number of polygon edges then the point lies outside the polygon. Contrariwise, if the ray crosses an odd number of polygon edges then the point lies inside the polygon. The easiest implementation of this simply runs the arbitrary ray parallel to one of the axes. A full discussion of the implementation

details of this, and other point-in-polygon algorithms, can be found in *Graphics Gems IV* pp. 24–46.

3.3.7 Disc

The disc is essential in ray tracers which implement cones and cylinders without end caps (such as *Rayshade* which was used to render the images in Figure 15). In Figure 15, the disc has been positioned so that it just catches the light – this illustrates how specular reflection varies across a completely flat surface. For comparison, the sphere, cone and torus images illustrate how specular reflection varies across three curved surfaces.

The disc is similar to the polygon. Both are planar objects. A disc can be defined by its centre, \mathbf{Q} , its radius, r , and a normal vector, \mathbf{N} . Finding the intersection point, \mathbf{P} , of the ray with the disc's plane proceeds as for a plane/ray or polygon/ray intersection. The routine to check if the intersection with the plane lies inside or outside of the disc is simpler than the equivalent routine for the polygon. Discovering whether \mathbf{P} lies inside the disc requires you to check only that:

$$(\mathbf{P} - \mathbf{Q}) \cdot (\mathbf{P} - \mathbf{Q}) \leq r^2 \quad (60)$$

3.3.8 Box

A box is essentially six polygons and could be ray traced as such. However, intersection with an axis-aligned box can be optimised. Any box can be axis-aligned by appropriate transformations. We can thus write a routine to intersect an arbitrary ray with an axis-aligned box and then transform the ray under consideration in exactly the same way as we transform the box which we are trying to intersect with it. This sort of idea generalises neatly to the concept of specifying any object in a convenient *object coordinate system* and then applying transforms to the whole object to place it at the appropriate point in the *world coordinate system*.

3.4 Intersections with arbitrarily positioned primitives

Of the above primitives, only the plane and polygon are arbitrarily defined⁴¹. The sphere, cylinder, cone, and torus are all defined as being centred at the origin, and all have other restrictions on their definition⁴². In order to ray trace one of these primitives in an arbitrary location we have two alternatives: (1) find general intersection algorithms between a ray and the arbitrarily located versions of the primitives; or (2) use geometric transforms to scale, rotate, and translate these primitives into the desired locations. This second option will be followed here.

⁴¹The disc is also arbitrarily defined if you want a circular disc, but it could be transformed to provide an ellipse.

⁴²The sphere has unit radius; the cylinder has unit radius and is aligned with the z -axis; the cone is similarly aligned and has unit slope; the torus is aligned with the z -axis.

3.4.1 Transforming the primitives

The basic idea is simple. We specify a scaling, a rotation, and a translation which, between them, transform the primitive from its standard position to the desired location. You will remember that this was covered in the IB *Computer Graphics & Image Processing* course. To perform the intersection we take the inverse transform of the ray, intersect this with the primitive in its standard position, and then transform the resulting intersection point to its correct location.

Remember that there is a difference between points and displacements. When transforming objects, displacements are scaled and rotated but *not* translated. Think of it this way: if you translate two points, the displacement between them stays exactly the same. But if you scale or rotate the two points, the displacement between them scales or rotates accordingly.

All this is by way of explaining how we transform a ray in order to intersect the appropriate ray with the primitive in its standard position. Let us assume that the primitive object in its standard position, $\hat{\mathbf{B}}$, undergoes transformation⁴³ TRS to get into the desired position $\mathbf{B} = \mathbf{TRS}\hat{\mathbf{B}}$. To intersect ray, $\mathbf{E} + t\mathbf{D}$, with \mathbf{B} we transform the point, \mathbf{E} , and the displacement \mathbf{D} as follows:

$$\hat{\mathbf{E}} = \mathbf{S}^{-1}\mathbf{R}^{-1}\mathbf{T}^{-1}\mathbf{E} \quad (61)$$

$$\hat{\mathbf{D}} = \mathbf{S}^{-1}\mathbf{R}^{-1}\mathbf{D} \quad (62)$$

The point is translated, rotated and scaled but the displacement is only rotated and scaled.

Now intersect ray, $\hat{\mathbf{E}} + t\hat{\mathbf{D}}$, with the object in its standard position, $\hat{\mathbf{B}}$, as described in the previous sections. This gives the value of t and consequently allows you to directly calculate $\mathbf{P} = \mathbf{E} + t\mathbf{D}$. You do not even have to transform back because t has the same value in both standard and transformed coordinate systems.

In addition to scaling, rotating and translating the standard primitives, this mechanism allows us to stretch and squash them by scaling them by different factors in the different dimensions. For example, an ellipsoid is simply a stretched sphere. However, such anisotropic scaling causes interesting problems with the normal vectors, as discussed in the following section.

3.4.2 Transforming normal vectors

In ray tracing we need to know not just the intersection point but also the normal vector to the surface at the intersection point. This normal vector is used in the illumination calculations. However, if we scale an object anisotropically (different amounts in each coordinate), the normal vector scales as the *inverse* of the object scaling, although it rotates in the *same* way as the object. Figure 16 graphically illustrates this phenomenon.

One way to understand this behaviour is through the following mathematics. Remember that the normal vector is related to the derivative of the surface. The normal vector is perpendicular to the surface, which means that it is perpendicular to

⁴³Where \mathbf{T} is a translation, \mathbf{R} a rotation, and \mathbf{S} a scaling.




Figure 16: At left is a circle and a normal vector to the circle at a particular point. If we scale the circle by a factor of two horizontally, the normal vector scales by a factor of a half in the same dimension.

any derivative vector. As an example, consider the ellipsoid created by scaling the unit sphere by the matrix:

$$\mathbf{S} = \begin{bmatrix} l & 0 & 0 & 0 \\ 0 & m & 0 & 0 \\ 0 & 0 & n & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (63)$$

The intersection point between the unit sphere and the inverse transformed ray will be at some point $\hat{\mathbf{P}} = (\hat{x}, \hat{y}, \hat{z})$. This equates to the spherical polar coordinate $(1, \theta, \phi)$ where⁴⁴ $(\hat{x}, \hat{y}, \hat{z}) = (\cos \phi \cos \theta, \cos \phi \sin \theta, \sin \phi)$. By virtue of the fact that the normal to every point on the sphere passes through the centre of the sphere, it is easy to see that the normal vector, $\hat{\mathbf{N}}$, is also $\hat{\mathbf{N}} = (\cos \phi \cos \theta, \cos \phi \sin \theta, \sin \phi)$.

Transforming $\hat{\mathbf{P}}$ to the true intersection point \mathbf{P} is simply a matter of applying $\mathbf{P} = \mathbf{S}\hat{\mathbf{P}}$. Thus the intersection point of the ray with the ellipsoid is at rectangular coordinate $(x, y, z) = (l \cos \phi \cos \theta, m \cos \phi \sin \theta, n \sin \phi)$.

To find the normal vector to the ellipsoid at this intersection point, we can find two non-parallel derivative vectors to the surface at the intersection point, and take their cross product to give the normal vector, \mathbf{N} . Two derivative vectors are:

$$\frac{\partial(\mathbf{P})}{\partial \phi} = \frac{\partial(x, y, z)}{\partial \phi} = (-l \sin \phi \cos \theta, -m \sin \phi \sin \theta, n \cos \phi) \quad (64)$$

$$\frac{\partial(\mathbf{P})}{\partial \theta} = \frac{\partial(x, y, z)}{\partial \theta} = (-l \cos \phi \sin \theta, m \cos \phi \cos \theta, 0) \quad (65)$$

This leads to the normal vector being⁴⁵:

$$\mathbf{N} = \frac{\partial(\mathbf{P})}{\partial \phi} \times \frac{\partial(\mathbf{P})}{\partial \theta} \quad (66)$$

$$= \left(\frac{1}{l} \cos \phi \cos \theta, \frac{1}{m} \cos \phi \sin \theta, \frac{1}{n} \sin \phi \right) \quad (67)$$

Thus, we conclude that while:

$$\mathbf{P} = \mathbf{S}\hat{\mathbf{P}} \quad (68)$$

⁴⁴If you are wondering where r has disappeared to, remember that, in a unit sphere, $r = 1$.

⁴⁵If you try this at home you will find that you will need to divide through by the factor $-1/(lmn \cos \phi)$.




Figure 17: Three ways to split a polygon into triangles: (a) define a central vertex and connect every edge to this vertex to make n isosceles triangles; (b) select one vertex on the n -gon, and make a triangle fan emanating from this vertex; or (c) start at one edge of the n -gon and make a triangle strip set which proceeds from this edge of the polygon to the opposite edge.

$$\mathbf{N} = \mathbf{S}^{-1}\hat{\mathbf{N}} \quad (69)$$

This leaves open the question of what to do about the rotation and translation components of the transformation. It is intuitively obvious that rotating an object causes the normal vector to rotate by the same amount and, because normal vectors are displacements rather than points, they should not be translated. So the final analysis is that:

$$\mathbf{P} = \mathbf{TRSP} \quad (70)$$

$$\mathbf{N} = \mathbf{RS}^{-1}\hat{\mathbf{N}} \quad (71)$$

This piece of mathematics is specific to the unit sphere, but the same result holds for any object.

3.5 Converting the primitives to polygons

We may be in a situation where we define objects in terms of the various primitives, but where we wish to draw the objects using polygon scan conversion. In this case we need to convert primitives into polygons. Some polygon scan conversion algorithms (notably their implementation in hardware in graphics cards) can only deal with triangles; in these cases we may need to do a little bit of extra work to ensure that all of the generated polygons are triangles.

Converting the curved primitives (sphere, cylinder, cone, torus, disc) involves approximating a curved profile by a series of straight line segments. The simplest example is the disc. This can be approximated by a regular n -gon, where n is chosen to give an adequate approximation to the disc. “Adequate” in this case will depend on the desired rendering resolution, the desired speed of rendering, and the desired quality of the final image. If necessary, this n -gon can be converted to triangles in one of a number of ways, as shown in Figure 17.

A cylinder or cone can be converted to polygons by polygonising the discs at either end into n -gons, and then connecting corresponding vertices on the two n -gons. In the special case of a cone with a point, one of the n -gons obviously degenerates to that point.

Spheres and tori can be most easily converted to polygons by considering their parameterisation in terms of θ and ϕ . For a sphere these are Equations 72–74 (below); for a torus they are Equations 53–55.

$$x = r \cos \phi \cos \theta \quad (72)$$

$$y = r \cos \phi \sin \theta \quad (73)$$

$$z = r \sin \phi \quad (74)$$

By selecting appropriate steps in the two parameters we can generate a set of quadrilaterals which approximates the curved primitive.

It should be noted that spheres can be polygonised with more uniform polygons by starting with one of the five Platonic solids, and subdividing its faces accordingly. The details of this are left as an exercise to the reader.

3.6 Conics, quadrics, and superquadrics

The ray tracing primitives, described above, have relatively simple mathematical definitions. This is what makes them attractive: the simple mathematical definition allows for simple ray-object intersection code. Following from this, it would seem logical to investigate other shapes with simple mathematical definitions. Spheres, cones and cylinders are part of a more general family of parametric surfaces called *quadrics* (N.B. tori are *not* quadrics). Quadrics are the 3D analogue of 2D conics. We describe these general families below, but it turns out that they are of little practical use. It would seem that the general quadrics are a “dead end” in graphics research.

3.6.1 Conics

A conic is a two dimensional curve described by the general equation:

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0$$

This general form can be rotated, scaled, and translated so that it is aligned along the axes of the coordinate system. It will then have the simpler equation:

$$ax^2 + by^2 = k \text{ or } ax^2 + by = k$$

The useful conics are the ellipse (of which the circle is a special case), the hyperbola, and the parabola. For more details see **R&A** section 4-10, especially Table 4-8 on page 242. Table 4-8 is included in the handout.

3.6.2 Quadrics

The quadrics are the three dimensional analogue of the conics. The general equation is:

$$Ax^2 + By^2 + Cz^2 + Dxy + Eyz + Fzy + Gx + Hy + Jz + K = 0$$

This general form can be rotated, scaled, and translated so that it is aligned along the axes of the coordinate system. It will then have the simpler equation:

$$ax^2 + by^2 + cz^2 = k \text{ or } ax^2 + by^2 + cz = k$$




Figure 18: Superquadric functions for $n = 0.5, 1, 1.5, 2, 4, 8, 100$.

The useful conics are the ellipsoid (of which the sphere is a special case), the infinite cylinder, and the infinite cone. Various hyperboloids, and paraboloids are also defined by these equations, but these have little real use unless one is designing satellite dishes (paraboloid), headlamp reflectors (also paraboloid), or power station cooling towers (hyperboloid). For more details see **R&A** section 6-4, especially Figure 6-18 on page 403. Figure 6-18 is included in the handout.

3.6.3 Superquadrics

These are an extension of quadrics, where the power on the coordinate does not have to be 2. The general form of a superquadric centred at the origin and aligned along the coordinate axes is:

$$(ax)^n + (by)^n + (cz)^n = k$$

Super-ellipsoids tend to be the only members of this family that are actually used, and even they are only used in very limited areas. The effect of n on a super-ellipsoid is roughly as follows: $n = 2$ is a standard ellipsoid; $n > 2$ becomes closer to a box as n increases; in the limit, as $n \rightarrow \infty$, the shape does become a box; $n < 2$ is a more pointy version, the “points” being along the main axes; $n = 1$ is an octahedral shape; and $n < 1$ has sharp points on the main axes. Figure 18 shows examples of the shapes of these curves.

The interested student may like to have a quick look at Alan Barr’s two papers on superquadrics. The papers can be found in the Computer Laboratory library in *IEEE Transactions on Computer Graphics and Applications* volume 1, number 1 (January 1981), pages 11-23, and volume 1, number 3 (July 1981), pages 41-47.

3.7 Exercises

1. Give mathematical equations which define a plane, a sphere, an infinitely long cylinder, an infinitely long cone, and a torus. You will find it helpful to centre each primitive at the origin and to align it in a sensible way with respect to the coordinate axes.

2. Show how to intersect a ray with each of the five primitives from Question 1. You may assume that you are provided with functions to find the roots of linear, quadratic, cubic and quartic equations. Also show how to compute the normal vector at the intersection point.
3. Show how to intersect a ray with a finite length closed cylinder. Ensure that you handle all special cases, including that of a ray which is parallel to the axis of a finite length cylinder. Give both intersection point and normal vector for all cases in which an intersection occurs.
4. Give a complete algorithm for intersecting a ray with a finite length closed cone, including calculation of both intersection point and normal vector.
5. Work out if there exists a faster intersection algorithm for an axis aligned $2 \times 2 \times 2$ unit box than just six separate polygon intersection calculations.
6. Show how to convert a cylinder into a polygon mesh. What changes do you have to make if the mesh may contain only triangles?
7. Show how to convert a torus into a polygon mesh.
8. Show how to convert a sphere into a triangle mesh. How can you get the most even distribution of triangle vertices across the sphere?
9. [1999/7/11] (a) Give a parametric definition of a torus centred at the origin and aligned with the coordinate axes. (b) Outline how you would find the first intersection point, if any, of a ray with the torus from the previous part.
10. [2000/9/4] (a) Show how you would calculate the first intersection point between an arbitrary ray and a finite length open cylinder of unit radius aligned along the x -axis. [**Note:** an open cylinder is one which has no end caps.] Having found the intersection point, how would you find the normal vector?
11. [2001/7/9] (b) (i) Show how to find the first intersection between a ray and a finite-length, open-ended cone, centred at the origin, aligned along the x -axis, for which both ends of the finite-length are on the positive x -axis (i.e. $0 < x_{\min} < x_{\max}$). [6 marks]
(ii) Extend this to cope with a closed cone (i.e. the same cone, but with end caps). Take care to consider any special cases. [5 marks]
(iii) Extend this further to give the normal vector at the intersection point. [3 marks]
12. [2002/7/9] (a) A disc is a finite, planar, circular object. Describe an algorithm to find the point of intersection of an arbitrary ray with an arbitrary disc in three dimensions. Ensure that you describe the parameters used to define both the ray and the disc. [6 marks]
(b) Given the above algorithm and an algorithm to find the intersection of an arbitrary ray with a finite-length open cylinder, a programmer has two choices for implementing an algorithm to find the intersection with a finite-length closed cylinder. She could simply use the finite-length open cylinder primitive and two disc

primitives. Alternatively she could implement the finite-length closed cylinder as a primitive in its own right by adding extra code to the open cylinder algorithm. Compare the two alternatives in terms of efficiency and accuracy. [4 marks]

4 Introduction to splines

While the above primitives allow us to specify particular types of curved surface, we find ourselves in need of some more general way of specifying arbitrary curved surfaces. We want some mechanism which allows us to specify any smooth curved surface which we desire. This problem was first faced in the 1960s for the design of aeroplanes and cars. We will look at three solutions: Bézier surfaces, B-spline surfaces (including NURBS) and subdivision surfaces. It transpires that one of the most important problems is getting different patches of surface to connect together smoothly, that is: with continuity of position ($C0$), slope ($C1$) and curvature ($C2$). These are continuity of the function, its first and its second derivatives, respectively. Much of the ensuing discussions consider how to achieve such continuity.

The course handout contains a slide presentation introducing the concepts in this part of the course.

5 Bézier curves

Bézier curves were covered in the Part IB *Computer Graphics and Image Processing* course. This section gives some of the mathematical details, as does **R&A** Section 5-8. Parts of this Section of **R&A** are included in the handout.

If you want to experiment with Bézier curves then there are a number of on-line tutorials. One such is available from the Technion in Israel⁴⁶.

A Bézier curve is a weighted sum of $n + 1$ control points, $\mathbf{P}_0, \mathbf{P}_1, \dots, \mathbf{P}_n$, where the weights are the Bernstein polynomials:

$$\mathbf{P}(t) = \sum_{i=0}^n \binom{n}{i} (1-t)^{n-i} t^i \mathbf{P}_i, 0 \leq t \leq 1 \quad (75)$$

The Bézier curve of order $n + 1$ (degree n) has $n + 1$ control points. Below are the first three orders of Bézier curve definitions.

$$\text{linear} \quad \mathbf{P}(t) = (1-t)\mathbf{P}_0 + t\mathbf{P}_1 \quad (76)$$

$$\text{quadratic} \quad \mathbf{P}(t) = (1-t)^2\mathbf{P}_0 + 2(1-t)t\mathbf{P}_1 + t^2\mathbf{P}_2 \quad (77)$$

$$\text{cubic} \quad \mathbf{P}(t) = (1-t)^3\mathbf{P}_0 + 3(1-t)^2t\mathbf{P}_1 + 3(1-t)t^2\mathbf{P}_2 + t^3\mathbf{P}_3 \quad (78)$$

5.1 Ways of thinking about Bézier curves

There are several useful ways in which you can think about Bézier curves. Here are the ones that I use.

⁴⁶<http://www.cs.technion.ac.il/~cs234325/Homepage/Applets/applets/bezier/html/>

Linear interpolation. Equation 76 is obviously a linear interpolation between two points. Equation 77 can be rewritten as a linear interpolation between linear interpolations between points:

$$\mathbf{P}(t) = (1-t)[(1-t)\mathbf{P}_0 + t\mathbf{P}_1] + t[(1-t)\mathbf{P}_1 + t\mathbf{P}_2] \quad (79)$$

Equation 78 can be rewritten as a linear interpolation between linear interpolations between linear interpolations between points. This is left as an exercise for the reader.

Weighted average. A Bézier curve can be seen as a weighted average of all of its control points. Because all of the weights are positive, and because the weights sum to one, the Bézier curve is guaranteed to lie within the convex hull of its control points.

Refinement of the control polygon. A Bézier curve can be seen as some sort of refinement of the polygon made by connecting its control points in order. The Bézier curve starts and ends at the two end points and its shape is determined by the relative positions of the $n-1$ other control points, although it will generally not pass through any of these other control points. The tangent vectors at the start and end of the curve pass through the end point and the immediately adjacent point.

Rogers and Adams list the properties of the Bézier curve on page 291.

5.2 Continuity

You should note that each Bézier curve is independent of any other Bézier curve. If we wish two Bézier curves to join with any type of continuity, then we must explicitly position the control points of the second curve so that they bear the appropriate relationship with the control points in the first curve.

