

Ridge Processor Reference Manual

RIDGE



9008 B

Ridge Processor Reference Manual

Ridge Computers
Santa Clara, California

Third Edition: 9008-B (JUN 84)

Copyright 1983, 1984, Ridge Computers.
All rights reserved.
Printed in the U.S.A.

PUBLICATION HISTORY

Manual Title: Ridge Processor Reference Manual

First Edition: 9008 (MAR 83)

Second Edition: 9008-A (FEB 84)

Third Edition: 9008-B (JUN 84)

NOTICE

No part of this document may be translated, reproduced, or copied in any form or by any means without the written permission of Ridge Computers.

The information contained in this document is subject to change without notice. Ridge Computers shall not be liable for errors contained herein, or for incidental or consequential damages in connection with the use of this material.

TABLE OF CONTENTS

Title -----	Page -----
CHAPTER 1: OVERVIEW	1
Key Features	1
Related Documents	1
Processor Overview	2
Instruction Formats	3
Processor Architecture	4
Internal Structure	6
Instruction Fetch Unit	9
Execution Unit	10
Memory Controller	12
Data Types	13
Integers	14
Real Numbers (Single Precision)	14
Real Numbers (Double Precision)	15
Syntax Conventions	16
CHAPTER 2: MEMORY REFERENCE INSTRUCTIONS	17
Instruction Formats	17
Instruction Descriptions	18
Load Instructions	19
Store Instructions	19
Load Address Instructions	20
CHAPTER 3: REGISTER FORMAT INSTRUCTIONS	21
Instruction Format	21
Instruction Descriptions	21
Integer Arithmetic Instructions	22
Logical Operator Instructions	23
Integer and Logical Immediate Instructions	24
Extended Precision Integer Instructions	25
Real Instructions	27
Double Real Instructions	28

Ridge Processor

Bit-Oriented Instructions	29
Test Instruction	30
Compare Instructions	31
Shift Instructions	32
Sign Extend Instructions	33
CHAPTER 4: PROGRAM CONTROL INSTRUCTIONS	35
Branch Instructions	35
Instruction Format	35
Instruction Descriptions	36
Branch Instructions	36
Loop Control Instruction	36
Subroutine Call and Return Instructions	37
Call Subroutine Instruction	37
Call Subroutine Register and Return Instructions	38
APPENDIX A: RIDGE OPCODE MAP	39

ILLUSTRATIONS

Figure		Page
-----		-----
1	Instruction Formats	3
2	Model of Processor Architecture	4
3	Processor Instruction Pipeline	7
4	Ridge 32 Internal Structure	8
5	Data Formats for Operand Types	13

CHAPTER 1: OVERVIEW

The Ridge 32 is an engineering workstation with a 32-bit, high performance processor implemented in MSI and LSI bipolar logic. This proprietary processor has a simple, general purpose, microcoded architecture that incorporates paged virtual memory. The Ridge 32 provides processing power equal to medium performance mainframes and high performance minicomputer systems. This manual describes the overall operation of the processor including its features, a block level description, and the instruction set.

KEY FEATURES

- reduced instruction set computer (RISC) architecture
- 125-nanosecond cycle time
- one-clock cycle minimum instruction time
- 4096-byte paged virtual memory
- four-gigabytes linear address space
- separated code and data
- branch prediction logic
- single and double real floating point instructions
- 16 general registers
- 375-nanosecond memory cycle time

RELATED DOCUMENTS

- | | |
|---|--|
| Ridge Assembler
Reference Manual
#9005 | - Gives instruction syntax and pseudo operations for the assembler program. Instructions are listed alphabetically. Descriptions include instruction exceptions. |
| Ridge Hardware
Reference Manual
#9007 | - Describes card cage, individual boards, cables, and operating specifications for I/O boards. |
| Ridge Architectural
Specification
(not yet available) | - Details of privileged instructions, traps, exceptions, interrupts, clock, timer, internal tables, internal registers, and virtual translation algorithm. |

PROCESSOR OVERVIEW

The Ridge 32 processor is a register-oriented, 16 general register computer. The processor provides virtual addressing using 4096-byte pages within a four-gigabyte address space. The Ridge 32 utilizes a 125-nanosecond machine cycle and can complete simple instructions in one cycle. The maximum instruction rate is eight-million instructions per second (8 MIPS).

The processor's style of architecture has become known as a reduced instruction set computer (RISC). The goals of a RISC architecture are to simplify the functions of the machine which reduces the amount of hardware necessary to implement the processor. The reduction in logic allows a faster cycle time and permits instructions to complete in one machine cycle. As a result, the Ridge 32 is a very fast and low cost computer.

The characteristics of a RISC architecture are:

Simple addressing modes. The Ridge 32 uses only three modes which reduces the amount of logic needed to perform memory references.

Simple instruction formats. The Ridge 32 uses three instruction formats that can be decoded with a minimum of logic.

Separated code and data. The Ridge 32 uses separated code and data eliminating the need for logic that detects and resolves self-modifying code.

Designed to support high level languages. The instructions provided are designed to match the code generation capabilities of such languages as FORTRAN, C, and Pascal. These languages tend to generate short sequences of instructions, using only a few instruction types to perform the required functions. Complex instructions and instructions not used by a compiler are eliminated. Thus, the Ridge 32 instruction set offers the "primitives" which will be assembled by a compiler.

Regularity. Data types and addressing modes are examples of regularity. For memory reference instructions there are four operand sizes and three addressing modes. Each of the addressing modes is available for all operands. To do otherwise complicates the compiler and may slow the overall operation of the machine.

Linear address space. Code and data space are each linear with a byte-addressable area that is four-gigabytes long. Segmentation schemes appear to save logic to support the full 32-bit address widths, but instead they complicate the hardware and compilers, and slow the processor's performance.

General registers. All registers are available for use as data, indexing, and addressing. If registers are specialized they complicate compilers, reduce the available fast storage area, and increase code size when data must be moved to the appropriate register type.

Instruction Formats

The processor contains 16 32-bit registers. The instruction set is of the two-operand form, and uses three instruction formats. The instruction formats are register-to-register (16-bits long), short displacement memory address (32-bits long), and long displacement memory address (48-bits long). The instruction formats are shown in Figure 1.

Instruction Format:

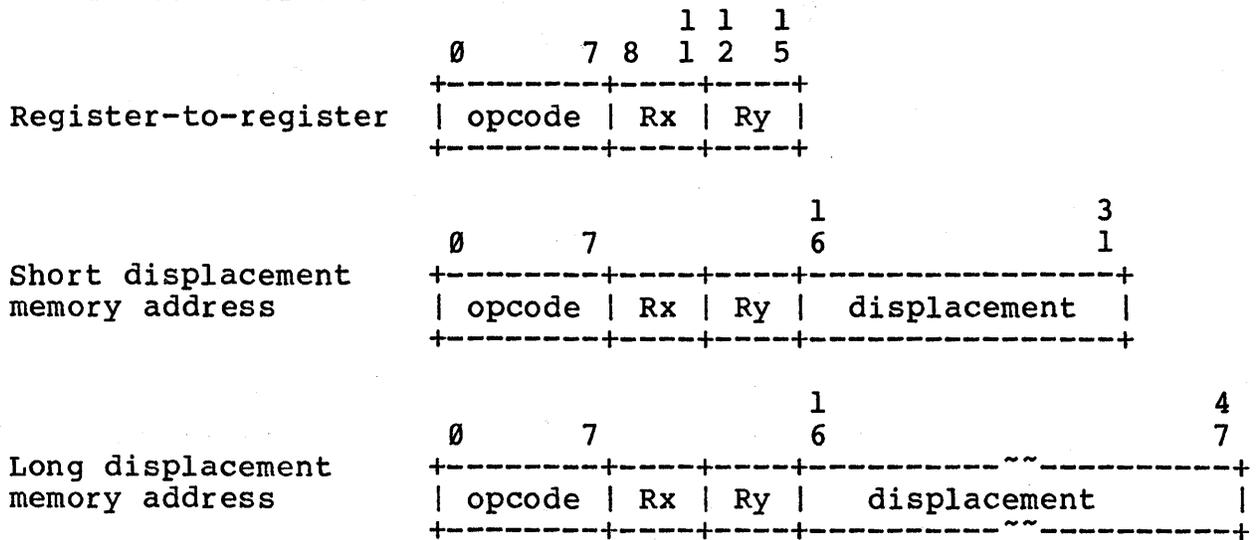


Figure 1. Instruction Formats

All instructions use an eight-bit opcode followed by two four-bit operands. The first operand always names a register or a register pair. The second operand names a register or is a four-bit constant. Instructions exist to operate on registers, load from memory, store to memory, and transfer program control.

The register-to-register format is used for instructions that operate on the contents of one or two registers and do not address memory. The short and long displacement memory address format instructions are used for memory-addressing instructions, such as storing and loading. The short displacement memory address format is used for referencing addresses that can be specified in 16 bits. The long displacement memory address format is used for referencing addresses that must be specified in 32 bits.

Any arithmetic or address operation can be performed on any register. Registers are not specialized for counting or indexing.

Processor Architecture

The model of the processor architecture is shown in Figure 2. The user-visible features of the processor are instructions, general registers and the program counter. Instructions operate on the general registers (register-to-register) or on a register and a memory location (load from memory or store to memory). The program counter is visible when using program control instructions such as subroutine call and branch.

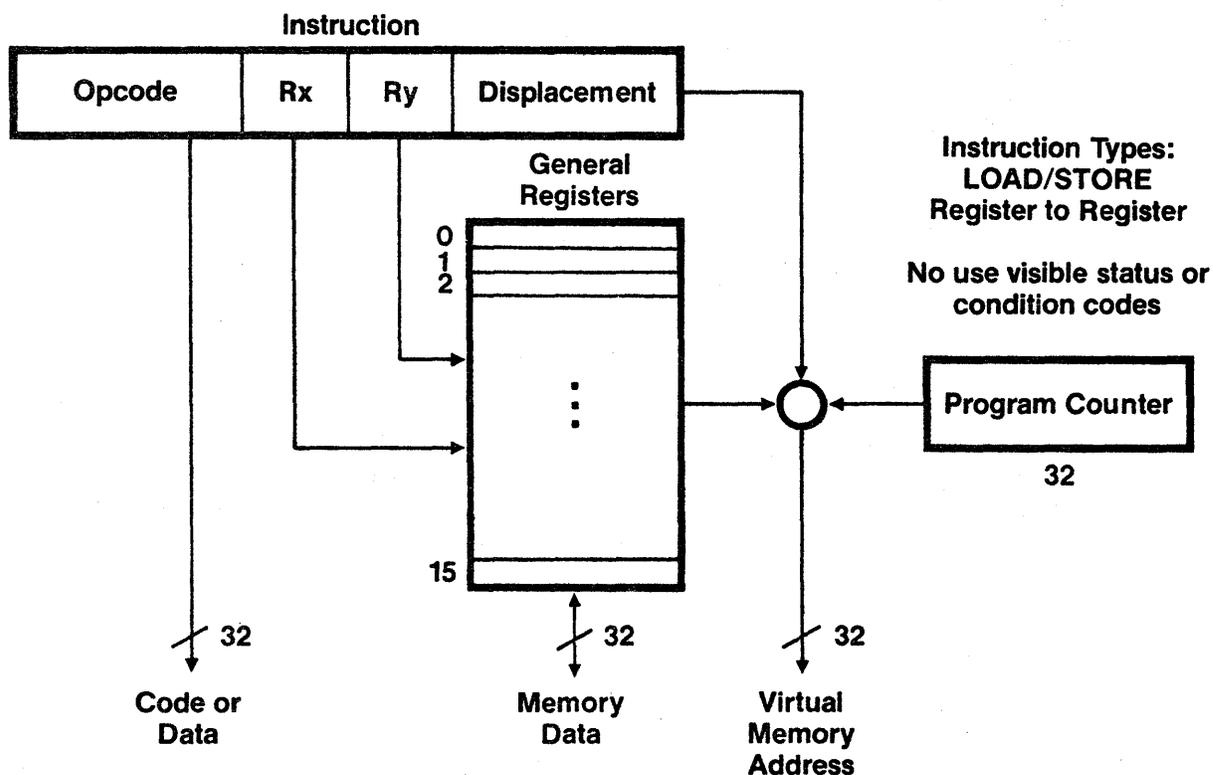


Figure 2 Model of Processor Architecture

Ridge Processor

Memory addresses can be one of the following forms: displacement field from the instruction; index register plus displacement; or program counter (PC) may be added to either of the other forms for program-relative locations. These memory address forms are shown in Figure 2.

All addresses generated by the processor are virtual addresses and are 32-bits wide. Memory reference instructions indicate code or data space by utilizing a bit in the instruction opcode. An individual program may access a maximum of four gigabytes of code space and a maximum of four gigabytes of data space.

A status register containing condition codes is purposely missing from this architecture. Status registers complicate and tend to slow down high speed processors. On high speed machines several instructions are in various stages of execution at any given moment. Condition codes tend to be generated at various times during these stages and must be properly propagated from stage to stage. In virtual machines, an additional problem occurs in preserving the condition codes throughout the stages when an instruction aborts due to a page fault.

The processor architecture includes the conditional branch instruction, that obviates the need for condition codes. This instruction combines the compare function and the conditional branch instruction. The compare function generates the condition code and the conditional branch instruction changes program flow of control based upon condition code values.

INTERNAL STRUCTURE

The Ridge 32 processor consists of two printed circuit boards. The first is the instruction fetch unit and the second is the execution unit. The processor has a private bus to the memory controller with separate 32-bit address and data lines. The instruction fetch unit and execution unit may each independently access main memory. Memory cycle time is 375 nanoseconds, which includes virtual-to-real memory translation and error correction.

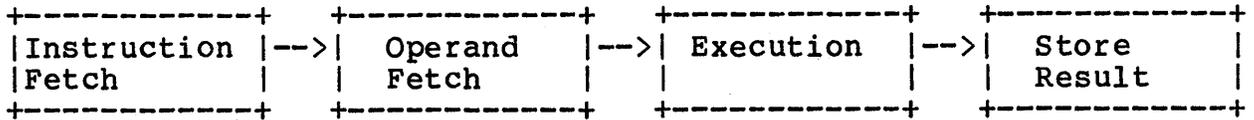
A block diagram of the processor, memory, and I/O system are shown in Figure 4. In the following text, the items in **bold type** are illustrated in Figure 4.

Pipelined Organization:

The Ridge processor uses a pipelined organization. The pipeline is composed of four stages: **instruction fetch**, **operand fetch**, **execution** and **store result**. Figure 3 illustrates the processor instruction pipeline. Each pipeline stage performs its function in one processor cycle.

