JACK DONGARRA
UNIVERSITY OF TENNESSEE
AND
OAK RIDGE NATIONAL LABORATORY

INTERNATIONAL CONFERENCE ON SCIENTIFIC COMPUTING AT EXTREME SCALES
ICSC 2014
MAY 8, 2014

INTERNATIONAL PANEL

Yutong Lu, NUDT

Satoshi Matsuoka, Tokyo Tech

Jack Dongarra, UTK

Thomas Shultess, ETH

Jacqueline Chen, SNL

QUESTIONS FOR THE PANELISTS

1. What are the driver-applications that motivate investment in technology towards achieving exascale computing?
2. What are the major challenges that need to be resolved before practical and effective exascale computation can be realized?
3. Briefly describe one key programs in which you are engaged that may contribute to useful exascale systems and methods.
4. In the broad international arena, what kinds of outreach can yield to cooperative sharing of advances? What would you like to see become widely available and common worldwide?

BROAD COMMUNITY SUPPORT AND DEVELOPMENT OF EXASCALE COMPUTING SINCE 2007

http://science.energy.gov/ascr/news-and-resources/program-documents/

- **Town Hall Meetings April-June 2007**
- **Scientific Grand Challenges Workshops Nov, 2008 – Oct, 2009**
  - **Climate Science (11/08)**
  - **High Energy Physics (12/08)**
  - **Nuclear Physics (1/09)**
  - **Fusion Energy (3/09)**
  - **Nuclear Energy (5/09)**
  - **Biology (8/09)**
  - **Material Science and Chemistry (8/09)**
  - **National Security (10/09)**
  - **Cross-cutting technologies (2/10)**
- **Exascale Steering Committee**
  - **"Denver" vendor NDA visits (8/09)**
  - **SC09 vendor feedback meetings**
  - **Extreme Architecture and Technology Workshop (12/09)**
  - **Applied Math at Exascale(3/14)**
  - **Top 10 Exascale Challenges (2/14)**
- **International Exascale Software Project**
  - **Santa Fe, NM (4/09); Paris, France (6/09); Tsukuba, Japan (10/09); Oxford (4/10); Maui (10/10); San Francisco (4/11); Cologne (10/11); Kobe (4/12)**
  - **BDEC Charleston (4/13) Kobe; (2/14)**

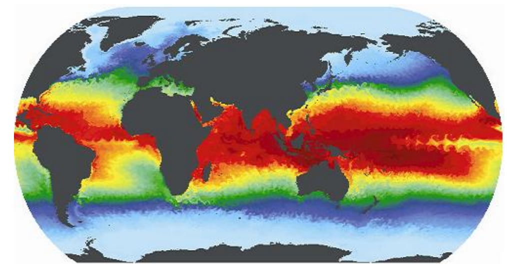

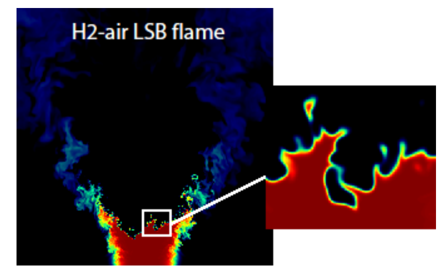

Mission Imperatives

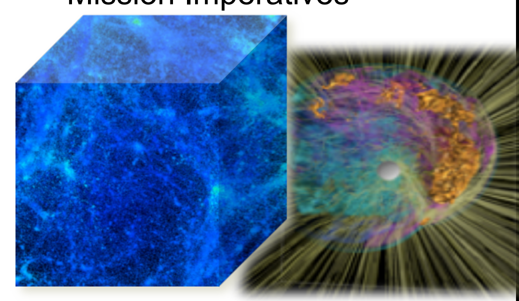

Fundamental Science

BDEC

Exec Committee
Pete Beckman
Jean-Yves Berthou
Jack Dongarra
Yutaka Ishikawa
Satoshi Matsuoka
Philippe Ricoux

