BLUE WATERS

SUSTAINED PETASCALE COMPUTING

Early Application Improvement Strategies and Experiences for Blue Waters

William Kramer
National Center for Supercomputing Applications,
University of Illinois at Urbana-Champaign

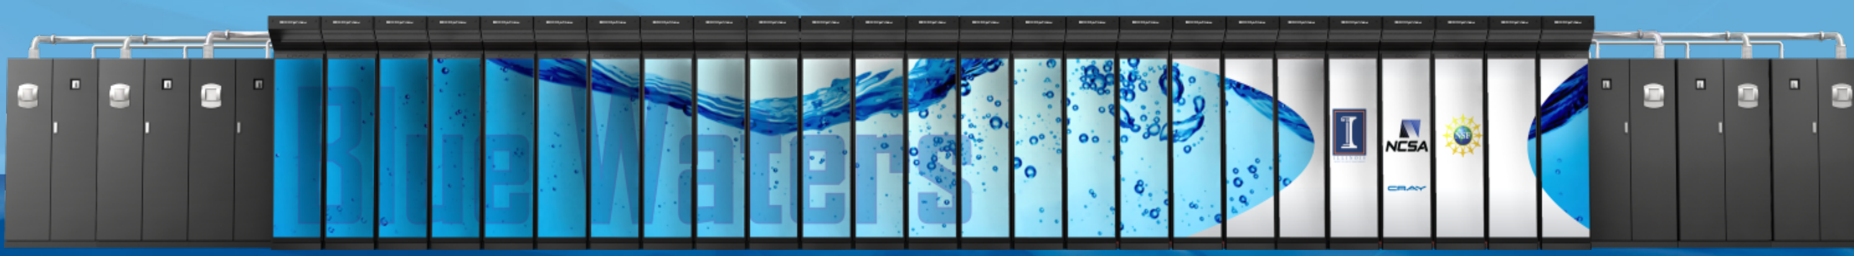

NCSA NSF GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION CRAY

Blue Waters Computing System

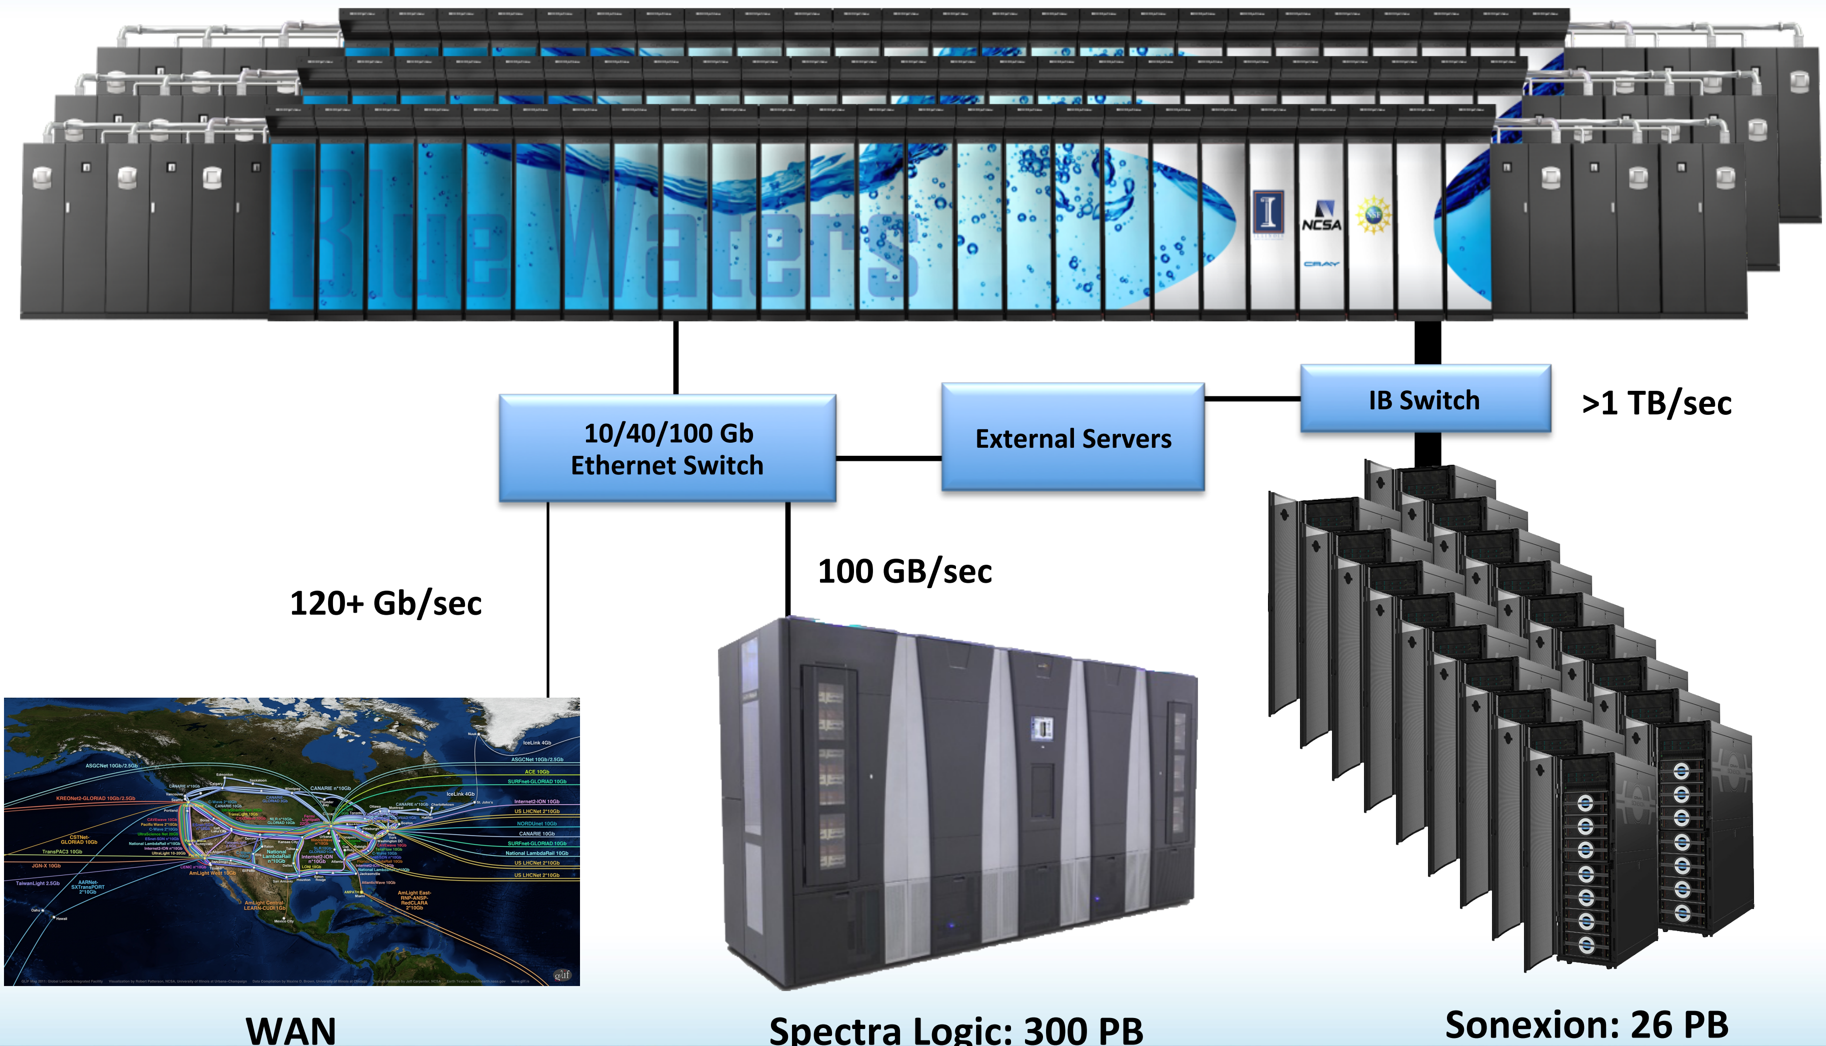

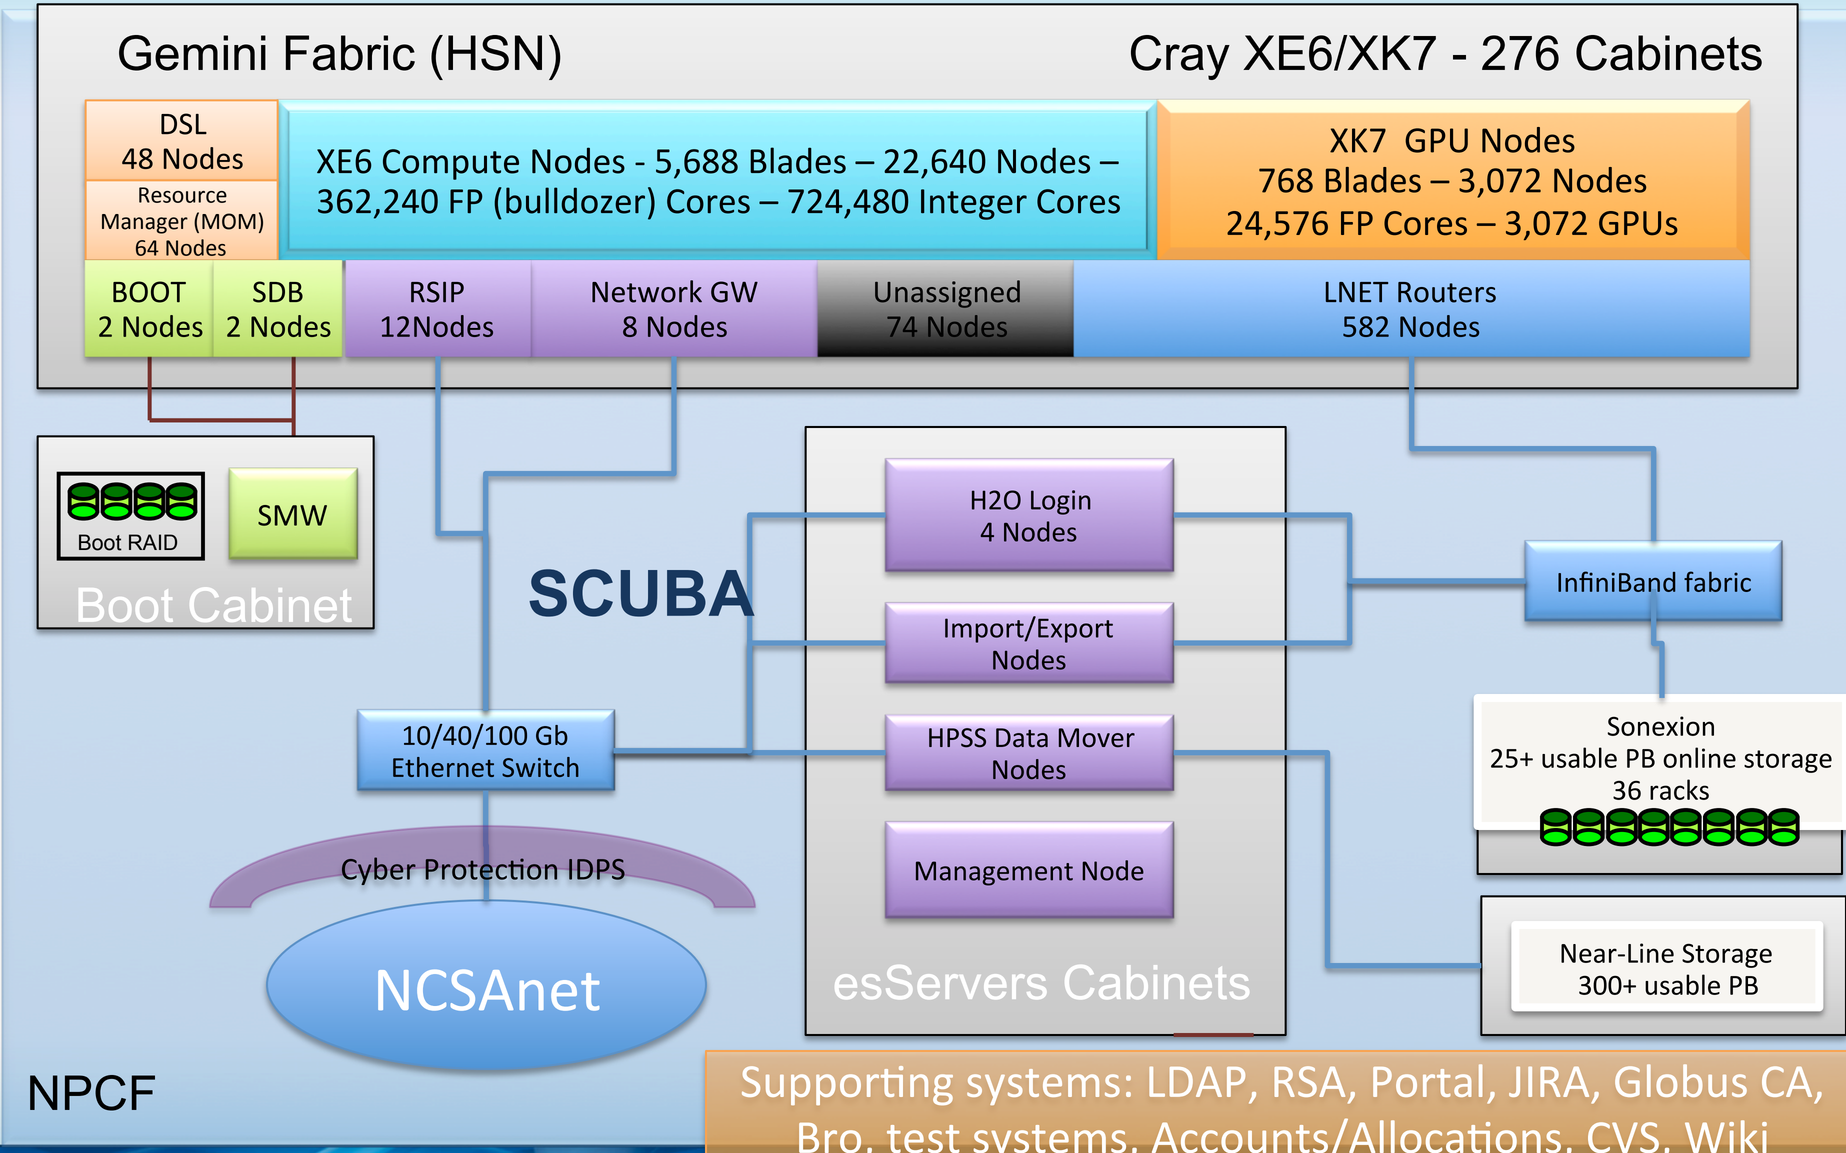
Gemini Fabric (HSN)
Cray XE6/XK7 - 276 Cabinets
DSL
48 Nodes
Resource Manager (MOM)
64 Nodes
XE6 Compute Nodes - 5,688 Blades – 22,640 Nodes – 362,240 FP (bulldozer) Cores – 724,480 Integer Cores
XK7 GPU Nodes
768 Blades – 3,072 Nodes
24,576 FP Cores – 3,072 GPUs
BOOT
2 Nodes
SDB
2 Nodes
RSIP
12Nodes
Network GW
8 Nodes
Unassigned
74 Nodes
LNET Routers
582 Nodes
Boot RAID
SMW
Boot Cabinet
SCUBA
H2O Login
4 Nodes
Import/Export
Nodes
HPSS Data Mover
Nodes
Management Node
esServers Cabinets
InfiniBand fabric
Sonexion
25+ usable PB online storage
36 racks
10/40/100 Gb
Ethernet Switch
Cyber Protection IDPS
NCSAnet
Near-Line Storage
300+ usable PB
NPCF
Supporting systems: LDAP, RSA, Portal, JIRA, Globus CA, Bro, test systems, Accounts/Allocations, CVS, Wiki

