An Adaptive, Scalable, and Secure I2-based Client-server Architecture for Interactive Analysis and Visualization of Volumetric Time Series Data

(The 4D Visible Mouse Project)

NLM Scaleable Information Infrastructure Awards

Contract # N01-LM-3-3510

NRBSC

PITTSBURGH SUPERCOMPUTING CENTER

Outline

- Goals of project
- Project members
- Infrastructure
- Software components
- Research Applications
- Demonstration
- Focus on mouse LV measurement
- Summary of impact

Project goals:

- Provide technology for serving and viewing large 4D datasets
- Provide secure data installation and access
- Provide networked tools for 4D data visualization and analysis
- Provide I2 accessible online data repository
- Evaluate the effectiveness of resulting tools and techniques for usability, effectiveness and applicability to other areas.

An Adaptive, Scalable, and Secure I2-based Client-Server Architecture for Interactive Analysis and Visualization of Volumetric-Time Series Data

- PSC TEAM
  - Arthur W. Wetzel - PI & System Architecture
  - David W. Deerfield II - co-PI and project management
  - Stuart Pomerantz - PSC Volume Browser Client
  - Démian Nave - Mesh/Model Construction
  - Chris Rapier - Network performance & programming
  - Matt Mathis - Web100/Networking
  - Anjana Kar - Systems Administration
  - Silvester Czanner - 3D Registration & Alignment

- Subcontract PIs
  - G. Allan Johnson - Duke Center for in Vivo Microscopy
  - Cynthia Gadd - University of Pittsburgh (CBMI)

- Additional CIVM Team
  - Cristian T. Badea, PhD, Assistant Research Professor, Radiology
  - Jeffrey Brandenburg, PhD, Software Engineer
  - Nilesh Mistry, PhD candidate, Duke Biomedical Engineering
  - Lucy Upchurch, Computer Systems and Network Manager
  - Sally Gewalt, MS, Software Applications & Visualization
  - Alexandra Badea, PhD, Post-doctoral Research Associate
  - Anjum Ali, PhD candidate: Duke Biomedical Engineering
  - Alexandra Petiet, PhD candidate, Duke Biomedical Engineering
  - Michael Fehnel, BS, Project Manager, IT Manager
  - Prachi Pandit, PhD candidate, Duke Biomedical Engineering
- Additional Evaluation Team
  - Valerie Monaco, PhD, Assistant Professor, Dept. of Biomedical Informatics
  - Robb Wilson, Evaluation Project Manager, Dept. of Biomedical Informatics

Client-Server System Architecture

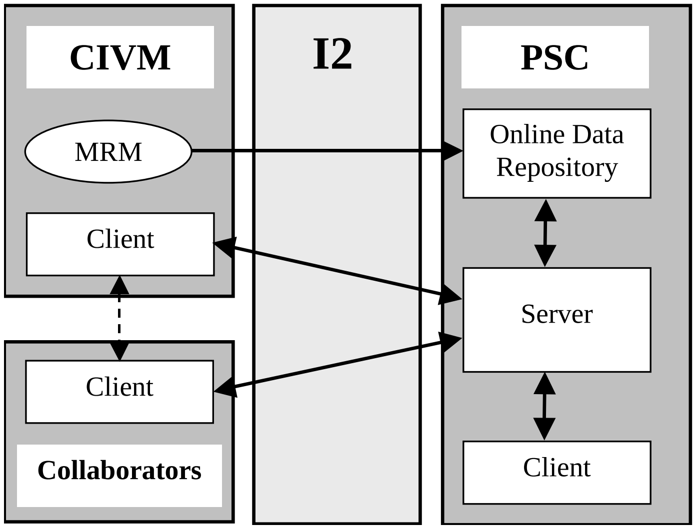

vs.psc.edu is located in PSC's machine room.

www.psc.edu

Major NLR routes

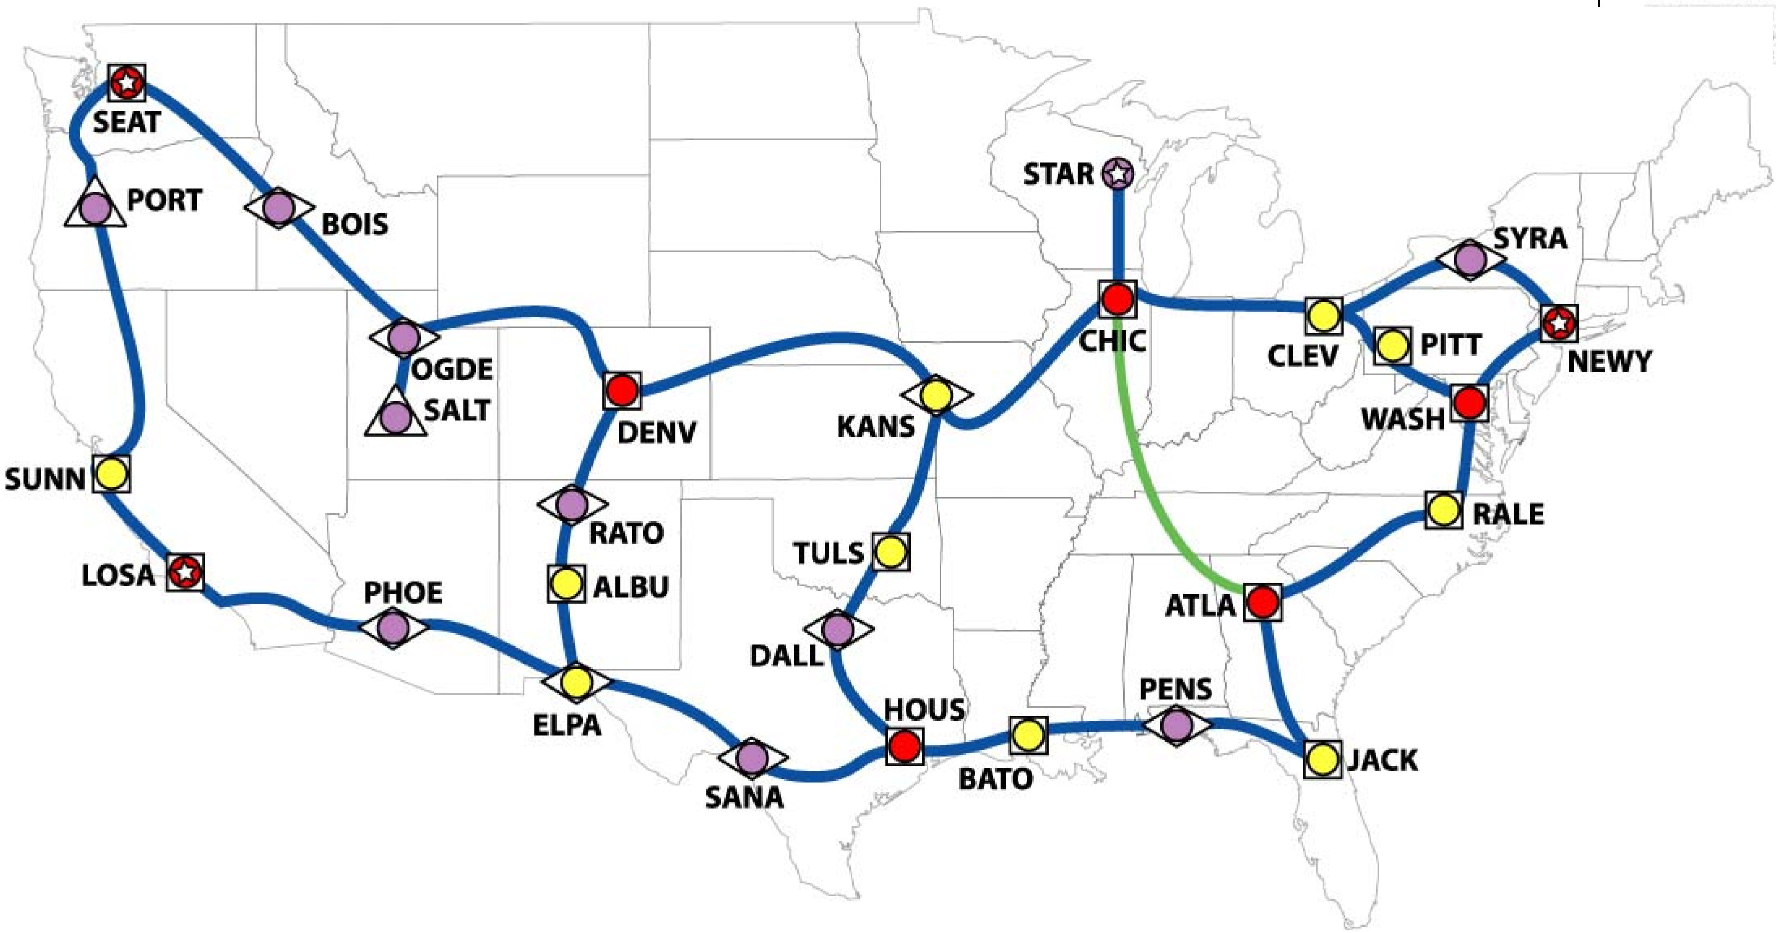

- NLR-owned fiber
- Managed Wave
- NLR WaveNet, FrameNet and PacketNet PoP
- NLR WaveNet and FrameNet PoP
- NLR WaveNet Pop
- PoP for primary connection point by a member (MetaPop)
- PoP needed for signal regeneration requirements, can also be used as secondary connection by a member
- PoP established by NLR for members' regional needs
- Pop established at exchange points

The PSC - CIVM pathway

- NLR is the main pathway
- Provides state-of-the-art performance
- RTTs have dropped from ~30 ms to 15 ms
- Jitter is now < 10 ms
- Center to center can reach 200 Mbits/sec
- Desktop end to end > 30 Mbits/sec

Network Parameter Tuning

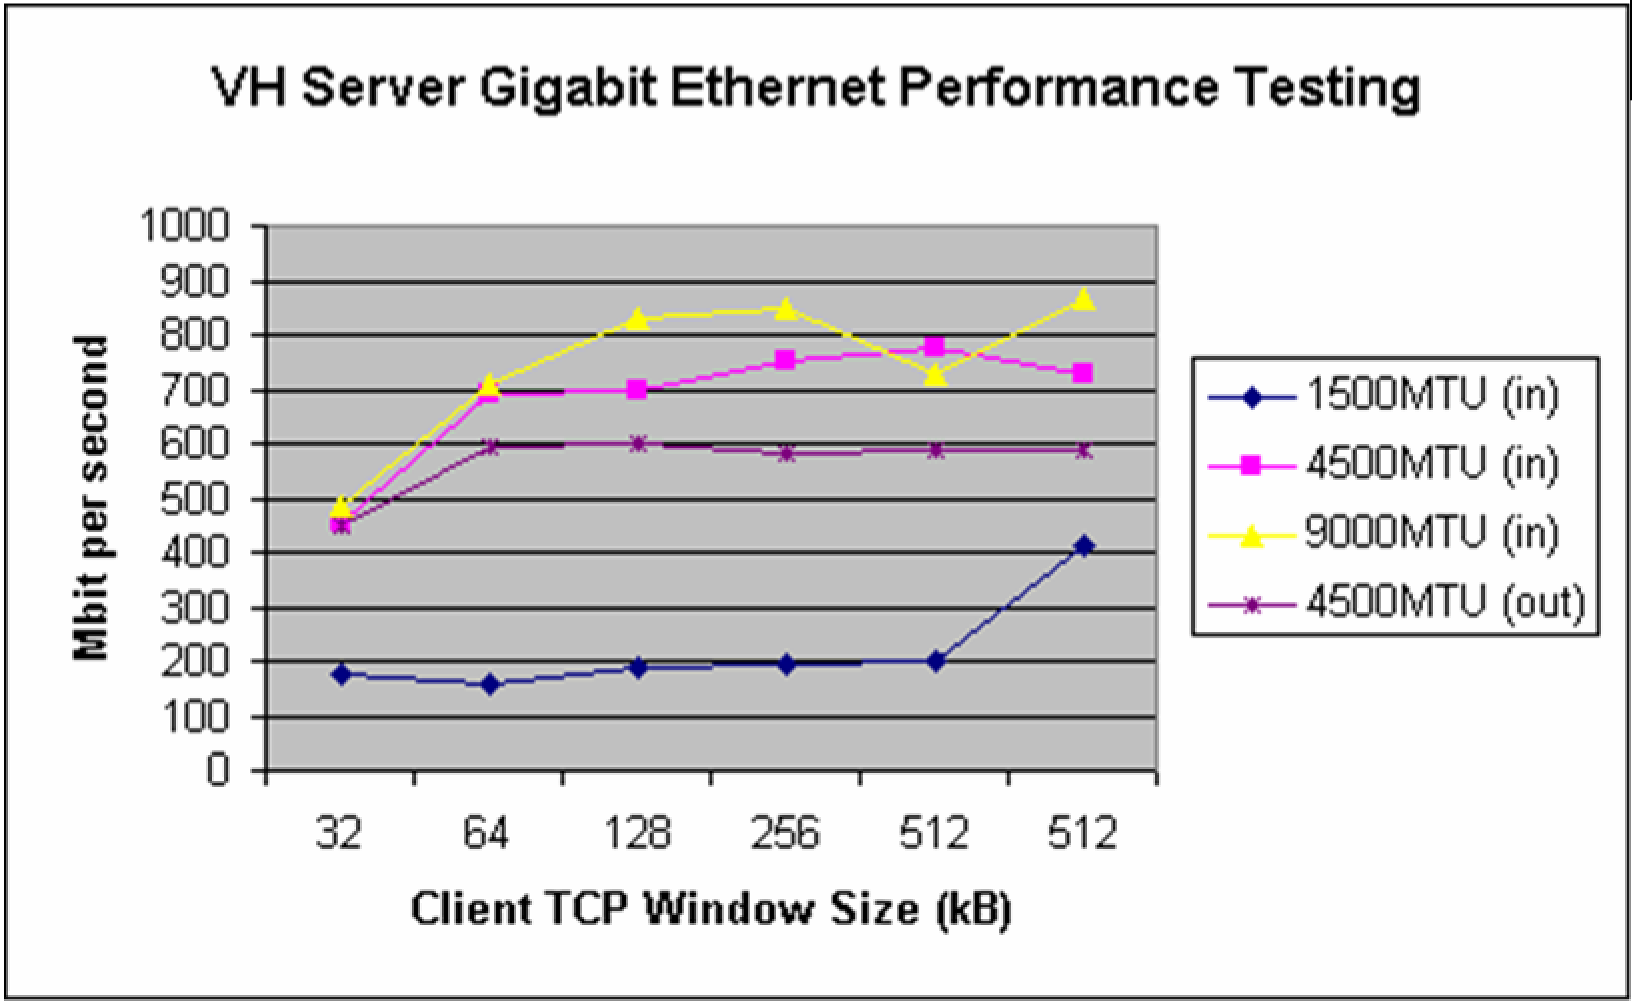

Back Forward Reload Stop http://www.psc.edu/networking/projects/hpn-ssh/

ADVANCED NETWORKING
PITTSBURGH SUPERCOMPUTING CENTER

Users Education Services Research News & Media

Home » Advanced Networking » Research Projects » High Performance Enabled SSH/SCP

High Performance SSH/SCP - HPN-SSH

On this page: **Abstract/Introduction**, **Patches**, **News and Updates**, **Theory and Implementation**, **Papers and Presenations**, **Contact**,

(PI) Chris Rapier PSC, Michael Stevens CMU
email: hpn-ssh@psc.edu

NEW! FAQ
Are you using HPN-SSH? Please let us know

Abstract

SCP and the underlying SSH2 protocol implementation in OpenSSH is network performance limited by statically defined internal flow control buffers. These buffers often end up acting as a bottleneck for network throughput of SCP, especially on long and high bandwith network links. Modifying the ssh code to allow the buffers to be defined at run time eliminates this bottleneck. We have created a patch that will remove the bottlenecks in OpenSSH and is fully interoperable with other servers and clients. In addition HPN clients will be able to download faster from non HPN servers, and HPN servers will be able to receive uploads faster from non HPN clients. However, the host receiving the data must have a properly tuned TCP/IP stack. Please refer to **this tuning page for more information**.

NRBSC

PITTSBURGH SUPERCOMPUTING CENTER

Hello,

I know this has come up before; but is the HPN patch (or elements thereof) currently being considered for integration in to the OpenSSH code base? Are there pending issues (buffer management, none cipher, etc) which still need to be addressed?

We have been using HPN-SSH for over a year now, and like others, have observed significant performance improvement over standard OpenSSH. I can scp a 1 GB test file between two HPN-SSH LAN hosts at 700 Mbps (<1 ms latency). And over a cross-country high-BDP WAN link, I'm able to achieve over 500 Mbps (85 ms latency). These single-stream scp transfers were run on well-tuned Linux kernels 2.6.15 (or higher) with the arcfour cipher. (I'll be happy to provide more details about these tests upon request.) I'm not sure how 'typical' my results are, but they represent an order of magnitude improvement over stock OpenSSH. While the improvement tends to vary among different platforms, I have never observed a performance degradation.

We recommend HPN SSH to our users who need to (securely) transfer their bulk scientific datasets ranging in size of hundreds megabytes to one terabyte; so naturally, performance is very important for them. But they (or their sysadmins) are often reluctant to deploy software which represents a deviation from a standard distribution...and the maintenance issues that follow.

