ORNL Applications Overview For Interconnects

Jeffrey Vetter

Doug Kothe - with input from dozens

Future Technologies Group

Managed by UT-Battelle
for the Department of Energy

Oak Ridge National Laboratory

Modeling and Simulation at the Exascale

- **Large number of applications drive DOE modeling and simulation efforts**
  - **Bio**
  - **Nanoscience**
  - **Materials**
  - **Climate**
  - **Fusion**
  - **Astro**
  - **Combustion**
  - **Others..**

- **Software**
- **Hardware**
- **Cyberinfrastructure**

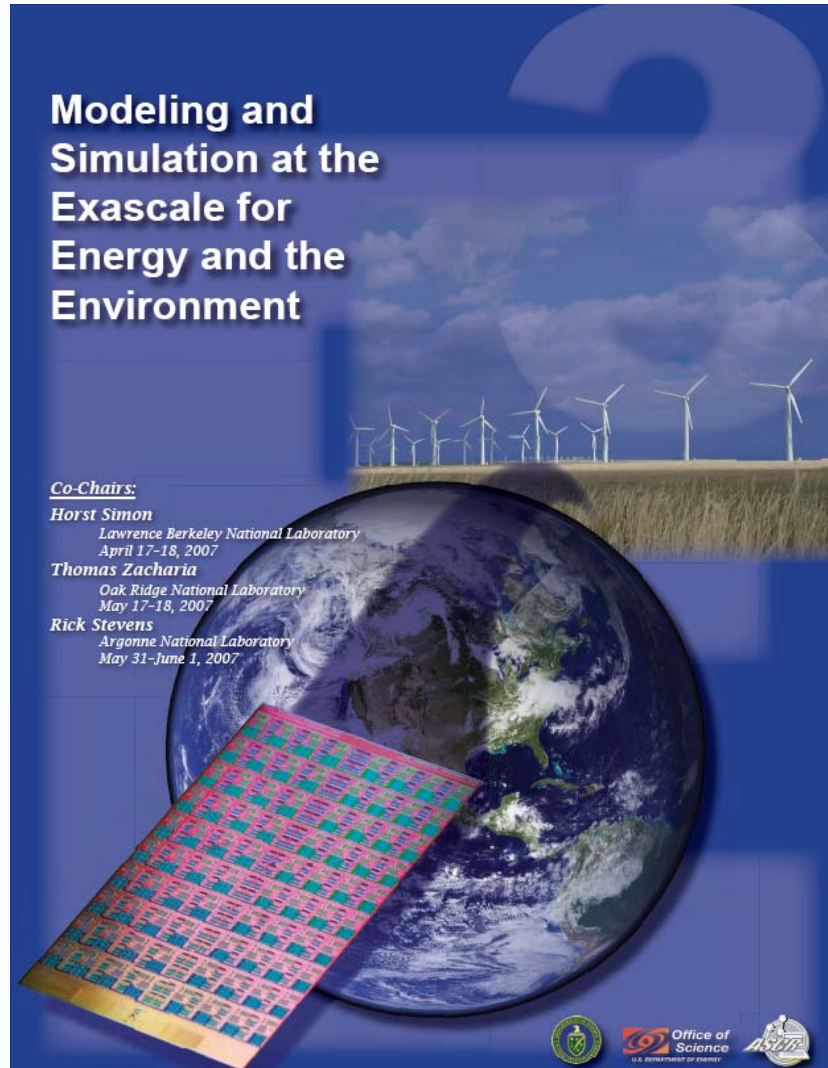

http://www.er.doe.gov/ASCR/ProgramDocuments/TownHall.pdf

Application Codes in 2008
An Incomplete List

- **Astrophysics**
  - CHIMERA, GenASiS, 3DHFEOS, Hahndol, SNe, MPA-FT, SEDONA, MAESTRO, AstroGK
- **Biology**
  - NAMD, LAMMPS
- **Chemistry**
  - CPMD, CP2K, MADNESS, NWChem, Parsec, Quantum Expresso, RMG, GAMESS
- **Nuclear Physics**
  - ANGFMC, MFDn, NUCCOR, HFODD
- **Engineering**
  - Fasel, S3D, Raptor, MFIX, Truchas, BCFD, CFL3D, OVERFLOW, MDOPT
- **High Energy Physics**
  - CPS, Chroma, MILC
- **Fusion**
  - AORSA, GYRO, GTC, XGC
- **Materials Science**
  - VASP, LS3DF, DCA++, QMCPACK, RMG, WL-LSMS, WL-AMBER, QMC
- **Accelerator Physics**
  - Omega3P, T3P
- **Atomic Physics**
  - TDCC, RMPS, TDL
- **Space Physics**
  - Pogorelov
- **Climate & Geosciences**
  - MITgcm, PFLOTRAN, POP, CCSM (CAM, CICE, CLM, POP)

