

Usage of DLX Instructions By SPEC92 Integer Code




FIGURE 2.28 Graphical display of instructions executed of the five programs from SPECint92 in Figure 2.26.

Usage of DLX Instructions By SPEC92 Floating-Point Code




FIGURE 2.29 Graphical display of instructions executed of the five programs from SPECfp92 in Figure 2.27.

```
beq  r1,$0 SKIP  ! If r1 = 0 goto SKIP
nop
add  r2, r2, r5
SKIP:
add  r2, r3, r4
```

Control transfer (control flow) instructions (CTIs) ...

... may cause next instruction to be fetched from ...

... somewhere other than $PC + 4$ (assuming 4-byte instructions).

Programs would be trivial without them.

Occur frequently in programs.

Often a performance bottleneck.

Names used below are common:

- *Branch*:
Conditional control transfer.
- *Jump*:
Unconditional control transfer.
Sometimes special case of branch.
- *Jump and Link*:
Unconditional, return address (PC) saved in register.
- *Call*:
Unconditional control transfer, PC, etc. saved.
Sometimes special case of jump and link.
- *Return*:
Unconditional, PC, etc. from most recent call restored.
Sometimes special case of jump and link.

CTI by type MIPS running SPEC92 benchmarks.

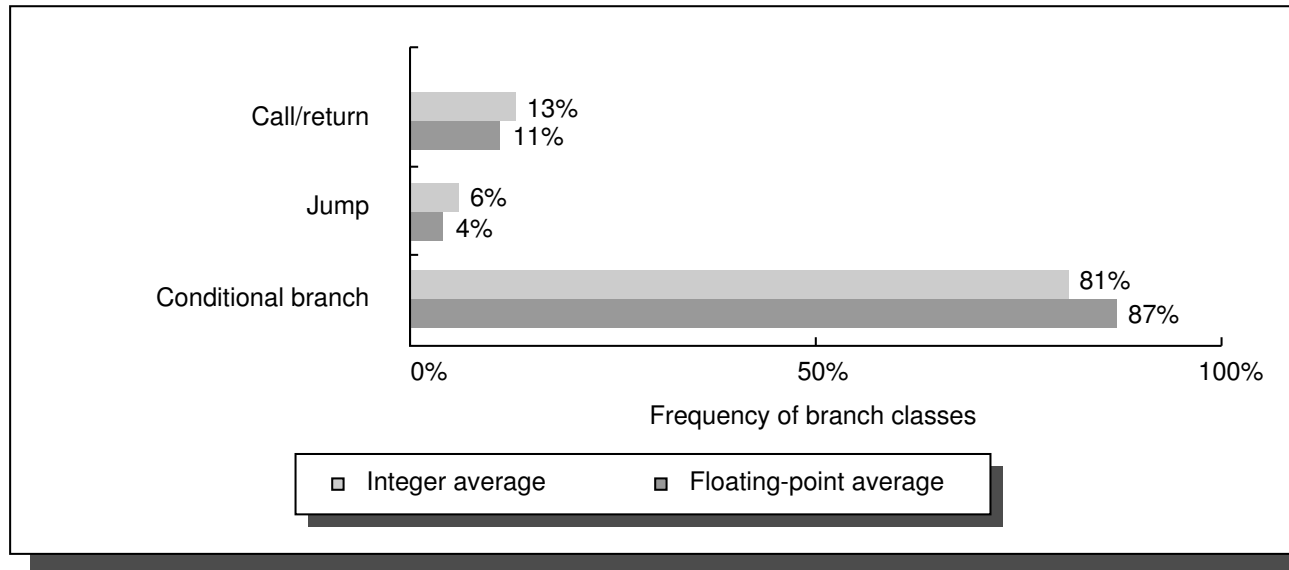


FIGURE 2.12 Breakdown of control flow instructions into three classes: calls or returns, jumps, and conditional branches.

Any addressing mode *could* be used for destination.

Several are common:

- *Absolute*

Destination address is an immediate.

Best for procedure calls. . .

. . .because destination can be far away.

Not used in MIPS or SPARC.

- *PC-Relative* or *Displacement*

Destination is *displacement* added to program counter.

ISA design choice: displacement size.

Good for conditional branches. . .

. . .because destination usually close by. . .

. . .and so small immediates suffice.

