

Fujitsu's Activities in Capability Computing

Ross Nobes

Fujitsu Laboratories of Europe

Contents



- Fujitsu's Renewed Interest in Capability Computing
- Open Initiatives
- Current Co-Design Vehicles
- Towards Exascale
- Summary

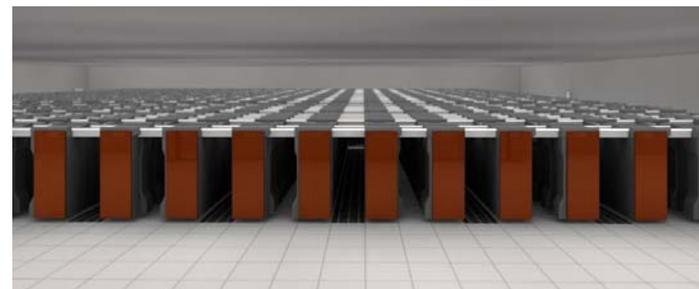
Fujitsu and Capability Computing



- RIKEN and Fujitsu are jointly developing the 'K Computer'
 - 10 PFLOP/s LINPACK
 - Installed at RIKEN's Advanced Institute for Computational Science, Kobe

Fujitsu Begins Shipping Japan's Next-Generation Supercomputer

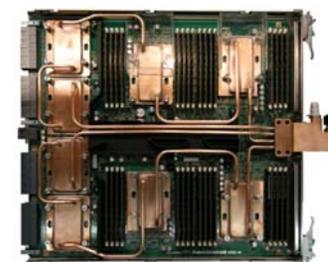
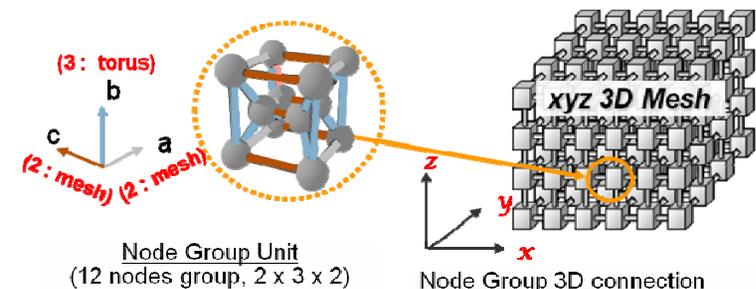
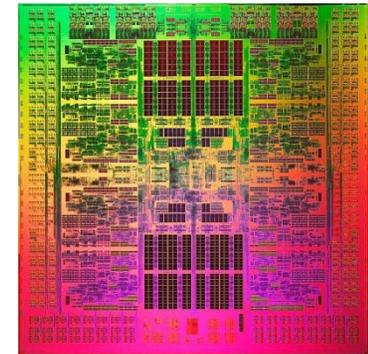
Tokyo, September 28, 2010 — Fujitsu announced that today it began shipping the computing units for Japan's Next-Generation Supercomputer, nicknamed the "K computer"



Hardware Technologies



- SPARC64™ VIIIfx: New HPC-enhanced CPU based on SPARCV9 architecture
 - 8 cores, 2 GHz, 2.2 GFLOP/s per watt
 - Extended HPC-ACE instruction set
 - Reciprocal, trig functions, max/min, SIMD, masked SIMD
 - Enlarged number of registers (Floating:256, Int:64)
 - User-controllable sector cache
- Tofu: 6D mesh/torus interconnect
 - Fast node to node communication, 5 GB/s (bi-directional)
 - Integrated MPI support for collective operations and global hardware barrier
 - Scalable to 100,000s of nodes
- Direct water cooling packaging

