Picture This: Visualizing Yesterday, Today and Tomorrow

Kelly Gaither

Director of Visualization/ Senior Research Scientist

Texas Advanced Computing Center

The University of Texas at Austin

What is Visualization?

- “To Envision information is to work at the intersection of image, word, number, art.”

  Edward R. Tufte, Envisioning Information

- “Visualization is any technique for creating from data images, diagrams, or animations to communicate a message.”

  (http://en.wikipedia.org/wiki/Visualization_(computer_graphics))

Computer Graphics versus Visualization

Computer Graphics versus Visualization

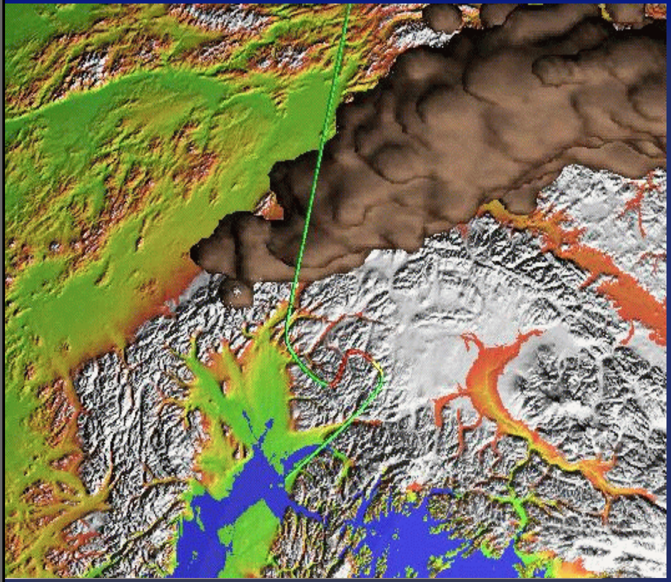

Where Does Technology Fit In?

- We have always used technology to create visualizations of what we see in our minds eye.
- What changes over time is the technology we use to do the visualization.

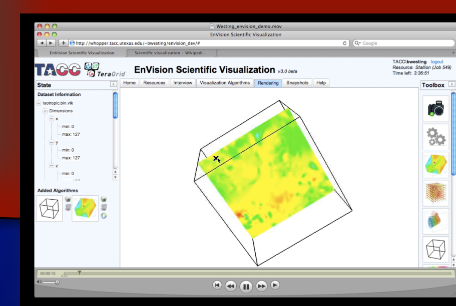

TACC

Visualization Over the Ages

15000 – 10000 BC

1137 AD

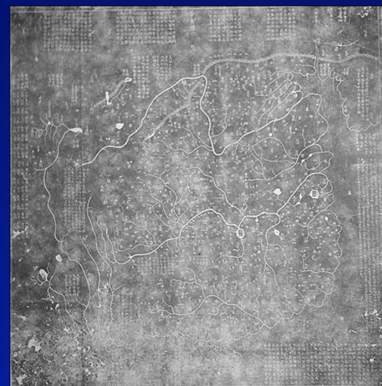

1510

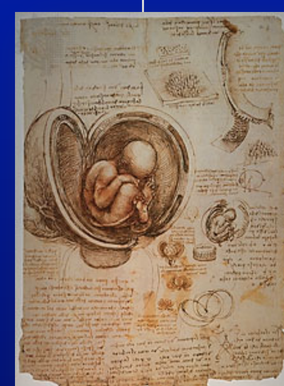

Statistical Graphics

1750 - 1800

1812

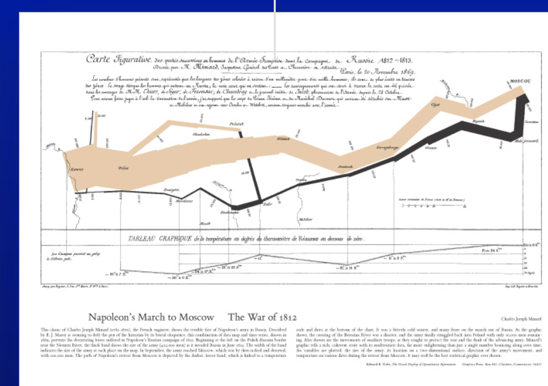

The Birth of Scientific Visualization

1987

The Birth of Visualization as a Field of Science

- The emphasis on visualization as a field of science started in 1987 with a special issue of Computer Graphics on Visualization in Scientific Computing.
- "*Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to observe their simulations and computations. Visualization offers a method for seeing the unseen. It enriches the process of scientific discovery and fosters profound and unexpected insights. In many fields it is revolutionizing the way scientists do science.*" Visualization in Scientific Computing, ACM SIGGRAPH 1987

What Was the Catalyst?

- It’s no coincidence that the first National Science Foundation funded supercomputing centers began in 1985.
- This began the tightly coupled relationship between high performance computing (hpc), data and visualization.



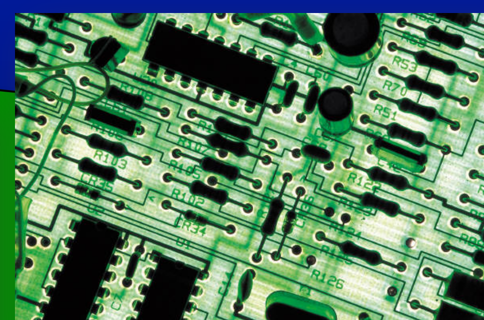


What are Today's Transformational Science Problems?

- What is Transformational Science?
  - In 2007, the Augustine report(*) provided the recommendation to sustain and strengthen the nation's traditional commitment to long-term basic research that has the potential to be transformational to maintain the flow of ideas that fuel the economy, provide security, and enhance the quality of life.

* Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future, National Academy Press, 2007.

Hurricane Prediction

- Forecasters predict the weather using advanced computing technologies.

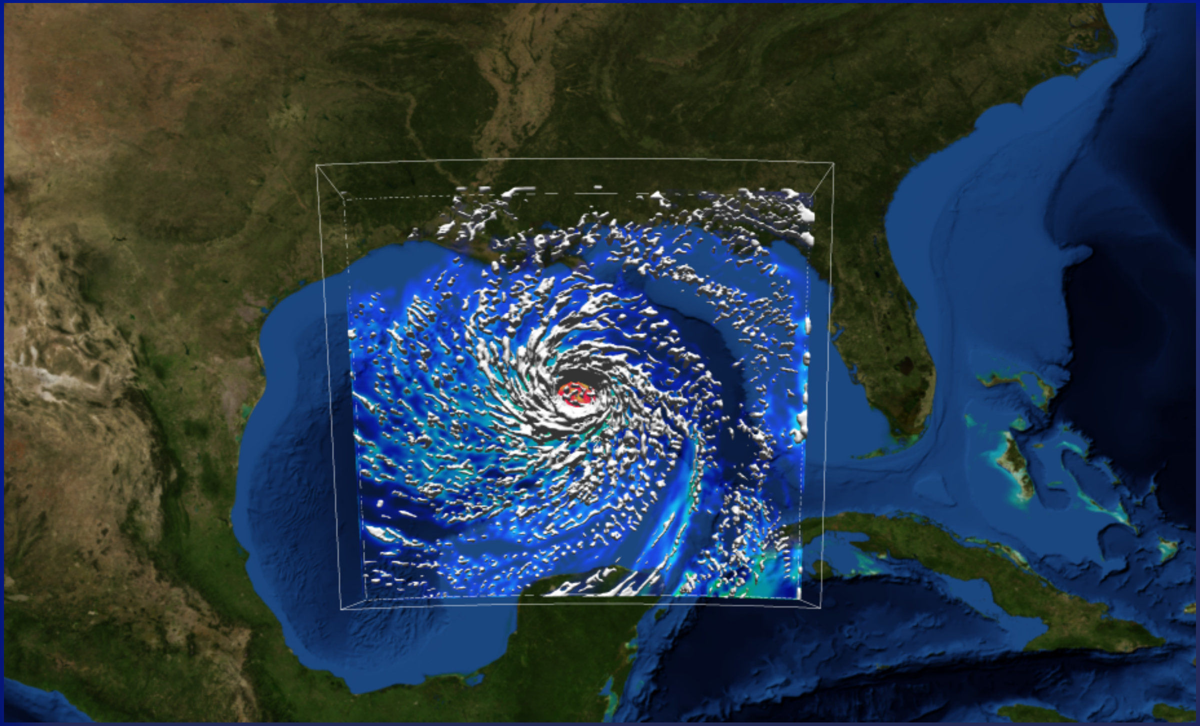

Image: Greg P. Johnson, Romy Schneider, TACC

Vehicle Design

- Your car was probably designed and crash tested in a supercomputer simulation.

