


Rasterization



MIT EECS 6.837
Frédo Durand and Barb Cutler

MIT EECS 6.837, Cutler and Durand

1

Last time?

- Point and segment Clipping
- Planes as homogenous vectors (duality)
- In homogeneous coordinates before division
- Outcodes for efficient rejection
- Notion of convexity
- Polygon clipping via walking
- Line rasterization, incremental computation

MIT EECS 6.837, Cutler and Durand

2

High-level concepts for 6.837

- Linearity
- Homogeneous coordinates
- Convexity
- Discrete vs. continuous
- Incremental computation

- Trying things on simple examples

MIT EECS 6.837, Cutler and Durand

3

Scan Conversion (Rasterization)

Modeling Transformations

Illumination (Shading)

Viewing Transformation (Perspective / Orthographic)

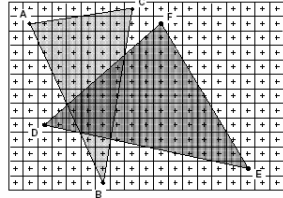
Clipping

Projection (to Screen Space)

Scan Conversion (Rasterization)

Visibility / Display

- Rasterizes objects into pixels
- Interpolate values as we go (color, depth, etc.)



MIT EECS 6.837, Cutler and Durand

4

Visibility / Display

Modeling Transformations

Illumination (Shading)

Viewing Transformation (Perspective / Orthographic)

Clipping

Projection (to Screen Space)

Scan Conversion (Rasterization)

Visibility / Display

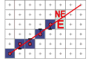
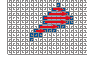
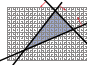
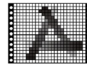
- Each pixel remembers the closest object (depth buffer)

- Almost every step in the graphics pipeline involves a change of coordinate system. Transformations are central to understanding 3D computer graphics.

MIT EECS 6.837, Cutler and Durand

5

Today


- Line scan-conversion 
- Polygon scan conversion
 - smart 
 - back to brute force 
- Visibility 

MIT EECS 6.837, Cutler and Durand

6

Scan Converting 2D Line Segments

- Given:
 - Segment endpoints (integers $x_1, y_1; x_2, y_2$)
- Identify:
 - Set of pixels (x, y) to display for segment




MIT EECS 6.837, Cutler and Durand

7

Line Rasterization Requirements

- Transform **continuous** primitive into **discrete** samples
- Uniform thickness & brightness
- Continuous appearance
- No gaps
- Accuracy
- Speed




MIT EECS 6.837, Cutler and Durand

8

Algorithm Design Choices

- Assume:
 - $m = dy/dx, 0 < m < 1$
- Exactly one pixel per column
 - fewer \rightarrow disconnected, more \rightarrow too thick




MIT EECS 6.837, Cutler and Durand

9

Naive Line Rasterization Algorithm

- Simply compute y as a function of x
 - Conceptually: move vertical scan line from x_1 to x_2
 - What is the expression of y as function of x ?
 - Set pixel $(x, \text{round}(y(x)))$



$$y = y_1 + \frac{x - x_1}{x_2 - x_1}(y_2 - y_1)$$

$$= y_1 + m(x - x_1)$$


$$m = \frac{dy}{dx}$$

MIT EECS 6.837, Cutler and Durand

10

Efficiency

- Computing y value is expensive
 - $y = y_1 + m(x - x_1)$
- Observe: $y \pm m$ at each x step ($m = dy/dx$)




MIT EECS 6.837, Cutler and Durand

11

Bresenham's Algorithm (DDA)

- Select pixel vertically closest to line segment
 - intuitive, efficient,
 - pixel center always within 0.5 vertically
- Same answer as naive approach




MIT EECS 6.837, Cutler and Durand

12

Bresenham's Algorithm (DDA)

- Observation:
 - If we're at pixel P (x_p, y_p), the next pixel must be either E (x_p+1, y_p) or NE (x_p+1, y_p+1)
 - Why?




