A RISC Processor for SR8000: Accelerating Large Scale Scientific Computing with SMP

Toshihiko Kurihara, Eiki Kamada,
Kentaro Shimada, Teruhisa Shimizu

Hitachi, Ltd.

HITACHI

Presentation Outline

- SR8000 Block Diagram
- Design Goals
- Pseudo Vector Processing
- Processor Core
- Chip Profile
- Summary

Block Diagram of SR8000

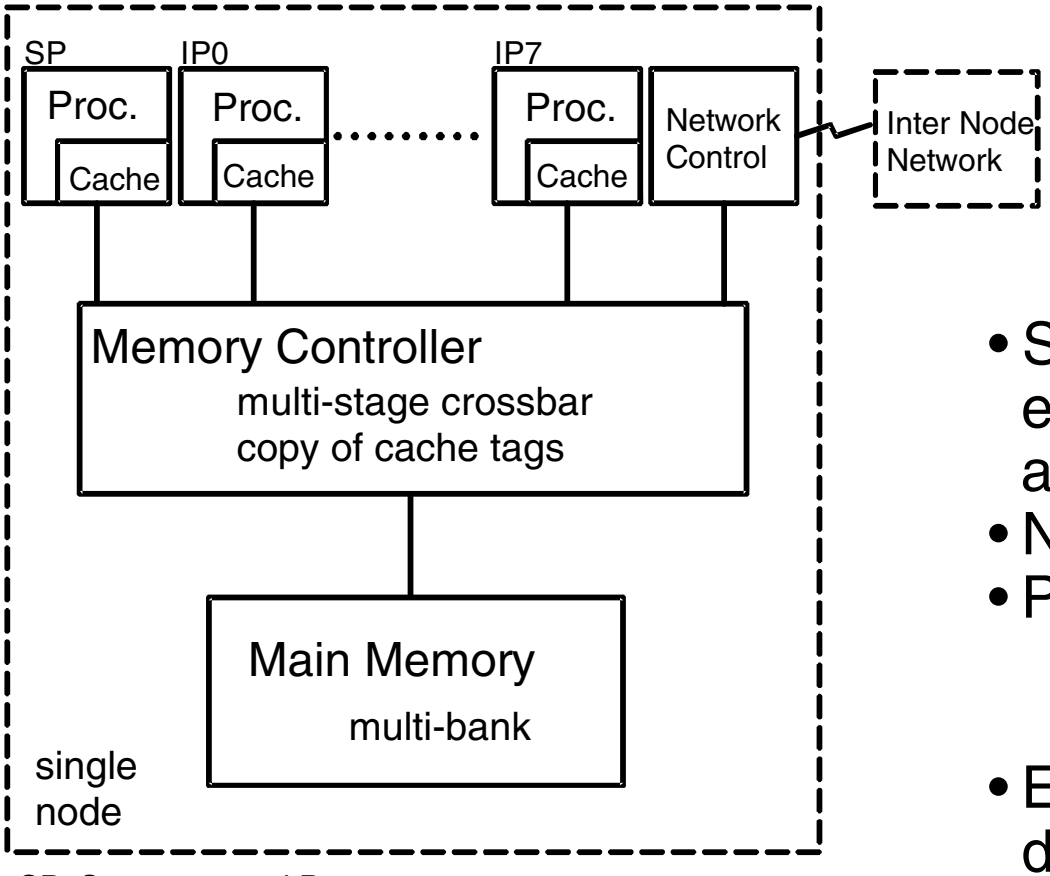

SP: System control Processor
IP: Instruction Processor

- System to efficiently execute large scale scientific applications.
- Number of nodes: Max 128
- Parallel execution of DO-loops with SMP.

- Efficient processing of data in main memory
- Pseudo Vector Processing

Design Goals

- High performance
  - High operating frequency
  - High memory throughput
- Enhancement of Pseudo Vector Processing (PVP) feature
- High reliability

Pseudo Vector Processing

- Vector Machine

example:
DO I = 1,N
A(I) = A(I) + B(I)