Regards,

Avnish Bhatnagar
NASA Ames Research Center

Initial PSC-VB Development during U.Michigan/NLM project

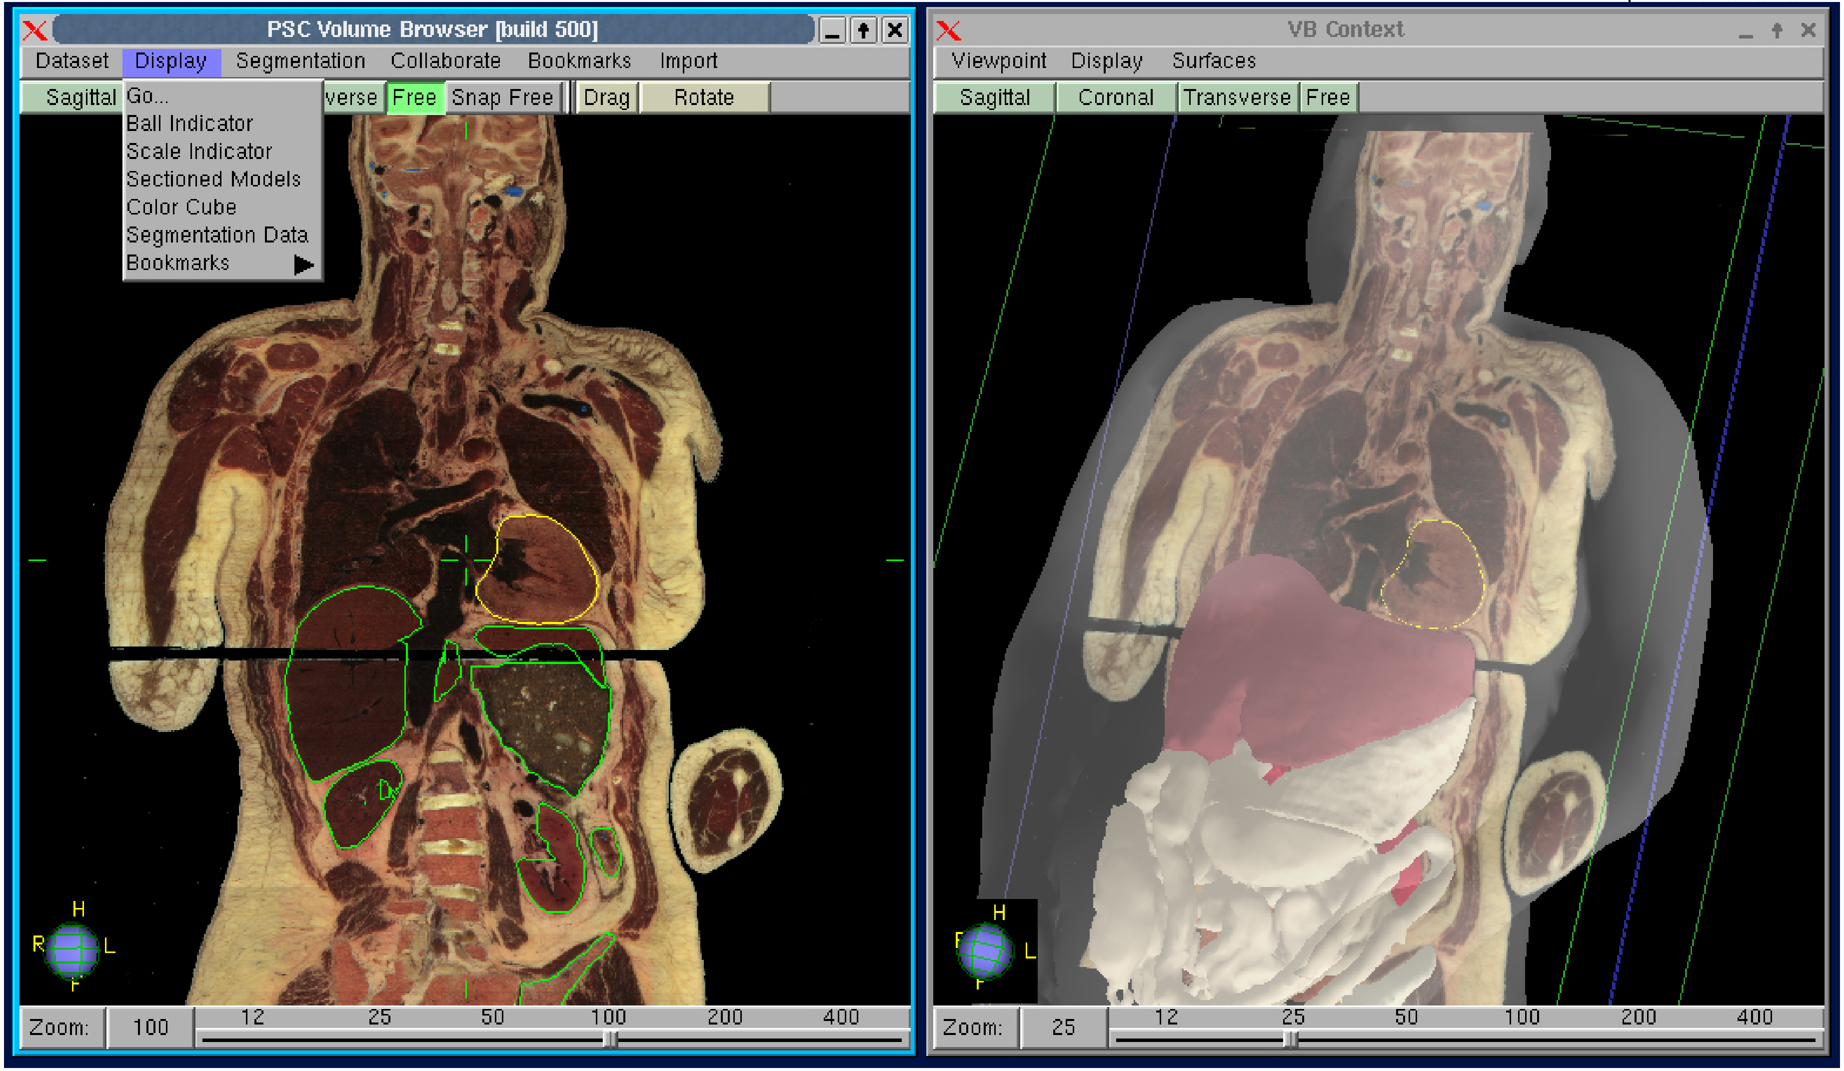

CIVM Segmented Mouse Brain

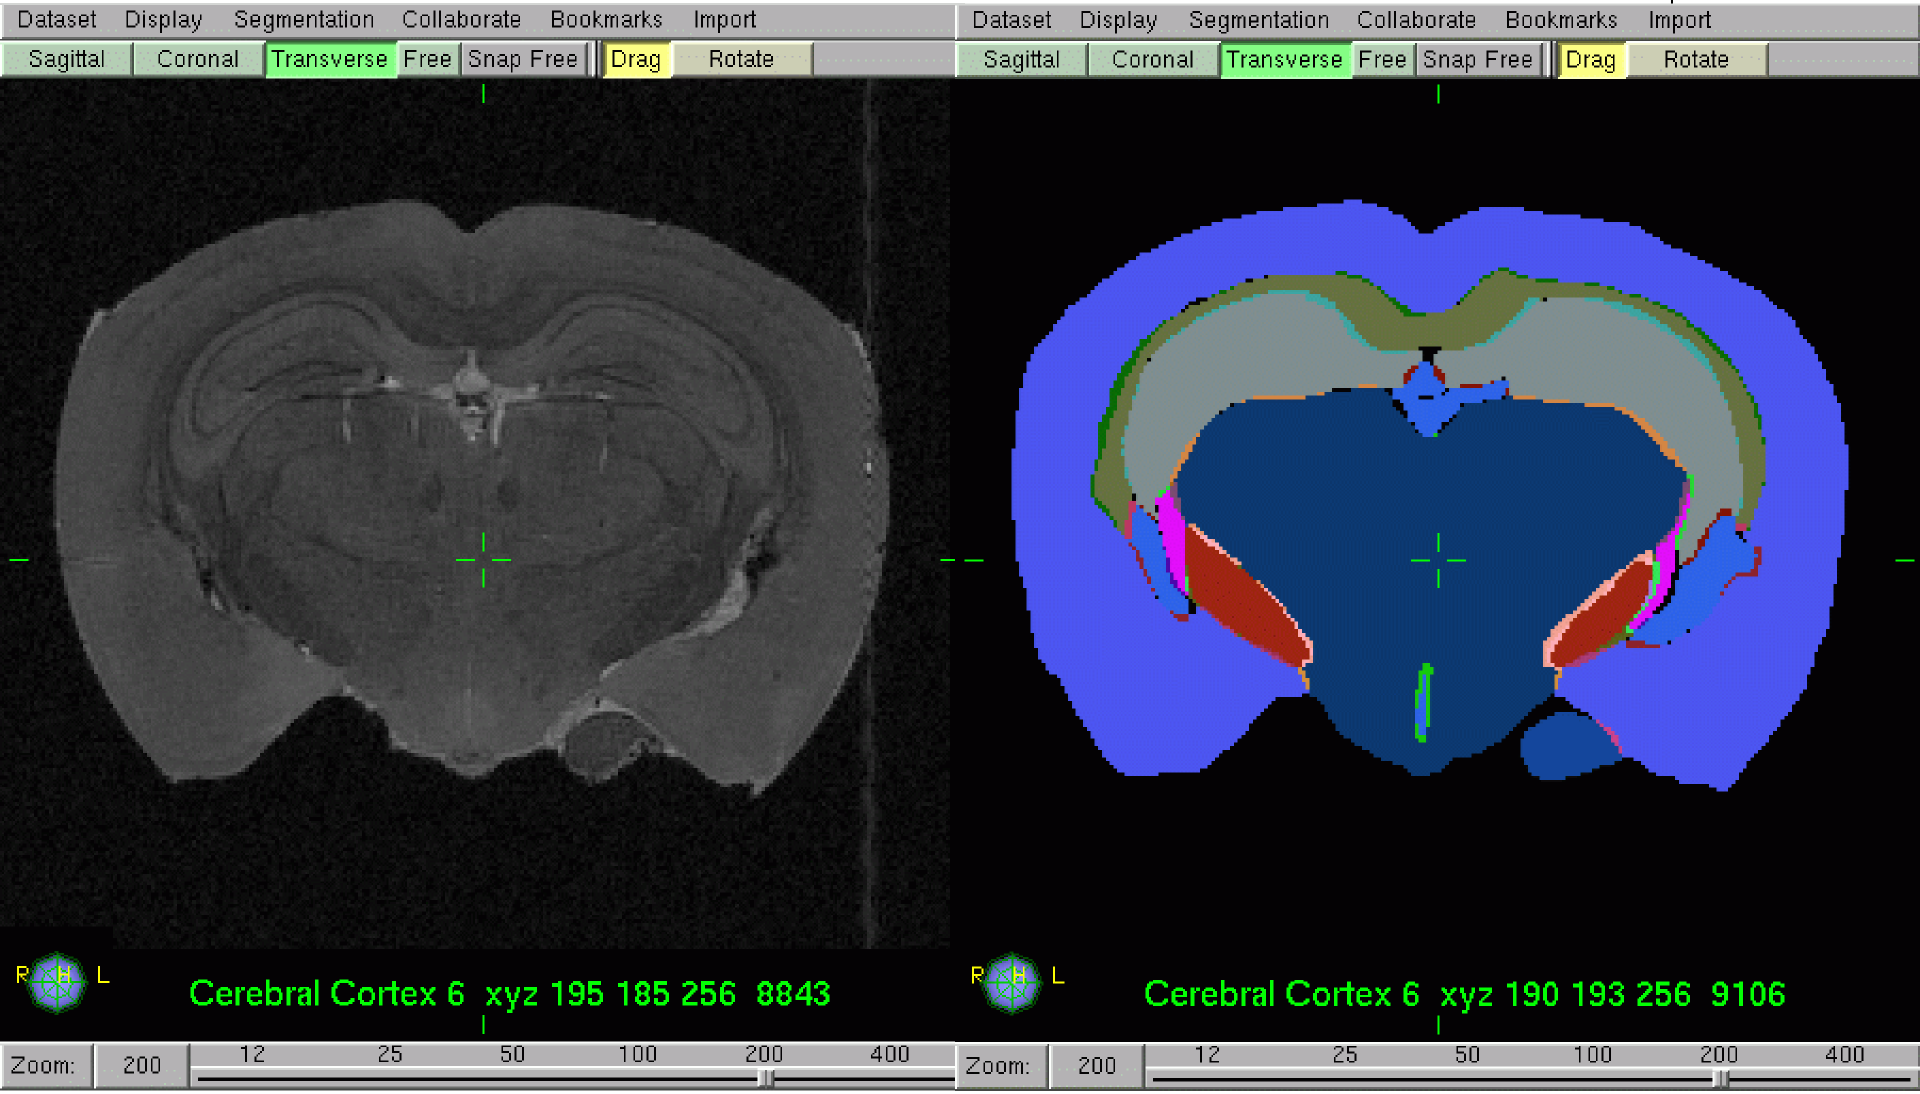

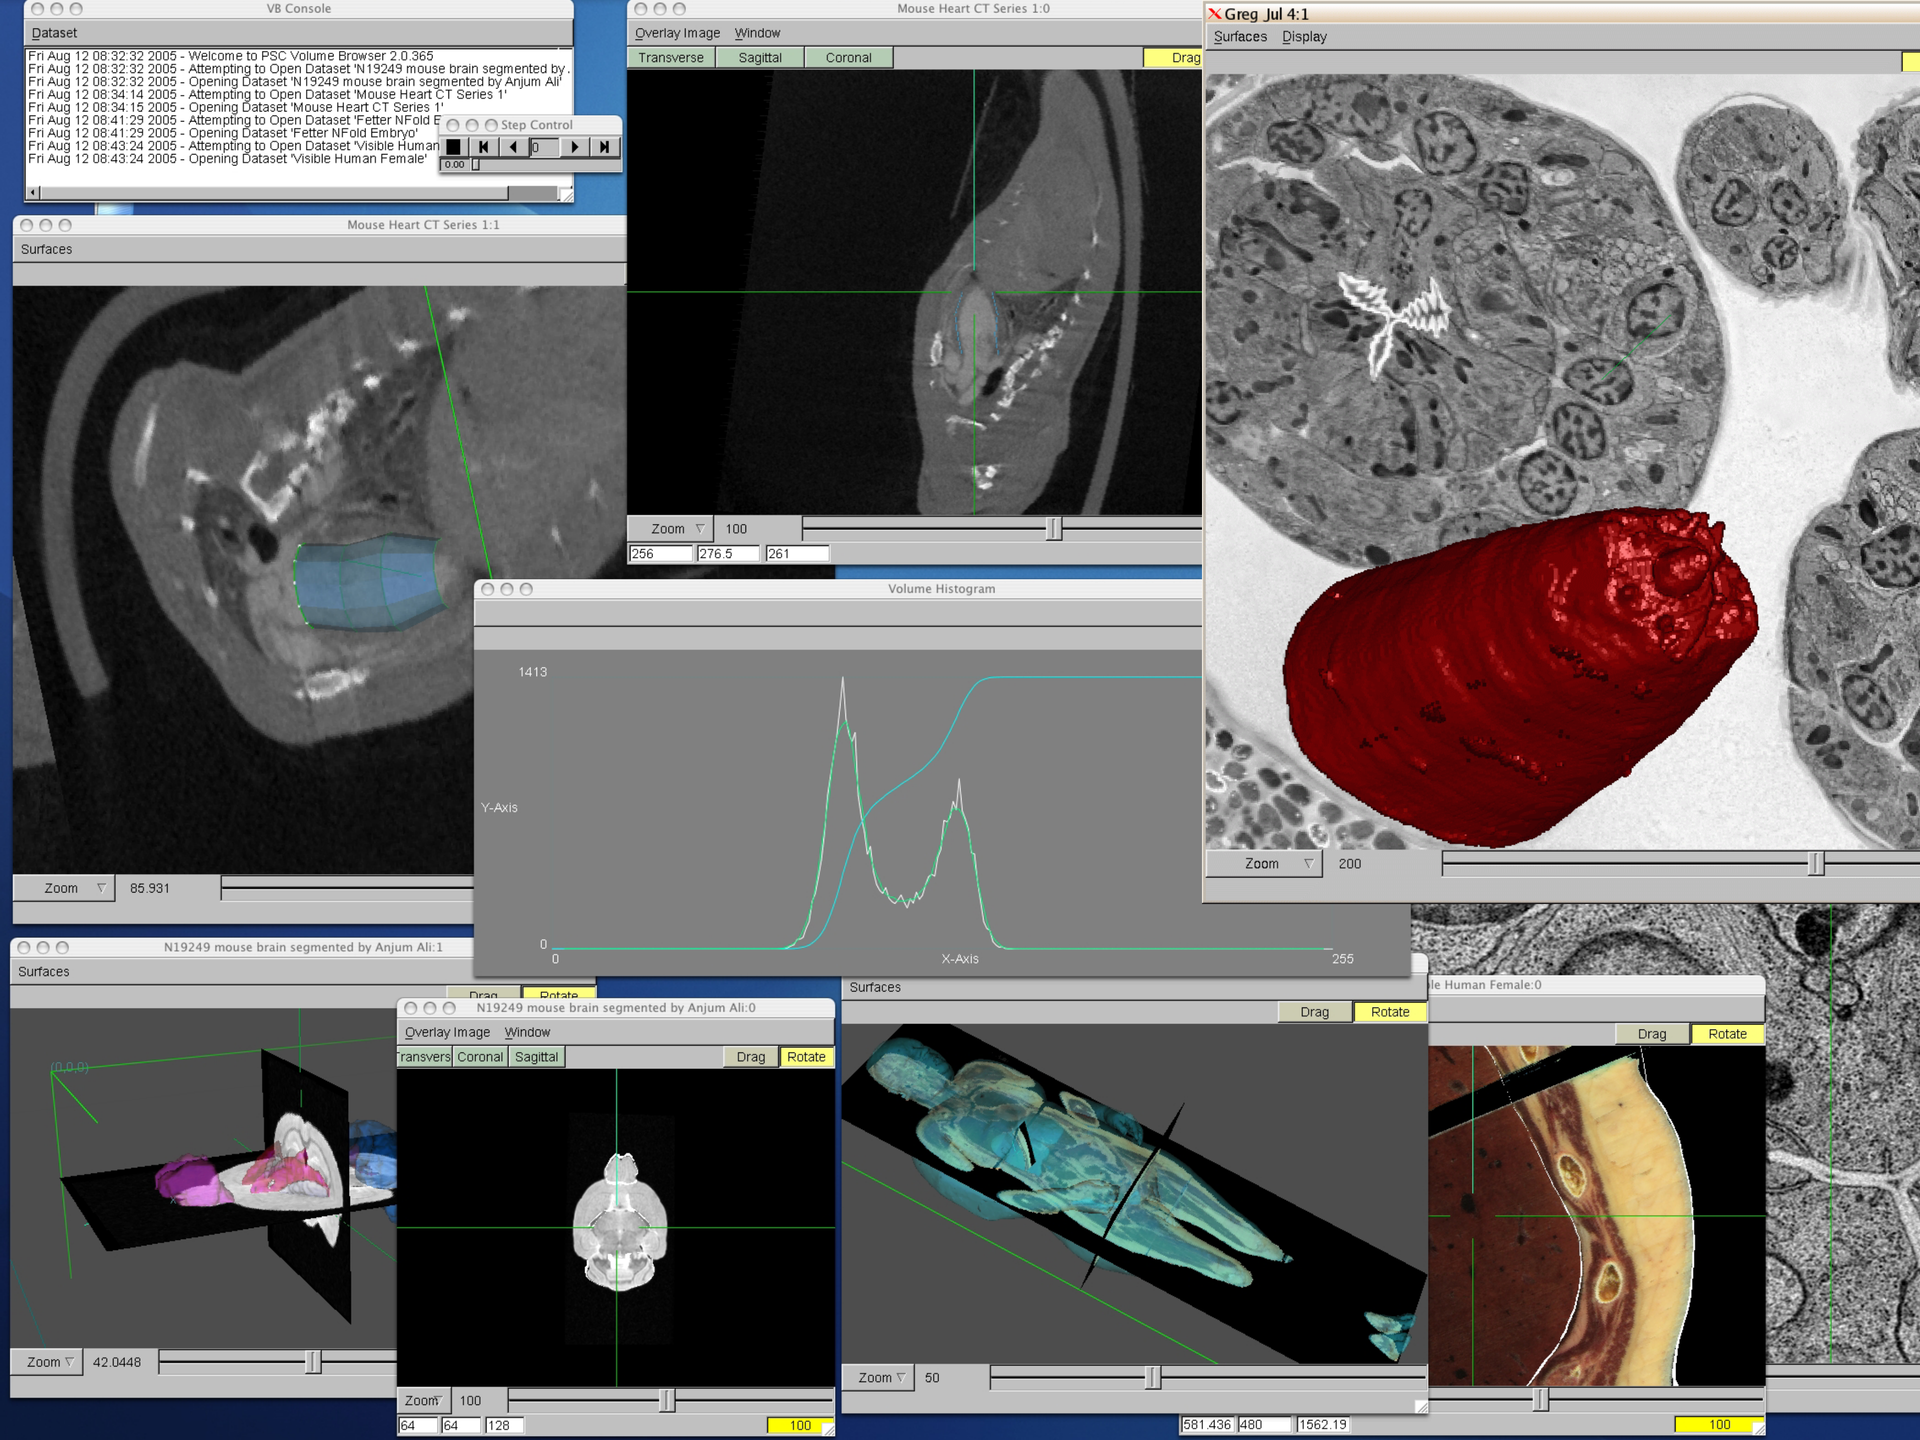

VB Console
Dataset
Fri Aug 12 08:32:32 2005 - Welcome to PSC Volume Browser 2.0.365
Fri Aug 12 08:32:32 2005 - Attempting to Open Dataset 'N19249 mouse brain segmented by
Fri Aug 12 08:32:32 2005 - Opening Dataset 'N19249 mouse brain segmented by Anjum Ali'
Fri Aug 12 08:34:14 2005 - Attempting to Open Dataset 'Mouse Heart CT Series 1'
Fri Aug 12 08:34:15 2005 - Opening Dataset 'Mouse Heart CT Series 1'
Fri Aug 12 08:41:29 2005 - Attempting to Open Dataset 'Fetter NFold E
Fri Aug 12 08:41:29 2005 - Opening Dataset 'Fetter NFold Embryo'
Fri Aug 12 08:43:24 2005 - Attempting to Open Dataset 'Visible Human
Fri Aug 12 08:43:24 2005 - Opening Dataset 'Visible Human Female'
Step Control
0.00
Mouse Heart CT Series 1:0
Overlay Image
Window
Transverse
Sagittal
Coronal
Drag
Zoom
100
256
276.5
261
Greg Jul 4:1
Surfaces
Display
Zoom
200
Mouse Heart CT Series 1:1
Surfaces
Zoom
85.931
Volume Histogram
1413
Y-Axis
0
0
X-Axis
255
N19249 mouse brain segmented by Anjum Ali:1
Surfaces
Drag
Rotate
(0,0,0)
Zoom
42.0448
N19249 mouse brain segmented by Anjum Ali:0
Overlay Image
Window
Transvers
Coronal
Sagittal
Drag
Rotate
Zoom
100
64
64
128
100
Surfaces
Drag
Rotate
Zoom
50
581.436
480
1562.19
le Human Female:0
Drag
Rotate
100

Relevant Technology Evolution

- Storage
