

CHAPTER 2

BASICS

2-1 Manipulating Rightmost Bits

Some of the formulas in this section find application in later chapters.

Use the following formula to turn off the rightmost 1-bit in a word, producing 0 if none (e.g., 01011000 \Rightarrow 01010000):

$$x \& (x - 1)$$

This may be used to determine if an unsigned integer is a power of 2; apply the formula followed by a 0-test on the result.

Similarly, the following formula can be used to test if an unsigned integer is of the form $2^n - 1$ (including 0 or all 1's):

$$x \& (x + 1)$$

Use the following formula to isolate the rightmost 1-bit, producing 0 if none (e.g., 01011000 \Rightarrow 00001000):

$$x \& (-x)$$

Use the following formula to isolate the rightmost 0-bit, producing 0 if none (e.g., 10100111 \Rightarrow 00001000):

$$-x \& (x + 1)$$

Use one of the following formulas to form a mask that identifies the trailing 0's, producing all 1's if $x = 0$ (e.g., 01011000 \Rightarrow 00000111):

$$-x \& (x - 1), \text{ or}$$

$$\neg(x \mid -x), \text{ or}$$

$$(x \& -x) - 1$$

The first formula has a little instruction-level parallelism.

Use the following formula to form a mask that identifies the rightmost 1-bit and the trailing 0's, producing all 1's if $x = 0$ (e.g., 01011000 \Rightarrow 00001111):

$$x \oplus (x - 1)$$

Use the following formula to right-propagate the rightmost 1-bit, producing all 1's if $x = \mathbf{0}$ (e.g., $0101\ 1000 \Rightarrow 0101\ 1111$):

$$x \mid (x - \mathbf{1})$$

Use the following formula to turn off the rightmost contiguous string of 1-bits (e.g., $0101\ 1000 \Rightarrow 0100\ 0000$):

$$((x \mid (x - \mathbf{1})) + \mathbf{1}) \& x$$

This may be used to see if a nonnegative integer is of the form $2^j - 2^k$ for some $j \geq k \geq 0$; apply the formula followed by a 0-test of the result.

These formulas all have duals in the following sense. Read what the formula does, interchanging 1's and 0's in the description. Then, in the formula, replace $x - \mathbf{1}$ with $x + \mathbf{1}$, $x + \mathbf{1}$ with $x - \mathbf{1}$, $\neg x$ with $\neg(x + \mathbf{1})$, $\&$ with \mid , and \mid with $\&$. Leave x and $\neg x$ alone. Then the result is a valid description and formula. For example, the dual of the first formula in this section reads as follows:

Use the following formula to turn on the rightmost 0-bit in a word, producing all 1's if none (e.g., $10100111 \Rightarrow 10101111$):

$$x \mid (x + \mathbf{1})$$

There is a simple test to determine whether or not a given function can be implemented with a sequence of *add*'s, *subtract*'s, *and*'s, *or*'s, and *not*'s [War]. We may, of course, expand the list with other instructions that can be composed from the basic list, such as *shift left* by a fixed amount (which is equivalent to a sequence of *add*'s), or *multiply*. However, we exclude instructions that cannot be composed from the list. The test is contained in the following theorem.

THEOREM. A function mapping words to words can be implemented with word-parallel add, subtract, and, or, and not instructions if and only if each bit of the result depends only on bits at and to the right of each input operand.

That is, imagine trying to compute the rightmost bit of the result by looking only at the rightmost bit of each input operand. Then, try to compute the next bit to the left by looking only at the rightmost two bits of each input operand, and so forth. If you are successful in this, then the function can be computed with a sequence of *add*'s, *and*'s, and so on. If the function cannot be computed in this right-to-left manner, then it cannot be implemented with a sequence of such instructions.

The interesting part of this is the latter statement, and it is simply the contrapositive of the observation that the functions *add*, *subtract*, *and*, *or*, and *not* can all be computed in the right-to-left manner, so any combination of them must have this property.

To see the “if” part of the theorem, we need a construction that is a little awkward to explain. We illustrate it with a specific example. Suppose that a function of two variables x and y has the right-to-left computability property, and suppose that bit 2 of the result r is given by

$$r_2 = x_2 \mid (x_0 \& y_1). \quad (1)$$

We number bits from right to left, 0 to 31. Because bit 2 of the result is a function of bits at and to the right of bit 2 of the input operands, bit 2 of the result is “right-to-left computable.”

Arrange the computer words x , x shifted left two, and y shifted left one, as shown below. Also, add a mask that isolates bit 2.

$$\begin{array}{cccccccc} x_{31} & x_{30} & \dots & x_3 & x_2 & x_1 & x_0 & \\ x_{29} & x_{28} & \dots & x_1 & x_0 & 0 & 0 & \\ y_{30} & y_{29} & \dots & y_2 & y_1 & y_0 & 0 & \\ 0 & 0 & \dots & 0 & 1 & 0 & 0 & \\ 0 & 0 & \dots & 0 & r_2 & 0 & 0 & \end{array}$$

Now, form the word-parallel *and* of lines 2 and 3, *or* the result with row 1 (following Equation (1)), and *and* the result with the mask (row 4 above). The result is a word of all 0’s except for the desired result bit in position 2. Perform similar computations for the other bits of the result, *or* the 32 resulting words together, and the result is the desired function.

This construction does not yield an efficient program; rather, it merely shows that it can be done with instructions in the basic list.

Using the theorem, we immediately see that there is no sequence of such instructions that turns off the leftmost 1-bit in a word, because to see if a certain 1-bit should be turned off, we must look to the left to see if it is the leftmost one. Similarly, there can be no such sequence for performing a right shift, or a rotate shift, or a left shift by a variable amount, or for counting the number of trailing 0’s in a word (to count trailing 0’s, the rightmost bit of the result will be 1 if there are an odd number of trailing 0’s, and we must look to the left of the rightmost position to determine that).

A novel application of the sort of bit twiddling discussed above is the problem of finding the next higher number after a given number that has the same number of 1-bits. You are forgiven if you are asking, “Why on earth would anyone want to compute that?” It has application where bit strings are used to represent subsets. The possible members of a set are listed in a linear array, and a subset is represented by a word or sequence of words in which bit i is on if member i is in the subset. Set unions are computed by the logical *or* of the bit strings, intersections by *and*’s, and so on.

You might want to iterate through all the subsets of a given size. This is easily done if we have a function that maps a given subset to the next higher number (interpreting the subset string as an integer) with the same number of 1-bits.

A concise algorithm for this operation was devised by R. W. Gosper [HAK item 175].¹ Given a word x that represents a subset, the idea is to find the rightmost contiguous group of 1's in x and the following 0's, and "increment" that quantity to the next value that has the same number of 1's. For example, the string $xxx0\ 1111\ 0000$, where xxx represents arbitrary bits, becomes $xxx1\ 0000\ 0111$. The algorithm first identifies the "smallest" 1-bit in x , with $s = x \& -x$, giving $0000\ 0001\ 0000$. This is added to x , giving $r = xxx1\ 0000\ 0000$. The 1-bit here is one bit of the result. For the other bits, we need to produce a right-adjusted string of $n - 1$ 1's, where n is the size of the rightmost group of 1's in x . This can be done by first forming the *exclusive or* of r and x , which gives $0001\ 1111\ 0000$ in our example.

This has two too many 1's, and needs to be right-adjusted. This can be accomplished by dividing it by s , which right-adjusts it (s is a power of 2), and shifting it right two more positions to discard the two unwanted bits. The final result is the *or* of this and r .

In computer algebra notation, the result is y in

$$\begin{aligned} s &\leftarrow x \& -x \\ r &\leftarrow s + x \\ y &\leftarrow r \mid (((x \oplus r) \ggg 2) \div s) \end{aligned} \tag{2}$$

A complete C procedure is given in Figure 2-1. It executes in seven basic RISC instructions, one of which is division. (Do not use this procedure with $x = 0$; that causes division by 0.)

```

unsigned snob(unsigned x) {
    unsigned smallest, ripple, ones;
    // x = xxx0 1111 0000
    smallest = x & -x;           // 0000 0001 0000
    ripple = x + smallest;       // xxx1 0000 0000
    ones = x ^ ripple;          // 0001 1111 0000
    ones = (ones >> 2)/smallest; // 0000 0000 0111
    return ripple | ones;       // xxx1 0000 0111
}

```

FIGURE 2-1. Next higher number with same number of 1-bits.

1. A variation of this algorithm appears in [H&S] section 7.6.7.

If division is slow but you have a fast way to compute the *number of trailing zeros* function $\text{ntz}(x)$, the *number of leading zeros* function $\text{nlz}(x)$, or *population count* ($\text{pop}(x)$ is the number of 1-bits in x), then the last line of Equation (2) can be replaced with one of the following:

$$y \leftarrow r \mid ((x \oplus r) \ggg (2 + \text{ntz}(x)))$$

$$y \leftarrow r \mid ((x \oplus r) \ggg (33 - \text{nlz}(s)))$$

$$y \leftarrow r \mid ((\mathbf{1} \ll (\text{pop}(x \oplus r) - 2)) - 1)$$

2-2 Addition Combined with Logical Operations

We assume the reader is familiar with the elementary identities of ordinary algebra and Boolean algebra. Below is a selection of similar identities involving addition and subtraction combined with logical operations:

- a. $-x = \neg x + \mathbf{1}$
- b. $\quad = \neg(x - \mathbf{1})$
- c. $\neg x = -x - \mathbf{1}$
- d. $-\neg x = x + \mathbf{1}$
- e. $\neg\neg x = x - \mathbf{1}$
- f. $x + y = x - \neg y - \mathbf{1}$
- g. $\quad = (x \oplus y) + 2(x \& y)$
- h. $\quad = (x \mid y) + (x \& y)$
- i. $\quad = 2(x \mid y) - (x \oplus y)$
- j. $x - y = x + \neg y + \mathbf{1}$
- k. $\quad = (x \oplus y) - 2(\neg x \& y)$
- l. $\quad = (x \& \neg y) - (\neg x \& y)$
- m. $\quad = 2(x \& \neg y) - (x \oplus y)$
- n. $x \oplus y = (x \mid y) - (x \& y)$
- o. $x \& \neg y = (x \mid y) - y$
- p. $\quad = x - (x \& y)$
- q. $\neg(x - y) = y - x - \mathbf{1}$
- r. $\quad = \neg x + y$
- s. $x \equiv y = (x \& y) - (x \mid y) - \mathbf{1}$
- t. $\quad = (x \& y) + \neg(x \mid y)$
- u. $x \mid y = (x \& \neg y) + y$
- v. $x \& y = (\neg x \mid y) - \neg x$

Equation (d) may be applied to itself repeatedly, giving $\neg\neg\neg\neg x = x + 2$, and so on. Similarly, from (e) we have $\neg\neg\neg\neg x = x - 2$. So we can add or subtract any constant, using only the two forms of complementation.

