SPARC64™ X:
Fujitsu's New Generation 16 core Processor for UNIX Server

19th April 2013

Toshio Yoshida

Processor Development Division
Enterprise Server Business Unit
Fujitsu Limited

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™

- SPARC64™ is Fujitsu's SPARC Processor series for UNIX servers and Supercomputers.
- SPARC64™ VIIIfx is Running on K computer Now!

- ✓ SPARC64™ X is the New Generation Processor for Fujitsu's UNIX Server

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Fujitsu's Processor Development

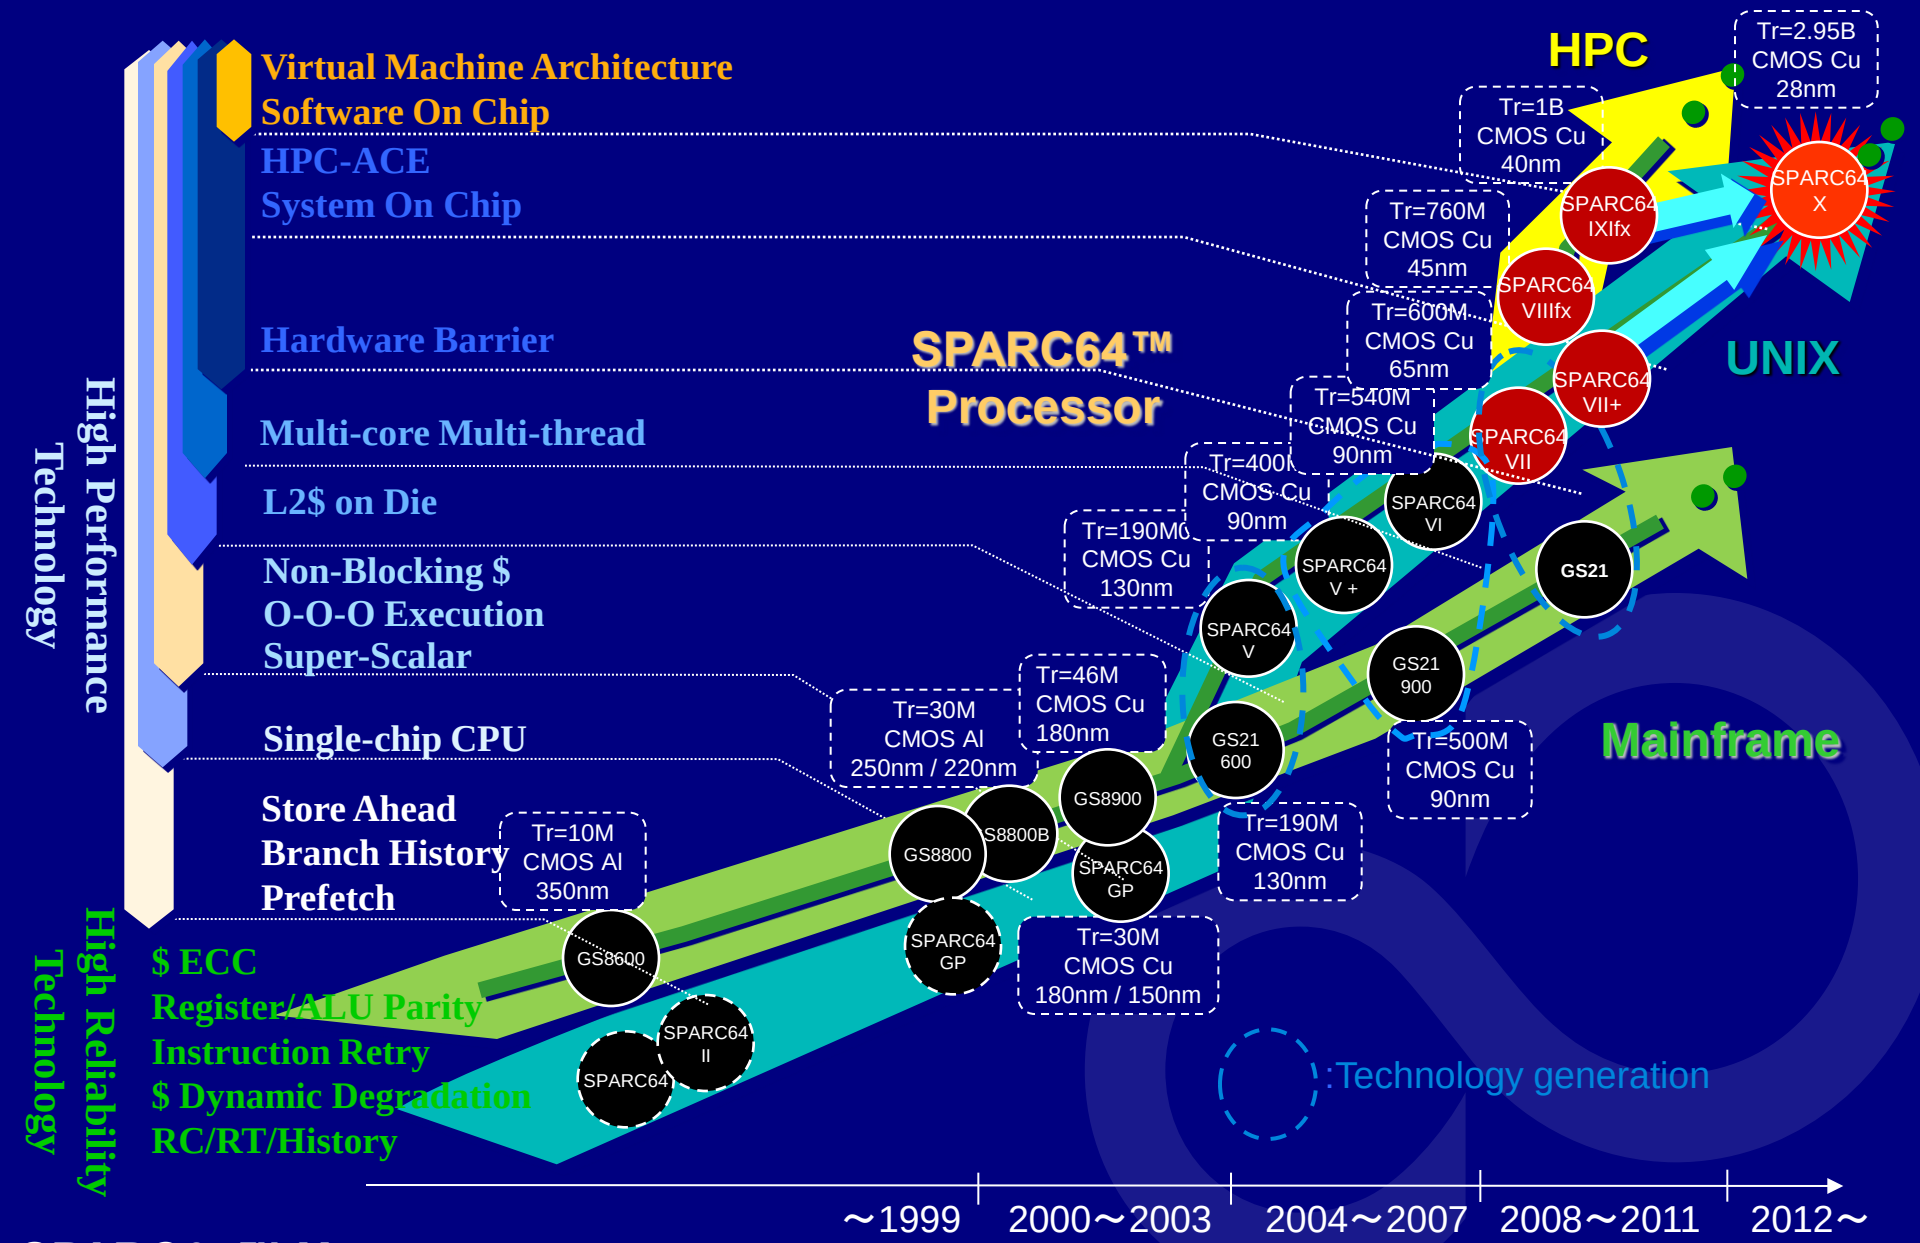

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Design Concept of SPARC64™ X

