Progress and Plans for FSP ASCR Activities

Presented by

Douglas B. Kothe

Director of Science
Oak Ridge Leadership Computing Facility
National Center for Computational Sciences

FSP Program Advisory Committee Meeting
Princeton, NJ
Mar 25-26, 2010

Outline

- ASCR overview
- Leadership Computing Facilities (LCFs)
  - Now and looking forward during the FSP execution period
- Fusion application examples: those on LCF platforms
- FSP challenges crosscutting ASCR mission and goals
  - Application “readiness”
    - Scalable, efficient algorithms and “Co-Design” teams
  - Mathematical formalisms
  - Data analytics, management, and visualization
  - V&V and UQ
  - Programming models, frameworks, and tools
- Putting together a coherent, cohesive, and executable plan

DOE Office of Science Advanced Scientific Computing Research (ASCR) Program

- ASCR delivers the tools to advance understanding in key DOE mission areas: energy security, nuclear security, scientific discovery and innovation, and environmental responsibility
  - *"to discover, develop, and deploy the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex phenomena important to the DOE … founded to develop the algorithms, computer programs and hardware that advance scientific research"*
- ASCR has three basic thrust areas
  - Research
    - Applied Mathematics, Computer Science, Integrated Network Environments
  - Facilities
    - NERSC, Leadership Computing, ESnet
  - Cross-cutting projects
    - Scientific Discovery Through Advanced Computing (SciDAC), Innovative and Novel Computational Impact on Experiment (INCITE), Multiscale Mathematics Initiative

FSP will directly benefit from ASCR activities – a partnership is essential

DOE/ASCR Strategic Directions*

2010 and beyond

*Dr. William F. Brinkman, Director, Office of Science, ASCAC presentation, August 11, 2009

- Science for National Need – Delivering forefront scientific knowledge and state-of-the-art tools to serve the nation
- Many areas of research require leadership computing power for discovery
- Even with the newly harnessed petascale resources, many critical simulations are limited by computing power
- ASCR's strategic directions for 2010 and beyond continues to be:
  - Advance the state of the art in computational capability
  - Develop tools and methods for harnessing that capability
  - Bring this capability to bear on scientific questions with national need:
    - **Climate Modeling** – half of the runs for the U.S. contribution to the IPPC AR4 were done on ASCR computing systems; the Earth Systems Grid is the primary mechanism for sharing this data; and SciDAC supports the development of next generation codes
    - **Combustion** – advances in computing power move combustion simulations closer to real world conditions and provides new insights into how to improve fuel efficiency and reduce emissions
    - **Bioenergy** – working with DOE Bioenergy Research Centers, simulations of enzymes breaking down cellulose will help make cellulosic ethanol (energy from non-food crops) an economically viable option
    - **Nuclear Energy** – the worlds largest simulation of a reactor core was achieved at the ALCF and will help engineers to enhance safety and reduce waste in next generation reactors
    - **Fusion** – advances in computing power move fusion simulations closer to real world conditions and help engineers improve efficiency and design control systems
    - **Advanced Materials** – including the worlds first petascale application modeling superconductivity

ASCR LCF Resource Capabilities

Increased over 750-fold in last 5 years & fusion science applications have been among the largest and most effective consumers

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes	Scaling applications and system software is the biggest challenge