Equation (f) is the dual of (j), where (j) is the well known relation that shows how to build a subtracter from an adder.

Equations (g) and (h) are from HAKMEM memo [HAK item 23]. Equation (g) forms a sum by first computing the sum with carries ignored ($x \oplus y$), and then adding in the carries. Equation (h) is simply modifying the addition operands so that the combination $0 + 1$ never occurs at any bit position; it is replaced with $1 + 0$.

It can be shown that in the ordinary addition of binary numbers with each bit independently equally likely to be 0 or 1, a carry occurs at each position with probability about 0.5. However, for an adder built by preconditioning the inputs using (g), the probability is about 0.25. This observation is probably not of value in building an adder, because for that purpose the important characteristic is the maximum number of logic circuits the carry must pass through, and using (g) reduces the number of stages the carry propagates through by only one.

Equations (k) and (l) are duals of (g) and (h), for subtraction. That is, (k) has the interpretation of first forming the difference ignoring the borrows ($x \oplus y$), and then subtracting the borrows. Similarly, Equation (l) is simply modifying the subtraction operands so that the combination $1 - 1$ never occurs at any bit position; it is replaced with $0 - 0$.

Equation (n) shows how to implement *exclusive or* in only three instructions on a basic RISC. Using only *and-or-not* logic requires four instructions ($(x | y) \& \neg(x \& y)$). Similarly, (u) and (v) show how to implement *and* and *or* in three other elementary instructions, whereas using DeMorgan's laws requires four.

2-3 Inequalities among Logical and Arithmetic Expressions

Inequalities among binary logical expressions whose values are interpreted as unsigned integers are nearly trivial to derive. Here are two examples:

$$(x \oplus y) \stackrel{u}{\leq} (x | y), \text{ and}$$

$$(x \& y) \stackrel{u}{\leq} (x \equiv y).$$

These can be derived from a list of all binary logical operations, shown in Table 2-1.

Let $f(x, y)$ and $g(x, y)$ represent two columns in Table 2-1. If for each row in which $f(x, y)$ is 1, $g(x, y)$ also is 1, then for all (x, y) , $f(x, y) \stackrel{u}{\leq} g(x, y)$. Clearly, this extends to word-parallel logical operations. One can easily read off such relations (most of which are trivial) as $(x \& y) \stackrel{u}{\leq} x \stackrel{u}{\leq} (x | \neg y)$, and so on. Furthermore, if two columns have a row in which one entry is 0 and the other 1, and another row in which the entries are 1 and 0 respectively, then no inequality relation exists between the corresponding logical expressions. So the question of

TABLE 2-1. THE 16 BINARY LOGICAL OPERATIONS

<i>x</i>	<i>y</i>	0	<i>x</i> & <i>y</i>	<i>x</i> & ¬ <i>y</i>	<i>x</i>	¬ <i>x</i> & <i>y</i>	<i>y</i>	<i>x</i> ⊕ <i>y</i>	<i>x</i> <i>y</i>	¬(<i>x</i> <i>y</i>)	<i>x</i> ≡ <i>y</i>	¬ <i>y</i>	<i>x</i> ¬ <i>y</i>	¬ <i>x</i>	¬ <i>x</i> <i>y</i>	¬(<i>x</i> & <i>y</i>)	1
0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
0	1	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1
1	0	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1

whether or not $f(x, y) \stackrel{u}{\leq} g(x, y)$ is completely and easily solved for all binary logical functions f and g .

Use caution when manipulating these relations. For example, for ordinary arithmetic, if $x + y \leq a$ and $z \leq x$, then $z + y \leq a$. But this inference is not valid if “+” is replaced with *or*.

Inequalities involving mixed logical and arithmetic expressions are more interesting. Below is a small selection.

- a. $(x | y) \stackrel{u}{\geq} \max(x, y)$
- b. $(x \& y) \stackrel{u}{\leq} \min(x, y)$
- c. $(x | y) \stackrel{u}{\leq} x + y$ if the addition does not overflow
- d. $(x | y) \stackrel{u}{\geq} x + y$ if the addition overflows
- e. $|x - y| \stackrel{u}{\leq} (x \oplus y)$

The proofs of these are quite simple, except possibly for the relation $|x - y| \stackrel{u}{\leq} (x \oplus y)$. By $|x - y|$ we mean the absolute value of $x - y$, which may be computed within the domain of unsigned numbers as $\max(x, y) - \min(x, y)$. This relation may be proven by induction on the length of x and y (the proof is a little easier if you extend them on the left rather than on the right).

2-4 Absolute Value Function

If your machine does not have an instruction for computing the absolute value, this computation can usually be done in three or four branch-free instructions. First, compute $y \leftarrow x \ggg 31$, and then one of the following:

abs	nabs
$(x \oplus y) - y$	$y - (x \oplus y)$
$(x + y) \oplus y$	$(y - x) \oplus y$
$x - (2x \& y)$	$(2x \& y) - x$

By “ $2x$ ” we mean, of course, $x + x$ or $x \ll 1$.

If you have a fast multiply by a variable whose value is ± 1 , the following will do:

$$((x \ggg 30) | 1) * x$$

2-5 Sign Extension

By “sign extension,” we mean to consider a certain bit position in a word to be the sign bit, and we wish to propagate that to the left, ignoring any other bits present. The standard way to do this is with *shift left logical* followed by *shift right signed*. However, if these instructions are slow or nonexistent on your machine, it may be done with one of the following, where we illustrate by propagating bit position 7 to the left:

$$\begin{aligned} & ((x + 0x00000080) \& 0x000000FF) - 0x00000080 \\ & ((x \& 0x000000FF) \oplus 0x00000080) - 0x00000080 \end{aligned}$$

The “+” above can also be “-” or “ \oplus .” The second formula is particularly useful if you know that the unwanted high-order bits are all 0’s, because then the *and* can be omitted.

2-6 Shift Right Signed from Unsigned

If your machine does not have the *shift right signed* instruction, it may be computed using the formulas below. The first formula is from [GM], and the second is based on the same idea. Assuming the machine has mod 64 shifts, the first four formulas hold for $0 \leq n \leq 31$, and the last holds for $0 \leq n \leq 63$. The last formula holds for any n if by “holds” we mean “treats the shift amount to the same modulus as does the logical shift.”

When n is a variable, each formula requires five or six instructions on a basic RISC.

$$\begin{aligned} & ((x + 0x80000000) \ggg n) - (0x80000000 \ggg n) \\ & t \leftarrow 0x80000000 \ggg n; \quad ((x \ggg n) \oplus t) - t \\ & t \leftarrow (x \& 0x80000000) \ggg n; \quad (x \ggg n) - (t + t) \\ & (x \ggg n) | (-x \ggg 31) \ll 31 - n \\ & t \leftarrow -(x \ggg 31); \quad ((x \oplus t) \ggg n) \oplus t \end{aligned}$$

In the first two formulas, an alternative for the expression $0x80000000 \ggg n$ is $1 \ll 31 - n$.

If n is a constant, the first two formulas require only three instructions on many machines. If $n = 31$, the function can be done in two instructions with $-(x \ggg 31)$.

2-7 Sign Function

The *sign*, or *signum*, function is defined by

$$\text{sign}(x) = \begin{cases} -1, & x < 0, \\ 0, & x = 0, \\ 1, & x > 0. \end{cases}$$

It may be calculated with four instructions on most machines [Hop]:

$$(x \overset{s}{\ggg} 31) \mid (-x \overset{u}{\ggg} 31)$$

If you don't have *shift right signed*, then use the substitute noted at the end of Section 2-6, giving the following nicely symmetric formula (five instructions):

$$-(x \overset{u}{\ggg} 31) \mid (-x \overset{u}{\ggg} 31)$$

Comparison predicate instructions permit a three-instruction solution, with either

$$\begin{aligned} (x > 0) - (x < 0), \text{ or} \\ (x \geq 0) - (x \leq 0). \end{aligned} \tag{3}$$

Finally, we note that the formula $(-x \overset{u}{\ggg} 31) - (x \overset{u}{\ggg} 31)$ almost works; it fails only for $x = -2^{31}$.

2-8 Three-Valued Compare Function

The *three-valued compare* function, a slight generalization of the *sign* function, is defined by

$$\text{cmp}(x, y) = \begin{cases} -1, & x < y, \\ 0, & x = y, \\ 1, & x > y. \end{cases}$$

There are both signed and unsigned versions, and unless otherwise specified, this section applies to both.