- Combine UNIX and HPC FJ Processor Features to Realize an Extremely High Throughput UNIX Processor
  - SPARC64 VII/VII+ (UNIX Processor) Feature
    - High clock frequency (up-to 3GHz)
    - Multicore/Multithread
    - Scalability : up-to 64 sockets
  - SPARC64 VIIIfx (HPC Processor) Feature
    - HPC-ACE: ISA extensions to SPARC-V9
      e.g. SIMD, Register # enhancement
    - High Memory B/W: Peak 64GB/s, Embedded memory controller

- Add New Features to UNIX Servers
  - Virtual Machine Architecture
  - Embedded IOC (PCI-GEN3 Controller)
  - Direct CPU-CPU Interconnect
  - Software On Chip

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ VII/VII+ Pipeline (2008)

- ✓SMT
- ✓4 Cores / Shared Level 2 Cache / Hardware Barrier (VISIMPACT)
- ✓64 Sockets per Node

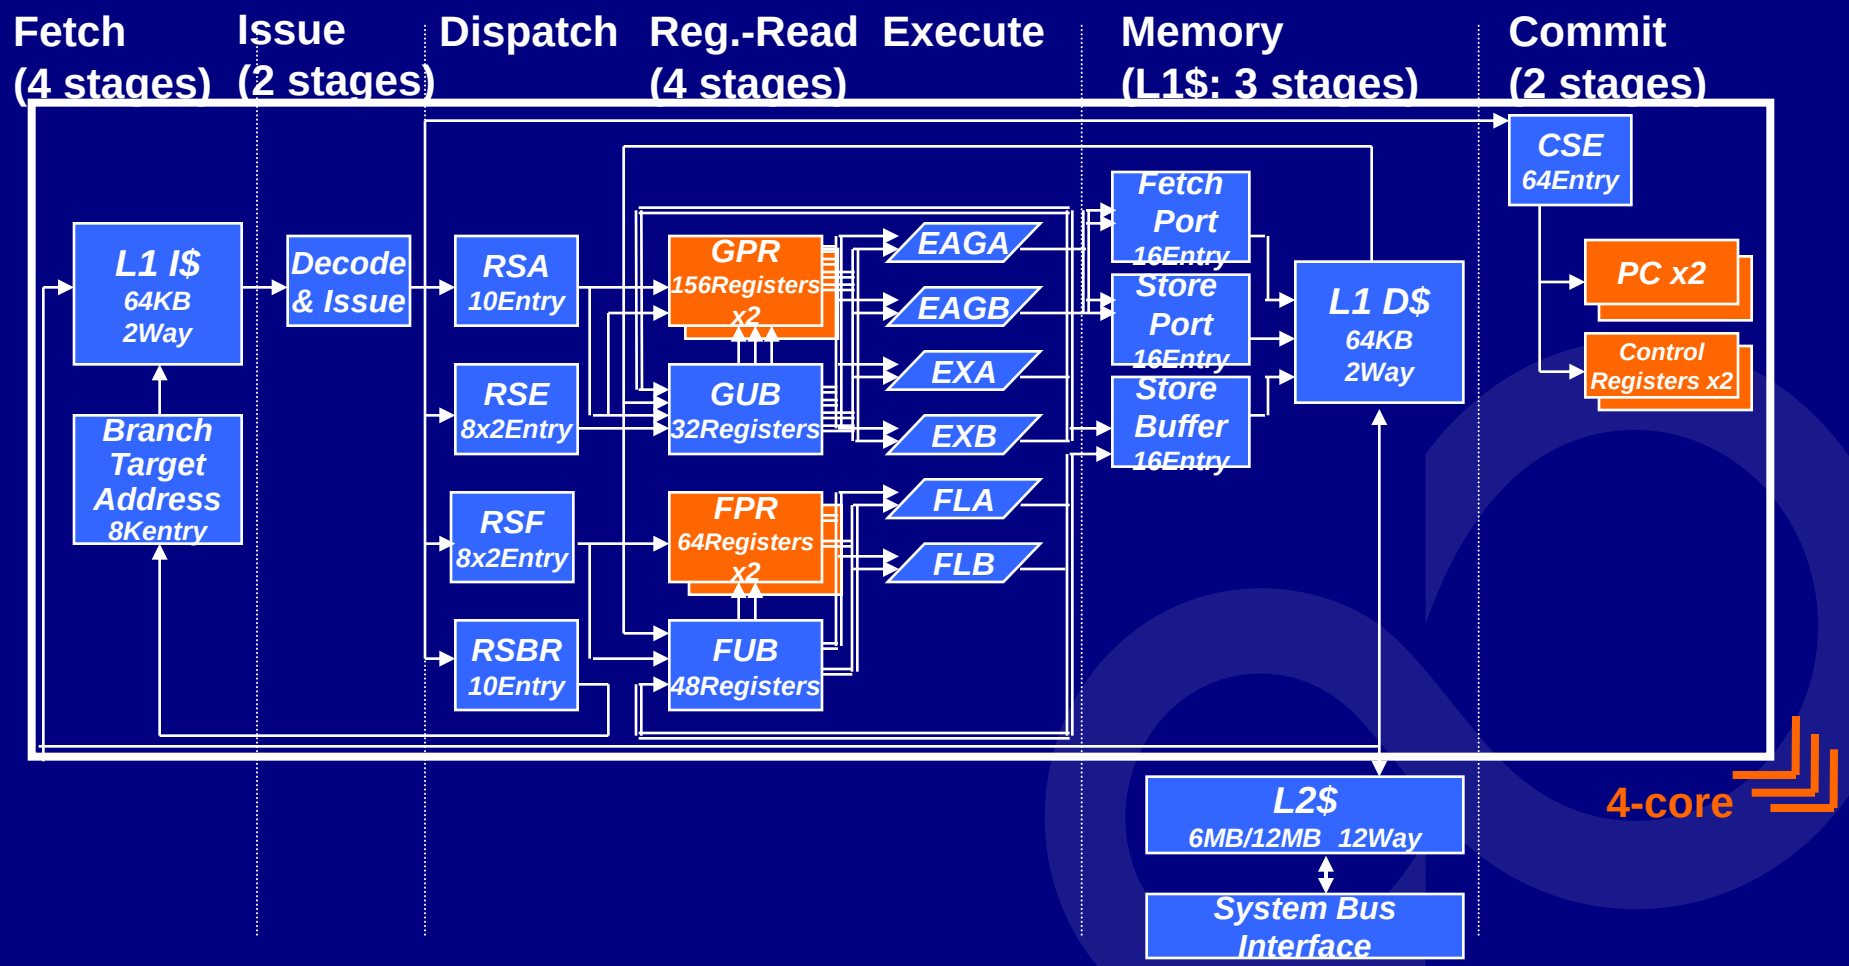

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ VIIIfx/IXfx Pipeline(2010)

- ✓8-16 Cores without SMT / HPC-ACE (SIMD, Register # Enhancement etc.)
- ✓High Memory Throughput
- ✓Single Socket per Node / Tofu Interconnect

