

## CHAPTER 3

## POWER-OF-2 BOUNDARIES

**3-1 Rounding Up/Down to a Multiple of a Known Power of 2**

Rounding an unsigned integer  $x$  down to, for example, the next smaller multiple of 8, is trivial:  $x \& -8$  does it. An alternative is  $(x \gg 3) \ll 3$ . These work for signed integers as well, provided “round down” means to round in the negative direction (e.g.,  $(-37) \& (-8) = -40$ ).

Rounding up is almost as easy. For example, an unsigned integer  $x$  can be rounded up to the next greater multiple of 8 with either of

$$(x + 7) \& -8, \text{ or}$$

$$x + (-x \& 7).$$

These expressions are correct for signed integers as well, provided “round up” means to round in the positive direction. The second term of the second expression is useful if you want to know how much you must add to  $x$  to make it a multiple of 8 [Gold].

To round a signed integer to the nearest multiple of 8 toward 0, you can combine the two expressions above in an obvious way:

$$t \leftarrow (x \gg 3) \& 7;$$

$$(x + t) \& -8$$

An alternative for the first line is  $t \leftarrow (x \gg 2) \gg 29$ , which is useful if the machine lacks *and immediate*, or if the constant is too large for its immediate field.

Sometimes the rounding factor is given as the  $\log_2$  of the alignment amount (e.g., a value of 3 means to round to a multiple of 8). In this case, code such as the following may be used, where  $k = \log_2(\text{alignment amount})$ :

$$\begin{array}{ll} \text{round down:} & x \& ((-1) \ll k) \\ & (x \gg k) \ll k \\ \text{round up:} & t \leftarrow (1 \ll k) - 1; (x + t) \& -t \\ & t \leftarrow (-1) \ll k; (x - t - 1) \& t \end{array}$$

### 3-2 Rounding Up/Down to the Next Power of 2

We define two functions that are similar to floor and ceiling, but which are directed roundings to the closest integral power of 2, rather than to the closest integer. Mathematically, they are defined by

$$\text{flp2}(x) = \begin{cases} \text{undefined, } x < 0, \\ 0, & x = 0, \\ 2^{\lfloor \log_2 x \rfloor}, & \text{otherwise;} \end{cases} \quad \text{clp2}(x) = \begin{cases} \text{undefined, } x < 0, \\ 0, & x = 0, \\ 2^{\lceil \log_2 x \rceil}, & \text{otherwise.} \end{cases}$$

The initial letters of the function names are intended to suggest “floor” and “ceiling.” Thus,  $\text{flp2}(x)$  is the greatest power of 2 that is  $\leq x$ , and  $\text{clp2}(x)$  is the least power of 2 that is  $\geq x$ . These definitions make sense even when  $x$  is not an integer (e.g.,  $\text{flp2}(0.1) = 0.0625$ ). The functions satisfy several relations analogous to those involving floor and ceiling, such as those shown below, where  $n$  is an integer.

$$\begin{aligned} \lfloor x \rfloor &= \lceil x \rceil \text{ iff } x \text{ is an integer} & \text{flp2}(x) &= \text{clp2}(x) \text{ iff } x \text{ is a power of 2 or is 0} \\ \lfloor x + n \rfloor &= \lfloor x \rfloor + n & \text{flp2}(2^n x) &= 2^n \text{flp2}(x) \\ \lceil x \rceil &= -\lfloor -x \rfloor & \text{clp2}(x) &= 1/\text{flp2}(1/x), \quad x \neq 0 \end{aligned}$$

Computationally, we deal only with the case in which  $x$  is an integer, and we take it to be unsigned, so the functions are well defined for all  $x$ . We require the value computed to be the arithmetically correct value modulo  $2^{32}$  (that is, we take  $\text{clp2}(x)$  to be  $0$  for  $x > 2^{31}$ ). The functions are tabulated below for a few values of  $x$ .

