



## The Massachusetts Green High Performance Computing Center

27 April 2012 updated from original presentation at  
7 February, 2012 Engineering Deans Conference, Washington DC

**Jack M. Wilson,**  
MGHPCC Founding Chair 2009-2011  
President Emeritus and Distinguished Professor of Higher Education, Emerging Technologies and Innovation  
The University of Massachusetts



# Agenda: Mass Green High Performance Computing Center

---

- MGHPCC Vision
- MGHPCC history, goals, and organization
- MGHPCC resources and services
- MGHPCC consortium for research and education

# VISION

---

- World-class Green High Performance Computing Center (1<sup>st</sup> in the Northeast)
- First multi-university HPC facility of its kind in the nation
- Platform for collaboration in R&D that will strengthen Mass R&D leadership in computing applications
- Catalyst for economic & workforce development in the state, region and Holyoke (Innovation District)
- Most significant state/industry/university partnership in state history



## MGHPCC History, Goals, and Organization



# MGHPCC

## A partnership between 5 universities



\$736.1M



\$468.7M



\$280.8M



\$61.3M



HARVARD  
UNIVERSITY

\$462.3

**Total Research Revenue in 2009 of \$2,009,078,000 ( \$2.0B )**

(Source NSF: [http://www.nsf.gov/statistics/nsf11313/content.cfm?pub\\_id=4065&id=2](http://www.nsf.gov/statistics/nsf11313/content.cfm?pub_id=4065&id=2))

**With additional support from the commonwealth and industrial sponsors**





- 1/6/09 Discussion: Boston Research University Presidents' Dinner
  - At MIT President, Susan Hockfield's, home.
- 1/9/09 Friday call from Susan Hockfield, MIT to Jack Wilson, UMass
  - The Scientific Fantasy: Boston Research Universities Build MGHPCC
- 1/27/09 Meet with Gov. Deval Patrick and Sec. of Econ. Dev.
- 1/28/09: MIT & UMass Teams first meet in my office -fantasy into reality
  - City closed by snow
    - Jack Wilson, President, UMass
    - Rafael Reif, Provost, MIT
    - Claude Canizares, Vice President for Research and Associate Provost, MIT
    - James Kurose, Dean of Natural Sciences and Mathematics, UMass Amherst
    - Rick Adrion, Professor and Past Chair of Computer Science, UMass Amherst
    - Tom Chmura, Vice President for Economic Development, UMass  
(Titles as of 2009)



- Susan and I decided to expand the group by chatting with other Presidents.
  - She called Drew Faust, Harvard
  - I called Bob Brown, Boston University
    - Brown: You called the wrong guy! I am the only person in the world to fail twice at building a HPCC! Brown was the former Provost at MIT
- CEO Joe Tucci, EMC and CEO John Chambers, CISCO agreed to help
- Accenture provides project management guidance
- Northeastern joins the group. Now the five largest research universities in Massachusetts are collaborating
- Many other companies involved in the conversation
  - Akamai, Google, Microsoft, IBM, etc.



- First steps are fundraising and developing a detailed scope of the project
- We were working in secrecy.
- Governor Patrick, Joe Tucci, John Chambers, I, and others do trade mission to California.
  - Governor speaks about this at every stop!
  - So much for secrecy.



## WHAT'S BEEN ACCOMPLISHED TO DATE?

- 5 leading research universities, EMC, Cisco and state committed
- 501C3 corporation formed to develop/manage facility
- \$95 M from universities, state, industry, tax credits
- Site selected/acquired, tech. requirements identified, planning/design completed
- John Goodhue (former VP at Cisco), is appointed as Executive Director
- Several joint proposals submitted –some granted



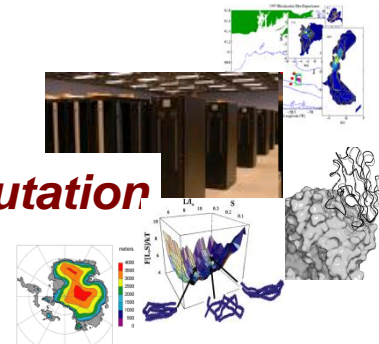
# WHAT IS HPC?

- Use of large clusters of computers or special-purpose computers in advanced applications – “3<sup>rd</sup> Leg” of science
- Increasingly important infrastructure to support R&D in broad array of fields

$$\begin{aligned}
 \oint \mathbf{E} \cdot d\mathbf{A} &= \frac{q_{enc}}{\epsilon_0} & \mathbf{E} &= mc^2 \frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho u_i) = S_m \\
 \oint \mathbf{B} \cdot d\mathbf{A} &= 0 & & \frac{\partial (\rho u_i) + \frac{\partial}{\partial x_j} (\rho u_i u_j) = \\
 \oint \mathbf{E} \cdot d\mathbf{s} &= -\frac{d\Phi_B}{dt} & \text{theory} & \frac{\partial \rho}{\partial t} + \rho \mathbf{g}_i + F_i \\
 \oint \mathbf{B} \cdot d\mathbf{s} &= \mu_0 \epsilon_0 \frac{d\Phi_E}{dt} + \frac{\partial}{\partial x_j} (\rho u_j)
 \end{aligned}$$




experimentation ——— computation



# Computing as “Third Leg” of Scientific Research

Run virtual experiments


Make sense of vast amounts of data



# Why a Dedicated Facility for Scientific Computing?

- Lower operating cost
- More competitive research proposals

**Less of this...**



**More of this...**



# Why Cooperate?

## Current Examples

Research Project  
BU Atlas


PI Coalition  
MIT Bates

Department  
MIT CSAIL

Campus/School  
UMass Medical  
Harvard FAS  
Harvard Medical  
BU Medical

University-wide  
Purdue, Princeton  
Stanford, NYU  
UNH, Colorado State

University Coalition  
MGHPCC facility  
Multi-university MRI proposal



## Goals

**Preserve Autonomy / simplify operation**  
Local administration and application support  
Faster install  
Space/power/cooling/security taken care of

**Lower Operating Cost**  
Leveraged capital investment  
Locate for low cost of power  
Design for cooling efficiency  
Amortize staff cost/expertise

**Greater Funding Leverage**  
Pooling of research grant funds  
Greater leverage for university subsidies  
Higher impact for major infrastructure grants