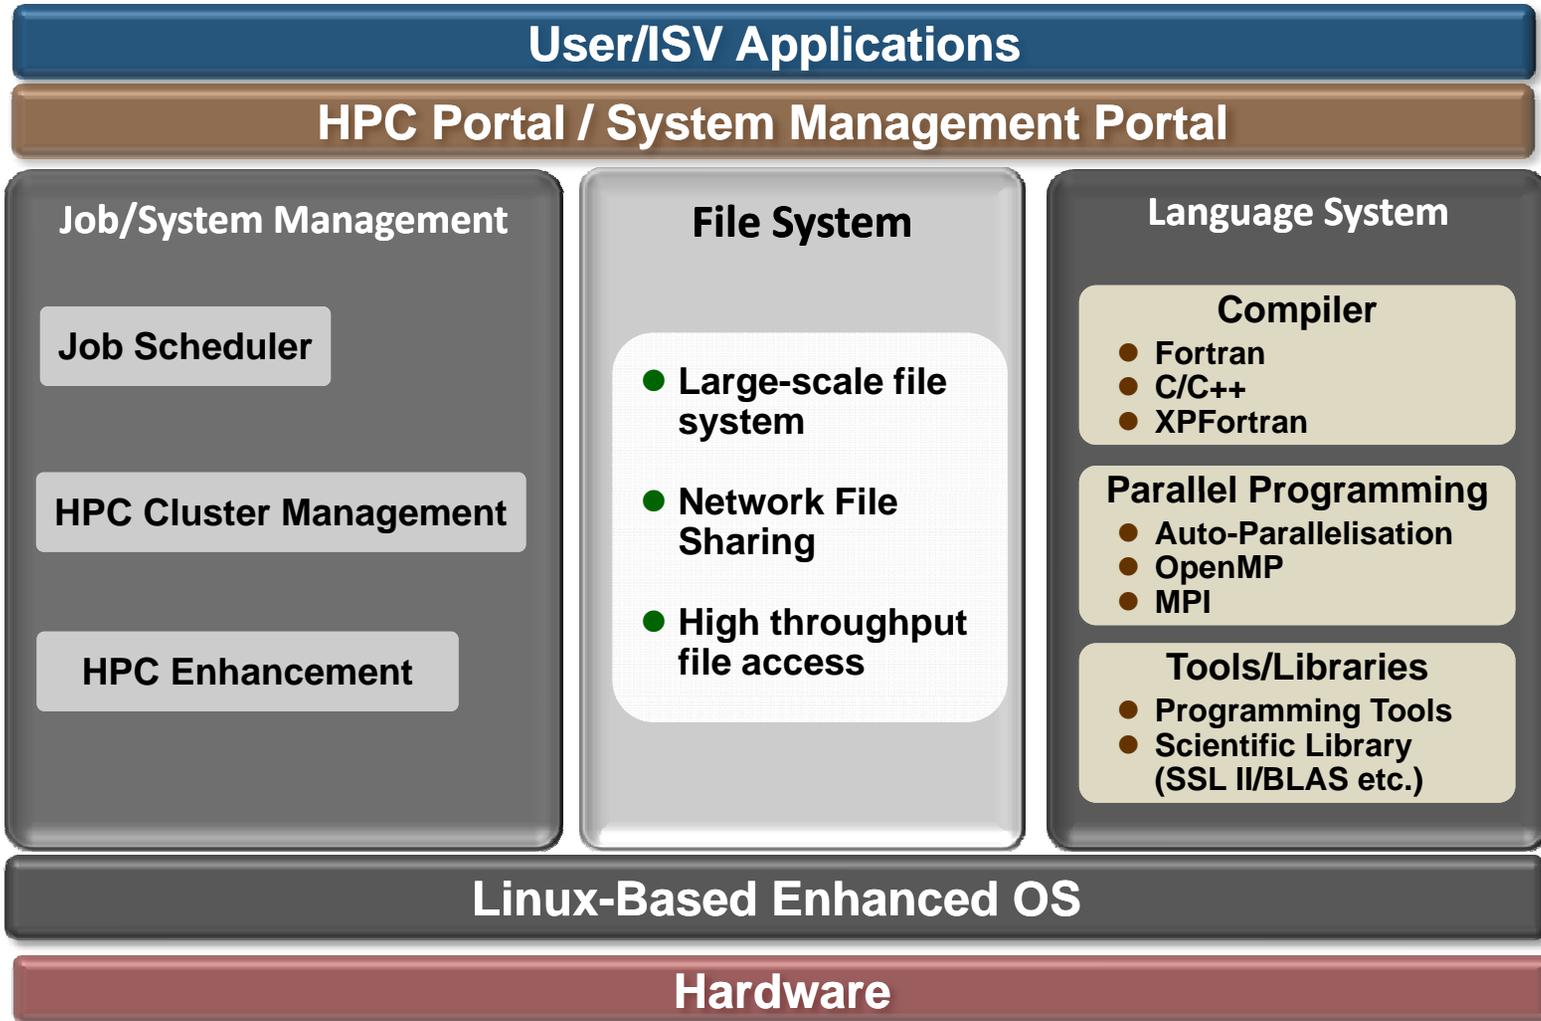


Direct water cooling System Board

System Software Stack



Based on open software, standards-conforming, and built for scalability

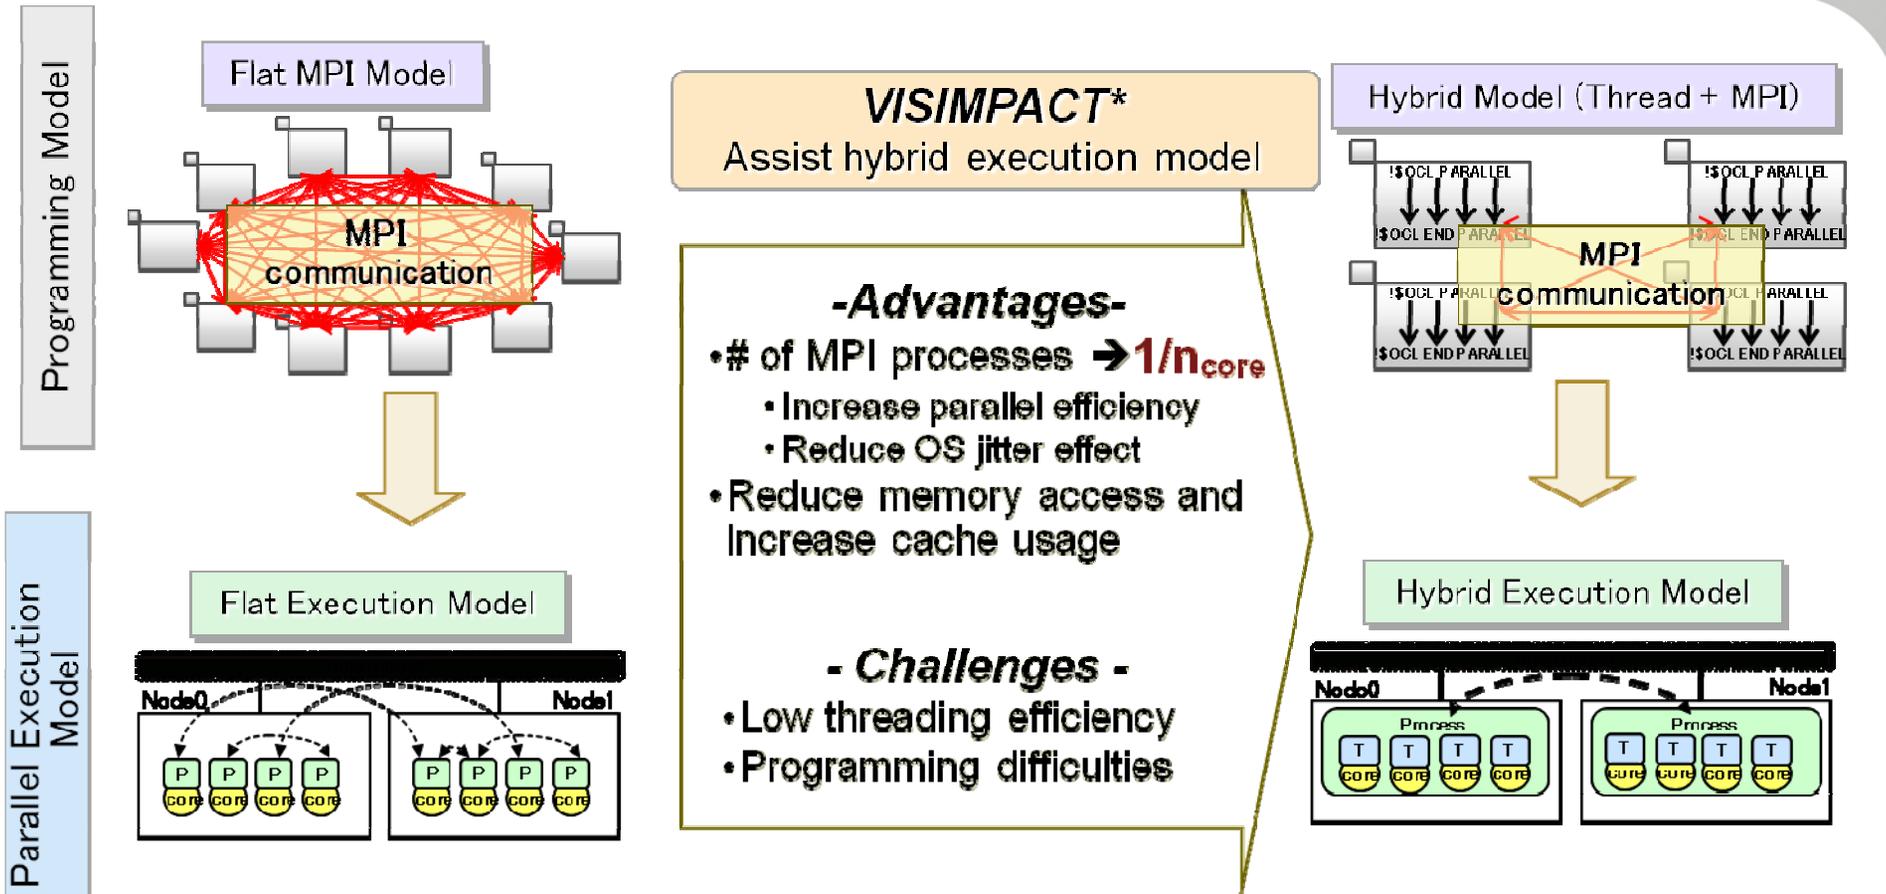


Programming Model



- Designed for a hybrid thread/task model
- Target fine-grained innermost loop parallelisation via multithreading
 - Makes use of Fujitsu's heritage in vectorisation technology
 - Automatic parallelisation or OpenMP
- Treats a multi-core chip as a single faster CPU with high performance
 - Programmers need not think of multiple cores
 - CPU architecture is designed for fine grain parallel execution

VISIMPACT



* : VISIMPACT realises highly efficient threading between cores

- Hardware barrier synchronisation between cores
- Shared L2 cache memory
- Highly efficient thread parallelisation (automatic parallelisation by vectorisation technology)