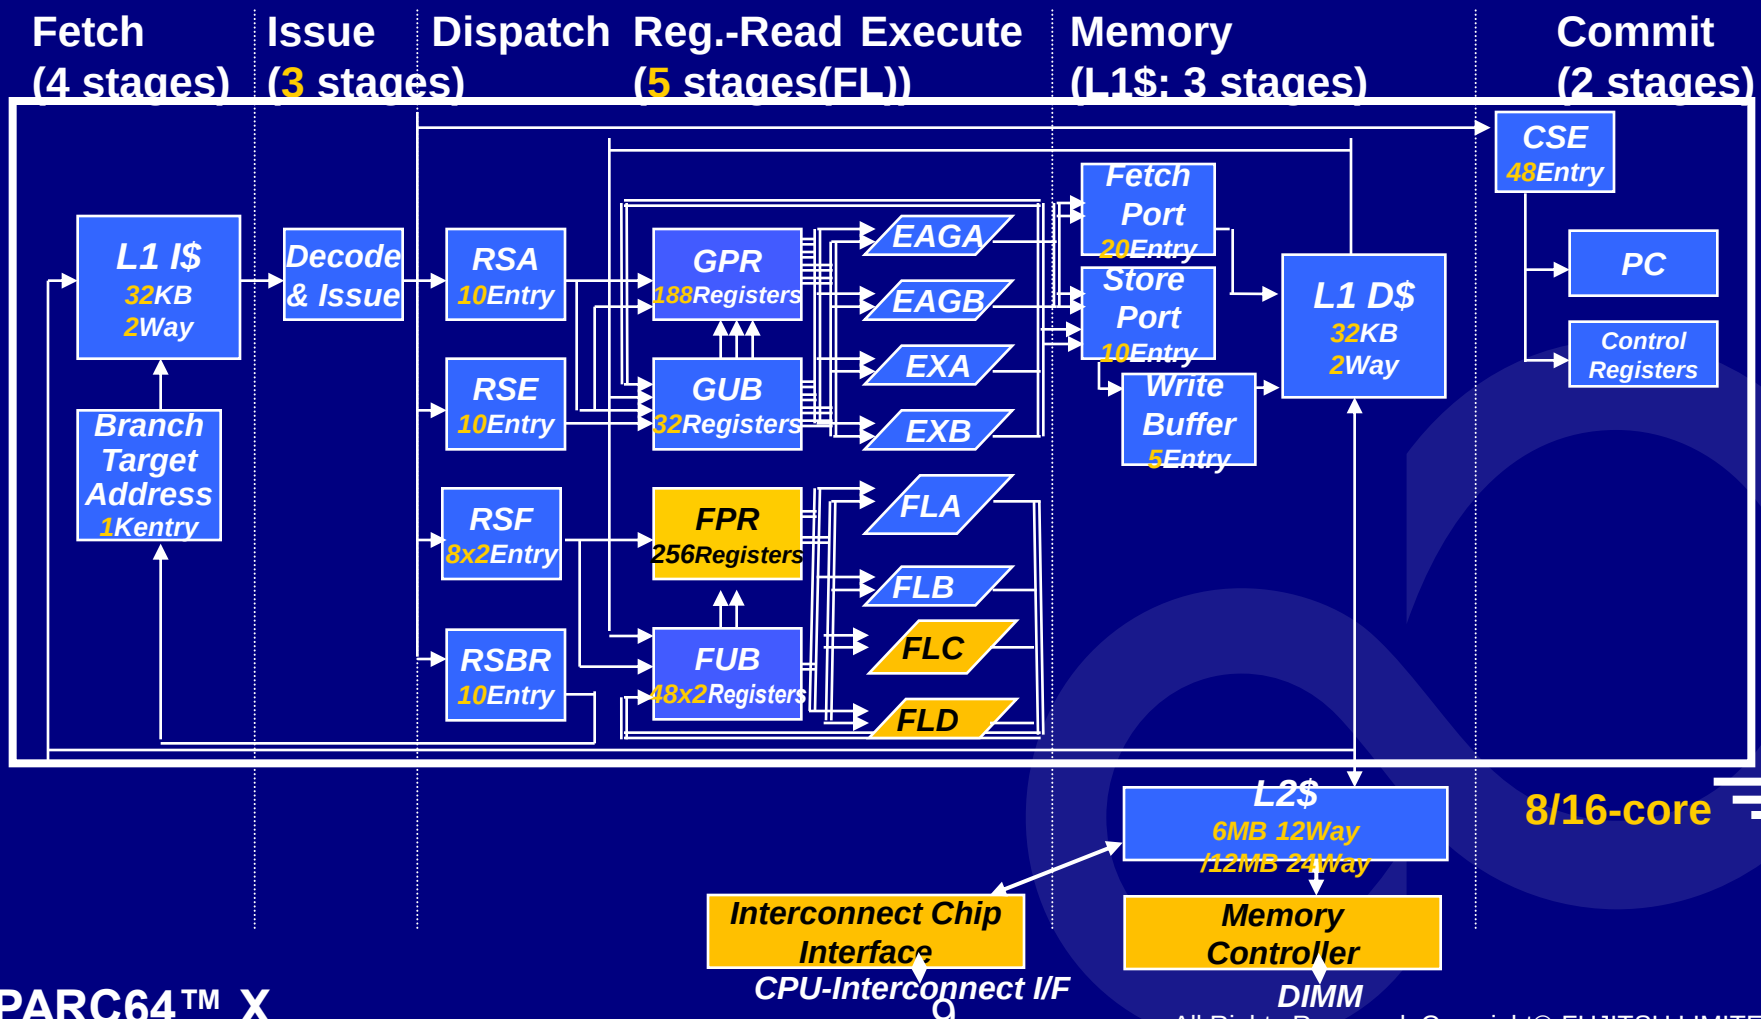

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ X Pipeline (2012)

- ✓16 Cores SMT / Integer Performance Enhancement
- ✓Software on Chip
- ✓System on Chip / 64 Sockets per Node