Any Bézier curve is infinitely differentiable within itself, and is therefore continuous to any degree (C^n -continuous, $\forall n$). We therefore only need concern ourselves with continuity across the joins between curves. Assume that we have two Bézier curves of the same order: $\mathbf{P}(t)$, defined by $(\mathbf{P}_0, \mathbf{P}_1, \dots, \mathbf{P}_n)$, and $\mathbf{Q}(t)$, defined by $(\mathbf{Q}_0, \mathbf{Q}_1, \dots, \mathbf{Q}_n)$. C^0 -continuity (continuity of position) can be achieved by setting $\mathbf{P}(1) = \mathbf{Q}(0)$. This gives a formula for \mathbf{Q}_0 in terms of the \mathbf{P}_i s:

$$\mathbf{Q}_0 = \mathbf{P}_n. \quad (80)$$

Similarly for C^1 -continuity, we need C^0 -continuity and $\mathbf{P}'(1) = \mathbf{Q}'(0)$, giving:

$$\mathbf{Q}_1 - \mathbf{Q}_0 = \mathbf{P}_n - \mathbf{P}_{n-1} \quad (81)$$

Combining Equations 81 and 80 gives a formula for \mathbf{Q}_1 in terms of the \mathbf{P}_i s:

$$\mathbf{Q}_1 = 2\mathbf{P}_n - \mathbf{P}_{n-1} \quad (82)$$

$$= \mathbf{P}_n + (\mathbf{P}_n - \mathbf{P}_{n-1}) \quad (83)$$

Continuing in this vein, we find that the requirements for C^2 -continuity (i.e. C^1 -continuity and $\mathbf{P}''(1) = \mathbf{Q}''(0)$) give:

$$\mathbf{Q}_2 - 2\mathbf{Q}_1 + \mathbf{Q}_0 = \mathbf{P}_n - 2\mathbf{P}_{n-1} + \mathbf{P}_{n-2} \quad (84)$$

Combining Equations 84, 81, and 80 gives a formula for \mathbf{Q}_2 in terms of the \mathbf{P}_i s:

$$\mathbf{Q}_2 = 4\mathbf{P}_n - 4\mathbf{P}_{n-1} + \mathbf{P}_{n-2} \quad (85)$$

$$= \mathbf{P}_{n-2} + 4(\mathbf{P}_n - \mathbf{P}_{n-1}) \quad (86)$$

5.3 Bézier surfaces

We learnt in the IB course that the simplest way to construct a Bézier surface is as the tensor product of Bézier curves. A tensor product Bézier surface of order $n + 1$ is defined by $(n + 1)^2$ control points. It is called a Bézier patch.

$$\mathbf{P}(s, t) = \sum_{i=0}^n \binom{n}{i} (1-s)^{n-i} s^i \sum_{j=0}^n \binom{n}{j} (1-t)^{n-j} t^j \mathbf{P}_{i,j} \quad (87)$$

You can think about this as moving the control points of one Bézier curve along a set of Bézier curves to sweep out a surface. Continuity across a boundary between two Bézier patches is only guaranteed if each of the Bézier curves across the join obey the curve continuity conditions. Again, this was covered in the IB course.

5.4 Exercises

1. Explain what C_0 -, C_1 -, C_2 -, C_n -continuity mean.
2. Derive the constraints on control point positions which ensure that two quartic Bézier curves join with (a) C_0 -continuity, (b) C_1 -continuity, and (c) C_2 -continuity.

6 B-splines

B-splines are covered in some detail below and also in **R&A** Section 5-9. Parts of this Section of **R&A** are included in the handout. Beware that none of the worked examples are in the handout. These may come in useful, and you will need to get hold of a real copy of **R&A** if you wish to work your way through them.

B-splines are a more general type of curve than Bézier curves. In a B-spline each control point is associated with a *basis function*, $N_{i,k}$.

$$\mathbf{P}(t) = \sum_{i=1}^{n+1} N_{i,k}(t) \mathbf{P}_i, t_{\min} \leq t < t_{\max} \quad (88)$$

There are $n + 1$ control points, $\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_{n+1}$. The $N_{i,k}$ basis functions are of order k (degree $k - 1$). k must be at least 2 (linear), and can be no more than $n + 1$ (the number of control points). The important point here is that the order of the curve (2 [linear], 3 [quadratic], 4 [cubic],...) is therefore not dependent on the number of control points (which it is for Bézier curves, where k must always equal $n + 1$).

Equation 88 defines a piecewise continuous function. The $N_{i,k}$ are defined by a *knot vector*, $(t_1, t_2, \dots, t_{k+(n+1)})$, must be specified. This determines the values of t at which the pieces of curve join, like knots joining bits of string. It is necessary that:

$$t_i \leq t_{i+1}, \forall i \quad (89)$$

The $N_{i,k}$ depend *only* on the value of k and the values, t_i , in the knot vector. $N_{i,k}$ is defined recursively as:

$$\begin{aligned} N_{i,1}(t) &= \begin{cases} 1, & t_i \leq t < t_{i+1} \\ 0, & \text{otherwise} \end{cases} \\ N_{i,k}(t) &= \frac{t - t_i}{t_{i+k-1} - t_i} N_{i,k-1}(t) + \frac{t_{i+k} - t}{t_{i+k} - t_{i+1}} N_{i+1,k-1}(t) \end{aligned} \quad (90)$$

This is essentially a modified version of the idea of taking linear interpolations of linear interpolations of linear interpolations. . .

At this point it would be instructive for you to work out $N_{1,1}$, $N_{2,1}$, $N_{3,1}$, $N_{1,2}$, $N_{2,2}$, $N_{1,3}$ for the knot vector $[0, 2, 3, 6]$. It helps if you draw the graphs for these functions.

There are several things that you should note about these equations. Each $N_{i,k}(t)$ depends only on the $k + 1$ knot values from t_i to t_{i+k} . $N_{i,k}(t) = 0$ for $t < t_i$ or $t \geq t_{i+k}$ so P_i only influences the curve for $t_i \leq t < t_{i+k}$. Formally, $P(t)$ is a polynomial of order k (degree $k - 1$) on each interval $t_i \leq t < t_{i+1}$. Across the knots $P(t)$ is C^{k-2} -continuous. $P(t)$ is, of course, continuous in all its derivatives between the knots. Remember from Equation 16 that a weighted sum of points only makes sense if the weights sum to one. $P(t)$ is therefore validly defined only where

$$\sum_{i=1}^{n+1} N_{i,k}(t) = 1. \quad (91)$$

This is the range $t_{\min} \leq t < t_{\max}$ where $t_{\min} = t_k$ and $t_{\max} = t_{n+2}$. Even more properties of B-splines are described in Rogers and Adams pp. 306–7.

6.1 The knot vector

The above explanation shows that the knot vector is very important. The knot vector can, by its definition, be any sequence of numbers provided that each one is greater than or equal to the preceding one. Some types of knot vector are more useful than others. Knot vectors are generally placed into one of three categories: uniform, open uniform, and non-uniform.

Uniform. These are knot vectors for which

$$t_{i+1} - t_i = \text{constant}, \forall i \quad (92)$$

For example:

$$\begin{aligned} &[1, 2, 3, 4, 5, 6, 7, 8] \\ &[0, 1, 2, 3, 4, 5] \\ &[0, 0.25, 0.5, 0.75, 1.0] \\ &[-2.5, -1.4, -0.3, 0.8, 1.9, 3.0] \end{aligned}$$

All of the basis functions are just shifted versions of one another and so the implementation is very easy.

Open Uniform. These are uniform knot vectors which have k equal knot values at each end:

$$\begin{aligned} &t_i = t_1, \quad i \leq k \\ &t_{i+1} - t_i = \text{constant}, \quad k \leq i < n + 2 \\ &t_i = t_{k+(n+1)}, \quad i \geq n + 2 \end{aligned} \quad (93)$$

For example:

$$\begin{aligned} [0, 0, 0, 0, 1, 2, 3, 4, 4, 4, 4] & \quad (k = 4) \\ [1, 1, 1, 2, 3, 4, 5, 6, 6, 6] & \quad (k = 3) \\ [0.1, 0.1, 0.1, 0.1, 0.1, 0.3, 0.5, 0.7, 0.7, 0.7, 0.7, 0.7] & \quad (k = 5) \end{aligned}$$

This is essentially just a simple modification to the uniform case which allows the curve to go through its two end points.

Non-uniform. This is the general case, the only constraint being the standard $t_i \leq t_{i+1}, \forall i$ (Equations 89). For example:

$$\begin{aligned} [1, 3, 7, 22, 23, 23, 49, 50, 50] \\ [1, 1, 1, 2, 2, 3, 4, 5, 6, 6, 6, 7, 7, 7] \\ [0.2, 0.7, 0.7, 0.7, 1.2, 1.2, 2.9, 3.6] \end{aligned}$$

The shapes of the $N_{i,k}$ basis functions are determined entirely by the *relative* spacing between the knots. Scaling ($t'_i = \alpha t_i, \forall i$) or translating ($t'_i = t_i + \Delta t, \forall i$) the knot vector has no effect on the shapes of the $N_{i,k}$ nor on the shape of the actual curve $\mathbf{P}(t)$.

The above gives a description of the various types of knot vector but it doesn't really give you any insight into how the knot vector determines the shape of the curve. The following subsections look at the different types of knot vector in more detail. However, the best way to get to feel for these is to derive and draw the basis functions yourself.

6.1.1 Uniform knot vector

For simplicity, let $t_i = i$ (this is allowable given that the scaling or translating the knot vector has no effect on the shapes of the $N_{i,k}$). The knot vector thus becomes $[1, 2, 3, \dots, k + (n + 1)]$ and Equation 90 simplifies to:

$$\begin{aligned} N_{i,1}(t) &= \begin{cases} 1, & i \leq t < i + 1 \\ 0, & \text{otherwise} \end{cases} \\ N_{i,k}(t) &= \frac{t - i}{k - 1} N_{i,k-1}(t) + \frac{i + k - t}{k - 1} N_{i+1,k-1}(t) \end{aligned} \quad (94)$$

You should be easily able to graph the first few of these for yourself. The principle thing to note about the uniform basis functions is that, for a given order k , the basis functions are all simply shifted versions of one another. See Rogers and Adams Figure 5-36.

6.1.2 Things you can change about a uniform B-spline

With a uniform B-spline, you obviously cannot change the basis functions (they are fixed because all the knots are equispaced). However you can alter the shape of the curve by modifying a number of things:

Moving control points. Moving the control points obviously changes the shape of the curve.

Multiple control points. Sticking two adjacent control points on top of one another causes the curve to pass closer to that point. Stick enough adjacent control points on top of one another and you can make the curve pass through that point (Rogers and Adams, Figure 5-45).

Order. Increasing the order k increases the continuity of the curve at the knots, increases the smoothness of the curve, and tends to move the curve farther from its defining polygon. (Rogers and Adams, Figure 5-44).

Joining the ends. You can join the ends of the curve to make a closed loop. Say you have M points, P_1, \dots, P_M . You want a closed B-spline defined by these points. For a given order, k , you will need $M + (k - 1)$ control points (repeating the first $k - 1$ points): $P_1, \dots, P_M, P_1, \dots, P_{k-1}$. Your knot vector will thus have $M + 2k - 1$ uniformly spaced knots.

6.1.3 Open uniform knot vector


The previous section intimated that uniform B-splines can be used to describe closed curves: all you have to do is join the ends as described above. If you do not want a closed curve, and you use a uniform knot vector, you find that you need to specify control points at each end of the curve which the curve doesn't go near (e.g. Rogers and Adams, Figure 5-44, the order 4 curve).

If you wish your B-spline to start and end at your first and last control points then you need an open uniform knot vector (e.g. Rogers and Adams, Figure 5-41). The only difference between this and the uniform knot vector being that the open uniform version has k equal knots at each end.

An order k open uniform B-spline with $n + 1 = k$ points is the Bézier curve of order k . It would be a useful exercise for you to prove this for $k = 3$. For ease of calculation take the knot vector to be $[0, 0, 0, 1, 1, 1]$.


6.1.4 The difference between uniform and open uniform

It may help, at this stage, to compare a particular uniform and an equivalent open uniform knot vector. This is a uniform knot vector for $n + 1 = 7$, $k = 3$:



The lines show the range of t over which each P_i is non-zero. The B-spline itself (the *overall* line in the diagram) is defined over the range $t_3 \leq t < t_8$, i.e. over the range $3 \leq t < 8$.

By comparison an open uniform knot vector for $n + 1 = 7$, $k = 3$ is:



The B-spline itself is defined over the range $t_3 \leq t < t_8$, i.e. over the range $1 \leq t < 6$. By the definition of an open uniform knot vector $t_3 = t_1$ and $t_8 = t_{10}$ and so an open uniform B-spline is defined over the *full* range of t from t_1 to $t_k + n + 1$.

6.1.5 Non-uniform knot vector

Any B-spline whose knot vector is neither uniform nor open uniform is non-uniform. Non-uniform knot vectors allow any spacing of the knots, including multiple knots (adjacent knots with the same value). We need to know how this non-uniform spacing affects the basis functions in order to understand where non-uniform knot vectors could be useful. It transpires that there are only three cases of any interest: (1) multiple knots (adjacent knots equal); (2) adjacent knots more closely spaced than the next knot in the vector; and (3) adjacent knots less closely spaced than the next knot in the vector. Obviously, case (3) is simply case (2) turned the other way round.

Multiple knots. A multiple knot reduces the degree of continuity at that knot value. Across a normal knot the continuity is C^{k-2} . Each extra knot with the same value reduces continuity at that value by one. This is the only way to reduce the continuity of the curve at the knot values. If there are $k - 1$ (or more) equal knots then you get a discontinuity in the curve.

Close knots. As two knots' values get closer together, relative to the spacing of the other knots, the curve moves closer to the related control point.

Distant knots. As two knots' values get further apart, relative to the spacing of the other knots, the curve moves further away from the related control point.

6.1.6 Use of non-uniform knot vectors

Standard procedure is to use uniform or open uniform B-splines unless there is a very good reason not to do so. Moving two knots closer together tends to move the curve only slightly and so there is usually little point in doing it. This leads to the conclusion that the main use of non-uniform B-splines is to allow for multiple knots, which adjust the continuity of the curve at the knot values.

However, non-uniform B-splines are the general form of the B-spline because they incorporate open uniform and uniform B-splines as special cases. Thus we will talk about *non-uniform B-splines* when we mean the general case, incorporating both uniform and open uniform.

6.1.7 What can you do to control the shape of a B-spline?

- Move the control points.
- Add or remove control points.
- Use multiple control points.
- Change the order, k .
- Change the type of knot vector.
- Change the relative spacing of the knots.
- Use multiple knot values in the knot vector.

6.1.8 What should the defaults be?

If there are no pressing reasons for doing otherwise, your B-spline should be defined as follows:

- $k = 4$ (cubic);
- no multiple control points;
- uniform (for a closed curve) or open uniform (for an open curve) knot vector.

6.2 B-spline patches

We generalise from B-spline curves to B-spline surfaces in the same way as we did for Bézier patches. Take a tensor product of two versions of Equation 88.

$$\mathbf{P}(s, t) = \sum_{i=1}^{m+1} \sum_{j=1}^{n+1} \mathbf{P}_{i,j} N_{i,k}(s) N_{j,l}(t), \quad s_{\min} \leq s < s_{\max}, t_{\min} \leq t < t_{\max} \quad (95)$$

where it is usual for the patch to have the same order (i.e. $k = l$) in both directions. Patches are thus defined by a quadrilateral grid of control points of size $(m+1) \times (n+1)$.

6.3 Why B-splines?

B-splines have many nice properties when compared to other families of curves which could be used. They:

- minimise the order of the polynomial pieces (order k)
- maximise the continuity between pieces (continuity $C(k-2)$)
- minimise the number of control points controlling a piece (k points)
- have positive basis functions
- have basis functions which partition unity, implying that each piece lies inside its control points' convex hull
- are invariant with respect to affine transforms

6.4 Exercises

1. How many control points are required for a quartic Bézier and how many for a quartic B-spline?
2. Why are cubics the default for B-spline use?
3. Explain the difference between Uniform, Open Uniform, and Non-Uniform knot vectors. What are the advantages of each type?
4. [2000/9/4] (b) A non-rational B-spline has knot vector $[1, 2, 4, 7, 8, 10, 12]$. Derive the first of the third order (second degree) basis functions, $N_{1,3}(t)$, and graph it. If this knot vector were used to draw a third order B-spline, how many control points would be required? [7 marks]
5. [2001/8/4] (a) For a given order, k , there is only one basis function for uniform B-splines. Every control point uses a shifted version of that one basis function. How many different basis functions are there for open-uniform B-splines of order k with $n + 1$ control points, where $n \geq 2k - 3$? [6 marks]
 (b) Explain what is different in the cases where $n < 2k - 3$ compared with the cases where $n \geq 2k - 3$. [3 marks]
 (c) Sketch the different basis functions for $k = 2$ and $k = 3$ (when $n \geq 2k - 3$). [4 marks]
 (d) Show that the open-uniform B-spline with $k = 3$ and knot vector $[0, 0, 0, 1, 1, 1]$ is equivalent to the quadratic Bézier curve. [7 marks]
6. [2002/7/9] (d) Derive the formula of and sketch a graph of $N_{3,3}(t)$, the third of the quadratic B-spline basis functions, for the knot vector $[0, 0, 0, 1, 3, 3, 4, 5, 5, 5]$. [6 marks]

6.5 NURBS

NURBS are covered below and in some detail in **R&A** Section 5-13. Parts of this Section of **R&A** are included in the handout.

Non-uniform rational B-splines are the curves that are currently used in any graphics application that requires curves and surfaces with more functionality than Bézier curves can offer. In most cases, you would actually use the special case of non-rational B-splines (those described in the previous section) but it is useful to have the more general rational versions available for certain types of curve and surface. In addition to the features listed above for B-splines, NURBS are invariant with respect to perspective transforms.

NURBS are generally rendered by converting them to lots of small polygons and then using polygon scan conversion. They can also be ray traced, but a general analytic ray-NURBS intersection algorithm is a nightmare, so numerical techniques are used to find the intersection point.

NURBS curves incorporate – as special cases – uniform B-splines, non-rational B-splines, Bézier curves, lines, and conics. NURBS surfaces incorporate planes, quadrics, and tori. Note that this does not quite mean what it says. It is tricky to get NURBS

to represent *infinite* surfaces, but they can certainly represent finite sections of infinite surfaces such as planes, paraboloids, and hyperboloids.

If you want to experiment with NURBS curves then there are a number of on-line tutorials. One such is available from the Technion in Israel⁴⁷.

Rational B-splines have all of the properties of non-rational B-splines plus the following two useful features:

- They produce the correct results under projective transformations (while non-rational B-splines only produce the correct results under affine transformations).
- They can be used to represent lines, conics, non-rational B-splines; and, when generalised to patches, can represent planes, quadrics, and tori.

In this case *rational* means “one polynomial divided by another” (see Equation 96). The antonym of *rational* is *non-rational* (i.e. a non-rational B-spline is just a polynomial (see Equation 88). Non-rational B-splines are a special case of rational B-splines, just as uniform B-splines are a special case of non-uniform B-splines. Thus, *non-uniform rational B-splines* encompass almost every other possible 3D shape definition. *Non-uniform rational B-spline* is a bit of a mouthful and so it is generally abbreviated to *NURBS*.

We have already learnt all about the the *B-spline* bit of *NURBS* and about the *non-uniform* bit. So now all we need to know is the meaning of the *rational* bit and we will fully(?) understand NURBS.

Rational B-splines are defined simply by applying the B-spline equation (Equation 88) to homogeneous coordinates, rather than normal 3D coordinates. We discussed homogeneous coordinates in the IB course. You will remember that these are 4D coordinates where the transformation from 4D to 3D is:

$$(x', y', z', w) \rightarrow \left(\frac{x'}{w}, \frac{y'}{w}, \frac{z'}{w} \right) \quad (96)$$

Last year we said that the inverse transform was:

$$(x, y, z) \rightarrow (x, y, z, 1) \quad (97)$$

This year we are going to be more cunning and say that:

$$(x, y, z) \rightarrow (xh, yh, zh, h) \quad (98)$$

Thus our 3D control point, $\mathbf{P}_i = (x_i, y_i, z_i)$, becomes the homogeneous control point, $\mathbf{C}_i = (x_i h_i, y_i h_i, z_i h_i, h_i)$.

A NURBS curve is thus defined as:

$$\mathbf{P}_H(t) = \sum_{i=1}^{n+1} N_{i,k}(t) \mathbf{C}_i, t_{\min} \leq t < t_{\max} \quad (99)$$

Compare Equation 99 with Equation 88 to see just how easy this is!