Leadership computing is advancing scientific discovery

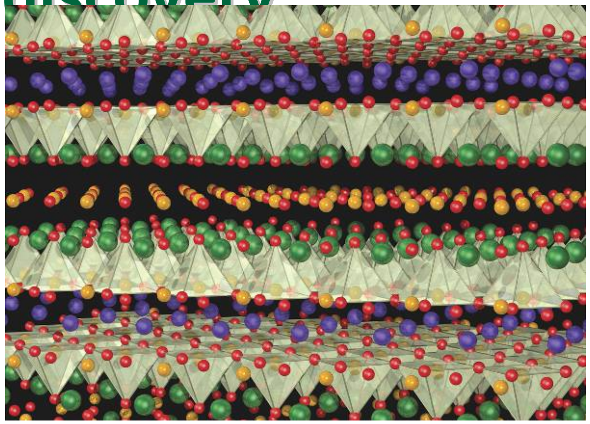

Resolved decades-long controversy about modeling physics of high temperature superconducting cuprates

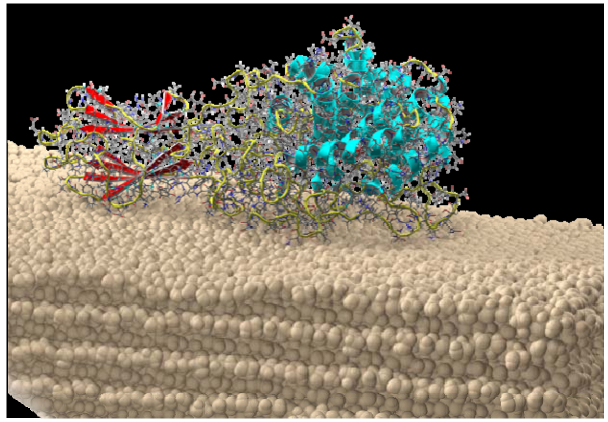

New insights into protein structure and function leading to better understanding of cellulose-to-ethanol conversion

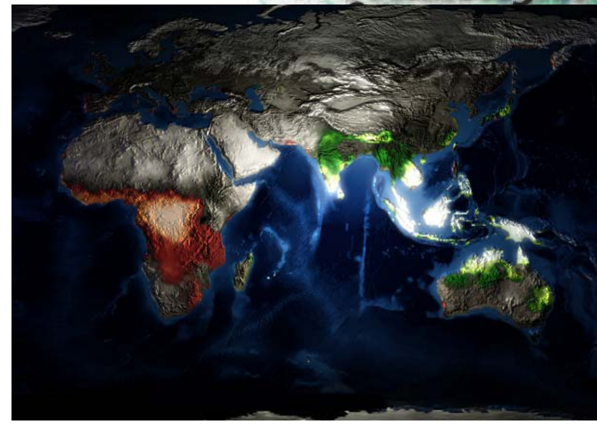

Addition of vegetation models in climate code for global, dynamic CO_2 exploration

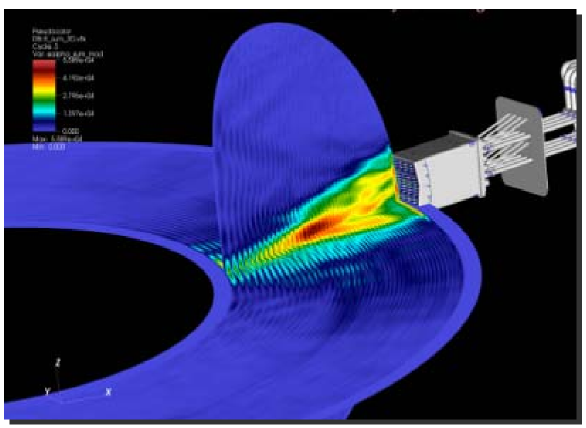

First fully 3D plasma simulations shed new light on engineering superheated ionic gas in ITER

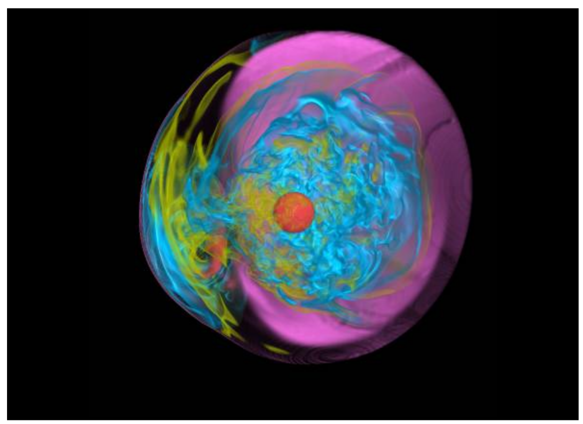

Fundamental instability of supernova shocks discovered directly through simulation

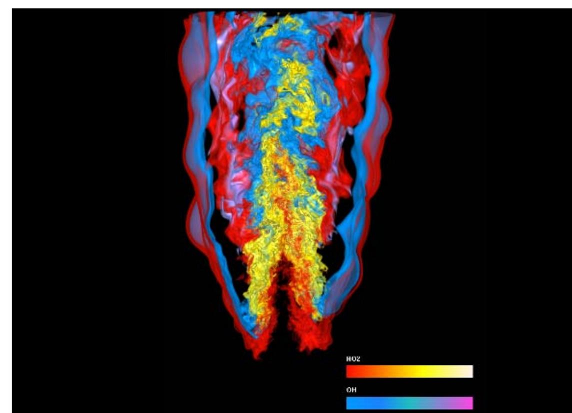

First 3-D simulation of flame that resolves chemical composition, temperature, and flow