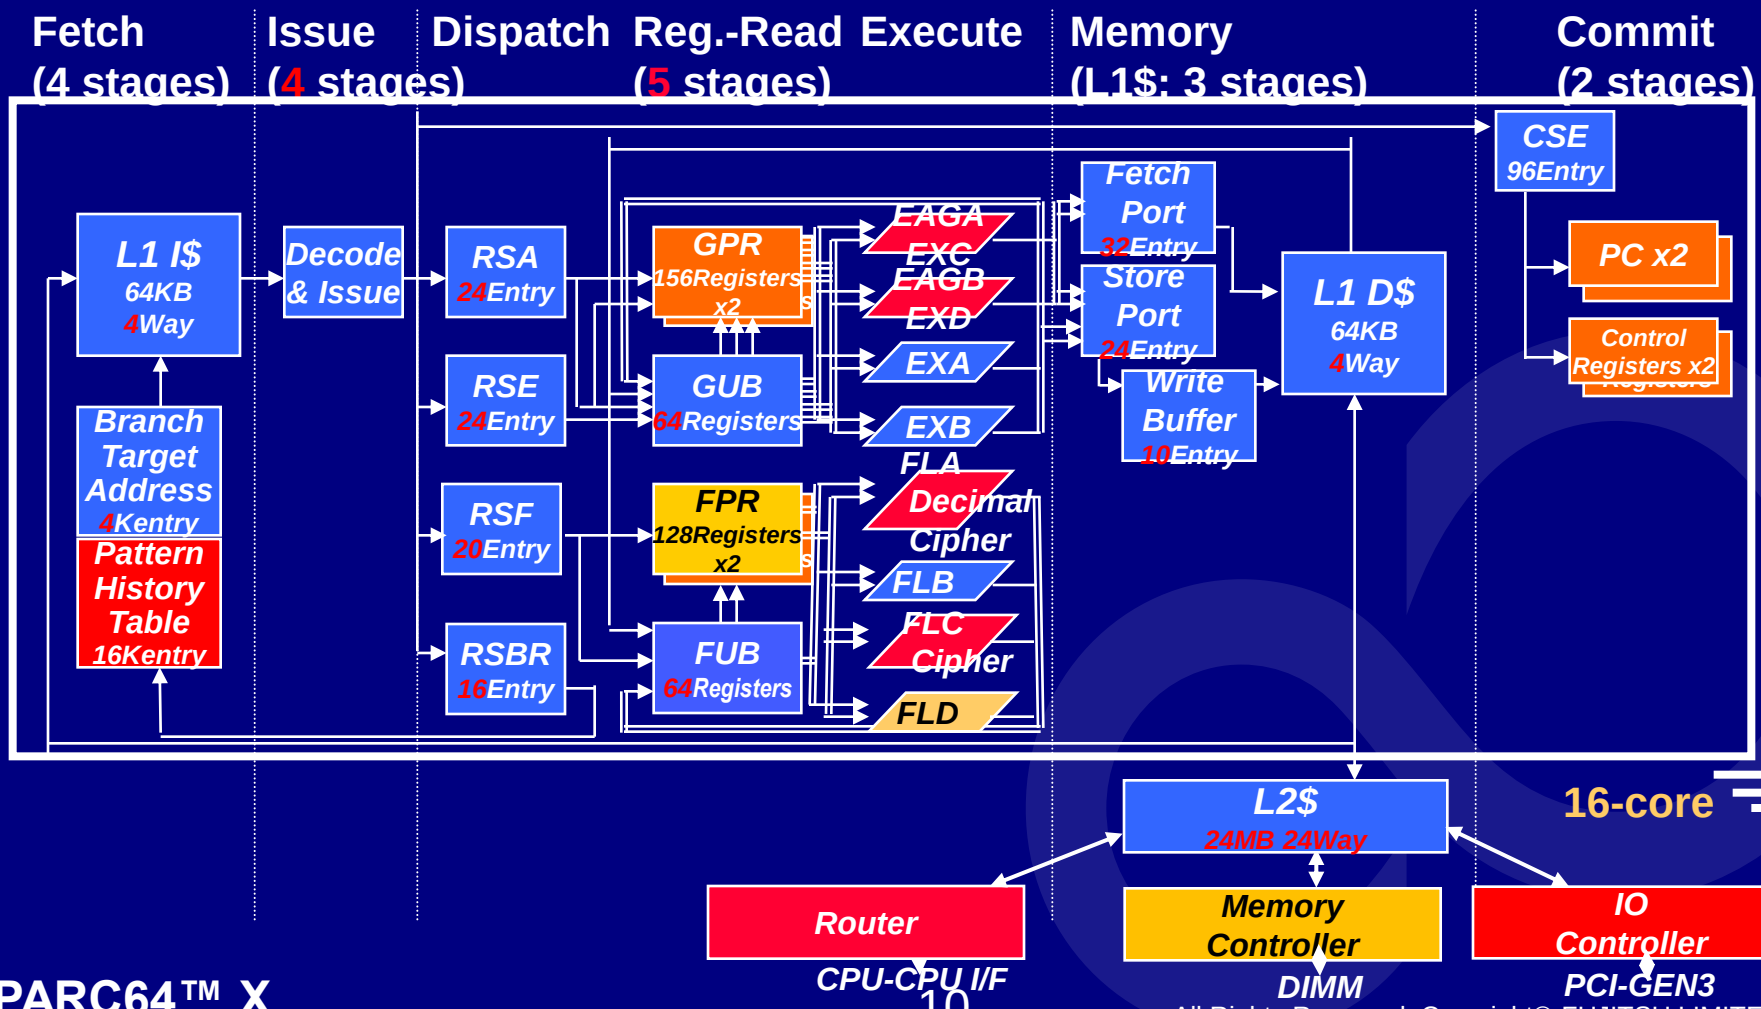

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ X Chip Overview

Photo of SPARC64™X chip

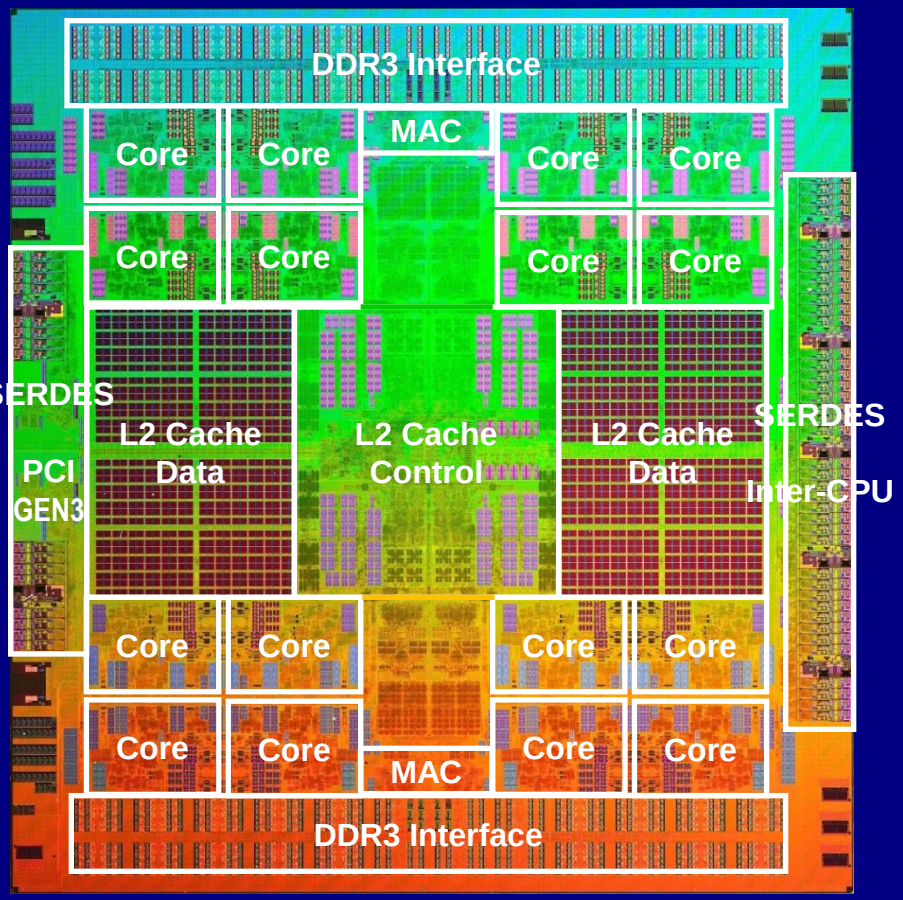

- **Architecture Features**
  - 16 cores x 2 threads
  - SWoC (Software on Chip)
  - Shared 24 MB L2$
  - Embedded Memory and IO Controller
- **28nm CMOS**
  - 23.5mm x 25.0mm
  - 2,950M Transistors
  - 1,500 Signal Pins
  - 3GHz
- **Performance (peak)**
  - 382GFlops
  - 102GB/s Memory Throughput

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ X Core Specification

Photo of SPARC64™X Core

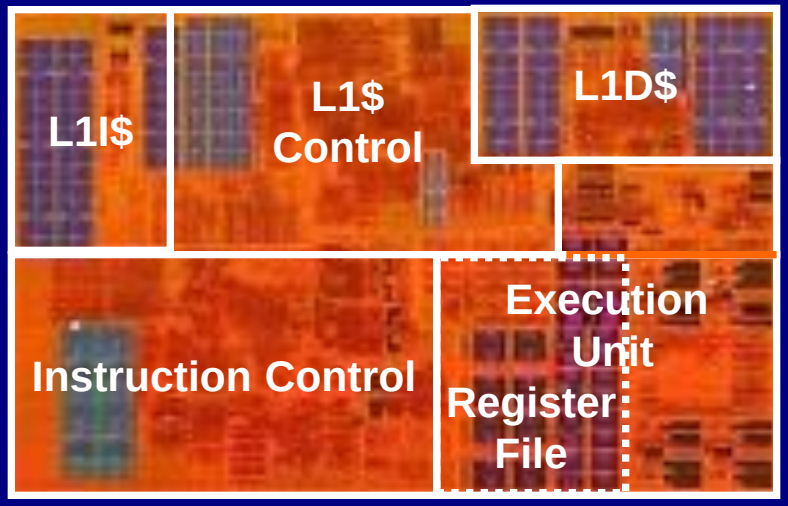

Instruction Set Architecture	SPARC-V9/JPS HPC-ACE VM SWoC
Branch Prediction	4K BRHIS 16K PHT
Integer Execution Units	156 GPR x 2 + 64 GUB
	ALU/SHIFT x2 ALU/AGEN x2 MULT/DIVIDE x1
Floating-Point Execution Units	128 FPR x 2 + 64 FUB
	FMA x4 (2-wide SIMD x2) FDIV x2 IMA/Logic x4 (2-wide SIMDx2) Decimal x1 / Cipher x2
L1$	L1I$ 64KB/4-way L1D$ 64KB/4-way

Underlined Parts Indicate Enhancement from SPARC64™ VII(+)