Comparison predicate instructions permit a three-instruction solution, an obvious generalization of Equations (3):

$$\begin{aligned} &(\mathbf{x} > \mathbf{y}) - (\mathbf{x} < \mathbf{y}), \text{ or} \\ &(\mathbf{x} \geq \mathbf{y}) - (\mathbf{x} \leq \mathbf{y}). \end{aligned}$$

A solution for unsigned integers on PowerPC is shown below [CWG]. On this machine, “carry” is “not borrow.”

```
subf  R5,Ry,Rx  # R5 <-- Rx - Ry.
subfc R6,Rx,Ry  # R6 <-- Ry - Rx, set carry.
subfe R7,Ry,Rx  # R7 <-- Rx - Ry + carry, set carry.
subfe R8,R7,R5  # R8 <-- R5 - R7 + carry, (set carry).
```

If limited to the instructions of the basic RISC, there does not seem to be any particularly good way to compute this function. The comparison predicates $\mathbf{x} < \mathbf{y}$, $\mathbf{x} \leq \mathbf{y}$, and so on., require about five instructions (see Section 2-11), leading to a solution in about 12 instructions (using a small amount of commonality in computing $\mathbf{x} < \mathbf{y}$ and $\mathbf{x} > \mathbf{y}$). On the basic RISC it’s probably preferable to use compares and branches (six instructions executed worst case if compares can be commoned).

2-9 Transfer of Sign

The *transfer of sign* function, called ISIGN in Fortran, is defined by

$$\text{ISIGN}(x, y) = \begin{cases} \text{abs}(x), & y \geq 0, \\ -\text{abs}(x), & y < 0. \end{cases}$$

This function may be calculated (modulo 2^{32}) with four instructions on most machines:

$$\begin{aligned} t \leftarrow y \ggg 31; & & t \leftarrow (x \oplus y) \ggg 31; \\ \text{ISIGN}(x, y) = (\text{abs}(x) \oplus t) - t & & \text{ISIGN}(x, y) = (x \oplus t) - t \\ = (\text{abs}(x) + t) \oplus t & & = (x + t) \oplus t \end{aligned}$$

2-10 Decoding a “Zero Means 2^{**n} ” Field

Sometimes a zero or negative value does not make much sense for a quantity, so it is encoded in an n -bit field with a zero value being understood to mean 2^n , and a nonzero value having its normal binary interpretation. An example is the length field of PowerPC’s *load string word immediate* (lswi) instruction, which occupies five bits. It is not useful to have an instruction that loads zero bytes, when the

length is an immediate quantity, but it is definitely useful to be able to load 32 bytes. The length field could be encoded with values from 0 to 31 denoting lengths from 1 to 32, but the “zero means 32” convention results in simpler logic when the processor must also support a corresponding instruction with a variable (in-register) length that employs straight binary encoding (e.g., PowerPC’s `lswx` instruction).

It is trivial to encode an integer in the range 1 to 2^n into the “zero means 2^n ” encoding — simply mask the integer with $2^n - 1$. To do the decoding without a test-and-branch is not quite as simple, but below are some possibilities (no doubt overdone), illustrated for a 3-bit field. They all require three instructions, not counting possible loads of constants.

$$\begin{array}{lll}
 ((x - 1) \& 7) + 1 & ((x + 7) | -8) + 9 & 8 - (-x \& 7) \\
 ((x + 7) \& 7) + 1 & ((x + 7) | 8) - 7 & -(-x | -8) \\
 ((x - 1) | -8) + 9 & ((x - 1) \& 8) + x &
 \end{array}$$

2-11 Comparison Predicates

A “comparison predicate” is a function that compares two quantities, producing a single bit result of 1 if the comparison is **true**, and 0 if the comparison is **false**. Below we show branch-free expressions to evaluate the result into the sign position. To produce the $1/0$ value used by some languages (e.g., C), follow the code with a *shift right* of 31. To produce the $-1/0$ result used by some other languages (e.g., Basic), follow the code with a *shift right signed* of 31.

These formulas are, of course, not of interest on machines such as MIPS, the Compaq Alpha, and our model RISC, which have comparison instructions that compute many of these predicates directly, placing a 0/1-valued result in a general purpose register.

A machine instruction that computes the negative of the absolute value is handy here. We show this function as “nabs.” Unlike absolute value, it is well defined in that it never overflows. Machines that do not have “nabs” but have the more usual “abs” may use $-\text{abs}(x)$ for $\text{nabs}(x)$. If x is the maximum negative number, this overflows twice, but the result is correct. (We assume that the absolute value and the negation of the maximum negative number is itself.) Because some machines have neither “abs” nor “nabs,” we give an alternative that does not use them.

The “nlz” function is the number of leading 0’s in its argument. The “doz” function (*difference or zero*) is described on page 37.

$x = y$:	$\text{abs}(x - y) - 1$ $\text{abs}(x - y + \mathbf{0x80000000})$ $\text{nlz}(x - y) \lll 26$ $-(\text{nlz}(x - y) \ggg 5)$ $\neg(x - y \mid y - x)$	
$x \neq y$:	$\text{nabs}(x - y)$ $\text{nlz}(x - y) - 32$ $x - y \mid y - x$	
$x < y$:	$(x - y) \oplus [(x \oplus y) \& ((x - y) \oplus x)]$ $(x \& \neg y) \mid ((x \equiv y) \& (x - y))$ $\text{nabs}(\text{doz}(y, x))$	[GSO]
$x \leq y$:	$(x \mid \neg y) \& ((x \oplus y) \mid \neg(y - x))$ $((x \equiv y) \ggg 1) + (x \& \neg y)$	[GSO]
$x \stackrel{u}{<} y$:	$(\neg x \& y) \mid ((x \equiv y) \& (x - y))$ $(\neg x \& y) \mid ((\neg x \mid y) \& (x - y))$	
$x \stackrel{u}{\leq} y$:	$(\neg x \mid y) \& ((x \oplus y) \mid \neg(y - x))$	

For $x > y$, $x \geq y$, and so on, interchange x and y in the formulas for $x < y$, $x \leq y$, and so on. The *add* of $\mathbf{0x80000000}$ may be replaced with any instruction that inverts the high order bit (in x , y , or $x - y$).

Another class of formulas can be derived from the observation that the predicate $x < y$ is given by the sign of $x/2 - y/2$, and the subtraction in that expression cannot overflow. The result can be fixed up by subtracting 1 in the cases in which the shifts discard essential information, as follows:

$$\begin{aligned}
 x < y: & \quad (x \ggg 1) - (y \ggg 1) - (\neg x \& y \& 1) \\
 x \stackrel{u}{<} y: & \quad (x \ggg 1) - (y \ggg 1) - (\neg x \& y \& 1)
 \end{aligned}$$

These execute in seven instructions on most machines (six if it has *and not*), which is no better than what we have above (five to seven instructions depending upon the fullness of the set of logic instructions).

The formulas above involving “nlz” are due to [Shep], and his formula for the $x = y$ predicate is particularly useful because a minor variation of it gets the predicate evaluated to a 1/0-valued result with only three instructions:

$$\text{nlz}(x - y) \ggg 5.$$

Signed comparisons to 0 are frequent enough to deserve special mention. Below are some formulas for these, mostly derived directly from the above. Again, the result is in the sign position.

$$\begin{array}{ll}
 \mathbf{x} = \mathbf{0}: & \text{abs}(\mathbf{x}) - \mathbf{1} \\
 & \text{abs}(\mathbf{x} + \mathbf{0x80000000}) \\
 & \text{nlz}(\mathbf{x}) \ll 26 \\
 & \neg(\text{nlz}(\mathbf{x}) \ggg 5) \\
 & \neg(\mathbf{x} \mid \neg\mathbf{x}) \\
 & \neg\mathbf{x} \ \& \ (\mathbf{x} - \mathbf{1}) \\
 \mathbf{x} \neq \mathbf{0}: & \text{nabs}(\mathbf{x}) \\
 & \text{nlz}(\mathbf{x}) - \mathbf{32} \\
 & \mathbf{x} \mid \neg\mathbf{x} \\
 & (\mathbf{x} \ggg 1) - \mathbf{x} \quad [\text{CWG}] \\
 \mathbf{x} < \mathbf{0}: & \mathbf{x} \\
 \mathbf{x} \leq \mathbf{0}: & \mathbf{x} \mid (\mathbf{x} - \mathbf{1}) \\
 & \mathbf{x} \mid \neg\neg\mathbf{x} \\
 \mathbf{x} > \mathbf{0}: & \mathbf{x} \oplus \text{nabs}(\mathbf{x}) \\
 & \neg\mathbf{x} \ \& \ \neg\mathbf{x} \\
 \mathbf{x} \geq \mathbf{0}: & \neg\mathbf{x}
 \end{array}$$

Signed comparisons can be obtained from their unsigned counterparts by biasing the signed operands upwards by 2^{31} and interpreting the results as unsigned integers. The reverse transformation also works. Thus we have

$$\begin{aligned}
 \mathbf{x} < \mathbf{y} &= \mathbf{x} + 2^{31} \lll \mathbf{y} + 2^{31}, \\
 \mathbf{x} \lll \mathbf{y} &= \mathbf{x} - 2^{31} < \mathbf{y} - 2^{31}.
 \end{aligned}$$

Similar relations hold for \leq , \ggg , and so on. Addition and subtraction of 2^{31} are equivalent, as they amount to inverting the sign bit.