We now want to see what a NURBS curve looks like in normal 3D coordinates, so we need to apply Equation 96 to Equation 99. In order to better explain what is going on, we

⁴⁷<http://www.cs.technion.ac.il/~cs234325/Homepage/Applets/applets/bspline/html/>

first write Equation 99 in terms of its individual components. Equation 99 is equivalent to:

$$x'(t) = \sum_{i=1}^{n+1} x_i h_i N_{i,k}(t) \quad (100)$$

$$y'(t) = \sum_{i=1}^{n+1} y_i h_i N_{i,k}(t) \quad (101)$$

$$z'(t) = \sum_{i=1}^{n+1} z_i h_i N_{i,k}(t) \quad (102)$$

$$h(t) = \sum_{i=1}^{n+1} h_i N_{i,k}(t) \quad (103)$$

Equation 96 tells us that, in 3D:

$$x(t) = x'(t)/h(t) \quad (104)$$

$$y(t) = y'(t)/h(t) \quad (105)$$

$$z(t) = z'(t)/h(t) \quad (106)$$

Thus the 4D to 3D conversion gives us the curve in 3D:

$$\mathbf{P}(t) = \frac{\sum_{i=1}^{n+1} N_{i,k}(t) \mathbf{P}_i h_i}{\sum_{i=1}^{n+1} N_{i,k}(t) h_i}, t_{\min} \leq t < t_{\max} \quad (107)$$

This looks a lot more fierce than Equation 99, but is simply the same thing written a different way.

So now, we need to define an additional parameter, h_i , for each control point, \mathbf{P}_i . The default is to set $h_i = 1, \forall i$. This results in the denominator of Equation 107 becoming one, and the NURBS equation (Equation 107) therefore reducing to the non-rational B-spline equation (Equation 88).

Increasing h_i pulls the curve closer to point \mathbf{P}_i . Decreasing h_i pushes the curve farther from point \mathbf{P}_i . Setting $h_i = 0$ means that \mathbf{P}_i has no effect on the curve at all. See Rogers and Adams Figure 5-58 for an example, and play with an on-line NURBS tutorials such as the one mentioned above.

6.6 An example: a circle defined by NURBS

This subsection provides an example of a shape which cannot be represented by non-rational B-splines: a circle. A non-rational B-spline or a Bézier curve cannot exactly represent a circle. An interesting exercise is to place a cubic Bézier curve's end points at $(0, 1)$ and $(1, 0)$, with the other control points at $(\alpha, 1)$ and $(1, \alpha)$. Now see how close this "quarter circle" comes to the real quarter circle defined by $x^2 + y^2 = 1$, i.e. what is the value of α for which the Bézier curve most closely matches the quarter circle. You will find that you can get a match which is almost, but not quite, circular.

NURBS *can* be used to represent circles, and all of the other conics. NURBS surfaces can be used to represent quadric surfaces. As an example, let us consider one way in which NURBS can be used to describe a true circle. Rogers and Adams cover this on

pages 371–375. The ways in which this is done require the designer to specify several things correctly at the same time, as we shall see. The details are so complicated that I would not expect you to remember it in an exam but I would expect you to remember that it can be done and have some idea of where to look it up if you needed it.

The method is as follows. Construct eight control points in a square. Let $\mathbf{P}_1, \mathbf{P}_3, \mathbf{P}_5,$ and \mathbf{P}_7 be the vertices of the square. Let $\mathbf{P}_0, \mathbf{P}_2, \mathbf{P}_4,$ and \mathbf{P}_6 be the midpoints of the respective sides, so that the vertices are numbered sequentially as you proceed around the square. Finally, you need a ninth point to join up the curve, so let $\mathbf{P}_8 = \mathbf{P}_0$.

Use a quadratic B-spline basis function with the knot vector $[0, 0, 0, 1, 1, 2, 2, 3, 3, 4, 4, 4]$. This means that the curve will pass through $\mathbf{P}_0, \mathbf{P}_2, \mathbf{P}_4, \mathbf{P}_6$ and \mathbf{P}_8 , and allows us to essentially treat each quarter of the circle independently. That is, we can just examine $\mathbf{P}_0, \mathbf{P}_1,$ and \mathbf{P}_2 , along with the knot vector $[0, 0, 0, 1, 1, 1]$. If this makes a quarter circle then the other three quarters will also be correct.

We finally need to specify the homogeneous co-ordinates. As a circle is symmetrical it should be obvious that that $h_1 = h_3 = h_5 = h_7 = \alpha$ and $h_0 = h_2 = h_4 = h_6 = h_8 = \beta$. As we would like the curve to pass through the even numbered points we know that $\beta = 1$. All we therefore need to determine is α , the value of the odd numbered homogeneous co-ordinates.

If $\alpha = 1$ then the NURBS curve will bulge out more than a circle. If $\alpha = 0$, it will bow in. This gives us limits on the value of α . To find the exact value we take the NURBS curve definition for the quarter circle:

$$\mathbf{P}(t) = \frac{(1-t)^2\mathbf{P}_0 + 2\alpha t(1-t)\mathbf{P}_1 + t^2\mathbf{P}_2}{(1-t)^2 + 2\alpha t(1-t) + t^2}, 0 \leq t < 1 \quad (108)$$

Assume now that $\mathbf{P}_0 = (0, 1)$, $\mathbf{P}_1 = (1, 1)$, and $\mathbf{P}_2 = (1, 0)$. Insert Equation 108 into the equation for the unit circle ($x(t)^2 + y(t)^2 = 1$). The resulting equation is:

$$\frac{((1-t)^2 + 2\alpha t(1-t))^2 + (2\alpha t(1-t) + t^2)^2}{((1-t)^2 + 2\alpha t(1-t) + t^2)^2} = 1, 0 \leq t < 1 \quad (109)$$

Now solve this for α . Equation 109 is essentially:

$$\frac{a_N t^4 + b_N t^3 + c_N t^2 + d_N t + e_N}{a_D t^4 + b_D t^3 + c_D t^2 + d_D t + e_D} = 1, 0 \leq t < 1 \quad (110)$$

From this we can conclude that we require $a_N = a_D, b_N = b_D, c_N = c_D, d_N = d_D,$ and $e_N = e_D$. The first three all solve to give the result that $\alpha = 1/\sqrt{2}$, while the last two cancel out totally to give the tautology $0 = 0$. Thus $\alpha = 1/\sqrt{2}$.

This derivation is not at all intuitive and similar cleverness is required to handle representations of other conics. The beauty of NURBS is that they allow us to do this sort of thing and unify all shapes into a single representation. The difficulty is that, in order to achieve this unification, we need to have this rather complicated but general mathematical mechanism.

6.7 Exercises

1. Review from IB: What are homogeneous coordinates and what are they used for in computer graphics?

2. Explain how to use homogeneous coordinates to get rational B-splines given that you know how to produce non-rational B-splines.
3. What are the advantages of NURBS over Bézier curves? (i.e. why have NURBS, in general, replaced Bézier curves in CAD?)
4. Show that you understand why NURBS includes Uniform B-splines, Non-Rational B-splines, Béziers, lines, conics, quadrics, and tori.
5. [1998/7/12] Consider the design of a user interface for a NURBS drawing system. Users should have access to the full expressive power of the NURBS representation. What things should users be able to modify to give them such access and what effect does each have on the resulting shape? [6 marks]
6. For each of the items (in the previous question) that the user can edit: (i) Give sensible default values; (ii) Explain how they would be constrained if a 'demo' version of the software was to be limited to cubic Uniform Non-rational B-Splines.
7. [1999/7/11] (c) Show how to construct a circle using non-uniform rational B-splines (NURBS). [8 marks]
 Note: this question is ludicrously hard unless you remember the worked example in these notes or **R&A** pages 371-375.
 (d) Show how the circle definition from the previous part can be used to define a NURBS torus. [4marks]
 Note: you need explain only the general principle and the location of the torus' control points.

7 Subdivision surfaces

Subdivision schemes work by taking a coarse polygon mesh and introducing new vertices to create a finer mesh. Iterating this process several times creates a very fine mesh of polygons. Given that we are interested in drawing things only to a certain level of accuracy (there is no point in having polygons that are much smaller than pixels), the easily understood subdivision idea has definite benefits over the mathematically complicated B-spline methods. In fact, two of the standard subdivision schemes (Doo-Sabin and Catmull-Clark) produce, in the limit, B-spline surfaces (uniform quadratic and uniform cubic respectively) except at their extraordinary points. Some of the mathematical detail of subdivision surfaces is given below. **W&W** survey the field and the related mathematical tools.

Subdivision schemes have been around for a long time. Subdivision methods for curves were first mathematically analysed in 1947. Their use in computer graphics dates from 1974 when Chaikin used them to derive a simple algorithm for generating curves quickly. In 1978 Doo & Sabin and Catmull & Clark generalised Chaikin's work from curves to surfaces. Much work has been done since then, but it seems that it is only since about 2000 that subdivision schemes have had widespread use owing to Pixar's adoption of them.

Subdivision schemes are increasingly being used as an alternative to NURBS. They combine mathematical elegance with an exceptionally simple implementation. For curves,

given an arbitrary control polygon, we use the positions of the current vertices to determine the location of the new vertices in a new, refined, more detailed, control polygon. Generally, each old vertex gives rise to two new vertices. For example, you could place new vertices one-quarter and three-quarters of the way between each adjacent pair of old vertices. Connecting all the new vertices together, in the appropriate order, produces a more refined control polygon. Repeat this process several times and you produce a very good approximation to the uniform quadratic B-spline curve defined by the original set of vertices. In the limit, the refined control polygon becomes this uniform quadratic B-spline curve. The Doo-Sabin subdivision method is the extension of this idea to surfaces, where the refined control polygon has four times as many vertices as the source control polygon. Given the simplicity of the implementation and the fact that you can stop whenever you like, you can see how attractive this method is for computer graphics.

The course handout contains a slide presentation that presents the concepts from this part of the course in an alternative way.

7.1 Mathematical details: curves

Take an arbitrary polygon defined by the sequence of control points:

$$\mathbf{P}^i = (\dots, \mathbf{p}_{-1}^i, \mathbf{p}_0^i, \mathbf{p}_1^i, \mathbf{p}_2^i, \dots)$$

Subdivision maps this sequence of control points to a new sequence, \mathbf{P}^{i+1} by applying subdivision rules. This process doubles⁴⁸ the number of points, and there is one rule for the odd numbered points and one for the even. For example, the subdivision rules on which the Doo-Sabin method is based are:

$$\mathbf{p}_{2j}^{i+1} = \frac{3}{4}\mathbf{p}_j^i + \frac{1}{4}\mathbf{p}_{j+1}^i \quad (111)$$

$$\mathbf{p}_{2j+1}^{i+1} = \frac{1}{4}\mathbf{p}_j^i + \frac{3}{4}\mathbf{p}_{j+1}^i \quad (112)$$

while the subdivision rules on which the Catmull-Clark method is based are:

$$\mathbf{p}_{2j}^{i+1} = \frac{1}{8}\mathbf{p}_{j-1}^i + \frac{6}{8}\mathbf{p}_j^i + \frac{1}{8}\mathbf{p}_{j+1}^i \quad (113)$$

$$\mathbf{p}_{2j+1}^{i+1} = \frac{4}{8}\mathbf{p}_j^i + \frac{4}{8}\mathbf{p}_{j+1}^i \quad (114)$$

As is the way with much mathematics, we can write it in a more compact, more general, but less obvious, form as:

$$\mathbf{p}_j^{i+1} = \sum_{k=-\infty}^{\infty} \alpha_{2k-j} \mathbf{p}_k^i \quad (115)$$

where the α_j are coefficients depending on the subdivision rules. Note that the index $2k - j$ alternately selects the even indexed α_j and the odd indexed α_j . So, the two schemes given above, can be compactly described as:

$$\alpha = \frac{1}{4}(\dots, 0, 0, 1, 3, 3, 1, 0, 0, \dots) \quad (116)$$

⁴⁸It doesn't quite double the number of points when the sequence is open and of finite length, but we will gloss over that at the moment.




Figure 19: Doo-Sabin subdivision. On left a mesh (solid dots and solid lines) that has been refined (open dots and dashed lines). At right the weights used to generated one of the refined vertices.

and

$$\alpha = \frac{1}{8}(\dots, 0, 0, 1, 4, 6, 4, 1, 0, 0, \dots) \quad (117)$$

respectively. You will recognise the sequences in parentheses as being two rows from Pascal's triangle.

It would now be constructive for you to draw an arbitrary control polygon and perform a couple of subdivision steps using the first of the two subdivision schemes. Once you feel happy that you understand what is going on, you may like to try the second scheme. For those for whom these two tasks seem simple, you may like to consider what happens if you try to use the previous row from Pascal's triangle (1,2,1) and, for even more excitement, what happens if you try to use the next row (1,5,10,10,5,1). Both produce valid subdivision methods, but you will find that (1,2,1) has a minimal effect on the *shape* of the control polygon.

7.2 Mathematical details: surfaces

The above subdivision methods can be easily extended from a control polygon to a quadrilateral mesh. This is a mesh where every polygon is a quadrilateral and every vertex is connected to four other vertices.

The Doo-Sabin subdivision method introduces four new vertices in each quadrilateral, and connects up vertices accordingly. The new vertices are blended mixtures of the old vertices in the proportions 9 : 3 : 3 : 1 (derived from the tensor product of the univariate case: $3 \times 3 : 3 \times 1 : 1 \times 3 : 1 \times 1$). This is illustrated in Figure 19.

Catmull-Clark subdivision is not much more difficult to understand. The only difference here is that is not all of the new vertices are created using the same weights. A vertex is introduced in the centre of each quadrilateral, in the centre of each edge, and near to each old vertex. Each of these three types of vertex has a different set of weights as illustrated in Figure 20.




Figure 20: Catmull-Clark subdivision. Above: a mesh (solid dots and solid lines) that has been refined (open dots and dashed lines). Below: the weights used to generated each type of refined vertex: centre, edge, and modified old vertex.

This all works beautifully for quadrilateral meshes. Now, suppose we have a quadrilateral mesh that contains extraordinary vertices, in other words a mesh that consists of quadrilaterals but has occasional vertices with other than four immediate neighbours. The Doo-Sabin scheme will still worked quite happily, because every polygon in the mesh is still quadrilateral. However, the Catmull-Clark subdivision scheme depends on every vertex having exactly four neighbours for the generation of the new vertex that is near to the old vertex position (the rightmost case in Figure 20). Catmull and Clark got around this problem by creating a new set of weights, one set of weights for each vertex valence (the valence of vertex is a number of other vertices to which it is connected). Instead of weights of $1/64$, $6/64$, and $36/64$ you can use weights of $1/4n^2$, $3/2n^2$, and $1 - 7/4n$, where n is the valence of the vertex. This particular set of weights was derived by Denis Zorin, other values can also be used.

However, this is not the only type of mesh with which we can deal. The Doo-Sabin scheme can be easily modified to cope with meshes in which some of the polygons are not quadrilateral, while still maintaining C^1 -continuity everywhere. For a k -sided polygon, the weights, α_k on the k vertices can be shown to be:

$$\alpha_0 = \frac{1}{4} + \frac{5}{4k} \quad (118)$$




Figure 21: Left: an example surface of revolution. Middle and right: two views of its generating polygon (the red quadrilateral), with the generated 3D object shown in semi-transparent cyan.

$$\alpha_i = \frac{1}{4k} \left(3 + 2 \cos \frac{2i\pi}{k} \right) \quad (119)$$

There are other schemes, notably the Loop scheme (named after Dr Loop) which works on triangular meshes. My research group at the Computer Laboratory has been working on the theory of subdivision since 2000 and has produced some interesting results including some rather whacky alternative subdivision schemes.

7.2.1 Exercises

1. Do the “constructive” exercises at the end of section 7.1.
2. Explain how Doo-Sabin subdivision works for an arbitrary polygon mesh.

8 Other 3D modelling mechanisms

8.1 Sweeps

These are three dimensional objects generated by *sweeping* a two dimensional shape along a path in 3D. Two special cases of the *sweep* are *surfaces of revolution*, where the path is a circle (see Figure 21); and *extrusions*, where the path is a straight line (see Figure 22). Some surfaces can be generated in more than one way, as illustrated in Figure 23. Surfaces of revolution are covered in **R&A** section 6-2. More general *sweeps* are covered in **R&A** section 6-3 and **FvDFH** section 12.4. Parts of Sections 6-2 and 6-3 of **R&A** are included in the handout.

If we push the idea of a sweep to its limit we can think of many things which could be modified to produce a three dimensional swept shape:

- **Cross section** This is the two dimensional shape that is to be swept along the sweep path. It does not have to be circular. Figure 24 shows two swept objects, one with a circular cross-section, one with a polygonal cross-section.




Figure 22: Left: an example extrusion. Right: its generating polygon (the red star), with the generated 3D object shown in semi-transparent cyan.




Figure 23: Left: an example object which could be an extrusion or a surface of revolution. Middle: generated as a surface of revolution from a rectangle. Right: generated as an extrusion from an annulus.




Figure 24: Two simple swept objects, one with a polygonal cross-section, one with a circular cross-section.




Figure 25: Two views of a polygonal cross-section swept along a path.




Figure 26: A polygonal cross-section swept along a path with (left) and without (right) twist.

- **Sweep path** This is the path along which the two dimensional cross section is swept to produce the three dimensional shape. It may be any curve. Figure 25 shows two views of the same swept object: a polygonal cross-section is swept along a convoluted path.
- **Twist** This is how the cross section twists (rotates) as it moves along the sweep path. The default would be to have no twist at all. Figure 26 shows a swept object with and without some twist.
- **Scale** This is how the cross section scales (changes size) as it moves along the sweep path. The default would be to have it stay the same size along the whole path. Figure 27 shows a cylinder, and the same cylinder with different scales along its length.
- **Normal vector direction** The normal vector of the 2D cross section will usually point along the sweep path at each point. Changing this will change the nature of the swept object. See **R&A** Figure 6-17 (in the handout) for an example.

You may be able to think of parameters, other than those in the list above, which could be modify. The examples given here were all generated using a generalised cylinder tool from Silicon Graphics would could vary cross-section, path, twist and scale. This is sufficient for most purposes. The input to such a system will either be NURBS curves or polylines. The output will be a NURBS surface or a polygon mesh.




Figure 27: A cylinder (top) swept along a path with scale changes (bottom).




Figure 28: A swept object with a circular cross-section, semi-circular path, and varying scale.




Figure 29: The object from the 1998 exam question.

8.1.1 Sweeping with NURBS

A NURBS surface is defined as the same two-dimensional extension to NURBS curves described in Equation 95, though obviously carried out in homogeneous co-ordinates. You can define sweeps using NURBS curves by using one NURBS curve as the sweep path, and another NURBS curve as the cross-section. You take the tensor product of the two curves. This is covered in **R&A**, pages 445–456, 465–477.

8.1.2 Converting swept objects to polygons

Swept objects are hard to ray trace. Imagine trying to write a ray/object intersection algorithm for a general swept object. This means that we generally need to polygonise swept objects in order to render them. For **polygon scan conversion** we obviously must convert them to polygons.

A swept surface may be easily converted to polygons by converting the outline of the 2D cross section to a polygon, and converting the sweep path to connected set of line segments. Moving the polygon to each vertex of the set of line segments, and connecting vertices accordingly, will produce a polygon mesh which approximates the swept surface.

8.1.3 Exercises

1. [1998/7/12] Show how the object in Figure 29 can be represented as a swept object. Show also how to convert the swept object into polygons. Explain what extra work would you need to do if you had to convert it into triangles.
2. Use the following different methods of specifying a geometrical model for the object in Figure 30 (assuming it is a three dimensional model and not a line drawing). Come as close as you can to the original for each of the methods, and describe the difficulties in using a particular method for this model.
 - (a) Extrusions
 - (b) Surfaces of revolution
 - (c) Sweeps along straight lines with scale changes.




Figure 30: A lamp.

3. For each of the following categories list five real-world objects which could be represented by the primitives in the category.
 - (a) The ray-tracing primitives in section 3.3
 - (b) Extrusions
 - (c) Surfaces of revolution
 - (d) General sweeps

4. A flume (water tunnel) at a swimming complex is modelled as a circle swept along a particular path. The designers also want to model the volume swept out by a person travelling down the flume. We can approximate the cross-section of a person with something roughly elliptical and we'll assume the 'virtual' person doesn't move legs or arms while hurtling along. Explain which parameters in the list would need to be modified to specify the shape of the flume and which would need to be modified to model the volume swept out by a person travelling down the flume (alternatively, specify which parameters would be held constant, in each case, for the entire length of the sweep).

8.2 Constructive Solid Geometry

Constructive solid geometry (CSG) essentially consists of Boolean set operations on closed primitives in 3D space. CSG is covered in **FvDFH** sections 12.7 and 15.10.3.

