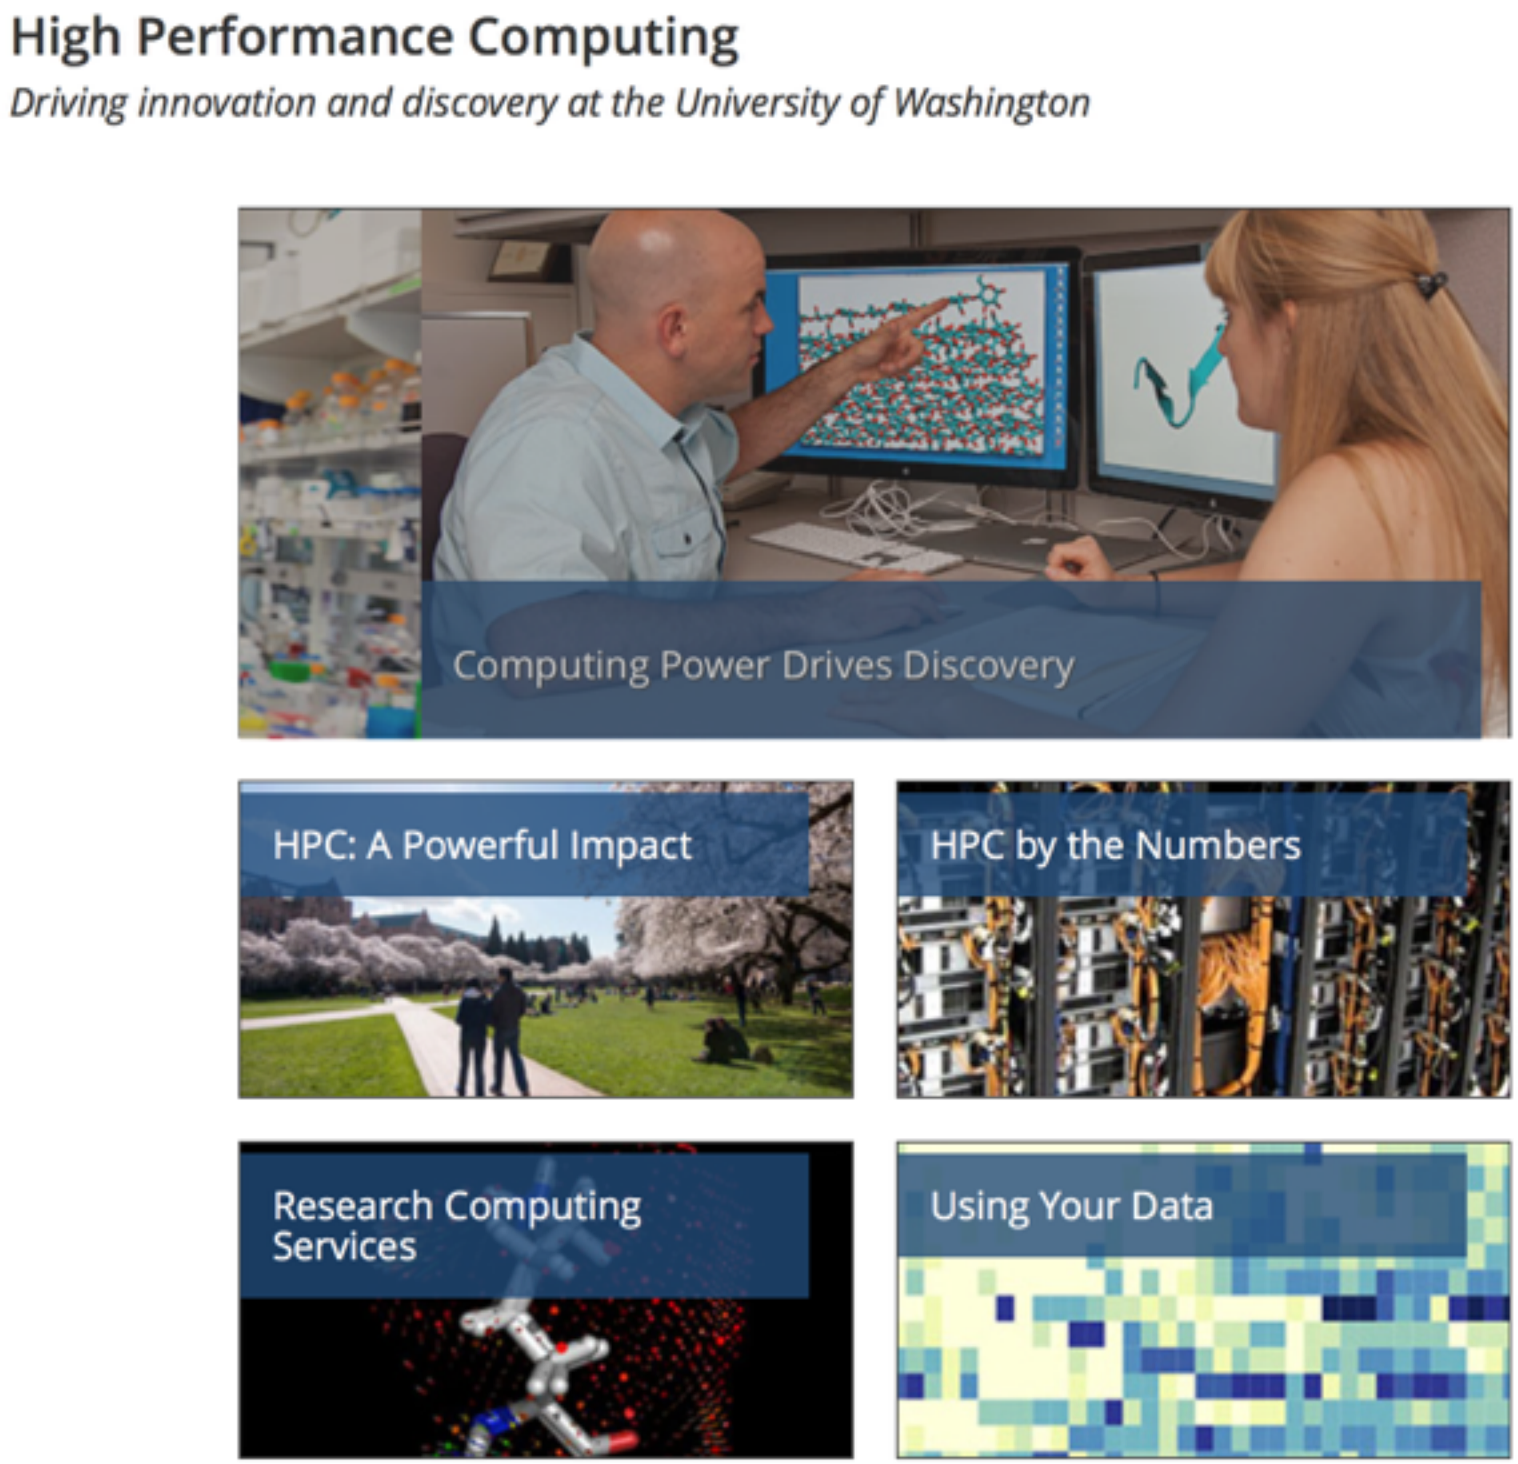

High Performance Computing

Driving innovation and discovery at the University of Washington

Scientific Computing Needs at UW

Exploratory Workshop
Role of High-Performance Computing in the Pacific Northwest
27 July 2015

Martin J Savage

INSTITUTE for NUCLEAR THEORY

Scientific Computing Initiates eScience @ UW 2007 : Astro/INT/Physics

http://staff.washington.edu/nfabbi/pdf/PerspAutumn08_final.pdf

Athena Unleashed

Athena doesn't look like much. Just a bunch of black computing towers, punctuated by cooling units, in a nondescript room. But looks can be deceiving.

Named for the Greek goddess of wisdom, Athena is currently the most powerful computer on the UW campus, helping physicists and astronomers tackle fundamental questions about our universe.

The push to acquire Athena was spearheaded by David Kaplan, director of the Institute for Nuclear Theory (INT), Tom Quinn, professor of astronomy, and Richard Coffey, director of IT for physics and astronomy. They were responding to what can be a frustrating Catch-22 for scientists whose research involves complex calculations—a field known as computational science. Scientists can apply for time on a huge computer at a national lab, but they will likely be turned down if they cannot demonstrate expertise at using such machines efficiently.

The team envisioned a UW supercomputer—not nearly as powerful as the machines at national labs, but about 1,000 times more powerful than an individual computer workstation—that could push existing research projects to new extremes and familiarize scientists with working with computers on a grand scale.

After several failed attempts to fund Athena through grants, Kaplan turned to the UW's Office of Research, which was looking for ways to support e-Science (science that links many computers). The Office of Research joined with the College of Arts and Sciences to provide $700,000 for the project.

"They gave a lot of money up front, believing this would work," says Kaplan. "That took vision."

The remaining funding came from departments willing to invest in the project. In the end, three A&S units ponied up: the Department of Astronomy, the Department of Physics, and the INT. The computer is housed in the UW's Center for Experimental Nuclear Physics and Astrophysics.

Brain Power to Harness Computing Power

Athena's power equals 133 high-end PC servers working in close communication. It calculates at nearly ten teraflops (Tflops), or about ten trillion calculations per second. But how that power and speed are harnessed, and how those PCs communicate, is complicated. So complicated, in fact, that the team would not consider purchasing Athena without hiring a computational science expert to program it.

"Scientists may know how to program their desktop, but programming a supercomputer requires a whole other set of skills because it is doing many tasks simultaneously," says Kaplan. "Careful choreography is required so that individual computations can be done in the right order and assembled into something useful, in a way that takes advantage of the speed of the machine."

That's where Jeff Gardner comes in. Gardner received a Ph.D. in astronomy at the UW and then spent five years at the Pittsburgh Supercomputing Center (PSC), programming one of those massive national computers, before returning to the UW as a senior research scientist to help with Athena.

"A truly amazing part of this collaboration was the hiring of Jeff," says Coffey, who led the national search for the position. "Three departments pooled their limited resources to hire an expert in the field, filling this often overlooked gap between the science and the computing."

Gardner is currently working with faculty on about a dozen research projects that use Athena. Some faculty meet with him sporadically, others weekly. All pay for his time through research grants.

"Every project, every scientific code is different, so it has to be parallelized in a different way," says Gardner. "You need a set of tools and experience to figure out how to go about it."

Gardner also assists in writing grant proposals, "because that's where plans formulate," he explains. "What we don't want is for faculty to propose something that our technology can't do."

Computing Complex Interactions

What Athena can do is impressive. Research projects range from the grandest scale—studying the universe—to the smallest, looking at atomic interactions. What all have in common is the complexity of interactions being studied.

One example is Tom Quinn's study of structure formation in the universe. Quinn looks at the creation of our galaxy and neighboring galaxies. He does this, in part, by gathering measurements of remote objects—dating back to when the universe was about 100,000 years old—and comparing them to the galaxies we see today, factoring in the role of gravity.

"That sort of calculation can't be done on the back of an envelope," says Quinn, massively understating the computational challenge. "With pencil and paper, you can figure out how three objects interact with gravity. But the universe has billions of objects in each galaxy."

Clearly Quinn needs tremendous computing power to handle such calculations. But just having multiple processors do the math simultaneously won't work. "The issue is how to get all those processors to work together," says Quinn. "It's not a problem I can easily divide up, because the calculations are all very interconnected."

Much of Quinn's research requires using a massive computer at a national center. But before he can tap into that resource, he needs to devise algorithms—with Gardner's help—that will work on a multi-processor system.

"I need something I can test on," says Quinn. With Athena, he is able to try different algorithms and compare results. He has the luxury of time for testing because his department has part ownership in the computer. And if a result leads to additional questions, he can pursue those immediately.

"It happens fluidly," says Quinn. "National centers aren't set up to do that sort of thing."

More Power, More Grants

When Athena was proposed, the units funding the project believed it would eventually pay for itself by helping to generate new grants. Within months of its arrival, that already started to happen. Several major NSF grants tied to Athena have been funded, with others pending.

Some challenges remain. While programming and maintaining the system have been manageable, the administration involved in sharing the computer and related personnel across departments has been daunting. Yet all agree that the benefits of Athena overshadow any administrative headaches.

"In 2001, the fastest computer on the planet for unclassified research was six teraflops," says Gardner, "and thousands of scientists across the country had to compete with one another for time on it. Athena is ten teraflops and is shared among just three departments.

"It is truly remarkable that we have access to so much computing right on campus." ◆

> "In 2001, the **fastest** computer on the planet for unclassified research was six teraflops.... **Athena** is ten teraflops and is **shared** among just three departments."

David Kaplan (left) and Richard Coffey stand in the "hot aisle" of Athena's two long rows of computer hardware. *Photo by Mary Levin.*

Jeff Gardner (left) and Tom Quinn discuss approaches for using Athena for Quinn's research. *Photo by Mary Levin.*

Athena By the Numbers

10 trillion calculations. In one second, Athena can compute 10 trillion calculations, compared to 10 billion on the average PC.

9 months. From conception to deployment, it took nine months to get the Athena cluster in place.

One quarter. That's the fraction of the purchase price used to cool and provide power to Athena.

1024 cores. Most modern computers have 2 cores; Athena has 1,024.

20.7 billion pages. That's the amount of data Athena is able to store.

40 hairdryers. Athena's heat output is equivalent to 40 hairdryers, all blowing at once.

15 minutes. Without an integrated cooling system, the room housing Athena would overheat in just 15 minutes. ◆

