FUJITSU

shaping tomorrow with you

Key Technologies for 100 PFLOPS

How to keep the HPC-tree growing

FUJITSU

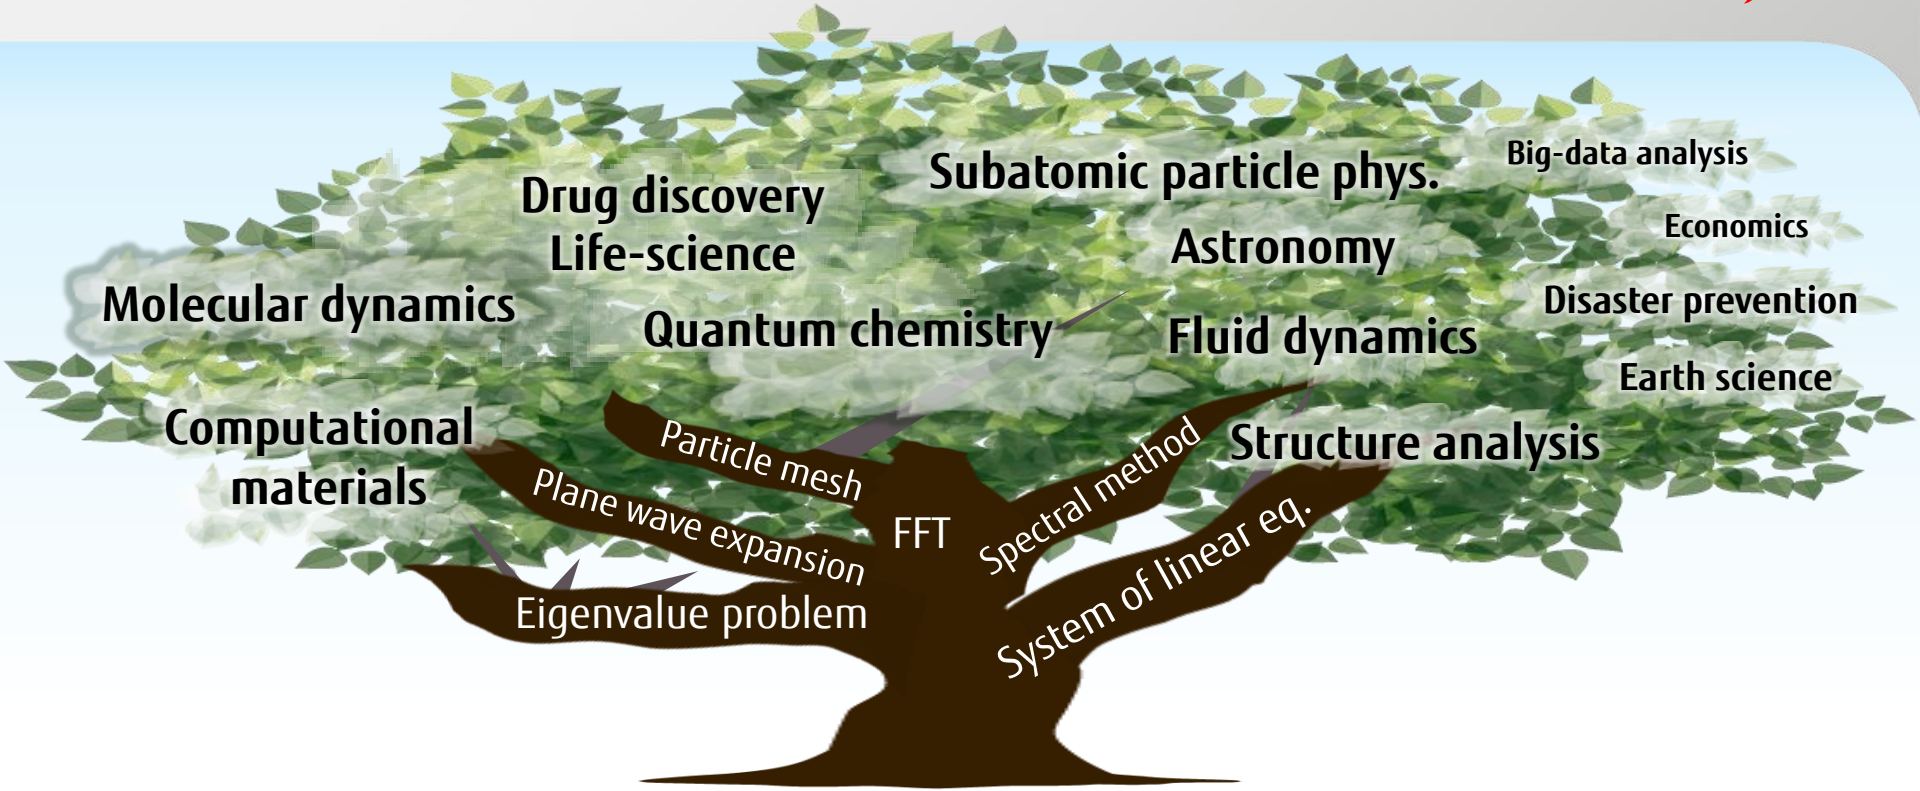

- A stable programming model is needed to secure the continuous software development.

Lasting programming model

FUJITSU

The MPI + (OpenMP or automatic-parallelization) programming model will stay important.

- VISIMPACT (Virtual Single Processor by Integrated Multi-core Parallel Architecture)
  - Automated multi-thread parallelization
    - Inherits vectorization technology
  - High performance hardware thread barrier
  - Easy hybrid parallelization reduces number of processes
    - Reducing communication overhead
    - Saving memory usage

POST-FX10

- 2011 **FUJITSU Supercomputer PRIMEHPC FX10**

- 2011 #1 on Top500 **K computer (*)**

(*) The K computer is a supercomputer jointly developed by RIKEN and Fujitsu.

- 2009 **FX1**
- 2000 **PRIMEPOWER**

- 1999 **VPP5000**
- 1993 #1 on Top500 **NUMERICAL WIND TUNNEL**

- 1974 **FACOM-M190**

Proven performance in various application fields

FUJITSU

Performance data on K computer are provided by Mr. Minami of RIKEN

Application	Research Field	# of nodes	Efficiency	
NICAM	Earth science	81,920	8%	
Seism3D	Disaster prevention	82,944	18%	
PHASE	Computational materials, Quantum chemistry	82,944	20%	
RSDFT	Computational materials	82,944	52%	
FrontFlow/blue	Fluid dynamics	80,000	3%	
Lattice QCD	Subatomic particle phys.	82,944	16%	
ZZ-EFSI	Life science, Fluid dynamics, Structure analysis	82,944	46%	*1
GreeM	Astronomy	82,944	42%	*2

*1: http://www.scls.riken.jp/eng/newsletter/Vol.7/report02.html
*2: T.Ishiyama, et.al., 4.45 Pflops astrophysical N-body simulation on K computer, SC'12

- **Post-FX10 maintains excellent performance efficiency with desirable technologies as well as a well-balanced system design.**