  - Disk (~4 GBytes/$, 1 TByte/drive, ~100 MB/sec/drive, ~100-200 accesses/sec
  - Main memory (~$40/GByte, 1-8 GBytes/CPU)
  - Solid state drives – very high cost
  - Flash drives (~1/2 cost of DRAM, 500us access, 30 MB/s bandwidth)
- Computation – Moore's law still holds
- Networking – improving rapidly but by jumps
- Graphics hardware – effective rates improving faster than Moore's law and now useful for general computation but still memory limited.

Disk access time is more critical than network latency

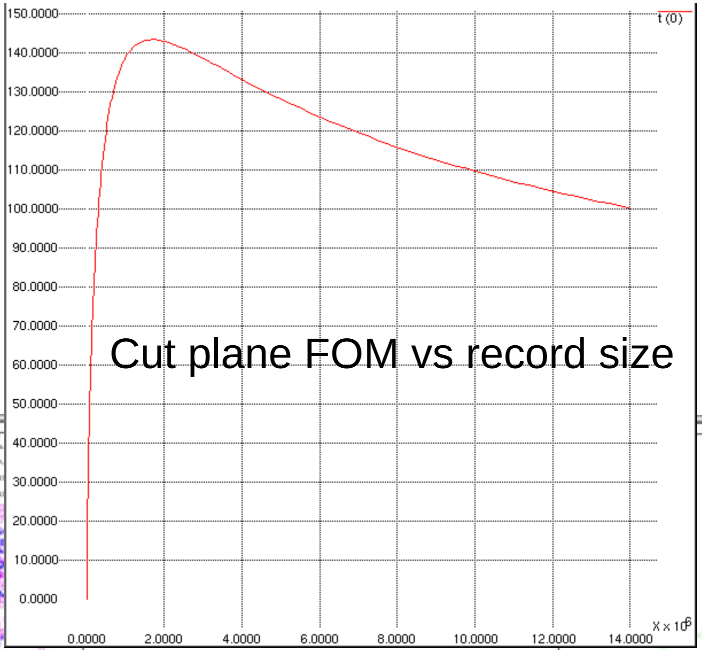

Cut plane FOM vs record size

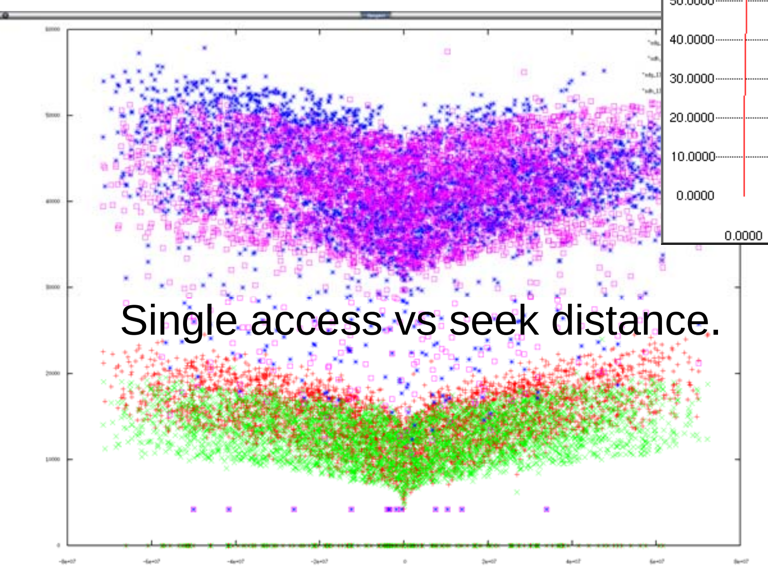

Single access vs seek distance.

Multi user elevator ordering effects.

Methodology items of note:

- Borrowed lossless H.264 transform
- Used a cooperative development strategy
- Communicate release instructions on web
- Evaluation team provided guidance user question forms/interviews
- Demonstrated CIVM database linkage by using PSC-VB as a helper application
- Tried to maintain broad applicability

Charlie Little & Brenda Rongish time-lapse avian development

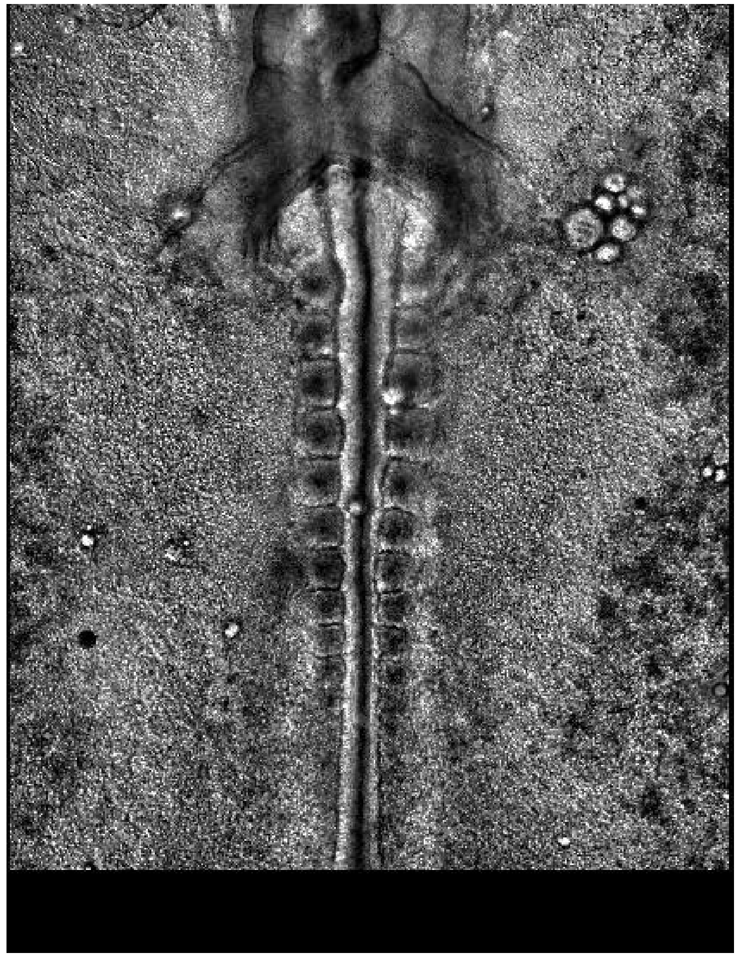

Duke Center for In Vivo Microscopy Imaging Systems

7T MRI

9.4 T MRI

2 T MRI

Optical Imaging

Ultrasound

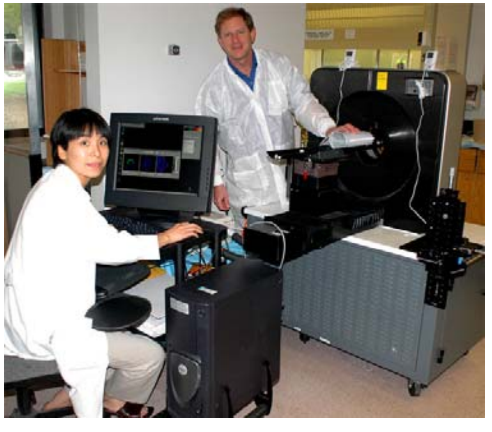

MicroPET

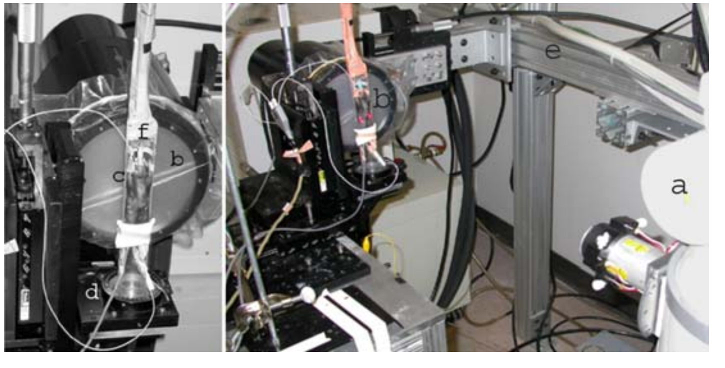

MicroCT


Virtual Human Data: National Library of Medicine
Human Image: Bill Lorensen, GE CR&D
Duke Center for In Vivo Microscopy NIH/NCRR

Large Buckets of Data

1. Very Large 3D Arrays

   - MR Histology: (1k x 1k x 4k)

2. 4D Arrays: (3D + time)