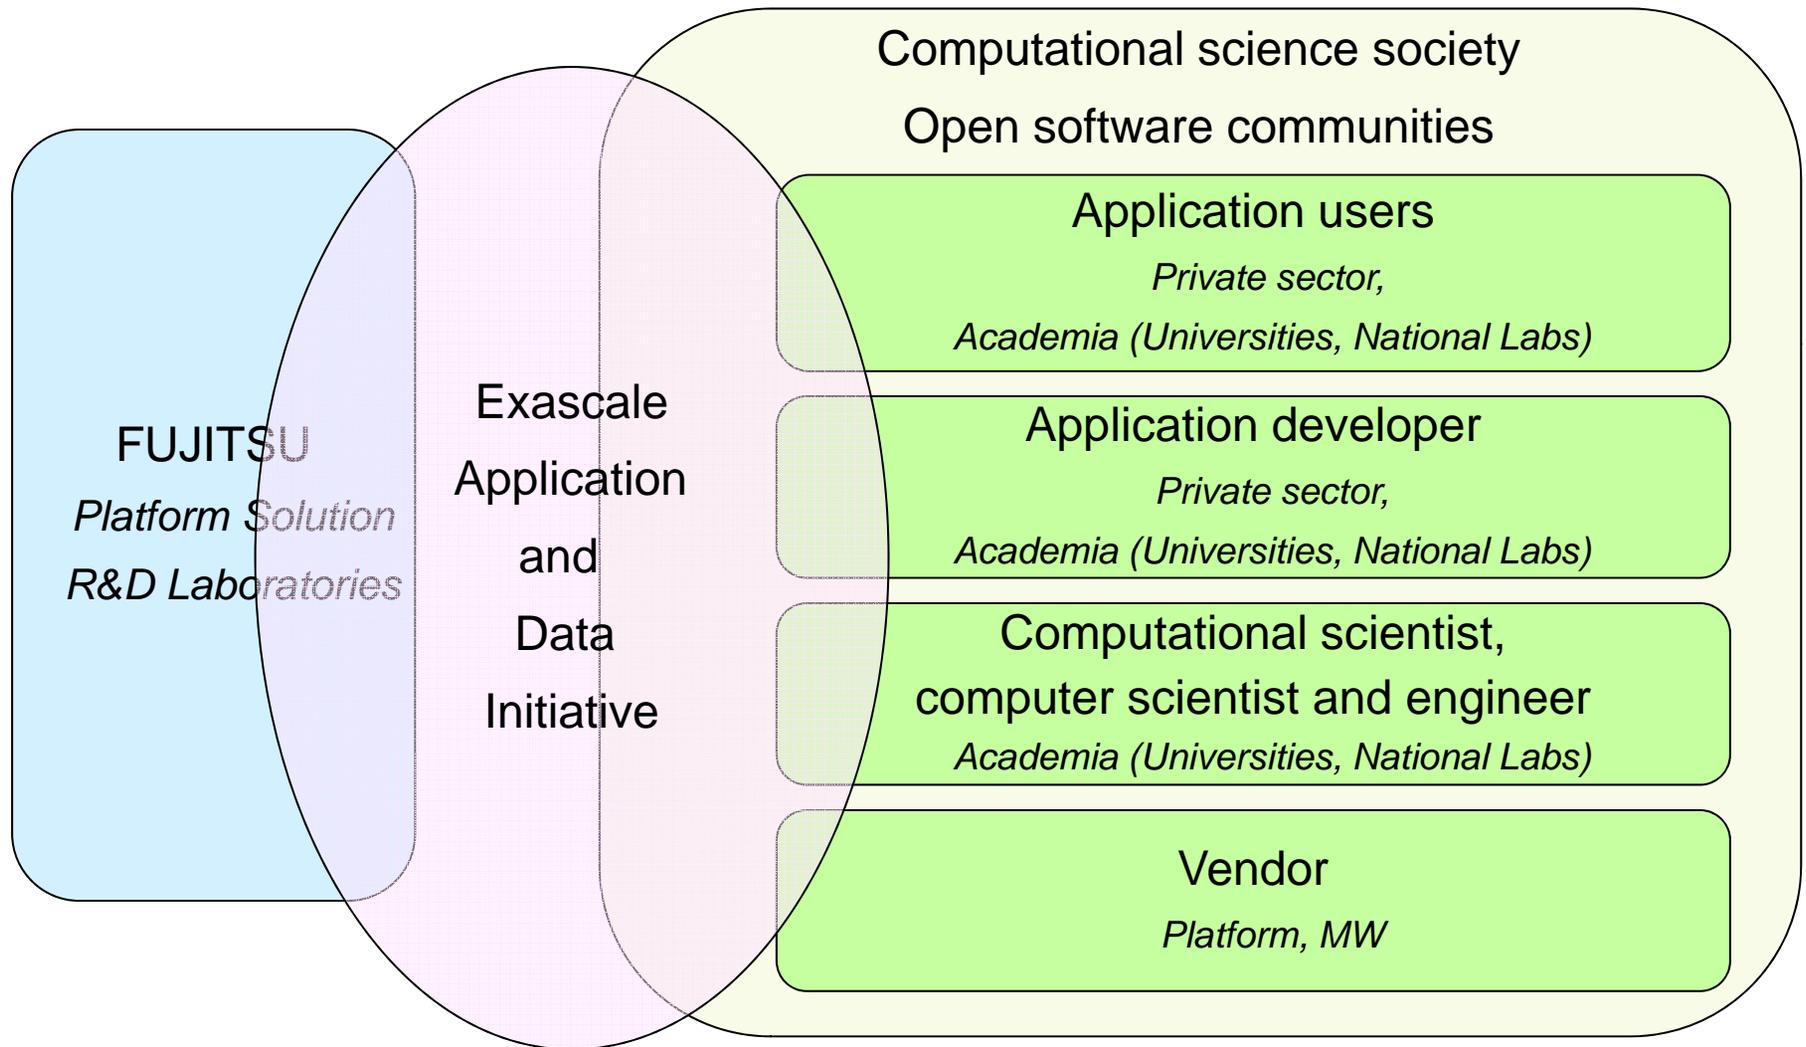
- Hardware and compiler technologies designed to enhance hybrid programming model
- Highly efficient threading
 - Hardware barrier synchronisation between cores
 - Shared L2 cache
 - Automatic parallelisation via vectorisation technology

Contents



- Fujitsu's Renewed Interest in Capability Computing
- Open Initiatives
- Current Co-Design Vehicles
- Towards Exascale
- Summary