SCUBA – Storage Configuration for User Best Access

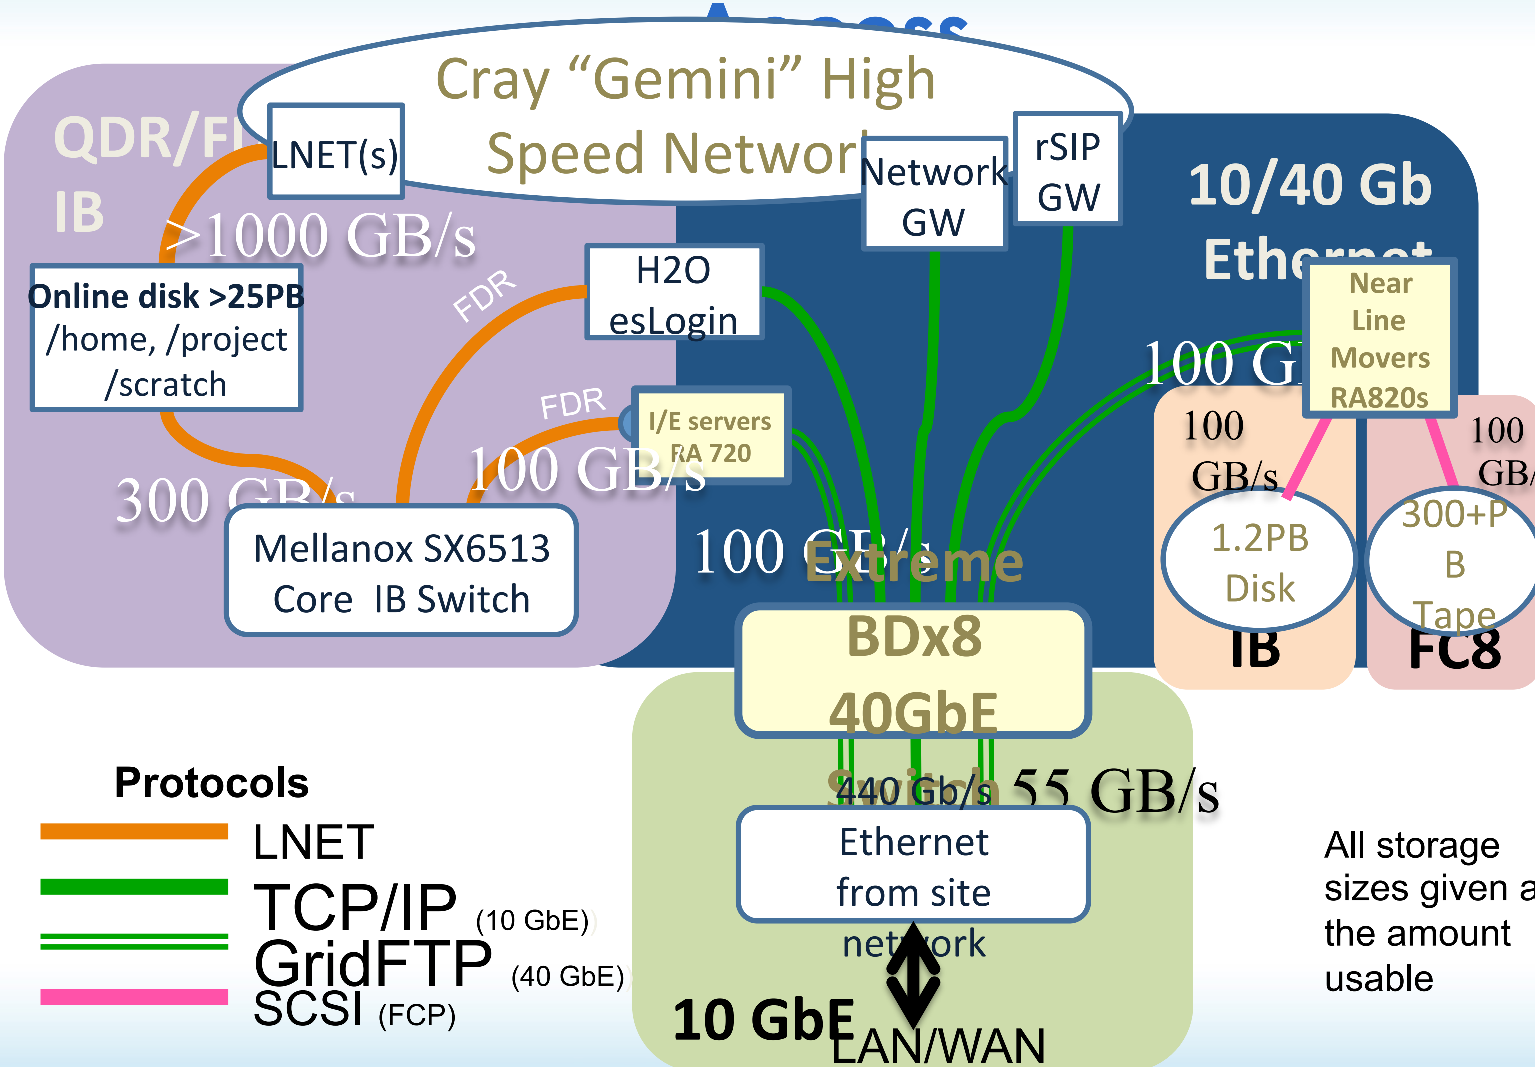

All storage sizes given as the amount usable

BLUE WATERS SCIENCE TEAMS

Expected Types and Number of Science Teams in Steady State

Type of Allocation	Percent of Allocation	Expected Number of Science Teams	Average Users per Science Team	Total Expected Users	Estimated Percent of Active Users per Month	Expected Active Users per Month
NSF PRAC	80	30-36	5-10	150-360	50%	75-180
NCSA industrial affiliates	5-7	2-4	3-5	6-20	50%	3-10
Great Lakes Consortium for Petascale Computation	2	10-20	3-10	30-400	25%	8-100
Educational efforts	1	20	50	500	10%	50
Illinois Allocations	5-7	10-20	5-10	50-200	40%	20-80
Innovation and exploration	5	5	5-10	25-50	50%	12-25
TOTAL (Max)	100			761-1,530		118-345

