Research Computing at BU: Introduction to SCV

www.bu.edu/tech/research/scv/

Erik Brisson
ebrisson@bu.edu

March 28, 2013

BOSTON UNIVERSITY

Introduction to IS&T / SCV (Scientific Computing and Visualization)

- Supports Research Computing
- Addresses high-performance computing and visualization needs of the BU community
- Manages BU's Scientific Computing Facility (SCF), a set of high-performance research computing and storage systems
- Offers consulting and training services
    - Live and online tutorials
    - Short and long-term consulting for computing
        - Code porting, optimization, parallelization
    - Short and long-term consulting for visualization
        - Data transformation, graphics coding and production

Systems and Storage

BU CCS / SCV Parallel Machines

- 1988: Thinking Machines CM-2
- 1992: Thinking Machines CM-5 (No. 59 on Top 500 list)
- 1996: SGI Origin2000 (No. 212 on Top 500 list)
- 1999: SGI Origin2000 modules interconnected to form 192-proc machine (No. 113)
- 2000: IBM RS/6000 SP
- 2002: IBM p690
- 2003: IBM p655
- 2005: IBM Blue Gene (No. 59 on Top 500 list)
- 2007: 54-processor Blade Center
- 2008: added 32 processors to Blade Center
    - Increased disk capacity to 75 TB, shared by all machines
- 2009: added 176 processors to Blade Center

Current Hardware - BladeCenter

- IBM BladeCenter
  - Katana
  - 173 nodes (blade servers)
  - heterogeneous mixture of blades
  - currently 1580 processors

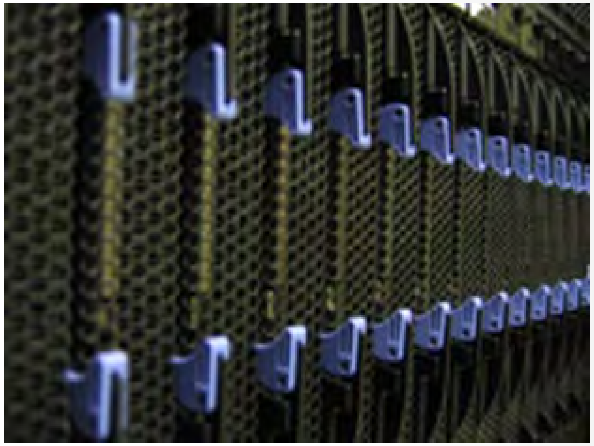

- Intel Xeon and AMD Opteron
  - 2.4-3.0 GHz
  - blades contain 4, 8, or 12 processors
  - May use up to 64 processors at a time
    - parallel or serial
  - 8 … 128 GB memory

Current Hardware – Blue Gene

- IBM Blue Gene
  - 1024 nodes, each with 2 processors
  - PowerPC 440 processors
    - 700 MHz
    - 32-bit
  - only runs MPI codes
  - each processor is relatively slow, so you need good scalability to make it worthwhile
  - May use 1024 processors during the day, all 2048 processors at night

Next Generation (MGHPCC)

- Dell C8000/C8220 servers
  - 64 node
  - 1024 cores
  - 9.26 TB memory
  - 10.6 TFLOP/S

Nodes	Network	Mem/Node	Total Cores	Total Mem	GFLOPS
36	FDR IB	128	576	4,608	5,990
20	GigE	128	320	2,560	3,328
8	10GigE	256	128	2,048	1,331

Computer Graphics Lab

- Six Workstations
  - Two 64-bit Linux, 8 cores, 16 GB, 2 Nvidia Quadro FX1700 boards
  - 64-bit Linux, 2 cores, 4 GB, Quadro FX1300
  - 32-bit Linux, 1 core, 2 GB, Quadro FX1100
  - 64-bit Windows 7, 8 cores, 7 GB, Quadro FX1700
  - 32-bit Windows XP, 1 core, 2 GB, Quadro

- Tiled Stereo Display Wall
  - Eight projectors (gives four tiles per eye)
  - Commodity hardware
  - Passive stereo using linear polarization

Accounts and Policies

- Accounts and allocations are based on “projects”
- Approximately 140 projects and 600 users
- Must be faculty or research staff to apply for projects
  - Project PI then adds accounts for students, collaborators, etc.
  - May assign post-doc or admin. staff as “administrative contact”
- Apply for project at

  http://www.bu.edu/tech/accounts/special/research/accounts/applications/
  - Click on “Boston University faculty and research staff may apply for a new project.”
  - If you’re not sure how much time to ask for, 1000 katana hours may be a good start
  - The form has fields for requesting time on each machine, but the time is actually all in one pot, and can be spent on any machine
  - This will give you accounts on all machines except for Blue Gene

Accounts and Policies (cont'd)

- For Blue Gene account, once you are awarded a "regular" account go back to
  - http://www.bu.edu/tech/accounts/special/research/accounts/applications/
- There is a paragraph starting with "Blue Gene accounts." Click on the "SCF User Information Page" link.
  - You will be requested for your user name and password
  - Click on the "Update Personal Information" link

Storage

- Storage for SCF systems is provided by a 84 terabyte disk array
- Backed-up, non-backed-up, and and scratch space available
- SCV archive service currently provided by a robotic tape system capable of storing 935 terabytes (current configuration provides 500 terabytes)
- Recently deployed IS&T archive service
  - http://www.bu.edu/tech/infrastructure/hosting/data-archiving/
- Will change as MGHPCC comes online