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Micro-Architecture Enhancements from SPARC64™ VII+

- CPU Core
  - Deeper Pipeline to Increase Frequency
  - Better Branch Prediction Scheme
  - Various Queue-Size and #Floating-Point Register Increase
  - Richer Execution Units, Including
    - 2 EX + 2 EAG → 2 EX + 2 EX/EAG
    - 2 FMA → 4 FMA to support 2-wide SIMD
    - SWoC engine (Decimal and Cipher)
  - More Aggressive O-O-O Execution of Load and Store
  - Multi-banked 2 Port L1-Cache

- System On Chip
  - # of Core and L2$ Size (4 core/12 MB→16 core/24 MB)
  - Memory Controller, IO Controller, and CPU-CPU I/F are Embedded to increase Performance and reduce Cost of Production.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Enhancement on Execution Units

- **Integer Execution Unit**
  - 2 EX + 2 EAG → 2 EX + 2 EX/EAG
  - 2 → 4 Write GPR
  - → 4 Integer Instructions Can be Executed per Cycle (Sustained)

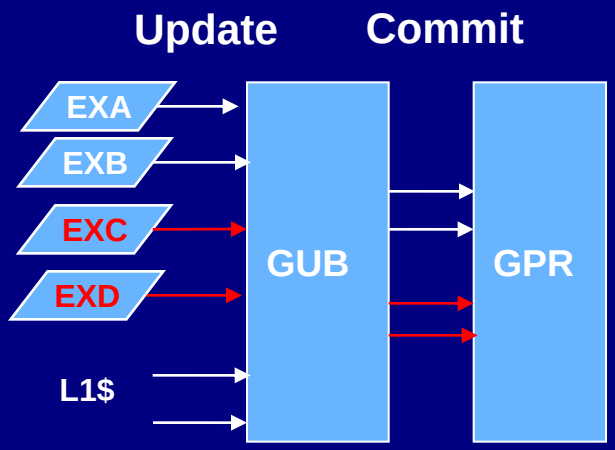

- **Load Store Unit**
  - Aggressive Load/Store O-O-O Execution:
    - Execute load without waiting for preceding store address calculation.
  - Multi-banked 2 Ports L1-cache to Execute 2 Load or 1 Load+1 Store in Parallel
  - Doubled L1$ Bandwidth
  - Doubled L1$ Associativity (2→4-way)
  - → Increase L1-cache Throughput and Hit-rate

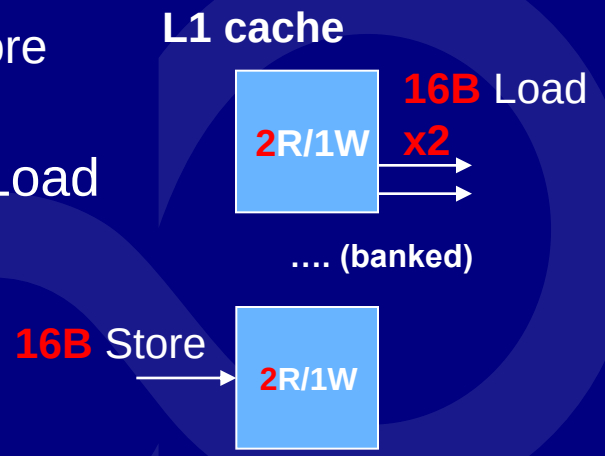

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ X interconnects

SPARC64™ VII/VII+ interconnects
(SPARC Enterprise M8000)

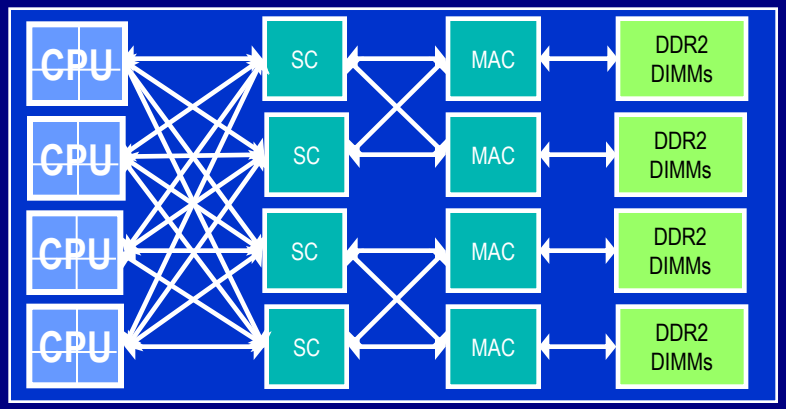

SPARC64™ X interconnects

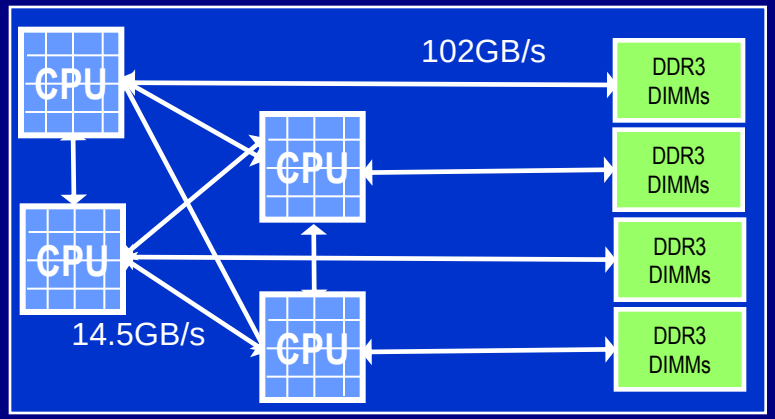

SPARC64™ VII/VII+ Interconnects

- 4 CPU Require 8 additional LSIs to be Connected with DIMM
- DIMM i/f: 4.35GB/s (STREAM triad)

SPARC64™ X interconnects

- No additional LSIs to be Connected with DIMM
- DIMM i/f: 65.6GB/s (STREAM triad)
- CPU i/f: 14.5GB/s x 5 ports (peak)
  - 3 Ports: Glueless 4-way CPU interconnect
  - 2 Ports: > 4-way CPU

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Reliability, Availability, Serviceability

Units	Error Detection and Correction Scheme
Cache (Tag)	ECC Duplicate & Parity
Cache (Data)	ECC Parity
Register	ECC (INT/FP) Parity(Others)
ALU	Parity/Residue
Cache Dynamic Degradation	Yes
HW Instruction Retry	Yes
History	Yes

SPARC64™ X RAS Capability Diagram

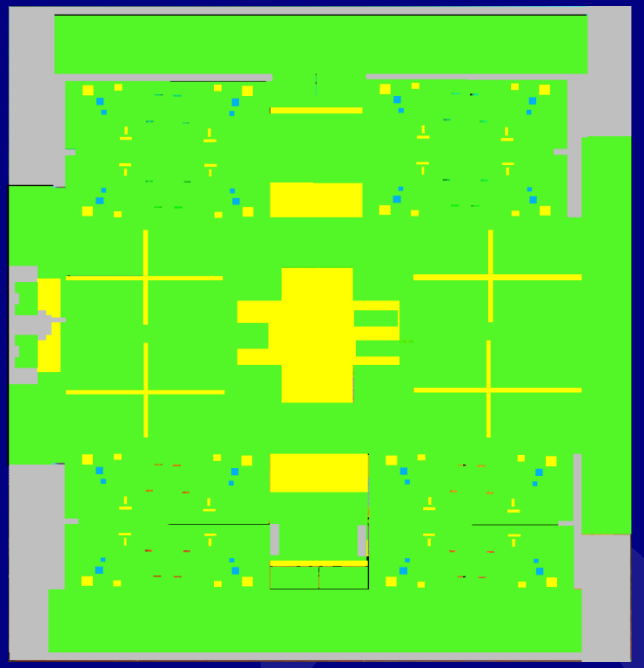

Green: 1bit Error Correctable
Yellow: 1bit Error Detectable
Gray: 1bit Error Harmless

- New RAS Features from SPARC64™ VII/VII+
  - Floating-Point registers are ECC protected
  - #Checkers increased to ~53,000 to identify a failure point more precisely

→ Guarantees Data Integrity

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Hardware Instruction Retry

- When an Error is Detected, Hardware Re-execute the Instruction Automatically to Remove the Transient Error by Itself.

