Everywhere you imagine. RENESAS

SH-X3

Flexible SuperH Multi-core for High-performance and Low-power Embedded Systems

Shinichi Shibahara[1], Masashi Takada[2], Tatsuya Kamei[1], Kiyoshi Hayase[1], Yutaka Yoshida[1], Osamu Nishii[1], Toshihiro Hattori[1]

[1] Renesas Technology Corp.
[2] Hitachi Ltd.

Hot Chips 19, 2007/8/20

Requirement for Embedded Systems

- Trend
  - Total scale is increasing by the introduction of advanced features.

Mobile Phone

2D / 3D	Mail	W-CDMA	GSM
Camera	Video	Sound	JAVA

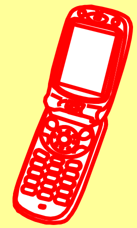

- Requirement
  - High Performance (for advanced features)
  - Small Area (for smaller gadgets)
  - Low Power (for long duration of battery)
- Solution: On-chip Multi-processor
  - Process technology allows to produce easily.
  - MP can be performance and power effective.

Multi-processor Approaches for Embedded Systems

- Features
  - Application Specific: Optimize hardware & software for each system
  - Low Power: Less than 1W (in the case of battery-run)
- Approaches
  - Heterogeneous / Homogeneous AMP
    - Integration of sub-systems / Deterministic behavior
  - Homogeneous SMP
    - Relatively easy programming model / Performance oriented
  - Hybrid (Mixed system of AMP and SMP)
  - ✓ Automatic Parallelizing Compiler