**Access to larger resource pool**  
Control access to what you own  
Opportunity to share idle resources

# Why Holyoke?

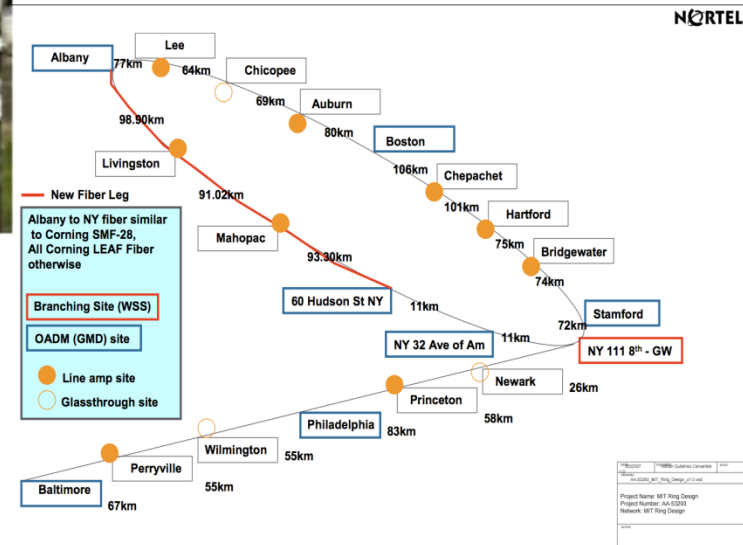


- Existing industrial site

- Green, low cost power



- Fiber Crossroads



# WHY HOLYOKE?

- Low-cost, clean energy (primarily hydro, potentially wind)
- Extremely low carbon footprint
- Superior IT connectivity, located at the IT cross-roads of New England
- Economic development incentives for developing in downtown “canal district”



# Hydropower that drove the mills drives the MGHPCC

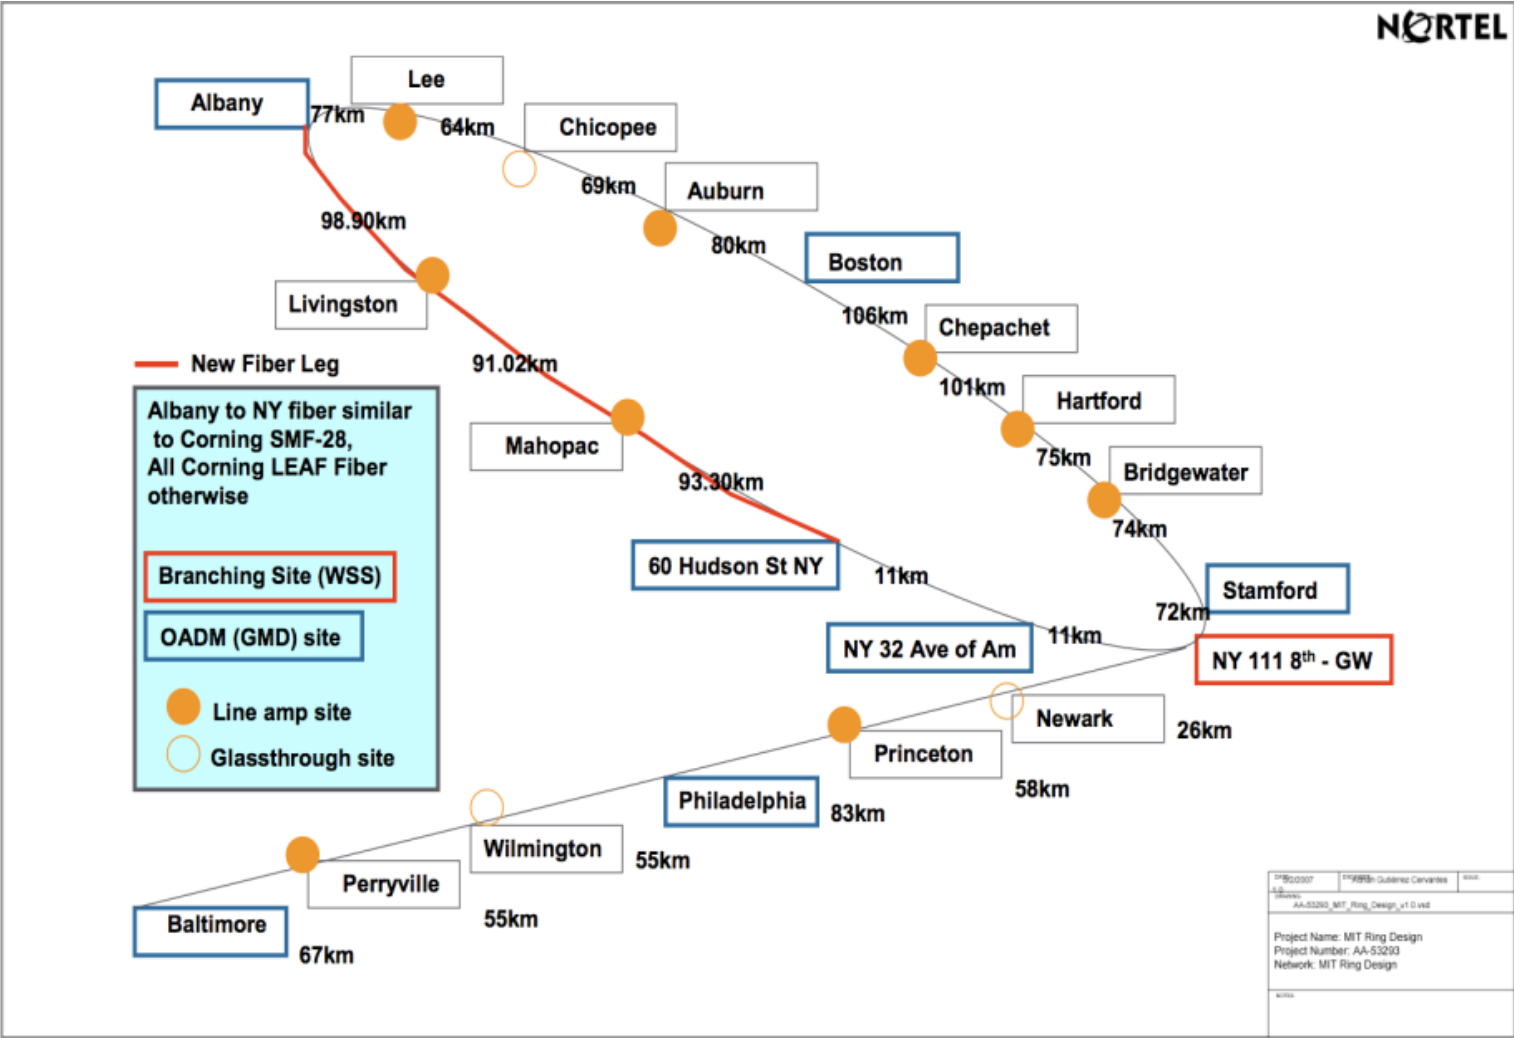
- Electricity supply
  - More than 78% of the local electricity supply is from renewable resources –primarily hydropower
    - (several nearby wind farms are permitted)
- Facility Design
  - Designed for LEED certification
  - Energy-efficient power distribution
  - Advanced cooling techniques that minimize electricity use
  - Green landscaping and storm water management
- Research
  - Collaborative research enabled by the MGHPCC will address fundamental questions in energy sciences, the environment, and green computing.