Disk Space

- The disk space in your home directory is minimal
- Most project PI's request a "project" directory
  - project directory contains disk space allocated to the specific project
  - once your account has been activated go back to

  http://www.bu.edu/tech/accounts/special/research/accounts/applications/

  and click "request a Project Disk Space allocation."
  - a request of a few GB will be rubber-stamped
  - large requests (hundreds of GB) will require stronger justification
- Allocations of non-backed-up work space available

Software and Services

Software

- All machines run Unix or Linux
- C, C++, Fortran, Java, Perl, Python, Tcl
- IDL, Matlab (includes PCT, parallel computing toolbox)
- Maple, Mathematica
- Gauss
- SAS, Stata, R (SPSS may be coming?)
- VTK, Paraview
- Maya, OpenGL, OpenSceneGraph

Scientific Programming and Visualization Consulting

- Scientific Programming Consulting
  - Software package usage
  - Code porting
  - Measurement and tuning
  - Numerical methods and algorithms
  - Parallelization
  - Statistics Programming
- Visualization Consulting
  - Software package usage
  - Visualization techniques
  - Algorithms and code development
  - Data formats and conversion

Training

- Tutorials
  - ~25 topics, multiple sessions each semester
  - Intro to Research Computing: Linux, Data Management
  - Basic Programming: C, Fortran, Python, Matlab
  - Parallelization: OpenMP, MPI, Matlab PCT
  - Viz: Paraview, VTK, OpenGL, SceneGraph, Matlab, Maya, Diagrams and Graphs
  - Statistics Programming: SAS, SPSS, R
- In-class presentations
- On-line (Web) help and training

Tutorial typical offerings (Fall, Spring, Summer)

- Introduction to Fortran Programming
- Introduction to C Programming
- Introduction to Python Programming
- Introduction to R Programming
- Introduction to MATLAB
- Code Tuning and Optimization
- Parallelization with OpenMP
- Tuning MATLAB Codes For Better Performance
- Introduction to MPI
- MATLAB Parallel Computing
- Introduction to Scientific Visualization
- Scientific Visualization Using ParaView
- Scientific Visualization Using VTK
- Scientific Visualization Using MATLAB
- Graphics Programming in C/C++: OpenGL and OpenSceneGraph
- Graphics and Images for Publication and Presentation
- Diagrams, Graphs, and PowerPoint: How to Create an Effective Presentation
- Introduction to Maya

Project Examples

Example Projects:

Lidar Imaging - Point Clouds

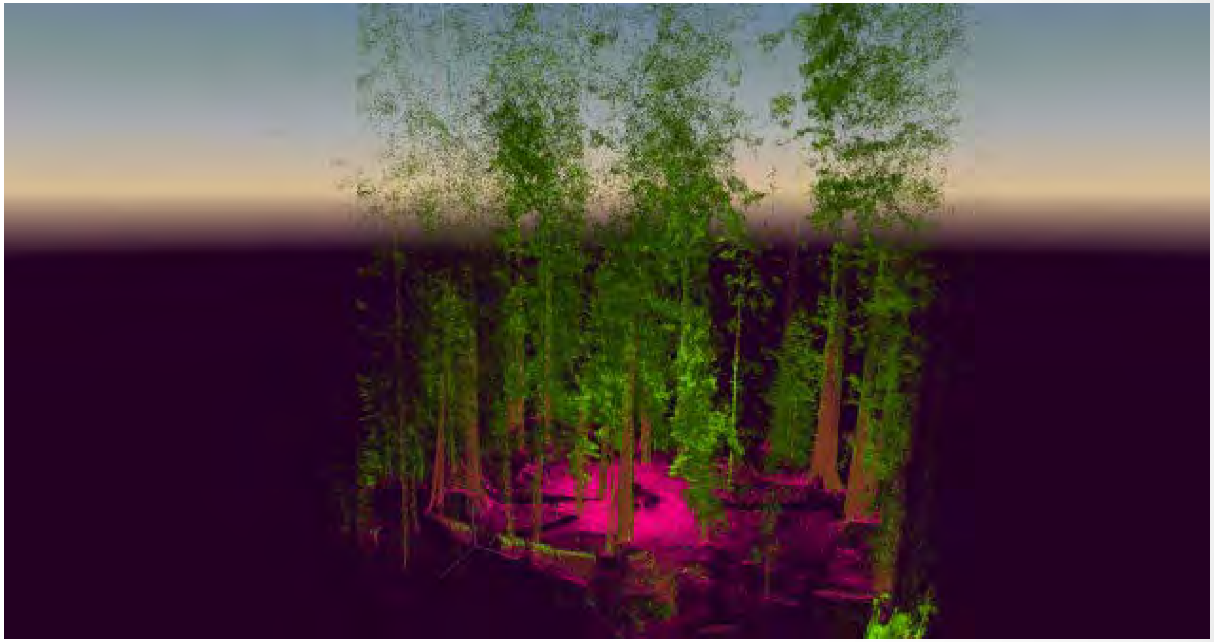

Tomographic Reconstruction of Forest Structure.
Xiaoyuan Yang1, Alan Strahler1, and Erik Brisson, BU (2008)

Example Projects:

Adding Functionality to OCAVS

Omphalitis Community Based Algorithm Validation Study (OCAVS)
BU Center for Global Health & Development