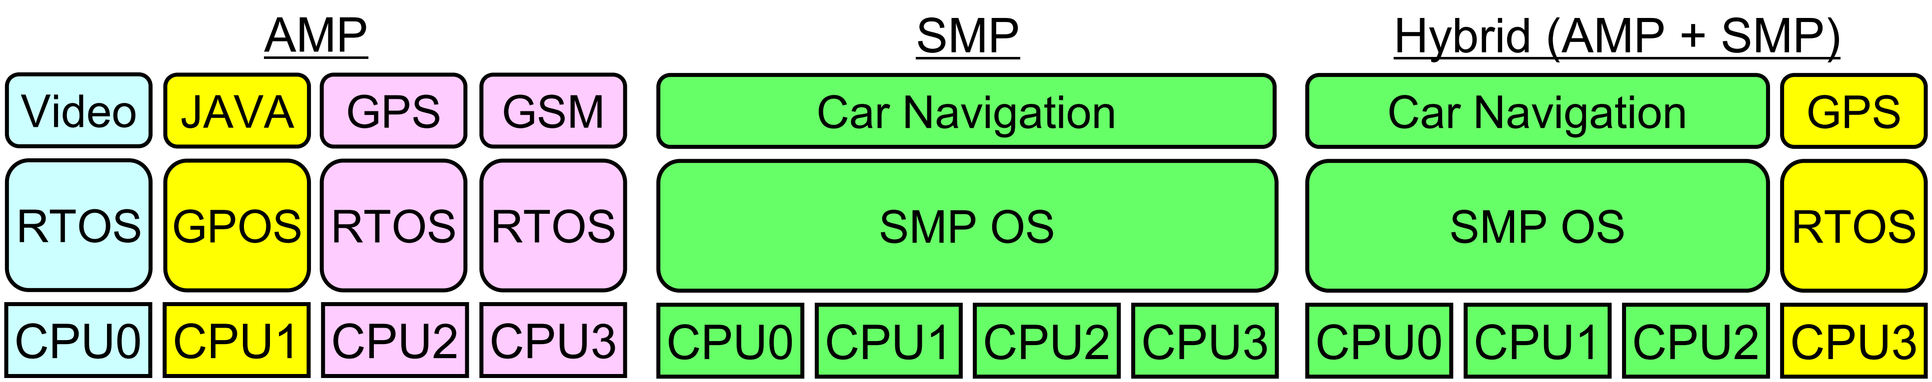

SuperH Processor Core Roadmap

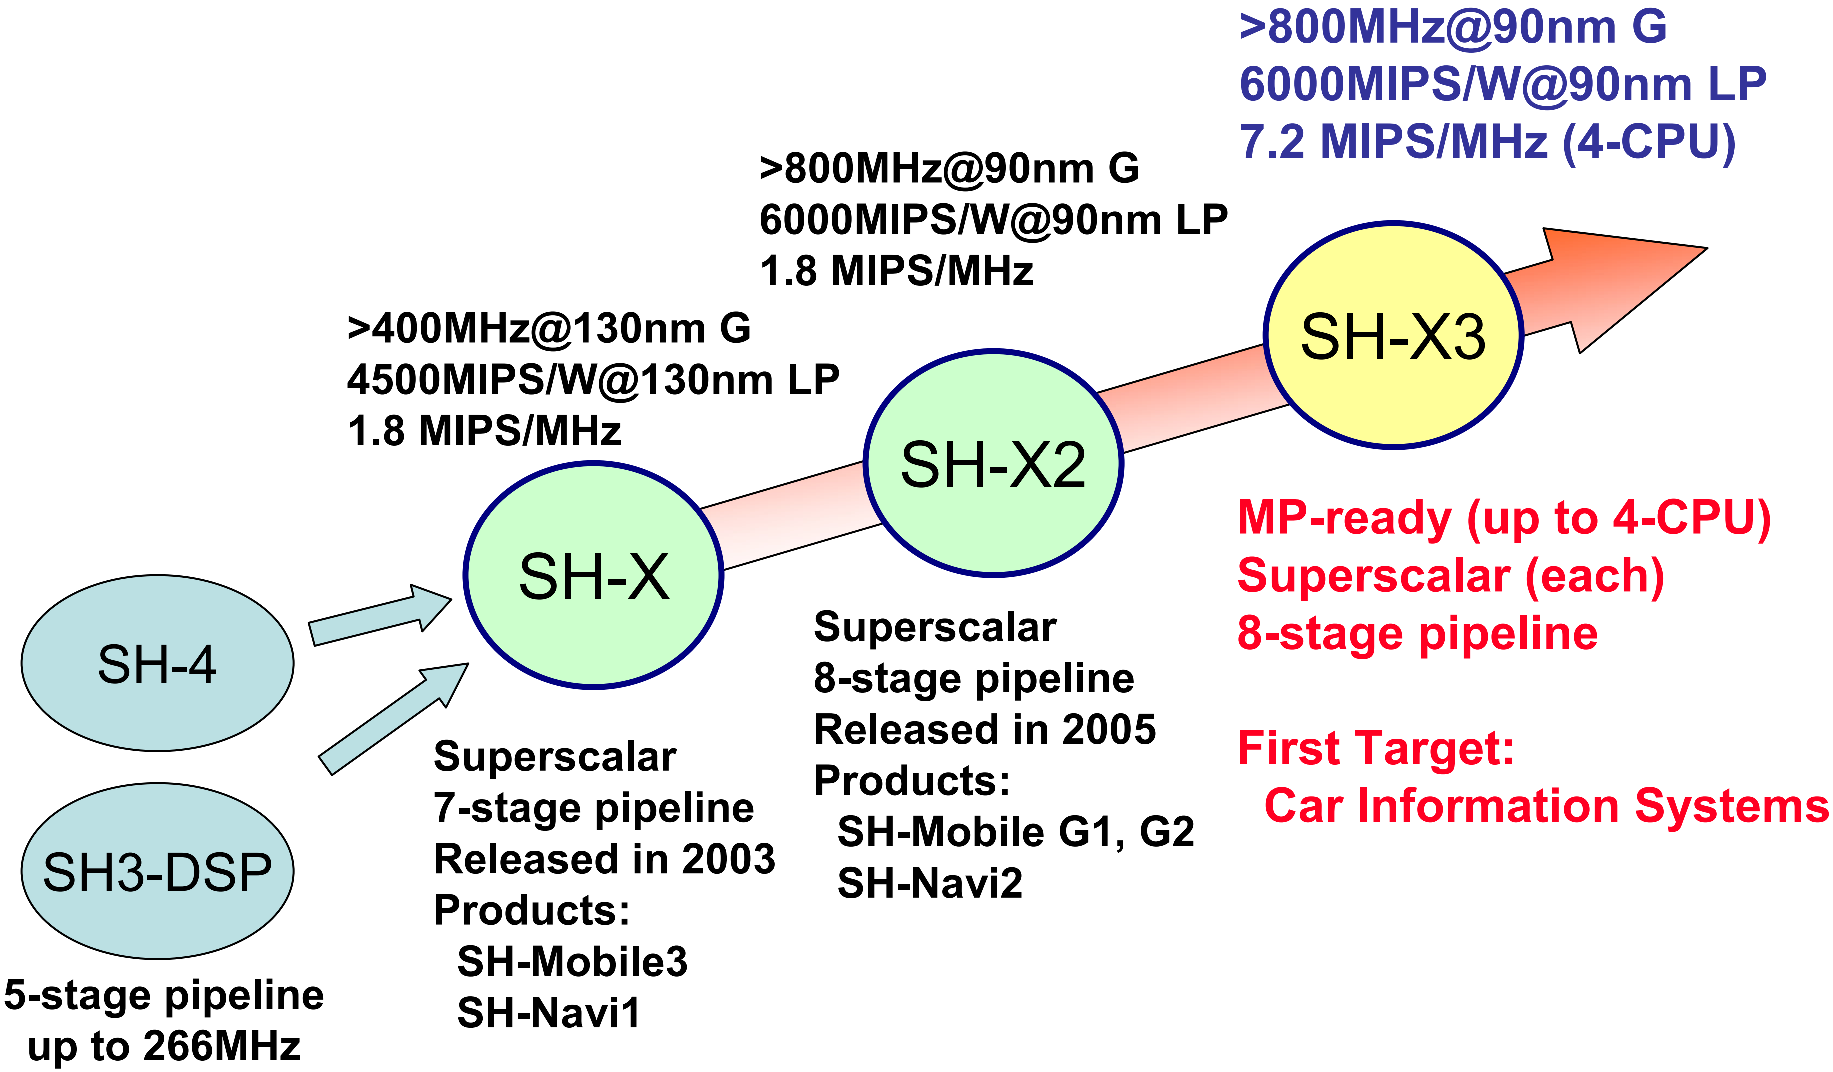

SH-X3 Block Diagram

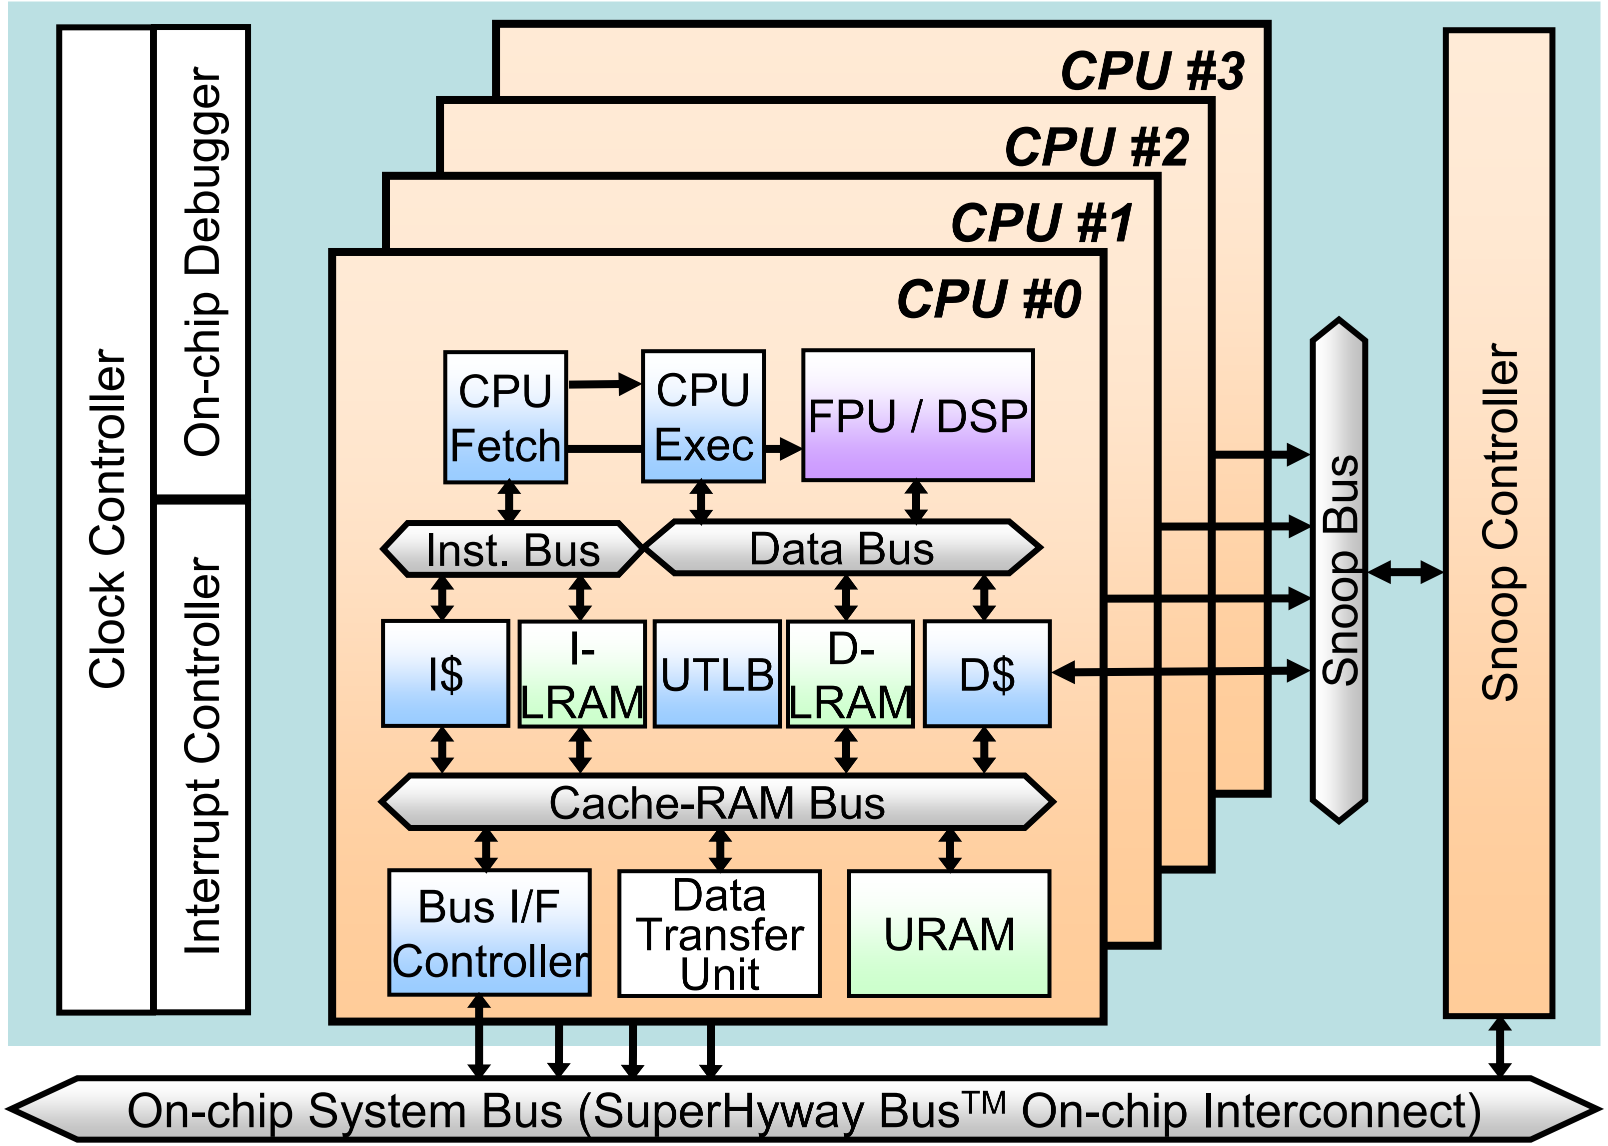

D-LRAM: Also called XY-RAM in SH4AL-DSP

Pipeline Structure

- Eight-stage dual-issue superscalar pipeline (Inherited from SH-X2)

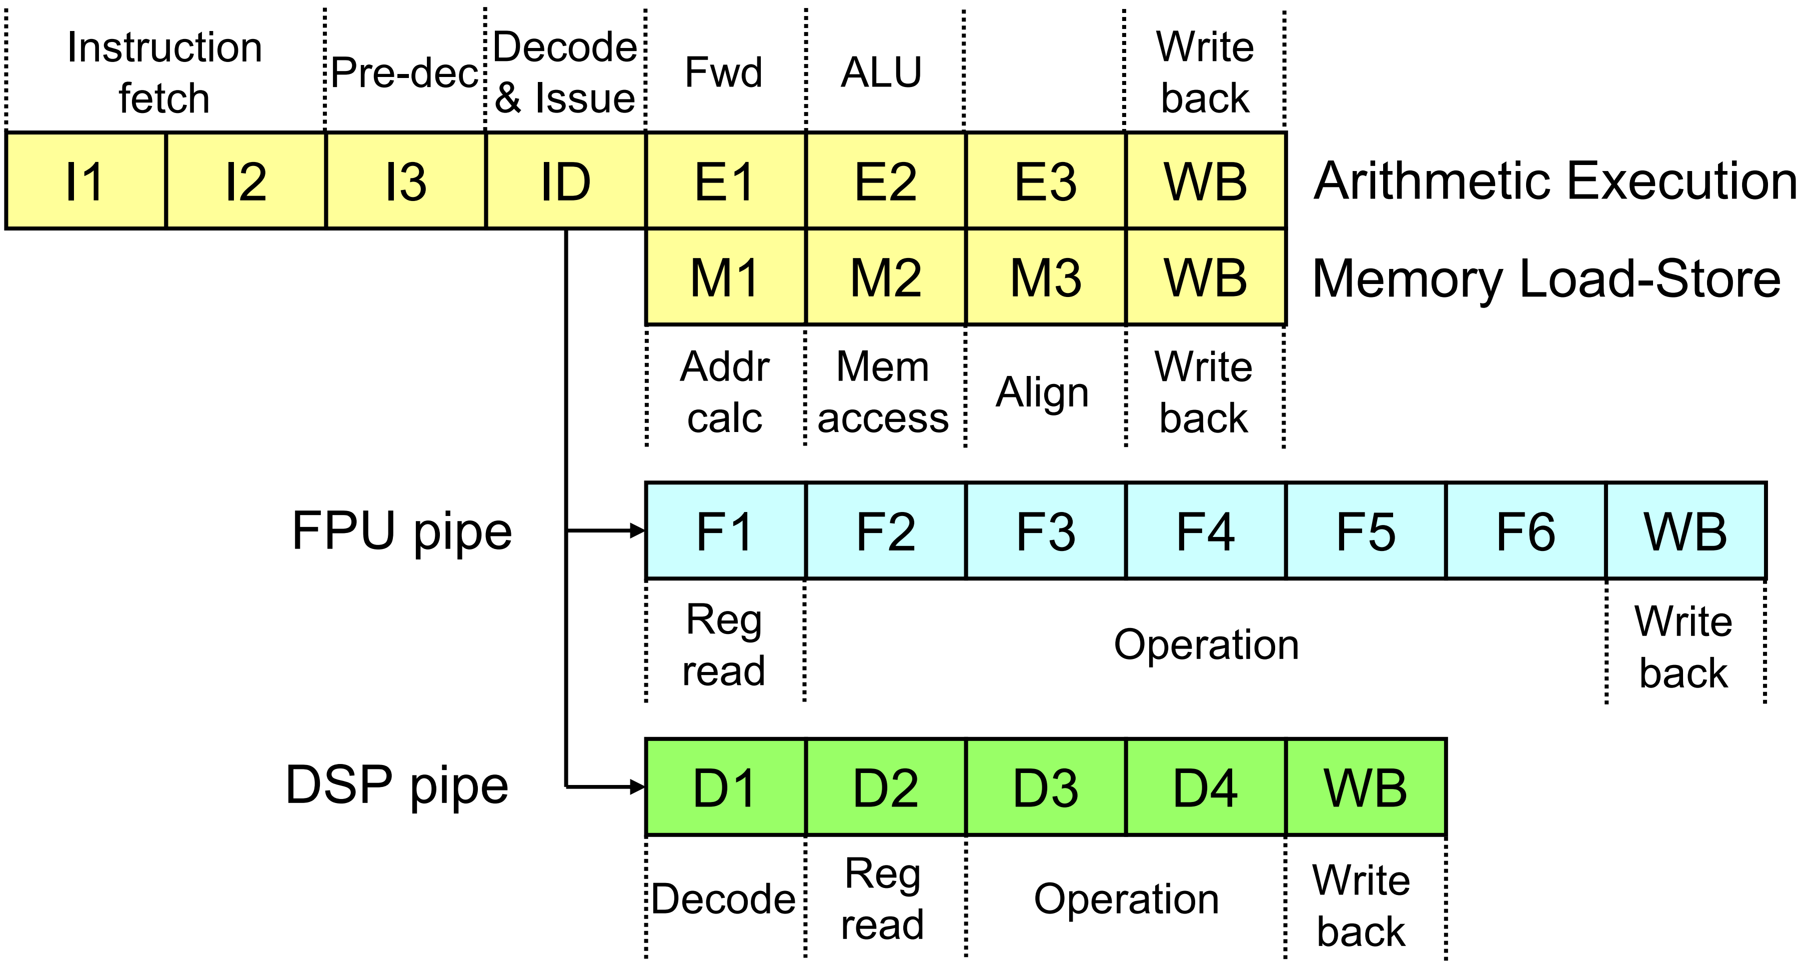

Specification Features

- Efficient for both SMP and AMP
  - Cache coherency (Snoop Controller) for SMP ← Today's Topic
  - Local memories (LRAM, URAM) and data transfer unit for AMP
  - Realization of hybrid MP model
- Fine power management for each CPU
  - Low-power modes according to workload (sleep, light sleep, standby etc.)
  - Flexible clock ratio (CPU Clock : System Bus Clock = m:n (m≥n), 1:n)
  - Hierarchical clock gating
- Configurable and synthesizable
  - Number of CPU (up to 4-CPU), Co-processor (DSP, FPU)
  - Cache (8KB~64KB/4way)
  - Local memory (LRAM: 4KB~128KB, URAM: 128KB~1MB)

Cache Coherency for Embedded Systems

- Problems in applying bus snooping (used in HPC servers)
  - Performance degradation by system bus occupation
  - Unnecessary power dissipation by snooping activity
  - Fixed write-cache mode: MESI (Copy-back) or ESI (Write-through)
    - MESI Fixed: HW accelerator on system bus cannot access the latest data.
    - ESI Fixed: CPU cannot run at the best performance due to store accesses.