| $x$          | $\text{flp2}(x)$ | $\text{clp2}(x)$ |
|--------------|------------------|------------------|
| 0            | 0                | 0                |
| 1            | 1                | 1                |
| 2            | 2                | 2                |
| 3            | 2                | 4                |
| 4            | 4                | 4                |
| 5            | 4                | 8                |
| ...          | ...              | ...              |
| $2^{31} - 1$ | $2^{30}$         | $2^{31}$         |
| $2^{31}$     | $2^{31}$         | $2^{31}$         |
| $2^{31} + 1$ | $2^{31}$         | 0                |
| ...          | ...              | ...              |
| $2^{32} - 1$ | $2^{31}$         | 0                |

Functions `flp2` and `clp2` are connected by the relations shown below. These can be used to compute one from the other, subject to the indicated restrictions.

$$\begin{aligned} \text{clp2}(x) &= 2 \text{flp2}(x - 1), & x \neq 1, \\ &= \text{flp2}(2x - 1), & 1 \leq x \leq 2^{31}, \\ \text{flp2}(x) &= \text{clp2}(x \ll 2 + 1), & x \neq 0, \\ &= \text{clp2}(x + 1) \ll 2, & x < 2^{31}. \end{aligned}$$

The round-up and round-down functions can be computed quite easily with the *number of leading zeros* instruction, as shown below. However, for these relations to hold for  $x = 0$  and  $x > 2^{31}$ , the computer must have its shift instructions defined to produce 0 for shift amounts of -1, 32, and 63. Many machines (e.g., PowerPC) have “mod 64” shifts, which do this. In the case of -1, it is adequate if the machine shifts in the opposite direction (that is, a shift left of -1 becomes a shift right of 1).

$$\begin{aligned} \text{flp2}(x) &= 1 \ll (31 - \text{nlz}(x)) \\ &= 1 \ll (\text{nlz}(x) \oplus 31) \\ &= 0x80000000 \gg \text{nlz}(x) \\ \text{clp2}(x) &= 1 \ll (32 - \text{nlz}(x - 1)) \\ &= 0x80000000 \gg (\text{nlz}(x - 1) - 1) \end{aligned}$$

### Rounding Down

Figure 3-1 illustrates a branch-free algorithm that might be useful if *number of leading zeros* is not available. This algorithm is based on right-propagating the leftmost 1-bit, and executes in 12 instructions.

Figure 3-2 shows two simple loops that compute the same function. All variables are unsigned integers. The loop on the right keeps turning off the rightmost 1-bit of  $x$  until  $x = 0$ , and then returns the previous value of  $x$ .

```

unsigned flp2(unsigned x) {
    x = x | (x >> 1);
    x = x | (x >> 2);
    x = x | (x >> 4);
    x = x | (x >> 8);
    x = x | (x >> 16);
    return x - (x >> 1);
}

```

FIGURE 3-1. Greatest power of 2 less than or equal to  $x$ , branch-free.

```

y = 0x80000000;          do {
while (y > x)            y = x;
    y = y >> 1;         x = x & (x - 1);
return                  } while(x != 0);
                        return y;

```

FIGURE 3-2. Greatest power of 2 less than or equal to  $x$ , simple loops.

The loop on the left executes in  $4\text{nlz}(x) + 3$  instructions. The loop on the right, for  $x \neq 0$ , executes in  $4\text{pop}(x)$  instructions,<sup>1</sup> if the comparison to 0 is zero-cost.

### Rounding Up

The right-propagation trick yields a good algorithm for rounding up to the next power of 2. This algorithm, shown in Figure 3-3, is branch-free and runs in 12 instructions.

An attempt to compute this with the obvious loop does not work out very well:

```

y = 1;

while (y < x)    // Unsigned comparison.
    y = 2*y;
return y;

```

This code returns 1 for  $x = 0$ , which is probably not what you want, loops forever for  $x \geq 2^{31}$ , and executes in  $4n + 3$  instructions, where  $n$  is the power of 2 of the returned integer. Thus, it is slower than the branch-free code, in terms of instructions executed, for  $n \geq 3$  ( $x \geq 8$ ).

```

unsigned clp2(unsigned x) {
    x = x - 1;
    x = x | (x >> 1);
    x = x | (x >> 2);
    x = x | (x >> 4);
    x = x | (x >> 8);
    x = x | (x >> 16);
    return x + 1;
}

```

FIGURE 3-3. Least power of 2 greater than or equal to  $x$ .