12

A&S PERSPECTIVES | AUTUMN 2008 13

We learned a lot about shared HPC resource infrastructure

eScience Institute @ UW Launched in Nov 2008 to enhance the Domain Sciences

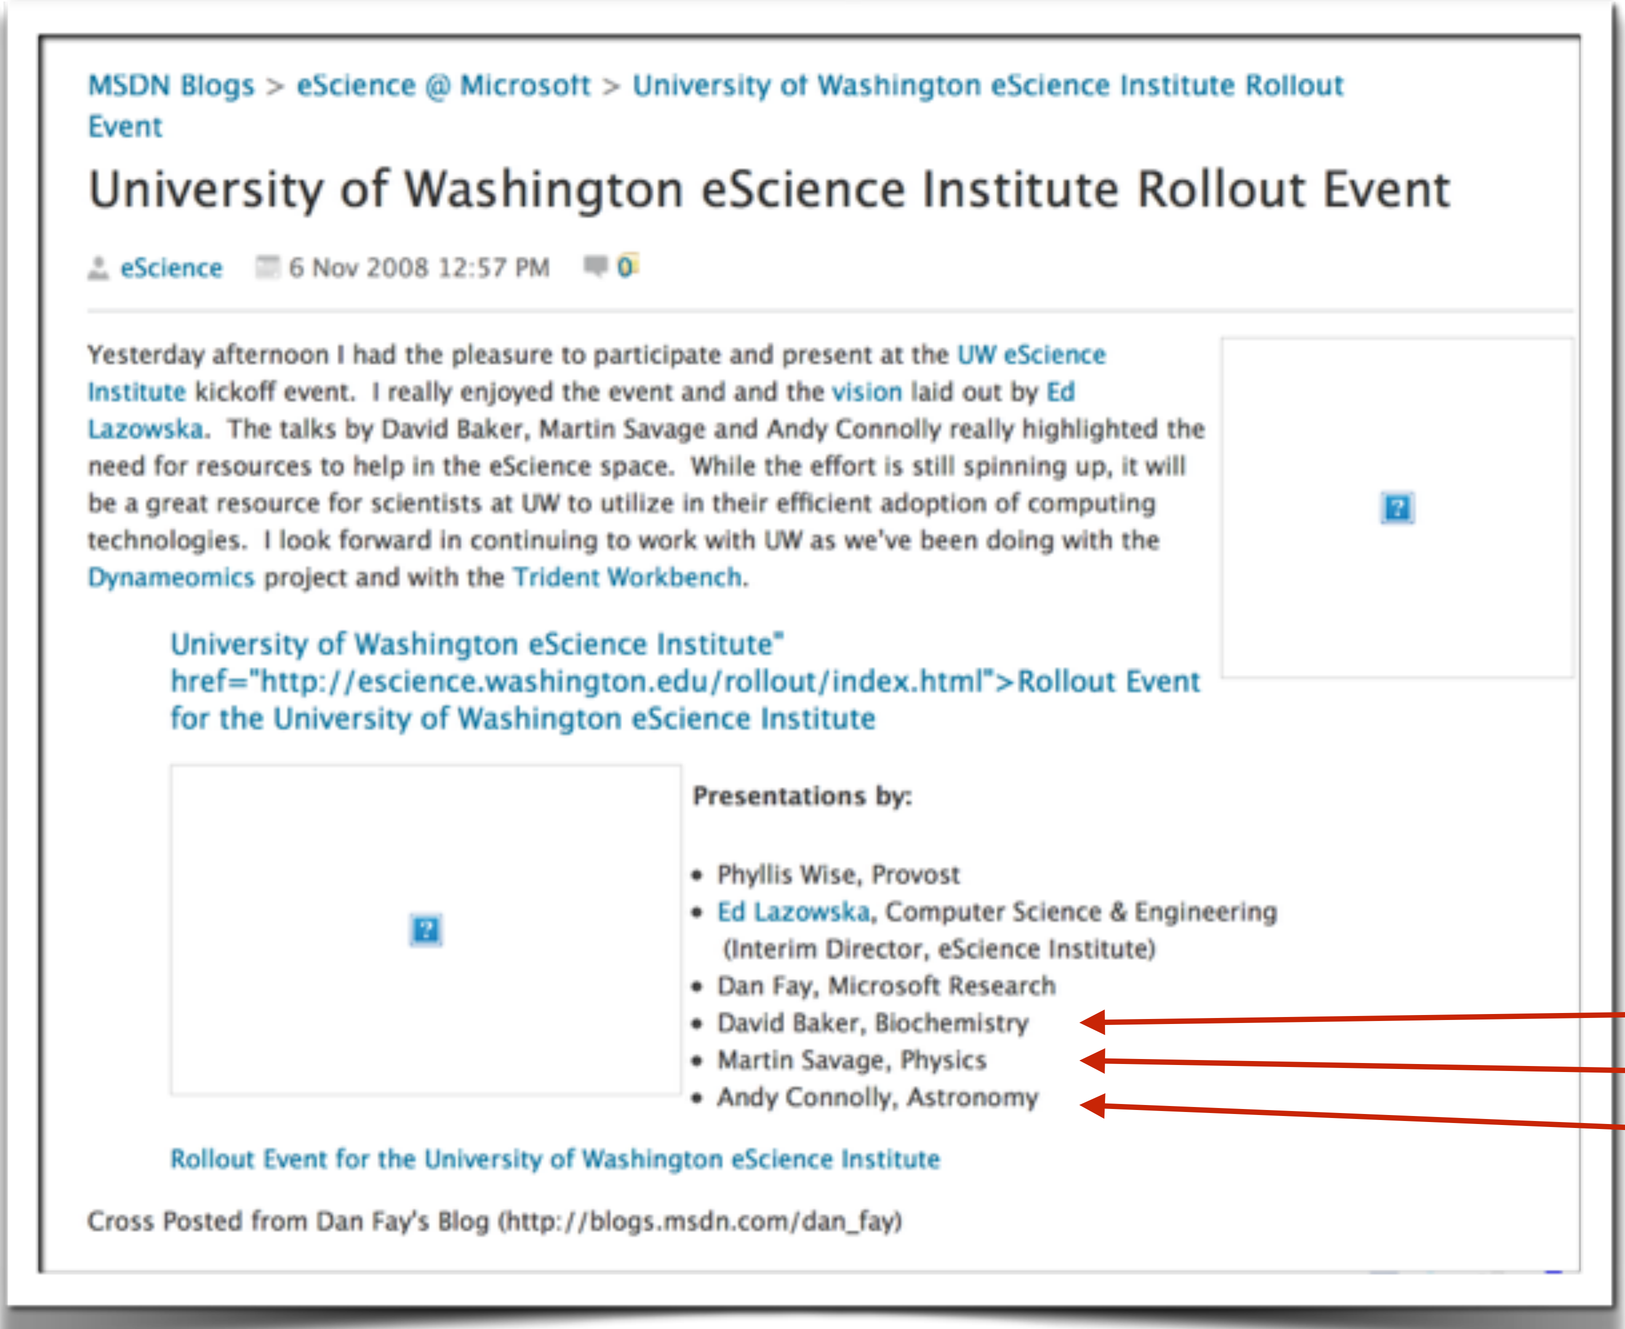

MSDN Blogs > eScience @ Microsoft > University of Washington eScience Institute Rollout Event

University of Washington eScience Institute Rollout Event

eScience 6 Nov 2008 12:57 PM 0

Yesterday afternoon I had the pleasure to participate and present at the UW eScience Institute kickoff event. I really enjoyed the event and and the vision laid out by Ed Lazowska. The talks by David Baker, Martin Savage and Andy Connolly really highlighted the need for resources to help in the eScience space. While the effort is still spinning up, it will be a great resource for scientists at UW to utilize in their efficient adoption of computing technologies. I look forward in continuing to work with UW as we've been doing with the Dynameomics project and with the Trident Workbench.

University of Washington eScience Institute" href="http://escience.washington.edu/rollout/index.html">Rollout Event for the University of Washington eScience Institute

Presentations by:

- Phyllis Wise, Provost
- Ed Lazowska, Computer Science & Engineering (Interim Director, eScience Institute)
- Dan Fay, Microsoft Research
- David Baker, Biochemistry
- Martin Savage, Physics
- Andy Connolly, Astronomy

Rollout Event for the University of Washington eScience Institute

Cross Posted from Dan Fay's Blog (http://blogs.msdn.com/dan_fay)

(All other Web sites/images/videos have been removed -including the video of the entire event)

Simulation
Calculation and Simulation
Simulation and Data

Lessons I Learned:

- Domain Scientists (DS) and CS have different agendas
- DS simply want to optimize scientific output - requires large HPC resources
- CS (and AM) need DS to justify/acquire large resources - little intrinsic need of their own

HYAK PIs

(=HYAK funder)

http://www.washington.edu/itconnect/research/hpc/

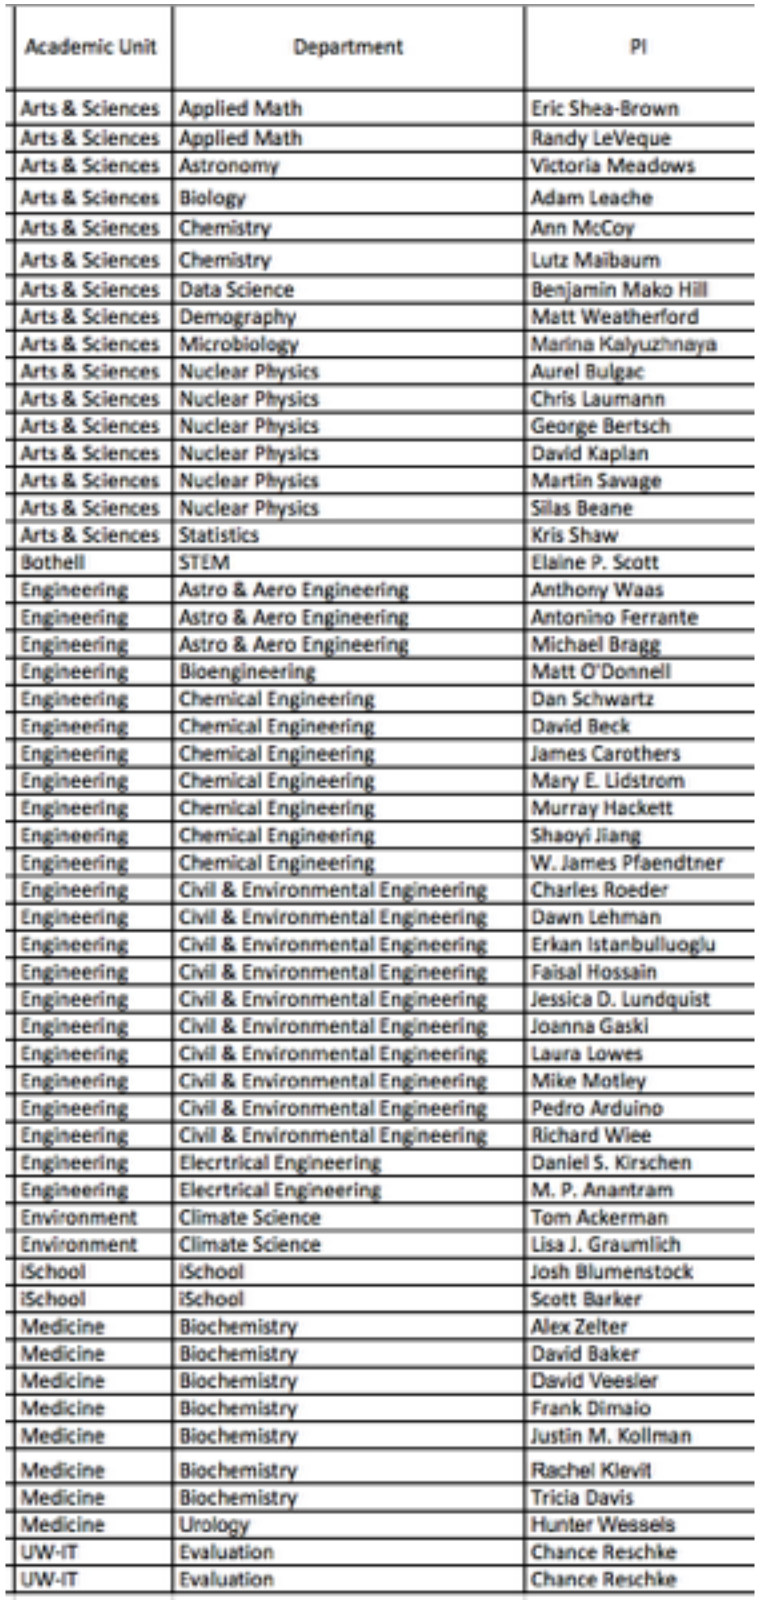

Academic Unit	Department	PI
Arts & Sciences	Applied Math	Eric Shea-Brown
Arts & Sciences	Applied Math	Randy LeVeque
Arts & Sciences	Astronomy	Victoria Meadows
Arts & Sciences	Biology	Adam Leache
Arts & Sciences	Chemistry	Ann McCoy
Arts & Sciences	Chemistry	Lutz Maibaum
Arts & Sciences	Data Science	Benjamin Mako Hill
Arts & Sciences	Demography	Matt Weatherford
Arts & Sciences	Microbiology	Marina Kalyuzhnaya
Arts & Sciences	Nuclear Physics	Aurel Bulgac
Arts & Sciences	Nuclear Physics	Chris Laumann
Arts & Sciences	Nuclear Physics	George Bertsch
Arts & Sciences	Nuclear Physics	David Kaplan
Arts & Sciences	Nuclear Physics	Martin Savage
Arts & Sciences	Nuclear Physics	Silas Beane
Arts & Sciences	Statistics	Kris Shaw
Bothell	STEM	Elaine P. Scott
Engineering	Astro & Aero Engineering	Anthony Waas
Engineering	Astro & Aero Engineering	Antonino Ferrante
Engineering	Astro & Aero Engineering	Michael Bragg
Engineering	Bioengineering	Matt O'Donnell
Engineering	Chemical Engineering	Dan Schwartz
Engineering	Chemical Engineering	David Beck
Engineering	Chemical Engineering	James Carothers
Engineering	Chemical Engineering	Mary E. Lidstrom
Engineering	Chemical Engineering	Murray Hackett
Engineering	Chemical Engineering	Shaoyi Jiang
Engineering	Chemical Engineering	W. James Pfaendtner
Engineering	Civil & Environmental Engineering	Charles Roeder
Engineering	Civil & Environmental Engineering	Dawn Lehman
Engineering	Civil & Environmental Engineering	Erkan Istanbulluoglu
Engineering	Civil & Environmental Engineering	Faisal Hossain
Engineering	Civil & Environmental Engineering	Jessica D. Lundquist
Engineering	Civil & Environmental Engineering	Joanna Gaski
Engineering	Civil & Environmental Engineering	Laura Lowes
Engineering	Civil & Environmental Engineering	Mike Motley
Engineering	Civil & Environmental Engineering	Pedro Arduino
Engineering	Civil & Environmental Engineering	Richard Wiee
Engineering	Elecrtrical Engineering	Daniel S. Kirschen
Engineering	Elecrtrical Engineering	M. P. Anantram
Environment	Climate Science	Tom Ackerman
Environment	Climate Science	Lisa J. Graumlich
iSchool	iSchool	Josh Blumenstock
iSchool	iSchool	Scott Barker
Medicine	Biochemistry	Alex Zelter
Medicine	Biochemistry	David Baker
Medicine	Biochemistry	David Veesler
Medicine	Biochemistry	Frank Dimaio
Medicine	Biochemistry	Justin M. Kollman
Medicine	Biochemistry	Rachel Klevit
Medicine	Biochemistry	Tricia Davis
Medicine	Urology	Hunter Wessels
UW-IT	Evaluation	Chance Reschke
UW-IT	Evaluation	Chance Reschke

Broad representation in sciences, engineering, bio

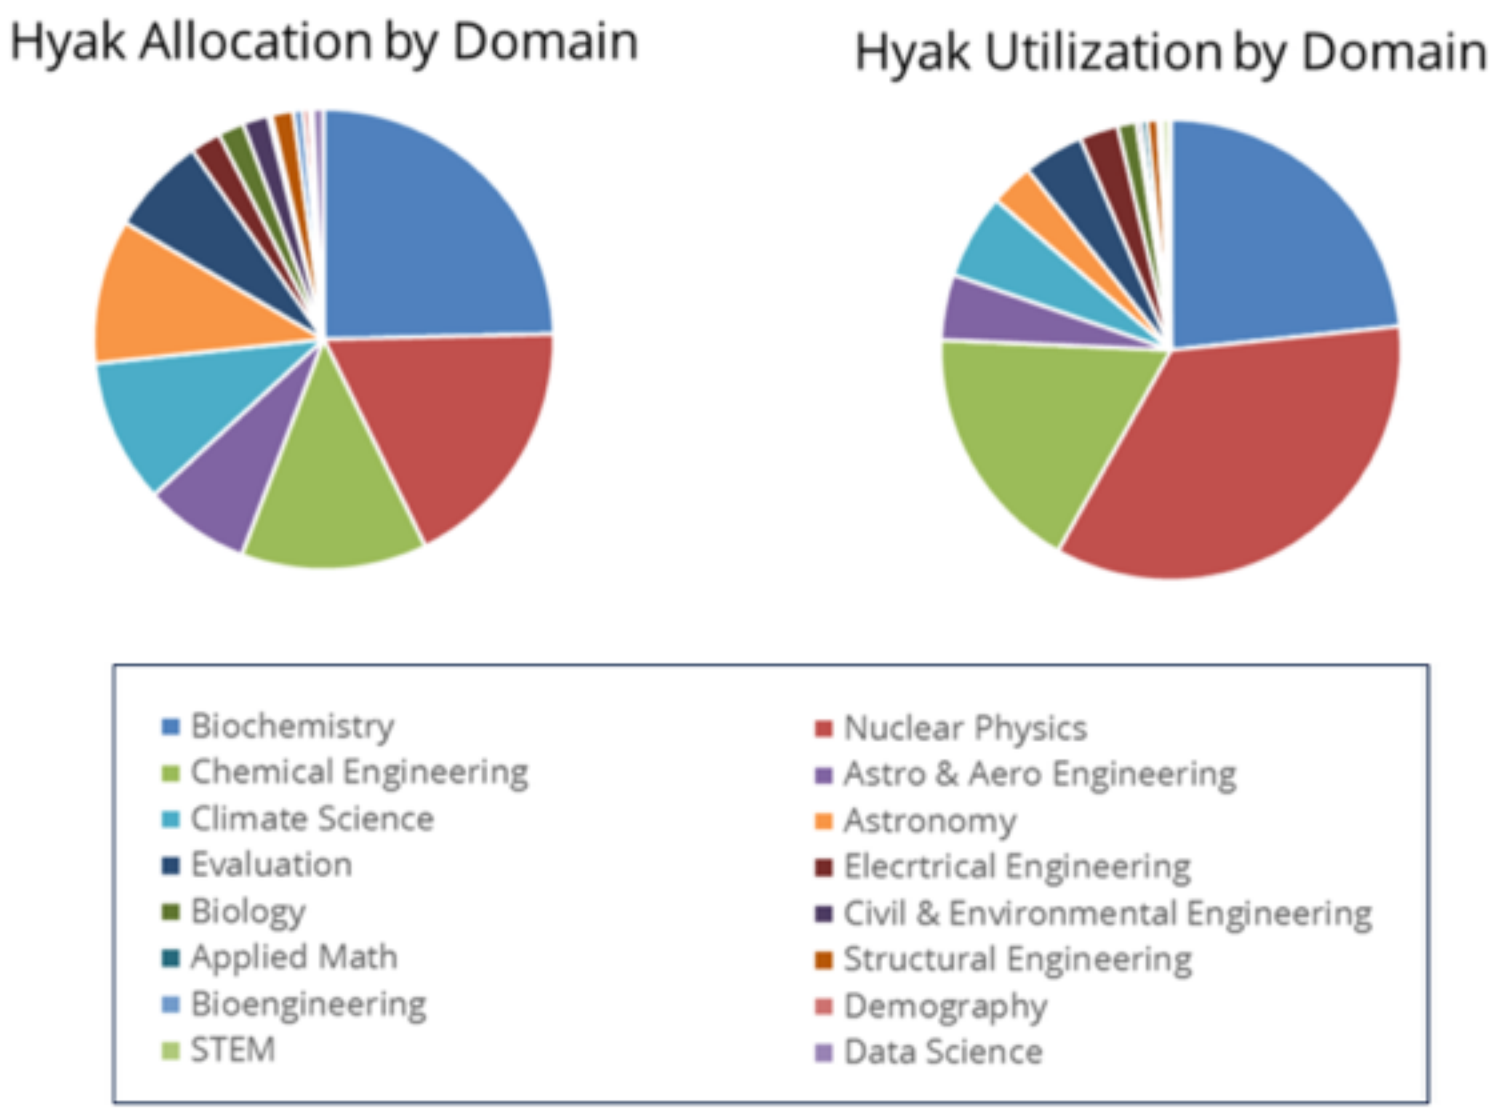

ist modified: July 1, 2015

Observe the ``usual'' ~Log-Normal distribution of use seen at other centers - long tail.