- Solution
  - Separation of system bus and snoop bus (Reduce bus occupation)
  - Centralized coherency control by snoop controller (Reduce bus activity)
  - Support of mixed cache coherency protocol (Each CPU can select mode)

Cache Coherency Maintenance (Conventional: Bus Snooping)

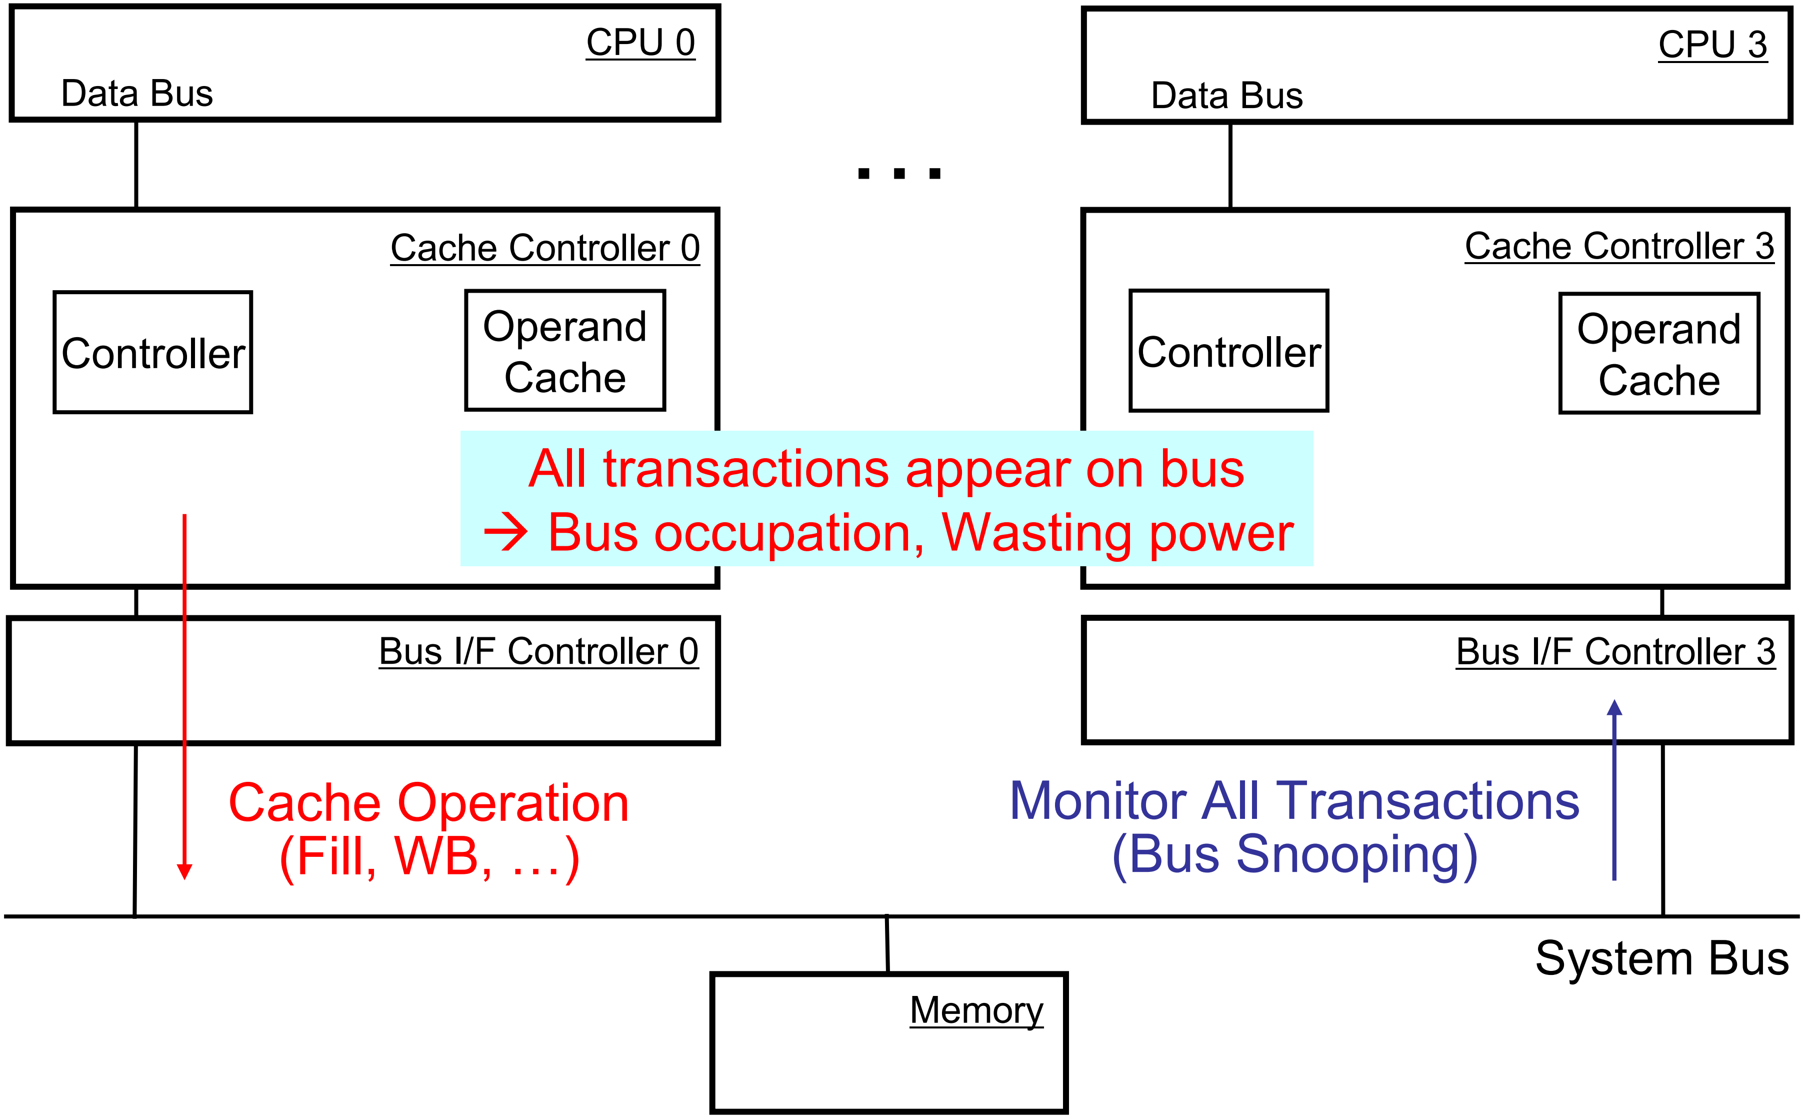

Cache Coherency Maintenance (MESI Protocol)