The purpose of the pipeline is to increase machine speed by using parallelism. Each stage of the pipe operates on a separate instruction. Instructions flow through each of the four stages of the pipe, one cycle at a time. Although complete execution of an instruction takes four machine cycles, one instruction completes each cycle, thus creating an effective processor speed that is four times the speed of a non-pipelined operation. The instruction pipeline includes all of the logic on the execution unit and part of the logic on the instruction fetch unit.

Processor Cycles



The operations performed during each processor cycle are as follows.

Instruction fetch. The instruction is fetched from the prefetch buffer. The opcode is used as an index into the control store, which controls instruction execution. The Rx and Ry operands in the instruction are used to enable the register select logic.

Operand fetch. Rx and Ry are fetched from the register files.

Execution. The ALU operates on Rx and Ry, the result passes through the barrel shifter and is stored into the result register.

Store result. The data is moved from the result register into the Rx and Ry register files.

Instruction Flow Through Pipeline Stages

Time

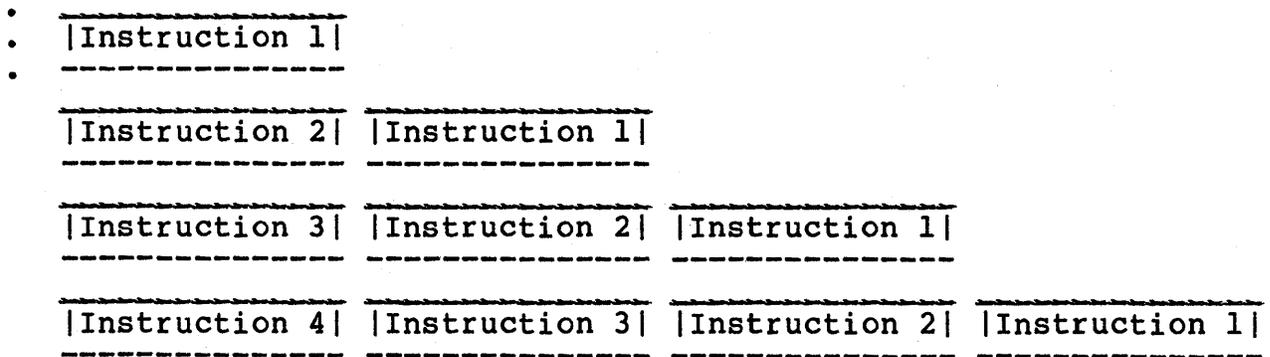


Figure 3. Processor Instruction Pipeline

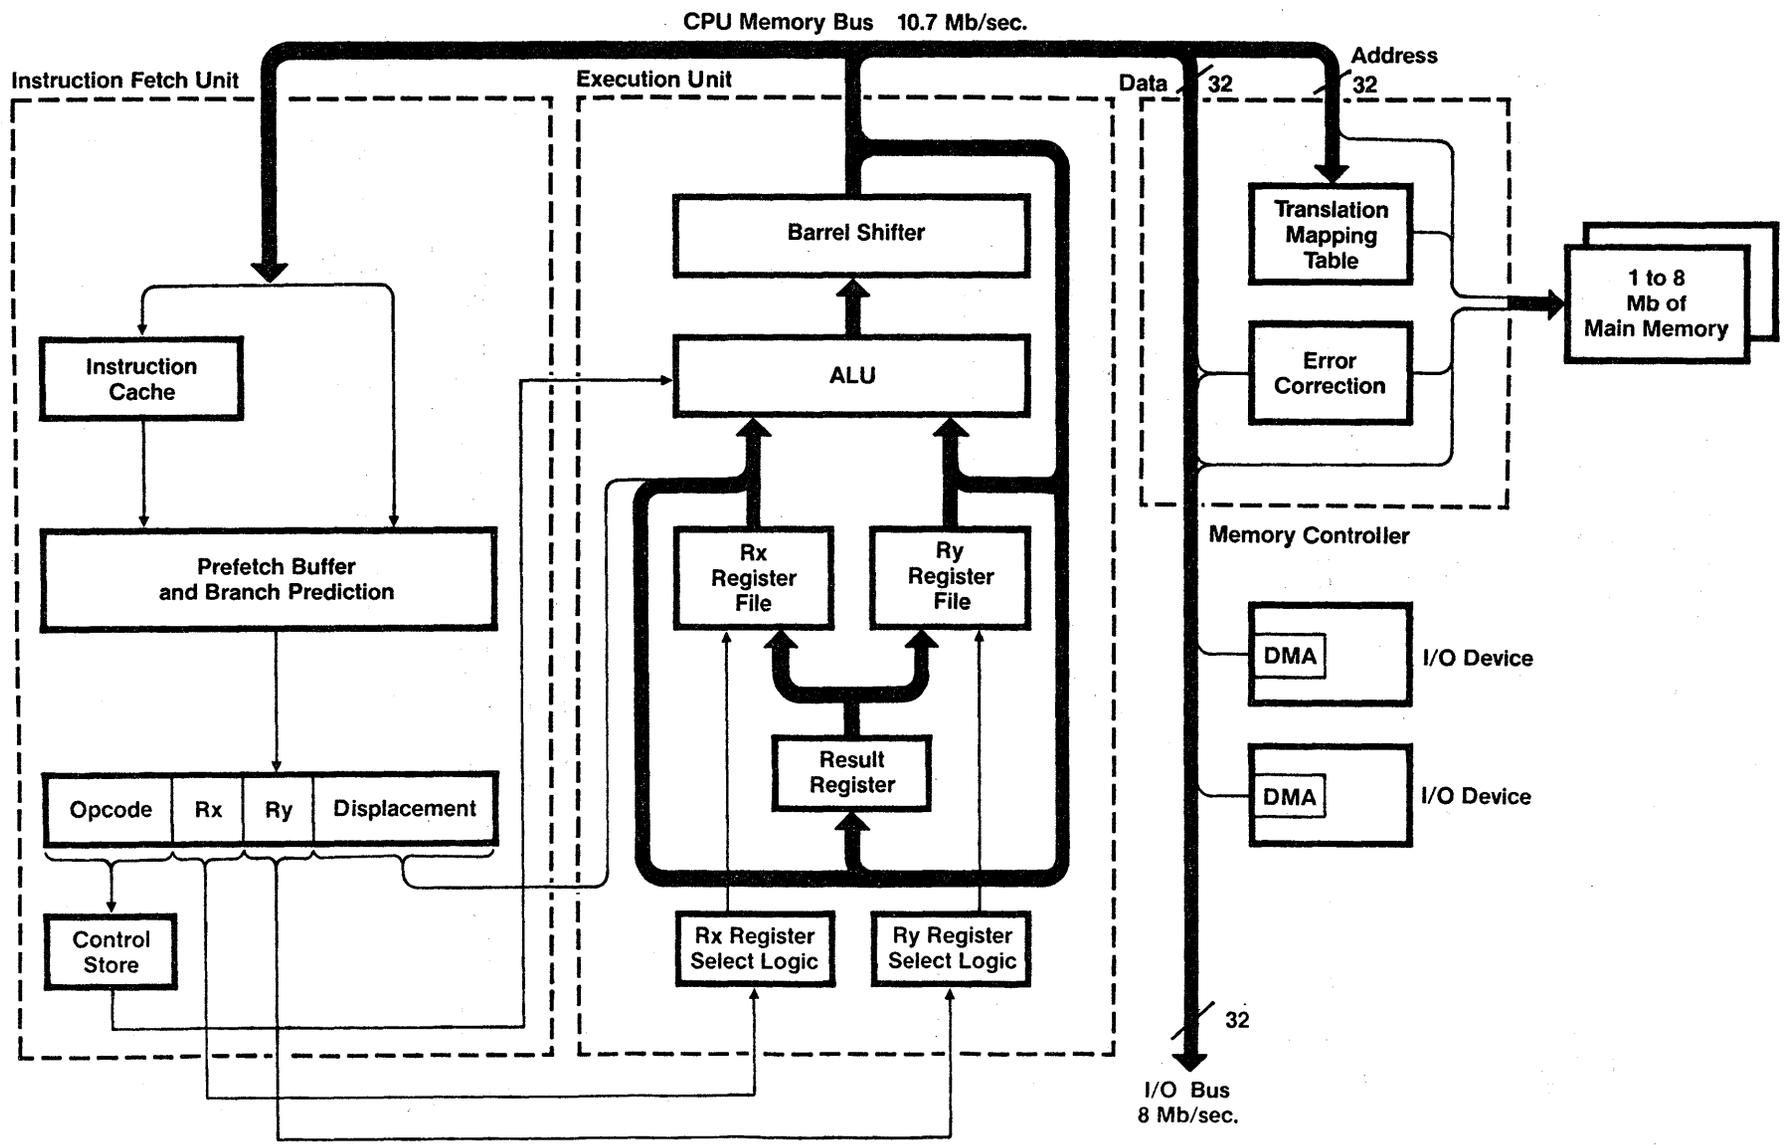


Figure 4. Ridge 32 Internal Structure

Instruction Fetch Unit

The instruction fetch unit performs instruction prefetch and decoding. It contains a 256-byte **instruction cache** and a maximum of 4096 words of 48-bit wide **control store**. The instruction fetch unit fetches instructions from the instruction cache or main memory ahead of the execution unit and stores them in its eight-byte **prefetch buffer**.

Branch Prediction:

The implementation of branch instructions is critical to the performance of pipelined machines. Without special handling, a conditional branch instruction empties the pipeline. This prevents the processor from prefetching the next instruction until the outcome of the branch has been determined.

For this reason, branches can be among the slowest instructions on high performance machines. The processor uses a technique to load the instruction into the pipe which is the most likely result of the branch, thus reducing the chance that the pipeline is loaded with instructions on the wrong path.

Conditional Branch Instructions:

Conditional branch instructions contain a static prediction bit in the instruction displacement field that can be set by a compiler. The **branch prediction** logic in the instruction fetch unit then fetches along the predicted path. This keeps the pipeline full and makes conditional branch instructions fast.

Branch Prediction Example:

For example, consider Pascal REPEAT ... UNTIL loops. The loop is constructed by the compiler as a linear section of code ended with a conditional branch. This branch is part of the UNTIL expression. Usually these loops are executed more than once, so the compiler marks the conditional branch at the bottom of the loop to be "predicted".

When the program is executed, the processor fetches and executes all the instructions in the linear portion of the loop. As the instruction fetch unit prefetches the conditional branch at the end of the loop, the prediction bit is detected. Instead of fetching the next sequential instruction as it normally would, the instruction fetch unit fetches the instruction at the top of the loop, which is the branch target. This prefetching the location of the branch target allows loops to execute at the same speed as linear sections of code.

Ridge Processor