1.  $\text{pop}(x)$  is the number of 1-bits in  $x$ .

### 3-3 Detecting a Power-of-2 Boundary Crossing

Assume memory is divided into blocks that are a power of 2 in size, starting at address 0. The blocks may be words, doublewords, pages, and so on. Then, given a starting address  $a$  and a length  $l$ , we wish to determine whether or not the address range from  $a$  to  $a + l - 1$ ,  $l \geq 2$ , crosses a block boundary. The quantities  $a$  and  $l$  are unsigned and any values that fit in a register are possible.

If  $l = 0$  or 1, a boundary crossing does not occur, regardless of  $a$ . If  $l$  exceeds the block size, a boundary crossing does occur, regardless of  $a$ . For very large values of  $l$  (wraparound is possible), a boundary crossing can occur even if the first and last bytes of the address range are in the same block.

There is a surprisingly concise way to detect boundary crossings on the IBM System/370 [CJS]. This method is illustrated below for a block size of 4096 bytes (a common page size).

```

O   RA, =A (-4096)
ALR RA, RL
BO  CROSSES

```

The first instruction forms the logical *or* of  $RA$  (which contains the starting address  $a$ ) and the number  $0xFFFFF000$ . The second instruction adds in the length, and sets the machine's 2-bit condition code. For the *add logical* instruction, the first bit of the condition code is set to 1 if a carry occurred, and the second bit is set to 1 if the 32-bit register result is nonzero. The last instruction branches if both bits are set. At the branch target,  $RA$  will contain the length that extends beyond the first page (this is an extra feature that was not asked for).

If, for example,  $a = 0$  and  $l = 4096$ , a carry occurs but the register result is 0, so the program properly does *not* branch to label CROSSES.

Let us see how this method can be adapted to RISC machines, which generally do not have *branch on carry and register result nonzero*. Using a block size of 8 for notational simplicity, the method of [CJS] branches to CROSSES if a carry occurred ( $(a \mid -8) + l \geq 2^{32}$ ) and the register result is nonzero ( $(a \mid -8) + l \neq 2^{32}$ ). Thus, it is equivalent to the predicate

$$(a \mid -8) + l > 2^{32}.$$

This in turn is equivalent to getting a carry in the final addition in evaluating  $((a \mid -8) - 1) + l$ . If the machine has *branch on carry*, this can be used directly, giving a solution in about five instructions counting a load of the constant  $-8$ .

If the machine does not have *branch on carry*, we can use the fact that carry occurs in  $x + y$  iff  $\neg x \stackrel{u}{<} y$  (see “Unsigned Add/Subtract” on page 29) to obtain the expression

$$\neg((a \mid -8) - 1) \stackrel{u}{<} l.$$

Using various identities such as  $\neg(x - 1) = -x$  gives the following equivalent expressions for the “boundary crossed” predicate:

$$\begin{aligned} & \neg(a \mid -8) \stackrel{u}{<} l \\ & \neg(a \mid -8) + 1 \stackrel{u}{<} l \\ & (\neg a \& 7) + 1 \stackrel{u}{<} l \end{aligned}$$

These can be evaluated in five or six instructions on most RISC computers.

Using another tack, clearly an 8-byte boundary is crossed iff

$$(a \& 7) + l - 1 \geq 8.$$

This cannot be directly evaluated because of the possibility of overflow (which occurs if  $l$  is very large), but it is easily rearranged to  $8 - (a \& 7) < l$ , which can be directly evaluated on the computer (no part of it overflows). This gives the expression

$$8 - (a \& 7) \stackrel{u}{<} l,$$

which can be evaluated in five instructions on most RISCs (four if it has *subtract from immediate*). If a boundary crossing occurs, the length that extends beyond the first block is given by  $l - (8 - (a \& 7))$ , which can be calculated with one additional instruction (*subtract*).