What is Needed Globally and Locally ?

- Needs are Exa-scale and beyond in essentially all areas
  - Some are Capability
  - Some are Capacity
  - Some are both

Capability Areas

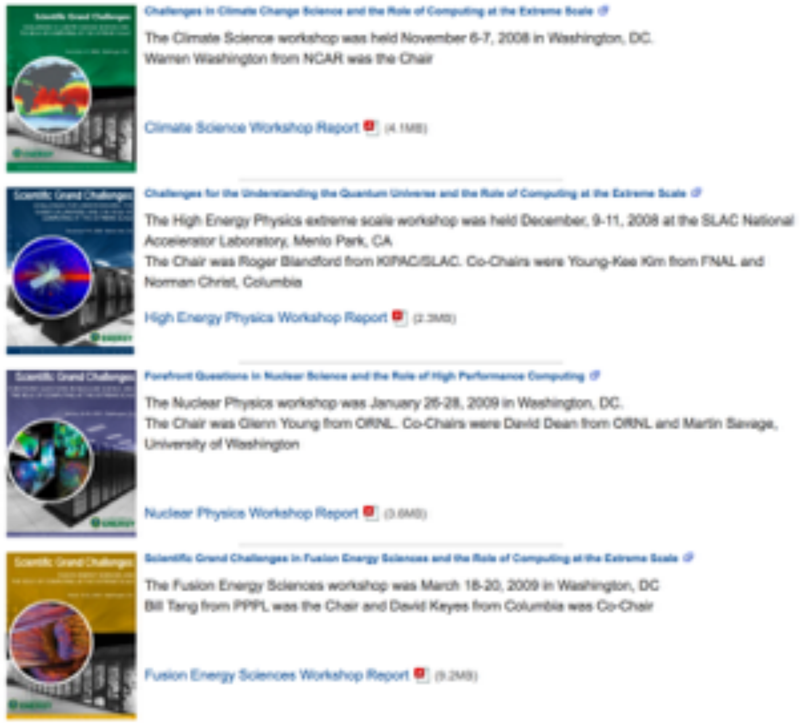

Challenges in Climate Change Science and the Role of Computing at the Extreme Scale

The Climate Science workshop was held November 6-7, 2008 in Washington, DC.
Warren Washington from NCAR was the Chair

Climate Science Workshop Report (4.1MB)

Challenges for the Understanding the Quantum Universe and the Role of Computing at the Extreme Scale

The High Energy Physics extreme scale workshop was held December, 9-11, 2008 at the SLAC National Accelerator Laboratory, Menlo Park, CA
The Chair was Roger Blandford from KIPAC/SLAC. Co-Chairs were Young-Kee Kim from FNAL and Norman Christ, Columbia

High Energy Physics Workshop Report (2.3MB)

Forefront Questions in Nuclear Science and the Role of High Performance Computing

The Nuclear Physics workshop was January 26-28, 2009 in Washington, DC.
The Chair was Glenn Young from ORNL. Co-Chairs were David Dean from ORNL and Martin Savage, University of Washington

Nuclear Physics Workshop Report (3.6MB)

Scientific Grand Challenges in Fusion Energy Sciences and the Role of Computing at the Extreme Scale

The Fusion Energy Sciences workshop was March 18-20, 2009 in Washington, DC
Bill Tang from PPPL was the Chair and David Keyes from Columbia was Co-Chair

Fusion Energy Sciences Workshop Report (9.2MB)

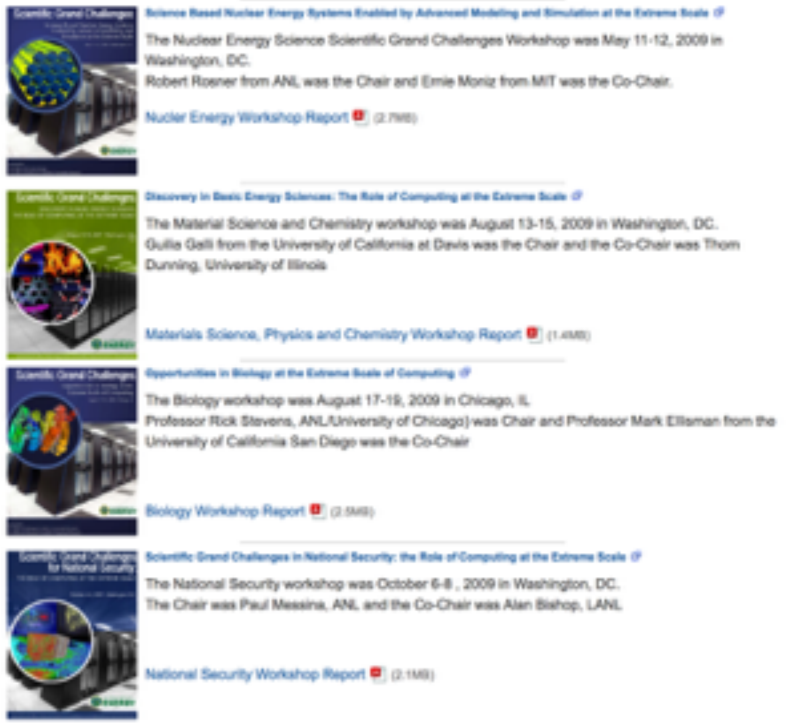

Science Based Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale

The Nuclear Energy Science Scientific Grand Challenges Workshop was May 11-12, 2009 in Washington, DC.
Robert Rosner from ANL was the Chair and Ernie Moniz from MIT was the Co-Chair.

Nucler Energy Workshop Report (2.7MB)

Discovery in Basic Energy Sciences: The Role of Computing at the Extreme Scale

The Material Science and Chemistry workshop was August 13-15, 2009 in Washington, DC.
Gulia Galli from the University of California at Davis was the Chair and the Co-Chair was Thom Dunning, University of Illinois

Materials Science, Physics and Chemistry Workshop Report (1.4MB)

Opportunities in Biology at the Extreme Scale of Computing

The Biology workshop was August 17-19, 2009 in Chicago, IL
Professor Rick Stevens, ANL/University of Chicago) was Chair and Professor Mark Ellisman from the University of California San Diego was the Co-Chair

Biology Workshop Report (2.5MB)

Scientific Grand Challenges in National Security: the Role of Computing at the Extreme Scale

The National Security workshop was October 6-8 , 2009 in Washington, DC.
The Chair was Paul Messina, ANL and the Co-Chair was Alan Bishop, LANL

National Security Workshop Report (2.1MB)

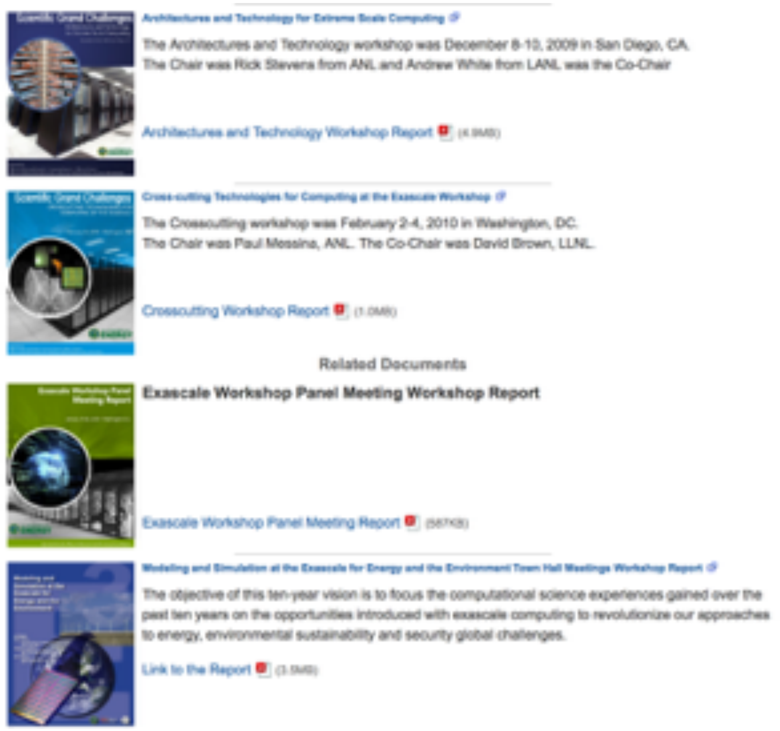

Architectures and Technology for Extreme Scale Computing

The Architectures and Technology workshop was December 8-10, 2009 in San Diego, CA.
The Chair was Rick Stevens from ANL and Andrew White from LANL was the Co-Chair

Architectures and Technology Workshop Report (4.9MB)

Cross-cutting Technologies for Computing at the Exascale Workshop

The Crosscutting workshop was February 2-4, 2010 in Washington, DC.
The Chair was Paul Messina, ANL. The Co-Chair was David Brown, LLNL.

Crosscutting Workshop Report (1.0MB)

Related Documents

Exascale Workshop Panel Meeting Workshop Report

Exascale Workshop Panel Meeting Report (587KB)

Modeling and Simulation at the Exascale for Energy and the Environment Town Hall Meetings Workshop Report

The objective of this ten-year vision is to focus the computational science experiences gained over the past ten years on the opportunities introduced with exascale computing to revolutionize our approaches to energy, environmental sustainability and security global challenges.

Link to the Report (3.5MB)

Capacity areas are a superset of capability areas,including what is called ``big-data'' (in reality: generally small data)

What is Needed Globally and Locally ?

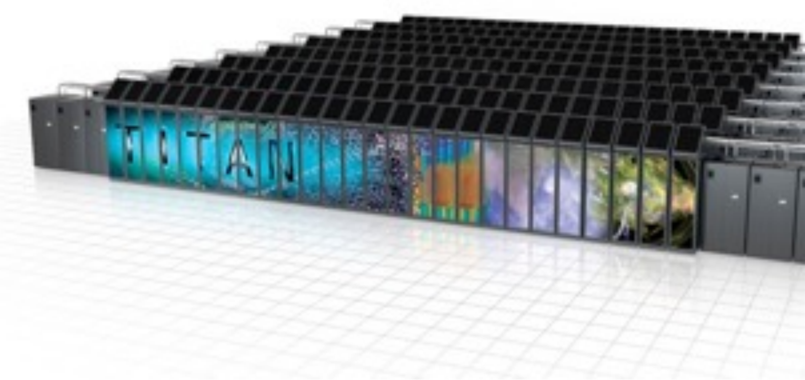

- **Capability resources provided by Leadership-Class facilities through peer-reviewed allocation process**

- **Capacity essentially not provided**
  - XSEDE, NERSC - but way too few cycles
  - US is not providing the necessary pyramid of resources required for projects

- **Need for Capability - BUT Bigger need for mid-scale Capacity** (e.g. 1K, 4K,16K cores)

- Limited ``big-data'' needs - just not seeing significant usage of the cloud by UW researchers - demand is at the few-nodes level - this will likely change, but not sure when or how

What is Needed Globally and Locally ?

``The Speed of Science”

- **Need access to hardware to allow for new ideas to be explored asap.**
  - e.g., if NERSC allocation runs out after 6 months, then have to wait 6 months to have more capacity computing time to develop idea
  - Hyak gives immediate access all-year round

- **Clear need for a local mid-scale resource that is on-demand for**
    - testing ideas
    - rapid development of new codes for Capability running
  - Needs to be big enough with necessary characteristics to be useful
  - Needs to be within a capable cyberinfrastructure
  - 40% of HYAK is used in this area

What is Needed Globally and Locally ?

``Preparing for Capability Production’’

- **Need on-demand resources for code development for production at scale.**
- Needs a ``representative environment’’
- Needs to be within a capable cyberinfrastructure
- 40% of HYAK is used in this area

What is Needed Globally and Locally ?

``Cloud Ain't Secure'' - despite what some might say!

- **Need on-demand resources for Cloud Computing that is secure**
  - local control over permissions etc - not NSA readable.

``Big Data Pipelines''

- **Need on-demand resources to handle large data sets generated with capability or capacity resources, or experiment.**
  - efficient data movement, storage and retrieval
    - fast connection to outside world, within the university, to leadership class and capacity computing centers.
- Semi-real-time processing data from high-throughput instruments
- Needs to be within a capable cyberinfrastructure
- Needs good network-infrastructure
- 20% of HYAK is used in this area

What is Needed Globally and Locally ?

UW Needs mid-scale Capacity Computing with Capability-Computing Hardware

More mid-scale capacity machines would meet the current need, e.g., 100K, 50K, 10K core machines
(Hyak currently has 11K cores and 48 GPUs)

Hardware needs to match the user characteristics at the UW

- small number of capability and large number of small job size
- CPUs at the 99% level - few GPU capable users
- science per unit time is the criterion
- configurable hardware options

What is Needed Globally and Locally ?

UW Needs more HPC-capable FTEs to be embedded in projects to port and continually evolve codes to rapidly changing architectures.

e.g., PNNL-UW collaboration in Nuclear Physics (Ken Roche) is a very successful demonstration of such a collaborative effort.

SciDAC projects work - need them for UW/WA research
(eScience Institute was supposed to do this - it does not)

What is Needed Globally and Locally ?