Science Workload

Job Sizes and Resource Usage of Key Applications

Code	2007 Resource Utilization (M core-hours)	Projected 2008 Resource Utilization (M core-hours)	Typical Job Size in 2006-2007 (K cores)	Anticipated Job Size in 2008 (K cores)
CHIMERA	2 (under development)	16	0.25 (under development)	>10
GTC	8	7	8	12
S3D	6.5	18	8-12	>15
POP	4.8	4.7	4	8
MADNESS	1 (under development)	4	0.25 (under development)	>8
DCA++	N/A (under development)	3-8	N/A (under development)	4-16 (w/o disorder) >40 (with disorder)
PFLOTRAN	0.37 (under development)	>2	1-2 (under development)	>10
AORSA	0.61	1	15-20	>20

Apps Teams Are Reasonably Adept at Using our Current Systems*

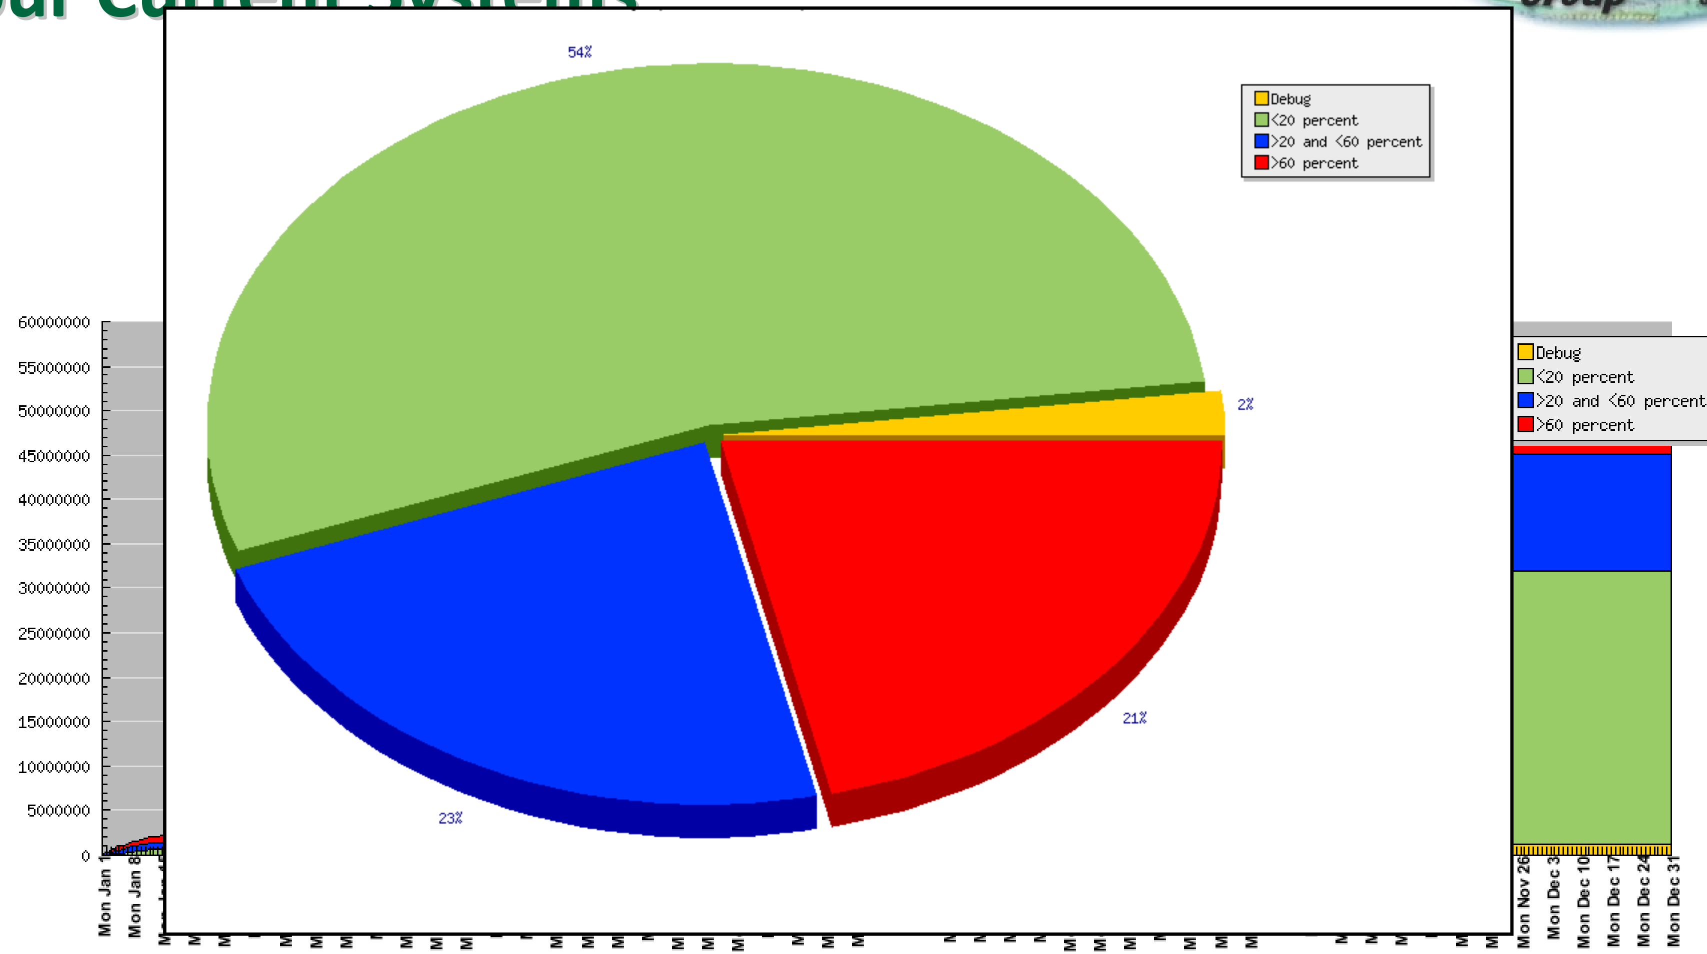

Preparing for the Exascale

Long-Term Science Drivers and Requirements

- We have recently surveyed, analyzed, and documented the science drivers and application requirements envisioned for exascale leadership systems in the 2020 timeframe
- These studies help to
  - Provide a roadmap for the ORNL Leadership Computing Facility
  - Uncover application needs and requirements
  - Focus our efforts on those disruptive technologies and research areas in need of our and the HPC community's attention

Oak Ridge National Laboratory
Managed by UT-Battelle for the Department of Energy

ORNL/TM-2007/232

Science Prospects and Benefits with Exascale Computing

Oak Ridge National Laboratory
Managed by UT-Battelle for the Department of Energy

ORNL/TM-2007/238

Scientific Application Requirements for Leadership Computing at the Exascale

December 2007

Prepared by
Computing Requirements Team
National Center for Computational Sciences

Science Prospects and Benefits with High End Computing (EF?) in the Next Decade

Opportunity	Key application areas	Goal and benefit
Materials science	Nanoscale science, manufacturing, and material lifecycles, response and failure	Design, characterize, and manufacture materials, down to the nanoscale, tailored and optimized for specific applications
Earth science	Weather, carbon management, climate change mitigation and adaptation, environment	Understand the complex biogeochemical cycles that underpin global ecosystems and control the sustainability of life on Earth
Energy assurance	Fossil, fusion, combustion, nuclear fuel cycle, chemical catalysis, renewables (wind, solar, hydro), bioenergy, energy efficiency, power grid, transportation, buildings	Attain, without costly disruption, the energy required by the United States in guaranteed and economically viable ways to satisfy residential, commercial, and transportation requirements