# Community and Economic Development


- Support for construction of a new power infrastructure in Holyoke
- Participation in the Holyoke Innovation District Task Force
- Deployment of advanced interactive textbook technology in the Holyoke school system
- “Youthforce 2020” introduction to construction engineering for area high school students
- College internship program in construction engineering
- Commonwealth Alliance for Information Technology Education –a University-Community College Partnership

# Fiber Crossroads




# WHAT DOES UMASS DO IN HPC?


- Faculty at all UMass campuses use HPC resources




Cosmology simulations,  
N. Katz, UMass/Amherst




Protein-Protein Docking  
Zhiping Weng, UMassMed




Ocean Mixing, Miles Sundermeyer  
UMass/Dartmouth



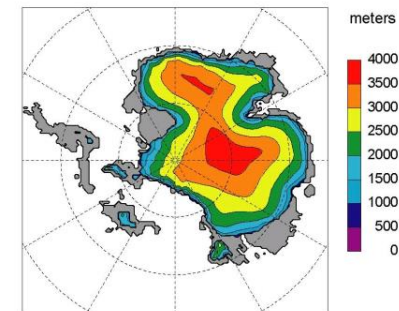
Proton Hopping in Organic Molecule  
S. Auerbach, UMassAmherst



S2-S2-S2  
Protein Unfolding, Valeri  
Barsegov, UMass/Lowell



Cooling a Single Atom in an Optical  
Cavity, Kurt Jacobs, UMass/Boston



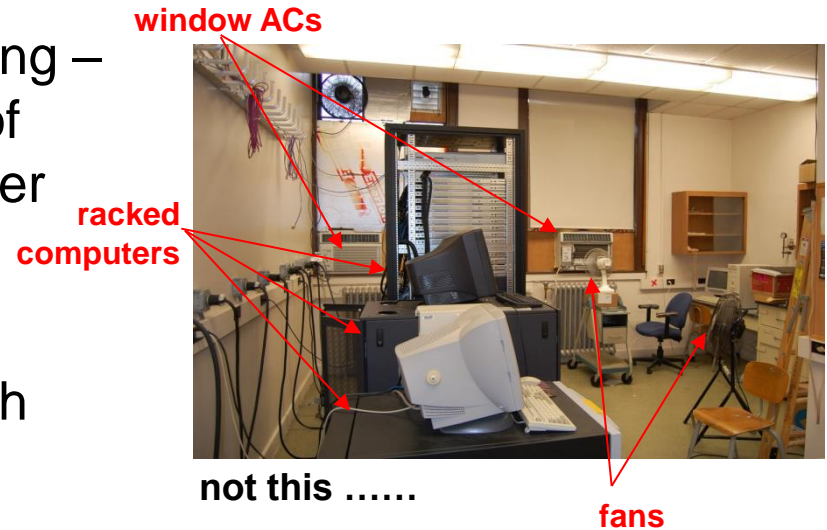
Climate change models, R. Decanto  
UMass/Amherst

# WHY SO IMPORTANT TO UMASS (ACADEMICALLY)

- HPC is no longer a “nice-to-have” but a competitive requirement for modern research universities
- Meets growing needs of our faculty for this kind of specialized infrastructure
- Helps to retain and recruit faculty
- Enables us to be more competitive for future R&D funding – both federal and industrial
- Creates new collaboration opportunities with private universities (e.g., \$130 M MIT-UMass proposal to US DOE)

# WHY IMPORTANT TO UMASS (BUSINESS-WISE)

- Cost savings – both capital and operating – thru shared infrastructure, economies of scale, centralized facility and much lower energy costs
- Frees up valuable space on campus
- Allows researchers to focus on research
- Meet “green” objectives thru facility design, operating plan, use of clean energy, and R&D in areas such as “green computing”
- Potential for carrying out our regional development mission



.... this

# Time Line

---

## Inception (Aug-Nov)

- Funding agreements
- Site purchase agreement
- Environmental testing
- Basis of Design
- Construction Manager Contract

## Design and Planning (Nov-Aug) + Site Remediation (Feb-Oct)

- Building and IT design
- Business model and operating agreements
- Environmental remediation and building demolition
- Site acquisition

We are Here




## Construction and Move-in Planning (Aug-Nov)

- Implement construction plan
- Develop and execute commissioning plan
- Prepare for occupancy
- Develop move-in plan


Move In (Nov+)