Julie Herlihy, Davidson Hamer, Katherine Semrau, Kojo Yeboah-Antwi, Arthur Mazimba, Caroline Grogan

Wrote PHP scripts to extend functionality beyond that provided by eMocha, allowing:

- US-based domain expert to evaluate photos of infants;
- Do remote diagnosis "blind" to other patient information;
- Record diagnosis and associated notes;
- Without changing underlying eMocha installation

Example Projects:

Flow Over a Cylinder

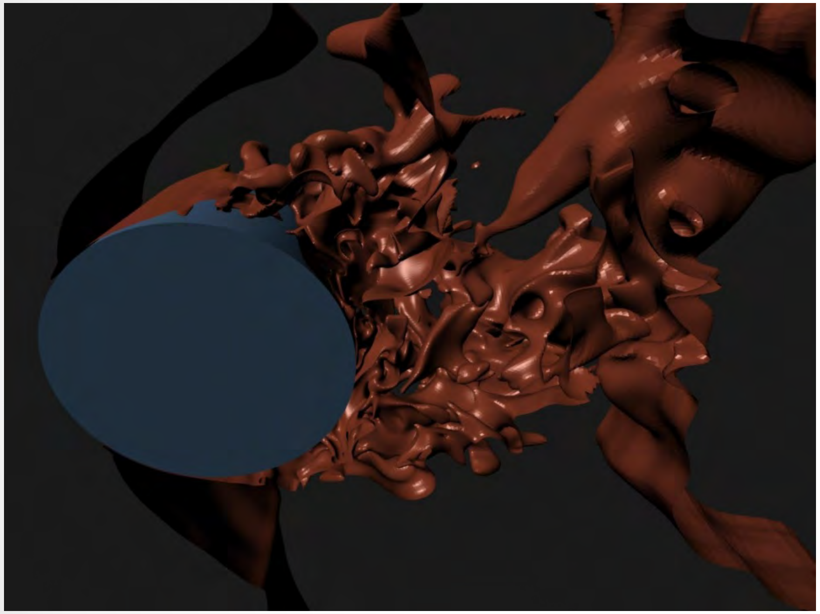

Simulation of vortex shedding over a right circular cylinder.
Douglas L. Sondak and Erik Brisson BU IS&T SCV (2009)