The three CSG operations are *union*, *intersection* and *difference*. Figure 31 illustrates the three CSG operations in use on simple three dimensional primitives. Figure 32, based on **FvDFH** Plate III.2, shows an object for which CSG is (probably) the only sensible modelling technique. The object rendered in the right-hand image is constructed from the primitives shown in the left-hand image. It is mostly made out of cylinders, but you will recognise the extruded star from Figure 22. The cover of Hofstadter's "Gödel, Escher, Bach"⁴⁹ has a carved wood shape which casts shadows of the letters G, E, and B onto three orthogonal planes. In Figure 33, I try this treatment on my initials. Unfortunately the letters N, A, and D are not as amenable to this as the letters G, E, and B: notice that the shadow of the N has a slight curve at its top right, owing to the N's

⁴⁹<http://www.amazon.com/exec/obidos/ASIN/0465026567/rpcman>




Figure 31: CSG example: two primitives (a sphere and a box) under the three CSG operations: union, intersection, and difference.




Figure 32: CSG example. Left: the primitives. Right: the finished object. It is difficult to see how this object could be produced easily in any other way.




Figure 33: A CSG shape casting three shadows of three different letters.

intersection with the curve on the D and the slope on the A. Various other arrangements of the three letters were tried, all of which gave more noticeable artefacts than this. Each of the letters is a CSG object (the D, for example, is constructed from cylinders and boxes). The final effect is produced simply by intersecting the three letters.

More information is required to fully understand how CSG can be raytraced. Some of this is given in **SSC** Chapter 18, some below, and more will be said in lectures.

8.2.1 Some algorithmic details

This section deals with how to raytrace a CSG object. A CSG object can be represented as a binary tree where each node is one of the three CSG operators: union (\cup), intersection (\cap), and difference (\setminus). The leaves of the tree are primitive objects for which a ray-object intersection test exists. The basic information which is passed up the tree is a list of *all* intersection point, t , between the ray and the object represented by the sub-tree. Combining two sub-trees requires you to sort through these lists of intersection points, keeping those from both sub-trees which are intersection points with the combined object and discarding the others.

CSG of two objects, A and B , can be driven by a simple state machine, where there are four states:

- In both A and B .
- In just A .

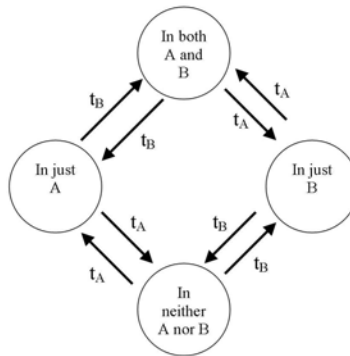


Figure 34: The simple state machine used in the CSG algorithm.

- In just B .
- In neither A nor B .

See Figure 34. As one progresses along the ray, each intersection point changes the state. Each of the four possible operations (\cup , \cap , $A \setminus B$, $B \setminus A$) requires you to keep the intersection points which transfer you into or out of just one of the four states. Thus $A \setminus B$ is defined by the transitions into and out of the “In just A ? state, while intersection is defined by the transitions into and out of the “In both A and B ? state.

A ray’s interaction with an object, A , can be stored as an ordered pair: (`inA` : Boolean, `tAs` : float list) where `inA` is `True` if the eye is inside object A and `False` otherwise. `tAs` is a list of t -values for which the ray intersects the surface of the object, sorted in increasing order. Of course, all t values are non-negative.

Figure 35 gives an ML function for finding the Constructive Solid Geometry intersection of two objects: $A \cap B$. You could try writing an equivalent function to find the difference of two objects, $A \setminus B$.

8.2.2 Exercises

1. Work out how to represent a Lego technic brick as a CSG Object (see Figure 36). You may assume that you have box and cylinder primitives.
2. [1998/7/12] Work out how to represent the object in Figure 29 using CSG. You may assume the following primitives: sphere, cylinder, cone, torus, box.
3. [1999/9/4] (c) List the three ways of combining objects using constructive solid geometry (CSG). Describe how an object built using CSG can be represented using a binary tree. Given the intersection points of a ray with each primitive in the tree, explain how these points are passed up the tree by each type of combination node to produce a list of intersection points for the whole CSG object. [8 marks]
4. [2002/8/4] (c) Describe how an object built using constructive solid geometry (CSG) can be represented using a binary tree. Given the intersection points of a ray with each primitive in the tree, show how to calculate the first intersection point of the ray with the entire CSG object. [6 marks]

```

fun inter( (inA, tA::tAs), (inB, tB::tBs) )
  =   if( tA < tB ) then
      if( inB ) then
        tA :: inter( (not inA, tAs), (inB, tB::tBs) )
      else
        inter( (not inA, tAs), (inB, tB::tBs) )
    else
      if( inA ) then
        tB :: inter( (inA, tA::tAs), (not inB, tBs) )
      else
        inter( (inA, tA::tAs), (not inB, tBs) )
|   inter( (inA, []), (inB, tB::tBs) )
  =   if( inA ) then
      tB :: tBs
    else
      []
|   inter( (inA, tA::tAs), (inB, []) )
  =   if( inB ) then
      tA :: tAs
    else
      []
|   inter( (inA, []), (inB, []) )
  =   [] ;

fun intersection( (inA, tAs), (inB, tBs) )
  =   ( inA andalso inB, inter( (inA, tAs), (inB, tBs) ) ) ;

```

Figure 35: An ML function for calculating the CSG intersection of two objects




Figure 36: Some CSG Lego technic bricks.

5. Write a function, similar to that in section 8.2.1 for $A \cap B$, to find the difference of two objects, $A \setminus B$.

8.3 Implicit surfaces, voxels and the marching cubes algorithm

8.3.1 Implicit surfaces

These will be described in lectures. Brian Wyvill has a good introduction to implicit surfaces⁵⁰ on his website at the University of Calgary⁵¹. Click on “Implicit Tutorial” for the introduction to implicit surfaces.

8.3.2 Voxels and the marching cubes algorithm

Voxels are the three dimensional analogue of pixels. Rather than storing a colour in a voxel, you will generally store a density value. Voxels and the marching cubes algorithm are both covered in Lorenson and Cline’s 1987 SIGGRAPH paper “Marching cubes: a high resolution 3D surface construction algorithm”, *Proc SIGGRAPH 87*, pages 163-169. This paper is included in the handout. One of the reasons for including the paper is to give you a taster of what a good graphics research paper looks like.

8.3.3 Exercises

1. Give a definition of an implicit surface and give three examples of where such things might be useful.
2. Explain how voxel data can be thought of as defining an implicit surface (or surfaces). Explain, conversely, an implicit surface can be converted into voxel data.
3. [2001/7/9] (a) The marching squares algorithm is a two-dimensional version of marching cubes, where you generate line segments in 2D rather than triangles in 3D. It could be used, for example, where you have a regular grid of height values and want to draw contours of constant height. Sketch an implementation of this two-dimensional marching squares algorithm. [6 marks]
4. Medical data is captured in slices. Each slice is a 2D image of density data. The distance between slices may be different to the distance between the pixels within a slice (for example, see Lorenson and Cline, Section 7.1, p. 167). What effect, if any, does this difference have on the voxel data? What effect, if any, does it have on the marching cubes algorithm?
5. [2002/8/4] (b) Implicit surfaces are normally combined by adding the field functions together to create a “blobby” blended surface. Describe an alternative mechanism (or mechanisms) for combining implicit surfaces which would produce results more akin to CSG union and intersection. Explain why it produces these results. Given this mechanism, suggest a way of combining implicit surfaces to produce a result similar to CSG difference. [4 marks]

⁵⁰http://pages.cpsc.ucalgary.ca/~blob/implicitall_files/frame.htm

⁵¹<http://pages.cpsc.ucalgary.ca/~blob/currentwork.html>


Slide 1

An introduction to the surface representations used in Computer Aided Design

Neil A. Dodgson
University of Cambridge
Computer Laboratory

with considerable help
from Sabin, Barthe,
Ivrissimtzis, Hassan,
Gérot, Kobbelt, Albat
& Müller



Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk




Slide 2

Applications

- ✦ Computer-aided design (CAD) of artistically attractive surfaces
- ✦ Principally the design of car bodies



Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk



Slide 3

Applications

- ✦ The design of industrial artefacts and animated characters



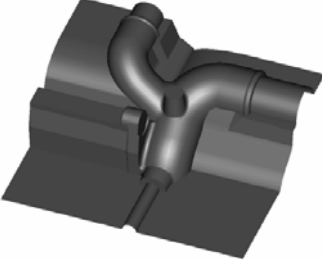

Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk




Slide 4

Applications

- ✦ Computer-aided design of industrial artefacts which do not need to be beautiful
 - ◆ Such as internal engine parts




Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk



Slide 5

Desirable features in CAD

- ✦ Need to handle *any* surface
- ✦ Need guaranteed continuity
 - ◆ Continuity of slope (C1)
 - Smooth surfaces
 - ◆ Continuity of curvature (C2)
 - Smoothly reflecting surfaces
 - Required for some aerodynamics
- ✦ Need to allow for discontinuities
 - ◆ Edges, creases and holes
- ✦ Needs to be easy to use

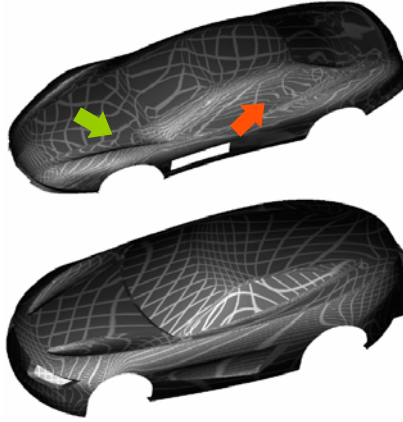


Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

Slide 6

Reflection patterns on car models

- ✦ The top car has problems
 - ◆ Some discontinuities in curvature ●
 - ◆ Some badly distorted areas owing to rapidly changing curvature ●
- ✦ Only the bottom car is a “Class A” surface
 - ◆ The reflection lines look nice



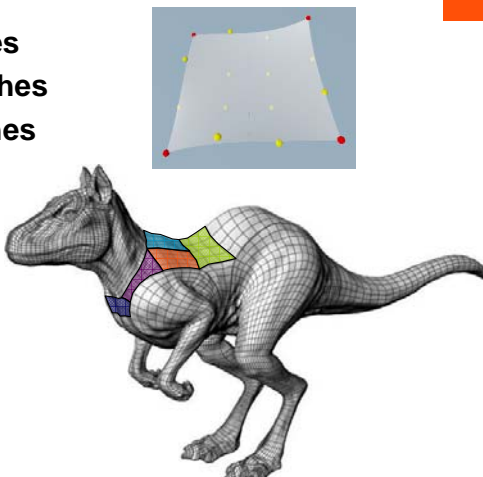
Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF
CAMBRIDGE

Slide 7

Traditional tools

- ✦ Bezier patches
- ✦ B-spline patches
- ✦ NURBS patches

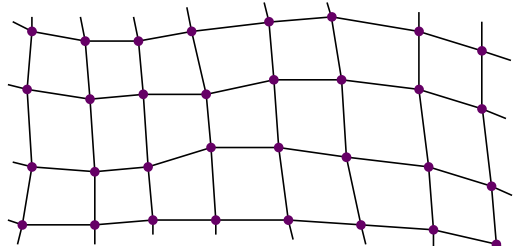


Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

Slide 8

Bezier & B-spline patches

- ✦ A rectangular array of control points
- ✦ A mathematical function determines where the surface goes based on those points
- ✦ Move a control point to change the surface



Part II
Advanced
Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF
CAMBRIDGE

Slide 9

The first problem

✦ Very few objects are made up of a single rectangular patch, so we need to join patches together

Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 10

The mathematics of joins

✦ We want to preserve certain types of mathematical continuity across joins

- ◆ C0: continuity of position
 - Prevents holes at the join
- ◆ C1: continuity of slope
 - Prevents a sharp edge at the join
- ◆ C2: continuity of curvature
 - Strongly related to aesthetics
 - Most often visible in reflections
 - Prevents sharp edges in reflected lines
- ◆ These are continuity of the zeroth, first and second derivatives

Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 11

Joining two Bezier patches

✦ C0 but not C1

- ◆ Four edge points are the same

✦ C0 and C1

- ◆ Four edge points are the same
- ◆ Next four points out in either direction are constrained

Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 12

Joining four Bezier patches

Constraints

- C0 ●
- C1 ● ●
- C2 ● ● ●

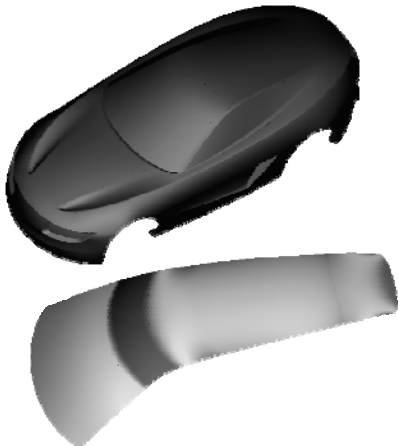
Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 13

An example: the car's roof

- ★ The car
- ★ Curvature plot of its roof



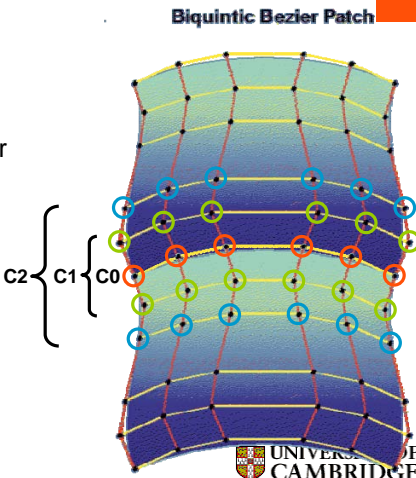
Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 14

Definition of the car's roof

- ★ 5x2 grid of biquintic Bezier patches
 - ◆ 36 control points per patch
 - ◆ 286 control points overall
 - ◆ Moving one point also moves several others to maintain C2 continuity



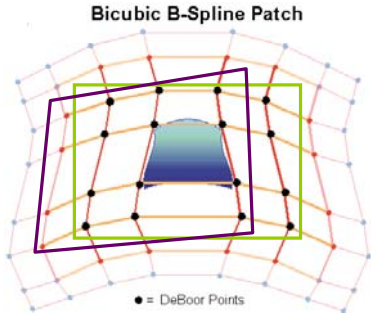
Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 15

B-spline patches

- ★ A rectangular array of points define a rectangular array of *automatically joined* patches
- ★ Example
 - ◆ The black points control the blue patch
 - ◆ All points together define a surface of many joined patches



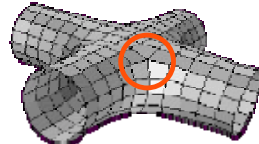
Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 16

The second problem

- ★ What do we do at special points where other than four patches meet?
 - ◆ Either we cannot get C2
 - Which means that curvature is not continuous
 - ◆ Or we get C2 by forcing curvature to be zero
 - Which produces a flat spot
 - ◆ Or we get C2 using very high degree patches
 - Which are very hard for a designer to control



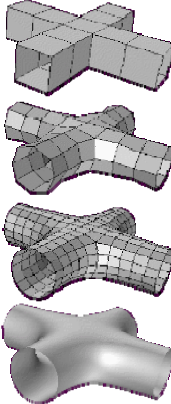
Part II
Advanced Graphics
2004
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 17

Subdivision surfaces

- ◆ Developed in the 1970s, adopted in computer animation in 1990s
- ◆ Replace the patch-based representation of B-splines and Beziers
- ◆ Base a curve or surface solely on its control points and their connectivity
- ◆ A simple mechanism produces a larger, more refined set of control points from the current set
- ◆ Iterate refinement until the appropriate level of detail is achieved



Part II
Advanced Graphics
2004

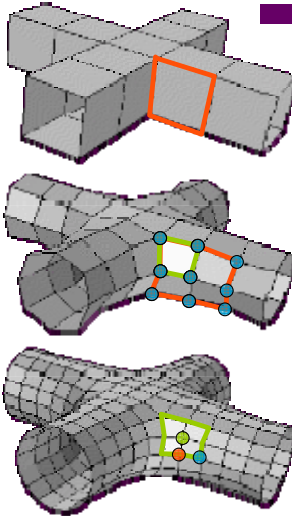
Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 18

Subdivision

- ✦ Introduce new points
 - ◆ At face-centres
 - ◆ At mid-edges
- ✦ Adjust positions of original points
- ✦ Repeat until sufficiently detailed



Part II
Advanced Graphics
2004

Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Slide 19

Subdivision

- ✦ Advantages
 - ◆ Reproduces everything which can be done by B-splines
 - ◆ Handles extraordinary points much more easily
- ✦ Disadvantages
 - ◆ Cannot get C2 unless you produce a flat spot
 - ◆ Generates other visual artefacts, not seen in B-spline surfaces
- ✦ Commercial position
 - ◆ Subdivision is replacing B-splines in computer animation
 - ◆ Subdivision is **not** replacing B-splines in CAD

Part II
Advanced Graphics
2004

Neil Dodgson nad@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Advanced Graphics 2006

- Subdivision curves & surfaces

Beware: some slides contain multi-layer animations, which do not print well.

© 2002,2003,2004,2006 Neil Dodgson

1

Modelling smooth 3D surfaces

- Where are smooth 3D surfaces used?
 - Computer Aided Design (CAD)
 - First developed for cars & aeroplanes
 - Adopted for other manufactured objects
 - Computer animation
- What mechanisms exist?
 - Bézier patches
 - NURBS surfaces
 - Subdivision surfaces

2

Desirable features

- Need to handle *any* surface
- Need guaranteed continuity
 - C1-continuity
 - Smooth surfaces
 - C2-continuity
 - Smoothly reflecting surfaces
 - Required for some aerodynamics
- Need to allow discontinuities
 - Edges, creases and holes
- Needs to be easy to use



History of 3D modelling 1/3

- Some mechanism was needed for modelling 3D surfaces
- Hermite interpolation was generalised to bivariate patches
 - ...but proved too difficult to use in practice
- Bézier patches
 - Developed for car design around 1960
 - Bézier (Renault), de Casteljau (Citroën), de Boor (GM)

4

History of 3D modelling 2/3

- B-spline theory
 - Developed in the 1960s and '70s, led to:
- NURBS (Non-Uniform Rational B-Splines)
 - More general than Bézier patches
 - Béziers are special cases of NURBS
 - NURBS quickly became the industry standard in CAD
 - ...and remain the industry standard today
 - Adopted by the computer animation industry when it began

5

History of 3D modelling 3/3


- Subdivision surfaces
 - Theory developed in 1970s and early '80s
 - Picked up by computer animation industry in late 1990s
 - Now replaced NURBS in computer animation
 - Solves one of the big problems of NURBS
 - Still under active research for use in CAD
 - Introduces new problems, not present in NURBS, which make it unsuitable for CAD in its present form

6

NURBS curve

- A curve is defined parametrically
- Its shape is determined by:
 - control points, P_i
 - and the NURBS basis functions, $N_{i,k}$

$$P(t) = \sum_{i=1}^{n+1} N_{i,k}(t)P_i$$



7

Basic properties of NURBS 1/3

$$P(t) = \sum_{i=1}^{n+1} N_{i,k}(t)P_i$$

- The basis functions must sum to 1 to produce a valid new point

$$\sum_{i=1}^{n+1} N_{i,k}(t) = 1, t_{\min} \leq t \leq t_{\max}$$

8

Basic properties of NURBS 2/3

$$P(t) = \sum_{i=1}^{n+1} N_{i,k}(t)P_i$$

- The basis functions are calculated from a *knot vector*
 - Just a non-decreasing sequence of real numbers
 - e.g. [0,0,0,1,1,1] or [1,2,3,4,5,6] or [1.2, 3.4, 5.6, 5.6, 7.2, 15.6]
 - See lecture notes or Rogers & Adams for details

9

Basic properties of NURBS 3/3


$$P(t) = \sum_{i=1}^{n+1} N_{i,k}(t)P_i$$

- If the basis functions are C_m -continuous at t , then $P(t)$ is guaranteed to be C_m -continuous at t
 - So continuity depends only on the basis functions, $N_{i,k}$
 - Continuity does *not* depend on the locations of the control points
 - you can sometimes get extra continuity by careful positioning of control points

10

NURBS surface

- A bivariate generalisation of the univariate NURBS curve




Curve $P(t) = \sum_{i=1}^{n+1} N_{i,k}(t)P_i$

Surface $P(s,t) = \sum_{i=1}^{m+1} \sum_{j=1}^{n+1} N_{i,k}(s)N_{j,k}(t)P_{i,j}$

11

The big constraint...


- NURBS surfaces require a quadrilateral mesh of $(m+1) \times (n+1)$ points



12

The first problem


- Very few objects are made up of a single rectangular patch, so we need to join patches together



13

The second problem

- What do we do at special points where other than four patches meet?