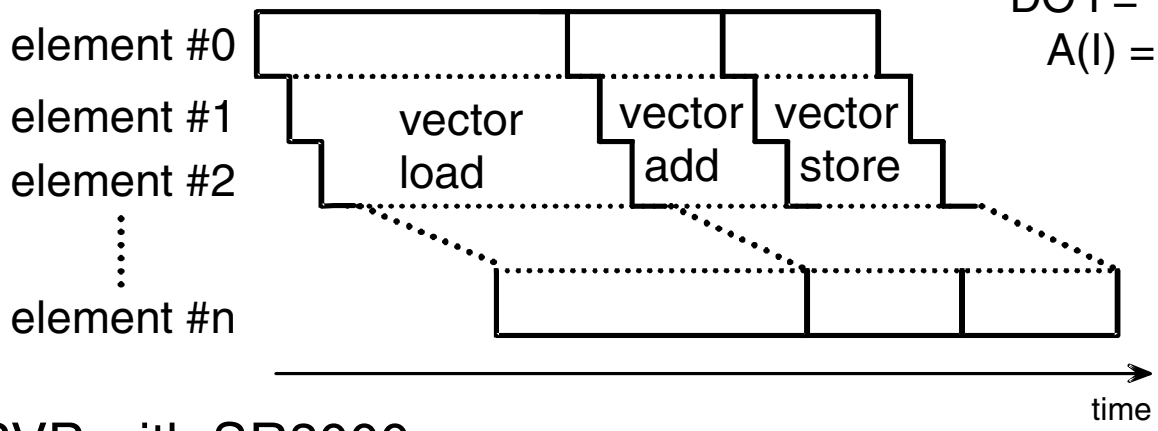

- PVP with SR8000

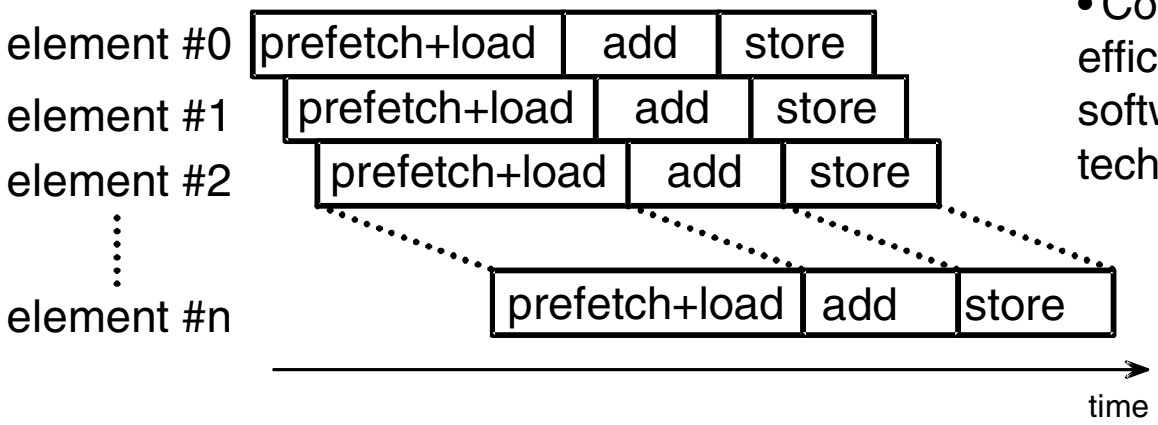

- Compiler generates efficient code using software pipeline technique

Pseudo Vector Processing(cont.)

- Architectural support
  - Prefetch instruction
  - Preload instruction
  - Slide-windowed floating-point registers
  - 128 logical floating-point registers
  - TLB supporting large pages
- Memory system support
  - Pipelined processing of memory requests
    Prefetch: 16, Preload: 128
  - High speed memory interface: 4GB/sec

Processor Core

- Architecture: 64bit PowerPC(*) + Hitachi Extension
- Out-of-order, superscalar execution
- 4 floating operation / cycle
- Instruction cache: 64KB, 2-way set associative, 128B / line
- Data cache: 128KB, 4-way set associative ,128B / line
- Quad-word (16B) load
- Fault detection logic
- Pseudo Vector Processing features

(*) PowerPC is a trade mark of International Business Machines Corporation.