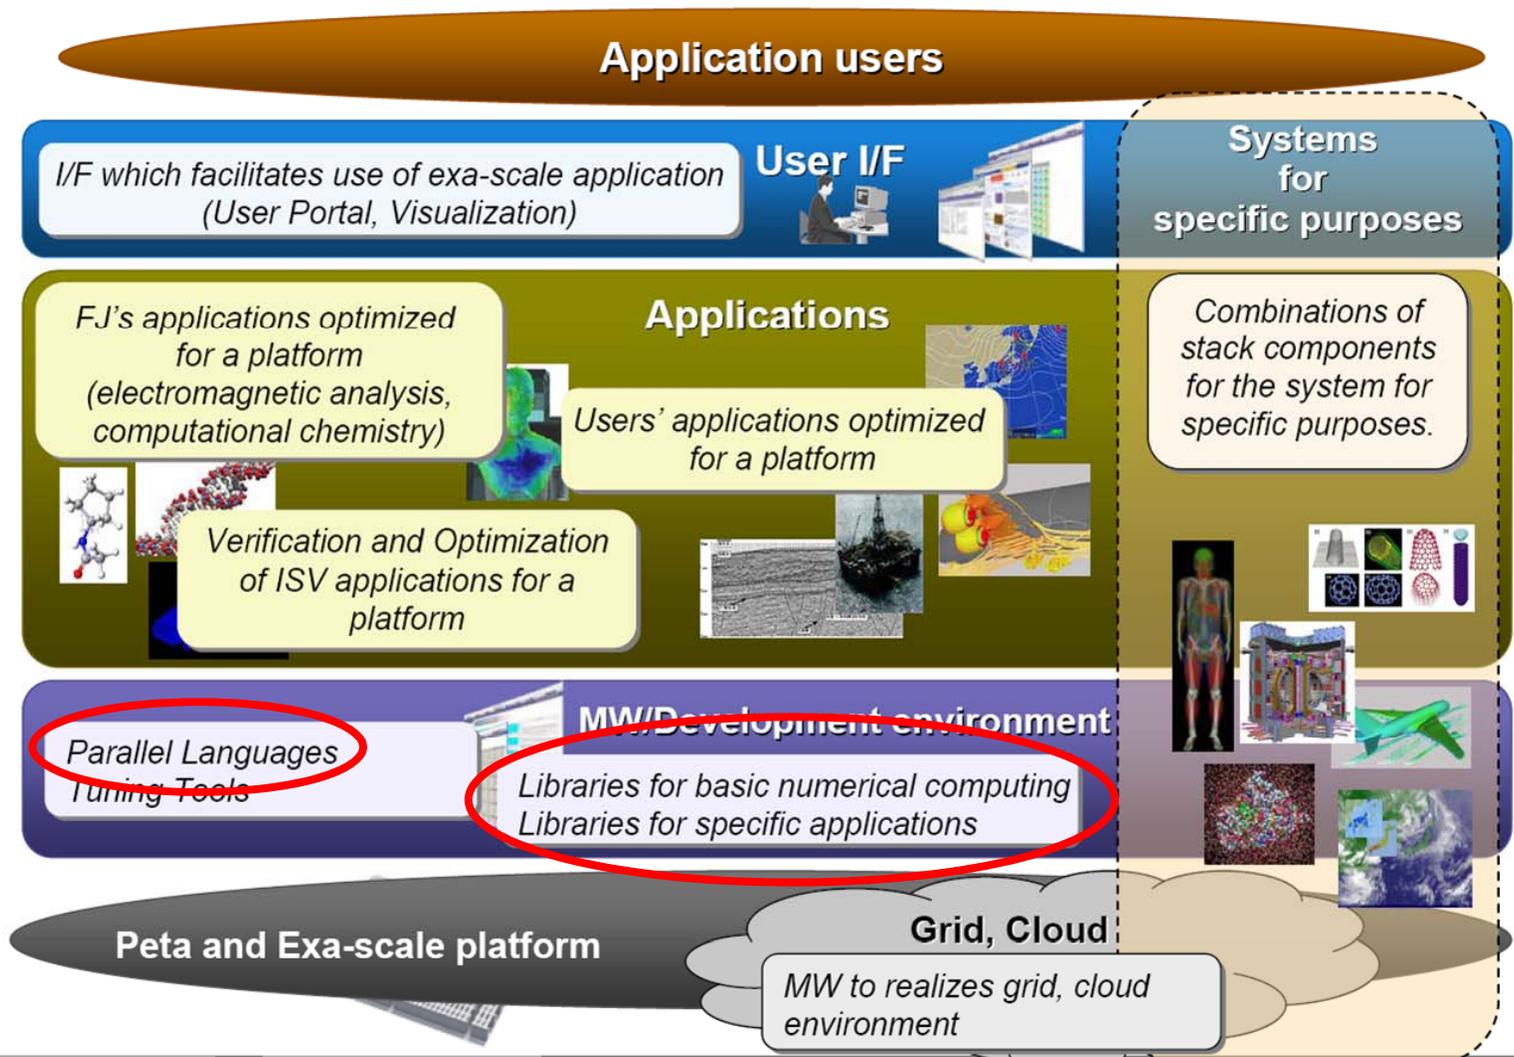
Exascale Application and Data Initiative



Exascale Application and Data Initiative



Solution Stack



XcalableMP Project



- Developed by XcalableMP Specification Working Group
- A draft “petascale” parallel language for “standard” parallel programming
 - To propose the draft to the world-wide community as a standard
- Academia: U. Tsukuba, U. Tokyo, Kyusyu U., Kyoto U.
- Research Lab.: RIKEN, NIFS, JAXA, JAMSTEC
- Industry: Fujitsu, NEC, Hitachi
- This research is carried out as a part of “Seamless and Highly-Productive Parallel Programming Environment for High-Performance computing” project funded by MEXT
- Center for Computational Sciences, University of Tsukuba, is organising the XcalableMP project

