

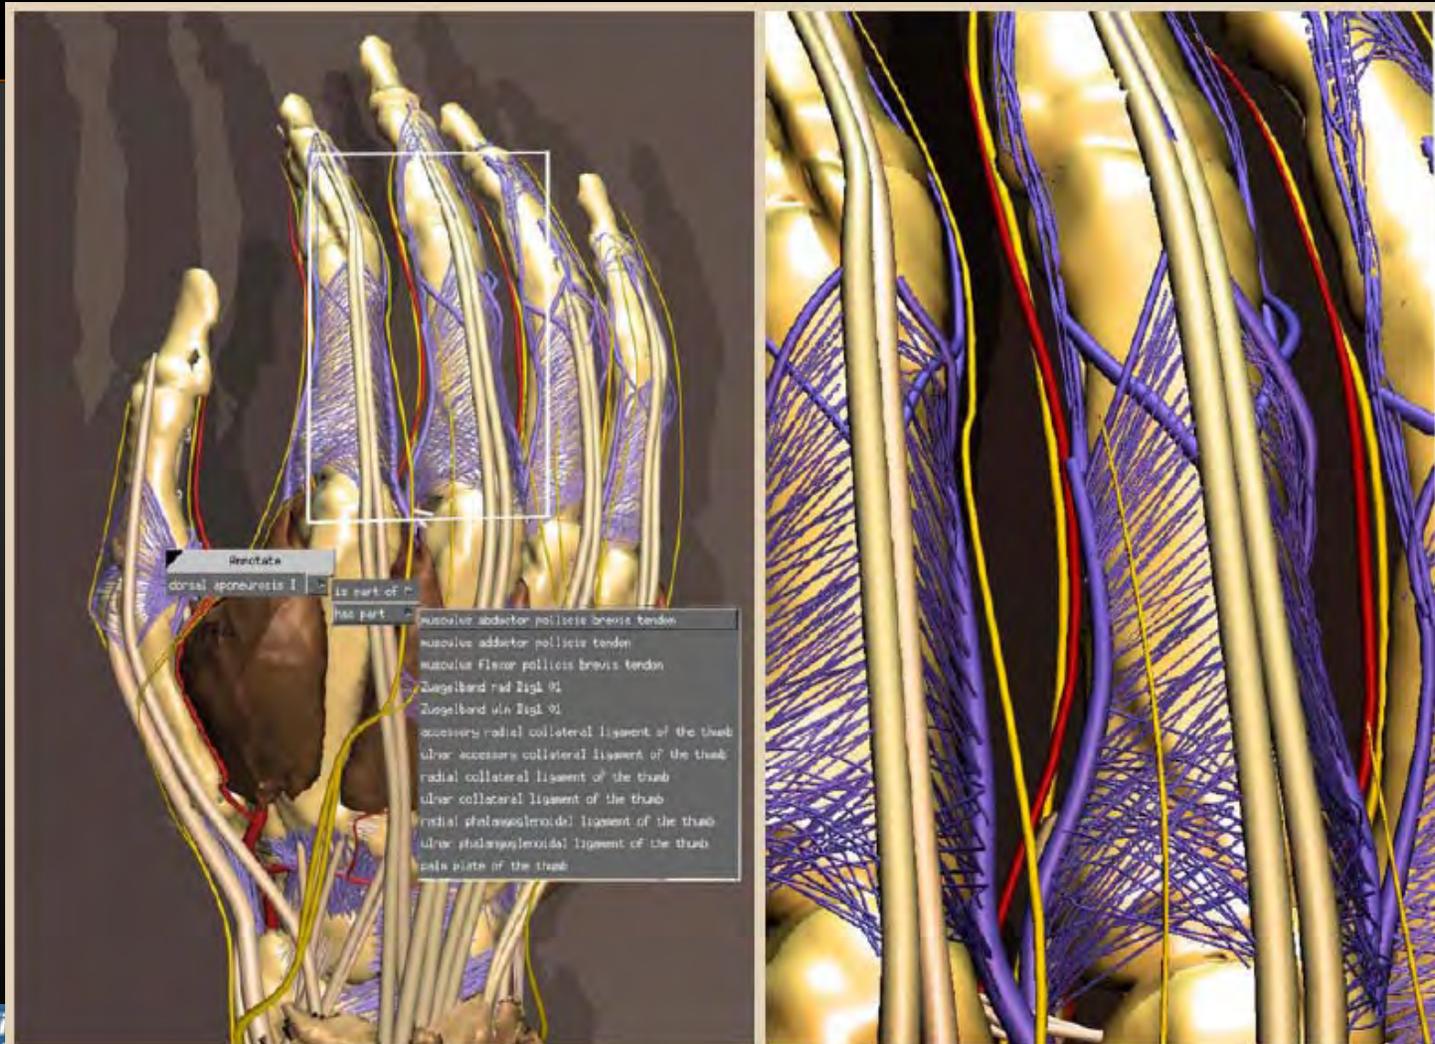
# Scientific Visualization: The Modern Oscilloscope for "Seeing the Unseeable"

E. Wes Bethel

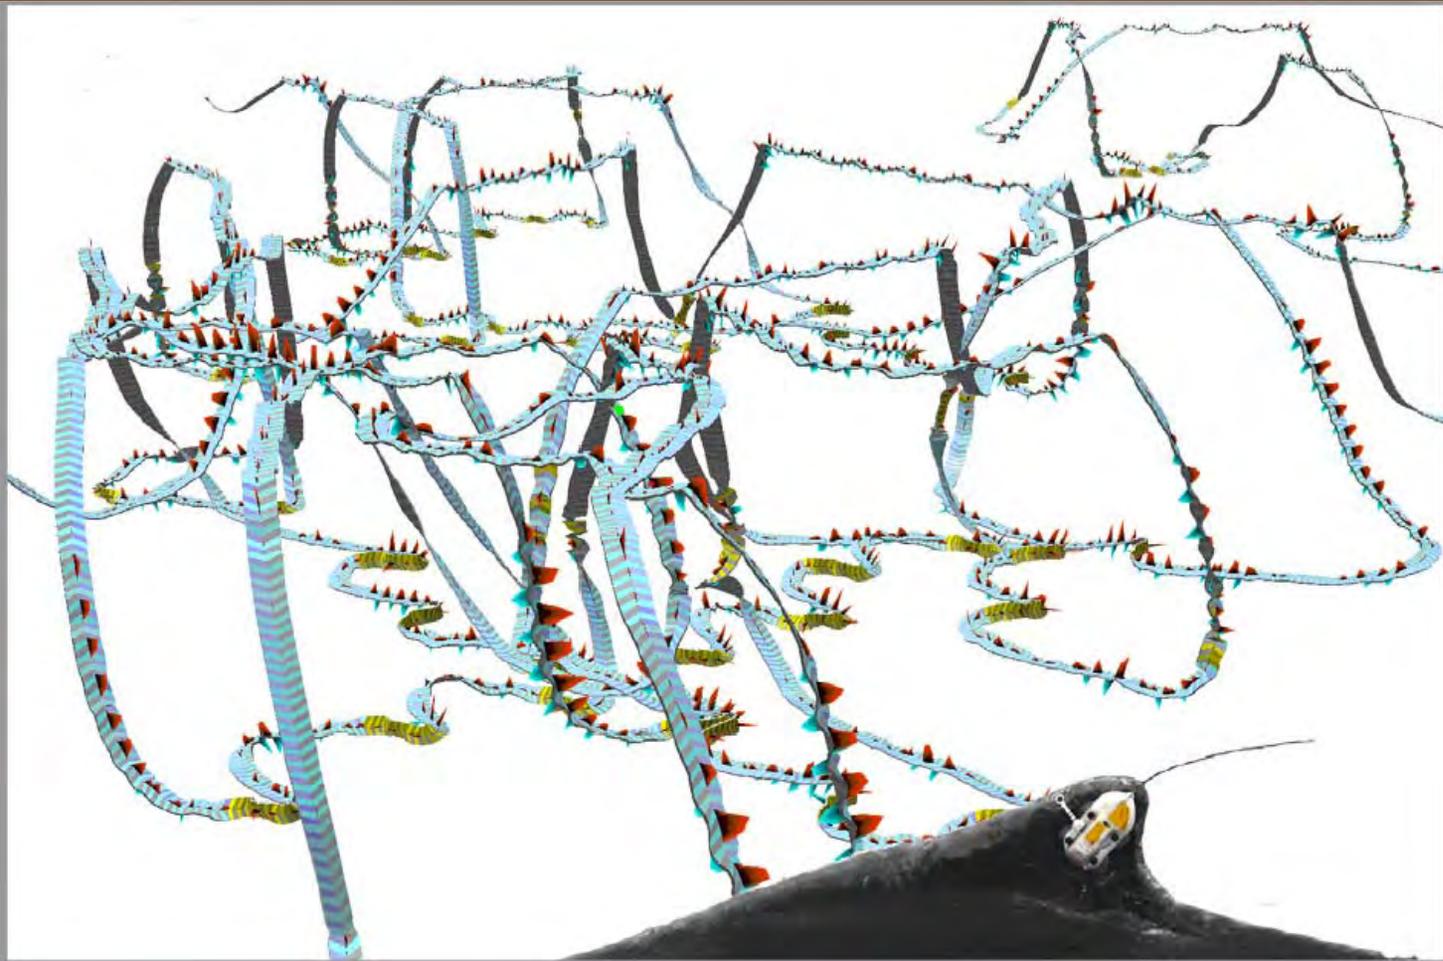
Lawrence Berkeley National Laboratory

24 June 2008

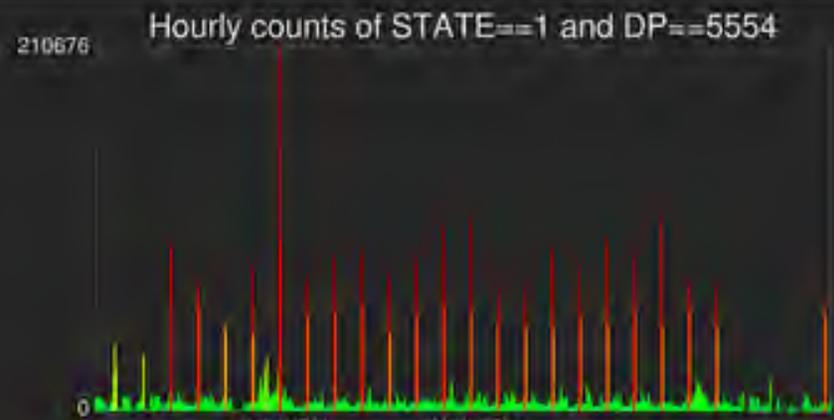
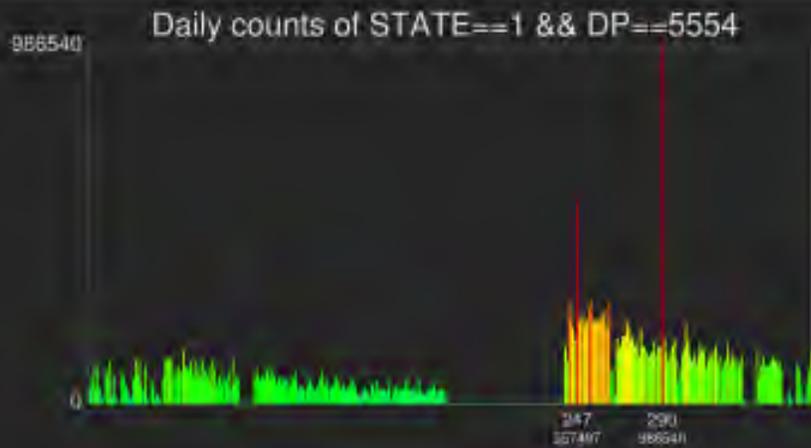
# Seeing the Unseeable - Anatomy



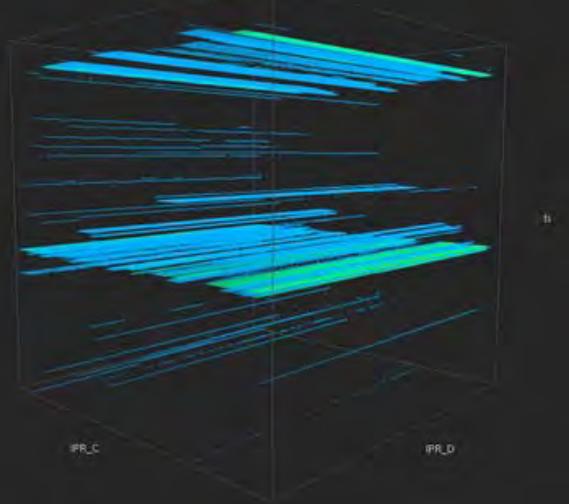
# Seeing the Unseeable – Whale Swim Path



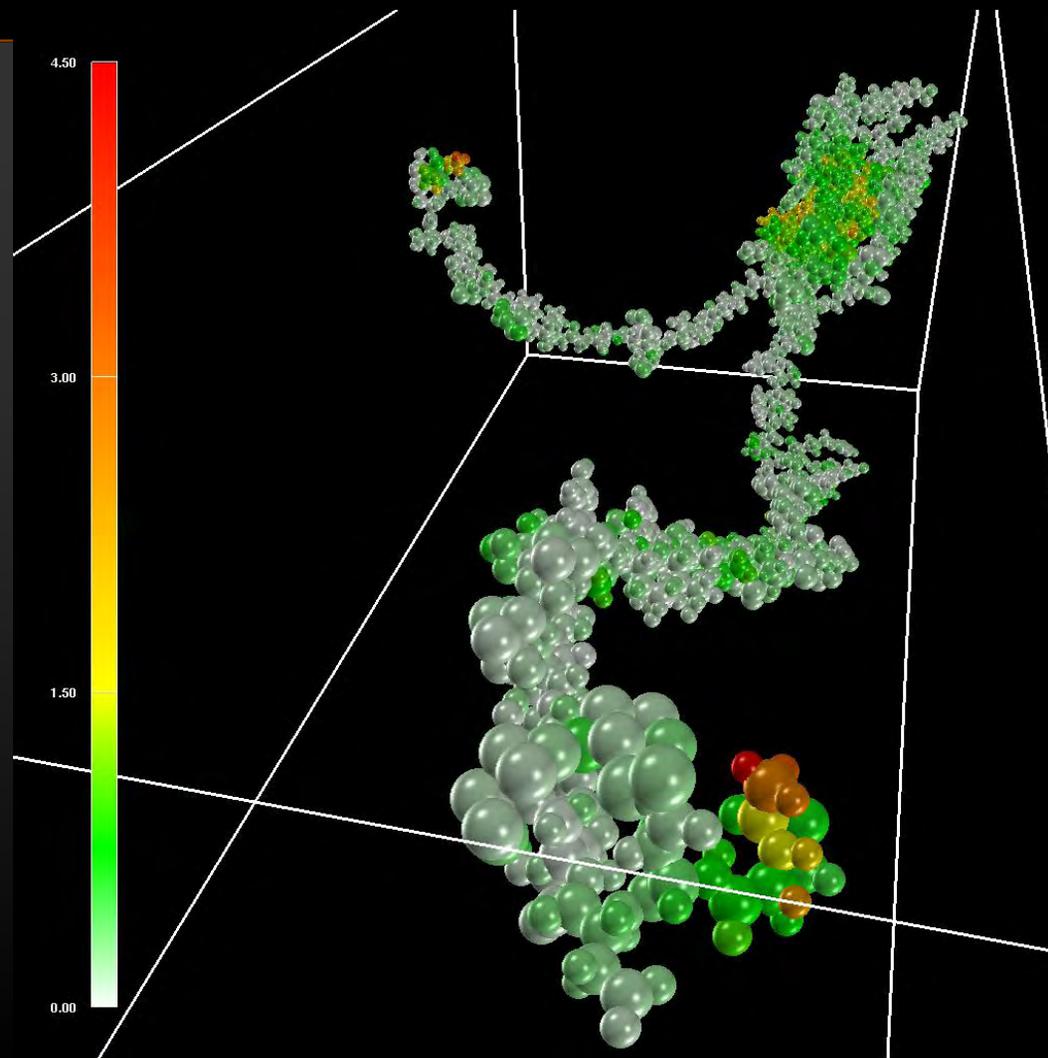
# Seeing the Unseeable – Cyberattack



IPR\_C vs IPR\_D vs ts



# Seeing the Unseeable – Protein Folding



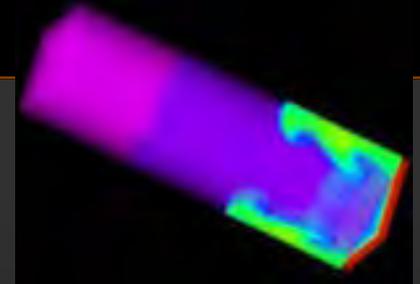
# Today's Discussion Topics

- **Objective: familiarize you with an interesting, problem-rich science domain.**
- **How?**
  - Intro to CG and Visualization
  - Ray Tracing and the Shading Equation.
  - Graphics APIs and the Graphics Pipeline.
  - Visualization: Examples, Challenges, Future.

# Introduction to CG and Visualization



# What is Computer Graphics and Visualization?



- **Graphics:**

- Transformation of data into images.
- Pen plotters (60's, 70's), storage tubes (70's), electrostatic plotters and raster displays (80's and beyond).

- **Visualization:**

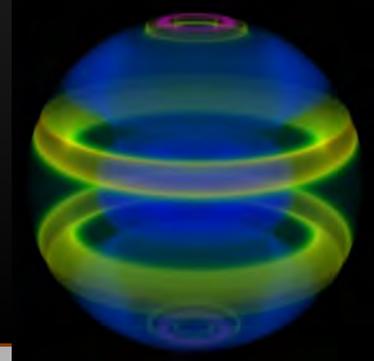
- Transformation of data into images.

- **Computer Graphics and Visualization are an alloy comprised of:**

- Computer science, cognitive science, mathematics, industrial engineering, etc.

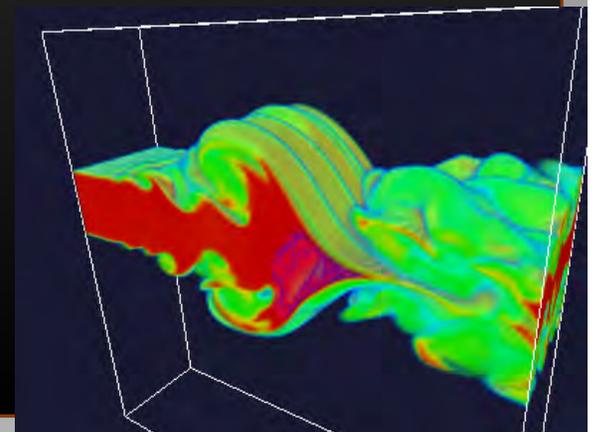
# What are Computer Graphics and Visualization?

