I. Rendering

A. There are four things needed to render a virtual scene:
   a. Surface Geometry
   b. Surface Material
   c. Lighting
   d. Camera Model

B. Scan-line Rendering – procedure by which most commercially available software renders a 3 dimensional scene.
   a. Images are made up of tiny units of color called pixels, and pixels are placed orderly in horizontal lines called, scan-lines.
   b. The rendering engine looks at every pixel, one after the other, scan-line by scan-line, and calculates the color that pixel should be rendered.

C. Raycasting – the method by which a rendering program renders a scene. (Maya default renderer)
   a. Ray casting procedure -
      i. From the camera's point of view, typically referred to as the eye," a ray is cast through the first pixel of the scan line.
      ii. Through the tiny pixel sized window it enters the three-dimensional world on the other side of the film gate.
      iii. The eye then follows the ray until it hits the surface of an object or exits the viewing volume.
      iv. Having hit an object, the computer calculates the color of the object, based on the shading model (*Blinn*), lighting, surface material and color.
      v. That resulting color becomes the color of the pixel in your final image.
      vi. Then the rendering program moves one pixel over along the scan-line.
   b. Pro's of Raycasting
      i. Fast rendering algorithm.
      ii. Ability to render very sophisticated scenes.
      iii. Depth sorting is an easy task.
   c. Con's of Raycasting
      i. They deal with surface shading of each object in the scene as if it existed in isolation.
      ii. With depth sorting, any objects deeper in the scene are immediately occluded, and not calculated in the appearance of the final render.
      iii. Reflections are problematic, since reflections reflect other objects onto on object in a scene. (*You must use environment maps to fake the reflections.*)
      iv. Transparency doesn't work with depth-sorted objects. (*Generally, objects are rendered on top of each other, and the colors blended to achieve fake transparency.*)

D. Raytracing – solves many of the problems of traditional ray casting rendering engines.
   a. Procedure:
      i. A ray is shot through a pixel in your film gate and enters the 3 dimensional world.
      ii. When the ray hits a surface, it bounces off in a pre-determined number of directions, gathering color, and surface information from adjacent surfaces.
      iii. The surface color of the original object is then calculated based on the inputs from all other objects in the scene.
      iv. For transparent objects, the ray continues through based on parameters of transparency and refraction. Each surface the ray passes through has rays fracturing top other objects in the scene, determining the color of the surface at that pixel.
      v. After figuring out the surface color and transparency of an object, the corresponding pixel in the final image is rendered.
   b. Pro's of raytracing –
      i. True and accurate reflections on shiny surfaces.
      ii. More accurate color rendering.
      iii. Cleaner shadow casting, with reflected color in dark shadows.
   c. Con's of raytracing
      i. Slow!
      ii. Doesn't handle diffuse reflection of light from one object to another object. (*Must use Radiosity.*)

E. Surface Tessellation – surface parameter that determines your objects smoothness.
   a. What's it all about?
      i. At render time, all Nurbs surfaces are broken down into smaller and smaller triangles (*triangulation*).

ii. Surfaces with fine detail require more triangles, while surfaces with little curvature or detail require fewer triangles.
iii. The surfaces are broken down into triangles because triangles are necessarily flat, a polygon of more then three points might not be flat.
iv. Flat triangular polygons are more easily rendered and require less computation.

b. Tessellation in Maya
i. In order to see your surfaces tessellation, you must be in Smooth Shade mode (*3*)
ii. By default, Maya tessellates on the surface as opposed to at isoparms or surface spans. (This achieves more uniform tessellation with lower triangle counts.
iii. When joining two surfaces together, or creating blends, you should switch to span-based tessellation.
1. Select Explicit Tessellation, and change attributes and test render until you achieve desired results.
2. Joined or blended surfaces have a tendency to crack at the isoparms or edges of the connecting points on the surface, explicit tessellation of the surface will help remove that.
3. Curve degree can help determine Mode U/Mode V.
4. Curve of degree 1, per span isoparms should be set to 1.
5. Curves of greater than 1 degree, should be set based on the following equation: **1.5*((#spans U/V) +1)**
iv. Geometry Anti-aliasing override – when your edge just isn't smooth enough.
1. Set under Render Globals -> Render Stats.
2. Only works if you have High or Highest Quality Anti-aliasing set in Render Globals -> Anti-aliasing Quality.
3. Only set this for surfaces that flicker, it adds time to the render process.