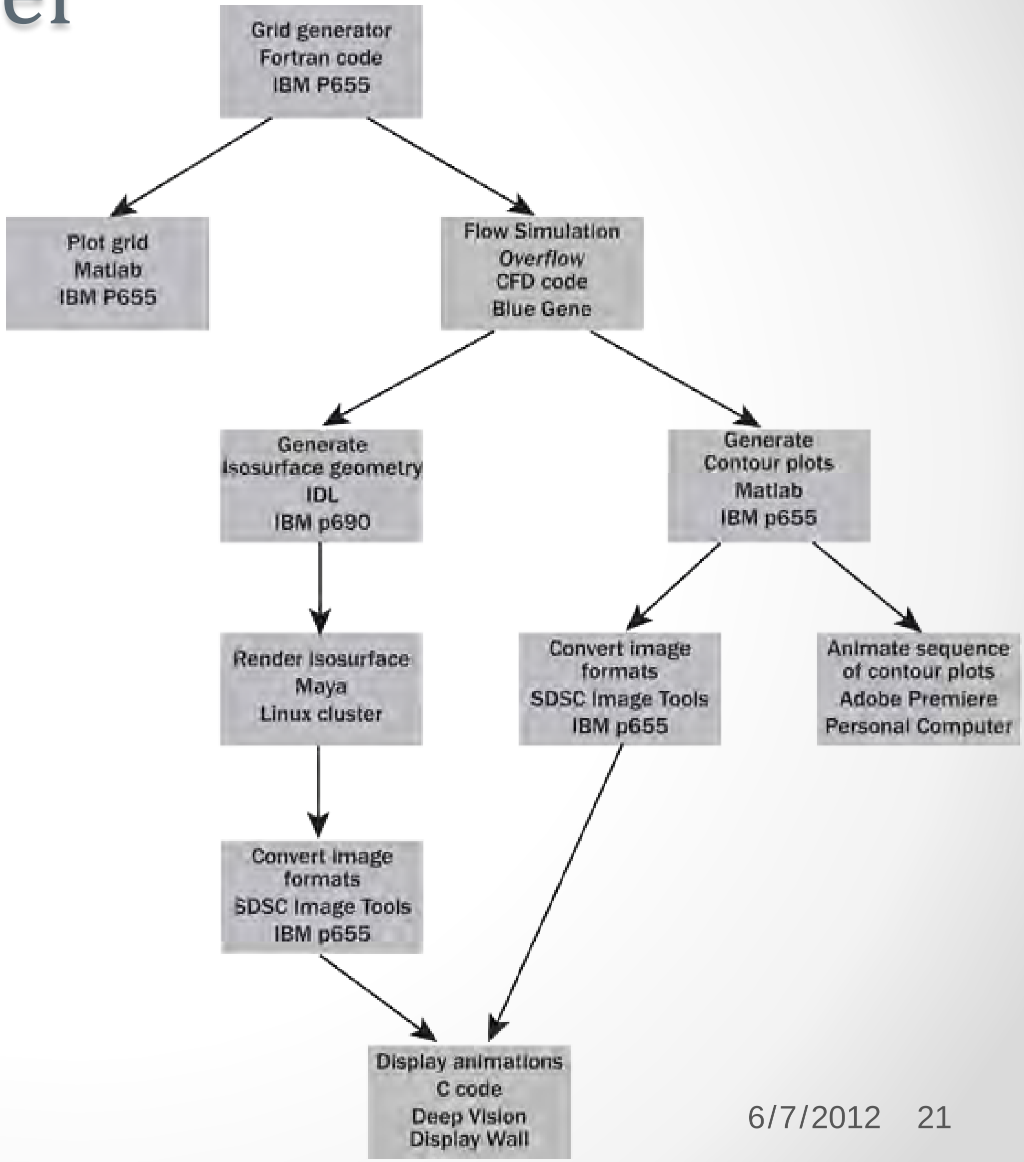

Example Projects - Brain EEG

- Nitin Bangera et al., Biomedical Engineering
- Model anisotropy of soft tissue and effects on electrical field in brain using MRI scans
- Compare finite-element simulation (ABAQUS) with EEG data

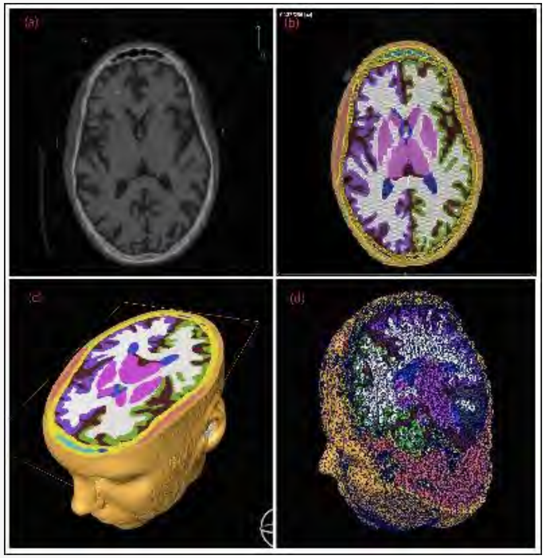

FEM Model Generation

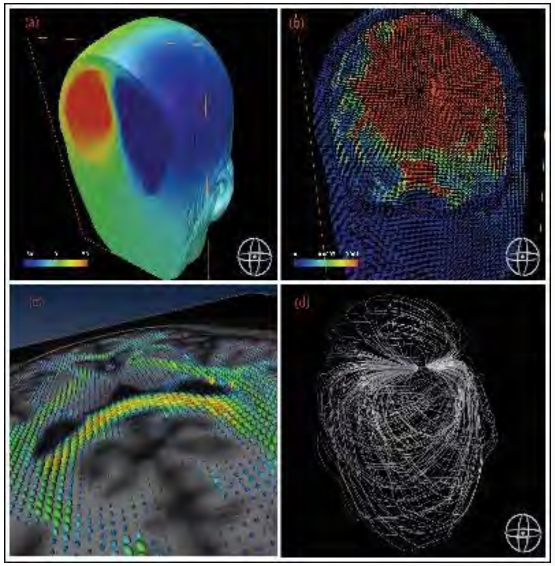

Visualization of (a) Scalar Field: Electric Potential (b) Vector Field: Electric Current Density Vector (c) Conductivity Tensor ellipsoids (d) Vector Field: Current Density

- www.bu.edu/tech/research/visualization/about/gallery/anisotropy/

Example Projects – Melt Reconstruction

Lab-produced “volcanic melt” using high-pressure piston
Polished/sliced and scanned using electron microscope
3D reconstruction from slice