Data intensive applications

FUJITSU

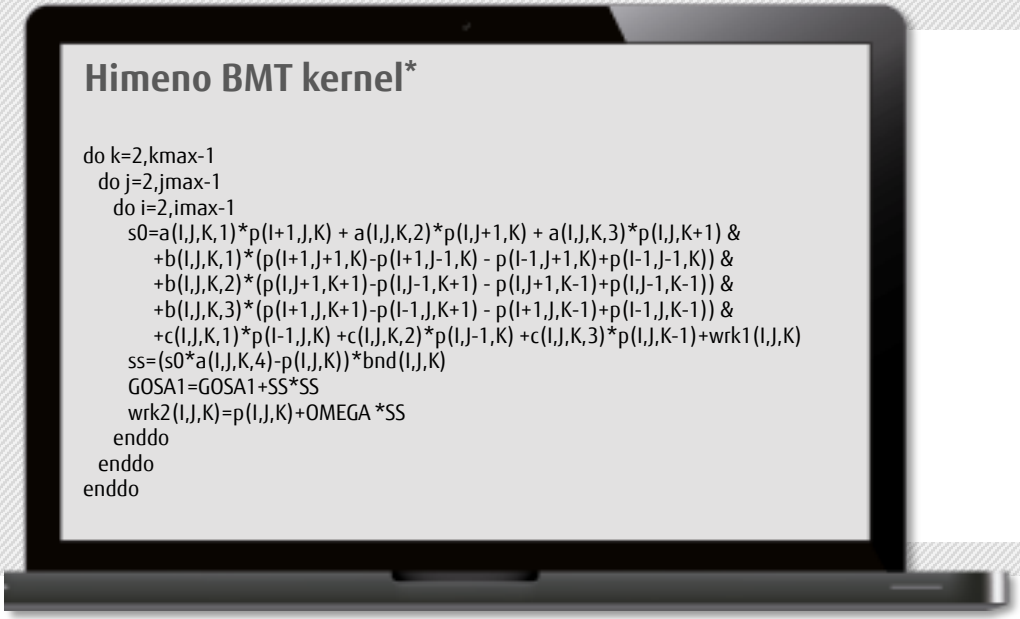

* Tentative result

- **HMC is effective for data intensive applications.**

* Incompressible fluid analysis code solving the Poisson's equation solution using the Jacobi iteration method. Many loads require high memory throughput.

Dense matrix problems

FUJITSU

- DGEMM performance of BLAS:

	Performance
FX10	95.3% (16core)
Post-FX10	97.9% (16core)

* Tentative result

- Enhanced sector-cache mechanism improve efficiency.
- DGEMM is the fundamental function for many applications.

Sparse matrix problems

FUJITSU

- NAS parallel benchmarks CG kernel

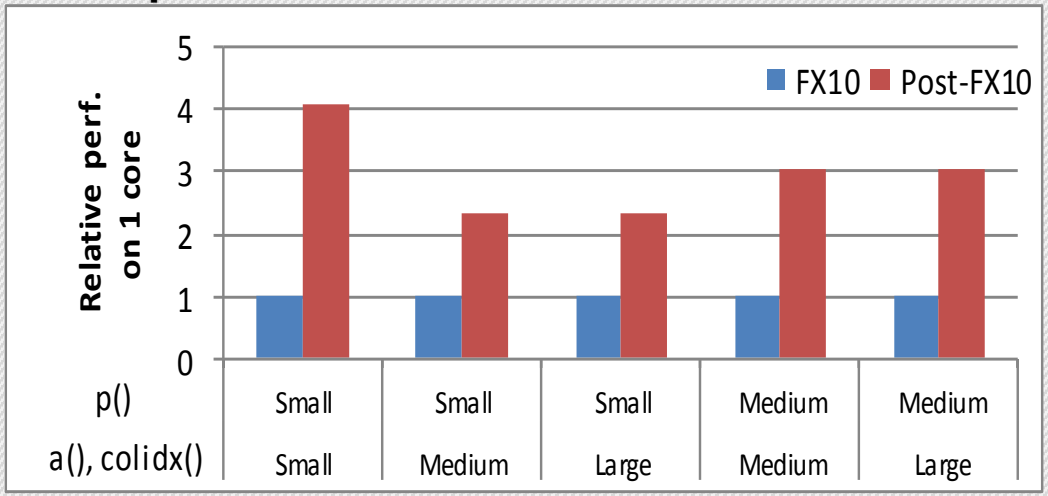

* Tentative result

```
do j=1,lastrow-firstrow+1
  sum = 0.d0
  do k=rowstr(j),rowstr(j+1)-1
    sum = sum + a(k) * p(colidx(k))
  enddo
  w(j) = sum
enddo
```

Small	Array size < 32KB
Medium	Array size < 12MB
Large	Array size > 20MB

- 2 to 4 times faster than FX10 (on 1 core)
- The new indirect-load of wider-SIMD operation contributes to performance improvement