Used in branch instructions by most ISAs, including MIPS and SPARC.

- *Register Indirect*

Destination in register.

Used for subroutine return addresses.

Used for function address passed as parameters, as with `qsort`.

Also used by ISAs in which immediates smaller than addresses.

Used in MIPS (`jr` and `jalr`) and SPARC (`jmp1`).

- *Indexed*

Destination is sum of two registers.

Useful for C `switch` and similar statements.

Used in SPARC `jmp1` (jump and link), not used in MIPS.

- *Region*

Destination is high few bits of PC concatenated with immediate.

Used in MIPS `j` and `jal` instructions.

Branch distances on DLX running SPEC92 programs.

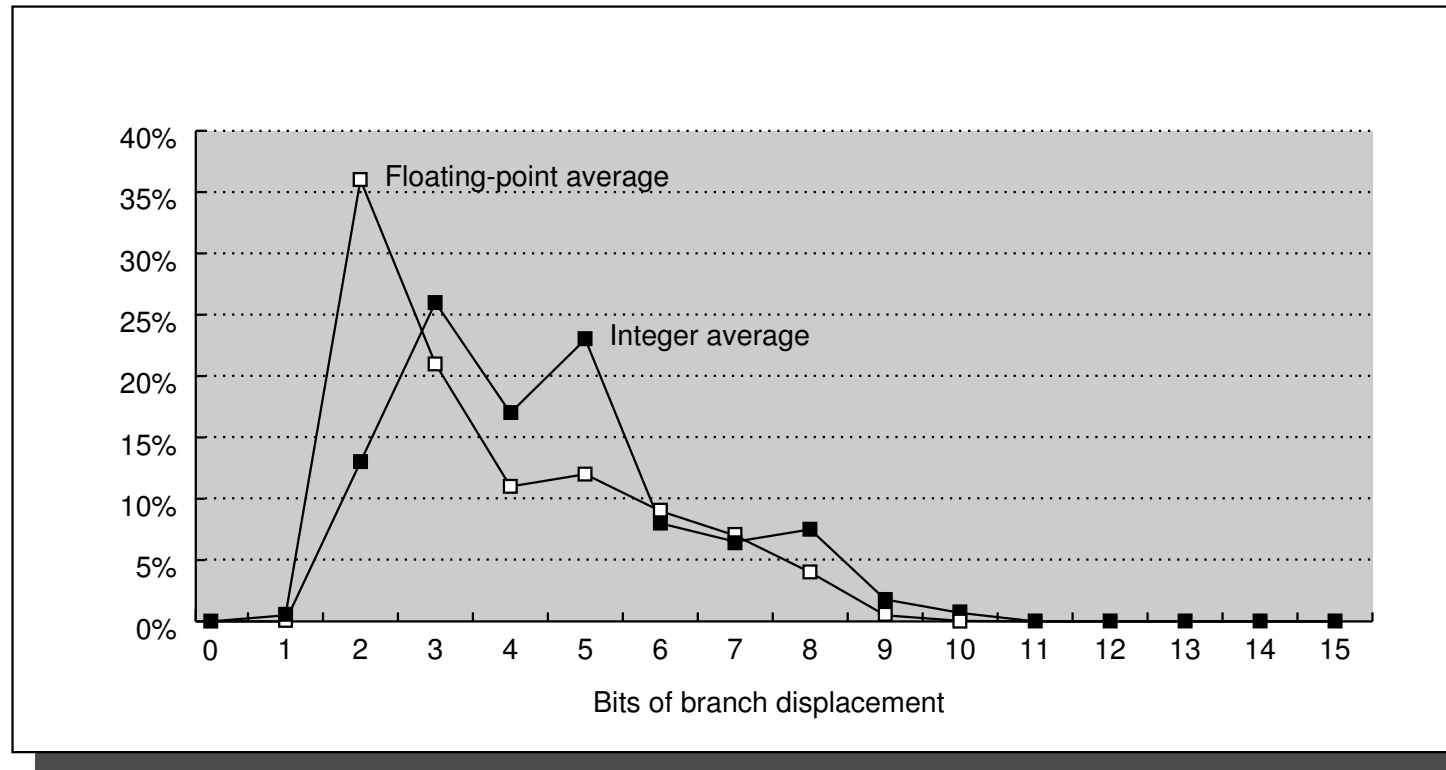


FIGURE 2.13 Branch distances in terms of number of instructions between the target and the branch instruction.