- If both graphics and visualization are “transformation of data into images,” what is the difference?
  - Graphics tends to focus on the process of rendering
    - Transformation of primitives into pixels
    - Shading, photorealism, etc.
  - Visualization tends to focus on the mapping of abstract data (e.g., velocity) into constructs that can be rendered (e.g., triangles).
  - The distinction occasionally becomes blurry.

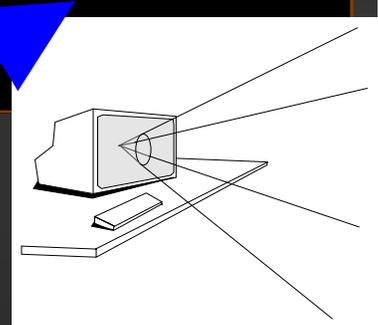


# Why Computer Graphics and Visualization?

- The visual cortex and associated machinery occupy more than half our brains – *we are innately visual creatures.*
- Vision is our principal means for understanding and interacting with the world.
- The best connection between humans and computers is through the high-bandwidth connection of our highly evolved visual system.

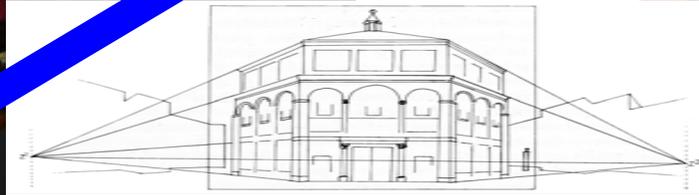


# The Path Towards Graphics and Visualization



Computers

Photography



Geometry,  
Perspective and models



Art and painting



U.S. DEPARTMENT OF ENERGY



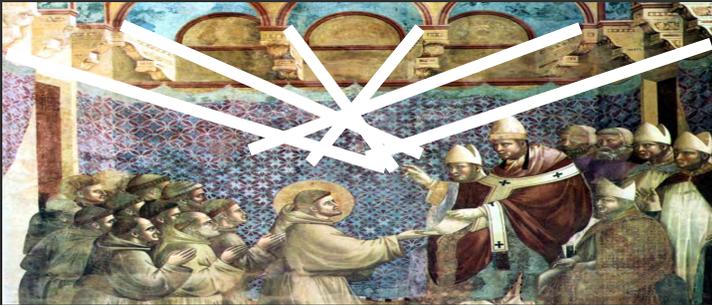
BERKELEY LAB

# The Invention of Drawing

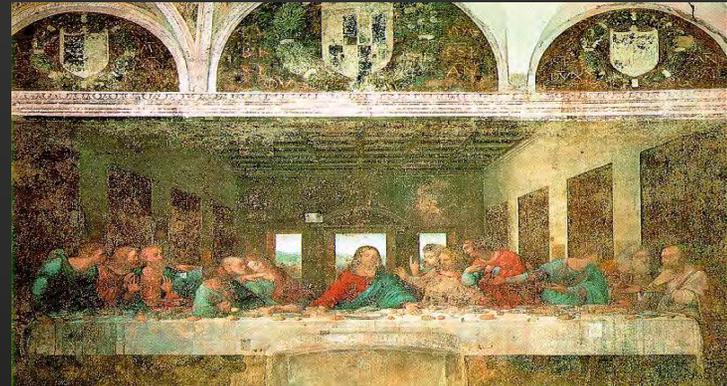
- **Painting based on mythical tale as told by Pliny the Elder: Corinthian man traces shadow of departing lover. Detail from The Invention of Drawing, 1830: Karl Friedrich Schinkel (Mitchell p.1)**



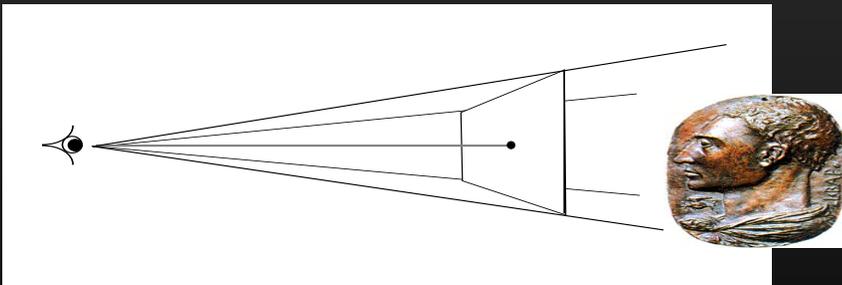
# Understanding Perspective



Incorrect perspective, Giotto, c.1295-1300



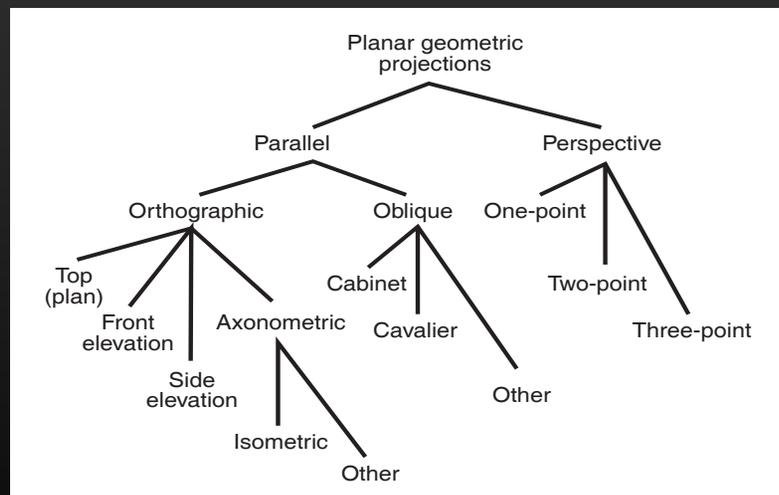
Correct perspective,  
Last Supper, Da Vinci, 1495



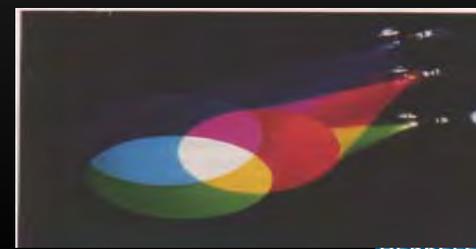
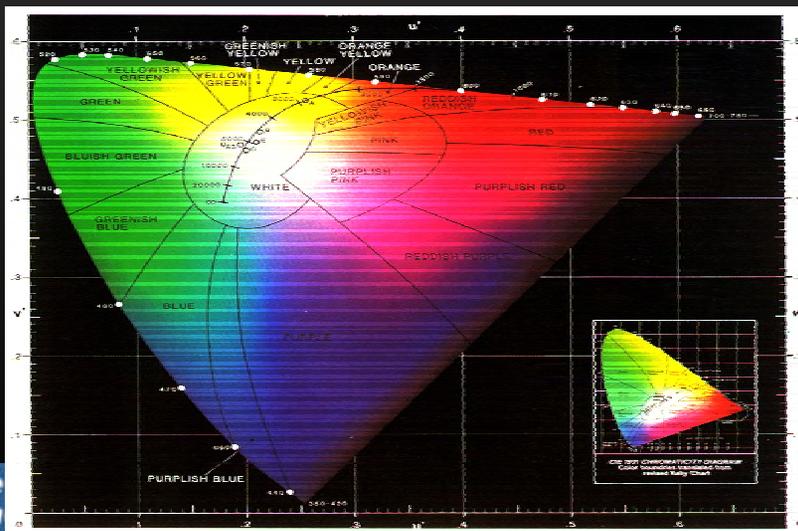
Alberti's 1435 treatise, *Della Pittura*, explained perspective for the first time

# Projections

Ender, c. 1855  
(detail of clockwork)

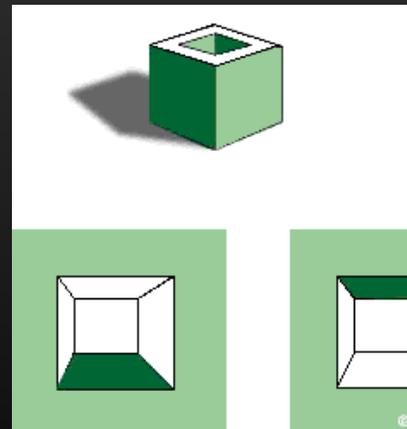


# Color and Perception



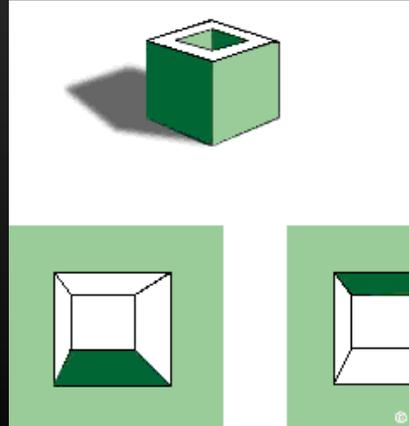
# Uses of Color

- To label, to indicate set membership.
- To indicate a value or magnitude.
- To represent reality.
- To decorate.

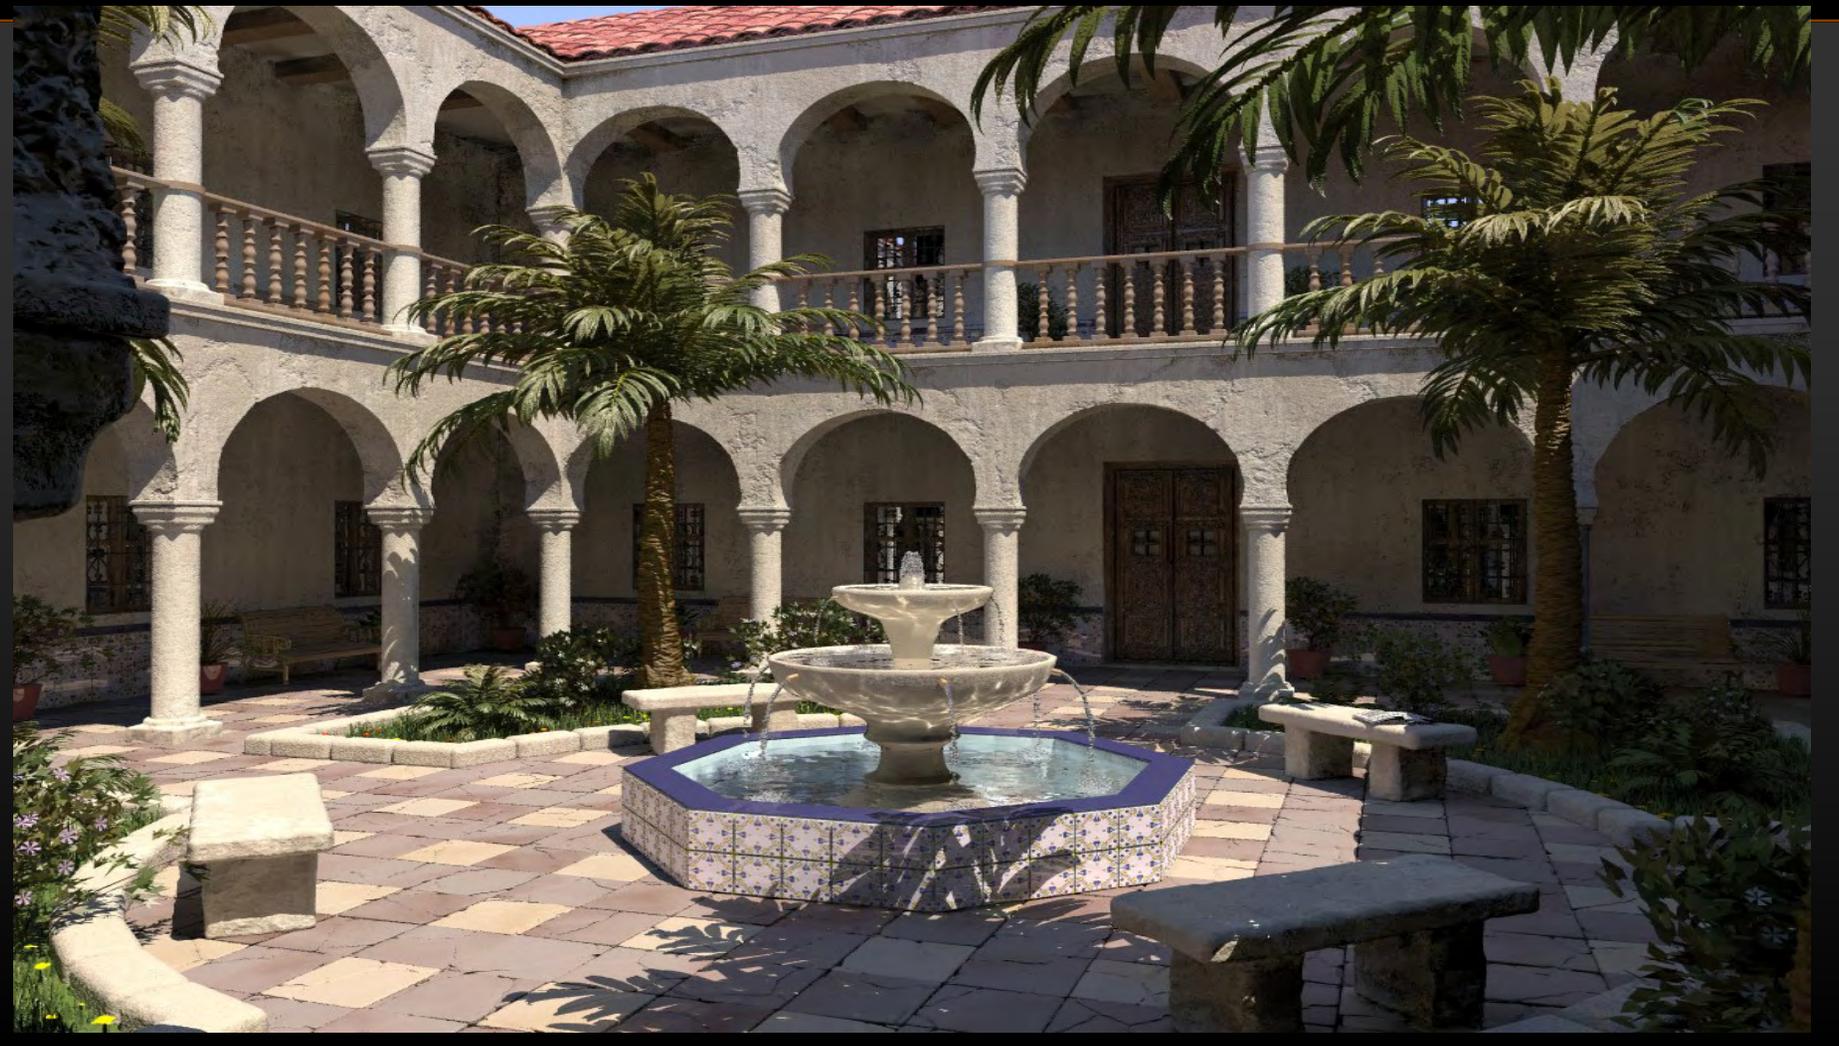


# Graphics vs. Visualization

- Top row: information encoded/presented using color (vis).
- Bottom row: color used decoratively or to represent reality.



# The Quest for Photorealism



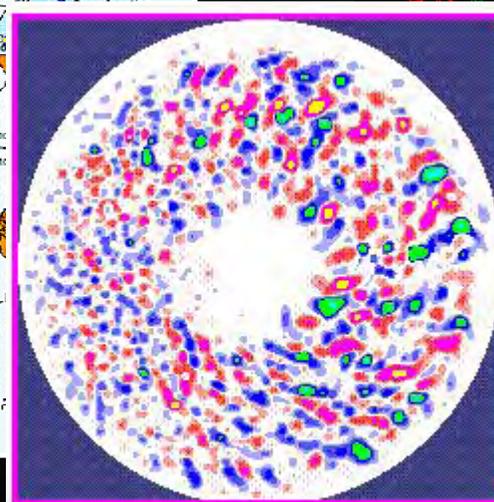
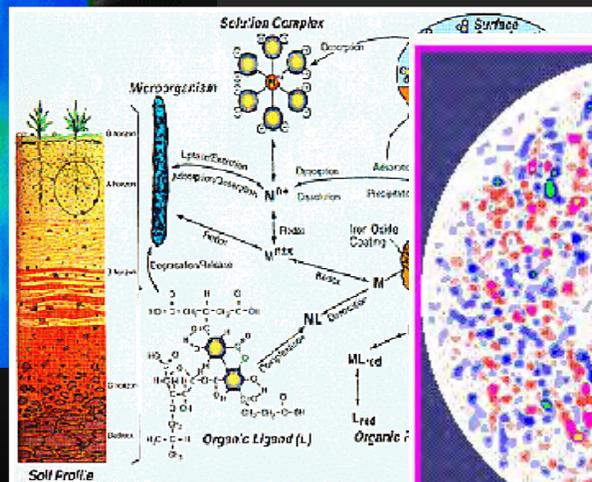
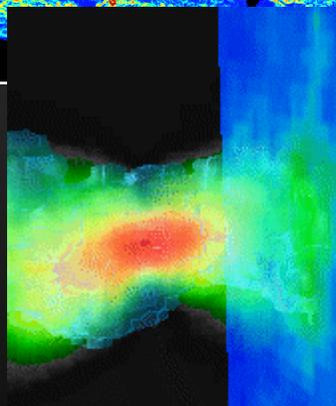
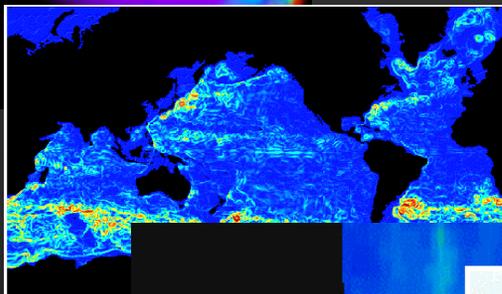
# Visualization – Enable Insight