MIT EECS 6.837, Cutler and Durand

13

Bresenham Step

- Which pixel to choose: E or NE?
 - Choose E if segment passes below or through middle point M
 - Choose NE if segment passes above M



MIT EECS 6.837, Cutler and Durand

14


Bresenham Step

- Use *decision function* D to identify points underlying line L:

$$D(x, y) = y - mx - b$$

- positive above L
- zero on L
- negative below L

$$D(p_x, p_y) = \text{vertical distance from point to line}$$



MIT EECS 6.837, Cutler and Durand

15

Bresenham's Algorithm (DDA)

- Decision Function:

$$D(x, y) = y - mx - b$$

- Initialize:

$$\text{error term } e = -D(x, y)$$


- On each iteration:

$$\text{update } x: \quad x' = x + 1$$

$$\text{update } e: \quad e' = e + m$$

$$\text{if } (e \leq 0.5): \quad y' = y \text{ (choose pixel E)}$$

$$\text{if } (e > 0.5): \quad y' = y + 1 \text{ (choose pixel NE)} \quad e' = e - 1$$




MIT EECS 6.837, Cutler and Durand

16

Summary of Bresenham

- initialize x, y, e
- for ($x = x1; x \leq x2; x++$)
 - plot (x, y)
 - update x, y, e




- Generalize to handle all eight octants using symmetry
- Can be modified to use only integer arithmetic

MIT EECS 6.837, Cutler and Durand

17

Line Rasterization


- We will use it for ray-casting acceleration
- March a ray through a grid



MIT EECS 6.837, Cutler and Durand

18

Line Rasterization vs. Grid Marching



Line Rasterization:
Best discrete approximation of the line

Ray Acceleration:
Must examine every cell the line touches

MIT EECS 6.837, Cutler and Durand

19


Questions?

MIT EECS 6.837, Cutler and Durand

20

Circle Rasterization

- Generate pixels for 2nd octant only
- Slope progresses from 0 \rightarrow -1
- Analog of Bresenham Segment Algorithm



MIT EECS 6.837, Cutler and Durand

21

Circle Rasterization

- Decision Function:

$$D(x, y) = x^2 + y^2 - R^2$$

- Initialize:

$$\text{error term } e = -D(x, y)$$


- On each iteration:

$$\text{update } x: \quad x' = x + 1$$

$$\text{update } e: \quad e' = e + 2x + 1$$

$$\text{if } (e \geq 0.5): \quad y' = y \quad (\text{choose pixel E})$$

$$\text{if } (e < 0.5): \quad y' = y - 1 \quad (\text{choose pixel SE}), \quad e' = e + 1$$



MIT EECS 6.837, Cutler and Durand

22

Philosophically

Discrete differential analyzer (DDA):

- Perform incremental computation
- Work on derivative rather than function
- Gain one order for polynomial
 - Line becomes constant derivative
 - Circle becomes linear derivative

MIT EECS 6.837, Cutler and Durand

23


Questions?

MIT EECS 6.837, Cutler and Durand

24

Antialiased Line Rasterization

- Use gray scales to avoid jaggies
- Will be studied later in the course




MIT EECS 6.837, Cutler and Durand

25

Today

- Line scan-conversion
 - smart
 - back to brute force
- Polygon scan conversion
- Visibility




MIT EECS 6.837, Cutler and Durand

26

2D Scan Conversion

- Geometric primitive
 - 2D: point, line, polygon, circle...
 - 3D: point, line, polyhedron, sphere...
- Primitives are continuous; screen is discrete




MIT EECS 6.837, Cutler and Durand

27

2D Scan Conversion

- Solution: compute discrete approximation
- Scan Conversion: algorithms for efficient generation of the samples comprising this approximation




MIT EECS 6.837, Cutler and Durand

28

Brute force solution for triangles

- For each pixel
 - Compute line equations at pixel center
 - “clip” against the triangle




MIT EECS 6.837, Cutler and Durand

29

Brute force solution for triangles

- For each pixel
 - Compute line equations at pixel center
 - “clip” against the triangle




MIT EECS 6.837, Cutler and Durand

30

Brute force solution for triangles

- Improvement: Compute only for the *screen bounding box* of the triangle
- How do we get such a bounding box?
 - Xmin, Xmax, Ymin, Ymax of the triangle vertices




MIT EECS 6.837, Cutler and Durand

31

Can we do better? Kind of!

- We compute the line equation for many useless pixels
- What could we do?