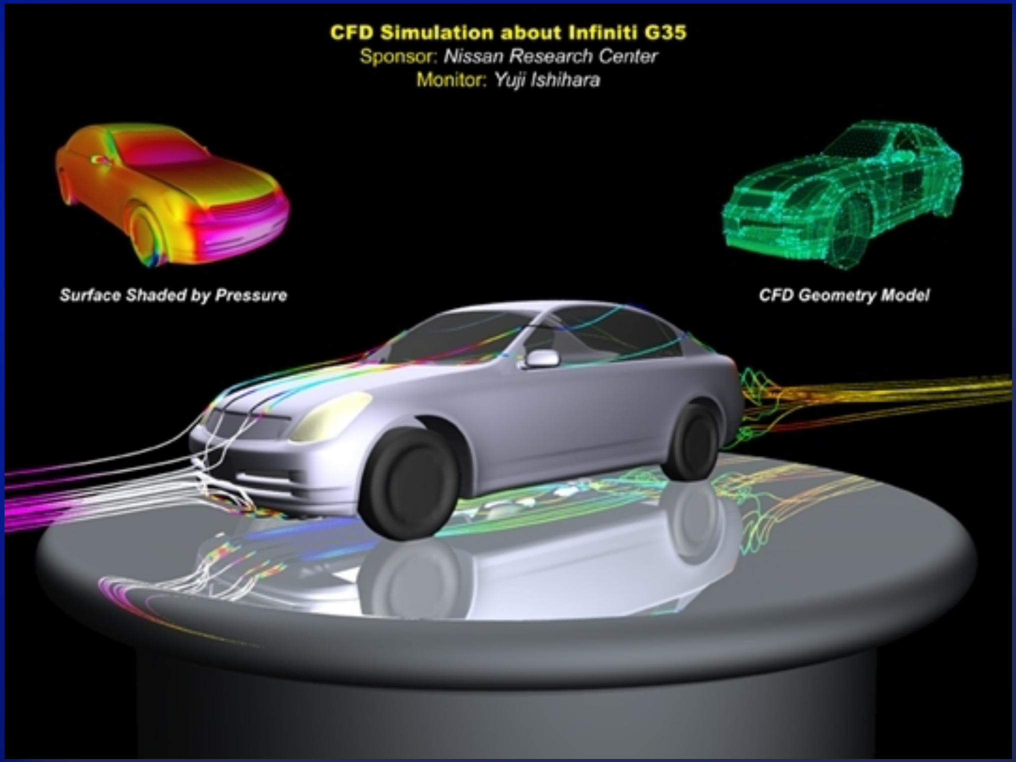

Image courtesy Mississippi State University Simulation and Design Center

BP Oil Spill Response

- Models generated with advanced computing are devoted to simulating the spread of the BP oil spill.

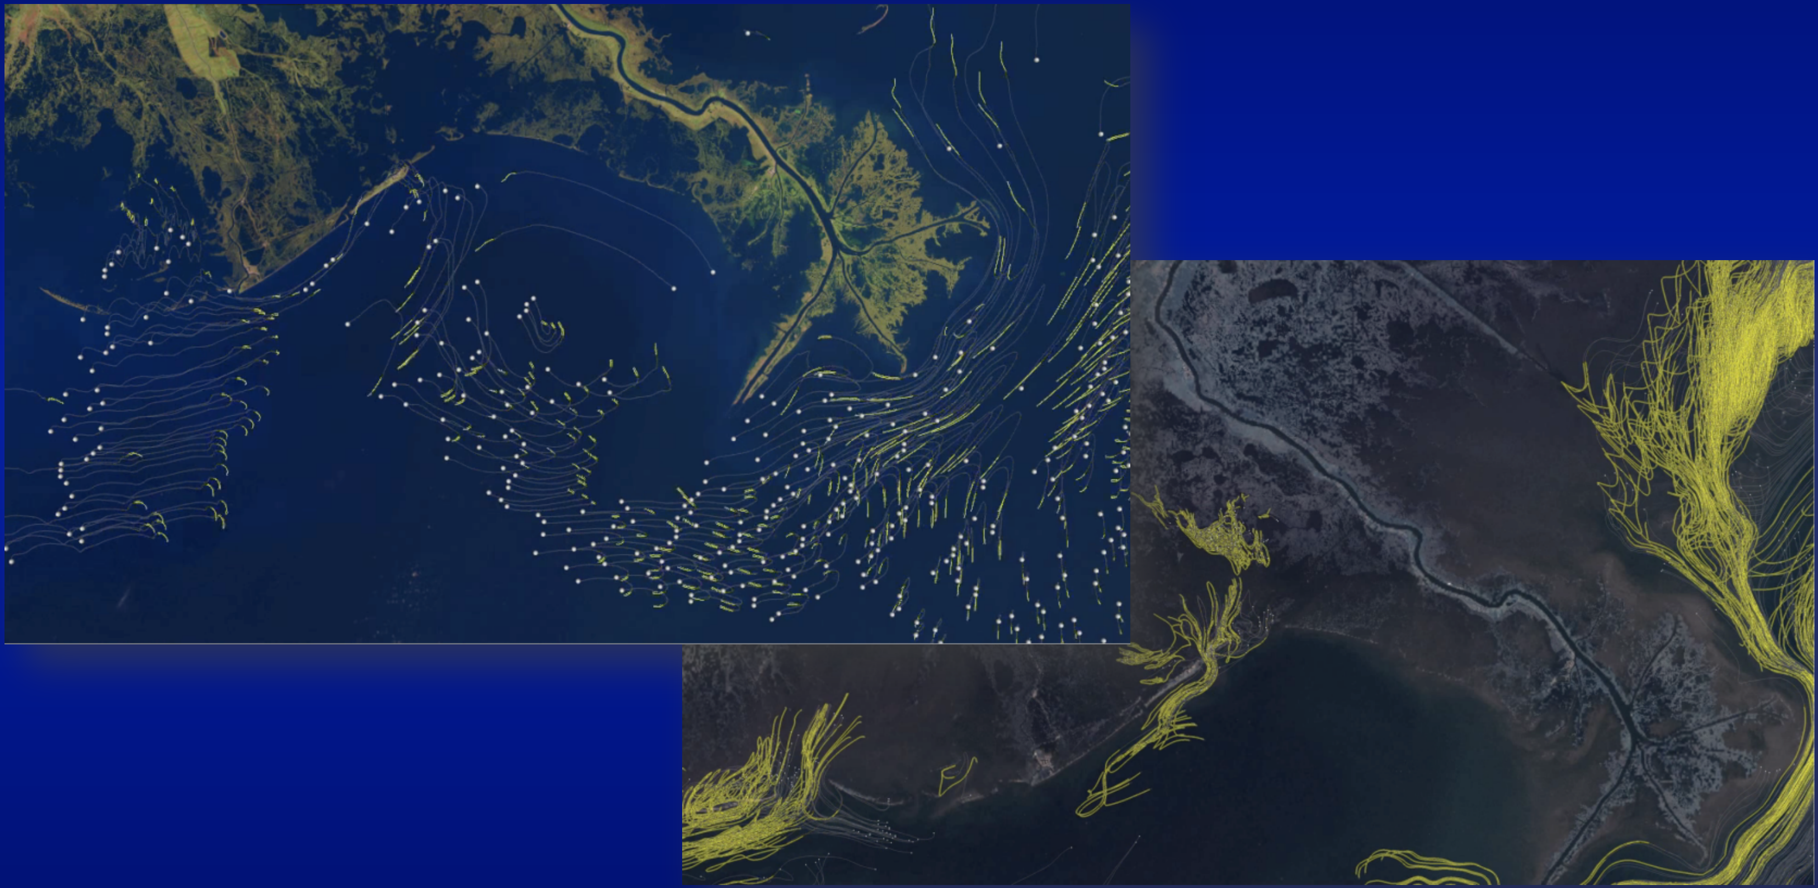

Images Courtesy Adam Kubach, Karla Vega, Clint Dawson

Emergency Planning and Response

- Models generated with advanced computing help plan emergency response to flooding.

TACC

Images Courtesy Greg P. Johnson, TACC and Gordon Wells, Center for Space Research

Coronary Artery Nano-particle Drug Delivery Visualization

Ben Urick, Jo Wozniak, Karla Vega, TACC; Erik Zumalt, FIC; Shaolie Hossain, Tom Hughes, ICES.