Data

Vis

Render

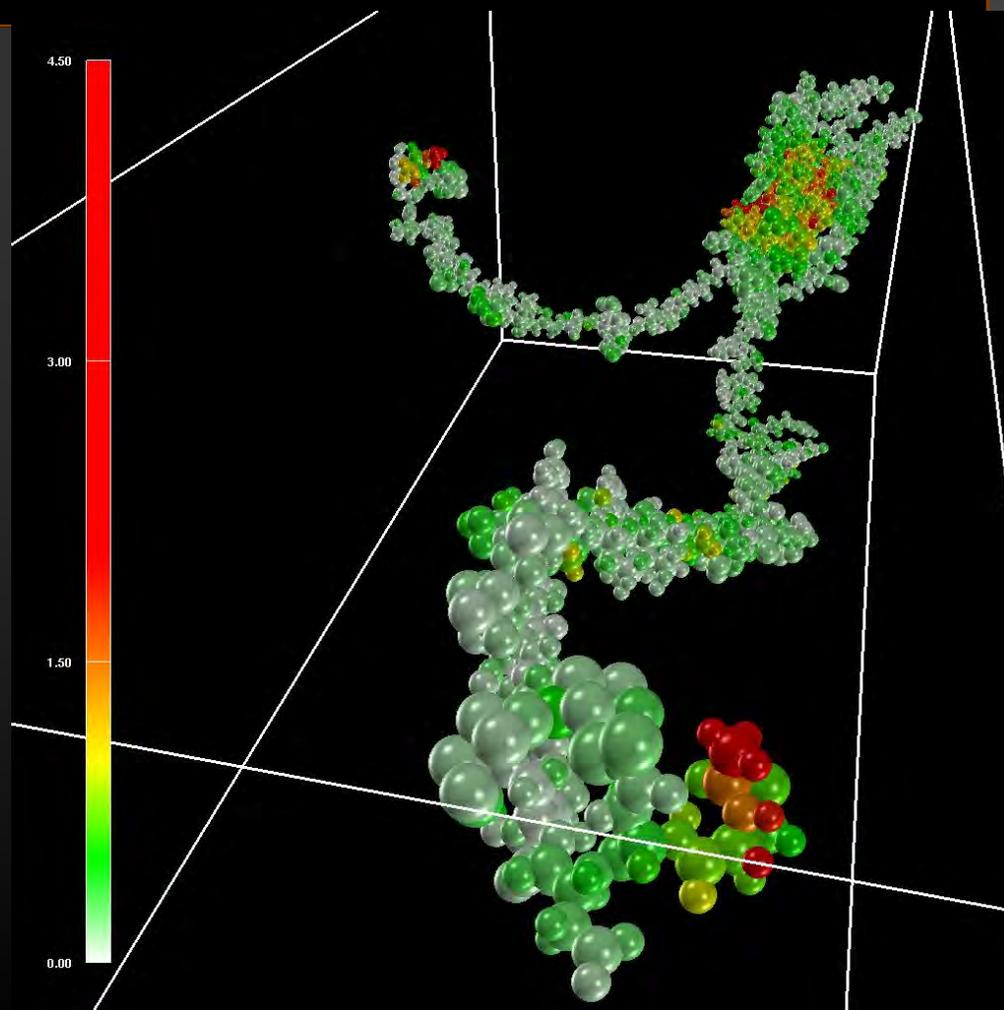
Display



SCIENCE 9/18/98  
Simulation of Plasma Turbulence

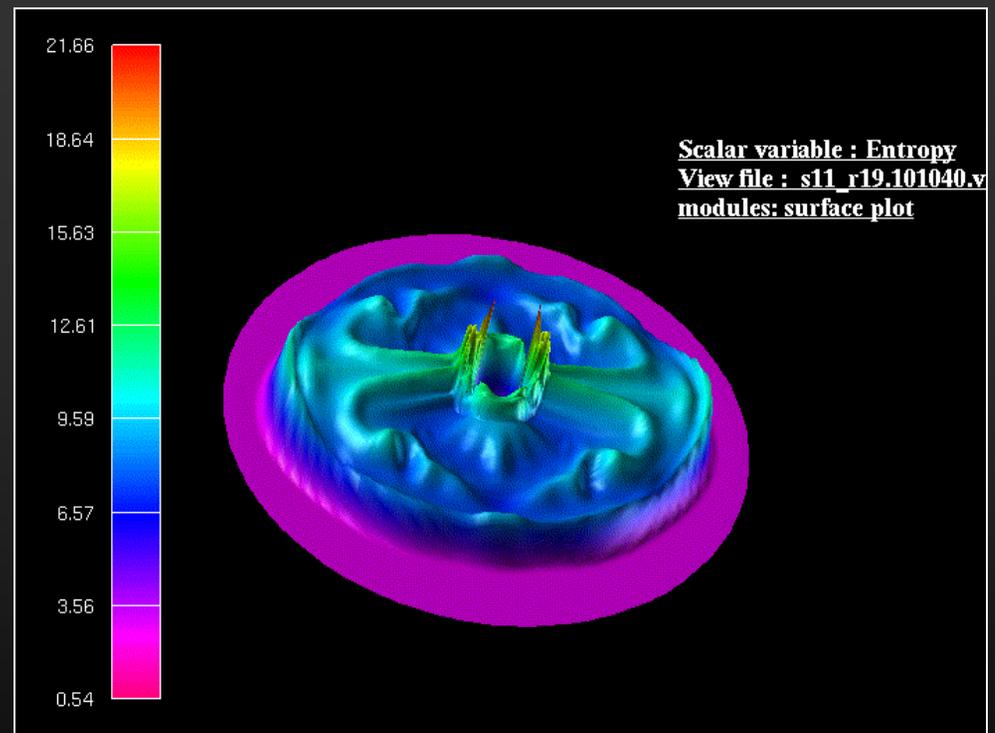
# Visualization – Good Use of Color

- Use “background colors from nature” for background elements.
- Use colors that produce a strong, innate biological response for features of interest.

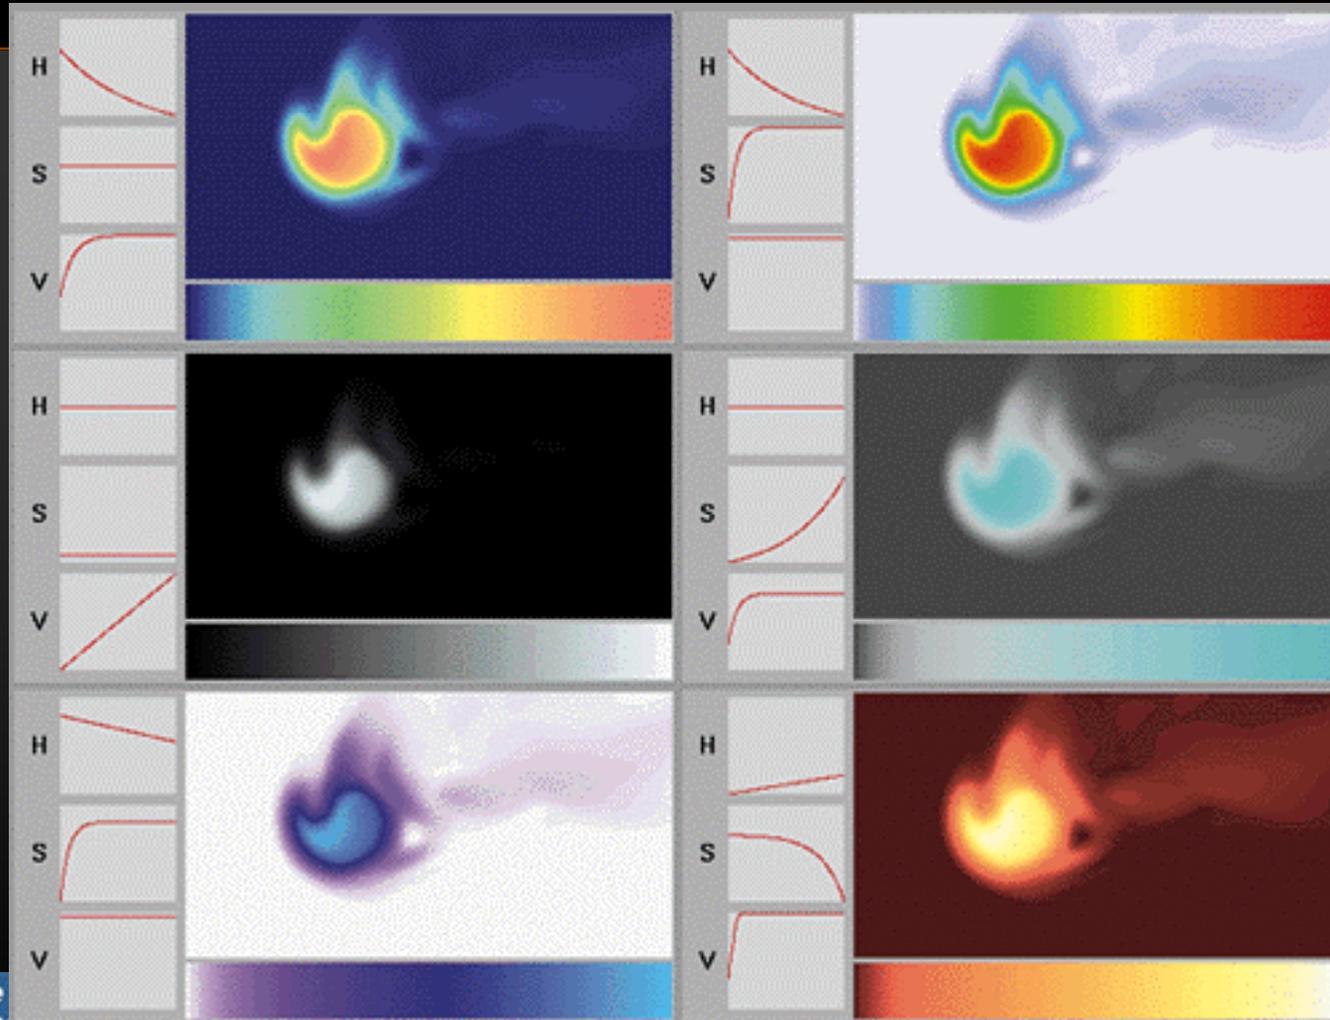


# Visualization – Bad Use of Color

- The “hue ramp” color table
- Induces 1+1=3 effect: eyes perceive a contour that doesn't really exist
- What are “important” aspects of a given visualization?



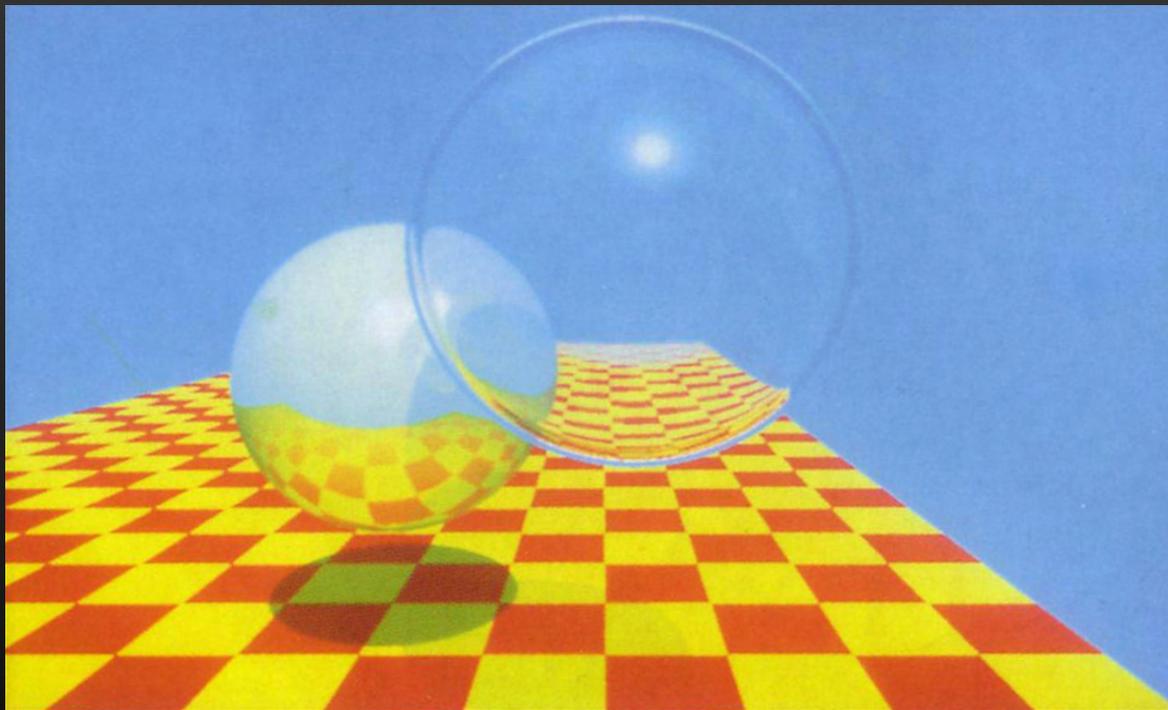
# Visualization – Use of Color



# Central Problem in Computer Graphics: The Shading Equation

- **Given a collection of objects and scene parameters, compute the color of each pixel on the screen.**
- **Two broad approaches:**
  - Object order: for each triangle (object), compute which screen pixels each one covers, compute color for all covered pixels.
  - Image order: for each image pixel, compute which objects contribute to the pixel color.

# Ray Tracing and the Rendering Equation – Image-Order CG



# Ray Tracing

- [An Improved Illumination Model for Shaded Display](#), Whitted, SIGGRAPH 1980.
- The role of the illumination model is to determine how much light is reflected to the viewer from a visible point on a surface as a function of light source direction and strength, viewer position, surface orientation, and surface properties.

# Phong's Illumination Equation (Not Ray Tracing – Yet)

$$I = I_a + k_d \sum_{j=1}^{j=ls} (\vec{N} \cdot \vec{L}_j) + k_s \sum_{j=1}^{j=ls} (\vec{N} \cdot \vec{L}'_j)^n, \quad (1)$$

where

$I$  = the reflected intensity,

$I_a$  = reflection due to ambient light,

$k_d$  = diffuse reflection constant,

$\vec{N}$  = unit surface normal,

$\vec{L}_j$  = the vector in the direction of the  $j$ th light source,

$k_s$  = the specular reflection coefficient,

$\vec{L}'_j$  = the vector in the direction halfway between the viewer and the  $j$ th light source,

$n$  = an exponent that depends on the glossiness of the surface.

# Whitted's Improvement (Ray Tracing)

- Retain diffuse term (computationally overwhelming to account for scene objects as light sources) from Phong Model.
- “Refine” specular term, add transmissive term.

$$I = I_a + k_d \sum_{j=1}^{j=ls} (\bar{N} \cdot \bar{L}_j) + k_s S + k_t T, \quad (2)$$

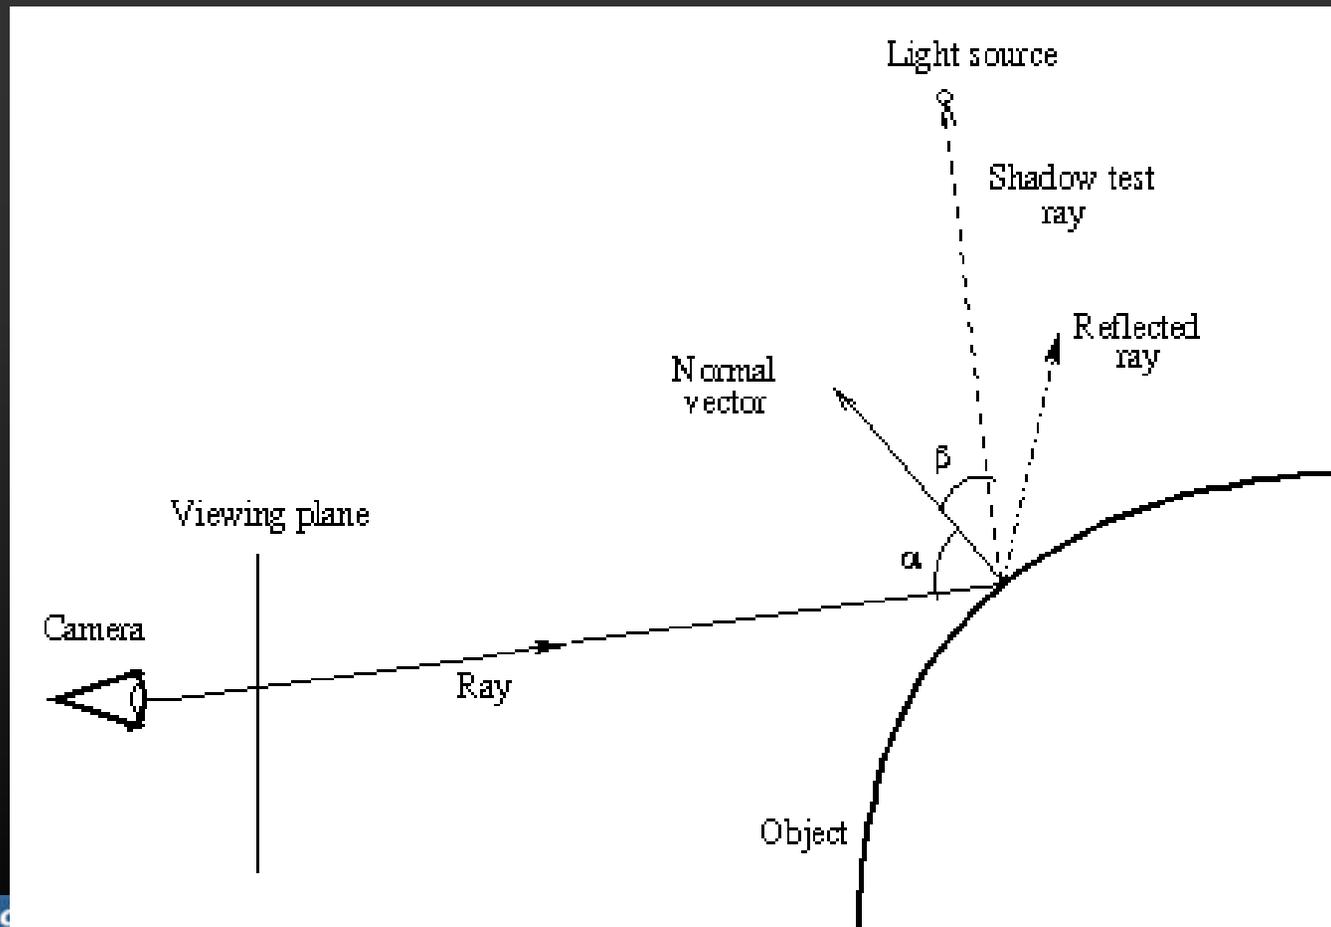
where

$S$  = the intensity of light incident from the  $\bar{R}$  direction,

$k_t$  = the transmission coefficient,

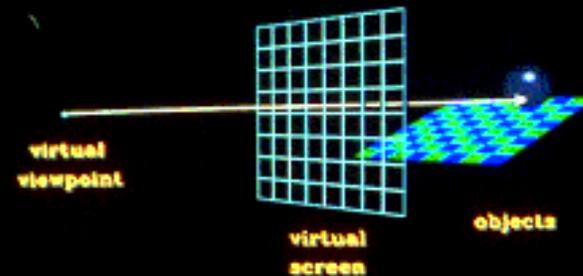
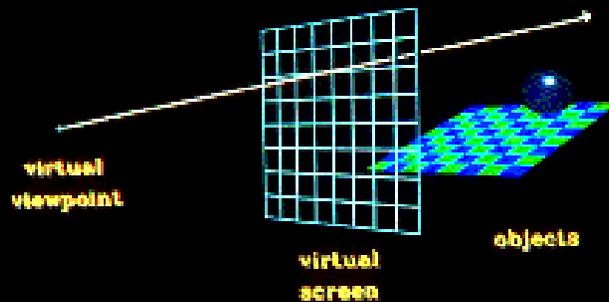
$T$  = the intensity of light from the  $\bar{P}$  direction.