Fundamental science	High-energy physics, nuclear physics, astrophysics, accelerator physics	Decipher and comprehend the core laws governing the Universe and unravel its origins
Biology and medicine	Proteomics, drug design, systems biology	Understand connections from individual proteins through whole cells into ecosystems and environments
National security	Disaster management, homeland security, defense systems, public policy	Analyze, design, stress-test, and optimize critical systems such as communications, homeland security, and defense systems; understand and uncover human behavioral systems underlying asymmetric operation environments
Engineering design	Industrial and manufacturing processes	Design, deploy, and operate safe and economical structures, machines, processes, and systems with reduced concept-to-deployment time

Science Case: Climate

Mitigation: Evaluate strategies and inform policy decisions for climate stabilization; 100-1000 year simulations
Adaptation: Decadal forecasts & region impacts; prepare for committed climate change; 10-100 year simulations

- **250 TF**
  - Mitigation: Initial simulations with dynamic carbon cycle and limited chemistry
  - Adaptation: Decadal simulations with high-resolution ocean (1/10°)
- **1 PF**
  - Mitigation: Full chemistry, carbon/nitrogen/sulfur cycles, ice-sheet model, multiple ensembles
  - Adaptation: High-resolution atmosphere (1/4°), land, and sea ice, as well as ocean
- **Sustained PF**
  - Mitigation: Increased resolution, longer simulations, more ensembles for reliable projections; coupling with socio-economic and biodiversity models
  - Adaptation: Limited cloud-resolving simulations, large-scale data assimilation
- **1 EF**
  - Mitigation: Multi-century ensemble projections for detailed comparisons of mitigation strategies
  - Adaptation: Full cloud-resolving simulations, decadal forecasts of regional impacts and extreme-

Resolve clouds, forecast weather & extreme events, provide quantitative mitigation strategies

Current requirements

Application	Structured	Unstructured	FFT	Dense	Sparse	Particles	Monte Carlo
Molecular		X	X	X		X	
Nanoscience	X			X		X	X
Climate	X		X		X	X	
Environment	X	X			X		
Combustion	X						
Fusion	X		X	X	X	X	X
Nuc. energy		X		X	X		
Astrophysics	X	X		X	X	X	
Nuc. physics				X			
Accelerator		X			X		
QCD	X						X
#X	7	5	3	6	6	5	3

Exascale requirements

Application	Structured	Unstructured	FFT	Dense	Sparse	Particles	Monte Carlo
Molecular		X	X	X		X	X
Nanoscience	X			X		X	X
Climate	X		X		X	X	X
Environment	X	X			X	X	X
Combustion	X			X		X	
Fusion	X	X	X	X	X	X	X
Nuc. energy		X		X	X		
Astrophysics	X	X		X	X	X	
Nuc. physics				X			
Accelerator		X			X		
QCD	X						X
#X	7	6	3	7	6	7	6