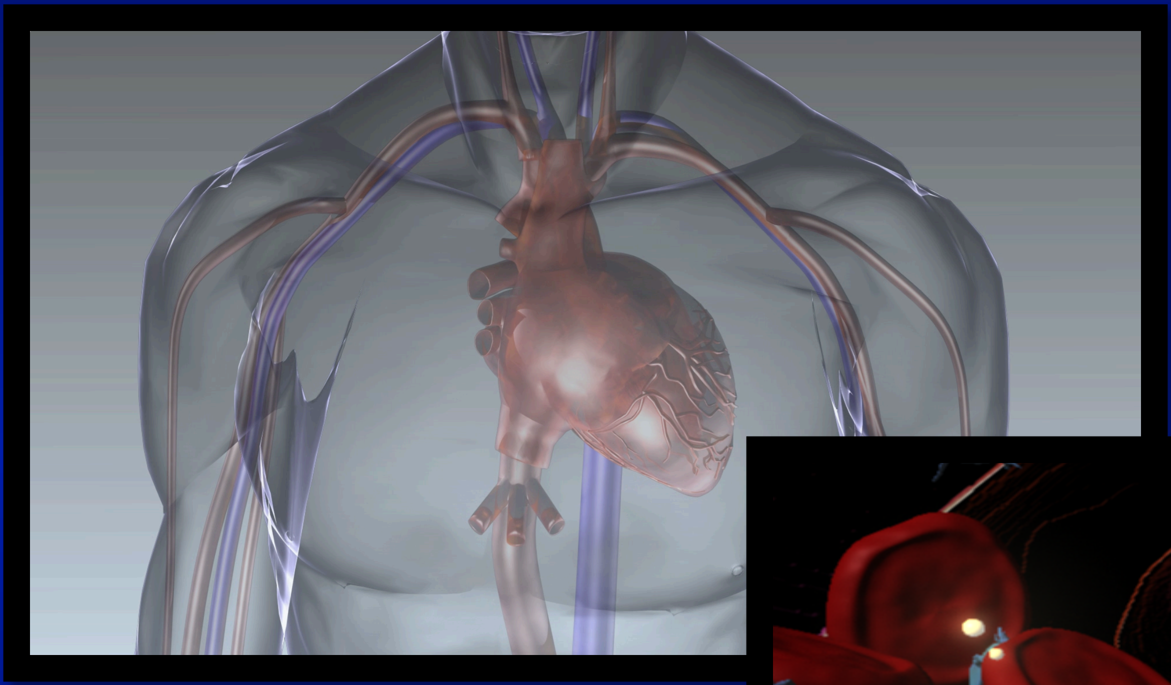

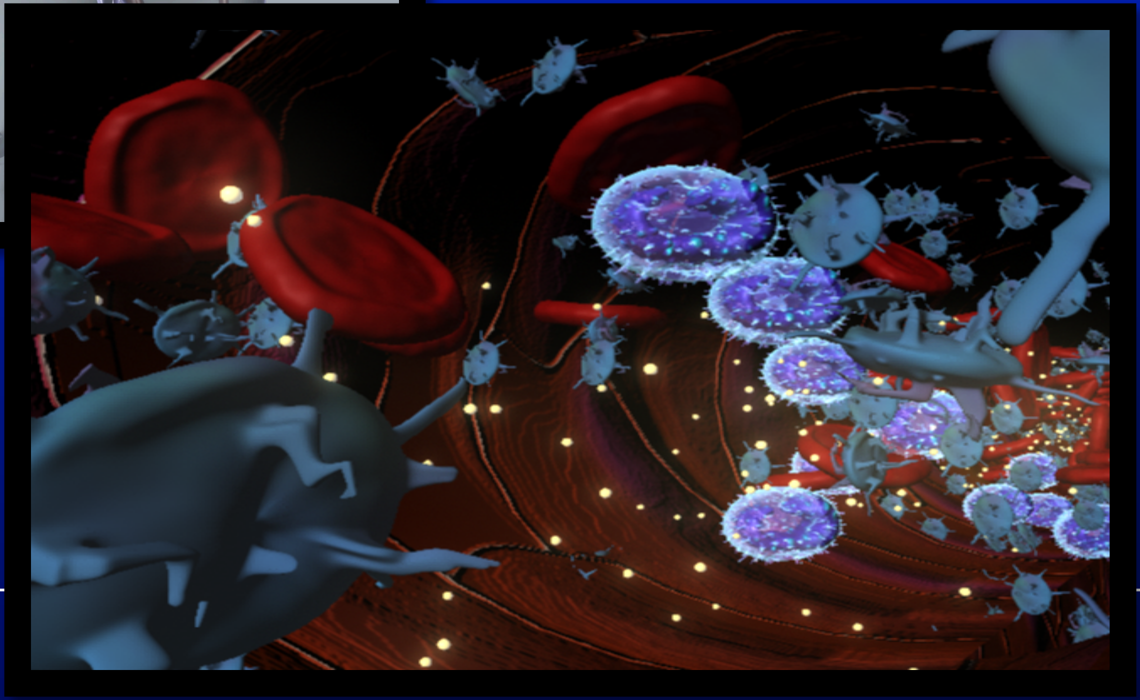

What Role is Visualization Playing in Understanding this Science?

Why Are Pictures So Powerful?

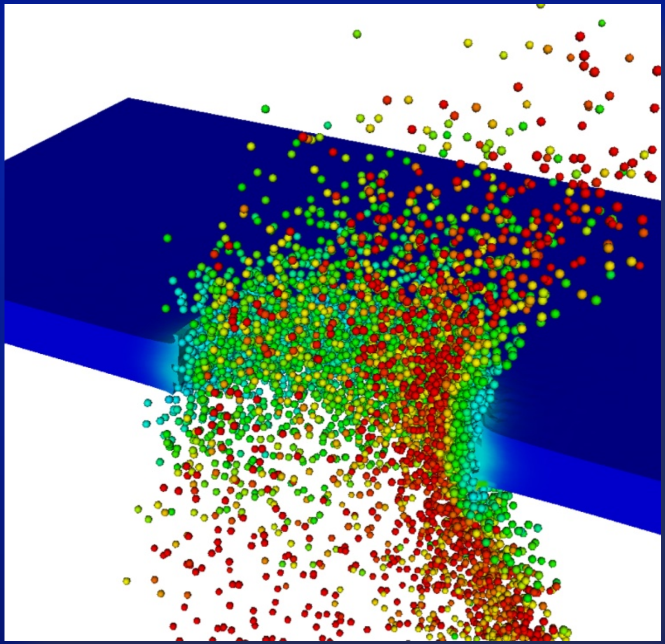

How Critical is Visualization to Science?

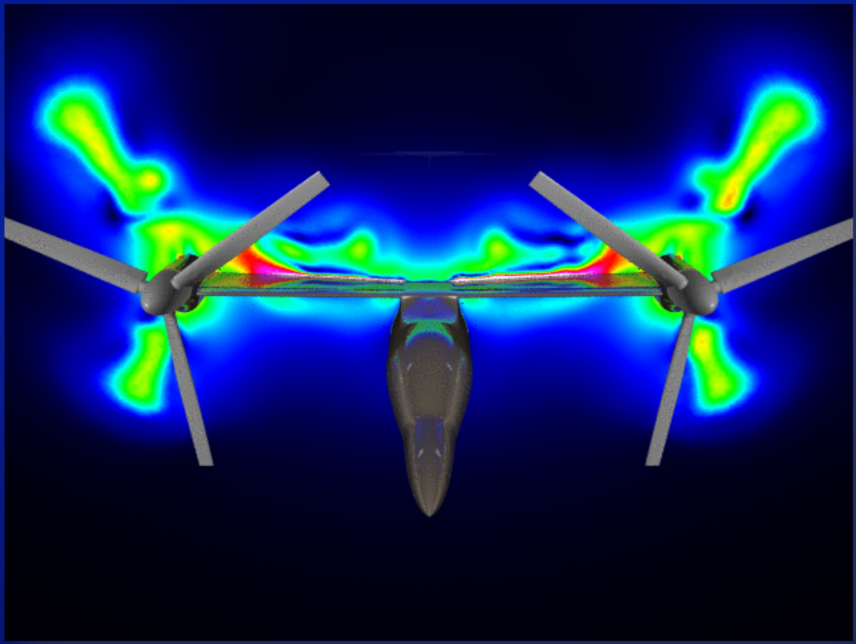

What Role Does the Human Brain Play in Successful Visualizations?