SUPERCOMPUTERS TOUCH EVERYONE WITH WEATHER FORECASTING

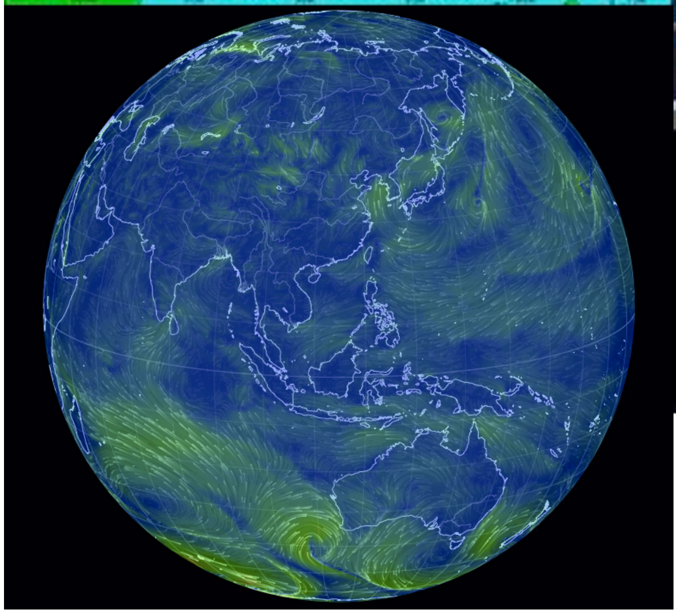

QUESTIONS FOR THE PANELISTS

1. What are the driver-applications that motivate investment in technology towards achieving exascale computing?
2. What are the major challenges that need to be resolved before practical and effective exascale computation can be realized?
3. Briefly describe one key programs in which you are engaged that may contribute to useful exascale systems and methods.
4. In the broad international arena, what kinds of outreach can yield to cooperative sharing of advances? What would you like to see become widely available and common worldwide?

LAST CENTURY'S DLA SOFTWARE

Software/Algorithms follow hardware evolution in time		
LINPACK (70's) (Vector operations)	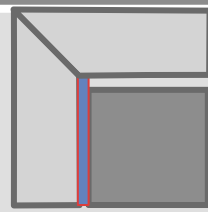	Rely on - Level-1 BLAS operations
LAPACK (80's) (Blocking, cache friendly)	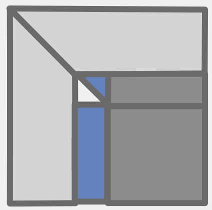	Rely on - Level-3 BLAS operations
ScaLAPACK (90's) (Distributed Memory)		Rely on - PBLAS Mess Passing