Communication intensive applications

FUJITSU

- 3-dimensional FFT using volumetric decomposition

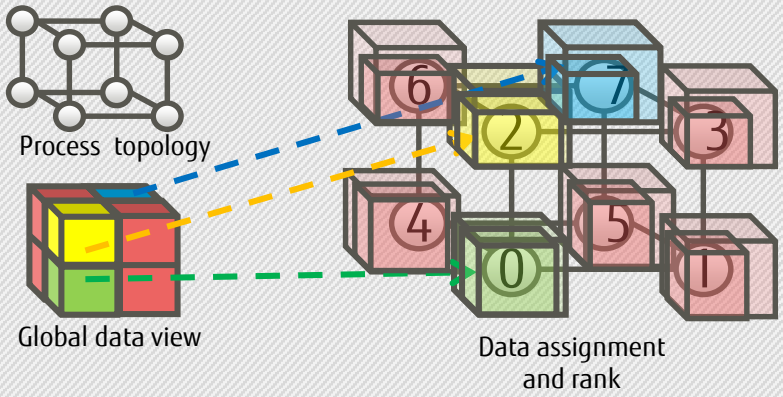

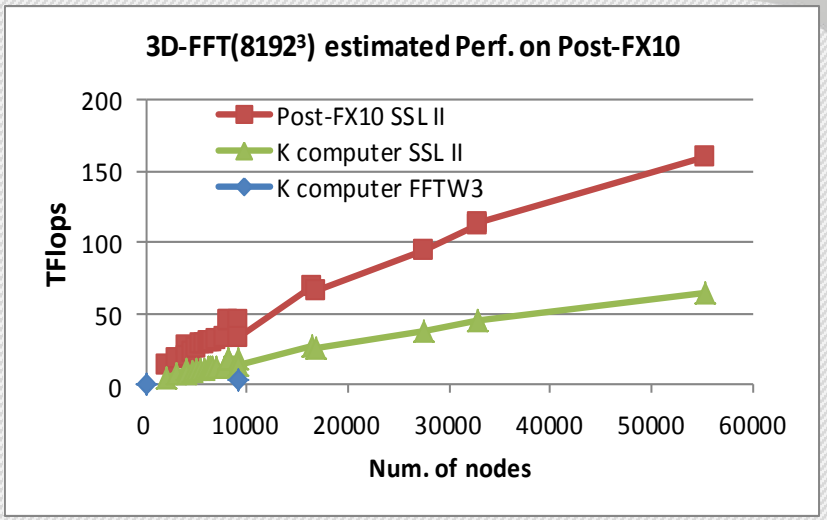

- **Fujitsu's math library provides scalable 3D-FFT functionality for massively parallel execution.**
  - Interconnect bandwidth improves significantly.
  - Tofu Interconnect and job manager enable various processing configuration.
  - Optimized MPI library provides high performance collective communications for Tofu Interconnect.

Conclusion

FUJITSU

- Fujitsu is developing the Post-FX10 system to be capable of 100 PFLOPS:
  - Enhanced SPARC64 CPU and Tofu Interconnect were designed exclusively for HPC.
  - Fujitsu's HPC software stack exploits maximum hardware performance.
- Post-FX10 employs best technologies for actual application performance:
  - 256 bit-wide SIMD processing and enhanced sector cache mechanism improve flops
  - New indirect-load/store SIMD operations are effective for complex data access
  - Micron's Hybrid Memory Cube fulfills high B/F requirements
  - Enhanced interconnect and assistant cores enable fast and scalable communication
- Post-FX10 strongly drives scientific progress forward.

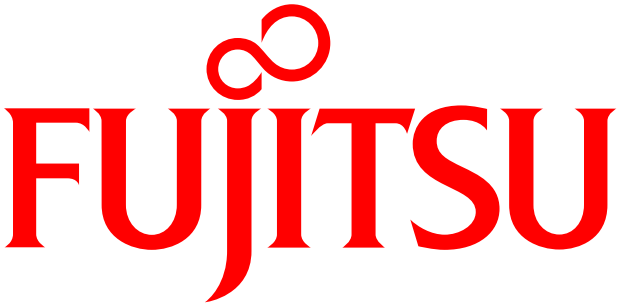

shaping tomorrow with you