# Whitted's Improvement, ctd.

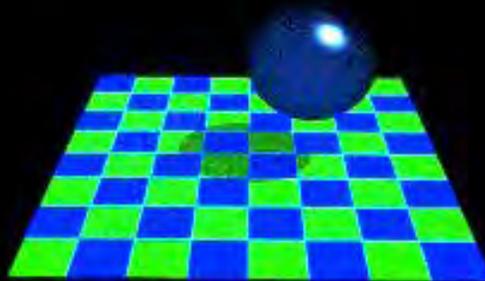
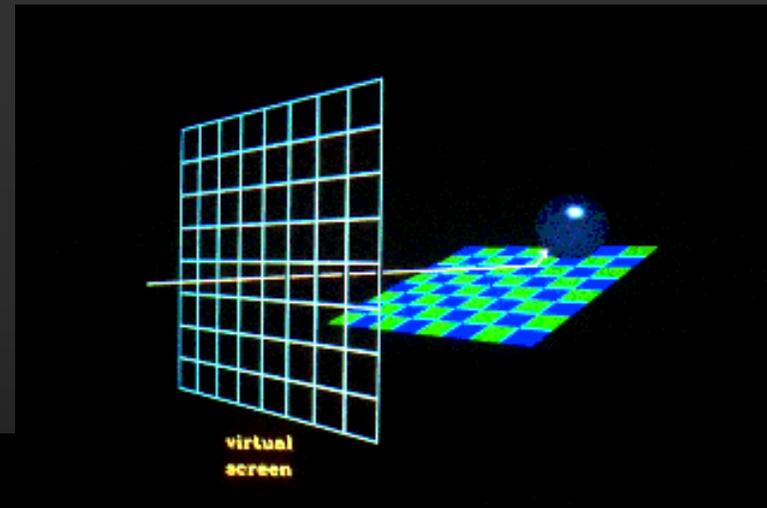
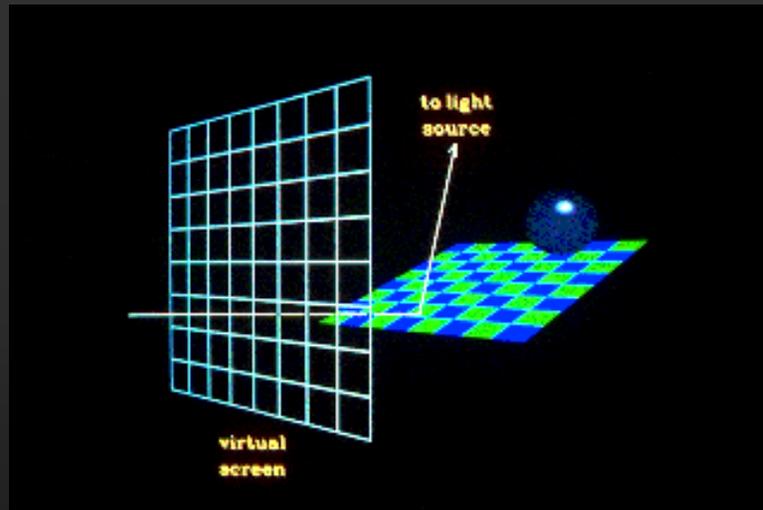


# Example Rays – Simple Intersections

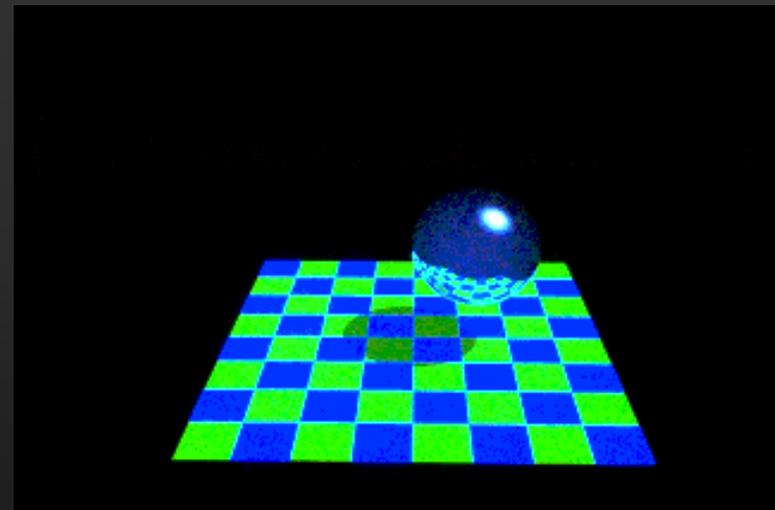
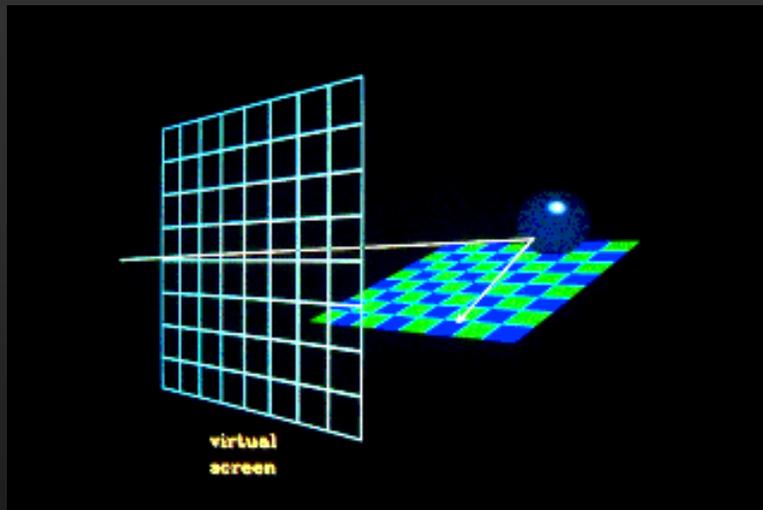
- (See <http://www.siggraph.org/education/materials/HyperGraph/raytrace/rtrace1.htm>)



# Shadow Rays

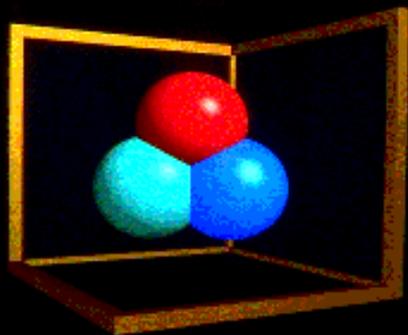


# Reflected Rays

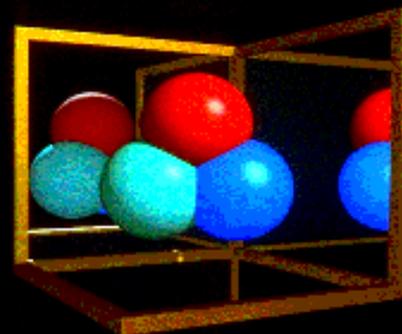


# More Reflected Rays

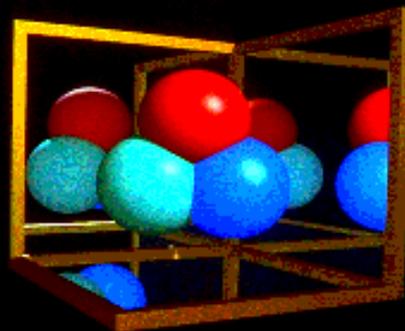
Zero bounces



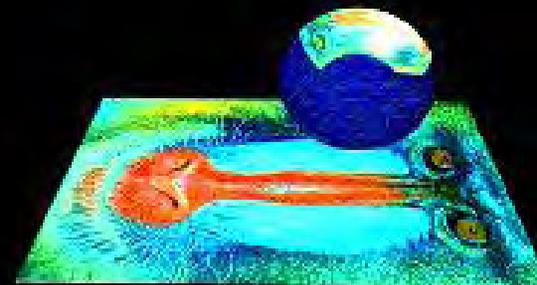
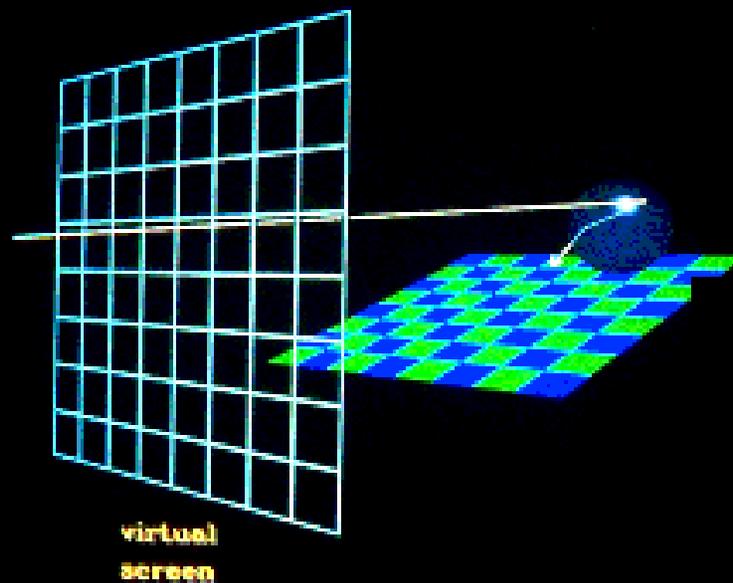
One bounce



Two bounces

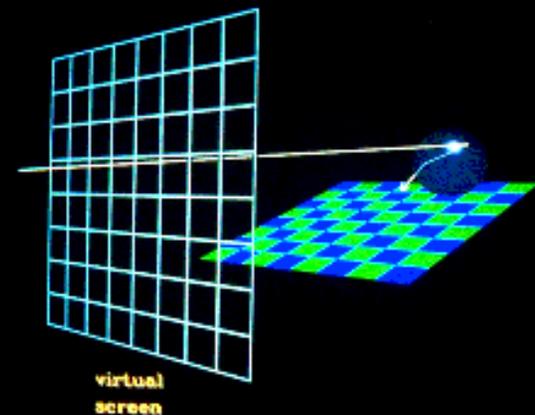


# Refracted Rays



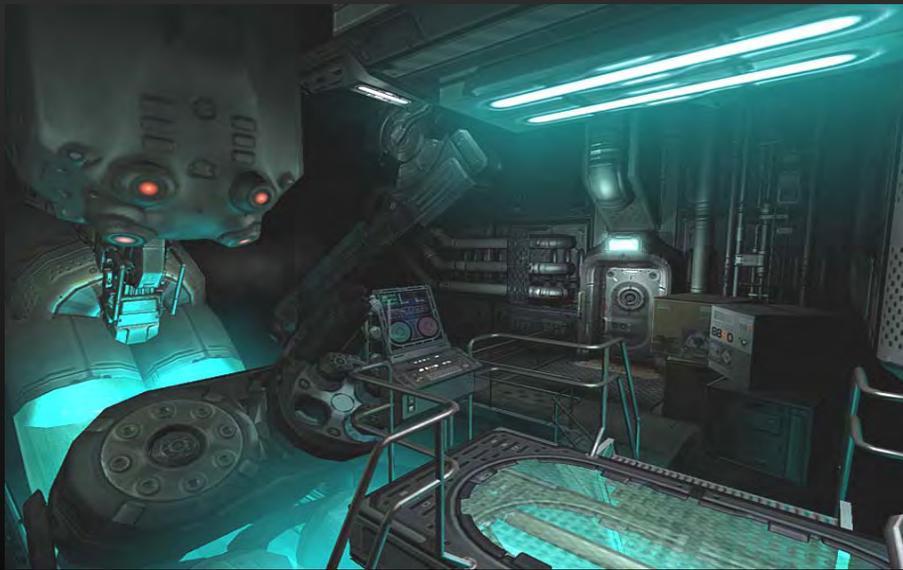
# Parallelizing Ray Tracing

- **Computation of color at each pixel is relatively independent.**
  - Image-order parallelization (data parallelism).
  - Requires (potential) duplication of the scene on each node.
- **Challenges**
  - Load balancing.
  - Data distribution.



# Graphics APIs and the Graphics Pipeline – Object-Order CG

(Images from Doom3)



# What is a Graphics API?

- **Declarative: what, not how.**
  - Scene description: viewer here, trees there, lights just so.
  - Can use a variety of rendering systems: Renderman, POVRay, Performer, OpenRM Scene Graph, etc.
- **Imperative: how, not what.**
  - A series of draw commands.
  - OpenGL, DirectX, D3D, Xlib, etc.

# OpenGL

- **Industry-standard graphics API.**
  - Supported on all major platforms.
  - Relationship between OpenGL and window system.
- **Basic primitives: lines, triangles, points.**
  - Higher level primitives (NURBS, quadrics) built using fundamental geometry.
- **Phong lighting model, but Gouroud shading.**

# Example OpenGL Code



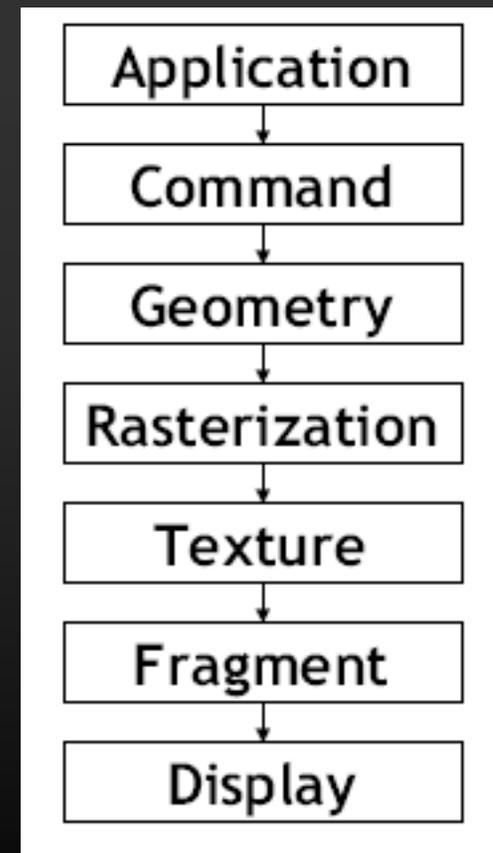
```
glBegin(GL_POLYGON) ;  
  glColor (RED) ;  
  glVertex3i (0,0,0) ;  
  glVertex3i (1,0,0) ;  
  glVertex3i (0,1,0) ;  
glEnd ()
```



```
glBegin(GL_POLYGON) ;  
  glColor (RED) ;  
  glVertex3i (0,0,0) ;  
  glColor (BLUE) ;  
  glVertex3i (1,0,0) ;  
  glColor (BLUE) ;  
  glVertex3i (0,1,0) ;  
glEnd ()
```

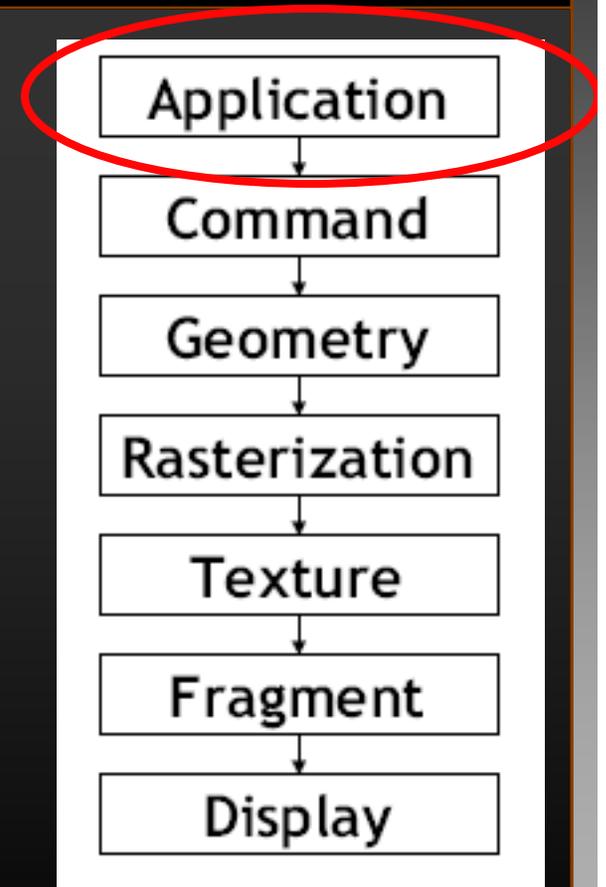
# General: Graphics Pipeline

- Application issues “draw” commands.
- Each command dispatches objects (e.g., geometry).
- Each object is scan-converted into pixel fragments.
- Each fragment is textured, etc.
- And finally displayed.



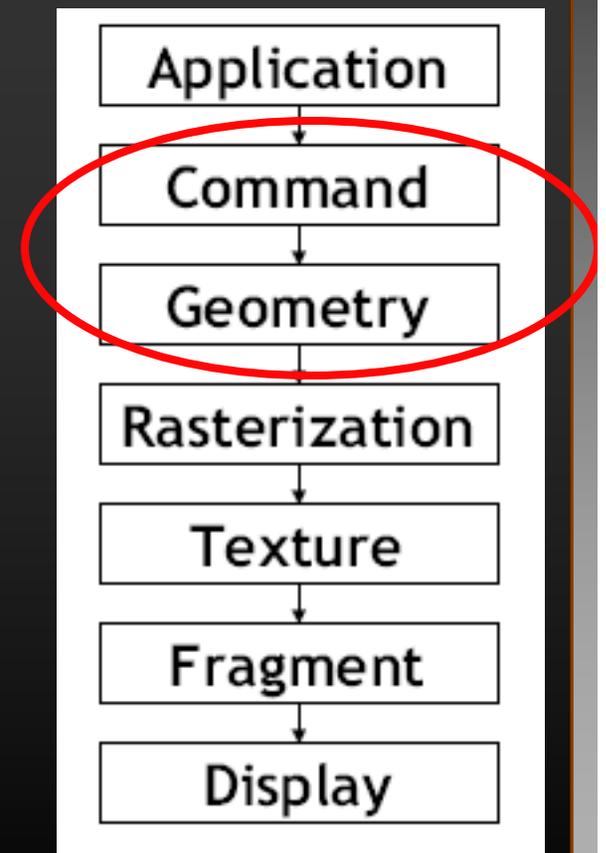
# The Graphics Application

- **Application**
  - Simulation
  - Visualization
  - Database Traversal
  - User interaction (games!)