   - cardiac micro-CT or MRM
   - perfusion micro-CT or MRM

3. Multimodality Data

   - combined micro-CT/DSA
   - combined micro-CT and micro-PET

The Visible Mouse @40,000X

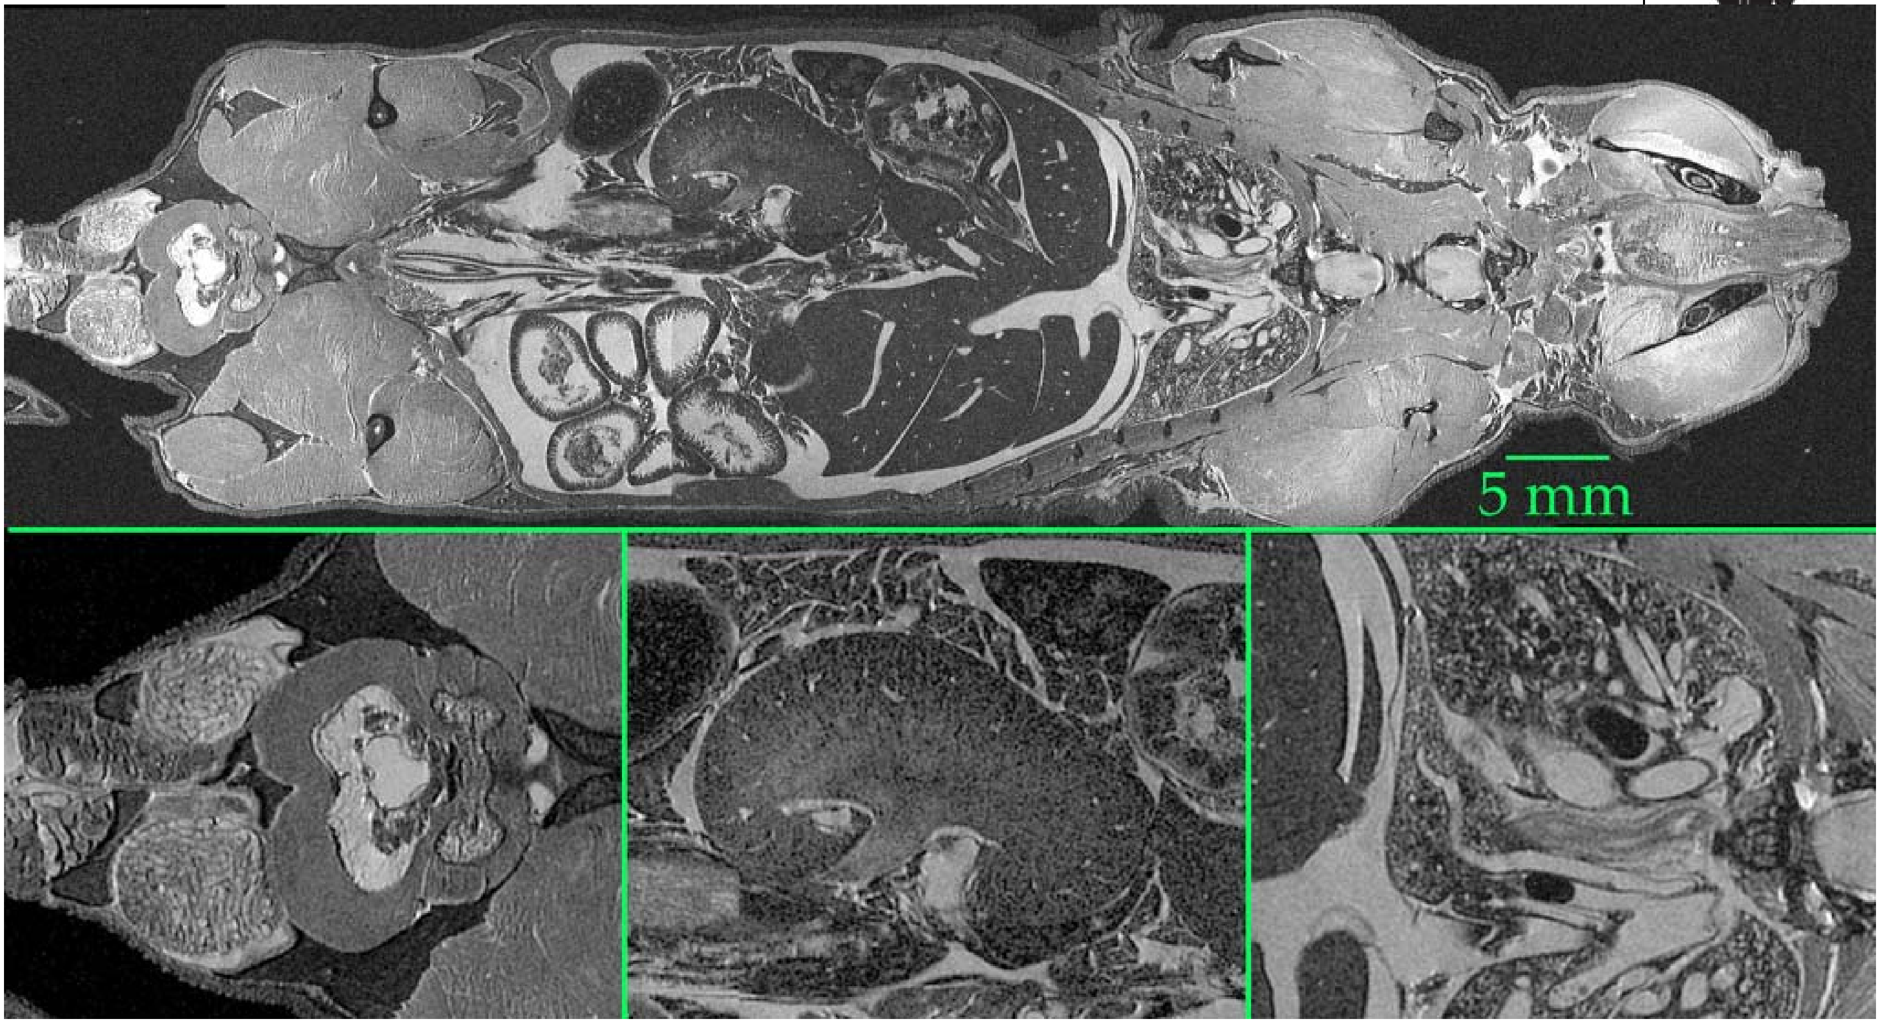

50x50x50 microns

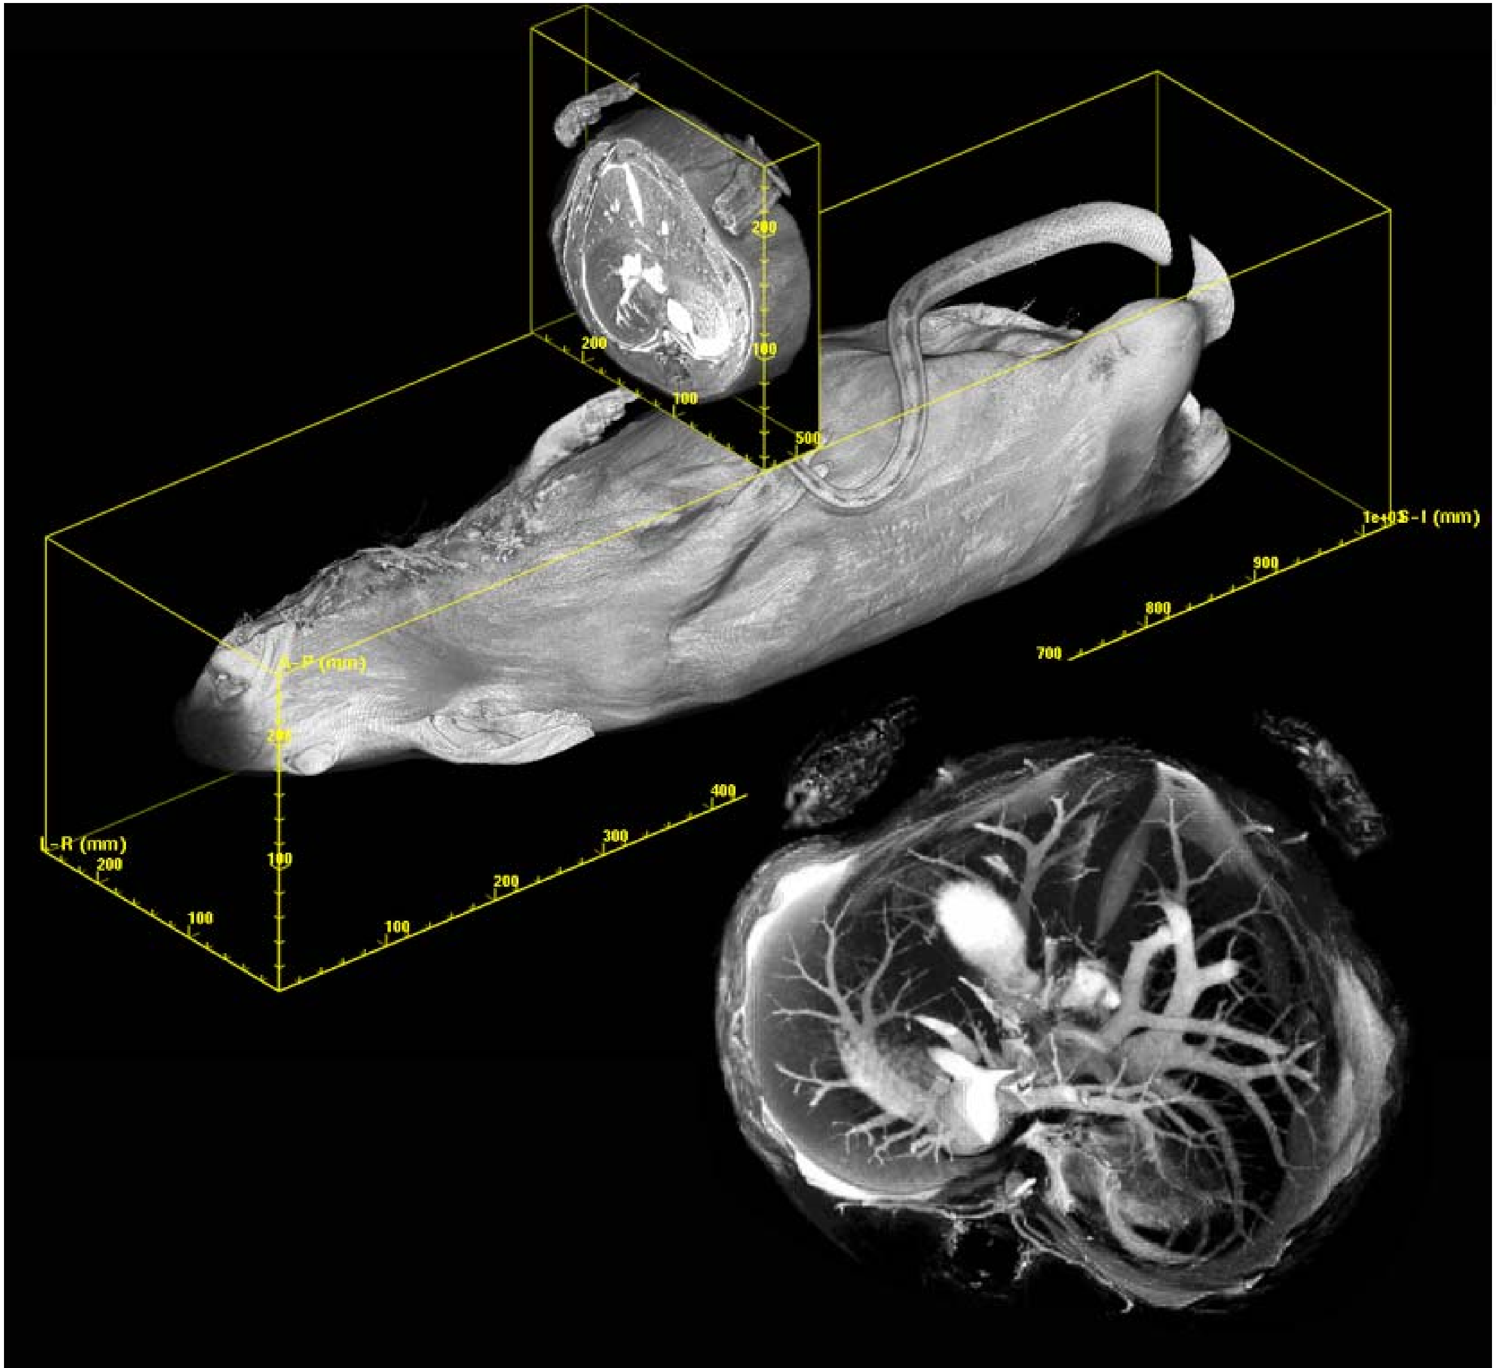

NRBSC

PITTSBURGH SUPERCOMPUTING CENTER

Results: Pre-natal Development

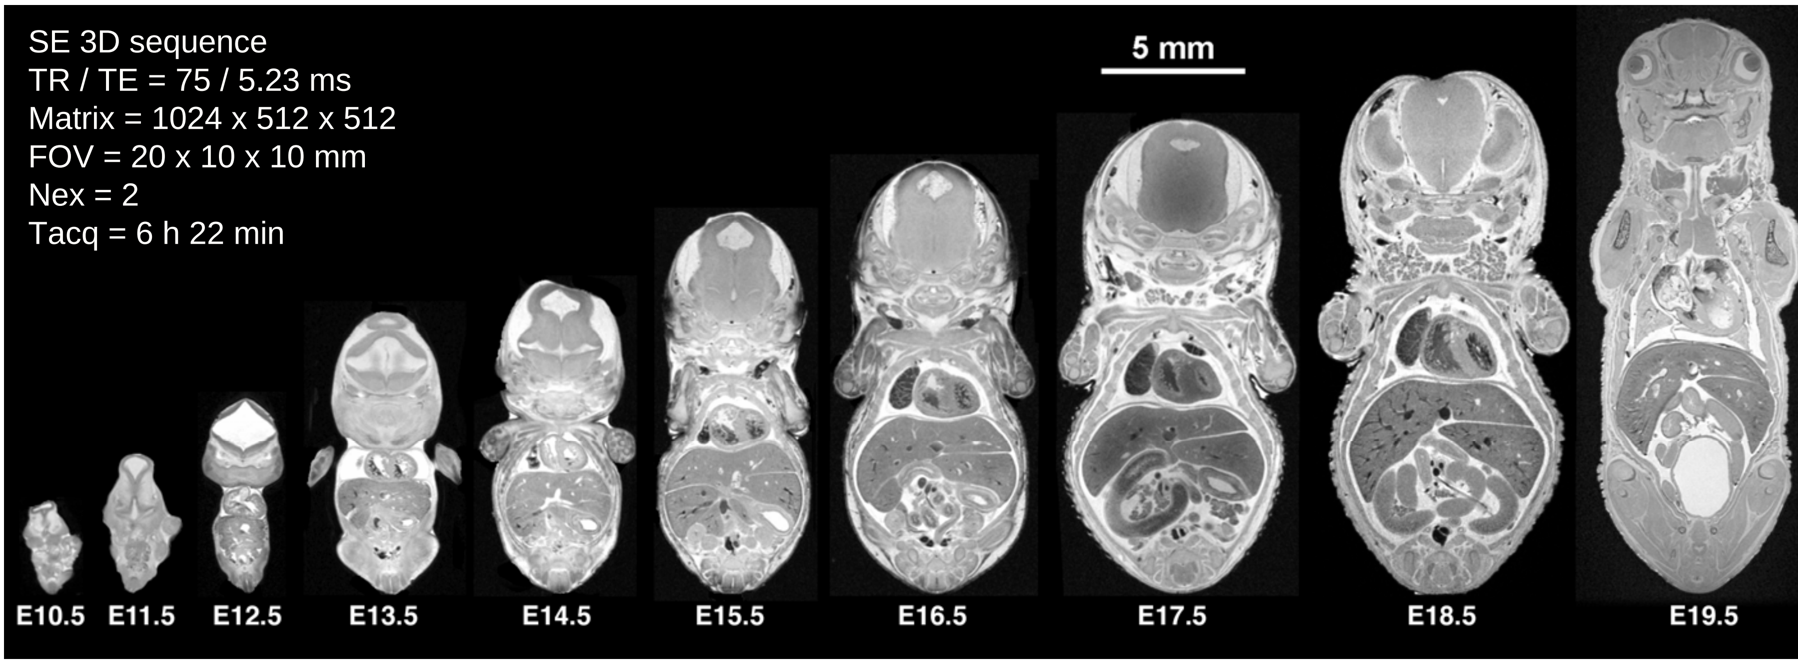

Mouse Embryo (E17.5) @ 20 μm

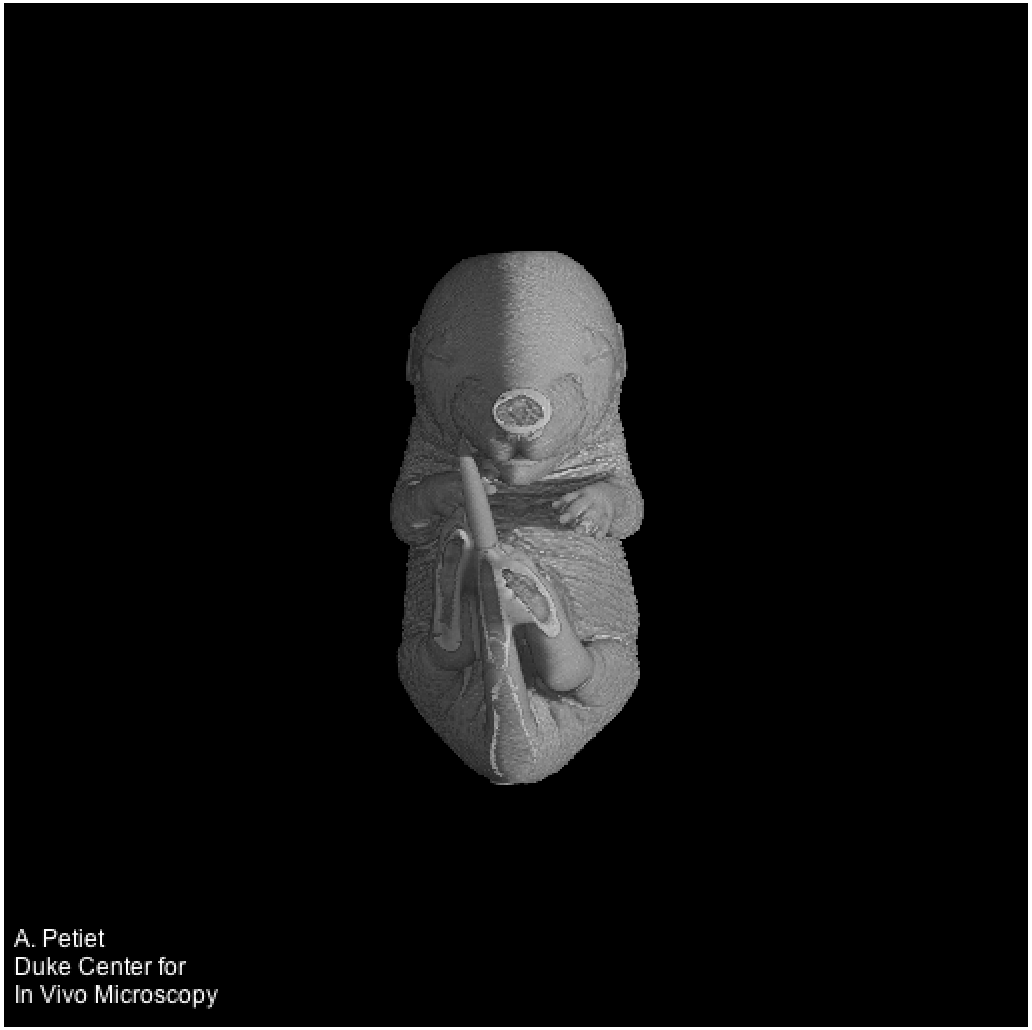

Duke CIVM Micro-CT*

2Kx2K Cooled CCD @ 50 um

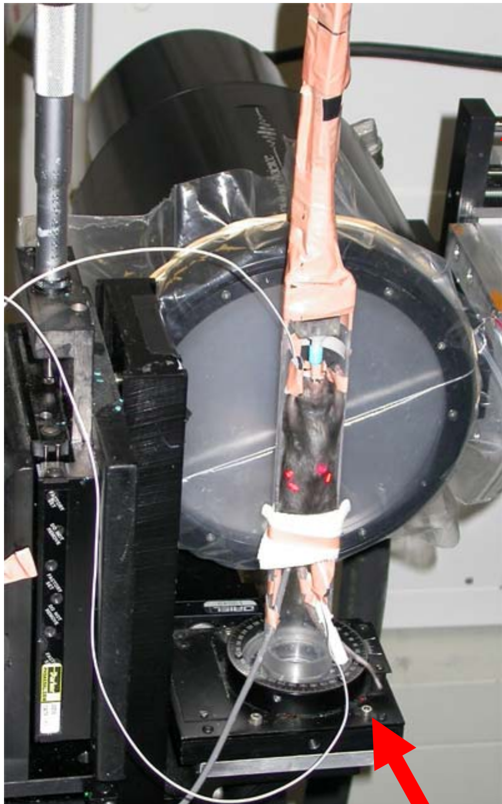

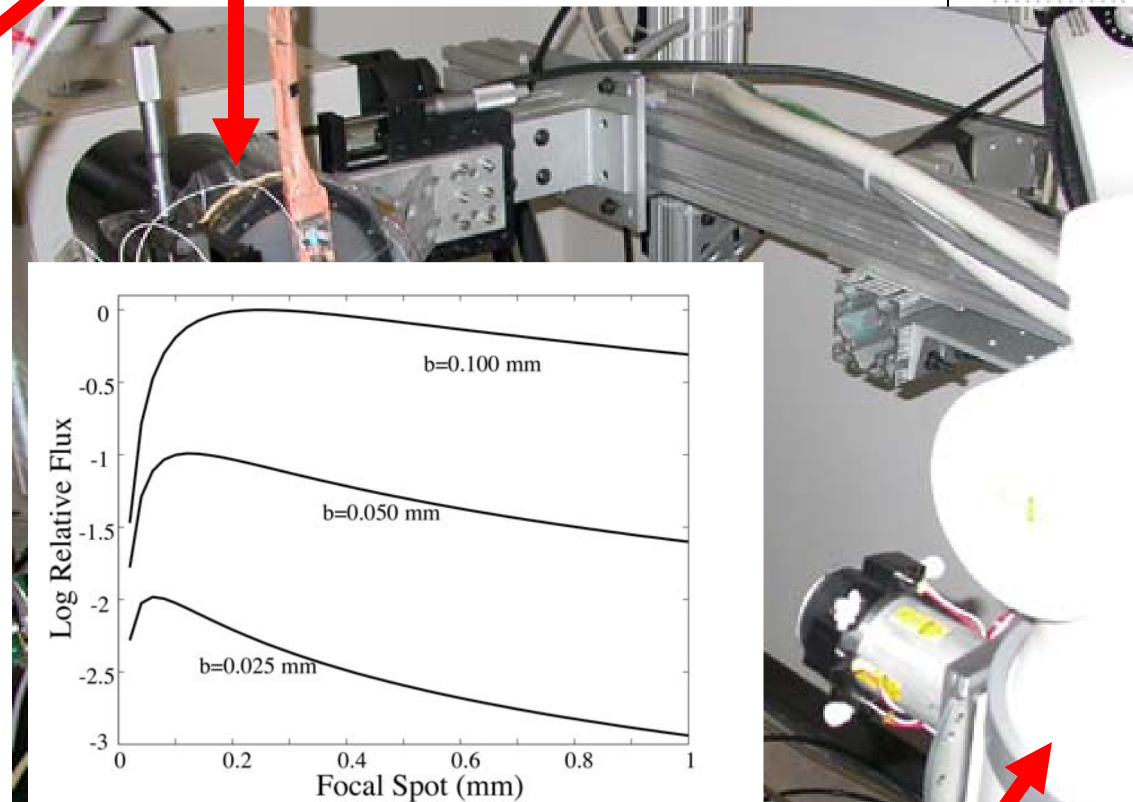

Computer Controlled Turntable

250 X Higher Flux !

Hi Flux X Ray Source

Badea CT, Hedlund LW, Johnson GA. Micro-CT with cardiac and respiratory gating. Med Phys 2004; 31(12):3324-9.

Triggered Acquisition

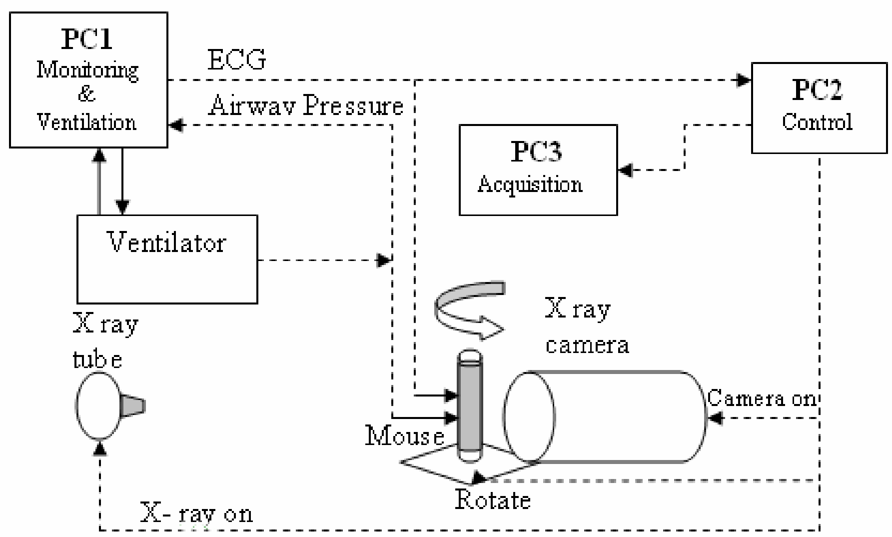

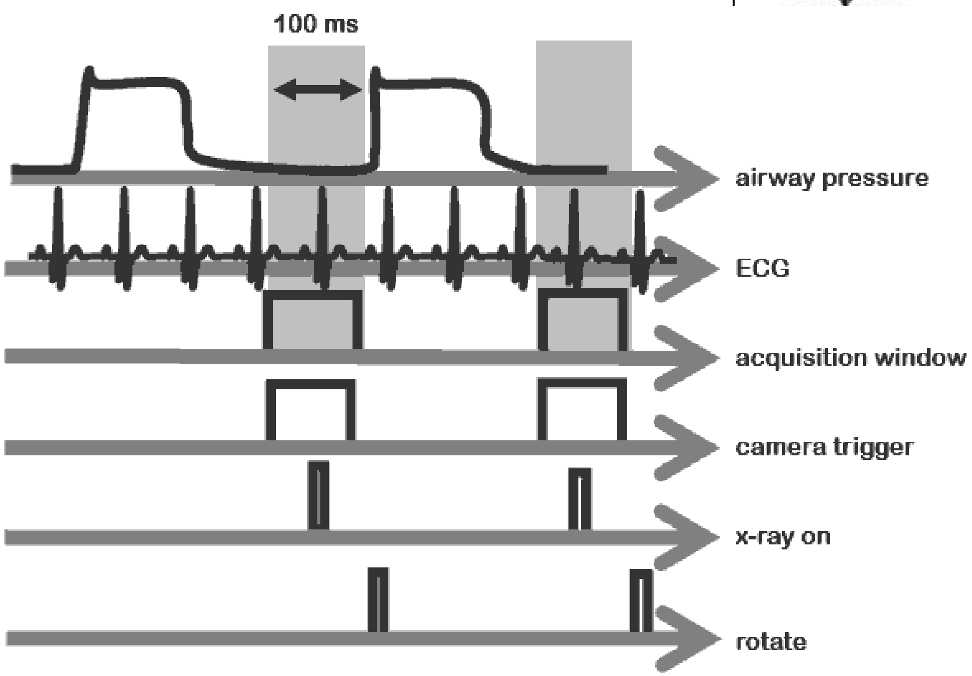

