Introduction to 2D and 3D Computer Graphics

Emphasizing both 2D and 3D Computer Graphics

Karla Steinbrugge Fant

Overview of 2D & 3D Graphics

Primitive Graphic Objects

- Make up a picture or a portion of a picture
- Allow the application or user...

  ...to create lines, arcs, symbols, character strings, polygons, circles, ellipses, images, and nonstandard objects (called GDPs)
- Have associated attributes that control the visible appearance

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Are formed from primitive graphic objects

  ...and may be defined hierarchically

- Are used to create complex pictures in 3D

  ...for example, a cube could be built from multiple polygons (i.e., squares)

- Are useful in determining...

  ...a product's manufacturability

  ...the best shape, color, layout, and orientation

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Are also useful in determining...
  - ...how mechanical parts will fit together
  - ...the quantity of material required
  - ...the internal components of a complex structure
  - ...the cost, area, and volume
  - ...how surfaces & materials will interact with light

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Associate their...

  …shape, geometry, color, reflectivity, transmittance, surface smoothness, and texture as part of the complex object

- • Can be defined as...

  ...surfaces

  ...implicit surfaces

  ...polyhedrals

  ...curved surfaces

  ...fractals

  ...graftals

  ...ellipsoids

  ...cylinders ...and much much more!

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Surfaces...

  ...are the simplest complex object

  ...consist of a collection of 3D points

  ...may be a simple list, contours, slices, or sections

  ...require a dense distribution of points for accurate rendering

- Implicit surfaces...

  ...are defined by algebraic formulas

  ...include quadrics

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Polyhedrals...

  ...are the most commmonly used complex graphic object

  ...consist of networks of polygons

  ...are used to form polygon mesh models

  ...where polygons are sized, shaped, and positioned to completely tile the surface of an area

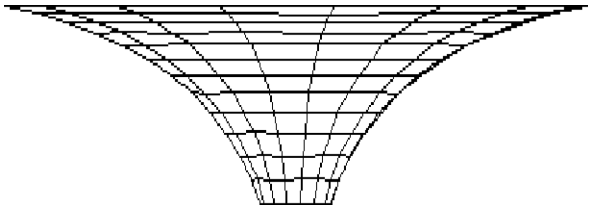

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Curved surfaces...

  ...are created by surface patches, obtained using numerical methods to approximate a coarse polygon grid or mesh

  ...allow the structure of the mesh to define the curvature of the surface

  ...require that only the vertices of the mesh be stored

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Curved surfaces include...

  ...Bezier

  ...Hermite

  ...bicubic

  ...B-spline

  ...ß-spline (Beta-spline)

  ...polynomial

  ...rational polynomial ...and many more!

Overview of 2D & 3D Graphics
Complex Graphic Objects

- The most common curves used to generate computer graphics pictures are...
  ...Bezier ...B-spline
- Bezier (curves, surfaces, patches)...
  ...were developed in the 1960s by Pierre Bezier
  ...were developed to design Renault car bodies!
  ...were probably preceeded by P. de Casteljau at Citroen in the early 1960s
  ...however his work was not discovered till 1975
  are specified by control points from

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Bezier (curves, surfaces, patches)...
  - ...are specified by control points from which a cubic polynomial is derived
  - ...allow the shape to be determined entirely from the position of four control points
  - ...allow the shape to be altered by changing the control points

Overview of 2D & 3D Graphics

Complex Graphic Objects

- **Bicubic Bezier curve examples...**

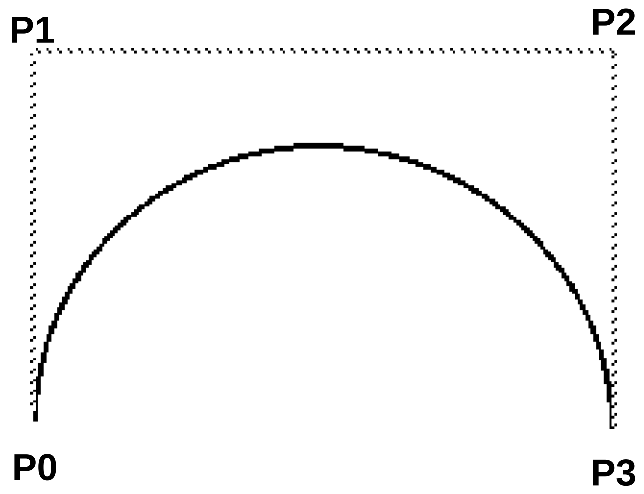

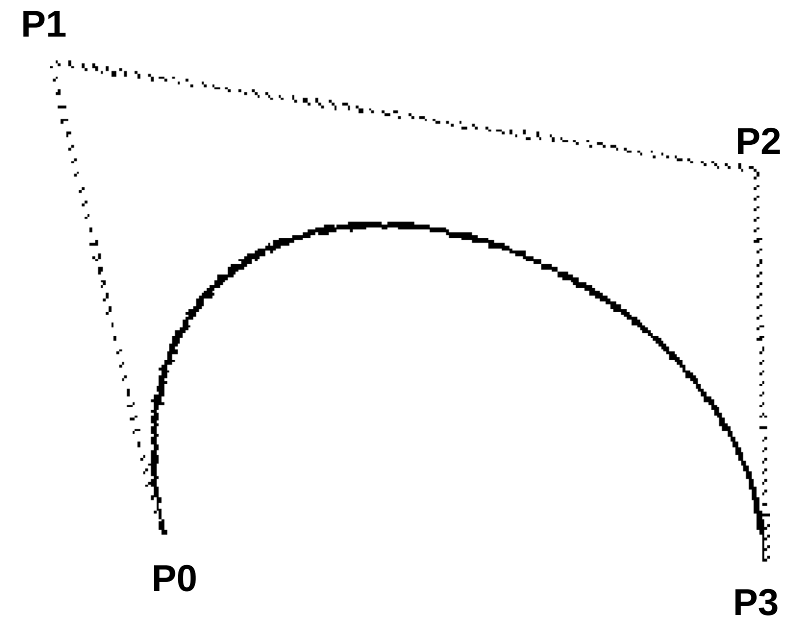

Overview of 2D & 3D Graphics

Complex Graphic Objects

- B-spline (curves, surfaces, patches)...

  ...are also common techniques for generating curves

  ...allow more than four points to be used

  ...in their simplest form are considered to be Bezier

  ...are specified using control points

  ...have localness...allowing control points to affect only a portion of a curve instead of an entire curve

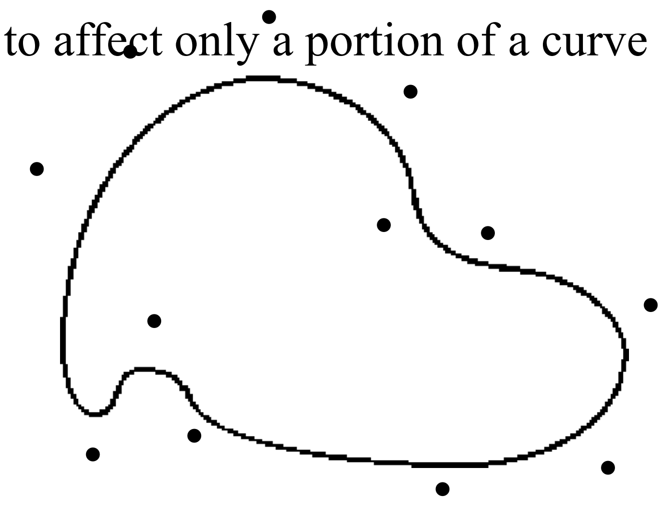

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Fractals...

  ...create surfaces using an implicit model that produces data when requested

  ...can model highly irregular and complex objects, such as clouds, hills, rocks, valleys, etc.

  ...do not require all of the data of each surface point to be stored (as in surface and polyhedral methods)

  ...are special "random" processes that achieve a natural looking "statistical" irregularity

  ...in theory, the frequency of shape distribution is supposed to be the same at all scales or zoom factors

Overview of 2D & 3D Graphics
Complex Graphic Objects

- Graftals...

  ...are like fractals, except they use a deterministic process to model more repetitive patterns

  ...can model trees, leaves, ice, snow, etc.