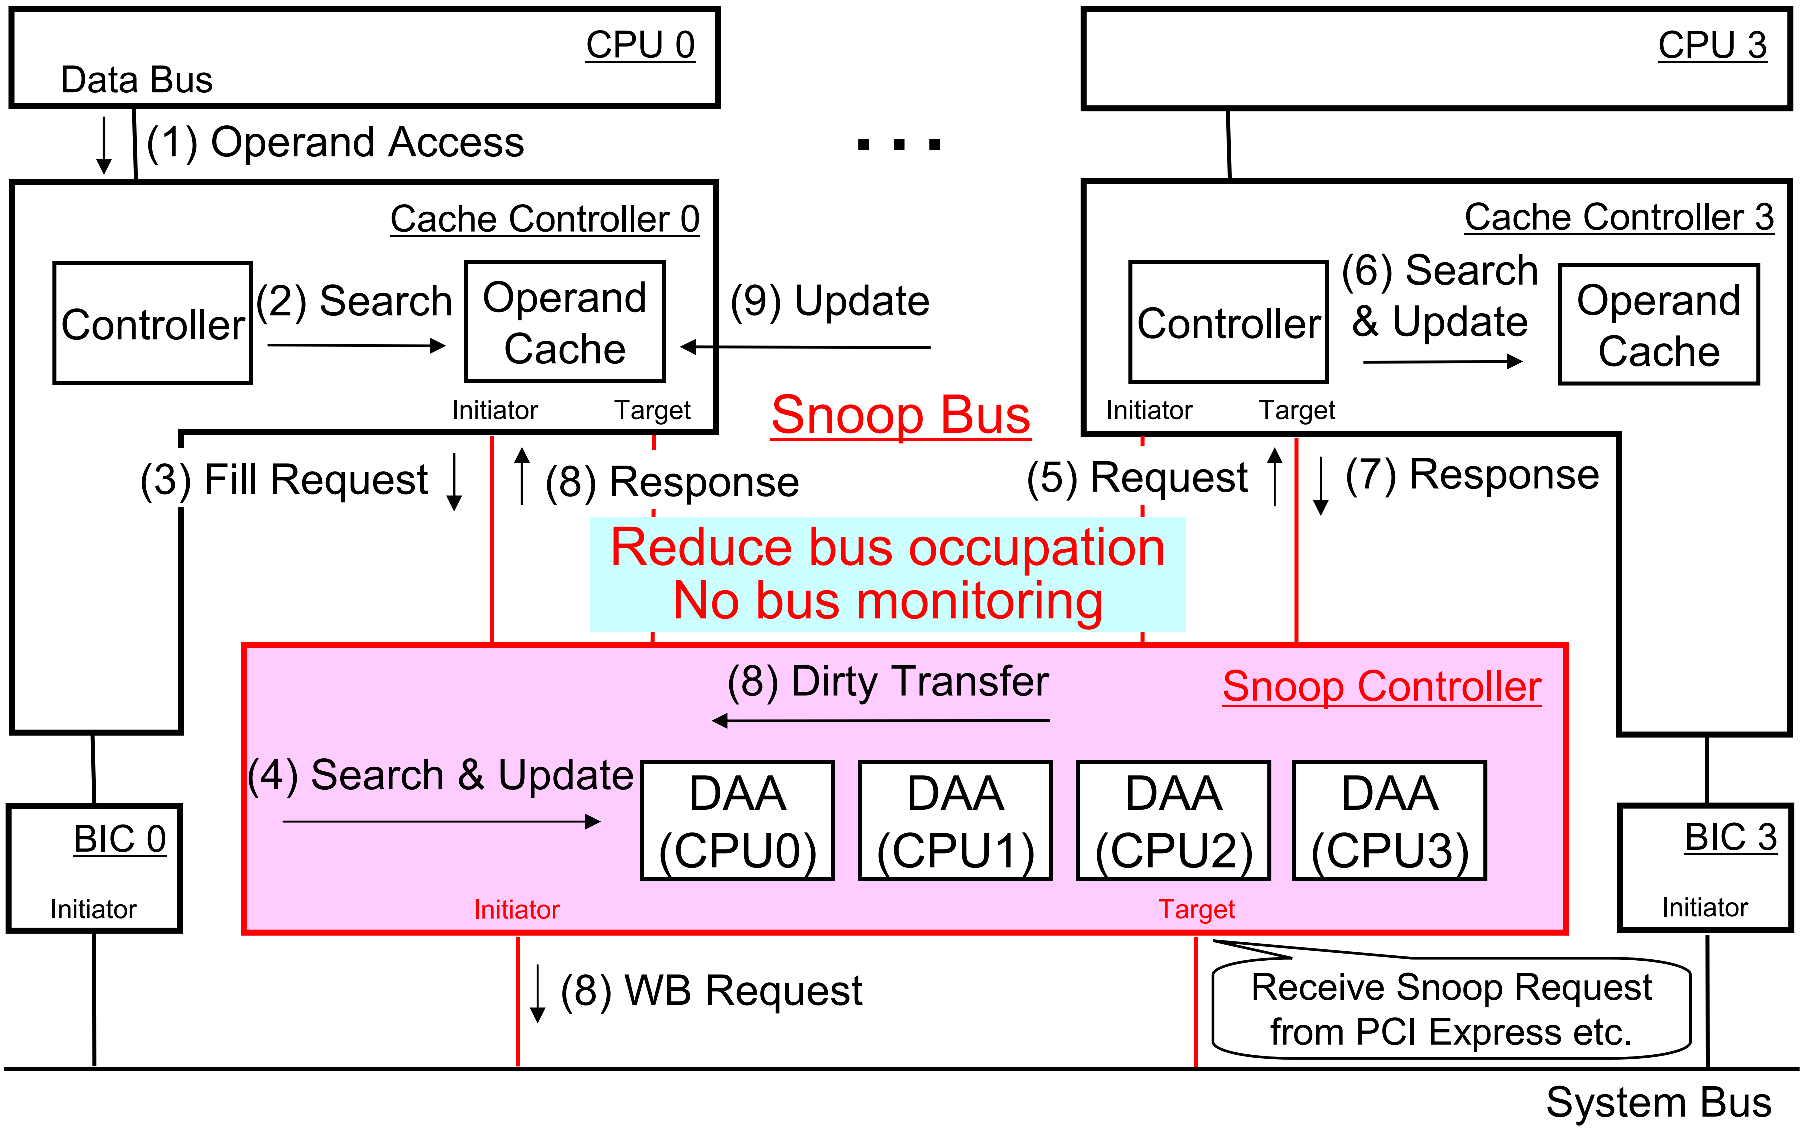

BIC: Bus I/F Controller, DAA: Duplicated Address Array

Snoop Latency Optimization (ESI Protocol)