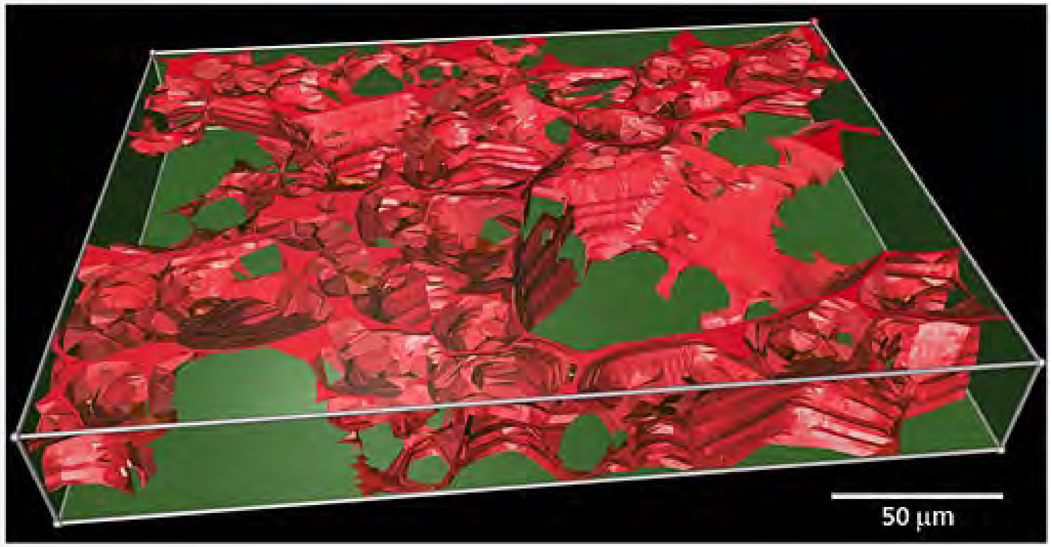

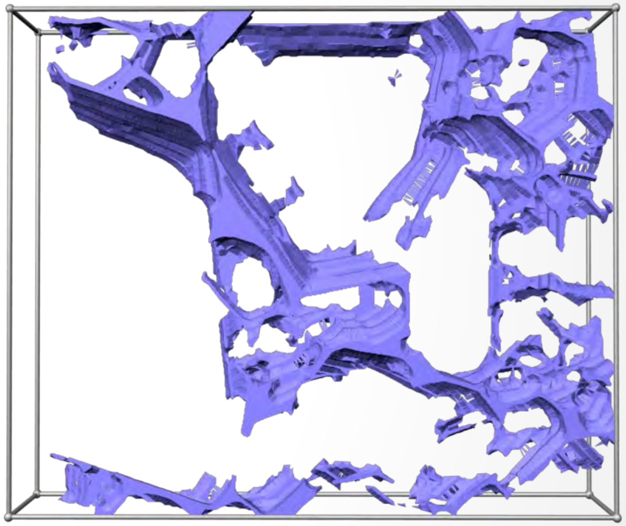

Gordana Garapić, Ulrich H. Faul, Dept of Earth and Environment
Erik Brisson, SCV

Massachusetts Green High Performance Computing Center (MGHPCC)

MGHPCC Collaboration

- University Partners
  - Boston University
  - Harvard
  - MIT
  - U. Mass
  - Northeastern