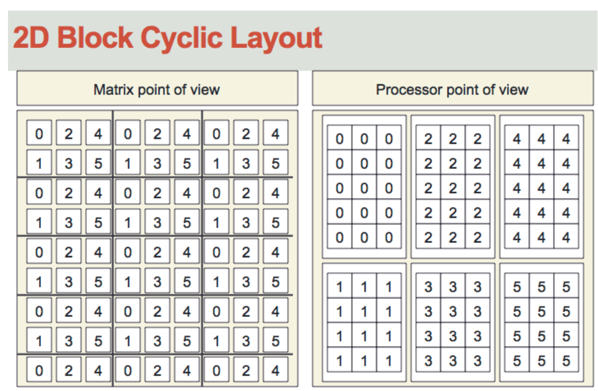

A NEW GENERATION OF DLA SOFTWARE

Office of Science
U.S. Department of Energy

Software/Algorithms follow hardware evolution in time		
LINPACK (70's) (Vector operations)	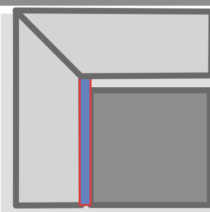	Rely on - Level-1 BLAS operations
LAPACK (80's) (Blocking, cache friendly)	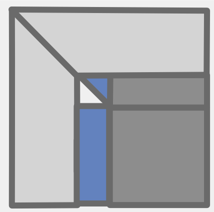	Rely on - Level-3 BLAS operations
ScaLAPACK (90's) (Distributed Memory)		Rely on - PBLAS Mess Passing
PLASMA New Algorithms (many-core friendly)	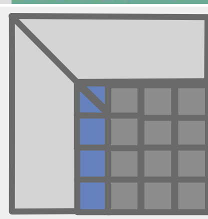	Rely on - a DAG/scheduler - block data layout - some extra kernels
MAGMA Hybrid Algorithms (heterogeneity friendly)	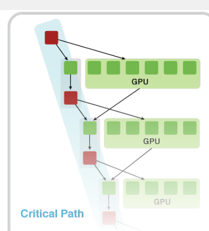	Rely on - hybrid scheduler - hybrid kernels

Parallelization of QR Factorization

Parallelize the update:

- Easy and done in any reasonable software.
- This is the $2/3n^3$ term in the FLOPs count.
- Can be done "efficiently" with LAPACK+multithreaded BLAS