- Teams and users will change over time – probably 10-20% a year.
- There is some ability to have assumed numbers change over time, dramatic shifts require reassessment of services and resources

PRAC PI Institutions – June 2012

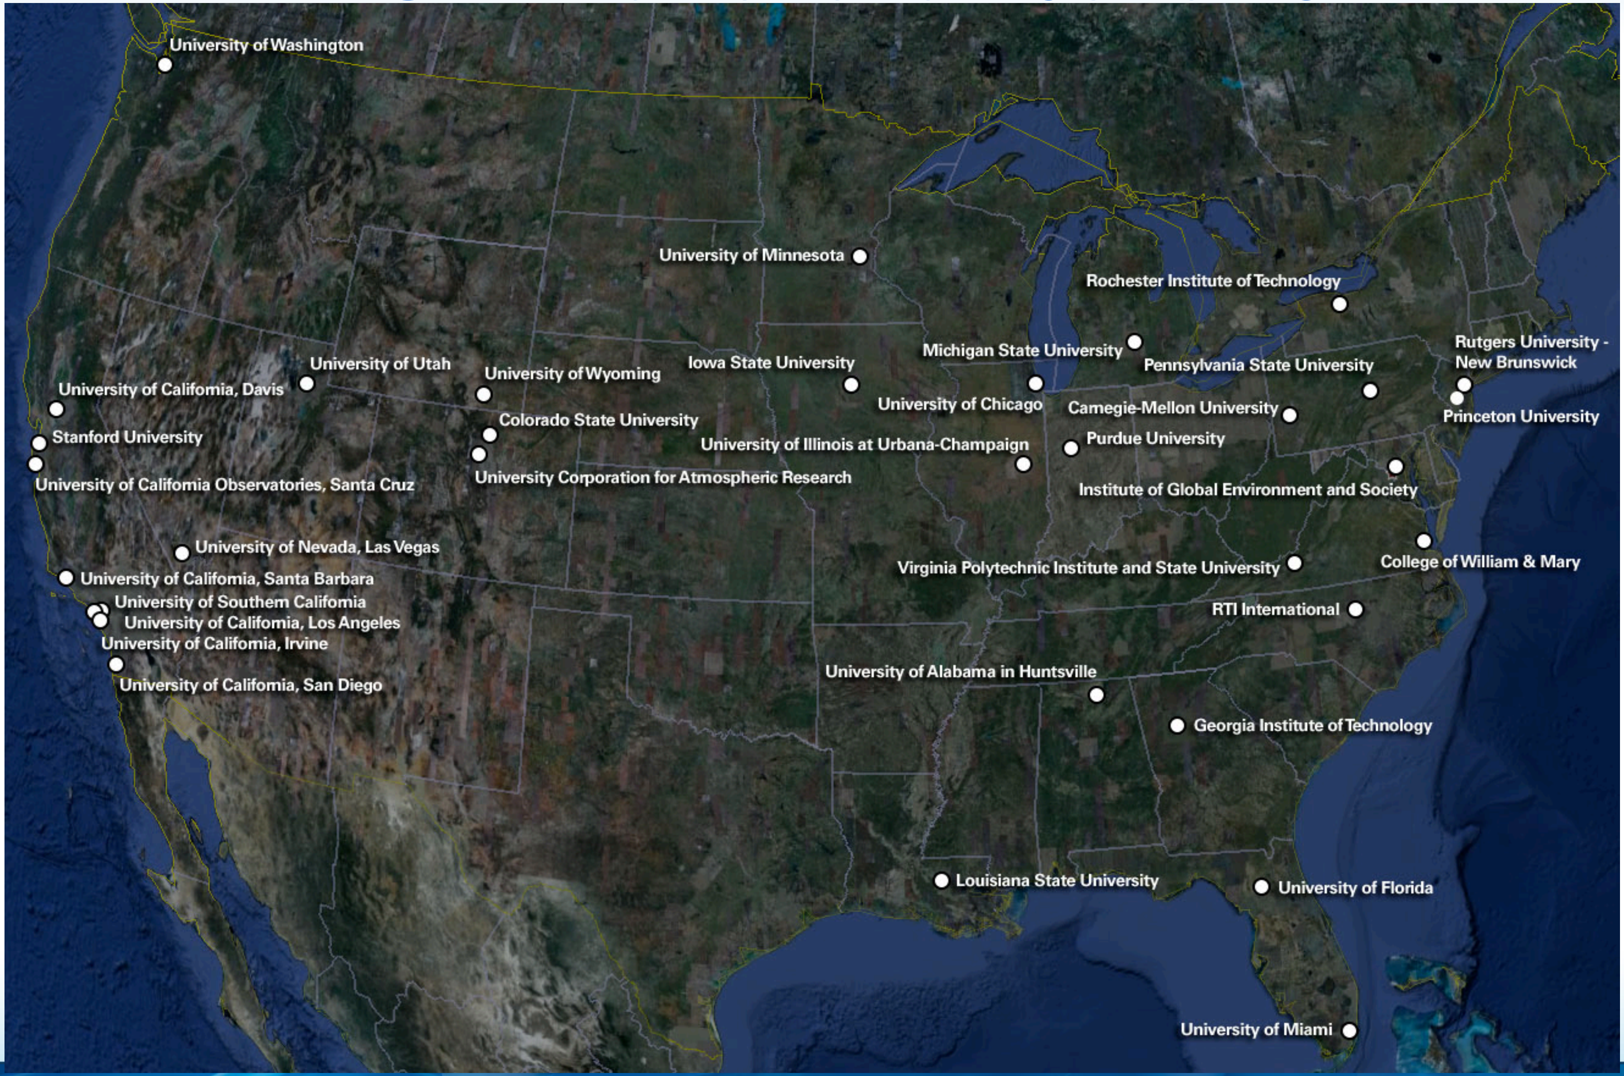

NSF PRAC Science Teams

PI	Award Date	Project Title
Sugar	04/15/2009	Lattice QCD on Blue Waters
Bartlett	04/15/2009	Super instruction architecture for petascale computing
Nagamine	04/15/2009	Peta-Cosmology: galaxy formation and virtual astronomy
Bissett	05/01/2009	Simulation of contagion on very large social networks with Blue Waters
O’Shea	05/01/2009	Formation of the First Galaxies: Predictions for the Next Generation of Observatories
Schulten	05/15/2009	The computational microscope
Stan	09/01/2009	Testing hypotheses about climate prediction at unprecedented resolutions on the NSF Blue Waters system
Campanelli	09/15/2009	Computational relativity and gravitation at petascale: Simulating and visualizing astrophysically realistic compact binaries
Yeung	09/15/2009	Petascale computations for complex turbulent flows
Schnetter	09/15/2009	Enabling science at the petascale: From binary systems and stellar core collapse To gamma-ray bursts
Woodward	10/01/2009	Petascale simulation of turbulent stellar hydrodynamics
Tagkopoulos	10/01/2009	Petascale simulations of Complex Biological Behavior in Fluctuating Environments
Wilhelmson	10/01/2009	Understanding tornadoes and their parent supercells through ultra-high resolution simulation/analysis
Wang	10/01/2009	Enabling large-scale, high-resolution, and real-time earthquake simulations on petascale parallel computers
Jordan	10/01/2009	Petascale research in earthquake system science on Blue Waters
Zhang	10/01/2009	Breakthrough peta-scale quantum Monte Carlo calculations
Haule	10/01/2009	Electronic properties of strongly correlated systems using petascale computing
Lamm	10/01/2009	Computational chemistry at the petascale