Another way to get signed comparisons from unsigned is based on the fact that if \mathbf{x} and \mathbf{y} have the same sign, then $\mathbf{x} < \mathbf{y} = \mathbf{x} \lll \mathbf{y}$, whereas if they have opposite signs, then $\mathbf{x} < \mathbf{y} = \mathbf{x} \ggg \mathbf{y}$ [Lamp]. Again, the reverse transformation also works, so we have

$$\begin{aligned}
 \mathbf{x} < \mathbf{y} &= (\mathbf{x} \lll \mathbf{y}) \oplus \mathbf{x}_{31} \oplus \mathbf{y}_{31} \quad \text{and} \\
 \mathbf{x} \lll \mathbf{y} &= (\mathbf{x} < \mathbf{y}) \oplus \mathbf{x}_{31} \oplus \mathbf{y}_{31},
 \end{aligned}$$

where x_{31} and y_{31} are the sign bits of x and y , respectively. Similar relations hold for \leq , $\stackrel{u}{\leq}$, and so on.

Using either of these devices enables computing all the usual comparison predicates other than $=$ and \neq in terms of any one of them, with at most three additional instructions on most machines. For example, let us take $x \stackrel{u}{\leq} y$ as primitive, because it is one of the simplest to implement (it is the carry bit from $y - x$). Then the other predicates can be obtained as follows:

$$x < y = \neg(y + 2^{31} \stackrel{u}{\leq} x + 2^{31})$$

$$x \leq y = x + 2^{31} \stackrel{u}{\leq} y + 2^{31}$$

$$x > y = \neg(x + 2^{31} \stackrel{u}{\leq} y + 2^{31})$$

$$x \geq y = y + 2^{31} \stackrel{u}{\leq} x + 2^{31}$$

$$x \stackrel{u}{<} y = \neg(y \stackrel{u}{\leq} x)$$

$$x \stackrel{u}{>} y = \neg(x \stackrel{u}{\leq} y)$$

$$x \stackrel{u}{\geq} y = y \stackrel{u}{\leq} x$$

Comparison Predicates from the Carry Bit

If the machine can easily deliver the carry bit into a general purpose register, this may permit concise code for some of the comparison predicates. Below are listed several of these relations. The notation $\text{carry}(expression)$ means the carry bit generated by the outermost operation in *expression*. We assume the carry bit for the subtraction $x - y$ is what comes out of the adder for $x + \bar{y} + 1$, which is the complement of "borrow."

$$x = y: \quad \text{carry}(\mathbf{0} - (x - y)), \text{ or } \text{carry}((x + \bar{y}) + \mathbf{1}), \text{ or } \text{carry}((x - y - \mathbf{1}) + \mathbf{1})$$

$$x \neq y: \quad \text{carry}((x - y) - \mathbf{1}), \text{ i.e., } \text{carry}((x - y) + (-\mathbf{1}))$$

$$x < y: \quad \neg \text{carry}((x + 2^{31}) - (y + 2^{31}))$$

$$x \leq y: \quad \text{carry}((y + 2^{31}) - (x + 2^{31}))$$

$$x \stackrel{u}{<} y: \quad \neg \text{carry}(x - y)$$

$$x \stackrel{u}{\leq} y: \quad \text{carry}(y - x)$$

$$x = \mathbf{0}: \quad \text{carry}(\mathbf{0} - x), \text{ or } \text{carry}(\bar{x} + \mathbf{1})$$

$$x \neq \mathbf{0}: \quad \text{carry}(x - \mathbf{1}), \text{ i.e., } \text{carry}(x + (-\mathbf{1}))$$

$$x < \mathbf{0}: \quad \text{carry}(x + x)$$

$$x \leq \mathbf{0}: \quad \text{carry}(2^{31} - (x + 2^{31}))$$

For $x > y$, use the complement of the expression for $x \leq y$, and similarly for other relations involving "greater than."

The GNU Superoptimizer has been applied to the problem of computing predicate expressions on the IBM RS/6000 computer and its close relative PowerPC [GK]. The RS/6000 has instructions for $\text{abs}(x)$, $\text{nabs}(x)$, $\text{doz}(x, y)$, and a number of forms of *add* and *subtract* that use the carry bit. It was found that the RS/6000 can compute all the integer predicate expressions with three or fewer elementary (one-cycle) instructions, a result that surprised even the architects of the machine. “All” includes the six two-operand signed comparisons, the four two-operand unsigned comparisons, all these with the second operand being zero, and all in forms that produce a 1/0 result or a -1/0 result. PowerPC, which lacks $\text{abs}(x)$, $\text{nabs}(x)$, and $\text{doz}(x, y)$, can compute all the predicate expressions in four or fewer elementary instructions.

How the Computer Sets the Comparison Predicates

Most computers have a way of evaluating the integer comparison predicates to a 1-bit result. The result bit may be placed in a “condition register” or, for some machines (such as our RISC model), in a general purpose register. In either case, the facility is often implemented by subtracting the comparison operands and then performing a small amount of logic on the result bits to determine the 1-bit comparison result.

Below is the logic for these operations. It is assumed that the machine computes $x - y$ as $x - \bar{y} + 1$, and the following quantities are available in the result:

C_o , the carry out of the high-order position

C_i , the carry into the high-order position

N , the sign bit of the result

Z , which equals 1 if the result, exclusive of C_o , is all-zero, and is otherwise 0

Then we have the following in Boolean algebra notation (juxtaposition denotes *and*, + denotes *or*):

$$\begin{array}{ll}
 V: & C_i \oplus C_o \quad (\text{signed overflow}) \\
 x = y: & Z \\
 x \neq y: & \bar{Z} \\
 x < y: & N \oplus V \\
 x \leq y: & (N \oplus V) + Z \\
 x > y: & (N \equiv V) \bar{Z} \\
 x \geq y: & N \equiv V \\
 x \stackrel{u}{<} y: & \bar{C}_o \\
 x \stackrel{u}{\leq} y: & \bar{C}_o + Z \\
 x \stackrel{u}{>} y: & C_o \bar{Z} \\
 x \stackrel{u}{\geq} y: & C_o
 \end{array}$$

2-12 Overflow Detection

“Overflow” means that the result of an arithmetic operation is too large or too small to be correctly represented in the target register. This section discusses methods that a programmer might use to detect when overflow has occurred, without using the machine’s “status bits” that are often expressly supplied for this purpose. This is important because some machines do not have such status bits (e.g., MIPS), and because even if the machine is so equipped, it is often difficult or impossible to access the bits from a high-level language.

Signed Add/Subtract

When overflow occurs on integer addition and subtraction, contemporary machines invariably discard the high-order bit of the result and store the low-order bits that the adder naturally produces. Signed integer overflow of addition occurs if and only if the operands have the same sign and the sum has sign opposite to that of the operands. Surprisingly, this same rule applies even if there is a carry into the adder—that is, if the calculation is $x + y + 1$. This is important for the application of adding multiword signed integers, in which the last addition is a signed addition of two fullwords and a carry-in that may be 0 or +1.

To prove the rule for addition, let x and y denote the values of the one-word signed integers being added, let c (carry-in) be 0 or 1, and assume for simplicity a 4-bit machine. Then if the signs of x and y are different,

$$\begin{aligned} -8 \leq x \leq -1, \text{ and} \\ 0 \leq y \leq 7, \end{aligned}$$

or similar bounds apply if x is nonnegative and y is negative. In either case, by adding these inequalities and optionally adding in 1 for c ,

$$-8 \leq x + y + c \leq 7.$$

This is representable as a 4-bit signed integer, and thus overflow does not occur when the operands have opposite signs.

Now suppose x and y have the same sign. There are two cases:

$$\begin{array}{ll} (a) & (b) \\ -8 \leq x \leq -1 & 0 \leq x \leq 7 \\ -8 \leq y \leq -1 & 0 \leq y \leq 7 \end{array}$$

Thus,

$$-16 \leq x + y + c \leq -1 \quad 0 \leq x + y + c \leq 15.$$

Overflow occurs if the sum is not representable as a 4-bit signed integer—that is, if

$$-16 \leq x + y + c \leq -9 \quad 8 \leq x + y + c \leq 15.$$

In case (a), this is equivalent to the high-order bit of the 4-bit sum being 0, which is opposite to the sign of x and y . In case (b), this is equivalent to the high-order bit of the 4-bit sum being 1, which again is opposite to the sign of x and y .

For subtraction of multiword integers, the computation of interest is $x - y - c$, where again c is 0 or 1, with a value of 1 representing a borrow-in. From an analysis similar to the above, it can be seen that overflow in the final value of $x - y - c$ occurs if and only if x and y have opposite signs and the sign of $x - y - c$ is opposite to that of x (or, equivalently, the same as that of y).

This leads to the following expressions for the overflow predicate, with the result being in the sign position. Following these with a *shift right* or *shift right signed* of 31 produces a 1/0- or a -1/0-valued result.

$x + y + c$	$x - y - c$
$(x \equiv y) \& ((x + y + c) \oplus x)$	$(x \oplus y) \& ((x - y - c) \oplus x)$
$((x + y + c) \oplus x) \& ((x + y + c) \oplus y)$	$((x - y - c) \oplus x) \& ((x - y - c) \equiv y)$

By choosing the second alternative in the first column, and the first alternative in the second column (avoiding the *equivalence* operation), our basic RISC can evaluate these tests with three instructions in addition to those required to compute $x + y + c$ or $x - y - c$. A fourth instruction (branch if negative) may be added to branch to code where the overflow condition is handled.

If executing with overflow interrupts enabled, the programmer may wish to test to see if a certain addition or subtraction will cause overflow, in a way that does not cause it. One branch-free way to do this is as follows:

$x + y + c$	$x - y - c$
$z \leftarrow (x \equiv y) \& \mathbf{0x80000000}$	$z \leftarrow (x \oplus y) \& \mathbf{0x80000000}$
$(x \equiv y) \& ((x \oplus z) + y + c) \equiv y$	$(x \oplus y) \& ((x \oplus z) - y - c) \oplus y$

The assignment to z in the left column sets $z = \mathbf{0x80000000}$ if x and y have the same sign, and sets $z = \mathbf{0}$ if they differ. Then, the addition in the second expression is done with x and y having different signs, so it can't overflow. If x and y are nonnegative, the sign bit in the second expression will be 1 if and only if $(x - 2^{31}) + y + c \geq \mathbf{0}$ —that is, iff $x + y + c \geq 2^{31}$, which is the condition for overflow in evaluating $x + y + c$. If x and y are negative, the sign bit in the second expression will be 1 iff $(x + 2^{31}) + y + c < \mathbf{0}$ —that is, iff $x + y + c < -2^{31}$, which again is the condition for overflow. The term $x \equiv y$ ensures the correct result (0 in the sign position) if x and y have opposite signs. Similar remarks apply to the case

of subtraction (right column). The code executes in nine instructions on the basic RISC.