-X-ray : 80 Kvp, 150 mA and 10ms (flux sufficient to fill the detector wells to ~ 25 %). Exposure 60 mR/proj

-Projections acquired on 190^0 with a step angle of 0.75^0.

-Scanning time about 8-10 mins

4D Cardiac Micro-CT in Mice

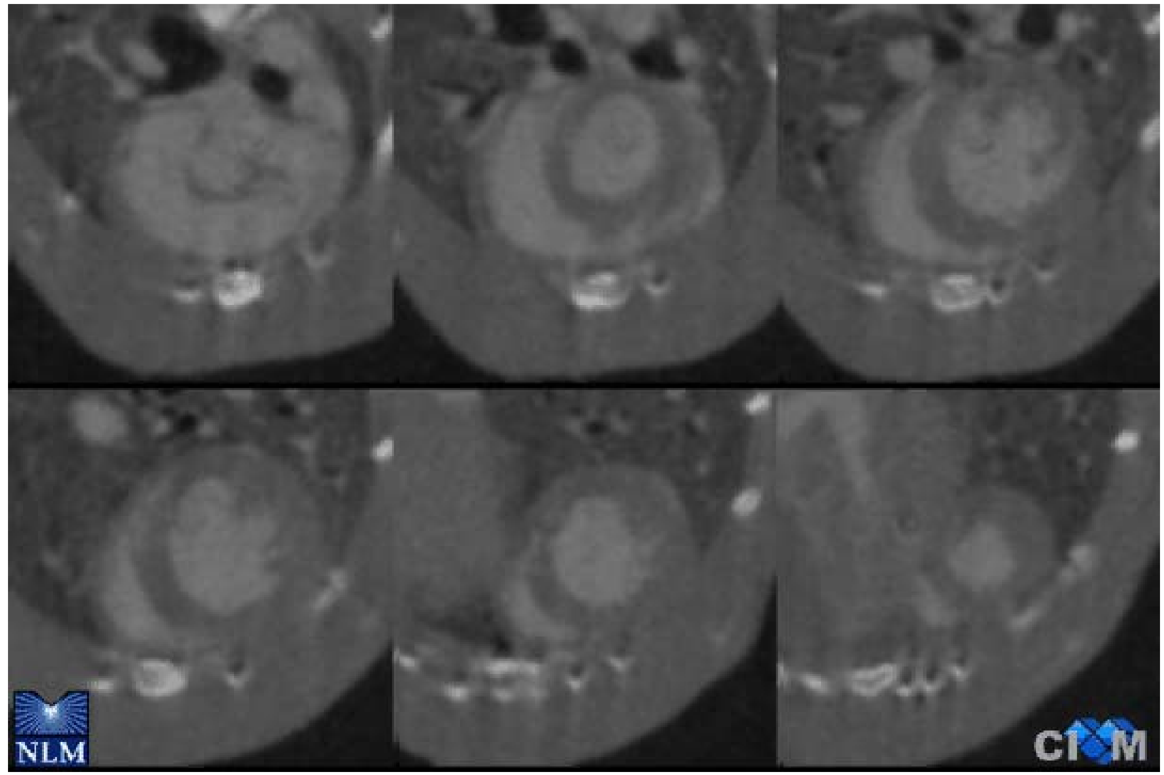

100x100x100 microns x 10 ms

Conventional Analysis of Cardiac Function

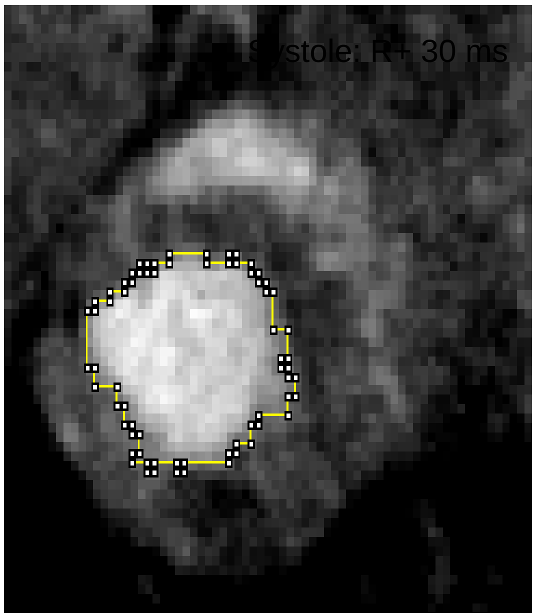

Cardiac measure	
LV diastole vol (micro l)	49 ±10.7
LV systole vol (micro l)	24 ± 3.8
Ejection fraction (%)	50.9±2.81
Stroke Volume (micro l)	25±6.3
Cardiac Output (ml /min)	12.05

Table1: Cardiac function estimation in (n=5) mice using the micro CT

Live Mouse Micro CT Goals

- Improve accuracy of LV measurement
- Reduce contrast agent dose
- Reduce radiation exposure
- Reduce analysis time
- Reduce manual analysis intervention
- Enable time studies of individual animals