- Either we cannot get C2
 - Which means that curvature is not continuous
- Or we get C2 by forcing curvature to be zero
 - Which produces a flat spot
- Or we get C2 using very high degree patches
 - Which are very hard for a designer to control

14

Drawing a NURBS curve

- NURBS curves and surfaces are always drawn on a pixelated surface
- NURBS curves can be approximated by straight lines
 - So long as each straight line deviates from the curve by less than half a pixel



15

Drawing a NURBS surface


- NURBS surfaces are subdivided and drawn as a series of planar polygons
- Each polygon is only one or two pixels in area on the screen
- Shading algorithms are used to ensure that the surfaces appear to be smoothly curved



16

Subdivision surfaces


- Do away with the explicit parametric representation
- Base a curve or surface solely on its control points and their connectivity
- Provide a simple mechanism which produces a larger, more refined set of control points from the current set
- Iterate refinement until the appropriate level of detail is achieved



17

History of subdivision schemes

- A univariate (curve) scheme was described by de Rahm in 1947
- Rediscovered by Chaikin in 1974
- Extended to bivariate (surfaces)
 - Doo-Sabin bi-quadratic patches (1978)
 - Catmull-Clark bi-cubic patches (1978)
- Flurry of mathematical work in the early 1980s
 - Dyn & Levin at Tel Aviv University



18

Use of subdivision schemes

- Pixar picked up the ideas and tested them in Geri's Game (1997)
- ...then discarded its NURBS based software in favour of subdivision schemes




- NURBS
 - Toy Story 1995
 - A Bug's Life 1998
- Subdivision surfaces
 - Toy Story II 1999
 - Monsters Inc. 2001
 - Finding Nemo 2003

19

Subdivision basics

- An example: Catmull-Clark subdivision
 - Introduce new points
 - At face-centres
 - At mid-edges
 - Adjust positions of original points
 - Repeat until sufficiently detailed



20

Chaikin curve subdivision

- Underlies Doo-Sabin surface subdivision
- C1-continuous in the limit
- Essentially just a $\frac{1}{4}$ - $\frac{3}{4}$ rule

21

The maths of Chaikin

$$P_{2i}^{n+1} = \frac{3}{4}P_i^n + \frac{1}{4}P_{i+1}^n$$

$$P_{2i+1}^{n+1} = \frac{1}{4}P_i^n + \frac{3}{4}P_{i+1}^n$$

$$h = [K, 0, 0, \frac{1}{4}, \frac{3}{4}, \frac{3}{4}, \frac{1}{4}, 0, 0, K]$$

$$P^n = [K, P_0^n, P_1^n, P_2^n, K]$$

$$\overleftrightarrow{P}^n = [K, P_0^n, 0, P_1^n, 0, P_2^n, 0, K]$$

$$P^{n+1} = h * \overleftrightarrow{P}^n$$

22

The limit curve

- It can be shown that the limit curve of the Chaikin scheme is the uniform quadratic B-spline, which is guaranteed to be C1
- When drawing curves in computer graphics, we draw a set of straight lines, so only need to subdivide until each segment is about a pixel long and we have a good enough approximation to the curve

23

C2 approximating scheme

- Underlies Catmull-Clark surface subdivision
- Can be described as: "Insert a midpoint and adjust the old control points"

24

The maths of the C2 scheme

$$P_{2i}^{n+1} = \frac{1}{8}P_{i-1}^n + \frac{6}{8}P_i^n + \frac{1}{8}P_{i+1}^n$$

$$P_{2i+1}^{n+1} = \frac{4}{8}P_i^n + \frac{4}{8}P_{i+1}^n$$

$$h = [K, 0, 0, \frac{1}{8}, \frac{4}{8}, \frac{6}{8}, \frac{4}{8}, \frac{1}{8}, 0, 0, K]$$

$$P^n = [K, P_0^n, P_1^n, P_2^n, K]$$

$$\vec{P}^n = [K, P_0^n, 0, P_1^n, 0, P_2^n, 0, K]$$

$$P^{n+1} = h * \vec{P}^n$$

25

Why this notation?

- Easy to analyse
- Allows use of the z-transform

$$h = [K, h_0, h_1, h_2, K]$$

↓

$$h(z) = \Lambda + h_0z^0 + h_1z^1 + h_2z^2 + \Lambda$$

vector
↓
polynomial

$$P^{n+1} = h * \vec{P}^n$$

↓

$$P^{n+1}(z) = h(z) \times P^n(z^2)$$

convolution
↓
multiplication

26

The analysis tools

- Generating function formalism
 - Use the z-transform on the kernel, h
 - Provides sufficient conditions for continuity
 - Essentially checks that the differences between adjacent points decrease fast enough at each refinement step to produce a smooth curve
- There is also a matrix formalism
 - Analyse stationary points
 - Provides necessary conditions for continuity
- For details see our research papers ☺

27

Useful subdivision kernels

- $h = \frac{1}{4}[1, 3, 3, 1]$
 - C1, approximating, limit curve is quadratic B-spline
- $h = \frac{1}{8}[1, 4, 6, 4, 1]$
 - C2, approximating, limit curve is cubic B-spline
- $h = \frac{1}{16}[-1, 0, 9, 16, 9, 0, -1]$
 - C1, interpolating, "four-point scheme"
 - There is also a C2 interpolating six-point scheme

28

From Chaikin to Doo-Sabin

- Doo-Sabin scheme is bivariate generalisation of Chaikin $1/4$ - $3/4$ scheme

29

Extraordinary polygons

- Need special co-efficients for these

$$\alpha_0 = \frac{1}{4} + \frac{5}{4K}$$

$$\alpha_i = \frac{1}{4K} (3 + 2 \cos \frac{2i\pi}{K})$$

(Doo-Sabin)

30

Catmull-Clark subdivision

- Catmull-Clark is based on the $1/8[1,4,6,4,1]$ univariate scheme

31

Catmull-Clark rules


face **edge** **vertex**

- This is easy: the rules are simply the tensor product of the univariate $1/8[1,4,6,4,1]$ rules.

32

Catmull-Clark special cases

- This is more difficult: we need to find co-efficients which maintain continuity
 - It is only possible to get C1 continuity at these extraordinary points.




Extraordinary polygons:
disappear after one step

$$\alpha = \frac{1}{4n^2}$$

$$\beta = \frac{3}{2n^2}$$


$$\gamma = 1 - \frac{7}{4n}$$



Extraordinary vertices:
remain in the mesh 33

Subdivision vs NURBS

- Extraordinary points
 - Subdivision handles them easily
 - NURBS requires the use of other types of surface to fill in the holes
- Memory requirements
 - Subdivision needs a lot (many MB)
 - NURBS is very compact
- Artifacts
 - Some artifacts present in both
 - Subdivision has extra artifacts



34

The future

- Computers now have enough memory to handle subdivision easily
- Subdivision now standard for computer animation
- NURBS still standard for CAD
- Subdivision will eventually replaced NURBS for CAD if we can sort out the artifact problems

35


Our work at Cambridge

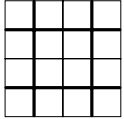
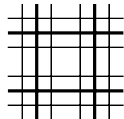
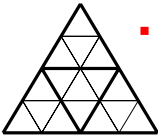
- Univariate schemes that are not binary
 - Ternary ($\times 3$) schemes
 - Sesquinary ($\times 1\frac{1}{2}$) schemes
- Towards a bestiary of bivariate schemes
 - Classification & analysis of all schemes
 - Identification & analysis of new schemes (especially ternary)
- Geometrically-sensitive subdivision
 - Modifying existing schemes to take account of geometric relationships



36

Principal subdivision schemes




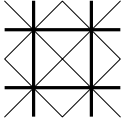
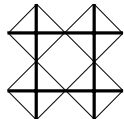
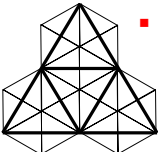
- 
 - Catmull-Clark
 - C2, approximating
 - only C1 at extraordinary points
 - Kobbelt (four-point)
 - C1, interpolating
- 
 - Doo-Sabin
 - C1, approximating
- 
 - Loop
 - C2, approximating
 - only C1 at extraordinary points
 - Butterfly
 - C1, interpolating

These are the five subdivision schemes which were thought to be the only useful ones.

37

$\sqrt{2}$ and $\sqrt{3}$ schemes



- 
 - Velho-Zorin
 - C4, approximating
 - only C1 at extraordinary points
- 
 - Reif-Peters
 - C1, approximating
 - The simplest possible scheme: there are no special cases!
- 
 - $\sqrt{3}$ (Kobbelt)
 - C2, approximating
 - only C1 at extraordinary points

These were discovered in the late 1990s

38

MARCHING CUBES: A HIGH RESOLUTION 3D SURFACE CONSTRUCTION ALGORITHM

William E. Lorensen
Harvey E. Cline

General Electric Company
Corporate Research and Development
Schenectady, New York 12301

Abstract

We present a new algorithm, called *marching cubes*, that creates triangle models of constant density surfaces from 3D medical data. Using a divide-and-conquer approach to generate inter-slice connectivity, we create a case table that defines triangle topology. The algorithm processes the 3D medical data in scan-line order and calculates triangle vertices using linear interpolation. We find the gradient of the original data, normalize it, and use it as a basis for shading the models. The detail in images produced from the generated surface models is the result of maintaining the inter-slice connectivity, surface data, and gradient information present in the original 3D data. Results from computed tomography (CT), magnetic resonance (MR), and single-photon emission computed tomography (SPECT) illustrate the quality and functionality of *marching cubes*. We also discuss improvements that decrease processing time and add solid modeling capabilities.

CR Categories: 3.3, 3.5

Additional Keywords: computer graphics, medical imaging, surface reconstruction

1. INTRODUCTION.

Three-dimensional surfaces of the anatomy offer a valuable medical tool. Images of these surfaces, constructed from multiple 2D slices of computed tomography (CT), magnetic resonance (MR), and single-photon emission computed tomography (SPECT), help physicians to understand the complex anatomy present in the slices. Interpretation of 2D medical images requires special training, and although radiologists have these skills, they must often communicate their interpretations to the referring physicians, who sometimes have difficulty visualizing the 3D anatomy.

Researchers have reported the application of 3D medical images in a variety of areas. The visualization of complex

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

acetabular fractures [6], craniofacial abnormalities [17,18], and intracranial structure [13] illustrate 3D's potential for the study of complex bone structures. Applications in radiation therapy [27,11] and surgical planning [4,5,31] show interactive 3D techniques combined with 3D surface images. Cardiac applications include artery visualization [2,16] and non-graphic modeling applications to calculate surface area and volume [21].

Existing 3D algorithms lack detail and sometimes introduce artifacts. We present a new, high-resolution 3D surface construction algorithm that produces models with unprecedented detail. This new algorithm, called *marching cubes*, creates a polygonal representation of constant density surfaces from a 3D array of data. The resulting model can be displayed with conventional graphics-rendering algorithms implemented in software or hardware.

After describing the information flow for 3D medical applications, we describe related work and discuss the drawbacks of that work. Then we describe the algorithm as well as efficiency and functional enhancements, followed by case studies using three different medical imaging techniques to illustrate the new algorithm's capabilities.

2. INFORMATION FLOW FOR 3D MEDICAL ALGORITHMS.

Medical applications of 3D consist of four steps (Figure 1). Although one can combine the last three steps into one algorithm, we logically decompose the process as follows:

1. Data acquisition.

This first step, performed by the medical imaging hardware, samples some property in a patient and produces multiple 2D slices of information. The data sampled depends on the data acquisition technique.




Figure 1. 3D Medical Information Flow.

X-ray computed tomography (CT) measures the spatially varying X-ray attenuation coefficient [3]. CT images show internal structure. For 3D applications, CT is frequently used to look at bone structure, although we have had success visualizing soft tissue.

Magnetic resonance (MR) measures three physical properties [20]. One property is the distribution of "mobile" hydrogen nuclei and shows overall structure within the slices. The other two properties measure relaxation times of the nuclei. MR, a recent technique, shows excellent contrast between a variety of soft tissues. However, the variety of surfaces presents a challenge to 3D surface construction and requires techniques for selective surface extraction and display.

A third acquisition technique, single-photon emission computed tomography (SPECT) measures the emission of gamma rays [24]. The source of these rays is a radioisotope distributed within the body. In addition to structure, SPECT can show the presence of blood in structures with a much lower dose than that required by CT.

2. Image processing.

Some algorithms use image processing techniques to find structures within the 3D data [1,32,30,29] or to filter the original data. MR data, in particular, needs image processing to select appropriate structure.

3. Surface construction.

Surface construction, the topic of this paper, involves the creation of a surface model from the 3D data. The model usually consists of 3D volume elements (voxels) or polygons. Users select the desired surface by specifying a density value. This step can also include the creation of cut or capped surfaces.

4. Display.

Having created the surface, the final step displays that surface using display techniques that include ray casting, depth shading, and color shading.

3. RELATED WORK.

There are several approaches to the 3D surface generation problem. An early technique [23] starts with contours of the surface to be constructed and connects contours on consecutive slices with triangles. Unfortunately, if more than one contour of surface exists on a slice, ambiguities arise when determining which contours to connect [14]. Interactive intervention by the user can overcome some of these ambiguities [8]; however, in a clinical environment, user interaction should be kept to a minimum.

Another approach, developed by G. Herman and colleagues [19] creates surfaces from cuberilles. A cuberille is "dissection of space into equal cubes (called voxels) by three orthogonal sets of parallel planes [7]." Although there are many ways to display a cuberille model, the most realistic images result when the gradient, calculated from cuberilles in a neighborhood, is used to find the shade of a point on the model [15]. Meagher [25] uses an octree representation to compress the storage of the 3D data, allowing rapid manipulation and display of voxels.

Farrell [12] uses ray casting to find the 3D surface, but rather than shade the image with a gray scale, uses hue lightness to display the surface. In another ray casting method, Hohne [22], after locating the surface along a ray, calculates the gradient along the surface and uses this gradient, scaled

by an "appropriate" value, to generate gray scales for the image.

A different approach, used at the Mayo Clinic [26], displays the density volume rather than the surface. This method produces, in effect, a conventional shadow graph that can be viewed from arbitrary angles. Motion enhances the three-dimensional effect obtained using the volume model.

Each of these techniques for surface construction and display suffer shortcomings because they throw away useful information in the original data. The connected contour algorithms throw away the inter-slice connectivity that exists in the original data. The cuberille approach, using thresholding to represent the surface as blocks in 3D space, attempts to recover shading information from the blocks. The ray casting methods either use depth shading alone, or try to approximate shading with an unnormalized gradient. Since they display all values and not just those visible from a given point of view, volume models rely on motion to produce a three-dimensional sensation.

Our approach uses information from the original 3D data to derive inter-slice connectivity, surface location, and surface gradient. The resulting triangle model can be displayed on conventional graphics display systems using standard rendering algorithms.

4. MARCHING CUBES ALGORITHM.

There are two primary steps in our approach to the surface construction problem. First, we locate the surface corresponding to a user-specified value and create triangles. Then, to ensure a quality image of the surface, we calculate the normals to the surface at each vertex of each triangle.

Marching cubes uses a divide-and-conquer approach to locate the surface in a logical *cube* created from eight pixels; four each from two adjacent slices (Figure 2).

The algorithm determines how the surface intersects this cube, then moves (or *marches*) to the next cube. To find the surface intersection in a cube, we assign a one to a cube's vertex if the data value at that vertex exceeds (or equals) the value of the surface we are constructing. These vertices are inside (or on) the surface. Cube vertices with values below the surface receive a zero and are outside the surface. The surface intersects those cube edges where one vertex is outside the surface (one) and the other is inside the surface (zero). With this assumption, we determine the topology of the surface within a cube, finding the location of the intersection later.




Figure 2. Marching Cube.

Since there are eight vertices in each cube and two states, inside and outside, there are only $2^8 = 256$ ways a surface can intersect the cube. By enumerating these 256 cases, we create a table to look up surface-edge intersections, given the labeling of a cube's vertices. The table contains the edges intersected for each case.

Triangulating the 256 cases is possible but tedious and error-prone. Two different symmetries of the cube reduce the problem from 256 cases to 14 patterns. First, the topology of the triangulated surface is unchanged if the relationship of the surface values to the cubes is reversed. Complementary cases, where vertices greater than the surface value are interchanged with those less than the value, are equivalent. Thus, only cases with zero to four vertices greater than the surface value need be considered, reducing the number of cases to 128. Using the second symmetry property, rotational symmetry, we reduced the problem to 14 patterns by inspection. Figure 3 shows the triangulation for the 14 patterns.

The simplest pattern, 0, occurs if all vertex values are above (or below) the selected value and produces no triangles. The next pattern, 1, occurs if the surface separates on vertex from the other seven, resulting in one triangle defined by the three edge intersections. Other patterns produce multiple triangles. Permutation of these 14 basic patterns using complementary and rotational symmetry produces the 256 cases.

We create an index for each case, based on the state of the vertex. Using the vertex numbering in Figure 4, the eight bit index contains one bit for each vertex.

This index serves as a pointer into an edge table that gives all edge intersections for a given cube configuration.

Using the index to tell which edge the surface intersects, we can interpolate the surface intersection along the edge. We use linear interpolation, but have experimented with higher degree interpolations. Since the algorithm produces at least one and as many as four triangles per cube, the higher degree surfaces show little improvement over linear interpolation.

The final step in *marching cubes* calculates a unit normal for each triangle vertex. The rendering algorithms use this normal to produce Gouraud-shaded images. A surface of constant density has a zero gradient component along the surface tangential direction; consequently, the direction of the gradient vector, \vec{g} , is normal to the surface. We can use this fact to determine surface normal vector, \vec{n} , if the magnitude of the gradient, $|\vec{g}|$, is nonzero. Fortunately, at the surface of interest between two tissue types of different densities, the gradient vector is nonzero. The gradient vector, \vec{g} , is the derivative of the density function

$$\vec{g}(x, y, z) = \nabla f(x, y, z). \quad (1)$$

To estimate the gradient vector at the surface of interest, we first estimate the gradient vectors at the cube vertices and linearly interpolate the gradient at the point of intersection. The gradient at cube vertex (i, j, k) , is estimated using central differences along the three coordinate axes by:

$$G_x(i, j, k) = \frac{D(i+1, j, k) - D(i-1, j, k)}{\Delta x} \quad (2)$$

$$G_y(i, j, k) = \frac{D(i, j+1, k) - D(i, j-1, k)}{\Delta y} \quad (3)$$

$$G_z(i, j, k) = \frac{D(i, j, k+1) - D(i, j, k-1)}{\Delta z} \quad (4)$$




Figure 3. Triangulated Cubes.




Figure 4. Cube Numbering.

where $D(i, j, k)$ is the density at pixel (i, j) in slice k and $\Delta x, \Delta y, \Delta z$ are the lengths of the cube edges. Dividing the gradient by its length produces the unit normal at the vertex required for rendering. We linearly interpolate this normal to the point of intersection. Note that to calculate the gradient at all vertices of the cube, we keep four slices in memory at once.

In summary, *marching cubes* creates a surface from a three-dimensional set of data as follows:

1. Read four slices into memory.
2. Scan two slices and create a cube from four neighbors on one slice and four neighbors on the next slice.
3. Calculate an index for the cube by comparing the eight density values at the cube vertices with the surface constant.
4. Using the index, look up the list of edges from a precalculated table.
5. Using the densities at each edge vertex, find the surface-edge intersection via linear interpolation.
6. Calculate a unit normal at each cube vertex using central differences. Interpolate the normal to each triangle vertex.
7. Output the triangle vertices and vertex normals.

5. ENHANCEMENTS TO THE BASIC ALGORITHM.

We have made several improvements to the original *marching cubes* that make the algorithm run faster and that add solid modeling capabilities.

5.1 Efficiency Enhancements.

The efficiency enhancements allow the algorithm to take advantage of pixel-to-pixel, line-to-line, and slice-to-slice coherence. For cubes interior to the original data limits (those not including slice 0, line 0, or pixel 0), only three new edges need to be interpolated for each cube. We can obtain the other nine edges from previous slices, lines, or pixels. In Figure 5, the shaded circles represent values available from prior calculations; only edges 6, 7, and 12 have to be calculated for the new cube.

Special cases are present along the boundaries of the data, but, by enumerating these cases, we can limit vertex calculations to once per vertex. In practice, we only save the previous pixel and line intersections because the memory required to save the previous slice's intersections is large. Using the coherence speeds up the algorithm by a factor of three.




Figure 5. Coherence.

Reducing the slice resolution, by averaging four pixels into one, decreases the number of triangles, improves the surface construction efficiency and smooths the image. Although there is some loss of detail in the averaged slices, the averaging makes the number of triangles more manageable for high-resolution slices.