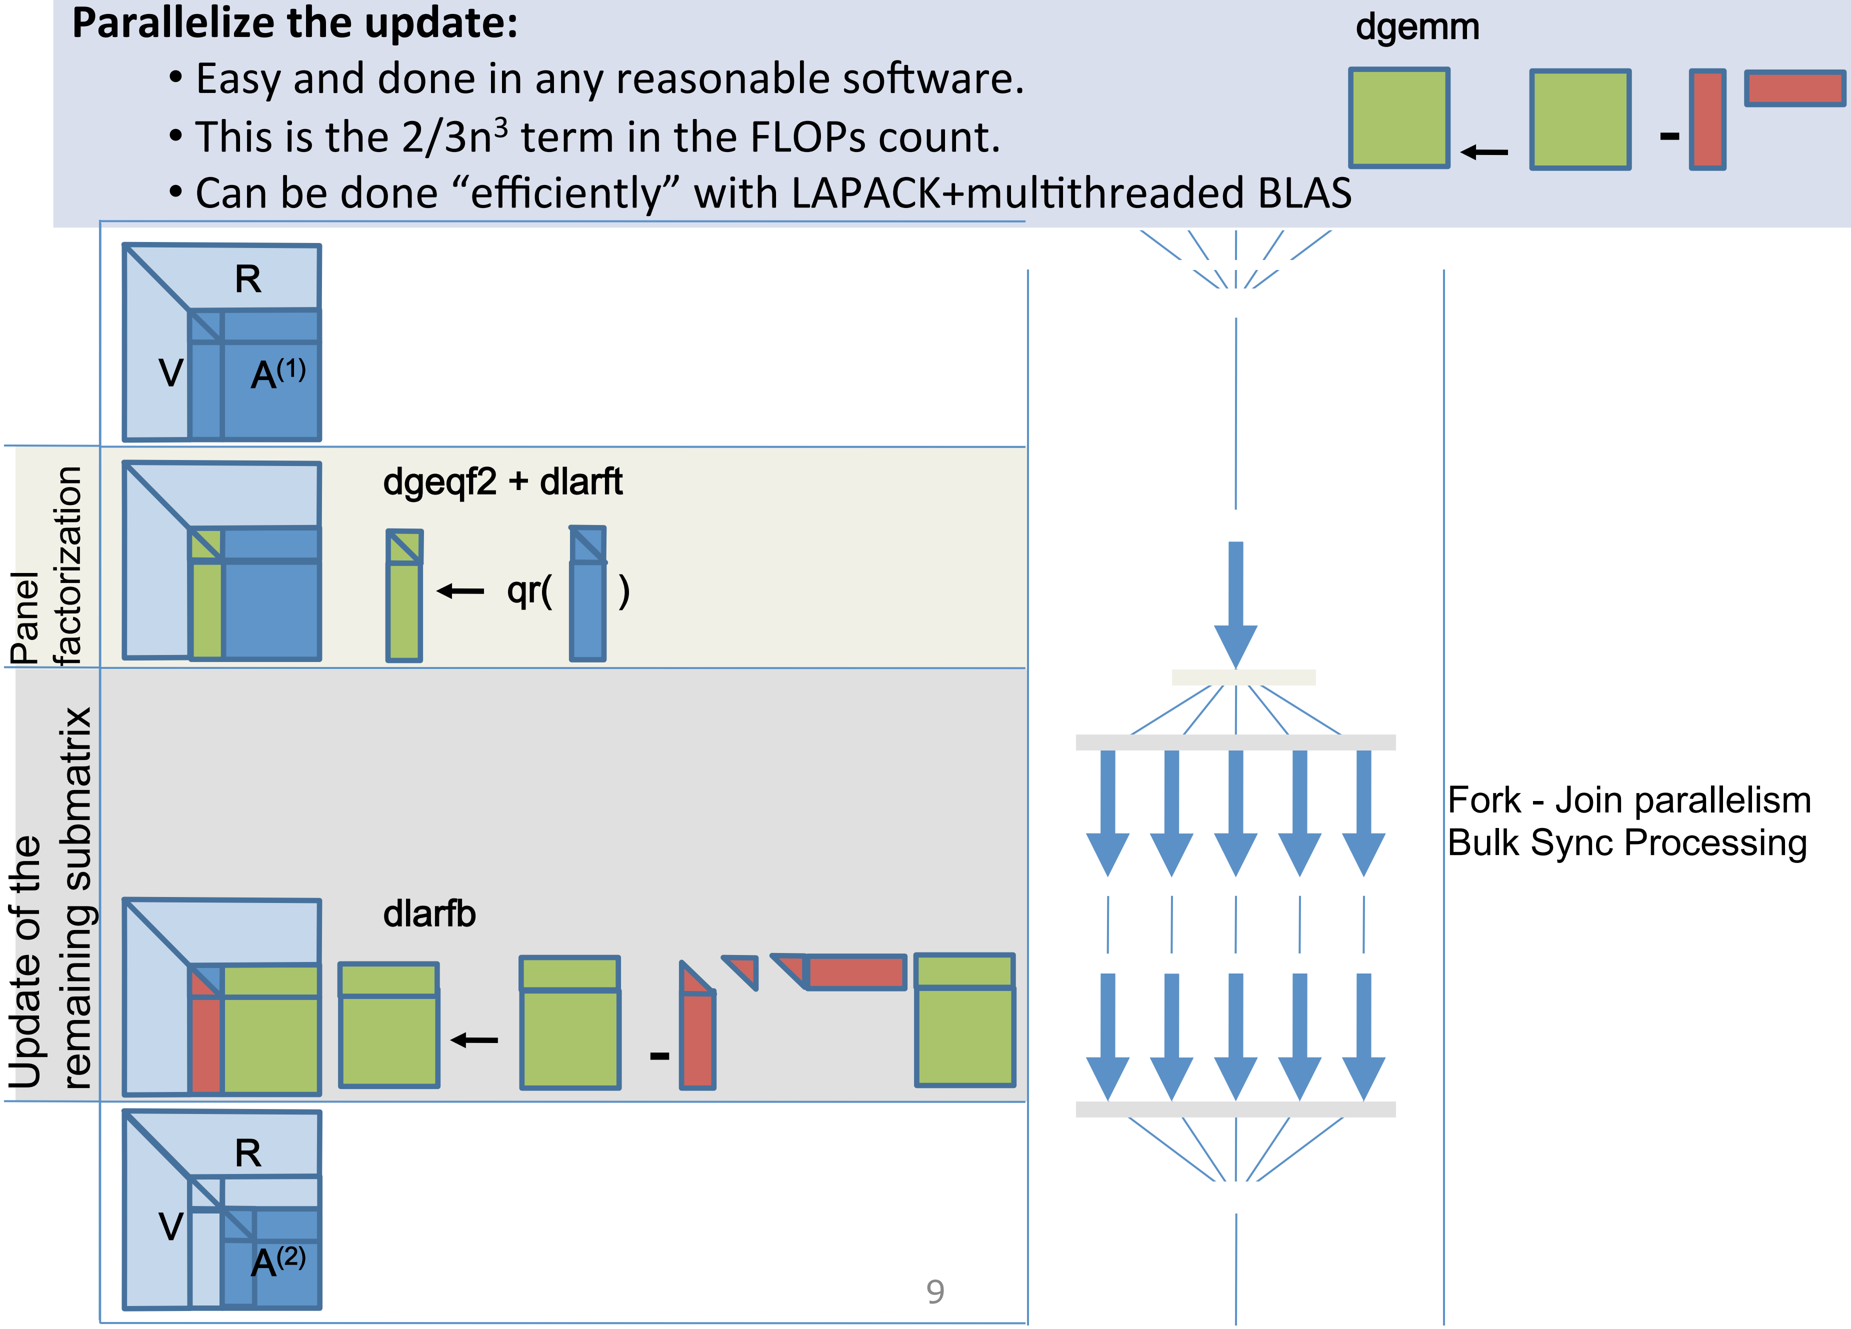

PLASMA: Parallel Linear Algebra s/w for Multicore Architectures

- **Objectives**
  - High utilization of each core
  - Scaling to large number of cores
  - Synchronization reducing algorithms
- **Methodology**
  - Dynamic DAG scheduling (QUARK)
  - Explicit parallelism
  - Implicit communication
  - Fine granularity / block data layout
- **Arbitrary DAG with dynamic scheduling**