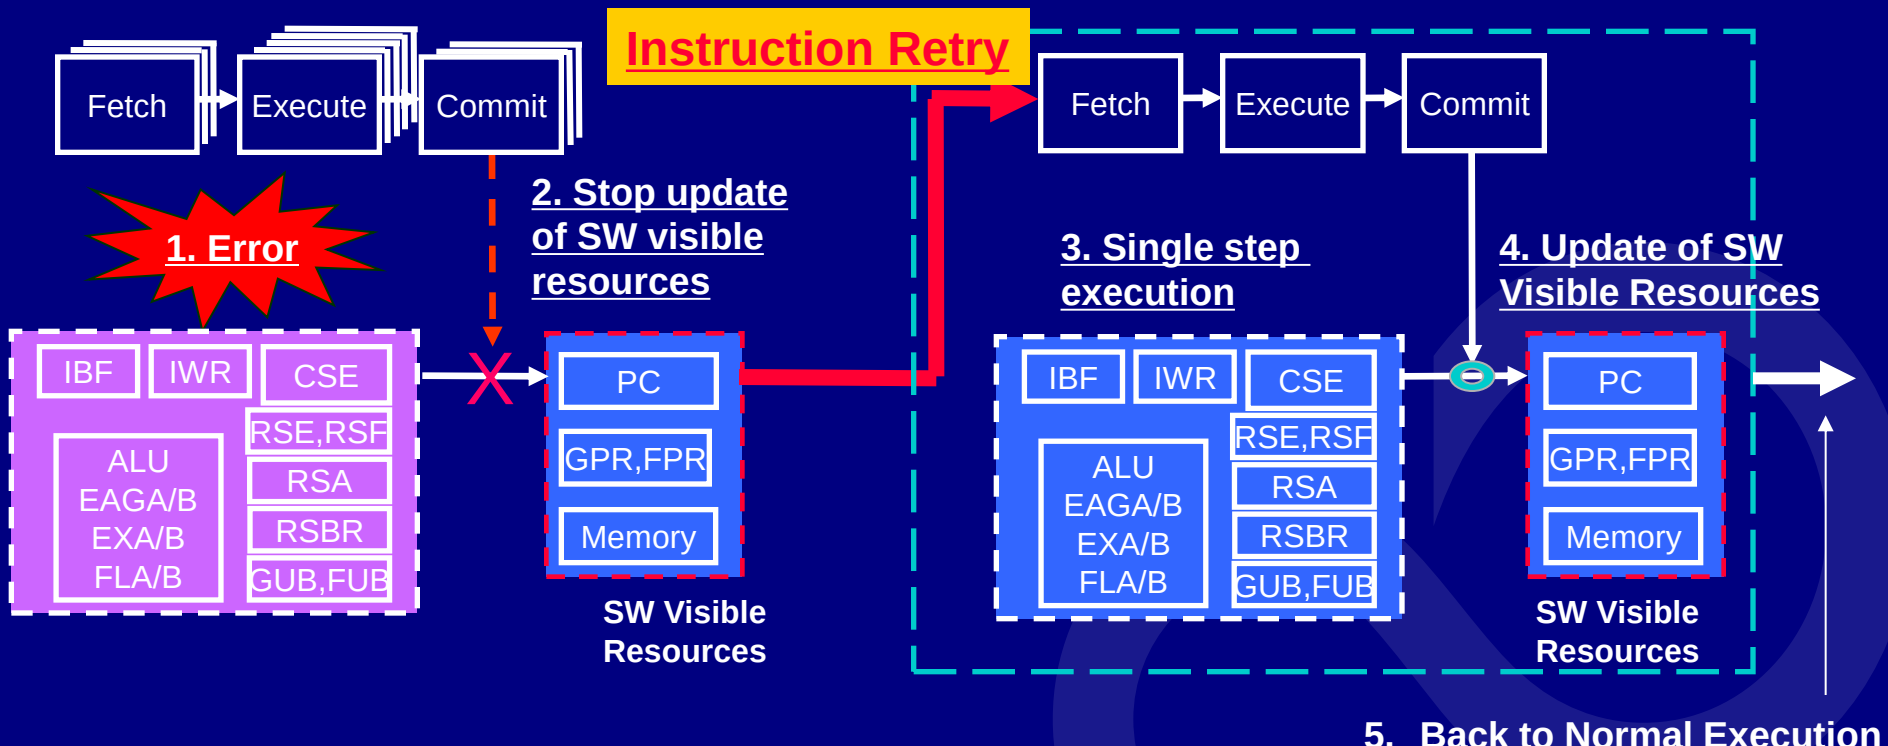

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

SPARC64™ X Performance @3GHz

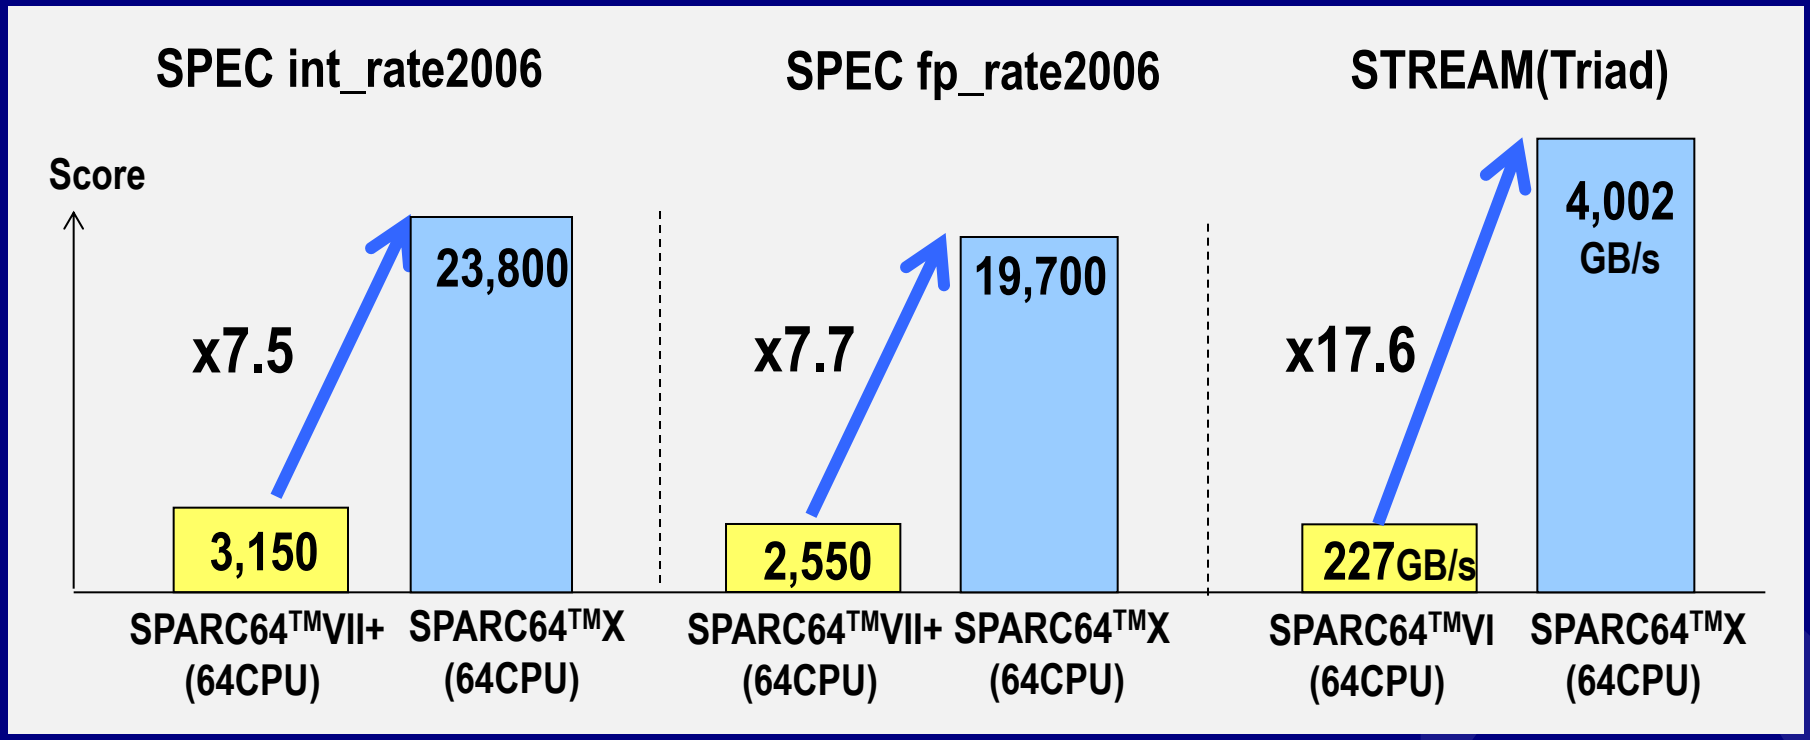

➔ SPARC64™ X Realizes over 7x INT/FP Throughput and 17.6x Memory Throughput of SPARC64™ VII+

SPEC®, SPECint® and SPECfp® are trademarks of SPEC.
Configuration used for measurement: SPARC M10-4S server with 64 SPARC64 X (3.0GHz) processors, Oracle Solaris 11.1, Oracle Solaris Studio 12.3, 1/13 Platform Specific Enhancement. The performance value of SPARC M10 has been submitted to SPEC and the latest information on this benchmark can be found at http://www.spec.org.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

CPI (Cycle Per Instruction) Analysis of SPARC64™ X

SPARC64™ VII+ v.s. SPARC64™ X
INT (single thread)

Hardware measured results

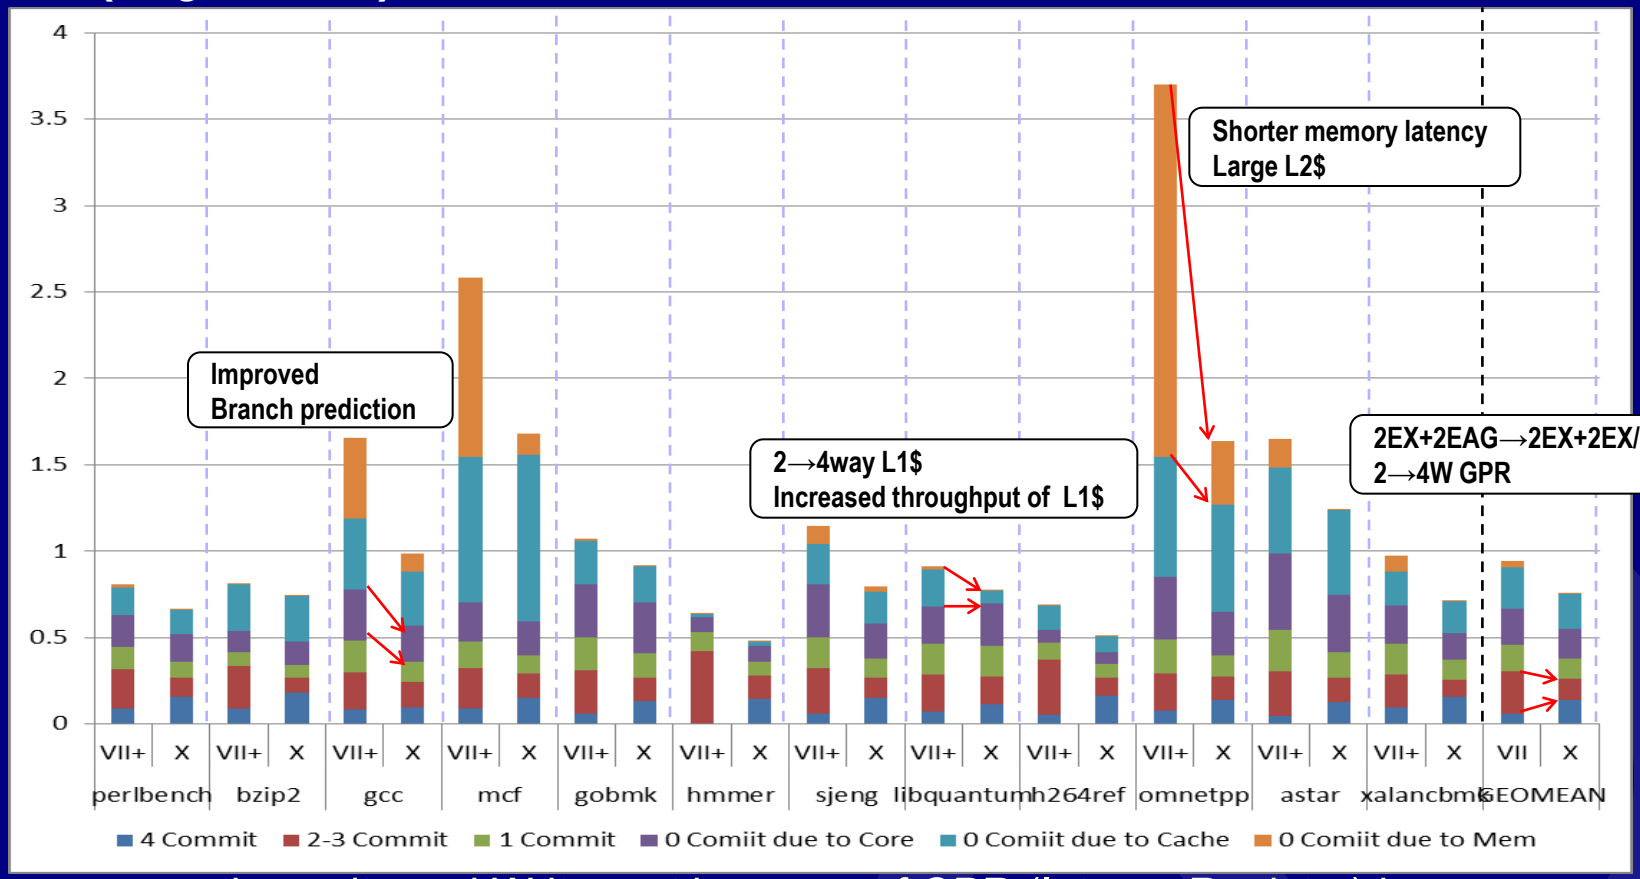

- 4 Integer execution units and Write port increase of GPR (Integer Register) improves overall Performance.
- Memory latency reduction, Large L2$, Branch prediction, and L1$ improvement also contribute to the High Performance dramatically.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Software on Chip 1/3

- Hardware for Software
  Accelerates Specific Software Function by Hardware

- The targets of SPARC64™ X for UNIX servers
  - ✓ Decimal Operation (IEEE754 Decimal and NUMBER)
  - ✓ Cipher Operation (AES/DES/SHA)
  - ✓ Database Acceleration

  - For HPC Apps, supports Trigonometric function, Exponent function.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Software on Chip 2/3

- HW implementation
  - ✓ The HW Engines for SWoC are Implemented in FPU
    - To fully utilize 128 FP registers & software pipelining
    - Area and # Gate is about 2% of Core
  - ✓ Implemented as Instructions rather than dedicated Co-Processor to Maximize Flexibility of SW.
    - 18 insts. for Decimal, and 10 insts. for Cipher operation