Block Diagram of Processor

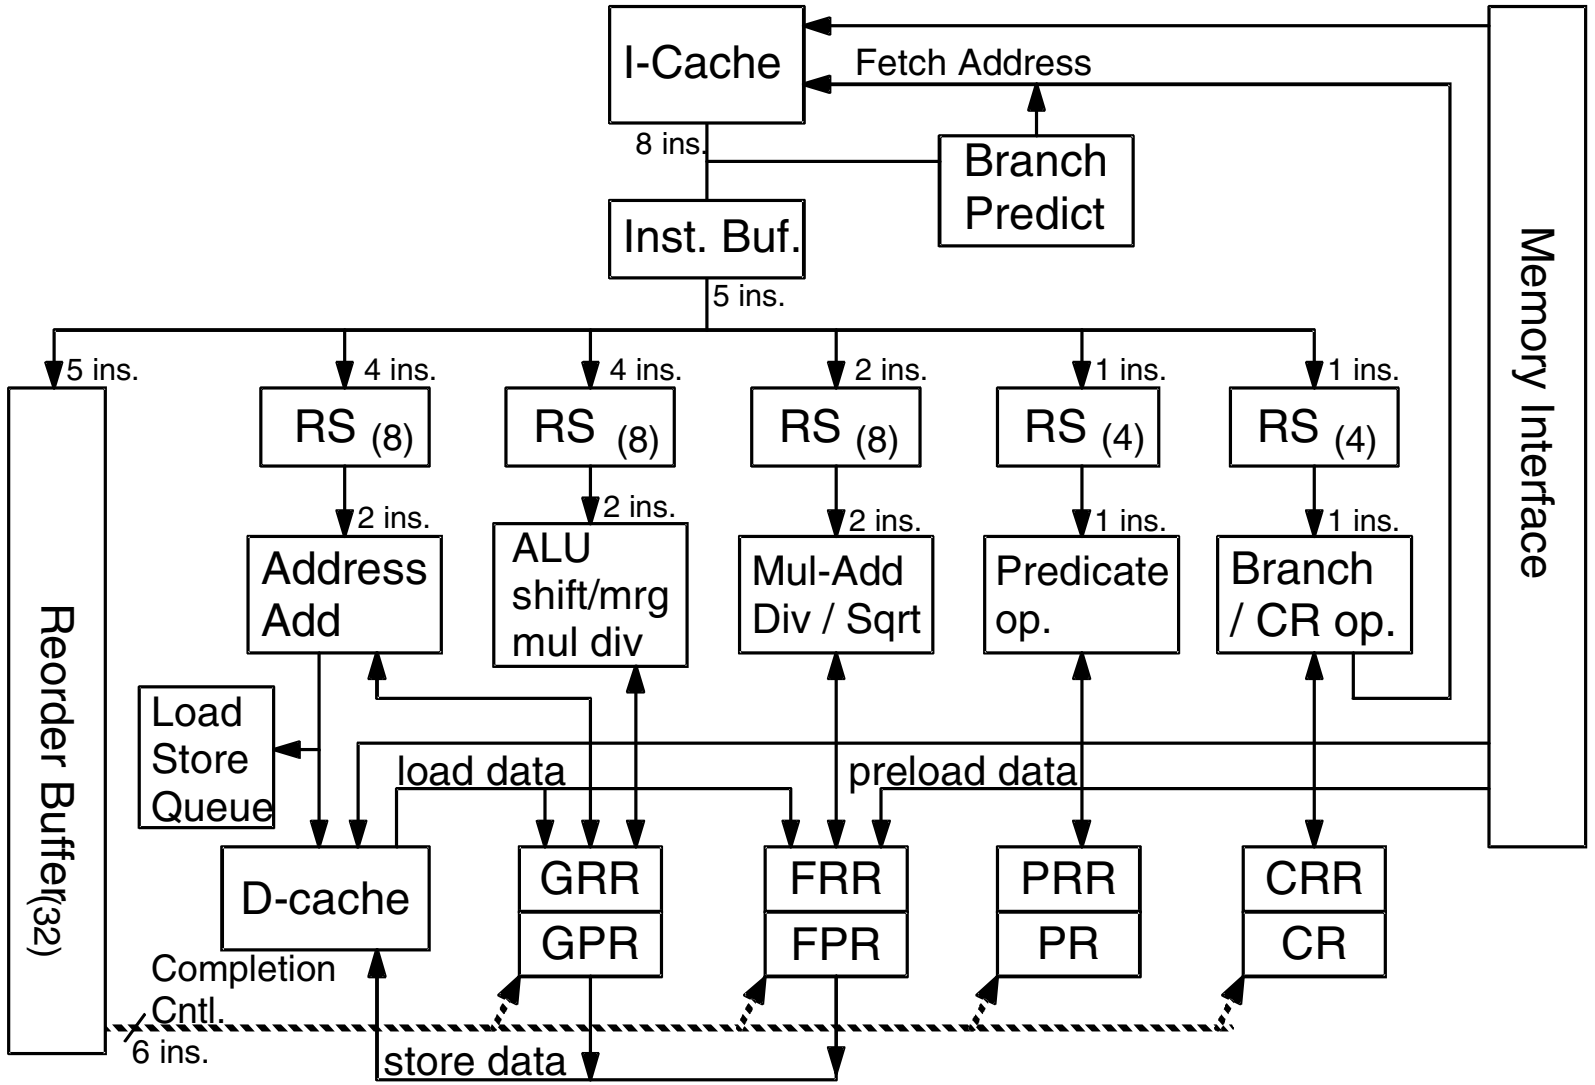

HITACHI

Summary of Superscalar Execution

- Number of Functional Units
  - Integer : 2 ALUs, 2 shifter / mergers,
    1 multiplier, 1 divider
  - Floating : 2 multiply-add units (4 floating operations / cycle) ,
    1 div / sqrt unit
  - 2 Load / Store units
  - 1 Branch / Condition Register operation unit
  - 1 Predicate unit
- Superscalar Execution
  - Fetch : 8 instructions / cycle
  - Dispatch : 5 instructions / cycle (5-th instruction must be a branch)
  - Execution : 8 instructions / cycle
  - Completion : 6 instructions / cycle