5.2 Functional Enhancements.

We have added a solid modeling capability to the algorithm. Boolean operations permit cutting and capping of solid models, as well as the extraction of multiple surfaces. In a medical application, cutting is analogous to performing surgery and capping (and texture mapping) is analogous to the medical imaging technique of reformatting.

We use the cube index described earlier to do Boolean operations on the surfaces. Here, just consider three values of the index:

- $index = 0$ for cubes outside the surface.
- $index = 255$ for cubes inside the surface.
- $0 < index < 255$ for cubes on the surface.

Solid modeling uses these notions of *inside*, *outside*, and *on* to create a surface. Analytic functions also provide the same information; so, for example the equation of a plane, $ax + by + cz - d$, tells where a given point lies with respect to the plane. Let $\sim S$, δS , and S represent sets of points that are outside, on, and inside a surface, respectively. Referring to Figure 6, we build a truth table, shown in Figure 7, for the Boolean intersection operation.

Nine entries in the truth table describe what to do when two surfaces have a given index. With x 's representing no operation, the entry for $(S, \sim P)$ shows that the cube in question is inside one surface but outside the other, resulting in no triangles. The $(\delta S, P)$ entry produces triangles from the S surface, while the $(S, \delta P)$ entry produces triangles from the P surface. The $(\delta S, \delta P)$ entry, created when a cube is on both surfaces, requires special processing. We clip




Figure 6. Point/Surface Relations.

	P	$\sim P$	dP
S	x	x	P
$\sim S$	x	x	x
dS	S	x	*

Figure 7. Truth Table.

each triangle from one surface against each triangle from the other, using the Sutherland-Hodgman clipping algorithm [28].

This technique applies to any surfaces that have inside/outside functions. We have used it with planes and with connectivity masks generated by separate image processing algorithms [9]. Application of a "logical or" truth table provides the capability for multiple surface extraction.

We implement texture mapping by finding the triangles on a plane's surface and attenuating the normal's length using the original slice data.

6. IMPLEMENTATION.

Marching cubes, written in C, runs on Sun Workstations¹ under Unix², VAX's under VMS³, and an IBM 3081 under IX/370⁴. We display the models using an in-house z-buffer program or a General Electric Graphicon 700⁵. For our models, the Graphicon displays at a rate of 10,000 triangles per second. In addition to surfaces of constant density, the software allows any number of planes that can be transparent, capped with triangles, or textured with interpolated density data. Medical practitioners refer to this texture mapping as reformatting. Execution times depend on the number of surfaces and resolution of the original data. Model creation times on a VAX 11/780 vary from 100 seconds for 64 by 64 by 48 SPECT data to 30 minutes for 260 by 260 by 93 CT studies. Times for the same studies on the IBM 3081 are twelve times faster. The number of triangles in a surface model is proportional to the area of the surface. This number can get large (over 500,000 in some cases), so we reduce it using cut planes and surface connectivity. Also, sometimes we reduce the resolution of the original data by filtering, producing a somewhat smoother surface with some loss of resolution.

7. RESULTS.

We have applied *marching cubes* to data obtained from CT, MR, and SPECT, as well as data generated from analytic functions. We present three case studies that illustrate the quality of the constructed surfaces and some modeling options. Each image was rendered at 512 by 512 resolution without antialiasing.

7.1 Computed Tomography.

The first case is a CT study of the head of a twelve year old male with a hole in the skull near the left side of the nose. The 93 axial slices are 1.5 mm thick, with pixel dimensions of 0.8 mm. This study by D.C. Hemmy, MD, of the Medical College of Wisconsin, illustrates the detail present in surfaces constructed by *marching cubes*. Figures 8 and 9 show the bone and soft tissue surfaces respectively. The tube in the patient's mouth is present to administer anesthetic during the scanning process. The soft tissue image shows fine detail that includes the patient's pierced ear and the impression of adhesive tape on the face. Although these details are not clinically significant, they do show the resolution present in the constructed surface. Figure 10 is a tilted view of the soft tissue surface that shows nasal and ear passages. In Figure 11, a sagittal cut, texture mapped with the original




Figure 8. Bone Surface.




Figure 9. Soft Tissue Surface.




Figure 10. Soft Tissue, Top View.

1 Sun Workstation is a trademark of Sun Microsystems.

2 Unix is a trademark of Bell Laboratories.

3 VAX and VMS are trademarks of Digital Equipment Corporation

4 IX/370 is a trademark of IBM.

5 Graphicon is a trademark of General Electric Company.




Figure 11. Sagittal Cut with Texture Mapping.

CT data, shows the slice data in relation to the constructed surface. The bone surface contains 550,000 triangles while the soft tissue surface has 375,000.

7.2 Magnetic Resonance.

The MR case of an adult male volunteer consists of 128 1.9 mm coronal slices. A 3D FT, flow compensated, fast sequence acquired the 128 slices in only 9 minutes. This pulse sequence, contrasting the unsaturated spins of the fresh blood flowing into the excited region of saturated spins, was produced by G. Glover of GE Medical Systems Group. Because of the complex anatomy present in the MR slices, we show, in Figure 12, the texture mapped cut surfaces intersected with the surface of the skin. Although the original slices are coronal, we show sagittal cuts to illustrate the algorithm's ability to interpolate texture on a cut plane. The largest surface model in the sequence contains 330,000 triangles, including triangles on the cut surface.

7.3 Single-Photon Emission Computed Tomography.

The SPECT study consisting of 29 coronal slices of the heart shows the algorithm's performance on low resolution data. D. Nowak from GE Medical Systems provided the 64 by 64 pixel data. Figure 13, showing the surface of the blood pool in the diastolic heart, contains 5,000 triangles. The descending aorta is the large vessel in the left of the picture.

8. CONCLUSIONS.

Marching cubes, a new algorithm for 3D surface construction, complements 2D CT, MR, and SPECT data by giving physicians 3D views of the anatomy. The algorithm uses a case table of edge intersections to describe how a surface cuts through each *cube* in a 3D data set. Additional realism is achieved by the calculation, from the original data, of the normalized gradient. The resulting polygonal structure can be displayed on conventional graphics display systems. Although these models often contain large numbers of triangles, surface cutting and connectivity can reduce this number. As CAD hardware increases in speed and capacity, we expect that *marching cubes* will receive increased use in practical, clinical environments.

Recently we developed another high-resolution surface construction algorithm called *dividing cubes* that generates points rather than triangles [10]. As the resolution of the 3D medical data increases, the number of triangles approaches



Figure 12. Rotated Sequence of Cut MR Brain.




Figure 13. Blood Pool in the Diastolic Heart.

the number of pixels in the displayed image. The density of surface points is chosen to cover the raster display. Both algorithms produce the same quality images, since the shading governs the perceived quality of the image.

9. ACKNOWLEDGMENT.

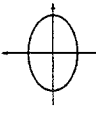
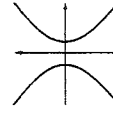
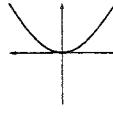

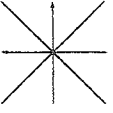
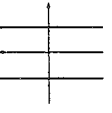

We thank C. Crawford from General Electric's Medical Systems for stimulating our work in this area. Throughout the project, he has provided us with data and encouragement to improve the algorithm. R. Redington from our laboratory's Medical Diagnostics Branch provided a stable research environment and insight into the practical applications of 3D in medicine. W. Leue assisted us in converting between the different medical data formats and provided interfaces to our MR equipment.

10. REFERENCES

- [1] Artzy, E., Frieder, G., and Herman, G. T. The Theory, Design, Implementation and Evaluation of a Three-Dimensional Surface Detection Algorithm. *Computer Graphics and Image Processing* 15, 1 (January 1981), 1-24.
- [2] Barillot, C., Gibaud, B., Scarabin, J., and Coatrieux, J. 3D Reconstruction of Cerebral Blood Vessels. *IEEE Computer Graphics and Applications* 5, 12 (December 1985), 13-19.
- [3] Bates, R. H., Garden, K. L., and Peters, T. M. Overview of Computerized Tomography with Emphasis on Future Developments. *Proc. of the IEEE* 71, 3 (March 1983), 356-372.
- [4] Bloch, P. and Udupa, J. K. Application of Computerized Tomography to Radiation Therapy and Surgical Planning. *Proc. of the IEEE* 71, 3 (March 1983), 351-355.
- [5] Brewster, L. J., Trivedi, S. S., Tut, H. K., and Udupa, J. K. Interactive Surgical Planning. *IEEE Computer Graphics and Applications* 4, 3 (March 1984), 31-40.
- [6] Burk, D. L., Mears, D. C., Kennedy, W. H., Cooperstein, L. A., and Herbert, D. L. Three-Dimensional Computed Tomography of Acetabula Fractures. *Radiology* 155, 1 (1985), 183-186.
- [7] Chen, L., Herman, G. T., Reynolds, R. A., and Udupa, J. K. Surface Shading in the Cuberille Environment. *IEEE Computer Graphics and Applications* 5, 12 (December 1985), 33-43.
- [8] Christiansen, H. N. and Sederberg, T. W. Conversion of Complex Contour Line Definitions into Polygonal Element Meshes. *Computer Graphics* 12, 3 (August 1978), 187-192.
- [9] Cline, H. E., Dumoulin, C. L., Lorensen, W. E., Hart, H. R., and Ludke, S. 3D Reconstruction of the Brain from Magnetic Resonance Images. *Magnetic Resonance Imaging* (1987, to appear).
- [10] Cline, H. E., Lorensen, W. E., Ludke, S., Crawford, C. R., and Teeter, B. C. High-Resolution Three-Dimensional Reconstruction of Tomograms. *Medical Physics* (1987, to appear).
- [11] Cook, L. T., Dwyer, S. J., Batnitzky, S., and Lee, K. R. A Three-Dimensional Display System for Diagnostic Imaging Applications. *IEEE Computer Graphics and Applications* 3, 5 (August 1983), 13-19.
- [12] Farrell, E. J. Color Display and Interactive Interpretation of Three-Dimensional Data. *IBM J. Res. Develop* 27, 4 (July 1983), 356-366.
- [13] Farrell, E. J., Zappulla, R., and Yang, W. C. Color 3D Imaging of Normal and Pathologic Intracranial Structures. *IEEE Computer Graphics and Applications* 4, 9 (September 1984), 5-17.
- [14] Fuchs, H., Kedem, Z. M., and Uselton, S. P. Optimal Surface Reconstruction from Planar Contours. *Comm. of the ACM* 20, 10 (October 1977), 693-702.
- [15] Gordon, D. and Reynolds, R. A. Image Space Shading of 3-Dimensional Objects. *Computer Graphics and Image Processing* 29, 3 (March 1985), 361-376.
- [16] Hale, J. D., Valk, P. E., and Watts, J. C. MR Imaging of Blood Vessels Using Three-Dimensional Reconstruction: Methodology. *Radiology* 157, 3 (December 1985), 727-733.
- [17] Hemmy, D. C., David, D. J., and Herman, G. T. Three-Dimensional Reconstruction of Craniofacial Deformity Using Computed Tomography. *Neurosurgery* 13, 5 (November 1983), 534-541.
- [18] Hemmy, D. C. and Tessier, P. L. CT of Dry Skulls with Craniofacial Deformities: Accuracy of Three-Dimensional Reconstruction. *Radiology* 157, 1 (October 1985), 113-116.
- [19] Herman, G. T. and Udupa, J. K. Display of 3D Digital Images: Computational Foundations and Medical Applications. *IEEE Computer Graphics and Applications* 3, 5 (August 1983), 39-46.
- [20] Hinshaw, W. S. and Lent, A. H. An Introduction to NMR Imaging: From the Bloch Equation to the Imaging Equation. *Proc. of the IEEE* 71, 3 (March 1983), 338-350.
- [21] Hoffman, E. A. and Ritman, E. L. Shape and Dimensions of Cardiac Chambers: Importance of CT Section Thickness and Orientation. *Radiology* 155, 3 (June 1985), 739-744.
- [22] Hohne, K. H. and Bernstein, R. Shading 3D-Images from CT Using Gray-Level Gradients. *IEEE Trans. on Medical Imaging MI-5*, 1 (March 1986), 45-47.
- [23] Keppel, E. Approximating Complex Surfaces by Triangulation of Contour Lines. *IBM J. Res. Develop* 19, 1 (January 1975), 2-11.
- [24] Knoll, G. F. Single-Photon Emission Computed Tomography. *Proc. of the IEEE* 71, 3 (March 1983), 320-329.
- [25] Meagher, D. J. Geometric Modeling Using Octree Encoding. *Computer Graphics and Image Processing* 19, 2 (June 1982), 129-147.
- [26] Robb, R. A., Hoffman, E. A., Sinak, L. J., Harris, L. D., and Ritman, E. L. High-Speed Three-Dimensional X-Ray Computed Tomography: The Dynamic Spatial Reconstructor. *Proc. of the IEEE* 71, 3 (March 1983), 308-319.
- [27] Sunguroff, A. and Greenberg, D. Computer Generated Images for Medical Application. *Computer Graphics* 12, 3 (August 1978), 196-202.
- [28] Sutherland, I. E. and Hodgman, G. W. Reentrant Polygon Clipping. *Comm. of the ACM* 17, 1 (January 1974), 32-42.
- [29] Trivedi, S. S., Herman, G. T., and Udupa, J. K. Segmentation Into Three Classes Using Gradients. *IEEE Trans. on Medical Imaging MI-5*, 2 (June 1986), 116-119.
- [30] Udupa, J. K. Interactive Segmentation and Boundary Surface Formation for 3-D Digital Images. *Computer Graphics and Image Processing* 18, 3 (March 1982), 213-235.
- [31] Vannier, M. W., Marsh, J. L., and Warren, J. O. Three Dimensional CT Reconstruction Images for Craniofacial Surgical Planning and Evaluation. *Radiology* 150, 1 (January 1984), 179-184.
- [32] Zucker, S. W. and Hummel, R. A. A Three-Dimensional Edge Operator. *IEEE Trans. on Pattern Analysis and Machine Intelligence PAMI-3*, 3 (May 1981), 324-331.

All the degenerate forms are central conics. Hence the parabola is the only noncentral conic. Table 4-8 summarizes the results.

Table 4-8 Summary of Conic Sections †

Name	Equation	Conditions	Type	Sketch
Ellipse	$\alpha x^2 + \beta y^2 = \kappa$	$\kappa, \alpha, \beta > 0$	Central	
Hyperbola	$\alpha x^2 + \beta y^2 = \kappa$	$\beta < 0 < \kappa, \alpha$	Central	
Parabola	$\alpha y^2 + \beta x = 0$ $\beta x^2 + \alpha y = 0$		Noncentral	
Empty set	$\alpha x^2 + \beta y^2 = \kappa$	$\alpha, \beta < 0 < \kappa$	(Central)	(No sketch)
Point	$\alpha x^2 + \beta y^2 = 0$	$\alpha, \beta > 0$	Central	
Pair of lines	$\alpha x^2 + \beta y^2 = 0$	$\beta < 0 < \alpha$	Central	
Parallel lines	$\alpha x^2 = \kappa$	$\alpha, \kappa > 0$	Central	
Empty set	$\alpha x^2 = \kappa$	$\alpha < 0 < \kappa$	(Central)	(No sketch)
'Repeated' line	$\alpha x^2 = 0$		Central	

†The ellipse, hyperbola and parabola are the nondegenerate conics, all the rest being called degenerate.

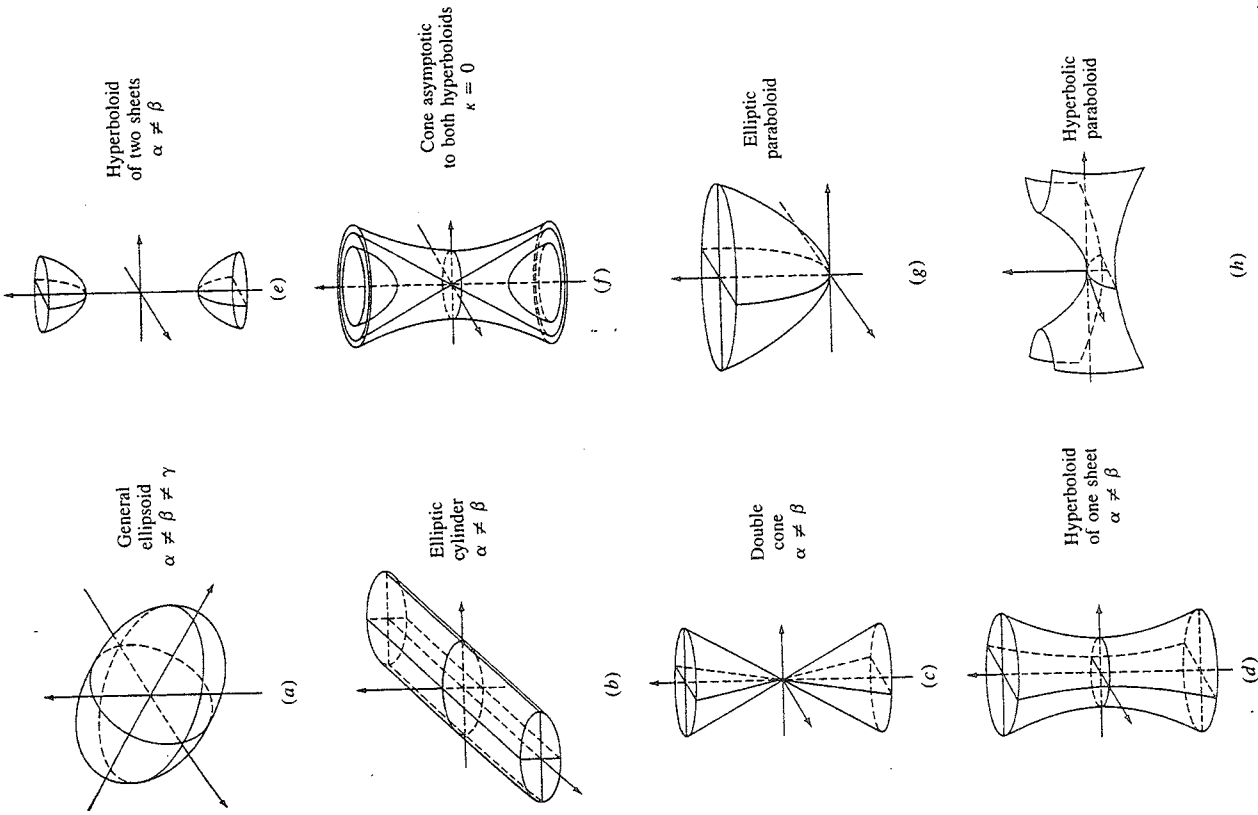


Figure 6-18 Quadric surfaces. (a) General ellipsoid $\alpha \neq \beta \neq \gamma$; (b) elliptic cylinder $\alpha \neq \beta$; (c) double cone $\alpha \neq \beta$; (d) hyperboloid of one sheet $\alpha \neq \beta$; (e) hyperboloid of two sheets $\alpha \neq \beta$; (f) cone asymptotic to both hyperboloids $\kappa = 0$; (g) elliptic paraboloid; (h) hyperbolic paraboloid.

6-2 SURFACES OF REVOLUTION

Perhaps the simplest method for generating a three-dimensional surface is to revolve a two-dimensional entity, e.g., a line or a plane curve, about an axis in space. Such surfaces are called surfaces of revolution. For simplicity, initially the axis of rotation is assumed coincident with the x -axis and in the positive direction. The point, line or plane curve to be rotated is assumed to lie in the xy plane. Later a procedure to remove these restrictions is developed.

The simplest entity that can be rotated about an axis is a point. Provided that the point does not lie on the axis, rotation through an angle of 2π (360°) yields a circle. Rotation through an angle less than 2π (360°) yields a circular arc.

Next in complexity is a line segment parallel to and not coincident with the axis of rotation. Rotation through an angle of 2π (360°) yields a circular cylinder. The radius of the cylinder is the perpendicular distance from the line to the rotation axis. The length of the cylinder is the length of the line segment. An example is shown in Fig. 6-1.




Figure 6-1 Cylindrical surface of revolution. (a) Schematic; (b) result.

If the line segment and the axis of rotation are coplanar and the line segment is not parallel to the rotation axis, then rotation about the axis through 2π (360°) yields a truncated right circular cone. The radius of the cone at each end is the perpendicular distance from the end points of the line segment to the axis of rotation. The length of the cone is the projected length of the line segment on the rotation axis. An example is shown in Fig. 6-2.

Again if the line segment and the axis of rotation are coplanar and the line segment is perpendicular to the axis of rotation, then rotation through 2π (360°) yields a planar disc. If the line segment intersects (or touches) the axis of rotation, a solid disc results; otherwise the disc has a circular hole in it.