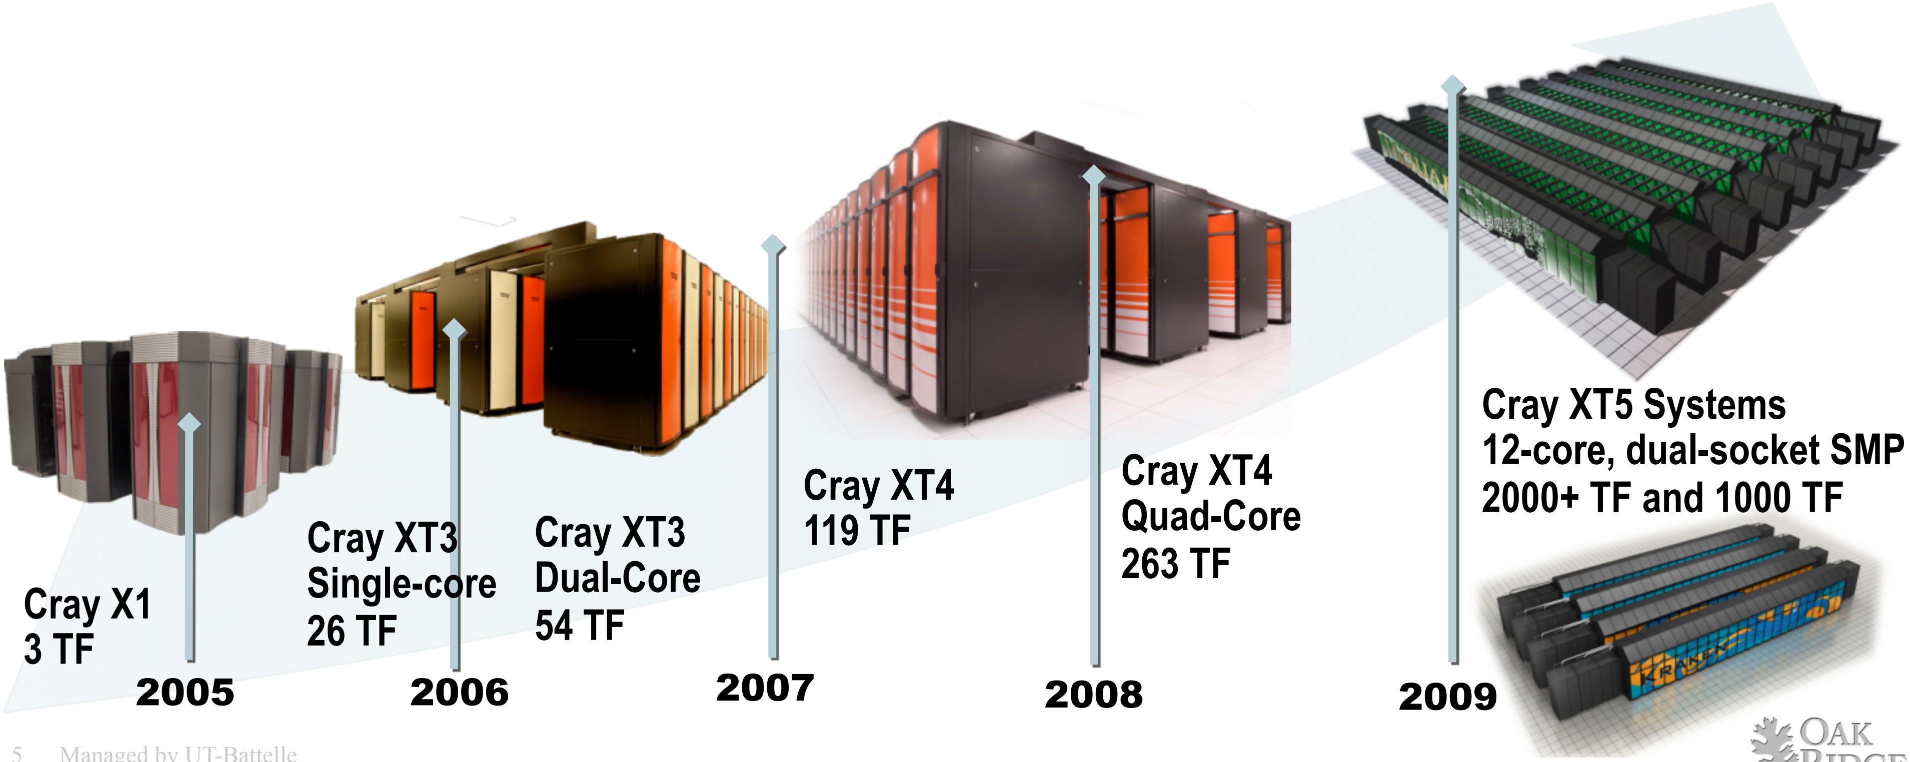

Cray X1
3 TF
2005

Cray XT3
Single-core
26 TF
2006

Cray XT3
Dual-Core
54 TF

Cray XT4
119 TF
2007

Cray XT4
Quad-Core
263 TF
2008

Cray XT5 Systems
12-core, dual-socket SMP
2000+ TF and 1000 TF
2009

The ORNL Jaguar Cray XT5 Leadership System

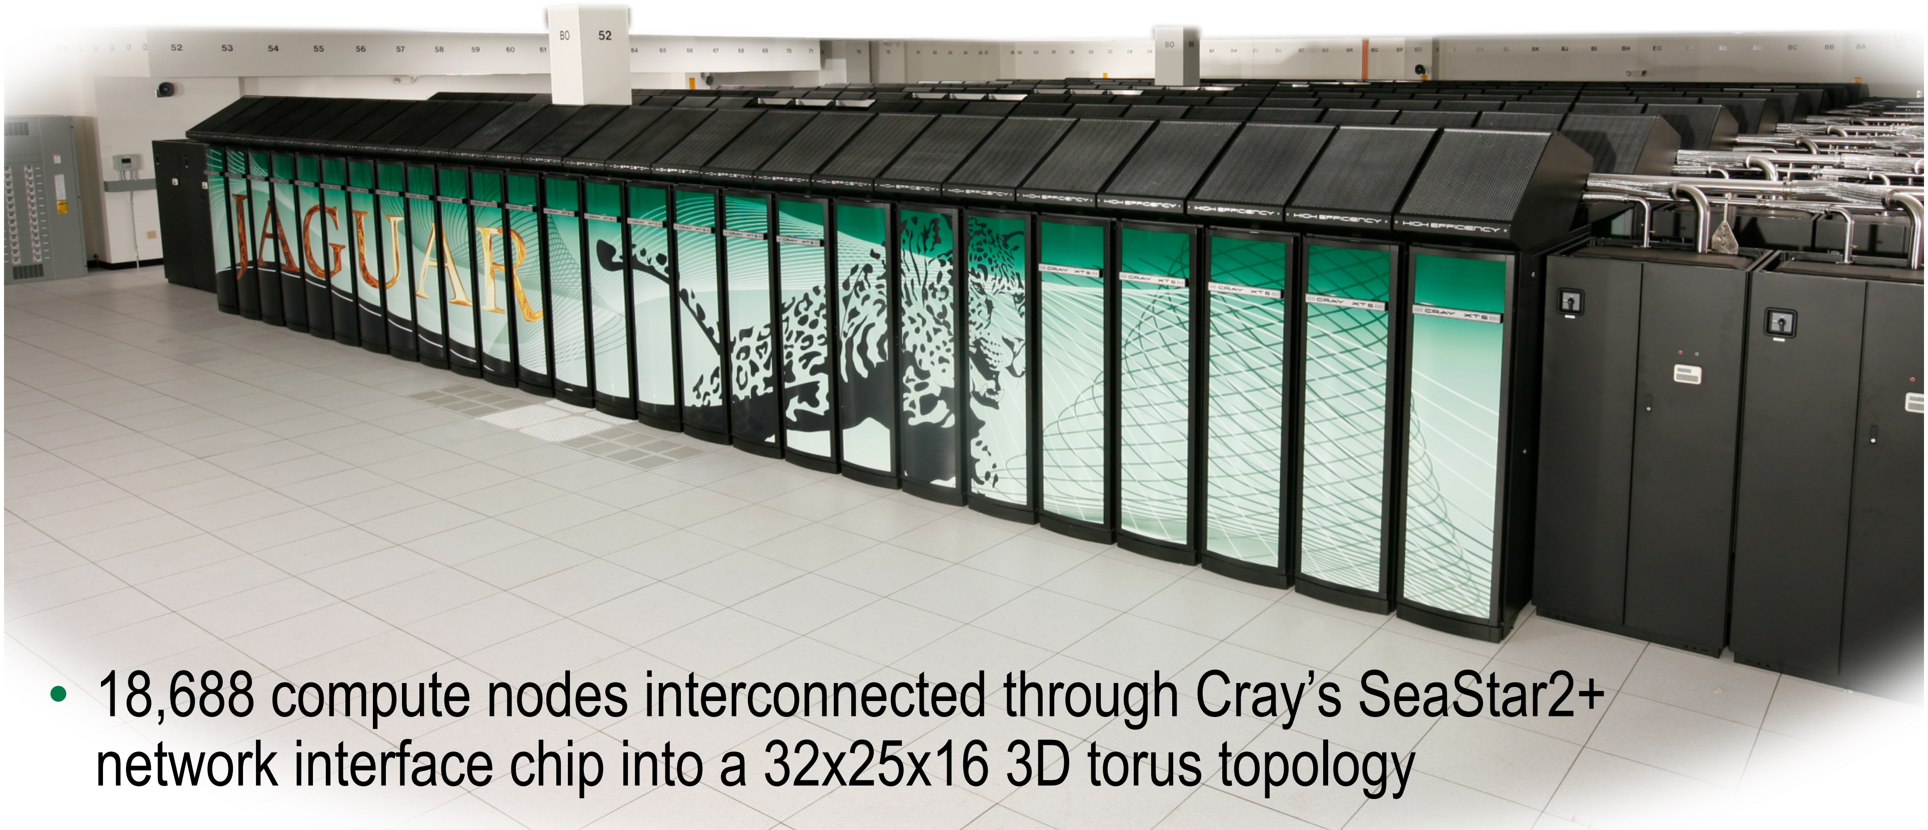

- 18,688 compute nodes interconnected through Cray's SeaStar2+ network interface chip into a 32x25x16 3D torus topology
- Compute node: dual 6-core 2.6 GHz AMD Opteron Istanbul processing sockets, each with 8 GB of 800 MHz DDR2 memory
  - With 4 flops/cycle, these nodes deliver 125 GF and 16 GB
- Compute partition aggregate: 2.3 PF, 224 GB memory, 224256 cores

LCF Compute Node Example

Already “many core”

- **Powerful node improves scalability**
- **Large shared memory**
- **OpenMP Support**
- **Low latency, High bandwidth interconnect**
- **Upgradable processor, memory, and interconnect**

GFLOPS	**124.8**
Memory (GB)	16
Cores	12
SeaStar2+	1

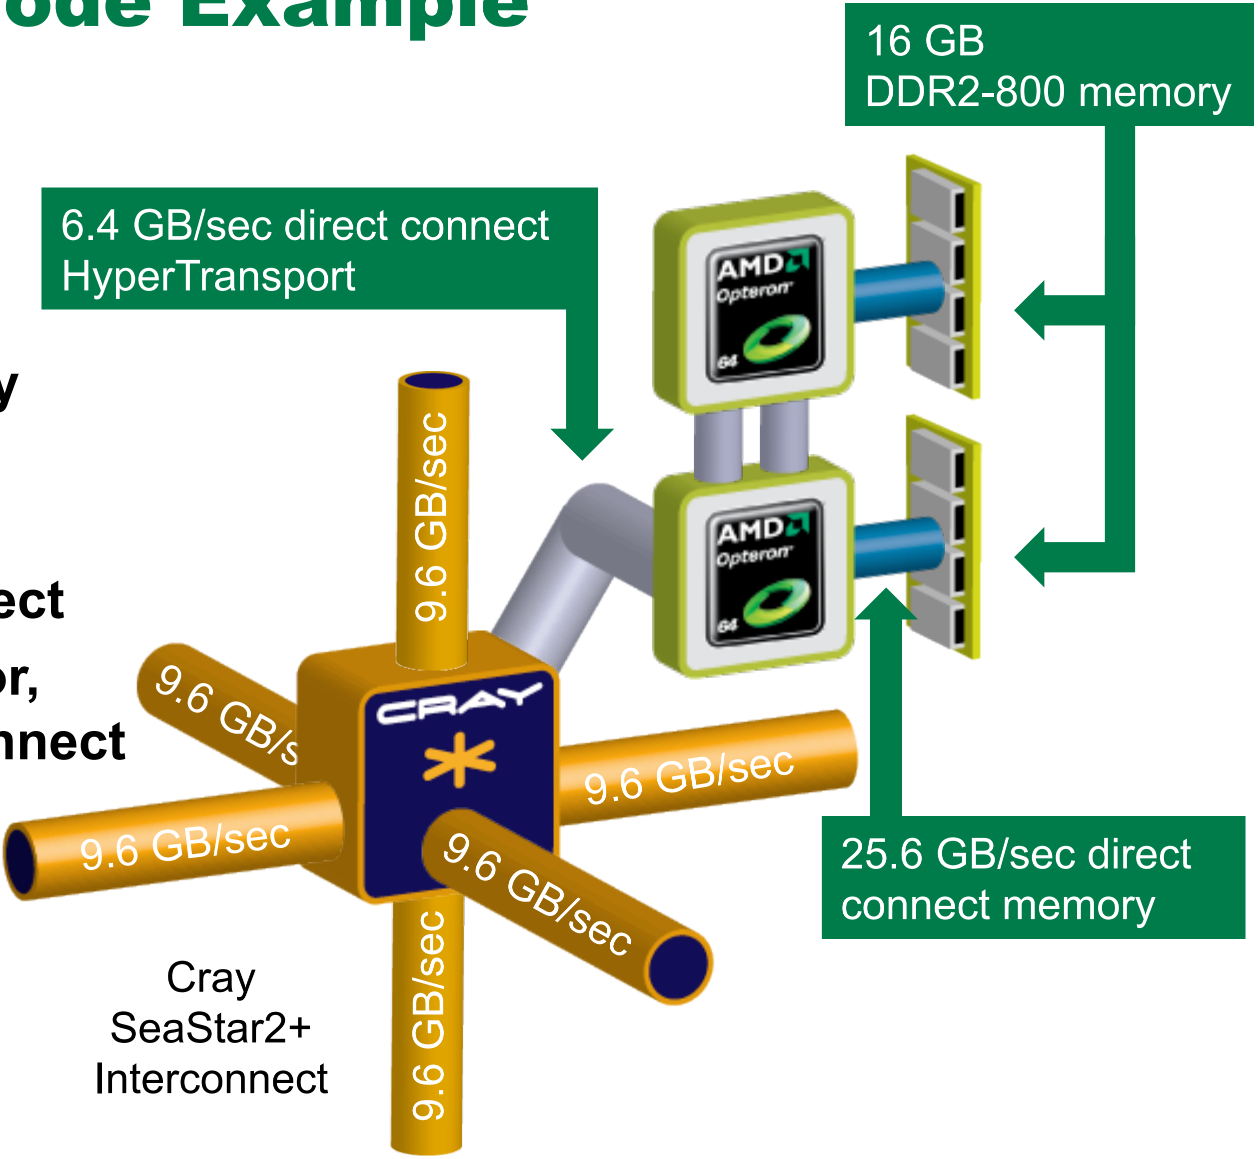

16 GB DDR2-800 memory

6.4 GB/sec direct connect HyperTransport

25.6 GB/sec direct connect memory

Cray SeaStar2+ Interconnect

Many fusion applications have nevertheless learned to exploit this node design for their science goals

60% of ASCR LCF Resources are Awarded Through the INCITE Program

Plasma physics received the largest fraction in 2010!

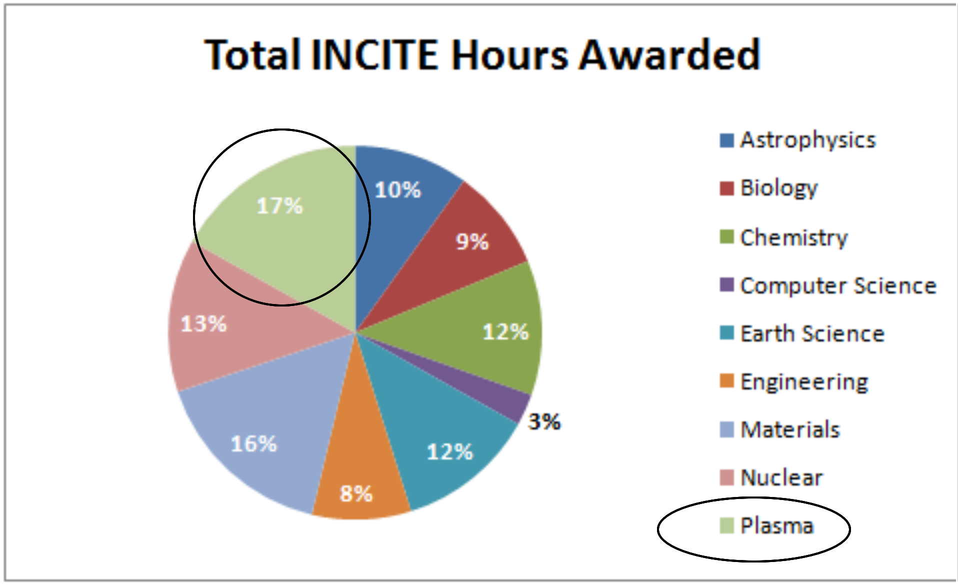

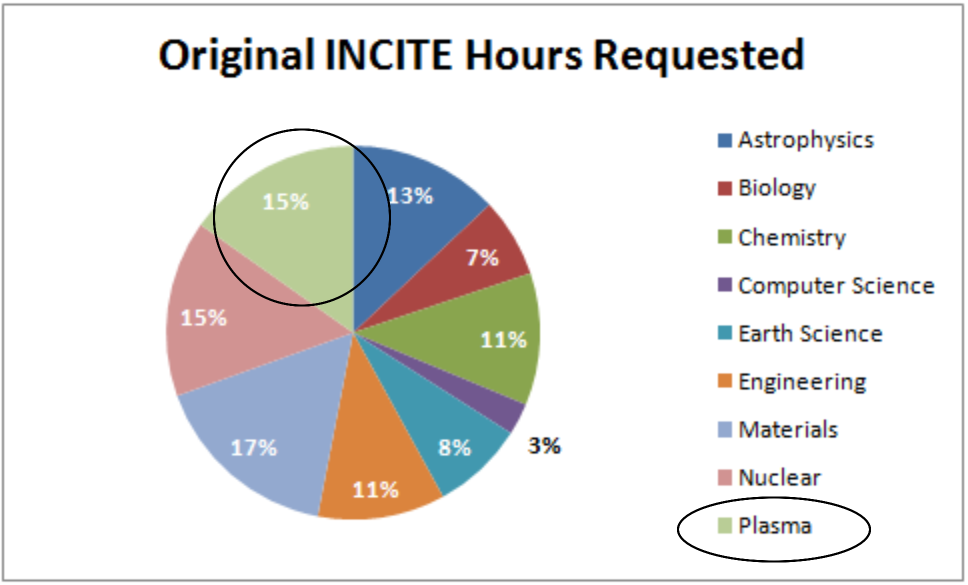

Fusion applications are some of the most “ready” to effectively exploit LCF resources to further their science goals