Overview of 2D & 3D Graphics

Complex Graphic Objects

- Complex three-dimensional graphic objects can be modeled, viewed, and displayed using...

  ...polyhedral -- *represents objects as planar polyhedra*

  ...free form -- *represents objects as patches*

  ...solid -- *represents objects as solid primitives*

  ...procedural -- *represents objects using construction rules and procedures for execution*

Overview of 2D & 3D Graphics

Modeling

- Modeling...

  ...allows information to be added to produce a better visual effect with finer detail

  ...allows information to be removed to produce a simpler and more efficient picture

  ...allows objects created to be manipulated

  ...provides rotation about an axis

  ...provides translation, keeping a constant orientation

  ...provides scaling to change an object's size

Overview of 2D & 3D Graphics
Modeling

- Modeling...

  ...supports curve fitting

  ...allows objects to be deformed: such as, skewed, stretched, folded, and randomly or fractally altered

Overview of 2D & 3D Graphics

Viewing

- Viewing...

  ...allows objects created to be prepared for display on a two-dimensional surface

  ...transforms primitive and complex objects using parallel or perspective projections

Overview of 2D & 3D Graphics
Viewing

- Parallel projection...

  ...collapses all objects to lie in a plane

  ...the placement of objects relative to the plane is preserved exactly

  ...the placement of objects perpendicular to the plane is ignored

  ... objects appear the same size regardless of how far away they really are

  ...useful for mechanical drawings

Overview of 2D & 3D Graphics

Viewing

- Perspective projection...

  ...causes objects to be drawn as if viewed from a particular point, called "camera point" or "viewpoint"

  ...causes objects farther away to be rendered smaller than objects closer

  ...creates depth perspective

  ...the closer the camera point is to the object, the more extreme the perspective...

  (like taking a picture with a wide angle lens)

Overview of 2D & 3D Graphics

Viewing -- Examples

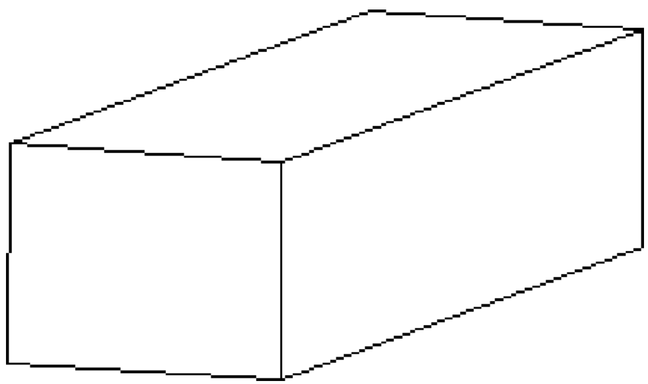

Parallel Projection

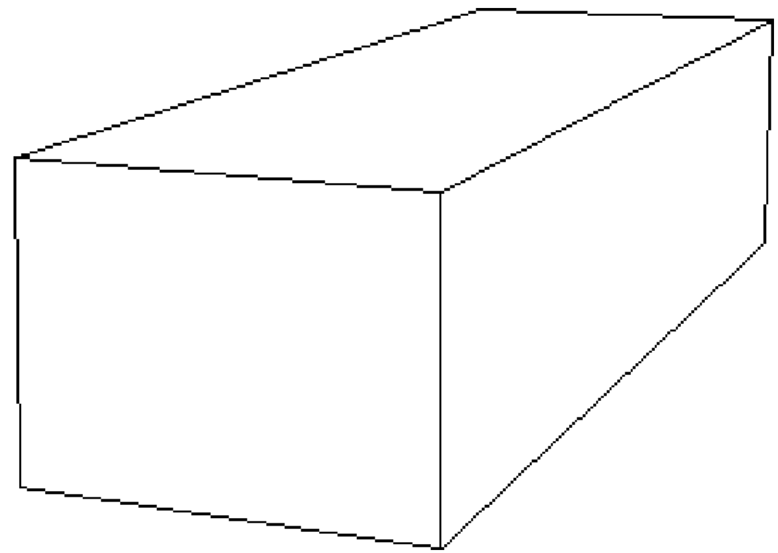

Perspective Projection

Overview of 2D & 3D Graphics

Rendering

- Rendering...

  ...transforms primitive and complex graphic objects into pictures on the drawing surface

  ...for 3D pictures it is the process of creating a 2D image on the drawing surface

  ...includes culling back-facing polygons

  ...scan converts (or rasterizes) objects into device-dependent units (usually called pixels)

  ...uses the appropriate light sources, atmosphere effects, shading, anti-aliasing, and color models

Overview of 2D & 3D Graphics

Rendering & Shading

- Rendering may be accomplished using...

  ...wireframe

  ...Z-buffer

  ...ray tracing

  ...radiosity methods

  ...or a combination of approaches

Overview of 2D & 3D Graphics

Rendering & Shading

- Shading...

  ...is the process of rendering to produce realistic graphic pictures

  ...helps to understand the three-dimensional nature of objects

  ...provides a means for viewing “curved” objects that in fact were generated by a series of flat planar polygons

Overview of 2D & 3D Graphics

Rendering & Shading

- Z-buffer techniques...

  ...allow for hidden surface removal

  ...shade pixels in the interior of a polygon, using an incremental shading scheme

  ...interpolate the depth of each pixel from the Z values after modeling and viewing transformations have been applied

  ...allow you to imagine a 3D screen space

Overview of 2D & 3D Graphics
Rendering & Shading

- Ray tracing techniques...
  ...provide the ultimate in picture quality
  ...take into account how light actually behaves
  ...can render reflective and transparent objects
  ...can render shadows

The Rendering Process
Rendering Pipeline: Abstract Rendering

- A Graphic Object changes from a parameterization of a primitive with associated attributes to a subset of real-valued two-dimensional space with associated attributes
- • For example, after abstract rendering, a line has width!

The Rendering Process

Rendering Pipeline: Abstract Rendering

- Attributes associated with graphic objects are either...

  ...geometric or ...cosmetic

- Geometric attributes are...

  ...applied during abstract rendering

  ...affected by subsequent modeling and viewing transformations

  if line width is a geometric attribute, when a line is scaled by 2 the width of the line would also be scaled by 2

The Rendering Process

Rendering Pipeline: Abstract Rendering

- Cosmetic attributes are...

  ...applied during later phases of rendering

  ...not affected by subsequent modeling and viewing transformations

  for example, the interior style is usually cosmetic; as a polygon rotates, the hatch pattern would stay fixed

The Rendering Process

Rasterization

- Is the process that converts transformed objects into pixel values...which may be stored in a frame buffer or video memory
- Consists of three operations...

  ...scan conversion

  ...image-precision hidden line & surface removal

  ...shading
- Requires calculating each object's contribution to each pixel

The Rendering Process

Rasterization: Polygon Example

- The first step for rasterizing a polygon is to determine the initial scan line intersected by the polygon...