Finally if the line segment is skew to the axis of rotation, i.e., not coplanar, then rotation through 2π (360°) yields a hyperboloid of one sheet (see Secs. 6-4 and 6-7).

Closed or open polygons can also be used to generate surfaces of revolution. An example representing a cone with a cylindrical hole in it is shown in Fig. 6-4.

The parametric equation for a point on a surface of revolution is developed by recalling that the parametric equation of the entity to be rotated, e.g.,

$$P(t) = [x(t) \quad y(t) \quad z(t)] \quad 0 \leq t \leq t_{max}$$

is a function of the single parameter t . Rotation about an axis causes the location of the point to also be a function of the rotation angle ϕ . Thus, a point on a surface of revolution is specified by two parameters t and ϕ . It is a biparametric function as shown in Fig. 6-5.




Figure 6-2 Conical surface of revolution. (a) Schematic; (b) result.




Figure 6-4 A surface of revolution from a closed polygon. (a) Schematic; (b) result.

Rotating plane curves also yields surfaces of revolution. A sphere is obtained by rotating an origin-centered semicircle in the xy plane about the x -axis,

An ellipsoid of revolution is obtained if the parametric equation of an origin centered semiellipse in the xy plane is substituted for that of the circle.

If the axis of rotation does not pass through the center of a complete circle or ellipse, then rotation about the axis generates a torus with a circular or elliptical cross section as appropriate.




Figure 6-5 Biperametric surface of revolution.




Figure 6-7 Tori. (a) Circular cross section; (b) elliptical cross section.

Any parametric curve can be used to create a surface of revolution. Obvious possibilities are cubic spline, parabolically blended, Bézier and B-spline curves. Figure 6-9 shows a surface of revolution created using a relatively simple parabolically blended curve. Figure 6-10 shows a handleless mug created as a surface of revolution using a complex open B-spline curve. Notice that the mug has both an inside and an outside. Here rotation is about the y -axis.




Figure 6-9 Parabolically blended surface of revolution. (a) Generating curve; (b) surface.




Figure 6-10 B-spline surface of revolution. (a) Polygon vertices; (b) B-spline curve; (c) surface.

6-3 SWEEP SURFACES

A three-dimensional surface is also obtained by traversing an entity, e.g., a line, polygon or curve, along a path in space. The resulting surfaces are called sweep surfaces. Sweep surface generation is frequently used in geometric modeling. The simplest sweep entity is a point. The result of sweeping a point along a path is, of course, not a surface but a space curve. However, it serves to illustrate the fundamental technique.

Consider the position vector $P[x \ y \ z \ 1]$ swept along the path represented by the sweep transformation $[T(s)]$. The position vector $Q(s)$ representing the resulting curve is given by

$$Q(s) = P[T(s)] \quad s_1 \leq s \leq s_2 \quad (6-14)$$

The transformation $[T(s)]$ determines the shape of the curve. For example, if the path is a straight line of length n parallel to the z -axis, then (see Eq. 3-14)

$$[T(s)] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & ns & 1 \end{bmatrix} \quad 0 \leq s \leq 1$$

If the path is an origin-centered circle in a $z = \text{constant}$ plane, then (see Eq. 3-8)

$$[T(s)] = \begin{bmatrix} \left(\frac{r}{x}\right) \cos\{2\pi(s + s_i)\} & 0 & 0 & 0 \\ 0 & \left(\frac{r}{y}\right) \sin\{2\pi(s + s_i)\} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad 0 \leq s \leq 1$$

where $s_i = (1/2\pi) \tan^{-1}(y_i/x_i)$ and for $P[x \ y \ z \ 1]$, $r = \sqrt{x^2 + y^2}$. Here, the subscript i is used to indicate the initial or starting point.

Complex paths can be developed by combining simple paths. For example, combining the two previous path transformations yields a single turn of a helical

path along the z -axis, i.e.,

$$[T(s)] = \begin{bmatrix} \left(\frac{r}{x}\right) \cos\{2\pi(s + s_i)\} & 0 & 0 & 0 \\ 0 & \left(\frac{r}{y}\right) \sin\{2\pi(s + s_i)\} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & ns & 1 \end{bmatrix} \quad 0 \leq s \leq 1$$

The simplest sweep surface is obtained by traversing a line segment along a path. Recalling that the parametric equation of a line segment is

$$P(t) = P_1 + (P_2 - P_1)t \quad 0 \leq t \leq 1 \quad (6-15)$$

The corresponding sweep surface is given by

$$Q(t, s) = P(t)[T(s)] \quad 0 \leq t \leq 1, \quad s_1 \leq s \leq s_2 \quad (6-16)$$

where again $[T(s)]$ is the sweep transformation. If the sweep transformation contains only translations and/or local or overall scalings, the resulting surface is planar. If the sweep transformation contains rotations, the resulting surface is non-planar. Figure 6-13 shows the helical sweep surface obtained by simultaneously translating along and rotating about the x -axis a line originally parallel to the y -axis.

Parametric curves, e.g., cubic splines, parabolically blended, Bézier and B-spline curves, are also used to generate sweep surfaces. The surface equation is identical to Eq. (6-15) where now $P(t)$ represents the parametric curve. Figure 6-14 shows a sweep surface generated from a single cubic spline curve segment swept parallel to the z -axis.




Figure 6-13 A helical sweep surface.

Care must be taken to avoid degenerate surfaces or parts of surfaces in generating sweep surfaces from lines and curves. An example is shown in Fig. 6-15. Here an *s*-shaped curve lying in the *xy* plane is swept parallel to the *x*-axis. Notice that the 'tails' at the left and right sides are degenerate (i.e., lines) of zero area. Such degenerate surface parts can cause difficulty in geometric modeling systems.

In addition to open curves, closed polygons and curves are used to create sweep surfaces. If the end surfaces are included, then the sweep surface encloses a finite volume. Many geometric modeling systems create primitive volumes in this way. A square or rectangle swept along a straight path yields a rectangular parallelepiped. A circle swept along a straight path yields a cylinder. A circle of decreasing radius swept along a straight path yields a cone. Rotation about the sweep axis is also possible. Figure 6-16 shows the sweep surface resulting from a plane square perpendicular to and centered on the *x*-axis being swept along and simultaneously rotated through 90° about the *x*-axis.

In sweeping a planar polygon or closed curve along an arbitrary path there are two important considerations. The first is, what point in the polygon continuously lies on the path? In general, any point in a polygon or on a closed curve can continuously lie on the path. For different points, the resulting surfaces are different.

The second is, what is the direction of the normal to the polygon or closed curve as the path is swept out? Here, two approaches are typically taken. The normal to the polygon or closed curve is in the direction of the instantaneous tangent to the path. Alternately, the normal direction is specified independent of the path. This latter alternate is extremely flexible. Two examples are shown in Fig. 6-17. Figure 6-17 shows the sweep surface generated by a square centered on the *x*-axis swept along the path $x = ls, y = \cos(\pi s) - 1$. In Fig. 6-17a the normal is maintained in the direction of the *x*-axis. In Fig. 6-17b the normal is maintained in the direction of the instantaneous tangent to the path. Notice the difference between the two sweep surfaces. A detailed example further illustrates this concept.




Figure 6-15 A sweep surface with degenerate parts.




Figure 6-16 Sweep surface generated by a square simultaneously sweeping along and rotating about the *x*-axis.




Figure 6-14 A cubic spline based sweep surface. (a) Curve; (b) surface.




Figure 6-17 A polygon swept along a path. (a) Normal in the *x* direction; (b) normal in the instantaneous tangent vector direction.




Figure 5-25 A Bézier curve and its defining polygon.

of Bézier curves are immediately known. For example:

The basis functions are real.

The degree of the polynomial defining the curve segment is one less than the number of defining polygon points.

The curve generally follows the shape of the defining polygon.

The first and last points on the curve are coincident with the first and last points of the defining polygon.

The tangent vectors at the ends of the curve have the same direction as the first and last polygon spans, respectively.

The curve is contained within the convex hull of the defining polygon, i.e., within the largest convex polygon obtainable with the defining polygon vertices. In Fig. 5-25 the convex hull is shown by the dashed and dotted lines.

The curve exhibits the variation diminishing property. Basically this means that the curve does not oscillate about any straight line more often than the defining polygon.

The curve is invariant under an affine transformation.

Several four-point Bézier polygons and the resulting cubic curves are shown in Fig. 5-26. With just the information given above a user quickly learns to predict the shape of a curve generated by a Bézier polygon.

Mathematically a parametric Bézier curve is defined by

$$P(t) = \sum_{i=0}^n B_i J_{n,i}(t) \quad 0 \leq t \leq 1 \tag{5-62}$$

where the Bézier or Bernstein basis or blending function is

$$J_{n,i}(t) = \binom{n}{i} t^i (1-t)^{n-i} \tag{5-63}$$

with

The previously discussed techniques for curve generation constrained the curves to pass through *existing* data points; i.e., they are curve fitting techniques. In many cases excellent results are achieved with these methods. They are particularly suited to shape description, where the basic shape is arrived at by experimental evaluation or mathematical calculation. Examples are aircraft wings, engine manifolds, mechanical and/or structural parts. However, there is another class of shape design problems that depends on *both* aesthetic *and* functional requirements. These problems are frequently termed *ab initio* design. Examples are the 'skin' of car bodies, aircraft fuselages, ship hulls, furniture and glassware. *Ab initio* design problems cannot be formulated entirely in terms of quantitative criteria but must be resolved by a judicious combination of computational and heuristic methods. These methods are most effective when implemented in a real time interactive graphics environment.

The curve fitting techniques, e.g., the cubic spline fitting method previously discussed, are rendered ineffective for interactive *ab initio* design. Numerical specification of both direction and magnitude of tangent derivatives to control the shape does not provide the intuitive 'feel' required for *ab initio* design. An obvious relationship between the numbers and the curve shape does not always exist.

An alternate method of shape description suitable for *ab initio* design of free-form curves and surfaces was developed by Pierre Bézier. Although Bézier originally derived the mathematical basis of the technique from geometrical considerations (Refs. 5-11 to 5-13), Forrest (Ref. 5-14) and Gordon and Riesenfeld (Ref. 5-15) have shown that the result is equivalent to the Bernstein basis or polynomial approximation function.

A Bézier curve is determined by a defining polygon, e.g., that shown in Fig. 5-25. Because the Bézier basis is also the Bernstein basis, several properties




Figure 5-26 Bézier polygons for cubics.

$$\binom{n}{i} = \frac{n!}{i!(n-i)!} \tag{5-64}$$

$J_{n,i}(t)$ is the i th n th-order Bernstein basis function. Here n , the degree of the defining Bernstein basis function and thus of the polynomial curve segment, is one less than the number of points in the defining Bézier polygon. The vertices of the Bézier polygon are numbered from 0 to n as shown in Fig. 5-25.† Also $(0)^0 \equiv 1$ and $0! \equiv 1$.

Figure 5-27 shows the blending functions for several values of n . Notice the symmetry of the functions. Each of the blending functions is of degree n . For

† This notation and polygon numbering scheme is chosen to be consistent with the vast body of existing literature on Bézier curves and Bernstein basis functions. For programming purposes an alternate formulation is somewhat more convenient, i.e.,

$$P(t) = \sum_{i=1}^N B_i J_{N,i}(t)$$

$$J_{N,i} = \binom{N-1}{i-1} t^{i-1} (1-t)^{N-i}$$

$$\text{and } \binom{N-1}{i-1} = \frac{(N-1)!}{(i-1)!(N-i)!}$$

The polygon points are numbered from 1 to N . The transformations $n = N - 1$ and $i = I - 1$ convert between the two notations.

example, each of the four blending functions shown in Fig. 5-27b for $n = 3$ is a cubic. The maximum value of each blending function occurs at $t = i/n$ and is given in Ref. 5-14 by

$$J_{n,i} \left(\frac{i}{n} \right) = \binom{n}{i} \frac{i^i (n-i)^{n-i}}{n^n} \tag{5-65}$$

For example, for a cubic $n = 3$. The maximum values for $J_{3,1}$ and $J_{3,2}$ occur at $1/3$ and $2/3$, respectively, with values

$$J_{3,1} \left(\frac{1}{3} \right) = \frac{4}{9} \quad \text{and} \quad J_{3,1} \left(\frac{2}{3} \right) = \frac{4}{9}$$

Figure 5-27b illustrates this result.




Figure 5-27 Bézier/Bernstein blending functions. (a) Three polygon points, $n = 2$; (b) four polygon points, $n = 3$; (c) five polygon points, $n = 4$; (d) six polygon points, $n = 5$.

Examining Eqs. (5-62) to (5-64) for the first point on the curve, i.e., at $t = 0$, shows that

$$J_{n,0}(0) = \frac{n!(1-0)^{n-0}}{n!} = 1 \quad i = 0$$

and
$$J_{n,i}(0) = \frac{n!(0)^i(1-0)^{n-i}}{i!(n-i)!} = 0 \quad i \neq 0$$

Thus,
$$P(0) = B_0 J_{n,0}(0) = B_0$$

and the first point on the Bézier curve and on its defining polygon are coincident as previously claimed.

Similarly, for the last point on the curve, i.e., at $t = 1$,

$$J_{n,n}(1) = \frac{n!(1)^n(0)^{n-n}}{n!(1)} = 1 \quad i = n$$

$$J_{n,i}(1) = \frac{n!}{i!(n-i)!} t^i (1-1)^{n-i} = 0 \quad i \neq n$$

Thus,
$$P(1) = B_n J_{n,n}(1) = B_n$$

and the last point on the Bézier curve and the last point on its defining polygon are coincident. The blending functions shown in Fig. 5-27 illustrate these results.

Further, it can be shown that for any given value of the parameter t , the summation of the basis functions is precisely one; i.e.,

$$\sum_{i=0}^n J_{n,i}(t) = 1 \tag{5-66}$$

5-9 B-SPLINE CURVES

From a mathematical point of view, a curve generated by using the vertices of a defining polygon is dependent on some interpolation or approximation scheme to establish the relationship between the curve and the polygon. This scheme is provided by the choice of basis function. As noted in Sec. 5-8, the Bernstein basis produces Bézier curves generated by Eq. (5-62). Two characteristics of the Bernstein basis, however, limit the flexibility of the resulting curves. First the number of specified polygon vertices fixes the order of the resulting polynomial which defines the curve. For example, a cubic curve must be defined by a polygon with four vertices and three spans. A polygon with six vertices always produces a fifth-degree curve. The only way to reduce the degree of the curve is to reduce the number of vertices, and conversely the only way to increase the degree of the curve is to increase the number of vertices.

The second limiting characteristic is due to the global nature of the Bernstein basis. This means that the value of the blending function $J_{n,i}(t)$ given by Eq. (5-63) is nonzero for all parameter values over the entire curve. Since any point on a Bézier curve is a result of blending the values of all defining vertices, a change in one vertex is felt throughout the entire curve. This eliminates the ability to produce a local change within a curve.

For example, since the end slopes of a Bézier curve are established by the directions of the first and last polygon spans, it is possible to change the middle vertex of a five-point polygon without changing the *direction* of the end slopes. However, the shape of the total curve is affected due to the global nature of the Bernstein basis. This lack of local control is detrimental in some applications.

There is another basis, called the B-spline basis (from Basis spline), which contains the Bernstein basis as a special case. This basis is generally nonglobal. The nonglobal behavior of B-spline curves is due to the fact that each vertex B_i is associated with a unique basis function. Thus, each vertex affects the shape of a curve only over a range of parameter values where its associated basis function is nonzero. The B-spline basis also allows the order of the basis function and hence the degree of the resulting curve to be changed without changing the number of defining polygon vertices. The theory for B-splines was first suggested by Schoenberg (Ref. 5-20). A recursive definition useful for numerical computation

was independently discovered by Cox (Ref. 5-21) and by de Boor (Ref. 5-22). Gordon and Riesenfeld (Refs. 5-15 and 5-23) applied the B-spline basis to curve definition.

Again letting $P(t)$ be the position vectors along the curve as a function of the parameter t , a B-spline curve is given by

$$P(t) = \sum_{i=1}^{n+1} B_i N_{i,k}(t) \quad t_{\min} \leq t < t_{\max}, \quad 2 \leq k \leq n+1 \quad (5-83)$$

where the B_i are the position vectors of the $n+1$ defining polygon vertices and the $N_{i,k}$ are the normalized B-spline basis functions.

For the i th normalized B-spline basis function of order k (degree $k-1$), the basis functions $N_{i,k}(t)$ are defined by the Cox-deBoor recursion formulas. Specifically,

$$N_{i,1}(t) = \begin{cases} 1 & \text{if } x_i \leq t < x_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (5-84a)$$

and

$$N_{i,k}(t) = \frac{(t - x_i) N_{i,k-1}(t)}{x_{i+k} - x_i} + \frac{(x_{i+k} - t) N_{i+1,k-1}(t)}{x_{i+k} - x_{i+1}} \quad (5-84b)$$

The values of x_i are elements of a knot vector satisfying the relation $x_i \leq x_{i+1}$. The parameter t varies from t_{\min} to t_{\max} along the curve $P(t)$.[†] The convention $0/0 = 0$ is adopted.

Formally a B-spline curve is defined as a polynomial spline function of order k (degree $k-1$) since it satisfies the following two conditions:

The function $P(t)$ is a polynomial of degree $k-1$ on each interval $x_i \leq t < x_{i+1}$.

$P(t)$ and its derivatives of order $1, 2, \dots, k-2$ are all continuous over the entire curve.

Thus, for example, a fourth-order B-spline curve is a piecewise cubic curve.

Because a B-spline basis is used to describe a B-spline curve, several properties in addition to those already mentioned above are immediately known:

The sum of the B-spline basis functions for any parameter value t can be shown (see Refs. 5-15 and 5-22) to be

$$\sum_{i=1}^{n+1} N_{i,k}(t) \equiv 1 \quad (5-85)$$

Each basis function is positive or zero for all parameter values, i.e., $N_{i,k} \geq 0$.

Except for $k=1$ each basis function has precisely one maximum value.

The maximum order of the curve is equal to the number of defining polygon vertices.

The curve exhibits the variation diminishing property. Thus the curve does not oscillate about any straight line more often than its defining polygon.

The curve generally follows the shape of the defining polygon.

Any affine transformation can be applied to the curve by applying it to the defining polygon vertices; i.e., the curve is transformed by transforming the defining polygon vertices.

The curve lies within the convex hull of its defining polygon.

In fact, the convex hull property of B-spline curves is stronger than that for Bézier curves. For a B-spline curve of order k (degree $k-1$) a point on the curve lies within the convex hull of k neighboring points. Thus, all points on a B-spline curve must lie within the union of all such convex hulls formed by taking k successive defining polygon vertices. Figure 5-32, where the convex hulls are shown shaded, illustrates this effect for different values of k . Notice in particular that for $k=2$ the convex hull is just the defining polygon itself. Hence, the B-spline curve is also just the defining polygon itself.

Using the convex hull property it is easy to see that if all the defining polygon vertices are colinear, then the resulting B-spline curve is a straight line for all orders k . Further, if l colinear polygon vertices occur in a noncolinear defining polygon, then the straight portions of the defining curve (if any) start and end at least $k-2$ spans from the beginning and end of the series of colinear polygon vertices. If the series of colinear polygon vertices is completely contained within a noncolinear defining polygon, then the number of colinear curve spans is at least $l-2k+3$. If the series of colinear polygon vertices occurs at the end of a noncolinear defining polygon, then the number of colinear curve spans is at least $l-k+1$. Figure 5-33 illustrates these results.

If at least $k-1$ coincident defining polygon vertices occur, i.e., $B_i = B_{i+1} \dots = B_{i+k-2}$, then the convex hull of B_i to B_{i+k-2} is the vertex itself. Hence, the resulting B-spline curve must pass through the vertex B_i . Figure 5-34 illustrates this point for $k=3$. Further, since a B-spline curve is everywhere C^{k-2} continuous, it is C^{k-2} continuous at B_i .

Finally, note that because of these continuity properties B-spline curves smoothly transition, with C^{k-2} continuity, into embedded straight segments as shown in Fig. 5-35.

Equations (5-84) clearly show that the choice of knot vector has a significant influence on the B-spline basis functions $N_{i,k}(t)$ and hence on the resulting B-spline curve. The only requirement for a knot vector is that it satisfy the relation $x_i \leq x_{i+1}$; i.e., it is a monotonically increasing series of real numbers. Fundamentally three types of knot vector are used: uniform, open uniform (or open) and nonuniform.

In a uniform knot vector, individual knot values are evenly spaced. Examples

are

$$[0 \quad 1 \quad 2 \quad 3 \quad 4]$$

[†] Here note that, in contrast to the Bézier curve, the defining polygon vertices or points are numbered from 1 to $n+1$.