Application Performance Limited by Interconnect Performance

- **Bandwidth limits GYRO performance**
  - Significant increase in bandwidth dramatically improved GYRO performance
- **Latency limits POP performance**
  - Strong scaling problem
  - CAF mapping exploited Globally addressable memory to provide dramatic scaling improvements

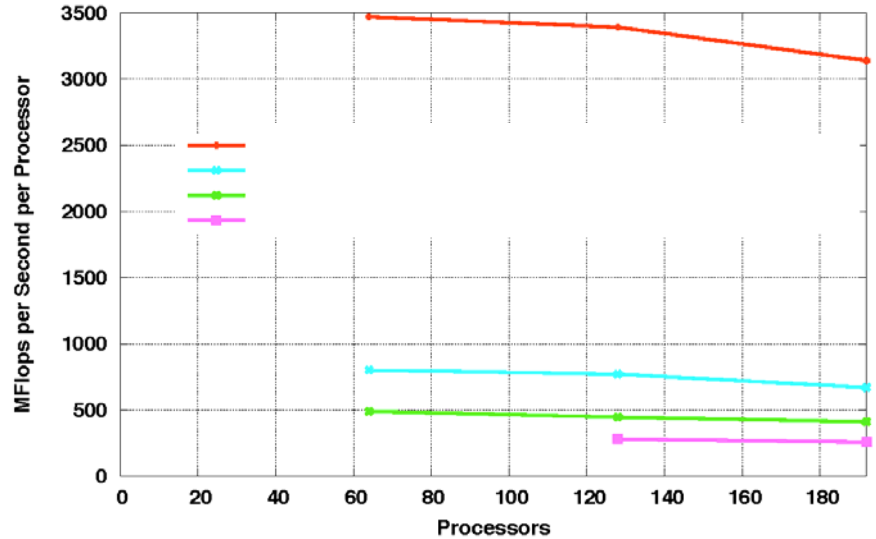

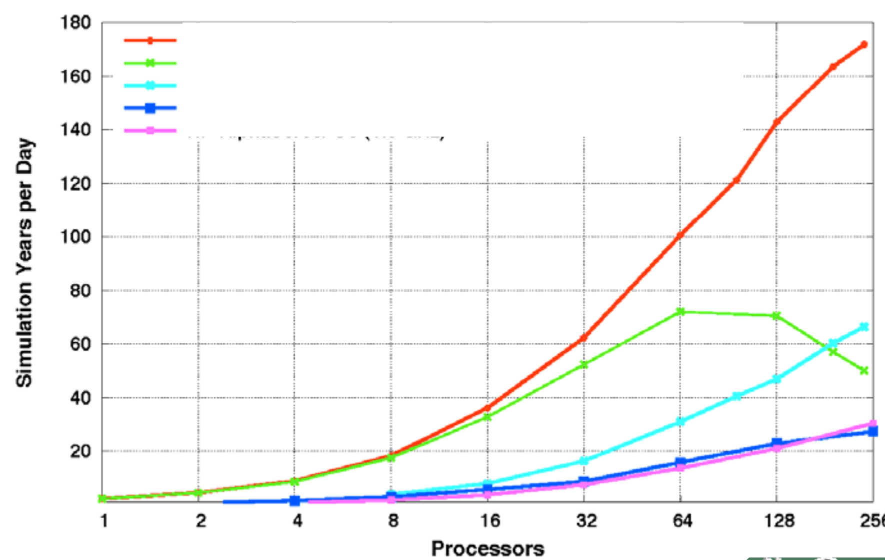

Empirical workload studies

Future Technologies Group ornl

- **Generated traces for HPCS studies**
  - http://ft.ornl.gov/~vetter/public/hpcs
- **Update underway**

ORNL Traces of HPCS Applications

1. Introduction, trace format, and application descriptions
2. Overview (short presentation)
3. Traces (hpc-traces.2005-07-15.tar.bz2) *LARGE FILE (~ 100MB compressed)*
4. UMT2K @ 1024, 2048 (~100MB)
5. GTC @ 1024 (~10MB)
6. Events definitions
7. INFO file
8. Errata file

Please send questions to Jeffrey Vetter (vetter@computer.org).

Application Characterization - Traffic

- DOE SC and NNSA applications have a wide range of requirements for interconnects
- Initial analysis suggests that both bandwidth and latency are equally important
- Logical application topologies range from static, regular structures to dynamic, irregular structures
- Requirements for new applications/models are not well understood
  - New algorithms at scale
  - One sided communication
  - Remote atomic operations