  ...this is determined by the vertex with the smallest y

  ...which is usually at a single pixel, with two edges projecting downward

  ...and calculate the delta values of x,z,R,G,B for the pixel

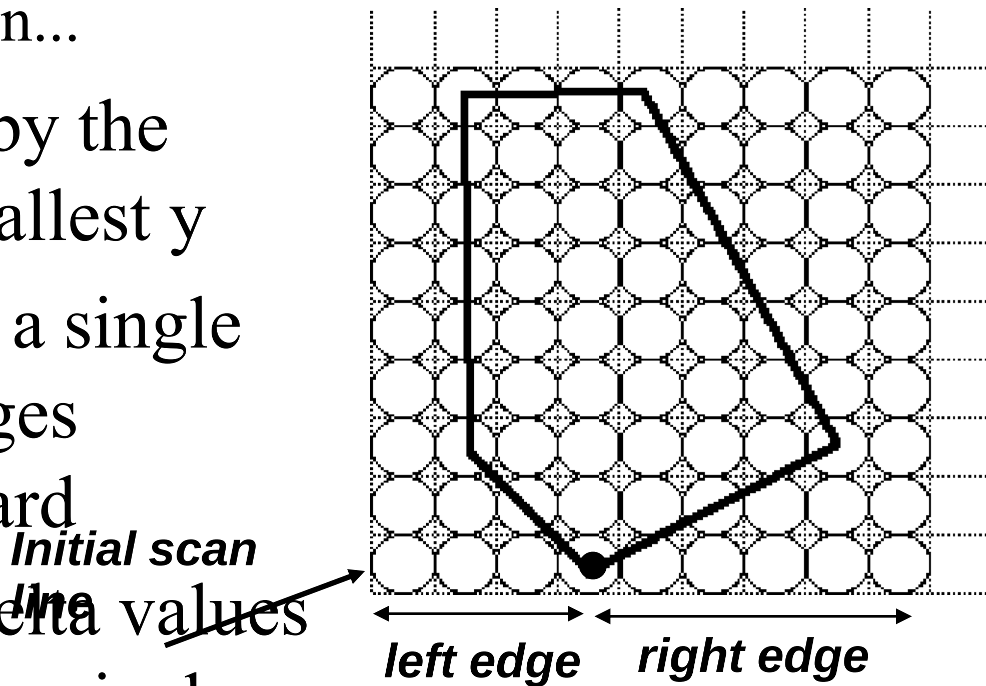

The Rendering Process
Rasterization: Polygon Example

- The next step is to process the edge
  ...by calculating the intersection points of the left and right edges with the current scan line
  ...calls a contiguous sequence of pixels on a scan line a span
  ...and calculates the delta values for incrementing x,z,R,G,B from pixel to pixel within the span ...this computation is done once per scan line

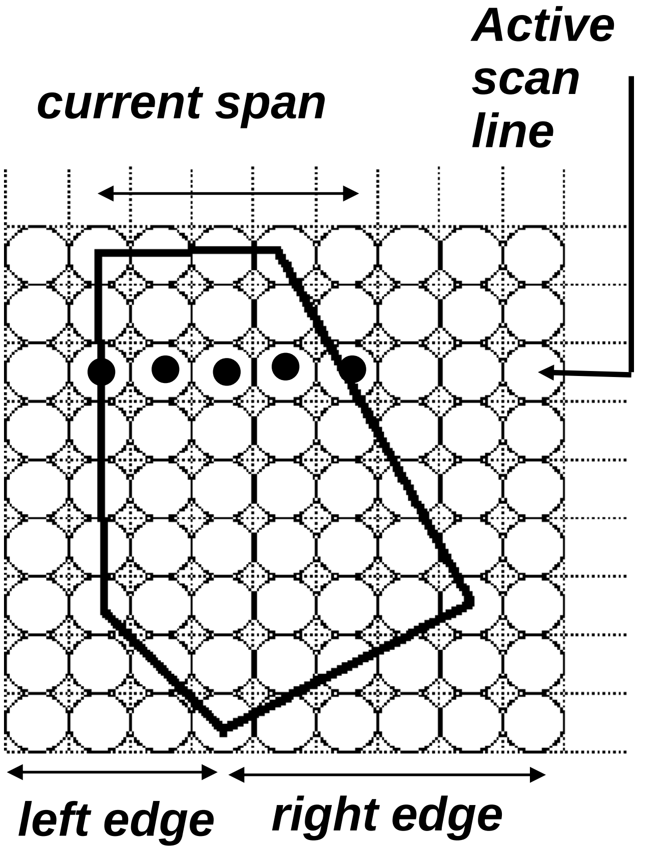

- The third step is to process the spans...

The Rendering Process

Rasterization: Example Block Diagram

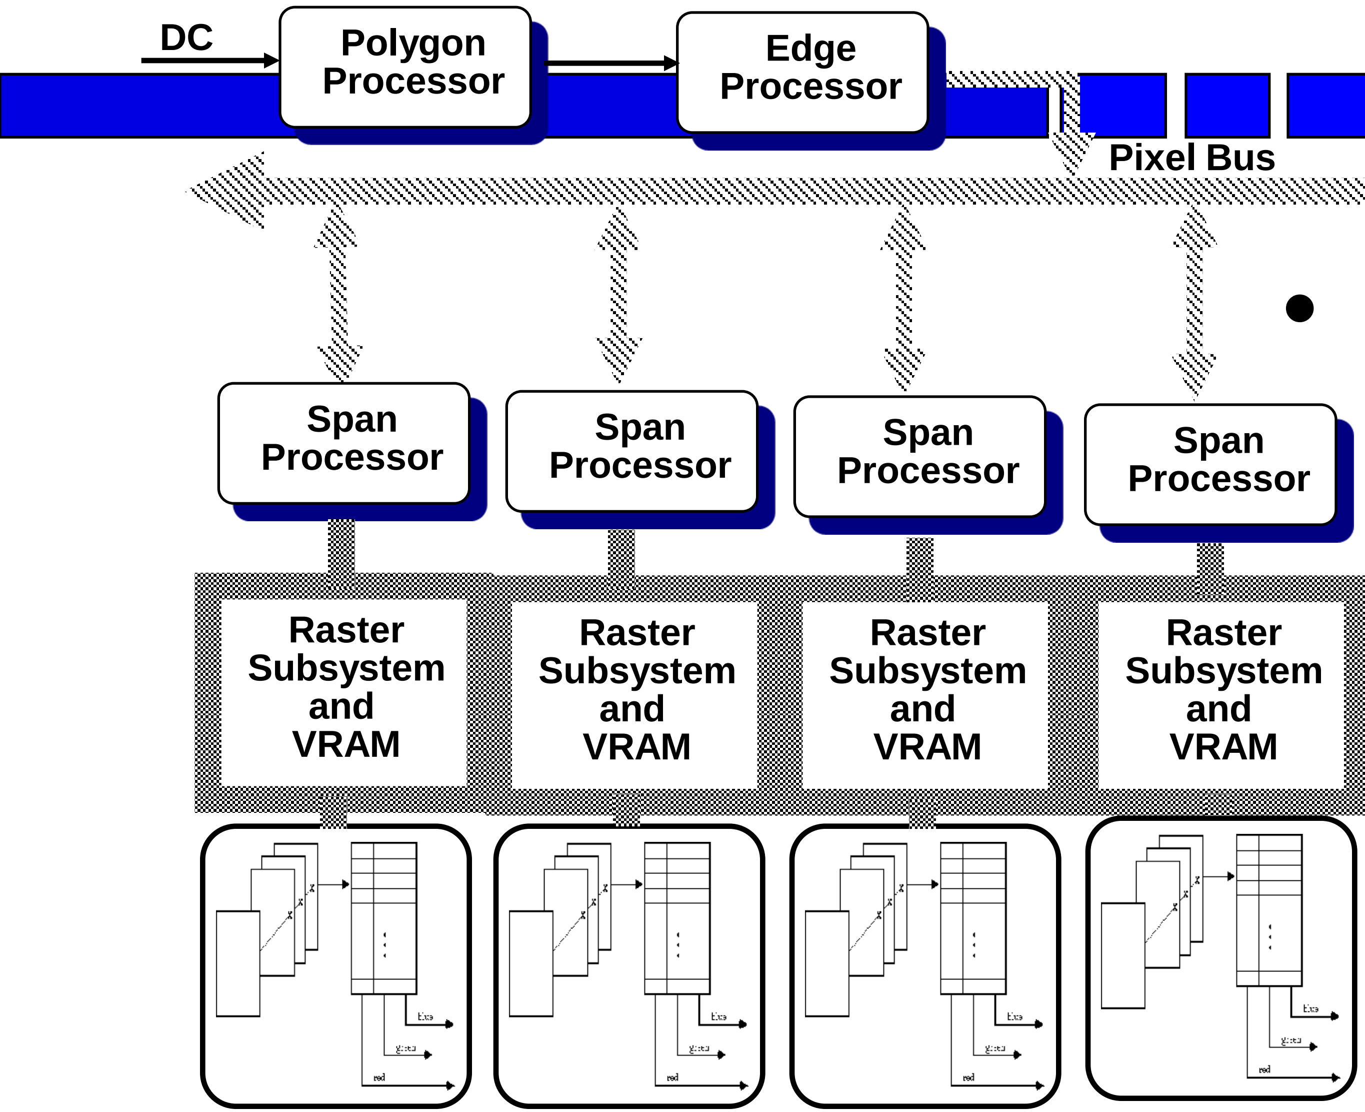

Span processors selectively update the bitmaps and z-buffer

Control and memory may be provided to allow for 32-bit images, 24-bit z-buffer, overlay planes, and window bitplanes

Overlay bitmaps support popup menus and windows

The Rendering Process
Rasterization: Scan Conversion

- Is the process that converts objects to pixels and...
  - ...depends on the underlying geometry of the object
  - ...depends on the attributes that alter the object's shape

The Rendering Process
Rasterization: Scan Conversion

- First, lets summarize the basics...
  - ...lines are rasterized per Bresenham's 1965 algorithm
  - ...lines are rasterized such that they generate a linear sequence of pixels with no gaps
  - ...boundaries of filled areas may allow gaps, since polygons are filled using horizontal line segments
  - ...for any given scan line, the right and left boundaries that intersect the scan line must be generated

The Rendering Process

Raserization: Scan Conversion Basics

- Pixel sequences...