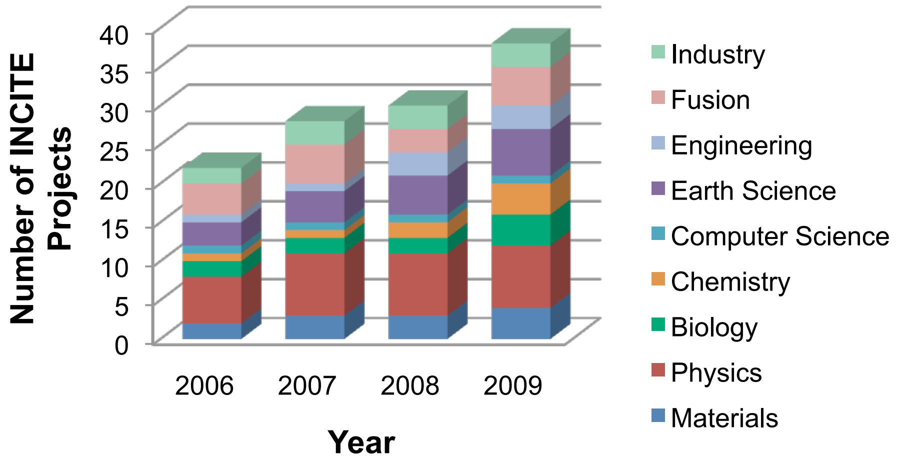

HPC Resource Requirements

ORNL LCF platform utilization for fusion projects in 2009

- INCITE Program
  - 5 projects (PIs – Berry; Diamond; Candy; Nevins; Chang)
  - 75.7M core-hours consumed
- Petascale early science program
  - 3 projects (PIs – Candy; Wang; Lin)
  - 55.2M core-hours consumed
- 2009 ASCR Joule Metric Program
  - 1 project (PI – Chang)
  - 33.6M core-hours consumed
- Leadership computing resource needs for FSP?
  - Need more analysis, but an estimate: **Y1 ~ 0.3 PF-years; Y5 ~ 30 PF-years**
- FSP will need substantial leadership computing resources (10-20%)
  - An "FSP Endstation" (much like climate) may be most appropriate

2010 INCITE Allocation for Plasma Physics

266M processor hours - 17% of total

Title (Renewals)	PI	Jaguar Award	Intrepid Award
High-fidelity tokamak edge simulation for efficient confinement of fusion plasma	Chang, C.S. (cschang@cims.nyu.edu)	50,000,000	
Validation of Plasma Microturbulence Simulations for Finite-Beta Fusion Experiments (2009)	Nevins, William (nevins@llnl.gov)	30,000,000	
Verification and validation of petascale simulation of turbulent transport in fusion plasmas (2008)	Diamond, Patrick (phd@mamacass.ucsd.edu)	35,000,000	
High Resolution Global Simulation of Plasma Microturbulence	Tang, William (wtang@princeton.edu)		12,000,000
Title (New Submittals)	**PI**	**Jaguar Award**	**Intrepid Award**
Petascale Particle-in-Cell Simulations of Plasma Based Accelerators	Mori, Warren (mori@physics.ucla.edu)	8,000,000	
Peta-Scale Particle-In-Cell Simulations of Fast Ignition	Tonge, John (tonge@physics.ucla.edu)		7,000,000
Unbalanced magnetohydrodynamic turbulence	Boldyrev, Stanislav (boldyrev@wisc.edu)		25,000,000
Gyrokinetic Simulation of Energetic Particle Turbulence in ITER Burning Plasmas	Lin, Zhihong (zhihongl@uci.edu)	20,000,000	
Simulations of laser-plasma interactions in targets for the National Ignition Facility and beyond	Hinkel, Denise (hinkel1@llnl.gov)		45,000,000
Investigation of Multi-Scale Transport Physics of Fusion Experiments Using Global Gyrokinetic Turbulence Simulations	Wang, Weixing (wwang@pppl.gov)	34,000,000	

Fusion science on LCF platforms

Core-edge nonlocal profile formation of tokamak plasma

C.S. Chang
New York Univ.

Science Objectives and Impact

- **Strategy:** Obtain first-principles understanding on the fusion efficiency of tokamak reactor plasma using extreme scale HPC
- **Driver:** Edge conditioning, which may be the only mean to control the reactor plasma, has been ubiquitously observed to improve the core plasma profile dramatically in experiments.
- **Objective:** Multiscale nonlocal core-edge simulation of the combined ion-temperature-driven turbulence dynamics and the background ion-temperature profile evolution in realistic DIII-D device geometry <u>in a day</u>.
- **Impact:** Open up a door to practical first-principles predictive simulation capability for ITER and DEMOs

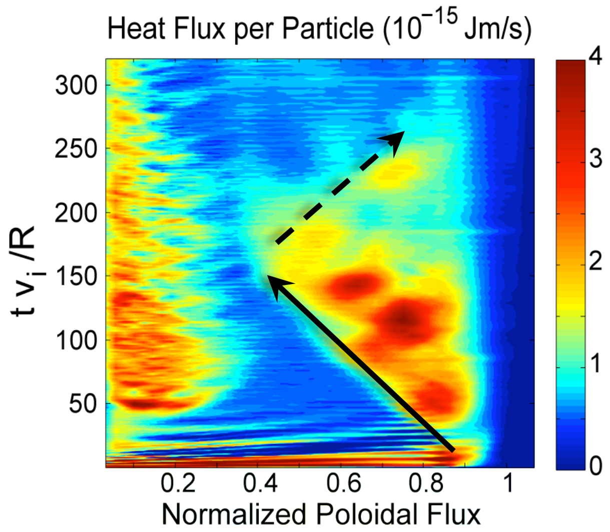

XGC1 Performance

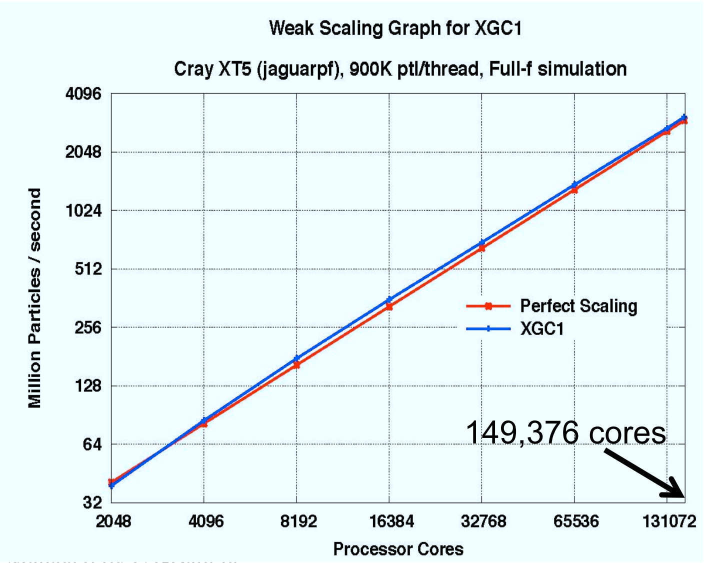

Science Results

- World's only whole volume turbulence simulation in realistic tokamak geometry.
- Core turbulence is sum of the incoming intensity from edge and the ambient local fluctuations, self-organizing the temperature gradient in turbulence propagation time scale (similar to experiments).

OLCF contribution

MPI/OpenMP architecture

Improved I/O speed in ADIOS

Image Caption

Shown above is the turbulent heat flux in time and radius (normalized flux). Turbulence propagates from edge to core (solid arrow), induces outward heat flux (dashes arrow), and leads to an eventual new self-organized nonlocal state.

Fusion science on LCF platforms

Gyrokinetic particle simulation of transport barrier dynamics in fusion plasmas

Zhihong Lin
UC Irvine

Science Objectives and Impact

- **Strategy:** Produce a net gain of energy, and produce a clean, intense and unlimited source of energy for the future.
- **Driver:** Understand the turbulent transport of energy, particle and momentum; and energetic particle physics in a tokamak burning plasma
- **Objective:** Use first principle simulations of the gyrokinetic equations, to study the Collisionless Trapped Electron Mode turbulence, and the Electron Temperature and Ion Temperature Gradient driven turbulence for ITER as well as the Toroidal Alfven Eigenmode turbulence excited by energetic particles.
- **Impact:** Understanding the nonlinear physics of the turbulent transport: improved plasma confinement required for steady state tokamak regime

GTC Performance

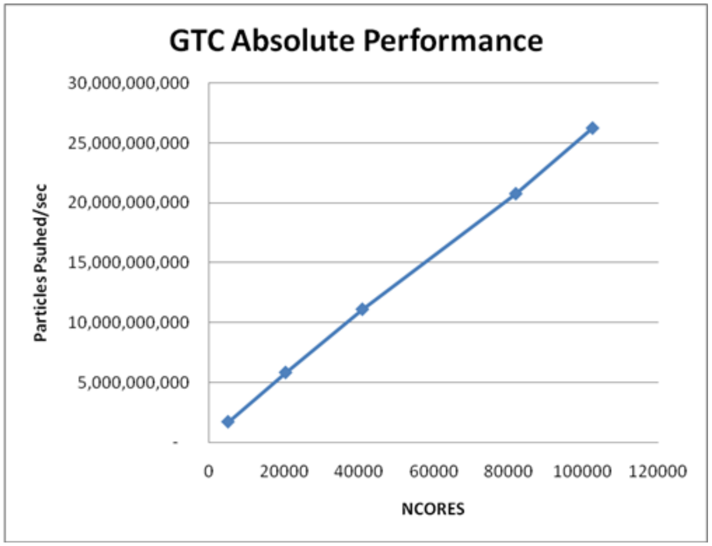

Science Results

- **CTEM**
  - Electron heat transport transits from Bohm to GyroBohm scaling when increasing system size.
  - Zonal flow is the dominant saturation mechanism for CTEM turbulence
  - The existence of mesoscale turbulent streamers is observed
  - Nondiffusive electron transport is found in the mesoscale streamers.
- **Momentum transport**
  - Toroidal momentum pinch flux is separated from the diffusive momentum flux
- **Energetic particles**
  - The energetic particle transport in the ITG turbulence is found to be diffusive.