Pipeline Stages

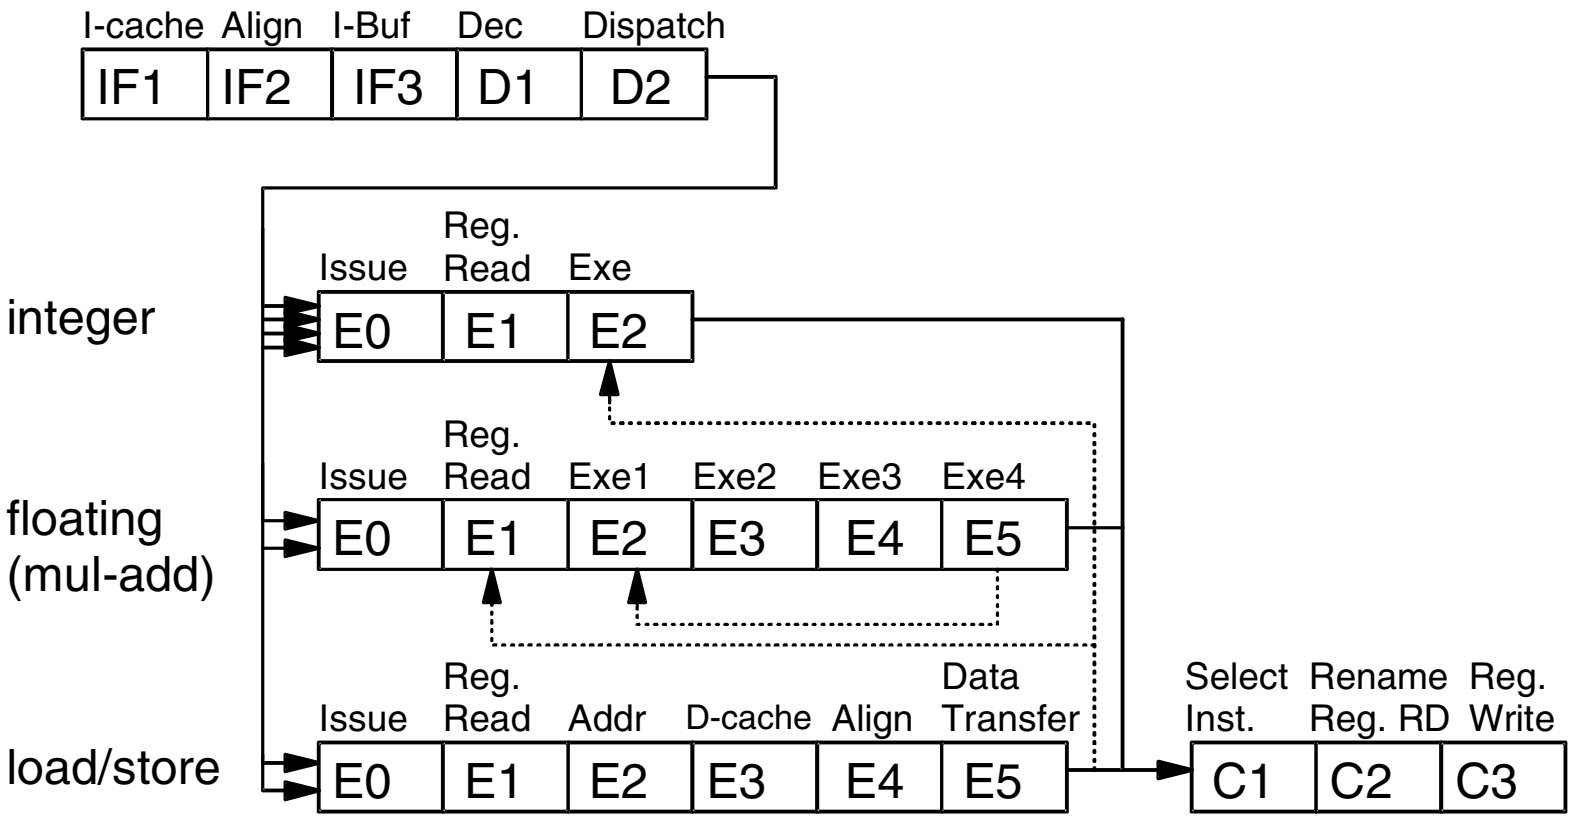

Instruction Fetch

- 64KB, 2-way set associative Instruction Cache
- Line Size : 128B
- 16B x 2 bank
- Up to 8 instructions are fetched per cycle.
- Branch prediction
  - BTAB : 0 cycle
  - BTAC : 1 cycle
  - BHT : 2 cycle

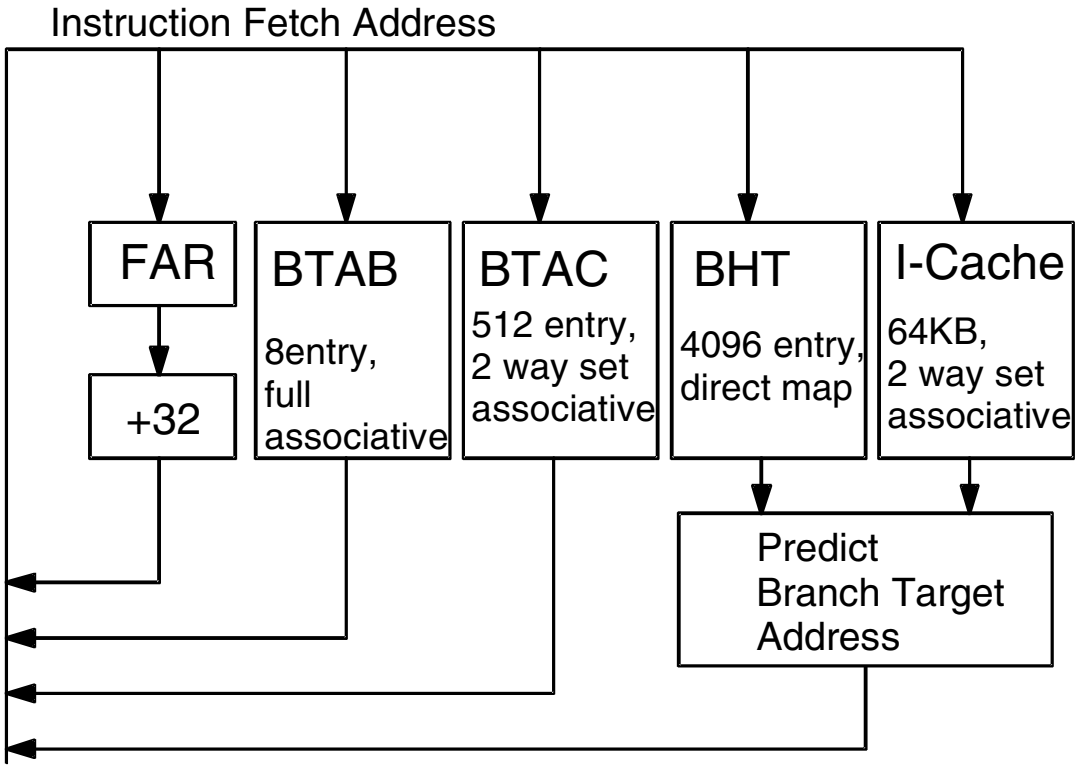

FAR: Fetch Address Register

BTAB: Branch Address Target Buffer

BTAC: Branch Address Target Cache

BHT: Branch History Table

Floating-Point Register

- Basic instructions (load/store, preload, fadd, ...) can directly specify 128FPRs.
- Slide-Windowed Floating-Point Register
  - Number of global registers: 4 / 8 / 16 / 32
  - Minimum slide pitch : 2
  - Instructions to control register window
- Large register file
  - 1.8KB