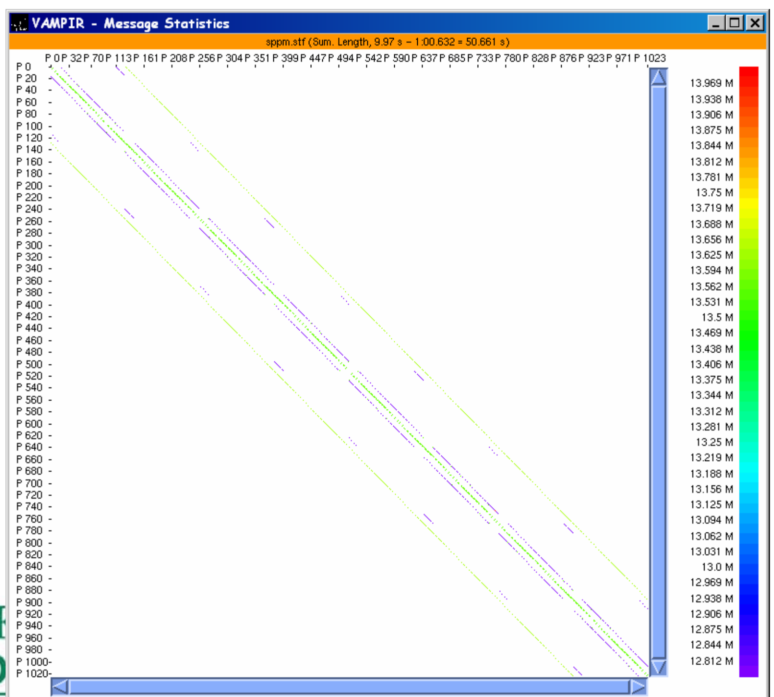

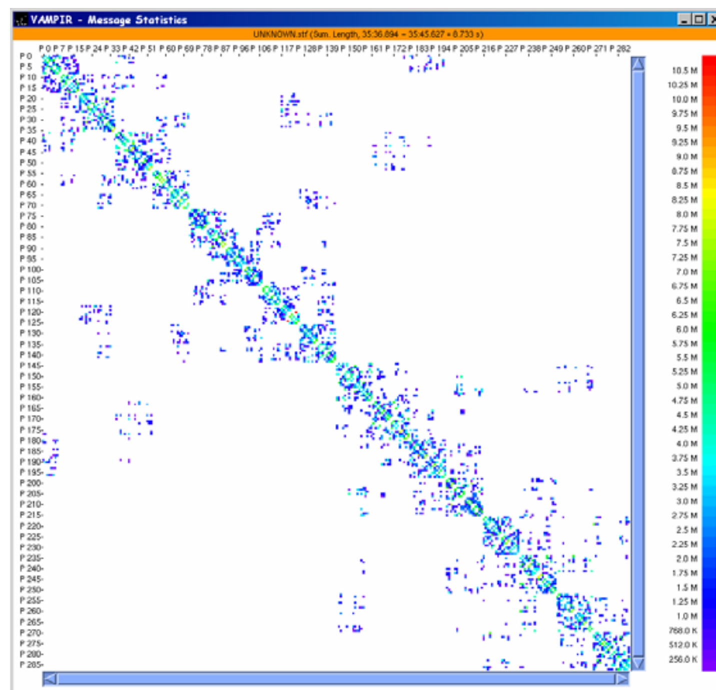

Collectives – sweet spot

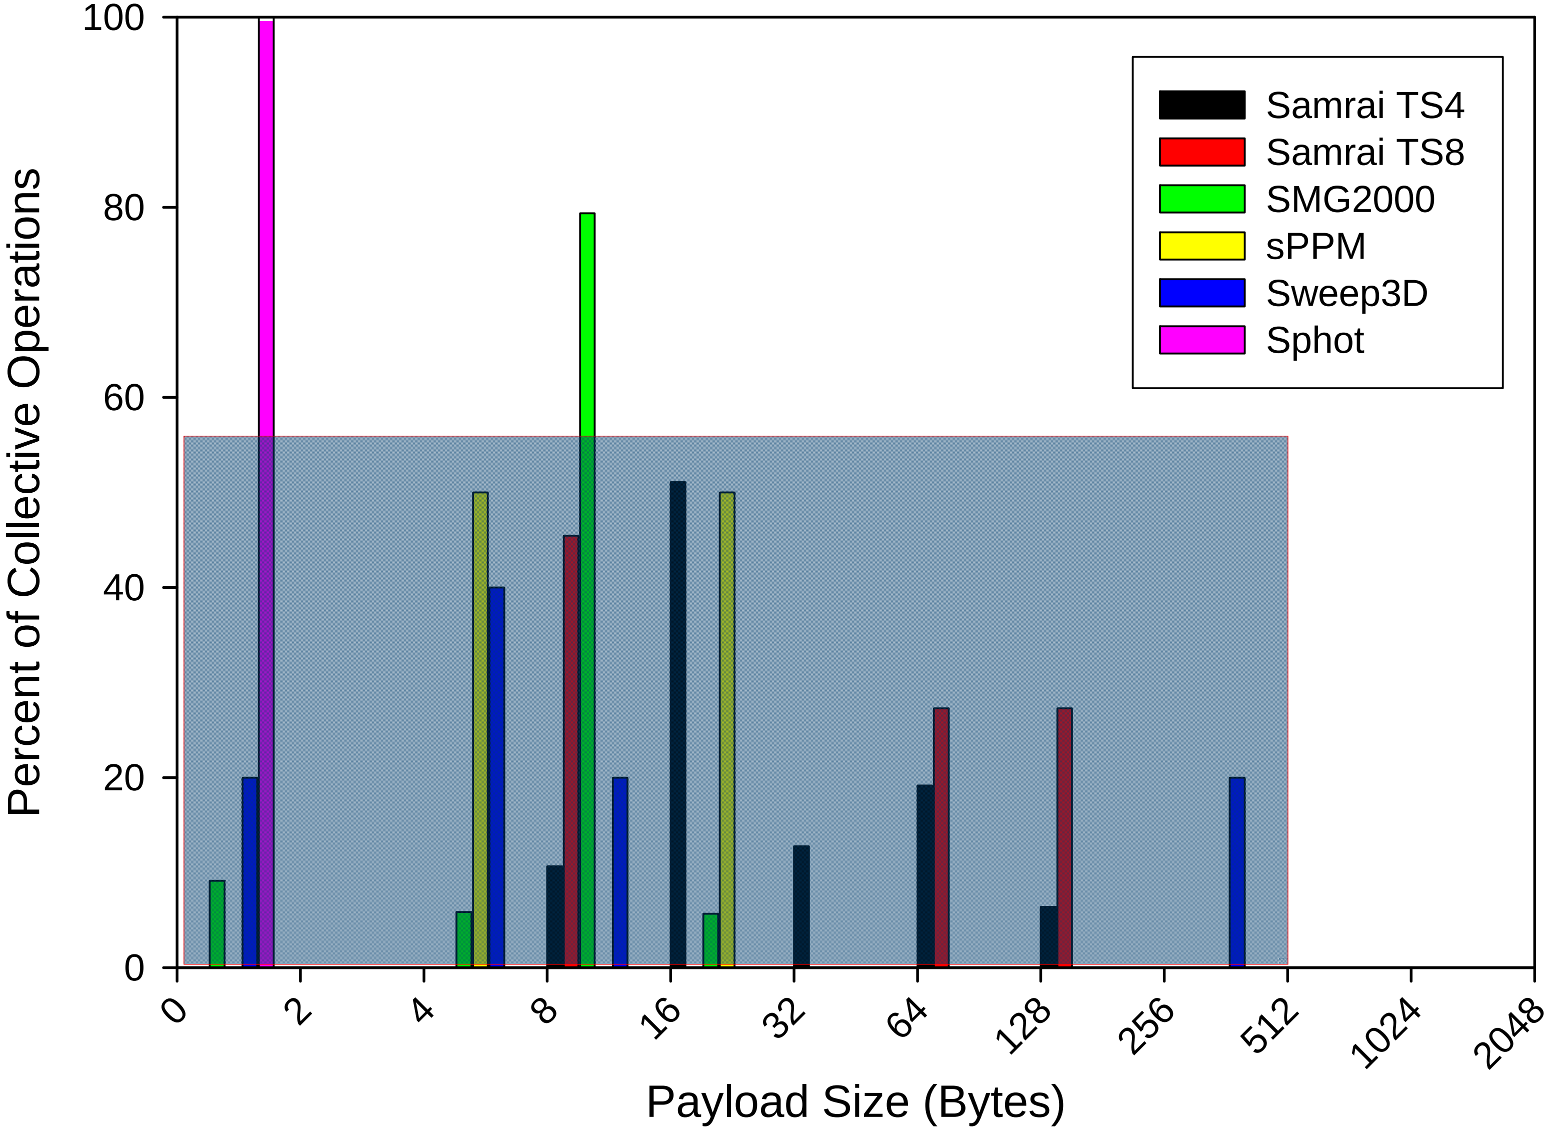

J.S. Vetter and F. Mueller, "Communication Characteristics of Large-Scale Scientific Applications for Contemporary Cluster Architectures," Proc. International Parallel and Distributed Processing Symposium (IPDPS), 2002.