- A directive-based language extension that is designed to be simpler and less time-consuming than coding using MPI
- Data parallel “global view programming model” enables parallelising sequential code using simple directives like OpenMP
- Also includes a Co-Array Fortran like PGAS feature as “local view” programming
- APIs are defined for C and Fortran 95
- Execution model allows combining with explicit MPI coding for more complicated and tuned parallel codes and libraries
- OpenMP directives can be combined with XcalableMP for thread programming inside each node (under discussion)
- XcalableMP is being designed based on experiences of HPF, Fujitsu XPF (VPP Fortran) and OpenMPD

Open Petascale Libraries Project

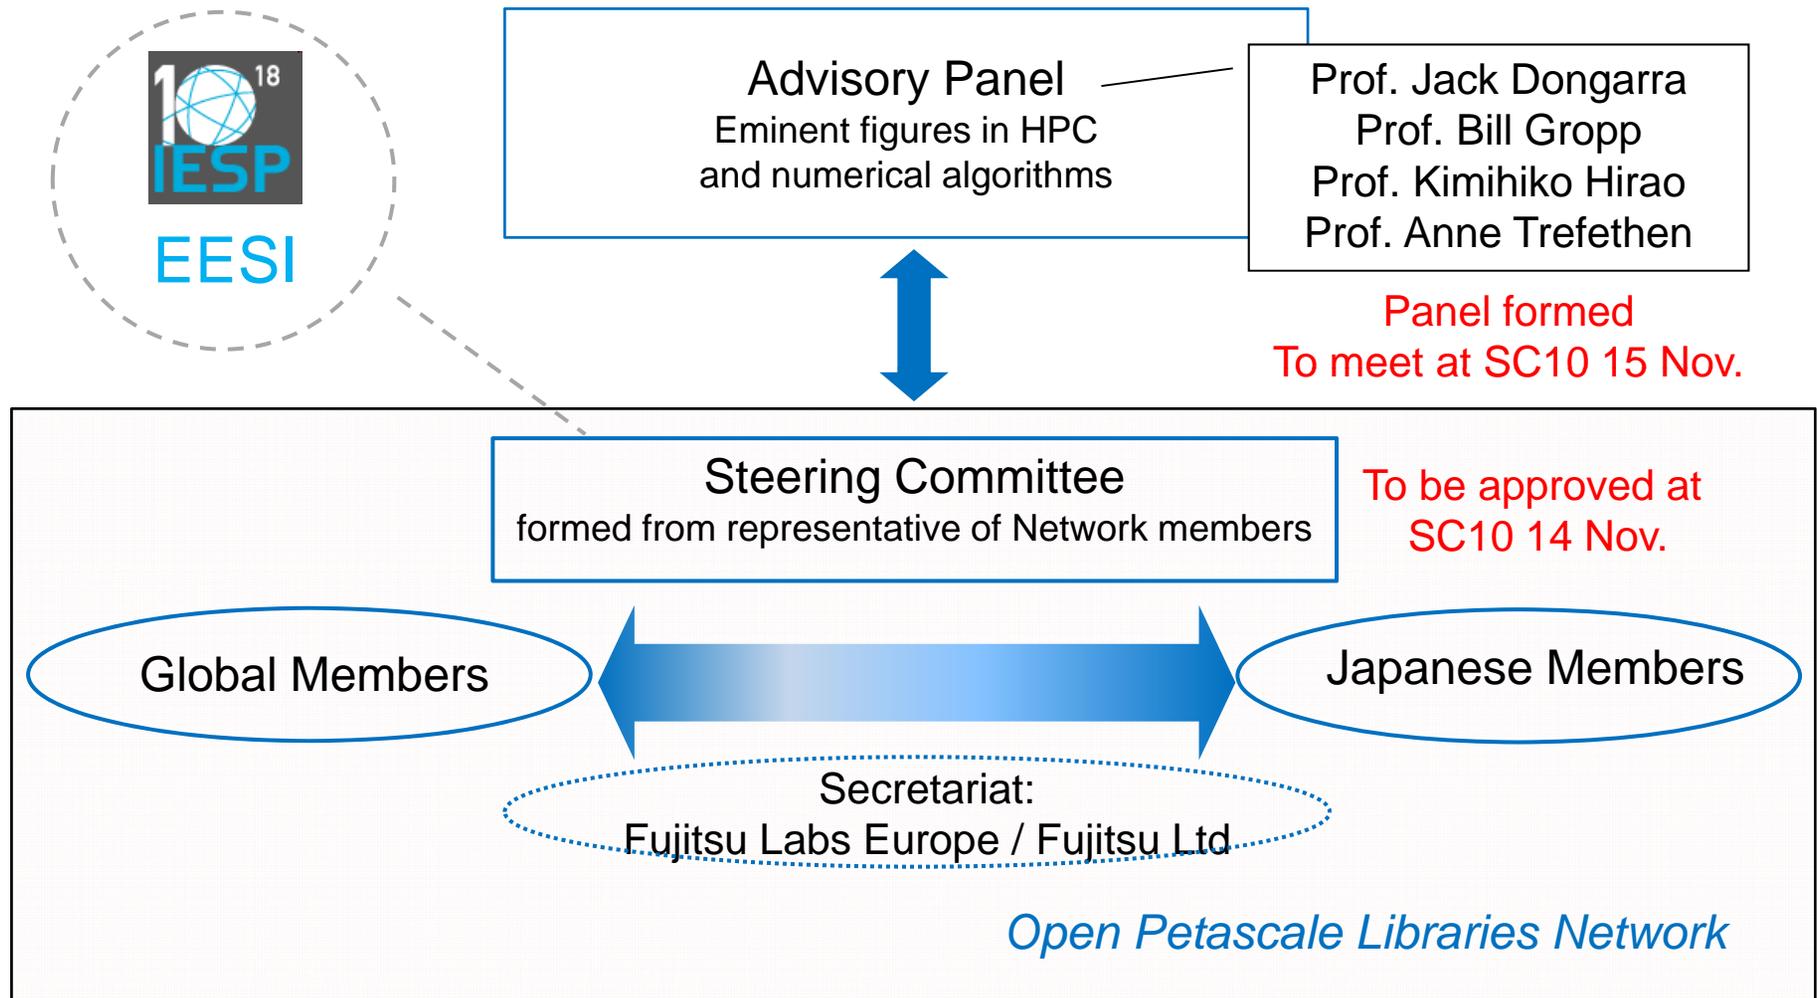


Global collaboration to develop advanced numerical software for supercomputing

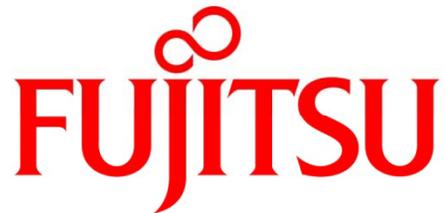
- Dedicated forum to promote the open exchange of ideas and the collaborative development of general-purpose and application-specific numerical libraries
- Targeted initially at parallel computers built from multicore processors
- All output available as open-source software
- Official launch of OPL in November



OPL: Organisation



OPL: Current Global Members



Imperial College
London



nag[®]



OPL: Proposed Work



- Survey of libraries and current research
- Application requirements
- Emerging libraries for dense linear algebra