Fusion science on LCF platforms

Weixing Wang &
Stephane Ethier, PPPL

Global gyrokinetic turbulence simulations of toroidal momentum and electron thermal transport and comparison with the national spherical torus (NSTX) experiments

Science Objectives and Impact

- **Strategy:** Validation of global gyrokinetic PIC simulations against existing fusion experiments to advance understanding
- **Driver:** i) Energy losses in ITER are expected to be dominated by electron transport, which remains poorly understood; ii) how ITER plasma rotates – critical for both MHD and transport behavior
- **Objective:** To understand ii) role of electron-temperature-gradient (ETG) turbulence driving electron transport in real NSTX experiments, ii) non-diffusive toroidal momentum transport
- **Impact:** Development of predictive capability for the assessment of confinement performance in ITER.

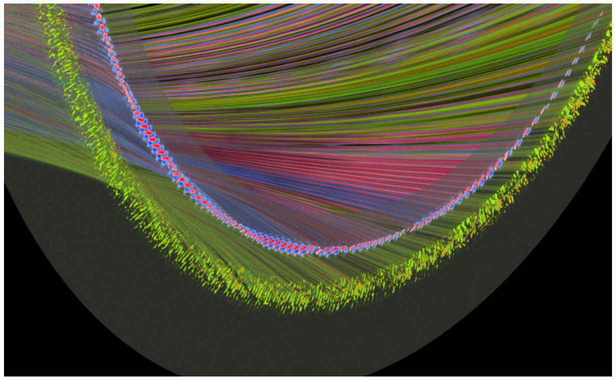

GTS Performance

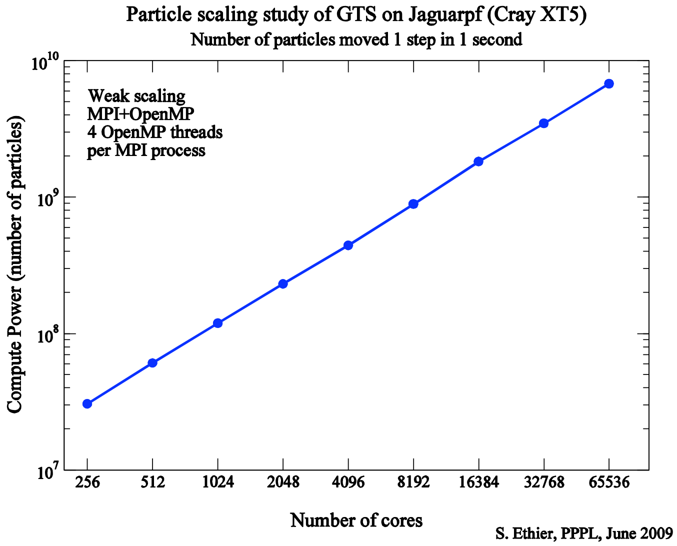

S. Ethier, PPPL, June 2009

Science Results

- GTS simulations i) confirmed the presence of electron-scale fluctuations associated with ETG turbulence; ii) but suggest its insignificant contribution to observed anomalous electron transport, being consistent with preliminary experimental evidence
- Trapped electron physics in turbulence plays a critical role in both producing proper ion heat transport and generating plasma rotation

OLCF contribution

Implementation of a fast parallel I/O routines and subsystem.

Image Caption

The image above shows electron-scale electrostatic potential fluctuations generated by ETG instabilities in a global gyrokinetic PIC simulation of the National Spherical Torus Experiment (NSTX). (Visualization by K.-L. Ma, UC Davis)

Fusion science on LCF platforms

Tokamak profile prediction using first-principles simulation

Jeff Candy
General Atomics

Science Objectives and Impact

- **Strategy:** Predict tokamak profiles inside pedestal radius using first-principles simulations rather than reduced models
- **Driver:** Fusion performance of ITER is sensitive to physics not well-captured in reduced models
- **Objective:** Proof of principles use of embedded gyrokinetic turbulence and neoclassical transport in massively-parallel transport solver
- **Impact:** Improved understanding of transport in existing experiments, and improved accuracy of ITER performance projections

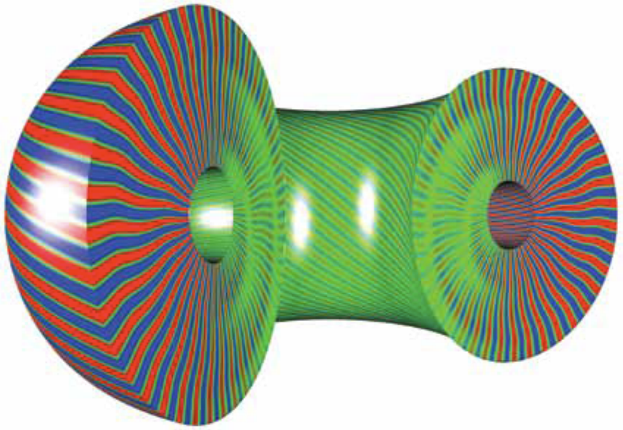

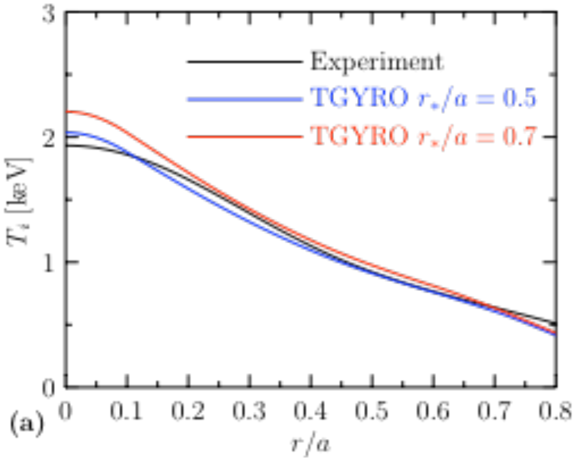

Science Results

- World's first direct gyrokinetic/neoclassical calculation of steady-state plasma profiles
- Determine electron and ion temperatures by solving the steady-state transport equations (a = i, e) given experimental sources S_{exp}:

$$\frac{1}{V'}\frac{\partial}{\partial r}\left[V'\left(Q_a^{\mathrm{NEO}}+Q_a^{\mathrm{GYRO}}\right)\right] = n_0\nu_{ab}\left(T_a-T_b\right)+S_a^{\mathrm{EXP}}$$

Image Captions

TGYRO prediction of DIII-D temperature profiles: colored curves show TGYRO predictions at different matching radii

Scientific Progress at the Petascale*

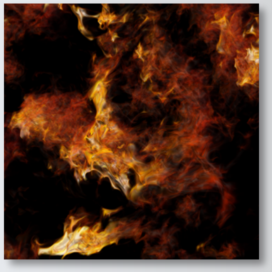

Turbulence
Understanding the statistical geometry of turbulent dispersion of pollutants in the environment.

Energy Storage
Understanding the storage and flow of energy in next-generation nanostructured carbon tube supercapacitors with potential storage capacity two to three times greater than conventional capacitors.

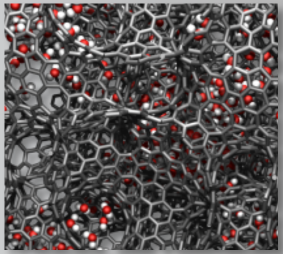

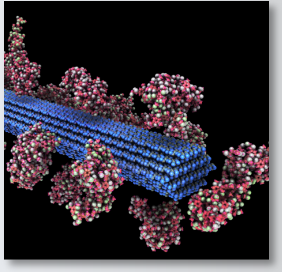

Biofuels
A comprehensive simulation model of lignocellulosic biomass to understand and overcome the physical basis of recalcitrance to hydrolysis, which is the bottleneck to sustainable and economical ethanol production.

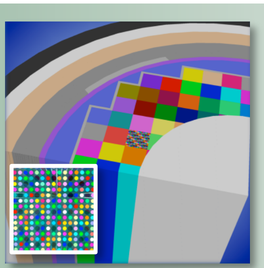

Nuclear Energy
High-fidelity predictive simulation tools for the design of next-generation nuclear reactors to safely increase operating margins.

Fusion Energy
Substantial progress in the understanding of anomalous electron energy loss in tokamak plasmas by validation with the National Spherical Torus Experiment (NSTX).

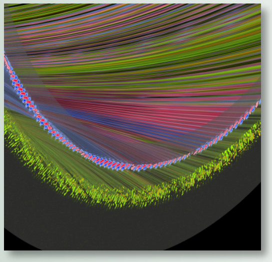

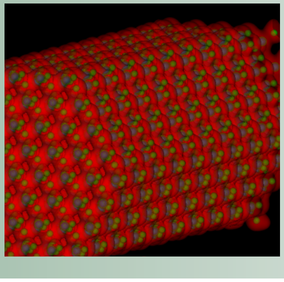

Nano Science
Understanding the atomic and electronic properties of nanostructures to serve as multiband semiconductors in next-generation photovoltaic solar cell materials.