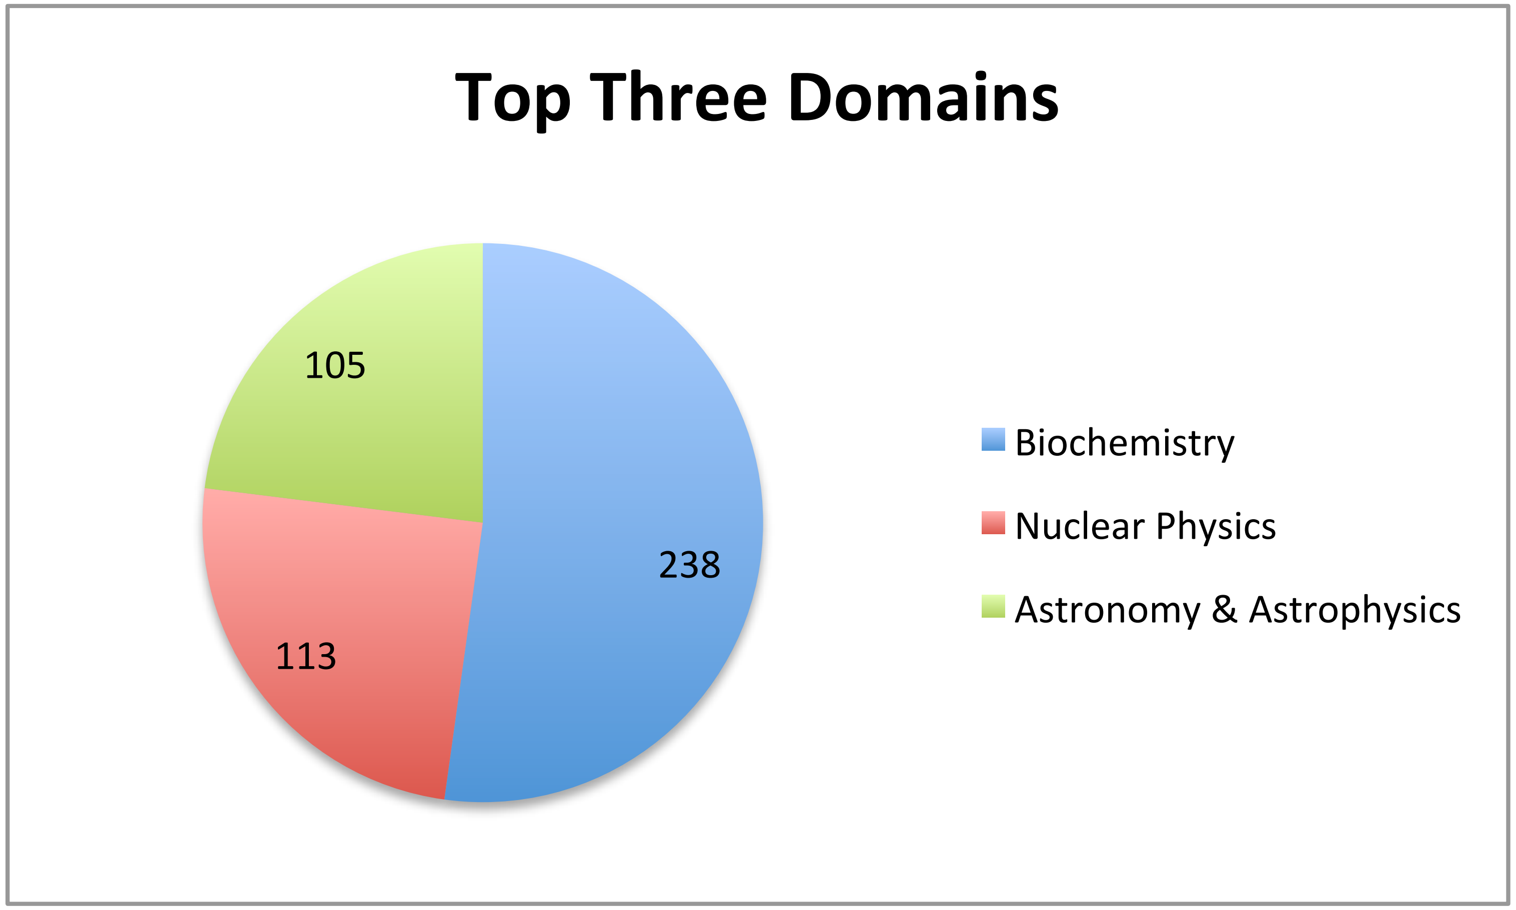

standardized core-hours

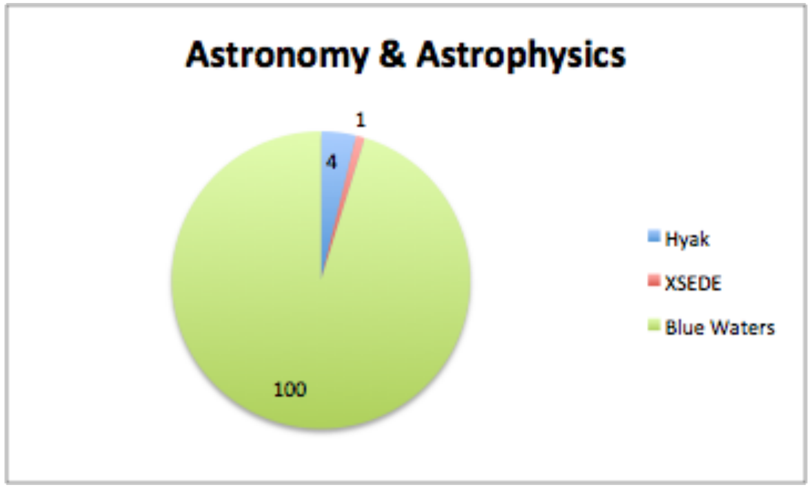

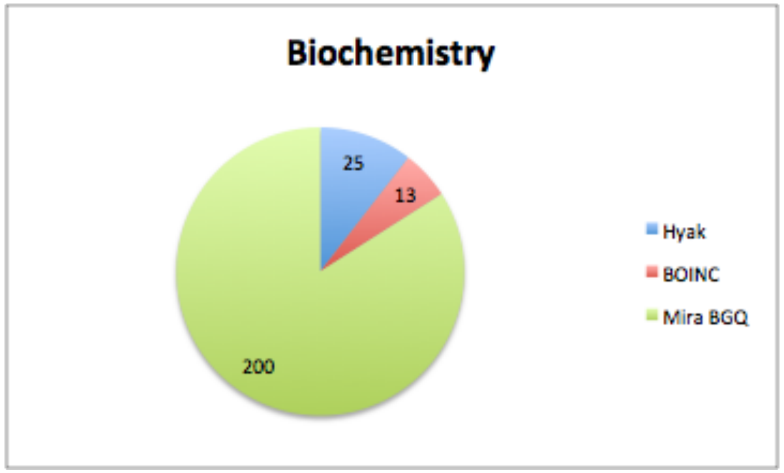

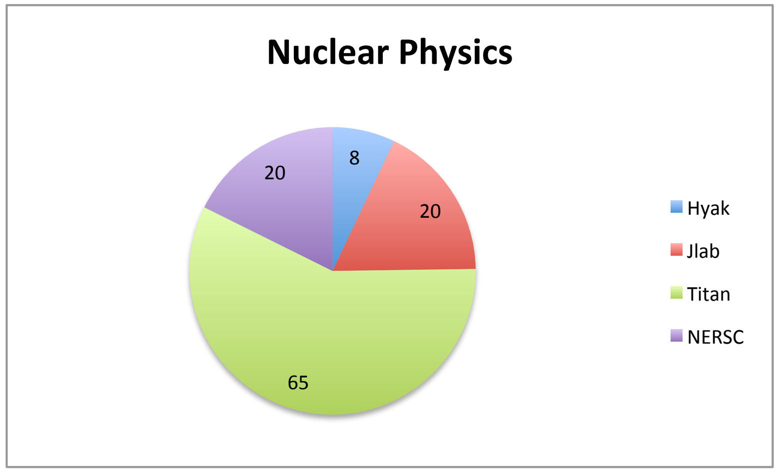

HPC by the Numbers

A (very) few Science Cases

Bioinformatics

Raw data is collected from the environment or *in vitro* mimics of environmental conditions, then collated, assembled and analyzed.

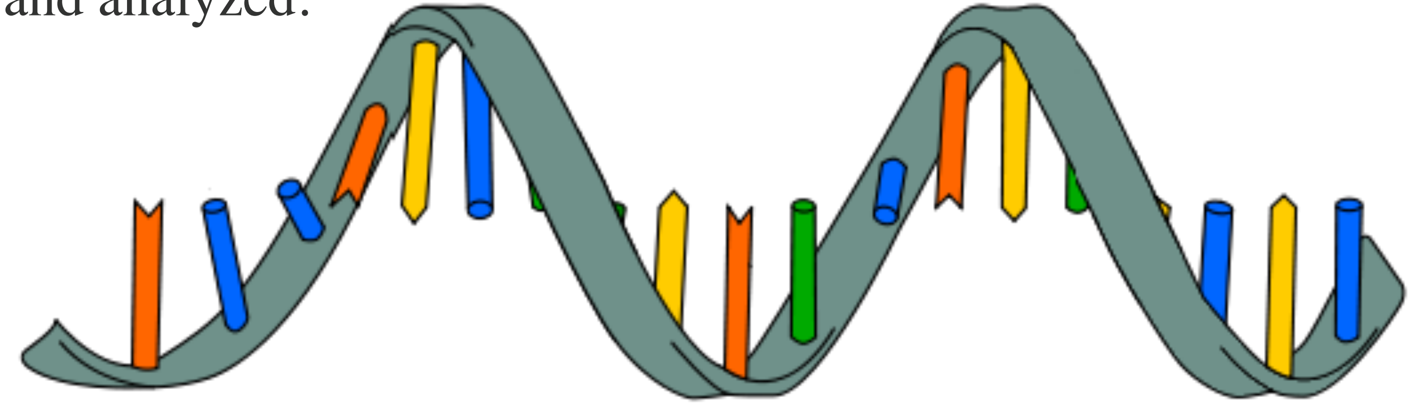

Lidstrom Lab

As an example, two data sets consisting of 60 million unattributed RNA sequences are run against 10 million known protein sequences.

The goal is to identify proteins that may be derived from the RNA sequences.

pdqBLAST (created by Dave Beck) provides an infrastructure for running such queries on Hyak, taking advantage of **Hyak's** significant storage capacity, bandwidth and computational power.

The query ran on 70 nodes, with 8 cores per node, for 20 days : requiring 268,800 core-hours.

pdqBLAST applied to queries of this sort is well suited to Hyak.
Specifically, Hyak offers the ability to perform week-long queries over large numbers of nodes. Tasks of this sort are difficult to arrange on existing systems at national supercomputer centers.

As bioinformatics workloads of this sort scale up, **Hyak** serves as a valuable development platform for the next generation of pre-exascale data intensive computers currently under development.

Genomics

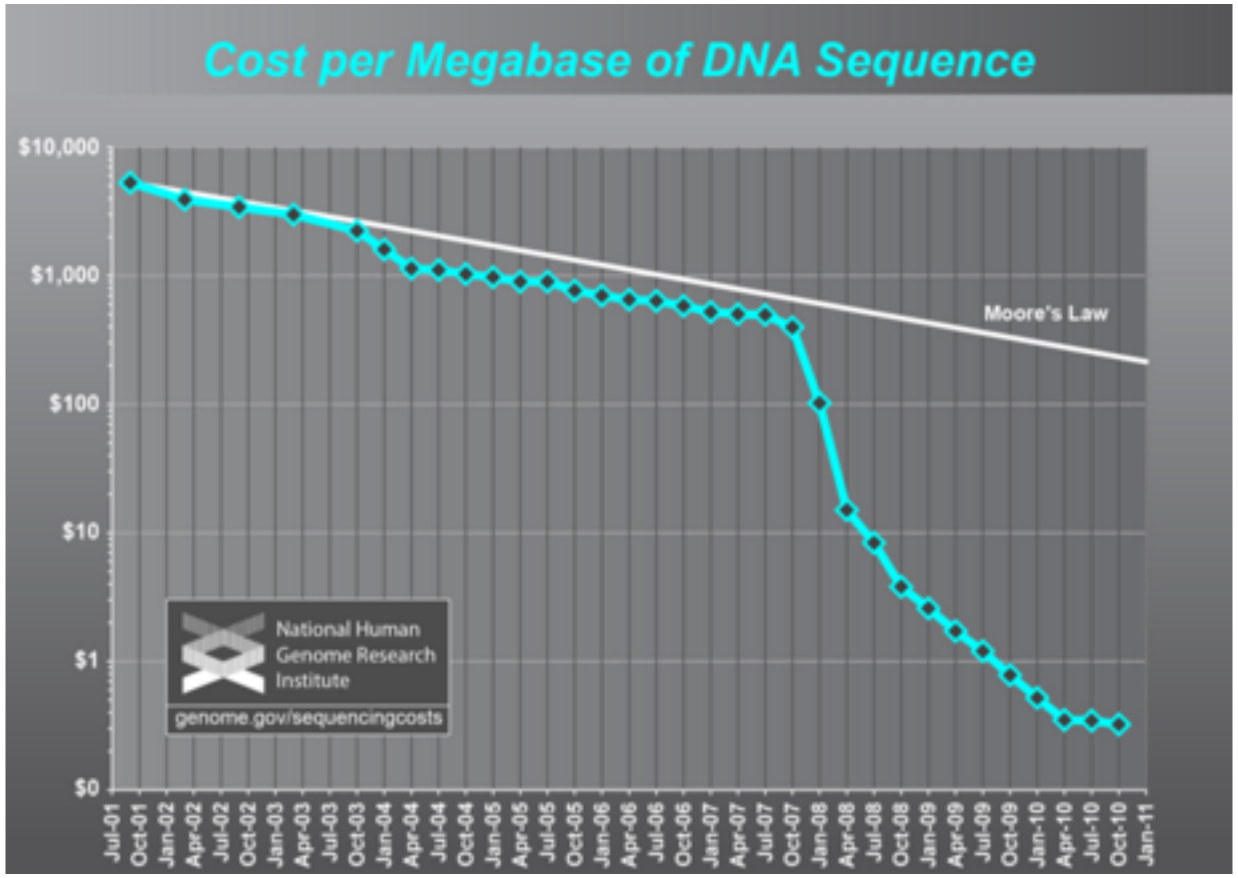

Cost per Megabase of DNA Sequence

Scatterplot of Project Talent Psychometric Test Scores (9th Grade)

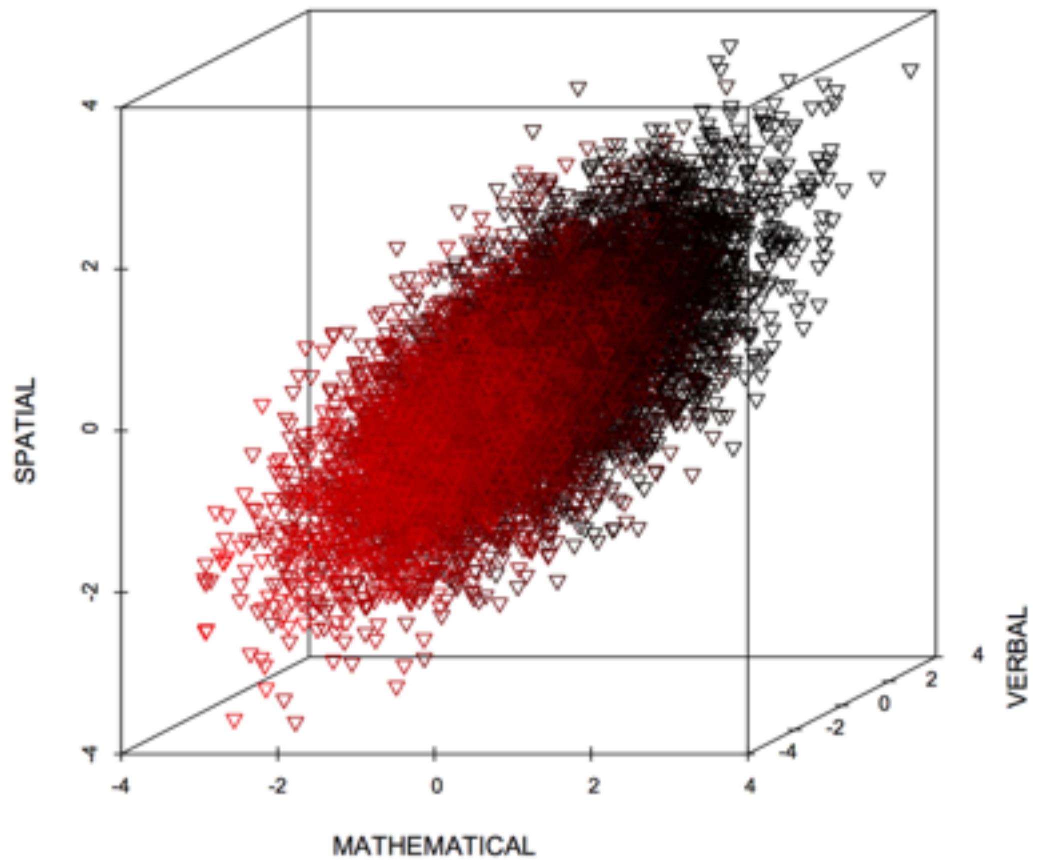

$$y = \sum_i g_i x_i + \sum_{ij} g_i g_j z_{ij} + \mathcal{O}(g^3) + \epsilon$$

- Linear Algebra common to many other fields
  - linear solvers - collaboration with AM/CS/Physics is crucial
- Requires Big-Data, Capacity and Capability Computing
  - predict height, intelligence with precision