Summarized Empirical Data Show Range of Application Requirements

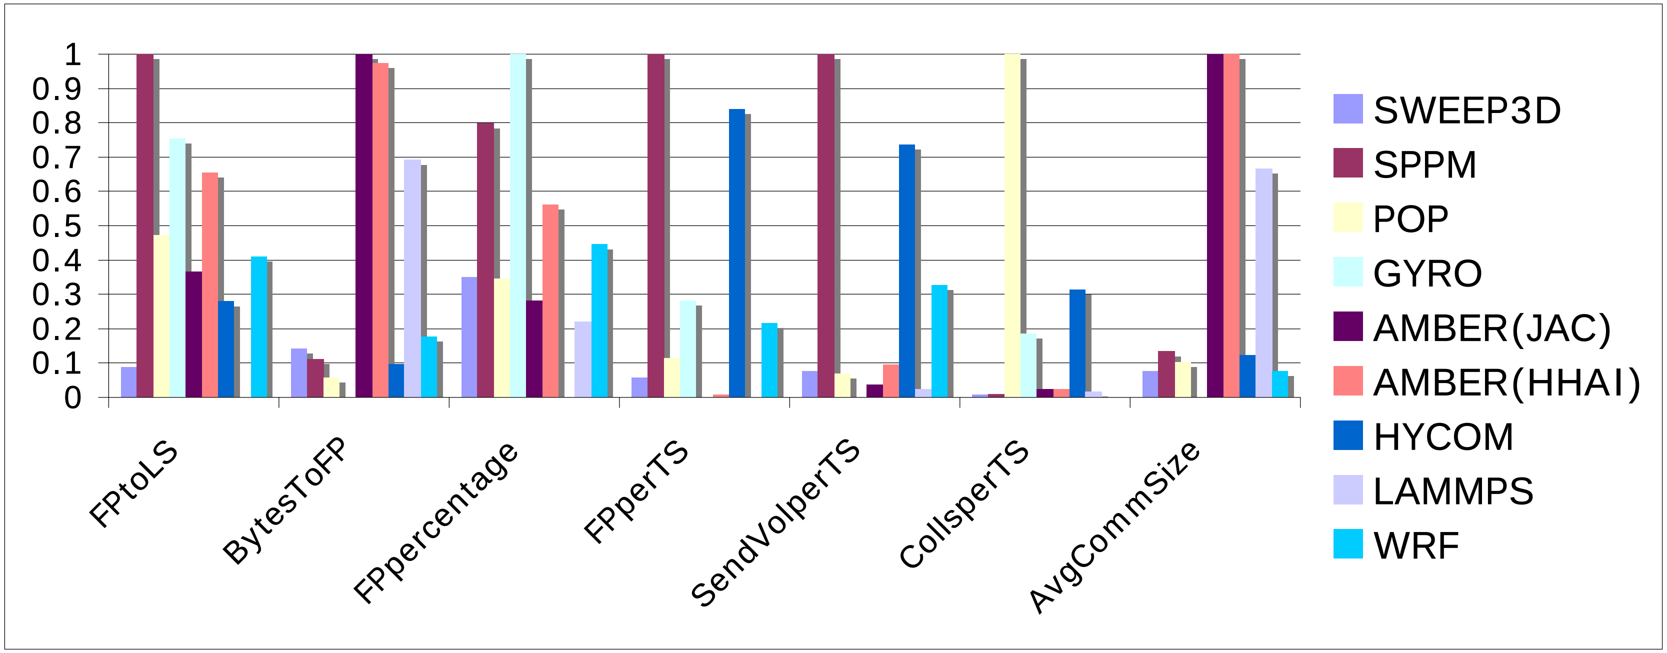

Each metric normalized to range of zero to one.

Symbolic Modeling

- **Analytical modeling of applications can provide useful analysis of application communication requirements**

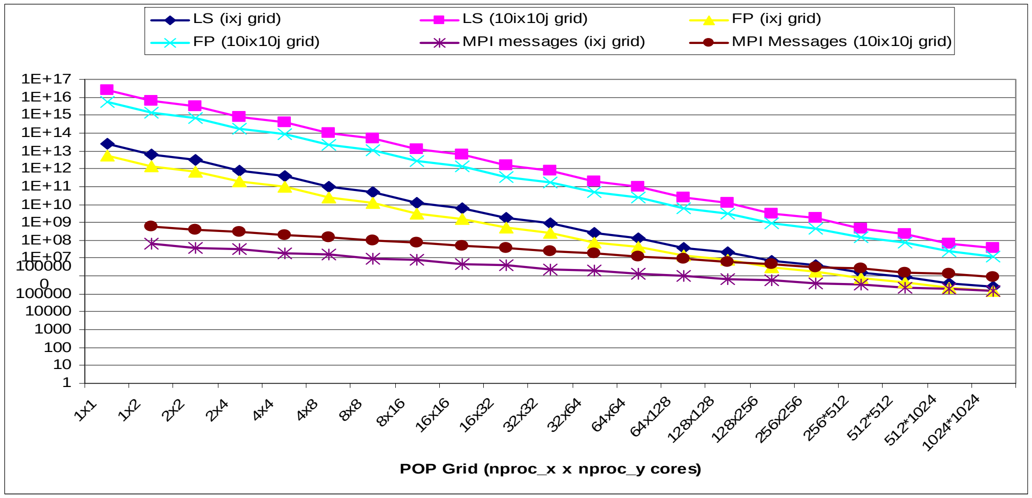

- **Extend PERI modeling assertions to generate scalable communication patterns of real workloads**
  - Generate workloads of any scale
  - Consumed by simulator

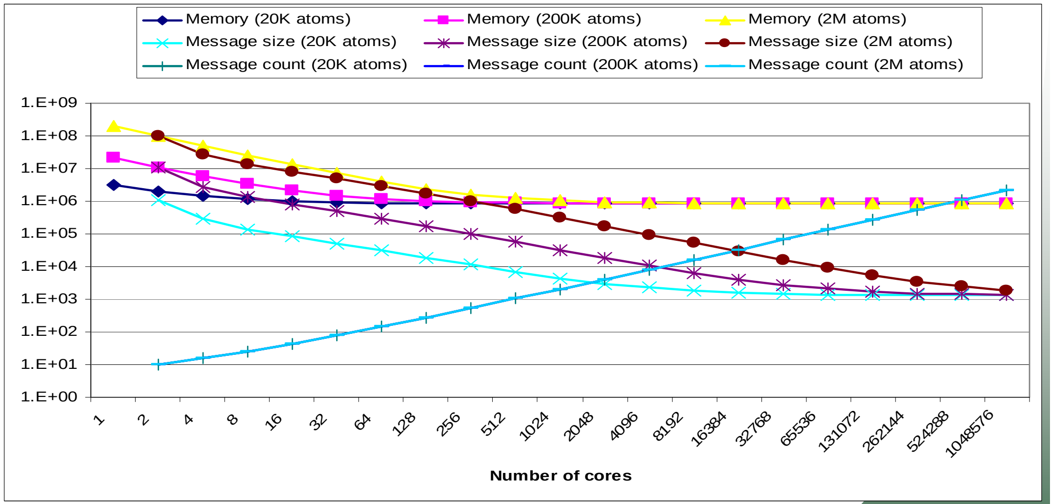

Random thoughts

- **What data is needed?**
  - **How do we distribute data to give architects 'real workloads?'**
- **In addition to better performance**
  - **Job scheduling and application placement**
  - **Performance diagnostics**