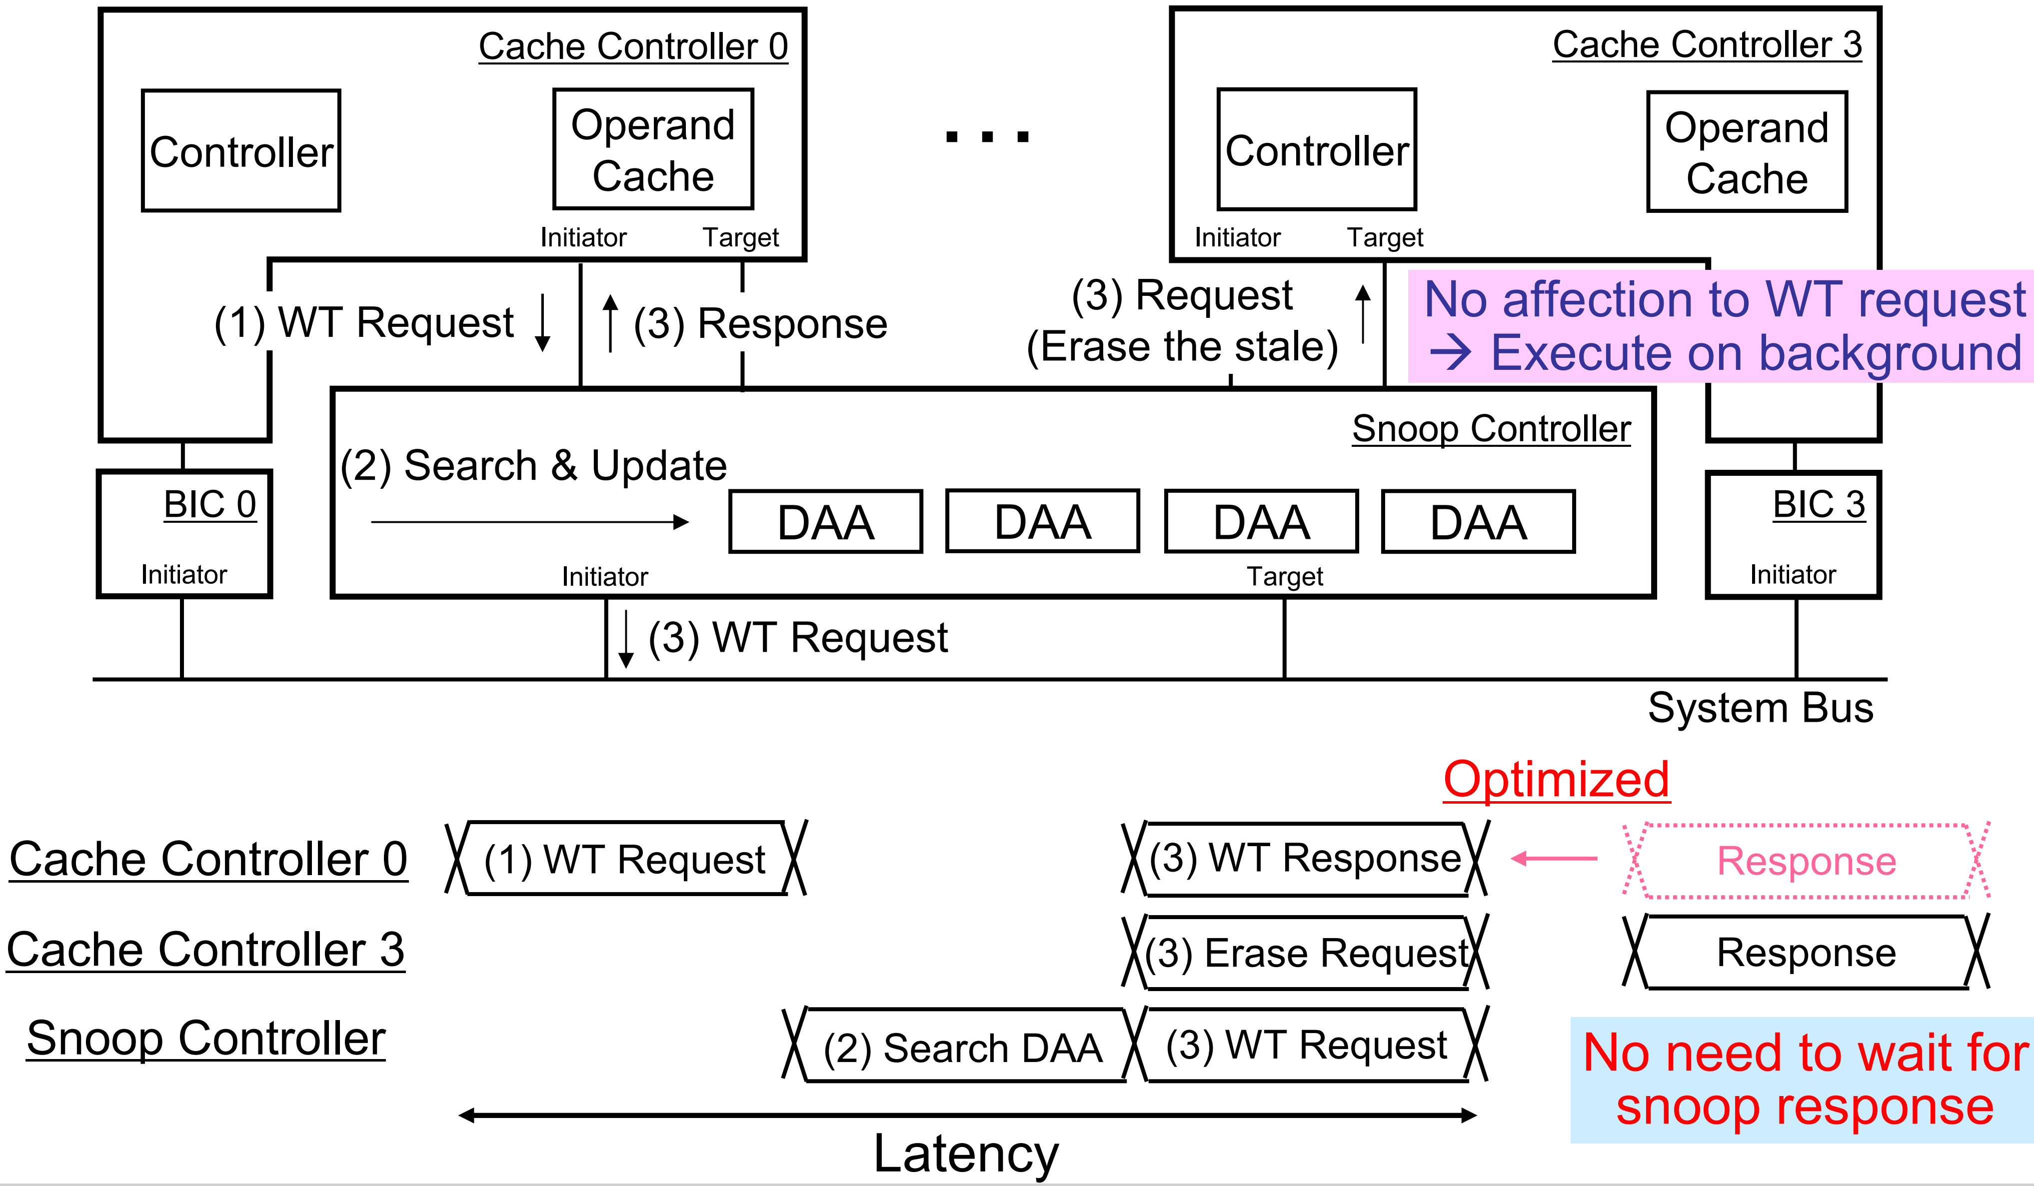

Mixed Coherency Protocol

Need to consider the latest in other CPUs

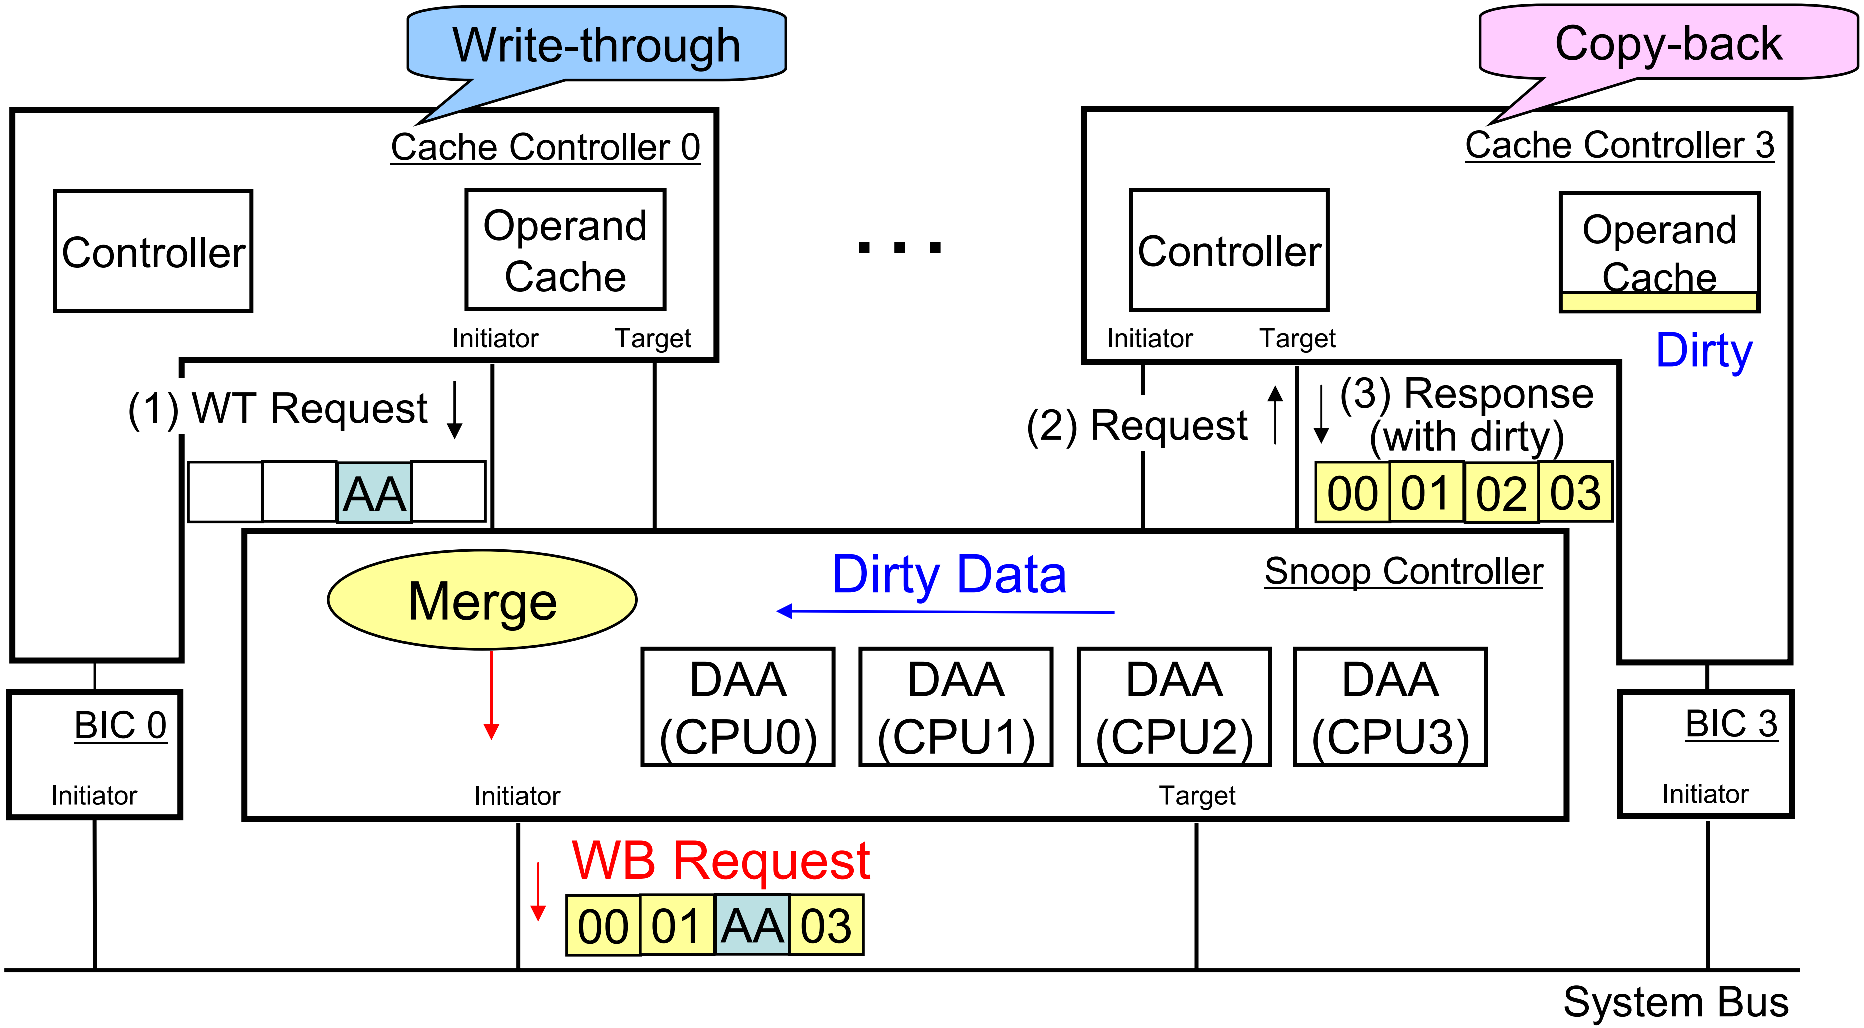

Everywhere you imagine. RENESAS

Difficulty of SMP OS Development

- Cache operation after process migration (Caused by time sharing processing)
  - Conventional measure
    - Flushing cache entries, accessed before, via inter-processor interrupt after process releases memory or virtual-physical address map is changed.
- Synonym problem (Caused by more than one virtual-physical address maps)
  - Conventional measure
    - Flushing the cache of synonym page during page allocation
    - Preventing the synonym occurrence by using page coloring

Problem in conventional measure

- Complicated to implement software
- Large software overhead

Solution

- Broadcast of operand cache operating instructions (OCBI, OCBP, OCBWB)
- Hardware implementation of synonym detection and eviction

Operand Cache Operating Instructions (Conventional)

Conventional Specification

Operate only my own cache line → Cannot operate other CPUs' cache line
→ Need inter-processor interrupt to operate others'

Example: Process Migration

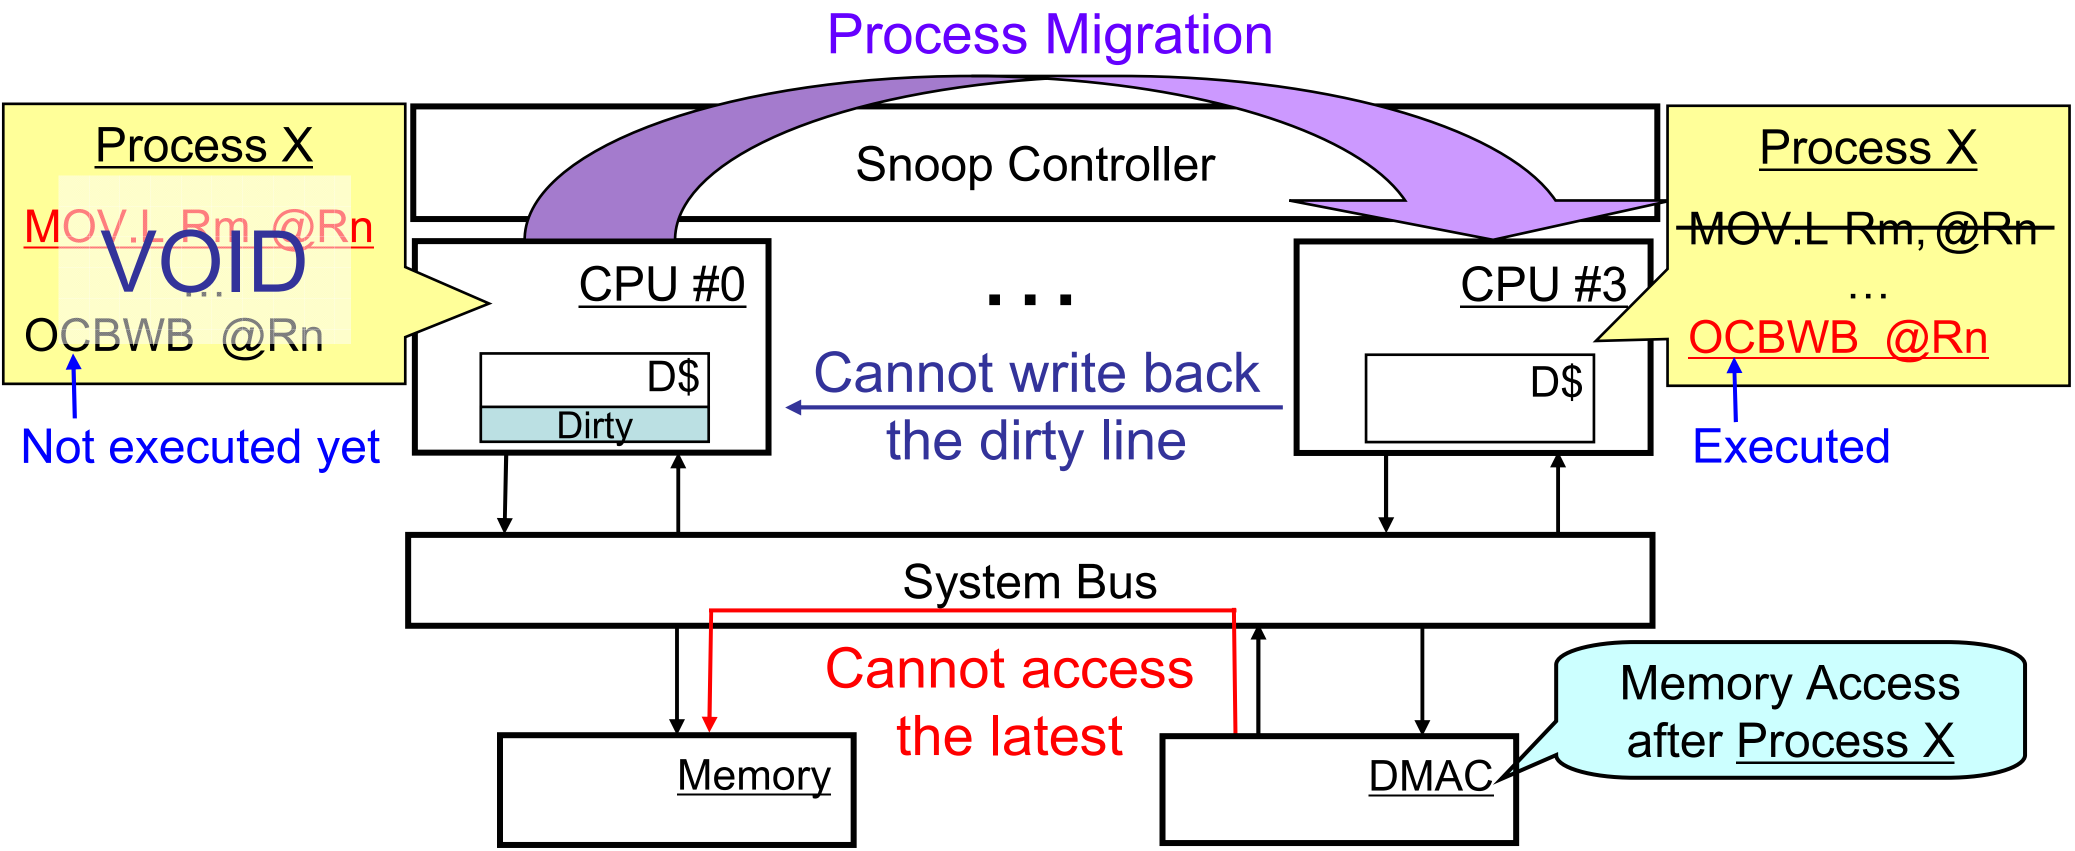

OCBWB: Write-back Cache Block

Broadcast of Operand Cache Operating Instructions

Extended Specification

Operate all CPUs' cache by broadcast → No need inter-processor interrupt to operate