Designer Proteins

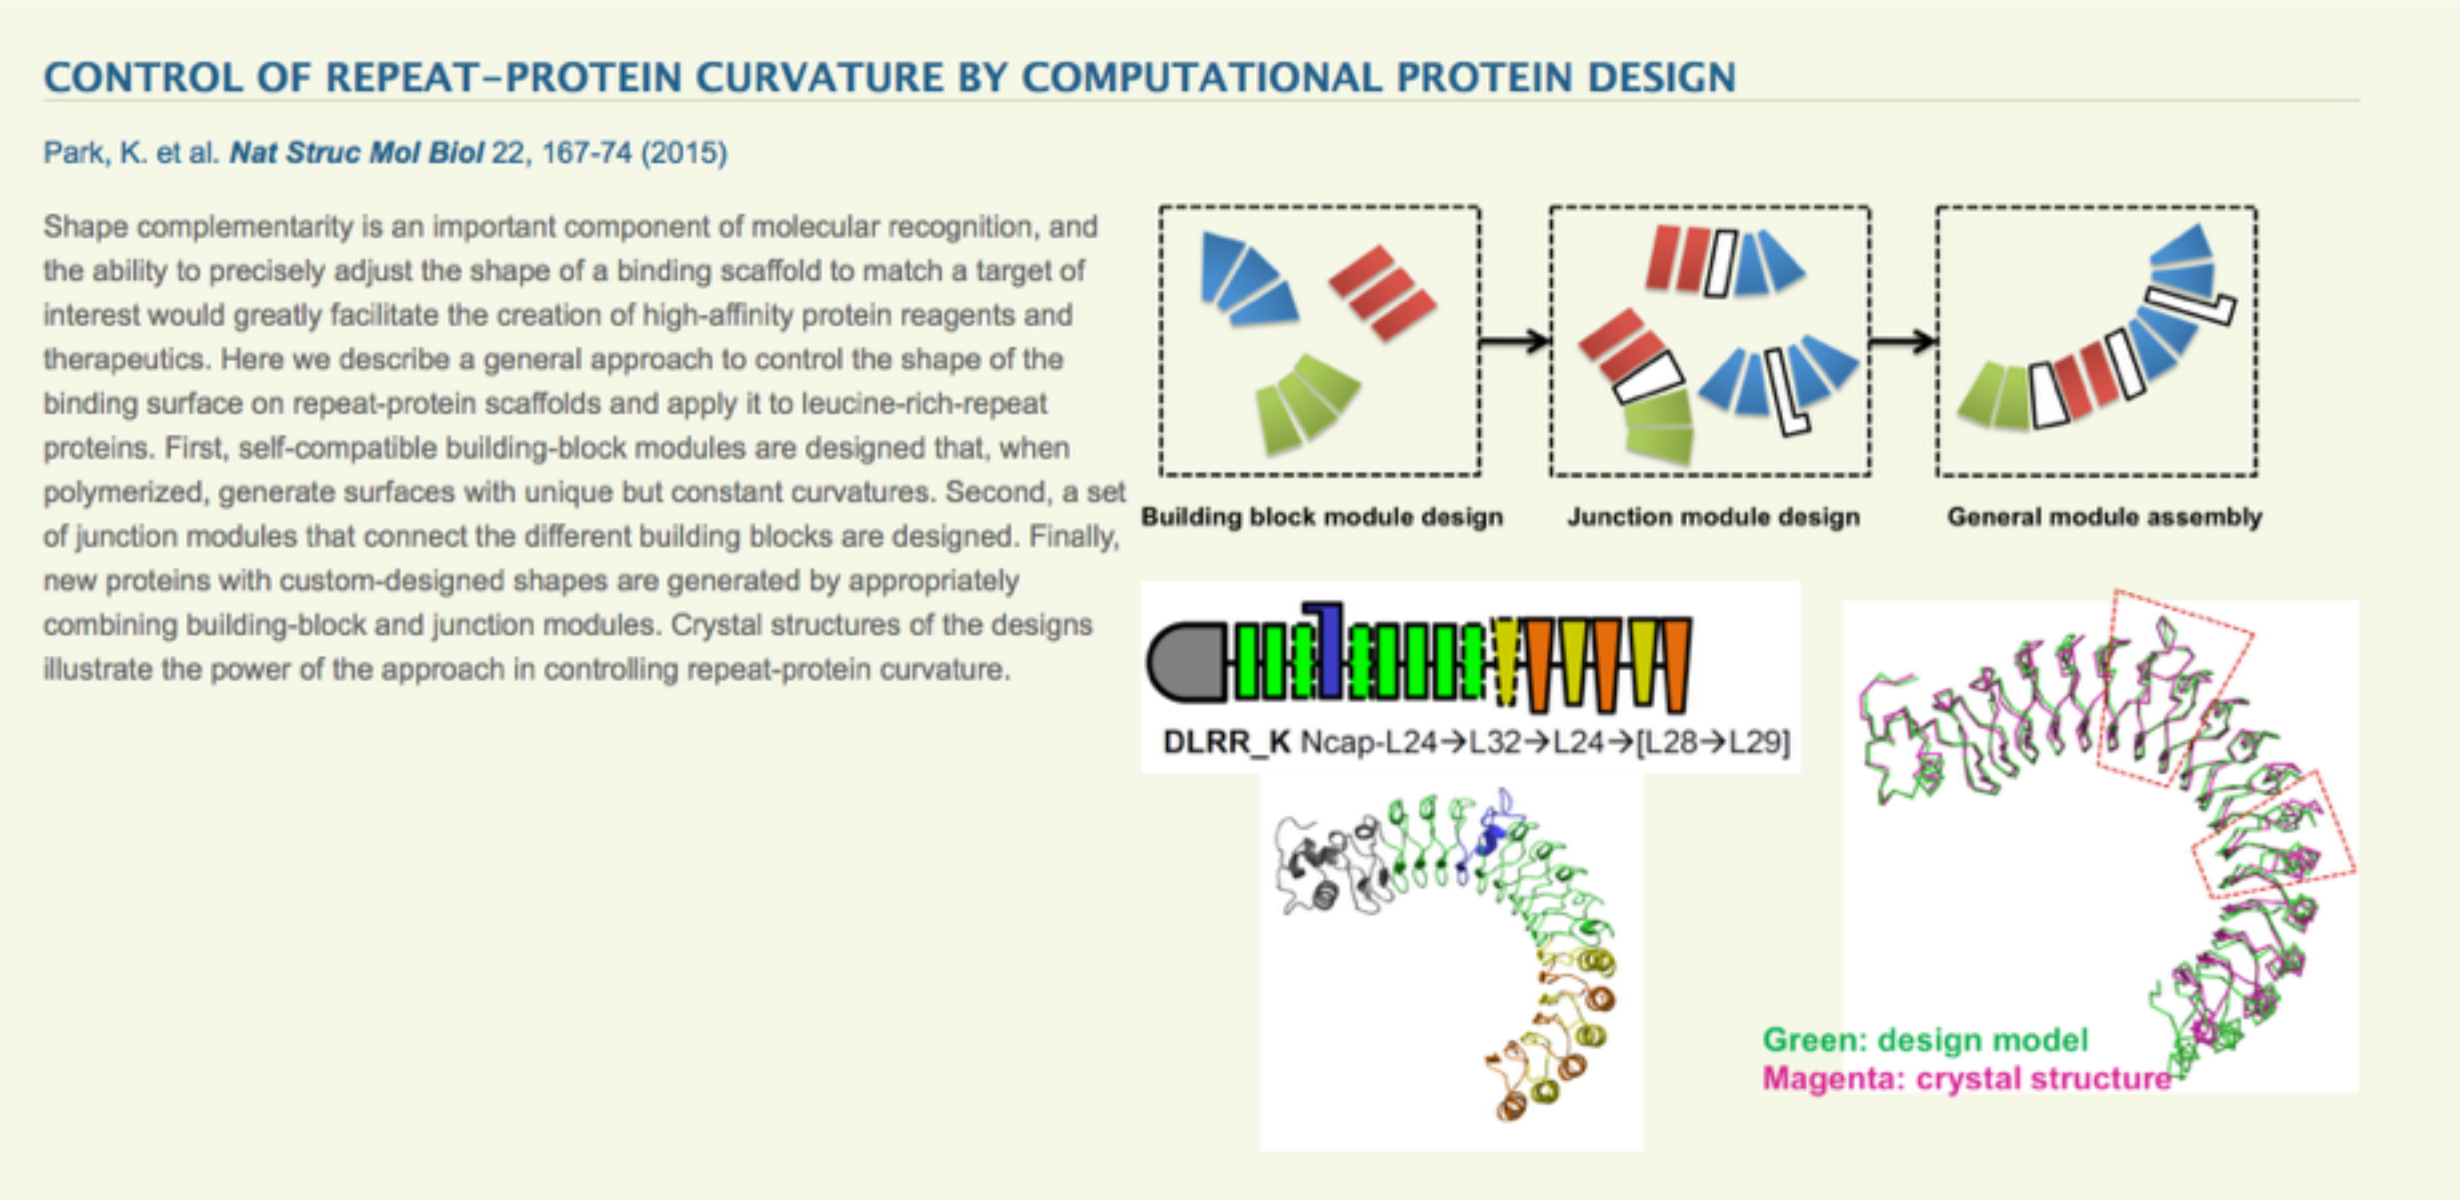

CONTROL OF REPEAT-PROTEIN CURVATURE BY COMPUTATIONAL PROTEIN DESIGN

Park, K. et al. ***Nat Struc Mol Biol*** 22, 167-74 (2015)

Shape complementarity is an important component of molecular recognition, and the ability to precisely adjust the shape of a binding scaffold to match a target of interest would greatly facilitate the creation of high-affinity protein reagents and therapeutics. Here we describe a general approach to control the shape of the binding surface on repeat-protein scaffolds and apply it to leucine-rich-repeat proteins. First, self-compatible building-block modules are designed that, when polymerized, generate surfaces with unique but constant curvatures. Second, a set of junction modules that connect the different building blocks are designed. Finally, new proteins with custom-designed shapes are generated by appropriately combining building-block and junction modules. Crystal structures of the designs illustrate the power of the approach in controlling repeat-protein curvature.

- Monte-Carlo strings of molecular subunits to minimize energy.
- Single core size jobs in general, runs on multi-cores on **HYAK**, also runs on BG/P and BG/Q through INCITE allocations.
- Not GPU/Phi-capable at present
- **Hyak** is MUCH cheaper compute platform than any other

Baker's Lab : http://www.bakerlab.org/index/index.html

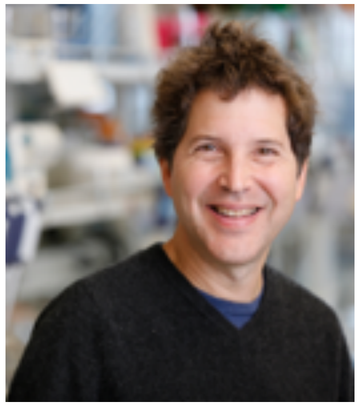

Designer Proteins

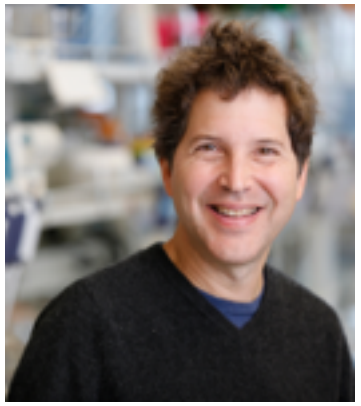

Kemp eliminase

Retro-aldolase

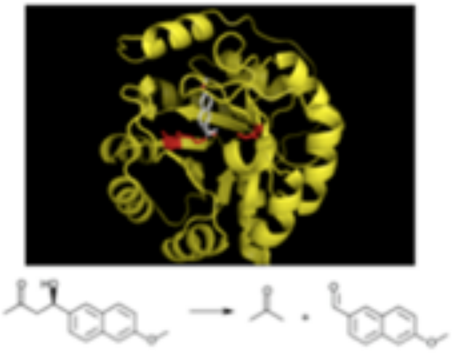

Esterase

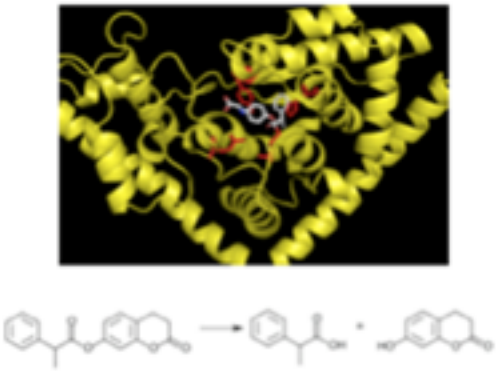

Diels-Alder enzyme

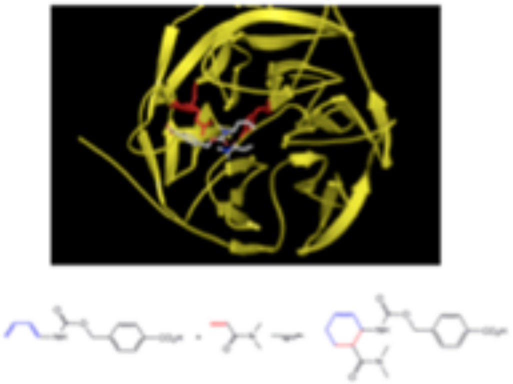

Rosetta@home

Protein Folding, Design, and Docking

Baker's Lab : http://www.bakerlab.org/index/index.html

Protein Design

Improve the physical model and the sampling methodology underlying the prediction and design calculations in **Rosetta**.

On the structure calculation side, strive for consistent near-atomic resolution *ab initio* structure prediction for small proteins, and work towards atomic level structure determination for proteins greater than 200 amino acids.

Focus on membrane proteins and other systems for which obtaining high resolution experimental data is difficult—this is where our approach are likely to contribute the most.

Extend data guided structure determination to biological assemblies.

Extend the methodology to non natural amino acids and cofactors to try to leapfrog over the limitations nature has faced with the limited set of twenty amino acids. Design a complete pathway for fuel production from CO2 using solar generated reducing equivalents.

Develop and test methods for designing high affinity binders/ inhibitors for any specified surface patch on a protein of known structure.

Develop new biomolecules with new functions—inhibitors, enzymes, endonucleases, and vaccines—that can have a positive impact on the world.