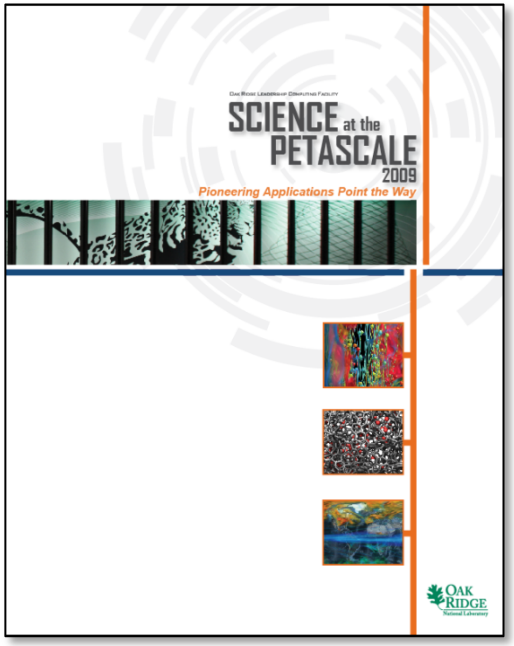

SCIENCE at the PETASCALE
2009
Pioneering Applications Point the Way

Computational fusion science projects on the ASCR LCF platforms are demonstrably leading the overall computational science community

*http://www.nccs.gov/wp-content/media/nccs_reports/PetaDocLowRes2-9-10.pdf

Moving to the Exascale

- The DOE (SC/NNSA) has held a series of workshops in 2009-2010 to assess the opportunities and challenges of exascale computing for the advancement of science, technology, and Office of Science missions.
- ASCR has been asked to identify strategies to address the challenges and deliver on such opportunities.
- As part of this strategy, ASCR will work with the LCFs in providing a series of increasingly powerful computer systems and work with the user community to scale applications to each of the new computer systems
  - FSP is a key driver and stakeholder

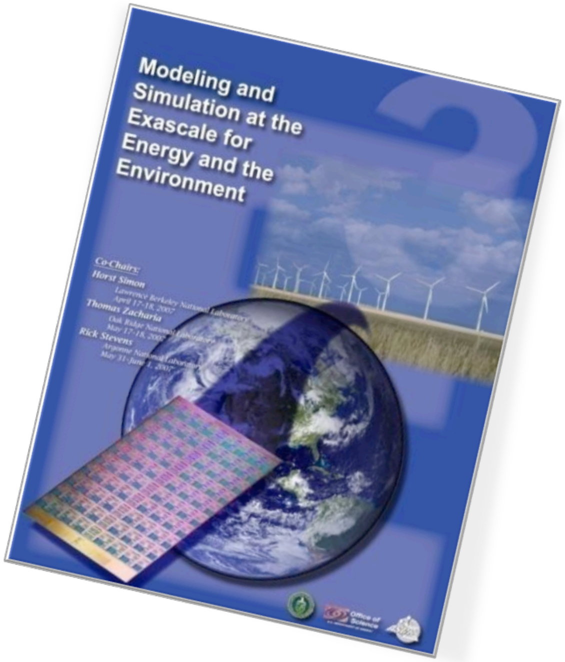

Modeling and Simulation at the Exascale for Energy and the Environment

Co-Chairs:
Horst Simon
Lawrence Berkeley National Laboratory
April 17-18, 2007
Thomas Zacharia
Oak Ridge National Laboratory
May 17-18, 2007
Rick Stevens
Argonne National Laboratory
May 31-June 1, 2007

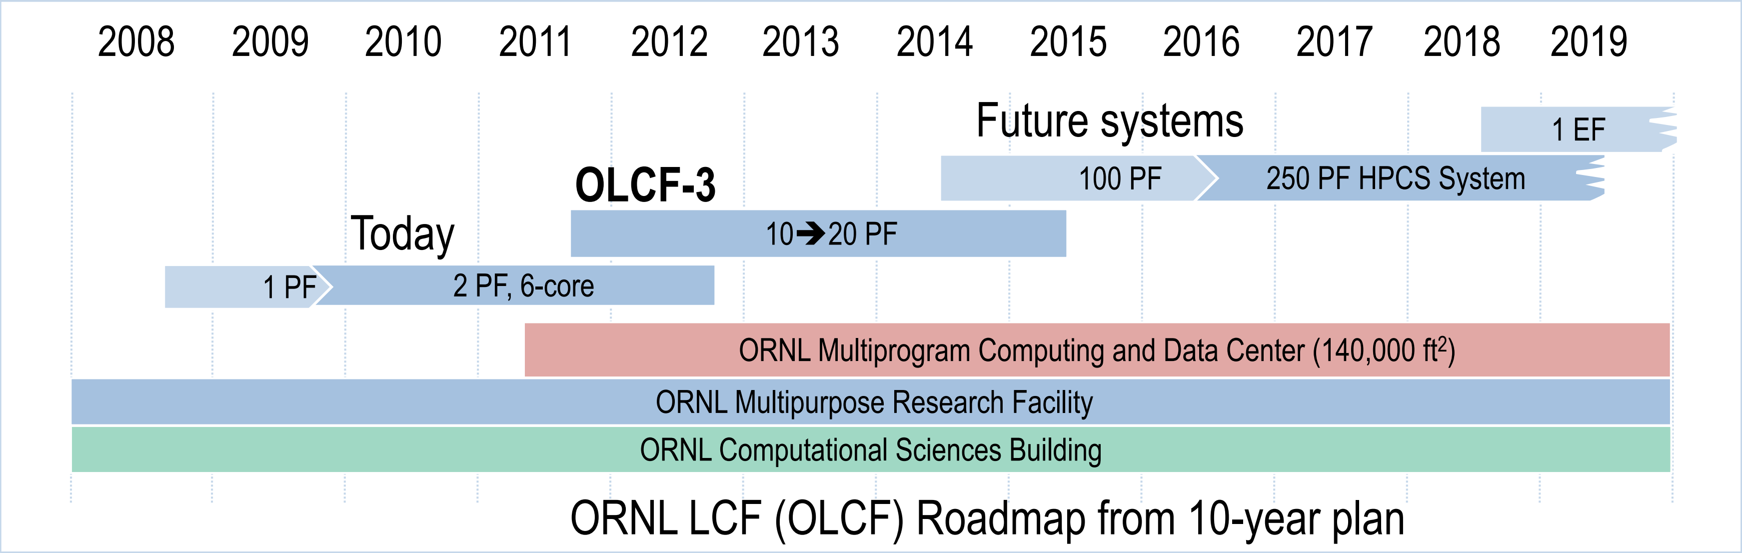

2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019

Future systems

1 EF

100 PF

250 PF HPCS System

OLCF-3

10➔20 PF

Today

1 PF

2 PF, 6-core

ORNL Multiprogram Computing and Data Center (140,000 ft^2)

ORNL Multipurpose Research Facility

ORNL Computational Sciences Building

ORNL LCF (OLCF) Roadmap from 10-year plan

FSP Applications Must be Prepared to Exploit Local Concurrency to Take Advantage of Coming Hybrid H/W

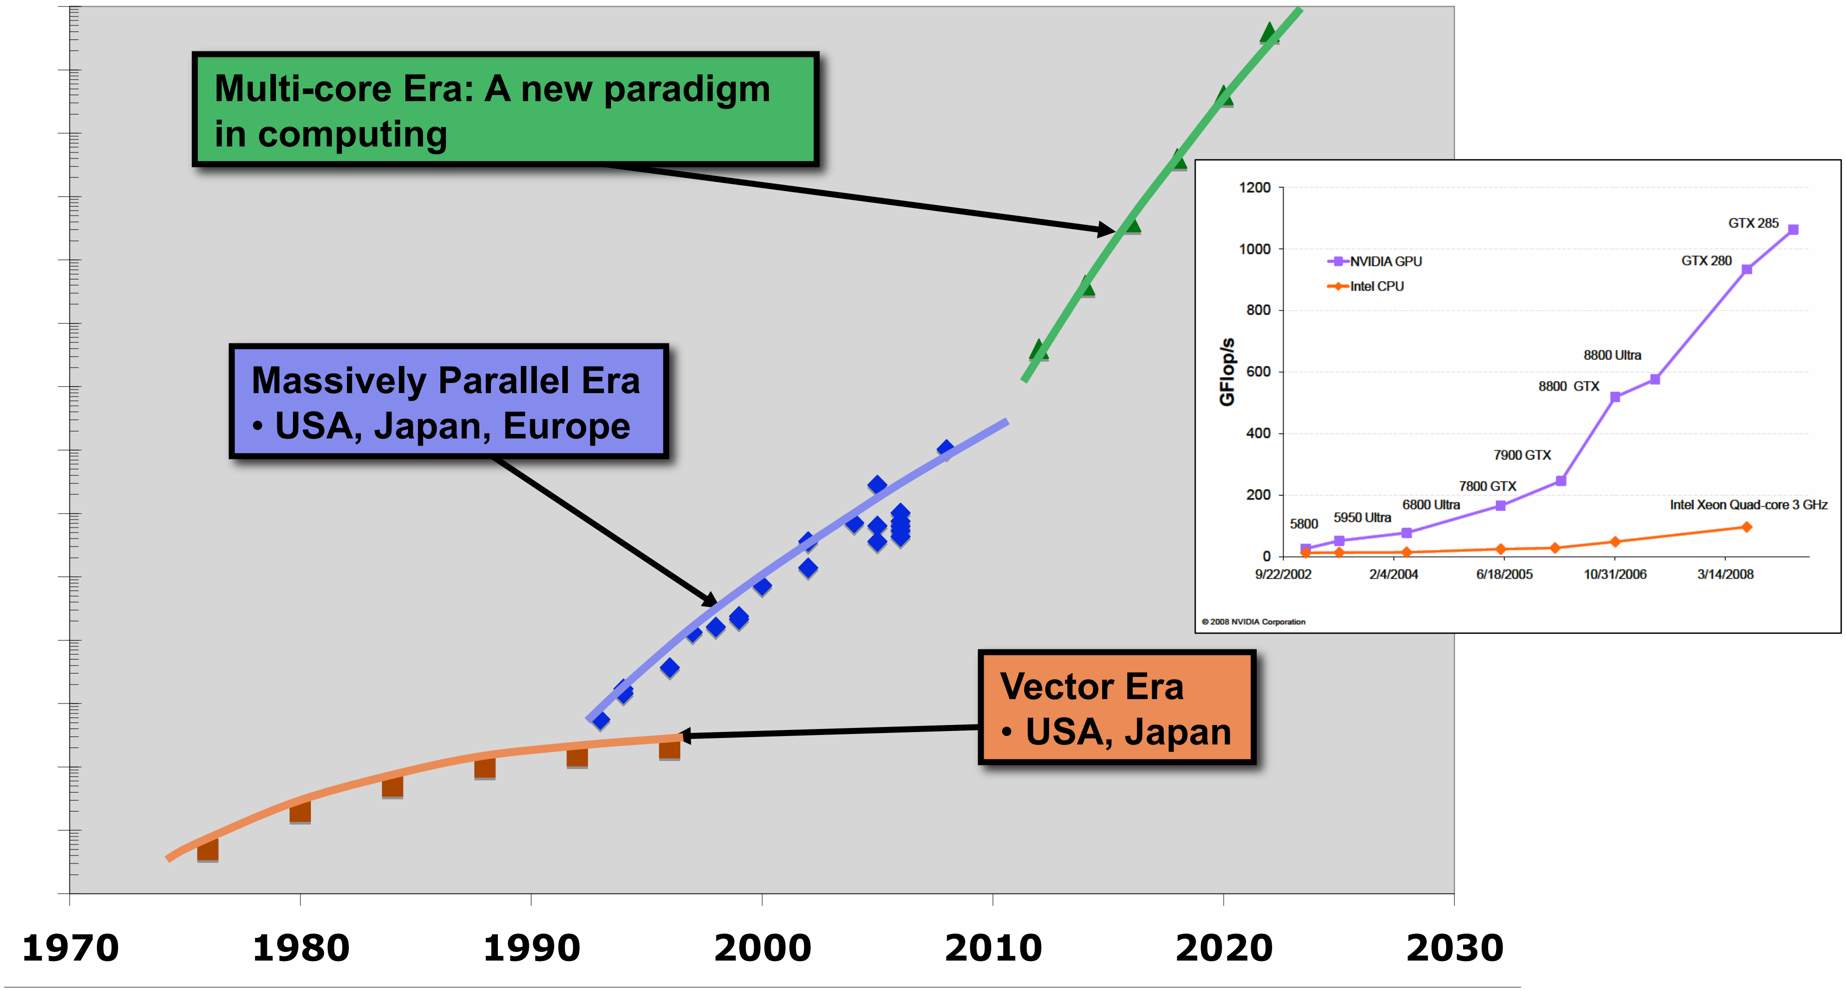

Multi-core Era: A new paradigm in computing

Massively Parallel Era
- **USA, Japan, Europe**

Vector Era
- **USA, Japan**

An Example of What FSP Applications Will Face Early On (2012-2013 Timeframe)