  ...for lines and boundaries may differ

  ...are calculated using DDA algorithms *(digital differential analyser)*

  ...use DDA to find out how much the x coordinate increases per scan line and then repeatedly adds this increment

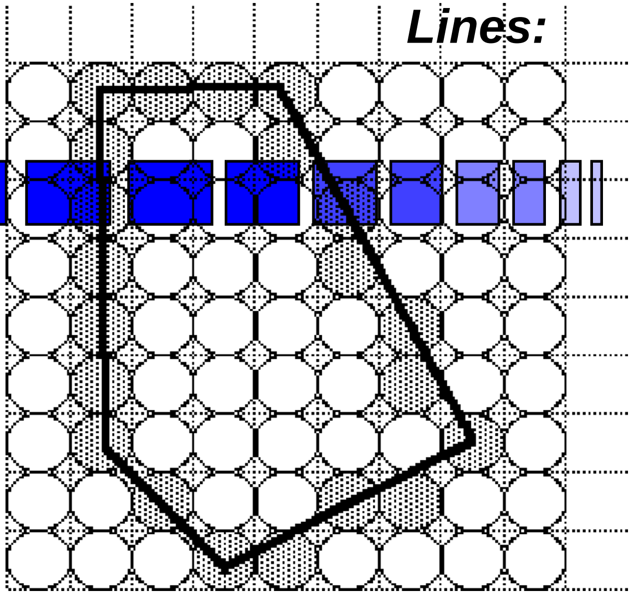

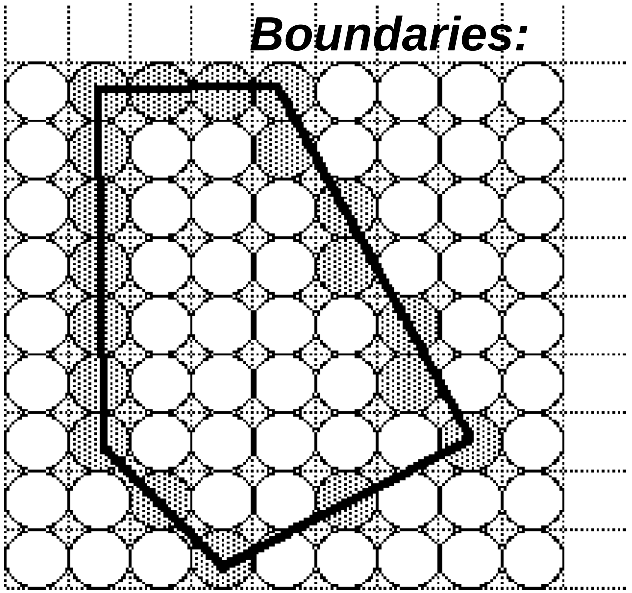

The Rendering Process

Rasterization: Scan Conversion

- Advanced polyline algorithms...

  ...are more complex than drawing lines repeatedly since the pixels at the end points should not be written more than once

  for example if a polyline is drawn in xor mode, pixels drawn by two connecting lines would be xored twice...which is usually NOT DESIRABLE!

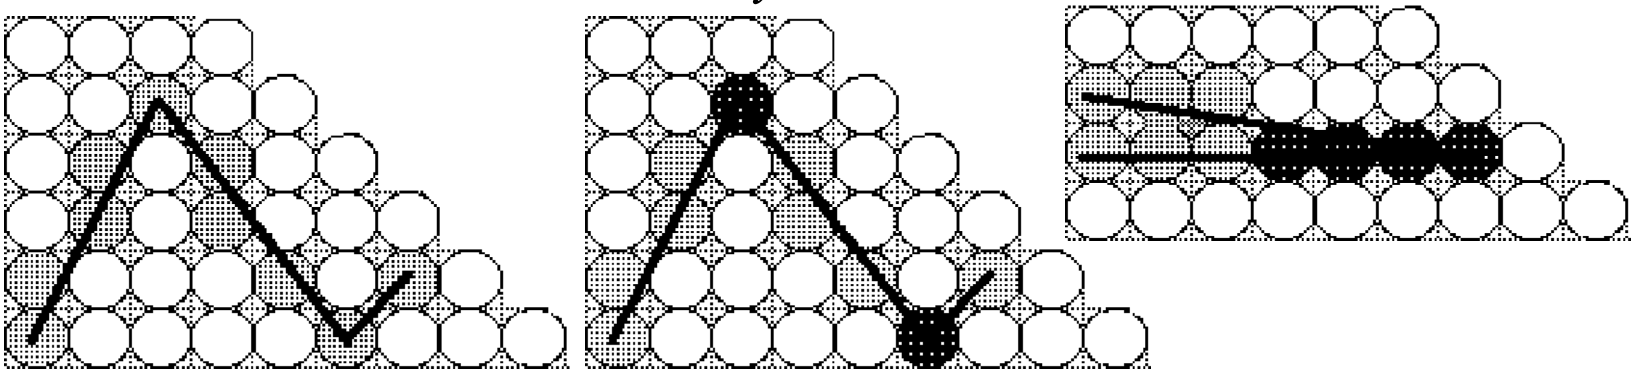

The Rendering Process

Raserization: Scan Conversion

- Scan converting thick objects...

  ...depends on how the ends are drawn:

  ...depends on how a wide line is "trimmed"

  ...if trim lines are perpendicular to the original line, the result is a rectangular thick line; if not the ends are slanted