MIT EECS 6.837, Cutler and Durand

32

Use line rasterization

- Compute the boundary pixels




Shirley page 55

MIT EECS 6.837, Cutler and Durand

33

Scan-line Rasterization

- Compute the boundary pixels
- Fill the spans




Shirley page 55

MIT EECS 6.837, Cutler and Durand

34

Scan-line Rasterization

- Requires some initial setup to prepare



Shirley page 55

MIT EECS 6.837, Cutler and Durand

35

Today

- Line scan-conversion
 - smart
- Polygon scan conversion
 - back to brute force
- Visibility




MIT EECS 6.837, Cutler and Durand

36

For modern graphics cards

- Triangles are usually very small
- Setup cost are becoming more troublesome
- Clipping is annoying
- Brute force is tractable




MIT EECS 6.837, Cutler and Durand

37

Modern rasterization

```

For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Compute line equations
    If all line equations>0 //pixel [x,y] in triangle
      Framebuffer[x,y]=triangleColor
    
```



MIT EECS 6.837, Cutler and Durand


38

Modern rasterization

```

For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Compute line equations
    If all line equations>0 //pixel [x,y] in triangle
      Framebuffer[x,y]=triangleColor
  
```

- Note that Bbox clipping is trivial




MIT EECS 6.837, Cutler and Durand

39

Can we do better?

```

For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Compute line equations
    If all line equations>0 //pixel [x,y] in triangle
      Framebuffer[x,y]=triangleColor
  
```



MIT EECS 6.837, Cutler and Durand


40

Can we do better?

```

For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Compute line equations ax+by+c
    If all line equations>0 //pixel [x,y] in triangle
      Framebuffer[x,y]=triangleColor
  
```

- We don't need to recompute line equation from scratch



MIT EECS 6.837, Cutler and Durand

41

Can we do better?

```

For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  Setup line eq
  compute aidx, bidy for the 3 lines
  Initialize line eq, values for bbox corner
  Li=aix0+biy0+ci
  For all scanline y in bbox
    For 3 lines, update Li
    For all x in bbox
      Increment line equations: Li+aidx
      If all Li>0 //pixel [x,y] in triangle
        Framebuffer[x,y]=triangleColor
  
```


- We save one multiplication per pixel

MIT EECS 6.837, Cutler and Durand

42

Adding Gouraud shading

- Interpolate colors of the 3 vertices
- Linear interpolation




MIT EECS 6.837, Cutler and Durand

43

Adding Gouraud shading

- Interpolate colors of the 3 vertices
- Linear interpolation, e.g. for R channel:
 - $R = a_R x + b_R y + c_R$
 - Such that $R[x_0, y_0] = R_0$; $R[x_1, y_1] = R_1$; $R[x_2, y_2] = R_2$
 - Same as a plane equation in (x, y, R)



MIT EECS 6.83

44

Adding Gouraud shading

```


Interpolate colors
For every triangle
  ComputeProjection
  Compute bbox, clip bbox to screen limits
  Setup line eq
  Setup color equation
  For all pixels in bbox
    Increment line equations
    Increment color equation
    If all  $L_i > 0$  //pixel [x,y] in triangle
      Framebuffer[x,y]=interpolatedColor
    
```

MIT EECS 6.837, Cutler and Durand

45

Adding Gouraud shading

- Interpolate colors of the 3 vertices
- Other solution: use barycentric coordinates
- $R = \alpha R_0 + \beta R_1 + \gamma R_2$
- Such that $P = \alpha P_0 + \beta P_1 + \gamma P_2$
-




MIT EECS 6.83

46

In the modern hardware

- Edge eq. in homogeneous coordinates $[x, y, w]$
- Tiles to add a mid-level granularity
 - Early rejection of tiles
 - Memory access coherence