- Government Partners
  - Commonwealth of Massachusetts
  - City of Holyoke

- Industry Partners
  - Cisco
  - EMC

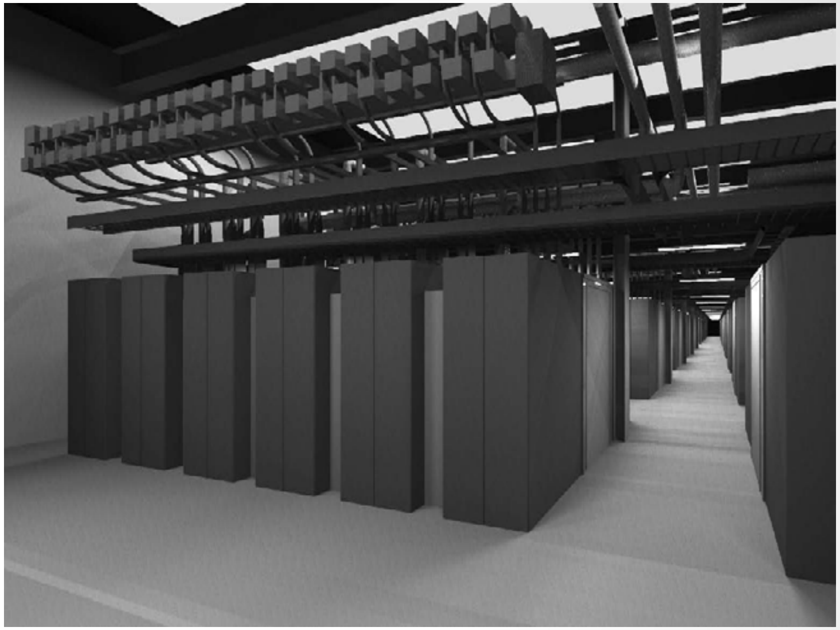

Holyoke Canal District Site

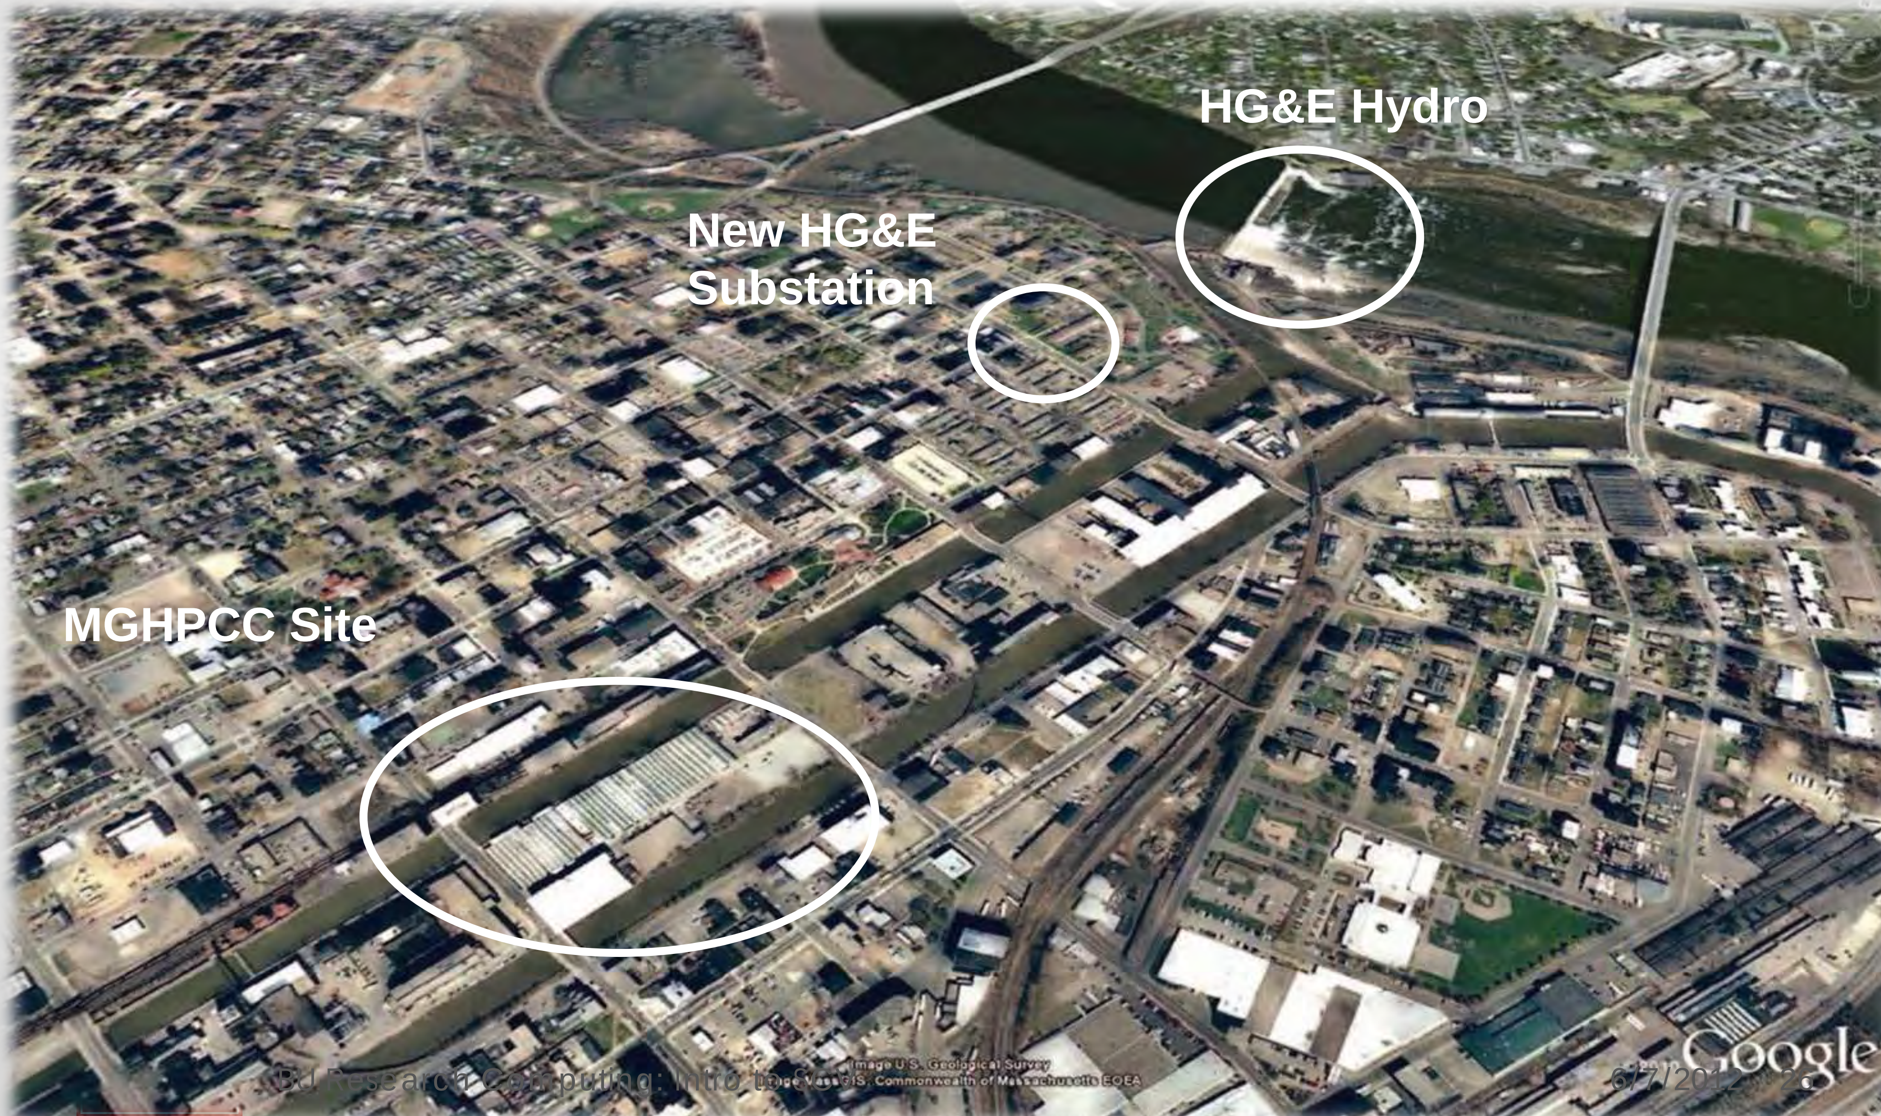

MGHPCC Benefits

- Inexpensive power
- Renewable energy source with low carbon footprint
- Inexpensive property; brownfield cleanup
- Revitalization of economically depressed region
- Low PUE green design
- Modern, controlled data center facility
- Flexible space within initial core & shell to accommodate 10 year growth
- Space on-site for building expansion (years 10-20)
- Opportunities for shared facilities and services
- Opportunities for collaboration with other institutions

Phasing

- Phase 1: years 1-10
  - Day 1
    - Core and shell for years 1-10
    - Building support systems for Day 2 build-out
    - Outfit for 388 IT racks; ~5 MW load
  - Day 2
    - Add chillers & generators
    - Complete outfit to 680 IT racks; 10 MW IT load
- Phase 2: years 10-20
  - New/expansion building on adjacent lot as needed

MGHPCC Operating Principles

- Equal share of resources for each of the university founding members
- Equal share of capital and common operating costs
- Lease arrangements for excess capacity
- 10% of IT capacity for non-members
- Operated by MGHPCC, Inc.