OLCF-3 System Description

- **10-20 PF peak performance, much larger memory**
- **Same number of cabinets, cabinet design, and cooling as Jaguar**
- **Operating system upgrade of today's Cray Linux Environment**
- **New Gemini interconnect**
  - 3-D Torus
  - Globally addressable memory
  - Advanced synchronization features
- **Same number of nodes (~19K) as current Jaguar!**
  - Accelerated node design
  - Next generation interconnect
  - Next generation AMD processor
  - Future NVIDIA accelerator
  - Fat nodes: 70 GB memory
  - Very high performance processors
  - Very high memory bandwidth
- **3x larger and 4x faster file system**

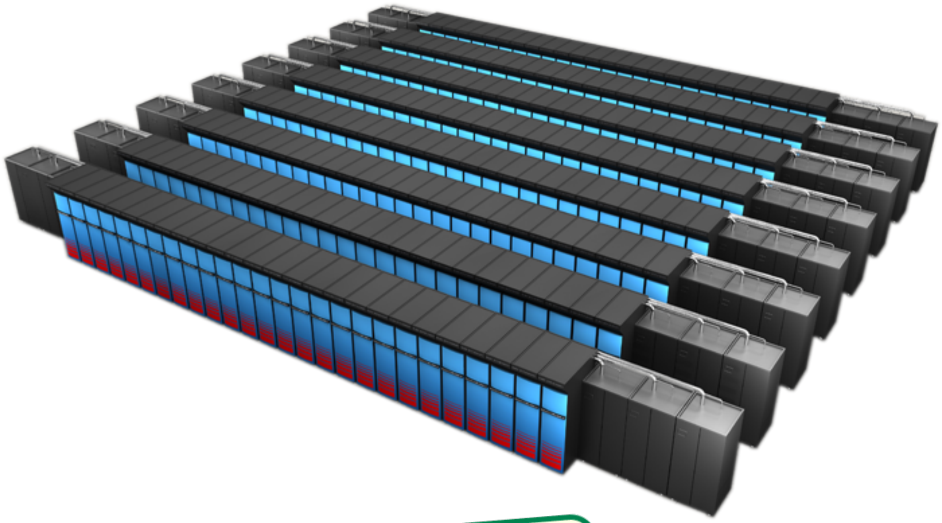

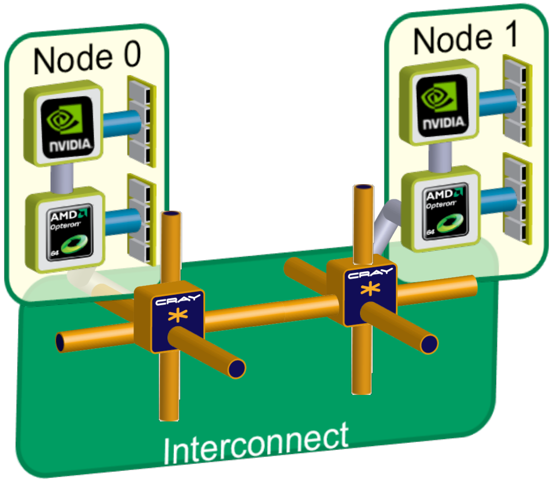

Contrary to popular belief, applications must work to "scale in" (*not* "scale out")

FSP Must Prepare for Hybrid Architectures*

General Purpose GPUs, Floating Point Accelerators, etc.

- **Large GPU-based systems springing up everywhere**
  - NSF Track 2D in negotiation with Georgia Tech/ORNL
  - Japan: “Tsubame” at Tokyo Institute of Technology (upgrade planned)
  - CEA: 300 TF Nehalem/NVIDIA cluster located at CCRT
  - “Orbit” ORNL 100 TF NVIDIA testbed
  - Oil and gas industry deploying large clusters
- **Features for computing on GPUs**
  - Added high-performance 64-bit arithmetic
  - Adding ECC and parity that other GPU vendors have not added
  - Critical for a large system
  - Larger memories
  - Dual copy engines for simultaneous execution and copy
  - S1070 has 4 GPUs exclusively for computing
  - No video out cables
  - Development of CUDA and recently announced work with PGI on Fortran CUDA
- **Large and growing pool of people who know how to program accelerators and who will develop tools**
  - Every laptop has a processor and GPU
    - Macintosh, PC, Linux ports of CUDA available
  - Most computer science programs now teach GPU programming

300+ accelerated applications listed on NVIDIA’s web site: http://www.nvidia.com/object/cuda_home.html

*See also the new Hybrid Multicore Consortium (HMC), http://computing.ornl.gov/HMC

DOE Exascale Initiative Technical Roadmap

Systems	2009	2015 +1/-0	2018 +1/-0
System peak	**2 Peta**	**100-200 Peta**	**1 Exa**
Power	**6 MW**	**~15 MW**	**~20 MW**
System memory	0.3 PB	5 PB	64 PB (+)
Node performance	125 GF	0.5 TF or 7 TF	1,2 or 15TF
Node memory BW	25 GB/s	1-2TB/s	2-4TB/s
Node concurrency	12	O(100)	O(1k) or 10k
Total Node Interconnect BW	3.5 GB/s	100-200 GB/s 10:1 vs memory bandwidth 2:1 alternative	200-400GB/s (1:4 or 1:8 from memory BW)
System size (nodes)	18,700	50,000 or 500,000	O(100,000) or O(1M)
Total concurrency	225,000	O(100,000,000) *O(10)-O(50) to hide latency	O(billion) * O(10) to O(100) for latency hiding
Storage	15 PB	150 PB	500-1000 PB (>10x system memory is min)
IO	0.2 TB	10 TB/s	60 TB/s (how long to drain the machine)
MTTI	days	O(1day)	O(0.1 day)

How Can FSP Ensure its Applications are "Ready" for Systems of the Future?

"Generic" application readiness process for a hybrid architecture (e.g., OLCF-3 system)

- **define test problem(s)**
- **baseline CPU performance**
- **identify kernel(s)**
- **understand kernel performance**
- **extract kernel(s)**
- **if appropriate…**
  - create standalone driver(s)
  - define standalone test problem(s)
- **provide kernel(s) to NVIDIA**
- **accelerate kernel(s)**
  - e.g. implement kernel(s) in CUDA
- **understand accelerated kernel performance & optimize**
- **implement accelerated kernel(s) in application**
- **understand accelerated application performance**

- **Kernel acceleration candidates**
- **PIC (GTC/GTS/XGC1)**
  - Particle pushing
  - Charge deposition
  - Field smoothing
- **RF heating (AORSA)**
  - Full spectrum of toroidal harmonics for specific antenna geometries
  - Dense linear solver
  - Mixed precision

Algorithms Moving Forward
What are application Teams Asking for?

- Automated diagnostics
  - Drivers: performance analysis, application verification, S/W debugging, H/W-fault detection and correction, failure prediction and avoidance, system tuning, and requirements analysis
- Hardware latency
  - Won't see improvement nearly as much as flop rate, parallelism, B/W in coming years
  - Can S/W strategies mitigate high H/W latencies?
- Hierarchical algorithms
  - Applications will require algorithms aware of the system hierarchy (compute/memory)
  - In addition to hybrid data parallelism, and file-based checkpointing, algorithms may need to include dynamic decisions between recomputing and storing, fine-scale task-data hybrid parallelism, and in-memory checkpointing
- Parallel programming models
  - Improved programming models needed to allow developer to identify an arbitrary number of levels of parallelism and map them onto hardware hierarchies at runtime
  - Models continue to be coupled into larger models, driving the need for arbitrary hierarchies of task and data parallelism
- Solver technology and innovative solution techniques
  - Global communication operations across 10^{6-8} processors will be prohibitively expensive, solvers will have to eliminate global communication where feasible and mitigate its effects where it cannot be avoided. Research on more effective local preconditioners will become a very high priority
  - If increases in memory B/W continue to lag the number of cores added to each socket, further research needed into ways to effectively trade flops for memory loads/stores
- Accelerated time integration
  - Are we ignoring the time dimension along which to exploit parallelism? (Ex: climate)
- Model coupling
  - Coupled models require effective methods to implement, verify, and validate the couplings, which can occur across wide spatial and temporal scales. The coupling requirements drive the need for robust methods for downscaling, upscaling, and coupled nonlinear solving
  - Evaluation of the accuracy and importance of couplings drives the need for methods for validation, uncertainty analysis, and sensitivity analysis of these complex models
- Maintaining current libraries

Co-Design Teams: A Better Way to Successfully Build Applications?