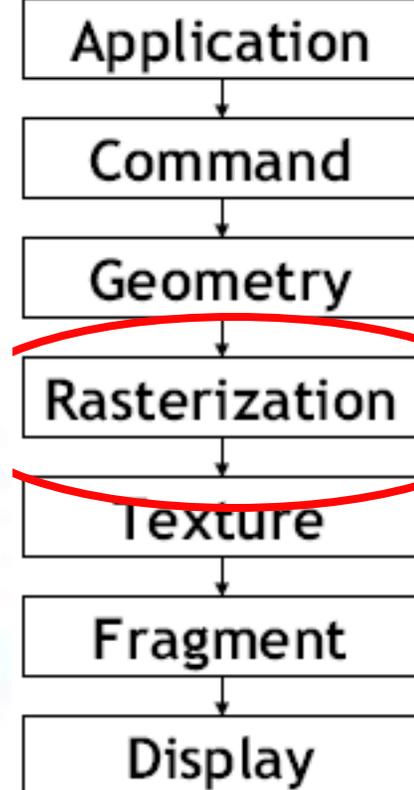
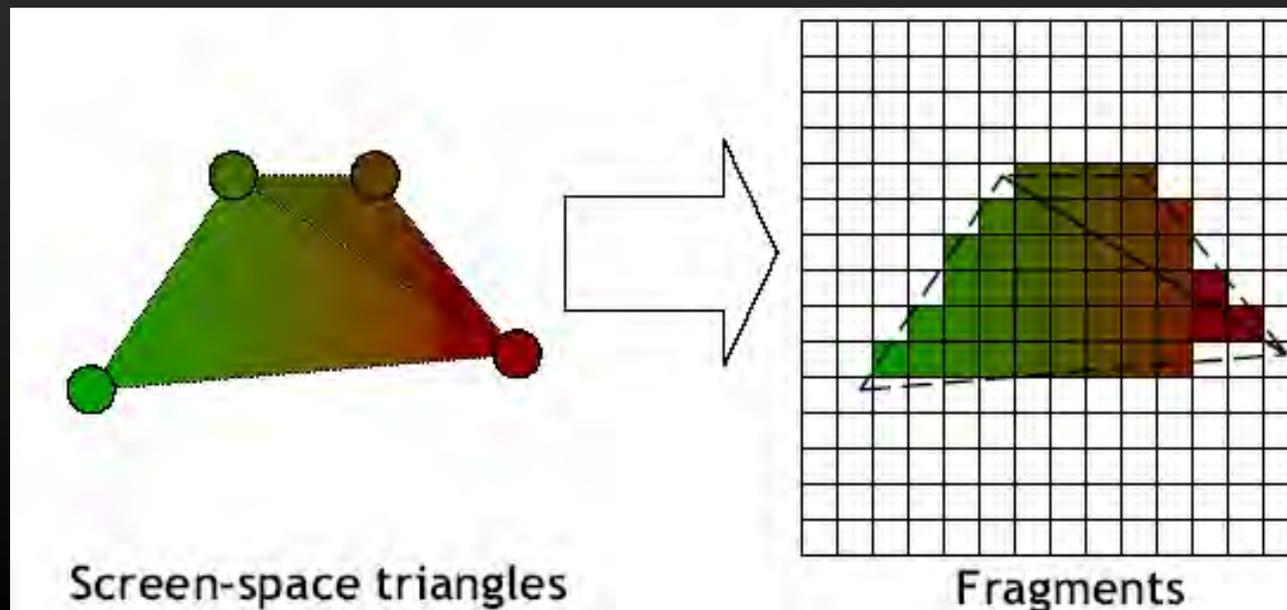
# Commands and Geometry

- **Command**
  - (What kinds of commands?)
  - Buffering, parsing.
- **Geometry**
  - Transform, light.
  - Automatic operations.
  - Culling, clipping.



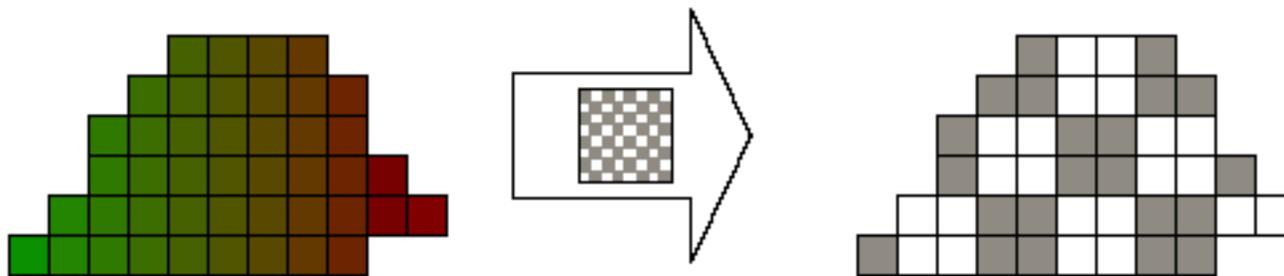
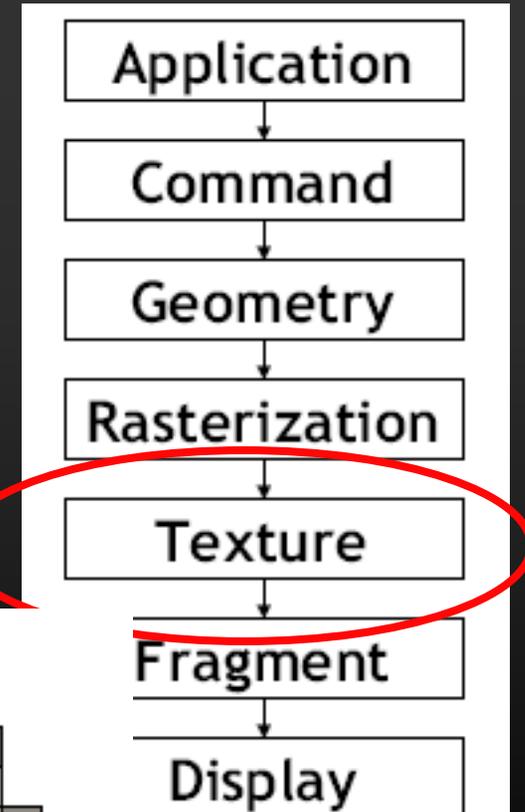
# Rasterization

- Setup, sampling (produces “fragments”, interpolation (color, depth).



# Texturing

- Texture transformation and projection.
- Texture address calculation.
- Texture filtering.



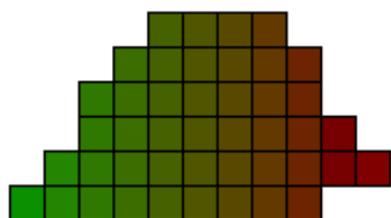
# Fragment Operations

Texture combiners and fog

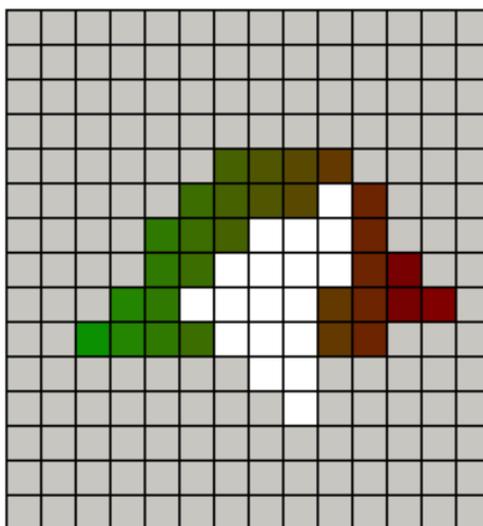
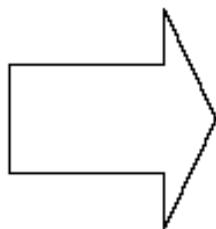
Owner, scissor, depth, alpha and stencil tests

Blending or compositing

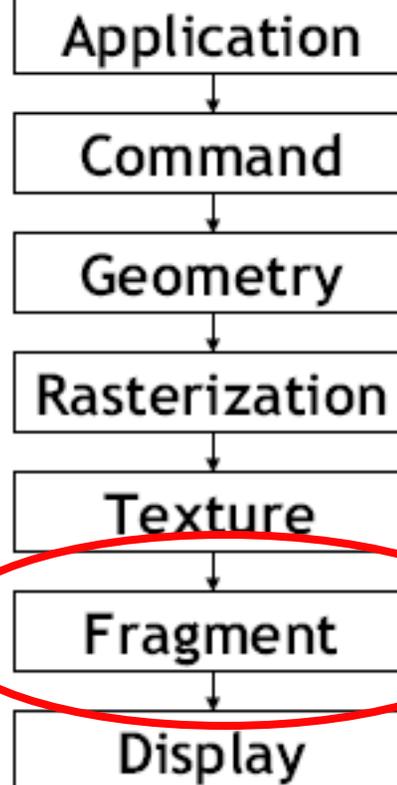
Dithering and logical operations



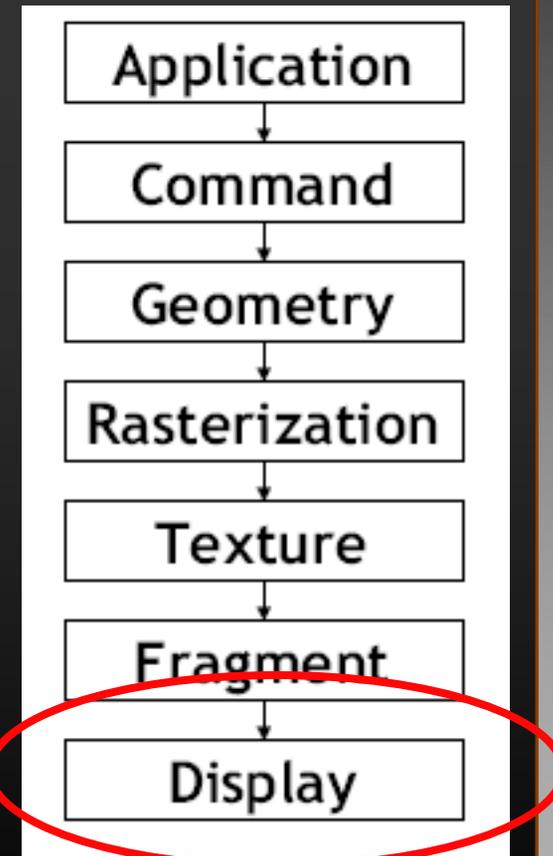
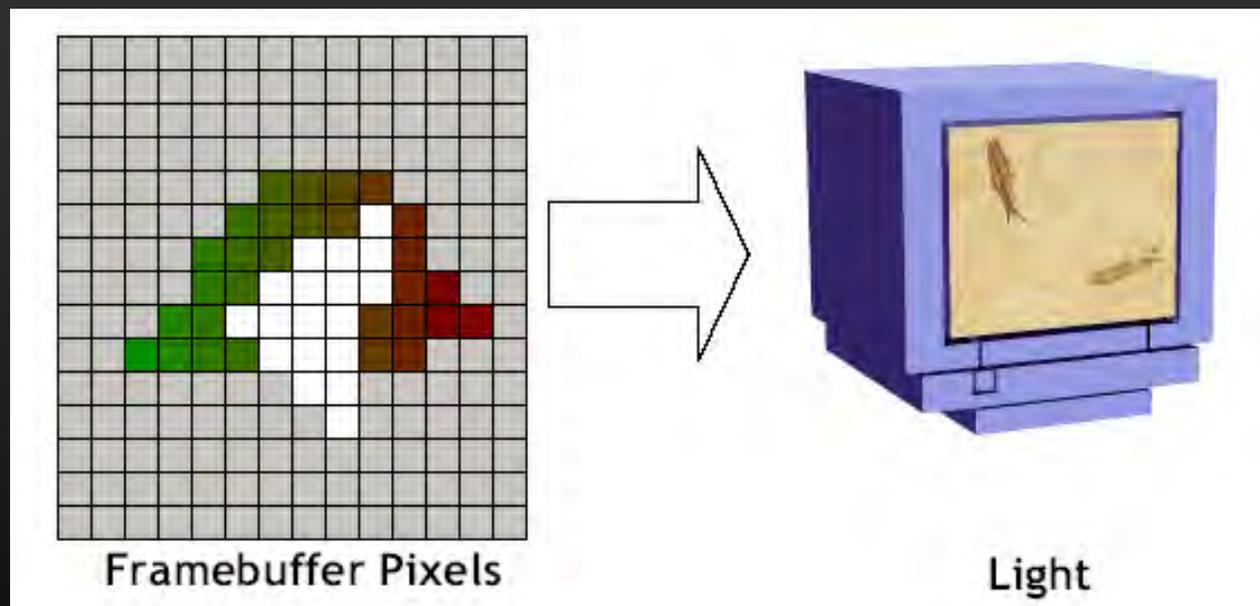
Textured Fragments



Framebuffer Pixels



# Display

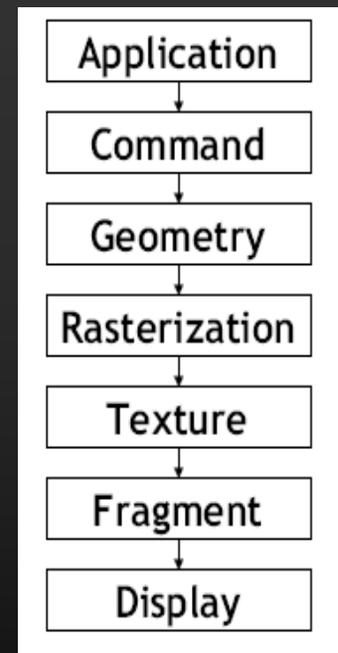


# Parallelism in Graphics



# Sources of Parallelism

- **Task parallelism**
  - Graphics Pipeline
    - Parallelizing operations within a single pipe.
    - Pipelining operations within a single pipe.
- **Data Parallelism**
  - Frame & image parallel (image order).
  - Object/geometry parallel (object order).



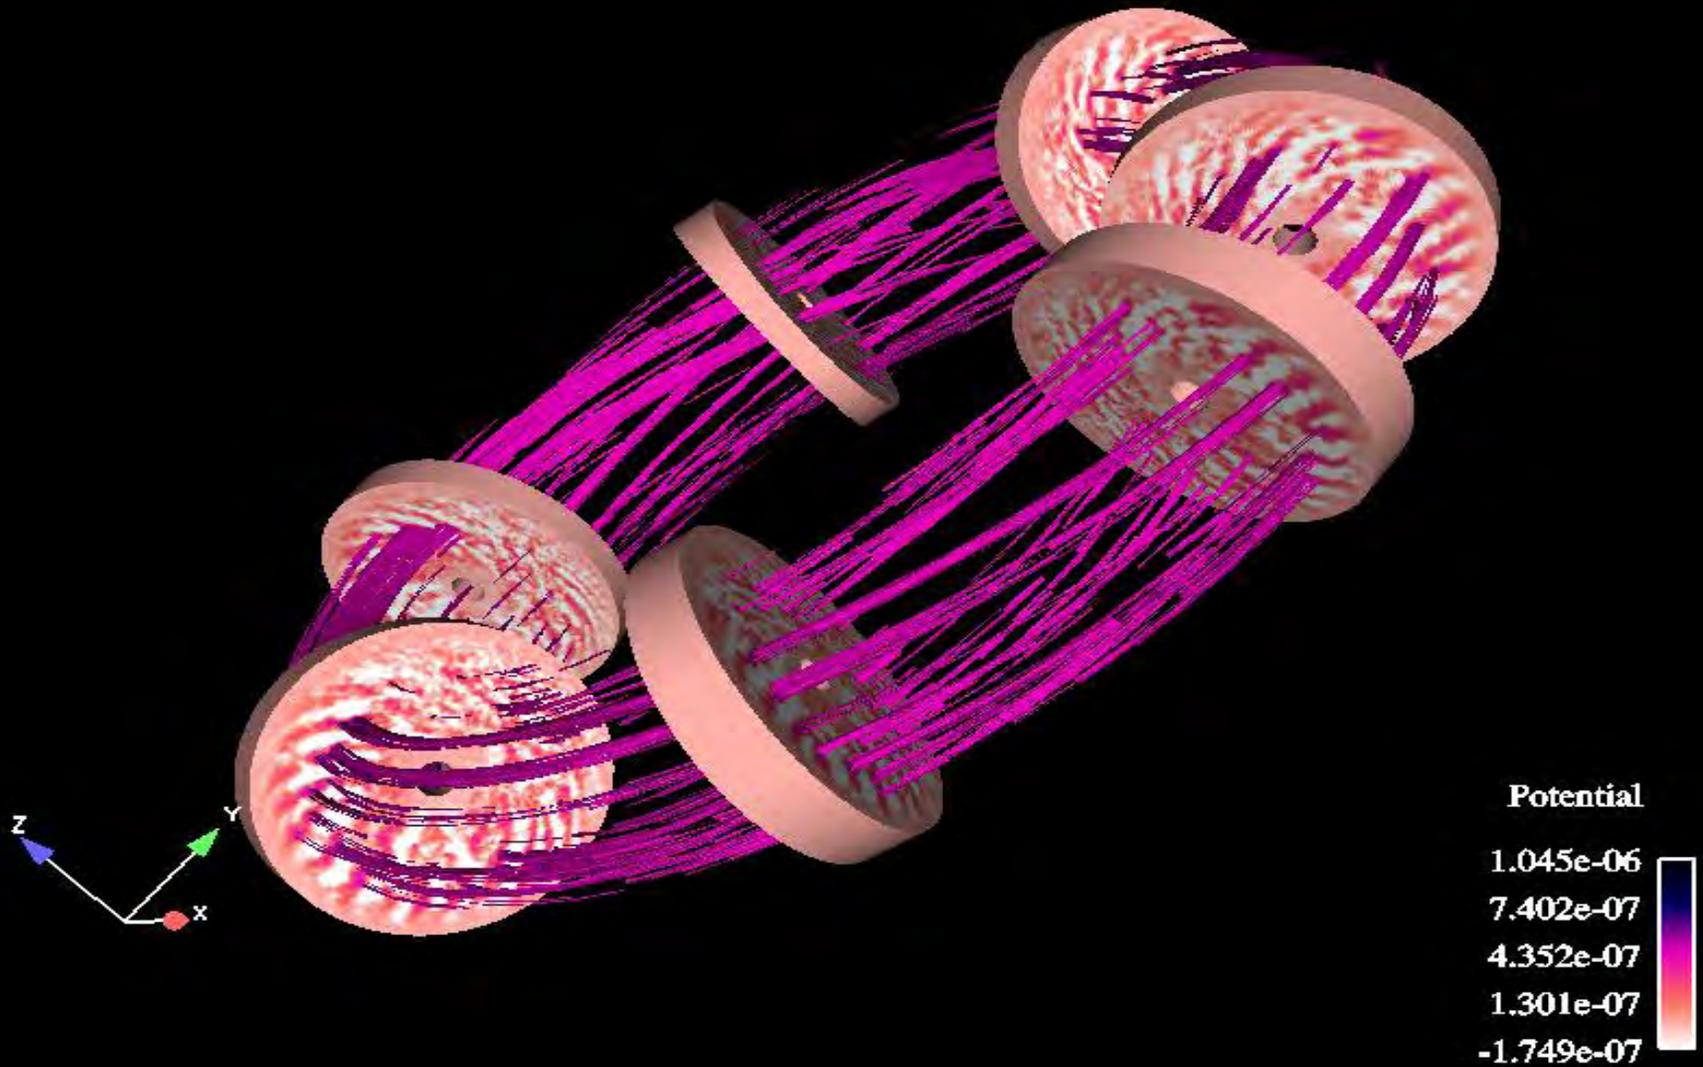
# Parallelism and GPUs

- **NVIDIA 8800 GTX Ultra**
  - 16 processors, 768 threads each.
  - 12K simultaneous threads!
- **These threads are simultaneously working on:**
  - Transform and light (vertex operations)
  - Rasterization
  - Texturing/fogging/alpha blending/stencil (fragment ops)