In DLX branch displacement limited to 16 bits.

How branch condition might be specified:

- Test value of general-purpose register (GPR).
- Test value of special-purpose *condition code register* (CCR).
- Test value of special-purpose *loop-counter register*.
- Comparison specified in branch instruction.

Note: *test* value means *test if value is zero* ...

...which is much faster than *test if value greater than constant*.

A typical ISA would use one or two of these methods.

Itanium has special *counted loop branches* and a *loop count* register (LC).

If LC register non-zero counted loop branch is taken and register decremented.

```
mov lc = 10           ! Set loop count to 10.
LOOP:
add r100 = r100, r101 ! There are 128 64-bit GPRs.
br.cloop LOOP;;      ! if( lc > 0 ) { lc--; goto LOOP; }
```

Factors

For compact code and programmer convenience:

⇒ Comparison in branch instruction.

For fast implementation:

⇒ Test GPR. (But may “waste” registers.)

⇒ Test CCR. (Maybe limited to one condition at a time.)

⇒ Test loop counter.

Procedures (A.k.a., subroutines, functions.)

Fundamental part of every nontrivial program.

Requires careful support in ISA.

Mandatory ISA Support

Call instruction saves PC in register.

Return restores saved PC.

Additional Support, Provided by ISA or Software (*ABI*).

Save and restore registers.

Prepare *stack frame* of called procedure.

CTI Behaviors Chosen to Speed Implementation

Delayed Transfer:

Control transfer occurs 1 instruction after CTI.

E.g., consider execution of instruction 1 of DLX code:

```
1  j TARGET          ! Jump to TARGET.
2  add r1,r2,r3      ! Executed after jump.
3  sub r4,r5,r6
```

TARGET:

Without delayed CT, instruction 2 not executed.

With delayed CT instruction 2 is executed, but not 3.

Speeds simple implementations, but complicates modern designs.

Delayed control transfers used on many RISC ISAs, including MIPS, SPARC.

Not included on Alpha.

Branch Instructions with *Prediction Hints*

Programmer indicates (hints) whether ...

... branch is expected to be taken most of the time or ...

... not taken most of the time.

Hint only affects execution speed, not what the program does.

If programmer correct, execution may be faster ...

... if programmer wrong execution is correct but slower.

A low-cost implementation might ignore the hint.

Example: **Hypothetical** prediction hints added to MIPS.

Hint indicated by following instruction mnemonic with ...

... **,pt** (predict taken) or ...

... **,pn** (predict not taken).

```
bltz,pn  r1, LINE1    ! Programmer expects r1 to be nonnegative most of the time. ■
nop
bgtz,pt  r2, LINE2    ! Programmer expects r2 positive most of the time.
nop
add      r3, r4, r5
LINE1:
addi     r3, r3, 1
LINE2:
add      r3, r6, r7
```

Predicated Execution

Predicated Instructions:

Arithmetic or register move instructions along with a condition.

Results written **only if** condition true.

If condition false instruction may still be executed ...
... though results not written.

When predicated execution used fewer branches needed ...
... though total number of instructions may be greater.

Because branches can slow implementations ...
... execution can be faster with predication even though more instructions executed.

Example: Add predication to MIPS (hypothetical). (Pay Attention)

Predicate indicated in parenthesis before mnemonic.

(⟨rp⟩) means write results if ⟨rp⟩ nonzero ...

... (!⟨rp⟩) means write results if ⟨rp⟩ is zero.

! High level code: if(r1 == 0) r2 = r2 + r5 else r2 = r2 + r4;

! Without predicated execution.

```
    beq    r1,$0 LINE1
```

```
    nop
```

```
    add    r2, r2, r4
```

```
    j     LINE2
```

```
    nop
```

```
LINE1: add    r2, r2, r5
```

```
LINE2: ...
```

! With predicated execution.

```
(r1) add    r2, r2, r4 ! Execute if r1 nonzero.
```

```
(!r1) add   r2, r2, r5 ! Execute if r1 zero.
```

Is it a coincidence that the destination and first source operands were the same in the `adds`?

Consider the way cars are classified:

Criteria: performance, comfort, off-road capabilities.