  ...slanted ends can be joined to make thick polylines (*called mitering)*

The Rendering Process

Raserization: Scan Conversion

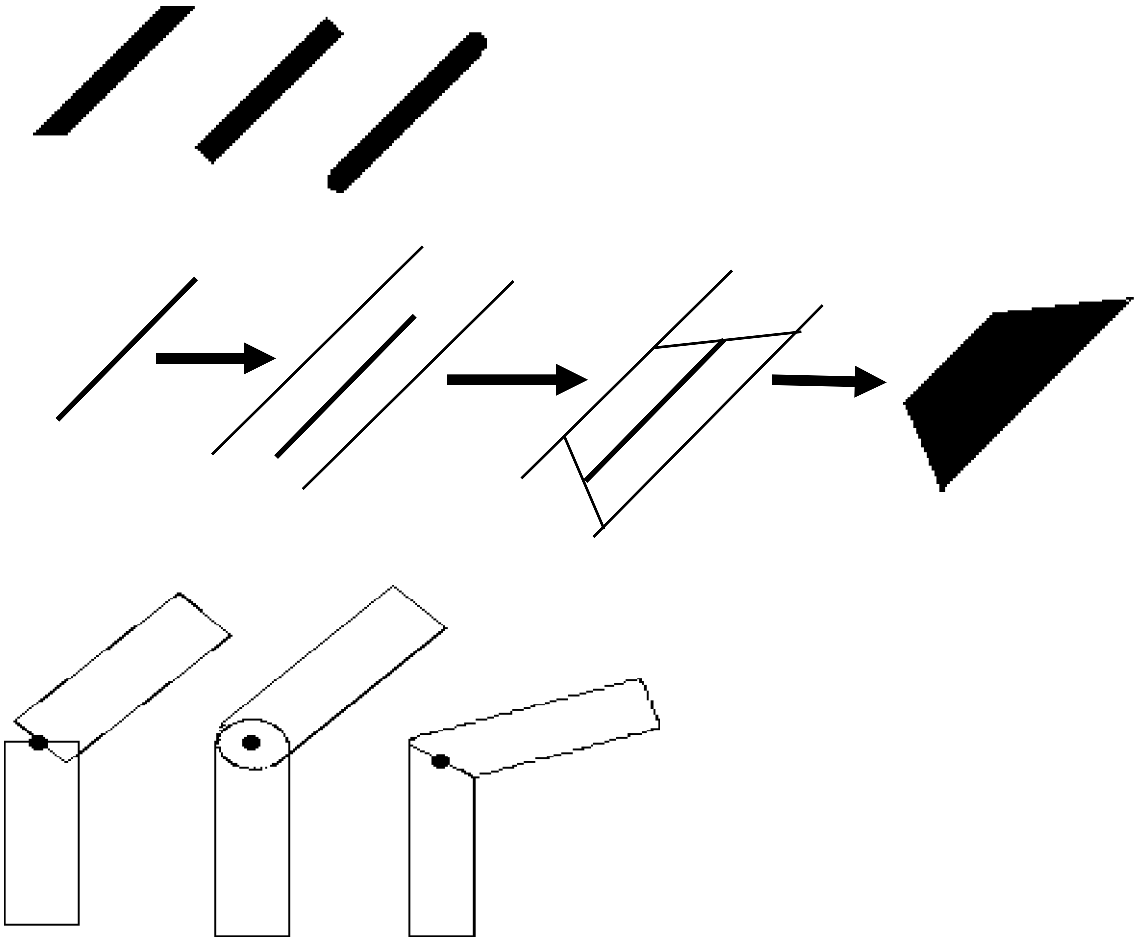

The Rendering Process
Raserization: Scan Conversion

- Scan converting thick curves...

  ...can be done either using multiple line segments, which may produce a very jagged circle unless the segments are short enough and the resolution high enough, or

  ...can be drawn using a "circular pen":

  ...first scan convert the curve into a list of pixels that are adjacent diagonally, vertically, or horizontally (*8-way stepping)*

  ...expand this list of pixels so that any two pixels are adjacent horizontally or vertically *(4-way stepping)*

  ...start by creating a filled circle (i.e., a disk)

  ...for each subsequent pixel, a half-circle (unfilled) is drawn which is centered at the new pixel with its diameter perpendicular to the line from the previous pixel to the current one

The Rendering Process

Raserization: Scan Conversion

- • Scan converting thick curves...

  ...can create gaps when using the circular pen approach without the 4-step method:

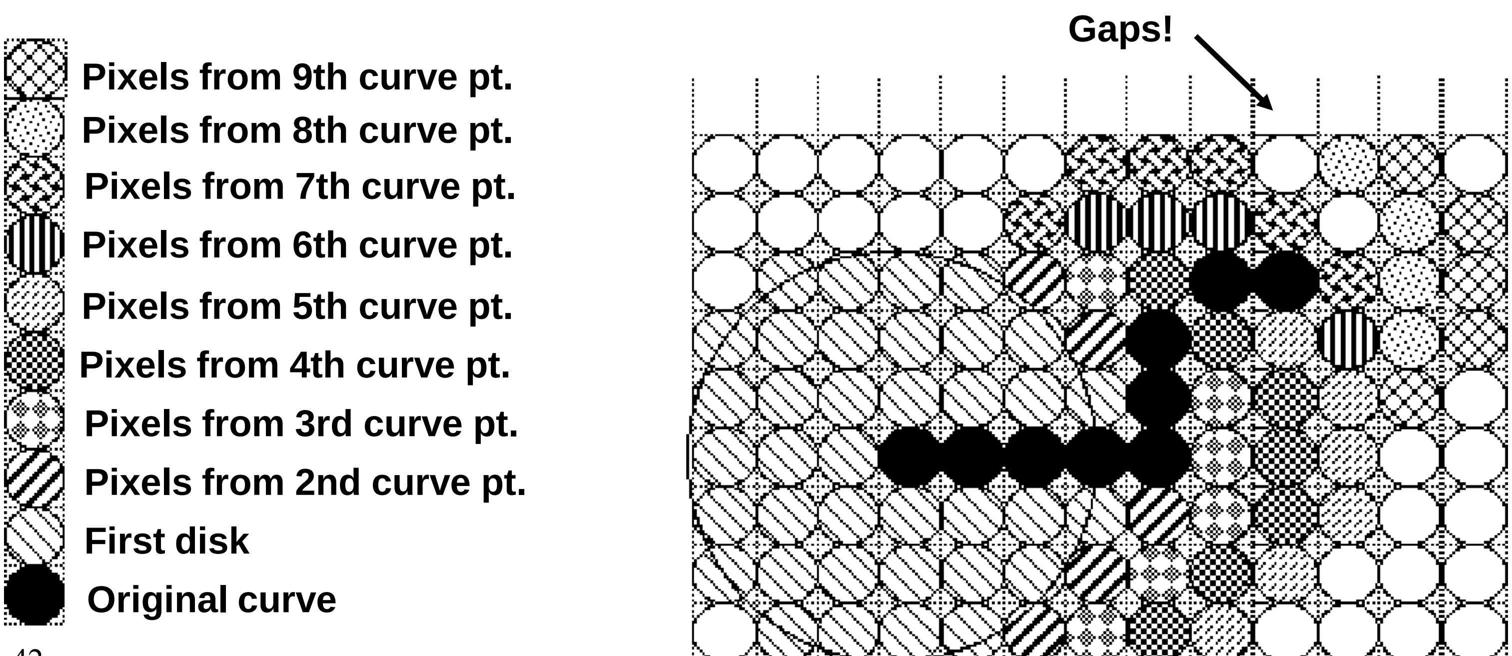

The Rendering Process

Raserization: Scan Conversion

- Scan converting filled areas with the...

  ...even-odd or parity rule: a line is drawn from a point to some other point distant from the area; if this line crosses an edge an odd number of times, the point is inside; if not, the point is outside

- *(note, the even-odd method should not pass the test line through vertices or results could be ambiguous!)*

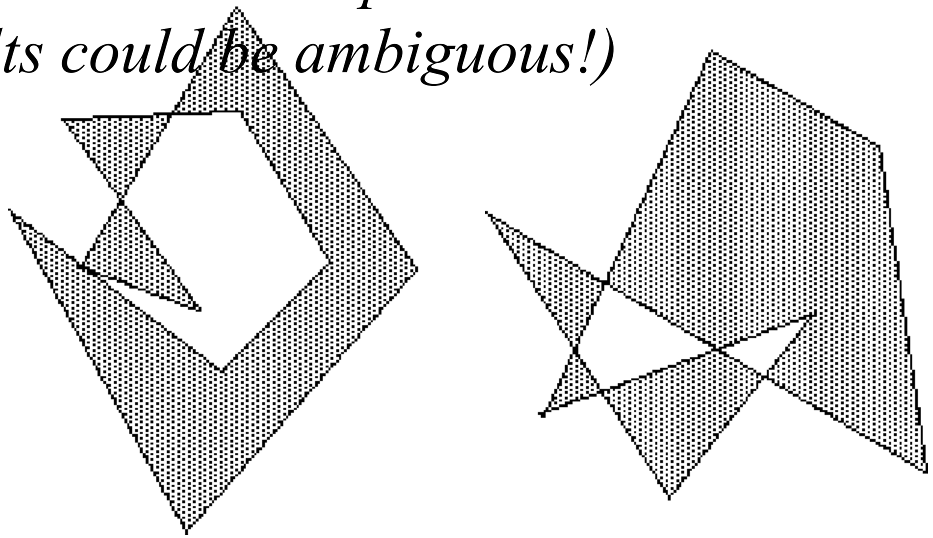

The Rendering Process

Raserization: Scan Conversion

- Scan converting filled areas with the... ...nonexterior rule: which uses a point distant from the area where any point that can be connected to this point by a path that does not intersect an edge is outside; the resulting interior region is everything else
- *(note, everything that can be filled inside the bounding lines is filled!)*