# Two Goals of the MGHPCC

---




Build and Operate  
a Data Center for  
Scientific Computing



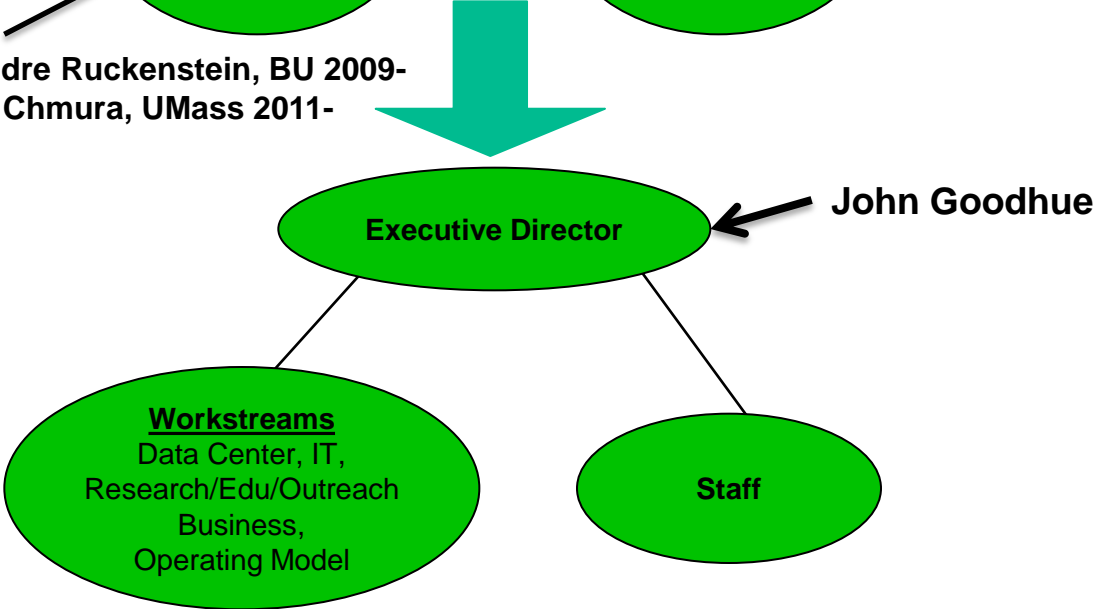
Encourage  
Novel and Collaborative  
Computationally Intensive  
Research

# MGHPCC Organization

Founding Chair: Jack M. Wilson, UMass, President 2009-2011  
Chair: Robert Brown, Boston U. , President 2011-



Chairs: Andre Ruckenstein, BU 2009-2011, Tom Chmura, UMass 2011-







## MGHPCC Resources and Services



# MGHPCC Resources

---

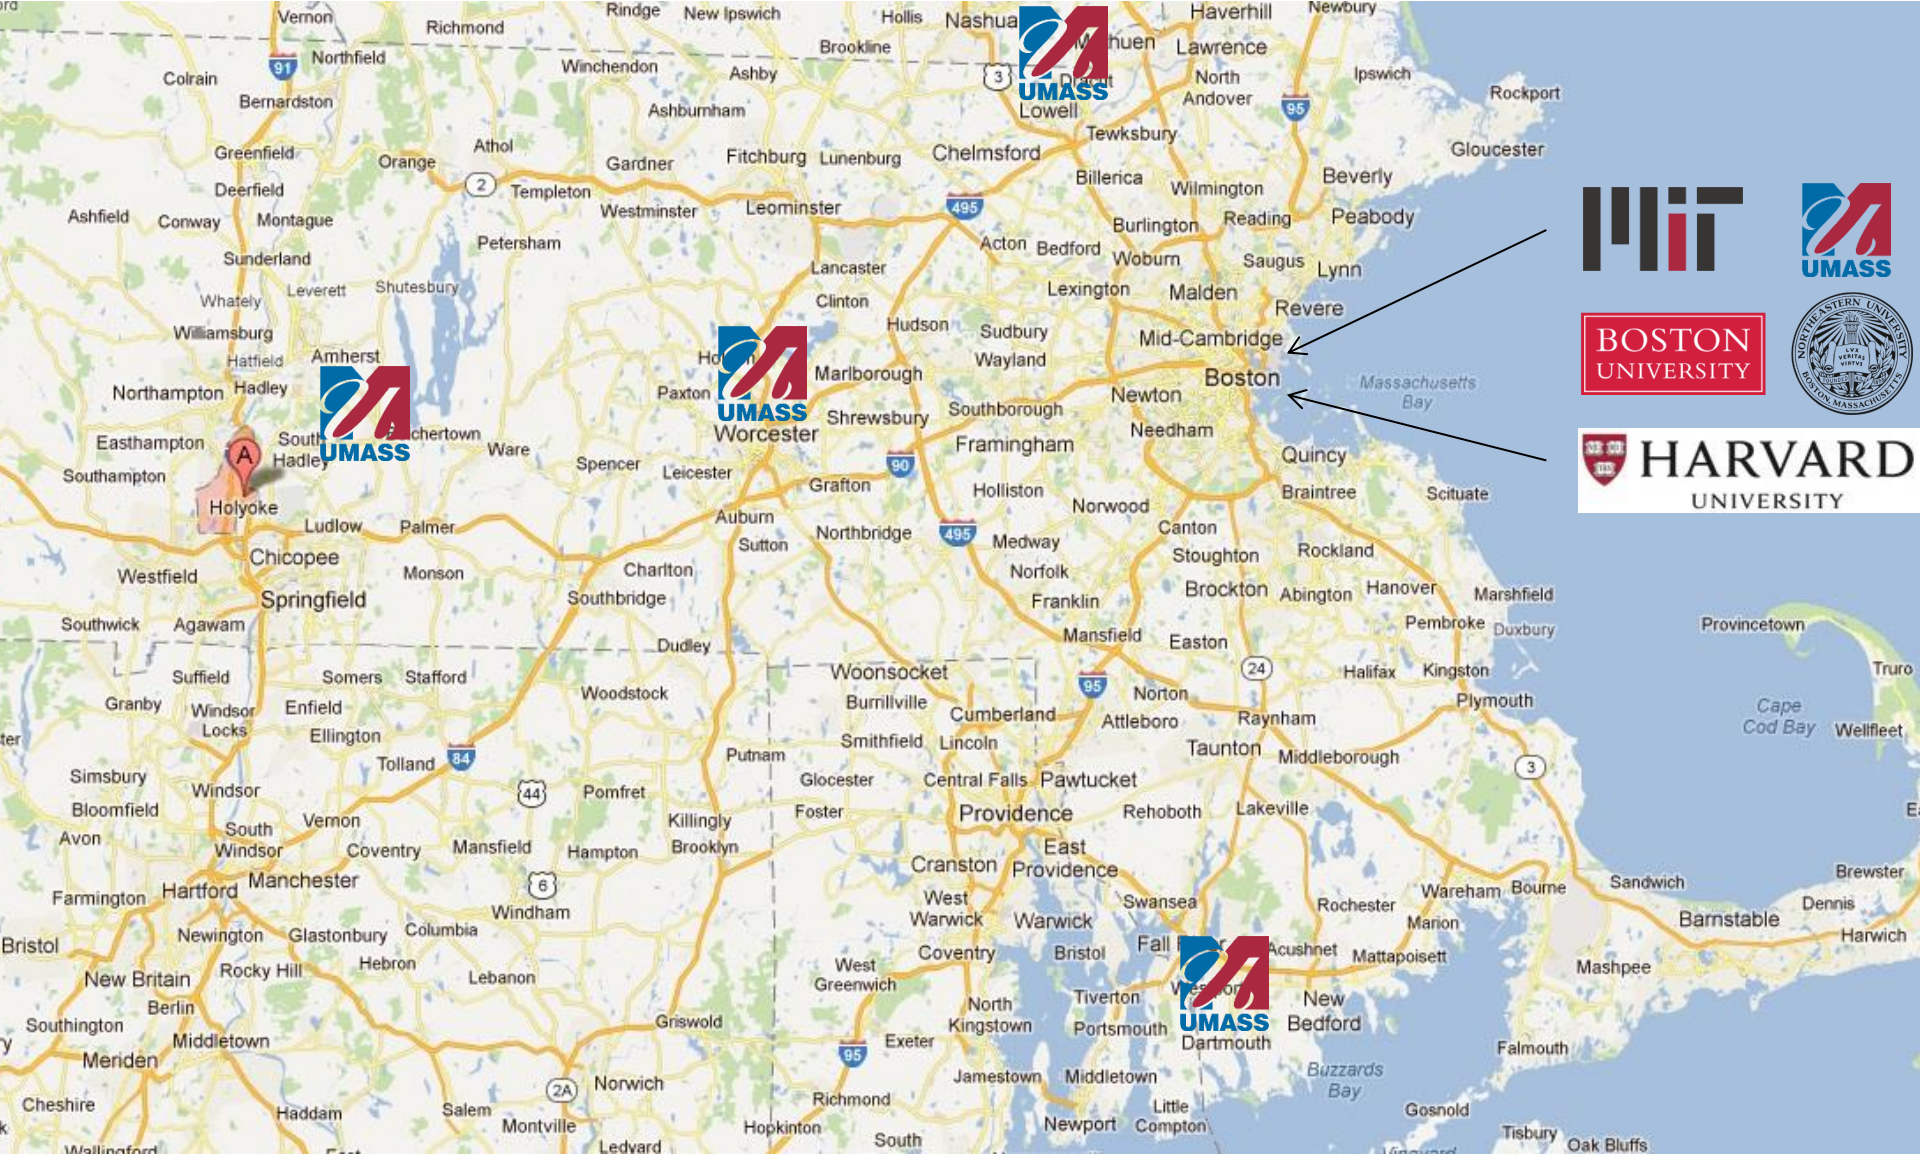
- Site
- Building
- Computer Room Floor
- Pod
- Network
- Services