ASCR is pushing to formalize this concept

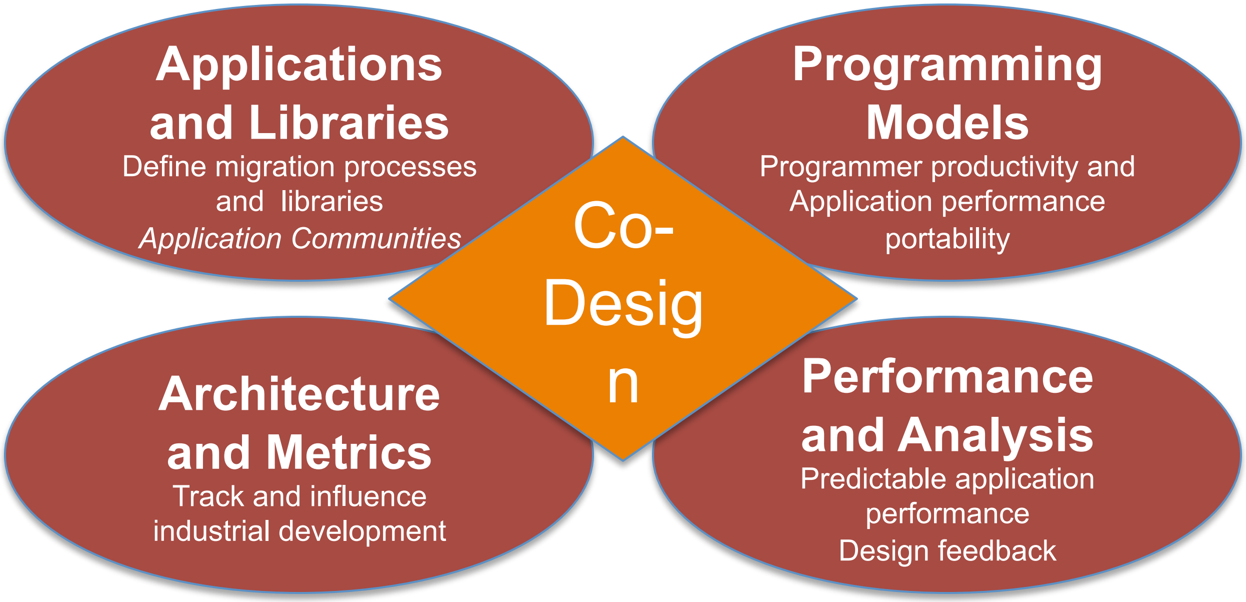

Applications and Libraries
Define migration processes and libraries
Application Communities

Programming Models
Programmer productivity and Application performance portability

Co-Design

Architecture and Metrics
Track and influence industrial development

Performance and Analysis
Predictable application performance
Design feedback

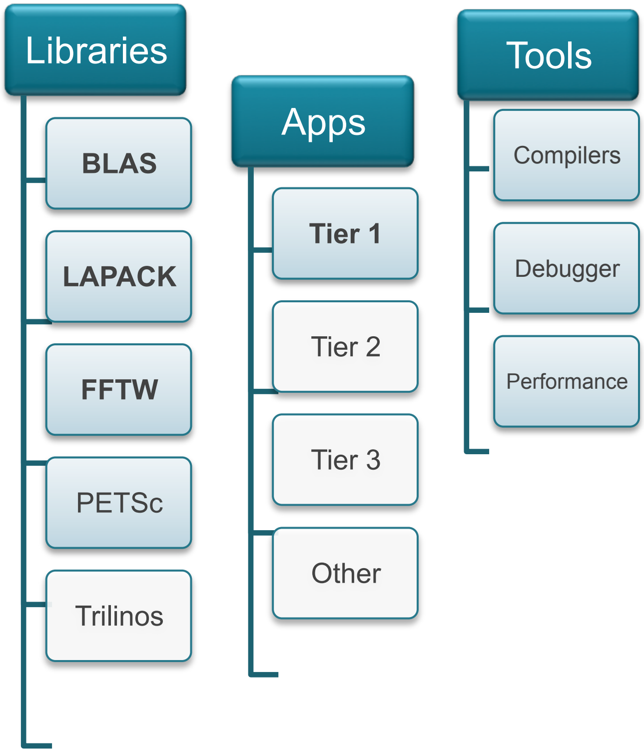

FSP can and must implement a co-design team structure to succeed. FSP will work with ASCR and OFES to communicate lessons learned and best practices to other co-design science teams within the DOE.

Fusion Application Computing Requirements

- We have elicited, analyzed, & validated science application requirements using a new comprehensive requirements questionnaire
  - Analyzed project overview, science motivation & impact, application models, algorithms, parallelization strategy, S/W, development process, SQA, V&V, usage workflow, performance
  - Results, analysis, and conclusions documented in 2009 OLCF application requirements document

PREPARING FOR EXASCALE

ORNL Leadership Computing
Application Requirements and Strategy

October 2009

OLCF

National Center for Computational Sciences

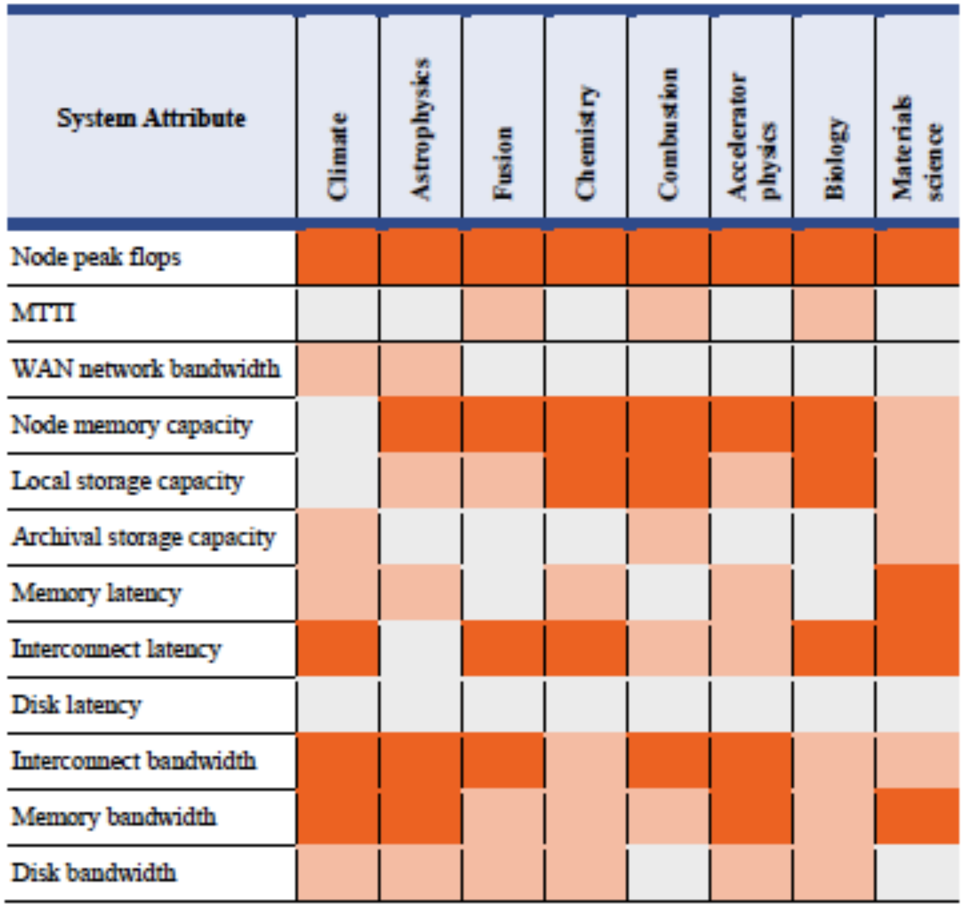

System Attribute	Climate	Astrophysics	Fusion	Chemistry	Combustion	Accelerator physics	Biology	Materials science
Node peak flops	🟧	🟧	🟧	🟧	🟧	🟧	🟧	🟧
MTTI	⬜	⬜	▨	⬜	▨	⬜	▨	⬜
WAN network bandwidth	▨	▨	⬜	⬜	⬜	⬜	⬜	⬜
Node memory capacity	⬜	🟧	🟧	🟧	🟧	🟧	🟧	▨
Local storage capacity	⬜	▨	▨	🟧	🟧	▨	🟧	▨
Archival storage capacity	▨	⬜	⬜	⬜	▨	⬜	⬜	▨
Memory latency	▨	▨	⬜	▨	⬜	▨	⬜	🟧
Interconnect latency	🟧	⬜	🟧	🟧	▨	▨	🟧	🟧
Disk latency	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜
Interconnect bandwidth	🟧	🟧	🟧	▨	🟧	🟧	▨	▨
Memory bandwidth	🟧	🟧	▨	▨	▨	🟧	▨	🟧
Disk bandwidth	▨	▨	▨	▨	⬜	▨	▨	⬜

Not surprisingly, fusion science applications are in need of increased node peak flops, node memory capacity, and reduced interconnect latency

http://www.nccs.gov/wp-content/media/nccs_reports/olcf-requirements.pdf

Community Library Requirements

Fusion Applications are among the most I/O intensive

Science domain	Code	I/O libraries	Math libraries
Accelerator design	T3P	NetCDF	MUMPS, ParMETIS, Zoltan
Astrophysics	CHIMERA	HDF5 (pNetCDF)	LAPACK
	VULCAN/2D	HDF5	PETSc
Biology	LAMMPS		FFTW
Chemistry	MADNESS		BLAS
	NWChem		BLAS, ScaLAPACK, FFTPACK
	OReTran		LAPACK
Climate	CAM	NetCDF	(SciLib)
	POP/CICE	NetCDF	
	MITgcm	NetCDF	
Combustion	S3D		
Fusion	AORSA	NetCDF	ScaLAPACK, FFTPACK
	GTC	MPI-IO, HDF5, NetCDF, XML	PetSC
	GYRO	MPI-IO, NetCDF	BLAS, LAPACK, UMFPACK, MUMPS, FFTW (SciLib, ESSL)
Geophysics	PFLOTRAN		BLAS, PetSC
Materials science	LSMS	HDF5, XML	BLAS, LAPACK
	QBOX	XML	LAPACK, ScaLAPACK, FFTW
	QMC		BLAS, LAPACK, SPRNG
Nanoscience	CASINO		BLAS
	VASP		BLAS, ScaLAPACK
Nuclear energy	NEWTRNX	HDF5	LAPACK, PARPACK
Nuclear physics	CCSD	MPI-IO	BLAS
QCD	MILC, Chroma		