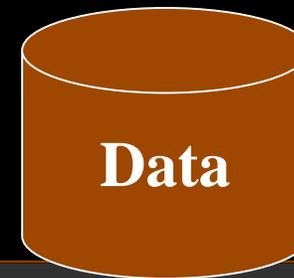
# Programming GPUs

- **“Old Way”**
  - Vertex and fragment programs.
  - “Assembly language” programming
- **New Way: Compute Unified Device Architecture – CUDA**
  - C language with extensions for multithreading
  - Lots of promise/interest for high performance computing.
  - Teraflop GPUs vs. Gigaflop CPUs

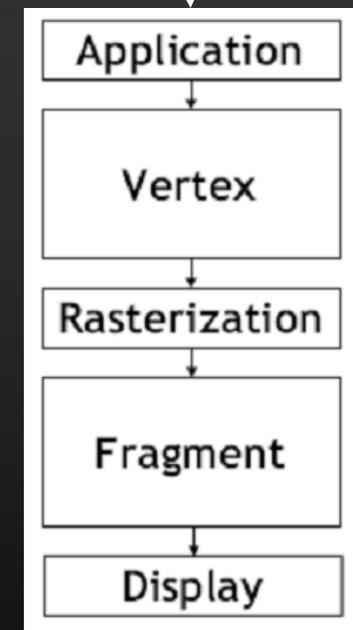
# Visualization



# Visualization



- **Visualization is the art/science of transforming abstract data into images.**
- **Visualization algorithms:**
  - Produce data that can be then processed by a traditional graphics pipeline, or
  - Are rendering algorithms that produce images.
  - Are highly data intensive (opportunity for performance gain through parallelism!!)

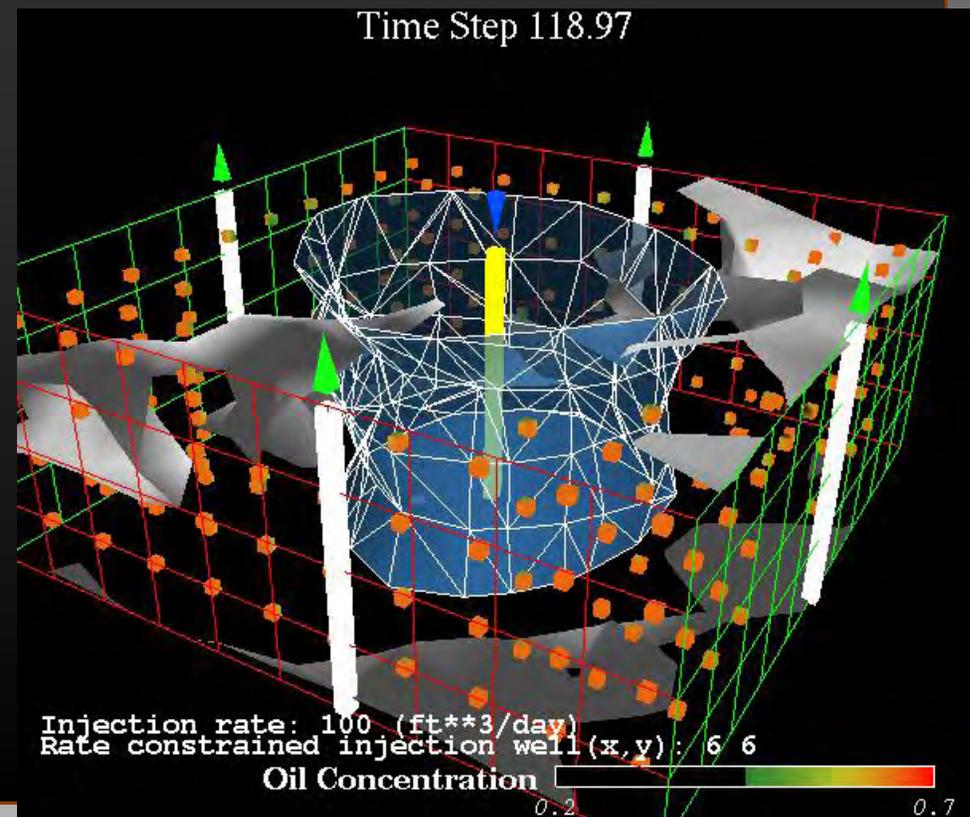


# Visualization Research Examples

- Petroleum Reservoir Modeling
- Protein structure prediction
- Quantitative analysis/comparison of Rayleigh-Taylor instability

# Petroleum Reservoir Modeling

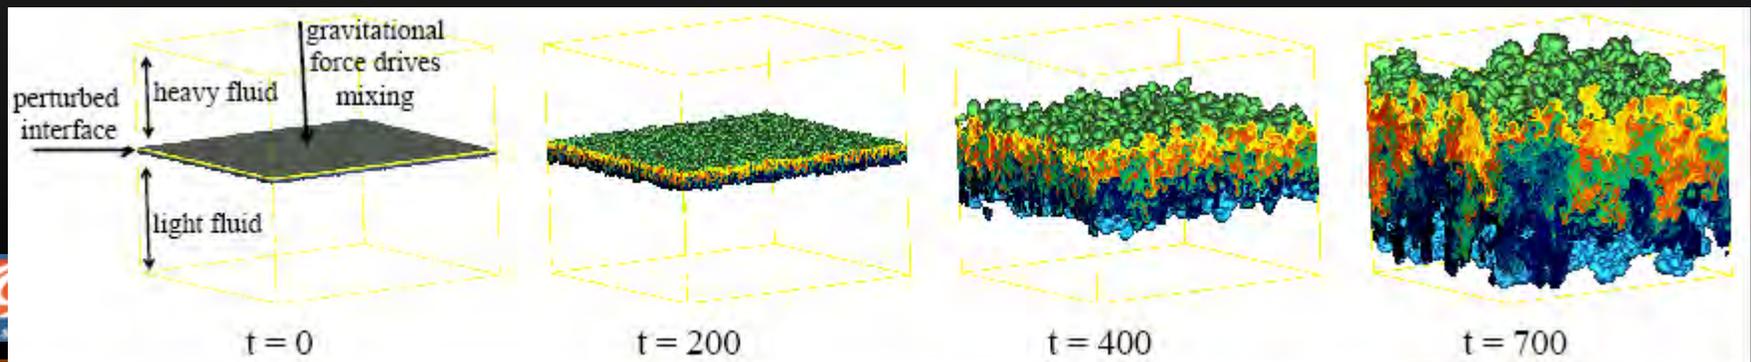
- **Problem:**
  - Find optimal location of injection and production wells.
  - No direct solution
- **Approach**
  - Leverage human intuition
  - Couple simulation and visualization



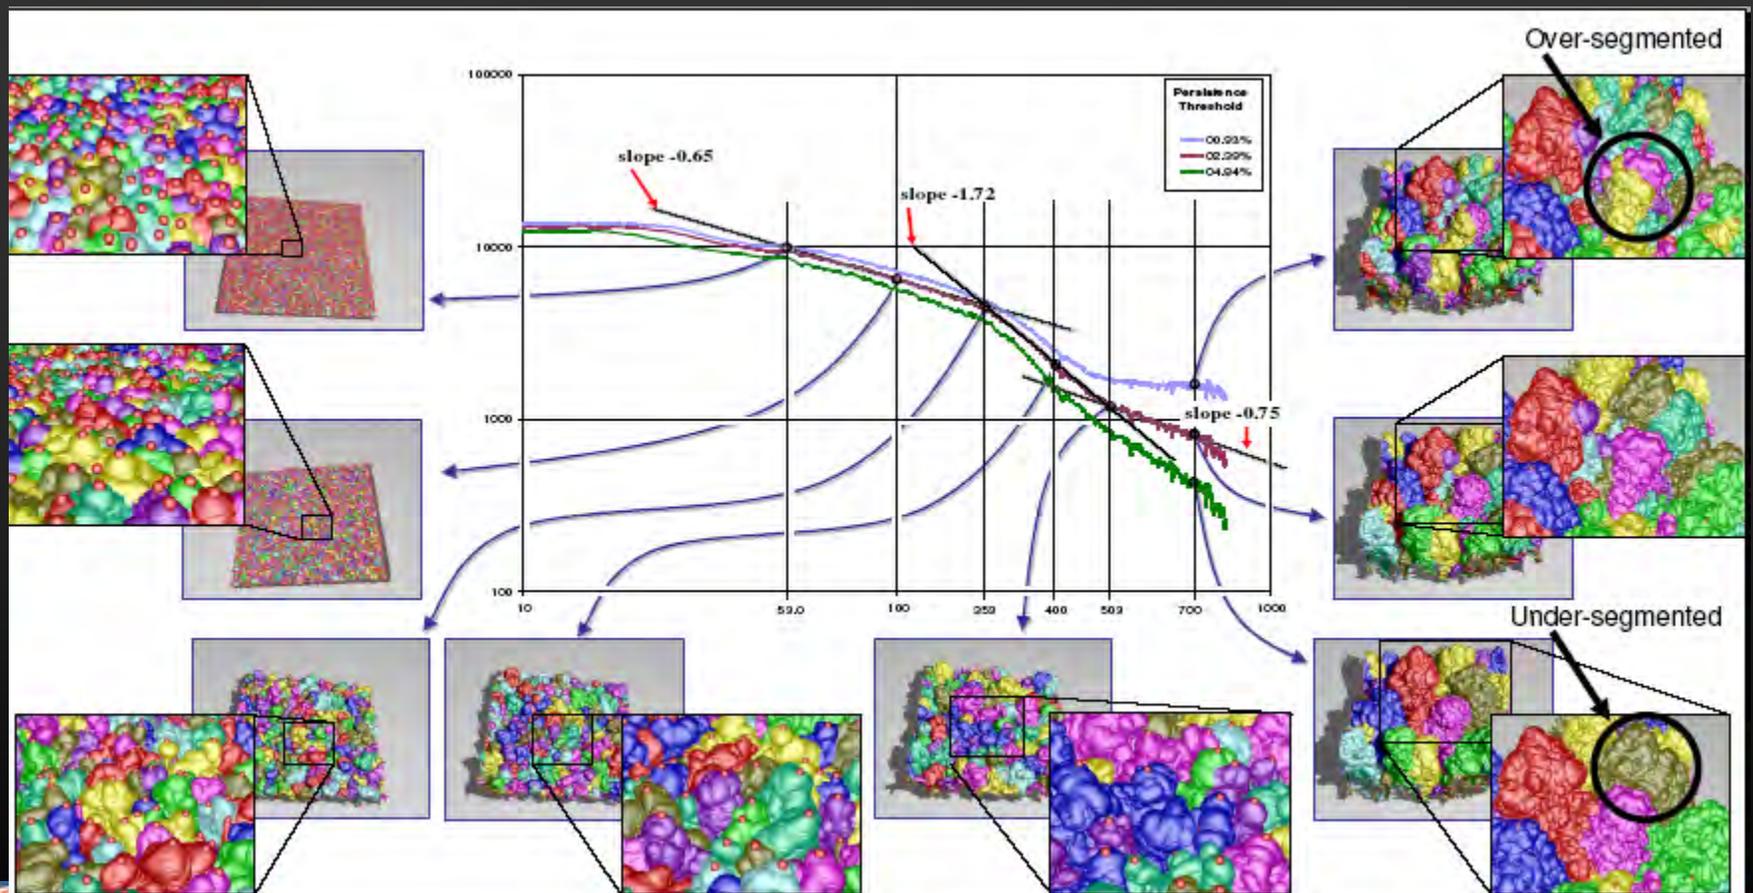


# Rayleigh-Taylor Instability

- **Problem: understand impact of various parameters on mixing.**
  - Want to compare results of different simulations.
- **Solution: topological analysis to identify, count, and track mixing regions.**