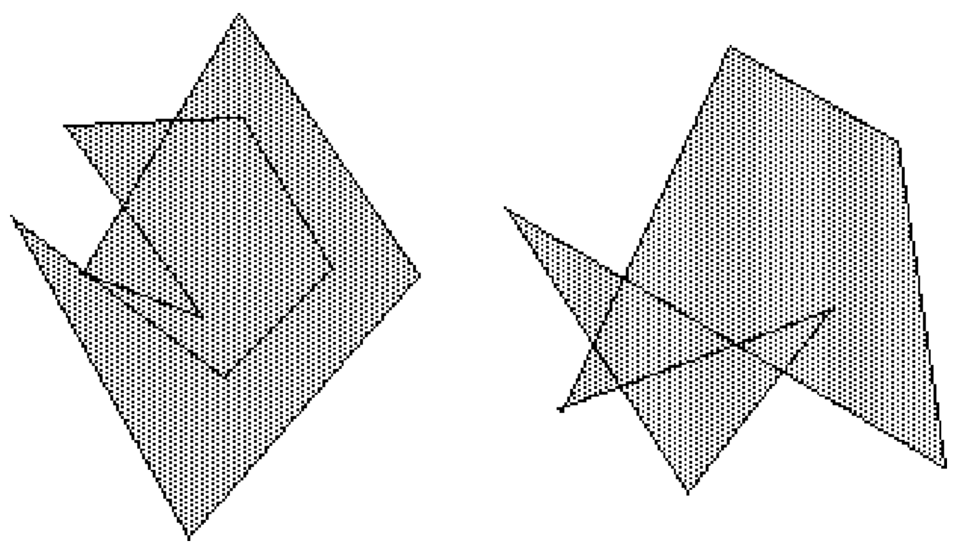

The Rendering Process
Raserization: Scan Conversion

- Scan converting filled areas with the...

  ...nonzero winding rule: a) draw a line from a point to infinity in any direction b) begin with a count of zero c) +1 if the line crosses a left to right edge; -1 if the line crosses a right to left edge d) zero means the point is outside; otherwise it is inside

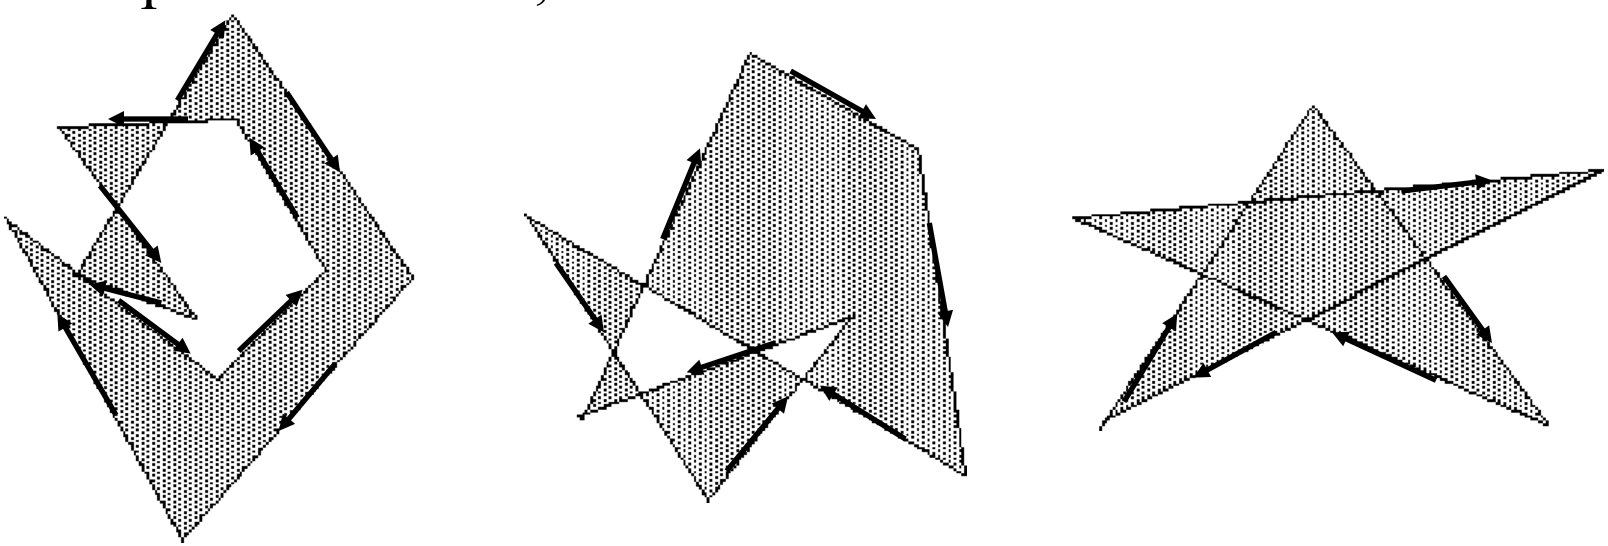

Realistic Rasterization
Antialiasing

- Scan conversion takes each pixel and either replaces it with the color of the graphic object or leaves it unchanged

  ...this process creates drawings with jagged edges

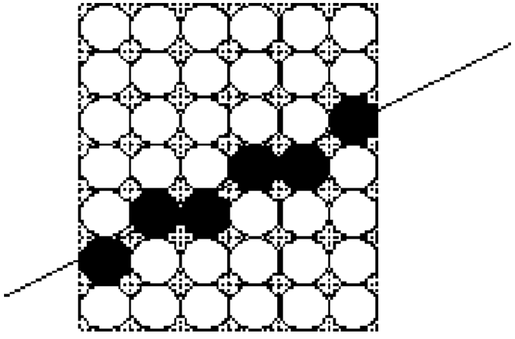

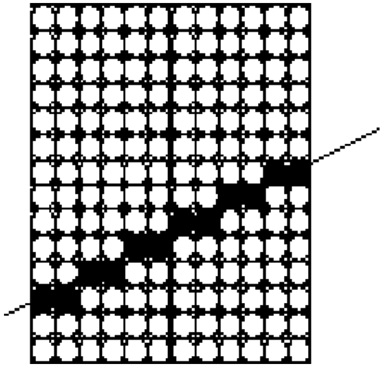

Realistic Rasterization

Antialiasing

- Jaggies are caused by *aliasing*

  ...*antialiasing* techniques reduce or eliminate aliasing

  ...increasing the resolution improves the picture, but is an expensive solution that only diminishes the problem!

Realistic Rasterization

Antialiasing

- Sampling is the process of choosing a color to represent the part of an image covering a pixel...
  - ...for example, sampling determines the color of the light at the center of a pixel
  - ...is closely related to the raster device's array of pixels

Realistic Rasterization
Antialiasing

- Aliasing is the fluctuation of color in adjacent pixels...

  ...created by fine detail that is too small to be resolved, depending on which part falls on the center of the pixel

  ...reduces realistic rendering

  ...causes staircasing or jaggies along edges

- Goal: completely integrate into each pixel all of the image that covers that pixel

