

Next-Generation PRIMEHPC

The K computer and the evolution of PRIMEHPC



	K computer	PRIMEHPC FX10	Post-FX10
CPU	SPARC64 VIIIfx	SPARC64 IXfx	SPARC64 XIfx
Peak perf.	128 GFLOPS	236.5 GFLOPS	1TFLOPS ~
# of cores	8	16	32 + 2
Memory	DDR3 SDRAM	←	HMC
Interconnect	Tofu Interconnect	←	Tofu Interconnect 2
System size	11PFLOPS	Max. 23PFLOPS	Max. 100PFLOPS
Link BW	5GB/s x bidirectional	←	12.5GB/s x bidirectional



Smaller, faster, more efficient



- Highly integrated components with high-density packaging.
- Performance of 1-chassis corresponds to approx. 1-cabinet of K computer.

➤ Efficient in space, time, and power

1 Fujitsu designed SPARC64™ XIfx




- 32 + 2 core CPU
- Tofu2 integrated

2 CPU Memory Board



- Three CPUs
- 3 x 8 HMCs (Hybrid Memory Cube)
- 8 optical modules (25Gbpsx12ch)

3 Chassis (12 CPUs)




- 1 CPU/1 node
- 12 nodes/2U chassis
- Water cooled

4 Cabinet



- Over 200 nodes/cabinet

5 Tofu 2



Architecture continuity for compatibility

■ Upper compatible CPU:

- Binary-compatible with the K computer & PRIMEHPC FX10
- Good byte/flop balance

■ New features:

- New instructions (stride load/store, indirect load/store, permutation, concatenation)
- Improved micro architecture (out-of-order, branch-prediction, etc.)

■ For distributed parallel executions:

- Compatible interconnect architecture
- Improved interconnect bandwidth

K computer



FX10



Post-FX10



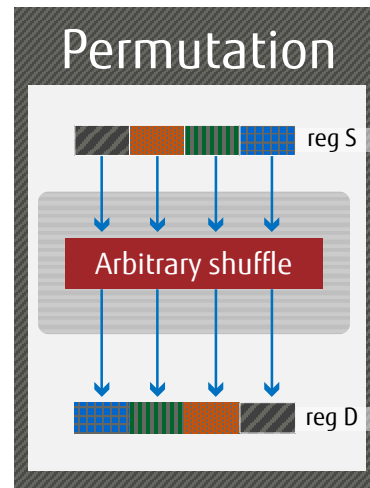


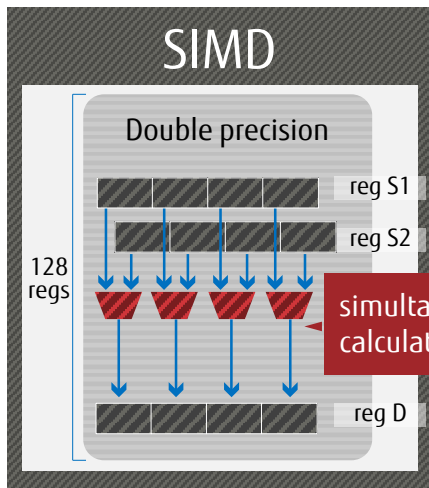
32 + 2 core SPARC64 XIfx

- Rich micro architecture improves single thread performance.
- HMC fulfills required bandwidth for multi-core high performance CPU.
- 2 additional, Assistant-cores for avoiding OS jitter and non-blocking MPI functions.

		K	FX10	Post-FX10	Note
Peak FP performance		128 GF	236.5 GF	1-TF class	Maintains similar architecture for compatibility with applications
Core config.	Execution unit	FMA × 2	FMA × 2	FMA × 2	
	SIMD	128 bit	128 bit	256 bit wide	Wider SIMD for better performance with small-sized additional hardware
	Dual SP mode	NA	NA	Double of DP	Accelerates SP-rich apps
	Integer SIMD	NA	NA	Support	Accelerates INT rich apps Assists SIMDization with list vector
	Single thread performance enhancement	-	-	Increase OOO resources, better branch prediction, larger cache	Application performance often limited by single thread performance and no FP calc.

Flexible SIMD operations

- New 256bit wide SIMD functions enable versatile operations
 - Four double-precision calculations
 - Stride load/store, Indirect (list) load/store, Permutation, Concatenation



Tofu Interconnect 2

■ Successor to Tofu Interconnect

- Highly scalable, 6-dimensional mesh/torus topology
- Increased link bandwidth by 2.5 times to 12.5 GB/s

■ Interconnect integrated into CPU

- System-on-chip (SoC) removes off-chip I/O
- Improved packaging density and energy efficiency


■ Optical cable connection between chassis



Flexible interconnect topology

- Tofu: Six-dimensional mesh/torus direct network
- Logical 3D, 2D or 1D torus network from the user's point of view

Scalable three-dimensional torus



Three-dimensional torus unit
 $2 \times 3 \times 2$

Well-balanced shape available

Entire software stack is enhanced for Post-FX10



Applications

HPC Portal / System Management Portal

Technical Computing Suite

System Management

- System management
- System control
- System monitoring
- System operation support

Job Management

- Job manager
- Job scheduler
- Resource management
- Parallel

High Performance File System *FEFS*

- Lustre based high performance distributed file system
- High scalability, high reliability and availability

Automatic parallelization compiler

- Fortran
- C
- C++

Tools and math libraries

- Programming support tools
- Mathematical libraries

Parallel languages and libraries

- OpenMP
- MPI
- XPFortran

Linux based OS (enhanced for FX series)

PRIMEHPC FX series

Summary

■ Successor to the PRIMEHPC FX10




- Fujitsu-developed, compatible SPARC64 CPU and Tofu Interconnect
- High-density packaging

■ SPARC64 XIfx


- Rich micro architecture (32 computing cores + 2 assistant cores)
- Richer SIMD operations supported
- High memory bandwidth (with HMC)

■ Interconnect

- Tofu Interconnect 2 is integrated into the CPU
- Optical connections between chassis



➤ 100 petaflops-capable system



shaping tomorrow with you