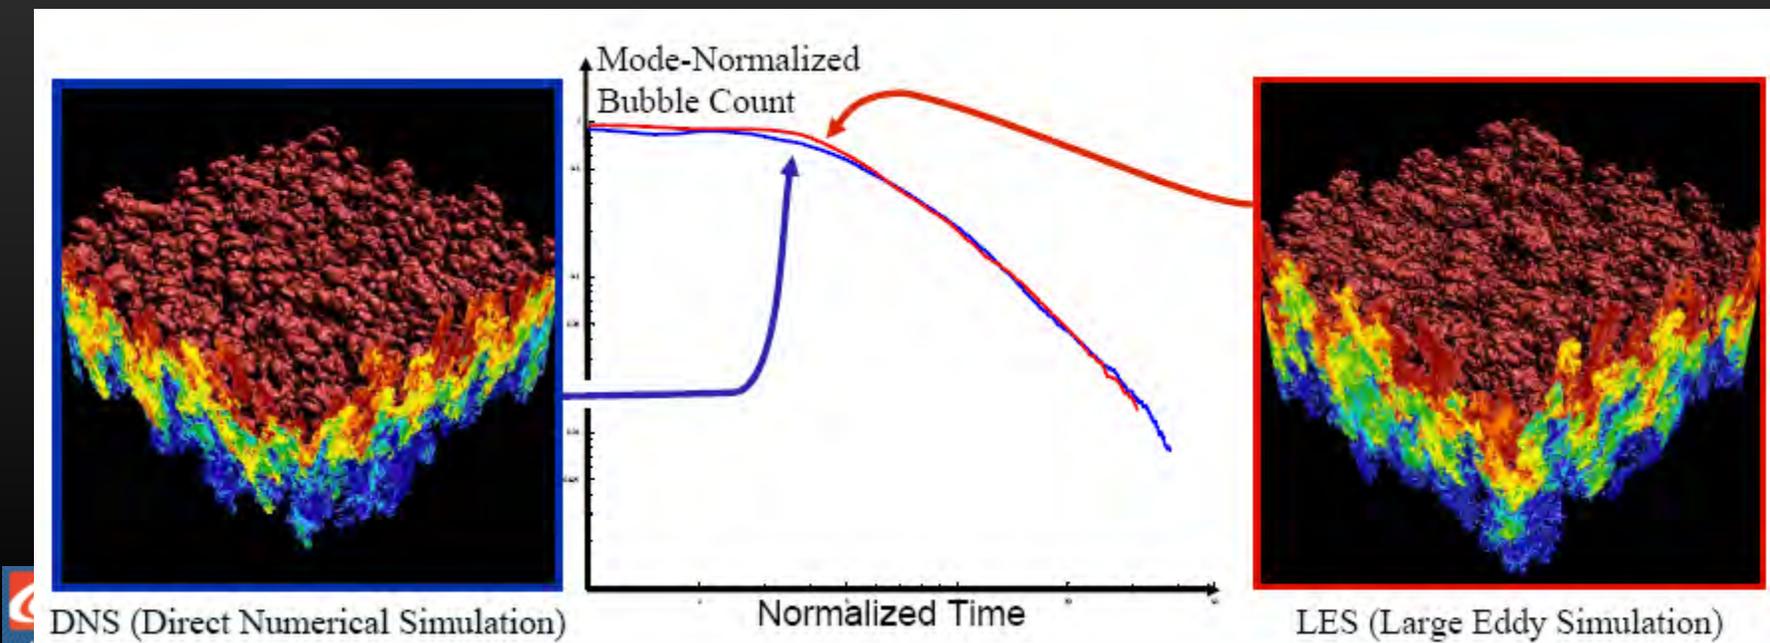


# Rayleigh-Taylor Instability, ctd.

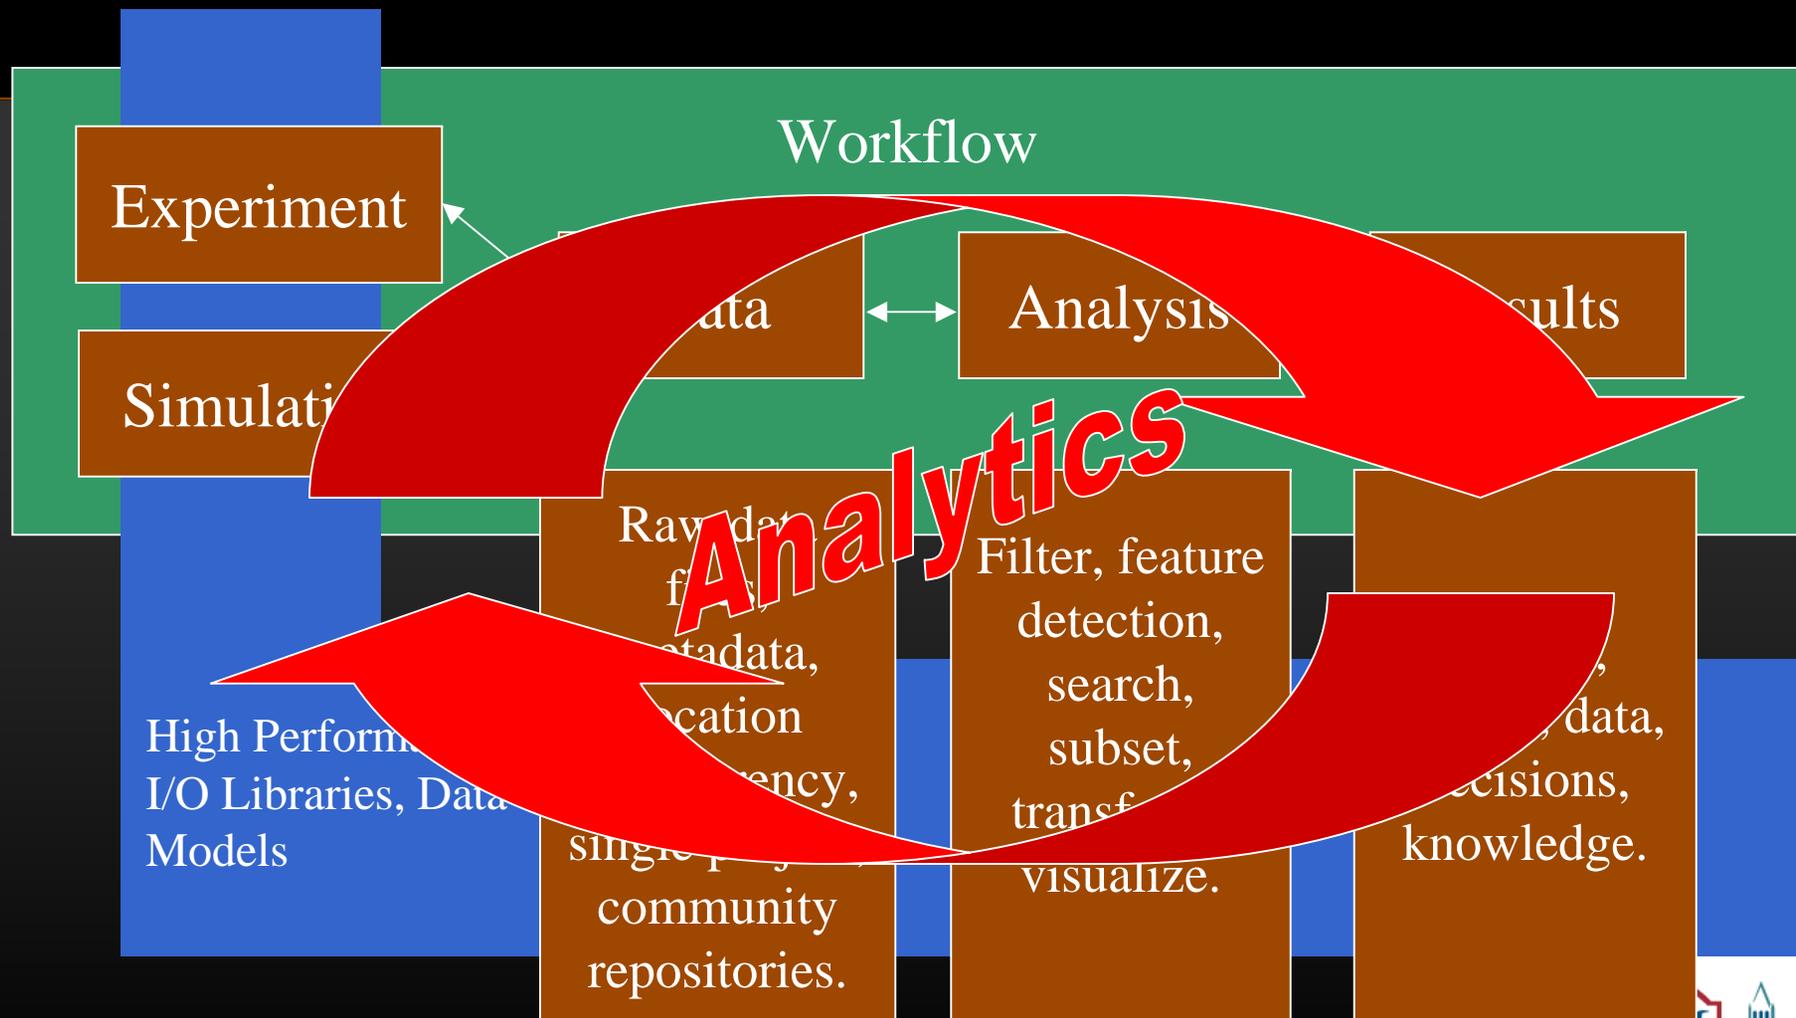


# Rayleigh-Taylor Instability, ctd.

- **Bubble count: quantitative basis for comparison (sim-sim, in this case)**

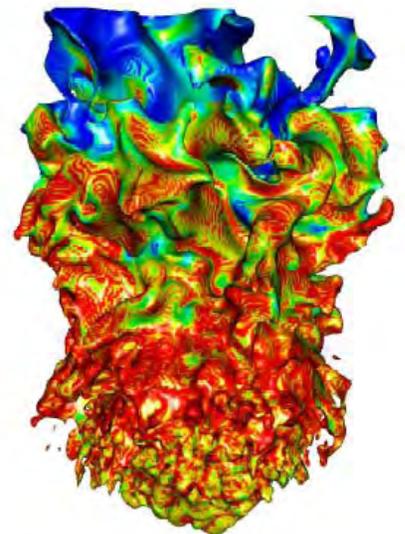


# Map of the Problem Space



# Challenges in Visualization/Graphics

- **Primary problems**
  - Computing platforms machines becoming more powerful and complex: increased parallelism, heterogeneous architectures.
  - Data becoming larger and more complex: increased spatial resolution, more variables, more timesteps.
- **Secondary problems**
  - Comparing datasets: sim-sim, experiment-sim.
  - Analysis to find/confirm correlations.
- **Really big problem**
  - Too much information!

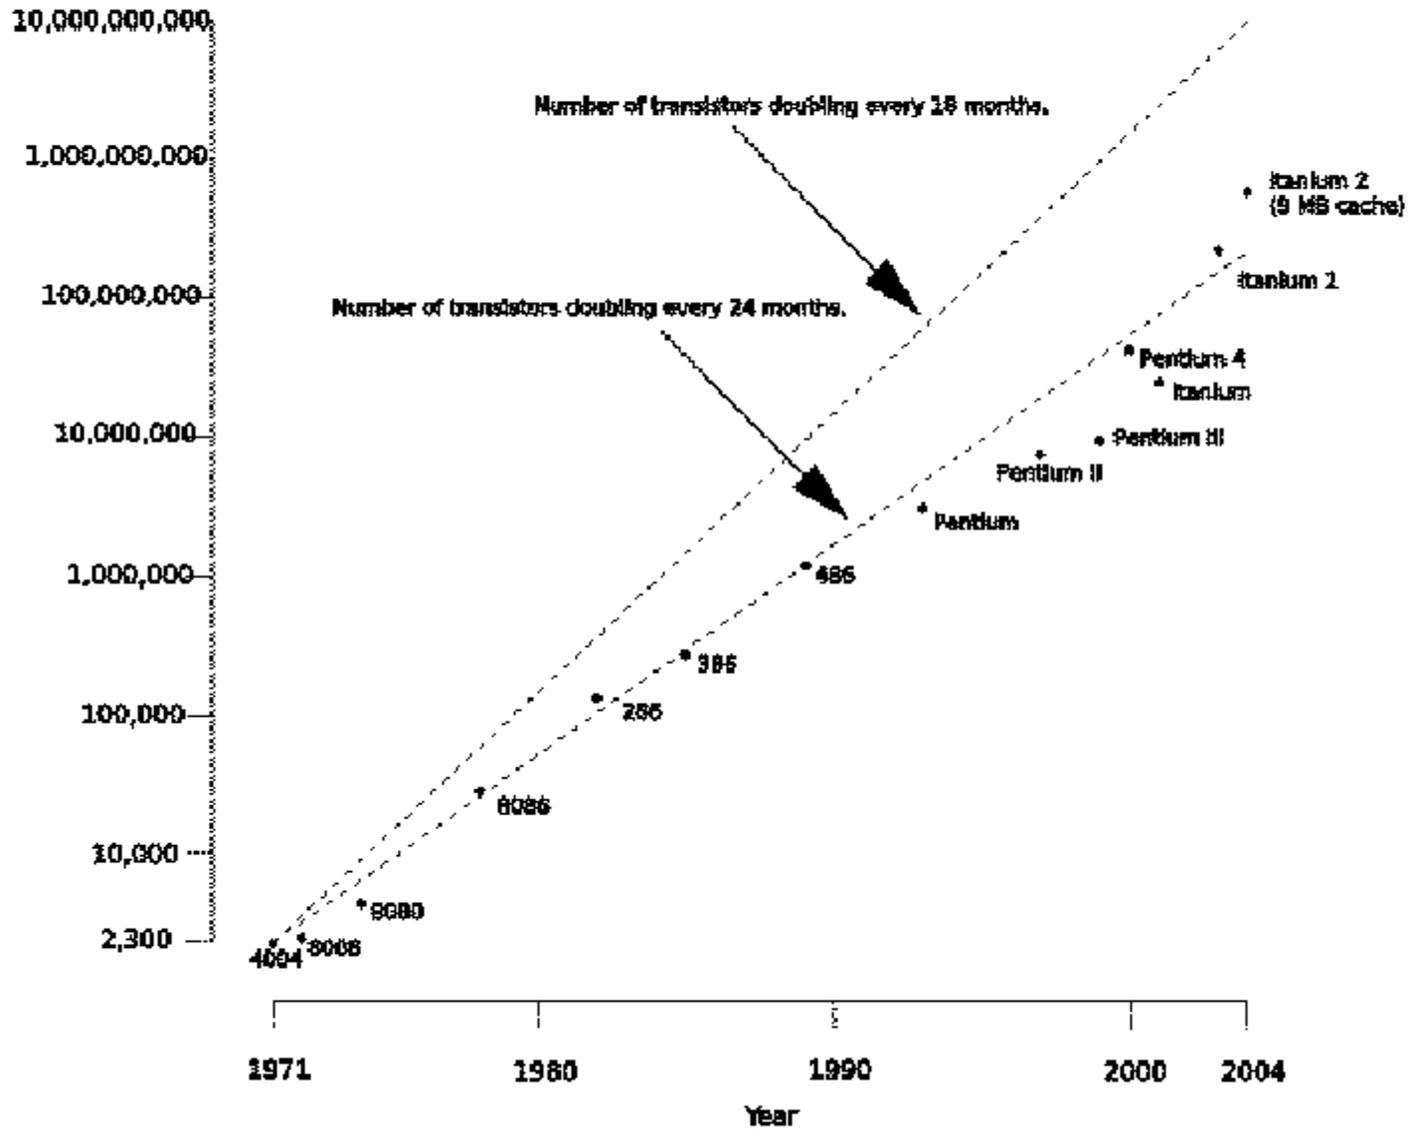


# Big Problem – Information Overload

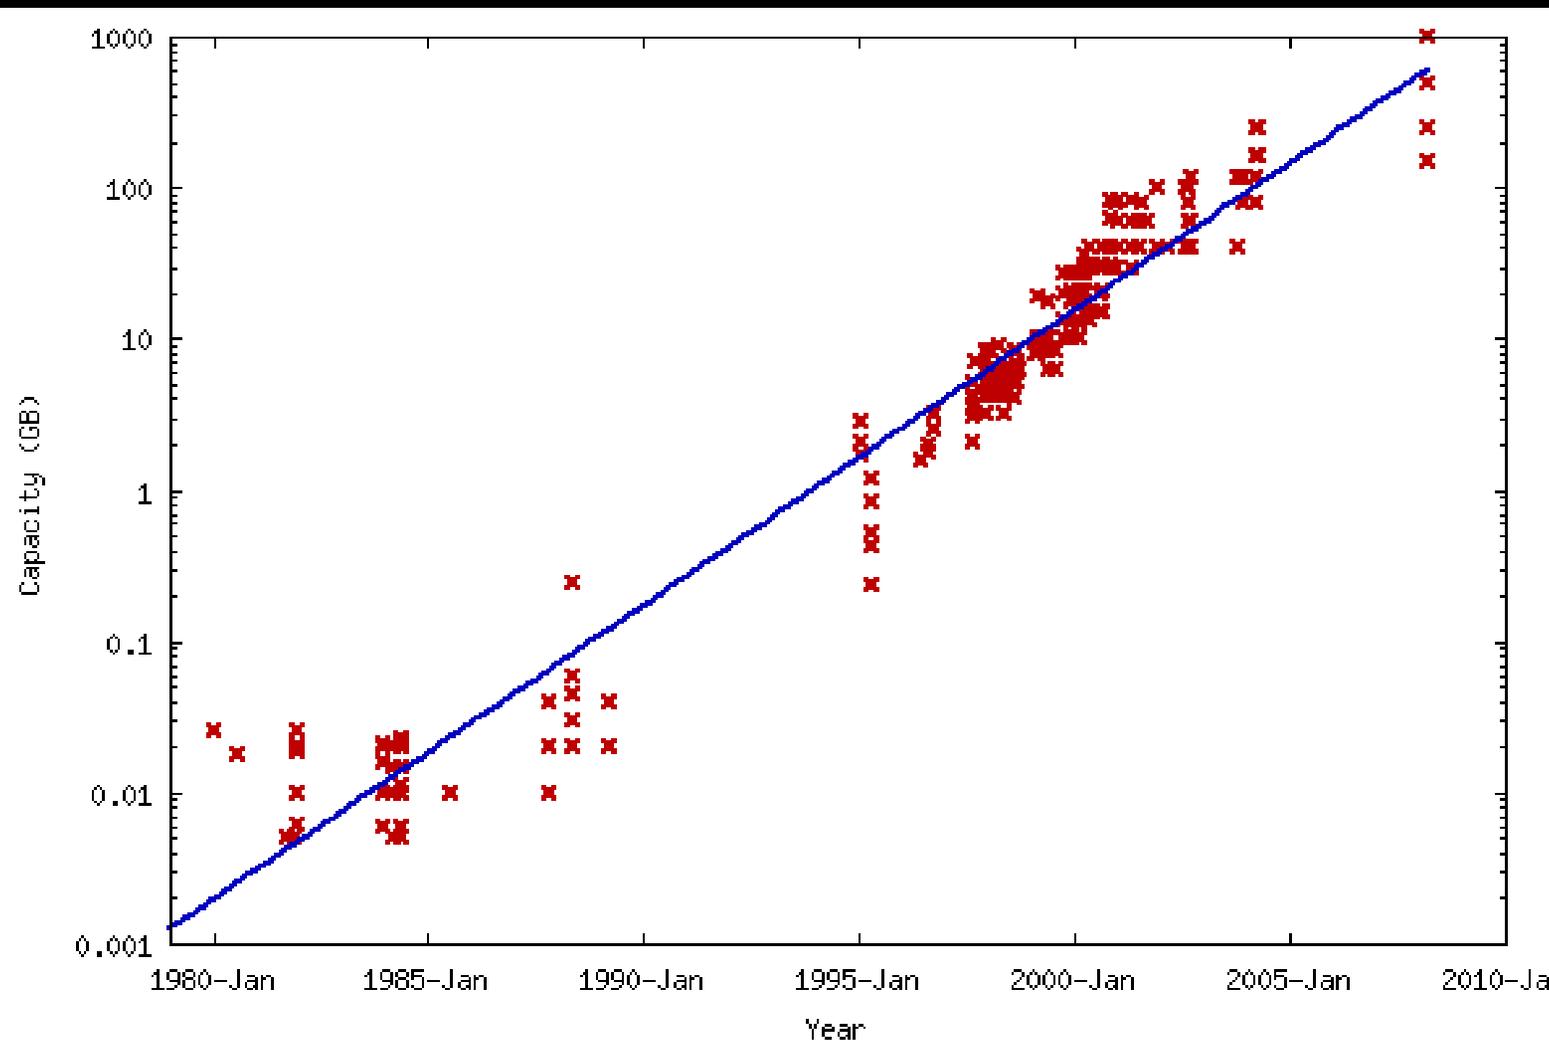
- **Our ability to create and store information exceeds our capacity to understand it.**
- **Information requires attention to process:**
  - “A wealth of information creates a poverty of attention.” – Hebert Simon, Nobel Prize, 1971.
- **Major challenge: gain insight from data.**