It might seem that if the carry from addition is readily available, this might help in computing the signed overflow predicate. This does not seem to be the case. However, one method along these lines is as follows.

If x is a signed integer, then $x + 2^{31}$ is correctly represented as an unsigned number, and is obtained by inverting the high-order bit of x . Signed overflow in the positive direction occurs if $x + y \geq 2^{31}$ —that is, if $(x + 2^{31}) + (y + 2^{31}) \geq 3 \cdot 2^{31}$. This latter condition is characterized by carry occurring in the unsigned add (which means that the sum is greater than or equal to 2^{32}) and the high-order bit of the sum being 1. Similarly, overflow in the negative direction occurs if the carry is 0 and the high-order bit of the sum is also 0.

This gives the following algorithm for detecting overflow for signed addition:

Compute $(x \oplus 2^{31}) + (y \oplus 2^{31})$, giving sum s and carry c .
 Overflow occurred iff c equals the high-order bit of s .

The sum is the correct sum for the signed addition, because inverting the high-order bits of both operands does not change their sum.

For subtraction, the algorithm is the same except in the first step a subtraction replaces the addition. We assume that the carry is that generated by computing $x - y$ as $x + \bar{y} + 1$. The subtraction is the correct difference for the signed subtraction.

These formulas are perhaps interesting, but on most machines they would not be quite as efficient as the formulas that do not even use the carry bit (e.g., overflow = $(x \equiv y) \& (s \oplus x)$ for addition, and $(x \oplus y) \& (d \oplus x)$ for subtraction, where s and d are the sum and difference, respectively, of x and y).

How the Computer Sets Overflow for Signed Add/Subtract

Machines often set “overflow” for signed addition by means of the logic “the carry into the sign position is not equal to the carry out of the sign position.” Curiously, this logic gives the correct overflow indication for both addition and subtraction, assuming the subtraction $x - y$ is done by $x + \bar{y} + 1$. Furthermore, it is correct whether or not there is a carry- or borrow-in. This does not seem to lead to any particularly good methods for computing the signed overflow predicate in software, however, even though it is easy to compute the carry into the sign position. For addition and subtraction, the carry/borrow into the sign position is given by the sign bit after evaluating the following expressions (where c is **0** or **1**):

$$\begin{array}{cc} \text{carry} & \text{borrow} \\ (x + y + c) \oplus x \oplus y & (x - y - c) \oplus x \oplus y \end{array}$$

In fact, these expressions give, at each position i , the carry/borrow into position i .

Unsigned Add/Subtract

The following branch-free code may be used to compute the overflow predicate for unsigned add/subtract, with the result being in the sign position. The expressions involving a right shift are probably useful only when it is known that $c = 0$. The expressions in brackets compute the carry or borrow generated from the least significant position.

$$\begin{aligned} & \mathbf{x + y + c, \text{ unsigned}} \\ & (\mathbf{x \& y}) \mid ((\mathbf{x} \mid \mathbf{y}) \& \neg(\mathbf{x + y + c})) \\ & (\mathbf{x} \ggg 1) + (\mathbf{y} \ggg 1) + [((\mathbf{x \& y}) \mid ((\mathbf{x} \mid \mathbf{y}) \& \mathbf{c})) \& \mathbf{1}] \end{aligned}$$

$$\begin{aligned} & \mathbf{x - y - c, \text{ unsigned}} \\ & (\neg \mathbf{x} \& \mathbf{y}) \mid ((\mathbf{x} \equiv \mathbf{y}) \& (\mathbf{x - y - c})) \\ & (\neg \mathbf{x} \& \mathbf{y}) \mid ((\neg \mathbf{x} \mid \mathbf{y}) \& (\mathbf{x - y - c})) \\ & (\mathbf{x} \ggg 1) - (\mathbf{y} \ggg 1) - [((\neg \mathbf{x} \& \mathbf{y}) \mid ((\neg \mathbf{x} \mid \mathbf{y}) \& \mathbf{c})) \& \mathbf{1}] \end{aligned}$$

For unsigned *add*'s and *subtract*'s, there are much simpler formulas in terms of comparisons [MIPS]. For unsigned addition, overflow (carry) occurs if the sum is less (by unsigned comparison) than either of the operands. This and similar formulas are given below. Unfortunately, there is no way in these formulas to allow for a variable c that represents the carry- or borrow-in. Instead, the program must test c , and use a different type of comparison depending upon whether c is 0 or 1 .

$$\begin{array}{cccc} \mathbf{x + y, \text{ unsigned}} & \mathbf{x + y + 1, \text{ unsigned}} & \mathbf{x - y, \text{ unsigned}} & \mathbf{x - y - 1, \text{ unsigned}} \\ \neg \mathbf{x} \overset{u}{<} \mathbf{y} & \neg \mathbf{x} \overset{u}{\leq} \mathbf{y} & \mathbf{x} \overset{u}{<} \mathbf{y} & \mathbf{x} \overset{u}{\leq} \mathbf{y} \\ \mathbf{x + y} \overset{u}{<} \mathbf{x} & \mathbf{x + y + 1} \overset{u}{\leq} \mathbf{x} & \mathbf{x - y} \overset{u}{>} \mathbf{x} & \mathbf{x - y - 1} \overset{u}{\geq} \mathbf{x} \end{array}$$

The first formula for each case above is evaluated before the add/subtract that may overflow, and it provides a way to do the test without causing overflow. The second formula for each case is evaluated after the add/subtract that may overflow.

There does not seem to be a similar simple device (using comparisons) for computing the signed overflow predicate.

Multiplication

For multiplication, overflow means that the result cannot be expressed in 32 bits (it can always be expressed in 64 bits, whether signed or unsigned). Checking for overflow is simple if you have access to the high-order 32 bits of the product. Let

us denote the two halves of the 64-bit product by $hi(x \times y)$ and $lo(x \times y)$. Then the overflow predicates can be computed as follows [MIPS]:

$$\begin{array}{ll} x \times y, \text{ unsigned} & x \times y, \text{ signed} \\ hi(x \times y) \neq 0 & hi(x \times y) \neq (lo(x \times y) \ggg 31) \end{array}$$

One way to check for overflow of multiplication is to do the multiplication and then check the result by dividing. But care must be taken not to divide by 0, and there is a further complication for signed multiplication. Overflow occurs if the following expressions are **true**:

$$\begin{array}{ll} \text{Unsigned} & \text{Signed} \\ z \leftarrow x * y & z \leftarrow x * y \\ y \neq 0 \ \& \ z \overset{u}{\div} y \neq x & (y < 0 \ \& \ x = -2^{31}) \mid (y \neq 0 \ \& \ z \overset{u}{\div} y \neq x) \end{array}$$

The complication arises when $x = -2^{31}$ and $y = -1$. In this case the multiplication overflows, but the machine may very well give a result of -2^{31} . This causes the division to overflow and thus any result is possible (for some machines). Thus this case has to be checked separately, which is done by the term $y < 0 \ \& \ x = -2^{31}$. The above expressions use the “conditional *and*” operator to prevent dividing by zero (in C, use the $\&\&$ operator).

It is also possible to use division to check for overflow of multiplication without doing the multiplication (that is, without causing overflow). For unsigned integers, the product overflows iff $xy > 2^{32} - 1$, or $x > ((2^{32} - 1)/y)$, or, since x is an integer, $x > \lfloor (2^{32} - 1)/y \rfloor$. Expressed in computer arithmetic, this is

$$y \neq 0 \ \& \ x \overset{u}{>} (0\text{x}\text{FFFFFFFF} \overset{u}{\div} y).$$

For signed integers, the determination of overflow of $x * y$ is not so simple. If x and y have the same sign, then overflow occurs iff $xy > 2^{31} - 1$. If they have opposite signs, then overflow occurs iff $xy < -2^{31}$. These conditions may be tested as indicated in Table 2-2, which employs signed division.

TABLE 2-2. OVERFLOW TEST FOR SIGNED MULTIPLICATION

	$y > 0$	$y \leq 0$
$x > 0$	$x > 0\text{x}7\text{FFFFFFF} \div y$	$y < 0\text{x}80000000 \div x$
$x \leq 0$	$x < 0\text{x}80000000 \div y$	$x \neq 0 \ \& \ y < 0\text{x}7\text{FFFFFFF} \div x$

This test is awkward to implement because of the four cases. It is difficult to unify the expressions very much because of problems with overflow and with not being able to represent the number $+2^{31}$.

The test can be simplified if unsigned division is available. One can use the absolute values of x and y , which are correctly represented under unsigned integer interpretation. The complete test can then be computed as shown below. The variable $c = 2^{31} - 1$ if x and y have the same sign, and $c = 2^{31}$ otherwise.

$$\begin{aligned} c &\leftarrow ((x \equiv y) \ggg 31) + 2^{31} \\ x &\leftarrow \text{abs}(x) \\ y &\leftarrow \text{abs}(y) \\ y \neq 0 \ \& \ x > \frac{c}{y} \end{aligned}$$

The *number of leading zeros* instruction may be used to give an estimate of whether or not $x * y$ will overflow, and the estimate may be refined to give an accurate determination. First, consider the multiplication of unsigned numbers. It is easy to show that if x and y , as 32-bit quantities, have m and n leading 0's, respectively, then the 64-bit product has either $m + n$ or $m + n + 1$ leading 0's (or 64, if either $x = 0$ or $y = 0$). Overflow occurs if the 64-bit product has fewer than 32 leading 0's. Hence,

$\text{nlz}(x) + \text{nlz}(y) \geq 32$: Multiplication definitely does not overflow.