- *"Visualization is so effective and useful because it utilizes one of the channels to our brain that have the highest bandwidths: our eyes. But even this channel can be used more or less efficiently."* [Kosara, 2002]

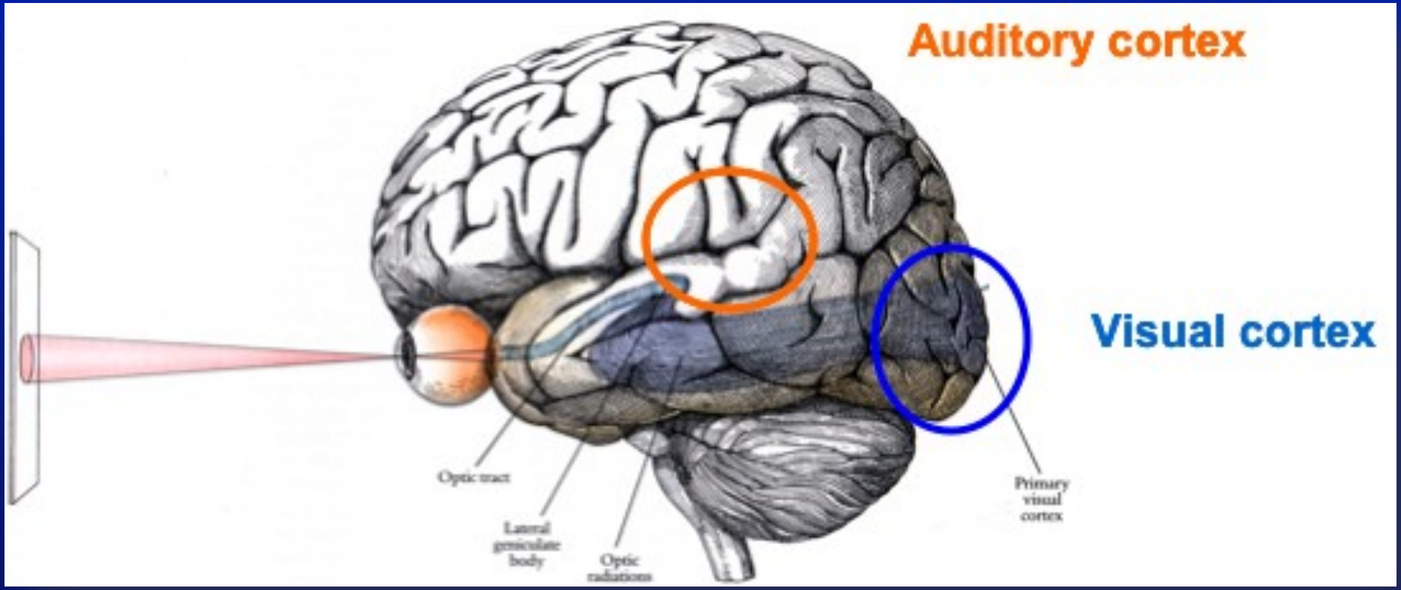

What Role Does the Human Brain Play in Successful Visualizations?

- Our vision encompasses:
  - seeing color
  - detecting motion
  - identifying shapes
  - gauging distance, speed and size
  - seeing objects in three dimensions
  - filling in blind spots
  - automatically correcting distorted information
  - erasing extraneous information that cloud our view.
- Visual cortex makes up 30% of cerebral cortex (77% by volume of human brain) *Trends in Neurosciences, 18:471-474, 1995*
- Touch makes up 8% of cerebral cortex
- Hearing makes up 3% of cerebral cortex

TACC

Discover Magazine, June 1993: "The Vision Thing: Mainly in the Brain"

Visualization at TACC

10 Years in the Making

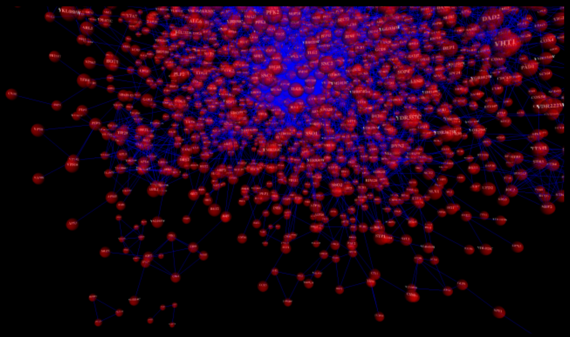

Bioinformatics

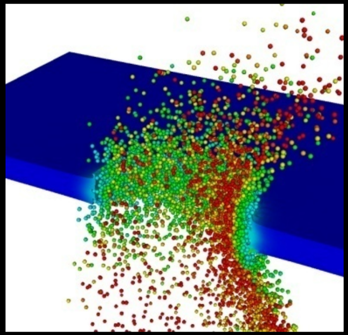

Orbital Debris

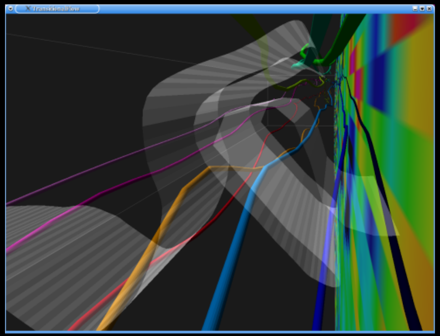

Turbulent Flow

CT Models

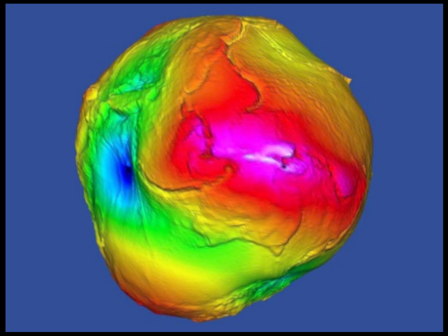

Gravity Map

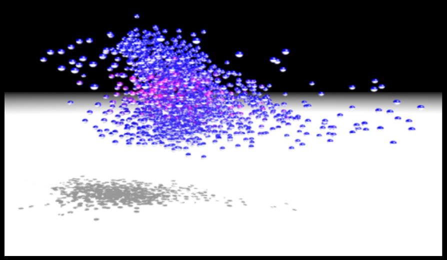

Quantum Chemistry

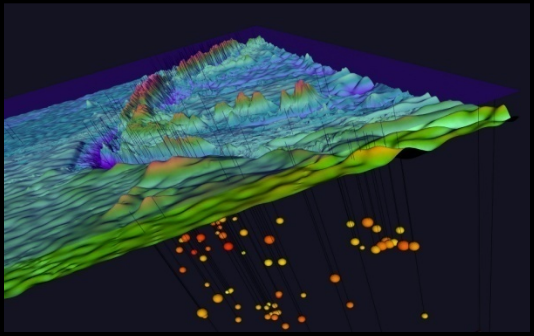

GeoSciences

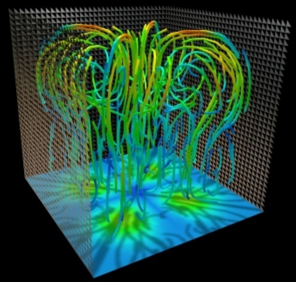

Natural Convection

TACC Visualization Group

- Provides resources/services to a growing local and national user community to enable scientific discovery and insight.

TACC Visualization Group

- Provides resources/services to a growing local and national user community to enable scientific discovery and insight.
- Researches and develops tools/techniques for the next generation of problems facing the user community.

TACC Visualization Group

- Provides resources/services to a growing local and national user community to enable scientific discovery and insight.
- Researches and develops tools/techniques for the next generation of problems facing the user community.
- Trains the next generation of scientists to visually analyze datasets of all sizes.