NSF PRAC Science Teams (cont)

PI	Award Date	Project Title
Karimabadi	11/01/2010	Enabling Breakthrough Kinetic Simulations of the Magnetosphere via Petascale Computing
Mori	01/15/2011	Petascale plasma physics simulations using PIC codes
Voth	02/01/2011	Petascale multiscale simulations of biomolecular systems
Woosley	02/01/2011	Type Ia supernovae
Cheatham	02/01/2011	Hierarchical molecular dynamics sampling for assessing pathways and free energies of RNA catalysis, ligand binding, and conformational change
Wuebbles	04/15/2011	Using petascale computing capabilities to address climate change uncertainties
Gropp	06/01/2011	System software for scalable applications
Klimeck	09/15/2011	Accelerating nano-scale transistor innovation
Pande	09/15/2011	Simulating vesicle fusion on Blue Waters
Elghobashi	05/18/2012	Direct Numerical Simulation of Fully Resolved Vaporizing Droplets in a Turbulent Flow
Quinn	05/18/2012	Evolutions of the Small Galaxy Populations From High Redshift to the Present
Wood/Reed	06/12/2012	Collaborative Research: Petascale Design and Management of Satellite Assets to Advance Space Based Earth Science
Pogorelov	0613/2012	Modeling Heliophysics and Astrophysics Phenomena with a Multi-Scale Fluid Kinetic Simulation Suite
Bernholc	07/15/2012	Petascale quantum simulations of nano systems and biomolucles
Stein	08/01/2012	Ab Initio Models of Solar Activity

Science Area	Number of Teams	Codes	Struct Grids	Unstruct Grids	Dense Matrix	Sparse Matrix	N-Body	Monte Carlo	FFT	PIC	Significant I/O
Climate and Weather	3	CESM, GCRM, CM1, HOMME	X	X		X		X			X
Plasmas/Magnetosphere	2	H3D(M), OSIRIS, Magtail/UPIC	X				X		X		X
Stellar Atmospheres and Supernovae	5	PPM, MAESTRO, CASTRO, SEDONA, ChaNGa, MS-FLUKSS	X			X	X	X		X	X
Cosmology	2	Enzo, pGADGET	X			X	X				
Combustion/Turbulence	2	PSDNS, DISTUF	X						X		
General Relativity	2	Cactus, Harm3D, LazEV	X			X					
Molecular Dynamics	4	AMBER, Gromacs, NAMD, LAMMPS			X		X		X		
Quantum Chemistry	2	SIAL, GAMESS, NWChem			X	X	X	X			X
Material Science	3	NEMOS, OMEN, GW, QMCPACK			X	X	X	X			
Earthquakes/Seismology	2	AWP-ODC, HERCULES, PLSQR, SPECFEM3D	X	X			X				X
Quantum Chromo Dynamics	1	Chroma, MILD, USQCD	X		X	X	X		X		
Social Networks	1	EPISIMDEMICS									
Evolution	1	Eve									
Engineering/System of Systems	1	GRIPS,Revisit						X			
Computer Science	1			X	X	X			X		X

Using GPUs

- While many science teams have GPUS activities
  - Few use GPUs on a production basis
  - Few use GPUs at scale
  - Few have their entire application using GPUs
  - Many still need to move to GPUs
- But
  - ***Nearly 1/3 of PRAC projects have active GPU efforts, including***
  - AMBER • LAMMPS • UCQCD/MILC • GAMESS • NAMD • QMCPACK • PLSQR/SPECFEM3D
  - Others are investigating use of GPUs (e.g., Cactus, PPM, AWP-ODC)

EARLY SCIENCE

ESS Exceeded Expectations

- System ran well for 13 weeks rather than 8
- We were able to provide more than 6.5 million node hours of computational to science teams (**>100M integer core hours**)
- We were able to support 10 science teams doing major work
- We were able to open the system to all interested PRAC teams for two weeks so they could port, tune and try the system
- We were able to verify major software components
- We were able to use ESS for benchmark performance improvements

ESS Use Purpose

- The primary purposes of science team use of the early science time were:
    - To provide a substantial new, interim resource to certain science and engineering teams who have the potential to accomplish a significant science result in an interim period of time.
    - To help the Blue Waters team test and evaluate the early system and prepare for full system testing.

ESS Selection Criteria

- Science team's **potential to accomplish a significant science result** with approximately one (1) aggregate peak PetaFlop-week of computational time (not dedicated, but shared) spread over a 4 to 8 week time period;
- **Computational readiness** of the team's code, including the team's ability to **predict performance via modeling** or other methods;
- Team's **experience running** on large scale systems **well before general availability**;
- Science team's **flexibility** in being able to use the early system between early March 2012 and at least early to mid May 2012 and their willingness to **engage in frequent communication with the Blue Waters** project on progress and problems;
- Degree to which the science team's research problem, code and institution **represent the diversity of the current PRAC portfolio**;
- Whether the science team **has a preliminary allocation** from NSF on the full Service Blue Waters system and is **current on all annual project reports**;
- Whether the codes the volunteering team **use overlap with the application benchmark codes** NCSA is using to test and accept the Blue Waters system (a list of those major application codes is attached).
    - Will probably add 1 or 2 limited use teams to help with benchmarks (~10-50,000 node hours (0.2-1M SUs))

Blue Waters Early Science System

- BW-ESS Configuration
  - 1.4+ PFs (peak)
  - 48 cabinets, 4,512 XE6 compute nodes, 96 service nodes ~18%
  - 2 PBs Sonexion Lustre storage appliance ~5%

- Current Projects
  - **Biomolecular Physics**—K. Schulten, University of Illinois at Urbana-Champaign
  - **Cosmology**—B. O'Shea, Michigan State University
  - **Climate Change**—D. Wuebbles, University of Illinois at Urbana-Champaign
  - **Lattice QCD**—R. Sugar, University of California, Santa Barbara
  - **Plasma Physics**—H. Karimabadi, University of California, San Diego
  - **Supernovae**—S. Woosley, University of California Observatories
  - **Severe Weather** –R Wilhelmson, University of Illinois
  - **High Resolution/Fidelity Climate** – C Stan, Center for Ocean-Land-Atmospheric Studies (COLA)
  - **Complex Turbulence** – P.K. Yeung, Georgia Tech
  - **Turbulent Stellar Hydrodynamics** – P Woodward, University of Minnesota

Basic ESS Usage

Group	Total #Jobs	#XK Jobs	Total Node*hrs
Bisset	382	0	8,558.9
Campanelli	580	0	30,650.8
Karimabadi	280	0	1,147,696.3
Lamm	38	0	9,507.3
O'Shea	729	0	427,700.8
Nagamine	3	0	16.1
Pande	120	0	5,403.7
Schulten	4,459	281	1,540,632.4
Stan	63	0	821.4
Sugar	3,744	0	1,540,630.7
Wilhelmson	151	0	180,267.2
Woodward	96	0	309,747.6
Woosley	579	0	834,147.6
Wuebbles	372	3	482,426.6
Yeung	287	0	52,678.6
bw_staff	5,576	160	723,291.8
vendor_cray	1,117	14	272,716.7

Access by 15 PRAC teams

Total Node-hours: 7,566,894.5

Total Node-hours for science teams: 6,570,866 (86% or wall clock)