$\text{nlz}(x) + \text{nlz}(y) \leq 30$: Multiplication definitely does overflow.

For $\text{nlz}(x) + \text{nlz}(y) = 31$, overflow may or may not occur. In this case, the overflow assessment may be made by evaluating $t = x \lfloor y/2 \rfloor$. This will not overflow. Since xy is $2t$ or, if y is odd, $2t + x$, the product xy overflows if $t \geq 2^{31}$. These considerations lead to a plan for computing xy but branching to "overflow" if the product overflows. This plan is shown in Figure 2-2.

```

unsigned x, y, z, m, n, t;

m = nlz(x);
n = nlz(y);
if (m + n <= 30) goto overflow;
t = x*(y >> 1);
if ((int)t < 0) goto overflow;
z = t*2;
if (y & 1) {
    z = z + x;
    if (z < x) goto overflow;
}
// z is the correct product of x and y.

```

FIGURE 2-2. Determination of overflow of unsigned multiplication.

For the multiplication of signed integers, we can make a partial determination of whether or not overflow occurs from the number of leading 0's of nonnegative arguments, and the number of leading 1's of negative arguments. Let

$$\begin{aligned} m &= \text{nlz}(x) + \text{nlz}(\bar{x}), \text{ and} \\ n &= \text{nlz}(y) + \text{nlz}(\bar{y}). \end{aligned}$$

Then we have

$m + n \geq 34$: Multiplication definitely does not overflow.

$m + n \leq 31$: Multiplication definitely does overflow.

There are two ambiguous cases: 32 and 33. The case $m + n = 33$ overflows only when both arguments are negative and the true product is exactly 2^{31} (machine result is -2^{31}), so it can be recognized by a test that the product has the correct sign (that is, overflow occurred if $m \oplus n \oplus (m * n) < 0$). When $m + n = 32$, the distinction is not so easily made.

We will not dwell on this further, except to note that an overflow estimate for signed multiplication can also be made based on $\text{nlz}(\text{abs}(x)) + \text{nlz}(\text{abs}(y))$, but again there are two ambiguous cases (a sum of 31 or 32).

Division

For the signed division $x \div y$, overflow occurs if the following expression is **true**:

$$y = 0 \mid (x = 0x80000000 \ \& \ y = -1)$$

Most machines signal overflow (or trap) for the indeterminate form $0 \div 0$.

Straightforward code for evaluating this expression, including a final branch to the overflow handling code, consists of seven instructions, three of which are branches. There do not seem to be any particularly good tricks to improve on this, but below are a few possibilities:

$$[\text{abs}(y \oplus 0x80000000) \mid (\text{abs}(x) \ \& \ \text{abs}(y \equiv 0x80000000))] < 0$$

That is, evaluate the large expression in brackets, and branch if it is less than 0. This executes in about nine instructions, counting the load of the constant and the final branch, on a machine that has the indicated instructions and that gets the “compare to 0” for free.

Some other possibilities are to first compute z from

$$z \leftarrow (x \oplus 0x80000000) \mid (y + 1)$$

(three instructions on many machines), and then do the test and branch on $y = 0 \mid z = 0$ in one of the following ways:

$$((y \mid -y) \& (z \mid -z)) \geq 0$$

$$(\text{nabs}(y) \& \text{nabs}(z)) \geq 0$$

$$((\text{nlz}(y) \mid \text{nlz}(z)) \gg 5) \neq 0$$

These execute in nine, seven, and eight instructions, respectively, on a machine that has the indicated instructions. The last line represents a good method for PowerPC.

For the unsigned division $x \overset{u}{\div} y$, overflow occurs if and only if $y = 0$.

2-13 Condition Code Result of *Add*, *Subtract*, and *Multiply*

Many machines provide a “condition code” that characterizes the result of integer arithmetic operations. Often there is only one *add* instruction, and the characterization reflects the result for both unsigned and signed interpretation of the operands and result (but not for mixed types). The characterization usually consists of the following:

- Whether or not carry occurred (unsigned overflow)
- Whether or not signed overflow occurred
- Whether the 32-bit result, interpreted as a signed two’s-complement integer and ignoring carry and overflow, is negative, 0, or positive

Some older machines give an indication of whether the infinite precision result (that is, 33-bit result for *add*’s and *subtract*’s) is positive, negative, or 0. However, this indication is not easily used by compilers of high-level languages, and so has fallen out of favor.

For addition, only nine of the 12 combinations of these events are possible. The ones that cannot occur are “no carry, overflow, result > 0 ,” “no carry, overflow, result = 0,” and “carry, overflow, result < 0 .” Thus four bits are, just barely, needed for the condition code. Two of the combinations are unique in the sense that only one value of inputs produces them: Adding 0 to itself is the only way to get “no carry, no overflow, result = 0,” and adding the maximum negative number to itself is the only way to get “carry, overflow, result = 0.” These remarks remain true if there is a “carry in”—that is, if we are computing $x + y + 1$.

For subtraction, let us assume that to compute $x - y$ the machine actually computes $x + \bar{y} + 1$, with the carry produced as for an *add* (in this scheme the meaning of “carry” is reversed for subtraction, in that carry = 1 signifies that the result fits in a single word, and carry = 0 signifies that the result does not fit in a single word). Then for subtraction only seven combinations of events are possible.

The ones that cannot occur are the three that cannot occur for addition, plus “no carry, no overflow, result = 0” and “carry, overflow, result = 0.”

If a machine’s multiplier can produce a doubleword result, then two *multiply* instructions are desirable: one for signed and one for unsigned operands. (On a 4-bit machine, in hexadecimal, $\mathbf{F} \times \mathbf{F} = \mathbf{01}$ signed, and $\mathbf{F} \times \mathbf{F} = \mathbf{E1}$ unsigned). For these instructions, neither carry nor overflow can occur, in the sense that the result will always fit in a doubleword.

For a multiplication instruction that produces a one-word result (the low-order word of the doubleword result), let us take “carry” to mean that the result does not fit in a word with the operands and result interpreted as unsigned integers, and let us take “overflow” to mean that the result does not fit in a word with the operands and result interpreted as signed two’s-complement integers. Then again there are nine possible combinations of results, with the missing ones being “no carry, overflow, result > 0,” “no carry, overflow, result = 0,” and “carry, no overflow, result = 0.” Thus considering addition, subtraction, and multiplication together, ten combinations can occur.

2-14 Rotate Shifts

These are rather trivial. Perhaps surprisingly, this code works for n ranging from 0 to 32 inclusive, even if the shifts are mod-32.

$$\text{Rotate left } n: \quad y \leftarrow (x \ll n) \mid (x \gg (32 - n))$$

$$\text{Rotate right } n: \quad y \leftarrow (x \gg n) \mid (x \ll (32 - n))$$

2-15 Double-Length Add/Subtract

Using one of the expressions shown on page 29 for overflow of unsigned addition and subtraction, we can easily implement double-length addition and subtraction without accessing the machine’s carry bit. To illustrate with double-length addition, let the operands be (x_1, x_0) and (y_1, y_0) , and the result be (z_1, z_0) . Subscript 1 denotes the most significant half, and subscript 0 the least significant. We assume that all 32 bits of the registers are used. The less significant words are unsigned quantities.

$$z_0 \leftarrow x_0 + y_0$$

$$c \leftarrow [(x_0 \& y_0) \mid ((x_0 \mid y_0) \& \neg z_0)] \gg 31$$

$$z_1 \leftarrow x_1 + y_1 + c$$

This executes in nine instructions. The second line can be: $c \leftarrow (z_0 \ll x_0)$, permitting a four-instruction solution on machines that have this comparison operator in

a form that gives the result as a **1** or **0** in a register, such as the “SLTU” (*Set on Less Than Unsigned*) instruction on MIPS [MIPS].

Similar code for double-length subtraction ($x - y$) is

$$\begin{aligned} z_0 &\leftarrow x_0 - y_0 \\ b &\leftarrow [(-x_0 \& y_0) \mid ((x_0 \equiv y_0) \& z_0)] \overset{u}{\gg} 31 \\ z_1 &\leftarrow x_1 - y_1 - b \end{aligned}$$

This executes in eight instructions on a machine that has a full set of logical instructions. The second line can be $b \leftarrow (x_0 \overset{u}{<} y_0)$, permitting a four-instruction solution on machines that have the “SLTU” instruction.

Double-length addition and subtraction can be done in five instructions on most machines by representing the multiple-length data using only 31 bits of the least significant words, with the high-order bit being 0 except momentarily when it contains a carry or borrow bit.

2-16 Double-Length Shifts

Let (x_1, x_0) be a pair of 32-bit words to be shifted left or right as if they were a single 64-bit quantity, with x_1 being the most significant half. Let (y_1, y_0) be the result, interpreted similarly. Assume the shift amount n is a variable ranging from 0 to 63. Assume further that the machine’s shift instructions are modulo 64 or greater. That is, a shift amount in the range 32 to 63 or -32 to -1 results in an all-0 word, unless the shift is a signed right shift, in which case the result is 32 sign bits from the word shifted. (This code will not work on the Intel x86 machines, which have mod-32 shifts.)

