



Panel: Many-Task Computing meets Exascales

Chairs

Ioan Raicu, Illinois Institute of Technology & Argonne National Laboratory

Ian Foster, University of Chicago & Argonne National Laboratory

Yong Zhao, University of Electronic Science and Technology of China, China

**ACM MTAGS 2011
November 14th, 2011**

Panelists

- **Dr. Dan Reed**

- Corporate Vice President of Technology Policy and Strategy, Microsoft Research
- Leader of the eXtreme Computing Group (XCG), Microsoft Research

- **Dr. Jack Dongarra**


- University Distinguished Professor of EECS, University of Tennessee
- Distinguished Research Staff of CSM, Oak Ridge National Laboratory
- Director, Innovative Computing Laboratory, University of Tennessee
- Director, Center for Information Technology Research, University of Tennessee

- **Dr. Dan S. Katz**

- Director of Science, TeraGrid GIG
- Area Co-director for Applications, Open Grid Forum
- Senior Fellow in the Computation Institute, University of Chicago

- **Dr. David Abramson**


- Professor of CS, Monash University, Australia
- Director, Monash eScience and Grid Engineering Lab, Monash University
- ARC Professorial Fellow, Faculty of Information Technology, Monash University



Australia
University, Australia

Dan S. Katz:

Many Task Computing for Exascale



- DOE X-stack project: ExM: System support for extreme-scale, many-task applications
- Extending the Swift programming model (<http://www.ci.uchicago.edu/swift>)
 - Script-based, no explicit DAG – control flow is based on data flow (“futures”)
 - Swift engine reads script, turns all function calls into parallel tasks; dispatches tasks when inputs become available, runs tasks on dynamically acquired resources
 - Execution is RPC-like: transfer input, run code, transfer output
- in these areas:
 - Tasks
 - JETS: dynamically run multiple MPI tasks on an HPC resource (treats MPI jobs as atomic functions)
 - Looking at tasks that are functions rather than executables: extending the programming model to in-memory activities; pass data by message in addition to via files
 - Execution
 - Developing Turbine distributed execution engine: fully parallelizes program *evaluation*
 - I/O
 - MosaStore: intermediate file storage; use files for message passing, but stripe them across RAMdisk (or fast local filesystem) on nodes (single distributed filesystem w/ shared namespace), backing store in shared file system, potentially cache in the middle
 - AME: intermediate file storage; use files for message passing, but store them in RAMdisk on nodes where written (multiple filesystems w/ multiple namespaces), copy to new nodes when needed for reading; integrates file exchange and task dispatching
 - Collective data management (build primitives for common operations to bundle I/O ops and take advantage of HPC/site hardware, e.g. BGP I/O nodes), hardware broadcast, fast scatter/gather

Jack Dongarra:

Critical Issues at Peta & Exascale for Algorithm and Software Design

- Synchronization-reducing algorithms
 - Break Fork-Join model
- Communication-reducing algorithms
 - Use methods which have lower bound on communication
- Mixed precision methods
 - 2x speed of ops and 2x speed for data movement
- Autotuning
 - Today's machines are too complicated, build "smarts" into software to adapt to the hardware
- Fault resilient algorithms
 - Implement algorithms that can recover from failures/bit flips
- Reproducibility of results
 - Today we can't guarantee this. We understand the issues, but some of our "colleagues" have a hard time with this.

Dan Reed

Many Task Computing, Clouds and Exascale

- Technical computing rides the commodity ecosystem
 - The tail follows the dog, not the other way around
- Commodities
 - Hardware (by volume): low power multicore and SoCs
 - Software: toolkits, domain specific languages, cloud services
 - Data: multidisciplinary, high volume, third parties
- Clouds, many task computing and toolkits
 - Let scientists be scientists, not programmers
 - Raise the level of abstraction
 - Trade performance for usability and expressiveness
 - Hide details of failures, scheduling, ...