# Facility Design and Construction

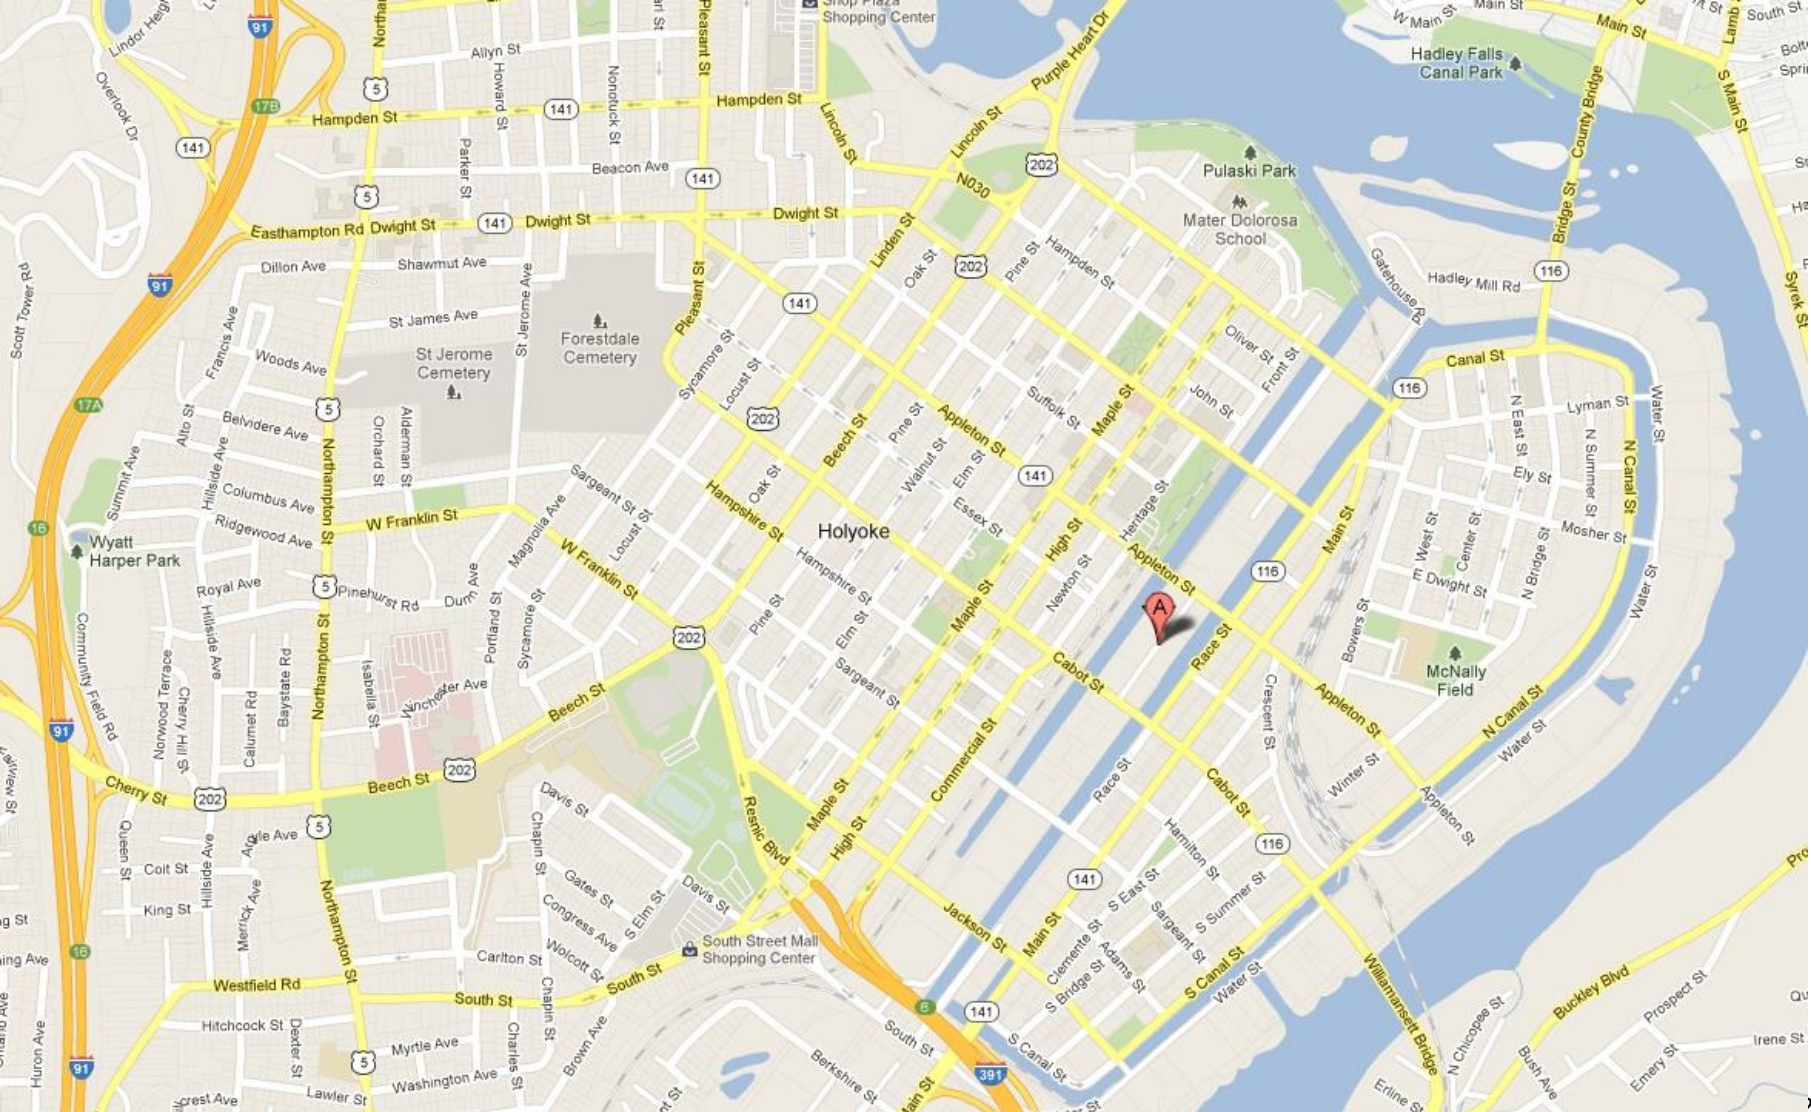
---

- Site – 8.6 Acres
- Location – Bigelow Street, Holyoke MA
- Facility size – 90,300 square feet
- Total Power Capacity: 15 Megawatts
- Power Capacity available for computing: 10 Megawatts
- Projected completion date: end of 2012
- Construction Manager – Turner Construction inc  
Architecture and Engineering – M+W Group
- Project Management – Multi-university management team, supported by Leggatt and McCall Properties