LV Volumetric Measurements

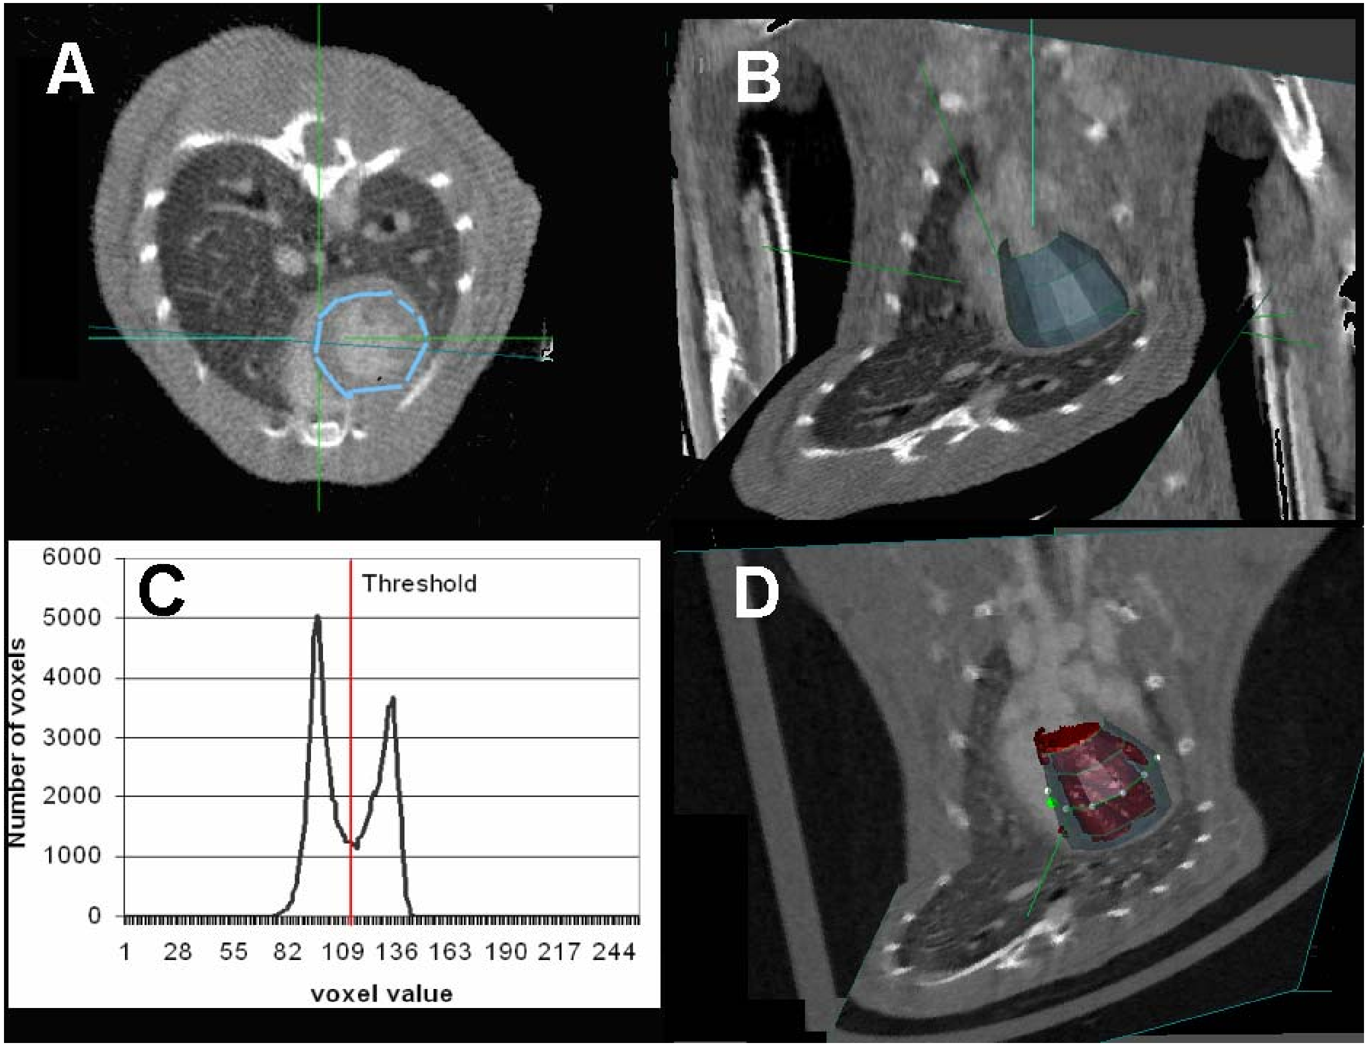

Volume Browser by **Pittsburgh Supercomputing Center**

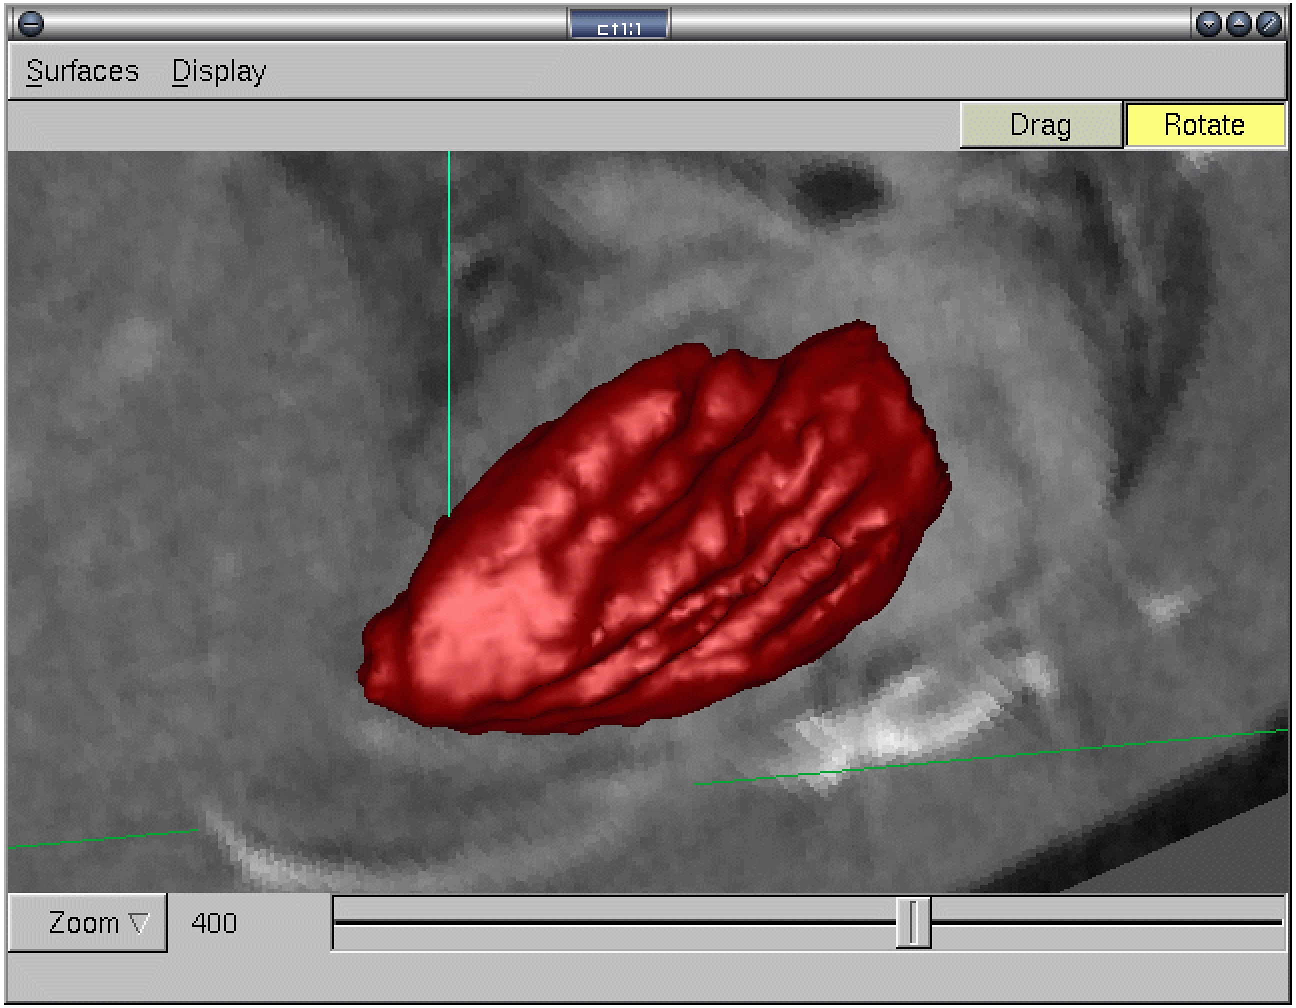

NRBSC

PITTSBURGH SUPERCOMPUTING CENTER

Micro-CT Study of Myocardial Infarction (MI) in Mice

- MI mouse model by LAD ligation
- Scanned at 5 days and 5 weeks post MI!
- Goal: MI size and cardiac function

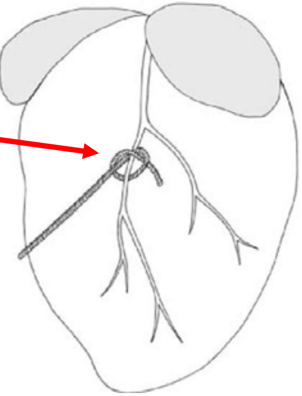

MI: Hyperenhancement

5 days

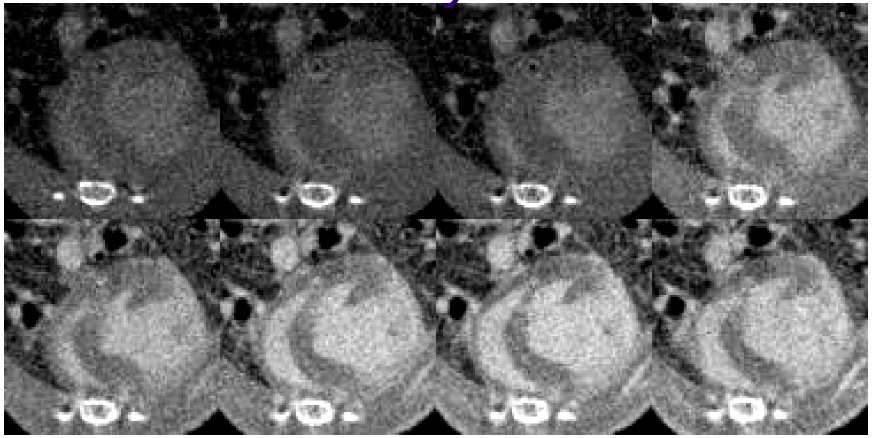

35 days

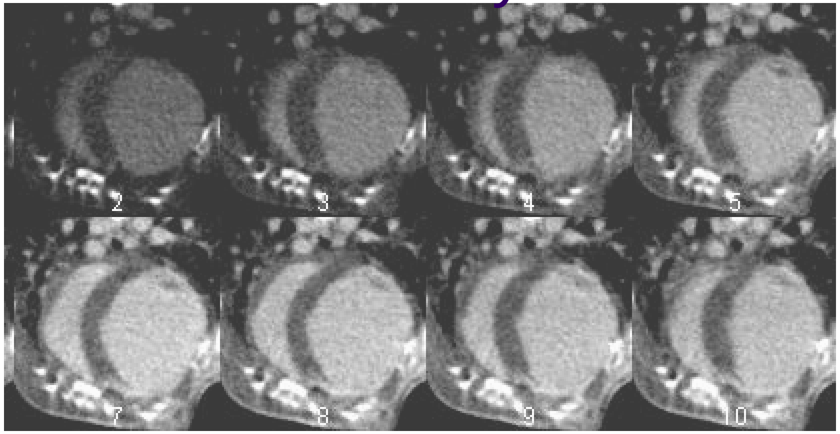

Nahrendorf M, Badea C, Hedlund LW, Figueiredo JL, Sosnovik DE, Johnson GA, Weissleder R. High Resolution Imaging of Murine Myocardial Infarction With Delayed Enhancement Cine Micro-CT. Am J Physiol Heart Circ Physiol. 2007