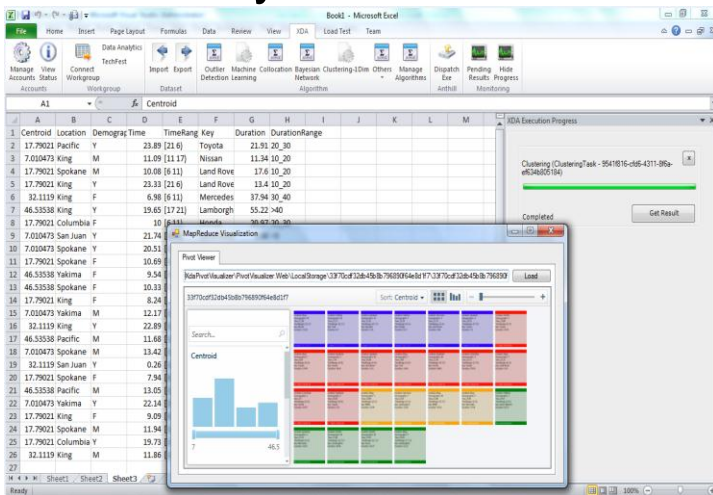
**Clouds ,
Data
and
Devices**

Dan Reed

Clouds and Many Task Technical Computing

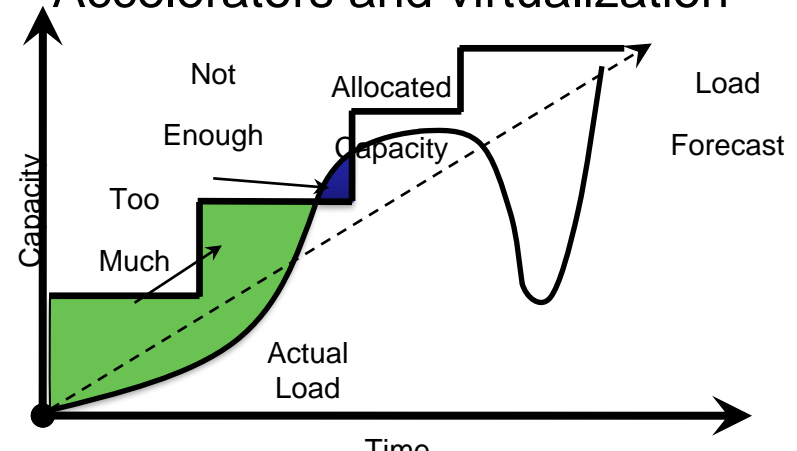
Similar technology issues

- Node and system architectures
- Communication fabrics
- Storage systems and analytics
- Physical plant and operations
- Programming models
- Reliability and resilience



Cultural Differences

- Resource scheduling
 - Gang scheduling and provisioning
- Network performance
 - Virtualization and networks
 - High performance interconnects
- Accelerators and virtualization



Panel Overview

- **What is Exascale Computing, and what are the major challenges?**
- **Is a Data-Centric Paradigm inevitable?**
- **Is Many-Task Computing a viable programming paradigm for exascale computing?**

What is Exascale Computing, and what are the major challenges?

- What can we expect exascale systems to look like?
- What are the major challenges of exascale computing?

Is a Data-Centric Paradigm inevitable?

- Exascale computing will bring new fundamental challenges in:
 - How we build computing systems (e.g. architecture),
 - How we manage them (e.g. operating systems, middleware, and storage)
 - How we program them (e.g. programming languages)
- **Process-centrics → Data-centric Computing**
 - Exposing data locality throughout the computing stack

Is Many-Task Computing a viable programming paradigm for exascale computing?

- MTC might address some of the HPC shortcomings at extreme scales (e.g. reliability, programmability)
- Loosely coupled computations (asynchronous vs synchronous communication),
- Reliability mechanisms (implicit independent checkpointing vs. explicit synchronized checkpointing),
- Data-centric vs. process-centric

More Information

- MTAGS 2011 Website:
 - <http://datasys.cs.iit.edu/events/MTAGS11/>
- Panel info:
 - <http://datasys.cs.iit.edu/events/MTAGS11/panel.html>
- Workshop program (7 exciting talks in the PM)
 - <http://datasys.cs.iit.edu/events/MTAGS11/program.html>
- Prize giveaway (win an Apple iPad 2):
 - <http://datasys.cs.iit.edu/events/MTAGS11/prize.html>
- Contact:
 - mtags11-chairs@datasys.cs.iit.edu