As the loop is executed for its last time, the instruction fetch unit incorrectly fetches the instruction at the top of the loop. This time, however, the UNTIL condition has been reached, and the loop has ended. Now the instruction fetch unit must flush this instruction and fetch the next sequential instruction, which will then be executed.

This flushing of the instruction pipeline causes a four-cycle delay for the incorrectly predicted conditional branch instruction. Measurements have shown this to be infrequent, and consequently program speed is increased by the use of the branch prediction logic.

Unconditional Branch Instructions:

Unconditional branch instructions also make use of the branch prediction and prefetch logic in the instruction fetch unit. In unconditional branches, the instruction is decoded, the target location is fetched and placed in the instruction stream, and the unconditional branch is flushed from the prefetch buffer. This effectively removes the unconditional branches from the program entirely, and if the instruction fetch unit is ahead of the execution unit, unconditional branches can be performed with zero instruction time.

Execution Unit

The execution unit contains the general registers and is responsible for instruction execution. The arithmetic logic unit (ALU) and barrel shifter are also found on this board. The barrel shifter can shift from zero to 31 positions, left, right, or circularly, in one clock cycle.

The general registers are found in the Rx register file. A duplicate copy of the registers is contained in the Ry register file. Duplicating the registers allows both Rx and Ry to be accessed in a single clock cycle.

The general data flow through the execution unit is as follows. Data is fetched from the Rx and Ry register files, operated on by the ALU, temporarily stored in the result register and then stored into the register files. Should data not yet stored into the register files be needed in a computation, the register select logic may bypass the register file and use the data on the bus as input into the ALU.

Memory Controller

The memory controller provides virtual-to-real address translation and **error correction**, and also handles all memory data for the processor and I/O devices. All memory accesses from the processor are virtual and go through the **translation mapping table** where they are converted to real addresses and presented to **main memory**. I/O devices on the **I/O bus** use real addresses and bypass the translation mapping table.

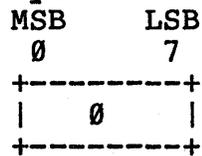
Main memory cycle time is 375 nanoseconds, and the memory controller processes four bytes per cycle. The **CPU memory bus** runs at full memory speed giving this bus a bandwidth of 10.7 megabytes per second. The I/O bus uses multiplexed **address** and **data** lines to minimize the use of connector pins on I/O boards. The I/O bus cycles in 500 nanoseconds and provides eight megabytes per second of direct memory access (**DMA**) bandwidth for I/O devices. Each board on the I/O bus contains its own DMA logic.

The memory controller can access from one to eight megabytes of main memory. All memory accesses are single-bit error corrected and double-bit error detected.

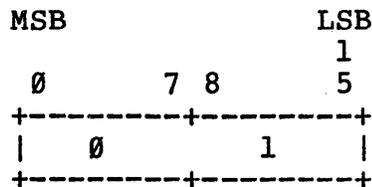
DATA TYPES

The processor has instructions to load and store four different sizes of operands. The basic addressable unit is the eight-bit byte. The other operand sizes are the halfword (16-bits), the word (32-bits) and the double word (64-bits). Data types are addressed from least significant bit (LSB) to most significant bit (MSB). Bytes are numbered from most significant byte to least significant byte. Figure 5 gives the notation and memory layout for each type of operand.