# Moore's Law

Number of transistors on an integrated circuit

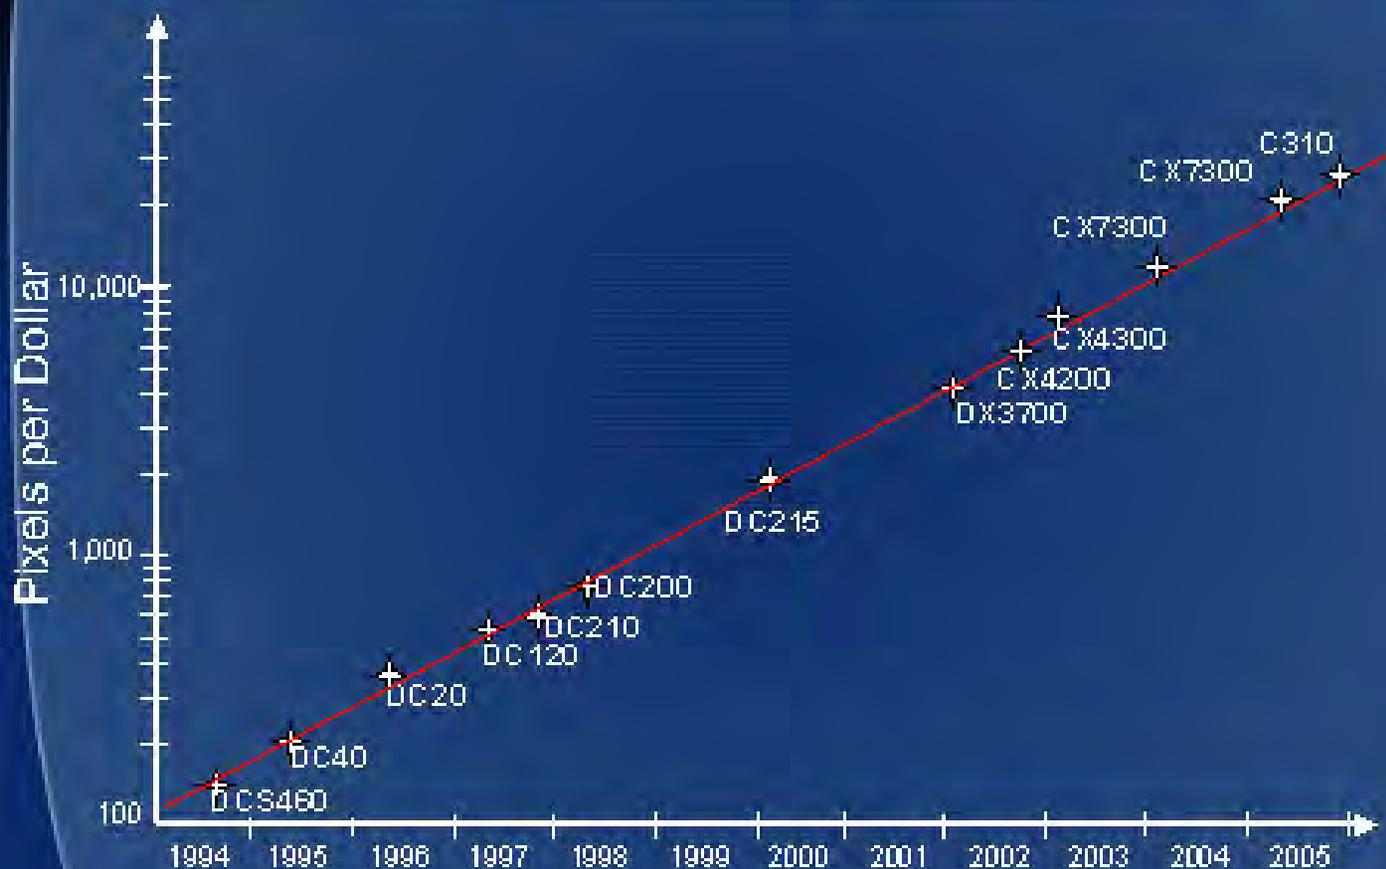


# Moore's Law – PC Hard Drive Capacity



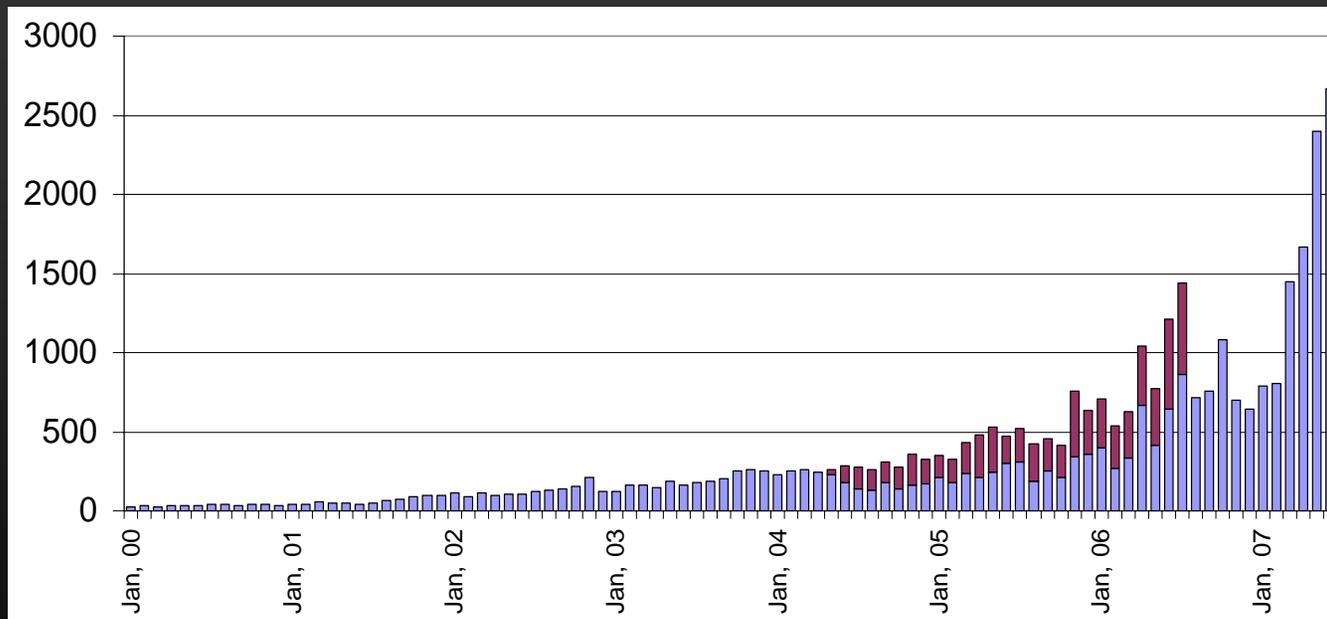
# Moore's Law – Digital Camera Resolution

## The Pixels per Dollar Projection

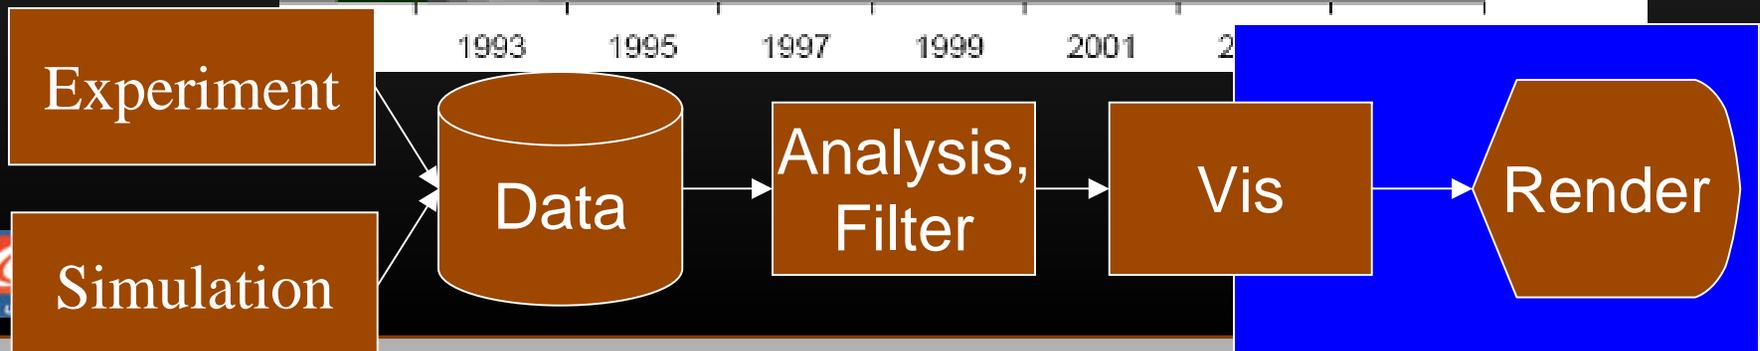
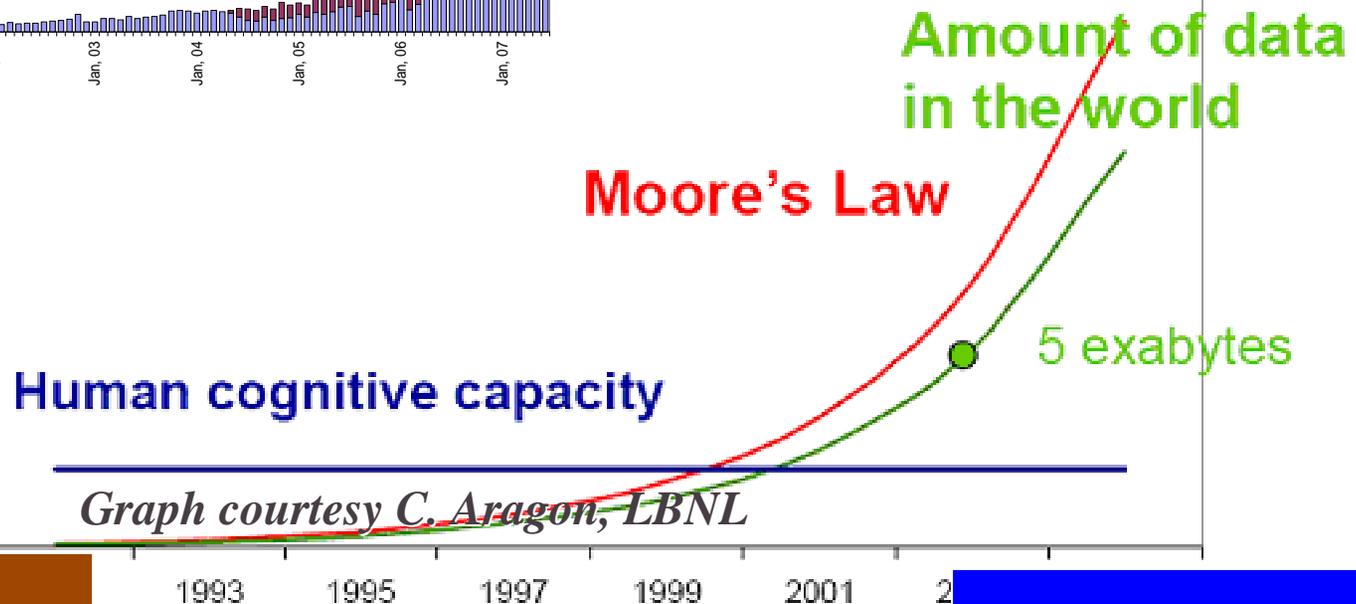
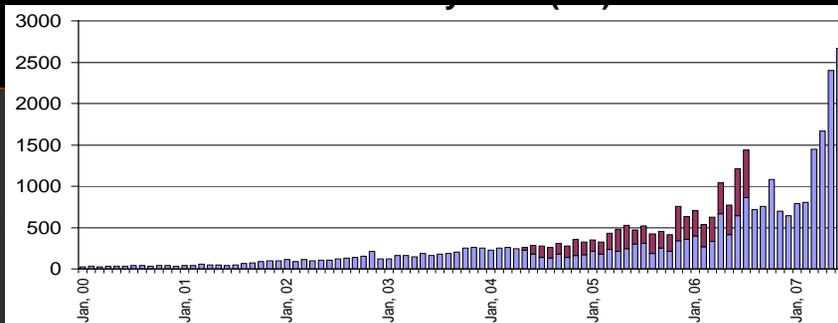


# ESnet: Shatters PB barrier

ESnet Monthly Data (TB)

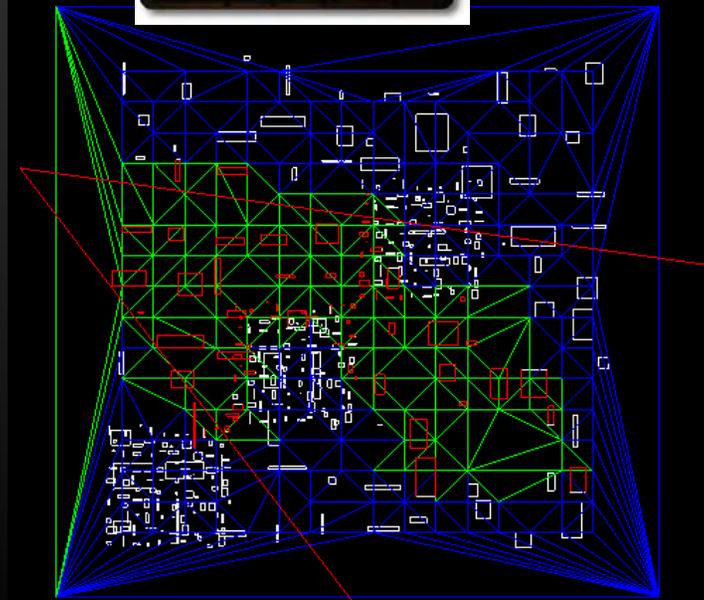
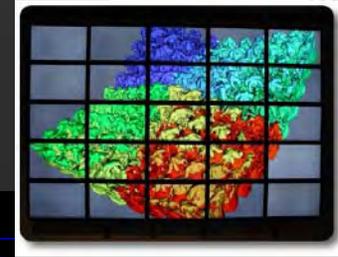


# Human Cognitive Limit



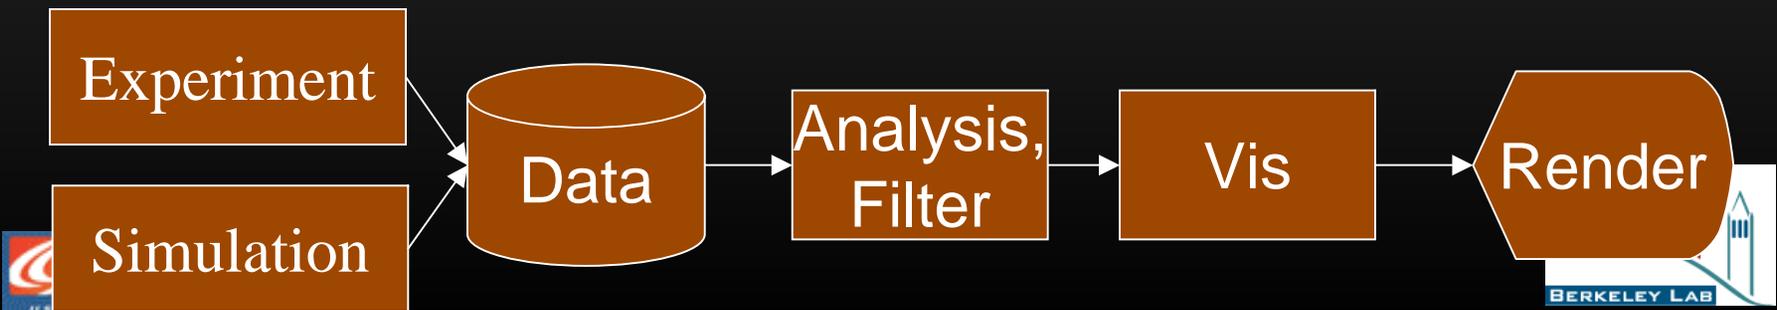
# How is Cognitive Bandwidth Abused?

- The “meat grinder” syndrome
  - Can’t “see” a TB, much less a PB
    - Depth complexity
      - $1k^3$  grid  $\rightarrow O(2-20)$ 
        - » Can’t see it
      - $4k^3$  grid  $\rightarrow O(20-200)$ 
        - » Can’t store it
    - Spatial complexity
      - Mesh cells map to sub-pixels



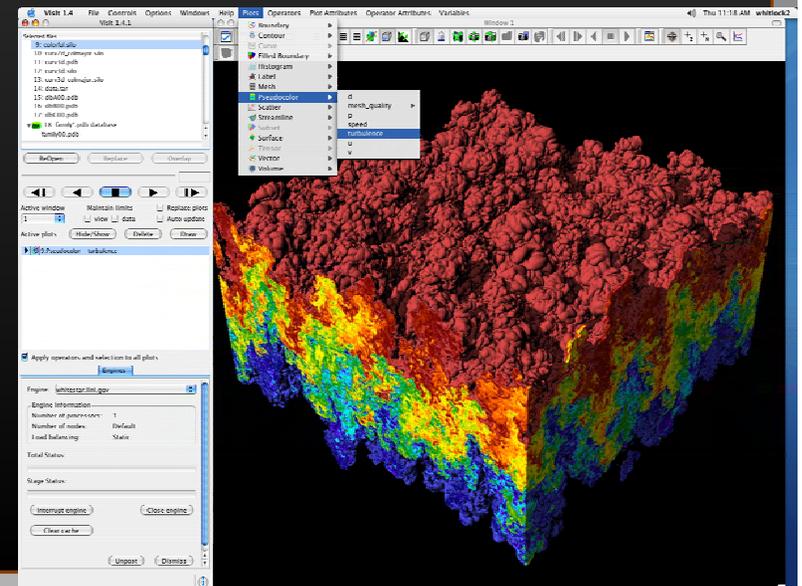
# LBNL Visualization Program

- DOE's SciDAC Visualization and Analytics Center for Enabling Technologies (VACET).
- NERSC Analytics.
- High performance visualization research.



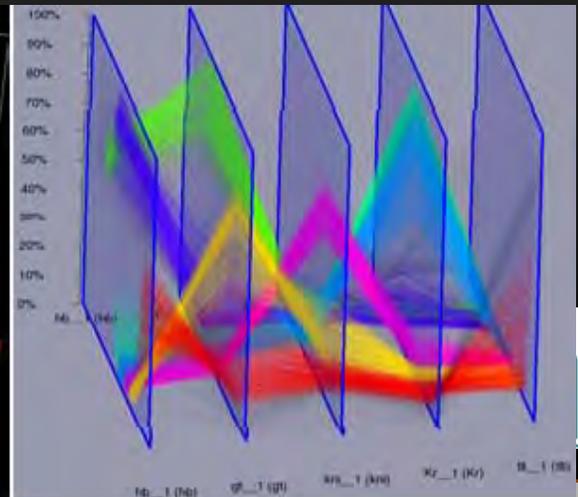
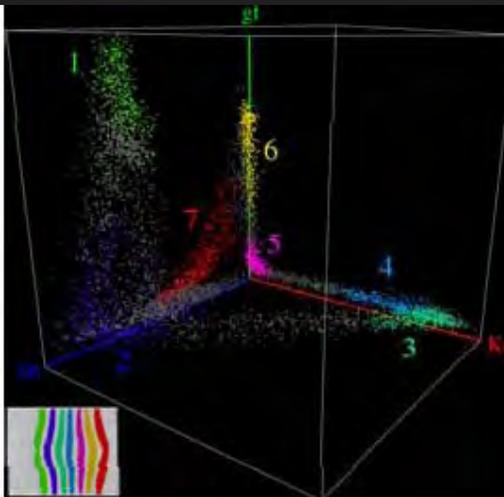
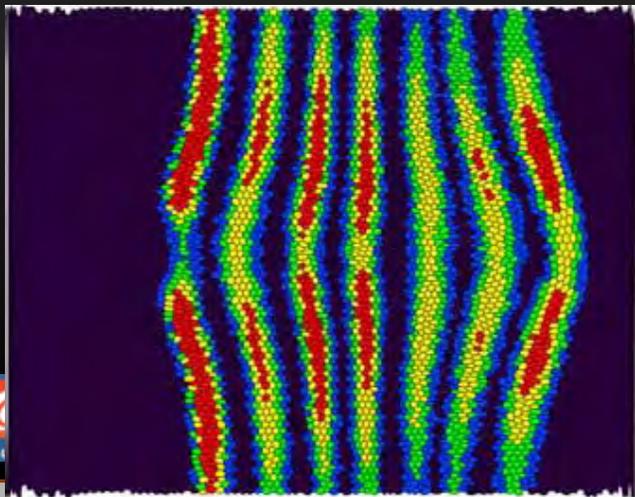
# The Road Ahead (< 5-10 years)

- Effectively leverage emerging computational platforms for visual data analysis.
  - MPP systems, hybrid/hierarchical parallelism.
  - “Exotic systems”:
    - CellBE, GPU, etc.



# The Road Ahead (< 5-10 years)

- **Increasing efficiency of data understanding techniques and methodologies.**
  - Combining visualization and traditional data analysis.



# The Road Ahead (< 5-10 years)

- **Usability engineering.**
  - Objective: simplify ease-of-use of complex machinery.
  - Factors: parallel machines, distributed teams, large and complex data, diverse software, ...
- **A very difficult set of problems!**
  - Data formats “Tower of Babel” problem.
  - Interface standards non-existent:
    - Programming interfaces
    - Data and “execution” interfaces
    - User interfaces



# The End