Classifications: sports cars, luxury cars, SUVs.

A classification may be based on one or more criteria . . .

. . . for example sports cars are high performance and comfortable (excluding ride).

Processors are also classified based on multiple criteria.

Processor Classification: Three Criteria

Address space size: 32-bit, 64-bit, etc.

Register use by instructions: Load/store, memory/memory, stack, etc.

Instruction format: Fixed length, variable length, *bundled* (VLIW).

Describes address space size.

n-bit processor indicates...

... *n-bit* address space and fast *n-bit* integer instructions for pointer arithmetic.

Terminology widely adopted but not universal.

Temptation to inflate *n*, *e.g.*, 128-bit processor (may refer to bus width).

Describes how instructions can use registers and addressing modes.

Load/Store, General Purpose Register

Memory only accessed by special load and store instructions.

Few special purpose registers.

Arithmetic and logic instructions use register or immediates.

Examples include most ISAs developed in 1980s and 1990s.

Memory/Memory, General Purpose Register

Memory can be read and written by arithmetic and logic instructions.

Arithmetic and logic instructions can also use register and immediate operands.

Few special purpose registers.

Characteristic of older ISAs.

Accumulator

Arithmetic instructions implicitly use accumulator.

May include other special-purpose registers.

Stack

Instructions refer to stack.

Used in several old and new machines.

Describes allowed instruction sizes and something called bundling.

Fixed Size

All instructions same size. (*E.g.*, 32 bits).

All instructions do about same amount of work.

Eases implementation.

Characteristic of newer ISAs.

Variable Size

Instruction size depends on instruction and operands.

Potential for compact code.

Implementation difficult.

Characteristic of older and special-purpose ISAs.

Bundled (VLIW,EPIC)

Instructions grouped into bundles, to be executed as a group.

Bundles also contain information to speed execution.

Used in Itanium ISA and older experimental machines.

Common Classifications and Status

RISC (Reduced Instruction Set Computing)...

... Current General Purpose.

CISC (Complex Instruction Set Computing)...

... Former and Legacy (still used) General Purpose.

VLIW (Very Long Instruction Word)...

... Future General Purpose?

Stack...

... Embedded (built in).

Accumulator...

... Past General Purpose.

CISC (*Complex Instruction Set Computing*)

Criteria

Register Use: Memory/Memory, General Purpose Register

Instruction Format: Variable Instruction Size

Characteristics

Programmer friendly, lots of instruction choices.

Programmer un-friendly, may contain many almost-useful special purpose instructions.

Average code size.

Implementation cost inflated by ISA complexity.

From mid 1980s to late 1990s ISA complexity slowed implementations but ...

... *fin-de-siècle*¹ CISC processors nearly as fast as RISC processors.

Examples: VAX, 68000, arguably IA-32 (popularly known as 80x86).

¹ End of century, 20th century here.

RISC (Reduced Instruction Set Computing)

Criteria

Register Use: Load/Store, General Purpose Register

Instruction Format: Fixed Instruction Sizes

Characteristics

Work performed by each instruction balanced to ease implementation.

Few special purpose, rarely used instructions.

Code size large. (Because of fixed instruction size.)

RISC Characteristics, continued.

Instructions designed to speed implementation ...

... though sometimes design biased towards contemporary technology ...

... ignoring future technology that might be used in later implementations.

Some implementations rely on compiler to arrange (schedule) instruction for fast execution (discussed later).

Most major ISAs since 1985 are RISC types.

Examples: SPARC V9, Alpha, MIPS, HP PA (Precision Architecture), Power4, PowerPC.

VLIW (Very Long Instruction Word)

Special case: EPIC (Explicitly Parallel Instruction Computing)

Criteria

Instruction Format: Bundled.

Characteristics

Extra information in bundles allows higher speed implementations.

Examples: Cray MTA (supercomputer), TI 320C6 (meant for signal processing), ...
... the new Intel Itanium ISA (implemented by the Itanium processor).

Discussed further later in the course.

Stack

Criteria

Register Use: Stack

Instruction Format: Variable Width

Characteristics

Small code size. Important for *embedded* [in cell phone, VCR, automobile engine, etc.] processors.

Limited speed because stack limits execution overlap opportunities.

Examples: Burroughs Corp. machines (1961-1977), INMOS Transputer (1983-).