MIT EECS 6.837, Cutler and Durand

47

Ref

- Henry Fuchs, Jack Goldfeather, Jeff Hultquist, Susan Spach, John Austin, Frederick Brooks, Jr., John Eyles and John Poulton, "Fast Spheres, Shadows, Textures, Transparencies, and Image Enhancements in Pixel-Planes", Proceedings of SIGGRAPH '85 (San Francisco, CA, July 22–26, 1985). In *Computer Graphics*, v19n3 (July 1985), ACM SIGGRAPH, New York, NY, 1985.
- Juan Pineda, "A Parallel Algorithm for Polygon Rasterization", Proceedings of SIGGRAPH '88 (Atlanta, GA, August 1–5, 1988). In *Computer Graphics*, v22n4 (August 1988), ACM SIGGRAPH, New York, NY, 1988. Figure 7: Image from the spinning teapot performance test.
- Triangle Scan Conversion using 2D Homogeneous Coordinates, Marc Olano Trey Greer
<http://www.cs.unc.edu/~olano/papers/2dh-tri/2dh-tri.pdf>

MIT EECS 6.837, Cutler and Durand

48

Take-home message

- The appropriate algorithm depends on
 - Balance between various resources (CPU, memory, bandwidth)
 - The input (size of triangles, etc.)
- Smart algorithms often have initial preprocess
 - Assess whether it is worth it
- To save time, identify redundant computation
 - Put outside the loop and interpolate if needed

MIT EECS 6.837, Cutler and Durand

49

Questions?

MIT EECS 6.837, Cutler and Durand

50

Today

- Line scan-conversion
- Polygon scan conversion
 - smart
 - back to brute force
- **Visibility**




MIT EECS 6.837, Cutler and Durand

51

Visibility

- How do we know which parts are visible/in front?




MIT EECS 6.837, Cutler and Durand

52

Ray Casting

- Maintain intersection with closest object




MIT EECS 6.837, Cutler and Durand

53

Painter's algorithm

- Draw back-to-front
- How do we sort objects?
- Can we always sort objects?




MIT EECS 6.837, Cutler and Durand

54

Painter's algorithm

- Draw back-to-front
- How do we sort objects?
- Can we always sort objects?
 - No, there can be cycles
 - Requires to split polygons



MIT EECS 6.837, Cutler and Durand

55

Painter's algorithm


- Old solution for hidden-surface removal
 - Good because ordering is useful for other operations (transparency, antialiasing)
- But
 - Ordering is tough
 - Cycles
 - Must be done by CPU
- Hardly used now
- But some sort of partial ordering is sometimes useful
 - Usual front-to-back
 - To make sure foreground is rendered first
 - For transparency

MIT EECS 6.837, Cutler and Durand

56

Visibility

- In ray casting, use intersection with closest t
- Now we have swapped the loops (pixel, object)
- How do we do?




MIT EECS 6.837, Cutler and Durand

57

Z buffer

- In addition to frame buffer (R, G, B)
- Store distance to camera (z-buffer)
- Pixel is updated only if new z is closer than z-buffer value



MIT EECS 6.837, Cutler and Durand

58

Z-buffer pseudo code


```

For every triangle
  Compute Projection, color at vertices
  Setup line equations
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Increment line equations
    Compute currentZ
    Increment currentColor
    If all line equations > 0 //pixel [x,y] in triangle
      If currentZ < zBuffer[x,y] //pixel is visible
        Framebuffer[x,y] = currentColor
        zBuffer[x,y] = currentZ
    
```

MIT EECS 6.837, Cutler and Durand

59

Works for hard cases!




MIT EECS 6.837, Cutler and Durand

60

What exactly do we store

- Floating point distance
- Can we interpolate z in screen space?
 - i.e. does z vary linearly in screen space?




MIT EECS 6.837, Cutler and Durand

61

Z interpolation

- $X' = x/z$
- Hyperbolic variation
- Z cannot be linearly interpolated



MIT EECS 6.837, Cutler and Durand

62

Simple Perspective Projection

- Project all points along the z axis to the $z = d$ plane, eyepoint at the origin

homogenize

$$\begin{pmatrix} x * d / z \\ y * d / z \\ d \\ 1 \end{pmatrix} = \begin{pmatrix} x \\ y \\ z \\ z / d \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1/d & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

63

Yet another Perspective Projection

- Change the z component
- Compute d/z
- Can be linearly interpolated*

homogenize

$$\begin{pmatrix} x * d / z \\ y * d / z \\ d/z \\ 1 \end{pmatrix} = \begin{pmatrix} x \\ y \\ 1 \\ z / d \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1/d & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