Molecular Science

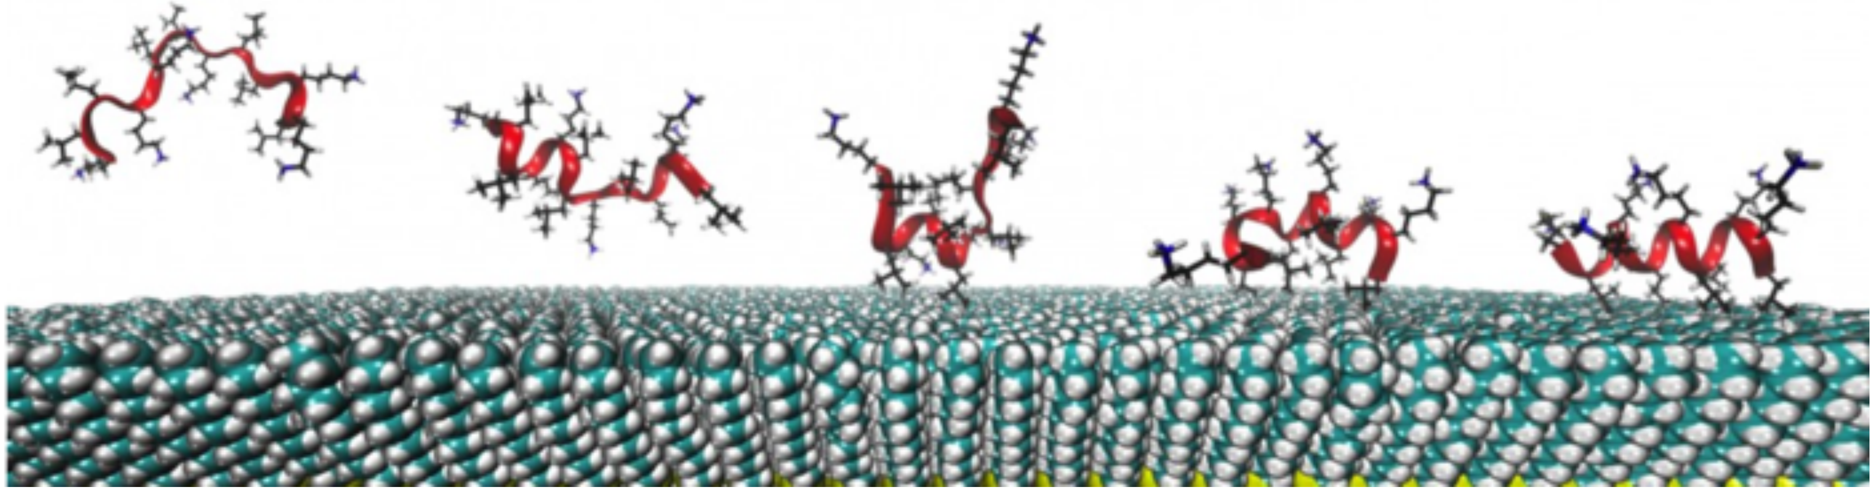

Jim Pfaendtner's research group in Chemical Engineering specializes in computational molecular science and engineering applied to a wide range of interesting systems and technological problems in soft matter, biophysics, biocatalysis, and reaction engineering.

Current areas of interest include protein/surface interactions and the interactions of enzymes with novel solvents.

Main computational tool is molecular dynamics.

One of the few efforts on UW campus that can utilize GPUs.

Without Hyak, Jim would not have come to the UW !

Hyak is important for recruiting and retaining for junior scientists.

Physics

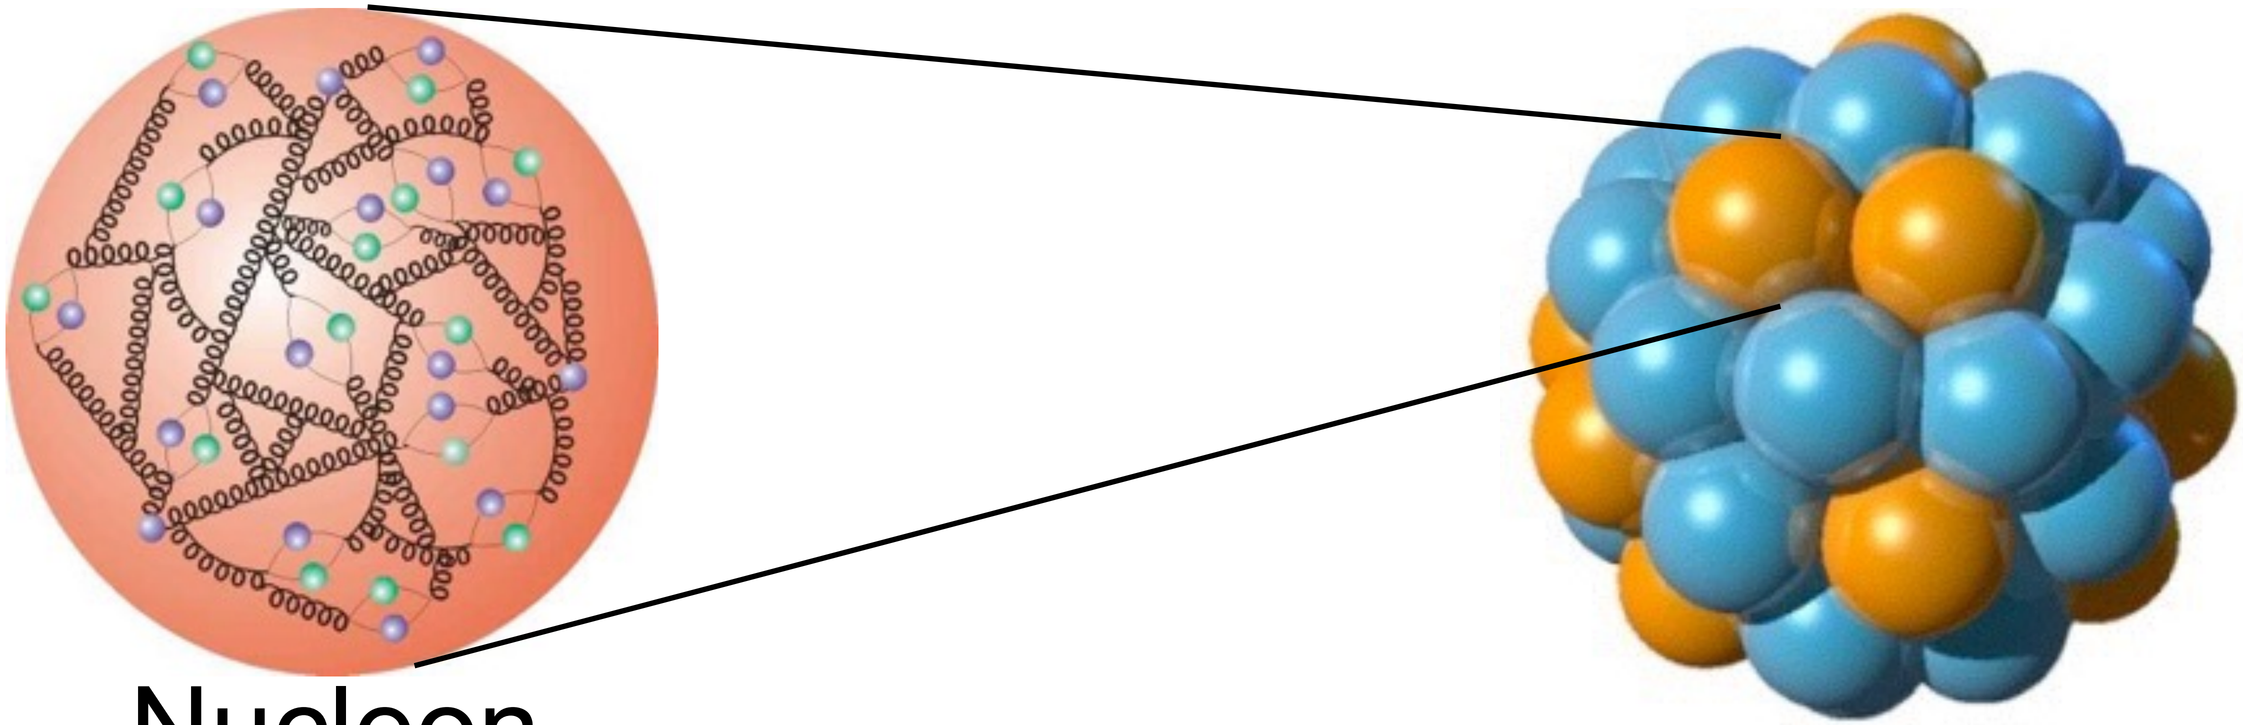

Nucleon

Nucleus

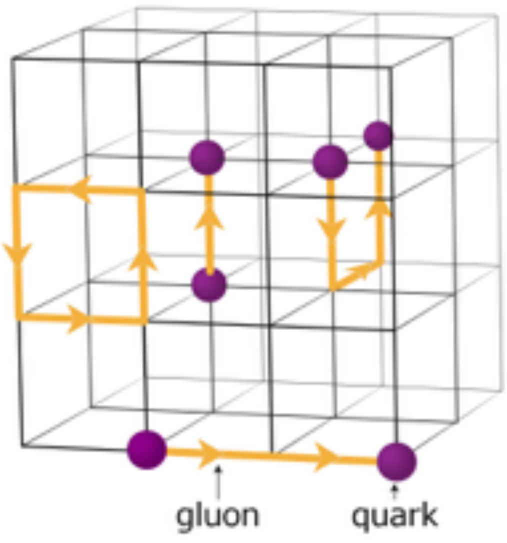

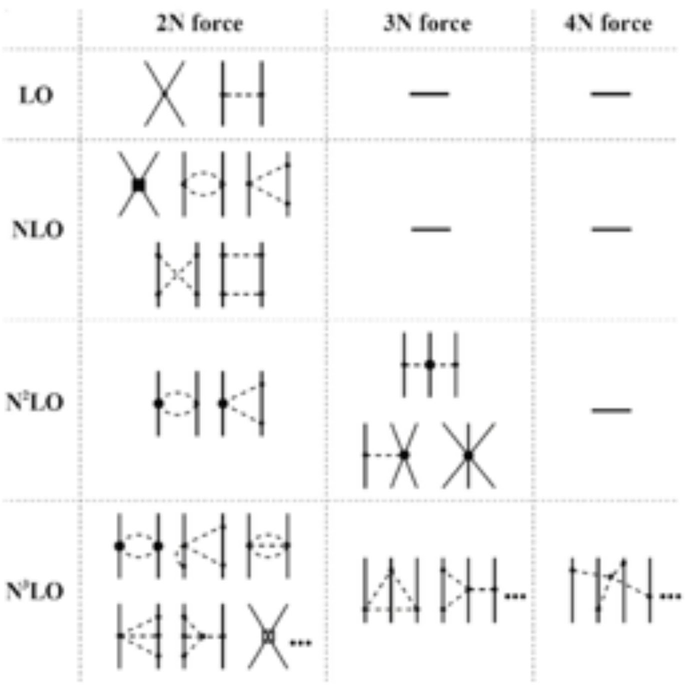

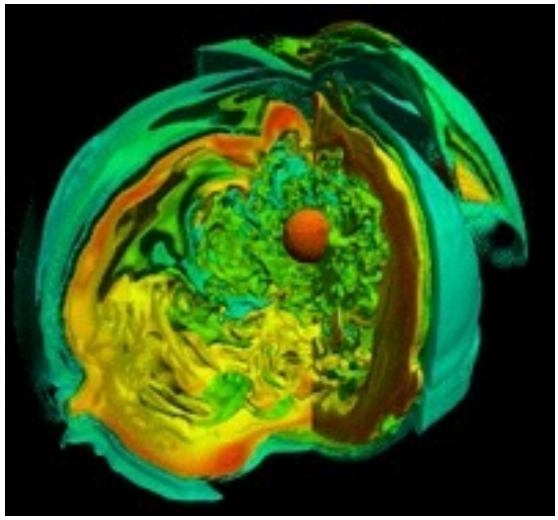

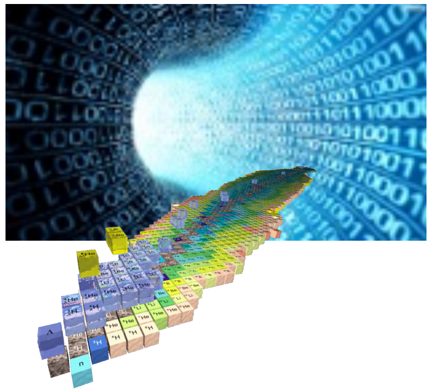

Nuclear Physics
Lattice QCD and Nuclear Many-Body

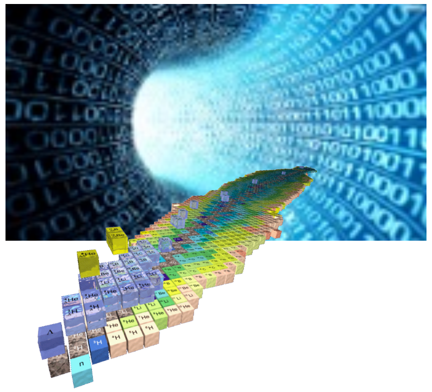

Lattice QCD - USQCD collaboration

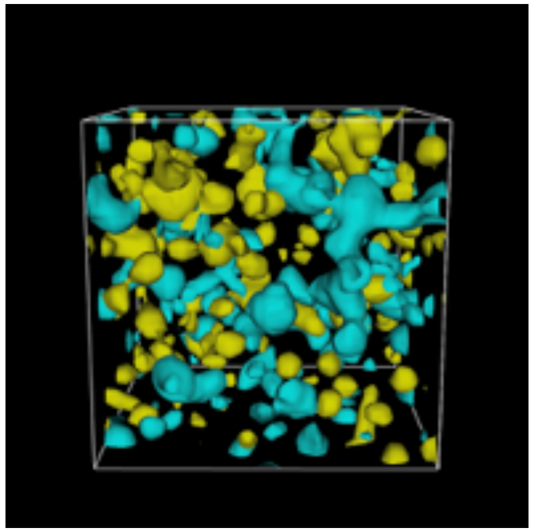

Configuration : e.g.,

Vol = 32 x 32 x 32 x 256 lattice sites
Vol x 4 x 8 = 268 Million independent real numbers
to define $U_\mu(x)$ (generally double precision)

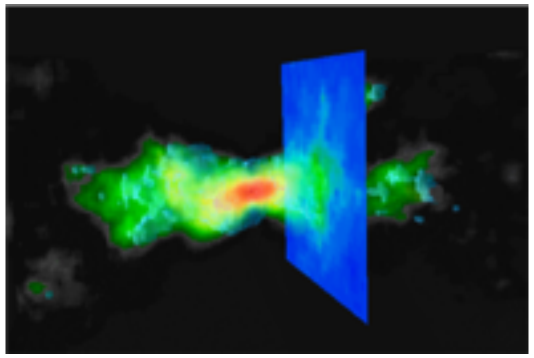

Propagator : e.g.,

32 x 32 x 32 x 256 lattice sites
~ 100 Million x 100 Million complex sparse matrix
(to invert and take determinant)

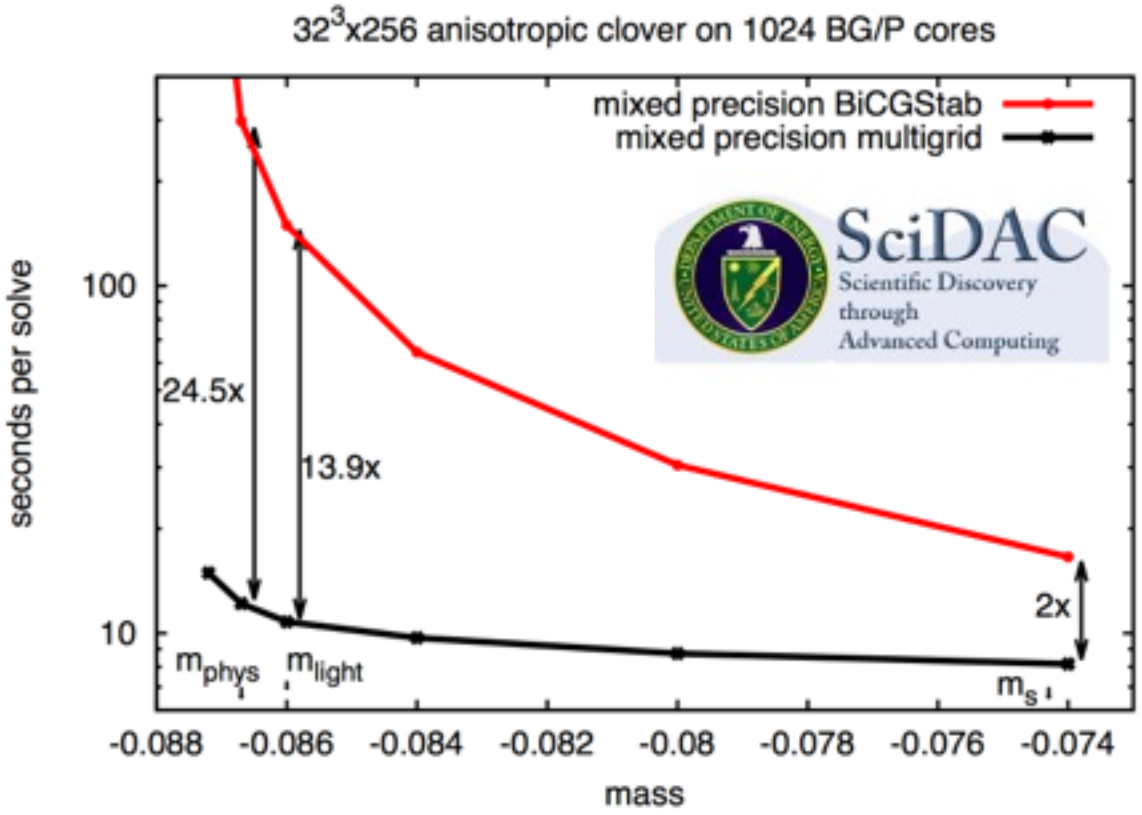

Nuclear Many-Body (Bulgac (UW)+Forbes(WSU)+Roche(PNNL))