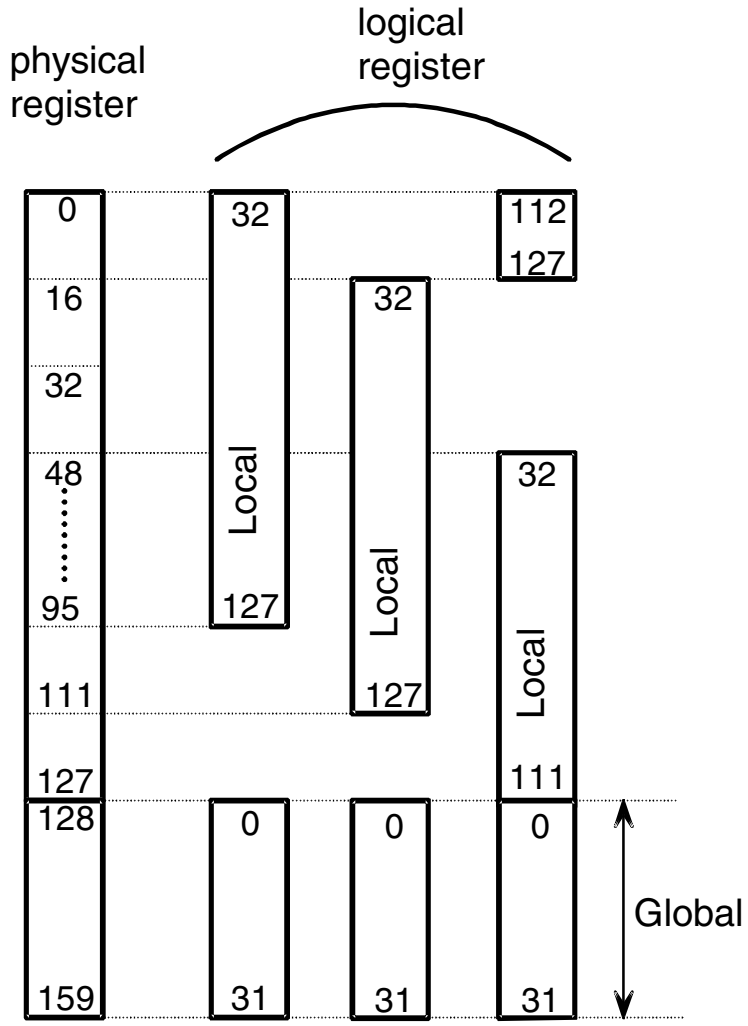

Load and Store

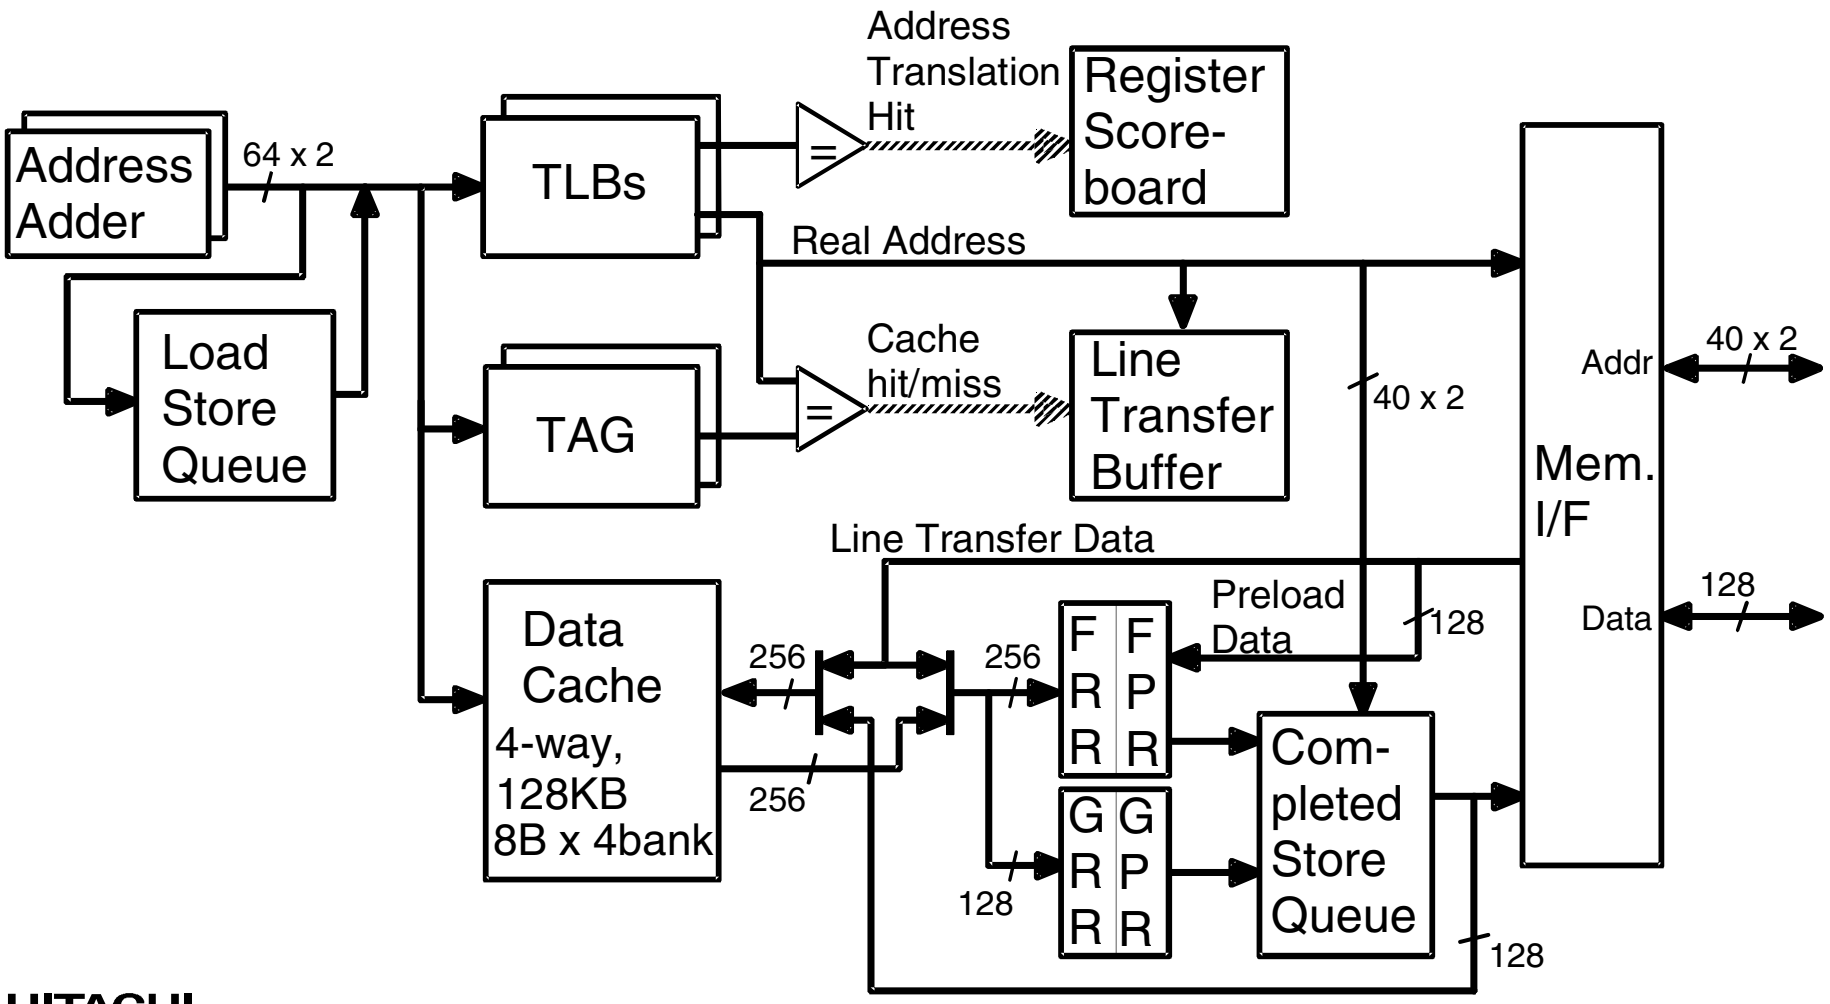

HITACHI

Prefetch and Preload

- Prefetch and preload releases resources (reorder buffers, rename registers,....) as soon as address translation completes.
  - Does not occupy resources for a long time.
- Prefetch
  - Transfers one cache line (128B) from main memory to cache.
  - Data arrival is checked by line transfer buffer.
  - Up to 16 transfers are processed in pipeline manner.
  - Line transfer buffer also handles cache miss of load instruction.
- Preload
  - Transfers 8B from main memory to a floating-point register.
  - Bypasses data cache.
  - Directly specifies 128 FPR.
  - Target register is not renamed.
  - Up to 128 transfers are processed in pipeline manner.
  - Data arrival is checked by scoreboard.