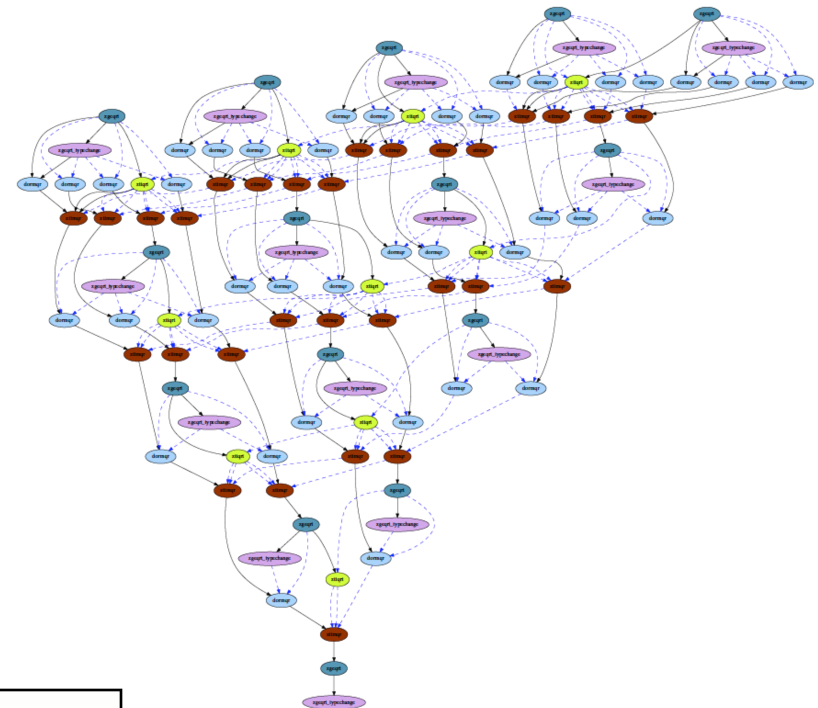

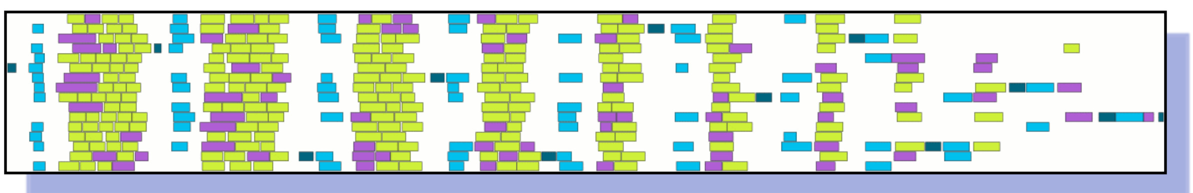

Fork-join parallelism
Notice the synchronization penalty in the presence of heterogeneity.

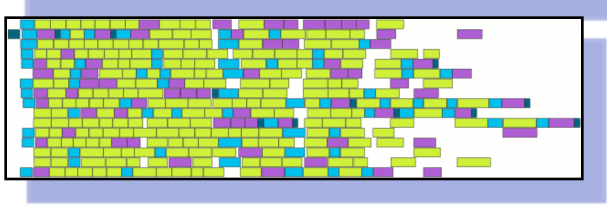

DAG scheduled parallelism

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.

QUESTIONS FOR THE PANELISTS

1. What are the driver-applications that motivate investment in technology towards achieving exascale computing?
2. What are the major challenges that need to be resolved before practical and effective exascale computation can be realized?
3. Briefly describe one key programs in which you are engaged that may contribute to useful exascale systems and methods.
4. **In the broad international arena, what kinds of outreach can yield to cooperative sharing of advances? What would you like to see become widely available and common worldwide?**

International Extreme Scale Projects

- The European Union, Japan, China and the United States have each launched extreme computing projects, each with differing emphases on hardware technologies, system software, algorithms and applications.
- Global competition for advanced computing leadership continues, there is active international collaboration.
  - The International Exascale Software Project (IESP) and Big Data and Extreme Computing (BDEC) efforts are two such example.
- Both competition and global research collaboration will be necessary to address design, test and deploy exascale class computing and data analysis capabilities.

International Exascale Software Project

INTERNATIONAL EXASCALE SOFTWARE PROJECT

- Overall goal:
  - To develop a plan for producing a software infrastructure capable of supporting exascale applications
- 8 Meetings:
  1. Santa Fe, NM, US, April 2009
  2. Paris, France, June 2009
  3. Tsukuba, Japan, October 2009
  4. Oxford, UK, April 2010
  5. Maui, HI, US, October 2010
  6. San Francisco, US, April 2011
  7. Cologne, Germany, October 2011
  8. Kobe, Japan April 2012
- SC08 (Austin), SC09 (Portland), ISC10 (Hamburg), SC10 (New Orleans), ISC11 (Hamburg), SC11 (Seattle)

International Community Effort

- We believe this needs to be an international collaboration for various reasons including:
  - The scale of investment
  - The need for international input on requirements
  - US, Europeans, Asians, and others are working on their own software that should be part of a larger vision for HPC.
  - No global evaluation of key missing components
  - Hardware features are coordinated with software development

IESP Makeup

Roadmap available at:
www.exascale.org



- Attendees from universities, research institutes, government, funding agencies, research councils, hardware and software vendors, industry
- 65 - 85 participants per workshop
- Rough distribution per meeting is rather constant:
  - 70% universities/research institutes
  - 15% vendors/industry
  - 15% government/funding agencies

Broad Goals

- Develop ways for EC and BD communities to work closely together
- Develop a ***shared*** vision for the future
  - Architecture & Facilities
  - Software
  - Applications
- Identify key research areas and develop a roadmap for building and extending BDEC capabilities in support of science

Exec Committee
Pete Beckman
Jean-Yves Berthou
Jack Dongarra
Yutaka Ishikawa
Satoshi Matsuoka
Philippe Ricoux

BIG DATA AND EXTREME-SCALE COMPUTING