Cardiac Function : MI vs. Controls

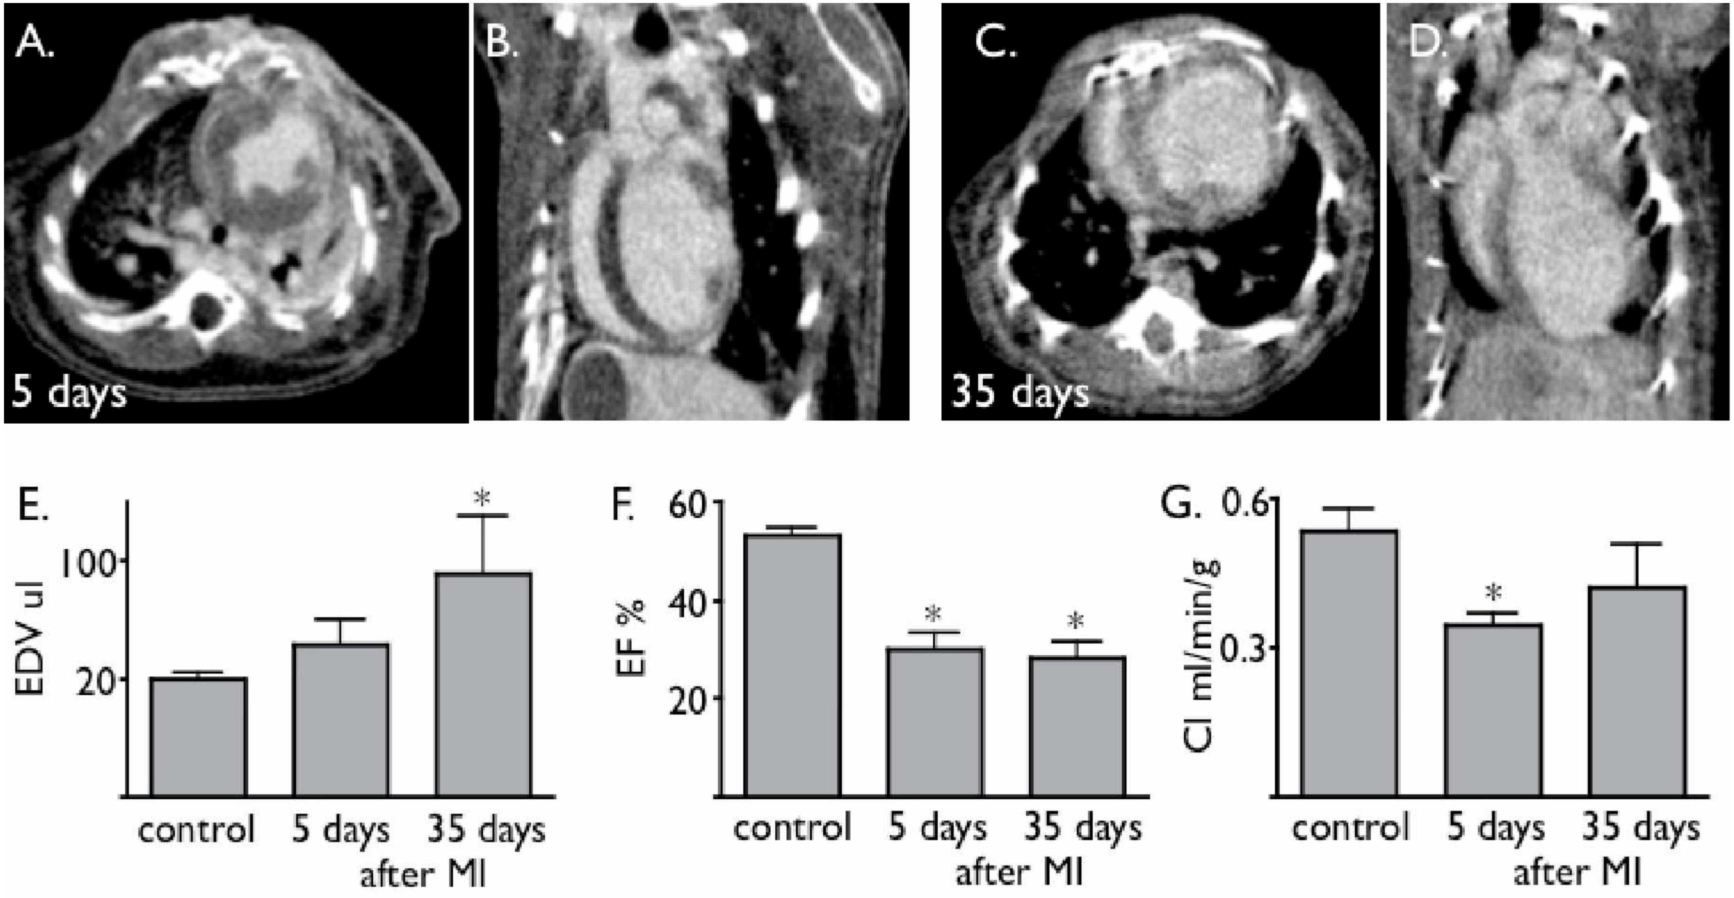

MICRO-CT FOR MORPHOLOGICAL AND FUNCTIONAL PHENOTYPING OF MLP NULL MICE

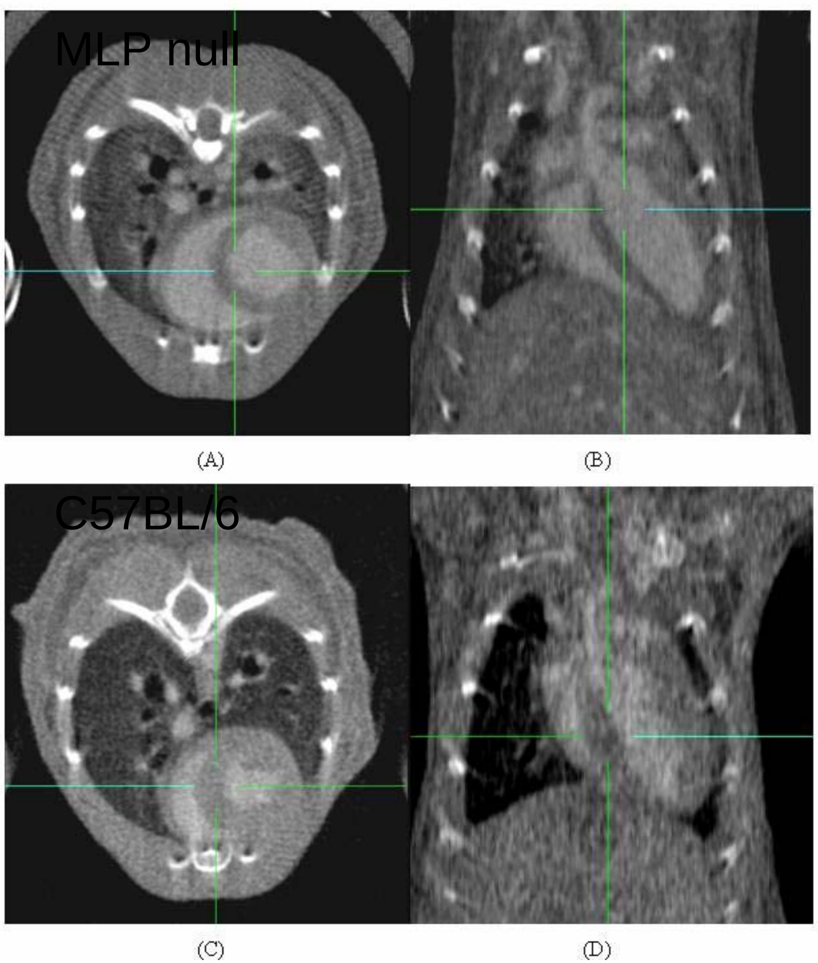

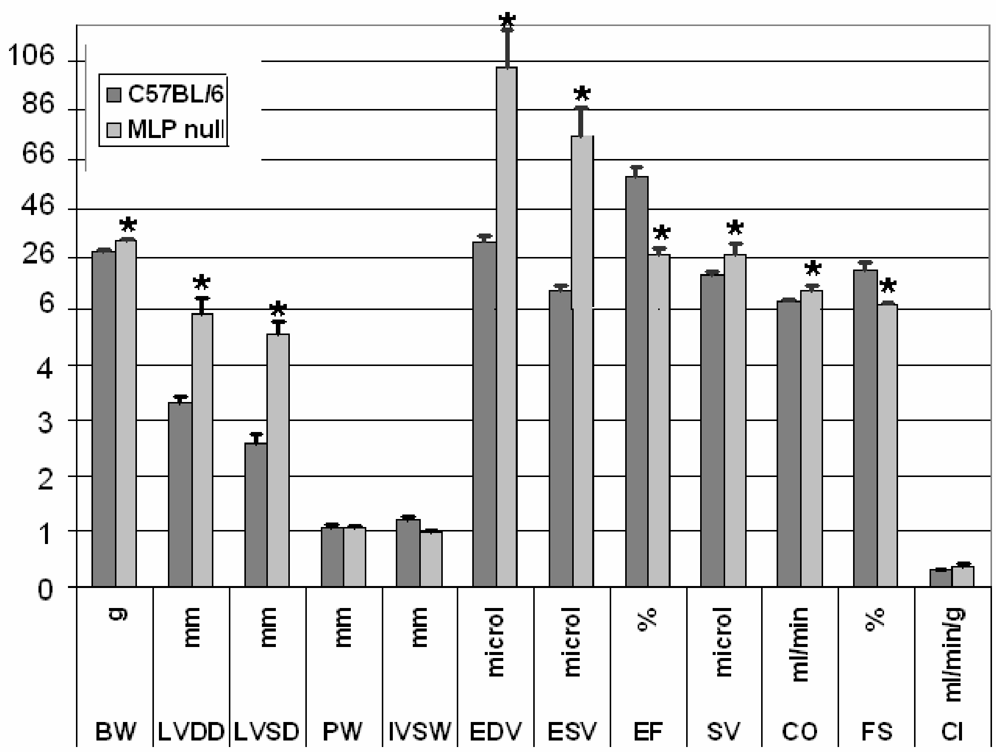

Effects of Contrast Agent and # of CT projections

0.125 ml Contrast Agent | **0.5 ml Contrast Agent**

63 projections | **380 projections** | **63 projections** | **380 projections**

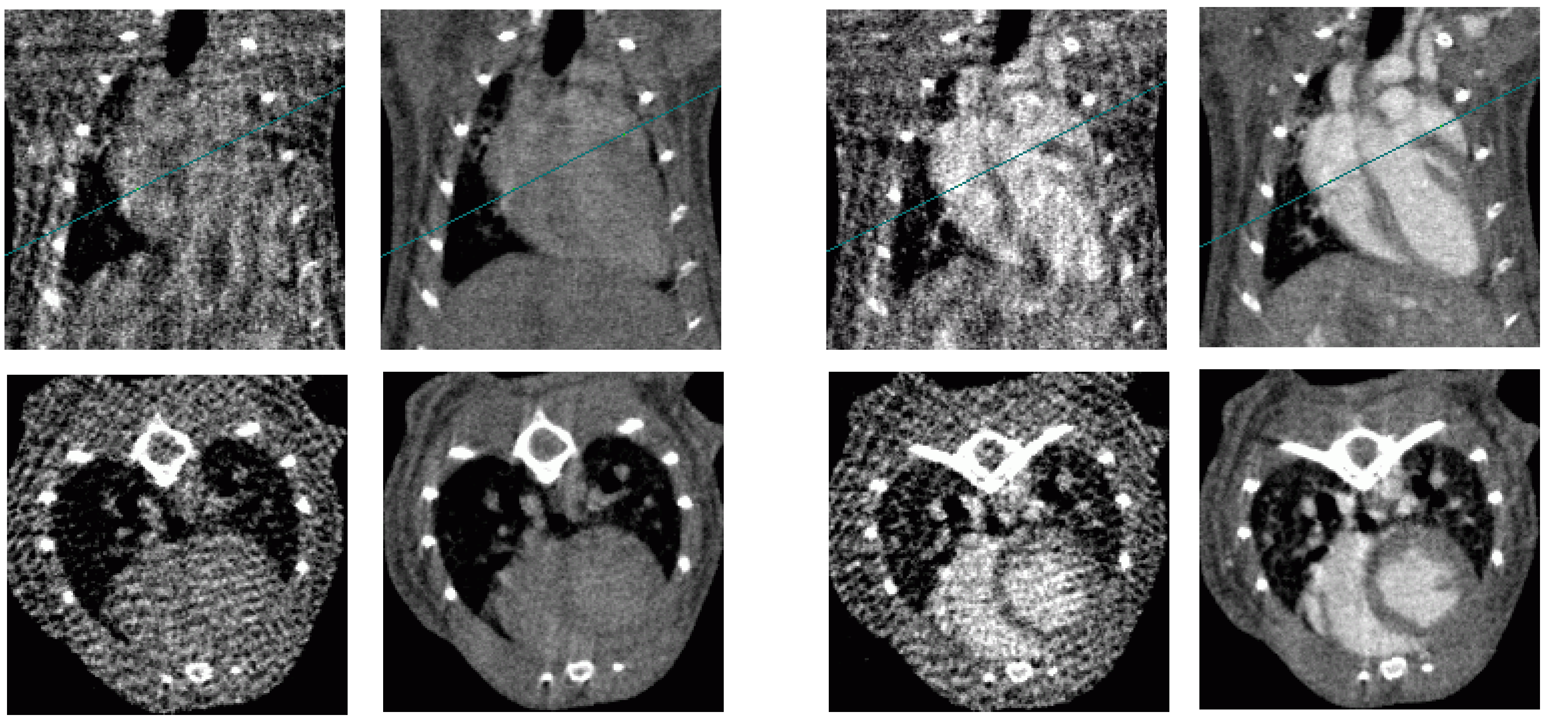

Live MicroCT vs. Fixed Micro MR

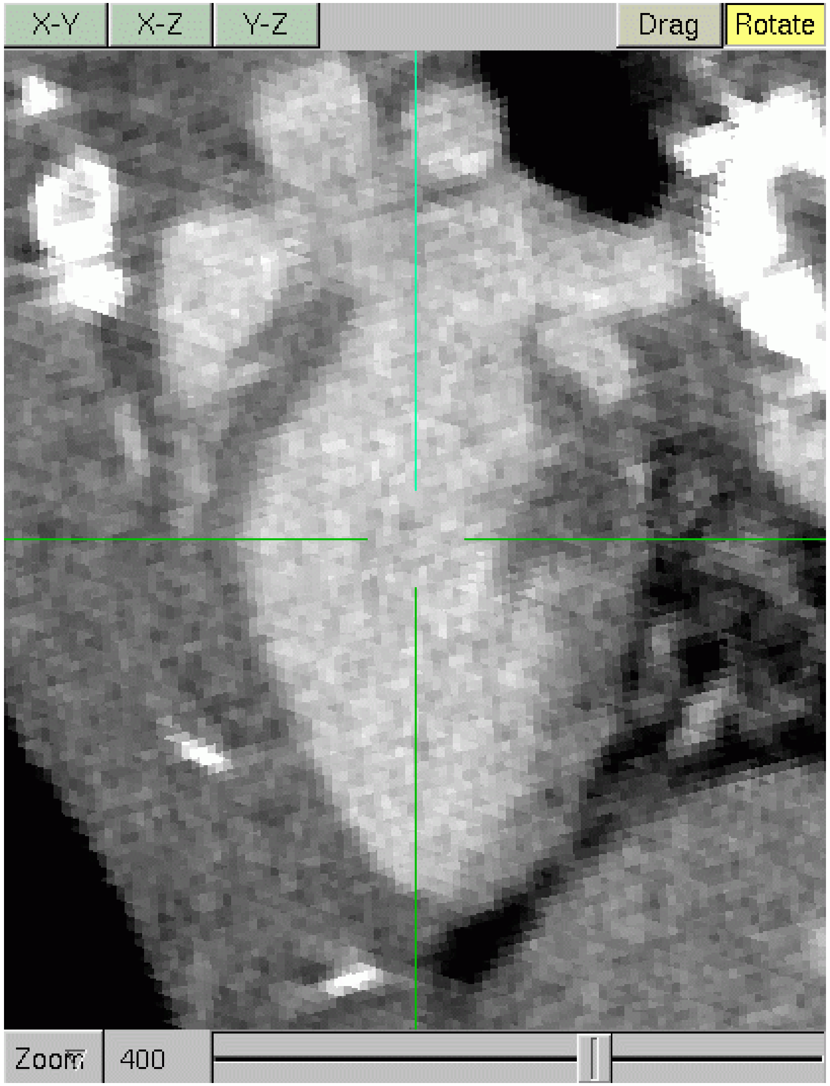

Interior detail from high resolution MRM fixed mouse heart shown using PSC-VB with 2X wavelet expansion.

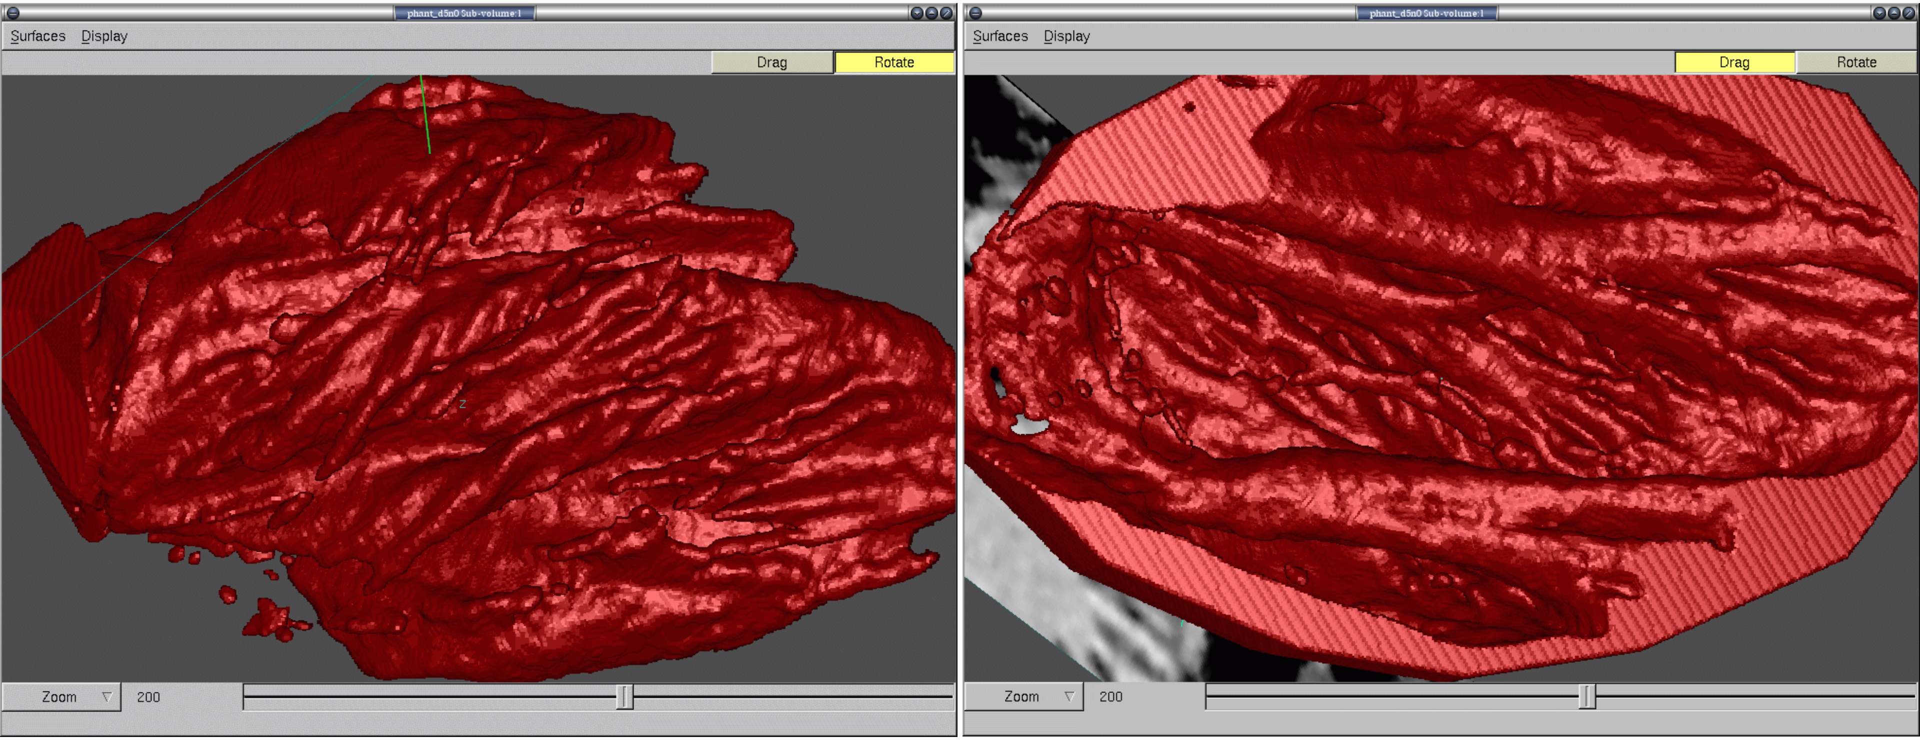

An alternative approach to LV volume measurement

- Take advantage of known binary mixture model
- Avoid difficulties of segmentation methods
- Account for unresolved detail and motion
- Tolerate high noise levels and artifacts
- Provide numerical error estimates
- Trade SNR against resolution - C. Shannon 1948
  - C = W log ((S+N)/N)
- Use all ROI data to form a simple ratio measurement
- Use measured result to constrain segmentations

Compute volume directly from gray values

FractBlood = (AvgROI – AvgMuscle) / (AvgBlood – AvgMuscle)

The resulting error is small despite high voxel noise.
The ROI volume is exactly known from its construction.

VolumeBlood = FractBlood * VolumeROI

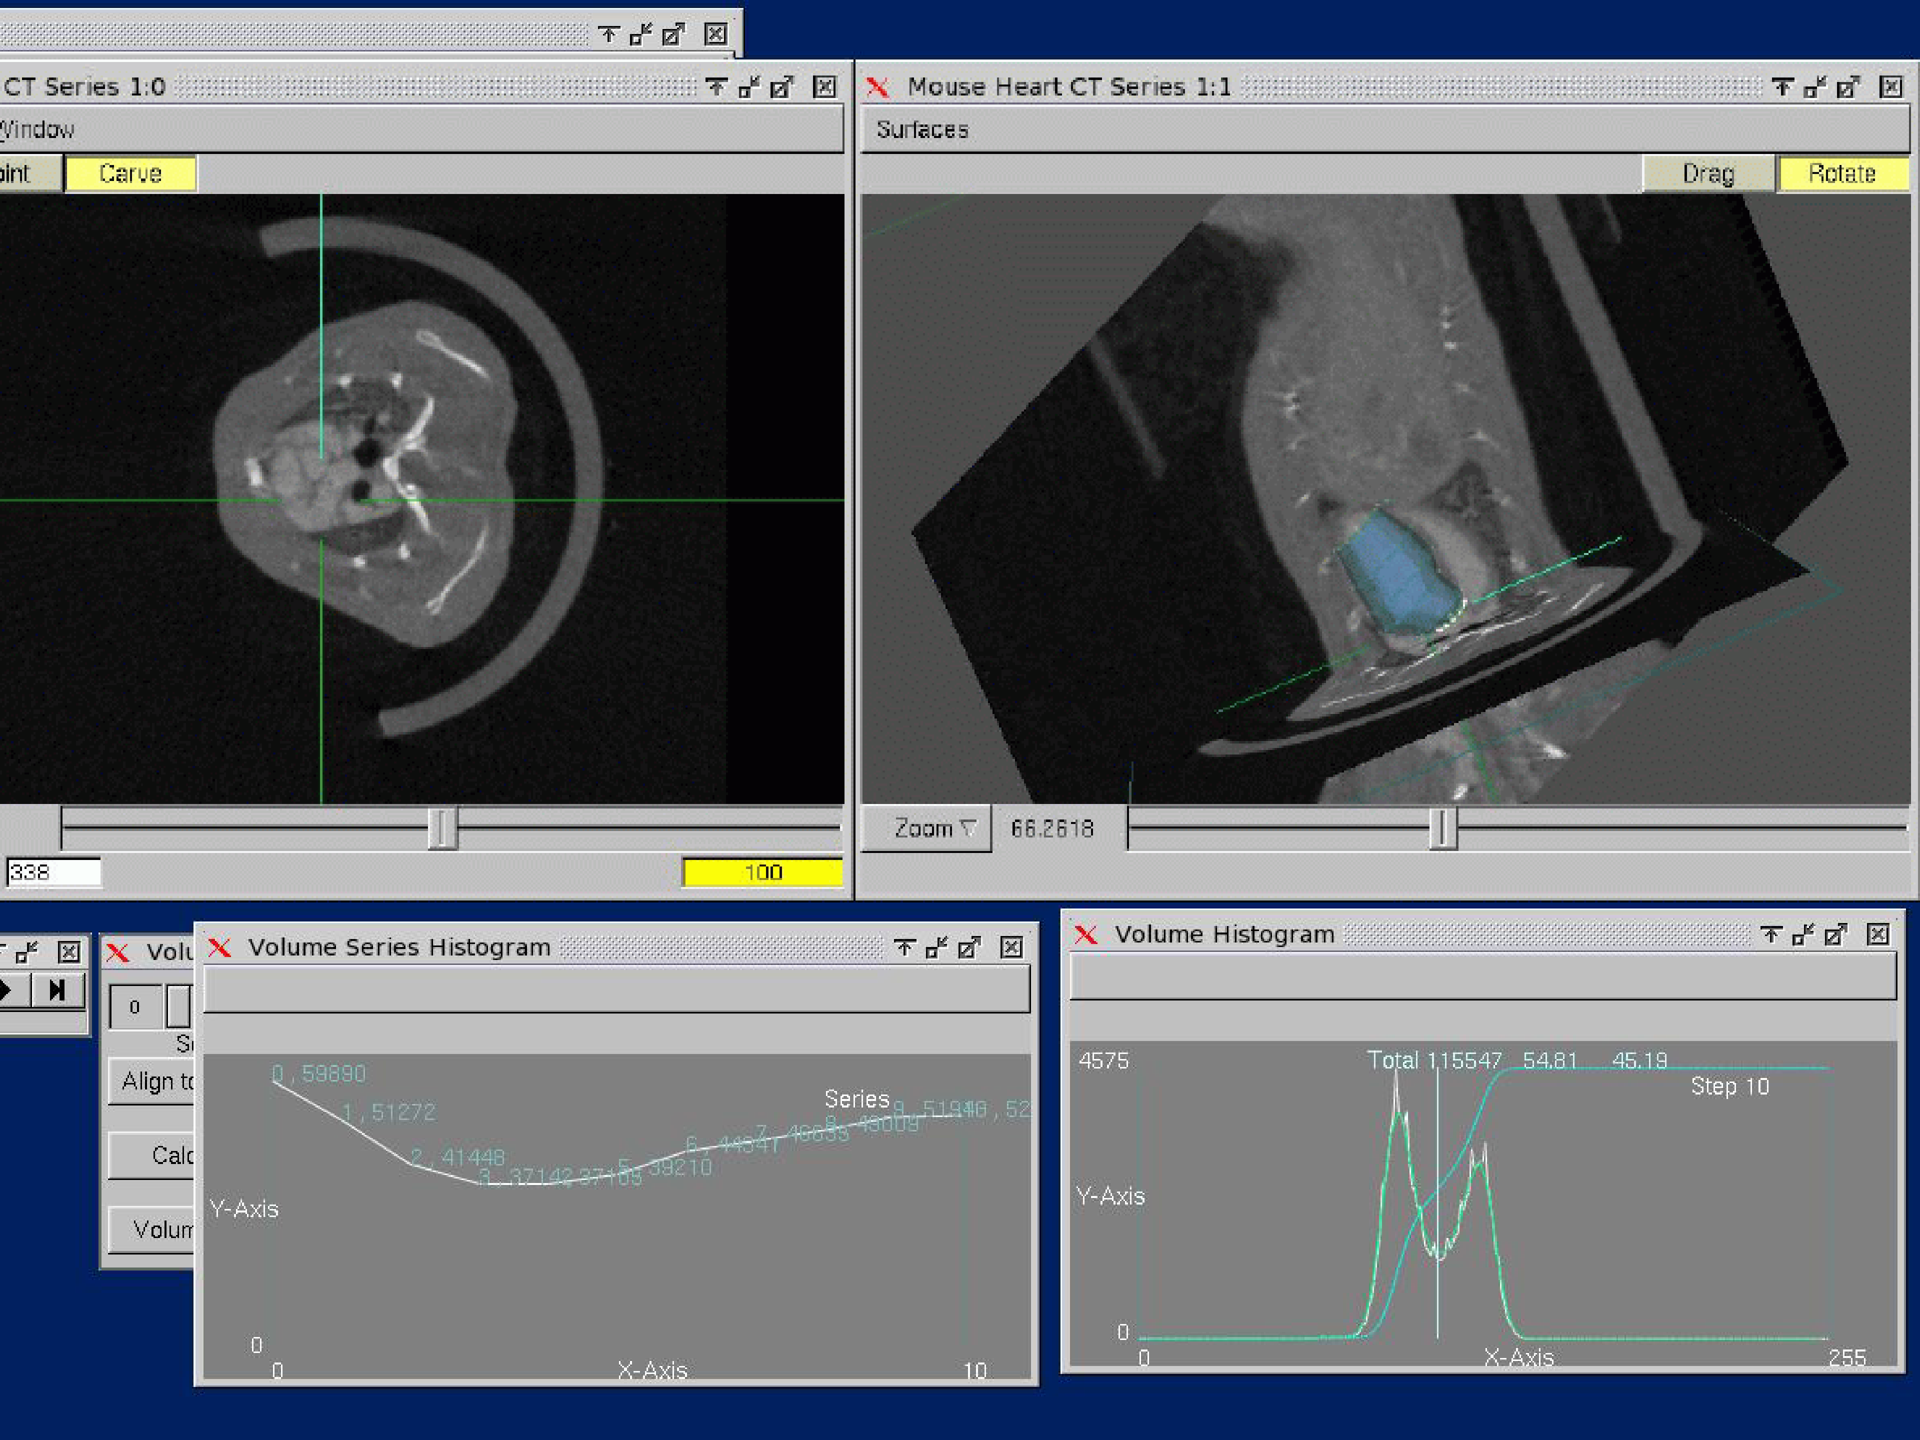
CT Series 1:0
Window
Carve
338
100
Mouse Heart CT Series 1:1
Surfaces
Drag
Rotate
Zoom
66.2618
Volume Series Histogram
Series
Y-Axis
X-Axis
0
10
Align to
Volume Histogram
4575
Total 115547
54.81
45.19
Step 10
Y-Axis
X-Axis
0
255