# Holyoke is an Old Mill Town in Western Mass

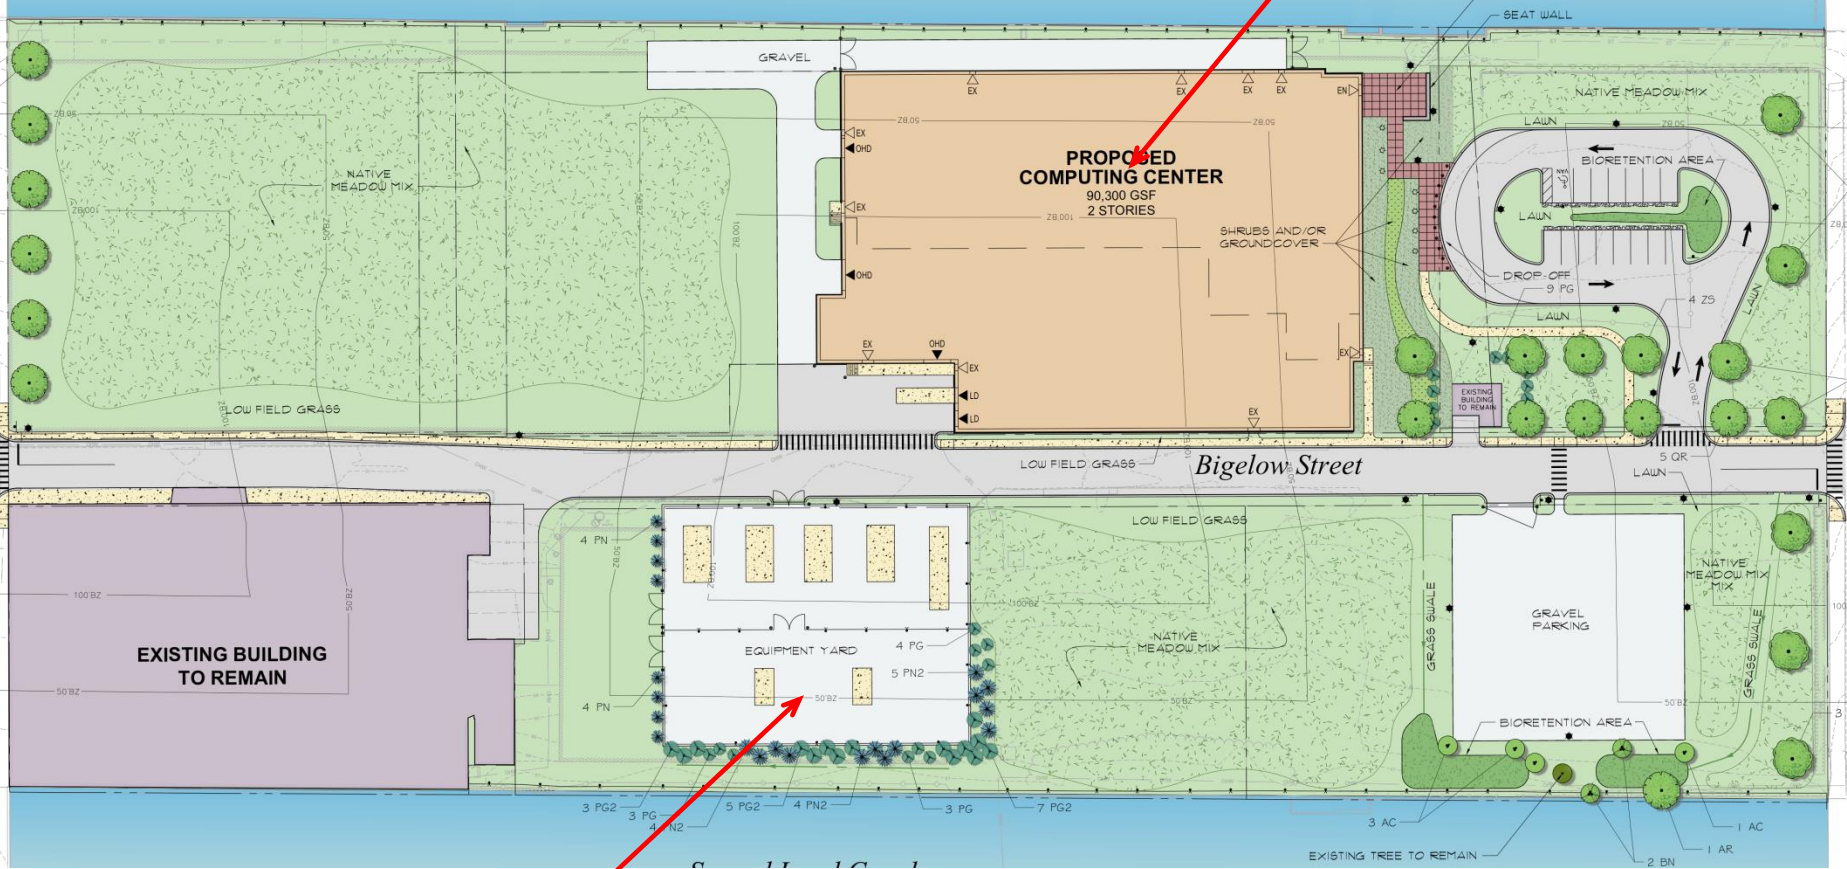


# Located on Bigelow Street Between Two Canals



# MGHPCC Site (8.6 acres total)

**Computer Center  
(90,300 Square Feet)**




**Transformer Yard  
15MW Initial capacity  