MGHPCC, Inc.

Responsibilities & Services

- Common shared infrastructure
  - Administration and business operations
  - Facilities operation & maintenance
  - Education, outreach & training spaces
  - Common networking
  - Security
- Power and cooling
  - Metered and billed by institution
- “Virtual hands” operational support
  - Minimal time included in shared cost basis
  - Additional support billed on per use basis

Compute Floor Layout

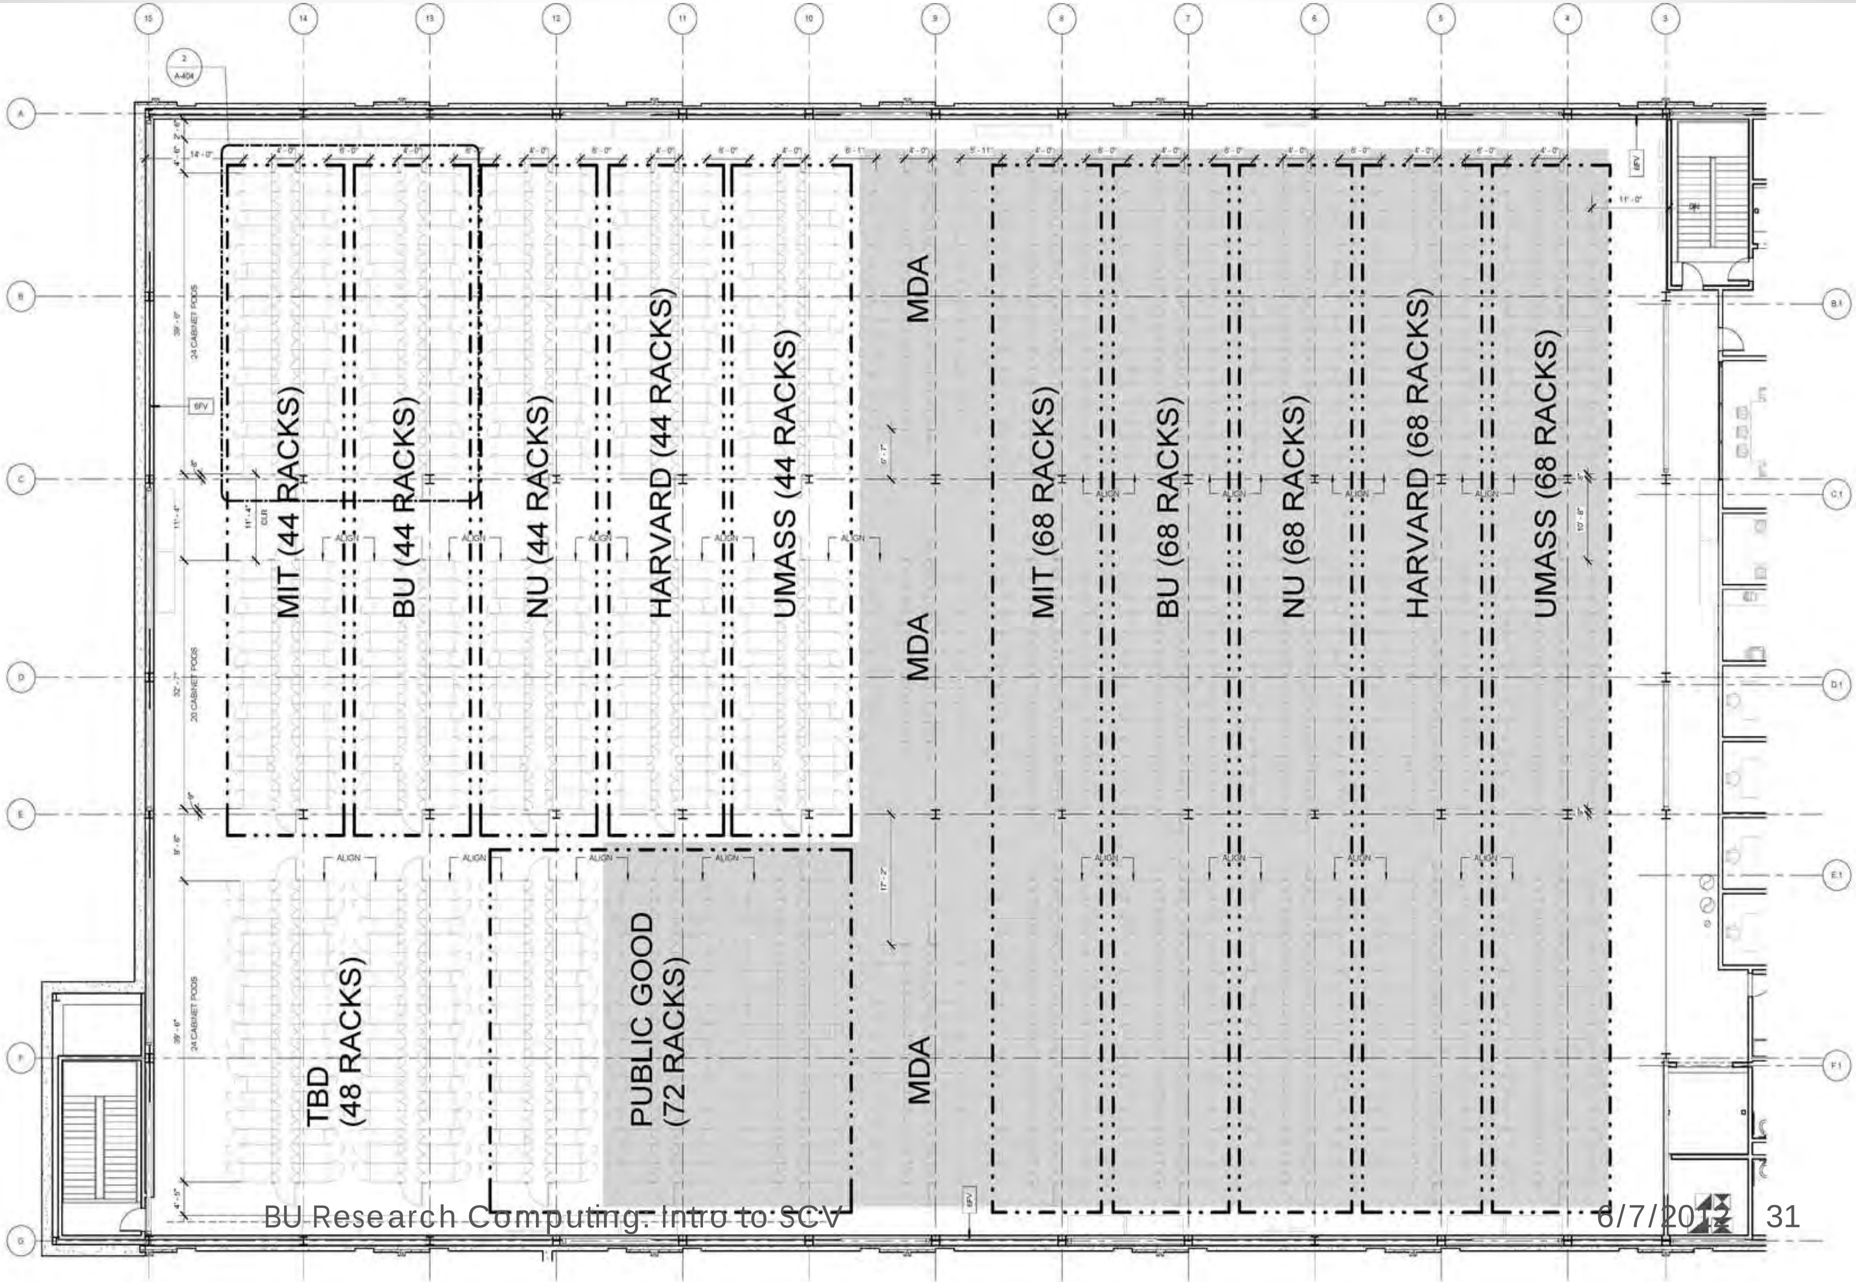

Proposed IS&T Service Models

- University pays for
  - MGHPCC shared operating costs (annual cost-basis)
  - Power
  - Cooling
  - Heat containment
  - *Racks, bus-plugs, PDUs? (TBD)*
  - WAN (shared) & zone networking
- Other services based on selected service model
  - ***Shared, Co-op/buy-in, Dedicated, Co-lo***

Subject to Research Computing Governance and budget approvals

IS&T Shared Service Models:

Shared, Buy-in/Co-op, Dedicated, Co-lo

- Shared
  - Acquired through central funding or institutional level infrastructure grants
  - Offered without charge to all faculty/research staff on an allocation basis; Allocations reviewed through Center for Computational Science
  - Jobs are fair-share scheduled
  - User consulting and support services
  - Shared storage facilities (including backup & archive)
  - Examples: SCF BG/L, SCF Linux (Katana)

IS&T Shared Service Models

- Buy-in (co-op/condo)
  - Standardized hardware which is integrated into the shared facility with priority for owner
    - Plan to provide multiple vendor options