Figure 5-32 Convex hull properties of B-spline curves.




Figure 5-33 B-spline convex hull properties for colinear curve segments. (a) Within the defining polygon vertices; (b) at the end of the defining polygon vertices.

$$[-0.2 \quad -0.1 \quad 0 \quad 0.1 \quad 0.2]$$

In practice, uniform knot vectors generally begin at zero and are incremented by 1 to some maximum value or are normalized in the range between 0 and 1, equal decimal intervals, e.g.,

$$[0 \quad 0.25 \quad 0.5 \quad 0.75 \quad 1.0]$$

For a given order k , uniform knot vectors yield periodic uniform basis functions for which

$$N_{i,k}(t) = N_{i-1,k}(t-1) = N_{i+1,k}(t+1)$$

Thus each basis function is a translate of the other. Figure 5-36 illustrates this.




Figure 5-34 Convex hull for coincident polygon vertices, $k = 3$.




Figure 5-35 Smooth (C^{k-2}) transition into straight segments.

An open uniform knot vector has multiplicity of knot values at the ends equal to the order k of the B-spline basis function. Internal knot values are evenly spaced. Some examples using integer increments are

- $k = 2$ [0 0 1 2 3 4 4]
- $k = 3$ [0 0 0 1 2 3 3 3]
- $k = 4$ [0 0 0 0 1 2 2 2 2]




Figure 5-36 Periodic uniform B-spline basis functions, $[X] = [0\ 1\ 2\ 3\ 4\ 5\ 6]$, $n + 1 = 4$, $k = 3$.

or for normalized increments

- $k = 2$ [0 0 1/4 1/2 3/4 1 1]
- $k = 3$ [0 0 0 1/3 2/3 1 1 1]
- $k = 4$ [0 0 0 0 1/2 1 1 1 1]

Formally, an open uniform knot vector is given by

- $x_i = 0 \quad 1 \leq i \leq k$
- $x_i = i - k \quad k + 1 \leq i \leq n + 1$
- $x_i = n - k + 2 \quad n + 2 \leq i \leq n + k + 1$

The resulting open uniform basis functions yield curves that behave most nearly like Bézier curves. In fact, when the number of defining polygon vertices is equal to the order of the B-spline basis and an open uniform knot vector is used, the B-spline basis reduces to the Bernstein basis. Hence, the resulting B-spline curve is a Bézier curve. In that case, the knot vector is just k zeros followed by k ones. For example, for four polygon vertices the fourth order ($k = 4$) open uniform knot vector is

$$[0\ 0\ 0\ 0\ 1\ 1\ 1\ 1]$$

A cubic Bézier/B-spline curve results. The corresponding open uniform basis functions are shown in Fig. 5-27b. Additional open uniform basis functions are shown in Fig. 5-37.




Figure 5-37 Open uniform B-spline basis functions, $[X] = [0\ 0\ 0\ 1\ 2\ 2\ 2]$, $k = 3$, $n + 1 = 4$.

312 MATHEMATICAL ELEMENTS FOR COMPUTER GRAPHICS

Nonuniform knot vectors may have either unequally spaced and/or multiple internal knot values. They may be periodic or open. Examples are

$$[0 \ 0 \ 0 \ 1 \ 1 \ 2 \ 2 \ 2 \ 2]$$

$$[0 \ 1 \ 2 \ 2 \ 3 \ 4]$$

$$[0 \ 0.28 \ 0.5 \ 0.72 \ 1]$$

Figures 5-38b to e show several nonuniform B-spline basis functions for order $k = 3$. The knot vectors used to generate the basis functions all have multiplicity of k equal values at the ends. Figure 5-38a gives the basis functions for an open uniform knot vector for comparison. Notice the symmetry of the basis functions in Figs. 5-38a and b and how that symmetry is lost for the nonuniform basis functions in Figs. 5-38c to e. Notice also that for multiple knot values within the knot vector a cusp occurs in one of the basis functions. Further, in Figs. 5-38d and e notice the shift of the location of the cusp corresponding to the change in location of the multiple knot value in the knot vector.




Figure 5-38 Nonuniform basis functions for $n + 1 = 5, k = 3$.
 (a) $[X] = [0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$;
 (b) $[X] = [0 \ 0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$;
 (c) $[X] = [0 \ 0 \ 0 \ 0 \ 1.8 \ 2.2 \ 3 \ 3 \ 3]$;
 (d) $[X] = [0 \ 0 \ 0 \ 1 \ 1 \ 3 \ 3 \ 3]$;
 (e) $[X] = [0 \ 0 \ 0 \ 2 \ 2 \ 3 \ 3 \ 3]$.

The above discussion shows the significant influence of the choice of knot vector on the shape of the B-spline basis functions and hence on the shape of any resulting B-spline curve.

Because of the flexibility of B-spline basis functions and hence of the resulting B-spline curves, different types of control 'handles' are used to influence the shape of the curve. Control is achieved by:

Changing the type of knot vector and hence basis function: periodic uniform, open uniform or nonuniform.

Changing the order k of the basis function.

Changing the number and position of the defining polygon vertices.

Using multiple polygon vertices.

Using multiple knot values in the knot vector.

These effects are illustrated first with open B-spline curves and then with periodic B-spline curves and finally with nonuniform B-spline curves.

The behavior of an open B-spline curve is in many respects analogous to that of a Bézier curve. In fact, as has previously been mentioned, when the order of a B-spline curve is equal to the number of defining polygon vertices, the B-spline basis reduces to the Bernstein basis. Consequently, the resulting B-spline curve is identical to a Bézier curve. For an open B-spline curve of any order ($k \geq 2$) the first and last points on the curve are coincident with the first and last polygon vertices. Further, the slope of the B-spline curve at the first and last polygon vertices is equal to the slope of the first and last polygon spans.

Figure 5-41 shows three open B-spline curves of different order, each defined by the same four polygon vertices. The fourth-order curve corresponds to the Bézier curve. This curve is a single cubic polynomial segment. The third-order curve is composed of two parabolic curve segments joined at the center of the second span with C^1 continuity. The second-order curve reproduces the defining polygon. It consists of three linear 'curve' segments joined at the second and third polygon vertices with C^0 continuity. Notice that all three curves have the same end slopes, determined by the slope of the first and last spans of the defining polygon. Notice also that as the order of the curve increases, the resulting curve looks less like the defining polygon. Thus, increasing the order 'tightens' or 'smooths' the curve.

Figure 5-42 illustrates the effect of multiple or coincident vertices on the defining polygon. The B-spline curves are all of the order $k = 4$. The lowest curve is defined by four polygon vertices as shown. Here the knot vector is $[0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1]$. The intermediate curve is defined by five polygon vertices with two coincident vertices at the second polygon vertex, i.e., $[3 \ 9]$. The knot vector for this curve is $[0 \ 0 \ 0 \ 0 \ 1 \ 2 \ 2 \ 2 \ 2]$. The highest curve is defined by six polygon vertices with three coincident vertices at $[3 \ 9]$. The knot vector is $[0 \ 0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3 \ 3]$. Thus, the defining polygons for the three curves are B_1, B_2, B_3, B_4 ; $B_1, \bar{B}_2, B_2, B_3, B_4$ and $B_1, \bar{B}_2, \bar{B}_2, B_2, B_3, B_4$, respectively.




Figure 5-41 Effect of varying order on B-spline curves.

The lowest curve is composed of a single cubic segment. The intermediate curve is composed of two segments joined midway between B_2 and B_3 . The highest curve is composed of three segments. The first is from B_1 to B_2 , the second from B_2 to midway between B_2 and B_3 . The final segment is from midway between B_2 and B_3 to B_4 . Notice that as the number of multiple vertices at B_2 increases, the curve is pulled closer to B_2 . If the number of multiple vertices is $k - 1$, then a sharp corner or cusp is created. This sharp corner is predicted by the convex hull properties of B-spline curves. Close examination of Fig. 5-42 shows that on both sides of the multiple vertex location a linear segment occurs.




Figure 5-42 Effect of multiple vertices at B_2 on a B-spline curve, $k = 4$.

322 MATHEMATICAL ELEMENTS FOR COMPUTER GRAPHICS

Although a cusp exists when $k - 1$ multiple vertices occur, the C^{k-2} differentiability of the curve is maintained. At first glance, this might seem contradictory. However, a cusp is defined by a zero tangent vector. But a zero tangent vector does not preclude the tangent vector varying continuously. The ability to include sharp corners or cusps within a continuously C^{k-2} differentiable curve is an important characteristic of B-spline curves.

Finally notice that each of the curves has the same slope at the ends. Figure 5-43 shows three fourth-order B-spline curves. The defining polygons each have eight vertices as indicated. The three curves shown are obtained by moving the polygon vertex B_5 successively to B_3' and B_5' . Note that moving B_5 influences the curve only over a limited region. Specifically, only the curve segments corresponding to the polygon spans B_3B_4 , B_4B_5 and B_5B_6 , B_6B_7 are affected by the movement of B_5 . In general, the curve is affected only over those curve segments corresponding to $\pm k/2$ polygon spans around the displaced point.

Turning now to periodic B-spline curves, Fig. 5-44 shows three periodic B-spline curves of different orders. Each of the curves is defined by the same polygon vertices as the open B-spline curves in Fig. 5-41. For $k = 2$ the B-spline curve again coincides with the defining polygon. However, notice that for periodic B-spline curves for $k > 2$ the first and last points on the B-spline curve do not correspond to the first and last defining polygon vertices. Nor in general is the slope at the first and last points the same as that of the first and last defining polygon spans. For $k = 3$ the B-spline curve starts at the midpoint of the first polygon span and ends at the midpoint of the last polygon span as indicated by the arrows. These effects are a result of the reduced parameter range for periodic B-spline basis functions. For $k = 2$ the periodic knot vector is $[0 \ 1 \ 2 \ 3 \ 4 \ 5 \ 6]$ with parameter range of $1 \leq t \leq 4$. For $k = 3$ the periodic knot vector is $[0 \ 1 \ 2 \ 3 \ 4 \ 5 \ 6]$ with parameter range of $2 \leq t \leq 4$. For $k = 4$ the periodic knot vector is $[0 \ 1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7]$ with parameter range of $3 \leq t \leq 4$. Comparing these results with those for open knot vectors permit the curve to be defined over the full range of parameter values. The effect is to 'pull' the curve out to the ends of the defining polygon.

Here again the fourth-order curve consists of a single cubic polynomial segment, the third-order curve of two parabolic segments joined at the center of the second polygon span with C^1 continuity, and the second-order 'curve' of three linear segments joined at the second and third polygon vertices with C^0 continuity. Notice that again increasing order has a 'smoothing' effect on the curve, but here it also decreases the curve length.

Figure 5-45 illustrates that the effect of multiple vertices in the defining polygon is similar for periodic and open B-spline curves. The small inset shows the details near B_2 .




Figure 5-44 Effect of varying order on periodic B-spline curves.




Figure 5-45 Effect of multiple vertices on a periodic B-spline curve, $k = 4$.

Turning now to nonuniform B-spline curves, Fig. 5-49 illustrates the effect of multiple interior knot values on the resulting curve. The upper third-order ($k = 3$) curve in Fig. 5-49 was generated with the open knot vector $[0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$. The basis functions for this curve are shown in Fig. 5-38a. The lower third-order curve in Fig. 5-49 was generated with the nonuniform knot vector $[0 \ 0 \ 0 \ 1 \ 1 \ 3 \ 3 \ 3]$. The basis functions are shown in Fig. 5-38d. An identical curve is obtained with the nonuniform knot vector $[0 \ 0 \ 0 \ 2 \ 2 \ 3 \ 3 \ 3]$ and the basis functions shown in Fig. 5-38e.

Figure 5-49 shows that the multiple interior knot value yields a sharp corner or cusp at B_3 . A multiple knot value introduces a span of zero length. Consequently, the width of support of the basis functions is reduced. Further, multiple interior knot values, in contrast to multiple polygon vertices, reduce the differentiability of the basis function at x_i to C^{k-m-1} where $m \leq k - 1$ is the multiplicity of the interior knot value. Locally, the nonuniform curve in Fig. 5-49 is C^0 ($k - m - 1 = 3 - 2 - 1 = 0$) continuous near B_3 . Thus, a position discontinuity or 'corner' in the curve occurs.

Figure 5-50 shows third-order ($k = 3$) open nonuniform B-spline curves generated using a knot vector with interior knot values proportional to the chord distances between polygon vertices. Specifically, the knot vector is given by

$$x_i = 0 \quad 1 \leq i \leq k$$

$$x_{i+k} = \left(\frac{i}{n-k+2} c_{i+1} + \sum_{j=1}^i c_j \right) (n-k+2) \quad 1 \leq i \leq n-k+1 \quad (5-86)$$

$$x_i = n-k+2 \quad n+1 \leq i \leq n+k$$

where $c_i = |B_{i+1} - B_i|$. For equally spaced polygon vertices the result reduces to equally spaced interior knot values, i.e., an open uniform knot vector. Hartley and Judd (Ref. 5-25) have suggested a similar scheme that yields distinct interior knot values. The curve generated with an open uniform knot vector is shown for comparison. Also shown is the curve obtained with a pair of coincident vertices at B_3 .

From these results it appears that nonuniform B-spline curves do not greatly differ from uniform B-spline curves unless the relative distances between polygon vertices vary radically.




Figure 5-49 Nonuniform B-spline curves, $k = 3$. (a) $[X] = [0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$; (b) $[X] = [0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 3 \ 3 \ 3]$.




Figure 5-50 Comparison of open nonuniform B-spline curves. (a) Uniform knot vector; (b) nonuniform chord distance proportional knot vector; (c) nonuniform chord distance proportional knot vector with double vertex at B_3 .

representing the common analytical shapes—lines, planes, conic curves including circles, free-form curves, quadric and sculptured surfaces—used in computer graphics and computer aided design.

Versprille (Ref. 5-39) was the first to discuss rational B-splines. The seminal papers by Tiller (Ref. 5-40) and Piegl and Tiller (Ref. 5-41) form the basis of the current discussion. Interestingly enough, nonuniform rational B-splines (NURBS) have been an Initial Graphics Exchange Specification (IGES) standard since 1983 (see Ref. 5-42). IGES is the standard for the interchange of design information between various computer aided design systems and between computer aided design and computer aided manufacturing systems. Rational B-splines have been incorporated into a number of geometric modeling systems. They have also been implemented in hardware (VLSI or microcode) by a number of graphics workstation manufacturers.

A rational B-spline curve is the projection of a nonrational (polynomial) B-spline curve defined in four-dimensional (4D) homogeneous coordinate space back into three-dimensional (3D) physical space. Specifically,

$$P(t) = \sum_{i=1}^{n+1} B_i^h N_{i,k}(t) \tag{5-121}$$

where the B_i^h 's are the 4D homogeneous defining polygon vertices for the non-rational 4D B-spline curve. $N_{i,k}(t)$ is the nonrational B-spline basis function previously given in Eq. (5-84).

Projecting back into three-dimensional space by dividing through by the homogeneous coordinate yields the rational B-spline curve

$$P(t) = \frac{\sum_{i=1}^{n+1} B_i h_i N_{i,k}(t)}{\sum_{i=1}^{n+1} h_i N_{i,k}(t)} = \sum_{i=1}^{n+1} R_i R_{i,k}(t) \tag{5-122}$$

where the B_i 's are the 3D defining polygon vertices for the rational B-spline curve and the

$$R_{i,k}(t) = \frac{h_i N_{i,k}(t)}{\sum_{i=1}^{n+1} h_i N_{i,k}(t)} \tag{5-123}$$

are the rational B-spline basis functions. Here, $h_i \geq 0$ for all values of i .†

As can be seen from Eqs. (5-121) to (5-123), rational B-spline basis functions and curves are a generalization of nonrational B-spline basis functions and curves. They carry forward nearly all the analytic and geometric characteristics of their nonrational B-spline counterparts. In particular:

Each rational basis function is positive or zero for all parameter values, i.e., $R_{i,k} \geq 0$.

† Note that rational B-spline basis functions for $h_i < 0$ are valid (see Ref. 5-39) but are not convenient in terms of the current discussion.




Figure 5-57 Rational B-spline basis functions for $n + 1 = 5$, $k = 3$ with open knot vector $[X] = [0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$, $[H] = [1 \ 1 \ h_3 \ 1 \ 1]$. (a) $h_3 = 0$; (b) $h_3 = 0.25$; (c) $h_3 = 1.0$; (d) $h_3 = 5.0$.

5-13 RATIONAL B-SPLINE CURVES

Rational curve and surface descriptions were first introduced into the computer graphics literature by Coons (Ref. 5-32). Rational forms of the cubic spline and Bézier curves previously discussed in this chapter are well known in the literature (see Refs. 5-33 to 5-37). Rational forms of the conic sections are also well known (see Ref. 5-38). Both because of space limitations and because they form a unifying foundation, the current discussion is limited to rational B-spline curves. Rational B-splines provide a single precise mathematical form capable of

The sum of the rational B-spline basis functions for any parameter value t is one, i.e.,

$$\sum_{i=1}^{n+1} R_{i,k}(t) \equiv 1 \quad (5-124)$$

Except for $k = 1$, each rational basis function has precisely one maximum.

A rational B-spline curve of order k (degree $k - 1$) is C^{k-2} continuous everywhere.

The maximum order of the rational B-spline curve is equal to the number of defining polygon vertices.

A rational B-spline curve exhibits the variation diminishing property.

A rational B-spline curve generally follows the shape of the defining polygon.

A rational B-spline curve lies within the union of convex hulls formed by k successive defining polygon vertices.

Any *projective* transformation is applied to a rational B-spline curve by applying it to the defining polygon vertices; i.e., the curve is invariant with respect to a *projective* transformation. Note that this is a stronger condition than that for a nonrational B-spline which is only invariant with respect to an *affine* transformation.

From Eqs. (5-85) and (5-123) it is clear that when all $h_i = 1$, $R_{i,k}(t) = N_{i,k}(t)$. Thus, nonrational B-spline basis functions and curves are included as a special case of rational B-spline basis functions and curves. Further, it is easy to show that an open rational B-spline curve with order equal to the number of defining polygon vertices is a rational Bézier curve. For the case of all $h_i = 1$, the rational Bézier curve reduces to a nonrational Bézier curve. Thus, both rational and nonrational Bézier curves are included as special cases of rational B-spline curves.

Since rational B-splines are a four-dimensional generalization of nonrational B-splines, algorithms for degree-raising (see Ref. 5-27 and Ex. 6-18), subdivision (see Sec. 5-12 and Refs. 5-28 to 5-31) and curve fitting (see Sec. 5-11) of nonrational B-spline curves are valid for rational B-splines simply by applying them to the 4D defining polygon vertices.

Open uniform, periodic uniform and nonuniform knot vectors can be used to generate rational B-spline basis functions and rational B-spline curves.

In Eqs. (5-122) and (5-123) the homogeneous coordinates h_i (occasionally called weights) provide additional blending capability. $h = 1$ is called the affine space. By convention it corresponds to physical space. The effect of the homogeneous coordinates h on the rational B-spline basis functions is shown in Fig. 5-57. Here, with an open uniform knot vector $[0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3] \ (n+1 = 5, k = 3)$ is used with a homogeneous coordinate vector $h_i = 1, i \neq 3$. Values of h_3 range from 0 to 5. The rational B-spline basis functions shown in Fig. 5-57c with $h = 1$ are identical to the corresponding nonrational B-spline basis functions. The rational B-spline curve for $h_3 = 1$, shown in Fig. 5-58, is also identical




Figure 5-58 Rational B-spline curves for $n + 1 = 5, k = 3$ with open knot vector $[X] = [0 \ 0 \ 0 \ 1 \ 2 \ 3 \ 3 \ 3]$ and $[H] = [1 \ 1 \ h_3 \ 1 \ 1]$.

with the corresponding nonrational B-spline curve. Notice that for $h_3 = 0$ (see Fig. 5-57a) $R_{3,3} = 0$ everywhere. Thus, the corresponding polygon vertex, B_3 , effectively has no influence on the shape of the corresponding B-spline curve. This effect is shown in Fig. 5-58 where the defining polygon vertices B_2 and B_4 are connected by a straight line. Figure 5-57 also shows that as h_3 increases $R_{3,3}$ also increases; but, as a consequence of Eq. (5-124), $R_{2,3}$ and $R_{4,3}$ decrease. The effects on the corresponding rational B-spline curves are shown in Fig. 5-58. Note, in particular, that as h_3 increases the curve is pulled closer to B_3 . Hence, as mentioned previously, the homogeneous coordinates provide additional blending capability.