- Enabling technologies for large-scale PDE-based simulations
- Scalable Fast Fourier Transforms
- Scalable random number generators
- Coding and documentation standards, testing and distribution of libraries

OPL: Fujitsu's Role



- Set up the required administrative infrastructure
- Contribute to the project through research in its global network of R&D laboratories
- Provide an environment to support the sharing of knowledge and software
 - Organise meetings, workshops, etc.
 - Establish a dedicated project website and code repository (www.openpetascale.org)
 - Provide an open development platform
 - Housed at STFC's Daresbury Laboratory

Contents



- Fujitsu's Renewed Interest in Capability Computing
- Open Initiatives
- **Current Co-Design Vehicles**
- Towards Exascale
- Summary

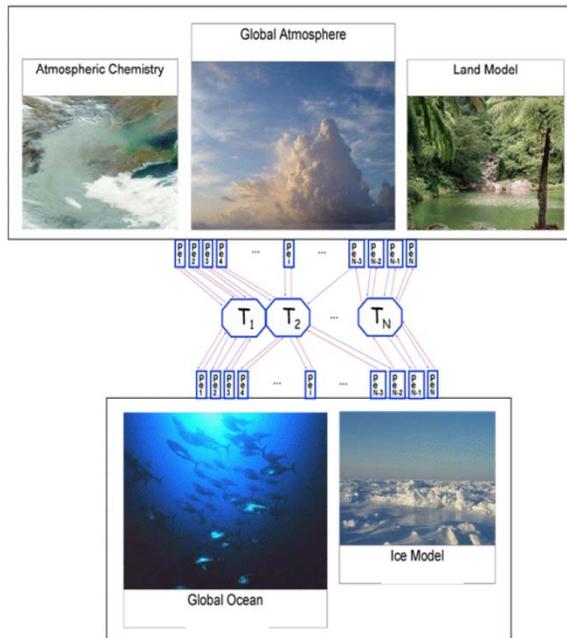
Co-Design Vehicles



- Petascale software development is using CDVs
 - Focused on reusable thread-parallel and hybrid kernels
 - Based on open-source software development
 - Environment and energy
 - Healthcare (computational physiology)
 - Computational chemistry
 - Fujitsu's in-house engineering requirements
- More extensive use of CDVs for hardware and system software development as we move towards exascale

Environmental Modelling

- Collaboration with Imperial College London and others
- Possible involvement in two G8 projects (Arctic Ocean modelling, exascale climate simulations)