  - Managed centrally by IS&T, but purchased by individual researchers
  - Scale-out to shared computing pool
  - Shared storage facilities (including backup & archive)
  - Priority access to owners; excess capacity shared
  - Standard services, including user support, are provided without charge
  - Examples: CISM (p655), CAS/Geography (Linux)

IS&T Shared Service Models

- Dedicated
  - Hosted and managed centrally by IS&T for dedicated use by owner
  - Systems purchased under individual grants
  - Systems administration by IS&T (% FTE) paid from grant
  - Physical infrastructure is provided without charge
  - Equipment, software licenses and other costs are paid directly by the researcher
  - Usage policies are set by the owner
  - User support services provided by owner
  - Example: ATLAS

IS&T Shared Service Models

- Co-lo
  - Intended for larger clusters (one or more full racks)
  - Hosted centrally, systems managed locally
  - Physical infrastructure is provided without charge
  - Computer equipment, cluster interconnects, in-rack networking purchased directly
  - Virtual hands services, including cabling, racking, etc charged back to owner; coordinated through IS&T/ MGHPCC
  - Systems administration, system security, software licenses, user support and other management provided by owner
  - Owner responsible for 7x24 emergency response

The End

SCV start page

www.bu.edu/tech/research/scv/

This presentation is online here:

scv.bu.edu/documentation/presentations/CTSI_2012/