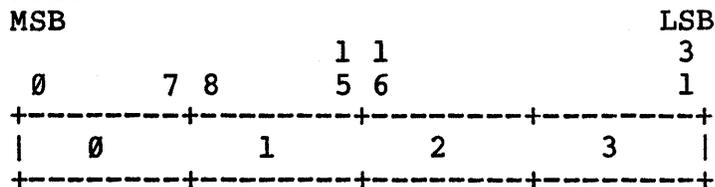
Byte



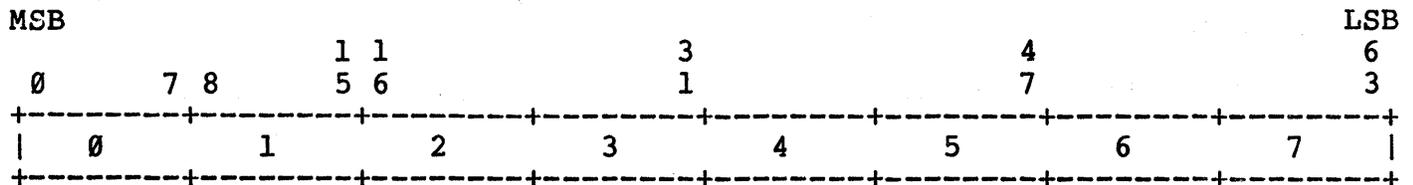
Half-word



Word



Double Word



MSB = most significant bit
LSB = least significant bit

Figure 5. Data Formats for Operand Types

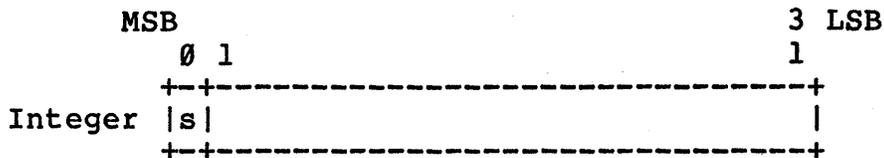
Ridge Processor

All data is manipulated in the processor's 16 32-bit general registers. There are instructions that manipulate these registers as 32-bit and 64-bit data types. There are three 32-bit data types: two's complement signed integers, unsigned integers, and real numbers. There is a single 64-bit data type which is double precision real numbers. Integer data types longer than 32 bits may be manipulated using extended precision integer arithmetic instructions.

Double words occupy register pairs. A register pair, RP_x, consists of R_x and R_{(x + 1) mod 16}. R_x holds the most significant bits (MSB) and R_{(x + 1) mod 16} holds the least significant bits (LSB).

Integers

The range of integers which can be represented in two's complement form is -2,147,483,648 through 2,147,483,647. The range of unsigned integers is 0 through 4,294,967,295. The MSB of any data type is referred to as the sign bit (s), as shown below.

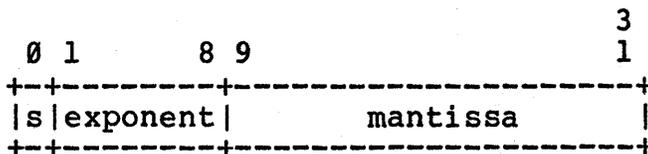


s = sign bit

Real Numbers (Single Precision)

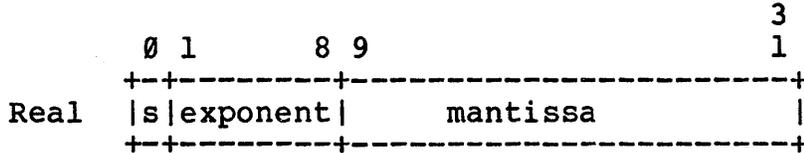
Real numbers (represented in floating-point form) consist of three parts: a sign, a power-of-two exponent, and a mantissa. The value of a real number is:

$$(-1)^{**s} \times 2^{**(\text{exponent}-127)} \times 1.\text{mantissa}$$



Ridge Processor

For positive numbers, the sign bit is 0. For negative numbers, the sign bit is 1. The exponent of a real number is 8 bits long, and is biased by +127. The eight bits of the exponent give a range of 0 through 255. Subtracting the bias yields an exponent range of -127 through +128. The mantissa has an implicit leading one, and is 23 bits long. Zero is represented by all zeros.

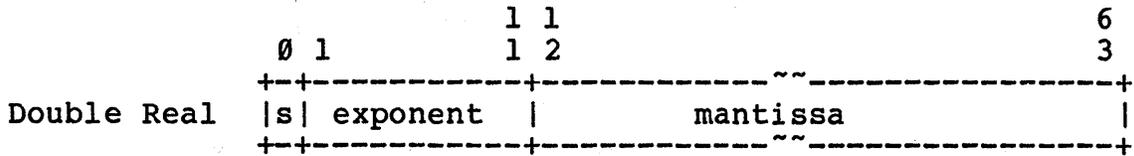


examples:

"1" = 0 01111111 000000000000000000000000 = 3F80 0000
 "-10" = 1 10000010 010000000000000000000000 = C120 0000

Real Numbers (Double Precision)

Double real numbers are similar to real numbers, except that the mantissa is 52 bits and the exponent is 11 bits. The exponent is biased by +1023. The eleven exponent bits give a range of 0 through 2047. Subtracting the bias yields an exponent range of -1023 through +1024.



examples:

"1" = 0 01111111111 0000000000000000...000000000000 = 3FF0 0000 0000 0000
 "-10" = 1 10000000010 0100000000000000...000000000000 = C024 0000 0000 0000

SYNTAX CONVENTIONS

In the descriptions of instructions, the 16 general registers are referred to as Rx or Ry. A register pair is referred to as RPx and consists of Rx and R(x + 1). Registers 0 through 15 are referred to as R0 through R15. The program counter is referred to as PC.

Some instructions can optionally specify the 4-bit value in the Ry register field instead of the contents of Ry. This is indicated by using "Ry field" instead of "Ry".

Specific bits of a register are enclosed in brackets. For example, bit 3 of a register is referred to as Rx[3]. The symbol ".." denotes a range of bits. For example, consecutive bits 6 through 9 of a register are referred to as Rx[6..9].

The instructions in the following sections are documented in the format shown below:

Name of Instruction or Instruction Class

Instruction Summary:

Instruction Mnemonic	Instruction Function	Syntactical Description
TYP	Typical	This is a typical instruction

Operation:

The TYP instruction has no operation, it is an example of syntax conventions.

Each effective address for a memory reference instruction is explained below.

Displacement. The memory address is the displacement field from the instruction. All memory references are 32-bit virtual addresses. This form references data space.

Ry + displacement. The contents of register Ry are added to the displacement field. Memory is then read or written at this location.

PC + displacement. Instructions that reference code space do so relative to the program counter (PC). PC is added to the displacement field and memory is read from this location. Code space is never written.

PC + Ry + displacement. PC is added to the displacement field, the result is added to the contents of Ry. Memory is then read at this location.

Indexing takes place with full 32-bit signed integers in two's complement notation. Displacements are also treated as 32-bit signed integers in two's complement notation. Short displacement memory addresses are sign extended to 32 bits by replicating the MSB into the upper 16 bits. The resulting effective address is an absolute displacement from location zero in the data space. Negative addresses (MSB set) are virtual addresses in the range of two to four billion.

These address computations allow indexes to be positive or negative relative to the displacement, or allow the displacement to be positive or negative relative to the index. Code space addresses are program counter (PC) relative and thus make relocatable constants.

All addressing formats have the same instruction execution time. Instructions referencing data space optionally add Ry to the displacement as the address is presented to memory. Instructions referencing code space optionally add Ry to the precomputed PC + displacement. The fetch unit contains logic that performs this function as part of the instruction prefetch.

INSTRUCTION DESCRIPTIONS

Descriptions of load, store, and load address memory instructions follow.

Load Instructions

Instruction Summary:

LOADB	Load Byte	Rx[24..31]	<- contents of (Ry + displacement)
		Rx[0..23]	<- 0
LOADH	Load Halfword	Rx[16..31]	<- contents of (Ry + displacement)
		Rx[0..15]	<- 0
LOAD	Load Word	Rx	<- contents of (Ry + displacement)
LOADD	Load Double Word	RPx	<- contents of (Ry + displacement)

Operation:

The register Rx is loaded with the data stored in memory at the effective address. Ry may optionally be used as an index register. The data element must be aligned on a boundary that is a multiple of the length of the data element.