64

Advantages of $1/z$


- Can be interpolated linearly in screen space
- Puts more precision for close objects
- Useful when using integers
 - more precision where perceptible

MIT EECS 6.837, Cutler and Durand

65

Integer z-buffer

- Use $1/z$ to have more precision in the foreground
- Set a near and far plane
 - $1/z$ values linearly encoded between $1/\text{near}$ and $1/\text{far}$
- Careful, test direction is reversed



MIT EECS 6.837, Cutler and Durand

66

Integer Z-buffer pseudo code

```


For every triangle
  Compute Projection, color at vertices
  Setup line equations, depth equation
  Compute bbox, clip bbox to screen limits
  For all pixels in bbox
    Increment line equations
    Increment current_lovZ
    Increment currentColor
    If all line equations > 0 //pixel [x,y] in triangle
      If current_lovZ > lovzBuffer[x,y] //pixel is visible
        Framebuffer[x,y] = currentColor
        lovzBuffer[x,y] = currentlovZ
    
```

MIT EECS 6.837, Cutler and Durand

67

Gouraud interpolation

- Gouraud: interpolate color linearly in screen space
- Is it correct?




MIT EECS 6.837, Cutler and Durand

68

Gouraud interpolation

- Gouraud: interpolate color linearly in screen space
- Not correct. We should use hyperbolic interpolation
- But quite costly (division)
- However, can now be done on modern hardware



MIT EECS 6.837, Cutler and Durand


69

Questions?



The infamous half pixel


- I refuse to teach it, but it's an annoying issue you should know about
- Do a line drawing of a rectangle from [top, right] to [bottom, left]
- Do we actually draw the columns/rows of pixels?



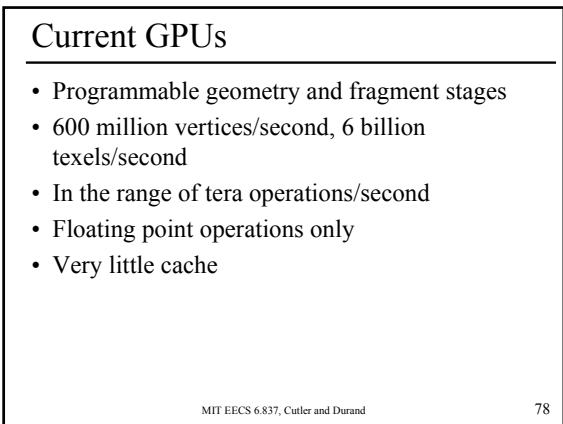
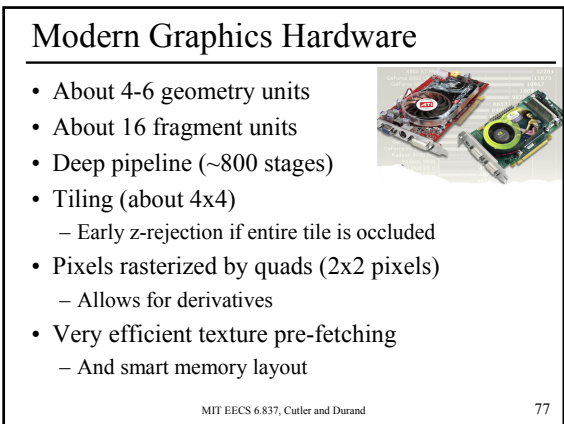
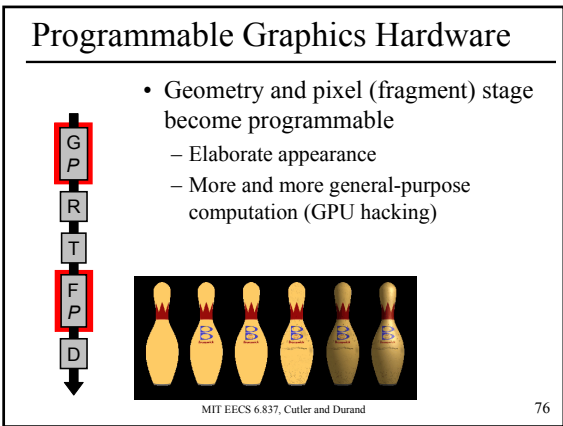
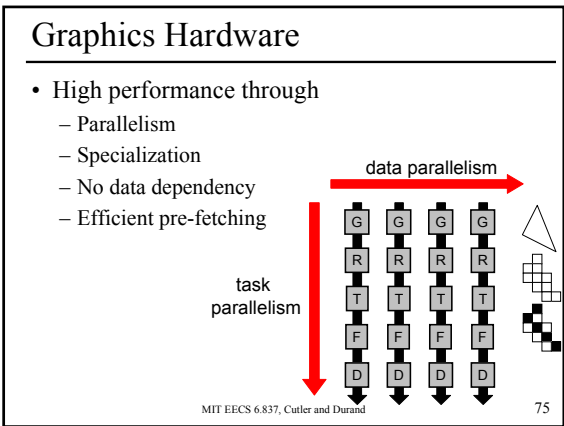
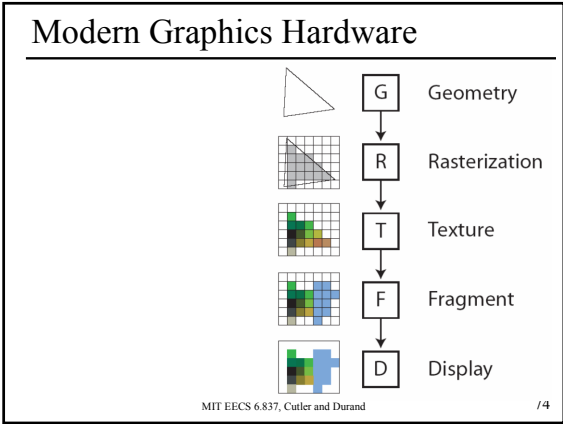
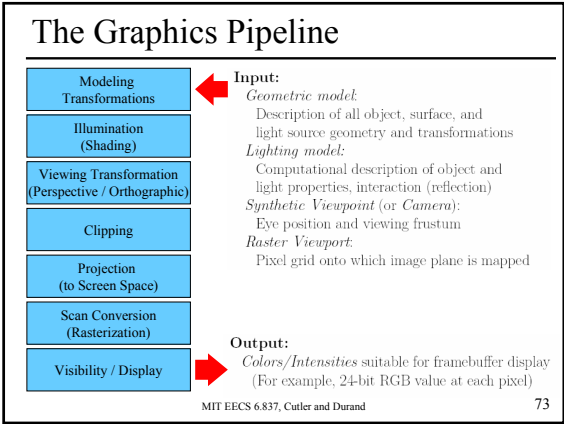
71