Realistic Rasterization

Antialiasing

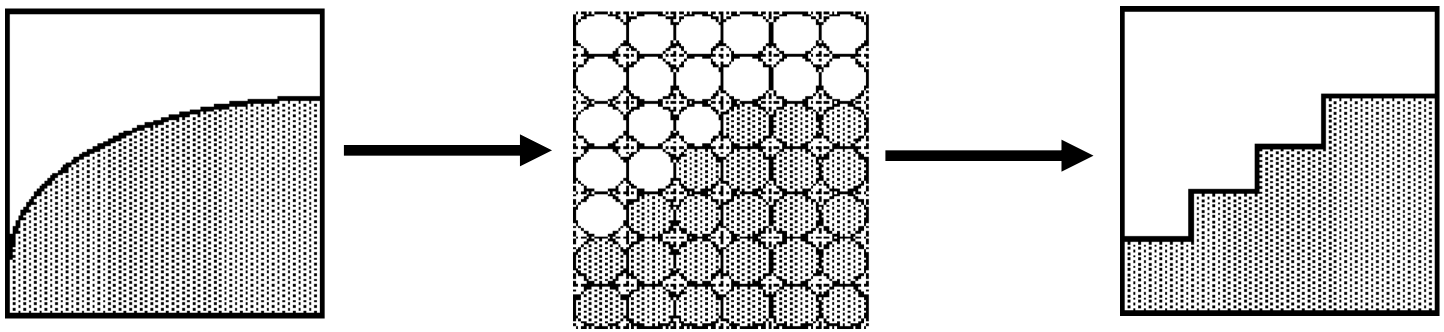

2D silhouette edge ***Pixel array and sample grid*** ***Sampled edge***

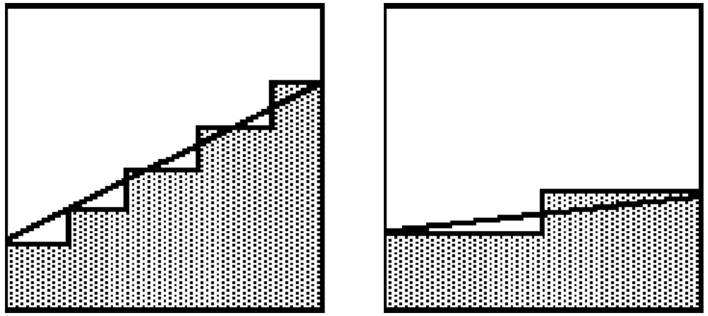

Jagged edges change depending on the line orientation

Realistic Rasterization

Antialiasing

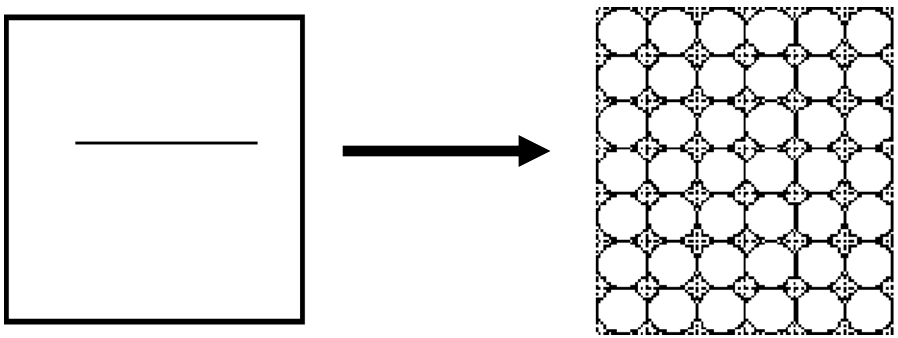

Long thin objects may completely disappear with aliasing

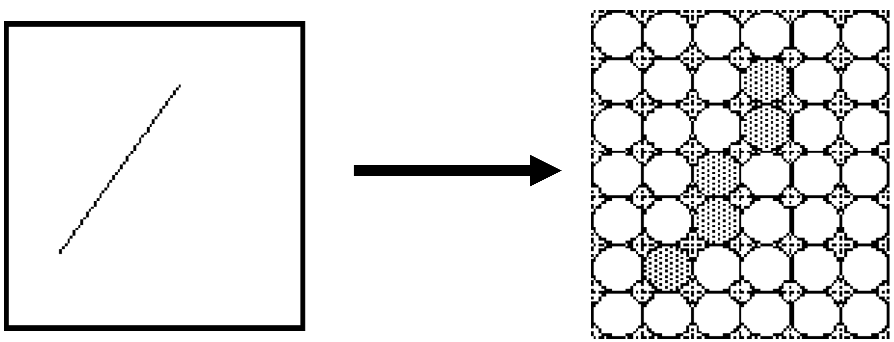

Or, they may break up unpredictably

Realistic Rasterization

Antialiasing

- Think of infinitesimally thin slices through images, representing the intensity along a single horizontal line

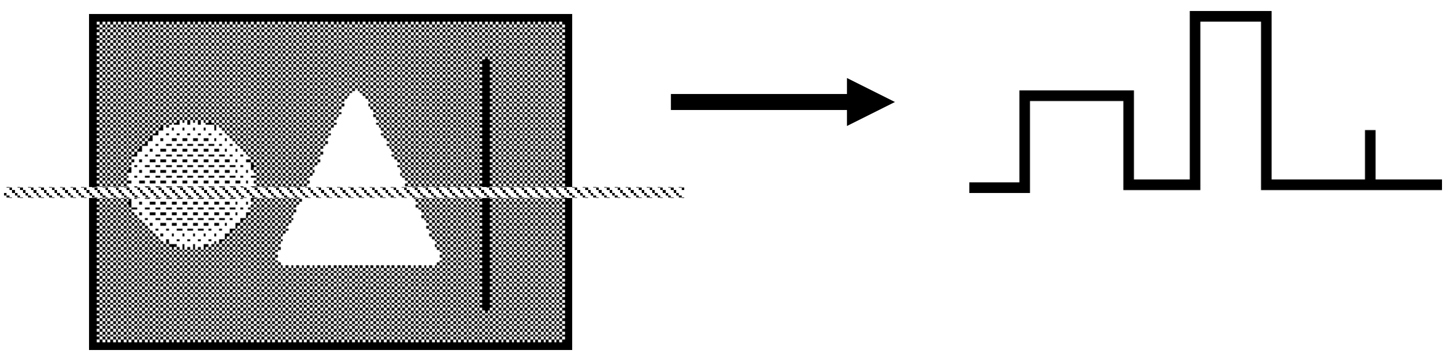

Realistic Rasterization

Antialiasing

- Continuous signals...

  ...are defined at a continuum of positions in space

  ...are generated by projecting a 3D object onto a view plane before scan conversion (creating a continuous 2D signal whose value at each point indicates an intensity)

Realistic Rasterization

Antialiasing

- Discrete signals...

  ...are defined as a set of discrete points in space

  ...are represented by a raster device's array of pixel values (creating a discrete 2D signal whose values are defined only at the positions in the array)

- The way in which values are mapped from continuous to discrete signals is based on either...

  ...*point sampling* or ...*area sampling*

Realistic Rasterization

Antialiasing -- Point Sampling

- Point sampling...

  ...is the most straightforward approach to select each pixel's value

  ...selects one point for each pixel, evaluates the original signal at this point, and assigns its value to the pixel

Realistic Rasterization

Antialiasing -- Point Sampling

- Key points to remember...

  ...projected vertices are not constrained to lie on integer grid points -- unlike scan conversion algorithms

  ...because only a finite set of points are sampled, important features of the signal may be missed (where the frequency of the signal is greanter than the frequency of the pixels)

  ...the more samples collected for a signal the more we will know about it.

  ...by sampling at a higher rate, more details can be shown

Realistic Rasterization

Antialiasing -- Point Sampling

- Sampling...

  ...relates the image to the resolution of the sampling grid...actually relates the spatial frequencies of the image to the resolution of the sampling grid

  ...the more complex and detailed the image is, the finer the sampling grid has to be in order to capture this detail

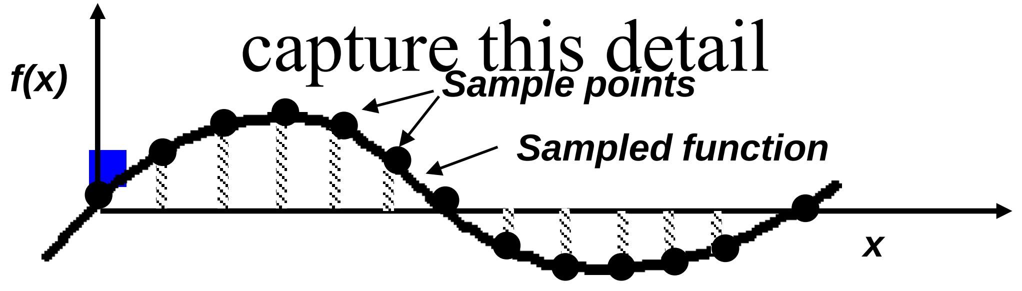

Space domain is represented as a sampling of a sine wave, where the interval is less than 1/2 the period of the sine wave... no information is lost!