Address Translation Resources

- Effective address: 64 bits; Real address: 40 bits
- Instruction TLB
  2-way set associative, 512 entries, 4KB / page
- Data TLB
  - 2-way set associative, 512 entries, 4KB / page
  - Covers only 2MB.
- TLB supporting large pages (LTLB)
  - Covers entire physical memory space.
  - Data reference only
  - Size of large page is 16MB-128MB (All entries have same mapping size at the same time.).
  - Direct map, 256 entries
- TLB miss handled by hardware

Memory Interface

- Pipelined processing of fetch requests
  - Prefetch or cache-miss load requests : 16
  - Preload : 128
  - Instruction fetch : 2
- 2 preload requests are issued per cycle.
- 2 store requests are issued per cycle.
- High speed interface --- 4GB/sec
- Invalidate request from memory controller
  - Data cache
  - Instruction cache
  - TLB

Reliability

- On chip memory
  - 100 % Single-bit error detection by parity
  - Erroneous data are not used for execution.
  - Detected error is reported as interrupt.
  - Interrupt handler clears memory.
- Data path
  - Variety of fault detection methods used
    ALU : parity prediction
    Floating-point unit : residue check
    Other important data paths (address, store data, ...): parity
  - Retry instruction if architecture resources maintain correct data

Chip Profile

- Process: 0.25-μm, CMOS
- Number of metal layers: 7
- Power supply: 1.8V
- Frequency: 400MHz(*)
- Die size: 18.5 mm X 18.5 mm
- Transistor: 23 million
- Number of signal pins : 560
- Static circuit except memories and register files

(*)Typical silicon processing and typical environment

Photograph of Chip

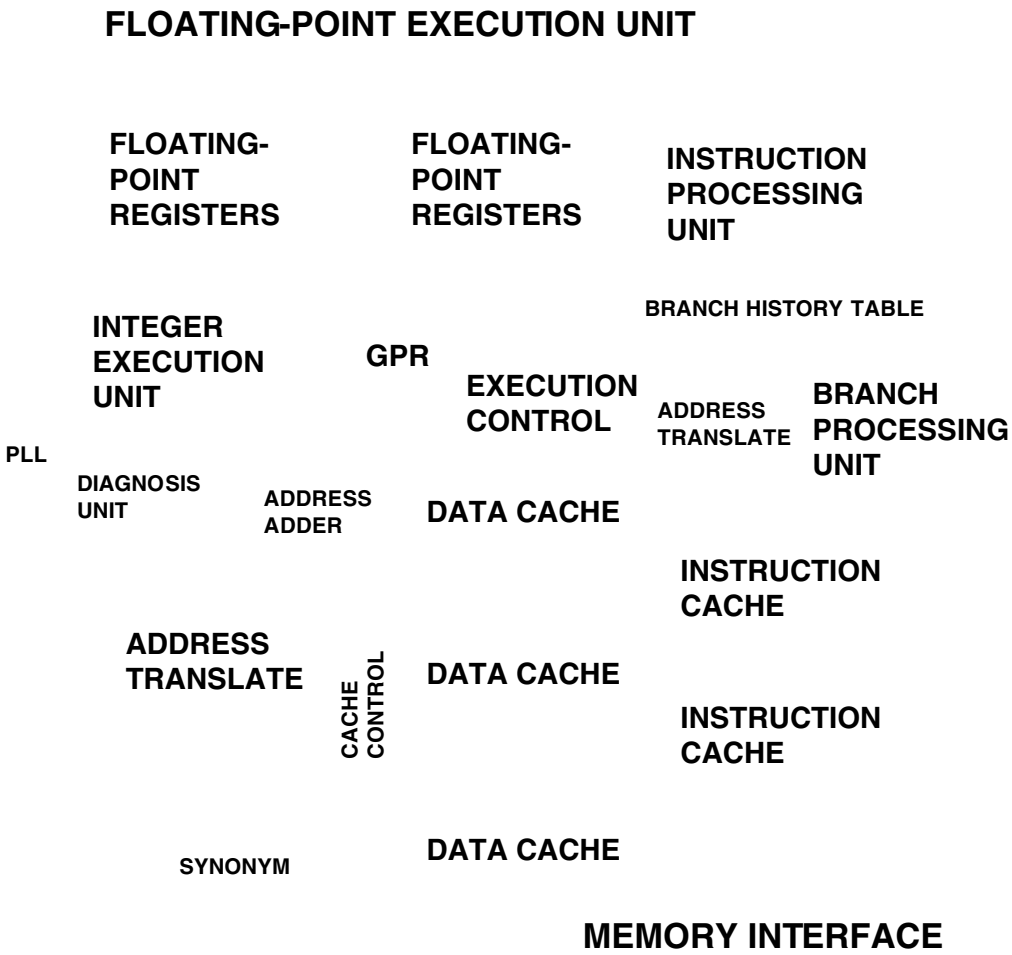

HITACHI

Summary

- High performance RISC processor for SR8000
- High memory throughput
- Enhanced Pseudo Vector Processing feature
- Aggressive superscalar execution
- High reliability
  - Fault detection of memory and data path