>1,000,000 integer core hours for science

The Computational Microscope: NAMD

- "Not in our wildest dreams could we have imagined the greatness" of Blue Waters

1. **Simulated flexibility of ribosome trigger factor complex at full length** and obtained better starting configuration of trigger factor model (simulated to 80ns)
2. **100ns simulation of cylindrical HIV 'capsule' of CA proteins** revealed it is stabilized by hydrophobic interactions between CA hexamers; maturation involves detailed remodeling rather than disassembly/re-assembly of CA lattice, as had been proposed.
    - 200ns simulation of CA pentamer surrounded by CA hexamers suggested interfaces in hexamer-hexamer and hexamer-pentamer pairings involve different patterns of interactions
3. **Simulated photosynthetic membrane** of a chromatophore in bacterium Rps. photometricum for 20 ns -- simulation of a few hundred nanoseconds will be needed

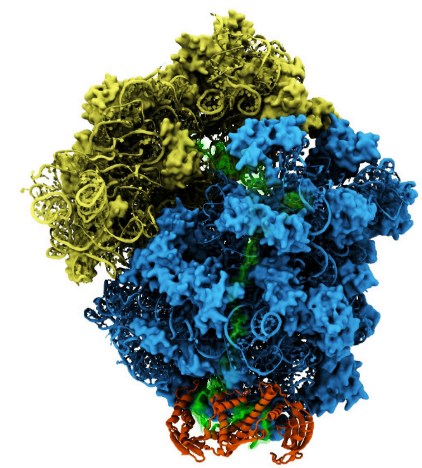

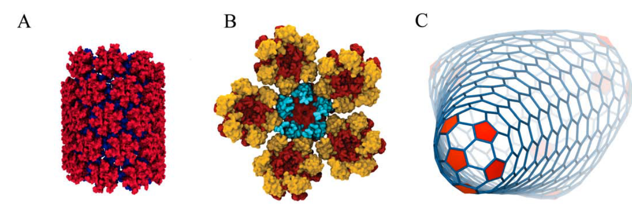

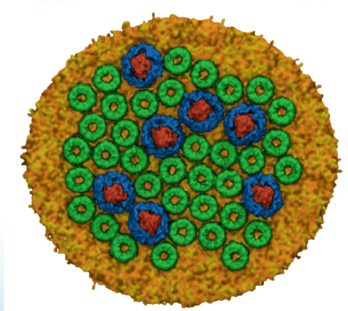

PI: Klaus Schulten, University of Illinois at Urbana-Champaign

Climate Change Uncertainties

- Ported and validated Community Earth System Model
- Performed initial tests of CAR system in CESM using low-resolution version
- Conducted **one-year test of CESM with 0.25° finite-volume dynamical core**. Validated by NCAR
- ESS time helped identify issues with CAM5-PROG at high resolution and I/O bug
- Results from the new prognostic aerosol run (CAM5-PROG) show significant differences in cloud radiative forcing compared to an earlier run with BAM prescribed aerosols, particularly in high latitudes. **Results clearly demonstrate the value of running CAM5-PROG at high resolution.**

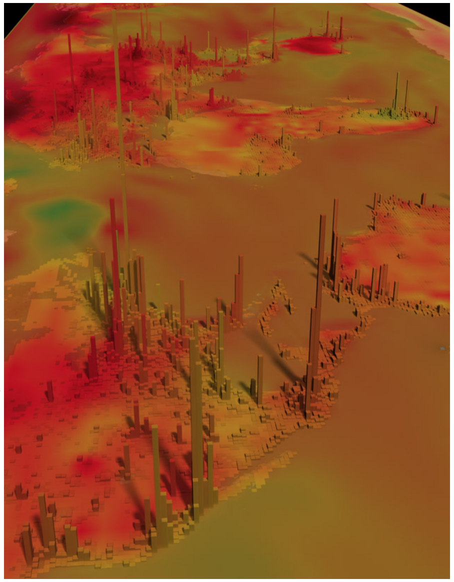

PIs: Donald Wuebbles and Xin-Zhong Liang, University of Illinois at Urbana-Champaign

Turbulent Stellar Hydrodynamics

- Obtained **12.02% of peak**, single precision flops for a very large (72 billion-cell) problem. 312 Tflops/s on ESS was counted
- Tested new inertial confinement fusion w. performance enhancements resulting from collaboration with Cray

PI: Paul Woodward, University of Minnesota-Twin Cities

Kinetic Simulations of the Magnetosphere

- Objective: understand 3D evolution of force-free current layers using recent theory describing tearing modes (plasma instability that produces magnetic reconnection while giving rise to topological changes in magnetic field)
- Performed 3 3D simulations with varied rotation in the magnetic field across the initial layer
    - 2 runs: 2048 x 2048 x 1024 = 4.3 billion cells / 1 trillion particles / run on 65536 cores
    - 3rd run: 2048 x 2048 x 1536 = 6.4 billion cells / 1.5 trillion particles / run on 98304 cores
    - Each run generated ~25-30 TB of grid-based data, and another 32TB of particle data
    - 8.3 million particle pushes per second for each core. Corresponds to ~0.2 petaflops for large run
- **New results are dramatically different than previous 2D simulations**
- Hope to have paper completed by end of summer
- “Both the stability and performance of the ESS were outstanding”

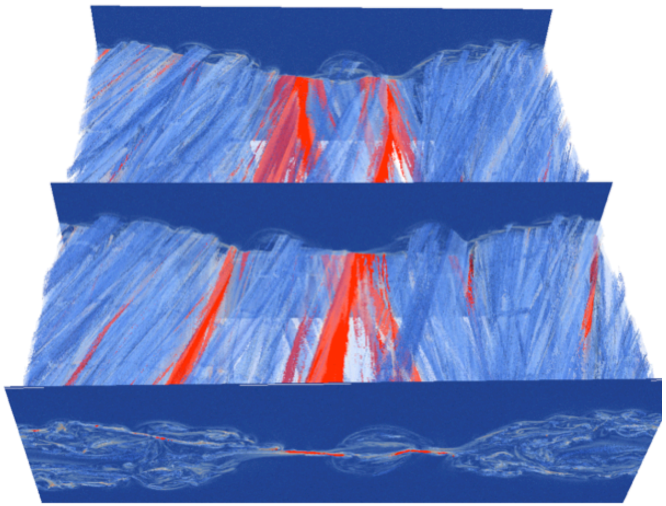

PIs: Homayoun Karimabadi, Kevin Quest, Amitava Majumdar, University of California-San Diego

Simulation of Contagion: EpiSimdemics

- Measured scaling w. **2 datasets: Michigan (pop. 9.M) and North Carolina/Tennessee/Texas (32.7M)**
- **Should efficiently run on 20k-30k cores**
- On full BW, plan to simulate spread of influenza across U.S., comparing intervention combinations. Problem has not been simulated with this level of detail and at this scale.