→ Reduce software overhead of using interrupt

Example: Process Migration

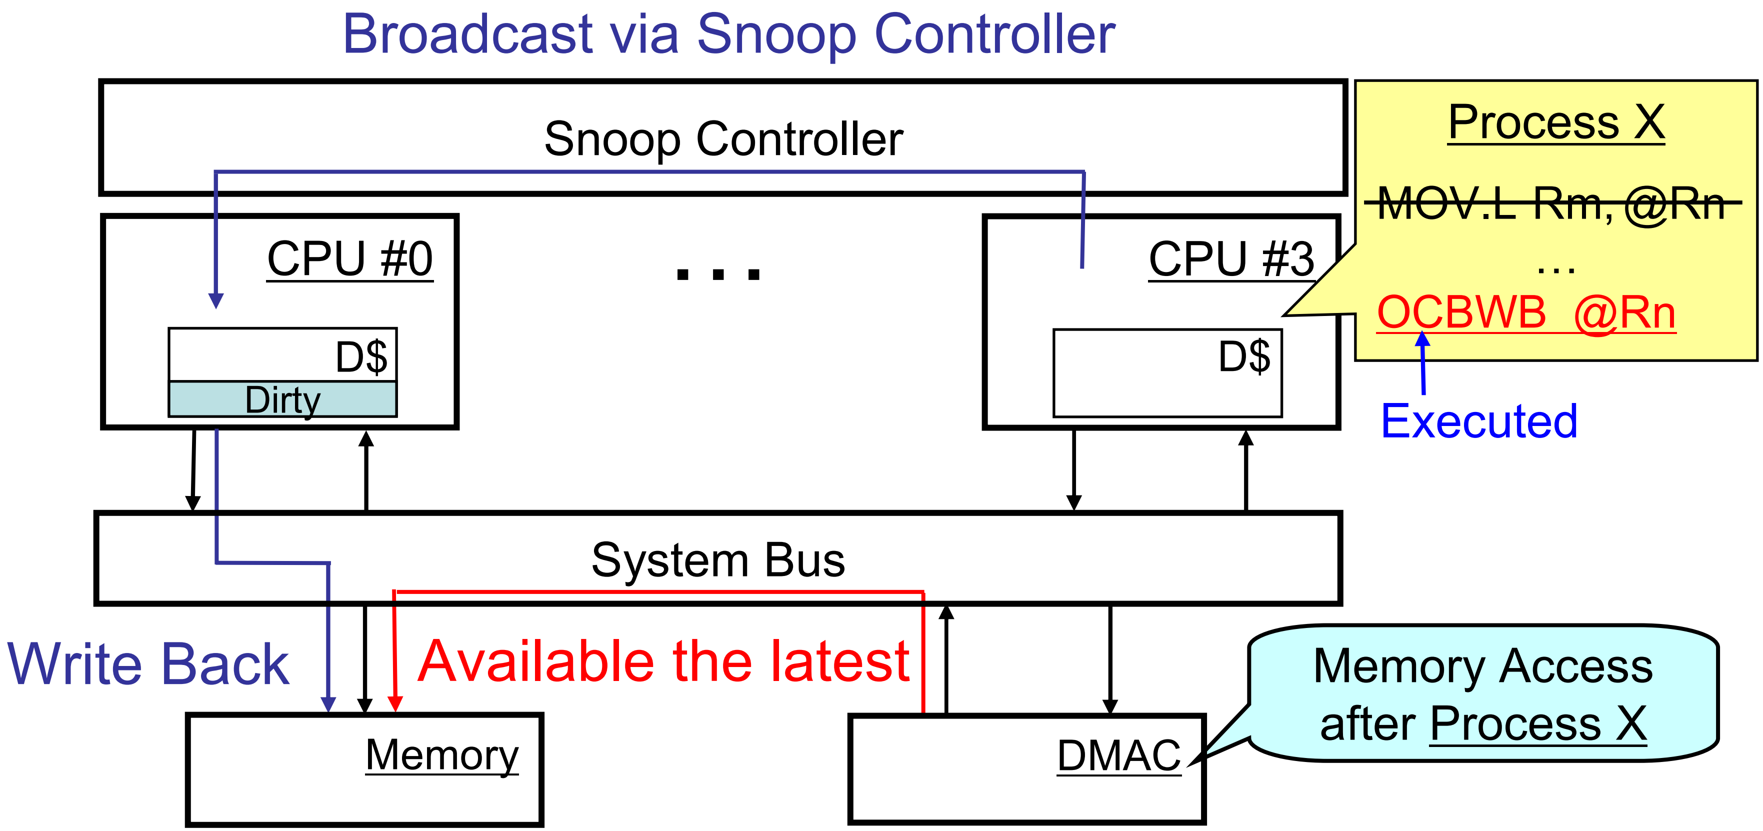

OCBWB: Write-back Cache Block

Synonym Detection and Eviction (In the case of 4KB/Page)

VPN[31:12] → PPN[31:12]