The infamous half pixel

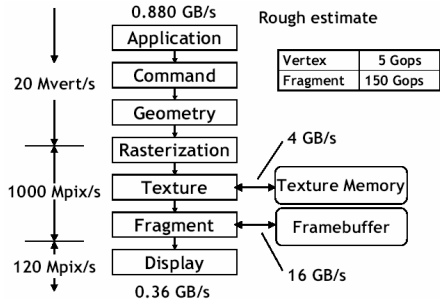
- Displace by half a pixel so that top, right, bottom, left are in the middle of pixels
- Just change the viewport transform



72



Computational Requirements



MIT EECS 6.837, Cutler and Durand

[Akeley, Hanrahan] 79

Questions?



Above:
Bonny, Case, *On Is*, 1990. Appropriated
longe chair, 92 x 89 x 76cm. Private
collection.

Above:
Clyde interacts with his sister's sculpture,
allowing his whole body to become
implicated in its heavily nuanced form.

Above:
Interpretive diagram by Peter MacLore:
1. Tail form.
2. Emphasis on edge.
3. Ovaloid aperture.
4. Restrictive vine form.

The synthetic fiber has been carefully
traced to resemble the texture and color
of a cat's tail in the upright, outstretched
position—twisting, yet guarding the
entrance beyond. However, this constricting
tail is itself compromised by restrictive
vines so that the whole ergonomically edged
aperture hints at pleasure tinged with the
possibility of entanglement.
MacLore, M. *Chaosmorphism*. Exhibition
catalog, Drexel Gallery of Non-Primate
Art, Philadelphia, 1992.

80

Next Week: Ray Casting Acceleration

MIT EECS 6.837, Cutler and Durand

81