Current Fusion Algorithms Exhibit Widely Varying Performance

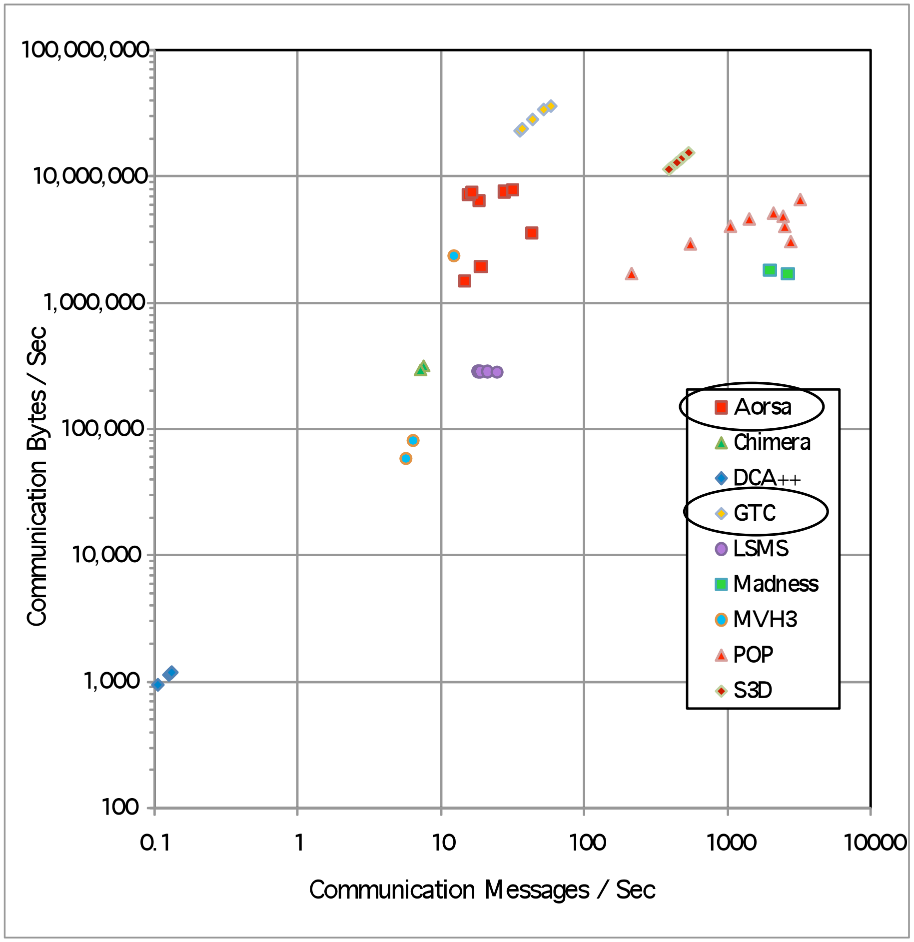

Application communication measures

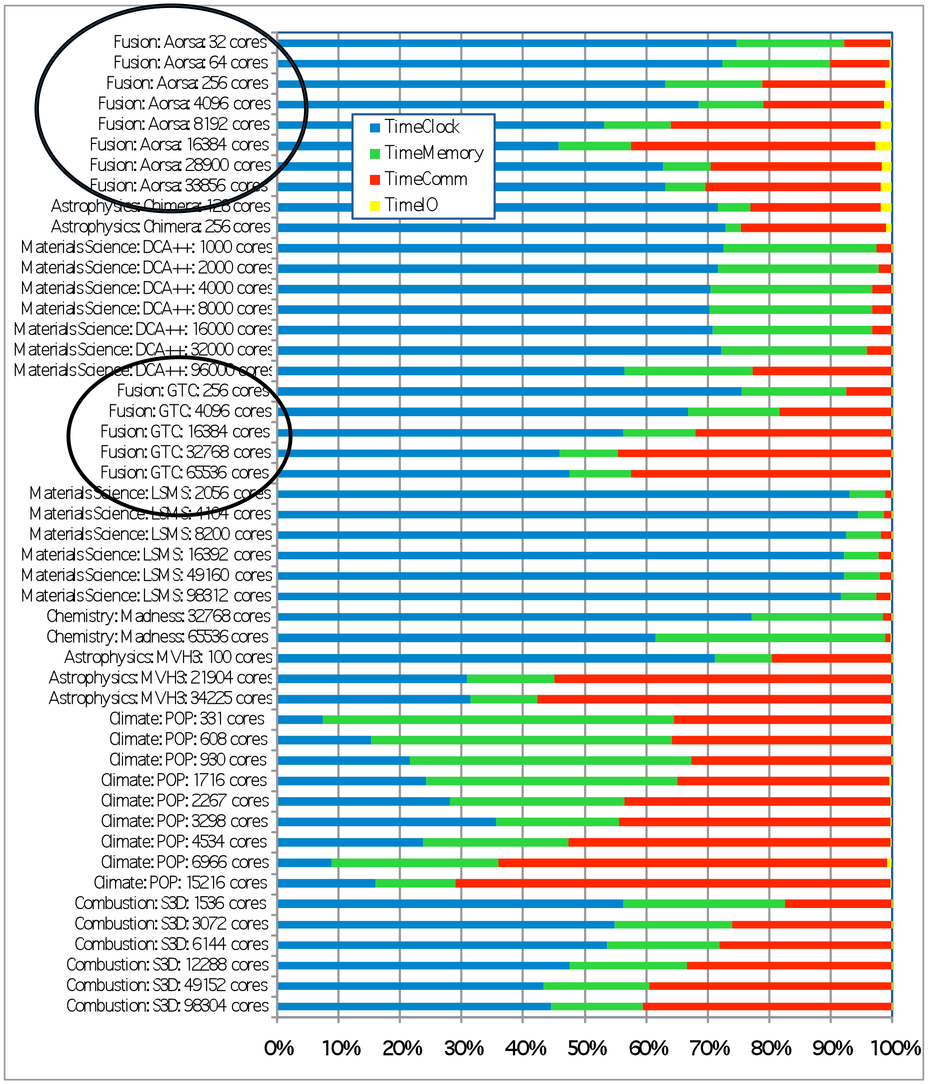

Application runtime fraction by H/W subsystem

Fusion Algorithms: Current and Trends

Science domain	Code	Structured grids	Unstructured grids	FFT	Dense linear algebra	Sparse linear algebra	Particles	Monte Carlo
Accelerator physics	T3P		X			X		
Astrophysics	CHIMERA	X			X	X	X	
	VULCAN/2D		X		X			
Biology	LAMMPS			X			X	
Chemistry	MADNESS		X		X			
	NWCHEM			X	X			
	OReTran	X		X	X			
Climate	CAM	X		X			X	
	POP/CICE	X				X	X	
	MITgcm	X				X	X	
Combustion	S3D	X						
Fusion	AORSA	X		X	X			
	GTC	X				X	X	X
	GYRO	X		X	X	X		
Geophysics	PFLOTRAN	X	X			X		
Materials science	QMC/DCA				X			X
	QBOX			X	X		X	
Nanoscience	CASINO						X	X
	LSMS	X						
Nuclear energy	NEWTRNX		X					
	Denovo	X						
Nuclear physics	CCSD							
QCD	MILC	X						

Algorithm Motifs: Current

Algorithm Motifs: Trends

Science Area	Application	Algorithm Motif Change
Astrophysics	CHIMERA	Structured grids, dense linear algebra (↑), sparse linear algebra (↑), particles, unstructured grids (↑)
Bioenergy/Biology	LAMMPS	FFT (↓), particles (↑)
Chemistry	MADNESS	Unstructured grids, dense linear algebra (↑)
Chemistry	NWCHEM	FFT, dense linear algebra
Climate	CAM-HOMME	Structured grids, particles, sparse linear algebra (↑)
Combustion	S3D	Structured grids (↑)
Fusion	AORSA	Structured grids, FFT, dense linear algebra (↑)
Fusion	GTC, GTS	Structured grids, sparse linear algebra, particles (↑), Monte Carlo
Materials Science	DCA++	Dense linear algebra (↑), Monte Carlo (↑)
Materials Science	WL-LSMS	Structured grids, dense linear algebra
Nuclear Energy	Denovo	Structured grids (↑), sparse linear algebra (↑), dense linear algebra (↑)

Scalable Algorithms
FSP priority research directions

- Optimal representations
- Multiphysics and multiscale algorithms
- Real-time algorithms
- Optimization
- Uncertainty quantification (UQ) and reduction
- Lower threshold of expertise required to use optimal algorithms on extreme architectures

Mathematical Formalisms

FSP priority research directions

- Many formulations for plasma physics exist
  - PDE-based and particle-based models for kinetic approaches
  - PDE-based models for moment closures, and many discretizations that are customized to asymptotic physical regimes in different devices or different subdomains of the same device
  - There are central problems in each of these areas for which new models and discretizations could play an important role
- Need for high-fidelity kinetics calculations, both in the core and edge
- More accurate gyrokinetic approximations
- Systematic methods for constructing nearly field-aligned coordinates
- Fundamental new numerical algorithms for particle-in-cell, the need for symplectic integrators for both particle-based and continuum-based methods, and treatments of kinetic electrons.
- Mathematically systematic treatment of coupled systems with vastly different spatial and/or temporal scales, including well-posedness, stability, and accuracy
  - A classic example is the coupled treatment of turbulence and transport.

Data Analytics, Management, & Visualization

- A big issue for FSP
  - Fusion applications on LCF platforms are the biggest generators of data along with climate and astrophysics
- A central repository for both experimental and simulation data is needed (Ex: ASCR-funded ESG project)
- In situ data analytics (compute and analyze simultaneously)
- Methods for analyzing large data sets and across many data sets - comparing experimental/simulation data for validation
- Strong analysis, visualization, and workflow are important
  - Strong data management will not just become important—it will become an absolute as we move forward
- FSP scientists must be able to ask “what-if” questions and have the software and hardware infrastructure capable of answering these questions in a timely fashion

Data Analytics, Management, & Visualization

FSP priority research directions

- Managing large-scale input/output volume and data movement
- Real-time monitoring of simulations and run-time metadata generation
- Data analysis at extreme scale
- Visualization of very large datasets
- Experiment-simulation data comparison

V&V and UQ

FSP priority research directions

- Generalized high-order data assimilation and model calibration methods for both data and modeling errors
- Time-dependent global, dynamically adaptable sensitivity analysis and model calibration methods
- UQ methods for multiple coupled models, bridging lower-length scale models into macroscale codes

Programming Models, Frameworks, Tools
FSP priority research directions

- Efficient algorithms and implementations that exploit new multicore, heterogeneous, massively parallel architectures
- New, productive approaches to writing, integrating, validating, and tuning complex applications
- Develop tools for understanding complex application program behavior at scale and for optimizing application performance
- Ensure a migration path from current fusion programming approaches to new ones
- Define common framework tools or components that can be reused in multiple fusion application domains
- Establish methods and systems that enable pervasive fault resilience

FSP Will Provide Synergistic Benefits to ASCR

- The FES community is well-positioned to be a major applications area for demonstrating the benefits of exascale computing.
- FES advances in models, algorithms, and software have and will continue to demonstrate “applications readiness” of benefit to the ASCR mission to “to develop the algorithms, computer programs and hardware that advance scientific research.”
- FES can “show the way” for other applications domains by its prominent role in cross-cutting ASCR-led HPC programs such as SciDAC (Scientific Discovery through Advanced Computing) and INCITE (Innovative and Novel Computational Impact on Experiment).
- Positive impact of availability of impressive suite of well-diagnosed FES experimental facilities & associated large foundational data sets aligned with FSP to help drive UQ (Uncertainty Quantification) R&D involving sensitivity analysis.