PIs: Keith Bisset, Virginia Tech; Shawn Brown, Carnegie-Mellon University; Douglas Roberts, Research Triangle Institute

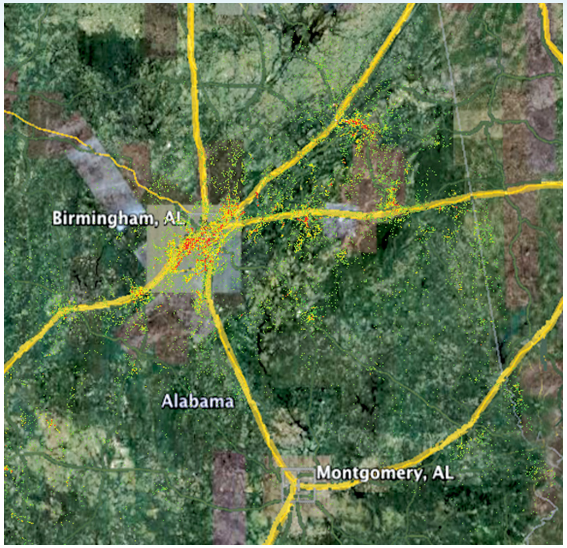

Modeling Type 1a Supernovae

- **Off-center ignition of Type 1a Supernovae, 1 second duration**
- Codes: MAESTRO and CASTRO
- Used 68 million core hours
- Produced 45 TBs of data

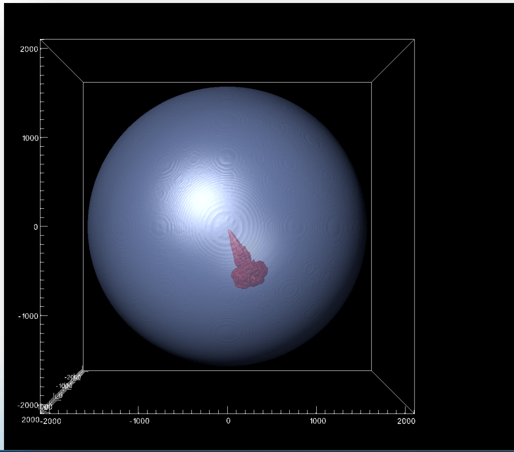

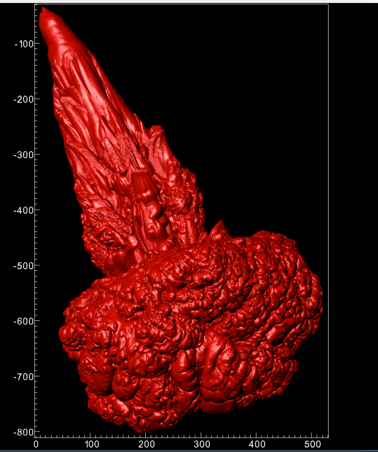

PIs: PI: S. Woosley, University of California Observatories

Lattice Gauge Theory on Blue Waters

- Calculation of the spectroscopy of charmonium, the positronium-like states of a charm quark and an anticharm quark.
- Also spent a limited amount of time preparing for the two large projects we hope to run on Blue Waters when the full system is available for production work.
- Able to reproduce these **mass splittings at a record precision of a couple MeV** is an impressive test of our methodology,
- Gives us confidence in our ability to make predictions of other levels where the experimental values are not known or the classification of the states is not understood.

PI: Robert Sugar, U C Santa Barbara

Simulations of Homogeneous Turbulence

- Performed **40963 simulation, run on 16k or 32k cores**, 20 sets of restart files were transferred to NICS
- Code performance obtained on BW-ESS was generally slightly better than on Jaguarpf at NCCS.
- Performance improvement using Co-Array Fortran (CAF) in place of mpi alltoall(v) with the help of Cray Team.
- In decent shape in terms of the viability of running an Re_{τ} ~ 5000 calculation (the next-generation channel flow target) on the full-scale Blue Waters.

PI: P. K. Yeung, Georgia Institute of Technology

Formations of galaxy complexes at high redshift

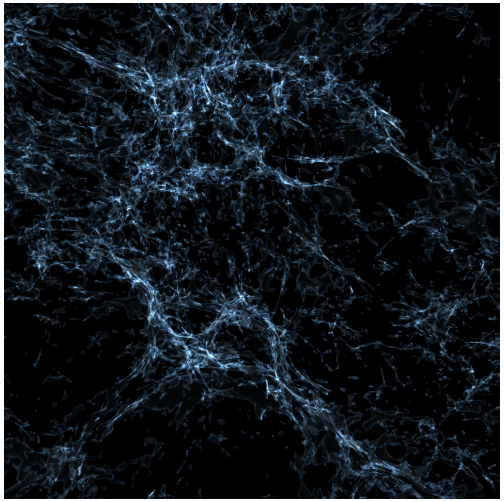

- Performed **8 cosmological simulations to understand how galaxies in the early Universe** (the first billion years or so) grow and evolve, in several statistically---dissimilar environments
- **Larger than any other AMR cosmological simulations ever done**.
- BW's high memory per core was crucial to the success.
- Happy with the performance of the system except I/O subsystem.
- System performed brilliantly and technical support was prompt and of exceptional quality.
- One of the Enzo simulations http://galactica.pa.msu.edu/~bwoshea/data/BlueWaters/ESSmovies/

PI: Brian O'Shea, Michigan State University, Co-PI: Michael Norman, UC San Diego/SDSC

Computational Chemistry at the Petascale

- Ported and obtained good initial timings for the Gamess code
- An energy + gradient calculation on a cluster of 512 water molecules was run using FMO at the MP2 level of theory using the aug-cc-pVDZ basis set.
- Timings for energy + gradient without the fully analytic gradient : **BW calculation on 4096 ~11.7 minutes, while the analogous BG/P calculation took 28.9 minutes on 8192 cores.**

PIs: Monica Lamm, Iowa State University

LESSONS ALREADY

Challenges to Come

- Increasing performance requires dramatic increases in parallelism that then generates complexity challenges for science and engineering teams
- Blue Waters is not only an unprecedented resource but also is a pointer to the future challenges for future systems
  - Application Functionality & Performance
    - Increase application efficiency and scaling in the face of limited bandwidth (interconnect, memory, etc.) and other architectural constraints with new methods and enhancements.
    - Enable the effective use of highly parallel heterogeneous (e.g. those with GPU accelerators or many core processors) computational units.
    - Enable application-based topology awareness to more effectively and efficiently utilize limited resources.
  - Application Flexibility
    - Enable the use of heterogeneous systems that have both general purpose CPU and acceleration units by single applications.
    - Enable and explore the use of application-based fault tolerant methods and algorithms to increase the effective and efficient use of resources.
  - Improve the use of advanced storage and data movement methods to increase the efficiency and time to solution of applications

Performance and Scalability

- The issue is fewer applications are able to scale in the face of limited bandwidths. Hence the need to work with science teams and technology providers to
  - Develop better process-to-node mapping using for graph analysis to determine MPI behavior and usage patterns.
  - Topology Awareness in Applications and in Resource Management
  - Improve use of the available bandwidth (MPI implementations, lower level communication, etc.).
  - Consider new algorithmic methods
  - Considering alternative programming models that improve efficiency of calculations

BW Benchmarks Benefit from Topology Awareness

- Original NSF Benchmarks
  - Full Size – **QCD (MILC)**, **Turbulence (PNSDNS)**, Molecular Dynamics (NAMD)
- SPP Application Mix (details and method available)
  - NAMD – molecular dynamics
  - **MILC**, Chroma – Lattice Quantum Chromodynamics
  - VPIC, SPECFEM3D – Geophysical Science
  - **WRF** – Atmospheric Science
  - **PPM** – Astrophysics
  - NWCHEM, GAMESS – Computational Chemistry
  - QMCPACK – Materials Science
- As we continue to work with the system, more applications will be identified that benefit from topology awareness
- **Bold** = codes known to benefit from topology optimization