Multi-Physics
Earth-system modelling

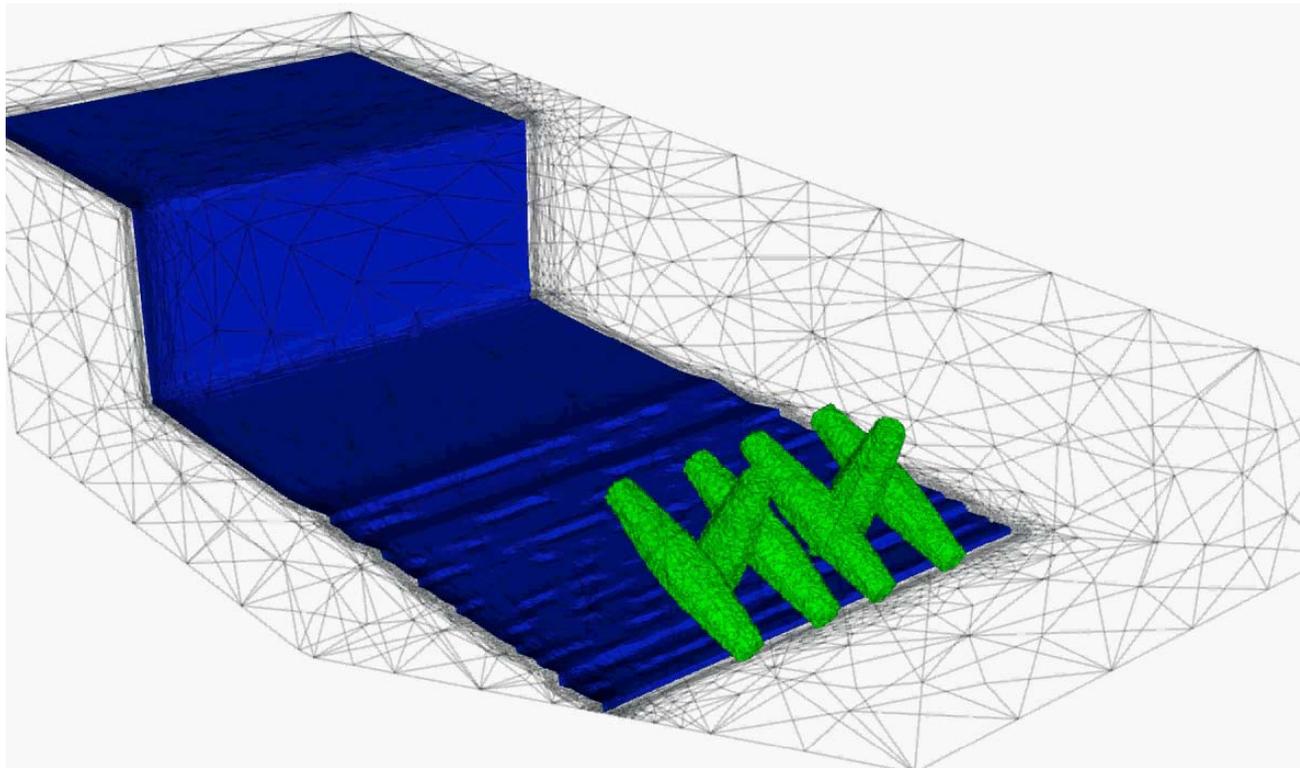


Multi-Application

From simulation of coastal flooding and tidal power generation to landslide-generated tsunamis and engineering applications

Technology

- Highly parallel load-balanced FEM using anisotropic adaptive meshing and combining OpenMP with MPI



Computational Physiology



Computational Prediction of Drug Cardiac Toxicity

High-performance finite-element solver for the electrophysiology of the heart

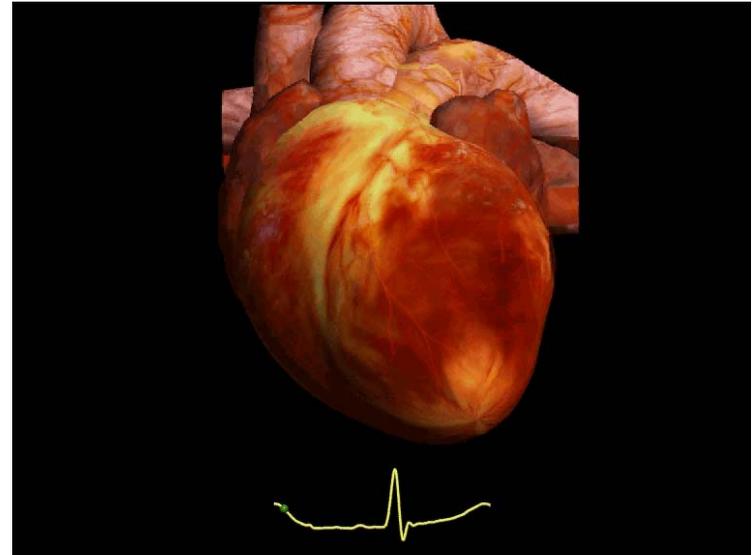


UNIVERSIDAD POLITÉCNICA DE VALENCIA



- Three-year €4 million European project
- Targeted at the pharmaceutical industry
- Part of the Virtual Physiological Human Initiative

- Improved early-stage safety testing for new drug candidates
- Need to streamline drug discovery process
 - Each new drug takes 13 years and costs on average € 1 billion to develop
- To be practical, high efficiency on petascale-class computers is essential



Computational Chemistry



- Lattice Boltzmann
- Molecular dynamics (GROMACS)
- Quantum chemistry (NWChem)

Contents



- Fujitsu's Renewed Interest in Capability Computing
- Open Initiatives
- Current Co-Design Vehicles
- Towards Exascale
- Summary

Towards Exascale



- Fujitsu is undertaking future technologies research and evaluation for exascale computing
 - Semiconductor technologies, future CPU architecture, packaging technologies, software technologies, increased memory bandwidths, high-performance and scalable interconnects
- Which applications need exascale performance?
 - What is the communication pattern, data size, etc.?
- The design of the system will be deeply affected by these characteristics and needs of applications
- What will be the role of disruptive technologies?

Summary



- Fujitsu has returned to capability computing with renewed vigour
- Fujitsu is committed to standards, open collaborations and open-source software
- The CDV mechanism is already in use for petascale software development and will have an increased importance for the full range of exascale system development



FUJITSU

shaping tomorrow with you