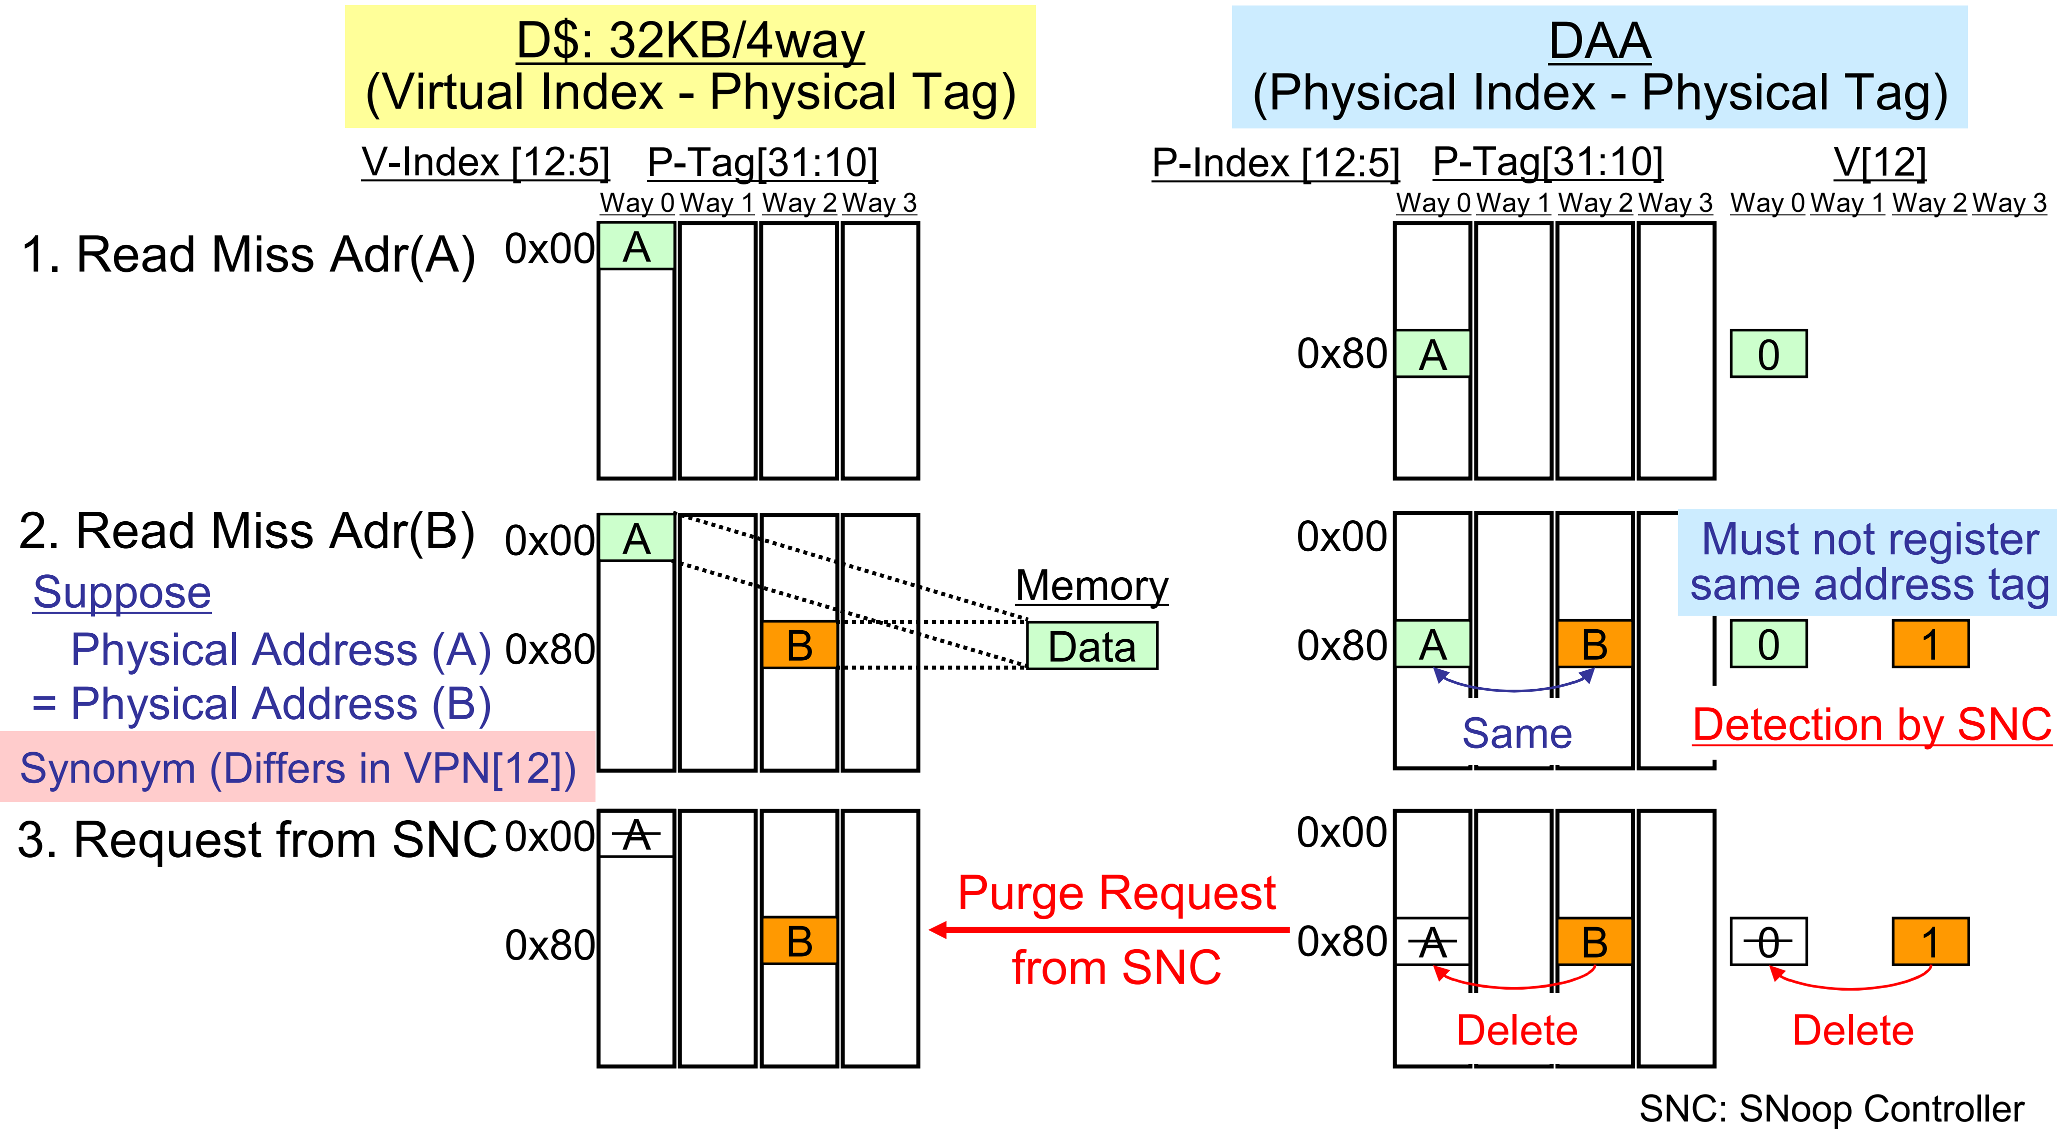

RP1: Experimental Chip

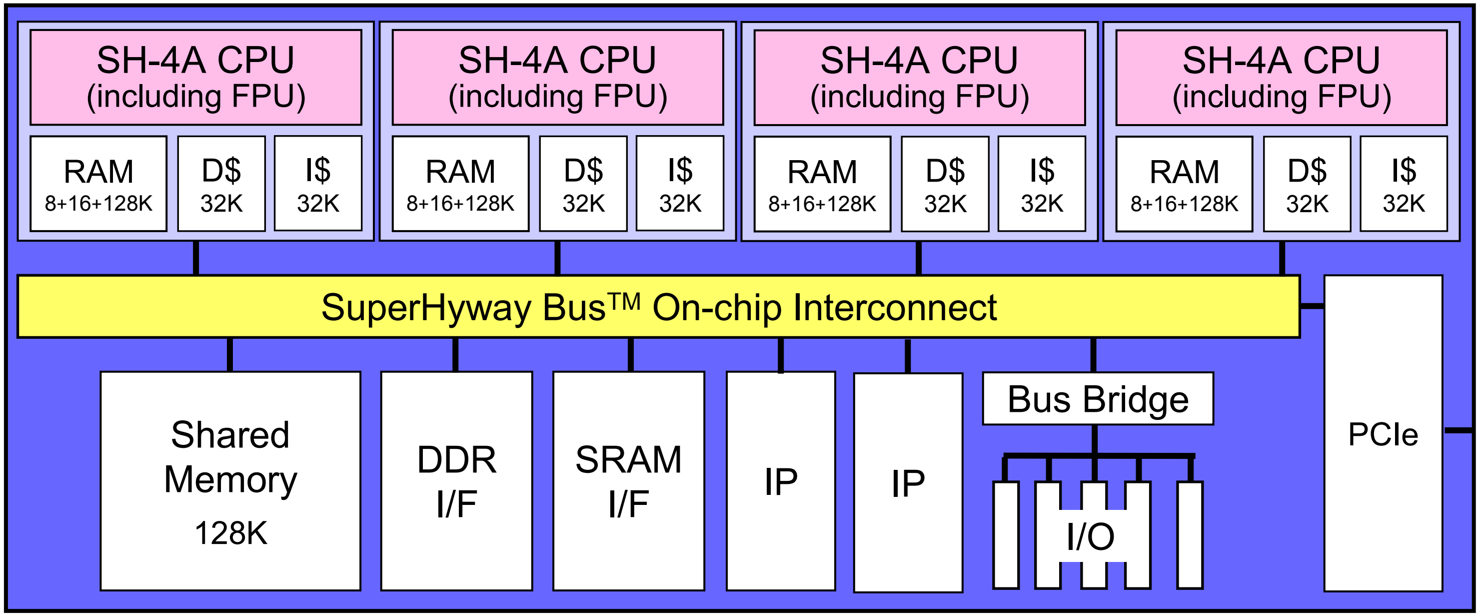

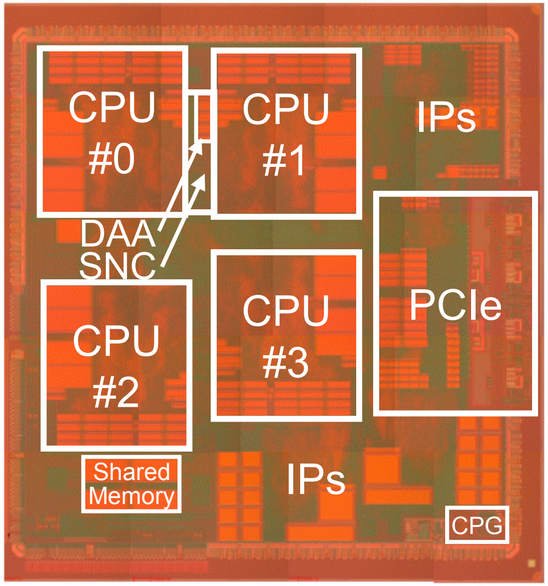

Process Technology	90-nm, 8-layer, Triple-Vth, Generic CMOS, 1.0V
Area	3.88mm² (Each CPU excluding all memories), 7.28mm² (Each CPU)
I/D Cache	32KB/4way set-associative (Each)
Local Memory	I-LRAM 8KB, D-LRAM 16KB, URAM 128KB (Each CPU)
Performance	1.8 MIPS/MHz/CPU (Dhrystone 2.1) 4320 MIPS @ 600MHz (4-CPU Total)
Power Consumption	0.6 mW/MHz/CPU @ 600MHz

Application of Synonym-related Function to SMP OS

- OS Enhancement for SMP
  - Applied: Linux-2.6.16 Kernel (Not MP-ready for SuperH multi-core)
  - Measurement: When kernel detects a synonym page, it flushes all entries of the page.
- Experiment (On evaluation board)
  - Enhanced for hardware implementation of synonym-related function
  - Executed shell command “find” for each CPU in parallel (Whole is stored in DDR)
    - Not Enhanced (Using original synonym measurement)

CPU	Execution Time (sec)
CPU0	31.88
CPU1	30.34
CPU2	29.31
CPU3	30.46

    - Enhanced

CPU	Execution Time (sec)
CPU0	14.39
CPU1	14.36
CPU2	14.10
CPU3	14.90

53%

Performance Improvement

Summary

- SH-X3: SuperH multi-core for high-performance and low-power systems
  - Efficient for both SMP and AMP
- Specification features for cache coherency
  - Separation of system bus and snoop bus
  - Centralized cache coherency by snoop controller
  - Support of mixed cache coherency protocol
- Specification features for SMP OS development
  - Broadcast of operand cache operating instructions (OCBI, OCBP, OCBWB)
  - Hardware implementation for synonym problem (53% performance improved)

Acknowledgement

This work was supported by NEDO (New Energy and Industrial Technology Development Organization) P03022, a joint project of Renesas Technology Corp., Hitachi Ltd., and Waseda University.

Renesas Technology

Backup Slides

Synonym Detection and Eviction (In the case of 4KB/Page)

VPN[31:12] → PPN[31:12]