Photo of SPARC64™X Core

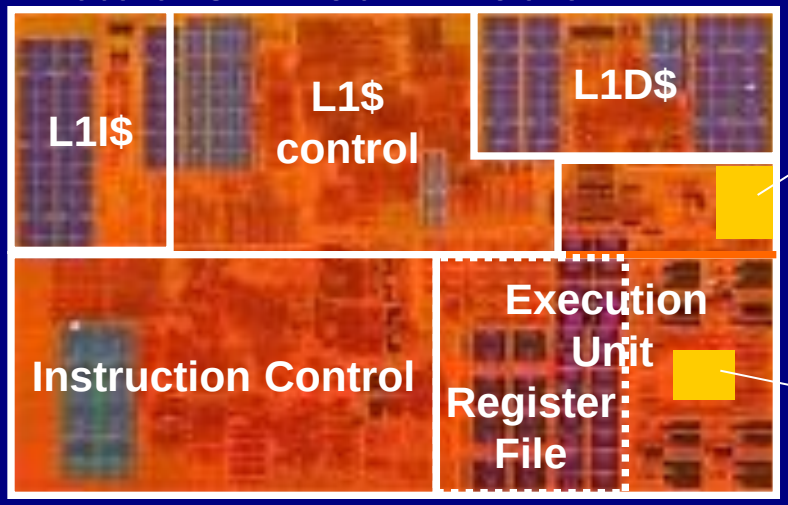

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Software on Chip 3/3

Decimal Instructions

- Supported data type
  - IEEE754 DPD (Densely Packed Decimal)
    8 Byte Fixed Length
  - NUMBER
    Variable Length (Max 21 Byte)

- Instructions
  - Both DPD/NUMBER Instructions are Defined as 8B Operation (Add/Sub/Mul/Div/Cmp) on FP Registers
    - To maximize performance with reasonable hardware cost
    - When the data length is > 8byte, Multiple such instructions will be used.
  - An Instruction for Special Byte-Shift on FP Registers is Newly Added to Support Unaligned NUMBER

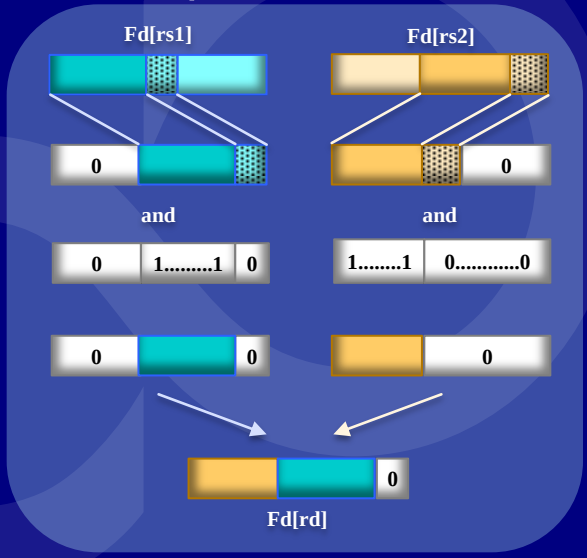

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Advantage of Software on Chip

- Perf. and Power Comparison among SWoC ISA and Non-SWoC ISA on SPARC64™ X
  - ✓ Non-SWoC ISA : Use SPARC-V9 ISA Instead of new SWoC ISA

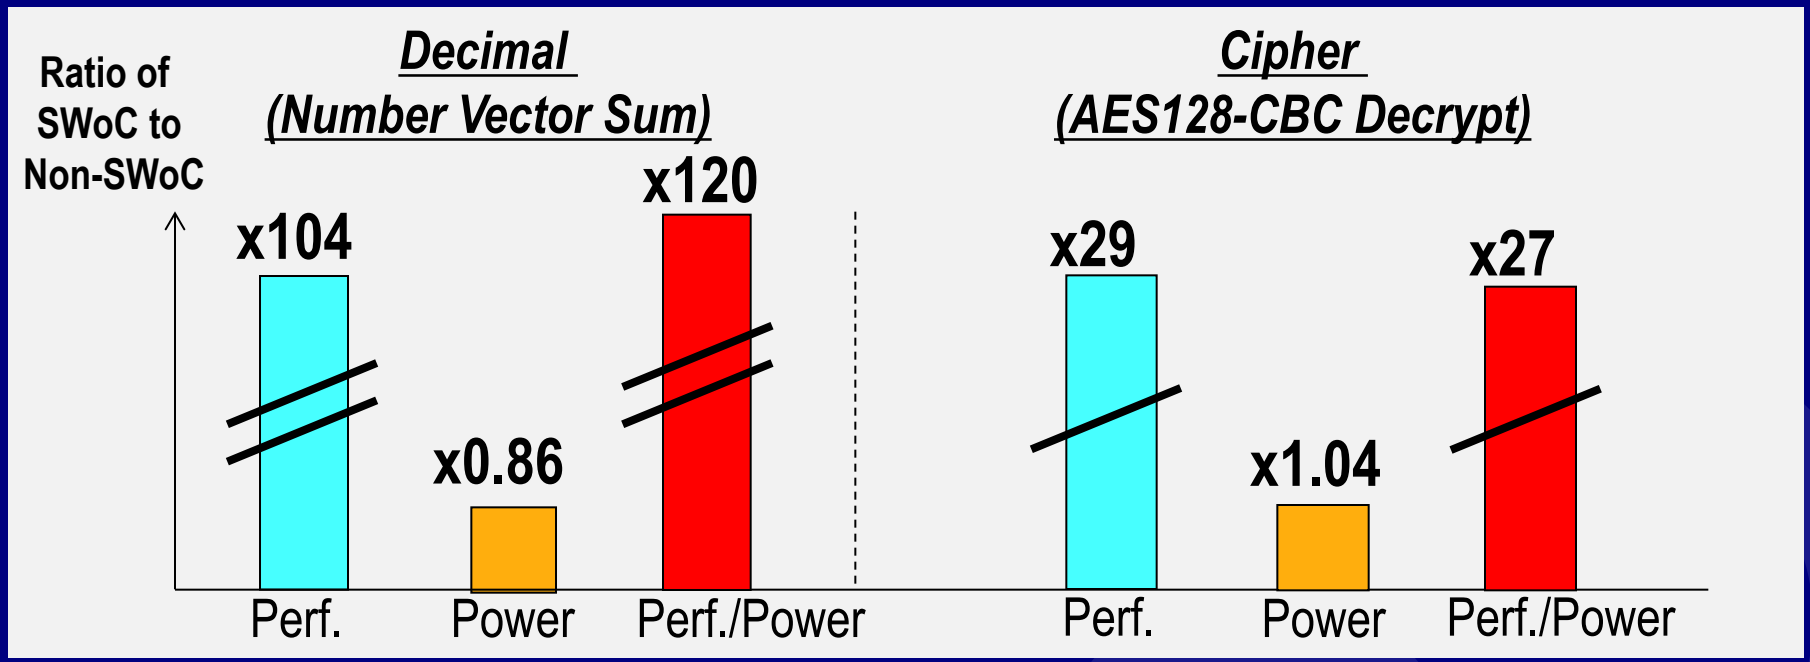

- SWoC Realizes a Significant Improvement on Performance with almost the Same Power Consumption.
  - ✓ SWoC reduces # insts. and pipeline usage for the same task.
- ➔ SWoC is a Power Efficient Technology.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Agenda

- Fujitsu's Processor Development History
- SPARC64™ X
  - Design Concept and Processor Chip Overview
  - Micro-Architecture
  - Performance
  - SWoC (Software on Chip)
- Summary

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Summary

- SPARC64™ X is Fujitsu's10th SPARC Processor for Fujitsu's New Generation UNIX Server.
- SPARC64™ X Integrates 16 Cores + 24MB L2 Cache with over 100GB/s(peak) Memory B/W.
- It Keeps Strong RAS Features.
- SPARC64™ X has Shown over 7x Throughput of SPARC64™ VII+.
- SWoC Accelerates Specific Software Functions with almost the Same Power Consumption.
- Fujitsu will Continue to Develop SPARC64™ series.

All Rights Reserved, Copyright© FUJITSU LIMITED 2013

Abbreviations

- SPARC64™ X
  - IB: Instruction Buffer
  - RSA: Reservation Station for Address generation
  - RSE: Reservation Station for Execution
  - RSF: Reservation Station for Floating-point
  - RSBR: Reservation Station for Branch
  - GUB: General Update Buffer
  - FUB: Floating point Update Buffer
  - GPR: General Purpose Register
  - FPR: Floating Point Register
  - CSE: Commit Stack Entry

All Rights Reserved, Copyright© FUJITSU LIMITED 2013