TACC Visualization Personnel

- 10 Full Time Staff
  - 4 Ph.D.
  - 2 Masters
  - 4 Bachelors
- 4 Students
  - 2 Undergraduate Students
  - 2 PhD Student

Visualization Group at TACC

- Areas of Expertise
  - Scientific & Information Visualization
  - Data Mining & Feature Detection
  - Large Scale GPU Clusters
  - Large Scale Tiled Displays
  - Remote & Collaborative Visualization Tools

Large-Scale GPU Clusters

Spur

Remote, Interactive Visualization System Directly Connected to Ranger

128 cores, 1 TB aggregate memory, 32 GPUs

- spur.tacc.utexas.edu
- 1 fat memory node
  - Sun Fire X4600 server
  - 8 AMD Opteron dual-core CPUs @ 3 GHz
  - 256 GB memory
  - 4 NVIDIA FX5600 GPUs
- 7 other nodes
  - Sun Fire X4440 server
  - 4 AMD Opteron quad-core CPUs @ 2.3 GHz
  - 128 GB memory
  - 4 NVIDIA FX5600 GPUs

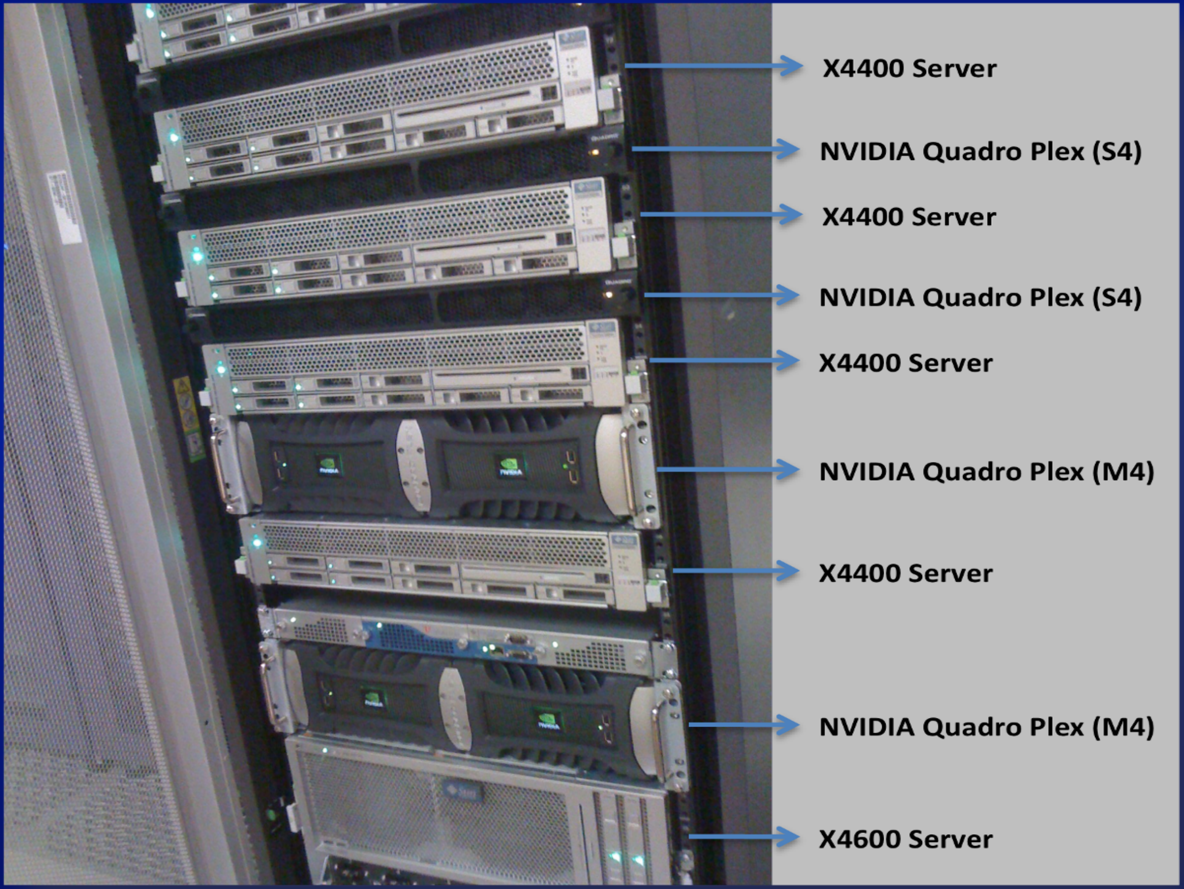

Longhorn

First NSF eXtreme Digital (XD) Visualization Resource

256 Nodes, 2048 Cores, 512 GPUs, 14.5 TB Memory

- 256 Dell Dual Socket, Quad Core Intel Nehalem Nodes
  - 240 with 48 GB shared memory/node (6 GB/core)
  - 16 with 144 GB shared memory/node (18 GB/core)
  - 73 GB Local Disk
  - 2 Nvidia GPUs/Node (FX 5800 – 4GB RAM)
- ~14.5 TB aggregate memory
- QDR InfiniBand Interconnect
- Direct Connection to Ranger's Lustre Parallel File System
- 10G Connection to 210 TB Local Lustre Parallel File System
- Jobs launched through SGE

Kelly Gaither (PI), Valerio Pascucci, Chuck Hansen, David Ebert, John Clyne (Co-PI), Hank Childs

TACC

What Was the Enabling Technology?

Graphics Performance

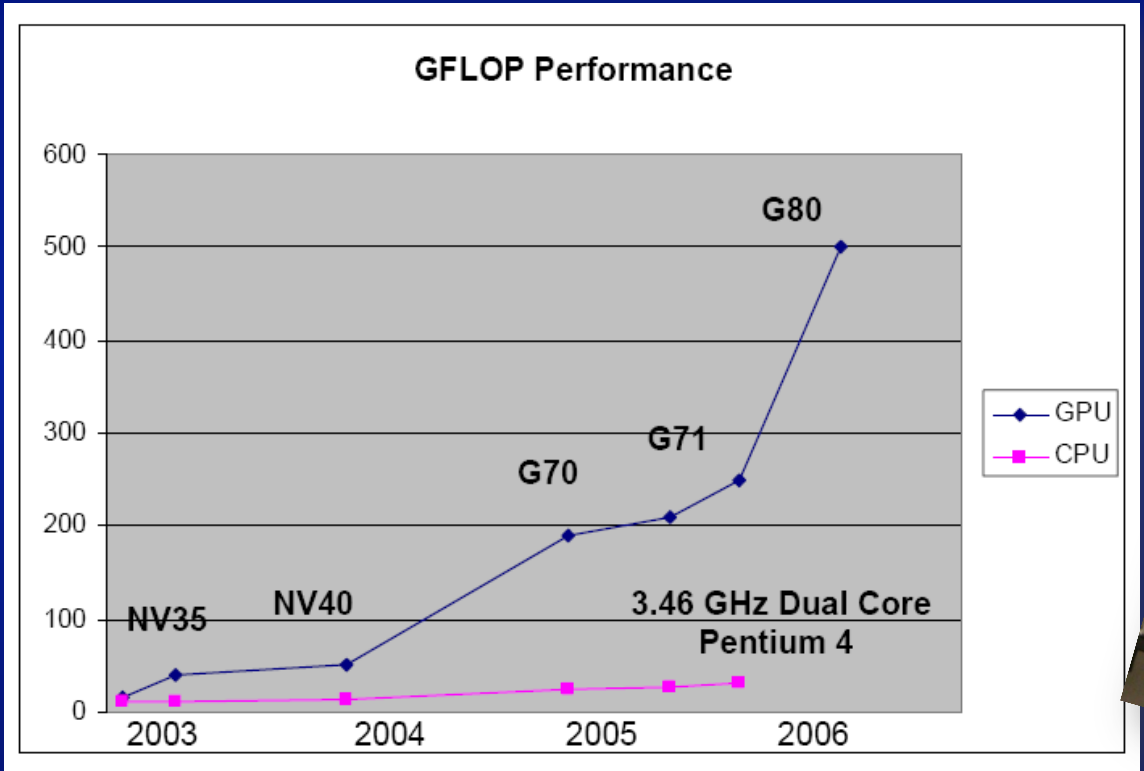

Image Courtesy Klaus Mueller, Stony Brook University

TACC

Longhorn Usage Modalities:

- Remote/Interactive Visualization
  - Highest priority jobs
  - Remote/Interactive capabilities facilitated through VNC
  - Run on 3 hour queue limit boundary
- GPGPU jobs
  - Run on a lower priority than the remote/interactive jobs
  - Run on a 12 hour queue limit boundary
- CPU jobs with higher memory requirements
  - Run on lowest priority when neither remote/interactive nor GPGPU jobs are waiting in the queue
  - Run on a 12 hour queue limit boundary

Software Available on Longhorn

- Programming APIs: OpenGL, vtk (Not natively parallel)
  - OpenGL – low level primitives, useful for programming at a relatively low level with respect to graphics
  - VTK (Visualization Toolkit) – open source software system for 3D computer graphics, image processing, and visualization
  - IDL
- Visualization Turnkey Systems
  - VisIt – free open source *parallel* visualization and graphical analysis tool
  - ParaView – free open source general purpose *parallel* visualization system
  - VAPOR – free flow visualization package developed out of NCAR
  - EnSight – commercial turnkey *parallel* visualization package targeted at CFD visualization
  - Amira – commercial turnkey visualization package targeted at visualizing scanned medical data (CAT scan, MRI, etc..)