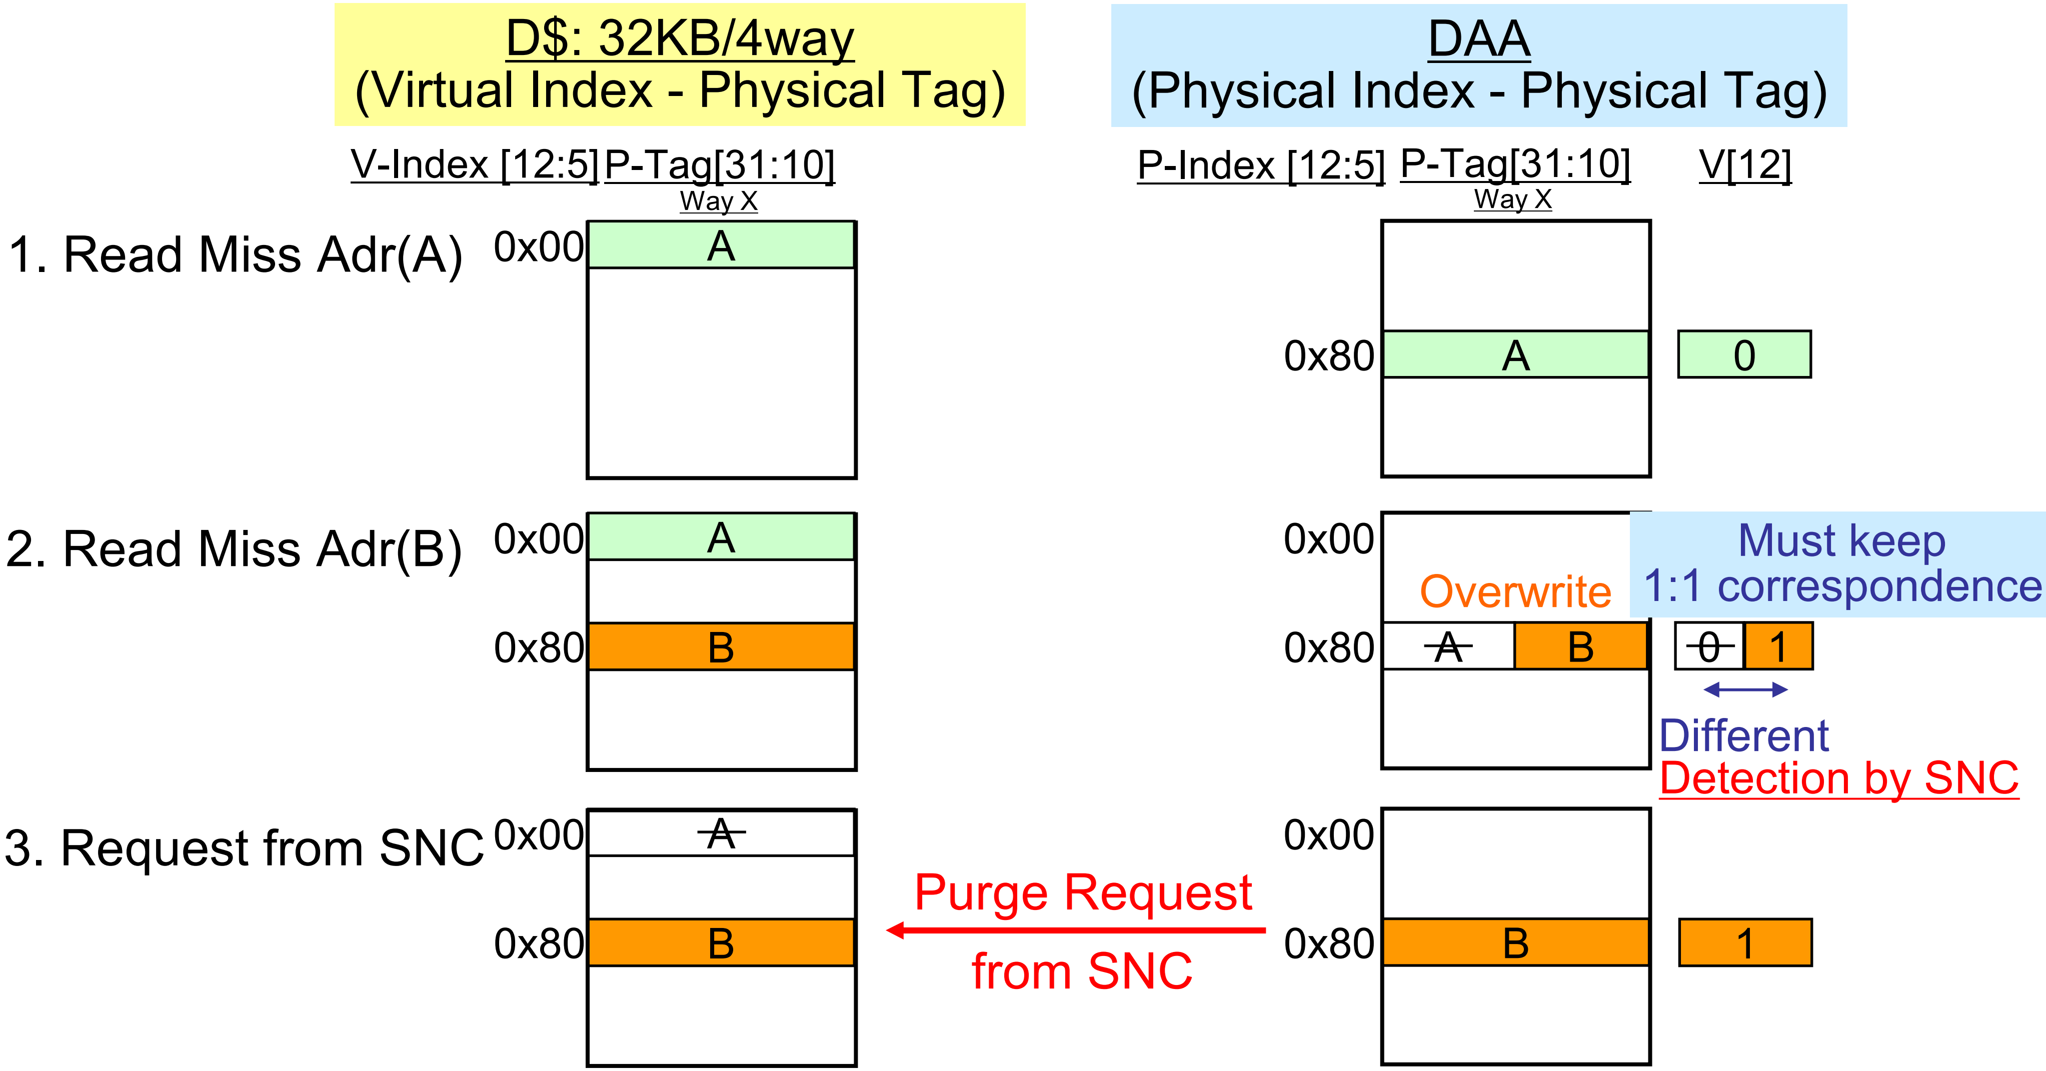

SNC: SNoop Controller

Synonym Detection and Eviction (In the case of 4KB/Page)

VPN[31:12] → PPN[31:12]

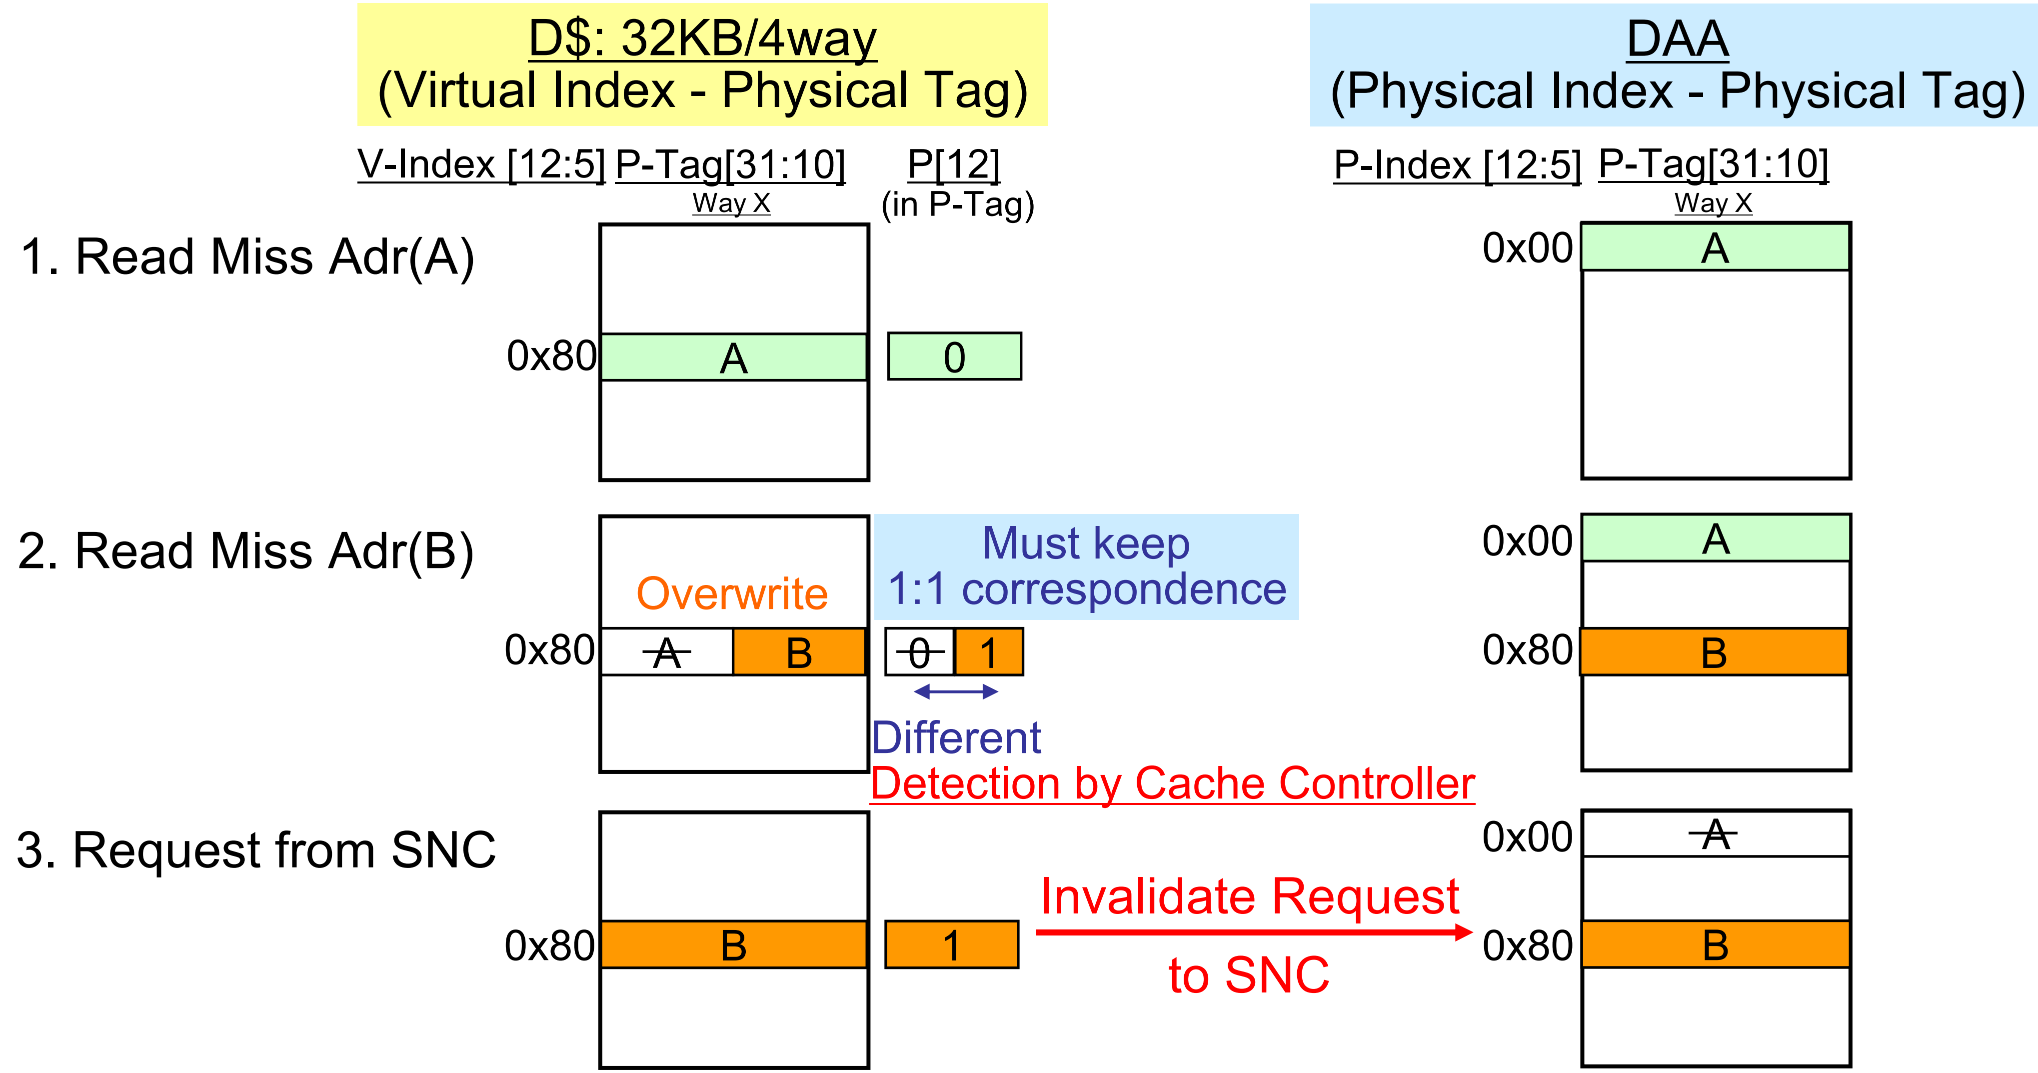

SNC: SNoop Controller

Application of Synonym-related Function to SMP OS

- OS Enhancement for SMP
  - Applied: Linux-2.6.16 Kernel (Not MP-ready for SuperH multi-core)
  - Measurement: When kernel detects a synonym page, it flushes all entries of the page.
- Experiment (On evaluation board) Freq. = (# of HW activation) / (# of request to SNC)
  - Enhanced for hardware implementation of synonym-related function
  - Executed shell command “find” for each CPU in parallel (Whole is stored in DDR)
    - Not Enhanced (Using original synonym measurement)

CPU	Activation Frequency of Function	Execution Time (sec)
CPU0	0.06%	31.88
CPU1	0.05%	30.34
CPU2	0.05%	29.31
CPU3	0.05%	30.46

- Enhanced Activated by VPN[12] != PPN[12] page mapping

CPU	Activation Frequency of Function	Execution Time (sec)
CPU0	0.16%	14.39
CPU1	0.13%	14.36
CPU2	0.17%	14.10
CPU3	0.13%	14.90

53% Performance Improvement