Performance and Scalability

- Use of heterogeneous computational units
  - While more than ½ of the science has some GPU based investigations, only a few are using GPUs in production science
    - Many applications are GPUized only in a very limited way
    - Few are using GPUs at scale (more GPU resources are relatively small scale with limited networks)
  - Help the science teams to make more effective use of GPUs consists of two major components.
    - Introduce compiler and library capabilities into the science team workflow to significantly reduce the programming effort and impact on code maintainability.
      - OpenACC support is the major path to more general acceptance
    - Load balancing at scale
- Storage Productivity
  - Interface with improved libraries and middle ware
  - Modeling of I/O
  - On-line and Near-line transparent interfaces

Application Flexibility

- Using both XE and XK nodes in single applications
  - For multi-physics applications that provide a natural decomposition into modules is to deploy the most appropriate module(s) different computational units.
  - For applications use the Charm++ adaptive runtime system, heterogeneity can be handled without significant changes to the application itself.
  - Some applications naturally involve assigning multiple blocks to individual processors include multiblock codes (typically in fluid dynamics), and the codes based on structured adaptive mesh refinement.
  - The application-level load balancing algorithms can be modified to deal with the performance heterogeneity created by the mix of nodes.
- Malleability
  - Understanding topology given and maximizing effectiveness
  - Being able to express desired topology based on algorithms
  - Mid ware support

Flexibility - Application Based Resiliency

- Application Based Resiliency
  - Multiple layers of Software and Hardware have to coordinate information and reaction
  - Analysis and understanding is needed before action
  - Correct and actionable messages need to flow up and down the stack to the applications so they can take the proper action with correct information
  - Application Situational Awareness - need to understand circumstances and take action
  - Flexible resource provisioning needed in real time
  - Interaction with other constraints so sub-optimization does not adversely impact overall system optimization

SUSTAINED PERFORMANCE MEANS TIME TO SOLUTION

Mis-leading Lists

Following the format of a popular U.S. talk show host uses for "top 10" lists, the issues with the TOP500 list are summarized as:

Number 10: The Linpack benchmark serves only one of the four purposes of a good benchmark.

Number 9: The TOP500 list disenfranchises many important application areas.

Number 8: There is no relationship between the TOP500 ranking and user productivity of system.

Number 7: The Linpack performance test is dominated by single-core peak performance.

Number 6: The TOP500 metric has not kept up with changing algorithmic methods.

Number 5: The TOP500 measure takes too long to run and does not represent strong scaling.

Number 4: The TOP500 is dominated by who has the most money to spend–not what system is the best.

Number 3: The TOP500 provides little historical value.

Number 2: The TOP500 encourages organizations to make poor choices.

Number 1: The TOP500 gives no indication of the cost of value of a system.

……

Blue Waters Is Pursuing a Quantative Method for "Sustained"

- Sustained Performance is accomplishing an amount of work in a elapsed time.
  - It is not a hardware rate
  - It is not the work needed to scale
  - It is reflection of the work needed completing meaningful problems

The Sustained Petascale Performance (SPP) Method

- Establish a set of application codes that reflect the intended work the system will do
  - Can be any number of tests as long as they have a common measure of the amount of work
- A test consists of a complete code and a problem set reflecting the science teams' intentions
- Establish the reference amount work (ops, atoms, years simulated, etc.) the problem needs to do for a fixed concurrency
- Time each test takes to execute
  - Concurrency and/or optimization can be fixed and/or varied as desired
- Determine the amount of work done for a given "schedulable unit" (node, socket, core, task, thread, interface, etc.)
  - Work = Total work (operations) /total time/number of scalable units
  - Work per unit= Total work/number of scalable units used for the test
- Composite the work per schedulable unit for all tests
  - Composite functions based on circumstances and test selection criteria
  - Can be weighed or not as desired
  - BW is using the Geometric mean – lowest of all means and reduces impact of outliers
- Determine the SPP of a system by multiplying the composite work per schedulable unit by the number of schedulable units in the system
- Determine the *Sustained Petascale Performance*

$$SSP_{s,k} = \sum_{\alpha=1}^{A_{s,k}} \left(\Phi \left(W \; , P_{s,k,\alpha} \right) * N_{s,k,\alpha} \right)$$

BW Sustained Petascale Performance Measures

- Original NSF Benchmarks
  - Full Size – QCD (MILC), Turbulence (PNSDNS), Molecular Dynamics (NAMD)
  - Modest Size – MILC, Paratec, WRF
- SPP – is a time to solution metric that is using the planned applications on representative parts of the Science team problems
  - Represents end to end problem run including I/O, pre and post phases, etc.
  - Coverage for science areas, algorthmic methods, scale
- SPP Application Mix (details and method available)
  - NAMD – molecular dynamics
  - MILC, Chroma – Lattice Quantum Chromodynamics
  - VPIC, SPECFEM3D – Geophysical Science
  - WRF – Atmospheric Science
  - PPM – Astrophysics
  - NWCHEM, GAMESS – Computational Chemistry
  - QMCPACK – Materials Science
- At least three SPP benchmarks run at full scale
- XK nodes have to add 15% more SPP

Sustained Petascale Performance Applications

- In addition to all of the NSF RPF Petascale benchmarks, NCSA is using the SPP to assess sustained performance
  - NAMD – molecular dynamics
  - MILC – lattice QCD
  - PPM – turbulent stellar atmospheres
  - QMCPACK – materials science
  - VPIC – Earth's magnetosphere and plasma physics
  - WRF – weather and climate
  - SPECFEM3D– geodynamics
  - NWChem– chemistry

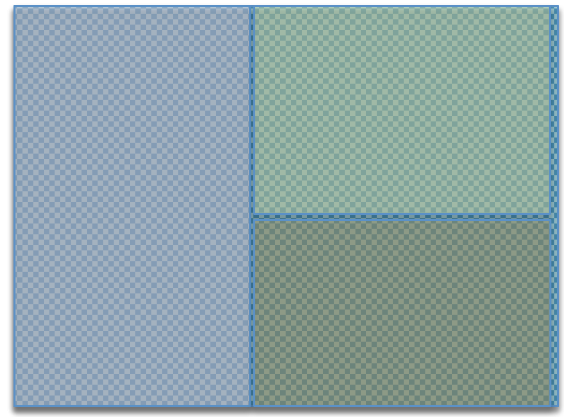

- The input, problem sizes, included physics, and I/O performed by each benchmark will be comparable to the simulations proposed by the corresponding science team for scientific discovery.
- Each benchmark will be sized to use one-fifth to one-half of the number of nodes in the full system.
  - Three SPP applications run at full system size
- **GPUs will quantitatively increase the SPP by 15%**

Summary

- Even a modest part of BW is highly useful for science
- Science teams produced notable accomplishments
- Petascale will present significant challenges for performance, flexibility and data
- Sustained Performance – we know it when we experience it – but now we can measure it and document it with SPP

The Blue Waters Team (partially)

Acknowledgements

This work is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (award number OCI 07-25070) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign, its National Center for Supercomputing Applications, Cray, and the Great Lakes Consortium for Petascale Computation.

The work described is achievable through the efforts of the Blue Waters Project.

Individual Contributions From

- Thom Dunning, Marc Snir, Wen-mei Hwu, Bill Gropp
- Cristina Beldica, Brett Bode, Michelle Butler, Greg Bauer, Mike Showerman, John Melchi, Scott Lathrop, Merle Giles
- Sanjay Kale, David Padua
- The Blue Waters Project Team and our partners
- NSF/OCI
- Cray, Inc, AMD, NVIDIA, Xyratex, Adaptive, Allinea