Large-Scale Tiled Displays

ACES Visualization Laboratory

Campus Presence for Collaborative Visualization

- Multi-user space with reserveable resources.
- Seamless environments from laptops to large-scale displays.
- Provides large pixel count displays and a collaboration room.
- Reconfigurable, flexible environment that can be used in a variety of ways.

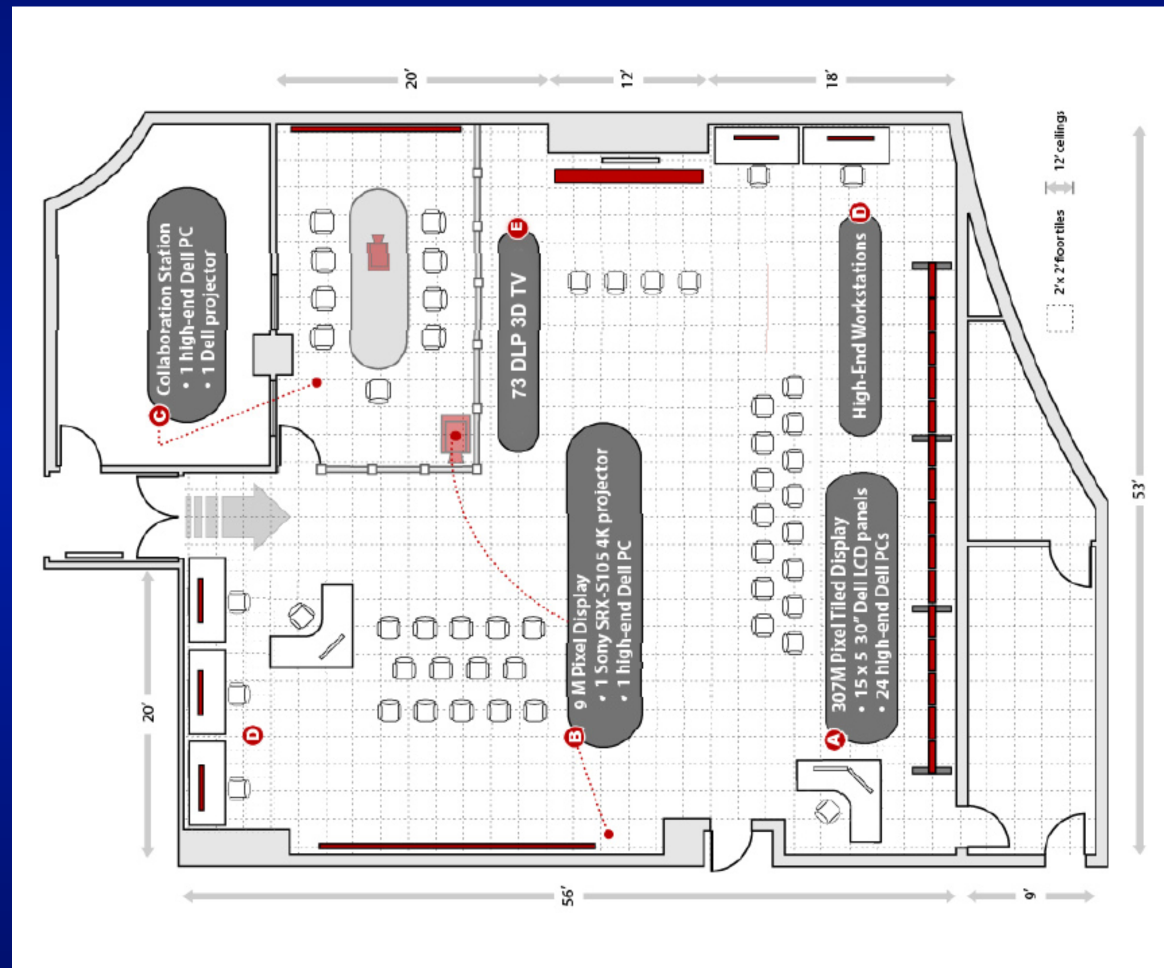

TACC

Stallion

- 15x5 tiled display of Dell 30-inch flat panel monitors
- 307M pixel resolution, 4.7:1 aspect ratio
- 100 processing cores with over 36GB of graphics memory and 108GB of system memory
- 6TB shared file system

TACC

What Was the Enabling Technology?

TACC

Display Performance

Stallion – currently world's highest-resolution tiled display

Dell 30" LCD

307 Megapixels
38400 x 8000 pixel resolution

Downtown Austin @ 1 Gigapixel (77263 x 14225)

HDR Photograph taken by Ricardo Mileschi, Austin TX

TACC

College of Education Visualization Cluster

- *Collaboration with the College of Education will help UT become a leader in education visualization and data analysis leading to improved curricula, instruction, testing, identification of issues and opportunities*

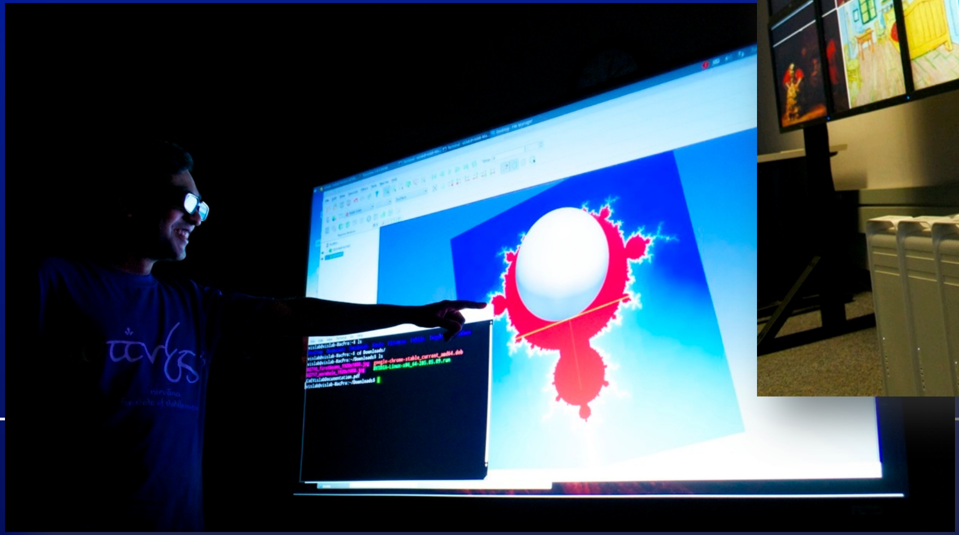

TACC

DisplayCluster

- A cross-platform software environment for interactively driving tiled displays
- Features:
  - Media display (Images (up to gigapixels in size, movies / animations
  - Pixel streaming (Real-time desktop streaming for collaboration / remote vis)
  - Scriptable via Python interface
  - Multi-user interaction (iPhone / iPad / Android devices, Joysticks, Kinect (in development))
  - Implementation (MPI, OpenGL, Qt, FFMPEG, Boost, TUIO, OpenNI, …)
- Short demonstration: http://www.youtube.com/watch?v=JwTwa46BhcU

Lasso
Multi-Touch Tiled Display

- 3x2 tiled display (1920x1600) – 12 MPixels
- PQ Labs multi-touch overlay with 32 point 5mm touch precision
- 11 mm bezels on the displays

TACC

Multi-touch Display Research

- Application of multi-touch gestures to visualizations
- Development of device-agnostic multi-touch protocols for interaction from all devices (phone, table, touchscreen, kinect)
- Evaluation of multi-touch device technologies: optical infrared detection, image processing detection