Realistic Rasterization
Antialiasing -- Point Sampling

- **Sampling examples...** ***(using the sampling theorem)***

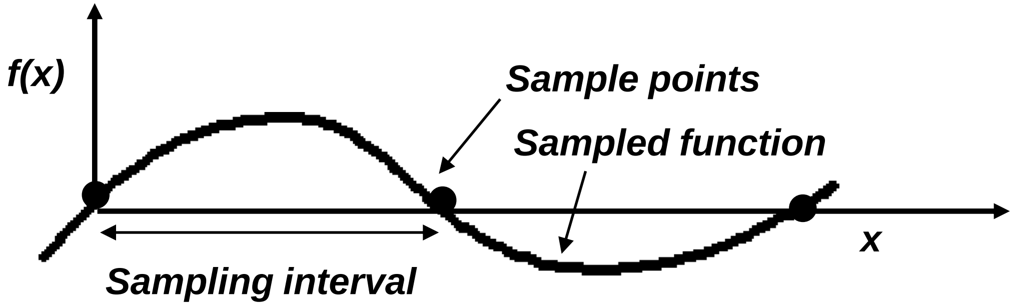

The sampling interval is equal to 1/2 the period of the sine wave... no information can be recovered from the samples concerning the sine wave

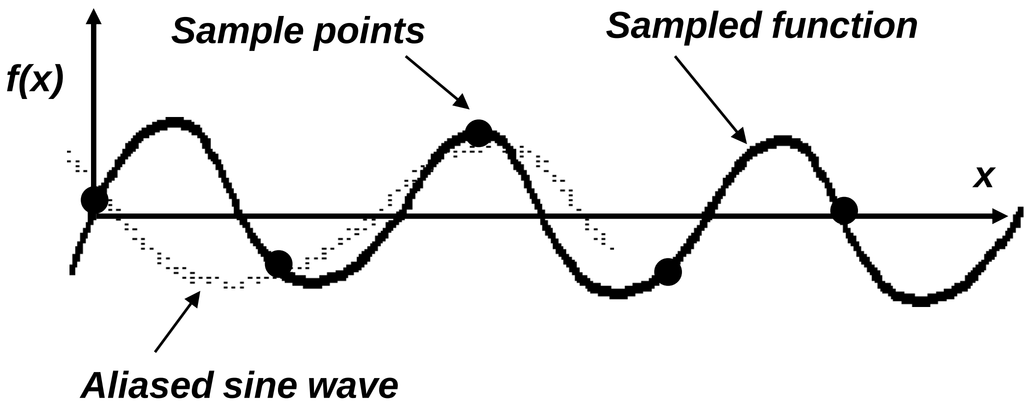

The sampling interval is greater than 1/2 the period of the sine wave... then sampled values are obtained at a lower frequency than the function being sampled. Think of this as two different waves out of sync; this causes aliasing!

Realistic Rasterization

Antialiasing -- Area Sampling

- Supersampling...

  ...is an easy approach and often achieves good results

  ...reduces aliasing...but does not eliminate it

  ...takes more than one sample for each pixel and combines them (e.g., by averaging)

  ...filters (e.g., weights) the sampled pixel values

  ...enables features present in a large image to contribute...they aren't just ignored!

Realistic Rasterization

Antialiasing -- Area Sampling

- A simple approach is unweighted area sampling...

  ...an object contributes to each pixel's intensity an amount proportional to the percentage of the pixel's tile it covers

  ...but objects do not affect pixels which they do not intersect and equal areas contribute equal intensity, regardless of the distance from the pixel's center

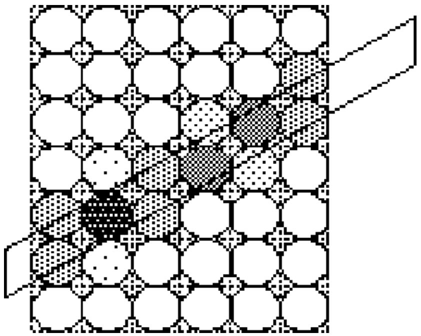

Realistic Rasterization

Antialiasing -- Area Sampling

- A more common approach is weighted area sampling...
  - ...is similar to the unweighted approach
  - ...decreases the intensity as the area decreases
  - ...allows areas to contribute unequally....a small area closer to the pixel center has greater influence than does one at a greater distance
  - ...uses a weighting function to determine the influence on the intensity of a pixel
  - ...such a weighting function is called a *filter*

Realistic Rasterization

Antialiasing -- Area Sampling

- •In more technical terms, filtering supersamples...

  ...creates a new signal by removing offending high frequencies ...

  called low-pass filtering: high frequencies are filtered out and only low frequencies are allowed to pass

  ...causes blurring since fine visual detail captured in the high frequencies is attenuated by low-pass filtering

Realistic Rasterization

Antialiasing -- Area Sampling

- Filtering supersamples allows...

  ...the final output value of a pixel to be a weighted average of supersamples

  ...the farther a sample is from the center, the less its weight

  ...a sample width to have x and y displacements

  ...the sample width to be greater than the size of a pixel...which means samples that go into computing a pixel's value can actually come from several different samples

Realistic Rasterization

Antialiasing -- Area Sampling

- NOTE: A side effect of filtering is blurring...

  ...which occurs because information is integrated from a number of neighboring pixels

  ...which means that the filter's extent is a compromise *(wider filters will be better at reducing aliasing artefacts but will blur the images more than a narrower filter which will be more expensive to implement)*

Realistic Rasterization

Antialiasing -- Area Sampling

- • A common approach is to weigh samples according to their distance from the pixel center:

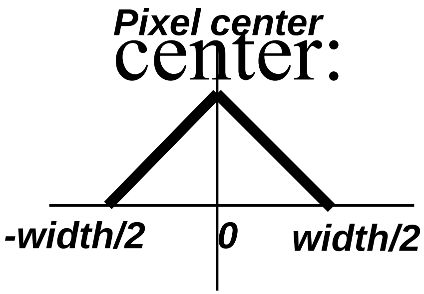

Triangle filter function

Samples affect a pixel from a two-dimensional area whose extent is xwidth and ywidth; a sample must be within this area to be used to calculate the pixel's value

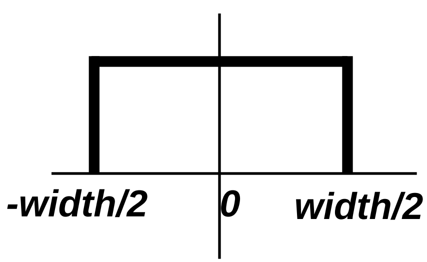

Box filter function

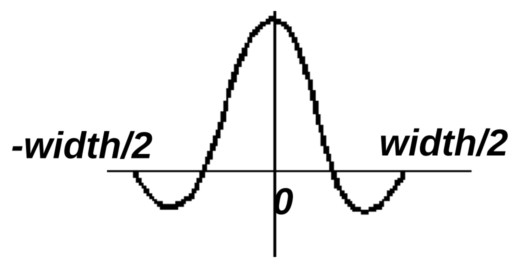

Catmull-Rom filter function

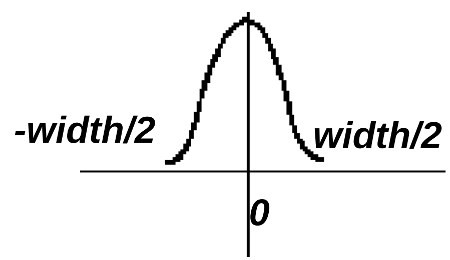

Gaussian filter function

Realistic Rasterization

Antialiasing -- Area Sampling

- • For example, another approach is to use Barlett windows...

1	2	1
2	4	2
1	2	1

1	2	3	2	1
2	4	6	4	2
3	6	9	6	3
2	4	6	4	2
1	2	3	2	1

3x3 and 5x5 Barlett windows

Using a 3x3 window means that nine supersamples are involved in the final pixel computation; using a 7x7 window means 49 integer multiplications... think about the implications!

1	2	3	4	3	2	1
2	4	6	8	6	4	2
3	6	9	12	9	6	3
4	8	12	16	12	8	4
3	6	9	12	9	6	3
2	4	6	8	6	4	2
1	2	3	4	3	2	1

7x7 Barlett window

Realistic Rasterization

Antialiasing Objects

- To antialias graphic objects...

  ...we have to drawn them with fuzzy edges

  ...we need to apply filtering since signals with high-frequencies generate aliases when sampled

  ...this filter is generally: $sinc(x)=sin(\pi x)/\pi x$

Realistic Rasterization

Antialiasing Objects

- For lines, for each pixel near a line, the distance to the line is used to determine the brightness of the pixel

  ...in efficient implementations, this distance is converted to an integer value between 0 and some small number (like 16) ===> this number is used as an index into a table of gray-scale values *(for example)*

  ...however, this does not address overlapping lines, causing the pixels where lines intersect to be overly bright

  ...in this case, a subdivision approach determines disjoint uncovered regions of each pixel and the areas of each region are combined and the resulting color assigned