II. Output Considerations
A. Setting your options.
a. Window -> Render Globals: used to determine all settings for render.
b. Camera Attribute Editor – used to set all parameters regarding film gate, focal length, depth of field, shutter angle, and environment maps.
B. Format - What you are going to deliver on?
a. Video - Dependent on your final tape format.
i. NTSC video - 720x486 - 2 field
ii. DV - 720x486 - 2 field
b. Transfer to film dependent on the hardware used. (Cineon, etc.)
i. 16mm - 1416x1062 - Full Cine, no sound track.
ii. 35mm - 4096x3122 – high-resolution 35mm
c. Print - Determined by the print resolution at the service bureau.
i. Pixels in a rendered image are 72 pixels per inch (*Monitor resolution*).
ii. Typical high-resolution print output is 1200 PPI (pixels per inch.)
iii. You do the math. (*This stuff takes a long time*.)
C. Pixel Aspect Ratio - the ratio of width:height of the rendered pixel.
a. Set in Render Globals
b. Old Abekas disk recorders use a rectangular pixel that is taller than it is wide, so the pixel aspect ratio is .9. Images rendered with this pixel aspect ratio appear squashed.
c. Most output renders will have a pixel aspect on 1.0, meaning the pixel on the monitor will match the final rendered pixel.
D. Output Image Format - Set in the Render Globals
a. Set it to TIFF as your default on Irix, or BMP on Windows NT machines.
b. Tiff format will allow you to store an 8 bit alpha channel in the rendered image for later compositing.
c. You must have the Alpha Channel (*mask*) button toggled ON for the alpha to be output.
d. Depth maps may also be output for compositing in applications like Composer or After Effects. *(3DS users can export depth information to be used in Edit and Effect to add effects without re-rendering your images.)*
E. Anti-aliasing - The act of removing the "jaggie" effect around curved or smooth edges that are caused by the rectangular pixels' inability to resolve the image to the screen.
a. Set in Render Globals -> Anti-aliasing Quality.
b. Shading and Max Shading determine
i. The number of super samples per pixel.

ii. Higher Minimum values will seriously slow your render.
iii. For best rendering speed use these settings – Shading 1, Max Shading 8

c. Contrast Threshold - uses the color of adjacent pixels to determine whether finer anti-aliasing is required.
d. Multi-pixel filtering – determines the type of anti-aliasing algorithm, and size used on your image.

III. Speeding it up.

A. Rules of Thumb

a. Reduce the subdivisions on insignificant surfaces, in an Objects Attribute Editor.
b. Reduce the number of super samples in your anti-aliasing parameters. (*1,4*)
c. Remove any excess geometry from the scene.
d. Use lower resolution textures on objects that are small, or which won't be viewed at a close distance.
e. Don't use backdrops. Use an external compositing program to combine your rendered images with your background plate.
f. Convert solid textures to parametric texture maps.
g. Use Spotlight Depth Maps

i. Spotlight depth maps are rendered once, on the first frame of an animation, or the first time the image is rendered.
ii. Subsequent renders will utilize the depth map, skipping the time consuming task of rendering shadows and lights.

h. Reduce the resolution of your shadow maps. Attribute Editor -> Dmap Resolution
i. Don’t use Raytracing.

B. Raycasting speedup.

a. Make as many surfaces single sided as possible. Window -> General Editors -> Attribute Spreadsheet.
b. Leave transparent surfaces double sided.
c. Turn Motion Blur off.
d. Turn off reflection visibility of objects that don’t need to be scene in reflections.

C. Raytracing speedup. Settings in Render Globals -> Raytracing

a. Reduce the max number of reflections for flat mirror like surfaces to one.
b. Reduce the max number of refractions.
c. Reduce the max Shadow level. 1 is a good setting to resolve the shadow in most images.

IV. Rendering your scene.

A. Render -> Save Batch Render – saves a copy of the current scene to a file and launches Render in the background.

a. Background Render on NT is very slow.
b. It is better to kill the render and **RMB+ Render** the file in the UI.

B. Manual Render – preferred method.

a. Set up your Render Globals, position your camera, and save your scene.
b. From Windows NT **RMB+Render** to launch the Render application.
c. In a command shell, or cshell type **Render <your file name here>**.

C. Folder structures

a. Render files are saved in the render folder.
b. Rendered image files are saved in the project’s images directory.
c. Scene files are kept in the project’s scene folder.

D. Render – command line interface to the renderer.

V. Render Flags and Attribute Spreadsheet

A. Render Flags

a. Left panel shows all the elements of your scene, determined by the scope of the show menu.
b. Right panel lists the attributes for the selected elements in the left panel.
c. Use the dialog to set render attributes for selected objects.
d. To change the attributes of more than one element in your scene, **shift+select** multiple elements on the right, and then change the attributes on the left.

B. Attribute Spreadsheet

a. A spreadsheet that displays the attributes for any node in your scene.
b. Much like the Render Flags dialog, it allows you to change attributes for all elements in your scene.

VI. Renderview – used to interactively tune rendering attributes and test render single frames.

- A. Maya gives priority to a render in progress, you will have to wait until it is finished to continue modeling. **ESC** will stop a Renderview test render in progress.
- B. Options
  - a. Render Globals – launches the render global dialog.
  - b. Test Resolution – Determines the resolution of Renderview rendering.
    - i. Render Globals renders uses the render resolution set in Render Globals.
    - ii. Camera panel – renders the image at its largest possible size in the Renderview.
  - **c. RMB+Render -> x** a quick way to test render a scene.
- C. IPR – Interactive Photo realistic Renderer.
  - a. IPR menu – used to setup and run IPR Renderviews.
  - b. IPR -> IPR Render Globals – used to set IPR render attributes.
  - c. IPR -> IPR Render runs an IPR render
    - i. IPR renders take a snapshot of all data in your scene at the time of the IPR Render.
    - ii. It makes a huge file, that tracks depth information and U/V Parameterization.
    - iii. Changes to U/V parameters, camera placement, and movement of geometry are not immediately reflected in the IPR render. You must restart the IPR so that it can pole the new data.
  - d. IPR will automatically update a region, or the entire image when surface properties are changed on shaders in the Hypershade/Visor.
  - e. Very nice interactive tool for troubleshooting surface texturing and shading.
- D. **File -> Keep Image in Render View** – keeps a copy of previous Renderview renders in the window for reference. Images are lost when you close Maya.