15MW expansion**

# MGHPCC Facility

**Cooling Towers**  
(~75% of cooling load)

**Computer Room**



**Shipping,  
Receiving,  
and Staging**

**Chilled water  
distribution**

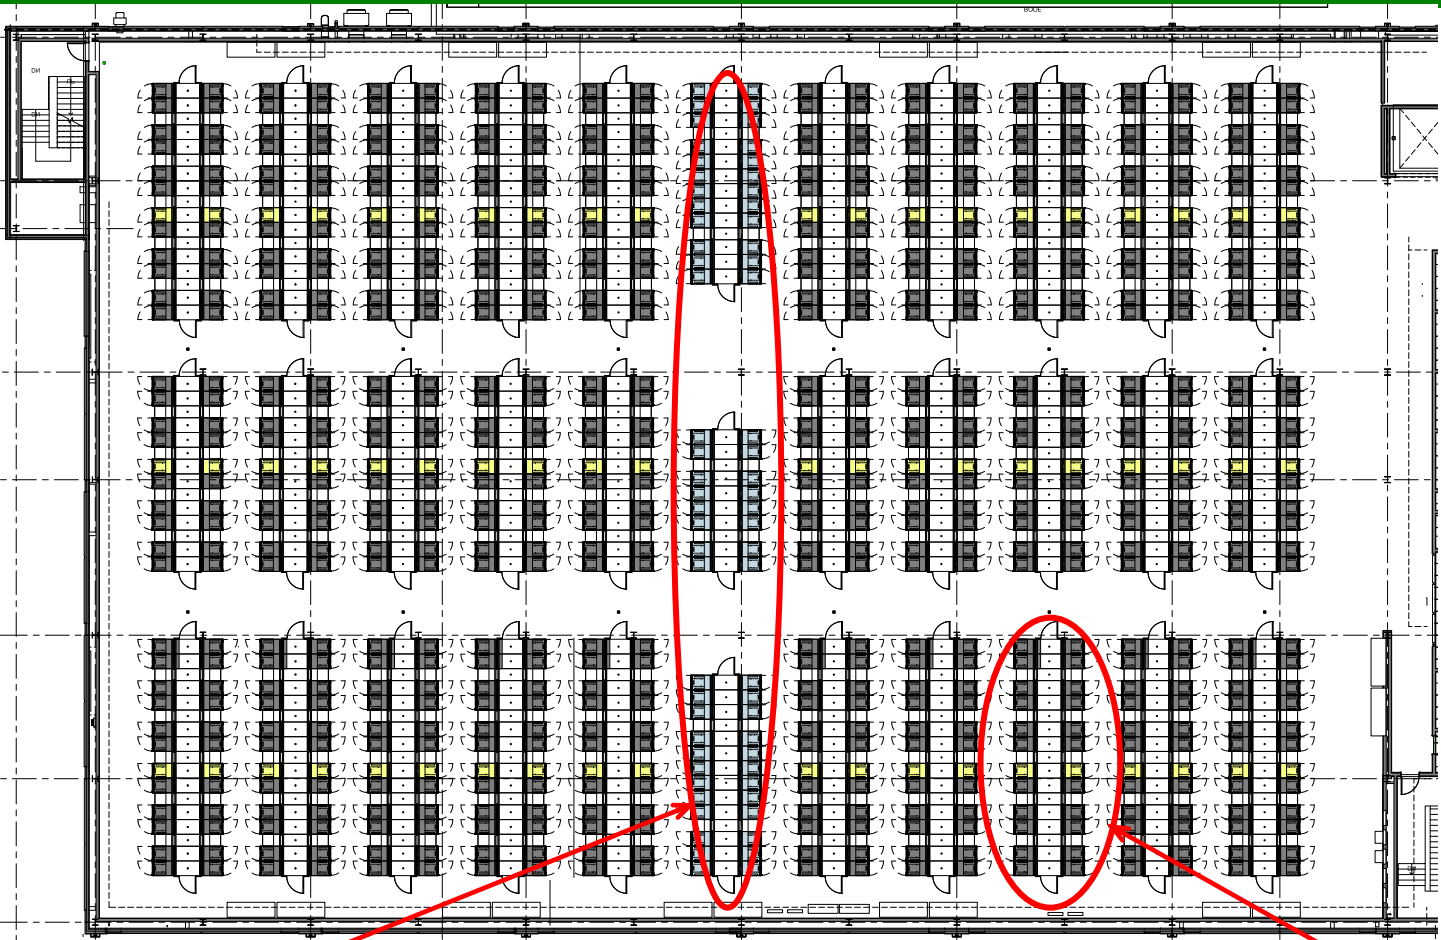
**Chillers**  
(~25% of cooling load)