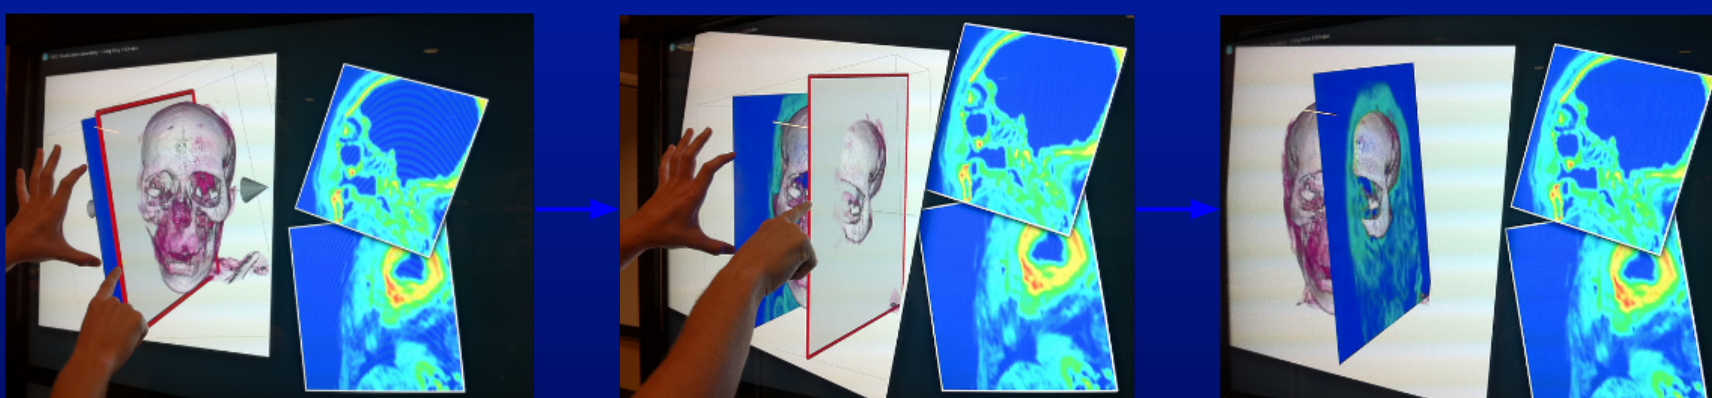

TACC

B. Westing, B. Urick, M. Esteva, F. Rojas, K. Vega, K. Gaither, and W. Xu., “Using Novel Touch Interface Techniques to Visualize and Interact with Digital Records Collections,” Submitted to EuroVis 2012.

Skeletal Tracking for Display Control

- Development of mapping algorithms to map skeletal gestures to large display systems
- Software development for creating effective multi-person skeletal tracking applications that communicate events over network protocols

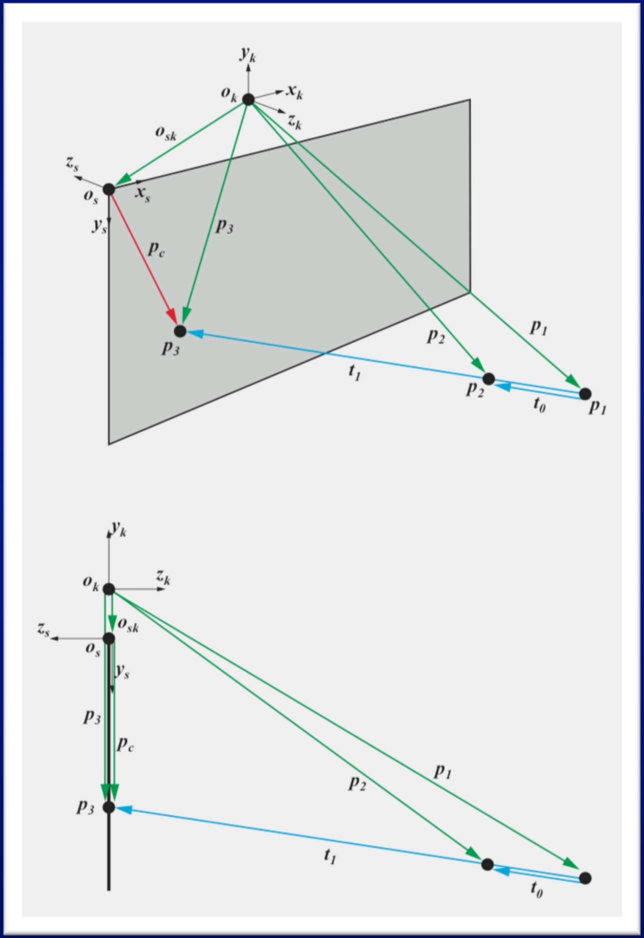

B. Westing, B. Urick, G. Johnson, J. Barbosa, and K. Gaither, "Control of Display Environments with Depth-based Skeletal Tracking," Submitted to CHI 2012.