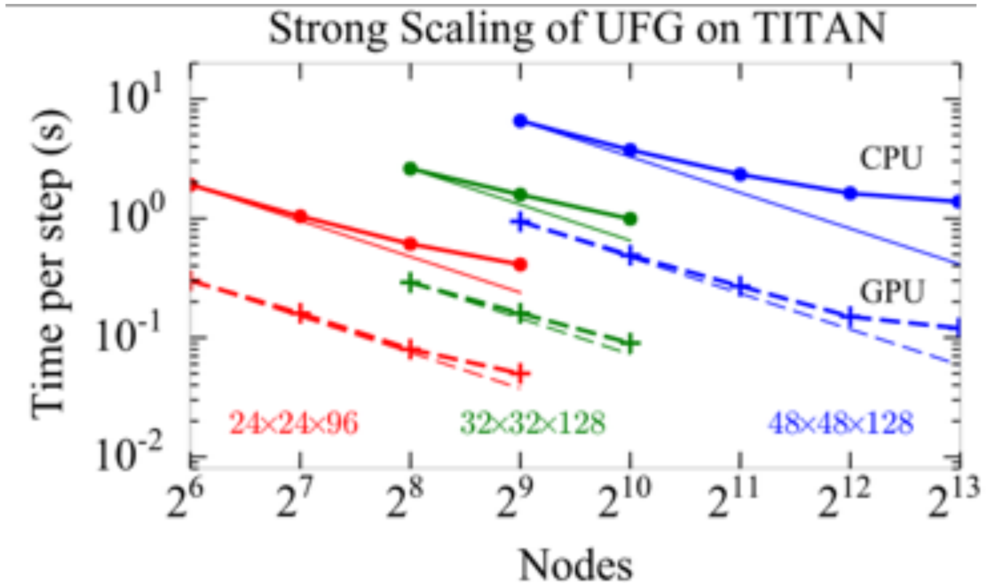

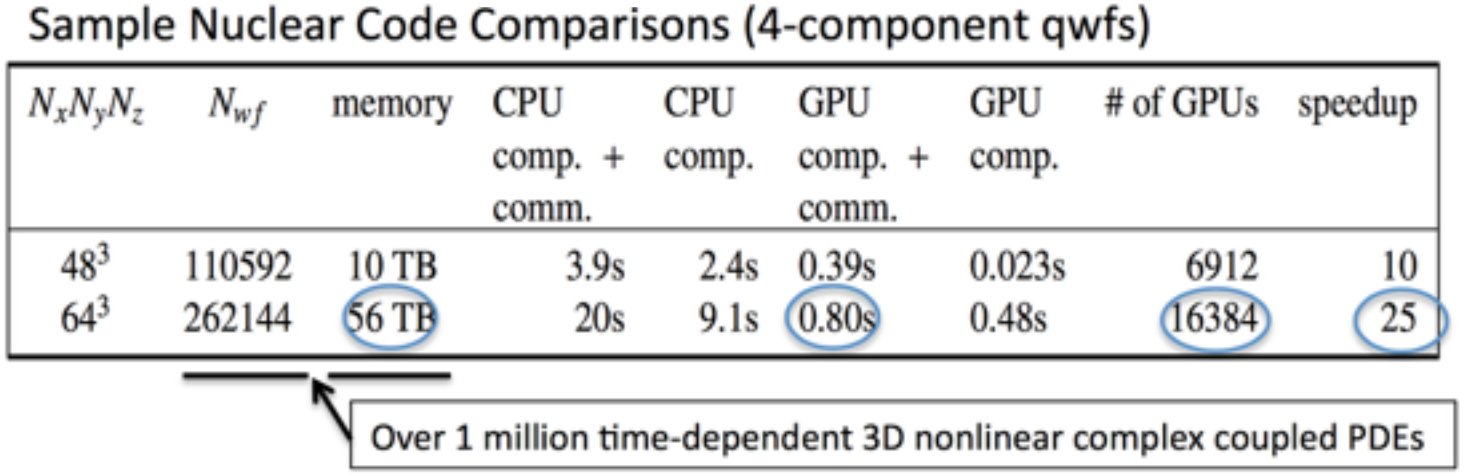

Sample Nuclear Code Comparisons (4-component qwfs)

$N_xN_yN_z$	N_{wf}	memory	CPU comp. + comm.	CPU comp.	GPU comp. + comm.	GPU comp.	# of GPUs	speedup
48^3	110592	10 TB	3.9s	2.4s	0.39s	0.023s	6912	10
64^3	262144	56 TB	20s	9.1s	0.80s	0.48s	16384	25

Over 1 million time-dependent 3D nonlinear complex coupled PDEs

Astronomy and Astrophysics

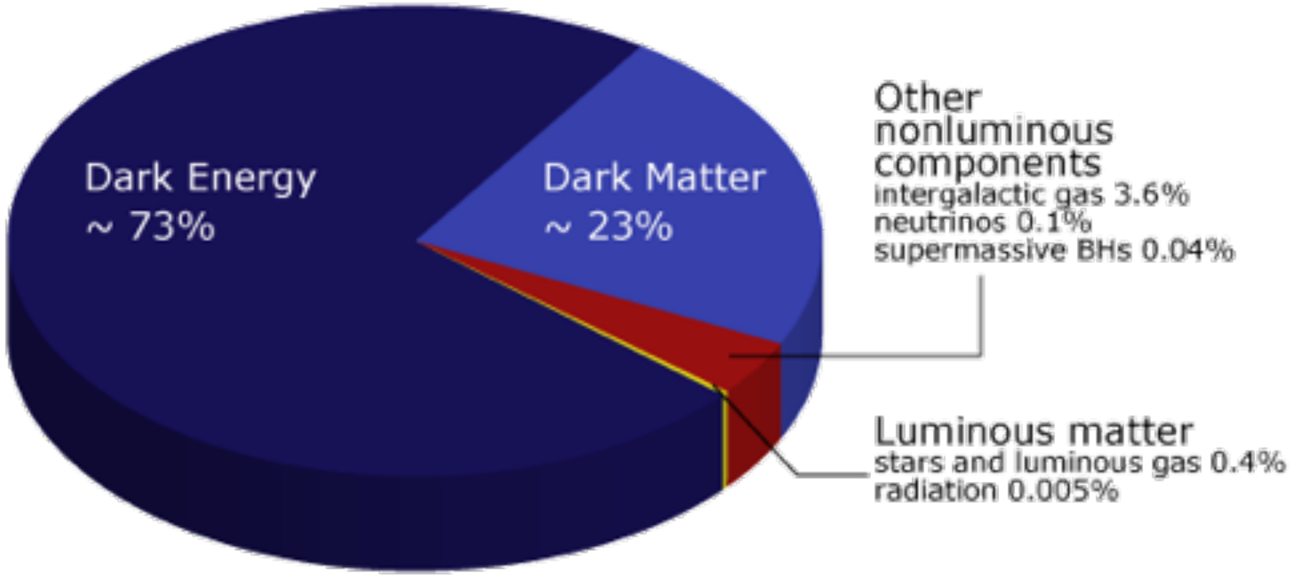

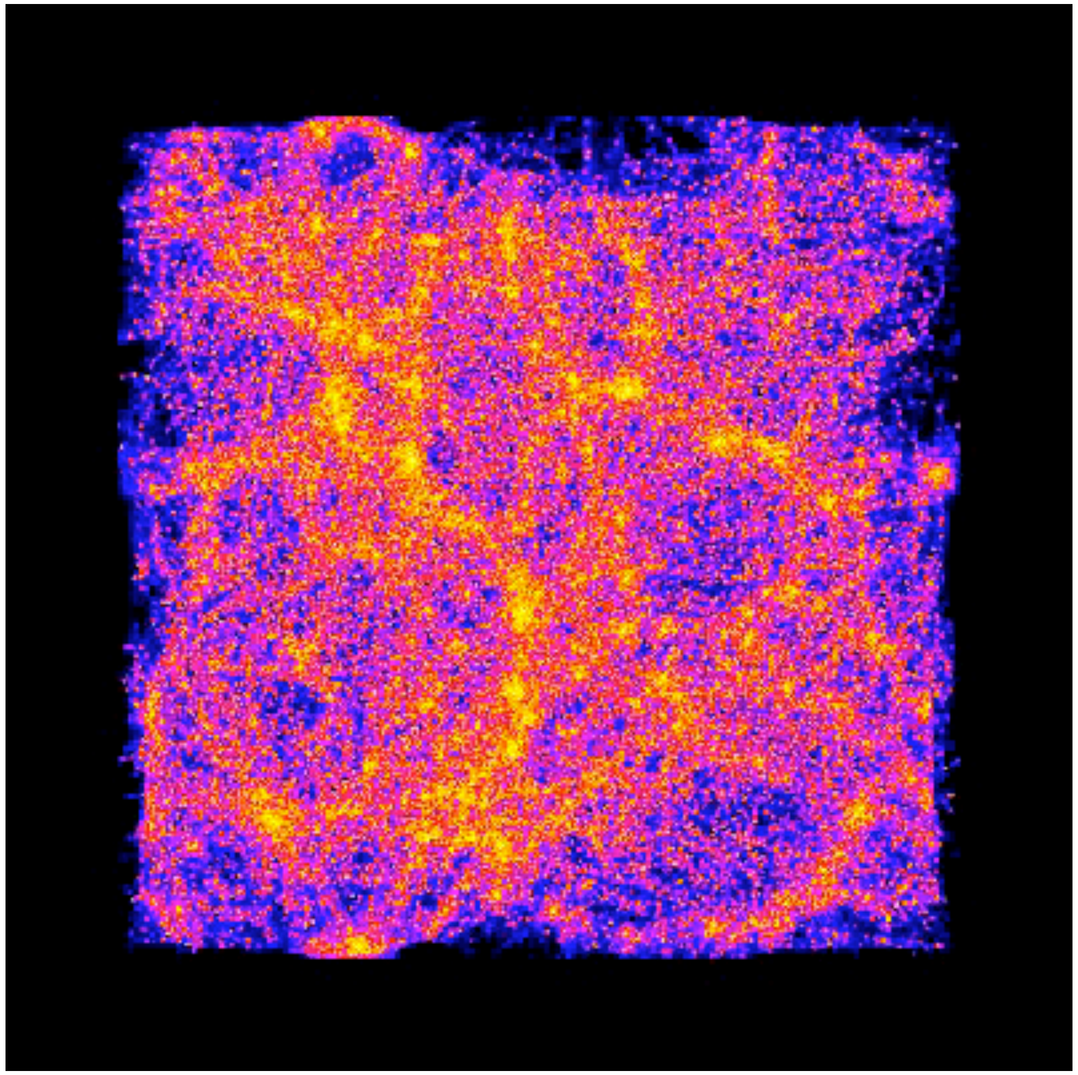

The N-Body Shop

Tom Quinn (UW)