Under these assumptions the *shift left double* operation may be accomplished as follows (eight instructions):

$$\begin{aligned} y_1 &\leftarrow x_1 \ll n \mid x_0 \overset{u}{\gg} (32 - n) \mid x_0 \ll (n - 32) \\ y_0 &\leftarrow x_0 \ll n \end{aligned}$$

The main connective in the first assignment must be *or*, not *plus*, to give the correct result when $n = 32$. If it is known that $0 \leq n \leq 32$, the last term of the first assignment may be omitted, giving a six-instruction solution.

Similarly, a *shift right double unsigned* operation may be done with

$$\begin{aligned} y_0 &\leftarrow x_0 \overset{u}{\gg} n \mid x_1 \ll (32 - n) \mid x_1 \overset{u}{\gg} (n - 32) \\ y_1 &\leftarrow x_1 \overset{u}{\gg} n \end{aligned}$$

Shift right double signed is more difficult, because of an unwanted sign propagation in one of the terms. Straightforward code follows:

```

if  $n < 32$  then  $y_0 \leftarrow x_0 \ggg^u n \mid x_1 \ll (32 - n)$ 
else  $y_0 \leftarrow x_1 \ggg^s (n - 32)$ 
 $y_1 \leftarrow x_1 \ggg^s n$ 

```

If your machine has the *conditional move* instructions, it is a simple matter to express this in branch-free code, in which form it takes eight instructions. If the conditional move instructions are not available, the operation may be done in ten instructions by using the familiar device of constructing a mask with the *shift right signed 31* instruction to mask the unwanted sign propagating term:

```

 $y_0 \leftarrow x_0 \ggg^u n \mid x_1 \ll (32 - n) \mid [(x_1 \ggg^s (n - 32)) \& ((32 - n) \ggg^s 31)]$ 
 $y_1 \leftarrow x_1 \ggg^s n$ 

```

2-17 Multibyte Add, Subtract, Absolute Value

Some applications deal with arrays of short integers (usually bytes or halfwords), and often execution is faster if they are operated on a word at a time. For definiteness, the examples here deal with the case of four 1-byte integers packed into a word, but the techniques are easily adapted to other packings, such as a word containing a 12-bit integer and two 10-bit integers, and so on. These techniques are of greater value on 64-bit machines, because more work is done in parallel.

Addition must be done in a way that blocks the carries from one byte into another. This can be accomplished by the following two-step method:

1. Mask out the high-order bit of each byte of each operand and *add* (there will then be no carries across byte boundaries).
2. Fix up the high-order bit of each byte with a 1-bit *add* of the two operands and the carry into that bit.

The carry into the high-order bit of each byte is of course given by the high-order bit of each byte of the sum computed in step 1. The subsequent similar method works for subtraction:

Addition

```

 $s \leftarrow (x \& \mathbf{0x7F7F7F7F}) + (y \& \mathbf{0x7F7F7F7F})$ 
 $(s \leftarrow ((x \oplus y) \& \mathbf{0x80808080}) \oplus s)$ 

```

Subtraction

```

 $d \leftarrow (x \mid \mathbf{0x80808080}) - (y \& \mathbf{0x7F7F7F7F})$ 
 $d \leftarrow ((x \oplus y) \mid \mathbf{0x7F7F7F7F}) \equiv d$ 

```

These execute in eight instructions, counting the load of **0x7F7F7F7F**, on a machine that has a full set of logical instructions. (Change the *and* and *or* of **0x80808080** to *and not* and *or not*, respectively, of **0x7F7F7F7F**.)

There is a different technique for the case in which the word is divided into only two fields. In this case, addition can be done by means of a 32-bit addition followed by subtracting out the unwanted carry. On page 28 we noted that the expression $(x + y) \oplus x \oplus y$ gives the carries into each position. Using this and similar observations about subtraction gives the following code for adding/subtracting two halfwords modulo 2^{16} (seven instructions):

Addition	Subtraction
$s \leftarrow x + y$	$d \leftarrow x - y$
$c \leftarrow (s \oplus x \oplus y) \& \mathbf{0x00010000}$	$b \leftarrow (d \oplus x \oplus y) \& \mathbf{0x00010000}$
$s \leftarrow s - c$	$d \leftarrow d + b$

Multibyte *absolute value* is easily done by complementing and adding 1 to each byte that contains a negative integer (that is, has its high-order bit on). The following code sets each byte of *y* equal to the absolute value of each byte of *x* (eight instructions):

```

a ← x & 0x80808080 // Isolate signs.
b ← a ≫ 7 // Integer 1 where x is negative.
m ← (a - b) | a // 0xFF where x is negative.
y ← (x ⊕ m) + b // Complement and add 1 where negative.

```

The third line could as well be $m \leftarrow a + a - b$. The addition of *b* in the fourth line cannot carry across byte boundaries, because the quantity $x \oplus m$ has a high-order 0 in each byte.

2-18 Doz, Max, Min

The “doz” function is “difference or zero,” defined as follows, for signed arguments:

$$\text{doz}(x, y) = \begin{cases} x - y, & x \geq y, \\ 0, & x < y. \end{cases}$$

It has been called “first grade subtraction,” because the result is 0 if you try to take away too much. We will use it to implement $\max(x, y)$ and $\min(x, y)$. In this connection it is important to note that $\text{doz}(x, y)$ can be negative; it is negative if the subtraction overflows. The *difference or zero* function can be used directly to implement the Fortran IDIM function, although in Fortran, results are generally undefined if overflow occurs.

There seems to be no very good way to implement $\text{doz}(x, y)$, $\text{max}(x, y)$, and $\text{min}(x, y)$ in a branch-free way that is applicable to most computers. About the best we can do is to compute $\text{doz}(x, y)$ using one of the expressions given on page 22 for the $x < y$ predicate, and then compute $\text{max}(x, y)$ and $\text{min}(x, y)$ from it, as follows:

$$\begin{aligned} d &\leftarrow x - y \\ \text{doz}(x, y) &= d \& [(d \equiv ((x \oplus y) \& (d \oplus x))) \ggg 31] \\ \text{max}(x, y) &= y + \text{doz}(x, y) \\ \text{min}(x, y) &= x - \text{doz}(x, y) \end{aligned}$$

This computes $\text{doz}(x, y)$ in seven instructions if the machine has *equivalence*, or eight if not, and it computes $\text{max}(x, y)$ or $\text{min}(x, y)$ in one more instruction.

The following are unsigned versions of these functions.

$$\begin{aligned} d &\leftarrow x - y \\ \text{dozu}(x, y) &= d \& \neg [((\neg x \& y) \mid ((x \equiv y) \& d)) \ggg 31] \\ \text{maxu}(x, y) &= y + \text{dozu}(x, y) \\ \text{minu}(x, y) &= x - \text{dozu}(x, y) \end{aligned}$$

The IBM RS/6000 computer, and its predecessor the 801, has $\text{doz}(x, y)$ provided as a single instruction. It permits computing the $\text{max}(x, y)$ and $\text{min}(x, y)$ of signed integers in two instructions, and is occasionally useful in itself. Implementing $\text{max}(x, y)$ and $\text{min}(x, y)$ directly is more costly because the machine would then need paths from the output ports of the register file back to an input port, bypassing the ALU.

Machines that have *conditional move* can get destructive² $\text{max}(x, y)$ and $\text{min}(x, y)$ in two instructions. For example, on our full RISC, $x \leftarrow \text{max}(x, y)$ can be calculated as follows (we write the target register first):

```

cmplt  z,x,y      Set z = 1 if x < y, else 0.
movne  x,z,y      If z is nonzero, set x = y.

```

2-19 Exchanging Registers

A very old trick is that of exchanging the contents of two registers without using a third [IBM]:

$$\begin{aligned} x &\leftarrow x \oplus y \\ y &\leftarrow y \oplus x \\ x &\leftarrow x \oplus y \end{aligned}$$

2. A destructive operation is one that overwrites one or more of its arguments.

This works well on a two-address machine. The trick also works if \oplus is replaced by the \equiv logical operation (complement of *exclusive or*), and can be made to work in various ways with *add*'s and *subtract*'s:

$$\begin{array}{lll} x \leftarrow x + y & x \leftarrow x - y & x \leftarrow y - x \\ y \leftarrow x - y & y \leftarrow y + x & y \leftarrow y - x \\ x \leftarrow x - y & x \leftarrow y - x & x \leftarrow x + y \end{array}$$

Unfortunately, each of these has an instruction that is unsuitable for a two-address machine, unless the machine has “reverse subtract.”

This little trick can actually be useful in the application of double buffering, in which two pointers are swapped. The first instruction can be factored out of the loop in which the swap is done (although this negates the advantage of saving a register):

$$\begin{array}{l} \text{Outside the loop: } t \leftarrow x \oplus y \\ \text{Inside the loop: } x \leftarrow x \oplus t \\ \qquad \qquad \qquad y \leftarrow y \oplus t \end{array}$$

Exchanging Corresponding Fields of Registers

The problem here is to exchange the contents of two registers x and y wherever a mask bit $m_i = 1$, and to leave x and y unaltered wherever $m_i = 0$. By “corresponding” fields, we mean that no shifting is required. The 1-bits of m need not be contiguous. The straightforward method is as follows:

$$\begin{array}{l} x' \leftarrow (x \& \bar{m}) \mid (y \& m) \\ y \leftarrow (y \& \bar{m}) \mid (x \& m) \\ x \leftarrow x' \end{array}$$

By using “temporaries” for the four *and* expressions, this can be seen to require seven instructions, assuming that either m or \bar{m} can be loaded with a single instruction and the machine has *and not* as a single instruction. If the machine is capable of executing the four (independent) *and* expressions in parallel, the execution time is only three cycles.