Remote and Collaborative Visualization

Connecting to Longhorn/Spur Using VNC

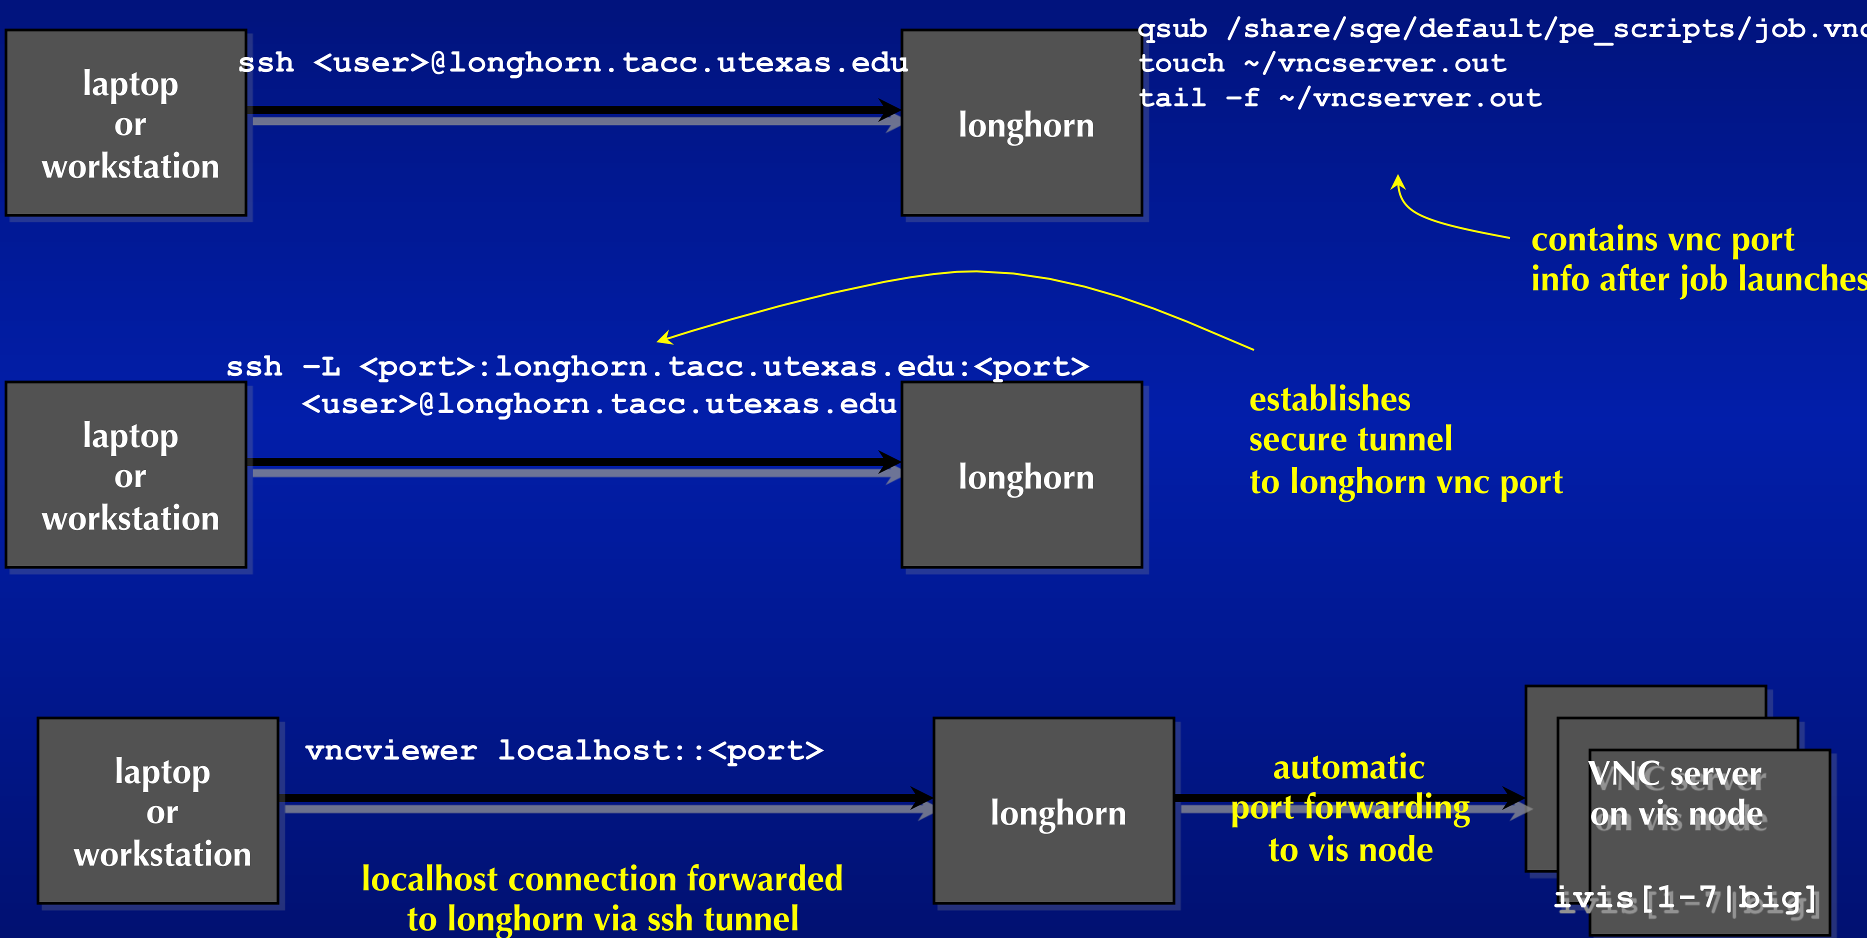

Longhorn Visualization Portal

portal.longhorn.tacc.utexas.edu

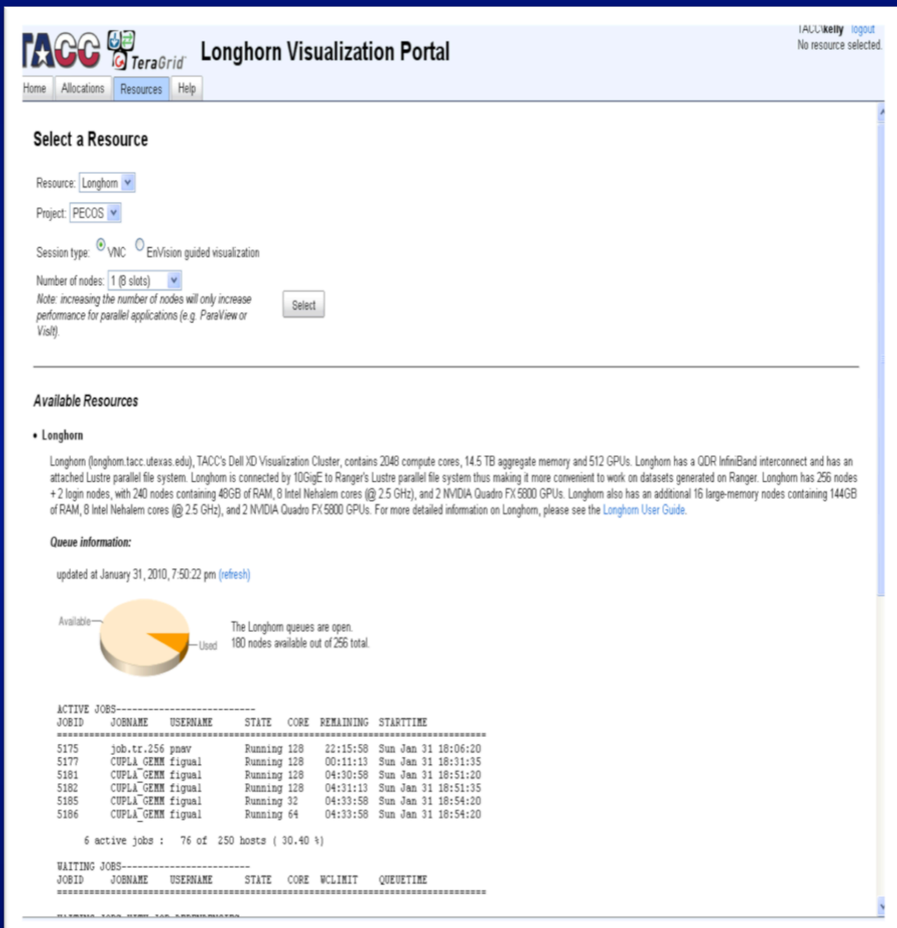

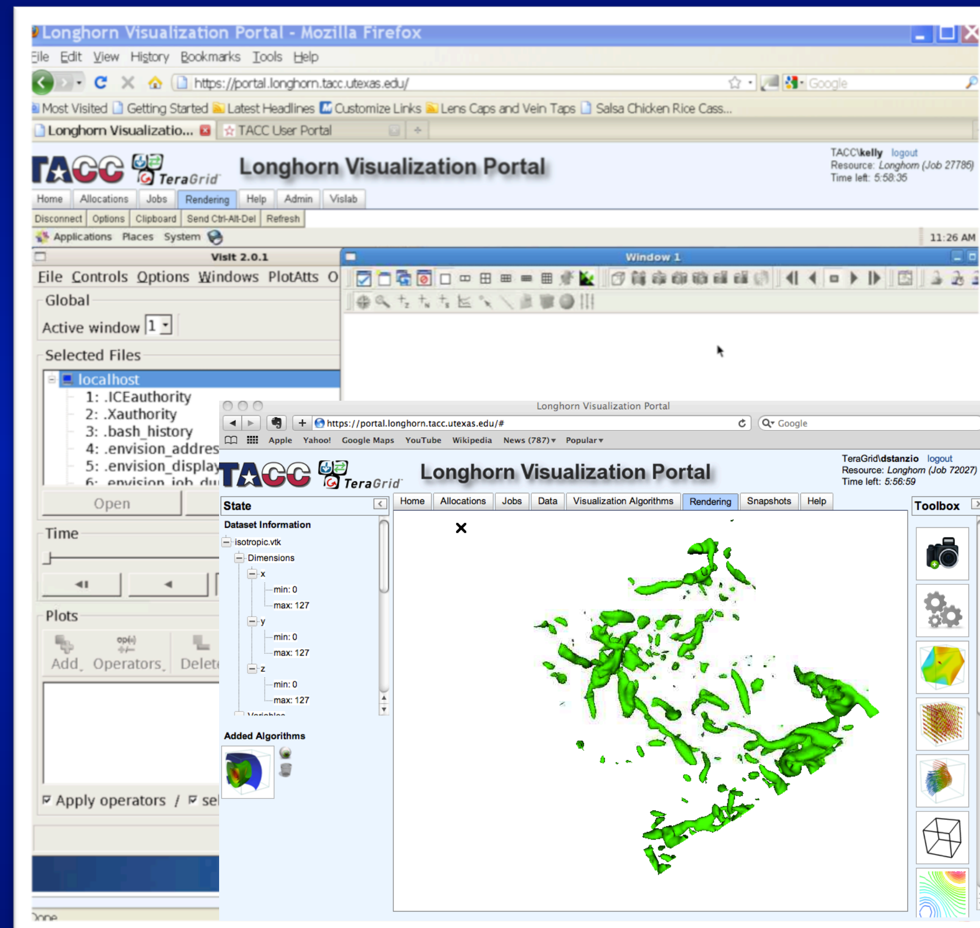

Lessons Learned Over the Past 10 Years

- Close collaborations with the science partners are key.
- Minimize data transfers if possible.
- Choose rendering techniques that best suit the application – is photorealism necessary?
- Scale resources effectively based on use cases.
- Easy accessibility to and interaction with technologies encourages participation from diverse communities.

Thoughts Towards Exascale:

- Data will get larger and more unwieldy – it will stop being moved around.
- High performance computing environments will become high performance science environments that provide computing and analytics.
- Computer graphics will be an equal guest at the high performance table.

Thoughts Towards Exascale:

- Rendering will continue to get less and less expensive.
- We will see a real blend in high performance environments of physical modeling and computer graphics.

Stampede

- 10PF+ Peak performance in initial system (2013)
  - 2PF conventional cluster - Intel Sandy Bridge processors, Dell dual-socket nodes w/32GB RAM, > 6,000 nodes, > 100,000 cores
  - 8PF accelerated co-processor system - Intel MIC processors, > 300,000 cores
- 14PB+ disk
- 200TB+ RAM
- 56Gb FDR InfiniBand interconnect
- Integrated shared memory and remote visualization
  - Stampede will include 16 1TB Sandy Bridge shared memory nodes with dual GPUs.
  - 128 of the compute nodes will be equipped with NVIDIA Kepler-class GPUs