Nested ROI and targeted blood/muscle sampling

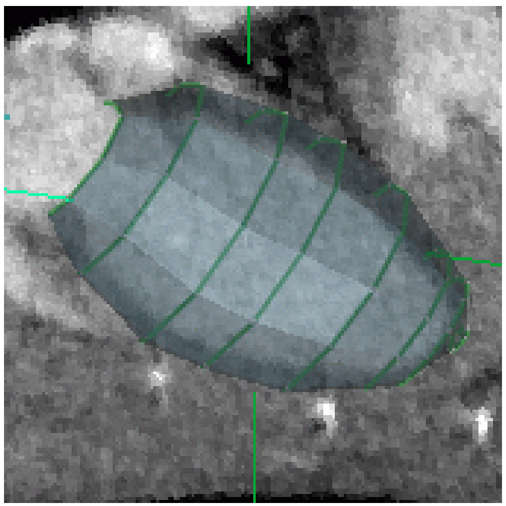

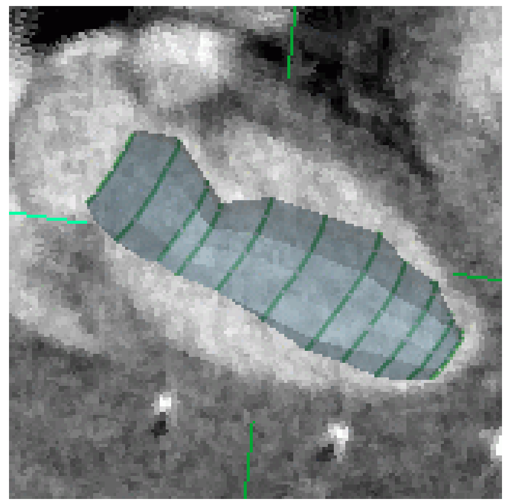

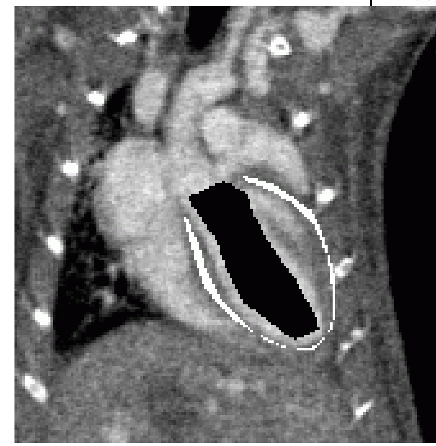

Histograms of isolated blood and muscle vs. contrast agent and #projections

0.125 ml contrast agent

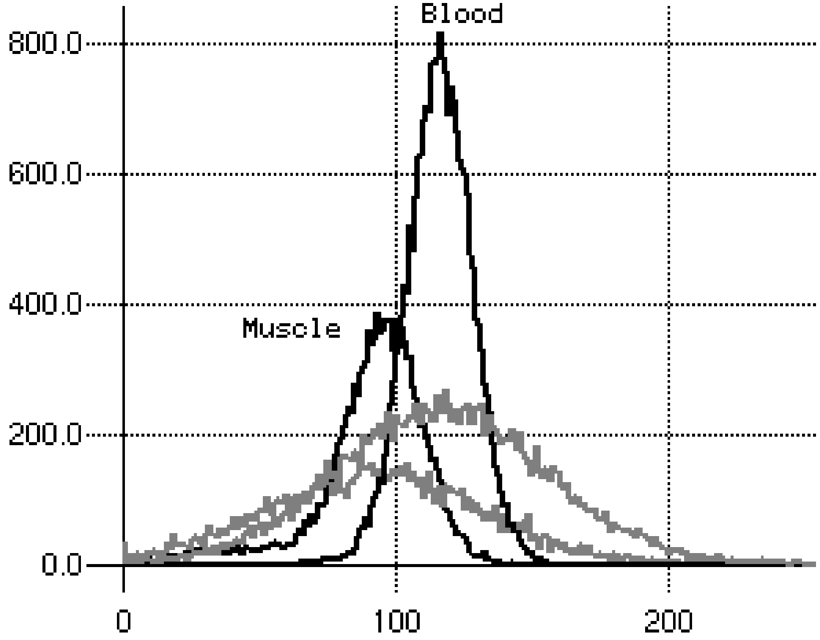

0.5 ml contrast agent

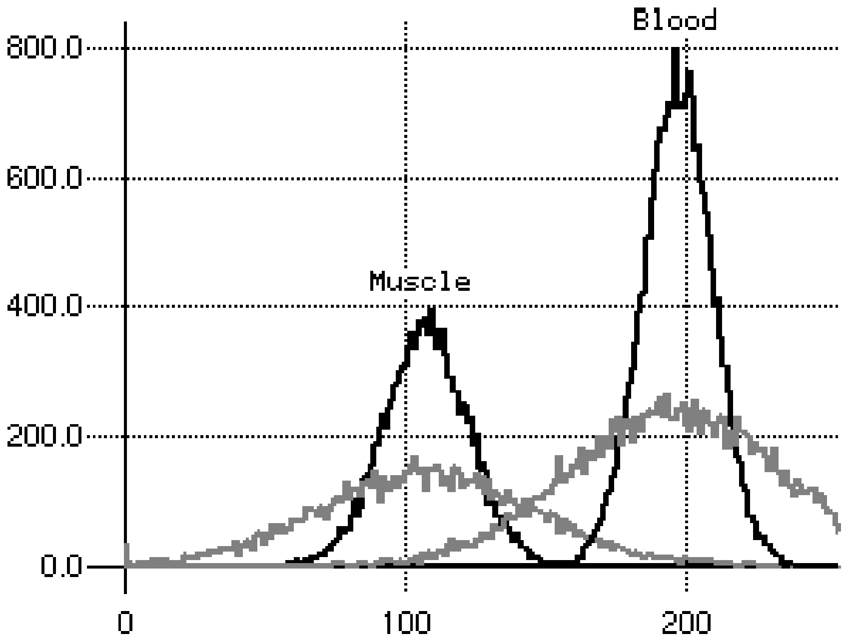

Black histograms are from 380 projections.

Gray histograms are from 63 projections.

Measured volumes and % standard error vs. contrast dose and #projections

Contrast 0.125 ml 0.250 ml 0.375 ml 0.5 ml

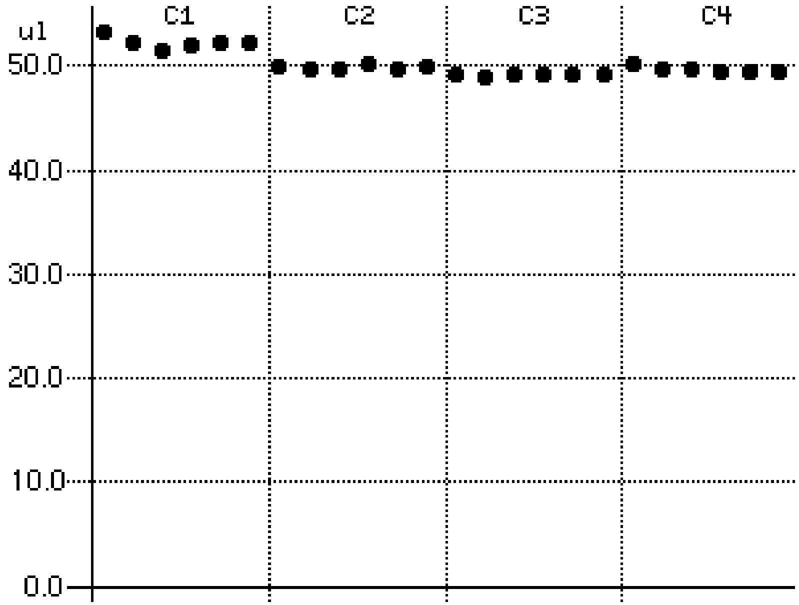

0.125 ml 0.250 ml 0.375 ml 0.5 ml

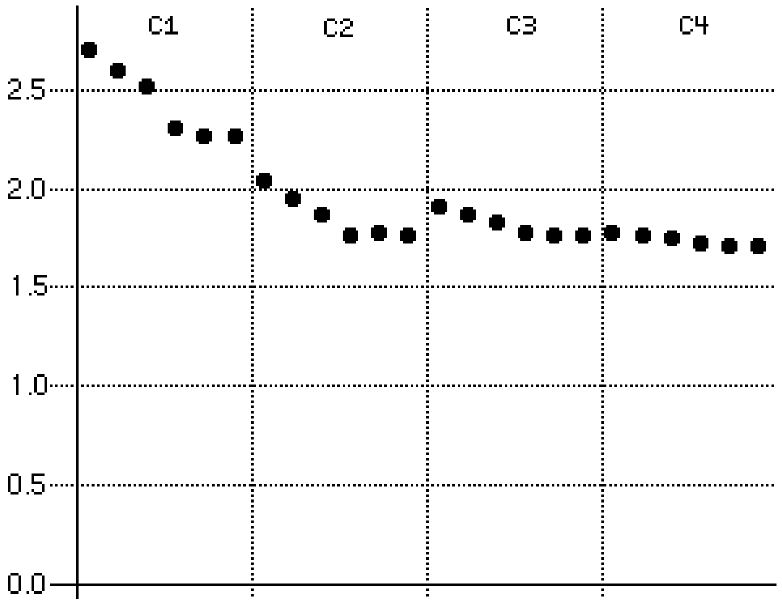

Within each column the 6 points are increasing # of projections through the series 63, 75, 95, 126, 190 and 380 resulting in a high to low noise progression.

LV volume changes vs. time

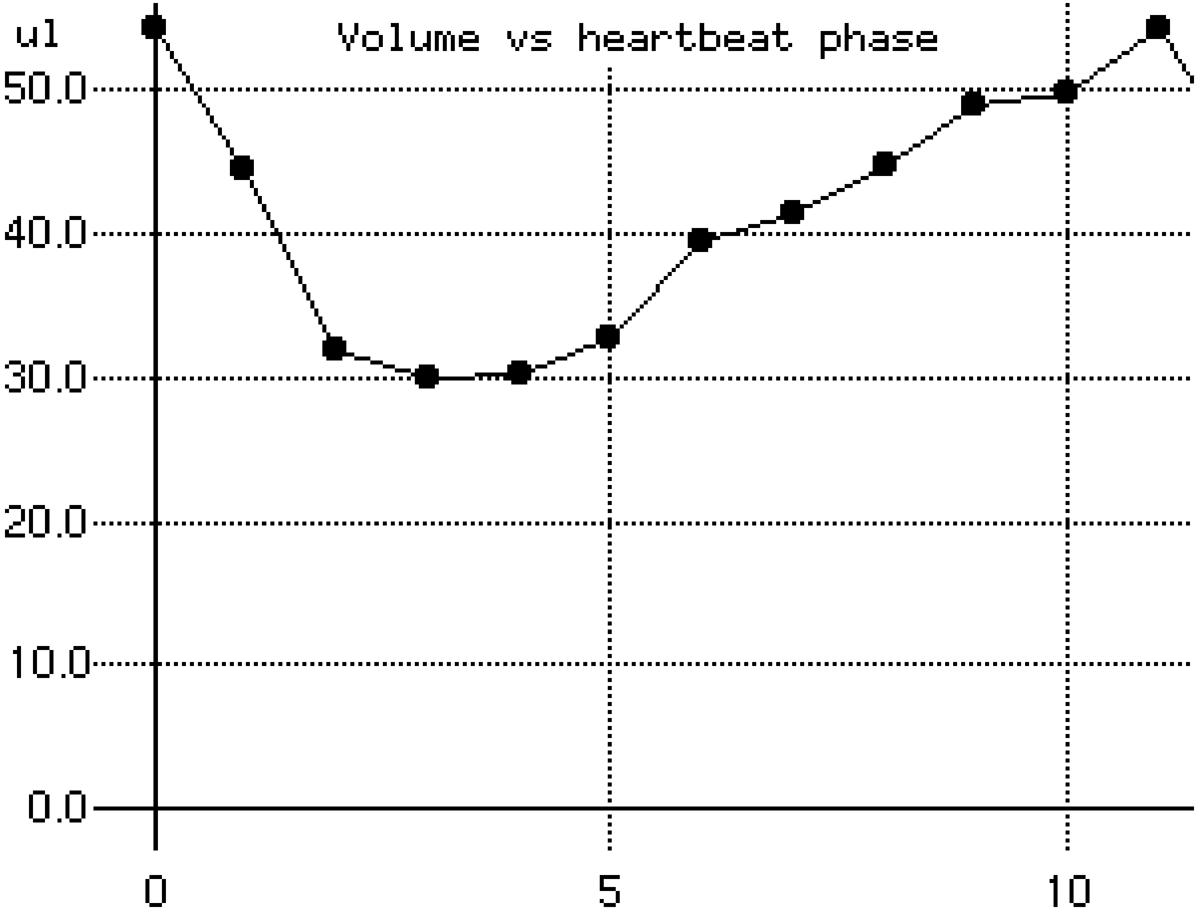

Dr. Allan Johnson, director of the Duke CIVM, discusses the impact of our collaborative project.

Acknowledgements

This work was primarily supported by NLM contract N01-LM-9-3531. National Resource grant RR006009 to the National Resource for Biomedical Supercomputing at the Pittsburgh Supercomputing Center also supported aspects of algorithm development. Portions of the experimental work performed at the Duke Center for In Vivo Microscopy, an NCRR/NCI National Biomedical Technology Resource Center were also supported by grants (P41 RR005959/ R24 CA-092656). All animal studies were conducted at the CIVM under protocols approved by the Duke University Institutional Animal Care and Use Committee.

We particularly appreciate the participation and encouragement of the late Dr. David W. Deerfield II who was instrumental in initiating the PSC/CIVM collaboration that lead to this work. We also thank all the members of the 4D mouse project team who have helped with evaluation, networking and other essential work.

Additional Support

NCRR P41 RR006009
PSC National Resource for Biomedical Supercomputing

NIH/NCRR P41 RR005959
Duke Center for in Vivo Microscopy

awetzel@psc.edu
Pittsburgh Supercomputing Center
300 South Craig St.
Pittsburgh, PA 15213
412-268-3912
FAX: 412-268-8200

NRBSC

PITTSBURGH SUPERCOMPUTING CENTER