A method that is probably better (five instructions, but four cycles on a machine with unlimited instruction-level parallelism) is shown in column (a) below. It is suggested by the “three *exclusive or*” code for exchanging registers.


$$\begin{array}{lll} \text{(a)} & \text{(b)} & \text{(c)} \\ x \leftarrow x \oplus y & x \leftarrow x \equiv y & t \leftarrow (x \oplus y) \& m \\ y \leftarrow y \oplus (x \& m) & y \leftarrow y \equiv (x \mid \bar{m}) & x \leftarrow x \oplus t \\ x \leftarrow x \oplus y & x \leftarrow x \equiv y & y \leftarrow y \oplus t \end{array}$$

The steps in column (b) do the same exchange as that of column (a), but column (b) is useful if m does not fit in an immediate field but \bar{m} does, and the machine has the *equivalence* instruction.

Still another method is shown in column (c) above [GLS1]. It also takes five instructions (again assuming one instruction must be used to load m into a register), but executes in only three cycles on a machine with sufficient instruction-level parallelism.

Exchanging Two Fields of the Same Register

Assume a register x has two fields (of the same length) that are to be swapped, without altering other bits in the register. That is, the object is to swap fields B and D , without altering fields A , C , and E , in the computer word illustrated below. The fields are separated by a shift distance k .



Straightforward code would shift D and B to their new positions, and combine the words with *and* and *or* operations, as follows:

$$t_1 = (x \& m) \ll k$$

$$t_2 = (x \gg k) \& m$$

$$x' = (x \& m') \mid t_1 \mid t_2$$

Here, m is a mask with 1's in field D (and 0's elsewhere), and m' is a mask with 1's in fields A , C , and E . This code requires nine instructions and four cycles on a machine with unlimited instruction-level parallelism, allowing for two instructions to load the two masks.

A method that requires only seven instructions and executes in five cycles, under the same assumptions, is shown below [GLS1]. It is similar to the code in column (c) on page 39 for interchanging corresponding fields of two registers. Again, m is a mask that isolates field D .

$$t_1 = [x \oplus (x \gg k)] \& m$$

$$t_2 = t_1 \ll k$$

$$x' = x \oplus t_1 \oplus t_2$$

The idea is that t_1 contains $B \oplus D$ in position D (and 0's elsewhere), and t_2 contains $B \oplus D$ in position B . This code, and the straightforward code given earlier, work correctly if B and D are “split fields”—that is, if the 1-bits of mask m are not contiguous.

Conditional Exchange

The exchange methods of the preceding two sections that are based on *exclusive or* degenerate into no-operations if the mask m is 0. Hence they can perform an exchange of entire registers, or of corresponding fields of two registers, or of two fields of the same register, if m is set to all 1's if some condition c is true, and to all 0's if c is false. This gives branch-free code if m can be set up without branching.

2-20 Alternating among Two or More Values

Suppose a variable x can have only two possible values a and b , and you wish to assign to x the value other than its current one, and you wish your code to be independent of the values of a and b . For example, in a compiler x might be an opcode that is known to be either *branch true* or *branch false*, and whichever it is, you want to switch it to the other. The values of the opcodes *branch true* and *branch false* are arbitrary, probably defined by a C `#define` or `enum` declaration in a header file.

The straightforward code to do the switch is

```
if (x == a) x = b;
else x = a;
```

or, as often seen in C programs,

```
x = x == a ? b : a;
```

A far better (at least more efficient) way to code it is either

$$x \leftarrow a + b - x, \text{ or}$$

$$x \leftarrow a \oplus b \oplus x.$$

If a and b are constants, these require only one or two basic RISC instructions. Of course, overflow in calculating $a + b$ can be ignored.

This raises the question: Is there some particularly efficient way to cycle among three or more values? That is, given three arbitrary but distinct constants a , b , and c , we seek an easy-to-evaluate function f that satisfies

$$f(a) = b,$$

$$f(b) = c, \text{ and}$$

$$f(c) = a.$$

It is perhaps interesting to note that there is always a polynomial for such a function. For the case of three constants,

$$f(x) = \frac{(x-a)(x-b)}{(c-a)(c-b)}a + \frac{(x-b)(x-c)}{(a-b)(a-c)}b + \frac{(x-c)(x-a)}{(b-c)(b-a)}c. \quad (4)$$

(The idea is that if $x = a$, the first and last terms vanish, and the middle term simplifies to b , and so on.) This requires 14 arithmetic operations to evaluate, and, for arbitrary a , b , and c , the intermediate results exceed the computer's word size. But it is just a quadratic; if written in the usual form for a polynomial and evaluated using Horner's rule,³ it would require only five arithmetic operations (four for a quadratic with integer coefficients, plus one for a final division). Rearranging Equation (4) accordingly gives

$$\begin{aligned} f(x) = & \frac{1}{(a-b)(a-c)(b-c)} \{ [(a-b)a + (b-c)b + (c-a)c]x^2 \\ & + [(a-b)b^2 + (b-c)c^2 + (c-a)a^2]x \\ & + [(a-b)a^2b + (b-c)b^2c + (c-a)ac^2] \}. \end{aligned}$$

This is getting too complicated to be interesting (or practical).

Another method, similar to Equation (4) in that just one of the three terms survives, is:

$$f(x) = ((-(x=c)) \& a) + ((-(x=a)) \& b) + ((-(x=b)) \& c).$$

This takes 11 instructions if the machine has the *equal* predicate, not counting loads of constants. Because the two addition operations are combining two 0 values with a nonzero, they can be replaced with *or* or *exclusive or* operations.

The formula be simplified by precalculating $a - c$ and $b - c$, and then using [GLS1]:

$$\begin{aligned} f(x) &= ((-(x=c)) \& (a-c)) + ((-(x=a)) \& (b-c)) + c, \text{ or} \\ f(x) &= ((-(x=c)) \& (a \oplus c)) \oplus ((-(x=a)) \& (b \oplus c)) \oplus c, \end{aligned}$$

3. Horner's rule simply factors out x . For example, it evaluates the fourth degree polynomial $ax^4 + bx^3 + cx^2 + dx + e$ as $x(x(x(ax+b) + c) + d) + e$. For a polynomial of degree n it takes n multiplications and n additions, and it is very suitable for the *multiply-add* instruction.

Each of these operations takes eight instructions. But on most machines these are probably no better than the straightforward C code shown below, which executes in four to six instructions for small a , b , and c .

```
if (x == a) x = b;
else if (x == b) x = c;
else x = a;
```

Pursuing this matter, there is an ingenious branch-free method of cycling among three values on machines that do not have comparison predicate instructions [GLS1]. It executes in eight instructions on most machines.

Because a , b , and c are distinct, there are two bit positions n_1 and n_2 where the bits of a , b , and c are not all the same, and where the “odd one out” (the one whose bit differs in that position from the other two) is different in positions n_1 and n_2 . This is illustrated below for the values 21, 31, and 20, shown in binary.

1 0 1 0 1	c
1 1 1 1 1	a
1 0 1 0 0	b
n_1	n_2

Without loss of generality, rename a , b , and c so that a has the odd one out in position n_1 and b has the odd one out in position n_2 , as shown above. Then there are two possibilities for the values of the bits at position n_1 , namely $(a_{n_1}, b_{n_1}, c_{n_1}) = (0, 1, 1)$ or $(1, 0, 0)$. Similarly, there are two possibilities for the bits at position n_2 , namely $(a_{n_2}, b_{n_2}, c_{n_2}) = (0, 1, 0)$ or $(1, 0, 1)$. This makes four cases in all, and formulas for each of these cases are shown below.

Case 1, $(a_{n_1}, b_{n_1}, c_{n_1}) = (0, 1, 1)$, $(a_{n_2}, b_{n_2}, c_{n_2}) = (0, 1, 0)$:

$$f(x) = x_{n_1} * (a - b) + x_{n_2} * (c - a) + b.$$

Case 2, $(a_{n_1}, b_{n_1}, c_{n_1}) = (0, 1, 1)$, $(a_{n_2}, b_{n_2}, c_{n_2}) = (1, 0, 1)$:

$$f(x) = x_{n_1} * (a - b) + x_{n_2} * (a - c) + (b + c - a).$$

Case 3, $(a_{n_1}, b_{n_1}, c_{n_1}) = (1, 0, 0)$, $(a_{n_2}, b_{n_2}, c_{n_2}) = (0, 1, 0)$:

$$f(x) = x_{n_1} * (b - a) + x_{n_2} * (c - a) + a.$$

Case 4, $(a_{n_1}, b_{n_1}, c_{n_1}) = (1, 0, 0)$, $(a_{n_2}, b_{n_2}, c_{n_2}) = (1, 0, 1)$:

$$f(x) = x_{n_1} * (b - a) + x_{n_2} * (a - c) + c.$$

In these formulas, the left operand of each multiplication is a single bit. A multiplication by 0 or 1 may be converted into an *and* with a value of 0 or all 1's. Thus the formulas can be rewritten as illustrated below for the first formula.

$$f(x) = ((x \ll (31-n_1)) \ggg 31) \& (a - b) + ((x \ll (31-n_2)) \ggg 31) \& (c - a) + b$$

Because all variables except x are constants, this can be evaluated in eight instructions on the basic RISC. Here again the additions and subtractions can be replaced with *exclusive or*.

This idea can be extended to cycling among four or more constants. The essence of the idea is to find bit positions n_1, n_2, \dots , at which the bits uniquely identify the constants. For four constants, three bit positions always suffice. Then (for four constants) solve the following equation for s, t, u , and v (that is, solve the system of four linear equations in which $f(x)$ is a, b, c , or d , and the coefficients x_{n_i} are 0 or 1):

$$f(x) = x_{n_1}s + x_{n_2}t + x_{n_3}u + v$$

If the four constants are uniquely identified by only two bit positions, the equation to solve is

$$f(x) = x_{n_1}s + x_{n_2}t + x_{n_1}x_{n_2}u + v.$$