The **LOADB** instruction loads the byte into bits 24-31 of the specified register and sets bits 0-23 to zero.

The **LOADH** instruction loads the halfword into bits 16-31 of the specified register and sets bits 0-15 to zero.

The **LOAD** instruction loads the word into the specified register.

The **LOADD** instruction loads two words into RPx.

The instructions shown above are for loading data from data space. A load-from-code-space form for each of the above instructions (**LOADBP**, **LOADHP**, **LOADP**, **LOADDP**) is described in the Ridge Assembler Manual.

Store Instructions

Instruction Summary:

STOREB	Store Byte	Rx[24..31]	-> contents of (Ry + displacement)
STOREH	Store Halfword	Rx[16..31]	-> contents of (Ry + displacement)
STORE	Store Word	Rx	-> contents of (Ry + displacement)
STORED	Store Double Word	RPx	-> contents of (Ry + displacement)

Operation:

The store instructions move data from the registers into memory. The effective address must be a multiple of the length of the data element.

The STOREB instruction places bits 24-31 of the specified register into memory at the effective address. Other bits (0-23) are ignored.

The STOREH instruction places bits 16-31 of the specified register into memory at the effective address. Other bits (0-15) are ignored.

The STORE instruction places the word into memory at the effective address.

The STORED instruction places the double words into memory at the effective address.

Load Address Instructions

Instruction Summary:

LADDR	Load Address	$Rx \leftarrow (\text{value of } Ry) + \text{displacement}$
LADDRP	Load Code Address	$Rx \leftarrow (\text{value of PC}) + Ry + \text{displacement}$

Operation:

The load address instructions store the effective address into Rx. These instructions do not perform a memory reference, but instead load a constant from the instruction stream into a register.

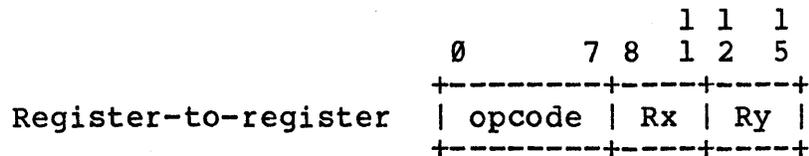
The LADDR instruction can be used to load two- or four-byte immediate values and, in indexed mode, can be used to add a constant to a register.

The LADDRP instruction is similar to the LADDR instruction except that PC is added to the displacement field.

CHAPTER 3: REGISTER FORMAT INSTRUCTIONS

INSTRUCTION FORMAT

Register-to-register format instructions process data taken from a specified general register. These instructions use the register-to-register instruction format shown below. Generally, two registers are specified and the result usually replaces Rx.



A few register-to-register format instructions also have an immediate mode. In immediate mode the 4-bit value of the Ry register field is used to specify an integer in the range from 0 to 15.

INSTRUCTION DESCRIPTIONS

Descriptions of the register-to-register format instructions follow.

Integer Arithmetic Instructions

Instruction Summary:

ADD	Integer add	$Rx \leftarrow Rx + Ry$
DIV	Integer divide	$Rx \leftarrow Rx/Ry$
MPY	Integer multiply	$Rx \leftarrow Rx * Ry$
NEG	Integer negate	$Rx \leftarrow 2's \text{ complement of } Ry$
REM	Integer remainder	$Rx \leftarrow Rx - ((Rx/Ry) * Ry)$
SUB	Integer subtract	$Rx \leftarrow Rx - Ry$

Operation:

The integer arithmetic instructions operate on 32-bit two's complement integers.

The ADD instruction adds Rx and Ry and puts the sum in Rx.

The DIV instruction divides Rx by Ry and puts the quotient in Rx.

The MPY instruction multiplies Rx and Ry and replaces the contents of Rx with the low order 32 bits of the product.

The NEG instruction puts the 2's complement of Ry in Rx.

The REM instruction divides Rx by Ry and puts the signed remainder in Rx. The sign of the remainder will be the sign of the divisor.

The SUB instruction subtracts Rx from Ry and puts the difference in Rx.

Logical Operator Instructions

Instruction Summary:

AND	Logical And	Rx ← Rx AND Ry
MOVE	Move Register	Rx ← Ry
NOT	Logical Not	Rx ← 1's complement of Ry
OR	Logical Or	Rx ← Rx OR Ry
XOR	Logical Xor	Rx ← Rx XOR Ry
NOP	No operation	Rx ← Rx

Operation:

The logical operator instructions operate on 32-bit unsigned integers in registers. The result replaces the contents of Rx.

The AND instruction performs logical AND on the contents of Rx and Ry and puts the result in Rx.

The MOVE instruction copies the contents of Ry into Rx.

The NOT instruction complements the contents of Ry and puts the result in Rx.

The OR instruction performs logical OR on the contents of Rx and Ry and puts the result in Rx.

The XOR instruction performs logical XOR on the contents of Rx and Ry and puts the result in Rx.

The NOP instruction performs no operation and is often used to fill instruction space thus consuming time. It supplies padding between modules to allow for proper alignment.

Extended Precision Integer Instructions

Instruction Summary:

EADD	Extended Integer Add	Rx <- Rx + Ry + R0[31] R0[31] <- carry R0[30] <- overflow
EDIV	Extended Integer Divide	Rx <- RPx/Ry Ry <- the remainder
EMPY	Extended Integer Multiply	RPx <- Rx*Ry
ESUB	Extended Integer Subtract	Rx <- Rx 1's complement + Ry + R0[31] R0[31] <- carry, R0[30] <- overflow

Operation:

The extended precision integer instructions can be used to implement multiple-word arithmetic.

The EADD instruction adds the two's-complement integers in Rx and Ry, and at the same time adds the carry-in from R0[31], and puts the least significant 32 bits of the sum in Rx. The carry-out (most significant) bit is put in R0[31]. Overflow is indicated in R0[30]. The upper 30 bits of R0 are set to zero.

The typical use of the EADD instruction to implement multiple-word arithmetic is used as follows: R0[31] is set to zero. The least significant words are EADDED, the next-most significant words are EADDED, and so on to the most significant words. Overflow can then be checked after the last EADD.

The EDIV instruction divides the 64-bit unsigned contents of RPx by the unsigned 32-bit contents of Ry, and places the unsigned quotient in Rx and the unsigned remainder in Ry.

The EMPY instruction takes two unsigned 32-bit integers and produces an unsigned 64-bit product and places it in RPx.

The ESUB instruction one's complement subtracts the two's-complement integers in Rx and Ry, and at the same time adds the carry-in from R0[31], then puts the least significant 32-bit two's complement difference in Rx. The carry-out (most significant) bit is put in R0[31]. Overflow is indicated in R0[30]. The upper 30 bits of R0 are set to zero.

Ridge Processor

The typical use of the ESUB instruction to implement multiple-word arithmetic is used as follows: R0[31] is set to one. The least significant words are ESUBed, the next-most significant words are ESUBed, and so on to the most significant words. Overflow can then be checked after the last ESUB.

Real Instructions

Instruction Summary:

FIXR	Round Real to Integer	Rx	<- ROUND Ry
FIXT	Truncate Real to Integer	Rx	<- TRUNC Ry
FLOAT	Convert Integer to Real	Rx	<- FLOAT Ry
MAKERD	Convert Real to Double Real	RPx	<- DOUBLE Ry
RADD	Real Add	Rx	<- Rx + Ry
RDIV	Real Divide	Rx	<- Rx/Ry
RMPY	Real Multiply	Rx	<- Rx*Ry
RNEG	Real Negate	Rx	<- -Ry
RSUB	Real Subtract	Rx	<- Rx - Ry

Operation:

These instructions operate on 32-bit real numbers.