Climate

UW’s Climates Impact Group
College of the Environment

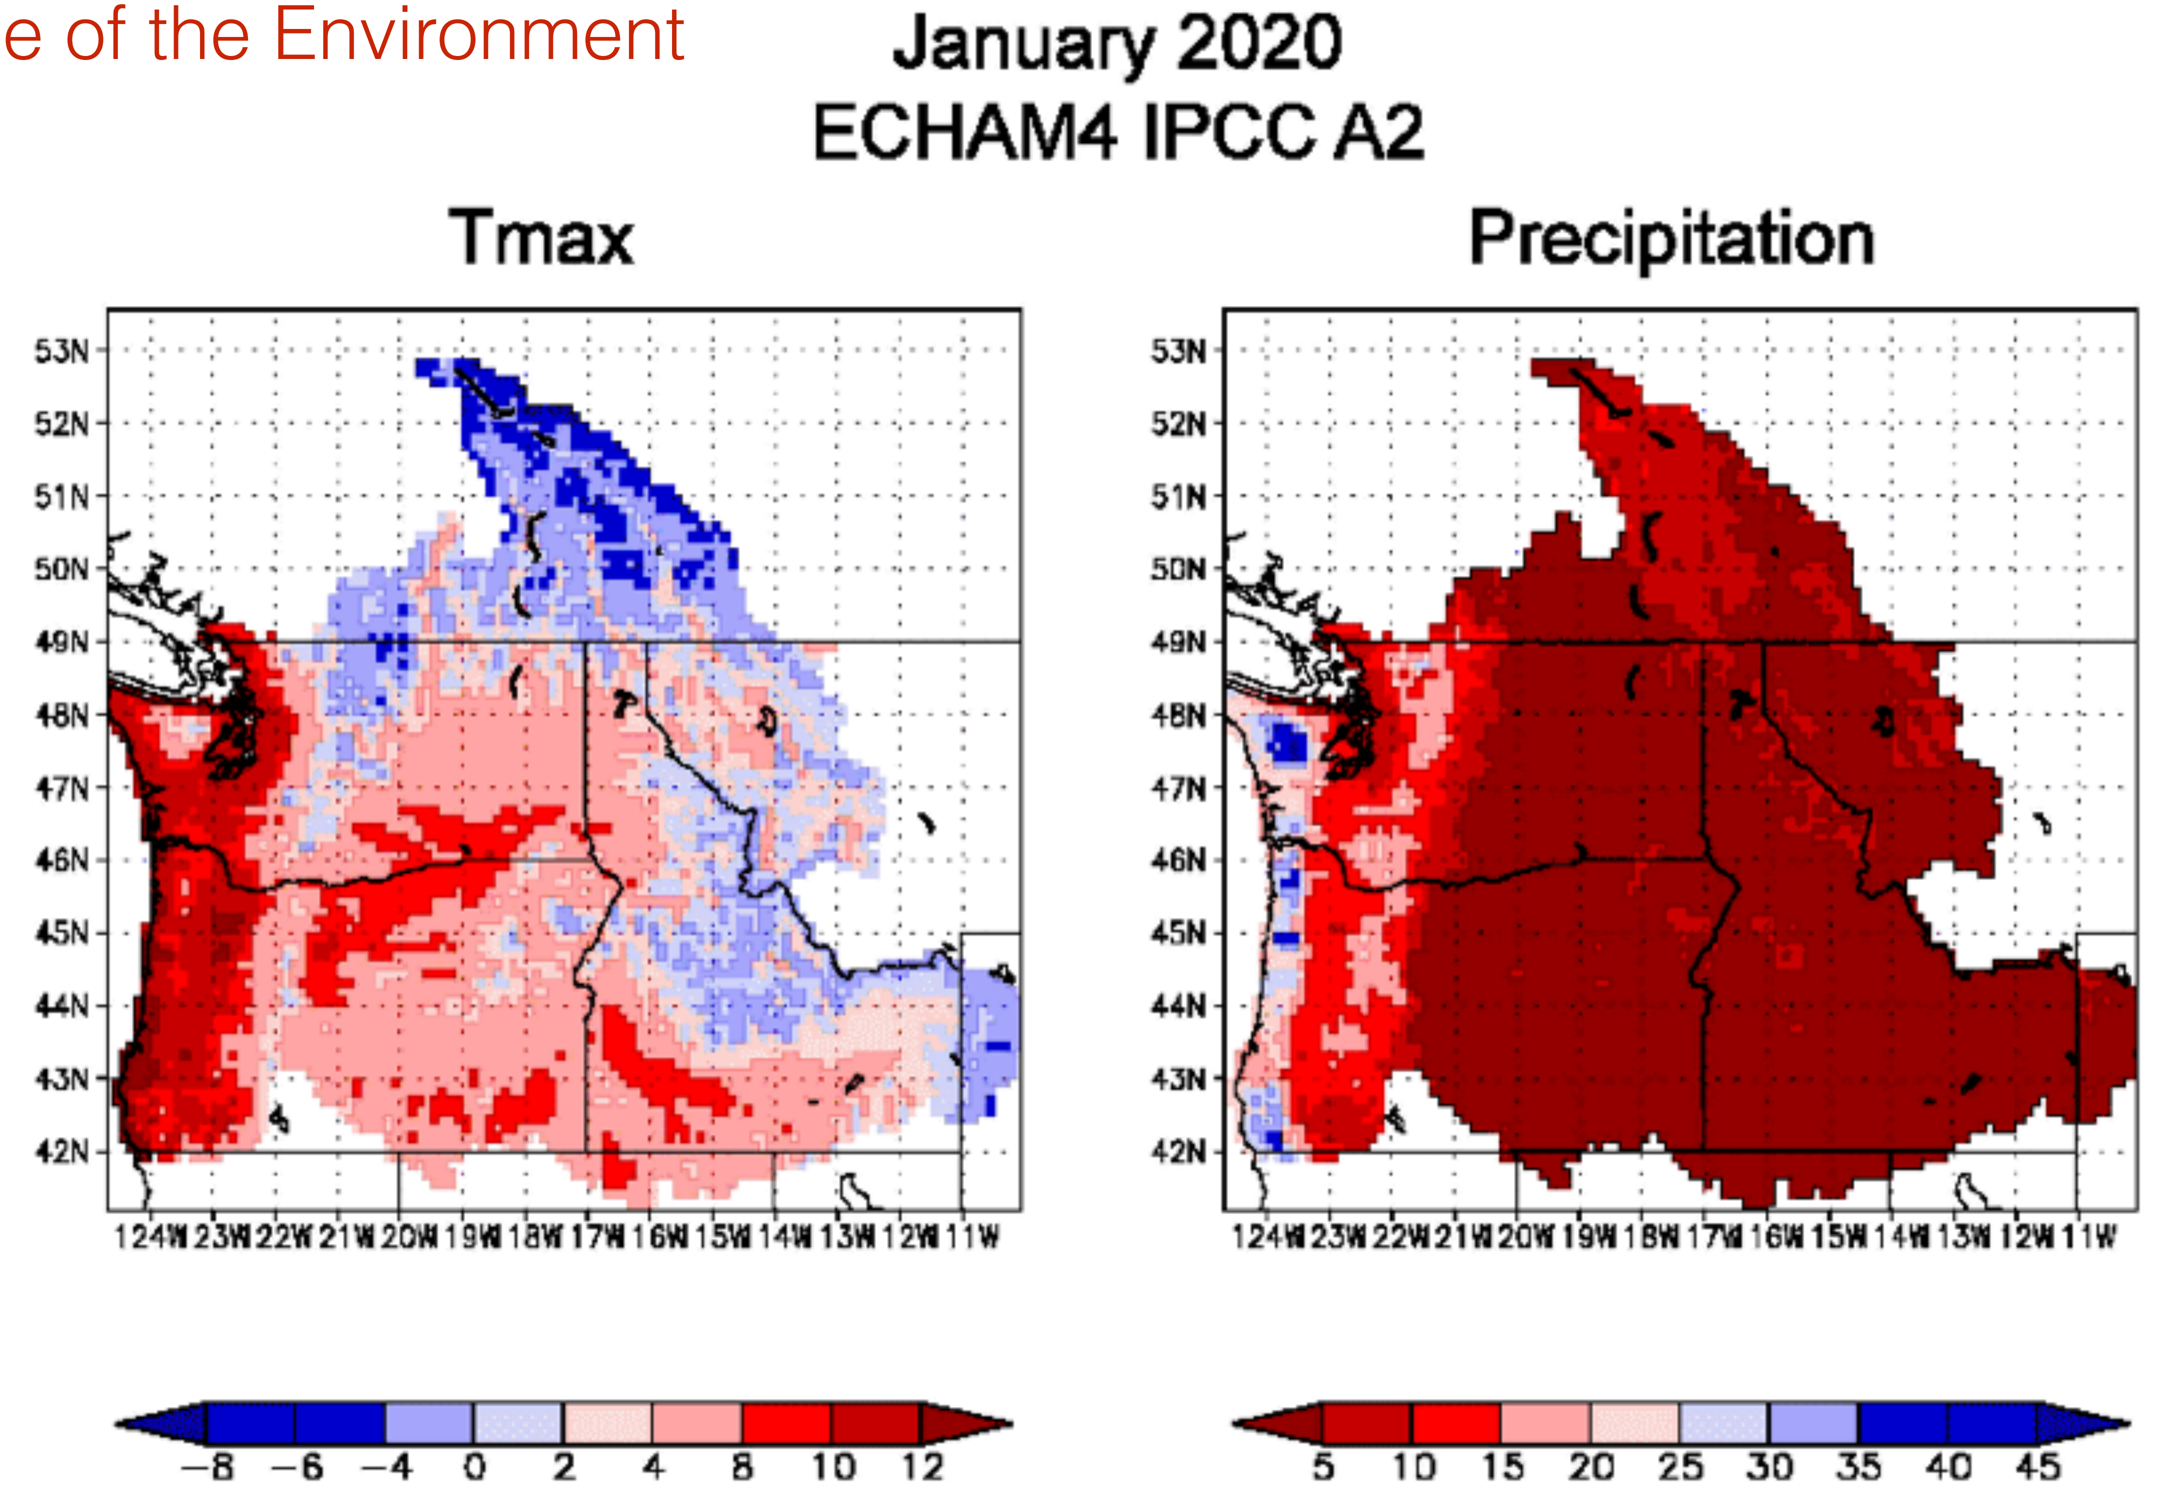

Summary

UW research programs have a long-standing, well-established and compelling need for larger capability and capacity computing resources.

- A small number of capability users, a large number of capacity users
- ~log-normal distribution of jobs (same as any platform)
- CPU are main hardware, accelerators have not been adopted at any significant level.
- UW research need more of both types of compute resources
- Limited, but specific, need for cloud-type computing

Since 2007

- infrastructure has been planned and deployed to address the mid-scale and small-scale capacity
- needed for code development and small-scale production
- support from UW admin. has been critical

More HPC-savvy people are needed at UW

Summary

UW should be a Center for Scientific Computing

WA State should be a leader in Scientific Computing

Talk Order

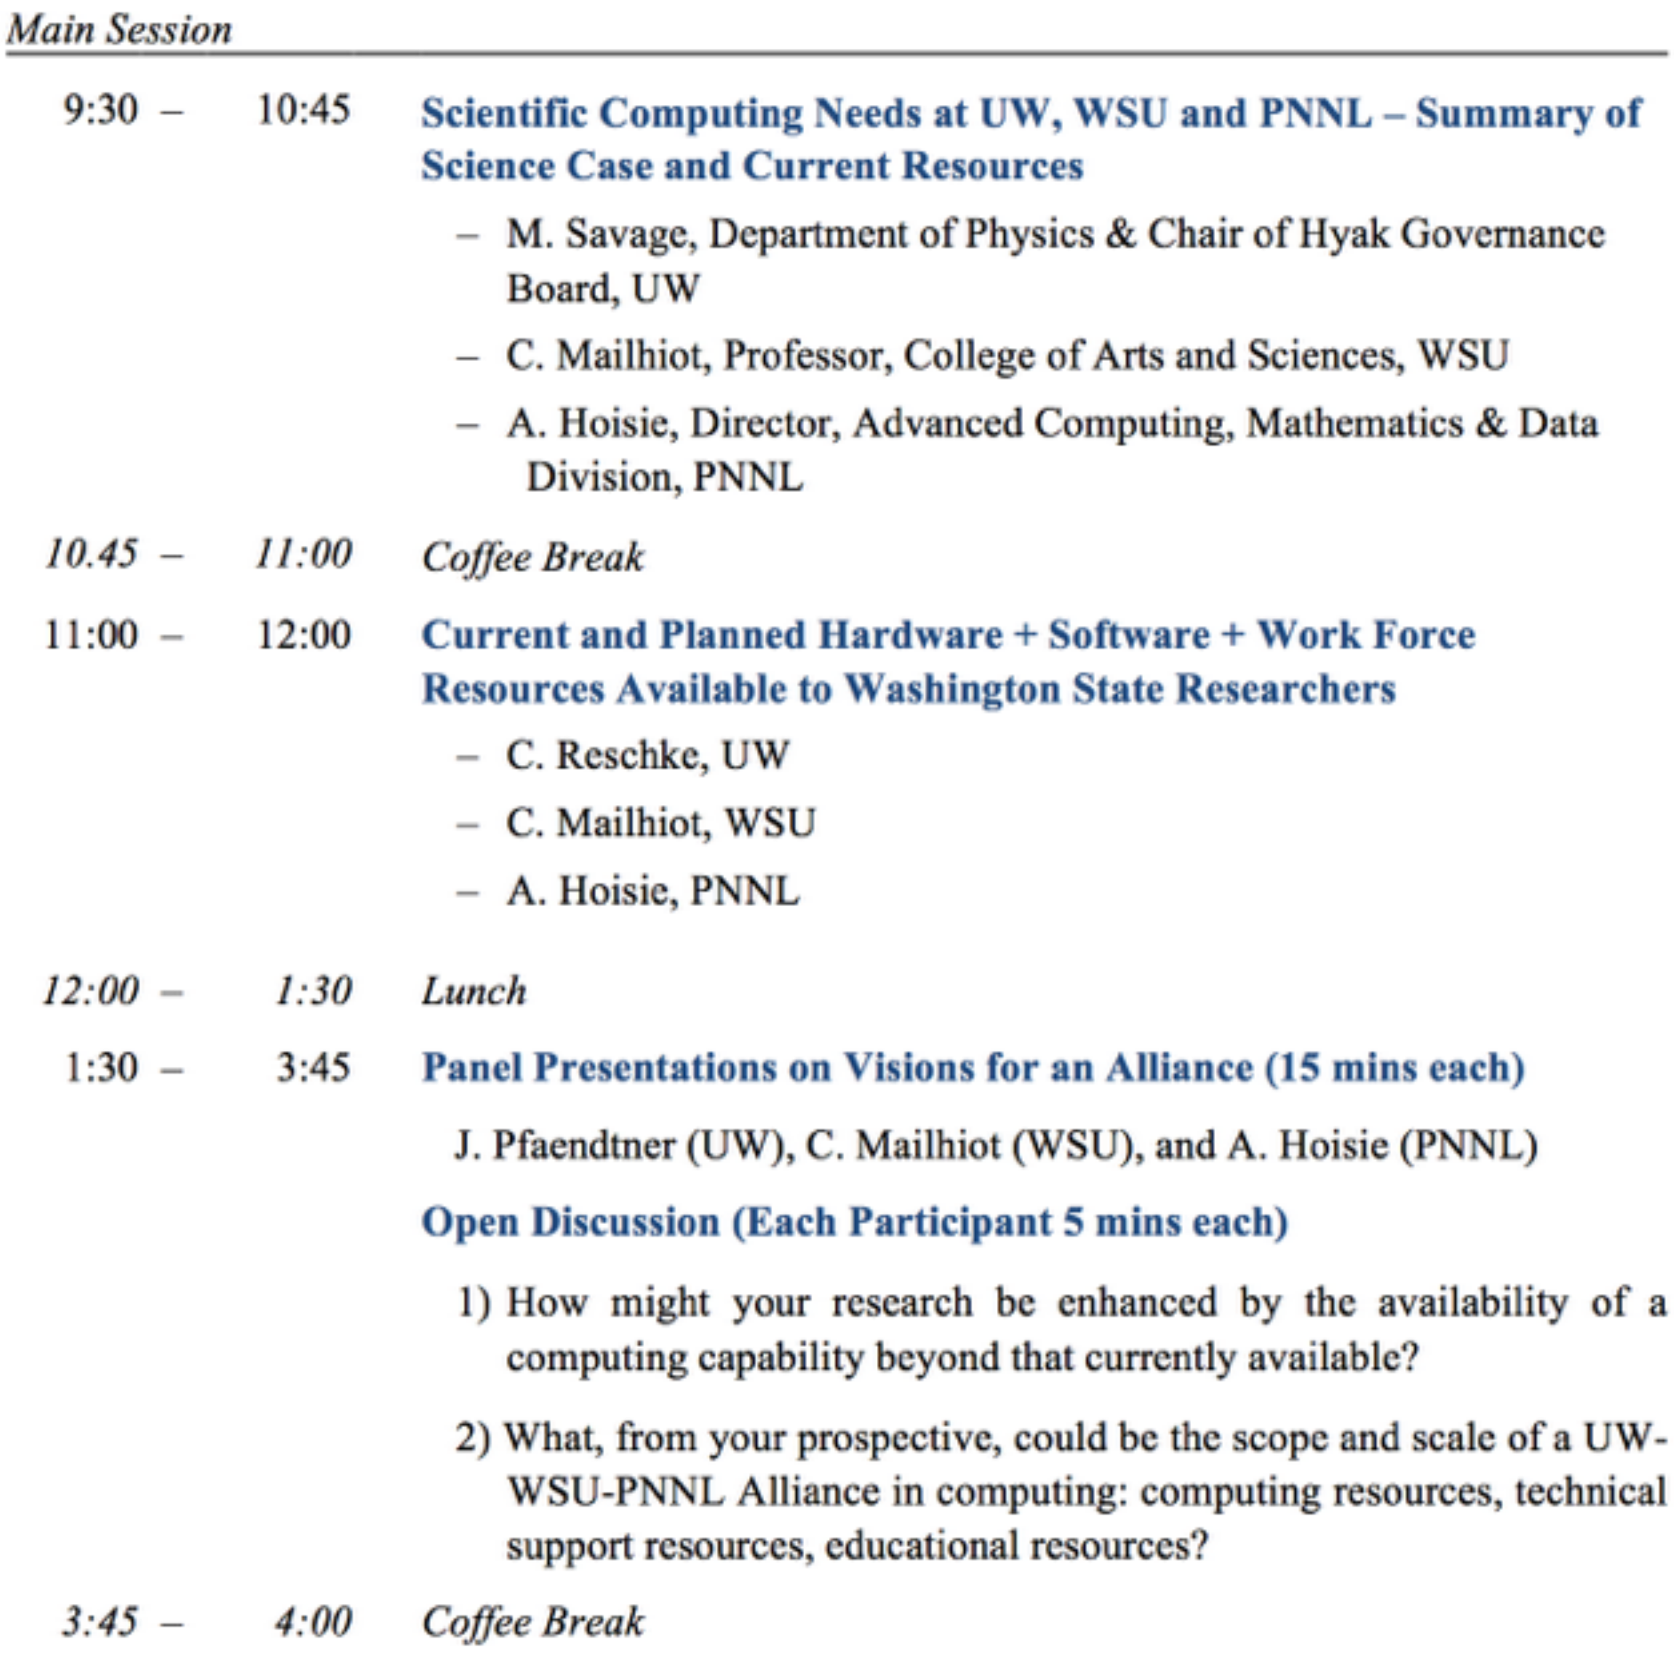

Main Session

9:30 – 10:45 **Scientific Computing Needs at UW, WSU and PNNL – Summary of Science Case and Current Resources**

- M. Savage, Department of Physics & Chair of Hyak Governance Board, UW
- C. Mailhiot, Professor, College of Arts and Sciences, WSU
- A. Hoisie, Director, Advanced Computing, Mathematics & Data Division, PNNL

10.45 – 11:00 Coffee Break

11:00 – 12:00 **Current and Planned Hardware + Software + Work Force Resources Available to Washington State Researchers**

- C. Reschke, UW
- C. Mailhiot, WSU
- A. Hoisie, PNNL

12:00 – 1:30 Lunch

1:30 – 3:45 **Panel Presentations on Visions for an Alliance (15 mins each)**

J. Pfaendtner (UW), C. Mailhiot (WSU), and A. Hoisie (PNNL)

Open Discussion (Each Participant 5 mins each)

1) How might your research be enhanced by the availability of a computing capability beyond that currently available?
2) What, from your prospective, could be the scope and scale of a UW-WSU-PNNL Alliance in computing: computing resources, technical support resources, educational resources?

3:45 – 4:00 Coffee Break

Savage - UW
Mailhiot - WSU
Hoisie - PNNL

Hoisie - PNNL
Reschke - UW
Mailhiot - WSU

Mailhiot - WSU
Hoisie - PNNL
Pfaendtner - UW