**Entrance  
Rooms**

**UPS**

**Power  
Distribution**

# Computer Room



**Space**  
33,500 sq ft  
680 racks

**Power and Cooling**  
10 MW for compute  
5 MW for cooling  
Target PUE: 1.3


**Network**  
10 Gb/s feeds to  
every campus

**Main Distribution Area**  
University Fiber Feeds  
University backbone switches  
Meet me switch

**Compute Pod**  
20 or 24 Racks  
14kW per rack average  
N x 10 GigE to the MDA



# Pod Design



# Ownership and Management of Computing Resources

<u>Ownership</u>		<u>Management</u>		<u>Allocation</u>
Research Project				
PI Coalition		Owner		Reservations
Department	X	Service Group	X	Job Scheduler
Campus/School		Outside Contractor		Virtual Machines
University-wide		etc		etc
University Coalition				

The MGHPCC is prepared to support all of these usage models

# MGHPCC Services

---

- Facility Operation
  - Site maintenance and security
  - Building management, including optimization of energy efficiency
  - Documentation and support
- Incident response
  - 24x7 coverage
  - First responder
  - Notification
  - Follow-through
- User-specific services
  - Shipping and receiving,
  - Vendor visit coordination and escort
  - Virtual Hands (button press, cable move, FRU replacement, etc)
  - Site preparation (when needed)



## MGHPCC Consortium for Research and Education



# MGHPCC Consortium

- Consortium of 5 universities organized as a parallel organization to the 501c3 to focus on research and education
- Research initiatives not achievable by any single institution
  - NSF STCs, ERCs, CISE Expeditions, DOE EFRCs, innovation-hub-like
  - An opportunity industry partners will want to participate in
- Collaborative effort to strengthen MA R&D leadership
  - Most significant state/industry/university partnership in MA history
  - Focus on computationally intensive research
  - Leverage the data center as a research tool
  - Stimulate collaboration in other areas ( e.g. cybersecurity)
- Catalyst for economic & workforce development
  - In the state, region and Holyoke (Innovation District)

# Research – Collaboration and Funding

- Goal:** large-scale collaborative proposals *not* achievable by any single institution alone
- NSF STCs, ERCs, CISE Expeditions, DOE EFRCs, innovation-hub-like
  - an opportunity industry partners will *want* to participate in

## ***Collaboration mechanisms:***

- joint seminar series, CRIBB seminar
- joint courses
- Joint proposals

## ***Activities to Date:***

- \$1M Research seed fund (\$500K/year for 2 years)
- \$2.3M MRI NSF Major Research Instrumentation proposal
- \$52M NSF Track 2 HPC center proposal
- Cross-university cybersecurity initiative

# Education and Outreach

- Goal: leverage academic and industry collaboration
  - Include universities and community colleges
  - inspire, educate computationally literate young generation
  - understanding impact of computing, “smart planet”
  - energy-aware, green technologies
- Activities to date
  - 2 HPC education workshops at HCC in 2010 (~70 people each)
  - \$600 K NSF education proposal (invited to resubmit)
  - \$250 K NSF MIT/UMass interactive textbook proposal (approved)
  - Holyoke Innovation District Task Force
  - October “IdeaMill” entrepreneurs’ workshop
  - November academic/Industry energy innovation workshop  
(with Holyoke Gas and & Electric and the Commonwealth)

# 2011 research seed fund results

- More than \$600K awarded to 7 projects
- Award criteria:
  - Cross-university team, addressing problems
  - Opens new avenues of research
  - Likely to lead to a larger follow-on project
- Projects spanned three key areas of research computing
  - Computers as a tool for scientific discovery
  - Application software to support new areas of scientific research
  - Laying foundations for the next 100X increase in computing capacity



# Research seed fund – example projects

- New algorithms and software techniques for exascale computing systems ( $10^{18}$  ops/second)
- Atomic-scale modeling to predict the properties of new materials combining metals, dielectrics, and magnetics
- Modeling the nutrient cycle near the edge of the continental shelf, supporting the new ocean observation platform off the Massachusetts coast
- Measure and improve the energy efficiency of large scale computer systems

# MGHPCC –Seed Grants 2012

- Lorena Barba of BU, Chris Cecka of Harvard and Hans Johnston of UMass Amherst
  - will work on software techniques for future generation “exascale” software platforms, which will be 1,000 times faster than current computing speeds of 1,000 trillion operations per second.
- Alfredo Alexander-Katz of MIT and Alan Aspuru-Guzik of Harvard
  - will develop high-performance computer code to simulate the kind of energy processing that occurs during photosynthesis in plants. Ultimately, the insights gained from such simulations will be transferred to the engineering domain to increase the energy-efficiency of such devices as solar panels.
- John Marshall and Pierre Lermusiaux of MIT, Amala Mahadevan of Woods Hole Oceanographic Institute and Amit Tandon of UMass Dartmouth
  - will create models to provide insights into the turbulent mixing that regulates nutrient cycle and ocean ecosystem dynamics off the New England coast. The project will be done in conjunction with the forthcoming deployment of a state-of-the-art underwater observation platform as part of the NSF-sponsored Ocean Observatories initiative.
- Hossein Mosallaei and David Kaeli of Northeastern and Efthimos Kaxiras of Harvard
  - will create computer models that simulate the behavior of metal, dielectric, and magnetic particles at extremely small scales, allowing insights into the behavior of important new materials.
- Patricia Ellen Grant of Harvard and Children’s Hospital and Jonathan Appavoo of BU
  - will use high-performance computing to automate medical imaging analysis in a way that will make it easier and less costly to use by radiological clinicians.
- Yanlei Diao and Li-Jun Ma of UMass Amherst, Samuel Madden of MIT, Bai-Lin Wu of Harvard and Children’s Hospital, Toby Bloom of the Broad Institute of Harvard and MIT, and James F. Gusella of Massachusetts General Hospital
  - will develop next-generation, on-demand services for managing and processing massive amounts of genome information.
- Ayse K. Coskun and Martin C. Herboldt of BU and Gunar Schirner of Northeastern
  - will explore ways to measure and improve the energy efficiency of large-scale-computing. Energy issues have become a barrier to progress on computer speed.



# Questions?