The FIXR instruction converts the single-precision real contents of Ry into a two's complement integer in Rx. Fractions of .5 or more are rounded up to the next higher absolute value.

The FIXT instruction converts the single-precision real number in Ry into a 32-bit integer in Rx. All bits to the right of the decimal point are lost.

The FLOAT instruction converts the integer in Ry into a real number in Rx and rounds if necessary.

The MAKERD instruction converts the real number in Ry into a double precision real number in RPx.

The RADD instruction adds the 32-bit real numbers in Rx and Ry and puts the sum in Rx.

The RDIV instruction divides the 32-bit real number in Rx by the 32-bit real number in Ry and puts the result in Rx.

The RMPY instruction multiplies the 32-bit real numbers in Rx and Ry and puts the product in Rx.

The RNEG instruction negates the real number in Ry and puts the result in Rx.

The RSUB instruction subtracts the real number in Ry from the real number in Rx and puts the difference in Rx.

Double Real Instructions

Instruction Summary:

DFIXR	Round Double Real to Integer	Rx <- ROUND RPy
DFIXT	Truncate Double Real to Integer	Rx <- TRUNC RPy
DFLOAT	Convert Integer to Double Real	RPx <- DOUBLE FLOAT Ry
DRADD	Double Real Add	RPx <- RPx + RPy
DRDIV	Double Real Divide	RPx <- RPx/RPy
DRMPY	Double Real Multiply	RPx <- RPx*RPy
DRNEG	Double Real Negate	RPx <- -RPy
DRSUB	Double Real Subtract	RPx <- RPx - RPy
MAKEDR	Round Double Real to Real	Rx <- REAL RPy

Operation:

The double real instructions perform the same operations as the real instructions previously described, except the double real instructions operate on double real format data, working on register pairs.

Bit-Oriented Instructions

Instruction Summary:

CBIT	Clear Bit	$RPx[Ry \bmod 64] \leftarrow 0$
SBIT	Set Bit	$RPx[Ry \bmod 64] \leftarrow 1$
TBIT	Test Bit	$Rx[31] \leftarrow RPx[Ry \bmod 64]$ $Rx[0..30] \leftarrow 0$

Operation:

The CBIT instruction specifies a bit number from 0-63 in Ry and the specified bit of RPx is set to zero.

The SBIT instruction specifies a bit number from 0-63 in Ry and the specified bit of RPx is set to 1.

In the TBIT instruction Ry specifies a bit number from 0-63 which is tested in RPx. The tested bit is duplicated in bit 31 of Rx, and bits 0-30 of Rx are set to zero.

Test Instruction

Instruction Summary:

TEST	Test Values	Rx ← 1 if Rx relop Ry is true, or Rx relop (4-bit Ry field) is true
		Rx ← 0 if Rx relop Ry is true, or Rx relop (4-bit Ry field) is false

Operation:

The TEST instruction uses a relational operator (relop) to compare two values and sets Rx to either 0 or 1, depending on the result of the test. The second operand is either the contents of the register Ry, or the 4-bit value of the Ry field. The comparison is done using signed two's complement arithmetic. The comparison relop may be one of the following: equal to (=), less than (<), greater than (>), not equal to (<>), less than or equal to (<=), or greater than or equal to (>=).

Compare Instructions

Instruction Summary:

LCOMP	Logical Compare	Rx ← -1, if Rx < Ry Rx ← 0, if Rx = Ry Rx ← 1, if Rx > Ry
DCOMP	Double Integer Compare	Rx ← -1, if RPx < RPy Rx ← 0, if RPx = RPy Rx ← 1, if RPx > RPy
RCOMP	Real Compare	Rx ← -1, if Rx < Ry Rx ← 0, if Rx = Ry Rx ← 1, if Rx > Ry
DRCOMP	Double Real Compare	Rx ← -1, if RPx < RPy Rx ← 0, if RPx = RPy Rx ← 1, if RPx > RPy

Operation:

The LCOMP instruction compares registers Rx and Ry using unsigned arithmetic. Register Rx is set to -1, 0, or +1, depending on whether Rx is less than, equal to, or greater than Ry, respectively.

The DCOMP instruction compares register pairs RPx and RPy using two's complement arithmetic. Register Rx is set to -1, 0, or +1, depending on whether RPx is less than, equal to, or greater than RPy, respectively.

The RCOMP instruction compares real numbers in registers Rx and Ry using sign magnitude form. Register Rx is set to -1, 0, or +1, depending on whether Rx is less than, equal to, or greater than Ry, respectively.

The DRCOMP instruction compares double real numbers in register pairs RPx and RPy using sign magnitude form. Register Rx is set to -1, 0, or +1, depending on whether RPx is less than, equal to, or greater than RPy, respectively.

Shift Instructions

The shift instructions take the shift count from the contents of register Ry or from the 4-bit value of the Ry field. All shift execution times are independent of the number of bits shifted due to the use of the barrel shifter.

Single register shifts shift the value in Rx from 0 to 31 bits. Double register shifts shift the value in RPx from 0 to 63 bits. Only the low order 5 bits (6 bits for double shifts) of Ry are used as the shift count. The immediate shift forms allow shifts from 0 to 15 bits using the four bits of Ry field as the shift count.

Instruction Summary:

CSL	Circular Shift Left	Rx circularly shifted left by Ry or 4-bit Ry field
LSL	Logical Shift Left	Rx shifted left by Ry or 4-bit Ry field
LSR	Logical Shift Right	Rx shifted right by Ry or 4-bit Ry field
ASL	Arithmetic Shift Left	Rx shifted left by Ry or 4-bit Ry field
ASR	Arithmetic Shift Right	Rx shifted right by Ry or 4-bit Ry field, filling with sign bit
DLSL	Double Logical Shift Left	RPx shifted left by Ry or 4-bit Ry field
DLSR	Double Logical Shift Right	RPx shifted right by Ry or 4-bit Ry field

Operation:

The CSL instruction circularly shifts bits left in Rx. Bits shifted out of bit 0 are shifted into bit 31.

The LSL instruction shifts bits left in Rx and fills emptied positions with zeros.

The LSR instruction shifts bits right in Rx and fills emptied positions with zeros.

The ASL instruction shifts left and preserves the sign bit.

The ASR instruction shifts right and fills the left bits with duplicates of the sign bit.

The DLSL and DLSR instructions correspond to LSL and LSR, except that RPx is treated as a single 64-bit register.

Sign Extend Instructions

Instruction Summary:

SEB	Sign Extend Byte	Rx[0..23] <- Ry[24],
		Rx[24..31] <- Ry[24..31]
SEH	Sign Extend	Rx[0..15] <- Ry[16],
	Halfword	Rx[16..31] <- Ry[16..31]

Operation:

The sign extend instructions change 8- or 16-bit integers into full word integers.

The SEB instruction makes bits 0-23 in register Rx the same as bit 24 in register Ry. Bits 24-31 in Ry are copied to Rx.

The SEH instruction makes bits 0-15 in register Rx the same as bit 16 in register Ry. Bits 16-31 in Ry are copied to Rx.

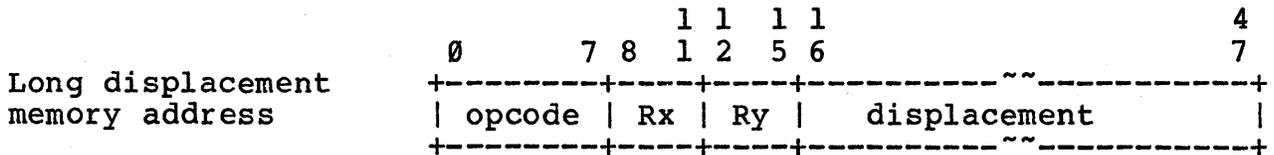
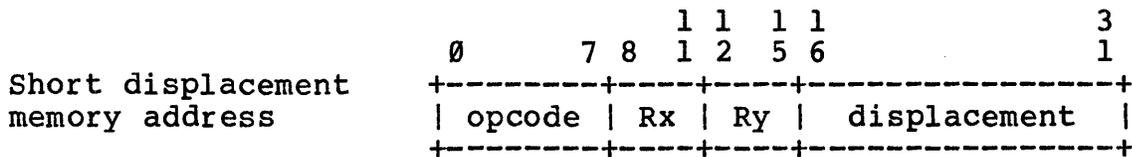
BLANK

CHAPTER 4: PROGRAM CONTROL INSTRUCTIONS

BRANCH INSTRUCTIONS

INSTRUCTION FORMAT

Branch instructions use either the short or long displacement memory address instruction formats shown below. When the least significant bit of the displacement is set, the branch is predicted to be taken.



Branch instructions either switch execution to the instruction at the branch target address, or have no effect. If the branch instructions have no effect then the next sequential instruction following the branch is executed. Branch instructions affect the value of the program counter (PC) as shown below.

$$\text{Next PC} \leftarrow \text{PC} + \text{branch instruction length} \quad (\text{next sequential instruction})$$

or

$$\text{Next PC} \leftarrow \text{PC} + \text{displacement} \quad (\text{branch target address})$$

The branch instructions use program counter (PC) relative addressing, which allows self-relocating code. The target address of the branch instruction is computed by adding the 32-bit signed displacement (sign extended to 32 bits in the short form case) to the PC at the beginning of the branch instruction.

The least significant bit of the displacement field is used by the processor to predict whether or not the branch will be taken. If the bit is one, the processor will prefetch the instruction at the target address. If the bit is zero, the processor will prefetch the next sequential instruction. If the bit is incorrect, the program will execute correctly, but the next instruction after the branch will be delayed by four cycles to fill the pipeline.

INSTRUCTION DESCRIPTIONS

Descriptions of the branch instructions follow.

Branch Instructions

Instruction Summary:

BR	Unconditional Branch	PC ← PC + displacement
BR	Conditional Branch	if Rx relop Ry, PC ← PC + displacement

Operation:

The unconditional branch instruction changes PC to the target address (PC + displacement). The branch prediction bit is ignored and the target instruction is always prefetched.

The conditional branch instruction compares Rx to the contents of Ry or to the 4-bit value of the Ry field, then may conditionally branch to the target location. The conditional branch instruction comparisons are made using two's complement arithmetic. The comparison uses the relational operator (relop), which may be: equal to (=), less than (<), greater than (>), not equal to (<>), less than or equal to (<=), or greater than or equal to (>=).

Loop Control Instruction

Instruction Summary:

LOOP	Increment and Branch	Rx ← Rx + Ry field, if Rx < 0, PC ← PC + displacement
------	-------------------------	---

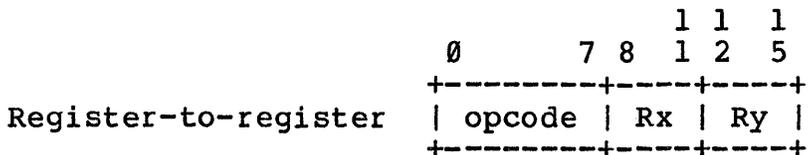
Operation:

The LOOP instruction is similar to the conditional branch described above. The LOOP instruction adds the 4-bit value of the Ry field to the contents of Rx and branches to the target location if the result is less than zero. If Rx is equal to or greater than zero, the next sequential instruction is executed.

Call Subroutine Register and Return Instructions

Instruction Format:

The CALLR and RET instructions use the register-to-register instruction format shown below.



Instruction Summary:

CALLR	Call Subroutine Register	Rx <- PC + 2,
		PC <- PC + Ry
RET	Return from Subroutine	Rx <- PC + 2,
		PC <- Ry

Operation:

The CALLR instruction stores the address of the next sequential instruction (PC + 2) in Rx, and branches to the location PC + Ry.

The RET instruction stores the address of the next sequential instruction (PC + 2) in Rx, and branches to the absolute address in Ry. The main use of RET is in returning from subroutines, but it can also be used as a call to a subroutine when the absolute rather than the relative address is known. Care must be taken in using the RET instruction for this purpose so that the code remains self-relocating.

APPENDIX A: RIDGE OPCODE MAP

RIDGE OPCODE MAP

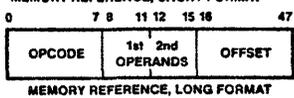
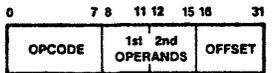
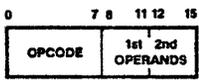
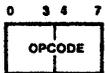
Least Significant Nibble (Hex), Opcode (4:7)

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0		MOVE	NEG	ADD	SUB	MPY	DIV	REM	NOT	OR	XOR	AND	CBIT	SBIT	TBIT	CHK
1	NOP	MOVE Immed		ADD Immed	SUB Immed	MPY Immed			NOT Immed			AND Immed				CHK Immed
2	FIXT	FIXR	RNEG	RADD	RSUB	RMPY	RDIV	MAKERD	LCOMP	FLOAT	RCOMP		EADD	ESUB	EMPTY	EDIV
3	DFIXT	DFIXR	DRNEG	DRADD	DRSUB	DRMPY	DRDIV	MAKEDR	DCOMP	DFLOAT	DRCOMP	TRAP				
4	SUS	LUS	RUM	LDREGS	TRANS	DIRT	MOVE SR ← R R ← SR						MAINT		READ	WRITE
5	>	TEST < =	CALLR	TEST Immediate > < =		RET	< =	TEST > = <>	KCALL	TEST Immediate < = > = <>						
6	LSL	LSR	ASL	ASR	DLSL	DLSR			CSL		SEB					
7	LSL Immed	LSR Immed	ASL Immed	DLSL Immed	DLSR Immed	ASR Immed			CSL Immed		SEH					
8	>		=	>	<	=	< =		<>				< =	> =	<>	
9	BR		BR	CALL	BR Immediate		LOOP				BR	BR	BR Immediate			
A																
B		STOREB X		STOREH X				STORE X		STORED X						
C			X		X				X		X					
D			LOADB X		LOADH X				LOAD X		LOADD X					LADDR X
E				X		X				X		X				
F			LOADBP X		LOADHP X				LOADP X		LOADDP X					LADDRP X
		X			X					X		X				X

X = Indexed (i.e., target address is further offset by a register named in the second operand field).
 Immediate (Immed) = the second operand field contains a value.

Register Format

INSTRUCTION FORMATS:



Segment Referenced	Length
Code	Short
Code	Long
Data	Short
Data	Long
Data	Short
Data	Long
Code	Short
Code	Long

Memory Reference Format

Ridge Computers
Corporate Headquarters

2451 Mission College Blvd.
Santa Clara, California 95054
Phone: (408) 986-8500
Telex: 176956

